

Arm[®] Forge

Version 21.0.1

User Guide

The logo for Arm, consisting of the word "arm" in a bold, lowercase, sans-serif font.

Arm® Forge

User Guide

Copyright © 2021 Arm Limited or its affiliates. All rights reserved.

Release Information

Document History

Issue	Date	Confidentiality	Change
2100-00	01 March 2021	Non-Confidential	Document update to version 21.0
2101-00	30 March 2021	Non-Confidential	Document update to version 21.0.1

Non-Confidential Proprietary Notice

This document is protected by copyright and other related rights and the practice or implementation of the information contained in this document may be protected by one or more patents or pending patent applications. No part of this document may be reproduced in any form by any means without the express prior written permission of Arm. **No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document unless specifically stated.**

Your access to the information in this document is conditional upon your acceptance that you will not use or permit others to use the information for the purposes of determining whether implementations infringe any third party patents.

THIS DOCUMENT IS PROVIDED “AS IS”. ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, third party patents, copyrights, trade secrets, or other rights.

This document may include technical inaccuracies or typographical errors.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

This document consists solely of commercial items. You shall be responsible for ensuring that any use, duplication or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws. Use of the word “partner” in reference to Arm’s customers is not intended to create or refer to any partnership relationship with any other company. Arm may make changes to this document at any time and without notice.

If any of the provisions contained in these terms conflict with any of the provisions of any click through or signed written agreement covering this document with Arm, then the click through or signed written agreement prevails over and supersedes the conflicting provisions of these terms. This document may be translated into other languages for convenience, and you agree that if there is any conflict between the English version of this document and any translation, the terms of the English version of the Agreement shall prevail.

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners. Please follow Arm’s trademark usage guidelines at <http://www.arm.com/company/policies/trademarks>.

Copyright © 2021 Arm Limited (or its affiliates). All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

(LES-PRE-20349)

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by Arm and the party that Arm delivered this document to.

Unrestricted Access is an Arm internal classification.

Product Status

The information in this document is Final, that is for a developed product.

Web Address

developer.arm.com

Progressive terminology commitment

Arm values inclusive communities. Arm recognizes that we and our industry have used terms that can be offensive. Arm strives to lead the industry and create change.

This document includes terms that can be offensive. We will replace these terms in a future issue of this document.

If you find offensive terms in this document, please contact terms@arm.com.

Contents

Arm® Forge User Guide

Preface

About this book	22

Part A

Arm Forge

Chapter A1

Introduction to Arm Forge

A1.1 Arm DDT	A1-30
A1.2 Arm MAP	A1-31
A1.3 Arm Performance Reports	A1-32
A1.4 Online resources	A1-33

Chapter A2

Installation

A2.1 Linux graphical install	A2-36
A2.2 Linux text-mode install	A2-39
A2.3 Mac remote client installation	A2-40
A2.4 Windows remote client installation	A2-41

Chapter A3

Licensing

A3.1 Workstation and evaluation licenses	A3-44
A3.2 Supercomputing and other floating licenses	A3-45
A3.3 Architecture licensing	A3-46
A3.4 Environment variables	A3-47

Chapter A4

Connecting to a remote system

A4.1	Connecting remotely	A4-52
A4.2	Remote connections dialog	A4-53
A4.3	Reverse Connect	A4-56
A4.4	Treeserver or general debugging ports	A4-58
A4.5	Using X forwarding or VNC	A4-59
A4.6	Starting Arm® Forge	A4-60

Part B

DDT

Chapter B1

Get started with DDT

B1.1	Prepare a program for debugging	B1-64
B1.2	Express Launch (DDT)	B1-65
B1.3	Run a program	B1-67
B1.4	remote-exec required by some MPIs (DDT)	B1-71
B1.5	Open core files	B1-72
B1.6	Attach to running programs	B1-73
B1.7	Debug single-process programs	B1-78
B1.8	Debug OpenMP programs	B1-80
B1.9	Debug MPMD programs	B1-81
B1.10	Manual launch of multi-process non-MPI programs	B1-82
B1.11	Start a job in a queue	B1-84
B1.12	Job scheduling with jsrun	B1-85
B1.13	Use custom MPI scripts	B1-86
B1.14	Start DDT from a job script	B1-89
B1.15	Numactl (DDT)	B1-90
B1.16	Python debugging	B1-91
B1.17	DDT user interface	B1-93
B1.18	Save and load sessions	B1-95

Chapter B2

Source code

B2.1	Source code viewer	B2-98
B2.2	Assembly debugging	B2-100
B2.3	Project Files	B2-101
B2.4	Finding lost source files	B2-102
B2.5	Find Files or Functions	B2-103
B2.6	Go To Line	B2-105
B2.7	Navigate through source code history	B2-106
B2.8	Static analysis	B2-107
B2.9	Version control information	B2-108

Chapter B3

Control program execution

B3.1	Process control and process groups	B3-113
B3.2	Focus control	B3-115
B3.3	Start, stop, and restart a program	B3-118
B3.4	Step through a program	B3-119
B3.5	Stop messages	B3-120
B3.6	Set breakpoints	B3-121
B3.7	Conditional breakpoints	B3-123
B3.8	Suspend breakpoints	B3-124

B3.9	Delete breakpoints	B3-125
B3.10	Load and save breakpoints	B3-126
B3.11	Default breakpoints	B3-127
B3.12	Synchronize processes	B3-128
B3.13	Set watchpoints	B3-129
B3.14	Tracepoints	B3-131
B3.15	Version control breakpoints and tracepoints	B3-133
B3.16	Examine the stack frame	B3-135
B3.17	Align stacks	B3-136
B3.18	View stacks in parallel	B3-137
B3.19	Browse source code	B3-139
B3.20	View multiple files simultaneously	B3-141
B3.21	Signal handling	B3-142

Chapter B4

Variables and data

B4.1	Variables window	B4-146
B4.2	Sparklines	B4-147
B4.3	Current line	B4-148
B4.4	Local variables	B4-149
B4.5	Arbitrary expressions and global variables	B4-150
B4.6	Edit variables	B4-152
B4.7	Help with Fortran modules	B4-153
B4.8	View complex numbers in Fortran	B4-154
B4.9	C++ STL support	B4-155
B4.10	Custom pretty printers	B4-156
B4.11	View array data	B4-157
B4.12	UPC support	B4-158
B4.13	Change data values	B4-159
B4.14	View numbers in different bases	B4-160
B4.15	Examine pointers	B4-161
B4.16	Multi-dimensional arrays in the Variable View	B4-162
B4.17	Multi-dimensional array viewer (MDA)	B4-164
B4.18	Cross-process and cross-thread comparison	B4-172
B4.19	Assign MPI ranks	B4-174
B4.20	View registers	B4-175
B4.21	Process details	B4-176
B4.22	Disassembly	B4-177
B4.23	Interact directly with the debugger	B4-178

Chapter B5

Program input and output

B5.1	View standard output and error	B5-180
B5.2	Save output	B5-181
B5.3	Send standard input	B5-182

Chapter B6

Logbook

B6.1	Usage	B6-184
B6.2	Annotation	B6-185
B6.3	Logbook comparison	B6-186

Chapter B7

Message queues

B7.1	View message queues	B7-188

B7.2	Interpret message queues	B7-189
B7.3	Deadlock	B7-191

Chapter B8

Memory debugging

B8.1	Enable memory debugging	B8-194
B8.2	CUDA memory debugging	B8-195
B8.3	PMDK memory debugging	B8-196
B8.4	Memory debugging options	B8-197
B8.5	Pointer error detection and validity checking	B8-200
B8.6	Current memory usage	B8-204
B8.7	Memory Statistics	B8-206

Chapter B9

Use and write plugins

B9.1	Supported plugins	B9-208
B9.2	Install a plugin	B9-209
B9.3	Use a plugin	B9-210
B9.4	Write a plugin	B9-212
B9.5	Plugin reference	B9-213

Chapter B10

CUDA GPU debugging

B10.1	CUDA licensing	B10-216
B10.2	Prepare to debug GPU code	B10-217
B10.3	Launch the program	B10-218
B10.4	Control GPU threads	B10-219
B10.5	Examine GPU threads and data	B10-220
B10.6	GPU devices information	B10-222
B10.7	Attach to running GPU programs	B10-223
B10.8	Open GPU core files	B10-224
B10.9	Known issues and limitations	B10-225
B10.10	GPU language support	B10-227

Chapter B11

Offline debugging

B11.1	Use offline debugging	B11-230
B11.2	Offline report HTML output	B11-232
B11.3	Offline report plain text output	B11-235
B11.4	Run-time job progress reporting	B11-236

Part C

MAP

Chapter C1

Get started with MAP

C1.1	Welcome page	C1-240
C1.2	Express Launch (MAP)	C1-242
C1.3	Prepare a program for profiling	C1-244
C1.4	Profile a program	C1-250
C1.5	remote-exec required by some MPIs (MAP)	C1-255
C1.6	Profile a single-process program	C1-256
C1.7	Send standard input (MAP)	C1-257
C1.8	Start a job in a queue (MAP)	C1-258
C1.9	Use custom MPI scripts (MAP)	C1-259
C1.10	Start MAP from a job script	C1-262

	C1.11	Numactl (MAP)	C1-263
	C1.12	MAP environment variables	C1-264
Chapter C2		Program output	
	C2.1	View standard output and error (MAP)	C2-268
	C2.2	Display selected processes	C2-269
	C2.3	Restrict output	C2-270
	C2.4	Save output (MAP)	C2-271
Chapter C3		Source code (MAP)	
	C3.1	View source code (MAP)	C3-274
	C3.2	OpenMP programs	C3-276
	C3.3	GPU programs	C3-278
	C3.4	Complex code: code folding	C3-279
	C3.5	Edit source code (MAP)	C3-280
	C3.6	Rebuild and restart (MAP)	C3-281
	C3.7	Commit changes (MAP)	C3-282
Chapter C4		Selected lines view	
	C4.1	Use Selected lines view	C4-284
	C4.2	Limitations	C4-286
	C4.3	GPU profiles	C4-287
Chapter C5		Stacks view	
	C5.1	Overview	C5-290
Chapter C6		OpenMP Regions view	
	C6.1	Overview	C6-292
Chapter C7		Functions view	
	C7.1	Overview	C7-296
Chapter C8		Project Files view	
	C8.1	Overview	C8-298
Chapter C9		Metrics view	
	C9.1	User interface	C9-300
	C9.2	CPU instructions	C9-302
	C9.3	CPU time	C9-305
	C9.4	I/O	C9-306
	C9.5	Memory	C9-307
	C9.6	MPI calls	C9-308
	C9.7	Detecting MPI imbalance	C9-309
	C9.8	Accelerator	C9-310
	C9.9	Energy	C9-311
	C9.10	Requirements	C9-312
	C9.11	Lustre	C9-313
	C9.12	Zooming	C9-314
	C9.13	View totals across processes and nodes	C9-315
	C9.14	Custom metrics	C9-316

Chapter C10	Configurable Perf metrics	
	C10.1 Performance events	C10-318
	C10.2 Permissions	C10-319
	C10.3 Probe target hosts	C10-320
	C10.4 Specify Perf metrics using the command line	C10-321
	C10.5 Specify Perf metrics using a file	C10-322
	C10.6 Specify Perf metrics using the Run window	C10-323
	C10.7 View events	C10-324
	C10.8 Advanced configuration (MAP)	C10-325
Chapter C11	PAPI metrics	
	C11.1 PAPI installation	C11-328
	C11.2 PAPI config file	C11-329
	C11.3 PAPI overview metrics	C11-330
	C11.4 PAPI cache misses	C11-331
	C11.5 PAPI branch prediction	C11-332
	C11.6 PAPI floating-point	C11-333
Chapter C12	Main-thread, OpenMP, and Pthread view modes	
	C12.1 Main thread only mode	C12-336
	C12.2 OpenMP mode	C12-337
	C12.3 Pthread mode	C12-338
Chapter C13	Processes and cores window	
	C13.1 Overview	C13-340
Chapter C14	Run MAP from the command line	
	C14.1 Command line arguments	C14-342
	C14.2 Profile MPMD programs	C14-344
	C14.3 Profile MPMD programs without Express Launch	C14-345
Chapter C15	Export profiler data in JSON format	
	C15.1 JSON format	C15-348
	C15.2 Activities	C15-350
	C15.3 Metrics	C15-352
	C15.4 Example JSON output	C15-354
Chapter C16	GPU profiling	
	C16.1 GPU Kernels tab	C16-358
	C16.2 Kernel analysis	C16-359
	C16.3 Compilation	C16-361
	C16.4 Performance impact	C16-362
	C16.5 Customize GPU profiling behavior	C16-363
	C16.6 Known issues	C16-364
Chapter C17	Python profiling	
	C17.1 Profile a Python script	C17-366
	C17.2 Known issues	C17-369
Chapter C18	Performance analysis with Caliper instrumentation	
	C18.1 Get Caliper	C18-372

C18.2	Annotate your program	C18-373
C18.3	Analyze your program	C18-374
C18.4	Guidelines	C18-376

Chapter C19

Arm Statistical Profiling Extension (SPE)

C19.1	SPE Prerequisites	C19-378
C19.2	Enable SPE from the command line	C19-379
C19.3	Enable SPE from the run dialog	C19-380
C19.4	Analyze your program	C19-381
C19.5	Guidelines	C19-383
C19.6	Memlock limit exceeded	C19-384
C19.7	SPE on Amazon Web Services	C19-385
C19.8	Known issues	C19-386

Part D

Performance Reports

Chapter D1

Get started with Performance Reports

D1.1	Compilers for example programs	D1-390
D1.2	Compile on Cray X-series systems	D1-391
D1.3	Run an example program	D1-392
D1.4	Generate a performance report for an example program	D1-393

Chapter D2

Run real programs

D2.1	Link with a program	D2-396
D2.2	Dynamic linking on Cray X-Series systems (Performance Reports)	D2-397
D2.3	Static linking (Performance Reports)	D2-398
D2.4	Static linking on Cray X-Series systems (Performance Reports)	D2-400
D2.5	Dynamic and static linking on Cray X-Series systems using the modules environment (Performance Reports)	D2-401
D2.6	map-link modules installation on Cray X-Series systems (Performance Reports)	D2-402
D2.7	Unsupported user applications (Performance Reports)	D2-403
D2.8	Express Launch mode (Performance Reports)	D2-404
D2.9	Compatibility Launch mode	D2-405
D2.10	Generate a performance report for a real program	D2-406
D2.11	Specify output locations	D2-407
D2.12	Support for DCIM systems	D2-408
D2.13	Enable and disable metrics	D2-409

Chapter D3

Summarize an existing MAP file

D3.1	Summarize an existing MAP file	D3-412

Chapter D4

Interpret performance reports

D4.1	HTML performance reports	D4-414
D4.2	Report summary	D4-416
D4.3	CPU breakdown	D4-417
D4.4	CPU metrics breakdown	D4-418
D4.5	MPI breakdown	D4-420
D4.6	I/O breakdown	D4-421
D4.7	OpenMP breakdown	D4-423

D4.8	Threads breakdown	D4-424
D4.9	Memory breakdown	D4-425
D4.10	Energy breakdown	D4-426
D4.11	Accelerator breakdown	D4-427
D4.12	Textual performance reports	D4-428
D4.13	CSV performance reports	D4-429
D4.14	Worked examples	D4-430

Chapter D5

Configure Perf metrics

D5.1	Permissions (perf_event_paranoid)	D5-432
D5.2	Probe target hosts	D5-433
D5.3	Specify Perf metrics via the command line	D5-434
D5.4	Specify Perf metrics via a file	D5-435
D5.5	View events	D5-436
D5.6	Advanced configuration	D5-437

Part E

Appendix

Appendix A

Configuration

A.1	Configuring Performance Reports	Appx-A-442
A.2	Configuration files	Appx-A-443
A.3	Integration with queuing systems	Appx-A-445
A.4	Template tutorial	Appx-A-446
A.5	Connecting to compute nodes and remote programs (remote-exec)	Appx-A-448
A.6	Optional configuration	Appx-A-449

Appendix B

Get support

B.1	Supporting information	Appx-B-452

Appendix C

Supported platforms

C.1	Reference table	Appx-C-454

Appendix D

Known issues

D.1	Arm® MAP	Appx-D-456
D.2	XALT Wrapper	Appx-D-457
D.3	SLURM support	Appx-D-458
D.4	See also	Appx-D-459

Appendix E

MPI distribution notes and known issues

E.1	Cray MPT	Appx-E-462
E.2	HP MPI	Appx-E-464
E.3	Intel MPI	Appx-E-465
E.4	MPICH 3	Appx-E-467
E.5	MVAPICH 2	Appx-E-468
E.6	Open MPI	Appx-E-469
E.7	SLURM	Appx-E-470
E.8	IBM Spectrum MPI	Appx-E-471

Appendix F

Compiler notes and known issues

F.1	AMD OpenCL compiler	Appx-F-474
F.2	Arm Fortran compiler	Appx-F-475

F.3	Cray compiler environment	Appx-F-476
F.4	GNU	Appx-F-477
F.5	IBM XLC/XLF	Appx-F-478
F.6	Intel compilers	Appx-F-479
F.7	Portland Group and NVIDIA HPC SDK compilers	Appx-F-480

Appendix G

Platform notes and known issues

G.1	Cray	Appx-G-482
G.2	GNU/Linux systems	Appx-G-483
G.3	Intel Xeon	Appx-G-484
G.4	NVIDIA CUDA	Appx-G-485
G.5	Arm	Appx-G-486
G.6	POWER8 and POWER9 (ppc64le)	Appx-G-487
G.7	Mac OS X	Appx-G-488

Appendix H

General troubleshooting and known issues

H.1	GUI cannot connect to an X Server	Appx-H-490
H.2	Licenses	Appx-H-491
H.3	F1 cannot display this document	Appx-H-495
H.4	MPI not detected	Appx-H-496
H.5	Starting a program	Appx-H-500
H.6	Attaching	Appx-H-503
H.7	Source Viewer	Appx-H-505
H.8	Input/Output	Appx-H-506
H.9	Controlling a program	Appx-H-507
H.10	Evaluating variables	Appx-H-508
H.11	Memory debugging	Appx-H-509
H.12	MAP specific issues	Appx-H-511
H.13	Excessive memory use	Appx-H-517
H.14	Obtaining support	Appx-H-518

Appendix I

Queue template script syntax

I.1	Queue template tags	Appx-I-520
I.2	Defining new tags	Appx-I-521
I.3	Specifying default options	Appx-I-523
I.4	Launching	Appx-I-524
I.5	Using PROCS_PER_NODE_TAG	Appx-I-526
I.6	Job ID regular expression	Appx-I-527
I.7	Arm IPMI Energy Agent	Appx-I-528

List of Figures

Arm® Forge User Guide

Figure A2-1	Installer--Install type	A2-37
Figure A2-2	Installer-Destination	A2-38
Figure A2-3	Mac Installer - Installation folder	A2-40
Figure A2-4	Installer-Destination	A2-41
Figure A4-1	Connect-remotely	A4-52
Figure A4-2	Remotelaunch--Configure	A4-52
Figure A4-3	Remote connections dialog	A4-53
Figure A4-4	Remote Launch options	A4-53
Figure A4-5	Reverse connect request	A4-57
Figure B1-1	Express Launch Run dialog box	B1-66
Figure B1-2	Run window	B1-67
Figure B1-3	MPI rank error	B1-70
Figure B1-4	The Open Core Files window	B1-72
Figure B1-5	Attach window	B1-73
Figure B1-6	Attaching with Open MPI	B1-74
Figure B1-7	MPI rank error	B1-75
Figure B1-8	Hosts Window	B1-76
Figure B1-9	Single-process Run dialog	B1-78
Figure B1-10	Manual Launch window	B1-82
Figure B1-11	Manual Launch Process Groups	B1-82
Figure B1-12	Using custom MPI scripts	B1-86
Figure B1-13	Using substitute MPI commands	B1-88
Figure B1-14	DDT main window	B1-93

Figure B2-1	File edited warning	B2-98
Figure B2-2	The '0101' button in the toolbar toggles assembly debugging mode	B2-100
Figure B2-3	The Disassembly viewer tab shows the disassembly of the current symbol with the source code interleaved	B2-100
Figure B2-4	Function Listing	B2-101
Figure B2-5	Find Files or Functions dialog	B2-103
Figure B2-6	Find in Files dialog	B2-104
Figure B2-7	Static Analysis error annotation	B2-107
Figure B2-8	DDT running with version control information enabled	B2-108
Figure B2-9	Version control information-tooltips	B2-109
Figure B2-10	Version control information-context menu	B2-109
Figure B3-1	The Process Group Detailed view	B3-113
Figure B3-2	The Process Group Summary view	B3-114
Figure B3-3	Focus options	B3-115
Figure B3-4	The Process Group Detailed View focused on a process	B3-115
Figure B3-5	The Stepping threads window	B3-117
Figure B3-6	The Add Breakpoint window	B3-122
Figure B3-7	The Breakpoints table	B3-123
Figure B3-8	Conditional breakpoints in Fortran	B3-123
Figure B3-9	Program stopped at watchpoint	B3-129
Figure B3-10	The Watchpoints table	B3-129
Figure B3-11	Output from tracepoints in a Fortran program	B3-131
Figure B3-12	DDT with version control tracepoints	B3-133
Figure B3-13	Enable version control using the View menu	B3-133
Figure B3-14	Version Control-Trace at this revision	B3-134
Figure B3-15	The Current Stack tab	B3-135
Figure B3-16	Parallel Stack View	B3-137
Figure B3-17	Current frame highlighting in Parallel Stack View	B3-138
Figure B3-18	Parallel Stack View tool tip	B3-138
Figure B3-19	Right-click menu-variable options	B3-140
Figure B3-20	Right-click menu-function options	B3-140
Figure B3-21	Horizontal alignment of multiple source files	B3-141
Figure B3-22	Signal Handling dialog	B3-143
Figure B4-1	Variables in the Locals tab	B4-146
Figure B4-2	Evaluating Expressions	B4-150
Figure B4-3	Edit variables	B4-152
Figure B4-4	Viewing the Fortran complex number $3+4i$	B4-154
Figure B4-5	Edit Type window	B4-157
Figure B4-6	2D Array in C: type of array is <code>int[4][3]</code>	B4-162
Figure B4-7	2D Array in Fortran: type of <code>twodee</code> is <code>integer(3,5)</code>	B4-163
Figure B4-8	Multi-Dimensional Array Viewer	B4-164
Figure B4-9	<code>myArray(\$i, \$j)</code> with $\$i = 0..3$ and $\$j = 0..4$	B4-166
Figure B4-10	The local array <code>myArray(\$i, \$j)</code> with $\$i = 0..3$ and $\$j = 0..4$ on ranks 0-2	B4-167
Figure B4-11	A three-dimensional distributed array comprised of the local array <code>myArray(\$i, \$j)</code> with $\$i = 0..3$ and $\$j = 0..4$ on ranks 0-2 with $\$p$ the distributed dimension	B4-168
Figure B4-12	A three-dimensional distributed array comprised of the local array <code>myArray(\$i, \$j)</code> with $\$i = 0..3$ and $\$j = 0..4$ on ranks 0-2 projected onto the Data Table with $\$p$ (the distributed dimension) and $\$j$ displayed as Columns and $\$i$ displayed as Rows.	B4-169
Figure B4-13	DDT visualization	B4-170
Figure B4-14	Cross-Process Comparison View-Compare view	B4-172

Figure B4-15	Registers tab	B4-175
Figure B4-16	Process Details dialog	B4-176
Figure B4-17	Disassembly dialog	B4-177
Figure B4-18	Raw Command window	B4-178
Figure B5-1	Standard output	B5-180
Figure B5-2	Sending input	B5-182
Figure B6-1	Logbook example of a debug session	B6-184
Figure B6-2	Logbook Files Comparison window with tracepoint difference selected	B6-186
Figure B7-1	Message Queue window	B7-189
Figure B8-1	Memory Debugging Options	B8-197
Figure B8-2	Memory error message	B8-200
Figure B8-3	Pointer details	B8-201
Figure B8-4	Invalid memory message	B8-201
Figure B8-5	Memory usage graphs	B8-204
Figure B8-6	Memory Statistics	B8-206
Figure B10-1	GPU Thread Selector	B10-220
Figure B10-2	CUDA threads in the Parallel Stack View	B10-220
Figure B10-3	Kernel Progress View	B10-221
Figure B10-4	GPU Devices tab	B10-222
Figure B11-1	Offline mode HTML output	B11-232
Figure B11-2	Memory leak report	B11-233
Figure B11-3	Memory leak report detail	B11-234
Figure C1-1	MAP Welcome page	C1-241
Figure C1-2	Express Launch Run dialog box	C1-243
Figure C1-3	Run window	C1-250
Figure C1-4	Running window	C1-253
Figure C1-5	Single-process Run dialog	C1-256
Figure C1-6	Sending input	C1-257
Figure C1-7	Using custom MPI scripts	C1-259
Figure C1-8	Using substitute MPI commands	C1-261
Figure C2-1	Standard output in the Input/Output tab	C2-268
Figure C3-1	Source code viewer	C3-274
Figure C3-2	OpenMP Source code view	C3-276
Figure C3-3	GPU kernel Source code view	C3-278
Figure C3-4	Typical Fortran code in MAP	C3-279
Figure C3-5	Folded Fortran code in MAP	C3-279
Figure C3-6	File edited warning	C3-280
Figure C4-1	Selected lines view	C4-284
Figure C4-2	Selected lines view (GPU kernel)	C4-287
Figure C5-1	Stacks view	C5-290
Figure C6-1	OpenMP Regions view	C6-292
Figure C7-1	Functions view	C7-296
Figure C8-1	Project Files view	C8-298
Figure C9-1	Metrics view	C9-300
Figure C9-2	MAP with a region of time selected	C9-301
Figure C10-1	Configure Perf Metrics window	C10-323
Figure C13-1	Processes and Cores window	C13-340
Figure C16-1	GPU Kernels tab	C16-358
Figure C16-2	Run dialog with CUDA kernel analysis enabled	C16-359
Figure C16-3	GPU Kernels tab (with CUDA kernel analysis)	C16-360

Figure C17-1	Profiling a Python script	C17-367
Figure C17-2	Viewing Python profiling results	C17-368
Figure C18-1	Selected regions chart	C18-374
Figure C18-2	Regions tab	C18-374
Figure C18-3	View all selected regions	C18-375
Figure C19-1	The Run dialog with SPE profiling enabled	C19-380
Figure C19-2	SPE Table tab with a flat list of SPE samples	C19-381
Figure C19-3	SPE Table tab with SPE samples grouped by function	C19-381
Figure C19-4	SPE Table tab with SPE samples grouped by function	C19-381
Figure C19-5	Message about hidden SPE code editor annotations	C19-382
Figure C19-6	Message explaining why SPE code editor annotations are hidden	C19-382
Figure D4-1	Performance report for the wave_openmp.c example	D4-414
Figure D4-2	Performance report for the wave_openmp.c example (continued)	D4-415
Figure D4-3	Energy metrics report	D4-426
Figure D4-4	Accelerator metrics report	D4-427
Figure A-1	Queuing systems	Appx-A-445
Figure H-1	License error notification	Appx-H-491
Figure H-2	License error message	Appx-H-492
Figure H-3	No licences found error notification in [forge]	Appx-H-493
Figure H-4	Setting the MPI implementation in the GUI	Appx-H-497
Figure H-5	Setting the MPI Implementation in System Settings	Appx-H-498
Figure I-1	Queue parameters window	Appx-I-522

List of Tables

Arm® Forge User Guide

Table C15-1	JSON data structure	C15-348
Table I-1	Queue template tags	Appx-I-520
Table I-2	Supported attributes	Appx-I-521
Table I-3	Default options	Appx-I-523
Table I-4	Regular expression characters	Appx-I-527

Preface

This preface introduces the *Arm® Forge User Guide*.

It contains the following:

- [About this book on page 22.](#)

About this book

This document describes how to install and use Arm® Forge.

Using this book

This book is organized into the following chapters:

Part A Arm Forge

Describes how to install and get started using Arm Forge.

Chapter A1 Introduction to Arm Forge

This chapter provides an overview of Arm Forge.

Chapter A2 Installation

This chapter describes how to install Arm Forge on Linux, Windows, and Mac operating systems.

Chapter A3 Licensing

This chapter describes how to install and manage licenses for Arm Forge.

Chapter A4 Connecting to a remote system

This chapter describes the Arm Forge Remote Client that allows you to debug and profile remote jobs, while running the user interface on your local machine.

Part B DDT

Describes how to use Arm DDT.

Chapter B1 Get started with DDT

Learn how to get started using Arm DDT.

Chapter B2 Source code

Describes source code viewing, editing, and rebuilding features.

Chapter B3 Control program execution

Whether debugging multi-process code or single process code, the mechanisms for controlling program execution are very similar.

Chapter B4 Variables and data

This chapter describes the **Variables** window.

Chapter B5 Program input and output

The **Input/Output** tab is at the bottom of the window (by default).

Chapter B6 Logbook

This chapter describes how to use the logbook.

Chapter B7 Message queues

The **Message Queues** window shows the status of the message buffers of MPI.

Chapter B8 Memory debugging

The powerful parallel memory debugging feature intercepts calls to the system memory allocation library, recording memory usage, and confirming correct usage of the library by performing heap and bounds checking.

Chapter B9 Use and write plugins

Plugins are a quick and easy way to preload a library into your program, and define where to set breakpoints and tracepoints. A plugin is an XML file.

Chapter B10 CUDA GPU debugging

Arm DDT can be used to debug programs that use NVIDIA CUDA devices. The code running on the GPU is debugged simultaneously with the code on the host CPU.

Chapter B11 Offline debugging

Offline debugging is when a program is run under the control of the debugger, but without user intervention and without a user interface.

Part C MAP

Describes how to use Arm MAP.

Chapter C1 Get started with MAP

Learn how to get started using Arm MAP.

Chapter C2 Program output

Arm MAP collects and displays output from all processes on the **Input/Output** tab. Both standard output and error are shown.

Chapter C3 Source code (MAP)

Arm MAP provides source code viewing, editing, and rebuilding features.

Chapter C4 Selected lines view

The **Selected lines** view enables you to get detailed information about how one or more lines of code are spending their time.

Chapter C5 Stacks view

The **Stacks** view offers a good top-down view of your program.

Chapter C6 OpenMP Regions view

The **OpenMP Regions** view gives insight into the performance of every significant OpenMP region in your program.

Chapter C7 Functions view

The **Functions** view shows a flat profile of the functions in your program.

Chapter C8 Project Files view

The **Project Files** view offers an effective way to browse and navigate through a large, unfamiliar code base.

Chapter C9 Metrics view

This section describes how the **Metrics** view works with the **Source code viewer**, **Stacks** view, and **Project Files** view to help you identify and understand performance problems.

Chapter C10 Configurable Perf metrics

The Perf metrics use the Linux kernel `perf_event_open()` system call to provide additional CPU related metrics available for Arm MAP.

Chapter C11 PAPI metrics

The PAPI metrics are additional metrics available for Arm MAP that use the Performance Application Programming Interface (PAPI). They can be used on any system supported by PAPI.

Chapter C12 Main-thread, OpenMP, and Pthread view modes

The percentage values and activity graphs shown alongside the source code and in the **Stacks**, **OpenMP Regions**, and **Functions** views can present information for multithreaded programs in a variety of different ways.

Chapter C13 Processes and cores window

The **Processes and Cores** window enables you to change the number of cores per process.

Chapter C14 Run MAP from the command line

Arm MAP can be run from the command line.

Chapter C15 Export profiler data in JSON format

You can export the profiler data in machine readable JSON format.

Chapter C16 GPU profiling

You can use the GPU profiling capabilities when working with CUDA programs.

Chapter C17 Python profiling

The Python profiling capabilities can be used to find and resolve bottlenecks for your Python codes.

Chapter C18 Performance analysis with Caliper instrumentation

Caliper is a program instrumentation and performance measurement framework. It is a performance analysis toolbox in a library, that enables you to insert performance analysis capabilities directly into programs, and activate them at runtime.

Chapter C19 Arm Statistical Profiling Extension (SPE)

The Arm Statistical Profiling Extension (SPE) is an optional feature in ARMv8.2 hardware that allows CPU instructions to be sampled and associated with the source code location where that instruction occurred.

Part D Performance Reports

Describes how to use Arm Performance Reports.

Chapter D1 Get started with Performance Reports

Learn how to get started using Arm Performance Reports.

Chapter D2 Run real programs

This chapter shows you how to compile and run your own programs.

Chapter D3 Summarize an existing MAP file

Arm Performance Reports can be used to summarize an application profile generated by Arm MAP.

Chapter D4 Interpret performance reports

This chapter explains how to interpret the reports produced by Arm Performance Reports.

Chapter D5 Configure Perf metrics

The Perf metrics use the Linux kernel `perf_event_open()` system call to provide additional CPU related metrics available for Performance Reports.

Part E Appendix

The Appendix contains lots of detailed information to help you when you are working with Arm Forge.

Appendix A Configuration

Arm Forge shares a common configuration file between Arm DDT, Arm MAP, and Arm Performance Reports. This makes it easy for you to switch between tools without reconfiguring your environment each time.

Appendix B Get support

This appendix describes how to get help if you need it.

Appendix C Supported platforms

Lists the supported platforms for Arm Forge.

Appendix D Known issues

Summarizes the most significant known issues for the latest release.

Appendix E MPI distribution notes and known issues

This appendix has brief notes on many of the MPI distributions supported by Arm Forge.

Appendix F Compiler notes and known issues

When compiling for a Arm DDT debugging session, always compile with a minimal amount of optimization, or no optimization. Some compilers reorder instruction execution and omit debug information when compiling with optimization enabled.

Appendix G Platform notes and known issues

This appendix provides details about specific issues affecting platforms. If a supported machine is not listed in this section, then there are no known issues.

Appendix H General troubleshooting and known issues

This appendix offers help with common issues you might encounter while using Arm Forge.

Appendix I Queue template script syntax

This appendix contains information to help you with your script syntax.

Glossary

The Arm® Glossary is a list of terms used in Arm documentation, together with definitions for those terms. The Arm Glossary does not contain terms that are industry standard unless the Arm meaning differs from the generally accepted meaning.

See the [Arm® Glossary](#) for more information.

Typographic conventions

italic

Introduces special terminology, denotes cross-references, and citations.

bold

Highlights interface elements, such as menu names. Denotes signal names. Also used for terms in descriptive lists, where appropriate.

`monospace`

Denotes text that you can enter at the keyboard, such as commands, file and program names, and source code.

monospace

Denotes a permitted abbreviation for a command or option. You can enter the underlined text instead of the full command or option name.

monospace italic

Denotes arguments to monospace text where the argument is to be replaced by a specific value.

monospace bold

Denotes language keywords when used outside example code.

<and>

Encloses replaceable terms for assembler syntax where they appear in code or code fragments. For example:

```
MRC p15, 0, <Rd>, <CRn>, <CRm>, <Opcode_2>
```

SMALL CAPITALS

Used in body text for a few terms that have specific technical meanings, that are defined in the *Arm® Glossary*. For example, IMPLEMENTATION DEFINED, IMPLEMENTATION SPECIFIC, UNKNOWN, and UNPREDICTABLE.

Feedback

Feedback on this product

If you have any comments or suggestions about this product, contact your supplier and give:

- The product name.
- The product revision or version.
- An explanation with as much information as you can provide. Include symptoms and diagnostic procedures if appropriate.

Feedback on content

If you have comments on content then send an e-mail to errata@arm.com. Give:

- The title *Arm Forge User Guide*.
- The number 101136_21.0.1_00_en.
- If applicable, the page number(s) to which your comments refer.
- A concise explanation of your comments.

Arm also welcomes general suggestions for additions and improvements.

Note

Arm tests the PDF only in Adobe Acrobat and Acrobat Reader, and cannot guarantee the quality of the represented document when used with any other PDF reader.

Other information

- [*Arm® Developer.*](#)
- [*Arm® Documentation.*](#)
- [*Technical Support.*](#)
- [*Arm® Glossary.*](#)

Part A

Arm Forge

Chapter A1

Introduction to Arm Forge

This chapter provides an overview of Arm Forge.

Arm Forge combines Arm DDT for parallel high-performance application debugging, Arm MAP for performance profiling and optimization advice, and Arm Performance Reports for summarizing and characterizing both scalar and MPI application performance.

Arm Forge supports many parallel architectures and models, including MPI, CUDA and OpenMP. Arm Forge is a cross-platform tool, with support for the latest compilers and C++ standards, and Intel, 64-bit Arm, AMD, OpenPOWER and Nvidia GPU hardware.

Arm Forge provides you with everything you need to debug, fix, and profile programs at any scale. One common interface makes it easy to move between Arm DDT and Arm MAP during code development.

Arm Forge provides native remote clients for Windows, Mac OS X, and Linux. Use a remote client to connect to your cluster, where you can run, debug, profile, edit, and compile your application files.

It contains the following sections:

- [A1.1 Arm DDT on page A1-30.](#)
- [A1.2 Arm MAP on page A1-31.](#)
- [A1.3 Arm Performance Reports on page A1-32.](#)
- [A1.4 Online resources on page A1-33.](#)

A1.1 Arm DDT

Arm DDT is a powerful graphical debugger suitable for many different development environments, including:

- Single process and multithreaded software.
- OpenMP.
- Parallel (MPI) software.
- Heterogeneous software, for example, GPU software.
- Hybrid codes mixing paradigms, for example, MPI with OpenMP, or MPI with CUDA.
- Multi-process software including client-server applications.

Arm DDT helps you to find and fix problems on a single thread or across hundreds of thousands of threads. It includes static analysis to highlight potential code problems, integrated memory debugging to identify reads and writes that are outside of array bounds, and integration with MPI message queues.

Arm DDT supports:

- C, C++, and all derivatives of Fortran, including Fortran 90.
- Limited support for Python (CPython 2.7).
- Parallel languages/models including MPI, UPC, and Fortran 2008 Co-arrays.
- GPU languages such as HMPP, OpenMP Accelerators, CUDA and CUDA Fortran.

Related references

[Chapter B1 Get started with DDT on page B1-63](#)

A1.2 Arm MAP

Arm MAP is a parallel profiler that shows you the longest running lines of code, and explains why. Arm MAP does not require any complicated configuration, and you do not need to have experience with profiling tools to use it.

Arm MAP supports:

- MPI, OpenMP, and single-threaded programs.
- Small data files. All data is aggregated on the cluster and only a few megabytes written to disk, regardless of the size or duration of the run.
- Sophisticated source code view, enabling you to analyze performance across individual functions.
- Both interactive and batch modes for gathering profile data.
- A rich set of metrics, that show memory usage, floating-point calculations, and MPI usage across processes, including:
 - Percentage of vectorized instructions, including AVX extensions, used in each part of the code.
 - Time spent in memory operations, and how it varies over time and processes, to verify if there are any cache bottlenecks.
 - A visual overview across aggregated processes and cores that highlights any regions of imbalance in the code.

Related references

Chapter C1 Get started with MAP on page C1-239

A1.3 Arm Performance Reports

Arm Performance Reports is a low-overhead tool that produces one-page text and HTML reports summarizing and characterizing both scalar and MPI application performance.

Arm Performance Reports provides the most effective way to characterize and understand the performance of HPC application runs.

One single page HTML report answers a range of vital questions for any HPC site:

- Is this application optimized for the system it is running on?
- Does it benefit from running at this scale?
- Are there I/O or networking bottlenecks affecting performance?
- Which hardware, software or configuration changes can be made to improve performance further?

Arm Performance Reports is based on the Arm MAP low-overhead adaptive sampling technology that keeps data volumes collected and application overhead low:

- Runs transparently on optimized production-ready codes by adding a single command to your scripts.
- Just 5% application slowdown even with thousands of MPI processes.

Related references

Chapter D1 Get started with Performance Reports on page D1-389

Chapter D4 Interpret performance reports on page D4-413

A1.4 Online resources

- Get Arm Forge at [Arm Forge downloads](#).
- Find tutorials, webinars and white papers on the [Arm Developer website](#).
- [Help and tutorials](#)
- [Known issues](#)
- If you require more support, [get in touch](#) with our dedicated support team.

Chapter A2

Installation

This chapter describes how to install Arm Forge on Linux, Windows, and Mac operating systems.

To learn how to install and manage licenses for Arm Forge tools, and environment variables for managing product behavior, see [Licensing on page A3-43](#).

It contains the following sections:

- [A2.1 Linux graphical install on page A2-36](#).
- [A2.2 Linux text-mode install on page A2-39](#).
- [A2.3 Mac remote client installation on page A2-40](#).
- [A2.4 Windows remote client installation on page A2-41](#).

A2.1 Linux graphical install

Install remotely using the Arm Forge Linux graphical installer.

Procedure

1. Download the installation package from [Arm Forge Downloads](#)
2. Untar the installation package and run the installer executable with these commands:

```
tar xf arm-forge-<version>-linux-<arch>.tar
```

```
cd arm-forge-<version>-linux-<arch>
```

```
./installer
```

Note

Replace <version> with the four digit version number of your installation package, using this format xx.x.x. Replace <arch> with the required architecture (AArch64, x86_64, ppc64le). For example:
arm-forge-xx.x.x-linux-aarch64.tar

3. On the **Install Type** page, specify who can use Arm Forge on this system.

Note

Only an administrator (root) can install Arm Forge for **All Users** in a common directory, such as /opt or /usr/local. Otherwise, the **Just For Me** option is selected by default, and Arm Forge is installed in the local directory.




Figure A2-1 Installer--Install type

4. On the installer **Destination** page, enter the directory where you want to install Arm Forge.

————— **Note** —————

For a cluster installation, choose a directory that is shared between the cluster login or frontend node, and the compute nodes. Alternatively, install or copy into the same location on each node.




Figure A2-2 Installer-Destination

5. Installation progress displays on the **Install** page.

Results

The installation adds icons for Arm DDT and Arm MAP to the **Development menu** in your desktop environment. When installation is complete, read the instructions in the `RELEASE-NOTES` file in the install package, for details about how to run Arm DDT, Arm MAP, and Arm Performance Reports.

Next steps

You can obtain an evaluation license at [Get software](#).

Arm Forge supports a large number of different site configurations and MPI distributions. You must ensure that you fully integrate all components into your environment. For example, ensure that you propagate environment variables to remote nodes, and that tool libraries and executables are available on the remote nodes.

Related references

[A2.2 Linux text-mode install](#) on page A2-39

[Chapter A3 Licensing](#) on page A3-43

[Chapter B1 Get started with DDT](#) on page B1-63

[Chapter C1 Get started with MAP](#) on page C1-239

[Chapter D1 Get started with Performance Reports](#) on page D1-389

A2.2 Linux text-mode install

Install remotely using the `textinstall.sh` text-mode install script.

Procedure

1. Download the required package from the Arm Forge downloads webpage: [Arm Forge Downloads](#).
2. Untar the installation package and run the `textinstall.sh` script by using these commands:

```
tar xf arm-forge-<version>-linux-<arch>.tar
cd arm-forge-<version>-linux-<arch>
./textinstall.sh
```

Note

- Replace `<version>` with the four digit version number of your installation package using this format `xx.x` for major product releases, and this format for support releases `xx.x.x`.
 - Replace `<arch>` with the required architecture (`aarch64`, `x86_64`, `ppc64le`).
3. When you are prompted, press **Return** to read the license, and enter the path of the installation directory. The directory must be accessible on all the nodes in your cluster.

Next steps

For details about how to run Arm DDT, Arm MAP, and Arm Performance Reports, see the `RELEASE-NOTES` file in the install package when the installation is complete.

You can obtain an evaluation license at [Get software](#).

Arm Forge supports a large number of different site configurations and MPI distributions, and therefore you must ensure that you fully integrate all components into your environment. For example, propagate environment variables to remote nodes, and ensure that the tool libraries and executables are available on them.

Related references

[A2.1 Linux graphical install](#) on page A2-36

[Chapter A3 Licensing](#) on page A3-43

[Chapter B1 Get started with DDT](#) on page B1-63

[Chapter C1 Get started with MAP](#) on page C1-239

[Chapter D1 Get started with Performance Reports](#) on page D1-389

A2.3 Mac remote client installation

The Mac OS X Arm Forge installation package is a remote client only for connecting to an Arm Forge installation. The Arm Forge remote client is supplied as an Apple Disk Image (.dmg) file.

You do not need to install a license file on a machine running the remote client when you connect remotely to Arm Forge.

Procedure

1. Download the required package from [Arm Forge Downloads](#).

The .dmg file includes the Documentation folder and the client application bundle icon. The Documentation folder contains a copy of this user guide and the release notes.

2. Drag and drop the client application bundle icon into the Applications directory.




Figure A2-3 Mac Installer - Installation folder

Related references

Chapter A3 Licensing on page A3-43

Chapter B1 Get started with DDT on page B1-63

Chapter C1 Get started with MAP on page C1-239

Chapter D1 Get started with Performance Reports on page D1-389

A2.4 Windows remote client installation

The Windows installation package is a remote client only for connecting to an Arm Forge installation. The Arm Forge remote client is supplied as a Windows executable (.exe) file.

You do not need to install a license file on a machine running Remote Client for connecting remotely to Arm Forge.

Procedure

1. Download the Remote Client for Windows installation package from the Arm Forge downloads webpage: [Arm Forge Downloads](#).
2. Run the Windows file executable to install the Arm Forge remote client.
3. On the installer **Destination** page, enter the directory where you want to install Arm Forge.




Figure A2-4 Installer-Destination

————— Note —————

If your user account has administrative privileges, then by default Arm Forge is installed under C:\Program Files. If your account does not have administrative privileges, then by default Arm Forge is installed under C:\Users\%USERNAME%\AppData\Local.

Related references

[Chapter A3 Licensing](#) on page A3-43

Chapter B1 Get started with DDT on page B1-63

Chapter C1 Get started with MAP on page C1-239

Chapter D1 Get started with Performance Reports on page D1-389

Chapter A3

Licensing

This chapter describes how to install and manage licenses for Arm Forge.

Before you begin:

- You must install a license file on a machine running Arm Forge tools.
- You do not need to install a license file on a machine running Arm Forge Remote Client for connecting remotely to Arm Forge tools on a remote system.

Note

If you do not have a valid license file, the user interface shows an alert in the lower-left corner, and you can not run, debug, or profile new programs.

Time-limited evaluation licenses are available from the Arm website: [Get software](#).

It contains the following sections:

- [A3.1 Workstation and evaluation licenses](#) on page A3-44.
- [A3.2 Supercomputing and other floating licenses](#) on page A3-45.
- [A3.3 Architecture licensing](#) on page A3-46.
- [A3.4 Environment variables](#) on page A3-47.

A3.1 Workstation and evaluation licenses

Workstation and evaluation license files for Arm Forge do not require Arm Licence Server.

Arm Forge supports separate license files for Arm DDT, Arm MAP, and Arm Performance Reports with a single installation of Arm Forge.

If there are multiple licenses installed for the same product, Arm Forge uses the license with the most tokens.

Do not edit license files because this prevents them from working.

Procedure

1. Copy your license files to {installation-directory}/licences. For example, /home/user/arm/forged/<version>/licences/Licence.ddt.

————— **Note** —————

Arm Forge accepts directories spelled `licence` or `license`. For example, both of these paths are valid license file locations: /opt/arm/licenses and /opt/arm/licences.

2. When Arm Forge starts, select Arm DDT or Arm MAP on the **Welcome** page.

If you would prefer to store the license files in an alternative location, use the environment variable ALLINEA_LICENCE_DIR to specify it. For example:

```
export ALLINEA_LICENCE_DIR=${HOME}/SomeOtherLicenceDir
```

————— **Note** —————

You can use ALLINEA_LICENSE_DIR, ALLINEA_LICENCE_DIR, ARM_LICENSE_DIR, and ARM_LICENCE_DIR interchangeably.

Related references

[A3.2 Supercomputing and other floating licenses on page A3-45](#)

Related information

[Get software](#)

[Arm support](#)

A3.2 Supercomputing and other floating licenses

If you are using floating licenses for an HPC cluster, you need to use Arm Licence Server. Download the required package from the Arm Forge downloads page on the [Arm Forge Downloads](#) page.

A floating license consists of two files:

- Server license (Licence.xxx)
- Client license (Licence)

You can find more information on the [Arm Licence Server](#) web page.

Procedure

1. Copy the client file (Licence) to <installation-directory>. For example,

```
/home/user/arm/forg/20.2.1/licences/Licence
```
2. Edit the hostname line to contain the host name or IP address of the machine running the Arm Licence Server.
3. See the Arm Licence Server user guide on the [Arm Licence Server](#) web page for instructions on how to install the server license.
4. Ensure that Arm Licence Server is running on the designated license server machine before you run Arm Forge.

Related information

A3.3 Architecture licensing

Licenses issued after the release of Arm Forge 6.1, specify the compute node architectures with which they can be used.

Licenses issued before this release enable the x86_64 architecture by default. If you are using another architecture, you will receive a new license to enable your architecture.

Contact [Arm support](#) if you have any licensing issues.

Using multiple architecture licenses

If you are using multiple license files to specify multiple architectures, Arm recommends that you follow these steps.

Procedure

1. Ensure that the default licenses directory is empty.
2. Create a directory for each architecture.
3. When you want to target a specific architecture, set ALLINEA_LICENSE_DIR to the relevant directory. Alternatively, set ALLINEA_LICENSE_FILE to specify the license file.

Example

On a site that targets two architectures, x86_64 and AArch64, create a directory for each architecture, and name them licenses_x86_64 and licenses_aarch64. Then, to target the architectures, set the license directories as follows:

To target AArch64:

```
export ALLINEA_LICENSE_DIR=/path/to/licenses/licenses_aarch64
```

To target x86_64:

```
export ALLINEA_LICENSE_DIR=/path/to/licenses/licenses_x86_64
```

A3.4 Environment variables

This section provides a reference to Arm Forge environment variables, with guidance on requirements and best practices for using them.

Performance Report customization

Environment variables to customize your reports.

Variable	Description
ALLINEA_NOTES	Any text in this environment variable is also included in all reports.
ALLINEA_MAP_TO_DCIM	Allows you to specify a <code>.map</code> file when using the <code>-dcim-output</code> argument.
ALLINEA_DCIM_SCRIPT	The path to the script to use for communicating with DCIM. The default is: <code>\\${ALLINEA_TOOLS_PATH}/performance-reports/ganglia-connector/pr-dcim</code>
ALLINEA_GMETRIC	Path to the <code>gmetric</code> instance to use. This is specific to the <code>pr-dcim</code> script. The default is <code>which gmetric</code> .

Warning suppression

Environment variables for warning suppression. Use these when autodetection generates erroneous messages.

Variable	Description
ALLINEA_NO_APPLICATION_PROBE	Do not attempt to auto-detect MPI or CUDA executables.
ALLINEA_DETECT_APRUN_VERSION	Automatically detect Cray MPT by passing <code>-version</code> to the <code>aprun</code> wrapper, and parsing the output.

I/O behavior

Environment variables for handling default I/O behavior.

Variable	Description
ALLINEA_NEVER_FORWARD_STDIN	Do not forward the <code>stdin</code> of the <code>perf-report</code> command <code>stdin</code> to the program under analysis, even if you are running without the user interface. By default, Arm Performance Reports only forwards <code>stdin</code> when running without the user interface.
ALLINEA_ENABLE_ALL_REPORTS_GENERATION	Enables the option in Arm Performance Reports to generate all types of results at the same time, using the <code>all</code> extension.

Licensing

Environment variables to handle licensing.

Variable	Description
ALLINEA_LICENCE_FILE	Location of the license file. The default is <code>\\${ALLINEA_TOOLS_PATH}/Licence</code> .
ALLINEA_FORCE_LICENCE_FILE	Location of the license file. This ensures that Arm Forge uses a specified license file.
ALLINEA_LICENCE_DIR	Location of the licenses directory. The default is <code>\\${ALLINEA_TOOLS_PATH}/licences</code> .
ALLINEA_MAC_INTERFACE	Specify the host name of the network interface the license is tied to.

Timeouts

Environment variables for handling timeouts.

Variable	Description
ALLINEA_NO_TIMEOUT	Do not time out if nodes do not connect after a specified length of time. This might be necessary if the MPI subsystem takes unusually long to start processes.
ALLINEA_PROCESS_TIMEOUT	Length of time (ms) to wait for a process to connect to the front end.
ALLINEA_MPI_FINALIZE_TIMEOUT_MS	Length of time (ms) to wait for MPI_Finalize to end, and the program to exit. The default is 300000 (5 minutes). 0 waits forever.

Sampler

Environment variables for handling sampler-related setup, runtime behavior, and backend processing.

Variable	Description
ALLINEA_SAMPLER_INTERVAL	Arm Performance Reports takes a sample in each 20 millisecond period, giving it a default sampling rate of 50Hz. This is automatically decreased as the run proceeds to ensure that a constant number of samples are taken. See ALLINEA_SAMPLER_NUM_SAMPLES. If your program runs for a very short period of time, you might benefit by decreasing the initial sampling interval. For example, ALLINEA_SAMPLER_INTERVAL=1 sets an initial sampling rate of 1000Hz, or once per millisecond. Higher sampling rates are not supported. Arm recommends that you avoid increasing the sampling frequency from the default if there are many threads or very deep stacks in the target program. An increase in the sampling frequency might not allow enough time to complete one sample before the next sample starts. ————— Note ————— Custom values for ALLINEA_SAMPLER_INTERVAL might be overwritten by values set from the combination of ALLINEA_SAMPLER_INTERVAL_PER_THREAD and the expected number of threads (from OMP_NUM_THREADS). For more information, see ALLINEA_SAMPLER_INTERVAL_PER_THREAD.
ALLINEA_SAMPLER_INTERVAL_PER_THREAD	To keep overhead low, Arm Performance Reports imposes a minimum sampling interval based on the number of threads. By default, the interval is 2 milliseconds per thread. For eleven or more threads, this increases the initial sampling interval to more than 20 milliseconds. To adjust this behavior, set ALLINEA_SAMPLER_INTERVAL_PER_THREAD to the minimum sample time per thread, in milliseconds. Arm recommends that you avoid lowering this value from the default if there are many threads. Lowering this value might not allow enough time to complete one sample before the next sample starts. ————— Note ————— - Whether OpenMP is enabled or disabled in Arm Performance Reports, the final script or scheduler values set for OMP_NUM_THREADS is used for calculating the sampling interval per thread ALLINEA_SAMPLER_INTERVAL_PER_THREAD. When configuring your job for submission, check that there is a default value for OMP_NUM_THREADS in your final submission script, scheduler, or the Arm Performance Reports GUI. - Custom values for ALLINEA_SAMPLER_INTERVAL are overwritten by values set from the combination of ALLINEA_SAMPLER_INTERVAL_PER_THREAD and the expected number of threads from OMP_NUM_THREADS.

(continued)

Variable	Description
ALLINEA_MPI_WRAPPER	To direct Arm Performance Reports to use a specific wrapper library, set <code>ALLINEA_MPI_WRAPPER=<path of shared object></code>. Arm Performance Reports ships with a number of precompiled wrappers. When your MPI is supported, Arm Performance Reports automatically selects and uses the appropriate wrapper. To manually compile a wrapper specifically for your system, set <code>ALLINEA_WRAPPER_COMPILE=1</code> and <code>MPICC</code>, and run: <pre><Path to Forge installation>/map/wrapper/build_wrapper</pre> This generates the wrapper library <code>~/.allinea/wrapper/libmap-sampler-<code>mpirun</code>-<hostname>.so</code> with symlinks to the following files:
ALLINEA_WRAPPER_COMPILE	To direct to fall back to creating and compiling a just-in-time wrapper, set <code>ALLINEA_WRAPPER_COMPILE=1</code>. To generate a just-in-time wrapper, ensure that an appropriate compiler is available on the machine where Arm Performance Reports is running, or on the remote host when using remote connect. Arm Performance Reports attempts to auto detect your MPI compiler. However, Arm recommends that you set the environment variable to the path of the correct compiler.
ALLINEA_MPIRUN	The path of <code>mpirun</code>, <code>mpiexec</code>, or an equivalent. If set, <code>ALLINEAMPIRUN</code> has higher priority than the priority set in the GUI, and the <code>mpirun</code> found in <code>PATH</code>.
ALLINEA_SAMPLER_NUM_SAMPLES	Arm Performance Reports collects 1000 samples per process by default. To avoid generating too much data on long runs, the sampling rate is automatically decreased as the run progresses. This ensures that only 1000 evenly spaced samples are stored. You can adjust this by setting <code>ALLINEA_SAMPLER_NUM_SAMPLES=<positive integer></code> ————— Note ————— Arm strongly recommends that you leave this value at the default setting. Higher values are not generally beneficial, and add extra memory overheads while running your code. With 512 processes, the default setting already collects half a million samples over the job, and the effective sampling rate can be very high indeed. —————
ALLINEA_KEEP_OUTPUT_LINES	Specifies the number of lines of program output to record in <code>.map</code> files. Setting this to 0 removes the line limit restriction. Arm recommends that you avoid doing this because it can result in very large <code>.map</code> files if the profiled program produces a lot of output.
ALLINEA_KEEP_OUTPUT_LINE_LENGTH	The maximum line length for program output that is recorded in <code>.map</code> files. Lines containing more characters than this limit are truncated. Setting this to 0 removes the line length restriction. Arm recommends that you avoid doing this because it can result in very large <code>.map</code> files if the profiled program produces a lot of output per line.

(continued)

Variable	Description
ALLINEA_PRESERVE_WRAPPER	To gather data from MPI calls, Arm Performance Reports generates a wrapper to the chosen MPI implementation. By default, Arm Performance Reports deletes the generated code and shared objects when they are no longer required. To prevent Arm Performance Reports from deleting these files, set <code>ALLINEA_PRESERVE_WRAPPER=1</code>. ————— Note ————— If you are using remote launch, you must export this variable in the remote script. —————
ALLINEA_SAMPLER_NO_TIME_MPI_CALLS	To prevent Arm Performance Reports from capturing the time spent in MPI calls, set <code>ALLINEA_SAMPLER_NO_TIME_MPI_CALLS</code> .
ALLINEA_SAMPLER_TRY_USE_SMAPS	To allow Arm Performance Reports to use <code>/proc/[pid]/smaps</code> to gather memory usage data, set <code>ALLINEA_SAMPLER_TRY_USE_SMAPS</code> . Arm recommends that you avoid using this because it slows down sampling significantly.
MPICC	To create the MPI wrapper Arm Performance Reports tries <code>MPICC</code> . If that fails, search for a suitable MPI compiler command in <code>PATH</code> . If the MPI compiler used to compile the target binary is not listed, or there are multiple MPI compilers listed, set <code>MPICC</code> .

Simple troubleshooting

Environment variables for simple troubleshooting.

Variable	Description
ALLINEA_DEBUG_HEURISTICS	To print the weights and heuristics used to autodetect which MPI is loaded, set to <code>1</code> .

Chapter A4

Connecting to a remote system

This chapter describes the Arm Forge Remote Client that allows you to debug and profile remote jobs, while running the user interface on your local machine.

This is faster than remote-X11 (particularly for slow connections) and provides a native user interface.

It contains the following sections:

- *A4.1 Connecting remotely* on page A4-52.
- *A4.2 Remote connections dialog* on page A4-53.
- *A4.3 Reverse Connect* on page A4-56.
- *A4.4 Treeserver or general debugging ports* on page A4-58.
- *A4.5 Using X forwarding or VNC* on page A4-59.
- *A4.6 Starting Arm® Forge* on page A4-60.

A4.1 Connecting remotely

The remote client is available for Windows, Mac, and Linux, and can also be used as a local viewer for collected Arm MAP profiles.

The Arm Forge Remote Client connects and authenticates using SSH (typically a login node), and uses existing licensing from your remote resource (compute nodes).




Figure A4-1 Connect-remotely

Before you begin

- For working remotely using Arm DDT or Arm MAP, ensure that the versions match between the locally Arm Forge Remote client and the Arm Forge tools installed on remote systems.
- You do not require a license file for the locally installed remote client. Arm Forge uses the license of the remote system when it connects.

Procedure

1. Use SSH to login from your desktop machine **mydesktop** to the login node **mycluster-login**.
 2. Start a job using the queue submission command **qsub**.
 3. Connect Arm Forge using Reverse Connect, typically to a batch compute node.
- See [Reverse Connect on page A4-56](#) for more information on Reverse Connect.
4. To connect to a remote system, click the **Remote Launch** drop-down list and select **Configure**.

OPTIONS

Remote Launch:



QUIT

Figure A4-2 Remotelaunch--Configure

The [Remote connections dialog on page A4-53](#) opens and allows you to edit the required settings.

A4.2 Remote connections dialog

The **Remote Connections Dialog** allows you to add, remove and edit connections to remote systems.




Figure A4-3 Remote connections dialog

To display the Remote launch settings, add or edit a host (or any duplicated hosts) in the list.

Remote launch settings

The **Remote launch settings** displays when you edit or add a remote host to the **Remote connections dialog**.




Figure A4-4 Remote Launch options

- **Connection Name:** Optional for this connection. If no name is specified, the **Host Name** is used.
- **Host Name** is the host name of the remote system that you are connecting to.

The syntax of the host name field is:

```
[username]@hostname[:port]...
```

- **username** - Optional. The name that you use on the remote system. If this is not specified, your local user name is used instead.
- **hostname** - The host name of the remote system.
- **port** - Optional. The port number that the SSH daemon on the remote host listens on. If not specified, the default of 22 is used.

To log in using one or more intermediate hosts (such as a gateway), enter the host names in order, separated by spaces. For example, `gateway.arm.com cluster.lan`.

————— **Note** —————

You must be able to log in to the third and subsequent hosts without a password.

————— **Note** —————

You can specify more SSH options in the `remote-exec` script covered in [Connecting to compute nodes and remote programs \(remote-exec\)](#) on page Appx-A-448.

- Remote Installation Directory is the full path to the installation on the remote system.
- Remote Script. This optional script runs before starting the remote daemon on the remote system. You can use this script to load the required modules for Arm DDT and Arm MAP, and your MPI and compiler.

See the following sections for more details.

————— **Note** —————

The optional script is usually not necessary when using **Reverse Connect**.

- Always look for source files locally. Select this option to use the source files on the local system instead of the remote system.
- KeepAlive Packets. Select this option to enable KeepAlive packets. These are dummy packets that are sent on regular intervals to keep some SSH connections from timing out. The interval can be configured from the spin box below.
- Proxy through login node. Select this to use the local RSA key for connecting to both the login and the batch nodes. This is equivalent to not setting `ALLINEA_NO_SSH_PROXYCOMMAND`.

When this option is not set, Arm DDT connects to the login node using your local RSA key, and then uses the key on the remote SSH configuration folder to connect to the batch node. This is the same as setting `ALLINEA_NO_SSH_PROXYCOMMAND=1`.

Remote script

The script can load modules using the `module` command, or set environment variables. Arm Forge sources this script before running its remote daemon (your script does not need to start the remote daemon itself).

Run the script using `/bin/sh` (usually a Bourne-compatible shell). If this is not your usual login shell, make allowances for the different syntax it might require.

You can install a site-wide script that is sourced for all users at `/path/to//remote-init`.

You can also install a user-wide script that is sourced for all of your connections at `$ALLINEA_CONFIG_DIR/remote-init`.

————— **Note** —————

`ALLINEA_CONFIG_DIR` defaults to `$HOME/.allinea` if it is not set.

Example Script

Create this script file on the remote system, and ensure that the full path to the file is entered in the Remote Script field box.

```
module load allinea-forge
module load mympi
module load mycompiler
```

A4.3 Reverse Connect

The Arm Forge Reverse Connect feature allows you to submit your job from a shell terminal as you currently do, with a minor adjustment to your `mpirun` (or equivalent), to allow that job to connect back to the Arm Forge user interface.

Reverse Connect makes it easy to debug and profile jobs with the correct environment. You can easily load the required modules and prepare all the setup steps that are necessary before launching your job.

Note

Node-locked licenses such as workstation or Arm DDT Cluster licenses do not include the Reverse Connect feature.

Procedure

1. Start Arm Forge and connect to your remote system (typically a login node) with SSH.
2. To enable Reverse Connect, modify your current (or equivalent) command line inside your interactive queue allocation, or queue submission script.

In most of the cases it is sufficient to prefix the script with `ddt/map --connect`. Almost all arguments beside `-offline` and `-profile` are supported by Reverse Connect.

Example

```
$ mpirun -n 512 ./examples/wave_f
```

1. To debug the job using Reverse Connect and [Express Launch \(DDT\)](#) on page B1-65, run:

```
$ ddt --connect mpirun -n 512 ./examples/wave_f
```

2. To profile the job using Reverse Connect and [Express Launch \(MAP\)](#) on page C1-242, run:

```
$ map --connect mpirun -n 512 ./examples/wave_f
```

Note

If your MPI is not yet supported by Express Launch mode you can use **Compatibility Mode**.

Debug:

```
$ ddt --connect -n 512 ./examples/wave_f
```

Profile:

```
$ map --connect -n 512 ./examples/wave_f
```

3. After a short period of time, the GUI shows the Reverse Connect request including the host from where the request was made (typically a batch compute node), and a command-line summary.




Figure A4-5 Reverse connect request

4. To accept the request, click **Accept**.

————— **Note** —————

Arm Forge connects to the specified host and executes what you specified with the command line. If you do not want to accept the request, click **Reject**.

Connection details

If a Reverse Connect is initiated, for example with `ddt --connect`, Arm Forge starts a server listening on a port in the range between 4201 and 4240.

If this port range is not suitable for some reason, such as when ports are already taken by other services, you can override the port range with the environment variable `ALLINEA_REMOTED_PORTS`.

```
$ export ALLINEA_REMOTED_PORTS=4400-4500
$ ddt --connect
```

The server selects a free port between 4400 and 4500 (inclusive).

This connection is between the batch or submit node (where `ddt -connect` is run from) and the login node. This connection can also be to a compute node if for example, you are running `ddt --connect mpirun` on a single node.

A4.4 Treeserver or general debugging ports

Connections are made in the following ways, depending on the use case.

Using a queue submission or X-forwarding

- A connection is made between the login node and the batch or submit node using ports 4242-4262.
- Connections are made between the batch or submit node and the compute nodes using ports 4242-4262.
- Connections are made from compute nodes to other compute nodes using ports 4242-4262.

Using reverse connect

- See **Connection details** in [Reverse Connect on page A4-56](#) for details about login node to batch/submit node ports.
- Connections are made between the batch or submit node and the compute nodes using ports 4242-4262.
- Connections are made from compute nodes to other compute nodes using ports 4242-4262.

A4.5 Using X forwarding or VNC

If you do not want to use the **Remote Launch** feature, try the following two methods for running Arm DDT or Arm MAP on a remote system:

- X forwarding

X forwarding is effective when the network connection is low latency, such as when the network spans a single physical site.

- VNC

Arm recommends using VNC (or similar Unix-supporting remote desktop software) when the network connection is moderate or slow.

Note

For users on Mac OS X with a single-button mouse, you can access a Linux or other Unix machine by pressing the **Command** key and mouse. This displays some important Arm DDT and Arm MAP features in the right-click menu.

Procedure

1. To use X forwarding from a Mac OS X system to access the instance running Arm Forge on a remote Linux/Unix system, start the X11 server (available in the `X11User.pkg`).
2. Set the display variable correctly to allow X applications to display by opening a terminal in Mac OS X and typing:

```
export DISPLAY=:0
```

3. Use ssh to connect to the remote system from the terminal, with ssh options `-X` and `-C` (X forwarding and compression).

For example:

```
ssh -CX username@login.mybigcluster.com
```

4. Start Arm DDT or Arm MAP on the remote system to display the window on your machine.

Note

A known issue with the free version of `XMing` prevents it working well on Arm Forge and on other Qt5 projects on Windows. Arm recommends using a more up to date X server such as `VcXsrv`.

A4.6 Starting Arm® Forge

To start Arm Forge, type one of the following commands into a terminal window:

```
forge  
forge program_name [arguments]
```

To start Arm Forge on Mac OS X, click the Arm Forge icon or type this command in the terminal window:

```
open /Applications/Arm\ Forge/Arm Forge.app [--args program_name [arguments]]
```

To launch additional instances of the Arm Forge application, right-click the Dock icon of a running instance of Arm Forge, and select **Launch a new instance of Arm Forge**.

Alternatively, you can use the following command in a terminal:

```
open -n /Applications/Arm\ Forge/Arm Forge.app [--args program_name [arguments]]
```

When Arm Forge starts, the **Welcome Page** displays, in which you can select the tool you would like to use (Arm DDT or Arm MAP). Arm Performance Reports can be run using the command line. Use the icons on the left-hand side to switch between Arm MAP and Arm DDT.

Select the tool you want to use, and click the buttons in the menu to select a debugging or profiling activity.

Note

To pipe input directly to the program, you must operate Arm Forge in Express Launch mode. For information about how to send input to your program, see [Program input and output on page B5-179](#) or [Run MAP from the command line on page C14-341](#).

Note

In Express Launch mode (see Arm DDT or Arm MAP), the Welcome Page does not display. Instead, Arm Forge directly displays the **Run** dialog. If no valid license is found, the program exits and the appropriate message displays in the console output.

Part B
DDT

Chapter B1

Get started with DDT

Learn how to get started using Arm DDT.

