



arm

# Regional Application Profiling with Caliper

[Florent.Lebeau@arm.com](mailto:Florent.Lebeau@arm.com)

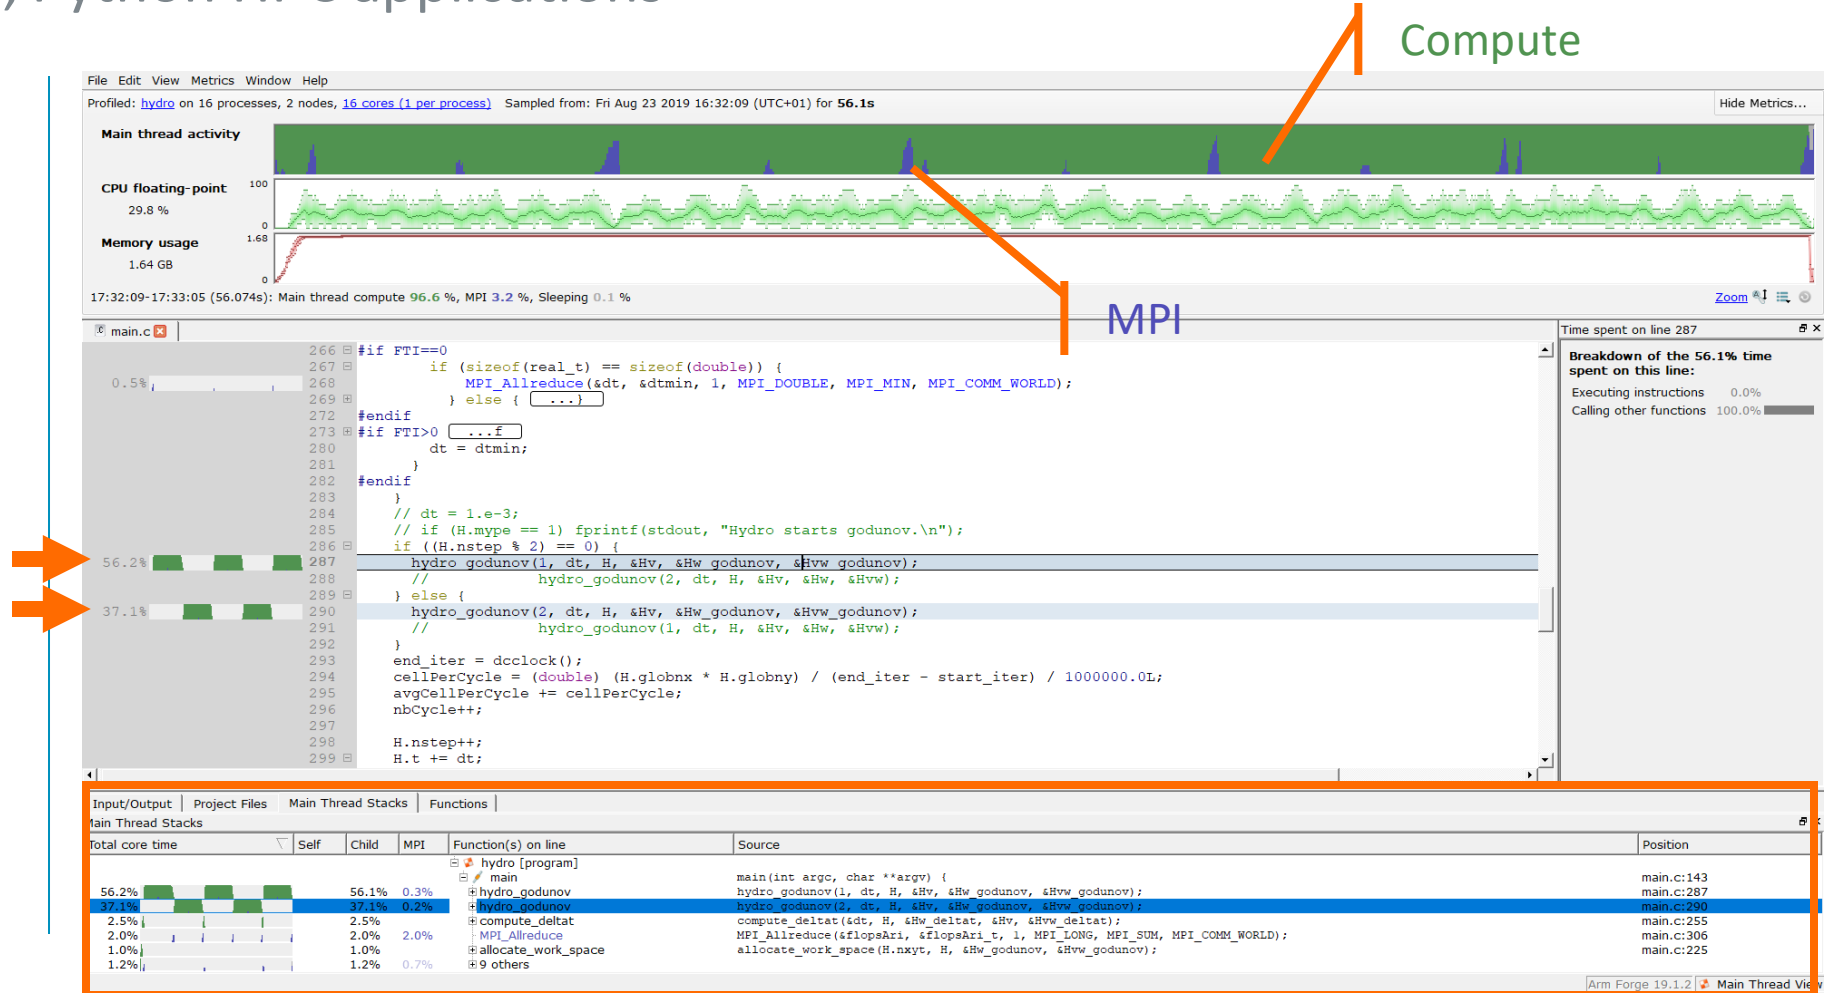
# Outline

- Introduction to Arm MAP
- Regional application profiling with Caliper
- Profiling Caliper-instrumented application with Arm MAP
- Analyzing profiling results
- Extracting regional profiling data for code optimization

# Performance Analysis with MAP, part of Arm Forge