It contains the following sections:

- *B1.1 Prepare a program for debugging* on page B1-64.
- *B1.2 Express Launch (DDT)* on page B1-65.
- *B1.3 Run a program* on page B1-67.
- *B1.4 remote-exec required by some MPIs (DDT)* on page B1-71.
- *B1.5 Open core files* on page B1-72.
- *B1.6 Attach to running programs* on page B1-73.
- *B1.7 Debug single-process programs* on page B1-78.
- *B1.8 Debug OpenMP programs* on page B1-80.
- *B1.9 Debug MPMD programs* on page B1-81.
- *B1.10 Manual launch of multi-process non-MPI programs* on page B1-82.
- *B1.11 Start a job in a queue* on page B1-84.
- *B1.12 Job scheduling with jsrun* on page B1-85.
- *B1.13 Use custom MPI scripts* on page B1-86.
- *B1.14 Start DDT from a job script* on page B1-89.
- *B1.15 Numactl (DDT)* on page B1-90.
- *B1.16 Python debugging* on page B1-91.
- *B1.17 DDT user interface* on page B1-93.
- *B1.18 Save and load sessions* on page B1-95.

B1.1 Prepare a program for debugging

When compiling the program that you want to debug, you must add the debug flag to your compile command. For most compilers this is `-g`.

We recommend that you turn off compiler optimizations as they can produce unexpected results when debugging. If your program is already compiled without debug information you will need to make the files that you are interested in again.

The Welcome page enables you to choose the kind of debugging you want to do, for example you can:

- Run a program from DDT and debug it.
- Debug a program you launch manually (for example, on the command line).
- Attach to an already running program.
- Open core files generated by a program that crashed.
- Connect to a remote system and accept a Reverse Connect request.

B1.2 Express Launch (DDT)

All of the Arm Forge products can be launched by typing their name in front of an existing `mpiexec` command:

```
$ ddt mpiexec -n 128 examples/hello memcrash
```

This startup method is called *Express Launch* and is the simplest way to get started.

The MPI implementations supported by Express Launch are:

- bullx MPI
- Cray X-Series (MPI/SHMEM/CAF)
- Intel MPI
- MPICH 3
- Open MPI (MPI/SHMEM)
- Oracle MPT
- Open MPI (Cray XT/XE/XK)
- Spectrum MPI
- Spectrum MPI (PMIx)
- Cray XT/XE/XK (UPC)

If your MPI is not supported by Express Launch, an error message will display:

```
$ 'Generic' MPI programs cannot be started using Express Launch syntax (launching with an  
mpirun command).  
  
Try this instead:  
  ddt --np=256 ./wave_c 20  
  
Type ddt --help for more information.
```

This is referred to as *Compatibility Mode*, in which the `mpiexec` command is not included and the arguments to `mpiexec` are passed via a `-mpiargs="args here"` parameter.

One advantage of Express Launch mode is that it is easy to modify existing queue submission scripts to run your program under one of the Arm Forge products. This works best for Arm DDT with Reverse Connect, `ddt -connect`, for interactive debugging or in offline mode (`ddt -offline`).

See [Reverse Connect on page A4-56](#) for more details.

If you cannot use Reverse Connect and want to use interactive debugging from a queue, you might need to configure DDT to generate job submission scripts for you. More details on this can be found in [Start a job in a queue on page B1-84](#) and in [Integration with queuing systems on page Appx-A-445](#).

Run dialog box

In Express Launch mode, the **Run** dialog has a restricted number of options:




Figure B1-1 Express Launch Run dialog box

B1.3 Run a program

When you click **Run** on the Welcome page, the **Run** window displays.

Run

Application: /home/user/ddt/examples/hello sleepy Details

Application: /home/user/ddt/examples/hello Folder icon

Arguments: sleepy Dropdown arrow

☐ **stdin file:** Folder icon

Working Directory: Folder icon

☒ **MPI:** 512 processes, Open MPI Details

☐ **OpenMP** Details

☐ **CUDA** Details

☐ **Memory Debugging** Details...

☐ **Submit to Queue** Configure... Parameters...

Environment Variables: none Details

Plugins: none Details

Help Options Run Cancel

Figure B1-2 Run window

The settings are grouped into sections. Click **Details** to expand a section.

Application

Application: The full file path to your application. If you specified one on the command line, this is automatically filled in. You can browse and select your application.

Note

Many MPIs have problems working with directory and program names that contain spaces. We recommend that you do not use spaces in directory and file names.

Arguments: (optional) The arguments passed to your application. These are automatically filled if you entered some on the command line.

Note

Avoid using quote characters such as ' and ", as these may be interpreted differently by DDT and your command shell. If you must use these characters but cannot get them to work as expected, contact [Arm support](#).

stdin file: (optional) This enables you to choose a file to be used as the standard input (stdin) for your program. Arguments are automatically added to mpirun to ensure your input file is used.

Working Directory: (optional) The working directory to use when debugging your program. If this is blank then DDT's working directory is used instead.

MPI

Note

If you only have a single process license or have selected **none** as your **MPI Implementation**, the **MPI** options will be missing. The **MPI** options are not available when DDT is in single process mode. See [Debug single-process programs on page B1-78](#) for more details about using DDT with a single process.

Number of processes: The number of processes that you want to debug. DDT supports hundreds of thousands of processes but this is limited by your license.

Number of nodes: This is the number of compute nodes that you want to use to run your program.

Processes per node: This is the number of MPI processes to run on each compute node.

Implementation: The MPI implementation to use. If you are submitting a job to a queue, the queue settings will also be summarized here. Click **Change** to change the MPI implementation.

Notes:

- The choice of MPI implementation is critical to correctly starting DDT. Your system will normally use one particular MPI implementation. If you are unsure which to choose, try generic, consult your system administrator or Arm support. A list of settings for common implementations is provided in [MPI distribution notes and known issues on page Appx-E-461](#).
- If the MPI command you want is not in your PATH, or you want to use an MPI run command that is not your default one, you can configure this using the **Options** window (See [System on Optional configuration on page Appx-A-449](#)).

mpirun arguments: (optional): The arguments that are passed to mpirun or your equivalent, usually prior to your executable name in normal usage. You can place machine file arguments here, if necessary. For most users this box can be left empty. You can also specify arguments on the command line (using the -mpiargs command-line argument) or using the ALLINEA_MPIRUN_ARGUMENTS environment variable if this is more convenient.

Notes:

- You should **not** enter the -np argument as DDT will do this for you.
- You should **not** enter the -task-nb or -process-nb arguments as DDT will do this for you.

OpenMP

Number of OpenMP threads: The number of OpenMP threads to run your application with. The OMP_NUM_THREADS environment variable is set to this value.

For more information on debugging OpenMP programs see [Debug OpenMP programs on page B1-80](#).

CUDA

If your license supports it, you can also debug GPU programs by enabling CUDA support. For more information on debugging CUDA programs see [CUDA GPU debugging on page B10-215](#).

Track GPU Allocations: Tracks CUDA memory allocations made using `cudaMalloc`, and similar methods. See [CUDA memory debugging on page B8-195](#) for more information.

Detect invalid accesses (memcheck): Turns on the CUDA-MEMCHECK error detection tool. See [CUDA memory debugging on page B8-195](#) for more information.

Memory debugging

Click **Details** to open the **Memory Debugging Options** window.

See [Memory debugging options on page B8-197](#) for details of the available settings.

Environment variables

The optional **Environment Variables** section should contain additional environment variables that should be passed to `mpirun` or its equivalent. These environment variables can also be passed to your program, depending on which MPI implementation your system uses. Most users will not need to use this section.

Note

On some systems it may be necessary to set environment variables for the DDT backend itself. For example, if `/tmp` is unusable on the compute nodes you may want to set `TMPDIR` to a different directory. You can specify such environment variables in `/path/to/ddt/lib/environment`. Enter one variable per line and separate the variable name and value with `=`. For example, `TMPDIR=/work/user`.

Plugins

The optional **Plugins** section lets you enable plugins for various third-party libraries, such as the Intel Message Checker or Marmot. See [Use and write plugins on page B9-207](#) for more information.

Run the program

Click **Run** to start your program, or **Submit** if working through a queue (see [Integration with queuing systems on page Appx-A-445](#)). This runs your program through the debug interface you selected and allows your MPI implementation to determine which nodes to start which processes on.

Note

If you have a program compiled with Intel `ifort` or GNU `g77` you may not see your code and highlight line when DDT starts. This is because those compilers create a pseudo `MAIN` function, above the top level of your code. To fix this you can either open your Source Code window, add a breakpoint in your code, then run to that breakpoint, or you can use the Step into function to step into your code.

When your program starts, DDT attempts to determine the MPI world rank of each process. If this fails, an error message displays:




Figure B1-3 MPI rank error

This means that the number DDT shows for each process may not be the MPI rank of the process. To correct this you can tell DDT to use a variable from your program as the rank for each process.

See [Assign MPI ranks on page B4-174](#) for details.

End the session

To end your current debugging session, select **File > End Session**. This closes all processes and stops any running code. If any processes remain you might have to clean them up manually using the `kill` command, or a command provided with your MPI implementation.

B1.4 remote-exec required by some MPIs (DDT)

When using *SGI MPT*, the MPMD variants of *MPICH 3*, or *Intel MPI*, you can use `mpirun` to start all the processes, then attach to them while they are inside `MPI_Init`.

This method is often faster than the generic method, but requires the `remote-exec` facility to be correctly configured if processes are being launched on a remote machine. For more information on `remote-exec`, see [Connecting to compute nodes and remote programs \(remote-exec\)](#) on page Appx-A-448 and *Choose hosts* on [Attach to running programs](#) on page B1-73.

Note

If DDT is running in the background (for example, `ddt &`) this process may get stuck. Some SSH versions cause this behavior when asking for a password. If this happens to you, go to the terminal and use the `fg` or similar command to make DDT a foreground process, or run DDT again, without using `&`.

If DDT cannot find a password-free way to access the cluster nodes then you will not be able to use the specialized startup options. Instead, you can use *generic*, although startup may be slower for large numbers of processes.

In addition to the listed MPI implementations above, all MPI implementations except for *Cray MPT* DDT require password-free access to the compute nodes when explicitly starting by attaching.

B1.5 Open core files

You can open and debug one or more core files generated by your application.

Procedure

1. On the Welcome page click **Open Core Files**. The **Open Core Files** window opens.




Figure B1-4 The Open Core Files window

2. Select an executable and a set of core files, then click **OK** to open the core files and start debugging them.

Note

While DDT is in this mode, you cannot play, pause, or step, because there is no process active. You are, however, able to evaluate expressions and browse the variables and stack frames saved in the core files.

Next steps

To end your current debugging session, select **File > End Session**.

B1.6 Attach to running programs

You can attach to running processes on any machine you can access, whether they are from MPI or scalar jobs, even if they have different executables and source pathnames.

Click **Attach to a Running Program** on the Welcome page to open the **Attach** window.




Figure B1-5 Attach window

There are two ways to select the processes you want to attach to. You can either choose from a list of automatically detected MPI jobs (for supported MPI implementations), or manually select from a list of processes.

Automatically-detected jobs

DDT can automatically detect MPI jobs started on the local host for selected MPI implementations. This also applies to other hosts you can access if an *Attach Hosts File* is configured. See *System on Optional configuration on page Appx-A-449* for more details.

The list of detected MPI jobs is shown on the **Automatically-detected jobs** tab of the **Attach** window. Click the header for a particular job to see more information about that job. When you have found the job you want to attach to, simply click the **Attach** button to attach to it.

Note

Non-MPI programs that were started using MPI may not appear in this window. For example `mpirun -np 2 sleep 1000`

If you only want to attach to a subset of ranks from your MPI job, you can choose this subset using **Attach to ranks** on the **Automatically-detected jobs** tab. You can change the subset later by selecting **File > Change Attached Processes**. The menu item is only available for jobs that were attached to, and not for jobs that were launched using DDT.

List of all processes

You can manually select which processes to attach to from the list of processes on the **List of all processes** tab of the **Attach** window.

If you want to attach to a process on a remote host, see [Connecting to compute nodes and remote programs \(remote-exec\)](#) on page Appx-A-448 first.

Initially the list of processes is blank while DDT scans the nodes, provided in your node list file, for running processes. When all the nodes have been scanned (or have timed out) the window appears as shown above. Use **Filter for process names containing** to find the processes you want to attach to. On non-Linux platforms you also need to select the application executable you want to attach to. Ensure that the list shows all the processes you wish to debug in your job, and no extra/unnecessary processes. You can modify the list by selecting and removing unwanted processes, or alternatively selecting the processes you want to attach to and clicking **Attach to selected processes**. If no processes are selected, DDT attaches to all processes in the list.

On Linux you can use DDT to attach to multiple processes running different executables. When you select processes with different executables the application box changes to read **Multiple applications selected**. DDT creates a process group for each distinct executable.

With some supported MPI implementations (for example, Open MPI) DDT shows MPI processes as children of the (or equivalent) command.

Process name	Host	PID	PPID	Forked	Executable
▼ mpirun	login1	10407	10240	No	/software/mpi/openmpi-3.
hello_c	login1	10412	0	No	/home/user/arm/forg
hello_c	login1	10413	0	No	/home/user/arm/forg
hello_c	login1	10414	0	No	/home/user/arm/forg
hello_c	login1	10415	0	No	/home/user/arm/forg
hello_c	login1	10416	0	No	/home/user/arm/forg
hello_c	login1	10419	0	No	/home/user/arm/forg
hello_c	login1	10423	0	No	/home/user/arm/forg
hello_c	login1	10425	0	No	/home/user/arm/forg
hello_c	login1	10429	0	No	/home/user/arm/forg
hello_c	login1	10431	0	No	/home/user/arm/forg
hello_c	login1	10433	0	No	/home/user/arm/forg

Figure B1-6 Attaching with Open MPI

If you click the command, this automatically selects all the MPI child processes.

When you click **Attach to selected/listed processes**, DDT uses `remote-exec` to attach a debugger to each process you selected and proceeds to debug your application as if you had started it with DDT. When you end the debug session, DDT detaches from the processes rather than terminating them. This means you can attach to them again later if you want.

DDT examines the processes it attaches to and tries to discover the MPI_COMM_WORLD rank of each process. If you have attached to two MPI programs, or a non-MPI program, you might see this message:




Figure B1-7 MPI rank error

If there is no rank, for example, if you have attached to a non-MPI program, you can ignore this message and use DDT as normal. If there is a rank, you can easily tell DDT the correct rank for each process via **Use as MPI Rank** in the **Cross-Process Comparison View**. See [Assign MPI ranks on page B4-174](#) for details.

Note that `stdin`, `stderr`, and `stdout` (standard input, standard error and standard output) are not captured by DDT if used in attaching mode. Any input/output continues to work as it did before DDT attached to the program, for example, from the terminal or perhaps from a file.

Choose hosts

To attach to remote hosts, click **Choose Hosts** in the **Attach** window. The **Hosts** window displays the list of hosts that you can use for attaching.




Figure B1-8 Hosts Window

In the **Hosts** windows you can add and remove hosts, and uncheck hosts that you want to temporarily exclude.

To import a list of hosts from a file, click **Import**.

The hosts list populates using the *Attach Hosts File*. To configure the hosts, click **File > Options** (Arm Forge > Preferences on Mac OS X) to open the **Options** window.

Each remote host is scanned for processes, and the result is displayed in the **Attach** window. If you have trouble connecting to remote hosts, see [Connecting to compute nodes and remote programs \(remote-exec\)](#) on page Appx-A-448.

Use command-line arguments

As an alternative to starting DDT and using the Welcome page, DDT can instead be instructed to attach to running processes from the command-line.

To do so, you need to specify a list of hostnames and process identifiers (PIDs). If a hostname is omitted then localhost is assumed.

The list of hostnames and PIDs can be given on the command-line using the `-attach` option:

```
user@holly:~$ ddt --attach=11057,node5:11352
```

Another command-line possibility is to specify the list of hostnames and PIDs in a file and use the `-attach-file` option:

```
user@holly:~$ cat /home/user/ddt/examples/hello.list
node1:11057
node1:11094
node2:11352
node2:11362
node3:12357

user@holly:~$ ddt --attach-file=/home/user/ddt/examples/hello.list
```

In both cases, if just a number is specified for a hostname:PID pair, then `localhost:` is assumed.

These command-line options work for both single- and multi-process attaching.

B1.7 Debug single-process programs

You can use the **Run** dialog to start debugging single-process programs.




Figure B1-9 Single-process Run dialog

Before you begin

- If you have a single-process license you will immediately see the **Run** dialog that is appropriate to run a single-process program.
- If you have a multi-process license you can clear the **MPI** checkbox to run a single-process program.

Procedure

1. Type the full file path to your application, or browse and select your application.
2. If required, type the arguments to pass to your program.
3. Click **Run** to start your program.

Note

If your program has been compiled with Intel ifort you may not see your code and highlight line when DDT starts. This is because this compiler creates a pseudo MAIN function, above the top level of your

code. To fix this you can either open your Source Code window, add a breakpoint in your code, then play to that breakpoint, or you can use the Step Into function to step into your code.

Next steps

To end your current debugging session select **File > End Session**. This will close all processes and stop any running code.

B1.8 Debug OpenMP programs

When you run an OpenMP program, set the **Number of OpenMP threads** value to the number of threads you require. DDT will run your program with the `OMP_NUM_THREADS` environment variable set to the appropriate value.

There are several important points to keep in mind when you debug OpenMP programs:

- Parallel regions created with `#pragma omp parallel (C)` or `!$OMP PARALLEL (Fortran)` will usually not be nested in the **Parallel Stack View** under the function that contained the `#pragma`. Instead they will appear under a different top-level item. The top-level item is often in the OpenMP runtime code, and the parallel region appears several levels down in the tree.
- Some OpenMP libraries only create the threads when the first parallel region is reached. It is possible you may only see one thread at the start of the program.
- You cannot step into a parallel region. Instead, select **Step threads together** and use the **Run to here** command to synchronize the threads at a point inside the region. These controls are discussed in more detail in their own sections of this document.
- You cannot step out of a parallel region. Instead, use the **Run to here** command to leave it. Most OpenMP libraries work best if you keep **Step threads together** selected until you have left the parallel region. With the Intel OpenMP library, this means you will see the **Stepping Threads** window and will have to click **Skip All** once.
- Leave **Step threads together** clear when you are outside a parallel region, as OpenMP worker threads usually do not follow the same program flow as the main thread.
- To control threads individually, use **Focus on Thread**. This allows you to step and play one thread without affecting the rest. This is helpful when you want to work through a locking situation or to bring a stray thread back to a common point. The Focus controls are discussed in more detail in their own section of this document.
- Shared OpenMP variables may appear twice in the **Locals** window. This is one of the many unfortunate side-effects of the complex way OpenMP libraries interfere with your code to produce parallelism. One copy of the variable may have a nonsense value, this is usually easy to recognize. The correct values are shown in the **Evaluate** and **Current Line** windows.
- Parallel regions may be displayed as a new function in the **Stacks** view. Many OpenMP libraries implement parallel regions as automatically-generated `outline` functions, and DDT shows you this. To view the value of variables that are not used in the parallel region, you may need to switch to thread 0 and change the stack frame to the function you wrote, rather than the `outline` function.
- Stepping often behaves unexpectedly inside parallel regions. Reduction variables usually require some sort of locking between threads, and may even appear to make the current line jump back to the start of the parallel region. If this happens step over several times and you will see the current line comes back to the correct location.
- Some compilers optimize parallel loops regardless of the options you specified on the command line. This has many strange effects, including code that appears to move backwards as well as forwards, and variables that are not displayed or have nonsense values because they have been optimized out by the compiler.
- The thread IDs displayed in the Process Group Viewer and Cross-Thread Comparison window will match the value returned by `omp_get_thread_num()` for each thread, but only if your OpenMP implementation exposes this data to DDT. GCC's support for OpenMP (GOMP) needs to be built with TLS enabled with our thread IDs to match the return `omp_get_thread_num()`, whereas your system GCC most likely has this option disabled. The same thread IDs will be displayed as tooltips for the threads in the thread viewer, but only your OpenMP implementation exposes this data.

If you are using DDT with OpenMP and would like to tell us about your experiences, please contact [Arm support](#), with the subject title *OpenMP feedback*.

B1.9 Debug MPMD programs

The easiest way to debug multiple program, multiple data (MPMD) programs is by using Express Launch to start your application. To use Express Launch, prefix your normal MPMD launch line with DDT, for example:

```
ddt mpirun -n 1 ./master : -n 2 ./worker
```

For more information on Express Launch, and compatible MPI implementations, see [Express Launch \(DDT\)](#) on page B1-65.

About this task

If you are using Open MPI, MPICH 3, or Intel MPI, DDT can be used to debug MPMD programs without Express Launch. This procedure shows how to start an MPMD program in DDT.

Procedure

1. (MPICH 3 and Intel MPI only) Select the MPMD variant of the MPI Implementation on the **System** page of the **Options** window. For example, for MPICH 3 select **MPICH 3 (MPMD)**.
2. On the Welcome page click **Run**.
3. In **Application**, choose one of the MPMD programs. It does not matter which executable you choose.
4. Enter the total amount of processes for the MPMD job in **Number of Processes**.
5. In the MPI section of the **Run** window, enter an MPMD style command line in **mpirun Arguments**. Make sure that the sum of processes in this command is equal to the number of processes set in **Number of Processes**. For example:

```
-np 4 hello : -np 4 program2
```

or:

```
--app /path/to/my_app_file
```

6. Click **Run**.

Example

If you are using Open MPI in Compatibility mode, for example, because you do not have SSH access to the compute nodes, then replace:

```
-np 2 ./progc.exe : -np 4 ./progf90.exe
```

in **mpirun Arguments** / appfile with:

```
-np 2 /path/to/ddt/bin/forg-client ./progc.exe : -np 4  
/path/to/ddt/bin/forg-client ./progf90.exe
```

B1.10 Manual launch of multi-process non-MPI programs

DDT can only launch MPI programs and scalar (single process) programs. You need to use the **Manual Launch (Advanced)** button on the **Welcome Page** to debug multi-process and multi-executable programs. These programs do not need to be MPI programs. You can debug programs that use other parallel frameworks, or both the client and the server from a client/server application in the same DDT session.

About this task

You must run each program you want to debug manually using the `forge-client` command, similar to debugging with a scalar debugger like the GNU debugger (gdb). However, unlike a scalar debugger, you can debug more than one process at the same time in the same session, as long as your license permits it. Each program you run will show up as a new process in the DDT window. In this example we show you how to debug both client and server in the same session.

Procedure

1. On the Welcome page click **Manual Launch (Advanced)**.
2. Set **Number of Processes** to 2.




Figure B1-10 Manual Launch window

3. Click **Listen**.
4. At the command line run:

```
forge-client server &
forge-client client &
```

The server process displays as process 0, and the client as process 1 in the DDT window.

All	0	1
client	1	
server	0	

Figure B1-11 Manual Launch Process Groups

Next steps

After you have run the initial programs you can add extra processes to the DDT session. For example, extra clients can be added using `forge-client` in the same way:

```
forge-client client2 &
```

If you select **Start debugging after the first process connects** you do not need to specify how many processes you want to launch in advance. You can start debugging after the first process connects and add extra processes later as above.

B1.11 Start a job in a queue

In most cases you can debug a job by putting `ddt -connect` in front of the existing `mpiexec` or equivalent command in your job script. If a GUI is running on the login node or it is connected to it via the remote client, a message is displayed prompting you with the option to debug the job when it starts.

See [Express Launch \(DDT\) on page B1-65](#) and [Reverse Connect on page A4-56](#) for more details.

If DDT has been configured to be integrated with a queue/batch environment, as described in [Integration with queuing systems on page Appx-A-445](#) then you can use DDT to submit your job directly from the user interface. In this case, a **Submit** button is displayed on the **Run** window, instead of a **Run** button. When you click **Submit** on the **Run** window the queue status is displayed until your job starts. DDT will execute the display command every second and show you the standard output. If your queue display is graphical or interactive you cannot use it here.

If your job does not start or you decide not to run it, click **Cancel Job**. If the regular expression you entered for getting the job id is invalid, or if an error is reported, DDT will not be able to remove your job from the queue. In this case we strongly recommend that you check the job has been removed before submitting another as it is possible for a forgotten job to execute on the cluster and either waste resources or interfere with other debug sessions.

When your job is running, it connects to DDT and you can debug it.

B1.12 Job scheduling with jsrun

Launching jobs with jsrun in a job scheduling system enables the topology of processes and threads on the node to be split into individual resource sets (the number of GPUs, CPUs, threads, and MPI tasks). You can specify the amount of computational resource allocated to a resource set.

How you decide to allocate resources has an impact on the runtime of Arm DDT and Arm MAP. For example, it is possible to allocate all of the CPUs on the node to just one resource set. Alternatively, you could allocate each CPU to its own resource set. In this case there are as many resource sets as there are CPUs on the node.

The more resource sets you have on each node, the longer the runtime is for Arm DDT and Arm MAP. To minimize runtime, we recommend that you aim to reduce the number of resource sets required.

For example, we recommend:

```
jsrun --rs_per_host=1 --gpu_per_rs=0 --cpu_per_rs=42 --tasks_per_rs=42 ...
```

to launch a job with 42 MPI processes per node in a single resource set, instead of:

```
jsrun --rs_per_host=42 --gpu_per_rs=0 --cpu_per_rs=1 --tasks_per_rs=1 ...
```

which launches 42 MPI processes per node, but uses 42 resource sets.

B1.13 Use custom MPI scripts

On some systems a custom mpirun replacement is used to start jobs, such as mpiexec. DDT normally uses whatever the default for your MPI implementation is, so for Open MPI it would look for mpirun and not mpiexec. Here we explain how to configure DDT to use a custom mpirun command for job start up.

There are typically two ways you might want to start jobs using a custom script, and DDT supports them both.

In the first way, you pass all the arguments on the command-line, like this:

```
mpiexec -n 4 /home/mark/program/chains.exe /tmp/mydata
```

There are several key variables in this command that DDT can fill in for you:

- The number of processes (4 in the above example).
- The name of your program (/home/mark/program/chains.exe).
- One or more arguments passed to your program (/tmp/mydata).

Everything else, like the name of the command and the format of its arguments remains constant. To use a command like this in DDT, you adapt the queue submission system described in [Job scheduling with jsrun on page B1-85](#). For this mpiexec example, the settings are:




Figure B1-12 Using custom MPI scripts

As you can see, most of the settings are left blank. There are some differences between the **Submit** command in DDT and what you would type at the command-line:

- The number of processes is replaced with `NUM_PROCS_TAG`.
- The name of the program is replaced by the full path to `forge-backend`.
- The program arguments are replaced by `PROGRAM_ARGUMENTS_TAG`.

Note

It is not necessary to specify the program name here. DDT takes care of that during its own startup process. The important thing is to make sure your MPI implementation starts `forge-backend` instead of your program, but with the same options.

In the second way, you start a job using a custom `mpirun` replacement with a settings file:

```
mpiexec -config /home/mark/myapp.nodespec
```

Where `myfile.nodespec` might contain something similar to the following:

```
comp00 comp01 comp02 comp03 : /home/mark/program/chains.exe /tmp/mydata
```

DDT can automatically generate simple configuration files like this every time you run your program, you need to specify a template file. For the above example, the template file `myfile.ddt` would contain:

```
comp00 comp01 comp02 comp03 : DDTPATH_TAG/libexec/forge-backend DDT_DEBUGGER_ARGUMENTS_TAG  
PROGRAM_ARGUMENTS_TAG
```

This follows the same replacement rules described above and in detail in [Integration with queuing systems on page Appx-A-445](#).

The settings in the **Options** window for this example might be:




Figure B1-13 Using substitute MPI commands

Note **Submit command** and **Submission template file** in particular. Arm DDT creates a new file and appends it to the submit command before executing it. In this case, `mpiexec -config /tmp/ddt-temp-0112` or similar is executed. Therefore, any argument like `-config` must be last on the line, because Arm DDT adds a file name to the end of the line. If there are any other arguments, they can come first.

If your system uses a non-standard startup command, Arm recommends that you read [Job scheduling with jsrun on page B1-85](#), which describes many features that might be useful to you.

If you do use a non-standard command, contact [Arm support](#).

B1.14 Start DDT from a job script

The usual way of debugging a program with Arm DDT in a queue/batch environment is with Reverse Connect where it connects back from inside the queue to the user interface (see [Reverse Connect](#)). You can also debug a program in a queue/batch environment by configuring Arm DDT to submit the program to the queue for you (see [Start a job in a queue](#)). This procedure describes another way - how to start Arm DDT from a job script that is submitted to the queue/batch environment.

Procedure

1. Configure Arm DDT with the correct MPI implementation.
2. Disable queue submission in the Arm DDT options.
3. Create a job script that starts Arm DDT using a command such as:

```
ddt --start MPIEXEC -n NPROCS PROGRAM [ARGUMENTS]
```

Or the following:

```
ddt --start --no-queue --once --np=NPROCS -- PROGRAM [ARGUMENTS]
```

In these examples MPIEXEC is the MPI launch command, NPROCS is the number of processes to start, PROGRAM is the program to run, and ARGUMENTS are the arguments to the program. The `-once` argument tells DDT to exit when the session ends.

4. Submit the job script to the queue.

B1.15 Numactl (DDT)

DDT supports launching programs via numactl for MPI programs, but has limited support for non-MPI programs.

MPI and SLURM

DDT can attach to MPI programs launched via numactl with or without SLURM. The recommended way to launch via numactl is to use [Express Launch \(DDT\)](#) on page B1-65.

```
$ ddt mpiexec -n 4 numactl -m 1 ./myMpiProgram.exe  
$ ddt srun -n 4 numactl -m 1 ./myMpiProgram.exe
```

It is also possible to launch via numactl using compatibility mode. When using compatibility mode, you must specify the full path to numactl in **Application**. You can find the full path by running:

```
which numactl
```

Enter the name of the required application in **Arguments**, after all arguments to be passed to numactl. It is not possible to pass any more arguments to the parallel job runner when using this mode for launching.

Note

When using memory debugging with a program launched via numactl, the Memory Statistics view will report all memory as ‘Default’ memory type unless allocated with memkind. See [Memory Statistics](#) on page B8-206.

Non-MPI Programs

There is a minor caveat to launching non-MPI programs via numactl. If you are using SLURM, set ALLINEA_STOP_AT_MAIN=1, otherwise DDT will not be able to attach to the program. For example, these two commands are examples of launching non-MPI programs via numactl:

```
$ ddt numactl -m 1 ./myNonMpiProgram.exe  
$ ALLINEA_STOP_AT_MAIN=1 ddt srun \  
  numactl -m 1 ./myNonMpiProgram.exe
```

When launched, the program stops in numactl main. To resume debugging as normal, set a breakpoint in your code (optional), then use the play and pause buttons to progress and pause the debugging respectively.

B1.16 Python debugging

Python debugging is supported with the following features:

- Debugs Python scripts running under the CPython interpreter (versions 3.5-3.8 only).
- Decodes the stack to show Python frames, function names, and line numbers.
- Displays both stacks in mixed Python/native programs where the script calls out into a native C library.
- Displays Python local variables when a Python frame is selected.
- Evaluations which can also include Python expressions and statements.
- Breakpoints and stepping in Python code.
- After a module is imported, you can see its Python source files listed in the **Project Files** tree.
- Debugs MPI programs written in Python using mpi4py.

This feature is useful when debugging a mixed C, C++, Fortran, and Python program.

Note

Python global variables are only shown in the **Locals** tab if the selected frame is at the module level. To see a global variable, you can add it in the **Evaluate** window. You can see all the global variables, if you add `globals()`.

These Arm DDT features are not supported in Python debugging:

- The **Multi-Dimensional Array Viewer** is not supported in Python frames.
- Offline Python debugging.
- Manual launch and attaching to a Python process.
- Python debugging by opening a core file.
- Watchpoints and Tracepoints.
- Python programs with multiple Python threads running concurrently as this causes Arm DDT to hang.
- Sub-processing libraries (such as `multiprocessing`) because it forks separate processes.
- Stepping from Python frames into native frames.
- Memory debugging only covers the Python interpreter.

Procedure

1. To debug Python scripts, specify the Python interpreter followed by `%allinea_python_debug%` and then the path to the script that you wish to debug. For example:

```
$ ddt python3 %allinea_python_debug% my-script.py
```

2. To debug Python scripts that use MPI, the same applies, except `mpirun` is also appended to the beginning:

```
$ ddt mpirun -np 4 python3 %allinea_python_debug% my-mpi-script.py
```

3. When passing arguments, they must appear after `%allinea_python_debug%` and the name of your script. To run the demo in the examples folder, change into the examples folder.
4. Run:

```
$ make -f python.makefile
```

5. Run:

```
$ ../bin/ddt python3 %allinea_python_debug% python-debugging.py
```

Note

On loading into DDT you will be inside the C code. This is normal as you are debugging the python binary. Depending on the interpreter that you are debugging you may also see a message about missing debug symbols. Clicking Play/Continue once after launching will bring you to the first line of your script.

6. Click **Run**.
7. Click **Play/Continue** to run to the first line of the script.
8. Set a breakpoint on a line inside the `call_out_to_a_library` function.
9. Use the **Add breakpoint** dialog to set a breakpoint on the function name *library_function*.
10. Click **Play/Continue** to run to the Python function and observe that local variables are visible.
11. Click **Play/Continue again** to run to the native function and observe how the stack appears when calling out of the interpreter.

Note

To disable Python debugging, omit `%allinea_python_debug%` from the command line.

B1.17 DDT user interface

Arm DDT uses a tabbed-document interface to display multiple documents. This means you can have many source files open. You can view one file in the full workspace area, or two if the **Source Code viewer** is 'split'.

Each component of Arm DDT is a dockable window that can be dragged around by a handle, usually on the top or left edge. Components can also be double-clicked, or dragged outside of Arm DDT, to form a new window. You can hide or show most of the components using the **View** menu. The screenshot shows the default layout.




Figure B1-14 DDT main window

The table shows the main components:

Key
1 Menu Bar
2 Process Controls
3 Process Groups
4 Find File or Function
5 Project Files
6 Source Code

(continued)

Key
7 Variables and Stack of Current Process/Thread
8 Parallel Stack, IO and Breakpoints
9 Evaluate Window
10 Status Bar

Note

On some platforms the default screen size might be insufficient to display the status bar. If this occurs, you should expand the Arm DDT window until it is completely visible.

B1.18 Save and load sessions

Most of the user-modified parameters and windows can be saved by right-clicking and selecting a save option in the corresponding window.

Arm DDT can also load and save all these options concurrently to minimize the inconvenience in restarting sessions. Saving the session stores such things as *Process Groups*, the contents of the **Evaluate** window, and more. This makes it easy to debug code with the same parameters set every time.

To save a session select **File > Save Session**, enter a file name for the save file (or select an existing file), then click **OK**. To load a session select **File > Load Session**, choose the correct file, then click **OK**.

Chapter B2

Source code

Describes source code viewing, editing, and rebuilding features.

Arm DDT integrates with the Git, Subversion, and Mercurial version control systems, and provides static analysis to automatically detect many classes of common errors.

It contains the following sections:

- [*B2.1 Source code viewer*](#) on page B2-98.
- [*B2.2 Assembly debugging*](#) on page B2-100.
- [*B2.3 Project Files*](#) on page B2-101.
- [*B2.4 Finding lost source files*](#) on page B2-102.
- [*B2.5 Find Files or Functions*](#) on page B2-103.
- [*B2.6 Go To Line*](#) on page B2-105.
- [*B2.7 Navigate through source code history*](#) on page B2-106.
- [*B2.8 Static analysis*](#) on page B2-107.
- [*B2.9 Version control information*](#) on page B2-108.

B2.1 Source code viewer

The source code editing and rebuilding capabilities are not designed for developing programs from scratch. They are designed to fit into existing debugging or profiling sessions that are running on a current executable.

The same capabilities are available for source code whether you are running remotely (using the remote client) or if you are connected directly to your system.

View source code

When you start a session, source code is automatically found from the information compiled in the executable.

Source and header files found in the executable are reconciled with the files present on the front-end server, and displayed in a simple tree view in the **Project Files** tab of the **Project Navigator** window. Click on the file name to view the source file.

When a selected process is stopped, the **Source Code viewer** will automatically move to the correct file and line, if the source is available.

The **Source code viewer** supports automatic color syntax highlighting for C and Fortran.

You can hide functions or subroutines you are not interested in by clicking the ‘-’ glyph next to the first line of the function. This will collapse the function. Click the ‘+’ glyph to expand the function again.

Edit source code

Source code can be edited in the **Source code viewer**. The actions **Undo**, **Redo**, **Cut**, **Copy**, **Paste**, **Select all**, **Go to line**, **Find**, **Find next**, **Find previous**, and **Find in files** are available from the **Edit** menu. Files can be opened, saved, reverted, and closed from the **File** menu.

Note

Information from Arm DDT will not match edited source files until the changes are saved, the binary is rebuilt, and the session restarted.

If the currently selected file has an associated header or source code file, you can open it by right-clicking in the editor and choosing **Open <filename>.<extension>**. There is a global shortcut on function key F4, or you can use **Edit > Switch Header/Source**.

To edit a source file in an external editor, right-click the editor for the file and choose **Open in external editor**. To change the editor used, or if the file does not open with the default settings, select **File > Options** to open the **Options** window (Arm Forge Preferences on Mac OS X) then enter the path to the preferred editor in **Editor**, for example `/usr/bin/gedit`.

If a file is edited, a warning will be displayed at the top of the editor:




Figure B2-1 File edited warning

This is to warn that the source code shown is not the source that was used to produce the currently executing binary. The source code and line numbers may not match the executing code.

Rebuilding and restarting

If source files are edited, the changes will not take effect until the binary is rebuilt and the session restarted. To configure the build command choose **File > Configure Build**, enter a build command and a directory in which to run the command, then click **Apply**.

To issue the build command choose **File > Build**, or press Ctrl+B (Cmd+B on Mac OS X). When a build is issued the **Build Output** view is shown. When a rebuild succeeds we recommend that you restart the session with the new build by choosing **File > Restart Session**.

Committing changes

Changes to source files can be committed using Git, Mercurial, or Subversion. To commit changes, choose **File > Commit**, enter a commit message in the **Commit changes** dialog then click **Commit**.

B2.2 Assembly debugging

You can view disassembly, step over instructions, step into instructions, and set breakpoints on instructions in the **Disassembly** viewer when you are in assembly debugging mode.

Enable assembly debugging

To enable assembly debugging mode click the '0101' button in the toolbar.



Figure B2-2 The '0101' button in the toolbar toggles assembly debugging mode

When you enable assembly debugging mode the **Disassembly** viewer opens where you can view the disassembly of the current symbol that contains the program counter and changes the behaviour of the step buttons to operate on the instruction level.




Figure B2-3 The Disassembly viewer tab shows the disassembly of the current symbol with the source code interleaved

The **Disassembly** viewer auto updates the disassembly when the current symbol that contains the program counter changes.

When assembly debugging mode is disabled, the **Disassembly** viewer closes and reverts the behavior of the step buttons back to stepping source lines.

Breakpoints

The **Disassembly** viewer enables you to set breakpoints on instructions and also on source lines.

Every breakpoint is listed on the **Breakpoints** tab towards the bottom of the window and can be further edited from the **Breakpoints** tab. See [Set breakpoints on page B3-121](#).

B2.3 Project Files

The **Project Files** tree shows a list of source files for your program. Click on a file in the tree to open it in the **Source Code viewer**. You can also expand a source file to see a list of classes, functions, and so on defined in that source file (C / C++ / Fortran only).




Figure B2-4 Function Listing

Click on a source code element (class, function, and so on) to display it in the **Source Code viewer**.

Application Code and External Code

Arm DDT automatically splits your source code into **Application Code**, which is source code from your application, and **External Code**, which is code from third party libraries. This helps you quickly distinguish between your own code and third party libraries.

You can control exactly which directories are considered to contain Application Code using the **Application / External Directories** window. Right-click on the **Project Files** tree to open the window.

The checked directories are the directories containing Application Code. When you have made your changes click **OK** to update the **Project Files** tree.

B2.4 Finding lost source files

Sometimes not all source files are found automatically. This can also occur, for example, if the executable or source files have been moved since compilation. To search for source files in additional directories, right-click in the **Project Files** tab, and select **Add/view Source Directory(s)**. You can also specify extra source directories on the command line using the `-source-dirs` command-line argument (separate each directory with a colon).

It is also possible to add an individual file, if this file has moved since compilation, or is on a different (but visible) file system. To do this right-click in the **Project Files** tab and select **Add File**.

Any directories or files you have added are saved and restored when you use **Save Session** and **Load Session** on the **File** menu. If DDT does not find the sources for your project, you might find these commands save you a lot of unnecessary clicking.

B2.5 Find Files or Functions

The **Find Files Or Functions** dialog is displayed above the sources file list in the **Project Files** tree.

You can type the name of a file, function, or other source code element (such as classes, Fortran modules, and so on) in this dialog to search for that item in the source tree. You can also type just part of a name to see all the items with name that contains the text you typed.

Double-click a result to jump to the corresponding source code location for that item.




Figure B2-5 Find Files or Functions dialog

Find code or variables

You can use **Edit > Find** to find occurrences of an expression in the currently open source file.

The search starts from the current cursor position for the next or previous occurrence of the search term. Click the magnifying glass icon for more search options:

Case sensitive: When selected, the search is case sensitive, so for example, Hello can not match hello.

Whole words only: When selected, the search only matches your search term against whole words in the source file. For example, Hello does not match HelloWorld.

Use Regular Expressions: When selected, your search can use Perl-style regular expressions.

Find code or variables in files

You can use **Edit > Find In Files** to search all source and header files associated with your program. Click the search results list to display the file and line number in the **Source Code viewer**. This can be used for setting a breakpoint at a function.




Figure B2-6 Find in Files dialog

Case sensitive: When selected, the search is case sensitive, so for example, Hello does not match hello.

Whole words only: When selected, the search only matches your search term against whole words in the source file. For example, Hello does not match HelloWorld.

Regular Expression: When selected, the search term is interpreted as a regular expression rather than a fixed string. The syntax of the regular expression is identical to that described in [Job ID regular expression](#) on page Appx-I-527.

B2.6 Go To Line

The Go To Line feature enables you to go directly to a particular line of source code.

Click **Edit > Go To Line** to display a dialog. Enter the line number in the source code that you want to view, then click **OK**. This will take you to that line in the code (if the line exists). You can also use the **CTRL+L** shortcut to open this dialog.

B2.7 Navigate through source code history

After jumping to a particular source code location or opening a new file, you can return to the previous location using **Edit > Navigate backwards in source code history** (or click **Navigate backwards in source code history** on the toolbar). You can navigate back to several previous locations in the source code.

After navigating backwards, you can use **Edit > Navigate forwards in source code history** (or click **Navigate forwards in source code history** in the toolbar) to return to the original location.

B2.8 Static analysis

Static analysis is a powerful companion to debugging. Arm DDT helps you discover errors by code and state inspection along with automatic error detection components such as memory debugging. Static analysis inspects the source code and attempts to identify errors that can be detected from the source alone, independently of the compiler and actual process state.

Arm DDT includes the static analysis tools `cppcheck` and `ftnchek`. These will by default automatically examine source files as they are loaded and display a warning symbol if errors are detected. Typical errors include:

- Buffer overflows. Accessing beyond the bounds of heap or stack arrays.
- Memory leaks. Allocating memory within a function and there being a path through the function which does not deallocate the memory and the pointer is not assigned to any externally visible variable, nor returned.
- Unused variables, and also use of variables without initialization in some cases.




Figure B2-7 Static Analysis error annotation

Static analysis is not guaranteed to detect all, or any, errors. An absence of warnings does not mean there are no bugs.

B2.9 Version control information

The version control integration in DDT and MAP enables you to see line-by-line information from Git, Mercurial, or Subversion next to the source file. Version control information is color-coded to indicate the age of the source line.




Figure B2-8 DDT running with version control information enabled

To view version control information, select **View > Version Control Information**. The column displayed to the left of the source code viewer shows how long ago the line was added or modified. Each line in the column is highlighted in a color that indicates its age:

- Lines modified in the current revision are highlighted in red.
- Lines that have been modified but not committed are highlighted in purple.
- All other lines are highlighted with a blend of transparent blue and opaque green, where blue indicates old changes and green indicates recent changes.

Note

Changes that have not been committed can only be viewed if you are using Git. Version control information for files with uncommitted changes is not displayed if you are using Mercurial or Subversion.

If you hover the cursor over the version control information column, a tooltip displays a preview of the commit message for the commit that last changed the line.




Figure B2-9 Version control information-tooltips

A folded block of code displays the annotation for the most recently modified line in the block.

To copy the commit message, right-click on the desired row in the column then select **Copy commit message**.




Figure B2-10 Version control information-context menu

Also see *Version control breakpoints and tracepoints* on page B3-133.

Chapter B3

Control program execution

Whether debugging multi-process code or single process code, the mechanisms for controlling program execution are very similar.

In multi-process mode, most of the features described in this chapter are applied using **Process Groups**.

For single-process mode the commands and behaviors are identical, but apply to only a single process. In this case you do not need to work with process groups.

It contains the following sections:

- *B3.1 Process control and process groups* on page B3-113.
- *B3.2 Focus control* on page B3-115.
- *B3.3 Start, stop, and restart a program* on page B3-118.
- *B3.4 Step through a program* on page B3-119.
- *B3.5 Stop messages* on page B3-120.
- *B3.6 Set breakpoints* on page B3-121.
- *B3.7 Conditional breakpoints* on page B3-123.
- *B3.8 Suspend breakpoints* on page B3-124.
- *B3.9 Delete breakpoints* on page B3-125.
- *B3.10 Load and save breakpoints* on page B3-126.
- *B3.11 Default breakpoints* on page B3-127.
- *B3.12 Synchronize processes* on page B3-128.
- *B3.13 Set watchpoints* on page B3-129.
- *B3.14 Tracepoints* on page B3-131.
- *B3.15 Version control breakpoints and tracepoints* on page B3-133.
- *B3.16 Examine the stack frame* on page B3-135.
- *B3.17 Align stacks* on page B3-136.
- *B3.18 View stacks in parallel* on page B3-137.

- *B3.19 Browse source code* on page B3-139.
- *B3.20 View multiple files simultaneously* on page B3-141.
- *B3.21 Signal handling* on page B3-142.

B3.1 Process control and process groups

MPI programs are designed to run as more than one process and can span many machines. Arm DDT lets you group these processes so that you can perform actions on more than one process at a time. The status of processes is displayed in the **Process Group viewer**.

The **Process Group viewer** is displayed (by default) at the top of the screen with multi-colored rows. Each row relates to a group of processes and operations that can be performed on the currently highlighted group (for example, playing, pausing, and stepping) by clicking on the toolbar buttons. You can switch between groups by clicking on them or their processes. The highlighted group is indicated by a lighter shade. Groups can be created, deleted, or modified at any time, with the exception of the *All* group, which cannot be modified.

To add a group click the **Create Group** button, or right-click in the **Process Group viewer** and choose **Create Group** from the context-sensitive menu that is displayed. This context menu can also be used to rename groups, delete individual processes from a group, and jump to the current position of a process in the source code viewer. You can load and save the current groups to a file, and you can create sub-groups from the processes currently playing, paused, or finished. You can even create a sub-group excluding the members of another group. For example, to take the complement of the *Workers* group, select the *All* group and choose *Copy, but without Workers*.

You can also use the context menu to switch between the two different methods of viewing the list of groups. These methods are called the *Detailed view* and the *Summary view*.

Detailed view

The Detailed view is ideal for working with small numbers of processes. If your program has 32 processes or less, DDT defaults to the Detailed view. You can use the context menu to switch to this view.




Figure B3-1 The Process Group Detailed view

In the Detailed view, each process is represented by a square containing its MPI rank (0 through n-1). The squares are color-coded; red for a paused process, green for a playing process, and gray for a finished/dead process. Selected processes are highlighted with a lighter shade of their color and the current process also has a dashed border.

When a single process is selected the local variables are displayed in the **Variable viewer** and displayed expressions are evaluated. You can make the **Source Code viewer** jump to the file and line for the current stack frame (if available) by double-clicking on a process.

To copy processes from one group to another, click and drag the processes. To delete a process, press the Delete key. When modifying groups it is useful to select more than one process by holding down one or more of the following keys:

Key	Description
Control	Click to add/remove process from selection
Shift	Click to select a range of processes
Alt	Click to select an area of processes

Note

Some window managers (such as KDE) use Alt and drag to move a window. You must disable this feature in your window manager if you want to use DDT's area select.

Summary view

The Summary view is ideal for working with moderate to large numbers of processes. If your program has 32 processes or more, DDT defaults to this view. You can use the context menu to switch to this view.

All	4 processes (0-3)	Paused: 4	Playing: 0	Finished: 0
Root	1 process (0)	Paused: 1	Playing: 0	Finished: 0
Workers	3 processes (1-3)	Paused: 3	Playing: 0	Finished: 0
Show processes	Currently selected:	1	(on , pid 19199, main thread IWP 19199)	
Create Group				

Figure B3-2 The Process Group Summary view

In the Summary view, individual processes are not shown. Instead, for each group, DDT shows:

- The number of processes in the group.
- The processes belonging that group. Here **1-2048** means processes 1 through 2048 inclusive, and **1-10, 12-1024** means processes 1-10 and processes 12-1024 (but not process 11). If this list becomes too long, it is truncated with a If you hover the mouse over the list you can see more details.
- The number of processes in each state (playing, paused, or finished). If you hover the mouse over each state you can see a list of the processes currently in that state.
- The rank of the currently selected process. You can change the current process by clicking here, typing a new rank, then pressing Enter. Only ranks belonging to the current group will be accepted.

Show processes lets you switch a single group into the Detailed view and back again. This is useful if you are debugging a 2048 process program, but have narrowed the problem down to just 12 processes, which you have put in a group.

B3.2 Focus control

The focus control allows you to focus on individual processes, or threads, or on process groups. When focused on a particular process or thread, actions such as stepping, playing/pausing, or adding breakpoints, will only apply to that process or thread rather than the entire group.

Focus on current: ☒ Group ☐ Process ☐ Thread

Figure B3-3 Focus options

Focusing affects a number of different controls in the main window. The user interface will change depending on whether you are focused on group, process, or thread. The relevant information about your currently focused object will be displayed.

Note

Focus control does not affect all windows. Windows such as the Multi-Dimensional Array Viewer, Memory Debugger, and Cross-Process Comparison are unchanged.

Process group viewer

The changes to the process group viewer are the most obvious in the user interface when using focus control. When focus on current group is selected, you see your process groups. If you switch to focus on current process or thread, the view changes to show the processes in the currently selected group, with their corresponding threads.



Figure B3-4 The Process Group Detailed View focused on a process

If there are 32 threads or more, by default the threads will be displayed using a Summary view (as in the Process Group View). You can change the view mode using the context menu.

If focused on process, a tooltip shows the OpenMP thread ID of each thread, if the value exists.

Breakpoints tab

The **Breakpoints** tab is filtered to only display breakpoints relevant to your current group, process, or thread. When focused on a process, the **Breakpoints** tab displays which thread the breakpoint belongs to. If you are focused on a group, the tab displays both the process and the thread the breakpoint belongs to.

Source Code viewer

The line highlight color in the Source Code viewer shows you a stack back trace of where each thread is in the call stack. This is filtered by the currently focused item, for example when focused on a particular process, you only see the back trace for the threads in that process.

Also, when you add breakpoints using the Source Code viewer, they are added for the group, process, or thread that is currently focused.

Parallel Stack View

The **Parallel Stack View** can also be filtered by focusing on a particular process group, process, or thread.

Playing and stepping

The behavior of playing, stepping, and the **Run to here** feature are also affected by your currently focused item:

- When focused on process group, the entire group is affected
- When focused on thread, only the current thread is executed
- When focused on process, only the current process is executed. In this case you can also use the **Step threads together** feature.

Step threads together

The **Step threads together** feature is only available when focused on process. If this option is enabled, the threads in the current process are synchronized when performing actions such as stepping, pausing, and using **Run to here**.

For example, if you have a process with two threads and you choose *Run to here*, your program will be paused when either of the threads reaches the specified line. If **Step threads together** is selected, both of the threads will be played to the specified line before the program is paused.

Note

You should always use **Step threads together** and **Run to here** to enter or move within OpenMP parallel regions. With many compilers it is also advisable to use **Step threads together** when leaving a parallel region, because threads can get "left behind" inside system-specific locking libraries and may not enter the next parallel region on the first attempt.

When using the **Step threads together** feature it is not always possible for all threads to synchronize at their target. There are two main reasons for this:

- One or more threads may branch into a different section of code (and never reach the target). This is especially common in OpenMP codes, where worker threads are created and remain in holding functions during sequential regions.
- As most of the supported debug interfaces cannot play arbitrary groups of threads together, this behavior is simulated by playing each thread in turn. This is usually not a problem, but can be if, for example, thread 1 is playing, but waiting for thread 2 (which is not currently playing). Arm DDT attempts to resolve this automatically but cannot always do so.

If either of these conditions occur, the **Stepping threads** window opens, displaying the threads which have not yet reached their target.




Figure B3-5 The Stepping threads window

The **Stepping threads** window also displays the status of threads, which may be one of the following:

- **Done:** The thread has reached its target (and has been paused).
- **Skipped:** The thread has been skipped and paused. Arm DDT no longer waits for it to reach its target.
- **Playing:** This is the thread that is currently being executed. Only one thread may be playing at a time while the **Stepping threads** window is open.
- **Waiting:** The thread is currently awaiting execution. When the currently *playing* thread is *done* or has been skipped, the highest *waiting* thread in the list is executed.

The **Stepping threads** window also lets you interact with the threads with the following options:

- **Skip:** Arm DDT skips and pauses the currently playing thread. If this is the last waiting thread the window is closed.
- **Try Later:** The currently playing thread is paused, and added to the bottom of the list of threads to be retried later. This is useful if you have threads which are waiting on each other.
- **Skip All:** This skips, and pauses, all of the threads and closes the window.

B3.3 Start, stop, and restart a program

If a program is running you can end it, run it again, or run another program using commands on the **File** menu. The **File** menu can be accessed at almost any time.

When the start up process is complete, your program should automatically stop either at the main function for non-MPI codes, or at the MPI-Init function for MPI.

Restart a program: When a job has run to the end, you are asked if you want to restart the job. If you select **Yes**, any remaining processes are killed, temporary files are cleared, and the session restarts from scratch with the same program settings.

Stop a program: When ending a job, DDT attempts to shut down all the processes and clear up any temporary files. If this fails for any reason you may have to manually `kill` your processes using `kill`, or a method provided by your MPI implementation such as `lamclean` for LAM/MPI.

B3.4 Step through a program

To continue a program playing click **Play/Continue**.

To stop a program playing click **Pause**.

In multi-process mode these commands will start/stop all the processes in the current process group (see [Process control and process groups](#) on page B3-113).

There are three different types of step available:

- **Step Into** moves to the next line of source code unless there is a function call, in which case it steps to the first line of that function.
- **Step Over** moves to the next line of source code in the bottom stack frame.
- **Step Out** executes the rest of the function and then stops on the next line in the stack frame above. The return value of the function is displayed in the **Locals** view. When using **Step Out** be careful not to step out of the main function, as doing this ends your program.

B3.5 Stop messages

There are five reasons why your program might be automatically paused:

- The program hit one of the default breakpoints, for example, `exit` or `abort`. See [Default breakpoints on page B3-127](#) for more information.
- The program hit a user-defined breakpoint, that is a breakpoint shown in the **Breakpoints** view.
- The value of a watched variable changed.
- The program was sent a signal. See [Signal handling on page B3-142](#) for more information.
- The program encountered a memory debugging error. See [Pointer error detection and validity checking on page B8-200](#) for more information.

The message displayed informs you why the program was paused. To copy the message text to the clipboard, select it, right-click, then select **Copy**.

If you want to suppress these messages, for example if you are playing from one breakpoint to another, use the **Control Messages** menu to enable or disable stop messages.

B3.6 Set breakpoints

Use the Source Code viewer

Right-click in the **Source Code viewer** where you want to place a breakpoint, then choose to add a breakpoint.

You can also add a breakpoint by clicking in the margin to the left of the line number.

If you have lots of source code files, use the **Find/Find In Files** window to search for a particular function.

In multi-process mode, this sets the breakpoint for every member of the current process group.

Every breakpoint is listed on the **Breakpoints** tab.

If you add a breakpoint at a location where there is no executable code, the line you selected as having a breakpoint is highlighted. However, when the breakpoint is reached, the program stops at the next executable line of code.

Use the Add Breakpoint window

You can also add a breakpoint by clicking the *Add Breakpoint* button in the toolbar. This opens the **Add Breakpoint** window.

Add Breakpoint

Location:

☒ **Line** **File:** /home/user/ddt/examples/hello.c

Line Number: 133

☐ **Function**

Applies To:

Process Group: All

Process: All

Thread: All

Hit Limits:

Start on the n-th pass: 0

Trigger every n-th pass: 1

Stop after n hits: Forever

☐ **Condition:**

Language: Auto

Help **Add** **Cancel**

Figure B3-6 The Add Breakpoint window

You can add a breakpoint in a function for which you do not have any source code, for example in `malloc`, `exit`, or `printf` from the standard system libraries. In this case, select **Function** and enter the name of the function in the field next to it.

You can specify which process group/process/thread you want the breakpoint to apply to using the **Applies To** section. You can also make the breakpoint conditional by selecting the **Condition** check box and entering a condition in the field next to it.

Pending breakpoints

Note

Pending breakpoints are not supported on all platforms.

If you try to add a breakpoint in a function that is not defined, you will be asked if you want to add it anyway. If you click **Yes** the breakpoint is applied to any shared objects that are loaded in the future.

B3.7 Conditional breakpoints

The **Breakpoints** tab displays all the breakpoints in your program. You can add a condition to any breakpoint by clicking on the **Condition** cell in the breakpoints table and entering an expression that evaluates to *true* or *false*.

Breakpoints									
	Processes	Threads	File	Line	Function	Condition	Start After	Trigger Every	Stop After
☒	process 0	all	hello.c	133			0	1	Forever
☒	All	all	hello.c	148		my_rank == 3	0	1	Forever

Figure B3-7 The Breakpoints table

Each time a process (in the group the breakpoint is set for) passes this breakpoint, it evaluates the condition and breaks only if it returns *true* (typically any non-zero value). You can drag an expression from the **Evaluate** window into the **Condition** cell for the breakpoint to automatically set it as the condition.

Conditions can be any valid expression for the language of the file containing the breakpoint. This includes other variables in your program and function calls.

Breakpoints									
	Processes	Threads	File	Line	Function	Condition	Start After	Trigger Every	Stop After
☒	process 0	all	hello.f	55			0	1	Forever
☒	All	all	hello.f	49		my_rank .EQ. 3	0	1	Forever

Figure B3-8 Conditional breakpoints in Fortran

We recommend that you avoid using functions with side effects as they will be executed every time the breakpoint is reached.

The expression evaluation may be more pedantic than your compiler. To ensure the correct interpretation of, for example, boolean operations, we recommend that you use brackets explicitly, to ensure correct evaluation.

B3.8 Suspend breakpoints

To deactivate a breakpoint, either:

- Clear the activated column in the **Breakpoints** tab. To reactivate a breakpoint, select the activated column.
- Right-click the breakpoint icon in the source code editor and choose **Disable**. To reactivate a breakpoint, right-click and choose **Enable**.
- Hold SHIFT and select a breakpoint icon in the source code editor.

Breakpoints that are disabled are grayed out.

B3.9 Delete breakpoints

There are a number of ways to delete breakpoints:

- Right-click on the breakpoint in the **Breakpoints** tab and select **Delete breakpoint**.
- Right-click in the file/line of the breakpoint while in the correct process group, then select **Delete breakpoint**.
- Click the breakpoint icon in the margin to the left of the line number in the **Source Code viewer**.

B3.10 Load and save breakpoints

To load or save the breakpoints in a session, right-click in the **Breakpoints** tab and select **Load/save**. Breakpoints are also loaded and saved as part of the load/save session.

B3.11 Default breakpoints

There are a number of default breakpoints that stop your program when particular conditions are met. You can enable/disable these while your program is running using **Control > Default Breakpoints**.

- **Stop at exit/_exit** (Disabled by default)
When enabled, your program is paused as it is about to end under normal exit conditions. Your program pauses both before and after any exit handlers have been executed.
- **Stop at abort/fatal MPI Error** (Enabled by default)
When enabled, your program is paused as it about to end after an error has been triggered. This includes MPI and non-MPI errors.
- **Stop on throw (C++ exceptions)** (Disabled by default)
When enabled, your program is paused whenever an exception is thrown (regardless of whether or not it will be caught). Due to the nature of C++ exception handling, you may not be able to step your program properly at this point. Instead, you should play your program or use the **Run to here** feature.
- **Stop on catch (C++ exceptions)** (Disabled by default)
As for **Stop on throw**, but triggered when your program catches a thrown exception. Again, you may have trouble stepping your program.
- **Stop at fork**
Your program stops when your program forks (that is, calls the fork system call to create a copy of the current process). The new process is added to your existing session, and can be debugged along with the original process.
- **Stop at exec**
When your program calls the exec system call, your program stops at the main function (or program body for Fortran) of the new executable.
- **Stop on CUDA kernel launch**
When debugging CUDA GPU code, this pauses your program at the entry point of each kernel launch.

B3.12 Synchronize processes

If the processes in a process group are stopped at different points in the code, you can resynchronize them to a particular line of code by right-clicking on the line where you want to synchronize them, and selecting **Run To here**. This effectively plays all the processes in the selected group, and puts a break point at the line where you choose to synchronize the processes, ignoring any breakpoints that the processes may encounter before they have synchronized at the specified line.

Notes:

- Breakpoints are ignored while the groups are synchronized, but they are not removed.
- If a process is already at the line where you choose to synchronize, the process will still be set to play. Make sure that your process will revisit this line, or alternatively, synchronize to the line immediately after the current line.
- If you choose to synchronize your code at a point where all processes do not reach, the processes that cannot get to this point will play to the end.

B3.13 Set watchpoints

A watchpoint is a variable or expression that is monitored, so when it is changed or accessed the program is paused.




Figure B3-9 Program stopped at watchpoint

To add a watchpoint:

- Right-click on the **Watchpoints** tab and select **Add Watchpoint**.
- Right-click over a variable in the **Source Code** viewer then select **Add Watchpoint**.
- Add watchpoints automatically by dragging a variable from the **Local Variables**, **Current Line**, or **Evaluate** views in to the **Watchpoints** tab.

Processes	Scope	Expression	Trigger On	Implemented in
☒ All		beingWatched	read and write	software

Figure B3-10 The Watchpoints table

When you add a watchpoint, the **Add Watchpoint** dialog is displayed. You can use this dialog to apply restrictions to the watchpoint:

- **Process Group** restricts the watchpoint to the chosen process group (see [Process control and process groups on page B3-113](#)).
- **Process** restricts the watchpoint to the chosen process.
- **Expression** is the variable name in the program to be watched.
- **Language** is the language of the portion of the program containing the expression.
- **Trigger On** allows you to select whether the watchpoint will trigger when the expression is read, written, or both.

You can set a watchpoint for either a single process, or every process in a process group.

Unlike breakpoints, watchpoints are not displayed in the **Source Code** viewer.

The automatic watchpoints are write-only by default.

A watchpoint is automatically removed when the target variable goes out of scope. If you are watching the value pointed to by a variable, that is, `*p`, you might want to continue watching the value at that address even after `p` goes out of scope. You can do this by right-clicking on `*p` in the **Watchpoints** tab and selecting **Pin to address**. This replaces the variable `p` with its address so the watch is not removed when `p` goes out of scope.

Modern processors have hardware support for a handful of watchpoints that are set to watch the contents of a memory location. Consequently, watchpoints can normally be used with no performance penalty.

Where the number of watchpoints used is over this quantity, or the expression being watched is too complex to tie to a fixed memory address, the implementation is through software monitoring, which imposes significant performance slowdown on the application being debugged.

The number of hardware watchpoints available depends on the system. The read watchpoints are only available as hardware watchpoints.

Consequently, watchpoints should, where possible, be a single value that is stored in a single memory location. While it is possible to watch the whole contents of non-trivial user defined structures or an entire array simultaneously, or complex statements involving multiple addresses, these can cause extreme application slow down during debugging.

B3.14 Tracepoints

Tracepoints enable you to see what lines of code your program is executing, and the variables, without stopping the program. When a thread reaches a tracepoint it will print the file and line number of the tracepoint to the **Input/Output** tab. You can also capture the value of any number of variables or expressions at that point.

Tracepoints can be particularly useful in these situations, for example:

- Recording entry values in a function that is called many times, but crashes only occasionally. If you set a tracepoint it is easier to correlate the circumstances that cause a crash.
- Recording entry to multiple functions in a library, enabling the user or library developer to check which functions are being called, and in which order. An example of this is the MPI History Plugin, which records MPI usage. See [Use a plugin on page B9-210](#).
- Observing progress of a program and variation of values across processes without having to interrupt the program.

Set tracepoints

To add a tracepoint:

- Right-click on a line in the **Source Code viewer** and select **Add Tracepoint**.
- Right-click in the **Tracepoints** tab and select **Add Tracepoint**.

When you right-click in the **Source Code viewer** a number of variables based on the current line of code are captured by default.

Tracepoints can lead to considerable resource consumption if they are placed in locations likely to generate a lot of passing. For example, if a tracepoint is placed inside a loop with N iterations, then N separate tracepoint passings will be recorded.

There is an attempt to merge such data in a scalable manner, but when alike tracepoints are passed in order between processes, where process behavior is likely to be divergent and unmergeable, then a considerable load would result.

If it is necessary to place a tracepoint inside a loop, set a condition on the tracepoint to ensure you only log what is of use to you. Conditions may be any valid expression in the language of the file the tracepoint is placed in and may include function calls, although you may want to be careful to avoid functions with side effects as these will be evaluated every time the tracepoint is reached.

Tracepoints also momentarily stop processes at the tracepoint location to evaluate the expressions and record their values. This means if they are placed inside (for example) a loop with a very large number of iterations, or a function executed many times per second, then you will see a slowdown in your program.

Tracepoint output

The output from tracepoints can be viewed on the **Tracepoint Output** tab.

Tracepoint	Processes	Values logged
subdomain (subdomain.f90:59)	16, ranks 0-15	ny: — 16 nx: — 16 nz: — 64
blts (blts.f90:58)	1, rank 0	m: 1 iend: 16 ldmz: 64 k: 2 ldmx: 16 i: 2 ldz: ist: 2 j: 2 ldmy: 16
blts (blts.f90:58)	1, rank 0	m: 2 iend: 16 ldmz: 64 k: 2 ldmx: 16 i: 2 ldz: ist: 2 j: 2 ldmy: 16
blts (blts.f90:58)	1, rank 0	m: 3 iend: 16 ldmz: 64 k: 2 ldmx: 16 i: 2 ldz: ist: 2 j: 2 ldmy: 16
blts (blts.f90:58)	1, rank 0	m: 4 iend: 16 ldmz: 64 k: 2 ldmx: 16 i: 2 ldz: ist: 2 j: 2 ldmy: 16

Figure B3-11 Output from tracepoints in a Fortran program

Where tracepoints are passed by multiple processes within a short interval, the outputs will be merged. Sparklines of the values recorded are shown for numeric values, along with the range of values obtained, showing the variation across processes.

As alike tracepoints are merged, this can lose the order/causality between *different* processes in tracepoint output. For example, if process 0 passes a tracepoint at time T, and process 1 passes the tracepoint at $T + 0.001$, this will be shown as one passing of both process 0 and process 1, with no ordering inferred.

Sequential consistency is preserved during merging, so for any process, the sequence of tracepoints for that process will be in order.

To find particular values or interesting patterns, use **Only show lines containing** at the bottom of the panel. Tracepoint lines matching the text entered here will be shown, the rest will be hidden. When searching for a particular value, for example `my_var: 34`, the space at the end helps distinguish between `my_var: 34` and `my_var: 345`.

For more detailed analysis, export the tracepoints by right-clicking and choosing **Export**. An HTML tracepoint log will be written using the same format as in offline mode.

B3.15 Version control breakpoints and tracepoints

Version control breakpoint and tracepoint insertion allows you to quickly record the state of the parts of the target program that were last modified in a particular revision. The resulting tracepoint output can be viewed in the **Tracepoint Output** tab or the **Logbook** tab, and can be exported or saved as part of a logbook or offline log.




Figure B3-12 DDT with version control tracepoints

Version control tracepoints can be inserted either in graphical interactive mode or in offline mode via a command-line argument.

In interactive mode, enable **View > Version Control Information**. The annotation column is displayed in the **Source Code viewer** for files that are tracked by a supported version control system.




Figure B3-13 Enable version control using the View menu

To find all the source files, detect the variables on the lines modified in the revision, and insert tracepoints (pending if necessary), right-click a line last modified by the revision of interest and choose **Trace variables at this revision**.




Figure B3-14 Version Control-Trace at this revision

A progress dialog may be displayed for lengthy tasks.

You can double-click on the tracepoints and tracepoint output in the **Tracepoints**, **Tracepoint Output**, and **Logbook** tabs during a session to jump to the corresponding line of code in the **Source Code viewer**.

In offline mode, supply the additional argument `-trace-changes` to apply the same process as in interactive mode using the current revision of the repository.

By default, version control tracepoints are removed after 20 hits. To change this hit limit set the environment variable `ALLINEA_VCS_TRACEPOINT_HIT_LIMIT` to an integer greater than or equal to 0. To configure version control tracepoints to have no hit limit set this to 0.

See also [Version control information on page B2-108](#).

B3.16 Examine the stack frame

The stack back trace for the current process and thread are displayed on the **Current Stack** tab of the **Variables** window.




Figure B3-15 The Current Stack tab

When you select a stack frame you jump to that position in the code, if it is available, and the local variables for that frame will be displayed. The toolbar can also be used to step up or down the stack, or jump straight to the bottom-most frame.

B3.17 Align stacks

Click **Align stacks** (Ctrl+Shift+A) to highlight the same stack-frame depth of the specific thread number across all processes.

This feature is particularly useful where processes are interrupted by the pause button, and are at different stages of computation. This enables tools such as the **Cross-Process Comparison** window to compare equivalent local variables, and also simplifies casual browsing of values.

B3.18 View stacks in parallel

You can use the **Parallel Stack View** to see the status of your program in one view.




Figure B3-16 Parallel Stack View

To find out the location of a group's processes, click on the group. The **Parallel Stack View** displays a tree of functions, merged from every process in the group (by default). If there is only one branch in this tree, one list of functions, then all your processes are at the same place.

If there are several branches, your group has split up and is in different parts of the code. Click on any branch to see its location in the **Source Code viewer**, or hover your mouse over it to view a list of the processes at that location in a popup. Right-click on any function in the list and select **New Group** to automatically gather the processes at that function together in a new group, labeled by the function's own name.

The **Parallel Stack View** can be used to create groups, to display and select large numbers of processes based on their location in your code (this is invaluable when dealing with moderate to large numbers of processes), and watch what happens as you step processes through your code.

The **Parallel Stack View** takes over much of the work of the **Stack** display, but instead of just showing the current process, this view combines the call trees (commonly called *stacks*) from many processes and displays them together. The call tree of a process is the list of functions (strictly speaking *frames* or locations within a function) that lead to the current position in the source code.

For example, if `main()` calls `read_input()`, and `read_input()` calls `open_file()`, and you stop the program inside `open_file()`, then the call tree looks like the following:

```

main()
  read_input()
    open_file()

```

If a function was compiled with debug information (typically using `-g`), extra information is added, displaying the exact source file and line number that your code is on.

Any functions without debug information are grayed-out and are not shown by default. Functions without debug information are typically library calls or memory allocation subroutines and are not generally of interest. To see the entire list of functions, right-click on one and choose **Show Children**.

You can click on any function to select it as the 'current' function in Arm DDT. If it was compiled with debug information, its source code is displayed in the main window, and its local variables and so on in the other windows.

One of the most important features of the **Parallel Stack View** is its ability to show the position of many processes at once. Right-click on the view to toggle between:

- Viewing all the processes in your program at once.
- Viewing all the processes in the current group at once (default).
- Viewing only the current process.

The function that is currently displayed and being used for the variable views is highlighted in dark blue. If you click on another function in the **Parallel Stack View** this selects another frame for the source code and variable views. It also updates the **Stack** display, since these two controls are complementary. If the

processes are at several different locations, only the location of the current process is displayed in dark blue. The locations of the other processes are displayed in light blue:

Processes	Threads	Function
16	16	main (hello.c:117)
16	16	func1 (hello.c:39)
16	16	progress_engine (minheap-internal.h:97)
16	16	progress_engine

Figure B3-17 Current frame highlighting in Parallel Stack View

In the example above, the processes of the program are at two different locations. One process is in the `main` function, at line 147 of `hello.c`. The other 15 processes are inside a function called `func1`, at line 39 of `hello.c`. To see the line of source code a function corresponds to, and display any local variable in that stack frame, click on the function.

There are two optional columns in the **Parallel Stack View**. **Processes** shows the number of processes at each location. **Threads** shows the number of threads at each location. By default, only the number of processes is shown. Right-click to turn these columns on and off. Note that in a normal, single-threaded MPI program, each process has one thread and these two columns will show identical information.

If you hover the mouse over any function in the **Parallel Stack View** this displays the full path of the filename, and a list of the process ranks that are at that location in the code:

Input/Output*	Breakpoints	Watchpoints	Stacks	Tracepoints	Tracepoint Output
Stacks					
Processes	Threads	Function			
4	4	main (hello.c:117)			
4	4	func1 (hello.c:39)			
4	4	func2 (hello.c:30)			
4	4	progress_engine (minheap-internal.h:97)			
4	4	progress_engine			

/home/user/arm/forge/examples/hello.c:30
4 Processes: ranks 0-3

Figure B3-18 Parallel Stack View tool tip

Arm DDT is at its most intuitive when each process group is a collection of processes doing a similar task. The **Parallel Stack View** is invaluable in creating and managing these groups.

To create a new process group that contains only the processes sharing that location in code, right-click on a function in the combined call tree and choose **New Group**. By default the name of the function is used for the group, or the name of the function with the file and line number if it is necessary to distinguish the group further.

B3.19 Browse source code

Source code is automatically displayed when a process is stopped, when you select a process, or position in the stack changed. If the source file cannot be found you are prompted for its location.

Lines of the source code are highlighted to show the current location of your program's execution. Lines that contain processes from the current group are shaded in that group's color. Lines only containing processes from other groups are shaded in gray.

This pattern is repeated in the focus on process and thread modes. For example, when you focus on a process, lines containing that process are highlighted in the group color, while other processes from that group are highlighted in gray.

Lines of code that are on the stack are also highlighted, functions that your program will return to when it has finished executing the current one. These are drawn with a faded look to distinguish them from the currently-executing lines.

You can hover the mouse over any highlighted line to see which processes/threads are currently on that line. This information is presented in a variety of ways, depending on the current focus setting:

Focus on Group

A list of groups that are on the selected line, along with the processes in them on this line, and a list of threads from the current process on the selected line.

Focus on Process

A list of the processes from the current group that are on this line, along with the threads from the current process on the selected line.

Focus on Thread

A list of threads from the current process on the selected line.

The tooltip distinguishes between processes and threads that are currently executing that line, and ones that are on the stack by grouping them under the headings **On the stack** and **On this line**.

Variables and Functions

Right-click on a variable or function name in the **Source Code viewer** to check whether there is a matching variable or function, and then to display extra information and options in a sub-menu.

In the case of a variable, the type and value are displayed, along with options to view the variable in the **Cross-Process Comparison View (CPC)** or the **Multi-Dimensional Array Viewer (MDA)**, or to drop the variable into the **Evaluate** window, each of which are described in the next chapter.




Figure B3-19 Right-click menu-variable options

In the case of a function, it is also possible to add a breakpoint in the function, or to the source code of the function when available.




Figure B3-20 Right-click menu-function options

B3.20 View multiple files simultaneously

Occasionally it may be useful to view two source files at the same time. For example, if you are tracking two different processes.

To view two files simultaneously, right-click in the **Source Code view** to split the view. This displays a second panel beneath the first panel. When viewing multiple files, the currently 'active' panel displays the file. Click on one of the panels to make it active.

To remove the split view and return to viewing one source file, right-click in the **Source Code viewer** and clear the split view option.




Figure B3-21 Horizontal alignment of multiple source files

B3.21 Signal handling

By default a process is stopped if it encounters one of the standard signals. The standard signals include:

- SIGSEGV - Segmentation fault

The process has attempted to access memory that is not valid for that process. Often this will be caused by reading beyond the bounds of an array, or from a pointer that has not been allocated yet. The [Memory debugging on page B8-193](#) feature might help to resolve this problem.

- SIGFPE - Floating Point Exception

This is raised typically for integer division by zero, or dividing the most negative number by -1. Whether or not this occurs is operating system dependent, and not part of the POSIX standard. Linux platforms will raise this.

Note that floating point division by zero will not necessarily cause this exception to be raised, behavior is compiler dependent. The special value Inf or -Inf may be generated for the data, and the process would not be stopped.

- SIGPIPE - Broken Pipe

A broken pipe has been detected while writing.

- SIGILL - Illegal Instruction

SIGUSR1, SIGUSR2, SIGCHLD, SIG63 and SIG64 are passed directly through to the user process without being intercepted by DDT.

Custom signal handling (signal dispositions)

You can change the way that individual signals are handled using the **Signal Handling** dialog. To open the dialog, select **Control > Signal Handling**.




Figure B3-22 Signal Handling dialog

To stop the process when it encounters a signal, set the **Action** to *Stop*.

To let the process receive the signal and continue playing without being stopped by the debugger, set the **Action** to *Ignore*.

Send signals

The **Send Signal** window lets you send a signal to the debugged processes. To send a signal, select **Control > Send Signal**, select the signal you want to send, then click **Send to process**.

Chapter B4

Variables and data

This chapter describes the **Variables** window.

It contains the following sections:

- *B4.1 Variables window* on page B4-146.
- *B4.2 Sparklines* on page B4-147.
- *B4.3 Current line* on page B4-148.
- *B4.4 Local variables* on page B4-149.
- *B4.5 Arbitrary expressions and global variables* on page B4-150.
- *B4.6 Edit variables* on page B4-152.
- *B4.7 Help with Fortran modules* on page B4-153.
- *B4.8 View complex numbers in Fortran* on page B4-154.
- *B4.9 C++ STL support* on page B4-155.
- *B4.10 Custom pretty printers* on page B4-156.
- *B4.11 View array data* on page B4-157.
- *B4.12 UPC support* on page B4-158.
- *B4.13 Change data values* on page B4-159.
- *B4.14 View numbers in different bases* on page B4-160.
- *B4.15 Examine pointers* on page B4-161.
- *B4.16 Multi-dimensional arrays in the Variable View* on page B4-162.
- *B4.17 Multi-dimensional array viewer (MDA)* on page B4-164.
- *B4.18 Cross-process and cross-thread comparison* on page B4-172.
- *B4.19 Assign MPI ranks* on page B4-174.
- *B4.20 View registers* on page B4-175.
- *B4.21 Process details* on page B4-176.
- *B4.22 Disassembly* on page B4-177.
- *B4.23 Interact directly with the debugger* on page B4-178.


B4.1 Variables window

The **Variables** window contains two tabs that provide different ways to list your variables. The **Locals** tab contains all the variables for the current stack frame. The **Current Line(s)** tab displays all the variables referenced on the currently selected lines.

Note

Several compilers and libraries (such as Cray Fortran, and OpenMP) generate extra code, including variables that are visible in Arm DDT's windows.

The right-click menu in the **Variables** window enables you to edit values, change the display base, compare data across processes and threads, and choose whether the fields in structures (classes or derived types) are displayed alphabetically by element name or not (which is useful for when structures have many different fields).



Name	Value
argc	1
▶ argv	0x7ffffff...
beingWatched	0
▶ bigArray	
dest	0
▶ dynamicArray	0x81802...
▶ environ	0x7ffffff...
i	0
message	""
my_rank	0
p	512
source	32767
▶ status	
▶ t2	0x60305...
▶ tables	
tag	50
▶ test	
x	10000
y	12

Figure B4-1 Variables in the Locals tab

B4.2 Sparklines

Numerical values may have sparklines displayed next to them. A sparkline is a line graph of process rank or thread index against value of the related expression. The exact behavior is determined by the focus control. See [Focus control](#).

When focused on process groups, process ranks are used. Otherwise, thread indices are used. The graph is bound by the minimum and maximum values found, or in the case that all values are equal the line is drawn across the vertical center of the highlighted region. Erroneous values such as *Nan* and *Inf* are represented as red vertical bars. If focus is on process groups, clicking on a sparkline displays the **Cross-Process Comparison** window for closer analysis. Otherwise, clicking on a sparkline displays the **Cross-Thread Comparison** window.

B4.3 Current line

You can select a single line by clicking on it in the **Source Code viewer**. You can select multiple lines by clicking and dragging. The variables are displayed in a tree view so that user-defined classes or structures can be expanded to view the variables contained within them. You can drag a variable from this tab into the **Evaluate** window. It is then evaluated in whichever stack frame, thread, or process you select.

B4.4 Local variables

The **Locals** tab contains local variables for the current process's currently active thread and stack frame.

For Fortran codes the amount of data reported as local can be substantial, as this can include many global or common block arrays. If this is a problem, it is best to conceal this tab underneath the **Current Line(s)** tab so it will not update after every step.

Note

Variables defined within common blocks might not appear in the **Locals** tab with some compilers. This is because they are considered to be global variables when defined in a common memory space.

The **Locals** tab compares the value of scalar variables against other processes. If a value varies across processes in the current group, the value is highlighted in green.

When stepping or switching processes, if the value of a variable is different from the previous position or process, it is highlighted in blue.

After stepping out of function the return value is displayed at the top of the **Locals** tab (for selected debuggers).

B4.5 Arbitrary expressions and global variables

The global variables and arbitrary expressions are not displayed with the local variables. You can click on the line in the **Source Code viewer** that contains a reference to the global variable and use the **Current Line(s)** tab.

Alternatively, the **Evaluate** window can be used to view the value of any arbitrary expression.




Figure B4-2 Evaluating Expressions

Right-click on the **Evaluate** window, click **Add Expression**, then type in the expression required in the current source file language. This value of the expression is displayed for the current process and stack/thread, and is updated after every step.

Notes:

- Arm DDT does not apply the usual rules of precedence to logical Fortran expressions, such as `x .ge. 32 .and. x .le. 45`. You need to bracket such expressions thoroughly, for example: `(x .ge. 32) .and. (x .le. 45)`.
- Although the Fortran syntax allows you to use keywords as variable names, Arm DDT is not able to evaluate such variables on most platforms. Contact [Arm support](#) if this issue affects you.

Expressions that contain function calls are only evaluated for the current process/thread, and sparklines are not displayed for those expressions. This is because of possible side effects caused by calling functions. Use the **Cross-Process Comparison** or **Cross-Thread Comparison** windows for functions instead. See [Cross-process and cross-thread comparison](#) on page B4-172.

Fortran intrinsics

The following Fortran intrinsics are supported by the default GNU debugger included with Arm DDT:

ABS	AIMAG	CEILING	CMPLX
FLOOR	IEEE_IS_FINITE	IEEE_IS_INF	IEEE_IS_NAN
IEEE_IS_NORMAL	ISFINITE	ISINF	ISNAN
ISNORMAL	MOD	MODULO	REALPART

Support in other debuggers, including the CUDA debugger variants, may vary.

Changing the language of an expression

By default, expressions in the **Evaluate** window, **Locals** tab, and **Current Line(s)** tab are evaluated in the language of the current stack frame. This might not always be appropriate. For example, a pointer to user-defined structure might be passed as value within a Fortran section of code, and you might want to

view the fields of the C structure. Alternatively, you might want to view a global value in a C++ class while your process is in a Fortran subroutine.

You can change the language of an expression by right-clicking on the expression, choosing **Change Type/Language**, and selecting the appropriate language for the expression. To restore the default behavior, change this back to **Auto**.

Macros and #defined constants

By default, many compilers do not output sufficient information to allow the debugger to display the values of **#defined** constants or macros, as including this information can greatly increase executable sizes.

For the GNU compiler, when you add the **-g3** option to the command line options this generates extra definition information which will then be displayed.

B4.6 Edit variables

You can edit the values of simple types such as scalars, pointers, and c-strings. To edit a value, right-click the value in a variable view and select **Edit Value**. Enter the new value in the **Edit Value** dialog then click **OK**.




Figure B4-3 Edit variables

B4.7 Help with Fortran modules

An executable that contains Fortran modules presents a special set of problems:

- If there are many modules, and they all contain many procedures and variables (each of which can have the same name as something else in a separate Fortran module), keeping track of which name refers to which entity can be difficult.
- When the **Locals** or **Current Line(s)** tabs (in the **Variables** window) display one of these variables, to which Fortran module does the variable belong?
- How do you refer to a particular module variable in the **Evaluate** window?
- How do you quickly jump to the source code for a particular Fortran module procedure?

To help with these problems, use the **Fortran Modules** tab in the **Project Navigator** window.

When a session starts, Fortran module membership is automatically found from the information compiled into the executable.

A list of Fortran modules found is displayed in a simple tree view in the **Fortran Modules** tab of the **Project Navigator** window.

To expand a module, click on the + symbol to the left of the module name. This displays a list of member procedures, member variables, and the current values of those member variables.

When you click on one of the procedure names the **Source Code viewer** jumps to that procedure's location in the source code. In addition, the return type of the procedure is displayed at the bottom of the **Fortran Modules** tab. Fortran subroutines will have a return type of VOID ().

When you click on one of the displayed variable names, the type of that variable is displayed at the bottom of the **Fortran Modules** tab.

A module variable can be dragged and dropped into the **Evaluate** window. Here, all of the usual **Evaluate** window functionality applies to the module variable. To help with variable identification in the **Evaluate** window, module variable names are prefixed with the Fortran module name and two colons ::.

If you right-click in the **Fortran Modules** tab, you can use the context menu to send the variable to the **Evaluate** window, **Multi-Dimensional Array Viewer**, or **Cross-Process Comparison View** window.

There are some limitations to the information displayed in the **Fortran Modules** tab:


- The **Fortran Modules** tab is not displayed if the underlying debugger does not support the retrieval and manipulation of Fortran module data.
- The **Fortran Modules** tab displays an empty module list if the debug data for the Fortran modules is not present or is not in a recognized format.
- The debug data compiled into an executable does not include any indication of the module USE hierarchy. For example, if module A USEs module B, the inherited members of module B are not shown under the data displayed for module A. Consequently, the **Fortran Modules** tab shows the module USE hierarchy in a flattened form, one level deep.

B4.8 View complex numbers in Fortran

When working with complex numbers, you might want to view only the real or imaginary elements of the number. This can be useful when evaluating expressions, or viewing an array in the **Multi-Dimensional Array Viewer**. See [Multi-dimensional array viewer \(MDA\)](#).

You can use the Fortran intrinsic functions `REALPART` and `AIMAG` to get the real or imaginary parts of a number, or their C99 counterparts `creal` and `cimag`.

Complex numbers in Fortran can also be accessed as an array, where element 1 is the real part, and element 2 is the imaginary part.



Name	Value
c	(3,4)
c(1)	3
c(2)	4

Figure B4-4 Viewing the Fortran complex number 3+4i

B4.9 C++ STL support

Arm DDT uses pretty printers for the GNU C++ STL implementation (version 4.7 and later), Nokia's Qt library, and Boost, designed for use with the GNU Debugger. These are used automatically to present such C++ data in a more understandable format.

For some compilers, the STL pretty printing can be confused by non-standard implementations of STL types used by a compiler's own STL implementation. In this case, and in the case where you want to see the underlying implementation of an STL type, you can disable pretty printing using the environment variable setting `ALLINEA_DISABLE_PRETTY_PRINT=1`.

Expanding elements in `std::map`, including `unordered` and `multimap` variants, is not supported when using object keys or pointer values.

B4.10 Custom pretty printers

In addition to the pre-installed pretty printers you can also use your own GDB pretty printers.

A GDB pretty printer consists of an auto-load script that is automatically loaded when a particular executable or shared object is loaded and the actual pretty printer Python classes themselves. To make a pretty printer available, copy it to `~/.allinea/gdb`.

An example pretty printer can be found in `{installation-directory}/examples`.

Compile the `fruit` example program using the GNU C++ compiler as shown:

```
cd {installation-directory}/examples
make -f fruit.makefile
```

Start Arm DDT with the example program:

```
ddt --start {installation-directory}/examples/fruit
```

When the program has started, right-click on line 20 then use the **Run to here** command. The internal variable of `myFruit` is displayed on the **Locals** tab.

Now install the `fruit` pretty printer by copying the files to `~/.allinea/gdb`:

```
cp -r {installation-directory}/examples/fruit-pretty-printer/* ~/.allinea/gdb/
```

Re-run the program and run to line 20, as before. Click on the **Locals** tab and you can see that now, instead of the internal variable `myFruit`, the type of `fruit` is displayed instead.

B4.11 View array data

Fortran users might find that it is impossible to view the upper bounds of an array. This is due to a lack of information from the compiler. In these circumstances, the array is displayed with a size of 0, or simply `<unknown_bounds>`. It is still possible to view the contents of the array using the **Evaluate** window to view `array(1)`, `array(2)`, and so on, as separate entries.

To input the size of the array, right-click on the array and select **Edit Type**. In the **Edit Type** window, enter the real type of the array in **New Type**.




Figure B4-5 Edit Type window

Alternatively, the **Multi-Dimensional Array Viewer** can be used to view the entire array.

B4.12 UPC support

Arm DDT supports the Cray UPC compiler.

When debugging UPC applications there are a few changes in the user interface:

- Processes are identified as **UPC Threads**. This is purely a terminology change for consistency with the UPC language terminology. UPC Threads have identical behavior to that of separate processes. For example, groups, process control, and cross-process data comparison apply across UPC Threads.
- The type qualifier `shared` is given for shared arrays or pointers to shared.
- Shared pointers are printed as a triple (address, thread, phase). For indefinitely blocked pointers the phase is omitted.
- Referencing shared items yield a shared pointer, and pointer arithmetic can be performed on shared pointers.
- Dereferencing a shared pointer (for example, dereferencing `*(&x[n] + 1)`) correctly evaluates and fetches remote data where required.
- Values in shared arrays are not automatically compared across processes: the value of `x[i]` is by definition identical across all processes. It is not possible to identify pending read/write to remote data. Non-shared data types such as local data or local array elements will still be compared automatically.
- Distributed arrays are handled implicitly by the debugger. There is no need to use the explicit *distributed dimensions* feature in the **Multi-Dimensional Array Viewer**.

All other components are identical to debugging any multi-process code.

B4.13 Change data values

To set the value of an expression, right-click in the **Evaluate** window and select **Edit Value**. This enables you to change the value of the expression for the current process, current group, or for all processes.

———— **Note** ————

The variable must exist in the current stack frame for each process you wish to assign the value to.

—————

B4.14 View numbers in different bases

When you are viewing an integer numerical expression you can right-click on the value and choose **View As** to change which base the value is displayed in. The **View As Default** option displays the value in its original (default) base.

B4.15 Examine pointers

You can examine pointer contents by clicking the + next to the variable or expression. This expands the item and dereferences the pointer.

In the **Evaluate** window, you can also use the **View As Vector**, **Get Address**, and **Dereference Pointer** menu items:

- **Dereference Pointer** wraps the expression in `*`.
- **Get Address** strips a single layer of `*` from the expression (if one exists).
- Both **Get Address** and **Dereference Pointer** only support raw pointers and not other pointer implementations, such as, C++11 smart pointers.

See also [Multi-dimensional array viewer \(MDA\)](#) on page B4-164.

B4.16 Multi-dimensional arrays in the Variable View

When you view a multi-dimensional array in the **Locals**, **Current Line(s)**, or **Evaluate** windows, it is possible to expand the array to view the contents of each cell.

In C/C++ the array expands from left to right, x, y, z will be seen with the x column first, then under each x cell a y column.

Name	Value
array	
[0]	
[0]	1
[1]	2
[2]	3
[1]	
[0]	2
[1]	4
[2]	6
[2]	
[0]	3
[1]	6
[2]	9
[3]	
[0]	4
[1]	8
[2]	12

Figure B4-6 2D Array in C: type of array is int[4][3]

In Fortran the opposite will be seen with arrays being displayed from right to left as you read it so x, y, z will have z as the first column with y under each z cell.

Current Line(s)	
Name	Value
▼ twodee	
▼ [1]	
[1]	1
[2]	2
[3]	3
[4]	4
[5]	5
▼ [2]	
[1]	2
[2]	4
[3]	6
[4]	8
[5]	10
▼ [3]	
[1]	3
[2]	6
[3]	9

Figure B4-7 2D Array in Fortran: type of twodee is integer(3,5)

The first thousand elements in an array are shown in the **Locals** or **Current Line(s)** tabs. Larger arrays are truncated. Elements after the first thousand can be viewed by evaluating an expression or by using the **Multi-Dimensional Array Viewer**.

B4.17 Multi-dimensional array viewer (MDA)

You can use the **Multi-Dimensional Array Viewer** (MDA) to view multi-dimensional arrays.

To open the **Multi-Dimensional Array Viewer**, right-click on a variable in the **Source Code** viewer, **Locals** tab, **Current Line(s)** tab or **Evaluate** window and select **View Array (MDA)**. You can also open the MDA directly by selecting **View > Multi-Dimensional Array Viewer**.




Figure B4-8 Multi-Dimensional Array Viewer

If you open the MDA by right-clicking on a variable, the **Array Expression** and other parameters will be automatically set based on the type of the variable. Click **Evaluate** to see the contents of the array in the **Data Table**.

Use **Full Window** to expand the table of values (and hide the settings at the top of the window). This enables you to make full use of your screen space. Click **Full Window** again to view the settings.

Array expression

The **Array Expression** is an expression containing a number of *subscript metavariables* that are substituted with the subscripts of the array. For example, the expression `myArray($i, $j)` has two metavariables, `$i` and `$j`. The metavariables are unrelated to the variables in your program.

The range of each metavariable is defined in the fields below the expression, for example **Range of \$i**. The **Array Expression** is evaluated for each combination of `$i`, `$j`, and so on. The results are shown in the **Data Table**. You can control whether each metavariable is shown in the **Data Table** using **Rows** or **Columns**.

By default, the ranges for these metavariables are integer constants entered using spin boxes. However, you can also specify these ranges as *expressions* in terms of program variables. These expressions are then evaluated in the debugger. To allow the entry of these expressions, select the **Staggered Array** checkbox. This converts all the range entry fields from spin boxes to line edits allowing the entry of freeform text.

The metavariables can be reordered by dragging and dropping them. For C/C++ expressions the major dimension is on the left and the minor dimension on the right. For Fortran expressions the major dimension is on the right and the minor dimension on the left. Distributed dimensions cannot be reordered, they must always be the most major dimensions.

Filter by value

You can configure the **Data Table** to only show elements that fit certain criteria, for example elements that are zero.

If the **Only show if** checkbox is selected, only elements that match the boolean expression in the adjacent field are displayed in the **Data Table**. For example, `$value == 0`. The special metavariable `$value` in the expression is replaced by the actual value of each element. The **Data Table** automatically hides rows or columns in the table where no elements match the expression.

Any valid expression for the current language can be used here, including references to variables in scope and function calls. We recommend that you avoid functions with side effects as these will be evaluated many times.

Distributed arrays

A distributed array is an array that is distributed across one or more processes as local arrays.

The **Multi-Dimensional Array Viewer** can display certain types of distributed arrays, namely UPC shared arrays (for supported UPC implementations), and general arrays where the distributed dimensions are the most major, that is, the distributed dimensions change the most slowly, and are independent from the non-distributed dimensions.

UPC shared arrays are treated the same as local arrays. Right-click on the array variable and select **View Array (MDA)**.

To view a non-UPC distributed array, create a process group that contains all the processes that the array is distributed over.

If the array is distributed over all processes in your job, select the **All** group when you right-click on the local array variable in the **Source Code viewer**, **Locals** tab, **Current Line(s)** tab or **Evaluate** window.

The **Multi-Dimensional Array Viewer** will open with the **Array Expression** already filled in.

Enter the number of **Distributed Array Dimensions**. A new subscript metavariable (such as `$p`, `$q`) will be automatically added for each distributed dimension.

Enter the ranges of the distributed dimensions so that the product is equal to the number of processes in the current process group, then click **Evaluate**.

Advanced: how arrays are laid out in the data table

The **Data Table** is two-dimensional, but the **Multi-Dimensional Array Viewer** may be used to view arrays with any number of dimensions, as the name implies. This section describes how multi-dimensional arrays are displayed in the two-dimensional table.

Each subscript metavariable (such as $\$i$, $\$j$, $\$p$, $\$q$) maps to a separate dimension on a hypercube. Usually the number of metavariables is equal to the number of dimensions in a given array, but this does not necessarily need to be the case. For example `myArray( $\$i$ ,  $\$j$ ) *  $\$k$` introduces an extra dimension, $\$k$, as well as the two dimensions corresponding to the two dimensions of `myArray`.

The figure below corresponds to the expression `myArray( $\$i$ ,  $\$j$ )` with $\$i = 0..3$ and $\$j = 0..4$.

	$\$j$				
	0	1	2	3	4
0	A	D	F		
1	B	E			
$\$i$					
2	C				
3					

Figure B4-9 `myArray( $\$i$ ,  $\$j$ )` with $\$i = 0..3$ and $\$j = 0..4$.

For example, imagine that `myArray` is part of a three-dimensional array distributed across three processes. The figure below shows what the local arrays look like for each process.




Figure B4-10 The local array `myArray($i, $j)` with `$i = 0..3` and `$j = 0..4` on ranks 0-2
And as a three-dimensional distributed array with `$p` the distributed dimension:




Figure B4-11 A three-dimensional distributed array comprised of the local array `myArray($i, $j)` with $i = 0..3$ and $j = 0..4$ on ranks 0-2 with p the distributed dimension

This cube is projected (just like 3D projection) onto the two dimensional **Data Table**. Dimensions marked **Display** as **Rows** are shown in rows, and dimensions marked **Display** as **Columns** are shown in columns, as you would expect.

More than one dimension may viewed as Rows, or more than one dimension viewed as Columns.

The dimension that changes fastest depends on the language your program is written in. For C/C++ programs the leftmost metavariable (usually `$i` for local arrays, or `$p` for distributed arrays) changes the most slowly (just like with C array subscripts). The rightmost dimension changes the most quickly. For Fortran programs the order is reversed, that is the rightmost is most major, the leftmost most minor.

The figure below shows how the three-dimensional distributed array above is projected onto the two-dimensional **Data Table**:

[illegible]

Figure B4-12 A three-dimensional distributed array comprised of the local array myArray(\$i, \$j) with \$i = 0..3 and \$j = 0..4 on ranks 0-2 projected onto the Data Table with \$p (the distributed dimension) and \$j displayed as Columns and \$i displayed as Rows.

Auto update

If you select the **Auto update** checkbox, the **Data Table** will automatically update as you switch between processes/threads and step through the code.

Compare elements across processes

When viewing an array in the **Data Table**, double-click an element or right-click an element and choose **Compare Element Across Processes** to open the **Cross-Process Comparison View** for the selected element.

See *Cross-process and cross-thread comparison* on page B4-172 for more information.

Statistics tab

The **Statistics** tab displays information which might be of interest, such as the range of the values in the table, and the number of special numerical values, such as `nan` or `inf`.

Export

You can export the contents of the table to a file in the comma-separated values (CSV) or HDF5 format so that it can be plotted or analyzed in your favourite spreadsheet or mathematics program.

There are two CSV export options: List (one row per value), and Table (same layout as the on screen table).

Note _____

If you export a Fortran array in HDF5 format, the contents of the array are written in column major order. This is the order expected by most Fortran code, but the arrays will be transposed if read with the default

settings by C-based HDF5 tools. Most HDF5 tools have an option to switch between row major and column major order.

Visualization

If your system is OpenGL-capable then a 2-D slice of an array, or table of expressions, can be displayed as a surface in 3-D space using the **Multi-Dimensional Array Viewer**.

You can only plot one or two dimensions at a time. If your table has more than two dimensions the **Visualize** button will be disabled.

After filling the table of the MDA viewer with values, click **Visualize** to open a 3-D view of the surface.

To display surfaces from two or more different processes on the same plot, select another process in the main process group window then click **Evaluate** in the MDA viewer. When the values are ready, click **Visualize** again.

The surfaces displayed on the graph can be hidden and shown using the checkboxes on the right side of the window.

The graph can be moved and rotated using the mouse and a number of extra options are available from the window toolbar.

The mouse controls are:

- Hold down the left button and drag to rotate the graph.
- Hold down the right button to zoom. Drag forwards to zoom in, and backwards to zoom out.
- Hold the middle button and drag to move the graph.

Notes:

- Arm DDT requires OpenGL to run. If your machine does not have hardware OpenGL support, software emulation libraries such as MesaGL are also supported.
- In some configurations OpenGL is known to crash. A workaround if the 3D visualization crashes is to set the environment variable `LIBGL_ALWAYS_INDIRECT` to 1. The precise configuration which triggers this problem is not known.




Figure B4-13 DDT visualization

The toolbar and menu have options to configure lighting and other effects, including a function to save an image of the surface as it currently displays.

B4.18 Cross-process and cross-thread comparison

The **Cross-Process Comparison View (CPC)** and **Cross-Thread Comparison View (CTC)** can be used to analyze expressions calculated on each of the processes in the current process group. Each window displays information in three ways: raw comparison, statistically, and graphically.

This is a more detailed view than the sparklines that are automatically drawn against a variable in the **Evaluate** window, **Locals** tab, and **Current Line(s)** tab for multi-process sessions.

To compare values across processes or threads, right-click on a variable in the **Source Code viewer**, **Locals** tab, **Current Line(s)** tab, or **Evaluate** window, then choose either **View Across Processes (CPC)** or **View Across Threads (CTC)**. You can also open the CPC or CTC directly from the **View** menu. Alternatively, clicking on a sparkline will bring up the CPC if focus is on process groups, and the CTC otherwise.




Figure B4-14 Cross-Process Comparison View-Compare view

Processes and threads are grouped by expression value when using the raw comparison. The precision of this grouping can be specified (for floating point values) by using the **Limit comparison to** fields.

If you are comparing across processes, you can turn each of these groupings of processes into a process group by clicking **Create Groups**. This creates several process groups, one for each line in the panel. Using this capability, large process groups can be managed with simple expressions to create groups. These expressions are any valid expression in the current language (that is, C/C++/Fortran).

For threaded applications, when using the CTC, if OpenMP thread IDs can be identified, a third column will display the corresponding OpenMP thread IDs for each thread that has each value. The value displayed in this third column for any non-OpenMP threads that are running depends on your compiler, but is typically -1 or 0

You can enter a second boolean expression in the **Only show if** field to control which values are displayed. Only values for which the boolean expression evaluates to `true` / `.TRUE.` are displayed in the results table. The special metavariable `$value` in the expression is replaced by the actual value. Click **See Examples** to view examples.

Select **Align stack frames** to ensure that the stack frame of the same depth in every thread (using the CTC) is examined when comparing the variable value. When using the CPC, **Align stack frames** ensures that the same thread number is also selected when examining variable values. This is very helpful for most programs, but you might want to disable it if different processes or threads run entirely different programs.

The **Use as MPI Rank** button is described in [Assign MPI ranks on page B4-174](#).

You can create a group for the ranks corresponding to each unique value by clicking **Create Groups**.

The **Export** button enables you to export the list of values and corresponding ranks as a comma-separated values (CSV) file.

The **Full Window** button hides the settings at the top of the window so that the list of values occupies the full window. This enables you to make full use of your screen space. Click **Full Window** again to display the settings.

The **Statistics** panel shows Maximum, Minimum, Variance, and other statistics for numerical values.

B4.19 Assign MPI ranks

Sometimes the MPI rank for each of your processes cannot be detected. This might be because you are using an experimental MPI version, because you have attached to a running program, or only part of a running program. Whatever the reason, it is easy to define what each process should be called:

1. Choose a variable that holds the MPI world rank for each process, or an expression that calculates it.
2. Use **Cross-Process Comparison View** to evaluate the expression across *all* the processes. If the variable is valid, **Use as MPI Rank** will be enabled.
3. Click **Use as MPI Rank**. All of its processes will be relabeled with these new values.

The criteria for a variable or an expression to be valid are:


- It must be an integer.
- Every process must have a unique number afterwards.

These are the only restrictions. As you can see, there is no need to use the MPI rank if you have an alternate numbering scheme that makes more sense in your application. In fact you can relabel only a few of the processes and not all, if you prefer, so long as afterwards *every* process still has a unique number.

B4.20 View registers

To view the values of machine registers on the currently selected process, select **View > Registers**. These values will be updated after each instruction, change in thread, or change in stack frame.

To edit the value of a register either double-click on the register, or right-click and select **Edit Value**.



Name	Value
rax	0x0 0
rbx	0x0 0
rcx	0x0 0
rdx	0x0 0
rsi	0x7fffffffbc88 1407...
rdi	0x1 1
rbp	0x7fffffffba0 0x7ff...
rsp	0x7fffffff1c00 0x7ff...
r8	0x2 2
r9	0x1 1
r10	0x603010 6303760
r11	0x1 1
r12	0x400ad0 4197072
r13	0x7fffffffbc80 1407...
r14	0x0 0
r15	0x0 0
rip	0x400d11 0x400d...
eflags	0x202 [IF 1]

Figure B4-15 Registers tab

B4.21 Process details

To view the **Process Details** dialog, select **Tools > Process Details**. Details can be sorted by any column, in ascending or descending order.




Figure B4-16 Process Details dialog

B4.22 Disassembly

To view the disassembly (assembly instructions) of a function, select **Tools > Disassemble**. By default you will see the disassembly of the current function. You can view the disassembly of a different function by typing the name in the **Function** field and clicking **Disassemble**.




Figure B4-17 Disassembly dialog

B4.23 Interact directly with the debugger

The **Raw Command** window enables you to send commands directly to the debugger interface. This window bypasses Arm DDT and its book-keeping. So, if you set a breakpoint in this window, it will not be listed in the **Breakpoints** tab.




Figure B4-18 Raw Command window

Be careful with this window. We recommend that you only use it if the graphical interface does not provide the information or control you require. If you send commands such as `quit` or `kill`, this may cause the interface to stop responding to Arm DDT.

Each command is sent to the current group or process, depending on the current focus. If the current group or process is running, you will be prompted to pause the group or process.

Chapter B5

Program input and output

The **Input/Output** tab is at the bottom of the window (by default).

It contains the following sections:

- [B5.1 View standard output and error](#) on page B5-180.
- [B5.2 Save output](#) on page B5-181.
- [B5.3 Send standard input](#) on page B5-182.

B5.1 View standard output and error

The output from all processes is collected and can be viewed on the **Input/Output** tab. Standard output and error are both displayed. However, on most MPI implementations, error is not buffered (output is buffered), so might be delayed.




Figure B5-1 Standard output

The output can be selected and copied to the clipboard.

MPI users should note that most MPI implementations place their own restrictions on program output. Some buffer it all until `MPI_Finalize` is called, and others may ignore it. If your program needs to emit output as it runs, try writing to a file.

————— **Note** —————

Many systems buffer `stdout` but not `stderr`. If you do not see your `stdout` appearing immediately, try adding `fflush(stdout)` or equivalent to your code.

—————

B5.2 Save output

Right-click on the text to either save it to a file or copy a selection to the clipboard.

B5.3 Send standard input

You can use the **stdin** field in the **Run** window to choose the file to use as the standard input (stdin) for your program. Arguments will be automatically added to mpirun to ensure your input file is used.

Alternatively, you can enter the arguments directly in the **mpirun Arguments** field. For example, if using MPI directly from the command-line you would normally use an option to the mpirun such as `-stdin filename`, then you can add the same options to the **mpirun Arguments** field when you start your session in the **Run** window.

It is also possible to enter input during a session. To do this, start your program as normal, then open the **Input/Output** tab. Type the input you want to send.

Click **More** to send input from a file, or send an EOF character.

————— **Note** —————

Although input can be sent while your program is paused, the program must then be played to read the input and act upon it.

The input you type will be sent to all processes.




Figure B5-2 Sending input

Chapter B6

Logbook

This chapter describes how to use the logbook.

It contains the following sections:

- [B6.1 Usage on page B6-184.](#)
- [B6.2 Annotation on page B6-185.](#)
- [B6.3 Logbook comparison on page B6-186.](#)

B6.1 Usage

The **Logbook** tab is located at the bottom of the main window.

The **Logbook** automatically generates a log of all of your actions. It does not require any additional configuration.

For example, if you add a breakpoint or a tracepoint, or play the program, these actions are logged. For each stop of the program, the reason and location is recorded together with the parallel stacks and local variables for one process.




Figure B6-1 Logbook example of a debug session

You can export the current logbook as HTML, or compare it to a previously exported one.

To export the logbook, click the disk icon on the right side of the **Logbook** tab and specify a filename.

You can open previously saved logbooks from the **Tools** menu.

This enables comparative debugging and repeatability. It is always clear how a certain situation in the debugger was caused as the previous steps are visible.

B6.2 Annotation

You can add annotations to the logbook using either the pencil icon on the right side of the **Logbook** tab, or by right-clicking the logbook and choosing **Add annotation**.

B6.3 Logbook comparison

Two logbooks can be compared side-by-side using the **Logbook Files Comparison** window.

To run a comparison, click the ‘compare’ icon on the right side of the **Logbook** tab. Compare the current logbook with another logbook file, or choose two different files to compare.

To easily find differences, align both logbook files to corresponding entries and click the ‘lock’ icon. This fixes the vertical and horizontal scrollbars of the logbooks so that they scroll together.




Figure B6-2 Logbook Files Comparison window with tracepoint difference selected

Chapter B7

Message queues

The **Message Queues** window shows the status of the message buffers of MPI.

You can use it to, for example:

- Show the messages that have been sent by a process but not yet received by the target.
- Detect common errors such as deadlock. This is where all processes are waiting for each other.
- Detect when messages are present that are unexpected, which can correspond to two processes disagreeing about the state of progress through a program.

This feature relies on the MPI implementation supporting it via a debugging support library: the majority of MPIs provide this. Not all implementations support this capability to the same extent, so a variation between the information provided by each implementation is to be expected.

It contains the following sections:

- [B7.1 View message queues on page B7-188.](#)
- [B7.2 Interpret message queues on page B7-189.](#)
- [B7.3 Deadlock on page B7-191.](#)

B7.1 View message queues

To open the **Message Queues** window, select **Tools > Message Queues**. This will query the MPI processes for information about the state of the queues.

When the window is open, click **Update** to refresh the current queue information. Note that this will stop all playing processes. A dialog might be displayed while the data is being gathered. You can cancel the request at any time.

The message queue support library from your MPI implementation (if one exists) will be automatically loaded. If it fails to load, an error message will display.

Common reasons for failing to load include:

- The support library does not exist, or its use must be explicitly enabled.

Most MPIs will build the library by default, without additional configuration flags. MPICH 3 must be configured with the `-enable-debuginfo` argument. MVAPICH 2 must be configured with the `-enable-debug` and `-enable-sharedlib` arguments. Some MPIs, notably Cray's MPI, do not support message queue debugging at all.

Intel MPI includes the library, but debug mode must be enabled. See [Intel MPI on page Appx-E-465](#) for details.

LAM and Open MPI automatically compile the library.

- The support library is not available on the compute nodes where the MPI processes are running.

Ensure the library is available, and set the environment variable `ALLINEA_QUEUE_DLL` if necessary to force using the library in its new location.

- The support library has moved from its original installation location.

Ensure the proper procedure for the MPI configuration is used. This might require you to specify the installation directory as a configuration option.

Alternatively, you can specifically include the path to the support library in the `LD_LIBRARY_PATH`. If this is not convenient you can set the environment variable, `ALLINEA_QUEUE_DLL`, to the absolute path of the library itself (for example, `/usr/local/mpich-3.3.0/lib/libtvmmpich.so`).

- The MPI is built to a different bit-size to the debugger.

In the unlikely case that the MPI is not built to the bit-size of the operating system, the debugger might not be able to find a support library that is the correct size. This is unsupported.

B7.2 Interpret message queues

To see the messages in a group, you must choose an option in **Select communicator**. The ranks displayed in the diagram are the ranks within the communicator (not MPI_COMM_WORLD), if **Show local ranks** is selected.




Figure B7-1 Message Queue window

To see the 'usual' ranks, select **Show global ranks**. The messages displayed can be restricted to particular processes or groups of processes. To restrict the display in the grid to a single process, select **Individual Processes** in **Display mode**, then select the rank of the process. To display a group of processes, select **Process Groups** in **Display mode** then select the ring arc corresponding to the required group. Both of these display modes support multiple selections.

There are three different types of message queues about which there is information. Different colors are used to display messages from each type of queue.

Label	Description
Send Queue	Calls to MPI send functions that have not yet completed.
Receive Queue	Calls to MPI receive functions that have not yet completed.
Unexpected Message Queue	Represents messages received by the system but the corresponding receive function call has not yet been made.

Messages in the *Send* queue are represented by a red arrow, pointing from the sender to the recipient. The line is solid on the sender side, but dashed on the recipient side (to represent a message that has been sent but not yet received).

Messages in the *Receive* queue are represented by a green arrow, pointing from the sender to the recipient. The line is dashed on the sender side, but solid on the recipient side, to represent the recipient being ready to receive a message that has not yet been sent.

Messages in the *Unexpected* queue are represented by a dashed blue arrow, pointing from sender of the unexpected message to the recipient.

A message to self is indicated by a line with one end at the center of the diagram.

Note that the quality and availability of message queue data can vary considerably between MPI implementations. Sometimes the data can therefore be incomplete.

B7.3 Deadlock

A loop in the graph can indicate deadlock. This is where every process is waiting to receive from the preceding process in the loop. For synchronous communications, such as with MPI_Send, this is a common problem.

For other types of communication it can be the case, with MPI_Send that messages get stuck, for example in an O/S buffer, and the send part of the communication is complete but the receive has not started. If the loop persists after playing the processes and interrupting them again, this indicates a deadlock is likely.

Chapter B8

Memory debugging

The powerful parallel memory debugging feature intercepts calls to the system memory allocation library, recording memory usage, and confirming correct usage of the library by performing heap and bounds checking.

Typical problems that can be resolved using memory debugging:

- Memory exhaustion due to memory leaks can be prevented using the **Current Memory Usage** display, which groups and quantifies memory according to the location at which blocks have been allocated.
- Persistent but random crashes caused by access of memory beyond the bounds of an allocation block can be diagnosed using the **Guard Pages** feature.
- Crashing due to deallocation of the same memory block twice, deallocation via invalid pointers, and other invalid deallocations, for example deallocating a pointer that is not at the start of an allocation.

It contains the following sections:

- [B8.1 Enable memory debugging](#) on page B8-194.
- [B8.2 CUDA memory debugging](#) on page B8-195.
- [B8.3 PMDK memory debugging](#) on page B8-196.
- [B8.4 Memory debugging options](#) on page B8-197.
- [B8.5 Pointer error detection and validity checking](#) on page B8-200.
- [B8.6 Current memory usage](#) on page B8-204.
- [B8.7 Memory Statistics](#) on page B8-206.

B8.1 Enable memory debugging

To enable memory debugging, in the **Run** window select the **Memory Debugging** checkbox.

The default options are usually sufficient, but you might need to configure extra options (described in the following sections) if you have a multithreaded application or multithreaded MPI, such as that found on systems using Open MPI with Infiniband, or a Cray XE6 system.

When memory debugging is enabled, start your application as normal. The settings will be propagated through your MPI or batch system when your application starts.

If it is not possible to load the memory debugging library, a message will be displayed. You should refer to [Memory debugging options on page B8-197](#) for possible solutions.

B8.2 CUDA memory debugging

There are two options for debugging memory errors in CUDA programs. They can be found in the **CUDA** section of the **Run** window.

See [Prepare to debug GPU code on page B10-217](#) before debugging the memory of a CUDA application.

When **Track GPU allocations** is enabled, CUDA memory allocations *made by the host* are tracked. That is, allocations made using functions such as `cudaMalloc`. You can find out how much memory is allocated, and where it was allocated from using the **Current Memory Usage** window.

Note

Currently only CUDA memory allocations made by calling `cudaMalloc`, `cudaMallocPitch`, and `cudaMalloc3D` are tracked. Memory allocations made for CUDA arrays with functions such as `cudaMallocArray` are not tracked.

Allocations are tracked separately for each GPU and the host. If you enable **Track GPU allocations**, host-only memory allocations made using functions such as `malloc` will be tracked as well. You can choose between GPUs using the drop-down list in the top-right corner of the **Memory Usage** and **Memory Statistics** windows.

The **Detect invalid accesses (memcheck)** option switches on the CUDA-MEMCHECK error detection tool. This tool can detect problems such as out-of-bounds and misaligned global memory accesses, and syscall errors, such as calling `free` () in a kernel on an already free'd pointer.

The other CUDA hardware exceptions (such as a stack overflow) are detected regardless of whether this option is selected or not.

For further details about CUDA hardware exceptions, see the [NVIDIA documentation](#).

Known issue: It is not possible to track GPU allocations created by an OpenACC compiler because it does not directly call `cudaMalloc`.

B8.3 PMDK memory debugging

Memory debugging can be used to track all allocations made by `libpmemobj`, an object store library that is part of the Persistent Memory Development Kit (PMDK).

To use PMDK memory debugging, enable memory debugging in the **Run** dialog. Optionally, a backtrace can be stored for each allocation. If memory debugging is not enabled, only the call site of the allocation is stored. No other configuration options have an effect on PMDK.

When the pool is opened with `pmemobj_open`, all the allocations that exist in the pool are tracked. The call site is where the pool is opened. The root object of the pool is not tracked. Allocation tracking persists after an aborted transaction.

When tracking allocations, to see if a pointer was allocated by PMDK, right-click on a pointer in a variable view and select **View Pointer Details**. In the pointer details, you also see the backtrace, or call site, of the allocation. In the main menu, **Tools > Current Memory Usage**, and **Tools > Overall Memory Stats**, are enabled. By default, allocations made by `libc` are shown. To see the **Memory Usage** window, **Allocation Table**, and **Memory Statistics** window that you see with regular memory debugging, select **PMDK** in **Allocations from**. The sizes displayed are the sizes returned by `pmemobj_alloc_usable_size`, not the sizes you request.

B8.4 Memory debugging options

Manual configuration is often unnecessary, but it can be used to adjust the memory checks and protection, or to alter the information which is gathered.

The current settings are displayed in the **Memory Debugging** section on the **Run** dialog.

To examine or change the settings, click **Details** adjacent to the **Memory Debugging** checkbox. The **Memory Debugging Options** window displays.




Figure B8-1 Memory Debugging Options

The available settings are:

- **Preload the memory debugging library.** When this is selected, the memory debugging library will be automatically loaded. The memory debugging library can only preload when you start a program. It uses shared libraries.

Preloading is not possible with statically-linked programs, or when attaching to a running process. See the *Static linking* section for more information.

When attaching, you can set the `DMALLOC_OPTIONS` environment variable before running your program. See the *Change settings at run time* section for more information.

- **Language.** Choose the option that best matches your program, for example C, Fortran, No threads, threads. It is often sufficient to leave this set to **C++**, **threads** rather than continually changing this value.
- The **Heap Debugging** section allows you to trade speed for thoroughness. Some important things to remember are:
 - Even the fastest (furthest left) setting will catch trivial memory errors such as deallocating memory twice.
 - The further right you go, the more slowly your program will execute. In practice, the **Balanced** setting is still fast enough to use and will catch almost all errors. If you come across a memory error that is difficult to pin down, choosing **Thorough** might expose the problem earlier, but you will need to be very patient for large, memory intensive programs. See also the *Change settings at run time* section.
- **Enabled Checks.** Shows which checks are enabled for each setting. See the *Available checks* section for a complete list of available checks.
- The **Heap Overflow/Underflow Detection** section can be used to detect out-of-bounds heap access. See *Writing beyond an allocated area* on [Pointer error detection and validity checking](#) on page B8-200 for more details.
- **Check heap consistency.** Almost all users can leave the heap check interval at the default value. It determines how often the memory debugging library will check the entire heap for consistency. This is a slow operation, so it is normally performed every 100 memory allocations. We recommend a higher value (1000 or above) if your program allocates and deallocates memory very frequently, for example, inside a computation loop.
- **Store stack backtraces for memory allocations.** If your program runs particularly slowly with memory debugging enabled, you might be able to get a modest speed increase by clearing this checkbox. This disables stack backtraces in the **View Pointer Details** and **Current Memory Usage** windows, support for custom allocators, and cumulative allocation totals.
- **Only enable for these processes.** When this is selected, you can enable memory debugging only for the specified MPI ranks. Note that when you enable this feature, the memory debugging library is still preloaded into the other processes, but no errors are reported. Furthermore, backtraces for memory allocation are not stored and guard pages are not added for the other processes.

Note

If you choose the wrong library to preload, or the wrong number of bits, your job might not start, or might make memory debugging unreliable. You should check these settings if you experience problems when memory debugging is enabled.

Static linking

If your program is statically linked you must explicitly link the memory debugging library with your program to use the memory debugging feature.

To link with the memory debugging library, you must add the appropriate flags from the table below at the *very beginning* of the link command. This ensures that all instances of allocators, in both user code and libraries, are wrapped. Any *definition* of a memory allocator preceding the memory debugging link flags can cause partial wrapping, and unexpected runtime errors.

Note

If in doubt use `libdmallocthcxx.a`.

Multi-thread	C++	Bits	Linker Flags
no	no	64	-Wl
yes	no	64	-Wl
no	yes	64	-Wl
yes	yes	64	-Wl

Where

- `-undefined=malloc` has the side effect of pulling in all libc-style allocator symbols from the library.
- `-undefined` works on a per-object-file level, rather than a per-symbol level, and the c++ and c allocator symbols are in different object files within the library archive. Therefore, you may also need to specify a c++ style allocator such as `_ZdaPv` below.
- `-undefined=_ZdaPv` has the side effect of pulling in all c++ style allocator symbols. It is the c++ mangled name of operator delete[].

To link the correct library, use the full path to the static library. This is more reliable than using the `-l` argument of a compiler.

See [Intel compilers on page Appx-F-479](#) or [Portland Group and NVIDIA HPC SDK compilers on page Appx-F-480](#) for compiler-specific information.

Available checks

The following heap checks are available in **Enable Checks**:

Name	Description
basic	Detect invalid pointers passed to memory functions (malloc, free, ALLOCATE, DEALLOCATE, etc.)
check-funcs	Check the arguments of addition functions (mostly string operations) for invalid pointers.
check-heap	Check for heap corruption, for example, due to writes to invalid memory addresses.
check-fence	Check the end of an allocation has not been overwritten when it is freed.
alloc-blank	Initialize the bytes of new allocations with a known value.
free-blank	Overwrite the bytes of freed memory with a known value.
check-blank	Check to see if space that was blanked when a pointer was allocated or when it was freed has been overwritten. Enables alloc-blank and free-blank.
realloc-copy	Always copy data to a new pointer when reallocating a memory allocation (for example, due to realloc).
free-protect	Protect freed memory where possible (using hardware memory protection) so subsequent read/writes cause a fatal error.

Change settings at run time

You can change most memory debugging settings while your program is running by selecting **Control > Memory Debugging Options**. This means you can enable memory debugging with a minimal set of options when your program starts, set a breakpoint at a place you want to investigate for memory errors, then switch on more settings when the breakpoint is hit.

B8.5 Pointer error detection and validity checking

When you have enabled memory debugging and started debugging, all calls to the allocation and deallocation routines of heap memory will be intercepted and monitored. This allows for automatic monitoring for errors, and for user-driven inspection of pointers.

Library usage errors

If the memory debugging library reports an error, an error message will display. This briefly reports the type of error detected, and gives the options to continue playing the program, or to pause execution.




Figure B8-2 Memory error message

If you choose to pause the program, the line of your code that was being executed when the error was reported will be highlighted.

Often this is enough to debug simple memory errors, such as freeing or dereferencing an unallocated variable, iterating past the end of an array and so on, as the local variables and variables on the current line will provide insight into what is happening.

If the cause of the issue is still not clear, it is possible to examine some of the pointers referenced to see whether they are valid, and which line they were allocated on. This is explained in the following sections.

View pointer details

When memory debugging is enabled, right-click on any of the variables or expressions in the **Evaluate** window and choose **View Pointer Details**. This will display the amount of memory allocated to the pointer, and which part of your code originally allocated and deallocated that memory:




Figure B8-3 Pointer details

Click on any of the stack frames to display the relevant section of your code, so that you can see where the variable was allocated or deallocated.

————— **Note** —————

Only a single stack frame will be displayed if **Store stack backtraces for memory allocations** on the **Memory Debugging Options** window is disabled.

This feature can also be used to check the validity of heap-allocated memory.

————— **Note** —————

Memory allocated on the heap refers to memory allocated by malloc, ALLOCATE, new and so on. A pointer may also point to a local variable, in which case a message will inform you it does not point to data on the heap. This can be useful, since a common error is taking a pointer to a local variable that later goes out of scope.




Figure B8-4 Invalid memory message

This is particularly useful for checking function arguments, and key variables when things seem to be going wrong. Of course, just because memory is valid does not mean it is the same type as you were expecting, or of the same size and dimensions, and so on.

Memory Type/Location

As well as invalid addresses, you can often get an indication of the type and location of the memory being pointed to. The different types are:

- Null pointer.
- Valid heap allocation.

- Fence-post area before the beginning of an allocation.
- Fence-post area beyond the end of an allocation.
- Freed heap allocation.
- Fence-post area before the beginning of a freed allocation.
- Fence-post area beyond the end a freed allocation.
- A valid GPU heap allocation.
- An address on the stack.
- The program's code section (or a shared library).
- The program's data section (or a shared library).
- The program's bss section or Fortran COMMON block (or a shared library).
- The program's executable (or a shared library).
- A memory mapped file.
- High bandwidth memory.

Note

It may only be possible to identify certain memory types with higher levels of memory debugging enabled. See [Memory debugging options on page B8-197](#) for more information.

For more information on fence post checking, see the *Fencepost checking* section.

Cross-process comparison of pointers

Memory debugging has an impact on the **Cross-Process Comparison View** and **Cross-Thread Comparison View**. See [Cross-process and cross-thread comparison on page B4-172](#).

If you are evaluating a pointer variable, the **Cross-Process Comparison View** shows a column with the location of the pointer.

Pointers to locations in heap memory are highlighted in green. Dangling pointers, that is pointers to locations in heap memory that have been deallocated, are shown in red.

The cross-process comparison of pointers helps you to identify:

- Processes with different addresses for the same pointer.
- The location of a pointer (heap, stack, .bss, .data, .text or other locations).
- Processes that have freed a pointer while other processes have not, null pointers, and so on.

If the cross-process comparison shows the value of what is being pointed at when the value of the pointer itself is wanted, then modify the pointer expression. For example, if you see the string that a char* pointer is pointing at when you actually want information concerning the pointer itself, then add (void *) to the beginning of the pointer expression.

Writing beyond an allocated area

Use the **Heap Overflow/Underflow Detection** section on the **Memory Debugging Options** window to detect reads and writes beyond or before an allocated block. Any attempts to read or write to the specified number of pages before or after the block will cause a segmentation violation that stops your program.

Add the guard pages after the block to detect heap overflows, or before to detect heap underflows. The default value of one page will catch most heap overflow errors, but if this does not work a good rule of thumb is to set the number of guard pages according to the size of a row in your largest array.

The exact size of a memory page depends on your operating system, but a typical size is 4 kilobytes. In this case, if a row of your largest array is 64 KiB, then set the number of pages to 64/4=16.

Note

Small overflows/underflows (for example less than 16 bytes) might not be detected. This is a result of maintaining correct memory alignment, and without this vectorized code may crash or generate false positives.

To detect small overflows or underflows, enable fencepost checking. See the *Fencepost checking* section.

Note

Your program will not be stopped at the exact location at which your program wrote beyond the allocated data, it only stops at the next heap consistency check.

On systems with larger page sizes (for example 2MB, 1GB) guard pages should be disabled or used with care as at least two pages will be used per allocation. On most systems you can check the page size with `getconf _PAGE_SIZE`.

Fencepost checking

‘Fence Post’ checking will also be performed when the **Heap Debugging** section on the **Memory Debugging Options** window is not set to **Fast**.

In this mode, an extra portion of memory is allocated at the start and/or end of your allocated block, and a pattern is written into this area.

If your program attempts to write beyond your data, say by a few elements, this will be noticeable. However, your program will not be stopped at the exact location that your program wrote beyond the allocated data, it will only be stopped at the next heap consistency check.

Suppress an error

If your program stops at an error, but you wish to ignore it (for example, if it is in a third party library that you cannot fix), you can select **Suppress memory errors from this line in future**. This will open the **Suppress Memory Errors** window, where you can select which function you want to suppress errors from.

B8.6 Current memory usage

Memory leaks can be a significant problem for software developers. If your program's memory usage grows faster than expected, or continues to grow through its execution, it is possible that memory is being allocated which is not being freed when it is no longer required.

This type of problem is typically difficult to diagnose, and particularly so in a parallel environment.

At any point in your program, select **Tools > Current Memory Usage** to display the currently allocated memory in your program for the currently selected group. For larger process groups, the processes displayed will be the ones that are using the most memory across that process group.




Figure B8-5 Memory usage graphs

To view graphical representations of memory usage, select the **Memory Usage** tab.

The pie chart gives an at-a-glance comparison of the total memory allocated to each process. This gives an indication of the balance of memory allocations. Any one process taking an unusually large amount of memory is identifiable here.

The stacked bar chart on the right is where the most interesting information starts. Each process is represented by a bar, and each bar broken down into blocks of color that represent the total amount of memory allocated by a particular function in your code. Say your program contains a loop that allocates one hundred bytes that is never freed. That is not a lot of memory. But if that loop is executed ten million times, you are looking at a gigabyte of memory being leaked! There are six blocks in total. The first five represent the five functions that allocated the most memory allocated, and the 6th (at the top) represents the rest of the allocated memory, wherever it is from.

As you can see, large allocations show up as large blocks of color. If your program is close to the end, or these grow, they are severe memory leaks.

Typically, if the memory leak does not make it into the top five allocations under any circumstances then it may not be significant. If you are still concerned you can view the data in the table view yourself.

For more information about a block of color, click on the block. This displays detailed information about the memory allocations comprising it in the bottom-left pane. Scanning down this list gives you a good

idea of what size allocations were made, how many, where from, and if the allocation resides in high bandwidth memory. If you double-click on any one of these the **Pointer Details** window will open, showing you exactly where that pointer was allocated in your code.

Note

Only a single stack frame will be displayed if **Store stack backtraces for memory allocations** on the **Memory Debugging Options** window is disabled.

To view the current memory usage in a tabular format, select the **Allocation Table** tab.

The table is split into five columns:

- **Allocated by:** Code location of the stack frame or function allocating memory in your program.
- **Count:** Number of allocations called directly from this location.
- **Total Size:** Total size (in bytes) of allocations directly from this location.
- **Count (including called functions):** Number of allocations from this location. This includes any allocations called indirectly, for example, by calling other functions.
- **Total Size (including called functions):** Total size (in bytes) of allocations from this location, including indirect allocations.

For example, if func1 calls func2 which calls malloc to allocate 50 bytes. An allocation of 50 bytes will be reported against func2 in the **Total Size** column. A cumulative allocation of 50 bytes will also be recorded against both functions func1 and func2 in the **Total Size (including called functions)** column.

Another valuable use of this feature is to play the program for a while, refresh the window, play it for a bit longer, refresh the window, and so on. If you pick the points at which to refresh, for example, after units of work are complete, you can watch as the memory load of the different processes in your job fluctuates and you will see any areas that continue to grow. These are problematic leaks.

Detect leaks when using custom allocators/memory wrappers

Some compilers wrap memory allocations inside many other functions. In this case you might find, for example, that all Fortran 90 allocations are inside the same routine. This can also happen if you have written your own wrapper for memory allocation functions.

In these circumstances you will see one large block in the **Memory Usage** tab. You can mark such functions as *Custom Allocators* to exclude them from the bar chart and table by right-clicking on the function and selecting **Add Custom Allocator**. Memory allocated by a custom allocator is recorded against its caller instead.

For example, if myfunc calls mymalloc and mymalloc is marked as a custom allocator, the allocation will be recorded against myfunc instead. You can edit the list of custom allocators by clicking **Edit Custom Allocators** at the bottom of the window.

B8.7 Memory Statistics

The **Memory Statistics** window (**Tools > Overall Memory Statistics**) shows a total of memory usage across the processes in a program. The processes using the most memory are displayed, along with the mean across all processes in the current group, which is useful for larger process counts.




Figure B8-6 Memory Statistics

The contents and location of the memory allocations themselves are not repeated here. Instead this window displays the total amount of memory allocated and freed since the program began, the current number of allocated bytes, and the number of calls to allocation and free routines.

These can help show if your program is unbalanced, if particular processes are allocating or failing to free memory, and so on. At the end of program execution you can usually expect the total number of calls per process to be similar (depending on how your program divides up work), and memory allocation calls should always be greater than deallocation calls. Anything else indicates serious problems.

If your application is using high bandwidth memory, the charts and tables in this window will be broken down into each type of memory in use.

Chapter B9

Use and write plugins

Plugins are a quick and easy way to preload a library into your program, and define where to set breakpoints and tracepoints. A plugin is an XML file.

Examples include MPI correctness checking libraries, or if you define a library that is preloaded with your program that performs your own monitoring of the program.

Plugins can also be used to display a message when breakpoints are hit, showing, for example, an error message where the message is provided by the library in a variable.

It contains the following sections:

- [B9.1 Supported plugins](#) on page B9-208.
- [B9.2 Install a plugin](#) on page B9-209.
- [B9.3 Use a plugin](#) on page B9-210.
- [B9.4 Write a plugin](#) on page B9-212.
- [B9.5 Plugin reference](#) on page B9-213.

B9.1 Supported plugins

Plugins for two MPI correctness-checking libraries:

- Intel Message Checker, part of the Intel Trace Analyser and Collector (Commercial with free evaluation: <http://software.intel.com/en-us/intel-trace-analyzer/>) version 7.1
- Marmot (Open source: <http://www.hlrs.de/organization/amt/projects/marmot>), support expected in version 2.2 and above.

Plugins for the GNU and LLVM compiler sanitizers:

- Address Sanitizer (also known as ASan)

This is a memory error detector for C/C++ code. It can be used to find various memory-related issues including use after free, buffer overflows, and use after return.

To enable the Address Sanitizer:

1. Compile your program whilst passing the `-fsanitize=address` compiler option to your compiler.
2. Enable the Address Sanitizer plugin in Arm DDT. For information how to enable a plugin, see [Use a plugin on page B9-210](#).

When compiling with GNU 7 you must disable leak detection due to a conflict with `ptrace` and this aspect of the plugin.

To disable leak detection, either:

1. Add the following piece of code into your program:

```
extern "C" int __lsan_is_turned_off() { return 1; }
```

2. Set the `LSAN_OPTIONS` environment variable at runtime, using:

```
LSAN_OPTIONS=detect_leaks=0
```

————— **Note** —————

ASan is not compatible with memory debugging.

- Thread Sanitizer (also known as TSan)

This is a data race detector for C/C++ code. A data race occurs when two different threads attempt to write to the same memory at the same time.

To enable the Thread Sanitizer:

1. Compile your application while you pass the `-fsanitize=thread` compiler option to your compiler.
2. Enable the Thread Sanitizer plugin in Arm DDT. For information how to enable a plugin, see [Use a plugin on page B9-210](#).

————— **Note** —————

TSan is not compatible with memory debugging.

B9.2 Install a plugin

To install a plugin, locate the XML plugin file provided by your application vendor and copy it to:

```
{arm-forge installation-directory}/plugins/
```

It will then be included in the list of available plugins on the **Run** dialog.

Each plugin is an XML file in this directory. These files are usually provided by third-party vendors to enable their application to integrate with Arm DDT.

B9.3 Use a plugin

To activate a plugin, select the checkbox next to it, then run your application.

Plugins can automatically perform one or more of the following actions:

- Load a particular dynamic library into your program
- Pause your program and show a message when a certain event such as a warning or error occurs
- Start extra, optionally hidden MPI processes. See [Write a plugin on page B9-212](#) for more details.
- Set tracepoints that log the variables during an execution.

If one of the plugins you have selected cannot be loaded, check that the program is correctly installed, and that the paths inside the XML plugin file match the installation path of the program.

Example Plugin: MPI History Library

The plugin directory contains a small set of files that make a plugin to log MPI communication.

- `Makefile` - Builds the library and the configuration file for the plugin.
- `README.wrapper` - Details the installation, usage, and limitations.
- `wrapper-config` - Used to create the plugin XML config file. Used to preload the library and set tracepoints to log the correct variables.
- `wrapper-source` - Used to automatically generate the source code for the library which will wrap the original MPI calls.

This plugin is designed to wrap around many of the core MPI functions and seamlessly intercept calls to log information which is then displayed. It is targeted at MPI implementations that use dynamic linking, as this can be supported without relinking the debugged program.

Static MPI implementations can be made to work as well, but this is outside the scope of this version.

This package must be compiled before first use so it is compatible with your MPI version. When compiled it will be listed in the user interface.

To install as a non-root user in your local directory, type:

```
make local
```

To install as root in the `plugins` directory, type:

```
make
```

To enable the plugin, click **Details** to expand the **Plugins** section of the **Run** window. Then select **History v1.0**, and start your job as normal. The library will preload and set default tracepoints.

This plugin records call counts, total sent byte counts, and the arguments used in MPI function calls. Function calls and arguments are displayed (in blue) in the **Input/Output** tab.

The function counts are available in the form of the variable `_MPIHistoryCount_{function}`.

The sent bytes counters are accumulated for most functions, but specifically they are not added for the vector operations such as `MPI_Gatherv`.

These count variables within the processes are available for use, in components such as the **Cross-Process Comparison View**, enabling a check that, for example, the count of `MPI_Barriers` is consistent, or primitive MPI bytes sent profiling information to be discovered.

The library does not record the received bytes, as most MPI receive calls in isolation only contain a maximum number of bytes allowed, rather than bytes received. The MPI status is logged, the `SOURCE` tag therein enables the sending process to be identified.

There is no per-communicator logging in this version.

This version is for demonstration purposes for the tracepoints and plugin features. It could generate excessive logged information, or cause your program to run slowly if it is a heavy communicator.

This library can be easily extended, or its logging can be reduced, by removing the tracepoints from the generated `history.xml` file (stored in `ALLINEA_FORGE_PATH` or `~/.allinea/plugins`). This would make execution considerably faster, but still retain the byte and function counts for the MPI functions.

B9.4 Write a plugin

XML plugin files must be structured like this example:

```
<plugin name="Sample v1.0" description="A sample plugin that demonstrates *(DDT)*'s plugin interface.">
  <preload name="samplelib1" />
  <preload name="samplelib2" />
  <environment name="SUPPRESS_LOG" value="1" />
  <environment name="ANOTHER_VAR" value="some value" />
  <breakpoint location="sample_log" action="log" message_variable="message" />
  <breakpoint location="sample_err" action="message_box" message_variable="message" />
  <extra_control_process hide="last" />
</plugin>
```

Only the surrounding plugin tag is required. All the other tags are optional.

A complete description of each tag can be found in [Plugin reference on page B9-213](#).

Note

If you are interested in providing a plugin as part of your application bundle, Arm can provide you with any assistance you need to get up and running. Contact [Arm support](#) for more information.

B9.5 Plugin reference

Tag	Attribute	Description
plugin	name	The plugin's unique name. This should include the program/library the plugin is for, and its version. This is shown in the Run dialog.
plugin	description	A short snippet of text to describe the purpose of the plugin/program to the user. This is shown in the Run dialog.
preload	name	Preloads a shared library of this name into the user's program. The shared library must be locatable using LD_LIBRARY_PATH, or the OS will not be able to load it.
environment	name	Sets a particular environment variable before running the user's program.
environment	value	The value that this environment variable should be set to.
breakpoint	location	Adds a breakpoint at this location in the code. The location can be in a preloaded shared library (see above). Typically this is a function name, or a fully-qualified C++ namespace and class name. C++ class members must include their signature and be enclosed in single quotes, for example, <code>'MyNamespace::DebugServer:: breakpointOnError(char*)'</code>
breakpoint	action	Only <code>message_box</code> is supported in this release. Other settings will stop the program at the breakpoint but take no action.
breakpoint	message_variable	A <code>char*</code> or <code>const char*</code> variable that contains a message to be shown to the user. Identical messages from different processes will be grouped together before displaying them to the user in a message box.
extra_control_process	hide	Starts one more MPI process than the user requested. The optional <code>hide</code> attribute can be <code>first</code> or <code>last</code> , and will hide the <code>first</code> or <code>last</code> process in <code>MPI_COMM_WORLD</code> from the user. This process will be allowed to execute whenever at least one other MPI process is executing, and messages or breakpoints (see above) occurring in this process will appear to come from all processes at once. This is only necessary for tools such as Marmot that use an extra MPI process to perform various runtime checks on the rest of the MPI program.
tracepoint	location	Similar to <code>breakpointlocation</code> .
tracepoint	variables	A comma-separated list of variables to log on every passing of the tracepoint location.

Chapter B10

CUDA GPU debugging

Arm DDT can be used to debug programs that use NVIDIA CUDA devices. The code running on the GPU is debugged simultaneously with the code on the host CPU.

Arm supports a number of GPU compilers that target CUDA devices:

- NVIDIA's CUDA Compiler
- Cray OpenACC
- NVIDIA HPC (formerly PGI) OpenACC and CUDA Fortran
- IBM XL OpenMP offloading

Arm supports version 8.0 and later of the CUDA toolkits and their drivers for toolkits.

It contains the following sections:

- [B10.1 CUDA licensing](#) on page B10-216.
- [B10.2 Prepare to debug GPU code](#) on page B10-217.
- [B10.3 Launch the program](#) on page B10-218.
- [B10.4 Control GPU threads](#) on page B10-219.
- [B10.5 Examine GPU threads and data](#) on page B10-220.
- [B10.6 GPU devices information](#) on page B10-222.
- [B10.7 Attach to running GPU programs](#) on page B10-223.
- [B10.8 Open GPU core files](#) on page B10-224.
- [B10.9 Known issues and limitations](#) on page B10-225.
- [B10.10 GPU language support](#) on page B10-227.

B10.1 CUDA licensing

To debug CUDA programs with Arm DDT you need a CUDA-enabled license key. This is an additional option. If CUDA is not included with a license, the CUDA options will be grayed-out on the **Run** dialog.

While debugging a CUDA program, an additional process from your license is used for each GPU. An exception to this is that single-process licenses continue to allow single GPU debugging.

Note

To serve a floating CUDA license you must use Arm Licence Server.

B10.2 Prepare to debug GPU code

You might need to add additional compiler command line options to enable GPU debugging.

For NVIDIA's `nvcc` compiler, kernels must be compiled with the `-g -G` flags. This enables generation of information for debuggers in the kernels, and also disables some optimizations that would hinder debugging. To use memory debugging with CUDA, `-cudart shared` must also be passed to `nvcc`.

For other compilers, see [GPU language support on page B10-227](#), and your vendor's own documentation.

Note

OpenCL debugging of GPUs is not supported.

B10.3 Launch the program

To launch a CUDA job, select **CUDA** on the **Run** dialog before you click **Run/Submit**. You can also enable memory debugging for CUDA programs in the **CUDA** section. See [CUDA memory debugging](#) for details.

It is not possible to attach to running CUDA programs if the program has already initialized the driver in some way, for example through having executed any kernel or called any functions from the CUDA library.

For MPI applications it is essential to place all CUDA initialization after the `MPI_Init` call.

B10.4 Control GPU threads

To control GPU threads use the standard play, pause, and breakpoints controls. They are all applicable to GPU kernels.

However, because GPUs have different execution models to CPUs, there are a few behavioral differences that are described below.

CUDA breakpoints

CUDA breakpoints can be set in the same way as other breakpoints. See [Set breakpoints on page B3-121](#).

Breakpoints affect all GPU threads, and cause the program to stop when a thread reaches the breakpoint. Where kernels have similar workload across blocks and grids, threads tend to reach the breakpoint together and the kernel pauses once per set of blocks that are scheduled, that is, the set of threads that fit on the GPU at any one time.

Where kernels have divergent distributions of work across threads, timing may be such that threads within a running kernel hit a breakpoint and pause the kernel. After continuing, more threads within the currently scheduled set of blocks will hit the breakpoint and pause the program again.

To apply breakpoints to individual blocks, warps, or threads, conditional breakpoints can be used. For example using the built-in variables `threadIdx.x` (and `threadIdx.y` or `threadIdx.z` as appropriate) for thread indexes and setting the condition appropriately.

Where a kernel pauses at a breakpoint, the currently selected GPU thread will be changed if the previously selected thread is no longer ‘alive’.

Stepping

The GPU execution model is noticeably different to that of the host CPU. In the context of stepping operations, that is, step in, step over, or step out, there are critical differences to note.

The smallest execution unit on a GPU is a warp, which on current NVIDIA GPUs is 32 threads. All threads in a warp execute in lockstep, which means that you cannot step each thread individually. All active threads in the warp execute step at the same time.

It is not currently possible to step over or step out of inlined GPU functions.

Note

GPU functions are often inlined by the compiler. This can be avoided (dependent on hardware) by specifying the `__noinline__` keyword in your function declaration.

Running and pausing

Click **Play/Continue** to run all GPU threads. It is not possible to run individual blocks, warps, or threads.

Click **Pause** to pause a running kernel. Note that the pause operation is not as quick for GPUs as for regular CPUs.

B10.5 Examine GPU threads and data

When working with GPUs, most of the user interface is unchanged from regular MPI or multithreaded debugging. However, there are a number of enhancements and additional features that have been added to help you understand the state of GPU programs:

Select GPU threads

The **Thread Selector** enables you to select your current GPU thread. The current thread is used for the variable evaluation windows, along with the various GPU stepping operations.



Figure B10-1 GPU Thread Selector

The first entries represent the block index. The subsequent entries represent the 3D thread index inside that block.

Changing the current thread updates the local variables, the evaluations, and the current line displays and source code displays to reflect the change.

The **Thread Selector** is also updated to display the current GPU thread if it changes as a result of any other operation. For example if:

- You change threads by selecting an item in the **Parallel Stack View**.
- A memory error is detected and is attributed to a particular thread.
- The kernel has progressed, and the previously selected thread is no longer present in the device.

The **Thread Selector** also displays the dimensions of the grid and blocks in your program.

It is only possible to inspect/control threads in the set of blocks that are actually loaded in to the GPU. If you try to select a thread that is not currently loaded, a message is displayed.

————— **Note** —————

The **Thread Selector** is only displayed when there is a GPU kernel active.

View GPU thread locations

The **Parallel Stack View** displays the location and number of GPU threads.




Figure B10-2 CUDA threads in the Parallel Stack View

Click an item in the **Parallel Stack View** to select the appropriate GPU thread, update the variable display components accordingly, and move the **Source Code viewer** to the appropriate location.

Hovering over an item in the **Parallel Stack View** also enables you to see which individual GPU thread ranges are at a location, as well as the size of each range.

It is not possible to collect the stack trace for all threads in a timely manner. The stack traces are gathered by collecting one for each thread that has stopped in a unique location.

Kernel Progress View

For a simple kernel that is to calculate an output value for each index in an array, it is not easy to check whether the value at position x in an array has been calculated, or whether the calculating thread has yet to be scheduled.

This contrasts sharply with scalar programming, where if the counter of a (up-)loop exceeds x then the value of index x can be taken as being the final value. If it is difficult to decide whether array data is fresh or stale, then clearly this will be a major issue during debugging.

You can use the **Kernel Progress View**, which is displayed at the bottom of the user interface by default when a kernel is in progress.




Figure B10-3 Kernel Progress View

This view identifies the kernels that are in progress. The number of kernels are identified and grouped by different kernel identifiers across processes. The identifier is the kernel name.

A colored progress bar is used to identify which GPU threads are in progress. The progress bar is a projection onto a straight line of the (potentially) 6-dimensional GPU block and thread indexing system, and is tailored to the sizes of the kernels operating in the program.

When you click in the color highlighted sections of the progress bar, a GPU thread will be selected that matches the click location as closely as possible. Selected GPU threads are colored blue. For deselected GPU threads, the ones that are scheduled are colored green, the unscheduled ones are white.

Source Code viewer

The **Source Code viewer** helps you visualize the program flow through your source code by highlighting lines in the current stack trace. When debugging GPU kernels, it will color highlight lines with GPU threads present and display the GPU threads in a similar manner to that of regular CPU threads and processes. Hovering over a highlighted line in the **Source Code viewer** will display a summary of the GPU threads on that line.

B10.6 GPU devices information

One of the challenges of GPU programming is in discovering device parameters, such as the number of registers, the device type, and whether a device is present.

The **GPU Devices** tab examines the GPUs that are present and in use across a program, and groups the information together scalably for multi-process systems.

Locals	Current Line(s)	Current Stack	GPU Devices
GPU Devices			
Attribute Name		Value	
Ranks 0			
GV100GL-A		2 Devices	
IDs		0-1	
Compute Capability		sm_70	
Number of SMs		80	
Warps per SM		64	
Lanes per Warp		32	
Registers per Lane		256	

Figure B10-4 GPU Devices tab

Note

GPU devices are only listed after initialization.

B10.7 Attach to running GPU programs

You can attach to a running GPU program, and then debug the GPU threads for all devices that support CUDA compute capability 2.0 and above.

For details how to attach to a running job, see [Attach to running programs on page B1-73](#), and select the **Debug CUDA** checkbox on the **Attach** window.

B10.8 Open GPU core files

NVIDIA GPU core files can be opened in exactly the same way as core files generated by CPU code.

See [Open core files on page B1-72](#) for details.

B10.9 Known issues and limitations

Debug multiple GPU processes

CUDA allows debugging of multiple CUDA processes on the same node. However, each process will still attempt to reserve all of the available GPUs for debugging.

This works for the case where a single process debugs all GPUs on a node, but not for multiple processes debugging a single GPU.

A temporary workaround when using Open MPI is to export the following environment variable before starting DDT:

```
ALLINEA_CUDA_DEVICE_VAR=OMPI_COMM_WORLD_LOCAL_RANK
```

This will assign a single device (based on local rank) to each process.

In addition:

- You must select **File > Options > Open MPI (Compatibility)** (**Arm Forge > Preferences** on Mac OS X). (Do not select **Open MPI**).
- The device selected for each process will be the only device visible when enumerating GPUs. This causes manual GPU selection code to stop working (due to changing device IDs, and so on).

Thread control

The focus on thread feature is not supported as it can lock up the GPU. This means that it is not possible to control multiple GPUs in the same process individually.

Pre SM_20 GPUs

For GPUs that have SM type less than `sm_20` (or when code is compiled targeting SM type less than `sm_20`), the following issues might apply.

- GPU code targeting less than SM type `sm_20` will inline all function calls. This can lead to behavior such as not being able to step over/out of subroutines.
- Debugging applications using textures is not supported on GPUs with SM type less than `sm_20`.
- If you are debugging code in device functions that get called by multiple kernels, setting a breakpoint in the device function will insert the breakpoint in only one of the kernels.

Debug multiple GPU processes on Cray

It is not possible to debug multiple CUDA processes on a single node on a Cray machine.

You must run with one process per node.

Notes

- NVIDIA CUDA toolkit version 8.0 and later are supported. We recommend using the most recent version of CUDA toolkit and driver.
- X11 cannot be running on any GPU used for debugging. (Any GPU running X11 will be excluded from device enumeration.)
- You must compile with `-g -G` to enable GPU debugging, otherwise your program will run through the contents of kernels without stopping.
- Debugging 32-bit CUDA code on a 64-bit host system is not supported.
- It is not possible to spot unsuccessful kernel launches or failures. An error code is provided by `getCudaLastError()` in the SDK which you can call in your code to detect this. The debugger cannot check this without resetting it, which is not desirable behavior.
- Device memory allocated via `cudaMalloc()` is not visible outside of the kernel function.
- Not all illegal program behavior can be caught in the debugger, for example, divide-by-zero.
- Device allocations larger than 100 MB on Tesla GPUs, and larger than 32 MB on Fermi GPUs, might not be accessible.

- Breakpoints in divergent code might not behave as expected.
- Debugging applications with multiple CUDA contexts running on the same GPU is not supported.
- If CUDA environment variable `CUDA_VISIBLE_DEVICES <index>` is used to target a particular GPU, make sure no X server is running on any of the GPUs.

———— **Note** ————

Any GPU running X will be excluded from enumeration, which can affect the device IDs.

- CUDA drivers require that programs be debugged in a mode matching their version. If your system is running with a toolkit version lower than the CUDA driver version, you should force DDT to use the correct CUDA version by setting the `ALLINEA_FORCE_CUDA_VERSION` environment variable. For example, if you are using the CUDA 8.0 driver, set `ALLINEA_FORCE_CUDA_VERSION=8.0`. Alternatively, you should consider upgrading your CUDA toolkit to match the CUDA driver.
- If memory debugging and CUDA support are enabled, only threaded memory preloads are available.

B10.10 GPU language support

In addition to the native nvcc compiler, a number of other compilers are supported.

Note

Debugging of OpenCL is not supported on the device.

Cray OpenACC

Cray OpenACC is fully supported. Code pragmas are highlighted, most variables are visible within the device, and stepping and breakpoints in the GPU code are supported. The compiler flag `-g` is required for enabling device (GPU-based) debugging. Do not use `-O0` as this disables use of the GPU and runs the accelerated regions on the CPU.

There are a few known issues:

- It is not possible to track GPU allocations created by the Cray OpenACC compiler as it does not directly call `cudaMalloc`.
- Pointers in accelerator code cannot be dereferenced in CCE 8.0.
- Memory consumption in debugging mode can be considerably higher than regular mode. If issues with memory exhaustion arise, consider using the environment variable `CRAY_ACC_MALLOC_HEAPSIZE` to set total heap size (bytes) used on the device, which can make more memory available to the program.

PGI/NVIDIA HPC OpenACC and CUDA Fortran

Arm DDT supports debugging both the host and CUDA parts of NVIDIA HPC/PGI OpenACC and CUDA Fortran programs when compiled with version 14.4 or later of the PGI compiler. PGI version 20.7 onwards is rebranded as the NVIDIA HPC SDK. Older versions of the PGI compiler support debugging only on the host.

IBM XLC/XLF with offloading OpenMP

We support debugging both the host and CUDA parts of OpenMP programs making use of offloading when compiled with version 13.1.7 or later of the IBM XLC/XLF compilers.

For the best debugging experience of offloading OpenMP regions, the following compiler flags are recommended `-g -O0 -qsmp=omp:noopt -qoffload -qfullpath -qnoinline -Xptxas -O0 -Xl1vm2ptx -nvvm-compile-options=-opt=0`.

Chapter B11

Offline debugging

Offline debugging is when a program is run under the control of the debugger, but without user intervention and without a user interface.

There are many situations where running offline will be useful, for example when access to a machine is not immediately available and may not be available during the working day. The program can run with features such as tracepoints and memory debugging enabled, and will produce a report at the end of the execution.

It contains the following sections:

- [*B11.1 Use offline debugging on page B11-230.*](#)
- [*B11.2 Offline report HTML output on page B11-232.*](#)
- [*B11.3 Offline report plain text output on page B11-235.*](#)
- [*B11.4 Run-time job progress reporting on page B11-236.*](#)

B11.1 Use offline debugging

To use offline debugging, specify the `-offline` argument. Optionally, specify an output filename with the `-output=<filename>` argument. A filename with a `.html` or `.htm` extension will generate an HTML version of the output. In other cases, a plain text report will be generated. If the `-output` argument is not used, an HTML output file will be generated in the current working directory and reports the name of that file upon completion.

```
ddt --offline mpiexec -n 4 myprog arg1 arg2
ddt --offline -o myjob.html mpiexec -n 4 myprog arg1 arg2
ddt --offline -o myjob.txt mpiexec -n 4 myprog arg1 arg2
ddt --offline -o myjob.html --np=4 myprog arg1 arg2
ddt --offline -o myjob.txt --np=4 myprog arg1 arg2
```

Additional arguments can be used to set breakpoints, at which the stack of the stopping processes will be recorded before they are continued. You can also set tracepoints at which variable values will be recorded, and set expressions to be evaluated on every program pause.

Settings from your current configuration file will be taken, unless over-ridden on the command line.

Command line options that are of the most significance for this mode of running are:

- `-session=SESSIONFILE` - run in offline mode using settings saved using **File > Save Session**.
- `-processes=NUMPROCS` or `-n NUMPROCS` - run with NUMPROCS processes.
- `-mem-debug[=(fast/balanced/thorough/off)]` - enable and configure memory debugging.
- `-snapshot-interval=MINUTES` - write a snapshot of the program's stack and variables to the offline log file every MINUTES minutes. See [Run-time job progress reporting on page B11-236](#).
- `-trace-at=LOCATION[,N:M:P],VAR1,VAR2,... [if CONDITION]` - set a tracepoint at location, beginning recording after the N'th visit of each process to the location, and recording every M'th subsequent pass until it has been triggered P times. Record the value of variable VAR2. The if clause allows you to specify a boolean CONDITION that must be satisfied for the tracepoint to trigger.

Example:

```
main.c:22,-:2-,x
```

This will record x every 2nd passage of line 22.

- `-break-at=LOCATION[,N:M:P][if CONDITION]` - set a breakpoint at LOCATION (either line or function), optionally starting after the N'th pass, triggering every M passes and stopping after it has been triggered P times. The if clause allows you to specify a boolean CONDITION that must be satisfied for the breakpoint to trigger. When using the if clause the value of this argument should be quoted. The stack traces of paused processes will be recorded, before the processes are then made to continue, and will contain the variables of one of the processes in the set of processes that have paused.

Examples:

```
--break-at=main
--break-at=main.c:22
--break-at=main.c:22 --break-at=main.c:34
```

- `-evaluate=EXPRESSION[;EXPRESSION2][;...]` - set one or more expressions to be evaluated on every program pause. Multiple expressions should be separated by a semicolon and enclosed in quotes. If shell special characters are present the value of this argument should also be quoted.

Examples:

```
--evaluate=i
--evaluate="i; (*addr) / x"
--evaluate=i --evaluate="i * x"
```

- `-offline-frames=(all/none/n)` - specify how many frames to collect variables for, where n is a positive integer. The default value is all.

Examples:

```
--offline-frames=all  
--offline-frames=none  
--offline-frames=1337
```

The program will run to completion, or to the end of the job.

When errors occur, for example a program crash, the stack back trace of crashing processes is recorded to the offline output file. In offline mode, it is as if you clicked **Continue** if the continue option was available in an equivalent ‘online’ debugging session.

Read a file for standard input

In offline mode, normal redirection syntax can be used to read data from a file as a source for the executable’s standard input.

Examples:

```
cat <input-file> | ddt --offline -o myjob.html ...  
ddt --offline -o myjob.html ... < <input-file>
```

Write a file from standard output

Normal redirection can also be used to write data to a file from the executable’s standard output:

```
ddt --offline -o myjob.html ... > <output-file>
```

B11.2 Offline report HTML output

The output file contains four sections: **Messages**, **Tracepoints**, **Memory Leak Report**, and **Output**. At the end of a job, the four sections of the log output are merged (tracepoint data, error messages, memory leak data, and program output) into one file. If the process is terminated abruptly, for example by the job scheduler, these separate files will remain and the final single HTML report might not be created.

- Note

The **Memory Leak Report** section is only created when memory debugging is enabled.




Figure B11-1 Offline mode HTML output

Timestamps are recorded with the contents in the offline log, and even though the file is neatly organized into four sections, it remains possible to identify ordering of events from the time stamp.

Messages

The **Messages** section contains the following:

- Error messages: for example if memory debugging detects an error, the message and the stack trace at the time of the error will be recorded from each offending process.
- Breakpoints: a message with the stopped processes and another one with the stacks, current stack and locals at this point.
- Additional information: after an error or a breakpoint has stopped execution, an additional information message is recorded. This message could contain the stacks, current stack, and local variables for the stopped processes at this point of the execution.

- The **Stacks** table displays the parallel stacks from all paused processes. Additionally, for every top-most frame the variables (locals and arguments) will be displayed by default. You can use the `-offline-frames` command-line option to display the variables for more frames or none. If `-offline-frames=none` is specified no variables at all will be displayed, instead a **Locals** table will show the variables for the current process. Clicking on a function expands the code snippet

and variables in one go. If the stop was caused by an error or crash, the stack of the responsible thread or process is listed first.

- The **Current Stacks** table shows the stack of the current process.
- The **Locals** table (if `-offline-frames=none`) and the **Variables** column of the **Stacks** table show the variables across the paused processes. The text highlighting scheme is the same as for the local variables in the user interface. The **Locals** table shows the local variables of the current process, whereas the **Variables** column shows the locals for a representative process that triggered the stop in that frame. In either case, a sparkline for each variable shows the distribution of values across the processes.

Tracepoints

The **Tracepoints** section contains the output from tracepoints, similar to that shown in the **Tracepoints** tab in an online debugging session. This includes sparklines displaying the variable distribution.

Memory Leak Report

The **Memory Leak Report** section displays a graphical representation of the largest memory allocations that were not freed by the end of the program.




Figure B11-2 Memory leak report

Each row corresponds to the memory still allocated at the end of a job on a single rank. If multiple MPI ranks are being debugged, only those with the largest number of memory allocations are shown. You can configure the number of MPI ranks shown with `-leak-report-top-ranks=X`.

The memory allocations on each rank are grouped by the source location that allocated them. Each colored segment corresponds to one location, identified in the legend. Clicking on a segment reveals a table of all call paths leading to that location along with detailed information about the individual memory allocations:




Figure B11-3 Memory leak report detail

By default all locations that contribute less than 1% of the total allocations are grouped together into the Other item in the legend.

This limit can be configured by setting the ALLINEA_LEAK_REPORT_MIN_SEGMENT environment variable to a percentage. For example, ALLINEA_LEAK_REPORT_MIN_SEGMENT=0.5 will only group locations with less than 0.5% of the total allocated bytes together.

In addition, only the eight largest locations are shown by default. This can be configured with the -leak-report-top-locations=Y command-line option.

The raw data can be exported to CSV format by clicking the export link.

Useful command line options:

Option	Description
-leak-report-top-ranks=X	Limit the memory leak report to the top X ranks (default 8, implies -mem-debug)
-leak-report-top-locations=Y	Limit the memory leak report to the top Y locations in each rank (default 8, implies -mem-debug)
-leak-report-top-call-paths=Z	Limit the memory leak report to the top Z call paths to each allocating function (default 8, implies -mem-debug)

Output

Output from the program is written to the **Output** section. For most MPIs this will not be identifiable to a particular process, but on those MPIs that do support it, which processes have generated the output will be reported.

Identical output from the **Output** and **Tracepoints** section is, if received in close proximity and order, merged in the output, where this is possible.

B11.3 Offline report plain text output

Unlike the offline HTML report, the plain text report does not separate the tracepoint, breakpoint, memory leak, and program output into separate sections.

Lines in the offline plain text report are identified as messages, standard output, error output, and tracepoints, as described in [Offline report HTML output on page B11-232](#).

For example, a simple report could look like the following:

```
message (0-3): Process stopped at breakpoint in main (hello.c:97).
message (0-3): Stacks
message (0-3): Processes Function
message (0-3): 0-3    main (hello.c:97)
message (0-3): Stack for process 0
message (0-3): #0 main (argc=1, argv=0x7fffffff378, \
    environ=0x7fffffff388) at /home/ddt/examples/hello.c:97
message (0-3): Local variables for process 0 \
    (ranges shown for 0-3)
message (0-3): argc: 1 argv: 0x7fffffff378 beingWatched: 0 \
    dest: 7 environ: 0x7fffffff388 i: 0 message: ",!\312\t" \
    my_rank: 0 (0-3) p: 4 source: 0 status: t2: 0x7ffff7ff7fc0 \
    tables: tag: 50 test: x: 10000 y: 12
```

B11.4 Run-time job progress reporting

In offline mode, you can compile a snapshot of a job, including its stacks and variables, and update the session log with that information. This includes writing the HTML log file, which otherwise is only written when the session has completed.

Snapshots can be triggered periodically via a command-line option, or at any point in the session by sending a signal to the front-end.

Periodic snapshots

Snapshots can be triggered periodically throughout a debugging session with the command-line option `-snapshot-interval=MINUTES`. For example, to log a snapshot every three minutes:

```
ddt --offline -o log.html --snapshot-interval=3 \  
mpiexec -n 8 ./myprog
```

Signal-triggered snapshots

Snapshots can also be triggered by sending a SIGUSR1 signal to the front-end process (called `forge.bin` in process lists), regardless of whether or not the `-snapshot-interval` command-line option was specified. For example, after running the following:

```
ddt --offline -o log.html mpiexec -n 8 ./myprog
```

A snapshot can be triggered by running (in another terminal):

```
# Find PID of DDT front-end:  
pgrep forge.bin  
> 18032  
> 18039  
  
# Use pstree to identify the parent if there are multiple PIDs:  
pstree -p  
  
# Trigger the snapshot:  
kill -SIGUSR1 18032
```

Part C

MAP

Chapter C1

Get started with MAP

Learn how to get started using Arm MAP.

It contains the following sections:

- *C1.1 Welcome page* on page C1-240.
- *C1.2 Express Launch (MAP)* on page C1-242.
- *C1.3 Prepare a program for profiling* on page C1-244.
- *C1.4 Profile a program* on page C1-250.
- *C1.5 remote-exec required by some MPIs (MAP)* on page C1-255.
- *C1.6 Profile a single-process program* on page C1-256.
- *C1.7 Send standard input (MAP)* on page C1-257.
- *C1.8 Start a job in a queue (MAP)* on page C1-258.
- *C1.9 Use custom MPI scripts (MAP)* on page C1-259.
- *C1.10 Start MAP from a job script* on page C1-262.
- *C1.11 Numactl (MAP)* on page C1-263.
- *C1.12 MAP environment variables* on page C1-264.

C1.1 Welcome page

Arm MAP is a source-level profiler that can show how much time was spent on each line of code. To see the source code in Arm MAP, compile your program with the debug flag. For most compilers this is `-g`. Do not use a debug build, you should always keep optimization flags turned on when profiling.

You can also use Arm MAP on programs without debug information. In this case, inlined functions are not shown, and the source code cannot be shown but other features should work as expected.

To start Arm MAP type one of the following shell commands into a terminal window:

```
map
map program_name [arguments]
map <profile-file>
```

Where `<profile-file>` is a profile file generated by a Arm MAP profiling run. It contains the program name and has a `.map` extension.

Notes:

- When starting Arm MAP to examine an existing profile file, a valid license is not needed.
- Unless you are using Express Launch mode (see [Express Launch \(MAP\) on page C1-242](#)), you should not attempt to pipe input directly to Arm MAP. For information about how to achieve the effect of sending input to your program, see [Program input and output on page B5-179](#).

We recommend that you add the `--profile` argument to Arm MAP. This runs without the interactive user interface and saves a `.map` file to the current directory, so is ideal for profiling jobs submitted to a queue.

When started in interactive mode, the *Welcome* page displays:




Figure C1-1 MAP Welcome page

The *Welcome* page enables you to choose the kind of profiling you want to do, for example you can:

- Profile a program.
- Load a profile from a previous run.
- Connect to a remote system and accept a Reverse Connect request.

Note

In Express Launch mode (see [Express Launch \(MAP\)](#) on page C1-242) the *Welcome* page is not shown and you will be brought directly to the **Run** dialog instead. If no valid license is found, the program exits, and a message is shown in the console output.

C1.2 Express Launch (MAP)

All of the Arm Forge products can be launched by typing their name in front of an existing `mpiexec` command:

```
$ map mpiexec -n 256 examples/wave_c 30
```

This startup method is called *Express Launch* and is the simplest way to get started.

The MPI implementations supported by Express Launch are:

- bullx MPI
- Cray X-Series (MPI/SHMEM/CAF)
- Intel MPI
- MPICH 3
- Open MPI (MPI/SHMEM)
- Oracle MPT
- Open MPI (Cray XT/XE/XK)
- Spectrum MPI
- Spectrum MPI (PMIx)
- Cray XT/XE/XK (UPC)

If your MPI is not supported by Express Launch, an error message will display:

```
$ 'Generic' MPI programs cannot be started using Express Launch syntax (launching with an  
mpirun command).  
  
Try this instead:  
  ddt --np=256 ./wave_c 20  
  
Type ddt --help for more information.
```

This is referred to as *Compatibility Mode*, in which the `mpiexec` command is not included and the arguments to `mpiexec` are passed via a `-mpiargs="args here"` parameter.

One advantage of Express Launch mode is that it is easy to modify existing queue submission scripts to run your program under one of the Arm Forge products. This works best for MAP, which gathers data without an interactive user interface (`map -profile`) or Reverse Connect (`map -connect`, see [Reverse Connect on page A4-56](#) for more details) for interactive profiling.

If you cannot use Reverse Connect and want to use interactive profiling from a queue, you might need to configure Arm MAP to generate job submission scripts for you. More details on this can be found in [Start a job in a queue \(MAP\) on page C1-258](#) and [Integration with queuing systems on page Appx-A-445](#).

Run dialog box

In Express Launch mode, the **Run** dialog has a restricted number of options:




Figure C1-2 Express Launch Run dialog box

C1.3 Prepare a program for profiling

In most cases, if your program is already compiled with debugging symbols (-g), you do not need to recompile your program to use it with MAP. However, in some cases it might need to be relinked, as explained in *Linking*.

Typically you should keep optimization flags enabled when profiling (rather than profiling a debug build). This will give more representative results.

The recommended set of compilation flags are:

- Arm Compiler for Linux: -g1 -O3 -fno-inline-functions -fno-optimize-sibling-calls
- Cray: -G2 -O3 -h ipa0
- GNU: -g1 -O3 -fno-inline -fno-optimize-sibling-calls
- IBM: -g -O3 -qnoinline
- Intel: -g1 -O3 -fno-inline -no-ip -no-ipo -fno-omit-frame-pointer -fno-optimize-sibling-calls
- PGI: -g -O3 -Meh_frame -Mnoautoinline
- NVIDIA HPC: -g -O3 -Meh_frame -Mnoautoinline

These flags preserve most performance optimizations whilst enabling file and line number information and maximizing stack trace readability by disabling features that might prevent MAP from obtaining stack traces (such as function inlining and tail call optimization). Minimal debug information is also used to reduce memory usage during profiling.

Debug symbols

If your compiler supports *minimal* debug info, consider using it (for file and line number information only) instead of full debug info. For GCC, Arm Compiler for Linux, and Intel, this means using -g1 for compiling instead of -g.

Although this can cause inlined functions not to be shown in profiles, it can significantly reduce the memory overhead when profiling.

This is particularly relevant for complex C++ codes, memory-constrained compute nodes, or when profiling many processes per node.

You can also use MAP on programs without debug information. In this case inlined functions are not shown, and the source code cannot be shown but other features will work as expected.

For some compilers, it is necessary to explicitly enable frame pointer information to ensure stack traces, particularly when debug information has been disabled. This is normally done with -fno-omit-frame-pointer (or -Meh_frame for PGI and NVIDIA HPC).

Cray compiler

For the Cray compiler we recommend using the -G2 option.

CUDA programs

When compiling CUDA kernels, do not generate debug information for device code (the -G or -device-debug flag) as this can significantly impair runtime performance. Use -lineinfo instead, for example:

```
nvcc device.cu -c -o device.o -g -lineinfo -O3
```

Disable function inlining

While compilers can inline functions, their ability to include sufficient information to reconstruct the original call tree varies between vendors and is not possible if compiling your program with minimal debug info (file & line info only) or without debug info.

To maximize readability of call trees, we recommend that you disable function inlining using the appropriate compiler-specific flags (see *Prepare a program for profiling*).

Note

Some compilers might still inline functions even when they are explicitly instructed not to do so.

There is typically a small performance penalty for disabling function inlining or enabling profiling information.

Disable tail call optimization

A function can return the result of calling another function, for example:

```
int someFunction()
{
    ...
    return otherFunction();
}
```

In this case, the compiler can change the call to `otherFunction` into a jump. This means that, when inside `otherFunction`, the calling function, `someFunction`, no longer appears on the stack.

This optimization, called tail recursion optimization, can be disabled by passing the `-fno-optimize-sibling-calls` argument to most compilers.

Linking

To collect data from your program, MAP uses two small profiler libraries, `map-sampler` and `map-sampler-mpi`. These profiler libraries must be linked with your program. On most systems MAP can do this automatically without any action by you. This is done via the system's `LD_PRELOAD` mechanism, which allows an extra library into your program when starting it.

Note

Although these libraries contain the word 'map' they are used for both Arm Performance Reports and Arm MAP.

This automatic linking when you start your program only works if your program is dynamically-linked. Programs can be dynamically-linked or statically-linked. For MPI programs this is normally determined by your MPI library. Most MPI libraries are configured with `-enable-dynamic` by default, and `mpicc/mpif90` produce dynamically-linked executables that Arm MAP can automatically collect data from.

The `map-sampler-mpi` library is a temporary file that is precompiled and copied or compiled at runtime in the directory `~/.allinea/wrapper`.

If your home directory will not be accessible by all nodes in your cluster you can change where the `map-sampler-mpi` library will be created by altering the shared directory as described in [No shared home directory on page Appx-H-501](#).

The temporary library will be created in the `.allinea/wrapper` subdirectory to this shared directory.

For Cray X-Series Systems the shared directory is not applicable, instead `map-sampler-mpi` is copied into a hidden `.allinea` sub-directory of the current working directory.

If Arm MAP warns you that it could not pre-load the sampler libraries, this often means that your MPI library was not configured with `-enable-dynamic`, or that the `LD_PRELOAD` mechanism is not supported on your platform. You now have three options:

- Try compiling and linking your code dynamically. On most platforms this allows MAP to use the `LD_PRELOAD` mechanism to automatically insert its libraries into your program at runtime.
- Link MAP's `map-sampler` and `map-sampler-mpi` libraries with your program at link time manually.

See *Dynamic linking on Cray X-Series systems*, or *Static linking and Static linking on Cray X-Series systems*.

- Finally, it may be that your system supports dynamic linking but you have a statically-linked MPI. You can try to recompile the MPI implementation with `-enable-dynamic`, or find a dynamically-linked version on your system and recompile your program using that version. This will produce a dynamically-linked program that MAP can automatically collect data from.

Dynamic linking on Cray X-Series systems

If the LD_PRELOAD mechanism is not supported on your Cray X-Series system, you can try to dynamically link your program explicitly with the MAP sampling libraries.

Compile the MPI Wrapper Library

1. Compile the Arm MPI wrapper library for your system using the `make-profiler-libraries -platform=cray -lib-type=shared` command.

————— **Note** —————

Performance Reports also uses this library.

```
user@login:~/myprogram$ make-profiler-libraries --platform=cray --lib-type=shared

Created the libraries in /home/user/myprogram:
libmap-sampler.so      (and .so.1, .so.1.0, .so.1.0.0)
libmap-sampler-mpi.so (and .so.1, .so.1.0, .so.1.0.0)

To instrument a program, add these compiler options:
compilation for use with MAP - not required for Performance Reports:
-g (or '-G2' for native Cray Fortran) (and -O3 etc.)
linking (both MAP and Performance Reports):
-dynamic -L/home/user/myprogram -lmap-sampler-mpi -lmap-sampler -Wl,--eh-frame-hdr

Note: These libraries must be on the same NFS/Lustre/GPFS filesystem as your program.

Before running your program (interactively or from a queue), set
LD_LIBRARY_PATH:
export LD_LIBRARY_PATH=/home/user/myprogram:$LD_LIBRARY_PATH
map ...
or add -Wl,-rpath=/home/user/myprogram when linking your program.
```

2. Link with the Arm MPI wrapper library

```
mpicc -G2 -o hello hello.c -dynamic -L/home/user/myprogram \
-lmap-sampler-mpi -lmap-sampler -Wl,--eh-frame-hdr
```

PGI Compiler

When linking OpenMP programs, you must pass the `-Bdynamic` command line argument to the compiler when linking dynamically.

NVIDIA HPC Compiler

When linking OpenMP programs, you must pass the `-Bdynamic` command line argument to the compiler when linking dynamically.

Static linking

If you compile your program statically, that is your MPI uses a static library or you pass the `-static` option to the compiler, you must explicitly link your program with the Arm sampler and MPI wrapper libraries.

Compile the Arm MPI Wrapper Library

1. Compile the Arm MPI wrapper library for your system using the `make-profiler-libraries -lib-type=static` command.

Note

Performance Reports also uses this library.

```
user@login:~/myprogram$ make-profiler-libraries --lib-type=static

Created the libraries in /home/user/myprogram:
    libmap-sampler.a
    libmap-sampler-pmpi.a

To instrument a program, add these compiler options:
  compilation for use with MAP - not required for Performance Reports:
    -g (and -O3 etc.)
  linking (both MAP and Performance Reports):
    -Wl,@/home/user/myprogram/allinea-profiler.ld ... EXISTING_MPI_LIBRARIES
  If your link line specifies EXISTING_MPI_LIBRARIES (e.g. -lmpi), then
  these must appear after the Arm sampler and MPI wrapper libraries in
  the link line. There's a comprehensive description of the link ordering
  requirements in the 'Preparing a Program for Profiling' section of either
  userguide-forge.pdf or userguide-reports.pdf, located in
  /opt/$(forgeInstallPathRaw)/doc/.
```

2. Link with the Arm MPI wrapper library. The `-Wl,@/home/user/myprogram/allinea-profiler.ld` syntax tells the compiler to look in `/home/user/myprogram/allinea-profiler.ld` for instructions on how to link with the Arm sampler. Usually this is sufficient, but not in all cases. The rest of this section explains how to manually add the Arm sampler to your link line.

PGI

The PGI C runtime static library contains an undefined reference to `__kmpc_fork_call`. This causes compilation to fail when linking `allinea-profiler.ld`. Add `-undefined __wrap__kmpc_fork_call` to your link line before linking to the Arm sampler.

NVIDIA HPC Compiler

The NVIDIA HPC C runtime static library contains an undefined reference to `__kmpc_fork_call`. This causes compilation to fail when linking `allinea-profiler.ld`. Add `-undefined __wrap__kmpc_fork_call` to your link line before linking to the Arm sampler.

Cray

When linking C++ programs you might encounter a conflict between the Cray C++ runtime and the GNU C++ runtime used by the Arm MAP libraries with an error similar to the one below:

```
/opt/cray/cce/8.2.5/CC/x86-64/lib/x86-64/libcray-c++-rts.a(rtti.o)
: In function '__cxa_bad_typeid':
/ptmp/ulib/buildslaves/cfe-82-edition-build/tbs/cfe/lib_src/rtti.c
:1062: multiple definition of '__cxa_bad_typeid'
/opt/gcc/4.4.4/snos/lib64/libstdc++.a(eh_aux_runtime.o):/tmp/peint
/gcc/repackage/4.4.4c/BUILD/snos_objdir/x86_64-suse-linux/libstdc++-v3/libsupc+
+/.//.//.//.//xt-gcc-4.4.4/libstdc++-v3/libsupc++/eh_aux_runtime.cc:46: first defined here
```

You can resolve this conflict by removing `-lstdc++` and `-lgcc_eh` from `allinea-profiler.ld`.

-lpthread

When linking `-Wl,@allinea-profiler.ld` must go before the `-lpthread` command-line argument if present.

Manual Linking

When linking your program you must add the path to the profiler libraries (`-L/path/to/profiler-libraries`), and the libraries themselves (`-lmap-sampler-pmpi`, `-lmap-sampler`).

The MPI wrapper library (`-lmap-sampler-pmpi`) must go:

1. *After* your program's object (`.o`) files.
2. *After* your program's own static libraries, for example `-lmylibrary`.
3. *After* the path to the profiler libraries (`-L/path/to/profiler-libraries`).

4. *Before* the MPI's Fortran wrapper library, if any. For example `-lmpichf`.
5. *Before* the MPI's implementation library usually `-lmpi`.
6. *Before* the sampler library `-lmap-sampler`.

The sampler library `-lmap-sampler` must go:

1. *After* the Arm MPI wrapper library.
2. *After* your program's object (`.o`) files.
3. *After* your program's own static libraries, for example `-lmylibrary`.
4. *After* `-Wl,-undefined,allinea_init_sampler_now`.
5. *After* the path to the profiler libraries `-L/path/to/profiler-libraries`.
6. *Before* `-lstdc++`, `-lgcc_eh`, `-lrt`, `-lpthread`, `-ldl`, `-lm` and `-lc`.

For example:

```
mpicc hello.c -o hello -g -L/users/ddt/arm \
-lmap-sampler-mpmi \
-Wl,--undefined,allinea_init_sampler_now \
-lmap-sampler -lstdc++ -lgcc_eh -lrt \
-Wl,--whole-archive -lpthread \
-Wl,--no-whole-archive \
-Wl,--eh-frame-hdr \
-ldl \
-lm

mpif90 hello.f90 -o hello -g -L/users/ddt/arm \
-lmap-sampler-mpmi \
-Wl,--undefined,allinea_init_sampler_now \
-lmap-sampler -lstdc++ -lgcc_eh -lrt \
-Wl,--whole-archive -lpthread \
-Wl,--no-whole-archive \
-Wl,--eh-frame-hdr \
-ldl \
-lm
```

Static linking on Cray X-Series systems

Compile the MPI Wrapper Library

1. On Cray X-Series systems, you can compile the MPI wrapper library using `make-profiler-libraries -platform=cray -lib-type=static`:

```
Created the libraries in /home/user/myprogram:
libmap-sampler.a
libmap-sampler-mpmi.a

To instrument a program, add these compiler options:
compilation for use with MAP - not required for Performance Reports:
-g (or -G2 for native Cray Fortran) (and -O3 etc.)
linking (both MAP and Performance Reports):
-Wl,@/home/user/myprogram/allinea-profiler.ld ... EXISTING_MPI_LIBRARIES
If your link line specifies EXISTING_MPI_LIBRARIES (e.g. -lmpi), then
these must appear after the Arm sampler and MPI wrapper libraries in
the link line. There's a comprehensive description of the link ordering
requirements in the 'Preparing a Program for Profiling' section of either
userguide-forge.pdf or userguide-reports.pdf, located in
/opt/*(\forgeInstallPathRaw)*/doc/.
```

2. Link with the MPI wrapper library using:

```
cc hello.c -o hello -g -Wl,@allinea-profiler.ld
ftn hello.f90 -o hello -g -Wl,@allinea-profiler.ld
```

Dynamic and static linking on Cray X-Series systems using the modules environment

If your system has the Arm module files installed, you can load them and build your program as usual. See *map-link modules installation on Cray X-Series*.

1. `module load forge` or ensure that `make-profiler-libraries` is in your `PATH`.
2. `module load map-link-static` or `module load map-link-dynamic`.
3. Recompile your program.

map-link modules installation on Cray X-Series

To facilitate dynamic and static linking of user programs with the Arm MPI Wrapper and Sampler libraries Cray X-Series System Administrators can integrate the map-link-dynamic and map-link-static modules into their module system. Templates for these modules are supplied as part of the Arm Forge package.

Copy files `share/modules/cray/map-link-_*` into a dedicated directory on the system.

For each of the two module files copied:

1. Find the line starting with **conflict** and correct the prefix to refer to the location the module files were installed. For example, `arm/map-link-static`. The correct prefix depends on the subdirectory (if any) under the module search path the `map-link-_*` modules were installed.
2. Find the line starting with **set MAP_LIBRARIES_DIRECTORY "NONE"** and replace "NONE" with a user writable directory accessible from the login and compute nodes.

After installed you can verify whether or not the prefix has been set correctly with ‘module avail’, the prefix shown by this command for the `map-link-*` modules should match the prefix set in the ‘conflict’ line of the module sources.

Unsupported user applications

Ensure that the program to be profiled does not set or unset the SIGPROF signal handler. This interferes with the MAP profiling function and can cause it to fail.

We recommend that you do not use MAP to profile programs that contain instructions to perform MPI profiling using MPI wrappers and the MPI standard profiling interface, PMPI. This is because MAP’s own MPI wrappers may conflict with those contained in the program, producing incorrect metrics.

C1.4 Profile a program

When you click **Profile** on the Welcome page, the **Run** window displays.

Run

Application: /home/user/ddt/examples/wave_c Details

Application: /home/user/ddt/examples/wave_c

Arguments:

☐ stdin file:

Working Directory:

Duration: Sampling entire program Details

Metrics Details

Perf Metrics: None selected, click *Details...* to configure. Details...

☐ **CUDA Kernel analysis** Details

☒ **MPI:** 16 processes, Open MPI Details

Number of Processes: 16

☐ Processes per Node 1

Implementation: Open MPI Change...

mpirun arguments

☐ Profile selected ranks: 0-15 100% Select All

☐ **OpenMP** Details

☐ **Submit to Queue** Configure... Parameters...

Environment Variables: none Details

Help Options Run Cancel

Figure C1-3 Run window

The settings are grouped into sections. Click **Details** to expand a section.

Application (MAP)

Application: The full path name to your application. If you specified one on the command line, this is automatically filled in. You can browse and select your application.

Note

Many MPIs have problems working with directory and program names that contain spaces. We recommend that you do not use spaces in directory and file names.

Arguments: (optional) The arguments passed to your application. These are automatically filled if you entered some on the command line.

Note

Avoid using quote characters such as ' and ", as these may be interpreted differently by MAP and your command shell. If you must use these characters but cannot get them to work as expected, contact [Arm support](#).

stdin file: (optional) This enables you to choose a file to be used as the standard input (`stdin`) for your program. Arguments are automatically added to `mpirun` to ensure your input file is used.

Working Directory: (optional) The working directory to use when running your program. If this is blank then MAP's working directory will be used instead.

Duration

Start profiling after: (optional) This enables you to delay profiling by a number of seconds into the run of your program.

Stop profiling after: (optional) This enables you to specify a number of seconds after which the profiler will terminate your program.

Metrics (MAP)

This section allows you to explicitly enable and disable metrics for which data is collected. Metrics are listed alphabetically with their display name and unique metric ID under their associated metric group. Select a metric to see a more detailed description, including the metric's default enabled/disabled state.

Only metrics that can be displayed in the **Metrics** view are listed. Metrics that are unlicensed, unsupported, or always disabled are not listed. Additionally, you cannot disable metrics that are always enabled.

The initial state of enabled/disabled metrics are the combined settings given by the metric XML definitions, the previous user interface session, and those specified with the `-enabled-metrics` and `-disable-metrics` command-line options. The command-line options take preference over the previous user interface session settings, and both take preference over the metric XML definitions settings. Of course, metrics that are always enabled or always disabled cannot be toggled.

All [PAPI metrics on page C11-327](#) displays if installed, and available for enabling/disabling. However, only metrics specified in the `PAPI.config` file are affected.

All [CPU instructions on page C9-302](#) displays if available for enabling/disabling.

MPI (MAP)

Note

If you only have a single process license or have selected none as your MPI Implementation, the MPI options will be missing. The MPI options are not available when in single process mode. See [Profile a single-process program on page C1-256](#) for more details about using a single process.

Number of Processes: The number of processes that you want to profile. MAP supports hundreds of thousands of processes, but this is limited by your license. This option might not be displayed if it is disabled on the **Job Submission** page in the **Options** window.

Number of nodes: This is the number of compute nodes that you want to use to run your program. This option is only displayed for certain MPI implementations, or if it is enabled on the **Job Submission** page in the **Options** window.

Processes per Node: This is the number of MPI processes to run on each compute node. This option is only displayed for certain MPI implementations, or if it is enabled on the **Job Submission** page in the **Options** window.

Implementation: The MPI implementation to use, for example, Open MPI, MPICH 3. Normally the Auto setting will detect the currently loaded MPI module correctly. If you are submitting a job to a queue, the queue settings will also be summarized here. Click **Change** to change the MPI implementation.

Note

The choice of MPI implementation is critical to correctly starting MAP. Your system will normally use one particular MPI implementation. If you are unsure which to pick, try generic, consult your system administrator, or Arm support. A list of settings for common implementations is provided in [MPI distribution notes and known issues on page Appx-E-461](#).

Note

If the MPI command you want is not in your PATH, or you want to use an MPI run command that is not your default one, you can configure this using the **Options** window. See [Optional configuration on page Appx-A-449](#).

mpirun arguments: (optional) The arguments that are passed to mpirun or your equivalent, usually prior to your executable name in normal mpirun usage. You can place machine file arguments, if necessary, here. For most users this field can be left empty.

Note

You should not enter the -np argument because MAP will do this for you.

Profile selected ranks: (optional) If you do not want to profile all the ranks, you can use the -select-ranks command-line option to specify a set of ranks to profile. The ranks should be separated by commas, and intervals are accepted. Example: 5,6-10.

OpenMP (MAP)

Number of OpenMP threads: The number of OpenMP threads to run your program with. This ensures the OMP_NUM_THREADS environment variable is set, but your program may override this by calling OpenMP-specific functions.

Environment variables (MAP)

The optional **Environment Variables** section should contain additional environment variables that should be passed to mpirun or its equivalent. These environment variables may also be passed to your program, depending on which MPI implementation your system uses. Most users will not need to use this section.

Profiling

Click **Run** to start your program, or **Submit** if working through a queue. See [Integration with queuing systems on page Appx-A-445](#). This will compile an MPI wrapper library that can intercept the MPI_INIT call and gather statistics about MPI use in your program. If you have problems, see [MPI wrapper libraries on page Appx-H-511](#). Then Arm MAP brings up the Running window and starts to connect to your processes.

The program runs inside MAP which starts collecting stats on your program through the MPI interface you selected and will allow your MPI implementation to determine which nodes to start which processes on.

MAP collects data for the entire program run by default. Arm's sampling algorithms ensure only a few tens of megabytes are collected even for very long-running jobs. You can stop your program at any time by using the **Stop and Analyze** button. MAP will then collect the data recorded so far, stop your program and end the MPI session before showing you the results. If any processes remain you may have to clean them up manually using the `kill` command, or a command provided with your MPI implementation, but this should not be necessary.




Figure C1-4 Running window

Profiling only part of a program

The easiest way to profile only part of a program in MAP is to use the **Start profiling after** and **Stop profiling after** settings in the **Run** dialog, or the equivalent `-start-after=TIME` and `-stop-after=TIME` command-line options. These options enable you to specify a range of wall-clock time (the job starts at 0 seconds) during which the job should be profiled. When the `-stop-after` time is reached, the job is terminated rather than letting it run to the end.

Alternatively, for more fine-grained control you can choose to start profiling programmatically at a later point by instrumenting your code. To do this you must set the `ALLINEA_SAMPLER_DELAY_START=1` environment variable before starting your program. For MPI programs it is important that this variable is set in the environment of all the MPI processes. It is not necessarily sufficient to simply set the variable

in the environment of the MPI command itself. You must arrange for the variable to be set or exported by your MPI command for all the MPI processes.

You can call `allinea_start_sampling` and `allinea_stop_sampling` once each. That is to say there must be one and only one contiguous sampling region. It is not possible to start, stop, start, stop. You cannot pause or resume sampling using the `allinea_suspend_traces` and `allinea_resume_traces` functions. This will not have the desired effect. You can only delay the start of sampling and stop sampling early.

C

To start sampling programmatically you should `#include "mapsampler_api.h"` and call the `allinea_start_sampling` function. You must point your C compiler at the MAP include directory, by passing the arguments `-I <install root>/map/wrapper` and also link with the MAP sampler library, by passing the arguments `-L <install root>/lib/64 -lmap-sampler`.

To stop sampling programmatically call the `allinea_stop_sampling`

Fortran

To start sampling programmatically call the `ALLINEA_START_SAMPLING` subroutine. You must also link with the MAP sampler library, for example by passing the arguments `-L <install root>/lib/64 -lmap-sampler`.

To stop sampling programmatically call the `ALLINEA_STOP_SAMPLING` subroutine.

C1.5 remote-exec required by some MPIs (MAP)

When using *SGI MPT*, or the MPMD variants of *MPICH 3* or *Intel MPI*, you can use `mpirun` to start all the processes, then attach to them while they are inside `MPI_Init`.

This method is often faster than the generic method, but requires the `remote-exec` facility to be correctly configured if processes are being launched on a remote machine. For more information on `remote-exec`, see [Connecting to compute nodes and remote programs \(remote-exec\)](#) on page Appx-A-448.

Note

If MAP is running in the background, for example using `map &`, this process might get stuck. Some SSH versions cause this behavior when asking for a password. If this happens to you, go to the terminal and use the `fg` or similar command to make MAP a foreground process, or run MAP again, without using `&`. If MAP cannot find a password-free way to access the cluster nodes, you will not be able to use the specialized startup options. Instead, you can use *generic*, although startup might be slower for large numbers of processes.

C1.6 Profile a single-process program

If you have a single-process license you will immediately see the **Run** dialog that is appropriate to run a single-process program.

If you have a multi-process license you can clear the **MPI** checkbox to run a single-process program.

Run

Application: /home/user/ddt/examples/simple busy Details

Application: /home/user/ddt/examples/simple ⌵ 📁

Arguments: busy ⌵

☐ **stdin file:** ⌵ 📁

Working Directory: ⌵ 📁

Duration: Sampling entire program Details

Metrics Details

Perf Metrics: None selected, click *Details...* to configure. Details...

☐ **CUDA Kernel analysis** Details

☐ **OpenMP** Details

☐ **Submit to Queue** Configure... Parameters...

Environment Variables: none Details

Help Options Run Cancel

Figure C1-5 Single-process Run dialog

1. Type the full file path to your application, or browse and select your application.
2. If required, type the arguments to pass to your program.
3. If required, select the **OpenMP** checkbox and select the **Number of OpenMP threads** to start your program with.
4. Click **Run** to start your program.

C1.7 Send standard input (MAP)

You can use the **stdin** field in the **Run** window to choose the file to use as the standard input (stdin) for your program. Arguments will be automatically added to mpirun to ensure your input file is used.

Alternatively, you can enter the arguments directly in the **mpirun Arguments** field. For example, if using MPI directly from the command-line you would normally use an option to the mpirun such as - **stdin filename**, then you can add the same options to the **mpirun Arguments** field when you start your session in the **Run** window.

It is also possible to enter input during a session. To do this, start your program as normal, then open the **Input/Output** tab. Type the input you want to send.

Click **More** to send input from a file, or send an EOF character.




Figure C1-6 Sending input

C1.8 Start a job in a queue (MAP)

If MAP is configured to integrate with a queue/batch environment, you can use MAP to submit your job directly from the user interface.

For details see [Integration with queuing systems on page Appx-A-445](#).

In this case, a **Submit** button is displayed on the **Run** window, instead of a **Run** button. When you click **Submit** on the **Run** window the queue status is displayed until your job starts. MAP will execute the display command every second and show you the standard output. If your queue display is graphical or interactive you cannot use it here.

If your job does not start or you decide not to run it, click **Cancel Job**. If the regular expression you entered for getting the job id is invalid, or if an error is reported, MAP will not be able to remove your job from the queue. In this case we strongly recommend that you check the job has been removed before submitting another as it is possible for a forgotten job to execute on the cluster and either waste resources or interfere with other profiling sessions.

After the sampling (program run) phase is complete, the analysis phase starts, collecting and processing the distinct samples. This can be a lengthy process depending on the size of the program. For very large programs, it could take up to 10 or 20 minutes.

You must ensure that your job does not hit its queue limits during the analysis process. Set the job time large enough to cover both the sampling and analysis phases.

MAP also requires extra memory, both in the sampling and in the analysis phases. If these fail and your program alone approaches one of these limits, you might need to run with fewer processes per node or a smaller data set, to generate a complete set of data.

When your job is running, it connects to MAP and you can profile it.

C1.9 Use custom MPI scripts (MAP)

On some systems, a custom mpirun replacement is used to start jobs, such as mpiexec. MAP typically uses whatever the default for your MPI implementation is, so for Open MPI it would look for mpirun and not mpiexec, for SLURM it would use srun, and so on. Here we explain how to configure MAP to use a custom mpirun command for job startup.

MAP supports two ways that you can use to start jobs using a custom script.

In the first way, you pass all the arguments on the command-line, like this:

```
mpiexec -n 4 /home/<user>/program/chains.exe /tmp/mydata
```

There are several key variables in this command that MAP can complete for you:

- The number of processes (4 in the above example).
- The name of your program (/home/<user>/program/chains.exe).
- One or more arguments passed to your program (/tmp/mydata).

Everything else, like the name of the command and the format of its arguments, remains constant.

To use a command like this in MAP, you adapt the queue submission system described in [Start a job in a queue \(MAP\)](#) on page C1-258. For this mpiexec example, the settings are :




Figure C1-7 Using custom MPI scripts

As you can see, most of the settings are left blank.

There are some differences between the **Submit** command in MAP and what you would type at the command-line:

- The number of processes is replaced with NUM_PROCS_TAG.
- The name of the program is replaced by the full path to `forge-backend`, used by both DDT and MAP.
- The program arguments are replaced by PROGRAM_ARGUMENTS_TAG.

Note

It is not necessary to specify the program name here. MAP takes care of that during its own startup process. The important thing is to make sure your MPI implementation starts `forge-backend` instead of your program, but with the same options.

In the second way, you start a job using a custom mpirun replacement with a settings file:

```
mpiexec -config /home/<user>/myapp.nodespec
```

Where `myfile.nodespec` might contain something similar to the following:

```
comp00 comp01 comp02 comp03 : /home/<user>/program/chains.exe /tmp/mydata
```

If you specify a template file, MAP can automatically generate simple configuration files like this every time you run your program. For the above example, the template file `myfile.template` would contain the following:

```
comp00 comp01 comp02 comp03 : DDTPATH_TAG/libexec/forge-backend DDT_DEBUGGER_ARGUMENTS_TAG  
PROGRAM_ARGUMENTS_TAG
```

This follows the same replacement rules described above and in detail in [Integration with queuing systems on page Appx-A-445](#).

The settings in the **Options** window for this example might be:




Figure C1-8 Using substitute MPI commands

Note **Submit command** and **Submission template file** in particular. MAP creates a new file and appends it to the submit command before executing it. In this case what would actually be executed might be `mpiexec -config /tmp/arm-temp-0112` or similar. Therefore, any argument like `-config` must be last on the line, because MAP will add a file name to the end of the line. Other arguments, if there are any, can come first.

We recommend that you read the section on queue submission, because there are many features described there that might be useful to you if your system uses a non-standard start-up command.

If you do use a non-standard command, contact [Arm support](#).

C1.10 Start MAP from a job script

When debugging it is common to submit runs from inside a debugger. When profiling the usual approach is to run the program offline, producing a profile file that can be inspected later.

To do this, replace your usual program invocation such as:

```
mpirun -n 4 PROGRAM [ARGUMENTS]...
```

With a MAP command such as:

```
map --profile mpirun -n 4 PROGRAM [ARGUMENTS]...
```

```
map --profile --np=4 PROGRAM [ARGUMENTS]...
```

MAP runs without a user interface, gathering data to a .map profile file. Its filename is based on a combination of program name, process count, and timestamp, such as `program_2p_2012-12-19_10-51.map`.

If you are using OpenMP, the value of `OMP_NUM_THREADS` is also included in the name after the process count, such as `program_2p_8t_2014-10-21_12-45.map`.

This default name can be changed with the `-output` argument. To examine this file, either run MAP and select **Load Profile Data File**, or access it directly with the command:

```
map program_2p_2012-12-19_10-51.map
```

Note

When you start MAP to examine an existing profile file, a valid license is not needed.

When running without a user interface, MAP prints a short header and footer to `stderr` with your program's output in between. The `-silent` argument suppresses this additional output so that your program's output is intact.

As an alternative to `-profile`, you can use Reverse Connect (see [Reverse Connect on page A4-56](#)) to connect back to the user interface if you wish to use interactive profiling from inside the queue. So the above example becomes either:

```
map --connect mpirun -n 4 PROGRAM [ARGUMENTS]...
```

or:

```
map --connect --np=4 PROGRAM [ARGUMENTS]...
```

C1.11 Numactl (MAP)

MAP supports launching programs via numactl for MPI programs. It works with or without SLURM. The recommended way to launch via numactl is to use [Express Launch \(MAP\)](#).

```
map mpiexec -n 4 numactl -m 1 ./myprogram.exe  
map srun -n 4 numactl -m 1 ./myprogram.exe
```

It is also possible to launch via numactl using compatibility mode. When using compatibility mode, you must specify the full path to numactl in **Application**. You can find the full path by running:

```
which numactl
```

Enter the name of the required application in **Arguments**, after all arguments to be passed to numactl. It is not possible to pass any more arguments to the parallel job runner when using this mode for launching.

C1.12 MAP environment variables

Variable	Description
ALLINEA_SAMPLER_INTERVAL	Arm MAP takes a sample in each 20ms period, giving it a default sampling rate of 50Hz. This is automatically decreased as the run proceeds to ensure a constant number of samples are taken. See ALLINEA_SAMPLER_NUM_SAMPLES. If your program runs for a very short period of time, you can benefit by decreasing the initial sampling interval. For example, ALLINEA_SAMPLER_INTERVAL=1 sets an initial sampling rate of 1000Hz, or once per millisecond. Higher sampling rates are not supported. Increasing the sampling frequency from the default is not recommended if there are lots of threads or very deep stacks in the target program, because this might not leave sufficient time to complete one sample before the next sample is started. ————— Note ————— Custom values for ALLINEA_SAMPLER_INTERVAL can be overwritten by values set from the combination of ALLINEA_SAMPLER_INTERVAL_PER_THREAD and the <i>expected</i> number of threads (from OMP_NUM_THREADS). For more information, see ALLINEA_SAMPLER_INTERVAL_PER_THREAD. —————
ALLINEA_SAMPLER_INTERVAL_PER_THREAD	To keep overhead low, Arm MAP imposes a minimum sampling interval based on the number of threads. By default, this is 2ms per thread. So for eleven or more threads Arm MAP increases the initial sampling interval to more than 20ms. You can adjust this behavior by setting ALLINEA_SAMPLER_INTERVAL_PER_THREAD to the minimum per-thread sample time in milliseconds. Arm recommends that you avoid lowering this value from the default if there are lots of threads, because this might not leave sufficient time to complete one sample before the next sample is started. ————— Note ————— - Whether OpenMP is enabled or disabled in Arm MAP, the final script or scheduler values set for OMP_NUM_THREADS are used to calculate the sampling interval per thread (ALLINEA_SAMPLER_INTERVAL_PER_THREAD). When configuring your job for submission, check whether your final submission script, scheduler, or the Arm MAP user interface has a default value for OMP_NUM_THREADS. - Custom values for ALLINEA_SAMPLER_INTERVAL are overwritten by values set from the combination of ALLINEA_SAMPLER_INTERVAL_PER_THREAD and the <i>expected</i> number of threads (from OMP_NUM_THREADS). —————
ALLINEA_MPI_WRAPPER	To direct Arm MAP to use a specific wrapper library, set ALLINEA_MPI_WRAPPER=<_path of shared object>. Arm MAP includes several precompiled wrappers. If your MPI is supported, Arm MAP automatically selects and uses the appropriate wrapper. To manually compile a wrapper specifically for your system, set ALLINEA_WRAPPER_COMPILE=1 and MPICC, and run <MAP-installation-directory>/map/wrapper/build_wrapper. This generates the wrapper library ~/.allinea/wrapper/libmap-sampler-_pmpi- <_hostname>.so with symlinks to these files: <pre> ~/.allinea/wrapper/libmap-sampler-_pmpi- <_hostname>.so.1 ~/.allinea/wrapper/libmap-sampler-_pmpi- <_hostname>.so.1.0 ~/.allinea/wrapper/libmap-sampler-_pmpi- <_hostname>.so.1.0.0 </pre>

(continued)

Variable	Description
ALLINEA_WRAPPER_COMPILE	To direct Arm MAP to fall back to creating and compiling a just-in-time wrapper, set <code>ALLINEA_WRAPPER_COMPILE=1</code> . To generate a just-in-time wrapper, an appropriate compiler must be available on the machine where Arm MAP is running, or on the remote host when using remote connect. Arm MAP attempts to auto detect your MPI compiler. However, we recommend that you set the <code>MPICC</code> environment variable to the path to the correct compiler.
ALLINEA_MPIRUN	The path of <code>mpirun</code> , <code>mpiexec</code> , or equivalent. If this is set, it has higher priority than the path set in the user interface and the <code>mpirun</code> found in <code>PATH</code> .
ALLINEA_SAMPLER_NUM_SAMPLES	By default, Arm MAP collects 1000 samples per process. To avoid generating too much data on long runs, the sampling rate is automatically decreased as the run progresses, to ensure that only 1000 evenly spaced samples are stored. You can adjust this by setting <code>ALLINEA_SAMPLER_NUM_SAMPLES=<_positive integer></code> . ————— Note ————— Arm recommends that you leave this value at the default setting. Higher values are not generally beneficial, and add extra memory overheads while running your code. Consider that with 512 processes, the default setting already collects half a million samples over the job, and the effective sampling rate can be very high indeed. —————
ALLINEA_KEEP_OUTPUT_LINES	Specifies the number of lines of program output to record in <code>.map</code> files. Setting to <code>0</code> removes the line limit restriction, although this is not recommended because it can result in very large <code>.map</code> files if the profiled program produces lots of output. See .
ALLINEA_KEEP_OUTPUT_LINE_LENGTH	The maximum line length for program output that is recorded in <code>.map</code> files - lines containing more characters than this limit are truncated. Setting to <code>0</code> removes the line length restriction. However, Arm does not recommend this because it can result in very large <code>.map</code> files if the profiled program produces lots of output per line. See .
ALLINEA_PRESERVE_WRAPPER	To gather data from MPI calls, Arm MAP generates a wrapper to the chosen MPI implementation. See . By default, the generated code and shared objects are deleted when Arm MAP no longer needs them. To prevent Arm MAP from deleting these files, set <code>ALLINEA_PRESERVE_WRAPPER =1</code> . ————— Note ————— If you use remote launch, this variable must be exported in the remote script. See Remote connections dialog on page A4-53. —————
ALLINEA_SAMPLER_NO_TIME_MPI_CALLS	Set this to prevent Arm MAP from timing the time spent in MPI calls.
ALLINEA_SAMPLER_TRY_USE_SMAPS	You can set this to allow Arm MAP to use <code>/proc/[pid]/smaps</code> to gather memory usage data. However, Arm does not recommend doing this because it slows down sampling significantly.
MPICC	To create the MPI wrapper Arm MAP attempts to use <code>MPICC</code> . If that fails, Arm MAP searches for a suitable MPI compiler command in <code>PATH</code> . If the MPI compiler used to compile the target binary is not in <code>PATH</code> (or if there are multiple MPI compilers in <code>PATH</code>), set <code>MPICC</code> .

Chapter C2

Program output

Arm MAP collects and displays output from all processes on the **Input/Output** tab. Both standard output and error are shown.

As the output is shown after the program has completed, there are not the problems with buffering that occur with Arm DDT.

It contains the following sections:

- [C2.1 View standard output and error \(MAP\) on page C2-268.](#)
- [C2.2 Display selected processes on page C2-269.](#)
- [C2.3 Restrict output on page C2-270.](#)
- [C2.4 Save output \(MAP\) on page C2-271.](#)

C2.1 View standard output and error (MAP)

The **Input/Output** tab is at the bottom of the window by default.




Figure C2-1 Standard output in the Input/Output tab

The output can be selected and copied to the X-clipboard.

C2.2 Display selected processes

You can choose whether to view the output for all processes, or just a single process.

————— **Note** —————

Some MPI implementations pipe stdin, stdout, and stderr from every process through mpirun or rank 0.

—————

C2.3 Restrict output

To keep file sizes within reasonable limits `.map` files will contain a summary of the program output limited to the first and last 500 lines (by default).

To change this number, profile with the environment variable `ALLINEA_KEEP_OUTPUT_LINES` set to the preferred total line limit (`ALLINEA_KEEP_OUTPUT_LINES=20` will restrict recorded output to the first 10 lines and last 10 lines).

Setting this to 0 will remove the line limit restriction, although this is not recommended as it may result in very large `.map` files if the profiled program produces lots of output.

The length of each line is similarly restricted to 2048 characters. This can be changed with the environment variable `ALLINEA_KEEP_OUTPUT_LINE_LENGTH`.

As before, setting this to a value of 0 will remove the restriction, although this is not recommended as it risks a large `.map` file if the profiled program emits binary data or very long lines.

C2.4 Save output (MAP)

Right-click on the text to either save it to a file or copy a selection to the clipboard.

Chapter C3

Source code (MAP)

Arm MAP provides source code viewing, editing, and rebuilding features.

It also integrates with most major version control systems, and provides static analysis to automatically detect many classes of common errors.

The source code editing and rebuilding capabilities are not designed for developing programs from scratch, but they are designed to fit into existing profiling sessions that are running on a current executable.

The same capabilities are available for source code whether you are running remotely (using the remote client) or are connected directly to your system.

It contains the following sections:

- [C3.1 View source code \(MAP\) on page C3-274.](#)
- [C3.2 OpenMP programs on page C3-276.](#)
- [C3.3 GPU programs on page C3-278.](#)
- [C3.4 Complex code: code folding on page C3-279.](#)
- [C3.5 Edit source code \(MAP\) on page C3-280.](#)
- [C3.6 Rebuild and restart \(MAP\) on page C3-281.](#)
- [C3.7 Commit changes \(MAP\) on page C3-282.](#)

C3.1 View source code (MAP)

Source and header files found in the executable are reconciled with the files present on the front-end server, and are displayed in a tree view in the **Project Files** tab of the **Project Navigator** window.

Click on the file name to view the source file.

The **Source code viewer** supports automatic color syntax highlighting for C and Fortran.

You can hide functions or subroutines you are not interested in by clicking the ‘-’ glyph next to the first line of the function. This will collapse the function. Click the ‘+’ glyph to expand the function again.




Figure C3-1 Source code viewer

The centre pane shows your source code, annotated with performance information. All the charts you see in Arm MAP share a common horizontal time axis. Your job starts from the left and ends on the right. Next to each line of source code are the sparkline charts. The sparkline charts show how the number of cores executing that line of code varies over time.

What does it mean to say a core is executing a particular line of code? In the **Source code viewer**, Arm MAP uses *inclusive* time, that is, time spent on this line of code or inside functions called by this line. The `main()` function of a single-threaded C or MPI program is typically at 100% for the entire run.

Only ‘interesting’ lines generate charts, that is, lines in which at least 0.1% of the selected time range was spent. In the previous figure, three different lines meet this criterion. The other lines were executed as well, but a negligible amount of time was spent on them.

The first line is a function call to `imbalance`, which runs for 30.4% of the wall-clock time. Looking closer, as well as a large block of green, there is a blue sawtooth pattern. Color identifies different kinds of time. In this single-threaded MPI code, there are three colors:

- **Dark green** Single-threaded computation time. For an MPI program, this is all computation time. For an OpenMP or multi-threaded program, this is the time the main thread was active and no worker threads were active.
- **Blue** MPI communication and waiting time. Time spent inside MPI calls is blue, regardless of whether that is in `MPI_Send` or `MPI_Barrier`. Typically you want to minimize the amount of blue, because the purpose of most codes is parallel *computation*, not communication.
- **Orange** I/O time. All time spent inside known I/O functions, such as reading and writing to the local or networked filesystem, is shown in orange. It is important to minimize the time spent in I/O. On many systems, the complex data storage hierarchy can cause unexpected bottlenecks to occur when scaling a code up. Arm MAP always shows the time from the program’s point of view, so all the underlying complexity is captured and represented as simply as possible.
- **Dark purple** Accelerator. The time the CPU is waiting for the accelerator to return the control to the CPU. Typically you want to minimize this time, to make the CPU work in parallel with the accelerator, using accelerator asynchronous calls.

In the above screenshot, you can see the following:

- First, a function called `imbalance` is called. This function spends around 55% of its time in computation (dark green) and around 45% of it in MPI calls (blue). Hovering the mouse over any

graph shows an exact breakdown of the time spent in it. There is a sawtooth pattern to the time spent in MPI calls that is investigated later.

- Then, the program moves on to a function called `stride`, which spends almost all of its time computing. You will see how to tell whether this time is well spent or not. At the end, you can also see an MPI synchronization. The triangle shape is typical of ranks finishing their work at different times, and spending varying amounts of time waiting at a barrier. Triangles in these charts indicate imbalance.
- Finally, a function called `overlap` is called, which spends almost all of its time in MPI calls.
- The other functions in this snippet of source code were active for <0.1% of the total runtime and can be ignored from a profiling point of view.

Because this example was an MPI program, the height of each block of color represents the percentage of MPI processes that were running each particular line at any moment in time. The sawtooth pattern showing MPI usage indicates that:

- The `imbalance` function goes through several iterations.
- In each iteration, all processes begin computing. There is more green than blue.
- As execution continues, more processes finish computing and transition to waiting in an MPI call. The transition causes the distinctive triangular sawtooth pattern which illustrates a workload imbalance.
- As each sawtooth pattern ends, all ranks finish communicating and the pattern begins again with the next iteration.

This is a classic sign of MPI imbalance. A sawtooth pattern in MAP graphs show that at first, only a few processes change to a different state of execution. Subsequently, more of these processes change until they all synchronize and move on to another state together. These are areas of interest that are worthwhile investigating.

To explore this example in more detail, open the `examples/slow.map` file and look at the `imbalance` function. Identify and investigate why some processes take longer to finish computing than others.

C3.2 OpenMP programs

For an OpenMP or multi-threaded program (or a mixed-mode MPI+OpenMP program) you will also see other colors used.




Figure C3-2 OpenMP Source code view

- **Light green** Multi-threaded computation time. For an OpenMP program this is time inside OpenMP regions. When profiling an OpenMP program you want to see as much light green as possible, because that is the only time you are using all available cores. Time spent in dark green is a potential bottleneck because it is serial code outside an OpenMP region.
- **Light blue** Multi-threaded MPI communication time. This is MPI time spent waiting for MPI communication while inside an OpenMP region or on a pthread. As with the normal blue MPI time you will want to minimize this, but also maximize the amount of multi-threaded computation (light green) that is occurring on the other threads while this MPI communication is taking place.
- **Dark gray** Time inside an OpenMP region in which a core is idle or waiting to synchronize with the other OpenMP threads. In theory, during an OpenMP region all threads are active all of the time. In practice there are significant synchronization overheads involved in setting up parallel regions and synchronizing at barriers. These will be seen as dark gray holes in the otherwise good light green of optimal parallel computation. If you see these there may be an opportunity to improve performance with better loop scheduling or division of the work to be done.
- **Pale blue** Thread synchronization time. Time spent waiting for synchronization between non-OpenMP threads (for example, a pthread_join). Whether this time can be reduced depends on the purpose of the threads in question.

In the screenshot above you can see that 12.8% of the time is spent calling `neighbor.build(atom)` and 75.3% of the time is spent calling `force->compute(atom, neighbor, comm, comm.me)`. The graphs show a mixture of light green indicating an OpenMP region and dark gray indicating OpenMP overhead. OpenMP overhead is the time spent in OpenMP that is not the contents of an OpenMP region (user code). Hovering the mouse over a line will show the exact percentage of time spent in overhead, but visually you can already see that it is significant but not dominant here.

Increasingly, programs use both MPI and OpenMP to parallelize their workloads efficiently. Arm MAP fully and transparently supports this model of working. It is important to note that the graphs are a reflection of the application activity over time:

- A large section of blue in a mixed-mode MPI code means that all the processes in the program were inside MPI calls during this period. Try to reduce these, especially if they have a triangular shape suggesting that some processes were waiting inside MPI while others were still computing.
- A large section of dark green means that all the processes were running single-threaded computations during that period. Avoid this in an MPI+OpenMP code, or you might as well leave out the OpenMP sections altogether.

- Ideally you want to achieve large sections of light green, showing OpenMP regions being effectively used across all processes simultaneously.
- It is possible to call MPI functions *from within* an OpenMP region. Arm MAP only supports this if the main thread (the OpenMP master thread) is the one that makes the MPI calls. In this case, the blue block of MPI time will be smaller, reflecting that *one* OpenMP thread is in an MPI function while the rest are doing something else such as useful computation.

C3.3 GPU programs

In a program using NVIDIA CUDA CPU, time spent waiting for GPU kernels to complete is shown in **Purple**.

When CUDA kernel analysis mode is enabled (see [GPU profiling on page C16-357](#)) Arm MAP will also display data for lines inside CUDA kernels. These graphs show when GPU kernels were active, and for each kernel a breakdown of the different types of warp stalls that occurred on that line. The different types of warp stalls are listed in [Kernel analysis on page C16-359](#). Refer to the tooltip or selected line display ([GPU profiles on page C4-287](#)) to get the exact breakdown, but in general:

- **Purple** Selected. Instructions on this line were being executed on the GPU.
- **Dark Purple** Not selected. This means warps on this line were ready to execute, but that there was no available SM to do the executing.
- **Red** (various shades) Memory operations. Warps on this line were stalled waiting for some memory dependency to be satisfied. Shade of red indicates the type of memory operation.
- **Blue** (various shades) Execution dependency. Warps on this line were stalled until some other action completes. Shade of blue indicates the type of execution dependency.

Note that warp stalls are only reported per-kernel, so it is not possible to obtain the times within a kernel invocation at which different categories of warp stalls occurred. As function calls in CUDA kernels are also automatically fully inlined it is not possible to see warp stalls for ‘time spent inside function(s) on line’ for GPU kernel code.




Figure C3-3 GPU kernel Source code view

In this screenshot a CUDA kernel involving this line was running on this line 13.1% of the time, with most of the warps waiting for a memory access to complete. The colored horizontal range indicates when any kernel observed to be using this source line was on the GPU. The height of the colored region indicates the proportion of sampled warps that were observed to be on this line. See the NVIDIA CUPTI documentation at http://docs.nvidia.com/cuda/cupti/r_main.html#r_pc_sampling for more information on how warps are sampling.

C3.4 Complex code: code folding

Real-world scientific code does not look much like the examples above. It looks more like the following:

```

138      ENDIF
139      mass_flux_x(j,k)=vol_flux_x(j,k)*(density1(donor,k)+limiter)
140
141      sigmam=ABS(mass_flux_x(j,k))/(density1(donor,k)*pre_vol(donor,k))
142      diffuw=energy1(donor,k)-energy1(upwind,k)
143      diffdw=energy1(downwind,k)-energy1(donor,k)
144      wind=1.0_8
145      IF(diffdw.LE.0.0) wind=-1.0_8
146      IF(diffuw*diffdw.GT.0.0) THEN
147         limiter=(1.0_8-sigmam)*wind*MIN(ABS(diffuw),ABS(diffdw)&
148         ,one_by_six*(sigma3*ABS(diffuw)+sigma4*ABS(diffdw)))
149      ELSE
150         limiter=0.0
151      ENDIF
152
153      ener_flux(i,k)=mass_flux_x(i,k)*(energy1(donor,k)+limiter)

```

Figure C3-4 Typical Fortran code in MAP

Here, small amounts of processing are distributed over many lines, and it is difficult to see which parts of the program are responsible for the majority of the resource usage.

To understand the performance of complex blocks of code like this, Arm MAP allows *code folding*. Each logical block of code such as an if-statement or a function call has a small [-] next to it. Clicking this *folds* those lines of code into one and shows one single sparkline for the entire block:

```

68      time = time + dt
69
70      IF(summary_frequency.NE.0) THEN [-]F
71      IF(visit_frequency.NE.0) THEN [-]F
72
73      ! Sometimes there can be a significant start up cost that appears in the first step.
74      ! Sometimes it is due to the number of MPI tasks, or OpenCL kernel compilation.
75      ! On the short test runs, this can skew the results, so should be taken into account
76      ! in recorded run times.
77
78      IF(step.EQ.1) first_step=(timer() - step_time)
79      IF(step.EQ.2) second_step=(timer() - step_time)
80
81      IF(time+q_small.GT.end_time.OR.step.GE.end_step) THEN [-]F
82
83      IF (parallel%boss) THEN [-]F
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208

```

Figure C3-5 Folded Fortran code in MAP

This helps you identify the conditional blocks where most of the processing occurs.

When exploring a new source file, a good way to understand its performance is to use the **View > Fold All** menu item to collapse all the functions in the file to single lines, then scroll through it looking for functions that take an unusual amount of time or show an unusual pattern of I/O or MPI overhead. These can then be expanded to show their most basic blocks, and the largest of these can be expanded again and so on.

C3.5 Edit source code (MAP)

Source code can be edited in the **Source code viewer**. The actions **Undo**, **Redo**, **Cut**, **Copy**, **Paste**, **Select all**, **Go to line**, **Find**, **Find next**, **Find previous**, and **Find in files** are available from the **Edit** menu.

Files can be opened, saved, reverted, and closed from the **File** menu.

Note

Information from Arm MAP will not match edited source files until the changes are saved, the binary is rebuilt, and a new profile is recreated.

If the currently selected file has an associated header or source code file, you can open it by right-clicking in the editor and choosing **Open <filename>.<extension>**. There is a global shortcut on function key F4, or you can use **Edit > Switch Header/Source**.

To edit a source file in an external editor, right-click the editor for the file and choose **Open in external editor**. To change the editor used, or if the file does not open with the default settings, select **File > Options** to open the **Options** window (Arm Forge > Preferences on Mac OS X) then enter the path to the preferred editor in **Editor**, for example `/usr/bin/gedit`.

If a file is edited, a warning will be displayed at the top of the editor.

 **This file has been edited.**

Figure C3-6 File edited warning

This is to warn that the source code shown is not the source that was used to produce the currently executing binary. The source code and line numbers may not match the executing code.

C3.6 Rebuild and restart (MAP)

To configure the build command choose **File > Configure Build**, enter a build command and a directory in which to run the command, then click **Apply**.

To issue the build command choose **File > Build**, or press Ctrl+B (Cmd+B on Mac OS X). When a build is issued the **Build Output** view is shown.

C3.7 Commit changes (MAP)

Changes to source files can be committed using one of Git, Mercurial, or Subversion. To commit changes, choose **File > Commit**, enter a commit message in the **Commit changes** dialog then click **Commit**.

Chapter C4

Selected lines view

The **Selected lines** view enables you to get detailed information about how one or more lines of code are spending their time.

Note

The **Selected lines** view is only available for profiles generated on x86_64 systems.

It contains the following sections:

- [C4.1 Use Selected lines view on page C4-284.](#)
- [C4.2 Limitations on page C4-286.](#)
- [C4.3 GPU profiles on page C4-287.](#)

C4.1 Use Selected lines view

To access the **Selected lines** view, open one of your program's source files in the **Source code viewer** and highlight a line.




Figure C4-1 Selected lines view

The **Selected lines** view is by default shown on the right side of the **Source code viewer**. It updates automatically to show a detailed breakdown of how the selected lines are spending their time.

You can select multiple lines, and Arm MAP will show information for all of the lines together.

You can also select the first line of a collapsed region to see information for the entire code block. See [View source code \(MAP\) on page C3-274](#) for more information.

If you use the **Metrics** view to select a region of time, the **Selected lines** view only shows details for the highlighted region. See [Metrics view on page C9-299](#) for more information.

The **Selected lines** view is divided into two sections.

The first section gives an overview of how much time was spent executing instructions on this line, and how much time was spent in other functions.

If the time spent executing instructions is low, consider using the **Stacks** view, or the **Functions** view to locate functions that are using a lot of CPU time. For more information on the **Stacks** view see [Stacks view on page C5-289](#). For more information on the **Functions** view see [Functions view on page C7-295](#).

The second section details the CPU instruction metrics for the selected line.

These largely show the same information as the global program metrics, described in [CPU instructions on page C9-302](#), but for the selected lines of source code.

Unlike the global program metrics, the line metrics are divided into separate entries for scalar and vector operations, and report time spent in implicit memory accesses.

On some architectures, computational instructions (such as integer or vector operations) are allowed to access memory implicitly. When these types of instruction are used, Arm MAP cannot distinguish between time performing the operation and time accessing memory, and therefore reports time for the instruction in both the computational category and the memory category.

The amount of time spent in ‘explicit’ and ‘implicit’ memory accesses is reported as a footnote to the time spent executing instructions.

Some guidelines are listed here:

- In general, aim for a large proportion of time in vector operations.
- If you see a high proportion of time in scalar operations, try checking to see if your compiler is correctly optimizing for your processor’s SIMD instructions.
- If you see a large amount of time in memory operations, look for ways to more efficiently access memory to improve cache performance.
- If you see a large amount of time in branch operations, look for ways to avoid using conditional logic in your inner loops.

[CPU instructions on page C9-302](#) offers detailed advice on what to look for when optimizing the types of instruction your program is executing.

C4.2 Limitations

Modern superscalar processors use instruction-level parallelism to decode and execute multiple operations in a single cycle, if internal CPU resources are free, and will retire multiple instructions at once, making it appear as if the program counter "jumps" several instructions per cycle.

Current architectures do not allow profilers such as MAP (or Intel VTune, Linux perftools, and others) to efficiently measure which instructions were "invisibly" executed by this instruction-level parallelism. This time is typically allocated to the last instruction executed in the cycle.

Most MAP users will not be affected by this for the following reasons:

- Hot lines in an HPC code typically contain rather more than a single instruction such as `nop`. This makes it unlikely that an entire source line will be executed invisibly via the CPU's instruction-level parallelism.
- Any such lines executed "for free" in parallel with another line by a CPU core will clearly show up as a "gap" in the **Source code view** (but this is unusual).
- Loops with stalls and mispredicted branches still show up highlighting the line containing the problem in all but the most extreme cases.

Key points:

- Expert users: those wanting to use MAP's per-line instruction metrics to investigate detailed CPU performance of a loop or kernel (even down to the assembly level) should be aware that instructions executed in parallel by the CPU will show up with time only assigned to the last one in the batch executed.
- Other users: MAP's statistical instruction-based metrics correlate well with where time is spent in the program and help to find areas for optimization. Feel free to use them as such. If you see lines with very few operations on them (such as a single add or multiply) and no time assigned to them inside your hot loops then these are probably being executed "for free" by the CPU using instruction-level parallelism. The time for each batch of such is assigned to the last instruction completed in the cycle instead.

C4.3 GPU profiles

When CUDA kernel analysis is enabled and the selected line is executed on the GPU, a breakdown of warp stall reasons on this line displays in this view.

For details about how to enable GPU profiling see [Kernel analysis on page C16-359](#).




Figure C4-2 Selected lines view (GPU kernel)

For a description of each of these warp stall reasons, refer to the tooltip for each of the entries or see [Kernel analysis on page C16-359](#).

Chapter C5

Stacks view

The **Stacks** view offers a good top-down view of your program.

It is easy to follow down from the main function to see which code paths took the most time. Each line of the **Stacks** view shows the performance of one line of your source code, including all the functions called by that line.

The sparkline graphs are described in detail in [Source code \(MAP\) on page C3-273](#)

It contains the following section:

- [C5.1 Overview on page C5-290](#).

C5.1 Overview

A short description of the main features of the **Stacks** view.




Figure C5-1 Stacks view

The **Stacks** view shows:

- The first line, `slow`, represents the entire program run.
- Nested below `slow` is a call to the `stride` function. Almost all of this call was in single-threaded compute (dark green).
- The `stride` function spent most of that time on the line `arr_out(i,j)=sqrt(...)` at `slow.f90`, line 114. 19.3% of the entire run was spent executing this line of code.
- The 0.3% MPI time inside `stride` comes from an `MPI_Barrier` on line 121.
- The next major function called from program `slow` is the `overlap` function, shown at the bottom of this figure. A more detailed breakdown is described in [Metrics view on page C9-299](#). This function ran for 30.0% of the total time, and almost all of this was in MPI calls.

When you click on any line of the **Stacks** view, the **Source Code viewer** jumps to that line of code. This makes it a very easy way to navigate and understand the performance of even complex codes.

The percentage MPI time gives an idea as to how well your program is scaling and shows the location of any communication bottlenecks. As discussed in [Source code \(MAP\) on page C3-273](#), a sawtooth pattern in the view represents imbalance between processes or cores.

In the figure above, the `MPI_Send` call inside the `overlap` function has a sawtooth pattern. This means that some processes took significantly longer to finish the call than others, perhaps because they were waiting longer for their receiver to become ready.

Stacks view shows which lines of code spend the most time running, computing, or waiting. As with most places in the user interface, you can hover over a line or chart for a more detailed breakdown.

Chapter C6

OpenMP Regions view

The **OpenMP Regions** view gives insight into the performance of every significant OpenMP region in your program.

Each region can be expanded just like in the **Stacks** view to see the performance of every line beneath it across every core in your job.

The sparkline graphs are described in detail in [Source code \(MAP\) on page C3-273](#).

It contains the following section:

- [C6.1 Overview on page C6-292](#).

C6.1 Overview

A short description of the main features of the **OpenMP Regions** view.

Input/Output	Project Files	OpenMP Stacks	OpenMP Regions	Functions
OpenMP Regions				
Total core time	MPI	Overhead	Function(s) on line	Source
			▼ wave_openmp [program]	
			▼ update [OpenMP region 0]	time_mpi_stop();
			▼ do_math	do_math(j);
43.2%				newval[i] = (2.0 * values[i]) - oldval[i]
0.2%			▼ 1 other	
11.5%		11.5%	<unknown> from /usr/lib/x86_64-linux...	+ (sqrtau * (values[i-1] - (2.0 * values[i]) + values[i+1]));
4.9%				<unknown> from /usr/lib/x86_64-linux-gnu/libgomp.so.1.0.0 (no debug info)
			▼ 2 others	if ((first + j - 1 == 1) (first + j - 1 == tpoints))
<0.1%		<0.1%	▼ GOMP_parallel	allinea_parallel_end(handle, 1);
				REAL(GOMP_parallel,

Figure C6-1 OpenMP Regions view

Note

If you are using MPI and OpenMP, this view summarizes all cores across all nodes, not just one node.

The **OpenMP Regions** view shows:

- The most time-consuming parallel region is in the `update` function at line 207. Clicking on this shows the region in the **Source Code viewer**.
- This region spends most of its time in the `do_math` function. Hovering on the line or clicking on the [-] symbol collapses the view down to show the figures for how much time.
- Of the lines of code inside `do_math`, the `(sqrtau * (values[i-1] ...))` one takes longest with 13.7% of the total core hours across all cores used in the job.
- Calculating `sqrtau = tau * tau` is the next most expensive line, taking 10.5% of the total core hours.
- Only 0.6% of the time in this region is spent on OpenMP overhead, such as starting/synchronizing threads.

From this you can see that the region is optimized for OpenMP usage, that is, it has very low overhead. If you want to improve performance you can look at the calculations on the lines highlighted in conjunction with the CPU instruction metrics, in order to answer the following questions:

- Is the current algorithm bound by computation speed or memory accesses? If the latter, you may be able to improve cache locality with a change to the data structure layout.
- Has the compiler generated optimal vectorized instructions for this routine? Small things can prevent the compiler doing this and you can look at the vectorization report for the routine to understand why.
- Is there another way to do this calculation more efficiently now that you know which parts of it are the most expensive to run?

See [Metrics view on page C9-299](#) for more information on CPU instruction metrics.

Click on any line of the **OpenMP Regions** view to jump to the **Source Code viewer** to show that line of code.

The percentage OpenMP synchronization time gives an idea as to how well your program is scaling to multiple cores and highlights the OpenMP regions that are causing the greatest overhead. Examples of things that cause OpenMP synchronization include:

- Poor load balancing, for example, some threads have more work to do or take longer to do it than others. The amount of synchronization time is the amount of time the fastest-finishing threads wait for the slowest before leaving the region. Modifying the OpenMP chunk size can help with this.
- Too many barriers. All time at an OpenMP barrier is counted as synchronization time. However, `omp atomic` does *not* appear as synchronization time. This is generally implemented as a locking modifier

to CPU instructions. Overuse of the `atomic` operator shows up as large amounts of time spent in *memory accesses* and on lines immediately following an `atomic` pragma.

- Overly fine-grained parallelization. By default OpenMP synchronizes threads at the start and end of each parallel region. There is also some overhead involved in setting up each region. In general, the best performance is achieved when outer loops are parallelized rather than inner loops. This can also be alleviated by using the `no_barrier` OpenMP keyword if appropriate.

When parallelizing with OpenMP it is extremely important to achieve good single-core performance first. If a single CPU core is already bottlenecked on memory bandwidth, splitting the computations across additional cores rarely solves the problem.

Chapter C7

Functions view

The **Functions** view shows a flat profile of the functions in your program.

It contains the following section:

- [C7.1 Overview on page C7-296.](#)

C7.1 Overview

A short description of the main features of the **Functions** view.

Input/Output	Project Files	OpenMP Stacks	OpenMP Regions	Functions	
Functions					
Self time	Total	MPI	Child	Overhead	Function
43.3%	43.3%				do_math
28.1%	88.3%		60.2%	11.6%	update
11.7%	11.7%			11.7%	[OpenMP overhead (no region active)]
11.5%	11.5%			11.5%	<unknown> from /usr/lib/x86_64-linux-gnu/libgomp.so.1.0.0
5.3%	48.6%		43.3%		update [OpenMP region 0]
0.1%	0.1%			0.1%	GOMP_parallel
<0.1%	88.3%		88.3%	11.6%	main

Figure C7-1 Functions view

The first three columns show different measures of the time spent in a given function:

- **Self** shows the time spent in code in the given function itself, but not its callees, that is, not in the other functions called by that function.
- **Total** shows the time spent in code in the given function itself, and all its callees.
- **Child** shows the time spent in the given function's callees only.

You can use the **Functions** view to find costly functions that are called from many different places.

Chapter C8

Project Files view

The **Project Files** view offers an effective way to browse and navigate through a large, unfamiliar code base.

It contains the following section:

- [C8.1 Overview on page C8-298](#).

C8.1 Overview

A short description of the main features of the **Project Files** view.



Figure C8-1 Project Files view

The **Project Files** view distinguishes between **Application Code** and **External Code**. You can choose which folders count as application code by right-clicking. **External Code** is typically system libraries that are hidden at startup.

Chapter C9

Metrics view

This section describes how the **Metrics** view works with the **Source code viewer**, **Stacks** view, and **Project Files** view to help you identify and understand performance problems.

It contains the following sections:

- [C9.1 User interface](#) on page C9-300.
- [C9.2 CPU instructions](#) on page C9-302.
- [C9.3 CPU time](#) on page C9-305.
- [C9.4 I/O](#) on page C9-306.
- [C9.5 Memory](#) on page C9-307.
- [C9.6 MPI calls](#) on page C9-308.
- [C9.7 Detecting MPI imbalance](#) on page C9-309.
- [C9.8 Accelerator](#) on page C9-310.
- [C9.9 Energy](#) on page C9-311.
- [C9.10 Requirements](#) on page C9-312.
- [C9.11 Lustre](#) on page C9-313.
- [C9.12 Zooming](#) on page C9-314.
- [C9.13 View totals across processes and nodes](#) on page C9-315.
- [C9.14 Custom metrics](#) on page C9-316.

C9.1 User interface

In the **Metrics** view, the horizontal axis is wall clock time. By default three metric graphs are shown.




Figure C9-1 Metrics view

The top graph is the **Main thread activity** chart, which uses the same colors and scales as the per-line sparkline graphs described in [Source code \(MAP\) on page C3-273](#). We recommend that you read that section to help understand the **Main thread activity** chart.

For CUDA programs profiled with CUDA kernel analysis mode enabled, a **Warp stall reasons** graph is also displayed. This shows the warp stalls for all CUDA kernels detected in the program, using the same colors and scales as the GPU kernel graphs described in [Kernel analysis on page C16-359](#). We recommend that you read that section to help understand the **Warp stall reasons** graph.

All of the other metric graphs show how single numerical measurements vary across processes and time. Two frequently used graphs are **CPU floating-point** and **Memory usage**. However, there are many other metric graphs available, and they can all be read in the same way. Each vertical slice of a graph shows the distribution of values across processes for that moment in time. The minimum and maximum are clear, and shading is used to display the mean and standard deviation of the distribution.

A thin line means all processes had very similar values. A ‘fat’ shaded region means there is significant imbalance between the processes. Extra details about each moment in time appear below the metric graphs as you move the mouse over them.

The **Metrics** view is at the top of the user interface as it ties all the other views together. Move your mouse across one of the graphs, and a black vertical line appears on every other graph in MAP, showing what was happening at that moment in time.

You can also click and drag to select a region of time within it. All the other views and graphs now redraw themselves to show just what happened during the selected period of time, ignoring everything else. This is a useful way to isolate interesting parts of your program’s execution. To reselect the entire time range double-click or click the **Select All** button.

In this figure, a short region of time has been selected around an interesting sawtooth pattern in time in `MPI_Barrier` because PE 1 is causing delays.



Figure C9-2 MAP with a region of time selected

The first block accepts data in PE order, and is severely delayed. The second block is more flexible, accepting data from any PE, so that PE 1 can compute in parallel. The **Source code viewer** shows how compute and comms are serialized in the first block, but overlap in the second.

There are many more metrics other than those displayed by default. Click the **Metrics** button or right-click on the metric graphs and you can choose one of the following presets or any combination of the metrics beneath them. You can return to the default set of metrics at any time by choosing the **Preset: Default** option.

C9.2 CPU instructions

The following sections describe the CPU instruction metrics available on each platform, x86_64, Armv8-A, Power 8, and Power 9 systems.

Note

Due to differences in processor models, not all metrics are available on all systems.

Tip:

When you select one or more lines of code in the **Source code viewer**, MAP will show a breakdown of the CPU instructions used on those lines. [Selected lines view on page C4-283](#) describes this view in more detail.

CPU instruction metrics available on x86_64 systems

These metrics show the percentage of time that the active cores spent executing different classes of instruction. They are most useful for optimizing single-core and OpenMP performance.

CPU floating-point:

The percentage of time each rank spends in floating-point CPU instructions. This includes vectorized instructions and standard x87 floating-point. High values here suggest CPU-bound areas of the code that are probably functioning as expected.

CPU integer:

The percentage of time each rank spends in integer CPU instructions. This includes vectorized instructions and standard integer operations. High values here suggest CPU-bound areas of the code that are probably functioning as expected.

CPU memory access:

The percentage of time each rank spends in memory access CPU instructions, such as move, load, and store. This also includes vectorized memory access functions. High values here may indicate inefficiently-structured code. Extremely high values (98% and above) almost always indicate cache problems. Typical cache problems include cache misses due to incorrect loop orderings but may also include more subtle features such as false sharing or cache line collisions.

CPU floating-point vector:

The percentage of time each rank spends in vectorized floating-point instructions. Optimized floating-point-based HPC code should spend most of its time running these operations. This metric provides a good check to see whether your compiler is correctly vectorizing hotspots. See [Controlling a program on page Appx-H-507](#) for a list of the instructions considered vectorized.

CPU integer vector:

The percentage of time each rank spends in vectorized and integer instructions. Optimized integer-based HPC code should spend most of its time running these operations. This metric provides a good check to see whether your compiler is correctly vectorizing hotspots. See [Controlling a program on page Appx-H-507](#) for a list of the instructions considered vectorized.

CPU branch:

The percentage of time each rank spends in test and branch-related instructions such as test, cmp and je. An optimized HPC code should not spend much time in branch-related instructions. Typically the only branch hotspots are during MPI calls, in which the MPI layer is checking whether a message has been fully-received or not.

CPU instruction metrics available on Armv8-A systems

Note

These metrics are not available on virtual machines. Linux perf events performance events counters must be accessible on all systems on which the target program runs.

The CPU instruction metrics available on Armv8-A systems are:

Cycles per instruction

The number of CPU cycles to execute an instruction. It is less than 1 when the CPU takes advantage of instruction-level parallelism.

L2 Data cache miss

The percentage of data L2 cache accesses that result in a miss.

Branch mispredicts

The rate of speculatively-executed instructions that do not retire due to incorrect prediction.

Stalled backend cycles

The percentage of cycles where no operation was issued because of the backend, due to a lack of required resources. Data-cache misses can be responsible for this.

Stalled frontend cycles

The percentage of cycles where no operation was issued because of the frontend, due to fetch starvation. Instruction-cache and i-TLB misses can be responsible for this.

CPU instruction metrics available on IBM Power 8 systems

Note

These metrics are not available on virtual machines. Linux perf events performance events counters must be accessible on all systems on which the target program runs.

The CPU instruction metrics available on IBM Power 8 systems are:

Cycles per instruction

The number of CPU cycles to execute an instruction when the thread is not idle. It is less than 1 when the CPU takes advantage of instruction-level parallelism.

CPU FLOPS lower bound

The rate at which floating-point operations completed.

Note

This is a lower bound because the counted value does not account for the length of vector operations.

CPU Memory Accesses

The processor's data cache was reloaded from local, remote, or distant memory due to a demand load.

CPU FLOPS vector lower bound

The rate at which vector floating-point instructions completed.

Note

This is a lower bound because the counted value does not account for the length of vector operations.

CPU branch mispredictions

The rate of mispredicted branch instructions. This counts the number of incorrectly predicted retired branches that are conditional, unconditional, branch and link, return or eret.

CPU instruction metrics available on IBM Power 9 systems

Note

These metrics are not available on virtual machines. Linux perf events performance events counters must be accessible on all systems on which the target program runs.

The CPU instruction metrics available on IBM Power 9 systems are:

Cycles per instruction

The number of CPU cycles to execute an instruction when the thread is not idle. It is less than 1 when the CPU takes advantage of instruction-level parallelism.

L3 cache miss per instruction

The ratio of completed L3 data cache demand loads to instructions.

Branch mispredicts

The rate of branches that were mispredicted.

Stalled backend cycles

The percentage of cycles where no operation was issued because of the backend, due to a lack of required resources. Data-cache misses can be responsible for this.

C9.3 CPU time

These metrics are particularly useful for detecting and diagnosing the impact of other system daemons on your program's run.

CPU time

This is the percentage of time that each thread of your program was able to spend on a core. Together with **Involuntary context switches**, this is a key indicator of oversubscription or interference from system daemons. If this graph is consistently less than 100%, check your core count and CPU affinity settings to make sure one or more cores are not being oversubscribed. If there are regular spikes in this graph, show it to your system administrator and ask for their help in diagnosing the issue.

User-mode CPU time

The percentage of time spent executing instructions in user-mode. This should be close to 100%. Lower values or spikes indicate times in which the program was waiting for a system call to return.

Kernel-mode CPU time

Complements the above graph and shows the percentage of time spent inside system calls to the kernel. This should be very low for most HPC runs. If it is high, show the graph to your system administrator and ask for their help in diagnosing the issue.

Voluntary context switches

The number of times per second that a thread voluntarily slept, for example while waiting for an I/O call to complete. This is normally very low for HPC code.

Involuntary context switches

The number of times per second that a thread was interrupted while computing and switched out for another one. This will happen if the cores are *oversubscribed*, or if other system processes and daemons start running and take CPU resources away from your program. If this graph is consistently high, check your core count and CPU affinity settings to make sure one or more cores are not being oversubscribed. If there are regular spikes in this graph, show it to your system administrator and ask for their help in diagnosing the issue.

System load

The number of active (running or runnable) threads as a percentage of the number of physical CPU cores present in the compute node. This value may exceed 100% if you are using hyperthreading, if the cores are *oversubscribed*, or if other system processes and daemons start running and take CPU resources away from your program. A value consistently less than 100% may indicate your program is not taking full advantage of the CPU resources available on a compute node.

C9.4 I/O

These metrics show the performance of the I/O subsystem from the application's point of view. Correlating these with the I/O time in the *Application Activity* chart helps to diagnose I/O bottlenecks.

POSIX I/O read rate:

The total I/O read rate of the application. This may be greater than **Disk read transfer** if data is read from the cache instead of the storage layer.

POSIX I/O write rate:

The total I/O write rate of the application. This may be greater than **Disk write transfer** if data is written to the cache instead of the storage layer.

Disk read transfer:

The rate at which the application reads data from disk, in bytes per second. This includes data read from network filesystems (such as NFS), but may not include all local I/O due to page caching.

Disk write transfer:

The rate at which the application writes data to disk, in bytes per second. This includes data written to network filesystems.

POSIX read syscall rate:

The rate at which the application invokes the `read` system call. Measured in calls per second, not the amount of data transferred.

POSIX write syscall rate:

The rate at which the application invokes the `write` system call. Measured in calls per second, not the amount of data transferred.

Notes:

- Disk transfer and I/O metrics are not available on Cray X-series systems as the necessary Linux kernel support is not enabled.
- I/O time in the *Application Activity* chart done via direct kernel calls will not be counted.
- Even if your application does not perform I/O, a non-zero amount of I/O will be recorded at the start of profile because of internal I/O performed by Arm MAP.

C9.5 Memory

Here the memory usage of your application is shown in both a per-process and per-node view. Performance degrades severely once all the node memory has been allocated and swap is required. Some HPC systems, notably Crays, will terminate a job that tries to use more than the total node memory available.

Memory usage:

The memory in use by the processes currently being profiled. Memory that is allocated and never used is generally not shown. Only pages actively swapped into RAM by the OS are displayed. This means that you will often see memory usage ramp up as arrays are initialized. The slopes of these ramps can be interesting in themselves.

Note

This means that if you malloc or ALLOCATE a large amount of memory but do not actually use it, the **Memory usage** metric will not increase.

Node memory usage:

The percentage of memory in use by all processes running on the node, including operating system processes and user processes not in the list of selected ranks when specifying a subset of processes to profile. If node memory usage is far below 100% then your code may run more efficiently using fewer processes or a larger problem size. If it is close to or reaches 100% then the combination of your code and other system daemons are exhausting the physical memory of at least one node.

C9.6 MPI calls

A detailed range of metrics offering insight into the performance of the MPI calls in your application. These are all per-process metrics and any imbalance here, as shown by large blocks with sloped means, has serious implications for scalability.

Use these metrics to understand whether the blue areas of the *Application Activity* chart are problematic or are transferring data in an optimal manner. These are all seen from the application's point of view.

An asynchronous call that receives data in the background and completes within a few milliseconds will have a much higher effective transfer rate than the network bandwidth. Making good use of asynchronous calls is a key tool to improve communication performance.

In multithreaded applications, Arm MAP only reports MPI metrics for MPI calls from main threads. If an application uses `MPI_THREAD_SERIALIZED` or `MPI_THREAD_MULTIPLE`, the *Application Activity* chart will show MPI activity, but some regions of the MPI metrics may be empty if the MPI calls are from non-main threads.

MPI call duration:

This metric tracks the time spent in an MPI call so far. PEs waiting at a barrier (MPI blocking sends, reductions, waits and barriers themselves) will ramp up time until finally they escape. Large areas show lots of wasted time and are prime targets for investigation. The PE with no time spent in calls is likely to be the last one to arrive, so should be the focus for any imbalance reduction.

MPI sent/received:

This pair of metrics tracks the number of bytes passed to MPI send/receive functions per second. This is not the same as the speed with which data is transmitted over the network, as that information is not available. This means that an MPI call that receives a large amount of data and completes almost instantly will have an unusually high instantaneous rate.

MPI point-to-point and collective operations:

This pair of metrics tracks the number of point-to-point and collective calls per second. A long shallow period followed by a sudden spike is typical of a late sender. Most processes are spending a long time in one MPI call (very low #calls per second) while one computes. When that one reaches the matching MPI call it completes much faster, causing a sudden spike in the graph.

Note

For more information about the MPI standard definitions for these types of operations, see chapters 3 and 5 in the [MPI Standard \(version 2.1\)](#).

MPI point-to-point and collective bytes:

This pair of metrics tracks the number of bytes passed to MPI send and receive functions per second. This is not the same as the speed with which data is transmitted over the network, as that information is not available. This means that an MPI call that receives a large amount of data and completes almost instantly will have an unusually high instantaneous rate.

Note

(for SHMEM users) MAP shows calls to `shmem_barrier_all` in **MPI collectives**, **MPI calls**, and **MPI call duration**. Metrics for other SHMEM functions are not collected.

C9.7 Detecting MPI imbalance

The **Metrics** view shows the distribution of their value across all processes against time, so any ‘fat’ regions are showing an area of imbalance in this metric. Analyzing imbalance in Arm MAP works like this:

1. Look at the **Metrics** view for any ‘fat’ regions. These represent imbalance in that metric during that region of time. This tells us (A) that there is an imbalance, and (B) which metrics are affected.
2. Click and drag on the **Metrics** view to select the ‘fat’ region, zooming the rest of the controls in to just this period of imbalance.
3. Now the **Stacks** view and the **Source code viewer** show which functions and lines of code were being executed during this imbalance. Are the processes executing different lines of code? Are they executing the same one, but with differing efficiencies? This tells us (C) which lines of code and execution paths are part of the imbalance.
4. Hover the mouse over the fattest areas on the metric graph and watch the minimum and maximum process ranks. This tells us (D) which ranks are most affected by the imbalance.

Now you know (A) whether there is an imbalance and (B) which metrics (CPU, memory, FPU, I/O) it affects. You also know (C) which lines of code and (D) which ranks to look at in more detail.

Often this is more than enough information to understand the immediate cause of the imbalance (for example, late sender, workload imbalance) but for a deeper view you can now switch to Arm DDT and rerun the program with a breakpoint in the affected region of code. Examining the two ranks highlighted as the minimum and maximum by Arm MAP with the full power of an interactive debugger helps get to the root cause of the imbalance behavior.

C9.8 Accelerator

If you have Arm Forge Professional, the NVIDIA CUDA accelerator metrics are enabled on x86_64. Contact Arm Sales for information on how to upgrade.

Note

Accelerator metrics are not available when linking to the static Arm MAP sampler library.

GPU utilization:

Percent of time that the GPU card was in use, that is, one or more kernels are executing on the GPU card. If multiple cards are present in a compute node this value is the mean across all the cards in a compute node. Adversely affected if CUDA kernel analysis mode is enabled (see [Kernel analysis on page C16-359](#)).

GPU memory usage:

The memory allocated from the GPU frame buffer memory as a percentage of the total available GPU frame buffer memory.

C9.9 Energy

The energy metrics are only available with Arm Forge Professional. All metrics are measured per node. If you are running your job on more than one node, Arm MAP shows the minimum, mean, and maximum power consumption of the nodes.

Note

Energy metrics are not available when linking to the static Arm MAP sampler library.

GPU power usage:

The cumulative power consumption of all GPUs on the node, as measured by the NVIDIA on-board sensor. This metric is available if the Accelerator metrics are present.

CPU power usage:

The cumulative power consumption of all CPUs on the node, as measured by the Intel on-board sensor (Intel RAPL).

System power usage:

The power consumption of the node as measured by the Intel Energy Checker or the Cray metrics.

C9.10 Requirements

CPU power measurement requires an Intel CPU with RAPL support, for example Sandy Bridge or newer, and the intel_rapl powercap kernel module to be loaded.

Node power monitoring is implemented through one of two methods: the Arm IPMI energy agent which can read IPMI power sensors, or the Cray HSS energy counters.

For more information on how to install the Arm IPMI energy agent, see [Arm IPMI Energy Agent on page Appx-I-528](#). The Cray HSS energy counters are known to be available on Cray XK6 and XC30 machines.

Accelerator power measurement requires an NVIDIA GPU that supports power monitoring. This can be checked on the command-line with `nvidia-smi -q -d power`. If the reported power values are reported as "N/A", power monitoring is not supported.

C9.11 Lustre

Lustre metrics are enabled if your compute nodes have one or more Lustre filesystems mounted. Lustre metrics are obtained from a Lustre client process running on each node. Therefore, the data presented gives the information gathered on a per-node basis. The data presented is also cumulative over all of the processes run on a node, not only the program being profiled. Therefore, there may be some data reported to be read and written even if the program itself does not perform file I/O through Lustre. However, an assumption is made that the majority of data read and written through the Lustre client will be from an I/O intensive program, not from background processes. This assumption has been observed to be reasonable. For generated program profiles with more than a few megabytes of data read or written, almost all of the data reported in Arm MAP is attributed to the program that is profiled.

The data that is gathered from the Lustre client process is the read and write rate of data to Lustre, as well as a count of some metadata operations. Lustre does not just store pure data, but associates this data with metadata, which describes where data is stored on the parallel file system and how to access it. This metadata is stored separately from data, and needs to be accessed whenever new files are opened, closed, or files are resized. Metadata operations consume time and add to the latency in accessing the data. Therefore, frequent metadata operations can slow down the performance of I/O to Lustre. Arm MAP reports on the total number of metadata operations, as well as the total number of file opens that are encountered by a Lustre client. With the information provided in Arm MAP, you can observe the rate at which data is read and written to Lustre through the Lustre client. You can also identify whether a slow read or write rate can be correlated to a high rate of expensive metadata operations.

Notes:

- For jobs run on multiple nodes, the reported values are the mean across the nodes.
- If you have more than one Lustre filesystem mounted on the compute nodes the values are summed across all Lustre filesystems.
- Metadata metrics are only available with Arm Forge Professional.

Lustre read transfer:

The number of bytes read per second from Lustre.

Lustre write transfer:

The number of bytes written per second to Lustre.

Lustre file opens:

The number of file open operations per second on a Lustre filesystem.

Lustre metadata operations:

The number of metadata operations per second on a Lustre filesystem. Metadata operations include file open, close, and create as well as operations such as readdir, rename, and unlink.

Note

Depending on the circumstances and implementation, ‘file open’ might count as multiple operations. For example, this might happen when a ‘file open’ creates a new file or truncates an existing one.

C9.12 Zooming

To examine a small time range in more detail you can horizontally zoom in the metric graphs by selecting the time-range you want to see then left-clicking inside that selected region.

All the metric graphs will then resize to display that selection in greater detail. This only effects the metric graphs, as the graphs in all the other views, such as the **Source code viewer**, will already have redrawn to display only the selected region when that selection was made.

A right-click on the metric graph zooms out on the metric graph.

This horizontal zoom is limited by the number of samples that were taken and stored in the Arm MAP file. The more you zoom in the more ‘blocky’ the graph becomes.

While you can increase the resolution by instructing Arm MAP to store more samples (see `ALLINEA_SAMPLER_NUM_SAMPLES` and `ALLINEA_SAMPLER_INTERVAL` in [MAP environment variables on page C1-264](#)), this is not recommended as it may significantly impact performance of both the program being profiled and of Arm MAP when displaying the resulting .map file.

You can also zoom in vertically to better see fine-grained variations in a specific metric’s values. The auto-zoom button beneath the metric graphs will cause the graphs to automatically zoom in vertically to fit the data shown in the currently selected time range. As you select new time ranges the graphs automatically zoom again so that you see only the relevant data.

If the automatic zoom is insufficient, you can take manual control of the vertical zoom applied to each individual metric graph. Hold down Ctrl (or CMD on Mac OS X), while either dragging on a metric graph or using the mouse-wheel while hovering over one, to zoom that graph vertically in or out, centered on the current position of the mouse.

A vertically-zoomed metric graph can be panned up or down by either holding down Shift while dragging on a metric graph or just using the mouse-wheel while hovering over it. Manually adjusting either the pan or zoom will disable auto-zoom mode for that graph, click the auto-zoom button again to reapply it.

Action	Usage	Description
Select	Drag a range in a metric graph.	Selects a time range to examine. Many components (but not the metric graphs) will rescale to display data for this time range only.
Reset	Click the Reset icon (under the metric graphs).	Selects the entire time range. All components (including the metric graphs) will rescale to display the entire set of data. All metric graphs will be zoomed out.
Horizontal zoom in	Left click a selection in a metric graph.	Zoom in (horizontally) on the selected time range.
Horizontal zoom out	Right-click a metric graph.	Undo the last horizontal zoom in action.
Vertical zoom in/out	Ctrl + mouse scroll wheel or Ctrl + Drag on a metric graph.	Zoom a single metric graph in or out.
Vertical pan	Mouse scroll wheel or Shift+drag on a metric graph.	Pan a single metric graph up or down.
Automatic vertical zoom	Toggle the Automatic Vertical Zoom icon (under the metric graphs).	Automatically change the zoom of each metric graph to best fit the range of values each graph contains in the selected time range. Manually panning or zooming a graph will disable auto vertical zoom for that graph only.

C9.13 View totals across processes and nodes

The metric graphs show the statistical distribution of the metric across ranks or compute nodes (depending on the metric). So, for example, the **Nodes power usage** metric graph shows the statistical distribution of power usage of the compute nodes.

If you hover the mouse over the name of a metric to the left side of the graph, a tooltip will display additional summary information. The tooltip will show you the *Minimum*, *Maximum*, and *Mean* of the metric across time and ranks or nodes.

For metrics which are not percentages, the tooltip will also show the peak sum across ranks / nodes. For example, the *Maximum ($\sum$ all nodes)* line in the tooltip for **Nodes power usage** shows the peak power usage summed across all compute nodes. This does not include power used by other components, for example, network switches.

For some metrics which are rates (for example, *Lustre read transfer*) the tooltip will also show the cumulative total across all ranks / nodes, for example, *Lustre bytes read ($\sum$ all nodes)*.

C9.14 Custom metrics

Custom metrics can be written to collect and expose additional data (for example, PAPI counters) in the **Metrics** view.

User custom metrics should be installed under the appropriate path in your home directory, for example, `/home/your_user/.allinea/map/metrics`. Custom metrics can also be installed for all users by placing them in the Arm MAP installation directory, for example, `/arm-installation-directory/map/metrics`. If a metric is installed in both locations, the user installation will take priority.

Detailed information about how to write custom metrics can be found in supplementary documentation bundled with the Arm Forge installation in `allinea-metric-plugin-interface.pdf`.

Chapter C10

Configurable Perf metrics

The Perf metrics use the Linux kernel `perf_event_open()` system call to provide additional CPU related metrics available for Arm MAP.

They can be used on any system supported by the Linux `perf` command (also called `perf_event`). These cannot be tracked on typical virtual machines.

Notes

- This feature is available to use only in Arm Forge Professional. Contact Arm Sales for details about how to upgrade.
- You cannot use configurable Perf metrics when collecting PAPI metrics. Use the PAPI installation script to uninstall PAPI metrics if required. Additionally, the following features are disabled when using configurable Perf metrics:
 - CPU instruction metrics on Armv8-A (see *CPU instruction metrics available on Armv8-A systems* in [CPU instructions on page C9-302](#)).
 - CPU instruction metrics on IBM Power systems (see *CPU instruction metrics available on IBM Power 8 systems* and *CPU instruction metrics available on IBM Power 9 systems* in [CPU instructions on page C9-302](#)).

It contains the following sections:

- [C10.1 Performance events on page C10-318](#).
- [C10.2 Permissions on page C10-319](#).
- [C10.3 Probe target hosts on page C10-320](#).
- [C10.4 Specify Perf metrics using the command line on page C10-321](#).
- [C10.5 Specify Perf metrics using a file on page C10-322](#).
- [C10.6 Specify Perf metrics using the Run window on page C10-323](#).
- [C10.7 View events on page C10-324](#).
- [C10.8 Advanced configuration \(MAP\) on page C10-325](#).

C10.1 Performance events

Perf metrics count the rate of one or more performance events that occur in a program. There are some software events that the Linux kernel provides, but most are hardware events tracked by the Performance Monitoring Unit (PMU) of the CPU. *Generalized* hardware events are event name aliases that the Linux kernel identifies.

The quantity and combinations (in some cases) of events that can be simultaneously tracked is limited by the hardware. This feature does not support multiplexing performance events.

If the set of events you requested cannot be tracked at the same time, Arm MAP ends the profiling session immediately with an error message. Try requesting fewer events, or a different combination. See the PMU reference manual for your architecture for more information on incompatible events.

C10.2 Permissions

On some systems, using Perf hardware counters can be restricted by the value of `/proc/sys/kernel/perf_event_paranoid`.

perf_event_paranoid	Description
3	Disable use of Perf events
2	Allow only user-space measurements
1	Allow kernel and user-space measurements
0	Allow access to CPU-specific data, but not raw trace-point samples.
-1	No restrictions

The value of `/proc/sys/kernel/perf_event_paranoid` must be 2 or lower to collect Perf metrics. To set this until the next reboot, run the following commands:

```
sudo sysctl -w kernel.perf_event_paranoid=2
```

To permanently set the paranoid level, add the following line to `/etc/sysctl.conf`:

```
kernel.perf_event_paranoid=2
```

C10.3 Probe target hosts

You must probe an example of a typical host machine before using these metrics. As well as other properties, this collects the CPU ID used to identify the set of potential hardware events for the host, and tests which generalized events are supported.

Ensure that `/proc/sys/kernel/perf_event_paranoid` is set to 2 or lower ([Permissions on page C10-319](#)) before performing the probe.

Note

It is not necessary to probe every potential host, a single compute node in a homogeneous cluster is sufficient.

If your home directory is writable, you can generate a probe file and install it in your config directory by running the following on the intended host:

```
/path/to/forge/bin/forge-probe --install=user
```

If the Forge installation directory is writable, you can generate and install the probe file for the current host with:

```
/path/to/forge/bin/forge-probe --install=global
```

To generate the probe file, but install it manually, execute:

```
/path/to/forge/bin/forge-probe
```

The probe is named `<hostname>_probe.json` and is generated in your current working directory. You must manually copy it to the location specified in the `forge-probe` output. This is typically only necessary when the compute node that you are probing does not have write access to your home file system.

Check that the expected probe files are correctly installed with:

```
/path/to/forge/bin/map --target-host=list
```

This shows something like:

```
0x00000000420f5160    (thunderx2)    e.g. node07.myarmhost.com
GenuineIntel-6-4E      (skylake)       e.g. node01.myintelhost.com
```

If you have exactly one probe file installed, this is automatically assumed to be the target host. If there are multiple installed probe files, you must specify the intended target whenever you use the configurable Perf metrics feature. When using the command line, use the `-target-host` argument. You can specify the intended target CPU ID (such as, `0x00000000420f5160`), family name (such as, `thunderx2`), or a unique substring of the hostname (`myarmhost`).

C10.4 Specify Perf metrics using the command line

You can list available events for a given probed host using:

```
/path/to/forge/bin/map --target-host=myarmhost --perf-metrics=avail
```

———— **Note** ————

Use `list` instead of `avail` to see the events listed on separate lines.

—————

Specify the events you want using a semicolon separated list:

```
/path/to/forge/bin/map --profile --target-host=myarmhost \  
--perf-metrics="cpu-cycles; bus-cycles; instructions" mpirun ...
```

C10.5 Specify Perf metrics using a file

The `-perf-metrics` argument can also take the name of a plain text file:

```
/path/to/forge/bin/map --profile --target-host=myhost \  
--perf-metrics=./myevents.txt mpirun ...
```

`myevents.txt` lists the events to track on separate lines, such as:

```
cpu-cycles  
bus-cycles  
instructions
```

`-perf-metrics=template` outputs a more complex template that lists all possible events with accompanying descriptions. Redirect this output to a file and uncomment the events to track, for example:

```
/path/to/forge/bin/map --target-host=myhost \  
--perf-metrics=template > myevents.txt  
  
vim myevents.txt  
  
/path/to/forge/bin/map --profile --perf-metrics=myevents.txt \  
mpirun ...
```

C10.6 Specify Perf metrics using the Run window

1. Click **Configure Perf metrics** in the **Run** window to open the **Configure Perf Metrics** window.




Figure C10-1 Configure Perf Metrics window

1. Select the target host from the drop-down list of installed hosts (see [Probe target hosts](#) on page C10-320).
2. Double-click an event, or use the arrow buttons to add or remove events from this list.

Note

On the left of the window is the list of Perf events available on the currently selected host, and on the right is the list of events you have selected for tracking.

3. Filter the list of available events by typing a substring of characters in the **Filter** field.

Note

The bottom of the window displays a preview of the section of the command line with the `-perf-metrics=` command, based on the currently selected list of events.

4. From the **File** menu, open the **Perf metric selection** dialog to help you construct a suitable `-perf-metrics=` command line without starting a job, that you can copy into a queue submission script.

C10.7 View events

You can view Perf event counts in the **Metrics** view ([Metrics view](#)) under the **Linux Perf CPU events** preset. All these metrics are reported as events per second with a suitable SI prefix (such as, K, M, G) that is automatically determined.

You can view the total number of events (over the entire program, or just within a selected time range) in the tooltip of the legend.

C10.8 Advanced configuration (MAP)

You can override the default settings used by Arm MAP when making `perf_event_open` calls. Specify one or more flags in a preamble section in square brackets at the start of the perf metrics definition string (either on the command line or at the top of a template file).

```
/path/to/forge/bin/map --profile --target-host=myarmhost \  
--perf-metrics="[optional,noinherit]; instructions; cpu-cycles"
```

Possible options are:

- **[optional]:** Do not abort the program if the requested metrics cannot be collected. Set this if you want to continue profiling even if the no Perf metric results is returned.
- **[noinherit]:** Disable multithreading support (new threads will not inherit the event counter configuration). If you specified events, they are only collected on the main thread (in the case of MPI programs, the thread that called `MPI_thread_init`).
- **[nopinned]:** Disable pinning events on the PMU. If you have specified this, event counting might be multiplexed. We recommend that you do not do this as it interacts poorly with the Arm Forge sampling strategy.
- **[noexclude=kernel]:** Do not exclude kernel events that happen in kernel space. This might require a more permissive `perf_event_paranoid` level.
- **[noexclude=hv]:** Do not exclude events that happen in the hypervisor. This is mainly for PMUs that have built-in support for handling this (such as IBM Power). Most machines require extra support for handling hypervisor measurements.
- **[noexclude=idle]:** Do not exclude counting software events when the CPU is running the idle task. This is only relevant for software events.

Chapter C11

PAPI metrics

The PAPI metrics are additional metrics available for Arm MAP that use the Performance Application Programming Interface (PAPI). They can be used on any system supported by PAPI.

Notes

- Arm Forge Professional is required to make use of this feature. Contact Sales for details on how to upgrade.
- PAPI metrics are collected from the main thread only.

Due to the limitations of PAPI, some metrics might not be available on your system. Arm MAP displays all available metrics. Where metrics are not available, error messages are displayed.

As there is a limit on the type and number of events that can be counted together, PAPI metrics have been split up into small groups of compatible events, so that you can choose which events to view.

To change which group of metrics Arm MAP uses, navigate to the directory indicated on completion of the installation process and modify the `PAPI.config` file.

It contains the following sections:

- [C11.1 PAPI installation](#) on page C11-328.
- [C11.2 PAPI config file](#) on page C11-329.
- [C11.3 PAPI overview metrics](#) on page C11-330.
- [C11.4 PAPI cache misses](#) on page C11-331.
- [C11.5 PAPI branch prediction](#) on page C11-332.
- [C11.6 PAPI floating-point](#) on page C11-333.

C11.1 PAPI installation

To use these metrics, download and install PAPI from <http://icl.cs.utk.edu/papi/index.html>. Then run the metrics installer `papi_install.sh` from the Arm Forge directory.

These metrics can be uninstalled by running `papi_uninstall.sh` from the Arm Forge directory.

C11.2 PAPI config file

When installation has completed, edit the `PAPI.config` file to set your configuration as required.

By default a template `PAPI.config` file is provided in your installation directory at `/arm_installation_directory/map/metrics`. Alternatively, the `PAPI.config` file can be located inside your configuration directory as set by the `ALLINEA_CONFIG_DIR` environment variable. By default your configuration directory is `\$HOME/.allinea`.

To use a `PAPI.config` file located elsewhere, set and export the `ALLINEA_PAPI_CONFIG` environment variable to point to your `PAPI.config` file. For example:

```
export ALLINEA_PAPI_CONFIG=/opt/arm/forge.20.2.1/map/metrics/PAPI.config
```

This needs to be set before running Arm MAP.

If you are using a queuing system, be sure that the `ALLINEA_PAPI_CONFIG` variable is set and exported to all the compute nodes, by adding the `ALLINEA_PAPI_CONFIG` export line to the job script before the Arm MAP command line.

The PAPI config file contains all the metrics sets that can be used. The location of it has been indicated at the end of the installation process. The default metric set is `Overview`. If you want to use another PAPI metrics set, modify the value of the variable called `set` to the desired PAPI metrics set of either `CacheMisses`, `BranchPrediction`, or `FloatingPoint`.

C11.3 PAPI overview metrics

This group of metrics gives a basic overview of the program which has been profiled.

DP FLOPS: The number of double precision floating-point operations performed per second. This uses the PAPI_DP_OPS (double precision floating-point operations) event. What it actually counts differs across architectures. Additionally, there are many caveats surrounding this PAPI preset on Intel architectures. See <http://icl.cs.utk.edu/projects/papi/wiki/PAPITopics:SandyFlops> for more details.

Cycles per instruction: The number of CPU cycles per instruction executed. This uses the PAPI_TOT_CYC (total cycles) and PAPI_TOT_INS (total instructions) events.

L2 data cache misses: The number of L2 data cache misses per second. This uses the PAPI_L2_DCM (L2 data cache misses) event. This metric is only available in this preset if the system has enough hardware counters (5 at least) to collect the required events.

C11.4 PAPI cache misses

This group of metrics focuses on cache misses at various levels of cache.

L1 cache misses: The number of L1 cache misses per second. This uses the PAPI_L1_TCM (L1 total cache misses) event. If this event is unavailable, the L1 data cache misses metric (using the PAPI_L1_DCM event) will be displayed instead.

L2 cache misses: The number of L2 cache misses per second. This uses the PAPI_L2_TCM (L2 total cache misses) event. If this event is unavailable, the L2 data cache misses metric (using the PAPI_L2_DCM event) will be displayed instead.

L3 cache misses: The number of L3 cache misses per second. This uses the PAPI_L3_TCM (L3 total cache misses) event. If this event is unavailable, the L3 data cache misses metric (using the PAPI_L3_DCM event) will be displayed instead.

C11.5 PAPI branch prediction

This group of metrics focuses on branch prediction instructions.

Branch instructions: The number of branch instructions per second. This uses the PAPI_BR_INS (branch instructions) event.

Mispredicted branch instructions: The number of conditional branch instructions that are mispredicted each second. This uses the PAPI_BR_MSP (mispredicted branch instructions) event.

Completed instructions: The number of completed instructions per second. This uses the PAPI_TOT_INS event, and is included to provide context for the above other metrics in this group.

C11.6 PAPI floating-point

This group of metrics focuses on floating-point instructions.

Floating-point scalar instructions: The number of scalar floating-point instructions per second. This uses the PAPI_FP_INS event.

Floating-point vector instructions: The number of vector floating-point instructions per second. This uses the PAPI_VEC_SP (single-precision vectorized instructions) and PAPI_VEC_DP (double-precision vectorized instructions) events. If these events are unavailable, the Vector Instructions metric will be displayed instead.

Vector instructions: The number of vector instructions (floating-point and integer) per second. This uses the PAPI_VEC_INS event. It is only displayed if the events needed for the Floating-point vector instructions metric are not available.

Completed instructions: The number of completed instructions per second. This uses the PAPI_TOT_INS event. It is included to provide context for the above other metrics in this group.

Chapter C12

Main-thread, OpenMP, and Pthread view modes

The percentage values and activity graphs shown alongside the source code and in the **Stacks**, **OpenMP Regions**, and **Functions** views can present information for multithreaded programs in a variety of different ways.

Arm MAP will initially choose the most appropriate view mode for your program. However, in some cases, for example such as when you have written a program to use raw pthreads rather than OpenMP, you might want to change the mode to get a different perspective on how your program is executing multiple threads and using multiple cores. You can switch between view modes using the **View** menu.

It contains the following sections:

- [C12.1 Main thread only mode on page C12-336.](#)
- [C12.2 OpenMP mode on page C12-337.](#)
- [C12.3 Pthread mode on page C12-338.](#)

C12.1 Main thread only mode

In this view mode only the main thread from each process is displayed. The presence of any other thread is ignored. A value of 100% for a function or line means that all the processes' main threads are at that location. This is the best mode to use when exploring single-threaded programs and programs that are unintentionally/indirectly multithreaded (recent implementations of both Open MPI and CUDA will start their own thread).

This is the default mode for all non-OpenMP programs. The **OpenMP Regions** view does not display in this mode.

Note that the **CPU instruction** metric graphs (showing the proportion of time in various classes of CPU instructions, such as integer, floating-point, and vector) are *not* restricted to the main thread when in the **Main thread only** view mode. These metric graphs always represent the data gathered from all the CPU cores.

C12.2 OpenMP mode

This view mode is optimized for interpreting programs where OpenMP is the primary source of multithreaded activity. Percentage values and activity graphs for a line or function indicate the proportion of the available resources that are being used on that line. For serial code on a main thread this is the proportion of processes at that location. For OpenMP code the contribution from each process is further broken down by the proportion of CPU cores running threads that are at that location in the code.

For example, a timeslice of an activity graph showing 50% dark green (serial, main-thread computation) and 50% light green (computation in an OpenMP region) means that half the processes were in serial code and half the processes were in an OpenMP region. Of the processes in an OpenMP region, 100% of the available cores (as determined by the cores per process value, see [Processes and cores window on page C13-339](#)) were being used for OpenMP.

This is the default mode for OpenMP programs. It is only available for programs where Arm MAP detected an OpenMP region.

C12.3 Pthread mode

This view mode is optimized for interpreting programs that make explicit use of pthreads. Percentage values and activity graphs reflect the proportion of CPU cores that are being used out of the maximum number of expected cores per process, see [Processes and cores window](#).

A value of 100% for a function or line means that 100% of the expected number of CPU cores per process were working at that location. The main thread's contribution gets no special attention so activity on the main thread(s) will appear the same height as activity from any other thread.

The advantage of this is that it makes it obvious when the program is not making full use of all the CPU cores available to it. The disadvantage is that it is harder to analyze the performance of the intentionally serial sections of code performed by each process. This is because activity occurring only on one thread per process will be restricted to at most 1/nth of a percentage value or height on an activity graph, where n is the number of cores per process.

This mode is not used by default so must be explicitly selected. It is only available for multithreaded programs.

The **OpenMP Regions** view is not displayed in this mode.

Chapter C13

Processes and cores window

The **Processes and Cores** window enables you to change the number of cores per process.

It contains the following section:

- [C13.1 Overview on page C13-340.](#)

C13.1 Overview

Most modern CPUs support hyperthreading and report multiple *logical* cores for each *physical* core. Some programs run faster when scheduling threads or processes to use these hyperthreaded cores, while most HPC codes run more slowly. Rather than show all of the sparklines at half-height simply because the hyperthreaded cores are (wisely) not being used, Arm MAP tries to detect this situation and will rescale its expectations to the number of physical cores used by your program.




Figure C13-1 Processes and Cores window

If this goes wrong for any reason, you will see large portions of unusual colors in your sparklines and the application activity chart (for example, bright red). When that happens, open this dialog and increase the *cores per process* value.

To open this dialog use the **View > Processes and Cores** menu option, or click the **X cores (Y per process)** hyperlinked text in the application details section above the metric graphs.

Chapter C14

Run MAP from the command line

Arm MAP can be run from the command line.

It contains the following sections:

- [C14.1 Command line arguments](#) on page C14-342.
- [C14.2 Profile MPMD programs](#) on page C14-344.
- [C14.3 Profile MPMD programs without Express Launch](#) on page C14-345.

C14.1 Command line arguments

Here are a number of command line arguments to use with Arm MAP.

Argument	Description
<code>--no-mpi</code>	Run Arm MAP with one process and without invoking <code>mpirun</code> , <code>mpiexec</code> , or equivalent.
<code>--queue</code>	Force Arm MAP to submit the job to the queueing system.
<code>--no-queue</code>	Run Arm MAP without submitting the job to the queueing system.
<code>--view=VIEW</code>	Start Arm MAP using <code>VIEW</code> as the default view. <code>VIEW</code> must be one of (<code>main</code> / <code>pthread</code> / <code>openmp</code>). If the selected view is not available, the <code>main</code> view will be displayed.
<code>--export=OUTPUT.json</code> <code>PROFILEDATA.map</code>	Export <code>PROFILEDATA.map</code> to <code>OUTPUT.json</code> in JSON format, without user interaction. For the format specification see JSON format on page C15-348 .
<code>--profile</code>	Generate an Arm MAP profile but without user interaction. This will not display the Arm MAP user interface. Messages are printed to the standard output and error. The job is not run using the queueing system unless used in conjunction with <code>--queue</code> . When the job finishes a map file is written and its name is printed.
<code>--export-functions=FILE</code>	Export all the profiled functions to <code>FILE</code> . Use this in conjunction with <code>--profile</code> . The output should be CSV file name. For example, <code>map -profile -export-functions=foo.csv ...</code>
<code>--start-after=TIME</code>	Start profiling <code>TIME</code> seconds after the start of your program. Use this in conjunction with <code>--stop-after=TIME</code> to focus Arm MAP on a particular time interval of the run of your program.
<code>--stop-after=TIME</code>	Stop profiling <code>TIME</code> seconds after the start of your program. This will terminate your program and proceed to gather the samples taken after the time given has elapsed.
<code>--enable-metrics=METRICS</code> and <code>--disable-metrics=METRICS</code>	Allows you to specify comma-separated lists which explicitly enable or disable metrics for which data is to be collected. If the metrics specified cannot be found, or if a metric is both enabled and disabled, an error message is displayed and Arm MAP exits. Metrics which are always enabled or disabled cannot be explicitly disabled or enabled, respectively. A metrics source library which has all its metrics disabled, either in the XML definition or via <code>--disable-metrics</code> , will not be loaded. Metrics which can be explicitly enabled or disabled can be listed using the <code>--list-metrics</code> option. The enabled/disabled metrics settings do not persist when running Arm MAP without the user interface, so they will need to be specified for each profiling session. When running Arm MAP in user interface mode, the effect of these settings will be displayed in the Metrics section of the Run dialog, where you can further refine the settings. These settings will then persist to the next user interface session.
<code>--cuda-kernel-analysis</code>	Enables CUDA kernel analysis mode, providing line level profiling information on CUDA kernels running on a GPU at the cost of potentially significant overhead. See GPU profiling on page C16-357 .

When running without the user interface, normal redirection syntax can be used to read data from a file as a source for the executable's standard input.

Examples:

```
cat <input-file> | map --profile ...
map --profile ... < <input-file>
```

Normal redirection can also be used to write data to a file from the executable's standard output:

```
map --profile ... > <output-file>
```

For OpenMP jobs, simply use the OMP_NUM_THREADS environment variable (or leave it blank) exactly as you usually would when running your program. There is no need to pass the number of threads to Arm MAP as an argument.

```
OMP_NUM_THREADS=8 map --profile ... > <output-file>
```

C14.2 Profile MPMD programs

The easiest way to profile MPMD programs is by using Express Launch to start your program.

To use Express Launch, prefix your normal MPMD launch line with `map`. For example, to profile an MPMD program without user interaction you can use:

```
map --profile mpirun -n 1 ./master : -n 2 ./worker
```

For more information about Express Launch, and compatible MPI implementations, see [Express Launch \(MAP\)](#) on page C1-242.

C14.3 Profile MPMD programs without Express Launch

The command to create a profile from an MPMD program using Arm MAP is:

```
map <map mode> --np=<#processes> --mpiargs=<MPMD command> <one MPMD program>
```

This example shows how to run Arm MAP without user interaction using the flag `-profile`:

```
map --profile --np=16 --mpiargs="-n 8 ./exe1 : -n 8 ./exe2" ./exe1
```

The number of processes used by the MPMD programs is set, in this case $8+8=16$. Then an MPMD style command as an mpi argument is specified, followed by one of the MPMD programs.

Chapter C15

Export profiler data in JSON format

You can export the profiler data in machine readable JSON format.

To export as JSON, open a `.map` file in Arm MAP. Then the profile data can be exported by selecting **File > Export Profile Data as JSON**.

For a command-line option, see [Run MAP from the command line](#) on page C14-341.

It contains the following sections:

- [C15.1 JSON format](#) on page C15-348.
- [C15.2 Activities](#) on page C15-350.
- [C15.3 Metrics](#) on page C15-352.
- [C15.4 Example JSON output](#) on page C15-354.

C15.1 JSON format

The JSON document contains a single JSON object containing two object members, `info` containing general information about the profiled program, and `samples` with the sampled information.

See an example of profile data exported to a JSON file in [Example JSON output on page C15-354](#).

Table C15-1 JSON data structure

Data	Type	Description
<code>info</code>	Object	If some information is not available, the value is null instead.
<code>command_line</code>	String	Command line call used to run the profiled program. For example, <code>aprun -N 24 -n 256 -d 1 ./my_exe</code> .
<code>machine</code>	String	Hostname of the node on which the executable was launched.
<code>notes</code>	String	A short description of the run or other notes on configuration and compilation settings. This is specified by setting the environment variable <code>ALLINEA_NOTES</code> before running Arm MAP.
<code>number_of_nodes</code>	Number	Number of nodes run on.
<code>number_of_processes</code>	Number	Number of processes run on.
<code>runtime</code>	Number	Runtime in milliseconds.
<code>start_time</code>	String	Date and time of run in ISO 8601 format.
<code>create_version</code>	String	Version of MAP used to create the <i>map</i> file.
<code>metrics</code>	Object	Attributes about the overall run, reported once per process, each represented by an object with <code>max</code> , <code>min</code> , <code>mean</code> , <code>var</code> , and <code>sums</code> fields, or <code>null</code> , when the metric is not available. The <code>sums</code> series contains the sum of the metric across all processes or nodes for each sample. In many cases, the values over all nodes is the same. This means that the <code>max</code> , <code>min</code> , and <code>mean</code> values are the same, with variance zero. For example, in homogeneous systems, <code>num_cores_per_node</code> is the same over all nodes.
• <code>wchar_total</code>	Object	The number of bytes written in total by I/O operation system calls (see <code>wchar</code> in the Linux Programmer's Manual page 'proc': <code>man 5 proc</code>).
• <code>rchar_total</code>	Object	The number of bytes read in total by I/O operation system calls (see <code>rchar</code> in the Linux Programmer's Manual page 'proc': <code>man 5 proc</code>).
• <code>num_cores_per_node</code>	Object	Number of cores available per node.
• <code>memory_per_node</code>	Object	RAM installed per node.
• <code>nvidia_gpus_count</code>	Object	Number of GPUs per node.
• <code>nvidia_total_memory</code>	Object	GPU frame buffer size per node.
• <code>num_omp_threads_per_process</code>	Object	Number of OpenMP worker threads used per process.
<code>samples</code>	Object	-
<code>count</code>	Number	Number of samples recorded.
<code>window_start_offset</code>	Array of Numbers	Offset of the beginning of each sampling window, starting from zero. The actual sample might have been taken anywhere in between this offset and the start of the next window, that is the window offsets <code>wi</code> and <code>wi+1</code> define a semi-open set (<code>wi,wi+1</code>) in which the sample was taken.

Table C15-1 JSON data structure (continued)

Data	Type	Description
activity	Object	Contains information about the proportion of different types of activity performed during execution, according to different view modes. The types of view modes possibly shown are OpenMP, PThreads and Main Thread, described in Main-thread, OpenMP, and Pthread view modes on page C12-335. Only available view modes are exported, for example, a program without OpenMP sections will not have an OpenMP activity entry. ————— Note ————— The sum of the proportions in an activity might not add up to 1, this can happen when there are fewer threads running than Arm MAP has expected. Occasionally the sum of the proportions shown for a sample in PThreads or OpenMP threads mode might exceed 1. When this happens, the profiled program uses more cores than Arm MAP assumes the maximum number of cores per process can be. This can be due to middleware services launching helper threads which, unexpectedly to Arm MAP, contribute to the activity of the profiled program. In this case, the proportions for that sample should not be compared with the rest of proportions for that activity in the sample set. —————
metrics	Object	Contains an object for each metric that was recorded. These objects contain four lists each, with the minimum, maximum, average, and variance of that metric in each sample. The format of a metrics entry is given in Metrics on page C15-352. All metrics recorded in a run are present in the JSON, including custom metrics. The names and descriptions of all core Arm MAP metrics are given in Metrics on page C15-352. It is assumed that a user including a custom metrics library is aware of what the custom metric is reporting. See the Arm Metric Plugin Interface documentation.

C15.2 Activities

Each exported object in an activity is presented as a list of fractional percentages (0.0 - 1.0) of sample time recorded for a particular activity during each sample window. Therefore, there are as many entries in these list as there are samples.

Description of categories

The following is the list of all of the categories. Only available categories are exported, see and below.

- `normal_compute`: Proportion of time spent on the CPU which is not categorized as any of the following activities. The computation can be, for example, floating point scalar (vector) addition, multiplication, or division.
- `point_to_point_mpi`: Proportion of time spent in point-to-point MPI calls on the main thread and not inside an OpenMP region.
- `collective_mpi`: Proportion of time spent in collective MPI calls on the main thread and not inside an OpenMP region.
- `point_to_point_mpi_openmp`: Proportion of time spent in point-to-point MPI calls made from any thread within an OpenMP region.
- `collective_mpi_openmp`: Proportion of time spent in collective MPI calls made from any thread within an OpenMP region.
- `point_to_point_mpi_non_main_thread`: Proportion of time spent in point-to-point MPI calls on a pthread, but not on the main thread nor within an OpenMP region.
- `collective_mpi_non_main_thread`: Proportion of time spent in collective MPI calls on a pthread, but not on the main thread nor within an OpenMP region.
- `openmp`: Proportion of time spent in an OpenMP region, that is compiler-inserted calls used to implement the contents of a OpenMP loop.
- `accelerator`: Proportion of time spent in calls to accelerators, that is, blocking calls waiting for a CUDA kernel to return.
- `threads`: Proportion of compute time on a non-main (worker) pthread.
- `openmp_overhead_in_region`: Proportion of time spent setting up OpenMP structures, waiting for threads to finish and so on.
- `openmp_overhead_no_region`: Proportion of time spent in calls to the OpenMP runtime from an OpenMP region.
- `synchronisation`: Proportion of time spent in thread synchronization calls, such as `pthread_mutex_lock`.
- `io_reads`: Proportion of time spent in I/O read operations, such as 'read'.
- `io_writes`: Proportion of time spent in I/O write operations. Also includes file open and close time as these are typically only significant when writing.
- `io_reads_openmp`: Proportion of time spent in I/O read operations from within an OpenMP region.
- `io_writes_openmp`: Proportion of time spent in I/O write operations from within an OpenMP region.
- `mpi_worker`: Proportion of time spent in the MPI implementation on a worker thread.
- `mpi_monitor`: Proportion of time spent in the MPI monitor thread.
- `openmp_monitor`: Proportion of time spent in the OpenMP monitor thread.
- `sleep`: Proportion of time spent in sleeping threads and processes.

Categories available in main_thread activity

- `normal_compute`
- `point_to_point_mpi`
- `collective_mpi`
- `point_to_point_mpi_openmp`
- `collective_mpi_openmp`
- `openmp`
- `accelerator`
- `openmp_overhead_in_region`
- `openmp_overhead_no_region`
- `synchronisation`

- io_reads
- io_writes
- io_reads_openmp
- io_writes_openmp
- sleep

Categories available in openmp and pthreads activities

- normal_compute
- point_to_point_mpi
- collective_mpi
- point_to_point_mpi_openmp
- collective_mpi_openmp
- point_to_point_mpi_non_main_thread
- collective_mpi_non_main_thread
- openmp
- accelerator
- pthreads
- openmp_overhead_in_region
- openmp_overhead_no_region
- synchronisation
- io_reads
- io_writes
- io_reads_openmp
- io_writes_openmp
- mpi_worker
- mpi_monitor
- openmp_monitor
- sleep

C15.3 Metrics

The following list contains the core metrics reported by Arm MAP.

Only available metrics are exported to JSON. For example, if there is no Lustre filesystem then the Lustre metrics will not be included. If any custom metrics are loaded, they will be included in the JSON, but are not documented here.

For more information on the metrics, see [Metrics view on page C9-299](#)

- CPU Instructions: see [CPU instructions on page C9-302](#)
 - `instr_fp`
 - `instr_int`
 - `instr_mem`
 - `instr_vector_fp`
 - `instr_vector_int`
 - `instr_branch`
 - `instr_scalar_fp`: The percentage of time each rank spends in standard x87 floating-point operations.
 - `instr_scalar_int`: The percentage of time each rank spends in standard integer operations.
 - `instr_implicit_mem`: Implicit memory accesses. The percentage of time spent executing instructions with implicit memory accesses.
 - `instr_other`: The percentage of time each rank spends in instructions which cannot be categorized as any of the ones given above.
- CPU Time: see [CPU time on page C9-305](#)
 - `cpu_time_percentage`: See *CPU time* in [CPU time on page C9-305](#)
 - `user_time_percentage`: See *User-mode CPU time* in [CPU time on page C9-305](#)
 - `system_time_percentage`: See *Kernel-mode CPU time* in [CPU time on page C9-305](#)
 - `voluntary_context_switches`: See *Voluntary context switches (1/s)* in [CPU time on page C9-305](#)
 - `involuntary_context_switches`: See *Involuntary context switches (1/s)* in [CPU time on page C9-305](#)
 - `loadavg`: See *System load* in [CPU time on page C9-305](#)
- I/O: see [I/O on page C9-306](#)
 - `rchar_rate`: See *POSIX I/O read rate (B/s)* in [I/O on page C9-306](#)
 - `wchar_rate`: See *POSIX I/O write rate (B/s)* in [I/O on page C9-306](#)
 - `bytes_read`: See *Disk read transfer* in [I/O on page C9-306](#)
 - `bytes_written`: See *Disk write transfer* in [I/O on page C9-306](#)
 - `syscr`: See *POSIX read syscall rate* in [I/O on page C9-306](#)
 - `syscw`: See *POSIX write syscall rate* in [I/O on page C9-306](#)
- Lustre
 - `lustre_bytes_read`: Lustre read transfer (B/s)
 - `lustre_bytes_written`: Lustre write transfer (B/s)
 - `lustre_rchar_total`: Lustre bytes read
 - `lustre_wchar_total`: Lustre bytes written
- Memory: see [Memory on page C9-307](#)
 - `rss`: See *Memory usage in bytes (Resident Set Size)* in [Memory on page C9-307](#)
 - `node_mem_percent`: See *Node memory usage* in [Memory on page C9-307](#)
- MPI: see [MPI calls on page C9-308](#)
 - `mpi_call_time`: See *MPI call duration (ns)* in [MPI calls on page C9-308](#)
 - `mpi_sent`: See *MPI sent* in [MPI calls on page C9-308](#)
 - `mpi_recv`: See *MPI received* in [MPI calls on page C9-308](#)
 - `mpi_calls`: Number of MPI calls per second per process
 - `mpi_p2p`: See *MPI point-to-point* in [MPI calls on page C9-308](#)
 - `mpi_collect`: See *MPI collective operations* in [MPI calls on page C9-308](#)

- `mpi_p2p_bytes`: See *MPI point-to-point bytes* in [MPI calls](#) on page C9-308
- `mpi_collect_bytes`: See *MPI collective bytes* in [MPI calls](#) on page C9-308
- Accelerator: see [Accelerator](#) on page C9-310
 - `nvidia_gpu_usage`: See *GPU utilization* in [Accelerator](#) on page C9-310
 - `nvidia_memory_used_percent`: See *GPU memory usage* in [Accelerator](#) on page C9-310
 - `nvidia_memory_used`: GPU memory usage in bytes
- Energy: see [Energy](#) on page C9-311
 - `nvidia_power`: See *GPU power usage (mW/node)* in [Energy](#) on page C9-311
 - `rapl_power`: See *CPU power usage (W/node)* in [Energy](#) on page C9-311
 - `system_power`: See *System power usage (W/node)* in [Energy](#) on page C9-311
 - `rapl_energy`: CPU energy, integral of `rapl_power` (J)
 - `system_energy`: CPU energy, integral of `system_power` (J)

C15.4 Example JSON output

In this section an example is given of the format of the JSON that is generated from a Arm MAP file. This illustrates the description that has been given in the previous sections. This is not a full file, but should be used as an indication of how the information looks after export.

```

{
  "info" : {
    "command_line" : "mpirun -np 4 ./exec",
    "machine" : "hal9000",
    "number_of_nodes" : 30,
    "number_of_processes" : 240,
    "runtime" : 8300,
    "start_time" : "2016-05-13T11:36:31",
    "create_version" : "6.0.4"
    "metrics": {
      wchar_total: {max: 384605588, min: 132, mean: 24075798, var: 546823},
      rchar_total: {max: 6123987, min: 63, mean: 9873, var: 19287},
      num_cores_per_node: {max: 4, min: 4, mean: 4, var: 0},
      memory_per_node: {max: 4096, min: 4096, mean: 4096, var: 0},
      nvidia_gpus_count: {max: 0, min: 0, mean: 0, var: 0},
      nvidia_total_memory: {max: 0, min: 0, mean: 0, var: 0},
      num_omp_threads_per_process: {max: 6, min: 6, mean: 6, var: 0},
    }
  },
  "samples" : {
    "count" : 4,
    "window_start_offsets" : [ 0, 0.2, 0.4, 0.6 ],
    "activity" : {
      "main_thread" : {
        "normal_compute" : [ 0.762, 0.996, 1, 0.971 ],
        "io_reads" : [ 0.00416, 0.00416, 0, 0.00416 ],
        "io_writes" : [ 0.233, 0, 0, 0 ],
        "openmp" : [ 0, 0, 0, 0.01667 ],
        "openmp_overhead_in_region" : [ 0, 0, 0, 0.1 ],
        "openmp_overhead_no_region" : [ 0, 0, 0, 0.00417 ],
        "sleep" : [ 0, 0, 0, 0 ]
      },
      "openmp" : {
        "normal_compute" : [ 0.762, 0.996, 1, 0.971 ],
        "io_reads" : [ 0.00416, 0.00416, 0, 0.00416 ],
        "io_writes" : [ 0.233, 0, 0, 0 ],
        "openmp" : [ 0, 0, 0, 0.01319 ],
        "openmp_overhead_in_region" : [ 0, 0, 0, 0 ],
        "openmp_overhead_no_region" : [ 0, 0, 0, 0 ],
        "sleep" : [ 0, 0, 0, 0 ]
      },
      "pthreads" : {
        "io_reads" : [ 0.00069, 0.00069, 0, 0.00069 ],
        "io_writes" : [ 0.0389, 0, 0, 0 ],
        "normal_compute" : [ 0.1270, 0.1659, 0.1666, 0.1652 ],
        "openmp" : [ 0, 0, 0, 0.01319 ],
        "openmp_overhead_in_region" : [ 0, 0, 0, 0.02153 ],
        "openmp_overhead_no_region" : [ 0, 0, 0, 0.00069 ],
        "sleep" : [ 0, 0, 0, 0 ]
      }
    },
    "metrics" : {
      "wchar_total" : {
        "mins" : [ 3957, 3957, 3958, 4959 ],
        "maxs" : [ 4504, 4959, 5788, 10059 ],
        "means" : [ 3965.375, 4112.112, 4579.149, 6503.496 ],
        "vars" : [ 2159.809, 49522.783, 169602.769, 2314522.699 ],
        "sums" : [ 15860, 16448, 18316, 26012 ]
      },
      "bytes_read" : {
        "mins" : [ 0, 0, 0, 0 ],
        "maxs" : [ 34647.255020415301, 0, 0, 0 ],
        "means" : [ 645.12988722358205, 0, 0, 0 ],
        "vars" : [ 9014087.0327749606, 0, 0, 0 ],
        "sums" : [ 2580, 0, 0, 0 ]
      },
      "bytes_written" : {
        "mins" : [ 0, 0, 0, 0 ],
        "maxs" : [ 123, 0, 0, 0 ],
        "means" : [ 32, 0, 0, 0 ],
        "vars" : [ 12, 0, 0, 0 ],
        "sums" : [ 128, 0, 0, 0 ]
      }
    }
  }
}

```

```
} } }
```


Chapter C16

GPU profiling

You can use the GPU profiling capabilities when working with CUDA programs.

Notes:

- CUDA kernels generated by OpenACC, CUDA Fortran, or offloaded OpenMP regions are not yet supported by Arm MAP.
- GPU profiling is only supported with CUDA 8.0 and later.
- GPU profiling is not supported if the CUDA driver and toolkit versions do not match (for example, profiling a CUDA 8.0 program with a CUDA 9.0 driver is not supported).
- GPU profiling is not supported when statically linking the Arm MAP sampler library.

It contains the following sections:

- [C16.1 GPU Kernels tab on page C16-358.](#)
- [C16.2 Kernel analysis on page C16-359.](#)
- [C16.3 Compilation on page C16-361.](#)
- [C16.4 Performance impact on page C16-362.](#)
- [C16.5 Customize GPU profiling behavior on page C16-363.](#)
- [C16.6 Known issues on page C16-364.](#)

C16.1 GPU Kernels tab

When profiling programs that use CUDA 8.0 and later, GPU kernels that can be tracked by NVIDIA's CUDA Profiling Tools Interface (CUPTI) will be displayed in the **GPU Kernels** tab.




Figure C16-1 GPU Kernels tab

This lists the CUDA kernels that were detected in the program alongside graphs indicating when those kernels were active. If multiple kernels were identified in a process within a particular sample they will have equal weighting in this graph.

C16.2 Kernel analysis

CUDA kernel analysis mode is an advanced feature that provides insight into the activity within CUDA kernels. This mode can be enabled from the **Run** dialog or from the command line with `-cuda-kernel-analysis`.




Figure C16-2 Run dialog with CUDA kernel analysis enabled

When enabled, the **GPU Kernels** tab is enhanced to show a line-level breakdown of warp stalls. The possible categories of warp stall reasons are as listed in the enum `Cupti_ActivityPCSamplingStallReason` in the CUPTI API documentation (http://docs.nvidia.com/cuda/cupti/group__CUPTI__ACTIVITY__API.html):

Selected No stall, instruction is selected for issue.

Instruction fetch Warp is blocked because next instruction is not yet available, because of an instruction cache miss, or because of branching effects.

Execution dependency Instruction is waiting on an arithmetic dependency.

Memory dependency Warp is blocked because it is waiting for a memory access to complete.

Texture sub-system Texture sub-system is fully utilized or has too many outstanding requests.

Thread or memory barrier Warp is blocked as it is waiting at `__syncthreads` or at a memory barrier.

__constant__ memory Warp is blocked waiting for `__constant__` memory and immediate memory access to complete.

Pipe busy Compute operation cannot be performed due to required resource not being available.

Memory throttle Warp is blocked because there are too many pending memory operations.

Not selected Warp was ready to issue, but some other warp issued instead.

Other Miscellaneous stall reason.

Dropped samples Samples dropped (not collected) by hardware due to backpressure or overflow.

Unknown The stall reason could not be determined. Used when CUDA kernel analysis has not been enabled (see above) or when an internal error occurred within CUPTI or Arm MAP.

Input/Output	Project Files	OpenMP Stacks	OpenMP Regions	Functions	GPU Kernels	
GPU Kernels						
Breakdown	Selected	Blocked	GPU Kernels		Source	Position
13.1%	<0.1%	13.1%	gpuprof.exe [program] mult_gpu		GPU: all functions in kernels are inlined	
1.0%	<0.1%	1.0%	14 others		res1 += sh_A1[k-ks] * tmpB;	device.cu:73
0.4%	<0.1%	0.4%			for(int k=ks; k<ks+BLOCK_X; k++)	device.cu:69
0.4%	<0.1%	0.3%			sh_A2[tx] = A[(i+1)*pitch_A_nbelem+ks+tx...	device.cu:66
<0.1%	<0.1%	<0.1%			res2 += sh_A2[k-ks] * tmpB;	device.cu:74
<0.1%	<0.1%	<0.1%			for(int ks=0; ks<size; ks=ks+BLOCK_X)	device.cu:62
<0.1%	<0.1%	<0.1%			double tmpB = B[k*pitch_B_nbelem+j];	device.cu:72
<0.1%	<0.1%	<0.1%			sh_A1[tx] = A[i*pitch_A_nbelem+ks+tx];	device.cu:65
<0.1%	<0.1%	<0.1%			C[i*pitch_C_nbelem+j] += res1;	device.cu:81

Figure C16-3 GPU Kernels tab (with CUDA kernel analysis)

Note

Warp stalls are only reported per-kernel, so it is not possible to obtain the times within a kernel invocation at which different categories of warp stalls occurred. As function calls in CUDA kernels are automatically fully inlined it is not possible to see a stack trace of code within a kernel on the GPU.

Warp stall information is also present in the **Source code editor** ([GPU programs on page C3-278](#)), the **Selected lines view** ([GPU profiles on page C4-287](#)), and in a **Warp stall reasons** graph in the **Metrics view** ([Metrics view on page C9-299](#)).

C16.3 Compilation

When compiling CUDA kernels, do not generate debug information for device code (the `-G` or `-device-debug` flag) as this can significantly impair runtime performance. Use `-lineinfo` instead, for example:

```
nvcc device.cu -c -o device.o -g -lineinfo -O3
```

C16.4 Performance impact

Enabling the CUPTI sampling will impact the target program in the following ways:

- A short amount of time will be spent post-processing at the end of each kernel. This will depend on the length of the kernel and the CUPTI sampling frequency.
- Kernels will be serialized. Each CUDA kernel invocation will not return until the kernel has finished and CUPTI post-processing has been performed. Without CUDA kernel analysis mode kernel invocation calls return immediately to allow CUDA processing to be performed in the background.
- Increased memory usage whilst in a CUDA kernel. This may manifest as fluctuations between two memory usage values, depending on whether a sample was taken during a CUDA kernel or not.

Taken together the above may have a significant impact on the target program, potentially resulting in orders of magnitude slowdown. To combat this profile and analyze CUDA code kernels (with `-cuda-kernel-analysis`) and non-CUDA code (no `-cuda-kernel-analysis`) in separate profiling sessions.

The NVIDIA GPU metrics will be adversely affected by this overhead, particularly the GPU utilization metric. See [Accelerator on page C9-310](#).

When profiling CUDA code it may be useful to only profile a short subsection of the program so time is not wasted waiting for CUDA kernels you do not need to see data for. See *Profiling only part of a program* in [Profile a program on page C1-250](#) for instructions.

C16.5 Customize GPU profiling behavior

The interval at which CUPTI samples GPU warps can be modified by the environment variable `ALLINEA_SAMPLER_GPU_INTERVAL`. Accepted values are `max`, `high`, `mid`, `low`, and `min`, with the default value being `high`. These correspond to the values in the enum `Cupti_ActivityPCSamplingPeriod` in the CUPTI API documentation (http://docs.nvidia.com/cuda/cupti/group__CUPTI__ACTIVITY__API.html).

Using CUDA 11.0+ on GPUs with compute capability 7.0+, the interval at which CUPTI samples GPU warps can also be modified by providing an integer value $5 \leq x \leq 31$ to the environment variable `ALLINEA_SAMPLER_GPU_INTERVAL`. This sets the interval in cycles to exactly $2x$.

Reducing the sampling interval means warp samples are taken more frequently. While this may be needed for very short-lived kernels, setting the interval too low can result in a very large number of warp samples being taken which then require significant post-processing time when the kernel completes. Overheads of twice as long as the kernel's normal runtime have been observed. We recommend that the CUPTI sampling interval is not reduced.

C16.6 Known issues

There are a few known issues for GPU profiling.

- GPU profiling is only supported using CUDA 8.0 and later.
- GPU profiling is not supported if the CUDA driver and toolkit versions do not match (for example, profiling a CUDA 8 program with a CUDA 9 driver is not supported).
- When preparing your program for profiling, we recommend that you match the version of the CUDA toolkit to that of the CUDA driver.
- CUPTI allocates a small amount of host memory each time a kernel is launched. If your program launches many kernels in a tight loop this overhead can skew the memory usage figures.
- CUDA kernels generated by OpenACC, CUDA Fortran, or offloaded OpenMP regions are not yet supported by Arm MAP.
- The graphs are scaled on the assumption that there is a 1:1 relationship between processes and GPUs, each process having exclusive use of its own CUDA card. The graphs may be of an unexpected height if some processes do not have a GPU, or if multiple processes share the use of a common GPU.
- Enabling CUDA kernel analysis mode can have a significant performance impact as described in [Performance impact on page C16-362](#).
- GPU profiling is not supported when statically linking the Arm MAP sampler library.
- Stopping GPU profiling mid-process can prevent the **GPU Kernels** tab displaying, and might not report the kernel samples. This occurs when using `--stop-after` or the **Stop and Analyze** button. For better results, run the process for a longer time period with longer running kernels. When kernel samples are reported, they can be truncated.

Chapter C17

Python profiling

The Python profiling capabilities can be used to find and resolve bottlenecks for your Python codes.

For the latest information about Python profiling in Arm MAP, see the [Python Profiling](#) web page.

It contains the following sections:

- [C17.1 Profile a Python script on page C17-366.](#)
- [C17.2 Known issues on page C17-369.](#)

C17.1 Profile a Python script

This task describes how to profile a Python script. This feature is useful when profiling a mixed C, C++, Fortran, and Python program.

Python profiling replaces main thread stack frames originating from the Python interpreter with Python stack frames of the profiled Python script. To disable this feature, set `ALLINEA_SAMPLER_DISABLE_PYTHON_PROFILING=1`.

Arm MAP supports Python profiling with the following features:

- Profiles Python scripts running under the CPython interpreter (versions 2.7, and 3.5-3.8).
- Profiles Python scripts running under the Intel Distribution for Python.
- Profiles Python scripts running under the Anaconda Python distribution (version 3.6 not supported).
- Profiles Python scripts running under virtual environments.
- Profiles Python scripts that import modules which perform MPI on the main thread, such as `mpi4py`.
- Profiles Python scripts that import modules which use OpenMP.
- Profiles Python scripts that use the `threading` module.

Note

Arm MAP will output warnings if the threading model of the MPI module is `MPI_THREAD_MULTIPLE`, such as in `mpi4py`. To prevent these warnings, change the default settings in `mpi4py` with the following: `mpi4py.rc.threaded = False` or `mpi4py.rc.thread_level = "funneled"`.

Note

If you are profiling on a system using ALPS or SLURM and the Python script does not use MPI, you can set environment variables ([Starting a program on page Appx-H-500](#)) or you can import the `mpi4py` module.

Procedure

1. Check that the Python script runs successfully:

```
$ python myscript.py
```

2. To profile the Python script with Arm MAP, prepend the run command with `map`:

```
$ map python myscript.py
```




Figure C17-1 Profiling a Python script

3. Click **Run** and wait for Arm MAP to finish profiling the Python script.
4. View the profiling results in Arm MAP.

Results

When Arm MAP finishes profiling the Python script, it saves a `.map` file in the current working directory and opens it for viewing in the user interface (unless you are using the offline feature).

Example: Profiling a simple Python script

This section demonstrates how to profile the Python example script `python-profiling.py` located in the `examples` directory.

1. Change into the `examples` directory and run the makefile to compile the example.

```
$ make -f python-profiling.makefile
```

2. Start Arm MAP.

```
$ ../bin/map python ./python-profiling.py --index 35
```

3. Click **Run**.

4. Wait for Arm MAP to finish analyzing samples after the Python script has completed.

Note

The Arm MAP user interface launches showing the Python script and the line in the script where the most time was spent is selected.




Figure C17-2 Viewing Python profiling results

5. Locate the first `fibonacci_c` stack frame in the **Main Thread Stacks** view. The callout to the `C` function is appended under `main` Python stack frame.

Next steps

- Examine the **Main Thread Activity** graph ([Metrics view on page C9-299](#)) for an overview of time spent in Python code compared with non-Python code.
- View source code lines ([View source code \(MAP\) on page C3-274](#)) on which time was spent executing Python code and non-Python code.
- Compare time spent on the selected line executing Python code with non-Python code in the **Selected lines** view ([Selected lines view on page C4-283](#)).
- View a breakdown of time spent in different code paths in the **Main Thread Stacks** view ([Stacks view on page C5-289](#)).

Related information

C17.2 Known issues

There are a few known issues for Python profiling.

- Arm MAP requires a significant amount of time to analyze samples when profiling a Python script that imports modules which use OpenBLAS, such as NumPy. This is caused by the lack of unwind information in OpenBLAS. This results in `partial trace` nodes being displayed in Arm MAP.
- `mpi4py` uses some MPI functions that were introduced in MPI version 3. For example `MPI_Mrecv`. Arm MAP does not collect metrics from these functions, therefore MPI metrics for `mpi4py` will be inaccurate. To workaround this, use a custom Python MPI wrapper that only uses functions that were available before MPI version 3.
- When using reverse connect (`-connect`) and quick start (`-start`) in conjunction, the full path to the Python application must be provided.
- The Anaconda Python 3.6.x interpreter is aggressively optimized. This causes multiple startup issues when profiling with Arm MAP and is not supported.
- Arm MAP is known to fail at startup with "python did not start properly" for some versions of Python like Intel Python distributed with the Intel Compiler. This can be worked around by setting the environment variable `ALLINEA_DISABLE_BREAK_BEFORE_MAIN_PROLOGUE=1`.

Chapter C18

Performance analysis with Caliper instrumentation

Caliper is a program instrumentation and performance measurement framework. It is a performance analysis toolbox in a library, that enables you to insert performance analysis capabilities directly into programs, and activate them at runtime.

Caliper is intended for use with HPC programs, but works for any C/C++/Fortran program on Unix/Linux.

When Arm MAP profiles a program instrumented with Caliper source code annotations, the stack of Caliper attributes with keys of interest is taken alongside regular samples. The keys of interest use the same attribute names as the Caliper high-level API (function, loop, statement, annotation). In Arm MAP, you can see where time was spent in any set of these key-attribute pairs.

It contains the following sections:

- [C18.1 Get Caliper on page C18-372.](#)
- [C18.2 Annotate your program on page C18-373.](#)
- [C18.3 Analyze your program on page C18-374.](#)
- [C18.4 Guidelines on page C18-376.](#)

C18.1 Get Caliper

Download Caliper (version 2.0.1 or later) from [GitHub](#).

Build and install Caliper, then use it to instrument programs of your choice as described in the Caliper documentation:

- [Read a summary](#)
- [Full documentation](#)
- [Pre-instrumented examples](#) (LULESH2 and Quicksilver)

C18.2 Annotate your program

Typically, we integrate Caliper into a program by marking source-code sections of interest with descriptive annotations. Arm MAP can connect to Caliper and access the information provided by these annotations.

Annotate in C/C++

Arm MAP supports Caliper's high-level C/C++ API for annotating functions, loops (although loop iterations are not recorded) and code regions. Refer to the [Annotation API](#) in the Caliper documentation, for details and examples.

Note

Using the low-level API for C/C++ applications is not recommended. Low-level API calls must be nested, otherwise they are not supported.

Annotate in Fortran

For Fortran programs, the low-level API must be used to emulate the high-level API. These label types are supported: function, loop, statement, annotation. They match the attribute names used by the high-level API.

Note

Caliper regions must be nested. Close all inner regions before closing an outer region.

```
use Caliper
call cali_begin_string_byname('function', 'myFunction')
...
call_end_byname('function')
```

C18.3 Analyze your program

Use Caliper with Arm MAP to get a quick idea of how much time is spent in the various phases of your program, to help you decide where to focus your optimization efforts.

Procedure

1. Instrument your program with Caliper annotations (or use one of the Caliper-provided examples).
2. Dynamically link your program against Caliper, and profile it with Arm MAP.

————— Note —————

Arm Forge automatically detects the Caliper regions of your code, and does not require you to set any specific flags or options.

When profiling, Caliper might prompt you to enable Caliper services. This is not required. You can enable Caliper services if you want to, but this increases overhead without providing any additional data to Arm MAP. Enabling it might produce Caliper output files you can manually examine.

In the **Metrics** view, the **Selected regions** chart (initially empty) displays below the **Main thread activity** chart.



Figure C18-1 Selected regions chart

3. Switch to the **Regions** tab at the bottom of the screen to view the regions of your code that you annotated with Caliper.

Input/Output Project Files Main Thread Stacks Functions Regions				
Regions				
	Activity	Label	Time	Legend
☐	100.0%	function:slow	100.0%	
☒	44.9%	function:stride	44.9%	
☒	36.8%	function:imbalance	36.8%	
☒	18.3%	function:overlap	18.3%	

Figure C18-2 Regions tab

4. Optionally, select the **Legend** column in the **Regions** tab to change the color that is used to represent a region.

————— Note —————

Forge assigns a color for each area of the code where Caliper instrumentation occurs.

5. Select one or more regions that you want to display in the **Selected regions** chart in the **Metrics** view.

————— Note —————

Press Ctrl (or Shift), and click to select multiple regions.

Periods of time where the program was inside that region show in the **Selected regions** chart.

6. Enable **View > Regions-focused view**.




Figure C18-3 View all selected regions

Results

In each of the charts, the horizontal axis indicates wall clock time.

The **Selected regions** chart, and the **Main thread activity** chart above it, use the same colors and scales as the per-line sparkline charts.

The **Application activity** chart can sometimes display in place of, or in addition to, the **Main thread activity** chart. This depends on the type of code that you profile, and the display mode you use (main-thread-only, pthread, or OpenMP).

For more information about how to read these charts, see [View source code \(MAP\) on page C3-274](#).

The button in the **Selected regions** chart toggles on and off the display of the selected region sparklines in the **Main Thread Stacks** view. These sparklines provide more detail about the percentage of the core time in the selected regions.

Activating the button can also modify the timeglyphs for sparklines in other views (**Thread Stacks**, **OpenMP Stacks**, **OpenMP Regions**, **Functions**)

Application activity timeglyphs in the PSVs, **Functions** view, and **Source code** view switch to showing the time in the currently selected set of regions.

C18.4 Guidelines

Next steps

Right-click on a region in the **Regions** tab to access further options. From here you can:

- Enable or disable all regions at once.
- Automatically reassign colors to regions based on the percentage of time in each in the current selected time range.
- Copy a text representation of the tab, or export it to a file.

Related information

Chapter C19

Arm Statistical Profiling Extension (SPE)

The Arm Statistical Profiling Extension (SPE) is an optional feature in ARMv8.2 hardware that allows CPU instructions to be sampled and associated with the source code location where that instruction occurred.

Arm MAP can use SPE to list the source code lines that frequently trigger certain hardware events such as:

- Branch mispredicts
- Level 1 data cache (L1D) refills
- Last level cache (LLC) misses
- Translation lookaside buffer (TLB) walks

It contains the following sections:

- [C19.1 SPE Prerequisites](#) on page C19-378.
- [C19.2 Enable SPE from the command line](#) on page C19-379.
- [C19.3 Enable SPE from the run dialog](#) on page C19-380.
- [C19.4 Analyze your program](#) on page C19-381.
- [C19.5 Guidelines](#) on page C19-383.
- [C19.6 Memlock limit exceeded](#) on page C19-384.
- [C19.7 SPE on Amazon Web Services](#) on page C19-385.
- [C19.8 Known issues](#) on page C19-386.

C19.1 SPE Prerequisites

The Arm Statistical Profiling Extension (SPE) is an optional feature in ARMv8.2 and is present in CPUs, such as the Neoverse N1. In addition to hardware, using this Arm MAP feature has the following prerequisites:

- Use Linux kernel 4.16 or later. To confirm that support is available, check for the path `/sys/bus/event_source/devices/arm_spe_XX`.
- For kernel versions 4.20 and later, apply the patch at <http://lkml.iu.edu/hypermail/linux/kernel/06760.html>, or upgrade to a kernel version greater than or equal to 5.1-RC5.
- Disable kernel page table isolation for the target. To ensure that kernel page table isolation is disabled, boot the machine with the command-line argument `kpti=off`.

C19.2 Enable SPE from the command line

Enable SPE profiling by using the `--spe` argument:

```
/path/to/forge/bin/map --profile --spe=<event-filter> ...
```

Replace `<event-filter>` with one of these values:

- `cache`: collect L1D refills and LLC miss events
- `mispredict`: collect branch mispredict events
- `tlb`: collect TLB walk events

C19.3 Enable SPE from the run dialog

Enable SPE profiling and select the appropriate event filter in the **Statistical Profiling Extension** section of the **Run** dialog. This section is only shown for aarch64 targets.




Figure C19-1 The Run dialog with SPE profiling enabled

A warning displays if SPE is not supported by the current host. This can be overridden and ignored if you know that SPE is available on the compute nodes.

C19.4 Analyze your program

Use Arm MAP with the Arm Statistical Profiling Extension (SPE) to find locations in your code where particular hardware events are most commonly triggered.

1. Profile your program using Arm MAP with SPE profiling enabled. Use the filter you are interested in (cache events, tlb walk events, or branch mispredict events).
2. Switch to the **SPE Tables** tab at the bottom of the screen to view the average number of hardware events sampled per process by the SPE feature, by source code line.

Function	Position	Source	Samples
slow::stride	slow.f90:114	arr_out(i,j) = sqrt(arr_in(i,j) - arr_in(i,j)) + sqrt(arr_in(i,j) - arr_in(i,j))	1,338.9
slow::stride	slow.f90:127	arr_out(i,j) = sqrt(arr_in(i,j) - arr_in(i,j)) + sqrt(arr_in(i,j) - arr_in(i,j))	51.5
slow::imbalance	slow.f90:87	a=sqrt(a)+1.1*j	18.5
open64		open64 (no debug info)	17.8
<unknown> from /home/jonbyr01/arm-forg...		<unknown> from /home/jonbyr01/arm-forg-21.0-linux-aarch64/map/me...	7.9
<unknown> from /home/jonbyr01/arm-forg...		<unknown> from /home/jonbyr01/arm-forg-21.0-linux-aarch64/map/me...	7.6
std::Rb_tree_increment(std::Rb_tree_node...		std::Rb_tree_increment(std::Rb_tree_node_base*) (no debug info)	3.6
clock_gettime		clock_gettime (no debug info)	3
_kernel_clock_gettime		_kernel_clock_gettime (no debug info)	2.4
memcpy@plt		memcpy@plt (no debug info)	2.4
strlen		strlen (no debug info)	1.5
mmmc_stack_trypop		mmmc_stack_trypop (no debug info)	1.5

Figure C19-2 SPE Table tab with a flat list of SPE samples

3. Optionally, toggle between a plain list of lines containing events, and group the rows by function.

Function	Position	Source	Samples
slow::stride	slow.f90:114	arr_out(i,j) = sqrt(arr_in(i,j) - arr_in(i,j)) + sqrt(arr_in(i,j) - arr_in(i,j))	1,338.9
slow::stride	slow.f90:127	arr_out(i,j) = sqrt(arr_in(i,j) - arr_in(i,j)) + sqrt(arr_in(i,j) - arr_in(i,j))	51.5
slow::imbalance	slow.f90:87	a=sqrt(a)+1.1*j	18.5
open64		open64 (no debug info)	17.8
<unknown> from /home/jonbyr01/arm-f...		<unknown> from /home/jonbyr01/arm-forg-21.0-linux-aarch64/map/me...	7.9
<unknown> from /home/jonbyr01/arm-f...		<unknown> from /home/jonbyr01/arm-forg-21.0-linux-aarch64/map/me...	7.6
std::Rb_tree_increment(std::Rb_tree_n...		std::Rb_tree_increment(std::Rb_tree_node_base*) (no debug info)	3.6
clock_gettime		clock_gettime (no debug info)	3
memcpy@plt		memcpy@plt (no debug info)	2.4
_kernel_clock_gettime		_kernel_clock_gettime (no debug info)	2.4
mmmc_stack_trypop		mmmc_stack_trypop (no debug info)	1.5

Figure C19-3 SPE Table tab with SPE samples grouped by function

4. Click a row to jump to that location in the source code editor. The average number of SPE samples per process on each line is also displayed here.

```

110 ! the inner loops - in that case recompile with -O0 instead of the default -O3
111 do i=1,82
112 do j=1,8000
113 do j=1,2000
114 arr_out(i,j) = sqrt(arr_in(i,j) - arr_in(i,j)) + sqrt(arr_in(i,j) - arr_in(i,j))
115 arr_out(i,j) = arr_out(i,j) * arr_out(i,j)
116 end do
117 end do
118 end do
119
120 ! on a busy workstation some processes often finish faster and wait here

```

Figure C19-4 SPE Table tab with SPE samples grouped by function

5. If the cache event filter is used, two tables are generated. One table is for all level 1 data cache refill events, and the other is for those for that caused a cache miss in the last level of cache.

Switch between these tables using the menu at the top of the **SPE Tables** tab. This also updates the count of SPE samples that the code editor displays.

Alternatively, identify bottlenecks using the regular Arm MAP stacks tabs, and check the surrounding lines in the code editor for SPE samples to draw conclusions about the cause of the bottleneck.

You can profile your application several times using the different SPE filters to collect information on the various types of event.

Note

For performance reasons, Arm MAP aggregates SPE samples over the lifetime of the program being profiled, and merges samples taken from all threads. Because it is not possible to correctly display which SPE samples apply to a selected time range, the SPE samples display in the code editor is hidden when a time range is selected. When this occurs, a dialog box displays that you can dismiss.




Figure C19-5 Message about hidden SPE code editor annotations

The **SPE Tables** tab also displays a warning in this case. When you view the **Main Thread Only** view mode, another warning displays because some of the SPE samples that are listed might have come from one of the non-main threads not normally included in this view.

Input/Output Project Files Main Thread Stacks Functions SPE Tables				
SPE Tables				
L1D refills				
⚠Displaying restricted time range, but SPE data always applies to entire runtime. ⚠In "Main Thread Only" View, but SPE data always applies to all threads.				
Function	Position	Source	Samples	
slow::stride	slow.f90:114	arr_out(i,j) = sqrt(arr_in(i,j) - arr_in(i,j)) + sqrt(arr_in(i,j) - arr_in(i,j))	1,338.9	
slow::stride	slow.f90:127	arr_out(i,j) = sqrt(arr_in(i,j) - arr_in(i,j)) + sqrt(arr_in(i,j) - arr_in(i,j))	51.5	
slow::imbalance	slow.f90:87	a=sqrt(a)+1.1*j	18.5	
open64		open64 (no debug info)	17.8	
<unknown> from /home/jonbyr01/arm-f...		<unknown> from /home/jonbyr01/arm-forge-21.0-linux-aarch64/map/me...	7.9	
<unknown> from /home/jonbyr01/arm-f...		<unknown> from /home/jonbyr01/arm-forge-21.0-linux-aarch64/map/me...	7.6	
std::Rb_tree_increment(std::Rb_tree_n...		std::Rb_tree_increment(std::Rb_tree_node_base*) (no debug info)	3.6	
clock_gettime		clock_gettime (no debug info)	3	
memcpy@plt		memcpy@plt (no debug info)	2.4	
kernel_clock_gettime		kernel_clock_gettime (no debug info)	2.4	

Figure C19-6 Message explaining why SPE code editor annotations are hidden

C19.5 Guidelines

Recommendations for correctly using and interpreting the results of SPE profiling:

- The bars showing the number of SPE samples for source code lines are for visually comparing the relative number of SPE samples between different lines of code within this profile. A full bar indicates that more events were sampled on a specific source code line than on the others.
- Arm MAP handles SPE data in a time-agnostic manner: the numbers reported in the **SPE Tables** tab are from the entire program run, and from across all threads. Selecting a time range or using a **Main Thread Only** view mode does not change what is reported by the **SPE Tables** tab or the SPE source code annotations.
- To keep the impact of enabling SPE profiling to reasonable levels, Arm MAP only utilizes a subset of the samples taken by SPE. Therefore, the number of SPE samples (hits) that are taken, depend on a number of factors that can vary between profiled applications, host machine configuration, and versions of Arm MAP.
- You can use SPE profiling in conjunction with the [Configurable Perf metrics on page C10-317](#) feature. Arm recommends that you enable the CPU instruction metrics which relate to the SPE filter you are using. For example, using the `branch-misses` CPU instruction metric with the `mispredict` SPE filter. The CPU instruction metrics are both time-based and accurate counts, mitigating some of the limitations mentioned here.

C19.6 Memlock limit exceeded

When using SPE profiling on a system with many logical cores, you can encounter an error regarding an inability to mlock any further memory. This can occur during Arm MAP profiling initialization or when an MPI or user program attempts to mlock memory.

Cause

In order to receive SPE data, Arm MAP must mmap a number of buffers per logical CPUs on the host system. This can sum to a significant amount of memory when the number of logical CPUs is high. In combination with the mlock needs of the target application the system mlock resource limit may be exceeded, preventing more memory from being locked.

Solution

Raise the system soft cap to some higher value using `ulimit -l`.

Solution

Modify the size of the largest buffers Arm MAP maps per core using the environment variable `ALLINEA_SAMPLER_SPE_AUX_BUFFER_SIZE`. This is the number of pages to mmap for each logical core on the system. It must be a positive power of 2.

Setting this environment variable too low reduces the amount of SPE data that can be processed by MAP, reducing the number of SPE samples obtained.

C19.7 SPE on Amazon Web Services

SPE is not available on AWS hosts.

Cause

SPE has been disabled for security reasons on potentially shared hosts.

Solution

Currently, you require a `.metal` Arm-64 instance for SPE profiling.

Solution

If enabled, Kernel Page Table Isolation (KPTI) blocks the use of SPE (see [SPE Prerequisites on page C19-378](#)). The stock Amazon Linux AMI boots with KPTI disabled but the standard Red Hat, Suse, and Ubuntu AMIs do not. Either use the Amazon Linux AMI or ensure that the AMI you are using has KPTI disabled.

C19.8 Known issues

These known issues affect SPE functionality:

- Neoverse N1 hardware has known a sampling bias in its SPE implementation (see errata 1694299, fixed in r4p1). SPE might see unexpectedly high sample counts for branch target instructions and unexpectedly low sample counts for some instructions closely following a branch target.

Part D

Performance Reports

Chapter D1

Get started with Performance Reports

Learn how to get started using Arm Performance Reports.

It contains the following sections:

- *D1.1 Compilers for example programs* on page D1-390.
- *D1.2 Compile on Cray X-series systems* on page D1-391.
- *D1.3 Run an example program* on page D1-392.
- *D1.4 Generate a performance report for an example program* on page D1-393.

D1.1 Compilers for example programs

Arm provides a number of example programs to help you get started using Arm Forge. One of the example programs is a simple 1-D wave equation solver that is useful as a profiling example program. Both C and Fortran variants are provided:

- examples/wave.c
- examples/wave.f90

Both of these variants can be built using the same makefile, wave.makefile. To navigate and run wave.makefile, use:

```
cd {installation-directory}/examples/  
make -f wave.makefile
```

There is also a mixed-mode MPI+OpenMP variant in examples/wave_openmp.c, which is built with the openmp.makefile makefile.

Note

The makefiles for all supplied examples are located in the {installation-directory}/examples directory.

Depending on the default compiler on your system you might see some errors when running the makefile, for example:

```
pgf90-Error-Unknown switch: -fno-inline
```

By default, this example makefile is set up for the GNU compilers. To setup the makefile for a different compiler, open the examples/wave.makefile file, uncomment the appropriate compilation command for the compiler you want to use, and comment those of the GNU compiler.

Notes:

- The compilation commands for other popular compilers are already present in the makefile, separated by compiler.
- Although the example makefiles include the -g flag, Arm Forge does *not* require this. Do not use them in your own makefiles.

In most cases Arm Forge can run on an unmodified binary with no recompilation or linking required.

D1.2 Compile on Cray X-series systems

On Cray X-series systems the example program must either be dynamically linked (using `-dynamic`) or explicitly linked with the Arm profiling libraries.

For example:

```
cc -dynamic -g -O3 wave.c -o wave -lm -lrt
ftn -dynamic -G2 -O3 wave.f90 -o wave -lm -lrt
```

About this task

This example shows you how to explicitly link the example program with the Arm profiling libraries.

Procedure

1. Create the libraries using the command `make-profiler-libraries -platform=cray -lib-type=static`:

```
Created the libraries in /home/user/examples:
libmap-sampler.a
libmap-sampler-mpi.a

To instrument a program, add these compiler options:
compilation for use with MAP - not required for Performance Reports:
-g (or -G2 for native Cray fortran) (and -O3 etc.)
linking (both MAP and Performance Reports):
-Wl,@/home/user/examples/allinea-profiler.ld ... EXISTING_MPI_LIBRARIES
If your link line specifies EXISTING_MPI_LIBRARIES (e.g. -lmpi), then
these must appear **after** the Arm sampler and MPI wrapper libraries in
the link line. There is a comprehensive description of the link ordering
requirements in *(\nameref{sec:MAPS:subsec:Preparing a program for profiling})*.
```

2. Follow the instructions in the output to link the example program with the Arm profiling libraries:

```
cc -g -O3 wave.c -o wave -g -Wl,@allinea-profiler.ld -lm -lrt
ftn -G2 -O3 wave.f90 -o wave -G2 -Wl,@allinea-profiler.ld -lm -lrt
```

D1.3 Run an example program

Describes how to run an example program with MPI.

Before you begin

Make sure that you have compiled the example program. See [Compilers for example programs on page D1-390](#).

About this task

This example uses MPI, so you must run it on a compute node on your cluster. The help pages and support staff on your site can tell you exactly how to do this on your machine. The simplest way when running small programs is often to request an interactive session.

Procedure

```
$ qsub -I
qsub: waiting for job 31337 to start
qsub: job 31337 ready
$ cd *(\forgeInstallPathRaw)*/examples
$ mpiexec -n 4 ./wave_c
Wave solution running with 4 processes
```

Results

If the output is similar to this, the example program is compiled and working correctly.

```
0: points = 1000000, running for 30 seconds
points / second: 63.9M (16.0M per process)
compute / communicate efficiency: 94% | 97% | 100%

Points for validation:
0:0.00 200000:0.95 400000:0.59 600000:-0.59 800000:-0.95 999999:0.00
wave finished
```

Next steps

[Generate a performance report for an example program on page D1-393](#)

D1.4 Generate a performance report for an example program

Describes how to load Performance Reports and then generate a report.

Before you begin

Make sure that the Arm Performance Reports component of Arm Forge installed on your system is loaded:

```
$ perf-report --version
Arm Performance Reports
Copyright (c) 2002-2020 Arm Limited (or its affiliates). All rights reserved.
...
```

Procedure

Type the `perf-report` command in front of your existing `mpiexec` command-line:

```
perf-report mpiexec -n 4 examples/wave_c
```

If your program is submitted through a batch queuing system, modify your submission script to load the Arm module and add the 'perf-report' line in front of the `mpiexec` command for which you want to generate a report.

Results

The program runs as usual, although startup and shutdown might take a few minutes longer while Arm Forge generates and links the appropriate wrapper libraries before running, and collects the data at the end of the run. The runtime of your code (between `MPI_Init` and `MPI_Finalize` is not expected to be affected by more than a few percent at most.

After the run finishes, a performance report is saved to the current working directory, using a name based on the application executable:

```
$ ls -lrt wave_c*
-rwx----- 1 mark mark 403037 Nov 14 03:21 wave_c
-rw----- 1 mark mark 1911 Nov 14 03:28 wave_c_4p_2013-11-14_03-27.txt
-rw----- 1 mark mark 174308 Nov 14 03:28 wave_c_4p_2013-11-14_03-27.html
```

Note

Both `.txt` and `.html` versions are automatically generated.

Chapter D2

Run real programs

This chapter shows you how to compile and run your own programs.

Arm Forge is designed to run on unmodified production executables, so in general no preparation step is necessary. However, there is one important exception: statically linked applications require additional libraries at the linking step.

It contains the following sections:

- [D2.1 Link with a program](#) on page D2-396.
- [D2.2 Dynamic linking on Cray X-Series systems \(Performance Reports\)](#) on page D2-397.
- [D2.3 Static linking \(Performance Reports\)](#) on page D2-398.
- [D2.4 Static linking on Cray X-Series systems \(Performance Reports\)](#) on page D2-400.
- [D2.5 Dynamic and static linking on Cray X-Series systems using the modules environment \(Performance Reports\)](#) on page D2-401.
- [D2.6 map-link modules installation on Cray X-Series systems \(Performance Reports\)](#) on page D2-402.
- [D2.7 Unsupported user applications \(Performance Reports\)](#) on page D2-403.
- [D2.8 Express Launch mode \(Performance Reports\)](#) on page D2-404.
- [D2.9 Compatibility Launch mode](#) on page D2-405.
- [D2.10 Generate a performance report for a real program](#) on page D2-406.
- [D2.11 Specify output locations](#) on page D2-407.
- [D2.12 Support for DCIM systems](#) on page D2-408.
- [D2.13 Enable and disable metrics](#) on page D2-409.

D2.1 Link with a program

To collect data from your program, Performance Reports uses two small profiler libraries, `map-sampler` and `map-sampler-mpi`. These profiler libraries must be linked with your program. On most systems Performance Reports can do this automatically without any action by you. This is done via the system's `LD_PRELOAD` mechanism, which allows an extra library into your program when starting it.

This automatic linking when you start your program only works if your program is dynamically-linked. Programs may be dynamically-linked or statically-linked. For MPI programs this is normally determined by your MPI library. Most MPI libraries are configured with `-enable-dynamic` by default, and `mpicc/mpif90` produce dynamically-linked executables that Arm Performance Reports can automatically collect data from.

The `map-sampler-mpi` library is a temporary file that is precompiled and copied or compiled at runtime in the directory `~/.allinea/wrapper`.

If your home directory will not be accessible by all nodes in your cluster you can change where the `map-sampler-mpi` library will be created by altering the shared directory as described in [No shared home directory on page Appx-H-501](#).

The temporary library will be created in the `allinea/wrapper` subdirectory to this shared directory.

Note

Although the profiler libraries contain the word 'map' they are used for both Arm Performance Reports and Arm MAP.

For Cray X-Series Systems the shared directory is not applicable, instead `map-sampler-mpi` is copied into a hidden `.allinea` sub-directory of the current working directory.

If Arm Performance Reports warns you that it could not pre-load the sampler libraries, this often means that your MPI library was not configured with `-enable-dynamic`, or that the `LD_PRELOAD` mechanism is not supported on your platform. You now have three options:

- Try compiling and linking your code dynamically. On most platforms this allows Arm Performance Reports to use the `LD_PRELOAD` mechanism to automatically insert its libraries into your application at runtime.
- Link MAP's `map-sampler` and `map-sampler-mpi` libraries with your program at link time manually.

See [Dynamic linking on Cray X-Series systems \(Performance Reports\) on page D2-397](#), or [Static linking \(Performance Reports\) on page D2-398](#) and [Static linking on Cray X-Series systems \(Performance Reports\) on page D2-400](#).

- Finally, it may be that your system supports dynamic linking but you have a statically-linked MPI. You can try to recompile the MPI implementation with `-enable-dynamic`, or find a dynamically-linked version on your system and recompile your program using that version. This will produce a dynamically-linked program that MAP can automatically collect data from.

D2.2 Dynamic linking on Cray X-Series systems (Performance Reports)

If the LD_PRELOAD mechanism is not supported on your Cray X-Series system, you can try to dynamically link your program explicitly with the Arm Performance Reports sampling libraries.

Procedure

1. Compile the Arm MPI wrapper library for your system using the `make-profiler-libraries -platform=cray -lib-type=shared` command.

```
user@login:~/myprogram$ make-profiler-libraries --platform=cray --lib-type=shared

Created the libraries in /home/user/myprogram:
libmap-sampler.so      (and .so.1, .so.1.0, .so.1.0.0)
libmap-sampler-mpi.so  (and .so.1, .so.1.0, .so.1.0.0)

To instrument a program, add these compiler options:
compilation for use with MAP - not required for Performance Reports:
  -g (or '-G2' for native Cray Fortran) (and -O3 etc.)
linking (both MAP and Performance Reports):
  -dynamic -L/home/user/myprogram -lmap-sampler-mpi -lmap-sampler -Wl,--eh-frame-hdr

Note: These libraries must be on the same NFS/Lustre/GPFS filesystem as your
program.

Before running your program (interactively or from a queue), set
LD_LIBRARY_PATH:
export LD_LIBRARY_PATH=/home/user/myprogram:$LD_LIBRARY_PATH
map ...
or add -Wl,-rpath=/home/user/myprogram when linking your program.
```

1. Link with the Arm MPI wrapper library.

```
mpicc -G2 -o hello hello.c -dynamic -L/home/user/myprogram \
-lmap-sampler-mpi -lmap-sampler -Wl,--eh-frame-hdr
```

PGI Compiler

When linking OpenMP programs, you must pass the `-Bdynamic` command line argument to the compiler when linking dynamically.

NVIDIA HPC Compiler

When linking OpenMP programs, you must pass the `-Bdynamic` command line argument to the compiler when linking dynamically.

D2.3 Static linking (Performance Reports)

If you compile your program statically, that is your MPI uses a static library or you pass the `-static` option to the compiler, then you must explicitly link your program with the Arm sampler and MPI wrapper libraries.

Procedure

1. Compile the Arm MPI wrapper library for your system using the `make-profiler-libraries --lib-type=static` command.

```
user@login:~/myprogram$ make-profiler-libraries --lib-type=static

Created the libraries in /home/user/myprogram:
libmap-sampler.a
libmap-sampler-pmpi.a

To instrument a program, add these compiler options:
compilation for use with MAP - not required for Performance Reports:
-g (and -O3 etc.)
linking (both MAP and Performance Reports):
-Wl,@/home/user/myprogram/allinea-profiler.ld ... EXISTING_MPI_LIBRARIES
If your link line specifies EXISTING_MPI_LIBRARIES (e.g. -lmpi), then
these must appear after the Arm sampler and MPI wrapper libraries in
the link line. There's a comprehensive description of the link ordering
requirements in the 'Preparing a Program for Profiling' section of either
userguide-forge.pdf or userguide-reports.pdf, located in
/opt/*(\forgeInstallPathRaw)*/doc/.
```

2. Link with the Arm MPI wrapper library. The `-Wl,@/home/user/myprogram/allinea-profiler.ld` syntax tells the compiler to look in `/home/user/myprogram/allinea-profiler.ld` for instructions on how to link with the Arm sampler. Usually this is sufficient, but not in all cases. The rest of this section explains how to manually add the Arm sampler to your link line.

PGI Compiler

The PGI C runtime static library contains an undefined reference to `__kmpc_fork_call`, which will cause compilation to fail when linking `allinea-profiler.ld`. Add `-undefined __wrap__kmpc_fork_call` to your link line before linking to the Arm sampler to resolve this.

NVIDIA HPC Compiler

The NVIDIA HPC C runtime static library contains an undefined reference to `__kmpc_fork_call`, which will cause compilation to fail when linking `allinea-profiler.ld`. To resolve this, add `-undefined __wrap__kmpc_fork_call` to your link line before linking to the Arm sampler.

Cray

When linking C++ programs you may encounter a conflict between the Cray C++ runtime and the GNU C++ runtime used by the Arm Performance Reports libraries with an error similar to the one below:

```
/opt/cray/cce/8.2.5/CC/x86-64/lib/x86-64/libcray-c++-rts.a(rtti.o): In function
'__cxa_bad_typeid':
/ptmp/ulib/buildslaves/cfe-82-edition-build/tbs/cfe/lib_src/rtti.c:1062: multiple definition
of '__cxa_bad_typeid'
/opt/gcc/4.4.4/snos/lib64/libstdc++.a(eh_aux_runtime.o):/tmp/peint/gcc/repackage/4.4.4c/
BUILD/snos_objdir/x86_64-suse-linux/libstdc++-v3/libsupc++/../../../../xt-gcc-4.4.4/libstdc++
-v3/libsupc++/eh_aux_runtime.cc:46: first defined here
```

You can resolve this conflict by removing `-lstdc++` and `-lgcc_eh` from `allinea-profiler.ld`.

-lpthread

When linking `-Wl,@allinea-profiler.ld` must go before the `-lpthread` command-line argument if present.

Manual Linking

When linking your program you must add the path to the profiler libraries (`-L/path/to/profiler-libraries`), and the libraries themselves (`-lmap-sampler-pmpi`, `-lmap-sampler`).

The MPI wrapper library `-lmap-sampler-mpmi` must go:

- *After* your program's object (.o) files.
- *After* your program's own static libraries, for example `-lmylibrary`.
- *After* the path to the profiler libraries (`-L/path/to/profiler-libraries`).
- *Before* the MPI's Fortran wrapper library, if any. For example `-lmpichf`.
- *Before* the MPI's implementation library usually `-lmpi`.
- *Before* the sampler library `-lmap-sampler`.

The sampler library `-lmap-sampler` must go:

- *After* the Arm MPI wrapper library.
- *After* your program's object (.o) files.
- *After* your program's own static libraries, for example `-lmylibrary`.
- *After* `-Wl,-undefined,allinea_init_sampler_now`.
- *After* the path to the profiler libraries `-L/path/to/profiler-libraries`.
- *Before* `-lstdc++`, `-lgcc_eh`, `-lrt`, `-lpthread`, `-ldl`, `-lm` and `-lc`.

For example:

```
mpicc hello.c -o hello -g -L/users/ddt/arm \
-lmap-sampler-mpmi \
-Wl,--undefined,allinea_init_sampler_now \
-lmap-sampler -lstdc++ -lgcc_eh -lrt \
-Wl,--whole-archive -lpthread \
-Wl,--no-whole-archive \
-Wl,--eh-frame-hdr \
-ldl \
-lm

mpif90 hello.f90 -o hello -g -L/users/ddt/arm \
-lmap-sampler-mpmi \
-Wl,--undefined,allinea_init_sampler_now \
-lmap-sampler -lstdc++ -lgcc_eh -lrt \
-Wl,--whole-archive -lpthread \
-Wl,--no-whole-archive \
-Wl,--eh-frame-hdr \
-ldl \
-lm
```

D2.4 Static linking on Cray X-Series systems (Performance Reports)

If you compile your program statically on a Cray X-Series system, you must explicitly link your program with the Arm sampler and MPI wrapper libraries.

Procedure

1. On Cray X-Series systems, you can compile the MPI wrapper library using `make-profiler-libraries -platform=cray -lib-type=static`:

```
Created the libraries in /home/user/myprogram:
  libmap-sampler.a
  libmap-sampler-pmpi.a

To instrument a program, add these compiler options:
  compilation for use with MAP - not required for Performance Reports:
    -g (or -G2 for native Cray Fortran) (and -O3 etc.)
  linking (both MAP and Performance Reports):
    -Wl,@/home/user/myprogram/allinea-profiler.ld ... EXISTING_MPI_LIBRARIES
If your link line specifies EXISTING_MPI_LIBRARIES (e.g. -lmpi), then
these must appear after the Arm sampler and MPI wrapper libraries in
the link line. There's a comprehensive description of the link ordering
requirements in the 'Preparing a Program for Profiling' section of either
userguide-forge.pdf or userguide-reports.pdf, located in
/opt/*(\forgeInstallPathRaw)*/doc/.
```

1. Link with the MPI wrapper library using:

```
cc hello.c -o hello -g -Wl,@allinea-profiler.ld
ftn hello.f90 -o hello -g -Wl,@allinea-profiler.ld
```


D2.5 Dynamic and static linking on Cray X-Series systems using the modules environment (Performance Reports)

If your system has the Arm module files installed, you can load them and build your application as usual. See [map-link modules installation on Cray X-Series systems \(Performance Reports\)](#).

Procedure

1. `module load forge` or ensure that `make-profiler-libraries` is in your `PATH`.
2. `module load map-link-static` or `module load map-link-dynamic`.
3. Recompile your program.

D2.6 map-link modules installation on Cray X-Series systems (Performance Reports)

To facilitate dynamic and static linking of your programs with the Arm MPI Wrapper and Sampler libraries, Cray X-Series System Administrators can integrate the map-link-dynamic and map-link-static modules into their module system. Templates for these modules are supplied as part of the Arm Forge package.

Procedure

1. Copy files `share/modules/cray/map-link-_*` into a dedicated directory on the system.
2. For each of the two module files copied, find the line starting with **conflict** and correct the prefix to refer to the location the module files were installed, for example, `arm/map-link-static`. The correct prefix depends on the subdirectory (if any) under the module search path the `map-link-_*` modules were installed.
3. For each of the two module files copied, find the line starting with **set**
MAP_LIBRARIES_DIRECTORY "NONE" and replace **"NONE"** with a user writable directory accessible from the login and compute nodes.

Results

After installation you can verify if the prefix has been set correctly using `'module avail'`. The prefix shown by this command for the `map-link-*` modules should match the prefix set in the `'conflict'` line of the module sources.

D2.7 Unsupported user applications (Performance Reports)

Ensure that the program to be profiled does not set or unset the SIGPROF signal handler. This interferes with the Arm Performance Reports profiling function and can cause it to fail.

We recommend that you do not use Arm Performance Reports to profile programs that contain instructions to perform MPI profiling using MPI wrappers and the MPI standard profiling interface, PMPI. This is because Performance Report's own MPI wrappers may conflict with those contained in the program, producing incorrect metrics.

D2.8 Express Launch mode (Performance Reports)

Arm Forge can be launched by typing its command name in front of an existing `mpiexec` command:

```
$ perf-report mpiexec -n 256 examples/wave_c 30
```

Compatible MPIs

The MPI implementations supported by Express Launch are:

- Bullx MPI
- Cray X-Series (MPI/SHMEM/CAF)
- Intel MPI
- MPICH 3
- Open MPI (MPI/SHMEM)
- Oracle MPT
- Open MPI (Cray XT/XE/XK)
- Spectrum MPI
- Spectrum MPI (PMIx)
- Cray XT/XE/XK (UPC)

If your MPI is not supported by Express Launch, an error message will display:

```
$ 'Generic' MPI programs cannot be started using Express Launch syntax (launching with an mpirun command).
```

```
Try this instead:  
perf-report --processes=256 ./wave_c 20
```

```
Type perf-report --help for more information.
```

This is referred to as Compatibility Mode, in which the `mpiexec` command is not included and the arguments to `mpiexec` are passed via a `-mpiargs="args here"` parameter.

One advantage of Express Launch mode is that it is easy to modify existing queue submission scripts to run your program under one of the Arm Forge products.

Normal redirection syntax may be used to redirect standard input and standard output.

D2.9 Compatibility Launch mode

Compatibility Launch mode must be used if Arm Forge does not support Express Launch mode for your MPI, or, for some MPIs, if it is not able to access the compute nodes directly (for example, using `ssh`).

To use Compatibility Launch mode, replace the `mpiexec` command with the `perf-report` command.

For example:

```
mpiexec --np=256 ./wave_c 20
```

becomes

```
perf-report --np=256 ./wave_c 20
```

Only a small number of `mpiexec` arguments are supported by `perf-report` (for example, `-n` and `-np`).

Other arguments must be passed using the `-mpiargs="args here"` parameter.

For example:

```
mpiexec --np=256 --nooversubscribe ./wave_c 20
```

becomes

```
perf-report --mpiargs="--nooversubscribe" --np=256 ./wave_c 20
```

Normal redirection syntax may be used to redirect standard input and standard output.

D2.10 Generate a performance report for a real program

Describes how to load Performance Reports and then generate a report.

Before you begin

Make sure that the Performance Reports component of Arm Forge installed on your system is loaded:

```
$ perf-report --version
Arm Performance Reports
Copyright (c) 2002-2020 Arm Limited (or its affiliates). All rights reserved.
...
```

If this command cannot be found, consult the site documentation to find the name of the correct module.

Procedure

Type the `perf-report` command in front of your existing `mpiexec` command-line:

```
perf-report mpiexec -n 4 examples/wave_c
```

If your program is submitted through a batch queuing system, then modify your submission script to load the Arm module and add the 'perf-report' line in front of the `mpiexec` command for which you want to generate a report.

Results

The program runs as usual, although startup and shutdown may take a few minutes longer while Arm Forge generates and links the appropriate wrapper libraries before running and collects the data at the end of the run. The runtime of your code (between `MPI_Init` and `MPI_Finalize` should not be affected by more than a few percent at most.

After the run finishes, a performance report is saved to the current working directory, using a name based on the application executable:

```
$ ls -lrt wave_c*
-rwx----- 1 mark mark 403037 Nov 14 03:21 wave_c
-rw----- 1 mark mark 1911 Nov 14 03:28 wave_c_4p_2013-11-14_03-27.txt
-rw----- 1 mark mark 174308 Nov 14 03:28 wave_c_4p_2013-11-14_03-27.html
```

Note

Both `.txt` and `.html` versions are automatically generated.

You can include a short description of the run or other notes on configuration and compilation settings by setting the environment variable `ALLINEA_NOTES` before running `perf-report`:

```
$ ALLINEA_NOTES="Run with inp421.dat and mc=1" perf-report mpiexec -n 512 ./parEval.bin --
use-mc=1 inp421.dat
```

The string in the `ALLINEA_NOTES` environment variable is included in all report files produced.

D2.11 Specify output locations

By default, performance reports are placed in the current working directory using an auto-generated name based on the application executable name, for example:

```
wave_f_16p_2013-11-18_23-30.html  
wave_f_2p_8t_2013-11-18_23-30.html
```

This is formed by the name, the size of the job, the date, and the time. If you are using OpenMP, the value of `OMP_NUM_THREADS` is also included in the name after the size of the job. The name will be made unique if necessary by adding a `_1/_2` suffix.

You can specify a different location for output files using the `-output` argument.

For example:

- `-output=my-report.txt` will create a plain text report in the file `my-report.txt` in the current directory.
- `-output=/home/mark/public/my-report.html` will create an HTML report in the file `/home/mark/public/my-report.html`.
- `-output=my-report` will create a plain text report in the file `my-report.txt` and an HTML report in the file `my-report.html`, both in the current directory.
- `-output=/tmp` will create reports with names based on the application executable name in `/tmp/`, for example, `/tmp/wave_f_16p_2013-11-18_2330.txt` and `/tmp/wave_f_16p_2013-11-18_2330.html`.

D2.12 Support for DCIM systems

Performance Reports includes support for Data Center Infrastructure Management (DCIM) systems.

You can output all the metrics generated by the Performance Reports to a script using the `-dcim-output` argument. By default, the `pr-dcim` script is called and the collected metrics are sent to Ganglia (a system monitoring tool).

The `pr-dcim` script looks for a `gmetric` implementation as part of the Ganglia software, and calls it as many times as there are metrics.

Customize your DCIM script

The default `pr-dcim` script is located in `{installation-directory}/performance-reports/ganglia-connector/pr-dcim`.

However, you can use your own custom script by specifying the `ALLINEA_DCIM_SCRIPT` environment variable.

This option is recommended if you are using a system monitoring tool other than Ganglia.

Use arguments as described below in your scripts. Each argument can be specified once per metric:

- `-V{METRIC}={VALUE}` (mandatory) specifies that the metric `METRIC` has the value `VALUE`.
- `-U{METRIC}={UNITS}` (optional) specifies that the metric `METRIC` is expressed in `UNITS`.
- `-T{METRIC}={TITLE}` (optional) specifies that the metric `METRIC` has title `TITLE`.
- `-t{METRIC}={TYPE}` (optional) specifies that the metric `METRIC` has `TYPE` data type.

Customize the gmetric location

You can specify the path to your `gmetric` implementation by using the `ALLINEA_GMETRIC` environment variable.

Your `gmetric` version must accept the following command-line arguments:

- `-n {NAME}` (mandatory) specifies the name of the metric (starts with `com.allinea`).
- `-t {TYPE}` (mandatory) specifies the type of the metric (for example, `double` or `int32`).
- `-v {VALUE}` (mandatory) specifies the value of the metric.
- `-g {GROUP}` (optional) specifies which groups the metric belongs to (for example `allinea`).
- `-u {UNIT}` (optional) specifies the unit of the metric. For example, `%`, `Watts`, `Seconds`, and so on.
- `-T {TITLE}` (optional) specifies the title of the metric.

D2.13 Enable and disable metrics

You can specify comma-separated lists which explicitly enable or disable metrics for which data is to be collected.

`-enable-metrics=METRICS -disable-metrics=METRICS`

If the metrics specified cannot be found, an error message is displayed and Performance Reports exits. Metrics which are always enabled or disabled cannot be explicitly disabled or enabled. A metrics source library which has all its metrics disabled, either in the XML definition or via `-disable-metrics`, will not be loaded. Metrics which can be explicitly enabled or disabled can be listed using the `-list-metrics` option.

Chapter D3

Summarize an existing MAP file

Arm Performance Reports can be used to summarize an application profile generated by Arm MAP.

It contains the following section:

- [D3.1 Summarize an existing MAP file on page D3-412.](#)

D3.1 Summarize an existing MAP file

To produce a performance report from an existing MAP output file called `profile.map`, simply run:

```
$ perf-report profile.map
```

Command-line options which would alter the execution of a program being profiled, such as specifying the number of MPI ranks, have no effect. Options affecting how Performance Reports produces its report, such as `-output`, work as expected.

For the best results, ensure that Performance Reports and MAP versions match, for example, Performance Reports 20.2.1 with MAP 20.2.1. Performance Reports can use MAP files from versions of MAP as old as 5.0.

Chapter D4

Interpret performance reports

This chapter explains how to interpret the reports produced by Arm Performance Reports.

Reports are generated in both HTML and textual formats for each run of your application, by default. The information presented in both of these formats is the same.

If you want to combine Arm Forge with other tools, consider using the CSV output format. See *CSV performance reports* for more details.

It contains the following sections:


- [D4.1 HTML performance reports](#) on page D4-414.
- [D4.2 Report summary](#) on page D4-416.
- [D4.3 CPU breakdown](#) on page D4-417.
- [D4.4 CPU metrics breakdown](#) on page D4-418.
- [D4.5 MPI breakdown](#) on page D4-420.
- [D4.6 I/O breakdown](#) on page D4-421.
- [D4.7 OpenMP breakdown](#) on page D4-423.
- [D4.8 Threads breakdown](#) on page D4-424.
- [D4.9 Memory breakdown](#) on page D4-425.
- [D4.10 Energy breakdown](#) on page D4-426.
- [D4.11 Accelerator breakdown](#) on page D4-427.
- [D4.12 Textual performance reports](#) on page D4-428.
- [D4.13 CSV performance reports](#) on page D4-429.
- [D4.14 Worked examples](#) on page D4-430.

D4.1 HTML performance reports


Viewing HTML files is best done on your local machine. Many sites have places you can put HTML files to be viewed from within the intranet. These directories are a good place to automatically send your performance reports. Alternatively, you can use `scp` or `sshfs` to make the reports available on your computer:

```
$ scp login1:*(\forgeInstallPathRaw)*/examples/wave_c_4p*.html .
$ firefox wave_c_4p*.html
```

The following report was generated by running the `wave_openmp.c` example program with 8 MPI processes and 2 OpenMP threads per process on a typical HPC cluster:



Summary: `wave_openmp` is **Compute-bound** in this configuration



This application run was **Compute-bound**. A breakdown of this time and advice for investigating further is in the **CPU** section below.

As little time is spent in **MPI** calls, this code may also benefit from running at larger scales.

CPU

A breakdown of the 72.6% CPU time:




The per-core performance is **memory-bound**. Use a profiler to identify time-consuming loops and check their cache performance.

No time is spent in **vectorized instructions**. Check the compiler's vectorization advice to see why key loops could not be vectorized.

MPI

A breakdown of the 27.4% MPI time:



Most of the time is spent in **point-to-point calls** with a very low transfer rate. This suggests load imbalance is causing synchronization overhead; use an MPI profiler to investigate.

Figure D4-1 Performance report for the `wave_openmp.c` example

I/O

A breakdown of the 0.0% I/O time:

Time in reads	0.0%	
Time in writes	0.0%	
Effective process read rate	0.00 bytes/s	
Effective process write rate	0.00 bytes/s	

No time is spent in I/O operations. There's nothing to optimize here!

Memory

Per-process memory usage may also affect scaling:

Mean process memory usage	38.6 MiB	
Peak process memory usage	53.7 MiB	
Peak node memory usage	17.0%	

The **peak node memory usage** is very low. Running with fewer MPI processes and more data on each process may be more efficient.

OpenMP

A breakdown of the 91.8% time in OpenMP regions:

Computation	9.9%	
Synchronization	90.1%	
Physical core utilization	100.0%	
System load	167.0%	

Significant time is spent **synchronizing** threads in parallel regions. Check the affected regions with a profiler.

The system load is high. Ensure background system processes are not running.

Energy

A breakdown of how energy was used:

CPU	not supported %	
System	not supported %	
Mean node power	not supported W	
Peak node power	0.00 W	

Energy metrics are not available on this system.

CPU metrics are not supported (no intel_rapl module)

Figure D4-2 Performance report for the wave_openmp.c example (continued)

When you run a report on this example program, the results might be different to the report shown here depending on the performance and network architecture of the machine on which you run it, but the basic structure of these reports is always the same. This common structure makes comparisons between reports simple, direct, and intuitive. Each section of the report is described in the following sections.

D4.2 Report summary

The Summary shows how the wall clock time of the application was spent. It is organized by Compute, MPI, and I/O.

In the example file, you can see that Arm Forge has identified that the program is compute-bound, which means that most of its time is spent inside application code rather than communicating or using the filesystem.

The pieces of advice that the program offers, such as `this code may benefit from running at larger scales`, are good starting points for future investigations. They are designed to be informative to scientific users with no previous MPI tuning experience.

The triangular radar chart in the top-right corner of the report, reflects the values of these three key measurements: compute, MPI and I/O. It is helpful to recognize and compare these triangular shapes when switching between multiple reports.

Compute

Time spent computing. This is the percentage of wall clock time spent in application and in library code, excluding time spent in MPI calls and I/O calls.

MPI

Time spent communicating. This is the percentage of wall clock time spent in MPI calls such as `MPI_Send`, `MPI_Reduce` and `MPI_Barrier`.

I/O (Input/Output)

Time spent reading from and writing to the filesystem. This is the percentage of wall clock time spent in system library calls such as `read`, `write` and `close`.

Note

All time spent in MPI-IO calls is included here, even though some communication between processes might also be performed by the MPI library. `MPI_File_close` is treated as time spent writing, which is often, but not always, correct.

D4.3 CPU breakdown

This section organizes the time spent in application and library code further by analyzing the kinds of instructions that this time was spent on.

Note

All of the metrics described in this section are only available on x86_64 systems.

Note

All percentages here are relative to the compute time, not to the entire application run. Time spent in MPI and I/O calls is not represented inside this section.

Single-core code

The percentage of wall clock in which the application executed using only one core per process, rather than multithreaded or OpenMP code. If you have a multithreaded or OpenMP application, a high value here indicates that your application is bound by Amdahl's law and that scaling to larger numbers of threads will not meaningfully improve performance.

OpenMP regions

The percentage of wall clock time spent in OpenMP regions. The higher this is, the better. This metric is only shown if the program spent a measurable amount of time inside at least one OpenMP region.

Scalar numeric ops

The percentage of time spent executing arithmetic operations such as `add`, `mul`, `div`. This does not include time spent using the more efficient vectorized versions of these operations.

Vector numeric ops

The percentage of time spent executing vectorized arithmetic operations such as Intel's SSE2 / AVX extensions.

Generally it is good if a scientific code spends most of its time in these operations because that is the only way to achieve anything close to the peak performance of modern processors.

If this value is low, you can check the vectorization report of the compiler to understand why the most time consuming loops are not using these operations. Compilers need a lot of help to efficiently vectorize non-trivial loops and the investment in time is often rewarded with 2x-4x performance improvements.

Memory accesses

The percentage of time spent in memory access operations, such as `mov`, `load`, `store`. A portion of the time spent in instructions that use indirect addressing is also included here. A high figure here shows the application is memory-bound and is not able to take full advantage of the CPU resources. Often it is possible to reduce this figure by analyzing loops for poor cache performance and problematic memory access patterns, improving performance significantly.

A high percentage of time spent in memory accesses in an OpenMP program is often a scalability problem. If each core spends most of its time waiting for memory, or the L3 cache, then adding further cores rarely improves matters. Equally, false sharing, in which cores block attempts to access the same cache lines, and the over-use of the `atomic` pragma, show up as increased time spent in memory accesses.

Waiting for accelerators

The percentage of time that the CPU is waiting for the accelerator.

D4.4 CPU metrics breakdown

This section presents key CPU performance measurements gathered using the Linux perf event subsystem.

Note

Metrics described in this section are only available on Armv8 and IBM Power systems. These metrics are not available on virtual machines. Linux perf events performance events counters must be accessible on all systems on which the target program runs. See *Armv8 (AArch64) known issues* or *POWER8 and POWER9 (POWER 64-bit)* in [Platform notes and known issues on page Appx-G-481](#).

Cycles per instruction

The average amount of CPU cycles lapsed for each retired instruction. This metric is affected by CPU frequency scaling and various issues, particularly hardware interrupt counts.

Stalled cycles

Note

This metric is available on Armv8 and IBM Power 9 systems only.

The percentage of CPU cycles that lapsed, on which operation instructions are not issued.

L2 cache misses

Note

This metric is available on Armv8 systems only.

The percentage of L2 data cache accesses which resulted in a miss.

L3 cache miss per instruction

Note

This metric is available on IBM Power 9 systems only.

The ratio of L3 data cache demand loads to instructions completed.

FLOPS scalar lower bound

This is a lower bound because its value is calculated from FLOPS vector lower bound, which does not account for the length of vector operations.

Note

This metric is available on IBM Power 8 systems only.

The rate at which floating-point scalar operations finished.

FLOPS vector lower bound

This is a lower bound because the counted value does not account for the length of vector operations.

Note

This metric is available on IBM Power 8 systems only.

The rate at which vector floating-point instructions completed.

Memory accesses (IBM Power 8)

Note

This metric is available on IBM Power 8 systems only.

The rate at which the processor's data cache reloaded from a memory location, including L4 from local, remote, or distant due to a demand load.

D4.5 MPI breakdown

This section organizes the time spent in MPI calls reported in the summary. It is only interesting if the program spends a significant amount of its time in MPI calls.

All the rates quoted here are inbound and outbound rates. This means that the rate of communication is being measured from the process to the MPI API, not of the underlying hardware directly.

This application-perspective is found throughout Arm Forge, and in this case allows the results to capture effects such as faster intra-node performance, zero-copy transfers, and other effects.

Note

For programs that make MPI calls from multiple threads (MPI is in `MPI_THREAD_SERIALIZED` or `MPI_THREAD_MULTIPLE` mode), Arm Forge only displays metrics for MPI calls made on the main thread.

Time in collective calls

The percentage of time spent in collective MPI operations such as `MPI_Scatter`, `MPI_Reduce`, and `MPI_Barrier`.

Time in point-to-point calls

The percentage of time spent in point-to-point MPI operations such as `MPI_Send` and `MPI_Recv`.

Effective process collective rate

The average transfer per-process rate during collective operations, from the perspective of the application code and not the transfer layer. For example, an `MPI_Alltoall` that takes 1 second to send 10 Mb to 50 processes and receive 10 Mb from 50 processes has an effective transfer rate of $10 \times 50 \times 2 = 1000$ Mb/s.

Collective rates can often be higher than the peak point-to-point rate if the network topology matches the application's communication patterns well.

Effective process point-to-point rate

The average per-process transfer rate during point-to-point operations, from the perspective of the application code and not the transfer layer. Asynchronous calls that allow the application to overlap communication and computation such as `MPI_IRecv` can achieve much higher effective transfer rates than synchronous calls.

Overlapping communication and computation is often a good strategy to improve application performance and scalability.

D4.6 I/O breakdown

This section organizes the amount of time spent in library and system calls relating to I/O, such as read, write and close. I/O that is generated by MPI network traffic is not included. In most cases, this should be a direct measure of the amount of time spent reading and writing to the filesystem, whether local or networked.

Some systems, such as the Cray X-series, do not have I/O accounting enabled for all filesystems. On these systems only Lustre I/O is reported in this section.

Even if your application does not perform I/O, a non-zero amount of I/O is reported because of internal I/O performed by Arm Performance Reports.

Time in reads

The percentage of time spent on average in read operations from the perspective of the application, not the filesystem. Time spent in the `stat` system call is also included here.

Time in writes

The percentage of time spent on average in write and sync operations from the perspective of the application, not the filesystem.

Opening and closing files is also included here, because measurements have shown that the latest networked filesystems can spend significant amounts of time opening files with create or write permissions.

Effective process read rate

The average transfer rate during read operations from the perspective of the application. A cached read has a much higher read rate than one that has to hit a physical disk. This is particularly important to optimize for because current clusters often have complex storage hierarchies with multiple levels of caching.

Effective process write rate

The average transfer rate during write and sync operations from the application's perspective. A buffered write will have a much higher write rate than one that has to hit a physical disk. However, unless there is significant time between writing and closing the file, the penalty will be paid during the synchronous close operation instead. All these complexities are captured in this measurement.

Lustre metrics

Lustre metrics are enabled if your compute nodes have one or more Lustre filesystems mounted. Lustre metrics are obtained from a Lustre client process that runs on each node. Therefore, the data gives the information gathered on a per-node basis. The data is also cumulative over all of the processes run on a node, not only the application being profiled. Consequently, there might be some data reported to be read and written, even if the application itself does not perform file I/O through Lustre.

However, an assumption is made that the majority of data that is read and written through the Lustre client will be from an I/O intensive application, not from background processes. This assumption has been observed to be reasonable. For generated application profiles with more than a few megabytes of data that is read or written, almost all of the data reported in Arm Performance Reports is attributed to the application being profiled.

The data that is gathered from the Lustre client process is the read and write rate of data to Lustre, and a count of some metadata operations. Lustre does not just store pure data, but associates this data with metadata, which describes where data is stored on the parallel filesystem and how to access it. This metadata is stored separately from data, and needs to be accessed whenever new files are opened, closed, or files are resized. Metadata operations consume time and add to the latency in accessing the data. Therefore, frequent metadata operations can slow down the performance of I/O to Lustre.

Arm Performance Reports reports on the total number of metadata operations, and also the total number of file opens that are encountered by a Lustre client. With the information provided in Arm Performance Reports you can observe the rate at which data is read and written to Lustre through the Lustre client, and also identify whether a slow read or write rate can be correlated to a high rate of expensive metadata operations.

Notes:

- For jobs run on multiple nodes, the reported values are the mean across the nodes.
- If you have more than one Lustre filesystem mounted on the compute nodes, the values are summed across all Lustre filesystems.
- Metadata metrics are only available if you have the Advanced Metrics Pack add-on for Arm Performance Reports.

Lustre read transfer:

The number of bytes read per second from Lustre.

Lustre write transfer:

The number of bytes written per second to Lustre.

Lustre file opens:

The number of file open operations per second on a Lustre filesystem.

Lustre metadata operations:

The number of metadata operations per second on a Lustre filesystem. Metadata operations include file open, close, and create, as well as operations such as readdir, rename, and unlink.

————— **Note** —————

Depending on the circumstances and implementation, ‘file open’ might count as multiple operations, for example, when it creates a new file or truncates an existing one.

D4.7 OpenMP breakdown

This section breaks down the time spent in OpenMP regions into computation and synchronization and includes additional metrics that help to diagnose OpenMP performance problems. It is only shown if a measurable amount of time was spent inside OpenMP regions.

Computation

The percentage of time threads in OpenMP regions that is spent computing rather than waiting or sleeping. Keeping this high is one important way to ensure that OpenMP codes scale well. If this is high, look at the CPU breakdown to see whether that time is being used optimally on floating-point operations for example, or whether the cores are mostly waiting for memory accesses.

Synchronization

The percentage of time threads in OpenMP regions spent waiting or sleeping. By default, each OpenMP region ends with an implicit barrier. If the workload is imbalanced and some threads finish sooner and wait, this value will increase. Also, there is some overhead associated with entering and leaving OpenMP regions and a high synchronization time might show that the threading is too fine-grained. In general, OpenMP performance is better when outer loops are parallelized, rather than inner loops.

Physical core utilization

Modern CPUs often have multiple *logical* cores for each *physical* core. This is often referred to as hyper-threading. These logical cores can share logic and arithmetic units. Some programs perform better when using additional logical cores, but most HPC codes do not.

If the value here is greater than 100, OMP_NUM_THREADS is set to a larger number of threads than physical cores that are available and performance can be impacted, usually appearing as a larger percentage of time in OpenMP synchronization or memory accesses.

System load

The number of active (running or runnable) threads as a percentage of the number of physical CPU cores that are present in the compute node. This value can exceed 100% if you are using hyper-threading, the cores are *oversubscribed*, or other system processes and daemons start running and take CPU resources away from your program. A value consistently less than 100% might indicate your program is not taking full advantage of the CPU resources available on a compute node.

D4.8 Threads breakdown

This section organizes the time spent by worker threads (non-main threads) into computation and synchronization, and includes additional metrics that help to diagnose multicore performance problems. This section is replaced by the OpenMP Breakdown if a measurable amount of application time was spent in OpenMP regions.

Computation (Threads)

The percentage of time that worker threads spend computing rather than waiting in locks and synchronization primitives. If this is high, look at the CPU breakdown to see whether that time is used optimally on floating-point operations for example, or whether the cores are mostly waiting for memory accesses.

Synchronization (Threads)

The percentage of time worker threads spend waiting in locks and synchronization primitives. This only includes time in which those threads were active on a core and does not include time spent sleeping while other useful work is being done. A large value here indicates a performance and scalability problem that can be detected with a multicore profiler such as Arm MAP.

Physical core utilization (Threads)

Modern CPUs often have multiple *logical* cores for each *physical* core. This is often referred to as hyper-threading. These logical cores can share logic and arithmetic units. Some programs perform better when using additional logical cores, but most HPC codes do not.

The value here shows the percentage utilization of physical cores. A value over 100% indicates that more threads are executing than there are physical cores, indicating that hyper-threading is in use.

Only threads actively and simultaneously consuming CPU time are included in this metric. A program can have many helper threads that do little except sleep, and are not shown.

System load (Threads)

The number of active (running or runnable) threads as a percentage of the number of physical CPU cores present in the compute node. This value can exceed 100% if you are using hyper-threading, if the cores are *oversubscribed*, or if other system processes and daemons start running and take CPU resources away from your program. A value consistently less than 100% might indicate your program is not taking full advantage of the CPU resources available on a compute node.

D4.9 Memory breakdown

Unlike the other sections, the memory section does not refer to one particular portion of the job. Instead, it summarizes memory usage across all processes and nodes over the entire duration. All of these metrics refer to RSS, meaning physical RAM usage, and not virtual memory usage. Most HPC jobs attempt to stay within the physical RAM of their node for performance reasons.

Mean process memory usage

The average amount of memory used per-process across the entire length of the job.

Peak process memory usage

The peak memory usage that is seen by one process at any moment during the job. If this varies a lot from the mean process memory usage, it might be a sign of either imbalanced workloads between processes or a memory leak within a process.

Note

This is not a true high-watermark, but rather the peak memory seen during statistical sampling. For most scientific codes, this is not a meaningful difference because rapid allocation and deallocation of large amounts of memory is generally avoided for performance reasons.

Peak node memory usage


The peak percentage of memory that is seen being used on any single node during the entire run. If this is close to 100%, swapping might be occurring, or the job might be likely to hit hard system-imposed limits. If this is low, it might be more efficient in CPU hours to run with a smaller number of nodes and a larger workload per node.

D4.10 Energy breakdown

This section shows the energy used by the job, organized by component, such as CPU and accelerators.

Energy

A breakdown of how the 1.77 Wh was used:



Significant time is spent on memory accesses. Reducing the CPU clock frequency could reduce the total energy usage.

Figure D4-3 Energy metrics report

CPU

The percentage of the total energy used by the CPUs.

CPU power measurement requires an Intel CPU with RAPL support, for example Sandy Bridge or newer, and the `intel_rapl` powercap kernel module to be loaded.

Accelerator (CUDA)

The percentage of energy used by the accelerators. This metric is only shown when a CUDA card is present.

System

The percentage of energy used by other components not shown above. If CPU and accelerator metrics are not available, the system energy will be 100%.

Mean node power

The average of the mean power consumption of all the nodes in Watts.

Peak node power

The node with the highest peak of power consumption in Watts.

Requirements

CPU power measurement requires an Intel CPU with RAPL support, for example Sandy Bridge or newer, and the `intel_rapl` powercap kernel module to be loaded.

Node power monitoring is implemented through one of two methods: the Arm IPMI energy agent which can read IPMI power sensors, or the Cray HSS energy counters.

For more information on how to install the Arm IPMI energy agent see *Arm IPMI Energy Agent* in [Queue template script syntax on page Appx-I-519](#). The Cray HSS energy counters are known to be available on Cray XK6 and XC30 machines.

Accelerator power measurement requires a NVIDIA GPU that supports power monitoring. This can be checked on the command-line with `nvidia-smi -q -d power`. If the reported power values are reported as N/A, power monitoring is not supported.

D4.11 Accelerator breakdown

This section shows the utilization of NVIDIA CUDA accelerators by the job.

Accelerators

A breakdown of how accelerators were used:

GPU utilization	47.8%	
Global memory accesses	1.6%	
Mean GPU memory usage	0.8%	
Peak GPU memory usage	0.8%	

GPU utilization is low; identify CPU bottlenecks with a profiler and offload them to the accelerator.

The **peak GPU memory usage** is low. It may be more efficient to offload a larger portion of the dataset to each device.

Figure D4-4 Accelerator metrics report

GPU utilization

The average percentage of the GPU cards working when at least one CUDA kernel is running.

Global memory accesses

The average percentage of time that the GPU cards were reading or writing to global (device) memory.

Mean GPU memory usage

The average amount of memory in use on the GPU cards.

Peak GPU memory usage

The maximum amount of memory in use on the GPU cards.

D4.12 Textual performance reports

Textual performance reports contain the same information as [HTML performance reports](#), but in a format better suited to automatic data extraction and reading from a terminal:

```
Command:      mpiexec -n 16 examples/wave_c 60
Resources:    1 node (12 physical, 24 logical cores per node, 2 GPUs per node available)
Memory:      15 GB per node, 11 GB per GPU
Tasks:       16 processes
Machine:      node042
Started on:   Tue Feb 25 12:14:06 2014
Total time:   60 seconds (1 minute)
Full path:    /home/user/*(\forgeInstallPathRaw)*/examples
Notes:
```

```
Summary: wave_c is compute-bound in this configuration
Compute:      82.4% |=====|
MPI:          17.6% |=|
I/O:          0.0%  |
This application run was compute-bound. A breakdown of this time and advice for
investigating further is found in the compute section below.
Because minimal time is spent in MPI calls, this code might also benefit from running at
larger scales.
...
```

You can use a combination of `grep` and `sed` for extracting and comparing values between multiple runs, or for automatically placing this data into a centralized database.

D4.13 CSV performance reports

A CSV (comma-separated values) output file can be generated using the `-output` argument and specifying a filename with the `.csv` extension:

```
perf-report --output=myFile.csv ...
```

The CSV file will contain lines in a NAME, VALUE format for each of the reported fields. This is convenient for passing to an automated analysis tool, such as a plotting program. It can also be imported into a spreadsheet for analyzing values among executions.

D4.14 Worked examples

The best way to understand how to use and interpret performance reports is by looking at examples. You can download several sets of real-world reports with analysis and commentary from the [Arm Developer website](#).

There are three tutorials available.

Code characterization and run size comparison

A set of runs from well-known HPC codes at different scales showing different problems:

[Characterizing HPC codes and problems](#)

Deeper CPU metric analysis

A look at the impact of hyper-threading on the performance of a code as seen through the CPU instructions breakdown:

[Exploring hyperthreading](#)

I/O performance bottlenecks

The open source MAD-bench I/O benchmark is run in several different configurations, including on a laptop, and the performance implications are analyzed:

[Understanding I/O behavior](#)

Chapter D5

Configure Perf metrics

The Perf metrics use the Linux kernel `perf_event_open()` system call to provide additional CPU related metrics available for Performance Reports.

They can be used on any system supported by the Linux `perf` command (also called `perf_event`). These cannot be tracked on typical virtual machines.

Perf metrics count the rate of one or more performance events occurring in a program. There are some software events provided by the Linux kernel but most are hardware events tracked by the Performance Monitoring Unit (PMU) of the CPU. *Generalized* hardware events are event name aliases that the Linux kernel identifies.

The quantity and combinations (in some cases) of events that can be simultaneously tracked is limited by the hardware. This feature does not support multiplexing performance events.

If the set of events you requested can not be tracked at the same time, Performance Reports ends the profiling session immediately with an error message. Try requesting fewer events, or a different combination. See the PMU reference manual for your architecture for more information on incompatible events.

It contains the following sections:

- [D5.1 Permissions \(*perf_event Paranoid*\) on page D5-432.](#)
- [D5.2 Probe target hosts on page D5-433.](#)
- [D5.3 Specify Perf metrics via the command line on page D5-434.](#)
- [D5.4 Specify Perf metrics via a file on page D5-435.](#)
- [D5.5 View events on page D5-436.](#)
- [D5.6 Advanced configuration on page D5-437.](#)

D5.1 Permissions (perf_event_paranoid)

On some systems, using the Perf hardware counters can be restricted by the value of `/proc/sys/kernel/perf_event_paranoid`.

perf_event_paranoid	Description
3	Disable use of Perf events
2	Allow only user-space measurements
1	Allow kernel and user-space measurements
0	Allow access to CPU-specific data but not raw trace-point samples
-1	No restrictions

The value of `/proc/sys/kernel/perf_event_paranoid` must be 2 or lower to collect Perf metrics. To set this until the next reboot, run the following commands:

```
sudo sysctl -w kernel.perf_event_paranoid=2
```

To permanently set the paranoid level, add the following line to: `/etc/sysctl.conf`.

```
kernel.perf_event_paranoid=2
```


D5.2 Probe target hosts

You must probe an example of a typical host machine before using these metrics. As well as other properties, this collects the CPU ID used to identify the set of potential hardware events for the host, and tests which generalized events are supported.

Ensure that `/proc/sys/kernel/perf_event_paranoid` is set to 2 or lower (see [Permissions \(perf_event_paranoid\)](#) on page D5-432) before performing the probe.

Note

It is not necessary to probe every potential host, a single compute node in a homogeneous cluster is sufficient.

If your home directory is writable you can generate a probe file and install it in your config directory by running the following on the intended host:

```
/path/to/forge/bin/forge-probe --install=user
```

If the Forge installation directory is writable, you can generate and install the probe file for the current host with:

```
/path/to/forge/bin/forge-probe --install=global
```

To generate the probe file, but install it manually, execute:

```
/path/to/forge/bin/forge-probe
```

The probe is named `<hostname>_probe.json` and is generated in your current working directory. You must manually copy it to the location specified in the `forge-probe` output. This is typically only necessary when the compute node that you are probing does not have write access to your home file system.

Check that the expected probe files are correctly installed with:

```
/path/to/forge/bin/map --target-host=list
```

This shows something like:

```
0x00000000420f5160    (thunderx2)    e.g. node07.myarmhost.com
GenuineIntel-6-4E    (skylake)      e.g. node01.myintelhost.com
```

If you have exactly one probe file installed, this is automatically assumed to be the target host. If there are multiple installed probe files, you must specify the intended target whenever you use the configurable Perf metrics feature. When using the command line, use the `-target-host` argument. You can specify the intended target CPU ID (such as, `0x00000000420f5160`), family name (such as, `thunderx2`), or a unique substring of the hostname (`myarmhost`).

D5.3 Specify Perf metrics via the command line

You can list available events for a given probed host using:

```
/path/to/forge/bin/perf-report map --target-host=myarmhost \  
--perf-metrics=avail
```

Note

Use `list` instead of `avail` to see the events listed on separate lines.

Specify the events you want using a semicolon separated list:

```
/path/to/forge/bin/perf-report --profile --target-host=myarmhost \  
--perf-metrics="cpu-cycles; bus-cycles; instructions" mpirun ...
```

D5.4 Specify Perf metrics via a file

The `-perf-metrics` argument can also take the name of a plain text file:

```
/path/to/forge/bin/perf-report --profile --target-host=myhost \  
--perf-metrics=./myevents.txt mpirun ...
```

`myevents.txt` lists the events to track on separate lines, such as:

```
cpu-cycles  
bus-cycles  
instructions
```

`-perf-metrics=template` outputs a more complex template that lists all possible events with accompanying descriptions. Redirect this output to a file and uncomment the events to track, for example:

```
/path/to/forge/bin/perf-report --target-host=myhost \  
--perf-metrics=template > myevents.txt  
  
vim myevents.txt  
  
/path/to/forge/bin/perf-report --profile \  
--perf-metrics=myevents.txt mpirun ...
```

D5.5 View events

You can view Perf event counts in the **CPU Metrics** section. All these metrics are reported as events per second with a suitable SI prefix (such as, K, M, G) that is automatically determined.

The default values that are reported are the mean of means:

- The mean value is taken across all processes for each sample (averaging across processes).
- The mean value is taken of those per-sample results (averaging across time).

D5.6 Advanced configuration

You can override the default settings used by Performance Reports when making `perf_event_open` calls. Specify one or more flags in a preamble section in square brackets at the start of the perf metrics definition string (either on the command line or at the top of a template file).

```
/path/to/forg/bin/perf-report --profile --target-host=myarmhost \  
--perf-metrics="[optional,noinherit]; instructions; cpu-cycles"
```

Possible options are:

- **[optional]:** Do not abort the program if the requested metrics cannot be collected. Set this if you wish to continue profiling even if the no Perf metric results is returned.
- **[noinherit]:** Disable multithreading support (new threads will not inherit the event counter configuration). If you specified events, they are only collected on the main thread (in the case of MPI programs, the thread that called `MPI_thread_init`).
- **[nopinned]:** Disable pinning events on the PMU. If you have specified this, event counting might be multiplexed. Arm does not recommend doing this as it interacts poorly with the Forge sampling strategy.
- **[noexclude=kernel]:** Do not exclude kernel events that happen in kernel space. This might require a more permissive `perf_event_paranoid` level.
- **[noexclude=hv]:** Do not exclude events that happen in the hypervisor. This is mainly for PMUs that have built-in support for handling this (such as IBM Power). Most machines require extra support for handling hypervisor measurements.
- **[noexclude=idle]:** Do not exclude counting software events when the CPU is running the idle task. This is only relevant for software events.

Part E

Appendix

Appendix A

Configuration

Arm Forge shares a common configuration file between Arm DDT, Arm MAP, and Arm Performance Reports. This makes it easy for you to switch between tools without reconfiguring your environment each time.

It contains the following sections:

- *A.1 Configuring Performance Reports* on page Appx-A-442.
- *A.2 Configuration files* on page Appx-A-443.
- *A.3 Integration with queuing systems* on page Appx-A-445.
- *A.4 Template tutorial* on page Appx-A-446.
- *A.5 Connecting to compute nodes and remote programs (remote-exec)* on page Appx-A-448.
- *A.6 Optional configuration* on page Appx-A-449.

A.1 Configuring Performance Reports

Arm Performance Reports generally requires no extra configuration before use. If you only intend to use Arm Performance Reports and you have verified that it works on your system, you can safely ignore most of the information in this section.

When Arm Performance Reports needs to access another machine as part of starting MPICH 3, Intel MPI, or SGI MPT, it attempts to use the `ssh` secure shell by default. However, this might not always be appropriate if `ssh` is disabled or running on a different port to the default port 22. If startup fails, see [Connecting to compute nodes and remote programs \(remote-exec\)](#) on page Appx-A-448.

A.2 Configuration files

Arm Forge uses two configuration files: the system-wide `system.config` file and the user specific `user.config`. The system-wide configuration file specifies properties such as MPI implementation. The user specific configuration file describes user's preferences such as font size. The files are controlled by environment variables:

Environment Variable	Default
ALLINEA_USER_CONFIG	\${ALLINEA_CONFIG_DIR}/user.config
ALLINEA_SYSTEM_CONFIG	\${ALLINEA_CONFIG_DIR}/system.config
ALLINEA_CONFIG_DIR	\${HOME}/.allinea

Sitewide configuration

If you are the system administrator, or have write-access to the installation directory, you can provide a configuration file which other users are automatically given a copy of the first time that they start Arm Forge. In this case, users no longer need to provide configuration for site-specific aspects such as queue templates and job submission.

Configure Arm Forge normally and run a test program to make sure all the settings are correct. When you are satisfied with your configuration, execute this commands

```
forge --clean-config
```

The `--clean-config` option removes any user-specific settings from your system configuration file and creates a `system.config` file that can provide the default settings for all users on your system. Instructions about how to do this are printed when `--clean-config` completes.

Note

Only the `system.config` file is generated. Arm Forge also uses a user-specific `user.config` which is not affected.

If you want to use to attach to running jobs, you must also create a file called `nodes` in the installation directory which lists the compute nodes to which you want to attach. See [Attach to running programs on page B1-73](#) for details.

Startup scripts

At startup, Arm Forge searches for a sitewide startup script called `allinearc` in the root of the installation directory. If this file exists, the software sources it and then starts the tool. When using the remote client, the software sources this startup script, and then starts any sitewide remote-init remote daemon startup script.

Similarly, you can also provide a user-specific startup script in `~/.allinea/allinearc`.

Note

If the environment variable is set, the software looks in `/allinearc` instead. When using the remote client, the software sources the user-specific startup script followed by the user-specific `~/.allinea/remote-init` remote daemon startup script.

Importing legacy configuration

If you have used a version of Arm DDT prior to version 4.0, your existing configuration is imported automatically. If the `DDTCONFIG` environment variable is set, or you use the `-config` command-line argument, the existing configuration is imported. However, the legacy configuration file is not modified, and subsequent configuration changes are saved as described in the previous sections.

Converting legacy sitewide configuration files

If you have existing sitewide configuration files from a version of Arm DDT prior to 4.0 you must convert them to the new 4.0 format. You can do this using the following command line:

Note

Ensure that `newconfig.ddt` does not exist before you submit the command.

```
forge --config=oldconfig.ddt --system-config=newconfig.ddt --clean-config
```

Using shared home directories on multiple systems

If your site uses the same home directory for multiple systems you might want to use a different configuration directory for each system.

You can do this by specifying the `ALLINEA_CONFIG_DIR` environment variable before starting Arm Forge. If you use the `module` system, you can set `ALLINEA_CONFIG_DIR` according to the system on which the module was loaded.

For example, if you have two systems: **harvester** with login nodes `harvester-login1` and `harvester-login2`, and **sandworm** with login nodes `sandworm-login1` and `sandworm-login2`, you can add something like the following code to your module file:

```
case $(hostname) in
  harvester-login*)
    ALLINEA_CONFIG_DIR=$HOME/.allinea/harvester
    ;;
  sandworm-login*)
    ALLINEA_CONFIG_DIR=$HOME/.allinea/sandworm
    ;;
esac
```

Using a shared installation on multiple systems

If you have multiple systems sharing a common installation, you can have a different default configuration for each system. You can use the `ALLINEA_DEFAULT_SYSTEM_CONFIG` environment variable to specify a different file for each system. For example, you can add something like the following code to your module file:

```
case $(hostname) in
  harvester-login*)
    ALLINEA_DEFAULT_SYSTEM_CONFIG=/sw/*(\forgeInstallPathRaw)*/harvester.config
    ;;
  sandworm-login*)
    ALLINEA_DEFAULT_SYSTEM_CONFIG=/sw/*(\forgeInstallPathRaw)*/sandworm.config
    ;;
esac
```

A.3 Integration with queuing systems

You can configure Arm Forge to interact with most job submission systems. This is useful if you wish to debug interactively but need to submit a job to the queue in order to do so.




Figure A-1 Queuing systems

Arm MAP is usually run as a wrapper around `mpirun` or `mpiexec`, using the `map -profile` argument. We recommend using this to generate `.map` files instead of configuring Arm MAP to submit jobs to the queue, but both usage patterns are fully-supported.

In the **Options** window (**Preferences** on Mac OS X), you should choose **Submit job through queue**. This displays extra options and switches the GUI into queue submission mode.

The basic stages in configuring to work with a queue are:

1. Making a template script.
2. Setting the commands used to submit, cancel, and list queue jobs.

Your system administrator can provide a configuration file containing the correct settings, and remove the need for individual users to configure their own settings and scripts.

In this mode, Arm Forge can use a template script to interact with your queuing system. The `templates` subdirectory contains some example scripts that can be modified to meet your needs. `<installation-directory>/templates/sample.qtf`, demonstrates the process of creating a template file in some detail.

A.4 Template tutorial

Typically, your queue script probably ends in a line that starts `mpirun` with your target executable. In most cases, you can replace that line with `AUTO_LAUNCH_TAG`. For example, if your script currently has this line:

```
mpirun -np 16 program_name myarg1 myarg2
```

Create a copy of it and replace that line with:

```
AUTO_LAUNCH_TAG
```

Select this file as the **Submission template file** on the **Job Submission Settings** page in the **Options** window.

..note::

You no longer need to explicitly specify the number of processes, and so on. Instead, specify the number of processes, program name, and arguments in the **Run** window.

Enter the **Submit command** with the command you usually use to submit your job, such as `qsub` or `sbatch`. Use the **Cancel command** with the command you usually use to cancel a job, for example `qdel` or `scancel`. Use the **Display command** with the command you usually use to display the current queue status, such as `qstat` or `squeue`.

You can usually use `( d+)` as the **Regex for job id**. This just scans for a number in the output from your **Submit command**.

When you have a simple template working, you can go on to make more things configurable from the GUI. For example, to be able to specify the number of nodes from the GUI, you could replace an explicit number of nodes with the `NUM_NODES_TAG`. In this case, replace:

```
#SBATCH --nodes=100
```

With:

```
#SBATCH --nodes=NUM_NODE_TAG
```

See [Queue template tags on page Appx-I-520](#) for a full list of tags.

The template script

The template script is based on the file you would typically use to submit your job. This is usually a shell script that specifies the resources needed, such as number of processes, output files, and executes `mpirun`, `vmirun`, `poe` or similar, with your application.

The most important difference is that job-specific variables, such as number of processes, number of nodes, and program arguments, are replaced by capitalized keyword tags, such as `NUM_PROCS_TAG`.

When Arm Forge prepares your job, it replaces each of these keywords with its value and then submits the new file to your queue.

To refer to tags in comments without detecting them as a required field, the comment line must begin with `##`.

Configure queue commands

When you have selected a queue template file, enter submit, display, and cancel commands.

When you start a session, Arm Forge generates a submission file and appends its file name to the submit command you give.

For example, if you normally submit a job by typing `job_submit -u myusername -f myfile`, you must enter `job_submit -u myusername -f` as the submit command.

To cancel a job, Arm Forge uses a regular expression that you provide to get a value for `JOB_ID_TAG`. This tag is found by using regular expression matching on the output from your submit command. See [Job ID regular expression on page Appx-I-527](#) for details.

Configure how job size is chosen

Arm Forge offers a number of flexible ways to specify the size of a job. You can choose whether **Number of Processes** and **Number of Nodes** options appear in the **Run** window, or whether these should be implicitly calculated. Similarly, you can choose to display **Processes per node** in the **Run** window, or set it to a Fixed value.

Note

If you choose to display **Processes per node** in the **Run** window and `PROCS_PER_NODE_TAG` is specified in the queue template file, the tag is always replaced by the **Processes per node** value from the **Run** dialog, even if the option is unselected there.

Quick restart

Arm DDT allows you reuse an existing queued job to quickly restart a run without resubmitting it to the queue, with the requirement that your MPI implementation supports this. Select the **Quick Restart** checkbox on the **Job Submission Options** page of the **Options** window. See [Optional configuration on page Appx-A-449](#).

When using the quick restart, your queue template file must use `AUTO_LAUNCH_TAG` to execute your job.

For more information on `AUTO_LAUNCH_TAG`, see **Using `AUTO_LAUNCH_TAG`** in [Launching on page Appx-I-524](#).

A.5 Connecting to compute nodes and remote programs (remote-exec)

Arm Forge attempts to use the `ssh` secure shell by default when it needs to access another machine for remote launch, or as part of starting some MPIs.

However, this might not always be appropriate, because `ssh` can be disabled or run on a different port to the default port 22. In this case, you can create a file called `remote-exec` in your `~/.allinea` directory, which Arm DDT uses instead.

Arm Forge looks for the script at `~/.allinea/remote-exec`, and it is executed as follows:

```
remote-exec HOSTNAME APPNAME [ARG1] [ARG2] ...
```

The script must start `APPNAME` on without further input (no password prompts). Standard output from `APPNAME` appears on the standard output of `remote-exec`.

For example:

```
SSH based remote-exec
```

This shows a `remote-exec` script using `ssh` running on a non-standard port.

```
#!/bin/sh
ssh -P {port-number} $*
```

For this to work without prompting for a password, generate a public and private SSH key, and ensure that the public key is added to the `~/.ssh/authorized_keys` file on machines you wish to use. See the `ssh-keygen` manual page for more information.

Testing

When you have set up your `remote-exec` script, Arm recommends that you test it from the command line. For example:

```
~/.allinea/remote-exec TESTHOST uname -n
```

This returns the output of `uname -n` on , without prompting for a password.

If you are having trouble setting up `remote-exec`, contact [Arm support](#) for assistance.

Windows

The functionality is also provided by the Windows remote client. However, there are two differences:

- The script is named `remote-exec.cmd` rather than `remote-exec`.
- The default implementation uses the `plink.exe` executable supplied with Arm Forge.

A.6 Optional configuration

This section includes details of the Arm Forge **Options** window (**Preferences** on Mac OS X), which allows you to edit the system, job submission, code viewer settings, and appearance.

Access the Arm® Forge Options or Preferences window

- To open the **Options** window, select **File > Options**.
- To open the **Preferences** window (Mac OS X), select **File > Preferences**.

System

This section provides details of the optional system settings that you can apply to Arm Forge.

Name	Description
MPI Implementation	Allows you to identify which MPI implementation you are using.
Override default mpirun path	Allows you to override the path to the mpirun (or equivalent) command.
Select Debugger	Tells Arm DDT which underlying debugger to use. Unless a specific debugger is required, leave this as Automatic . On Linux systems, Arm Forge ships with the following versions of the GNU GDB debugger: GDB 7.6.2, GDB 7.12.1, GDB 8.2 and GDB 10.0. Arm recommends the GDB 8.2 debugger for Arm DDT. This recommended default is selected automatically when you select Automatic (recommended) in Options > System Settings . The GDB 10.0 debugger is optional for Arm DDT, but has limitations: GDB 10.0 does not support UPC, MPC and PGI compiler versions prior to 19.1.
Create Root and Workers groups automatically	If this option is selected, Arm DDT automatically creates a Root group for rank 0, and a Workers group for ranks 1-n when you start a new MPI session.
Heterogeneous system support	Arm DDT has support for running heterogeneous MPMD MPI applications where some nodes use one architecture and other nodes use another architecture. This requires a little preparation of your installation. • You must have a separate installation of Arm DDT for each architecture. The architecture of the machine running the Arm Forge GUI is called the host architecture. • You must create symbolic links from the host architecture installation of Arm to the other installations for the other architectures. For example, with a 64-bit x86_64 host architecture (running the GUI) and some compute nodes running the 32-bit i686 architecture: <pre>ln -s /{installation-directory(i686)}/libexec/forg-backend \ /{installation-directory(x86_64)}/bin/forg-backend.i686</pre>
Enable CUDA software pre-emption	Allows debugging of CUDA kernels on a workstation with a single GPU.
Default groups file	Entering a file here allows you to customize the groups displayed by Arm DDT when starting an MPI job. If you do not specify a file, Arm DDT creates the default Root and Workers groups if the previous option is selected. .. note:: You can create a groups file while your program is running by right-clicking the Process groups panel and selecting Save groups .
Attach hosts file	When attaching, Arm DDT fetches a list of processes for each of the hosts listed in this file. See Attach to running programs on page B1-73 for more details.

Job submission

This section of the **Options** window allows you to configure to use a custom command, or submit your jobs to a queuing system. For more information on this, see [Integration with queuing systems](#) on page Appx-A-445.

Code viewer settings

This section allows you to configure the appearance of the code viewer, which is used to display your source code while debugging.

Name	Description
Tab size	Sets the width of a tab character in the source code display. A width of 8 means that a tab character has the same width as 8 space characters.
Font name	The name of the font used to display your source code. Arm recommends that you use a fixed-width font.
Font size	The size of the font used to display your source code.
External Editor	This is the program Arm Forge executes if you right-click in the code viewer and choose Open file in external editor . This command launches a graphical editor. If no editor is specified, Arm Forge attempts to launch the default editor that is configured in your desktop environment.
Colour Scheme	Color palette to use for the code viewer background, text, and syntax highlighting. Defined in Kate syntax definition format in the <code>resource/styles</code> directory of the install.
Visualize Whitespace	Enables or disables this display of symbols to represent whitespace. Useful for distinguishing between space and tab characters.
Warn about potential programming errors	This setting enables or disables the use of static analysis tools that are included with the installation. These tools support F77, C, and C++, and analyze the source code of viewed source files to discover common errors, but they can cause heavy CPU usage on the system running the Arm Forge user interface. You can disable this by clearing this option.

Appearance

This section allows you to configure the graphical style of Arm Forge, as well as fonts and tab settings for the code viewer.

Name	Description
Look and Feel	This determines the general graphical style of Arm Forge. This includes the appearance of buttons, context menus.
Override System Font Settings	This setting can be used to change the font and size of all components in (except the code viewer).

Appendix B

Get support

This appendix describes how to get help if you need it.

It contains the following section:

- [*B.1 Supporting information on page Appx-B-452.*](#)

B.1 Supporting information

This user guide attempts to cover as many parts of the installation, features, and uses of Arm Forge as possible. However, there are scenarios or configurations that are not covered, or are only briefly mentioned, or you might on occasion experience a problem using the product. If the solution to your problem is not in this guide, contact [Arm support](#).

Provide as much detail as you can about the scenario, such as:

- Version number of Arm Forge. For example, `forge -version` and your operating system, and the distribution, such as Red Hat Enterprise Linux 6.4.

This information is all available by using the `-version` option on the command line of any Arm tool:

```
bash$ forge --version

Arm DDT
Part of Arm Forge.
Copyright (c) 2002-2021 Arm Limited (or its affiliates). All rights reserved.

Version: 18.0.2
Build ID: 556f23c4895e
Build Platform: Ubuntu 16.04 x86_64
Build Date: Jan 25 2018 21:15:53

Frontend OS: Ubuntu 16.04.2 LTS
Nodes' OS: unknown
Last connected forge-backend: unknown
```

- The compiler in use and its version number.
- The MPI library and CUDA toolkit version, if appropriate.
- A description of the issue: what you expected to happen and what actually happened.
- An exact copy of any warning or error messages that you have encountered.

Appendix C

Supported platforms

Lists the supported platforms for Arm Forge.

It contains the following section:

- [C.1 Reference table on page Appx-C-454.](#)

C.1 Reference table

This table describes the architectures, operating systems, MPI distributions, compilers, and accelerators that are supported by Arm Forge, including Arm DDT, Arm MAP, and Arm Performance Reports.

Architecture	Operating systems	MPI	Compilers	Accelerators
Armv8 (AArch64)	Red Hat Enterprise Linux/ CentOS 7 and 8 SuSE Linux Enterprise Server 12 and 15 Ubuntu 16.04 to 20.04	Cray MPT HPE MPI a MPICH MVAPICH2 a Open MPI 3 to 4	Arm Compiler for Linux Cray Compiling Environment GNU C/C++/Fortran Compiler NVIDIA HPC (PGI) Compiler	Nvidia CUDA Toolkit 11.0 to 11.1
Intel and AMD (x86_64)	Red Hat Enterprise Linux/ CentOS 7 and 8 SuSE Linux Enterprise Server 12 and 15 Ubuntu 16.04 to 20.04	Cray MPT HPE MPI a Intel MPI MPICH MVAPICH2 a Open MPI 3 to 4	Cray Compiling Environment GNU C/C++/Fortran Compiler Intel Parallel Studio NVIDIA HPC (PGI) Compiler	Nvidia CUDA Toolkit 9.0 to 11.1
IBM Power (ppc64le)	Red Hat Enterprise Linux/ CentOS 7 and 8	IBM Spectrum MPI Open MPI 3 to 4	GNU C/C++/Fortran Compiler IBM XL Compiler NVIDIA HPC (PGI) Compiler	Nvidia CUDA Toolkit 9.2 to 11.1

a These MPIs do not support Express Launch. See [Express Launch \(DDT\) on page B1-65](#) for more details.

Note

See [SLURM on page Appx-E-470](#) for more details about slurm support.

Forge

- DDT:
 - Pretty printing of C++ types is not supported for the NVIDIA HPC (PGI) or Cray compilers.
 - Message queue debugging is supported for Intel MPI, MPICH, MVAPICH2, and Open MPI.
- MAP and Performance Reports:
 - Arm profiling libraries are pre-compiled for these MPIs: Open MPI 3.x.x to 4.0.x, Intel MPI 5.x.x, 2017.x, 2018.x and 2019.x, Cray MPT, and MVAPICH 2.x.x. For other MPIs, these will be created dynamically at run time.
 - The Arm Forge profiler libraries must be explicitly linked with statically-linked programs. This mostly applies to the Cray X-Series.

Forge Remote Client

The Arm Forge Remote Client is available for the following x86_64 platforms:

- MacOS 10.13 (High Sierra), and above.
- Windows 7, and above.
- Any of the Linux platforms listed in **Supported platforms**.

Appendix D

Known issues

Summarizes the most significant known issues for the latest release.

It contains the following sections:

- *D.1 Arm® MAP* on page Appx-D-456.
- *D.2 XALT Wrapper* on page Appx-D-457.
- *D.3 SLURM support* on page Appx-D-458.
- *D.4 See also* on page Appx-D-459.

D.1 Arm® MAP

The following known issues affect Arm MAP and Arm Performance Reports.

- I/O metrics are not available on some systems, including Cray systems.
- CPU instruction metrics are only available on x86_64 systems.
- Thread activity is not sampled while a process is inside an MPI call with a duration spanning multiple samples. This can appear as "uncategorized" (white) time in the **Application activity** bar, when in the **Pthread view**. The uncategorized time coincides with long running MPI calls.
- Arm MAP and Arm Performance Reports do not support code that spawns new processes, such as `fork`, `exec` and `MPI_Comm_spawn`. In these cases, Arm MAP and Arm Performance Reports only profile the original process.
- Arm Performance Reports might fail to finalize a profiling session if the cores are oversubscribed on AArch64 architectures. For example, this occurs when attempting to profile a 64 process MPI program on a machine with only 8 cores. This appears as a hang after finishing a profile.
- Arm MAP displays a warning when running OpenMP programs compiled with PGI 20.3 or earlier. "Another OpenMP runtime library has been detected in this application..." This warning relates to earlier versions of PGI implementing a subset of OpenMP. Upgrade to PGI 20.4 or later. .. note:: PGI is renamed to NVIDIA HPC after version 20.7. or silence the warning by setting the `NV_OMP_DISABLE_WARNINGS=true` environment variable.

D.2 XALT Wrapper

XALT MPI job launcher wrappers are known to cause several issues when used in conjunction with Arm Forge, such as:

- MPI programs cannot be debugged due to a hang during start up.
- Error messages are reported relating to the permissions on `qstat`.

A workaround is to remove any XALT MPI launcher wrappers from your PATH. If these are provided by a module, this can be achieved by unloading the module.

D.3 SLURM support

On Cray X-series systems, only native SLURM is supported. Hybrid mode is not supported.

With newer versions of SLURM, you might receive this startup error:

```
srun: error: Unable to create step for job <N>: Requested nodes are busy
```

To work around this issue, set `ALLINEA_DEBUG_SRUN_ARGS="%default% --oversubscribe"`. If that does not solve the issue, then set `ALLINEA_USE_SSH_STARTUP=1`.

D.4 See also

See also these additional known issues:

Category	Known Issues
MPI Distribution	MPI distribution notes and known issues on page Appx-E-461
Compiler	Compiler notes and known issues on page Appx-F-473
Platform	Platform notes and known issues on page Appx-G-481
General	General troubleshooting and known issues on page Appx-H-489

Appendix E

MPI distribution notes and known issues

This appendix has brief notes on many of the MPI distributions supported by Arm Forge.

Advice on settings and problems particular to a distribution are given here.

Note that Arm MAP supports fewer MPI distributions than Arm DDT. See [Supported platforms on page Appx-C-453](#) for more details.

It contains the following sections:

- [E.1 Cray MPT](#) on page Appx-E-462.
- [E.2 HP MPI](#) on page Appx-E-464.
- [E.3 Intel MPI](#) on page Appx-E-465.
- [E.4 MPICH 3](#) on page Appx-E-467.
- [E.5 MVAPICH 2](#) on page Appx-E-468.
- [E.6 Open MPI](#) on page Appx-E-469.
- [E.7 SLURM](#) on page Appx-E-470.
- [E.8 IBM Spectrum MPI](#) on page Appx-E-471.

E.1 Cray MPT

This section only applies when using `aprun`.

For `srun` (Native SLURM mode), see [SLURM on page Appx-E-470](#).

Arm DDT and Arm MAP have been tested with Cray XT 5/6, XE6, XK6/7, and XC30 systems. Arm DDT can launch and support debugging jobs in more than 700,000 cores. Arm Performance Reports has been tested with Cray XK7 and XC30 systems.

Several template files are included in the distribution for launching applications from within the queue, using the Arm job submission interface. These might require some minor editing to cope with local differences on your batch system.

To attach to a running job on a Cray system, the MOM nodes where `aprun` is launched, must be accessible using `ssh` from the node where Arm DDT is running. You can either specify the `aprun` host manually in the **Attach** dialog when scanning for jobs, or configure a hosts list containing all nodes.

Preloading Arm DDT memory debugging libraries is not supported with `aprun`. If the program is dynamically linked, Arm MAP and Arm Performance Reports support preloading the profile libraries with `aprun` (with `aprun/ALPS 4.1` or later). Preloading is not supported in MPMD mode and requires that sampling libraries are linked with the application before running on this platform.

See the **Linking** section in [Prepare a program for profiling on page C1-244](#) for more information.

By default, scripts wrapping Cray MPT are not detected. However, you can force the detection before starting Arm DDT, Arm MAP, or Arm Performance Reports, by setting the environment variable to `yes`.

Using DDT with Cray ATP (the Abnormal Termination Process)

Arm DDT is compatible with the Cray ATP system, which is the default on some XE systems. This runtime addition to applications automatically gathers crashing process stacks, and can be used to let Arm DDT attach to a job before it is cleaned up during a crash.

To debug after a crash when an application is run with ATP but without a debugger, initialize the environment variable before launching the job. For a large Petascale system, a value of 5 is sufficient, giving 5 minutes for the attach to complete.

The following example shows the typical output of an ATP session:

```
n10888@kaibab:~> aprun -n 1200 ./atploop
Application 1110443 is crashing. ATP analysis proceeding...
Stack walkback for Rank 23 starting:
_start@start.S:113
__libc_start_main@libc-start.c:220
main@atploop.c:48
__kill@0x4b5be7
Stack walkback for Rank 23 done
Process died with signal 11: 'Segmentation fault'
View application merged backtrace tree file 'atpMergedBT.dot'
with 'statview'
You may need to 'module load stat'.

atpFrontend: Waiting 5 minutes for debugger to attach...
```

To debug the application at this point, launch Arm DDT.

Arm DDT can attach using the **Attaching** dialogs described in [Attach to running programs on page B1-73](#), or given the PID of the `aprun` process, the debugging set can be specified from the command line.

For example, to attach to the entire job:

```
ddt --attach-mpi=12772
```

If a particular subset of processes are required, then the subset notation could also be used to select particular ranks.

```
ddt --attach-mpi=12772 --subset=23,100-112,782,1199
```

E.2 HP MPI

Select **HP MPI** as the MPI implementation.

A number of HP MPI users have reported a preference to using `-f jobconfigfile` instead of `-np 10 a.out` for their particular system. It is possible to configure Arm DDT to support this configuration using the support for batch (queuing) systems.

The role of the queue template file is comparable to the `-f jobconfigfile`.

If your job config file normally contains:

```
-h node01 -np 2 a.out  
-h node02 -np 2 a.out
```

Then your template file must contain:

```
-h node01 -np PROCS_PER_NODE_TAG /usr/local/ddt/libexec/forge-backend  
-h node02 -np PROCS_PER_NODE_TAG /usr/local/ddt/libexec/forge-backend
```

You must also enter this in the **Submit Command** field:

```
mpirun -f
```

Select the **Template uses** NUM_NODES_TAG and PROCS_PER_NODE_TAG radio button. When you click **Ok** to configure this, and you can start jobs.

————— **Note** —————

The **Run** button is replaced with **Submit**, and the number of processes box is replaced by **Number of Nodes**.

E.3 Intel MPI

Select **Intel MPI** from the MPI implementation list.

Arm DDT, Arm MAP, and Arm Performance Reports have been tested with Intel MPI 4.1.x, 5.0.x, and later.

Known issue: Arm Forge is unable to interoperate with the Intel MPI 2018 and 2019 message queue debugging.

Make sure that you are familiar with the changes in the `mpivars.sh` script with Intel MPI 5.0.

You can pass it an argument to say whether you want to use the debug or release version of the MPI libraries. If you omit the argument, the default is the release version. However, message queue debugging does not work with this version. To use the debug version, it must be explicitly specified in the argument.

Arm DDT also supports the Intel Message Checker tool that is included in the Intel Trace Analyser and Collector software. A plugin for the Intel Trace Analyser and Collector version 7.1 is provided in the Arm DDT plugins directory. Once you have installed the Intel Trace Analyser and Collector, you must make sure that the following directories are in your `LD_LIBRARY_PATH`:

```
<path to intel install directory>/itac/7.1/lib  
<path to intel install directory>/itac/7.1/slib
```

Note

The Intel Message Checker only works if you are using the Intel MPI.

Make sure that the Intel `mpiexec` is in your path, and that your application is compiled against the Intel MPI. Launch Arm DDT, select the plugin checkbox, and debug your application.

If one of the above steps is missed out, Arm DDT might report an error and say that the plugin could not be loaded.

When you are debugging with the plugin loaded, Arm DDT automatically pauses the application whenever Intel Message Checker detects an error. The Intel Message Checker log can be seen in the standard error (`stderr`) window.

Note

The Intel Message Checker aborts the job after 1 error by default. You can modify this by adding `-genv VT_CHECK_MAX_ERRORS 0` to the **mpirun arguments** field in the **Run** window. See the Intel documentation for more details on this and other environment variable modifiers.

Attach dialog

Arm DDT cannot automatically discover existing running MPI jobs that use Intel MPI if the processes are started using the `mpiexec` command (which uses the MPD process starting daemon). To attach to an existing job, you must list all potential compute nodes individually in the dialog.

Note

The `mpiexec` method of starting MPI processes is deprecated by Intel. Arm recommends using `mpirun` or `mpiexec.hydra` (which use the newer scalable Hydra process starting daemon). All processes that are started by either `mpirun` and `mpiexec.hydra` are discovered automatically by Arm DDT.

If you use Spectrum LSF as workload manager in combination with Intel MPI, you might need to set/export `I_MPI_LSF_USE_COLLECTIVE_LAUNCH=1` before executing the job, under these two conditions:

- You get one of these errors:

```
<target program> exited before it finished starting up. One or more  
processes were killed or died without warning
```

```
<target program> encountered an error before it initialised the MPI  
environment. Thread 0 terminated with signal SIGKILL
```

- The job is killed during launching/attaching.

Related information

Using IntelMPI under LSF quick guide

Resolve the problem of the Intel MPI job ...hang in the cluster

E.4 MPICH 3

MPICH 3.0.3 and 3.0.4 do not work with Arm Forge. MPICH 3.1 is supported.

There are two MPICH 3 modes, Standard and Compatibility. If the standard mode does not work on your system, select **MPICH 3 (Compatibility)** as the **MPI Implementation** on the **System Settings** page of the **Options** window.

The message queue data provided by the MPICH debugging support library is limited and results in unreliable information in the **Message Queue** debugging feature.

E.5 MVAPICH 2

Memory debugging in Arm DDT interferes with the on-demand connection system used by MVAPICH2 above a threshold process count and applications fails to start.

The default threshold value is 64.

Workaround

Set the environment variable `MV2_ON_DEMAND_THRESHOLD` to the maximum job size you expect on your system and Arm DDT can work with memory debugging enabled for all jobs.

Do not set this as a system-wide setting because it might increase startup times for jobs and memory consumption.

MVAPICH 2 now offers `mpirun_rsh` instead of `mpirun` as a scalable launcher binary.

1. To use this with Arm DDT, go to the **System** page (:menuselection: *File > Options*, or :menuselection: *Arm Forge > Preferences* on Mac OS X).
2. Select **Override default mpirun path** and enter `mpirun_rsh`.
3. Add `-hostfile <hosts>` in the :menuselection: *mpirun_rsh arguments* field in the **Run** window. `<hosts>` is the name of your hosts file.

To enable message Queue Support, compile MVAPICH 2 with the flags `-enable-debug -enable-sharedlib`. These are not set by default.

E.6 Open MPI

There are three different Open MPI choices in the list of MPI implementations to choose from in Arm Forge when debugging or profiling for Open MPI.

- **Open MPI** - the job is launched with a custom launch agent that, in turn, launches the Arm daemons.
- **Open MPI (Compatibility)** - mpirun launches the Arm daemons directly. This startup method does not take advantage of the Arm scalable tree.
- **Open MPI for Cray XT/XE/XK/XC** - for Open MPI running on Cray XT/XE/XK/XC systems. This method is fully capable of using the Arm scalable tree infrastructure.

To launch with aprun (instead of mpirun), enter one of these commands:

```
ddt --mpi="OpenMPI (Cray XT/XE/XK)" --mpiexec aprun [arguments]
map --mpi="OpenMPI (Cray XT/XE/XK)" --mpiexec aprun [arguments]
```

Known issues

- Message queue debugging does not work with the UCX or Yalla PML, due to UCX and Yalla not storing the required information.
- The version of Open MPI packaged with Ubuntu has the Open MPI debug libraries stripped. This prevents the **Message Queues** feature of Arm DDT from working.
- On Infiniband systems, Open MPI and CUDA can conflict in a manner that results in failure to start processes, or a failure for processes to be debugged. To enable CUDA interoperability with Infiniband, set the CUDA environment variable to 1.

These versions of Open MPI do not work with Arm Forge because of bugs in the Open MPI debug interface:

- Open MPI 3.0.0 when compiled with the Arm Compiler for Linux on Armv8 (AArch64) systems.
- Open MPI 3.0.x when compiled with some versions of the GNU compiler on Armv8 (AArch64) systems.
- Open MPI <= 3.x.4 and <= 4.0.1 when compiled with some versions of IBM XLC/XLF, NVIDIA HPC, or PGI compilers.
- Open MPI 3.1.x or 4.0.x when compiled with -O2 or -O3 optimization flags, and PGI 19.x or 20.1 compilers.
- Open MPI 3.1.0 and 3.1.1.
- Open MPI 3.x with any version of PMIx earlier than 2.0.
- Open MPI 4.0.1 with PMIx 3.1.2.
- Open MPI up to version 3.1.4 and 3.0.4 on Ubuntu 20.04 x86_64 systems.

To resolve any of these issues, select **Open MPI (Compatibility)** for the MPI Implementation.

Open MPI 3.x on IBM Power with the GNU compiler

To use Open MPI versions 3.0.0 to 3.0.4 (inclusive) and Open MPI versions 3.1.0 to 3.1.3 (inclusive) with the GNU compiler on IBM Power systems, you must configure the Open MPI build with `CFLAGS=-fasynchronous-unwind-tables`. This fixes a startup bug where Arm Forge is unable to step out of `MPI_Init` into your main function.

The startup bug occurs because of missing debug information and optimization in the Open MPI library. If you already configure with `-g`, you do not need to add this extra flag.

An example configure command is:

```
./configure --prefix=/software/openmpi-3.1.2 CFLAGS=-fasynchronous-unwind-tables
```

An alternative workaround if you do not have the option to recompile your MPI, is to select **Open MPI (Compatibility)** for the MPI Implementation.

This issue is fixed in later versions.

E.7 SLURM

To start MPI programs, use the `srun` command instead of your typical MPI `mpirun` command (or equivalent).

Select **SLURM (MPMD)** as the **MPI Implementation** on the **System Settings** page of the **Options** window.

This option works with most, but not all, MPIs. On Cray, Hybrid SLURM mode (that is, SLURM + ALPS) is not supported. Instead, you must start your program with the Cray `aprun` command. See [Cray MPT on page Appx-E-462](#).

You can use SLURM as a job scheduler with Arm DDT and Arm MAP by using a queue template file. See an example of this in `templates/slurm.qtf` supplied with the installation. See [Integration with queuing systems on page Appx-A-445](#) for more information on how to customize the template.

E.8 IBM Spectrum MPI

IBM Spectrum MPI 10.2 is supported on IBM PowerPC little-endian (ppc64le).

Arm Forge launches its backend components using the file system by default. If Arm Forge fails when starting a job because the filesystem is not available on the compute nodes, you can try with the environment variable `ALLINEA_DISABLE_SPECTRUMMPI_SCALABLE_SHIPOUT` set to `no`.

This launches the backend components using a shipout mechanism provided by IBM Spectrum MPI. The shipout mechanism is only available when using the `jsrun` command.

Arm DDT supports debugging jobs launched using **Spindle** (`jsrun -use_spindle=1`), but you must ensure that Spindle does not strip debug information from shared object files. The option to configure debug information stripping is `JSMD_SPINDLE_OPT_STRIP` and must be set to `0`. This can be set either in the system configuration (usually `/opt/ibm/spectrum_mpi/jsm_pmix/etc/jsm.conf`), or in your user configuration (`~/jsm.conf`).

Note

Changing the configuration only takes effect in subsequent `bsub` jobs.

Arm MAP does not support Spindle.

Appendix F

Compiler notes and known issues

When compiling for a Arm DDT debugging session, always compile with a minimal amount of optimization, or no optimization. Some compilers reorder instruction execution and omit debug information when compiling with optimization enabled.

For a list of supported compiler versions, see [Supported platforms](#) on page Appx-C-453.

It contains the following sections:

- [F.1 AMD OpenCL compiler](#) on page Appx-F-474.
- [F.2 Arm Fortran compiler](#) on page Appx-F-475.
- [F.3 Cray compiler environment](#) on page Appx-F-476.
- [F.4 GNU](#) on page Appx-F-477.
- [F.5 IBM XLC/XLF](#) on page Appx-F-478.
- [F.6 Intel compilers](#) on page Appx-F-479.
- [F.7 Portland Group and NVIDIA HPC SDK compilers](#) on page Appx-F-480.

F.1 AMD OpenCL compiler

Not supported by Arm MAP and Arm Performance Reports.

The AMD OpenCL compiler can produce debuggable OpenCL binaries. However, the target must be the CPU rather than the GPU device. The build flags `-g -O0` must be used when building the OpenCL kernel, typically by setting the environment variable:

```
AMD_OCL_BUILD_OPTIONS_APPEND="-g -O0"
```

Run the example codes in the AMD OpenCL toolkit on the CPU by adding the parameter `-device cpu`. With the above environment variable set, this results in debuggable OpenCL.

F.2 Arm Fortran compiler

Debugging of Fortran code might be incomplete or inaccurate. For more information, check the known issues section in [*ARM Compiler for Linux release notes*](#).

F.3 Cray compiler environment

Arm DDT supports Cray fast-track debugging with GDB 8.2, 8.1 and 7.12.1.

To enable the supported versions of GDB, access the Systems Settings options. Select **File > Options > System**, or select **Options > System** on the **Welcome page**, then choose from the Debugger options.

To enable fast-track debugging, compile your program with `-Gfast` instead of `-g`.

See **Using Cray Fast-track Debugging** in the *Cray Programming Environment User's Guide* for more information.

Call-frame information can also be incorrectly recorded, which can sometimes lead to Arm DDT stepping into a function instead of stepping over it. This might also result in time being allocated to incorrect functions in MAP.

C++ pretty printing of the STL is not supported by Arm DDT for the Cray compiler.

Known Issue: If you are compiling static binaries, then linking in the Arm DDT memory debugging library is not straightforward for F90 applications. You must do the following:

1. Manually rerun the compiler command with the `-v` (verbose) option to get the linker command line. Ensure that the object files are already created.
2. Run `ld` manually to produce the final statically linked executable. For this, the following path modifications are required in the previous `ld` command: Add `-L{ddt-path}/lib/64 -ldmalloc` immediately prior to where `-lc` is located. For multi-threaded programs you have to add `-ldmallocth -lpthread` before the `-lc` option.

See CUDA/GPU debugging notes for details of Cray accelerator support.

Arm DDT supports UPC using the Cray UPC compiler, but Arm MAP and Arm Performance Reports do not.

Compile scalar programs on Cray

To launch scalar code with `aprun`, using Arm Forge on Cray, you must link your program with Cray PMI. With some configurations of the Cray compiler drivers, Cray PMI is discarded during linking. For static executables, consider using the `-wl, -u, PMI_Init` compilation flags to preserve Cray PMI.

- To use Arm MAP or Arm Performance Reports, see **Linking** in [Prepare a program for profiling on page C1-244](#).
- To use `aprun` to launch your program, see [Starting scalar programs with aprun on page Appx-H-500](#).
- To use SLURM, see [Starting scalar programs with srun on page Appx-H-501](#).

F.4 GNU

Do not use the compiler flag `-fomit-frame-pointer` because this can prevent Arm Forge identifying the lines of code that your program has stopped.

For GNU C++, large projects can result in large amounts of debug information, which can lead to high memory.

The `-foptimize-sibling-calls` optimization (used in `-O2`, `-O3` and `-Os`) interferes with the detection of some OpenMP regions. If your code is affected by this issue, add `-fno-optimize-sibling-calls` to disable it, and allow Arm MAP and Arm Performance Reports to detect all the OpenMP regions in your code.

Using the `-dwarf-2` flag together with the `-strict-dwarf` flag can cause problems with stack unwinding, and result in a `cannot find the frame base` error. DWARF 2 does not provide all the information necessary for unwinding the call stack, so many compilers add DWARF 3 extensions with the missing information. Using the `-strict-dwarf` flag prevents compilers from doing so, and the error message is reported. Removing `-strict-dwarf` fixes this problem.

F.5 IBM XLC/XLF

Arm recommends using the `-qfullpath` option to prevent issues with Arm Forge locating source files.

Using IBM XL compilers with optimization level `-O2` or higher can lead to partial traces in Arm MAP, because insufficient information is available to fully unwind the call stack.

For the best OpenMP debug experience, compile your code with `-qsmp=omp:noopt` instead of `-qsmp=omp`.

For more information about debugging OpenMP, see [Debug OpenMP programs on page B1-80](#).

To view Fortran assumed size arrays in DDT, right-click on the variable, select **Edit Type** and enter the type of the variable with its bounds, for example `integer arr(5)`.

For the best experience when debugging GPU code built with IBM XL, disable all GPU optimizations. For example:

```
xlc -g -O0 -qsmp=omp:noopt -qoffload -qfullpath -qnoinline -Xptxas -Xllvm2ptx -nvvm-compile-  
options=-opt=0 target_example1.c -o target_example1.exe
```

F.6 Intel compilers

Arm supports Intel compilers.

For details see [Reference table on page Appx-C-454](#).

- You might experience issues with missing or incomplete stack traces. For example, [partial trace] entries display in Arm MAP, or no stack traces for allocations display in the Arm DDT **View Pointer Details** window.

Try recompiling your program with the `-fno-omit-frame-pointer` argument.

- Some optimizations that are performed when you specify `-ax` options in IFC/ICC can result in programs which cannot be debugged or profiled due to lack of frame pointer information.
- Some optimizations that are performed using Interprocedural Optimization (IPO), (implicitly enabled by the `-O3` flag) can interfere with the ability of Arm MAP to display call stacks. This makes it more difficult to understand what the program is doing.

To prevent this disable IPO by adding `-no-ip` or `-no-ipo` to the compiler flags. The `-no-ip` flag disables IPO within files, and `-no-ipo` disables IPO between files.

- The Intel compiler does not always provide enough information to correctly determine the bounds of some Fortran arrays when they are passed as parameters, in particular the lower-bound of assumed-shape arrays.
- The Intel OpenMP compiler always optimizes parallel regions, regardless of `-O0` parameters. This can cause issues with reordered instructions or variables that can not be read.
- Files with a `.F` or `.F90` extension are automatically preprocessed by the Intel compiler. This can also be turned on with the `-fpp` command-line option. Unfortunately, the Intel compiler does not include the correct location of the source file in the executable produced when preprocessing is used.

If your Fortran file does not make use of macros and does not need preprocessing, you can rename its extension to `.f` or `.f90` and/or remove the `-fpp` flag from the compile line instead.

Alternatively, to help Arm DDT discover the source file, right-click on the **Project Files** window, select **Add/view source directory**, and add the correct directory.

- Some versions of the compiler emit incorrect debug information for OpenMP programs which might cause some OpenMP variables to show as `<not allocated>`.
- By default Fortran `PARAMETERS` are not included in the debug information output by the Intel compiler. You can force them to be included by passing the `-debug-parameters all` option to the compiler.
- If you are compiling static binaries, for example on a Cray XT/XE machine, then linking in the Arm DDT memory debugging library is not straightforward for F90 applications. You need to manually rerun the last `ld` command (as seen with `ifort -v`) to include `-L{ddt-path}/lib/64 -ldmalloc` in two locations:
 - Include immediately prior to where `-lc` is located.
 - Include the `-zmuldefs` option at the start of the `ld` line.
- STL sets, maps and multi-maps cannot be fully explored, because only the total number of items is displayed. Other data types are unaffected.
- To disable pretty printing, set the environment variable `ALLINEA_DISABLE_PRETTY_PRINTING` to 1 before starting Arm DDT. This enables you to manually inspect the variable in the case of, for example, the incomplete `std::set` implementations.

F.7 Portland Group and NVIDIA HPC SDK compilers

PGI 20.7 has been re-branded to the NVIDIA HPC SDK. A result of this change is that `pgcc`, `pgc++`, and `pgfortran` have been aliased to `nvc`, `nvc++` and `nvfortran` respectively. This has not affected the functionality of the compilers with Arm Forge. The old compiler directives can still be used directly without making use of the `nv` prefixed directives.

If you experience problems with missing or incomplete stack traces (that is [partial trace] entries in Arm MAP or no stack traces for allocations in the **View Pointer Details** window of Arm DDT), compile your program with the `-Mframe` and `-Meh` frame arguments.

Some known issues are listed here:

- Included files in Fortran 90 generate incorrect file and line information in debug information. The information provides refers to **including** file but displays the line numbers from the **included** file.
- The PGI compiler might emit incorrect line number information, particularly for optimized code.

This can cause Arm DDT and Arm MAP to show your program on a different line to the one expected.

- When using memory debugging with statically-linked PGI executables (-Bstatic) you must add a localrc file to your PGI installation because of the in-built ordering of library linkage for F77/F90. The localrc file defines the correct linkage when using Arm DDT and (static) memory debugging. Append the following to <{>pgi-path>/bin/localrc:

```
# Add the path to the localrc file to the PATH environment variable
export PATH=$PATH:$(pwd)

# Add the path to the localrc file to the LD_LIBRARY_PATH environment variable
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$(pwd)
```
- When you pass an array splice as an argument to a subroutine that has an assumed shape array argument, the offset of the array splice might be ignored by Arm DDT. If this affects you, contact [Arm support](#).
- Arm DDT might show extra symbols for pointers to arrays and some other types. For example, if your program uses the variable ia1loc2d, then the symbol ia1loc2d\$sd might also be displayed. The extra symbols are added by the compiler and can be ignored.
- The Portland compiler also wraps F90 allocations in a compiler-handled allocation area, rather than directly using the systems memory allocation libraries directly for each allocate statement. This means that bounds protection (Guard Pages) cannot function correctly with this compiler.

Appendix G

Platform notes and known issues

This appendix provides details about specific issues affecting platforms. If a supported machine is not listed in this section, then there are no known issues.

It contains the following sections:

- [G.1 Cray on page Appx-G-482.](#)
- [G.2 GNU/Linux systems on page Appx-G-483.](#)
- [G.3 Intel Xeon on page Appx-G-484.](#)
- [G.4 NVIDIA CUDA on page Appx-G-485.](#)
- [G.5 Arm on page Appx-G-486.](#)
- [G.6 POWER8 and POWER9 \(ppc64le\) on page Appx-G-487.](#)
- [G.7 Mac OS X on page Appx-G-488.](#)

G.1 Cray

This section describes several issues when using Arm Forge tools on Cray and offers workarounds steps you can take to avoid them.

- If you are using Arm MAP on Cray, Arm recommends you read the following topics in [Prepare a program for profiling on page C1-244](#).
 - **Debugging symbols**
 - **Static linking on Cray X-Series systems**
 - **Dynamic and static linking on Cray X-Series systems using the modules environment**

Arm supplies module files in `FORGE_INSTALLATION_PATH/share/modules/cray`.

- Compilers may link statically on this platform by default. See [Compiler notes and known issues on page Appx-F-473](#) for compiler-specific information on static linking.
- See [MPI distribution notes and known issues on page Appx-E-461](#) for information on Cray MPT.
- Cray GPU debugging requires a working `TMPDIR` to be available, if `/tmp` is not available. You must ensure that this directory is not a shared filesystem such as NFS or Lustre.

To set `TMPDIR` for the compute nodes only, use the `DDT_BACKEND_TMPDIR` environment variable instead. Arm DDT automatically propagates this environment variable to the compute nodes.

- An extra step is necessary if you run single-process scalar codes (that is, non-MPI/SHMEM/UPC applications) on the compute nodes. These are required to be executed by `aprun`. However, `aprun` does not execute these applications using the typical debug-supporting protocols.
- Running a dynamically-linked, single-process, non-MPI program that runs on a compute node (that is, non-MPI CUDA or OpenACC code) requires that you add the `-target=native` flag to the compiler. This flag prevents the compiler linking in the MPI job launch routines, which can interfere with debuggers on this platform. Alternatively, you can convert the program to an MPI program by adding `MPI_Init` and `MPI_Finalize` statements, and run it as a one-process MPI job.

G.2 GNU/Linux systems

General

These are several known issues to be aware of when using Arm Forge tools on GNU/Linux.

- If you see the following when launching Arm Forge:

```
Unable to load the Qt Plugins.
```

This can mean that the `libx11-xcb1` package required by Qt-5 is not installed on your system. The package is available for installation on all of the Arm Forge supported platforms. For more information, see [Supported platforms on page Appx-C-453](#).

Ubuntu

On Ubuntu 19.04 and later, you must install the `libncurses5` and `libtinfo5` packages. These packages can be installed using the command:

```
$ sudo apt install libncurses5 libtinfo5
```

G.3 Intel Xeon

Intel Xeon processors, starting with Sandy Bridge, include Running Average Power Limit (RAPL) counters. Arm MAP can use the RAPL counters to provide energy and power consumption information for your programs.

Enabling RAPL energy and power counters when profiling

To enable the RAPL counters to be read by Arm MAP, you must load the `intel_rapl` kernel module.

The `intel_rapl` module is included in Linux kernel releases 3.13 and later. For testing purposes, Arm has backported the `powercap` and `intel_rapl` modules for older kernel releases.

G.4 NVIDIA CUDA

There are a number of issues you should be aware of:

- Arm DDT memory leak reports do not track GPU memory leaks.
- Debugging paired CPU/GPU core files is possible but is not yet fully supported.
- CUDA metrics in Arm MAP and Arm Performance Reports are not available for statically-linked programs.
- CUDA metrics in Arm MAP are measured at the node level, not the card level.
- NVIDIA Linux driver 418.43 or later might restrict GPU profiling to users with administrative privileges (CAP_SYS_ADMIN capability set). See the following NVIDIA page for details and instructions for disabling this restriction: [NVIDIA Development Tools Solutions - ERR_NVGPUCTRPERM: Permission issue with Performance Counters](#).
- Cray CCE 8.1.2 OpenACC and previous releases fail to generate debug information for local variables in accelerated regions. Install CCE 8.1.3 to address this issue.
- When debugging a CUDA application, adding watchpoints on either host or kernel code is not supported.
- When debugging a CUDA application, using the **Step threads together** box and **Run to here** to step into OpenMP regions is not supported. Use breakpoints to stop at the required line.
- When CUDA is set to **Detect invalid accesses (memcheck)**, placing breakpoints in CUDA kernels is only supported in CUDA 10.1 or later.
- A driver issue in CUDA 9.1 prevents Arm DDT from debugging CUDA GPU applications on Cray machines using Cray MPT (aprun). As a workaround, launch the CUDA application outside of Arm DDT and attach to it.

G.5 Arm

Armv8 (AArch64) known issues

- For best operation, Arm DDT requires that debug symbols are installed for the runtime libraries, in addition to debug symbols for the program itself. Otherwise, Arm DDT might show the incorrect values for local variables in program code if the program is currently stopped inside a runtime library. Arm recommends, as a minimum requirement, that you install the debug symbols for `glibc` and `OpenMP` (if applicable).
- For best operation, Arm MAP and Arm Performance Reports require that debug symbols are installed for the runtime libraries, in addition to debug symbols for the program itself. Otherwise, Arm MAP might report time in partial traces or unknown locations without debug symbols. Arm recommends, as a minimum requirement, that you install the debug symbols for `glibc` and `OpenMP` (if applicable).
- Arm MAP and Arm Performance Reports might fail to finalize a profiling session if the cores are oversubscribed on AArch64 platforms. For example, this issue is likely to occur when attempting to profile a 64 process MPI program on a machine with only 8 cores. This issue will appear as a hang after finishing a profile, or after pressing the **Stop and analyze** button in Arm MAP.

G.6 POWER8 and POWER9 (ppc64le)

Supported features

Split DWARF (Fission) and compressed DWARF are supported by Arm DDT and Arm MAP. Benefits include smaller debug information size, and potentially less memory consumption in Arm DDT, due to the ability to load debug symbols on demand. For example, if you use these flags with GCC, (which requires using the Binutils Gold linker):

```
gcc -gdwarf-4 -gsplit-dwarf -fdebug-types-section -Wl,-fuse-ld=gold,--gdb-index,--compress-debug-sections=zlib myprogram.c
```

IBM XLC 13.1.7 requires these flags. Configure the compiler to use the gold linker:

```
qdebug=NDWFSTR -gsplit-dwarf -Wl,--gdb-index, --compress-debug-sections=zlib
```

Known issues

See these POWER platform known issues:

- For best performance, Arm DDT, Arm MAP, and Arm Performance Reports require that debug symbols for the runtime libraries are installed, in addition to debug symbols for the program itself.

Without debug symbols, Arm DDT might show the incorrect values for local variables in program code if the program is currently stopped inside a runtime library.

Similarly, without debug symbols, Arm MAP might report time in partial traces or unknown locations. Arm recommends as a minimum requirement, that the debug symbols are installed for `glibc` and `OpenMP` (if applicable).

See the documentation for your operating system for instructions on how to install debug symbols.
- If a program is paused or sampled inside a system call or library function without debug or unwind information, Arm Forge may not be able to determine the call stack. To work around this in Arm DDT, try selecting a known line of code and choose **Run to here**. If this issue affects you, please contact [Arm support](#).
- Due to issues with the Data Address Watchpoint Register (DAWR) on POWER9, hardware watchpoints are not available from Linux kernel version 4.17. In this case, Arm DDT falls back to (much slower) software watchpoints. For more information, see [this GitHub page](#).
- On very rare occasions, Arm Forge can trigger a kernel bug on the POWER8. The bug is a soft lockup and has been observed with kernel version 3.10.0-327.36.3.el7. If the soft lockup occurs, stale `forge-backend` processes are left running at high CPU usage and affected nodes might eventually lockup completely. It is not possible to interact with these stale processes. For example, this can happen when sending kill or terminate signals. The only solution to this issue is to reboot affected nodes.

G.7 Mac OS X

The following menu items are not supported:

- **Edit Special Characters**
- **Edit Start Dictation**
- **View Enter Full Screen**
- **View Show Tab Bar**

Appendix H

General troubleshooting and known issues

This appendix offers help with common issues you might encounter while using Arm Forge.

Also, check the support pages on the [Arm Developer website](#), and check that you have the latest version of the product on the [Arm Forge Downloads](#) page.

It contains the following sections:

- [H.1 GUI cannot connect to an X Server](#) on page Appx-H-490.
- [H.2 Licenses](#) on page Appx-H-491.
- [H.3 F1 cannot display this document](#) on page Appx-H-495.
- [H.4 MPI not detected](#) on page Appx-H-496.
- [H.5 Starting a program](#) on page Appx-H-500.
- [H.6 Attaching](#) on page Appx-H-503.
- [H.7 Source Viewer](#) on page Appx-H-505.
- [H.8 Input/Output](#) on page Appx-H-506.
- [H.9 Controlling a program](#) on page Appx-H-507.
- [H.10 Evaluating variables](#) on page Appx-H-508.
- [H.11 Memory debugging](#) on page Appx-H-509.
- [H.12 MAP specific issues](#) on page Appx-H-511.
- [H.13 Excessive memory use](#) on page Appx-H-517.
- [H.14 Obtaining support](#) on page Appx-H-518.

H.1 GUI cannot connect to an X Server

Arm Forge does not open and you cannot connect to it when running on a remote server. "Map cannot connect to X server" displays.

The DISPLAY variable is not set correctly

The DISPLAY variable ensures that Arm Forge can connect to the X server.

Solution

- Check that the DISPLAY variable is correctly set. If connecting via SSH, ensure X11 forwarding is enabled.
- Ensure that you can run simple X applications, such as `xterm` from the same command line.

If you continue to experience problems, contact [Arm support](#).

Related references

[Chapter A4 Connecting to a remote system on page A4-51](#)

H.2 Licenses

If you are using Arm Licence Server, but Arm Forge products cannot connect to it, see the [user guide](#) for more troubleshooting information.

This section contains the following subsections:

- [H.2.1 License error on page Appx-H-491.](#)
- [H.2.2 No licenses found on page Appx-H-492.](#)
- [H.2.3 Related information on page Appx-H-493.](#)

H.2.1 License error

The Arm Forge user interface opens, but shows the message "Licence error ?" at the bottom of the sidebar.




Figure H-1 License error notification

Invalid license file

Your license file has been edited, or you are not able to connect to the license server.

Solution

- Click ? to see more information about the error.




Figure H-2 License error message

- Verify that you have a license file for the correct product in the license directory.
- Check the expiry date inside the license to ensure that the license is still valid.

If you continue to experience problems, contact [Arm support](#).

H.2.2 No licenses found

The user interface opens, but it is unresponsive and shows the message "No licences found" at the bottom of the sidebar.




Figure H-3 No licences found error notification in |forge|

Invalid license file (No licenses)

Arm Forge requires a license file so that it can run, debug, and profile your programs.

Solution

- Buy a license, or get a free trial license from the [Arm website](#).
- If you continue to experience problems, contact [Arm support](#).

Related references

[Chapter A2 Installation](#) on page A2-35

Related information

[Using Arm Licence Server](#)

[Arm Allinea Studio Licensing](#)

H.2.3 Related information

Related references

Chapter A2 Installation on page A2-35

Related information

Using Arm Licence Server

Arm Allinea Studio Licensing

H.3 F1 cannot display this document

A blank screen displays instead of this document when you press F1.

Corrupt files prevent Qt Assistant starting

There might be corrupt files that are preventing the documentation system (Qt Assistant) from starting.

Solution

Remove the stale files in `$HOME/.local/share/data/Allinea`.

H.4 MPI not detected

When you run an Arm Forge product (Arm DDT, Arm MAP, or Arm Performance Reports), you are notified about a failure. The nature of the failure is dependent on which product you are running and whether you are running it offline using command-line instructions, or using the Arm Forge GUI.

This section contains the following subsection:

- [H.4.1 MPI settings not configured on page Appx-H-496.](#)

H.4.1 MPI settings not configured

When you first run an Arm Forge product, the `system.config` file is created in your home directory under the `.allinea` folder. This file contains settings for enabling the product to auto detect the correct MPI implementation, and for specifying the default implementation. The failure to run MPI can arise if either of these settings are not configured.

Solution

To permanently enable auto detect, and to specify a default MPI type, edit the `system.config` file MPI section in your home directory and ensure that your product can run applications using MPI.

```
[mpi]
auto_detect = yes
type = openmpi
```

Solution

To permanently set the default MPI type, launch the Arm Forge GUI from the command-line interface, and select an implementation from the list in **MPI implementation** in the **Run** dialog.

1. Get a list of supported MPIs using `-list-mpis`:

```
$ map --list-mpis
```

2. Specify an MPI type for running your program using `-mpi`.

```
$ map --mpi=openmpi-compatible -n 1 ./wave_c
```

The GUI launches to display the MPI type you specified on the command line.




Figure H-4 Setting the MPI implementation in the GUI

Solution

To change the MPI type, use the Arm Forge GUI and set the MPI implementation in the **Run** dialog:

1. Select an MPI implementation in the **Run** dialog.
2. Click **MPI > Details > Implementation: Change**.
3. Select an implementation from the menu in **Options > System Settings > MPI/UPC implementation**.




Figure H-5 Setting the MPI Implementation in System Settings

Note

This permanently modifies the MPI section in the `system.config` file with the MPI type you select, and persists for future sessions. However, if auto detect is not set to `yes` in `system.config`, you can still encounter a problem using MPI in subsequent sessions.

Solution

To set the MPI type for the current session only, use the command-line interface.

1. Get a list of supported MPIS using `-list-mpis`:

```
$ map --list-mpis
```

2. Specify an MPI type for running your program using `-mpi`.

```
$ map --profile --mpi=openmpi -n 8 ./hello_c
```

Note

This change persists only for the current session and does not modify the `system.config` file.

Related references

A4.6 Starting Arm® Forge on page A4-60

H.5 Starting a program

This section contains the following subsections:

- [H.5.1 Starting scalar programs](#) on page Appx-H-500.
- [H.5.2 Starting scalar programs with aprun](#) on page Appx-H-500.
- [H.5.3 Starting scalar programs with srun](#) on page Appx-H-501.
- [H.5.4 Problems when you start an MPI program](#) on page Appx-H-501.
- [H.5.5 Starting multi-process programs](#) on page Appx-H-501.
- [H.5.6 No shared home directory](#) on page Appx-H-501.
- [H.5.7 Arm® DDT or Arm® MAP cannot find your hosts or the executable](#) on page Appx-H-502.
- [H.5.8 The progress bar does not move and Arm® Forge times out](#) on page Appx-H-502.

H.5.1 Starting scalar programs

There are several potential sources for issues.

MPI problem

The most common issue arises when the software reports a problem with MPI and you know your program is not using MPI.

Solution

1. Before you attempt to start a program, check [Compiler notes and known issues](#) on page Appx-F-473 and ensure that it is compiled correctly.
2. Select the **Run Without MPI Support** checkbox.

If you have selected this option and the software still refers to MPI, contact [Arm support](#).

Other issues starting a program

Other potential problems are:

- A previous Arm session is still running, or has not released resources required for the new session. Usually this can be resolved by killing stale processes. The most obvious symptom of this is a delay of approximately 60 seconds and a message stating that not all processes are connected. You might also see a QServerSocket message in the terminal.
- The target program does not exist or is not executable.
- The backend daemon of the Arm Forge products, `forge-backend`, is missing from the `bin` directory. In this case, check your installation and contact [Arm support](#).

H.5.2 Starting scalar programs with aprun

For compilation, see **Compile scalar programs on Cray** in: `doc:../compiler_notes_and_known_issues/cray-compiler_environment`. Ensure that you export the following environment variables:

```
export ALLINEA_MPI_INIT=main
export ALLINEA_HOLD_MPI_INIT=1
```

Instead of setting a breakpoint in the default `MPI_Init` location, these environment variables set a breakpoint in `main`, and hold the program there.

If using compatibility launch with a scalar program, the **Run** dialog automatically detects Cray MPI even though it is a non-MPI program. Keep MPI selected, set one process, then click **Run**.

If the above environment variables do not work, try an alternative solution by exporting:

```
export ALLINEA_STOP_AT_MAIN=1
```

`ALLINEA_STOP_AT_MAIN` holds the program wherever it was when it attached. This can be before `main`. For Arm DDT, set a breakpoint in the `main` of your program, then select **Play/Continue** to run to this breakpoint.

H.5.3 Starting scalar programs with srun

Export the following environment variables:

```
export ALLINEA_MPI_INIT=main
export ALLINEA_HOLD_MPI_INIT=1
```

Instead of setting a breakpoint in the default MPI_Init location, these environment variables set a breakpoint in main, and hold the program there.

If you are using compatibility launch mode with a scalar program, the **Run** dialog automatically detects SLURM. Keep **MPI** selected, set one process, then click **Run**.

If the above environment variables do not work, try an alternative solution by exporting:

```
export ALLINEA_STOP_AT_MAIN=1
```

ALLINEA_STOP_AT_MAIN holds the program wherever it was when it attached. This can be before main. For Arm DDT, set a breakpoint in the main of your program, then select **Play/Continue** to run to this breakpoint.

H.5.4 Problems when you start an MPI program

You encounter problems when you start an MPI program.

Solution

- Check whether you can run a single-process (non-MPI) program such as a trivial Hello, World! program, resolve any issues that arise, and repeat the attempt to run a multi-process job. Use any issues that you encounter as the starting point for diagnosing the problem.
- Verify that MPI is working correctly by running a job without any of Arm Forge products applied, such as the example in the **examples** directory of the installation.

```
mpirun -np 8 ./a.out
```

- Verify that mpirun is in the PATH, or the environment variable

ALLINEA_MPIRUN is set to the full pathname of mpirun.

H.5.5 Starting multi-process programs

If the progress bar does not report that at least process 0 has connected, the remote forge-backend daemons cannot be started or cannot connect to the GUI.

Sometimes problems are caused by environment variables not propagating to the remote nodes while starting a job. To a large extent, the solution to these problems depends on the MPI implementation that is being used.

Solution

- If only one, or very few, processes connect, it might be because you have not chosen the correct MPI implementation. Examine the list and look carefully at the options. If you cannot find another suitable MPI, contact [Arm support](#).
- If a large number of processes are reported by the status bar to have connected, it is possible that some have failed to start because of resource exhaustion, timing out, or, unusually, an unexplained crash.

To check for time-out problems, set the ALLINEA_NO_TIMEOUT environment variable to 1 before launching the GUI and see if further progress is made. This is not a solution, but aids the diagnosis. If all processes can start, contact [Arm support](#).

H.5.6 No shared home directory

Your home directory is not accessible to all the nodes in your cluster, and your jobs might fail to start.

Solution

1. Open the `~/allinea/system.config` file in a text editor.
2. Change the shared directory option in the `[startup]` section so that it points to a directory which is available and shared by all the nodes. If no such directory exists, change the `use session cookies` to `no` instead.

H.5.7 Arm® DDT or Arm® MAP cannot find your hosts or the executable

This can happen when attempting to attach to a process running on other machines. Ensure that the host names that reports issues with can be reached, using `ping`.

Solution

If Arm DDT fails to find the executable, ensure that it is available in the same directory on every machine.

See [Connecting to compute nodes and remote programs \(remote-exec\)](#) on page Appx-A-448 for more information on configuring access to remote machines.

H.5.8 The progress bar does not move and Arm® Forge times out

This can occur because of one of these issues:

- The program `forge-backend` has not been started by `mpirun` or has aborted.

You can log onto your nodes and confirm this by looking at the process list before clicking **Ok** when Arm Forge times out.

Ensure that `forge-backend` has all the libraries it needs and that it can run successfully on the nodes using `mpirun`.

- One or more processes (`forge-backend`, `mprun`, `rsh`) could not be terminated.

This can happen if Arm Forge is killed during its startup or due to MPI implementation issues.

You must kill the processes manually, using `ps -x` to get the process ids, then `kill` or `kill -9` to terminate them.

This issue can also arise for `mpich-p4mpd`, and the solution is explained in [MPI distribution notes and known issues](#) on page Appx-E-461.

If your intended command is not in your `PATH`, you can either add it or set the environment variable `ALLINEA_MPIRUN` to contain the full pathname of the correct `mpirun`.

Your home directory must be accessible to all the nodes in your cluster. If not, jobs might fail to start by this method.

Related references

[H.5.6 No shared home directory](#) on page Appx-H-501

H.6 Attaching

This section provides troubleshooting information about issues when attaching Arm DDT to running processes.

This section contains the following subsections:

- [H.6.1 The system does not allow connecting debuggers to processes \(Fedora, Ubuntu\) on page Appx-H-503.](#)
- [H.6.2 The system does not allow connecting debuggers to processes \(Fedora, Red Hat\) on page Appx-H-503.](#)
- [H.6.3 Running processes not shown in the attach window on page Appx-H-503.](#)

H.6.1 The system does not allow connecting debuggers to processes (Fedora, Ubuntu)

The Ubuntu ptrace scope control feature only allows a process to attach to other processes that it has launched directly.

See <http://wiki.ubuntu.com/Security/Features#ptrace> for more details.

Solution

- To disable this feature until the next reboot, run the following command:

```
echo 0 | sudo tee /proc/sys/kernel/yama/ptrace_scope
```

- To disable the feature permanently, add this line to `/etc/sysctl.d/10-ptrace.conf` or `/etc/sysctl.conf`:

```
kernel.yama.ptrace_scope = 0
```

This takes effect after the next reboot.

Note

On Fedora, ptrace might be blocked by SELinux and to Yama. See section **The system does not allow connecting debuggers to processes (Fedora, Red Hat)**.

H.6.2 The system does not allow connecting debuggers to processes (Fedora, Red Hat)

The `deny_ptrace` boolean in SELinux, used by Fedora and Red Hat, only allows a process to attach to other processes that it has launched directly.

See <http://fedoraproject.org/wiki/Features/SELinuxDenyPtrace> for more details.

Solution

- To disable this feature until the next reboot, run the following command:

```
setsebool deny_ptrace 0
```

- To disable this feature permanently, run this command:

```
setsebool -P deny_ptrace 0
```

Note

From Fedora version 22 and later, ptrace might be blocked by Yama and also to the SELinux boolean. See [The system does not allow connecting debuggers to processes \(Fedora, Ubuntu\) on page Appx-H-503](#) for more information.

H.6.3 Running processes not shown in the attach window

Running processes that do not show up in the **Attach** window is usually a problem with either the `remote-exec` script or the node list file.

Workaround

Ensure that the entry in your node list file corresponds with either localhost, if you are running on your local machine, or with the output of `hostname` on the desired machine.

Try running `/path/to/arm/forge/20.2/libexec/remote-exec` manually and check the output.

For example,

```
/path/to/arm/forge/20.2/libexec/remote-exec <hostname> ls
```

If this manual run fails, there is a problem with your `remote-exec` script. If `rsh` is still used in your script, check that you can `rsh` to the desired machine. Otherwise, check that you can attach to your machine in the way specified in the `remote-exec` script.

For more information, see **Connecting to compute nodes and remote programs (remote-exec)** in [Configuration on page Appx-A-441](#).

If you still experience problems with your script, contact [Arm support](#) for assistance.

H.7 Source Viewer

No variables or line number information

You must compile your programs with debug information included. You can usually do this, depending on your compiler, by adding the `-g` option to your compile command.

Source code does not appear when you start Forge

If you cannot see any text, the default selected font might not be installed on your system. If not, you can resolve the issue, go to **File > Options** (**Arm Forge > Preferences** on Mac OS X) and select a fixed width font such as Courier.

If you see a screen of text telling you that Arm Forge could not find your source files, follow the instructions given. If you still cannot see your source code, check that the code is available on the machine you are running the software on, and that the correct file and directory permissions are set. If some files are missing and others found, try adding source directories and rescanning for further instruction.

If the problem persists, contact [Arm support](#).

Code folding does not work for OpenACC/OpenMP pragmas

This is a known issue. If an OpenACC or OpenMP pragma is associated with a multi-line loop, the loop block might be folded instead.

H.8 Input/Output

This section contains the following subsections:

- [H.8.1 Output to stderr does not display](#) on page Appx-H-506.
- [H.8.2 Unwind errors](#) on page Appx-H-506.

H.8.1 Output to stderr does not display

Arm Forge automatically captures anything written to `stdout/stderr` and displays it.

Some shells, such as `csh`, do not support this feature and you might see your `stderr` mixed with `stdout`, or it might not display at all.

Solution

Arm strongly recommends writing program output to files instead, because the MPI specification does not cover `stdout/stderr` behavior.

H.8.2 Unwind errors

When using Arm MAP, you might see these errors reported in the output in the form:

```
Arm Sampler: 3 libunwind: Unspecified (general) error (4/172 samples)
Arm Sampler: 3 Maximum backtrace size in sampler exceeded, stack too deep. (1/172 samples)
```

These indicate that Arm MAP was only able to obtain a partial stack trace for the sample.

Solution

If the proportion of samples that generate such errors is low, they can safely be ignored.

If a large proportion of samples exhibit these errors, consult the advice on partial traces in [Intel compilers](#) on page Appx-F-479 or [Portland Group and NVIDIA HPC SDK compilers](#) on page Appx-F-480 if you are using these compilers.

If this does not resolve your issue, contact [Arm support](#).

H.9 Controlling a program

This section provides troubleshooting information about issues controlling programs.

This section contains the following subsections:

- [H.9.1 Program jumps forwards and backwards when stepping through](#) on page Appx-H-507.
- [H.9.2 Arm® DDT might stop responding when using the Step Threads Together option](#) on page Appx-H-507.

H.9.1 Program jumps forwards and backwards when stepping through

The program behaves this way when a program is compiled with optimizations. The compiler shuffles your program instructions into a more efficient order and changes the running sequence in the code.

Workaround

Arm recommends compiling with `-O0` when debugging, which disables the reordering behavior, and other optimizations.

Note

If you use the Intel OpenMP compiler, the compiler generates code that appears to jump in and out of the parallel blocks regardless of your `-O0` setting. Therefore, Arm recommends that you do not step inside parallel blocks.

H.9.2 Arm® DDT might stop responding when using the Step Threads Together option

Arm DDT might stop responding if a thread exits when the **Step Threads Together** option is enabled. This is most likely to occur on Linux platforms using NPTL threads. This might happen if you try to **Play to here** to a line that is never reached. In this case, your program would run to the end and then exit.

Workaround

Set a breakpoint at the last statement executed by the thread and turn off **Step Threads Together** when the thread stops at the breakpoint.

If this problem affects you, contact [Arm support](#).

H.10 Evaluating variables

This section provides troubleshooting information about issues with variables.

This section contains the following subsections:

- [H.10.1 Some variables cannot be viewed when the program is at the start of a function](#) on page Appx-H-508.
- [H.10.2 Incorrect values printed for Fortran array](#) on page Appx-H-508.
- [H.10.3 Evaluating an array of derived types, containing multiple-dimension arrays](#) on page Appx-H-508.
- [H.10.4 C++ STL types are not pretty printed](#) on page Appx-H-508.

H.10.1 Some variables cannot be viewed when the program is at the start of a function

Some compilers produce faulty debug information, forcing Arm DDT to enter a function during the *prologue*, or the variable might not yet be in scope.

In this region, which appears to be the first line of the function, some variables have not been initialized yet.

Solution

To view all the variables with their correct values, it might be necessary to play or step to the next line of the function.

H.10.2 Incorrect values printed for Fortran array

Pointers to non-contiguous array blocks (allocatable arrays using strides) are not supported.

Workaround

If this issue affects you, contact [Arm support](#) for a workaround or fix.

There are also many compiler limitations that can cause this issue. See Appendix [Compiler notes and known issues](#) on page Appx-F-473 for details.

H.10.3 Evaluating an array of derived types, containing multiple-dimension arrays

The **Locals**, **Current Line** and **Evaluate** views might not show the contents of these multi-dimensional arrays inside an array of derived types.

Solution

You can view the contents of the array by clicking the array name and dragging it into the evaluate window as an item on its own, or by using the Multi-dimensional array viewer (MDA). For more information about the MDA viewer, see [Multi-dimensional array viewer \(MDA\)](#) on page B4-164.

H.10.4 C++ STL types are not pretty printed

The pretty printers provided with are compatible with GNU compilers version 4.7 and above, and Intel C++ version 12 and above.

H.11 Memory debugging

This section contains the following subsections:

- [H.11.1 The View Pointer Details window says a pointer is valid but does not show you which line of code it was allocated on](#) on page Appx-H-509.
- [H.11.2 mprotect fails error when using memory debugging with guard pages](#) on page Appx-H-509.
- [H.11.3 Allocations made before or during MPI_Init show up in Current Memory Usage but have no associated stack back trace](#) on page Appx-H-509.
- [H.11.4 Deadlock when calling printf or malloc from a signal handler](#) on page Appx-H-509.
- [H.11.5 Program runs more slowly with Memory Debugging enabled](#) on page Appx-H-510.

H.11.1 The View Pointer Details window says a pointer is valid but does not show you which line of code it was allocated on

The **View Pointer Details** window says a pointer is valid but does not show you which line of code it was allocated on. The Pathscale compilers have known issues that can cause this. See the [Compiler notes and known issues](#) for more details.

The Intel compiler might need the `-fp` argument to allow you to see stack traces on some machines. If this happens with another compiler, contact [Arm support](#) with the vendor and version number of your compiler.

H.11.2 mprotect fails error when using memory debugging with guard pages

This can happen if your program makes more than 32768 allocations; a limit in the kernel prevents Arm DDT from allocating more protected regions than this.

Solution

Try these options to resolve the issue:

- Run `echo 123456 >/proc/sys/vm/max_map_count` (requires root) which increases the limit to 61728 (123456/2, because some allocations use multiple maps).
- Disable guard pages completely. This hinders the ability of Arm DDT to detect heap over/underflows.
- Disable guard pages temporarily. You can disable them at program start, add a breakpoint before the allocations you wish to add guard pages for, and then reenable the feature.

See **Configuration** in [Memory debugging](#) on page B8-193 for information on how to disable guard pages.

H.11.3 Allocations made before or during MPI_Init show up in Current Memory Usage but have no associated stack back trace

Memory allocations that are made before or during `MPI_Init` appear in the **Current Memory Usage** window along with any allocations made subsequently.

However, the call stack at the time of the allocation is not recorded for these allocations and does not show up in the **Current Memory Usage** window.

H.11.4 Deadlock when calling printf or malloc from a signal handler

The memory allocation library calls (such as, `malloc`) that are provided by the memory debugging library are not `async-signal-safe`, unlike the implementations in recent versions of the GNU C library.

POSIX does not require `malloc` to be `async-signal-safe` but some programs might expect this behavior.

For example, a program that calls `printf` from a signal handler can deadlock when memory debugging is enabled in Arm DDT, because the C library implementation of `printf` might call `malloc`.

Solution

See a table of the functions that can be safely called from an asynchronous signal handler, in the OpenGroup reference to [Signal Concepts](#)

H.11.5 Program runs more slowly with Memory Debugging enabled

The Memory Debugging library performs more checks than the memory allocation routines of normal runtime. However, these checks make the library slower.

Solution

If your program runs too slowly when Memory Debugging is enabled, there are several options you can change to speed it up.

- Try reducing the **Heap Debugging** option to a lower setting. For example, if it is currently on **High**, try changing it to **Medium** or **Low**.
- Increase the heap check interval from the default of 100 to a higher value. The heap check interval controls how many allocations might occur between full checks of the heap, which can take some time.
- A higher setting (1000 or above) is recommended if your program allocates and deallocates memory very frequently, for example from inside a computation loop.
- You can disable the **Store backtraces for memory allocations** option, at the expense of losing backtraces in the **View Pointer Details** and **Current Memory Usage** windows.

H.12 MAP specific issues

This section includes details about known issues with Arm MAP.

This section contains the following subsections:

- [H.12.1 MPI wrapper libraries](#) on page Appx-H-511.
- [H.12.2 Thread support limitations](#) on page Appx-H-511.
- [H.12.3 No thread activity while blocking on an MPI call](#) on page Appx-H-512.
- [H.12.4 I am not getting enough samples](#) on page Appx-H-512.
- [H.12.5 I just see main \(external code\) and nothing else](#) on page Appx-H-512.
- [H.12.6 Arm® MAP reports time spent in a function definition](#) on page Appx-H-513.
- [H.12.7 Arm® MAP does not correctly identify vectorized instructions](#) on page Appx-H-513.
- [H.12.8 Linking with the static Arm® MAP sampler library fails with FDE overlap errors](#) on page Appx-H-513.
- [H.12.9 Linking with the static Arm® MAP sampler library fails with an undefined references to `\_\_real\_dlopen`](#) on page Appx-H-514.
- [H.12.10 Arm® MAP adds unexpected overhead to my program](#) on page Appx-H-514.
- [H.12.11 Arm® MAP takes an extremely long time to gather and analyze my OpenBLAS-linked application](#) on page Appx-H-514.
- [H.12.12 Arm® MAP over-reports MPI, Input/Output, accelerator or synchronization time](#) on page Appx-H-515.
- [H.12.13 Arm® MAP collects very deep stack traces with `boost::coroutine`](#) on page Appx-H-515.

H.12.1 MPI wrapper libraries

Unlike Arm DDT, Arm MAP wraps MPI calls in a custom shared library. A precompiled wrapper is copied that is compatible with your system, or one is built for your system each time you run Arm MAP.

See [Supported platforms](#) on page Appx-C-453 for the list of supported MPIs.

You can also try setting `ALLINEA_WRAPPER_COMPILE=1` and `MPICC` directly:

```
$ MPICC=my-mpicc-command bin/map -n 16 ./wave_c
```

If you have problems, contact [Arm support](#) for more information.

H.12.2 Thread support limitations

Arm MAP and Arm Performance Reports provide limited support for programs when threading support is set to `MPI_THREAD_SERIALIZED` or `MPI_THREAD_MULTIPLE` in the call to `MPI_Init_thread`.

MPI activity on non-main threads contributes towards the MPI-time of the program, but not the MPI metric graphs.

Additionally, MPI activity on a non-main thread can result in additional profiling overhead due to the mechanism employed by Arm MAP and Arm Performance Reports for detecting MPI activity.

Solution

Arm recommends using the pthread view mode for interpreting MPI activity instead of the OpenMP view mode, because OpenMP view mode scales MPI-time depending on the resources requested. As a result, non-main thread MPI activity might provide non-intuitive results when detected outside of OpenMP regions.

Warnings display when the user initiates and completes profiling a program that sets `MPI_THREAD_SERIALIZED` or `MPI_THREAD_MULTIPLE` as the required thread support.

Arm MAP and Arm Performance Reports both fully support making calls to `MPI_Init_thread` using either `MPI_THREAD_SINGLE` or `MPI_THREAD_FUNNELED` to specify the required thread support.

Note

The MPI specification requirements for programs using `MPI_THREAD_FUNNELED` are the same as the Arm MAP and Arm Performance Reports requirements: all MPI calls must be made on the thread that is called `MPI_Init_thread`.

In many cases, multi-threaded MPI programs can be refactored so that they comply with this restriction.

H.12.3 No thread activity while blocking on an MPI call

Arm MAP and Arm Performance Reports are currently unable to record thread activity on a process where a long-duration MPI call is in progress.

If you have an MPI call that takes a significant amount of time to complete, as indicated by a sawtooth on the **MPI call duration** metric graph, Arm MAP displays no thread activity for the process executing that call for most of the duration of that MPI call.

Solution

See **MPI call duration** in [MPI calls](#) on page C9-308.

H.12.4 I am not getting enough samples

By default, the starting sample interval is every 20ms. You can change the sampling rate if you get warnings about too few samples on a fast run, or want more detail in the results.

Solution

To increase the interval to every 10ms, set environment variable `ALLINEA_SAMPLER_INTERVAL= 10`.

Note

Sampling frequency automatically decreases over time to ensure a manageable amount of data is collected, and does not depend on the length of the run.

Arm recommends that you avoid increasing the sampling frequency if there are lots of threads or very deep stacks in the target program. This might not leave sufficient time to complete one sample before the next sample is started.

Note

Whether OpenMP is enabled or disabled in Arm MAP or Arm Performance Reports, the final script or scheduler values set for `OMP_NUM_THREADS` is used to calculate the sampling interval per thread (`ALLINEA_SAMPLER_INTERVAL_PER_THREAD`). When configuring your job for submission, check whether your final submission script, scheduler or the Arm MAP GUI has a default value for `OMP_NUM_THREADS`.

Note

Custom values for `ALLINEA_SAMPLER_INTERVAL` are overwritten by values set from the combination of `ALLINEA_SAMPLER_INTERVAL_PER_THREAD` and the expected number of threads (from `OMP_NUM_THREADS`).

H.12.5 I just see main (external code) and nothing else

This can happen if you compile without `-g`. It can also happen if you move the executable out of the directory it was compiled in.

Solution

To check that your compile line includes `-g`, right-click the **Project Files** panel in Arm MAP, and select **Add Source Directory**.

If you have any further issues, contact [Arm support](#) for more information.

H.12.6 Arm® MAP reports time spent in a function definition

Any overheads involved in setting up a function call (such as pushing arguments to the stack) are usually assigned to the function definition. Some compilers might assign them to the opening brace (`{`) and closing brace (`}`) instead.

If this function is inlined, the situation becomes more complicated and any setup time, such as time for allocating space for arrays, is often assigned to the definition line of the enclosing function.

H.12.7 Arm® MAP does not correctly identify vectorized instructions

The instructions identified as vectorized (packed) are listed here:

- Packed floating-point instructions:

```
addpd addps addsubpd addsubps andnpd andnps andpd andps divpd divps dppd dpps haddpd
haddps hsubpd hsubps maxpd maxps minpd minps mulpd mulps rcpps rsqrtps sqrtpd sqrtps
subpd subps
```

- Packed integer instructions:

```
mpsadbw pabsb pabsd pabsw paddb paddd paddq paddsb paddsw paddusb paddusw paddw paligr
pavgb pavgw phaddb phaddsw phaddw phminposw phsubd phsubsw phsubw pmaddusb pmaddsw
pmaddw pmaxsb pmaxsd pmaxsw pmxub pmxud pmxuw pminsb pminsd pminsw pminub pminud pminuw
pmuldb pmulhrsw pmulhw pmulld pmullw pmuludq pshufb pshufw psignb psignd psignw
pslld psllq psllw psrad psraw psrld psrlq psrlw psubb psubd psubq psubsb psubsw psubusb
psubusw psubw
```

Arm also identifies the AVX-2 variants of these instructions, with a `V` prefix.

If you think that your code contains vectorized instructions that have not been listed and are not being identified in the CPU floating-point/integer vector metrics, contact [Arm support](#).

H.12.8 Linking with the static Arm® MAP sampler library fails with FDE overlap errors

When linking with the static sampler library, you might get FDE overlap errors similar to:

```
ld: .eh_frame_hdr table[791] FDE at 0000000000822830 overlaps table[792] FDE at
0000000000825788
```

This can occur when the version of `binutils` on a system has been upgraded to 2.25 or later and is most commonly seen on Cray machines using CCE 8.5.0 or later.

Solution

To fix this issue, rerun `make-profiler-libraries -lib-type=static` and use the freshly generated static libraries and `allinea-profiler.ld` to link these with your program.

See **Static Linking** in [Prepare a program for profiling on page C1-244](#) for more details.

If you do not use a Cray or SUSE build of Arm Forge and you require a `binutils` 2.25 compatible static library, contact [Arm support](#).

The error message occurs because the version of `libmap-sampler.a` that you attempted to link was not compatible with the version of `ld` in `binutils` versions 2.25 or greater.

For Cray machines, there is a separate library called `libmap-sampler-binutils-2.25.a` which is provided for use with this updated linker.

The `make-profiler-libraries` script automatically selects the appropriate library to use based on the version of `ld` found in your `PATH`.

If you erroneously attempt to link `libmap-sampler-binutils-2.25.a` with your program using a version of `ld` prior to 2.25, the following errors can occur:

```
/usr/bin/ld.x: libmap-sampler.a(dl.o): invalid relocation type 42
```

If this happens, check that the correct version of `ld` is in your `PATH` and rerun `make-profiler-libraries --lib-type=static`.

H.12.9 Linking with the static Arm® MAP sampler library fails with an undefined references to `__real_dlopen`

When linking with the static sampler library, you might get undefined reference errors similar to the following:

```
../lib/64/libmap-sampler.a(dl.o): In function '__wrap_dlopen':  
/build/overnight/ddt-2015-01-28-12322/code/ddt/map/sampler/build64-static/./src/dl.c:21:  
undefined reference to '__real_dlopen'  
../lib/64/libmap-sampler.a(dl.o): In function '__wrap_dlclose':  
/build/overnight/ddt-2015-01-28-12322/code/ddt/map/sampler/build64-static/./src/dl.c:28:  
undefined reference to '__real_dlclose'  
collect2: ld returned 1 exit status
```

Solution

To avoid these errors, follow the instructions in [Prepare a program for profiling on page C1-244](#).

Note

Look at the use of the `-Wl,@/home/user/myprogram/allinea-profiler.ld` syntax.

H.12.10 Arm® MAP adds unexpected overhead to my program

The Arm MAP sampler library adds a small overhead to the execution of your program. Usually, this is less than 5% of the wall clock execution time.

However, under some circumstances the overhead can exceed this, especially for short runs. This is particularly likely if your program has high OpenMP overhead, for example, if it is greater than 40%.

In this case, the measurements reported by Arm MAP are affected by this overhead and therefore less reliable. Increasing the run time of your program, for example, by changing the size of the input, decreases the overall overhead, although the initial few minutes still incur the higher overhead.

At high per-process thread counts, the Arm MAP sampler library can incur more significant overhead.

By default, when Arm MAP detects a large number of threads, it automatically reduces the sampling interval to limit the performance impact.

Sampling behavior can be modified by setting the `ALLINEA_SAMPLER_INTERVAL` and `ALLINEA_SAMPLER_INTERVAL_PER_THREAD` environment variables.

For more information on the use of these environment variables, see [MAP environment variables on page C1-264](#).

H.12.11 Arm® MAP takes an extremely long time to gather and analyze my OpenBLAS-linked application

OpenBLAS versions 0.2.8 and earlier incorrectly stripped symbols from the `.symtab` section of the library, causing binary analysis tools such as Arm MAP and `objdump` to see invalid function lengths and addresses.

This causes Arm MAP to take an extremely long time disassembling and analyzing apparently overlapping functions containing millions of instructions.

Solution

A fix for this was accepted into the OpenBLAS codebase on October 8th 2013. Version 0.2.9 and later are not affected.

To work around this problem without updating OpenBLAS, run `strip libopenblas*.so` to remove the incomplete `.symtab` section without affecting the operation or linkage of the library.

H.12.12 Arm® MAP over-reports MPI, Input/Output, accelerator or synchronization time

Arm MAP employs a heuristic to determine which function calls to consider as MPI operations.

Any function defined in your code defines that starts with `MPI_` (case insensitive), is treated as part of the MPI library. This causes the time spent in MPI calls to be over-reported by the activity graphs, and the internals of those functions are omitted from the **Parallel Stack View**.

Solution

Do not append the prefix `MPI_` to your function names. This is explicitly forbidden by the MPI specification. This is described on page 19, sections 2.6.2 and 2.6.3 of the [MPI 3 specification document](#):

"All MPI names have an `MPI_` prefix, and all characters are capitals. Programs must not declare names, for example, for variables, subroutines, functions, parameters, derived types, abstract interfaces, or modules, beginning with the prefix `MPI_`."

Similarly categorizes I/O functions and accelerator functions by name.

Other prefixes to avoid starting your function names with include `PMPI_`, `_PMI_`, `OMPI_`, `omp_`, `GOMP_`, `shmem_`, `cuda_`, `__cuda_`, `cu[A-Z][a-z]` and `allinea_`.

All of these prefixes are case-insensitive.

Do not name a function `start_pes` or any name also used by a standard I/O or synchronization function, such as `write`, `open`, `pthread_join`, and `sem_wait`.

H.12.13 Arm® MAP collects very deep stack traces with `boost::coroutine`

A known bug in Boost (<https://svn.boost.org/trac/boost/ticket/12400>) prevents from unwinding the call stack correctly.

This can be worked around by applying the patch attached to the bug report to your boost installation, or by specifying a manual stack allocator that correctly initializes the stack frame.

Add the following custom stack allocator:

```
#include <boost/coroutine/coroutine.hpp>
#include <boost/coroutine/stack_context.hpp>

struct custom_stack_allocator {
    void allocate(
        boost::coroutines::stack_context & ctx,
        std::size_t size ) {

        void * limit = std::malloc( size);
        if ( ! limit)
            throw std::bad_alloc();

        //Fix. RBP in the 1st frame of the stack will contain 0
        const int fill=0;

        std::size_t stack_hdr_size=0x100;
        if (size<stack_hdr_size)
            stack_hdr_size=size;
        memset(static_cast< char * >(limit)+size-stack_hdr_size,
            fill,
            stack_hdr_size);

        ctx.size = size;
        ctx.sp = static_cast< char * >( limit) + ctx.size;
    }

    void deallocate( boost::coroutines::stack_context & ctx) {
        void * limit = static_cast< char * >( ctx.sp) - ctx.size;
        std::free( limit);
    }
};
```

Modify your program to use the custom allocator whenever a coroutine is created:

```
boost::coroutines::coroutine<int()> my_coroutine(<func>,  
boost::coroutines::attributes(), custom_stack_allocator());
```

For more information, see the `boost::coroutine` documentation on stack allocators for your version of Boost.

H.13 Excessive memory use

If you are running out of memory when you are using an Arm Forge tool, consider the suggestions in this section.

This section contains the following subsections:

- [H.13.1 Reduce processes per node on page Appx-H-517.](#)
- [H.13.2 Reduce debug information on page Appx-H-517.](#)
- [H.13.3 Profiling settings on page Appx-H-517.](#)

H.13.1 Reduce processes per node

If your code allows it, run with fewer MPI processes per node. You can reduce the number of MPI processes in total or spread the processes out over more nodes.

Most tool-based memory usage is incurred per-process not per-thread, so you might want to increase the number of threads to compensate.

This impacts where your application spends its time so this might not be applicable when using Arm MAP or Arm Performance Reports.

H.13.2 Reduce debug information

Reduce the memory load by compiling some or all of your application with reduced debug information.

Depending on your use case, it might be appropriate to use minimal (file and line information only), or no debug information for some or all of your code.

Solution

To enable minimal debug, use the compiler option `-g1` (GCC, and Intel). To disable debug information, do not specify any `-g` option to the compiler or by use `-g0`.

- Arm DDT - Full debug information is required for code that you are debugging. Consider using split DWARF (`-gsplit-dwarf`) if your compiler supports it. Split DWARF allows full debug information to be lazily loaded with finer granularity, potentially saving memory.

Alternatively, parts of the code that you are sure do not need to be debugged can be compiled with minimal debug information.

- Arm MAP - Only minimal debug info (file and line numbers only) is required for most functionality. Stack frames that have been inlined can only be displayed if full debug information for that area of code is available).
- Arm Performance Reports - It is not required to compile with debug information enabled, but avoid stripping binaries of debug information after they have been compiled because this might remove required information.

H.13.3 Profiling settings

If your program appears to run correctly during profiling, but runs out of memory while Arm MAP or Arm Performance Reports collects the results, you can reduce memory usage by changing the processing order.

Solution

Define `ALLINEA_REDUCE_MEMORY_USAGE=1` in your environment and then rerun Arm MAP.

This environment variable causes the results for each process on a node to be processed sequentially instead of processed in parallel.

This reduces the amount of free memory needed on each node, but takes longer to complete.

H.14 Obtaining support

If you are unable to find the information that you need in this guide, contact [Arm support](#) with a detailed report.

If possible, obtain a log file for the problem. To generate a log file, either select the **Help > Logging > Automatic** menu option or start Arm Forge with the `-debug` and `-log` arguments:

```
$ ddt --debug --log=<logfilename>  
$ map --debug --log=<logfilename>
```

<logfilename> is the name of the log file to generate.

Reproduce the problem using as few processors and commands as possible, and when complete, close the program as usual.

On some systems, this file might be quite large. If so, compress it using a program such as `gzip` or `bzip2` before sending it to support.

If your problem can only be replicated on large process counts, do not use the **Help > Logging > Debug** menu item or the `-debug` argument because these generate very large log files. Instead, use the **Help > Logging > Standard** menu option or the `-log` argument.

If you are connecting to a remote system, the log file is generated on the remote host and copied back to the client when the connection is closed. The copy does not happen if the target application crashes or the network connection is lost.

In these cases, the remote copy of the log file is in the `tmp` subdirectory of the Arm configuration directory for the remote user account. The directory is `~/.allinea`, unless overridden by the `ALLINEA_CONFIG_DIR` environment variable.

Sometimes it might be helpful to illustrate your problem with a screenshot of the Arm Forge main window. To take a screenshot, choose the **Take Screenshot** option under the **Window** menu. You are prompted for a file name when you save the screenshot.

Appendix I

Queue template script syntax

This appendix contains information to help you with your script syntax.

It contains the following sections:

- [*I.1 Queue template tags*](#) on page Appx-I-520.
- [*I.2 Defining new tags*](#) on page Appx-I-521.
- [*I.3 Specifying default options*](#) on page Appx-I-523.
- [*I.4 Launching*](#) on page Appx-I-524.
- [*I.5 Using PROCESSES_PER_NODE_TAG*](#) on page Appx-I-526.
- [*I.6 Job ID regular expression*](#) on page Appx-I-527.
- [*I.7 Arm IPMI Energy Agent*](#) on page Appx-I-528.

I.1 Queue template tags

Each of the tags that are replaced is listed in the following table, and an example of the text that will be generated when submits your job is given for each.

————— **Note** —————

It is often sufficient to use AUTO_LAUNCH_TAG. See an example in **The template script** in [Template tutorial on page Appx-A-446](#).

Table I-1 Queue template tags

Tag	Description	After Submission Example
AUTO_LAUNCH_TAG	This tag expands to the entire replacement for your command line.	forge-mpirun -np 4 myexample.bin
DDTPATH_TAG	The path to the installation	/opt/
WORKING_DIRECTORY_TAG	The working directory was launched in	/users/ned
NUM_PROCS_TAG	Total number of processes	16
NUM_PROCS_PLUS_ONE_TAG	Total number of processes + 1	17
NUM_NODES_TAG	Number of compute nodes	\S
NUM_NODES_PLUS_ONE_TAG	Number of compute nodes + 1	\S
PROCS_PER_NODE_TAG	Processes per node	\S
PROCS_PER_NODE_PLUS_ONE_TAG	Processes per node + 1	\S
NUM_THREADS_TAG	Number of OpenMP threads per node (empty if OpenMP is off)	\S
OMP_NUM_THREADS_TAG	Number of OpenMP threads per node (empty if OpenMP is off)	\S
MPIRUN_TAG	mpirun binary (can vary with MPI implementation)	/usr/bin/mpirun
AUTO_MPI_ARGUMENTS_TAG	Required command-line flags for mpirun (can vary with MPI implementation)	-np 4
EXTRA_MPI_ARGUMENTS_TAG	Additional mpirun arguments specified in the Run window	-partition DEBUG
PROGRAM_TAG	Target path and filename	/users/ned/a.out
PROGRAM_ARGUMENTS_TAG	Arguments to target program	-myarg myval
INPUT_FILE_TAG	The stdin file specified in the Run window	/users/ned/input.dat

Additionally, any environment variables in the GUI environment ending in _TAG are replaced throughout the script by the value of those variables.

1.2 Defining new tags

In addition to the pre-defined tags listed in **Queue template tags**, you can also define new tags in your template script, and you can specify their values in the GUI.

Tag definitions have the following format:

```
EXAMPLE_TAG: { key1=value1, key2=value2, ... }
```

key1 and key2 are tag attribute names. value1 and value2 are the corresponding values.

The tag is replaced with the value specified in the GUI, as shown in this example.

```
#PBS -option EXAMPLE_TAG
```

These are the supported attributes:

Table I-2 Supported attributes

Attribute	Purpose	Examples
type	text - General text input. select - Select from two or more options. check - Boolean. file - File name. number - Real number. integer - Integer number.	type=text
label	The label for the user interface widget.	label="Account"
default	Default value for this tag	default="interactive"
text-type	-	-
mask	Input masks 0 ASCII digit permitted but not required. 9 ASCII digit required. 0-9. N ASCII alphanumeric character required. A-Z, a-z, 0-9. n ASCII alphanumeric character permitted but not required.	mask="09:09:09"
options type	-	-
Options	Options to use, separated by the pipe () character.	options="not_shared shared"
checked	Value of a check tag if checked.	checked="enabled"
unchecked	Value of a check tag if unchecked.	unchecked="enabled"
min	Minimum value.	min="0"
max	Maximum value.	max="100"
step	Amount to step by when the up or down arrows are clicked.	step="1"
decimals	Number of decimal places.	decimals="2"
suffix	Display only suffix (not included in tag value).	suffix="s"
prefix	Display only prefix (not included in tag value).	prefix="\$"
file type	-	-
mode	open-file an existing file. save-file a new or existing file. existing-directory an existing directory. open-files one or more existing files, separated by spaces.	mode="open-file"
caption	Window caption for file chooser.	caption="Select File"
dir	Initial directory for file chooser.	dir="/work/output"
filter	Restrict the files displayed in the file selector to a certain file pattern.	filter="Text files (*.txt)"

Examples

```
# JOB_TYPE_TAG: {type=select,options=parallel| \
serial,label="Job Type",default=parallel}

# WALL_CLOCK_LIMIT_TAG: {type=text,label="Wall Clock Limit", \
default="00:30:00",mask="09:09:09"}

# NODE_USAGE_TAG: {type=select,options=not_shared| \
shared,label="Node Usage",default=not_shared}

# ACCOUNT_TAG: {type=text,label="Account",global}
```

See the template files in <installation-directory>/templates for more examples.

Specifying tag values

To specify values for these tags, click the **Edit Template Variables** button on the **Job Submission Options** page (see [Integration with queuing systems](#) on page Appx-A-445) or the **Run** window.

A window displays which is similar to this:




Figure I-1 Queue parameters window

The values you specify are substituted for the corresponding tags in the template file when you run a job.

I.3 Specifying default options

A queue template file can specify defaults for the options on the **Job Submission** page so that when a user selects the template file, these options are automatically completed.

Table I-3 Default options

Name	Job Submission Setting	Example
submit	Submit command The command might include tags.	qsub -n NUM_NODES_TAG -t WALL_CLOCK_LIMIT_TAG -mode script -A PROJECT_TAG
display	Display command The output from this command is shown while waiting for a job to start.	qstat
job regexp	Job regexp	(\d+)
cancel	Cancel command	qdel JOB_ID_TAG
submit scalar	Also submit scalar jobs through the queue	yes
show num_procs	Number of processes: Specify in Run window	yes
show num_nodes	Number of nodes: Specify in Run Window	yes
show procs_per_node	Processes per node: Specify in Run window	yes
procs_per_node	Processes per node: Fixed	16

Example

```
# submit: qsub -n NUM_NODES_TAG -t WALL_CLOCK_LIMIT_TAG \
--mode script -A PROJECT_TAG
# display: qstat
# job regexp: (\d+)
# cancel: qdel JOB_ID_TAG
```

I.4 Launching

Usually, your queue script ends in a line that starts with your target executable.

In a template file, this must be modified to run a command that also launches the Arm Forge backend agents.

This section refers to some of the methods for doing this.

Using AUTO_LAUNCH_TAG

This is the easiest method to launch Arm Forge backend agents, and caters for the majority of cases. Replace your command line with AUTO_LAUNCH_TAG. Arm Forge replaces this with a command appropriate for your configuration (one command on a single line).

For example an `mpirun` line that looks like this:

```
mpirun -np 16 program_name myarg1 myarg2
```

Becomes:

```
AUTO_LAUNCH_TAG
```

AUTO_LAUNCH_TAG is roughly equivalent to:

```
DDT_MPIRUN_TAG DDT_DEBUGGER_ARGUMENTS_TAG \\  
MPI_ARGUMENTS_TAG PROGRAM_TAG ARGS_TAG
```

A typical expansion is:

```
/opt/*(\forgeInstallPathRaw)*/bin/forgempirun --ddthost login1,192.168.0.191 \  
--ddtport 4242 --ddtsession 1 \  
--ddtsessionfile /home/user/.allinea/session/login1-1 \  
--ddtshreddirectory /home/user --np 64 \  
--npernode 4 myprogram arg1 arg2 arg3
```

Using forge-mpirun

If you need more control than is available using AUTO_LAUNCH_TAG, Arm Forge also provides a drop-in `mpirun` replacement that can be used to launch your job.

Note

This is only suitable for use in a queue template file when Arm Forge is submitting to the queue itself.

Replace `mpirun` with `DDTPATH_TAG/bin/forgempirun`.

For example, if your script currently has the line:

```
mpirun -np 16 program_name myarg1 myarg2
```

Then (for illustration only) the equivalent that Arm Forge must use would be:

```
DDTPATH_TAG/bin/forgempirun -np 16 program_name myarg1 myarg2
```

For a template script, you use tags in place of the program name, arguments and so on, so that they can be specified in the user interface rather than editing the queue script each time:

```
DDTPATH_TAG/bin/forgempirun -np NUM_PROCS_TAG \  
EXTRA_MPI_ARGUMENTS_TAG DDTPATH_TAG/libexec/forgempirun \  
DDT_DEBUGGER_ARGUMENTS_TAG PROGRAM_TAG PROGRAM_ARGUMENTS_TAG
```

See [Queue template tags on page Appx-I-520](#) for more information on template tags.

Scalar programs

If AUTO_LAUNCH_TAG is not suitable, you can also use the following method to launch scalar jobs with your template script:

```
DDTPATH_TAG/bin/forge-client DDT_DEBUGGER_ARGUMENTS_TAG \\  
PROGRAM_TAG PROGRAM_ARGUMENTS_TAG
```

I.5 Using PROCS_PER_NODE_TAG

Some queue systems allow you to specify the number of processes, others require you to select the number of nodes and the number of processes per node.

The software caters for both of these but it is important to know whether your template file and queue system expect to be told the number of processes (NUM_PROCS_TAG) or the number of nodes and processes per node (NUM_NODES_TAG and PROCS_PER_NODE_TAG).

See `sample.qtf` for an explanation of the queue template system in the Arm Forge installation directory, at `<installationdirectory>/templates`.

I.6 Job ID regular expression

The **Regexp for job id** regular expression is matched on the output from your submit command. The first bracketed expression in the regular expression is used as the job ID. The elements listed in the table are available in addition to the conventional quantifiers, range and exclusion operators.

Table I-4 Regular expression characters

Element	Matches
C	A character represents itself
\t	A tab
.	Any character
\d	Any digit
\D	Any non-digit
\s	White space
\S	Non-white space
\w	Letters or numbers (a word character)
\W	Non-word character

For example, your submit program might return the output `job id j1128 has been submitted`. One possible regular expression for retrieving the job ID is `id\s(.+)\shas`.

If you would normally remove the job from the queue by typing `job_remove j1128`, you must enter `job_remove JOB_ID_TAG` as the cancel command.

I.7 Arm IPMI Energy Agent

The Arm IPMI Energy Agent allows Arm MAP and Arm Performance Reports to measure the total energy consumed by the compute nodes in a job.

Note

Measuring energy with IPMI Energy agent requires Arm Forge Professional.

The IPMI Energy Agent is available to download from the Arm website: [IPMI Energy Agent](#).

Requirements

- The compute nodes must support IPMI.
- The compute nodes must have an IPMI exposed power sensor.
- **The compute nodes must have an OpenIPMI compatible kernel module**
installed, such as `ipmi_devintf`.
- **The compute nodes must have the corresponding device node in /dev,**
for example `/dev/ipmi0`.
- The compute nodes must run a supported operating system.
- The IPMI Energy Agent must be run as root.

To list the names of possible IPMI power sensors on a compute node use the following command:

```
ipmitool sdr | grep 'Watts'
```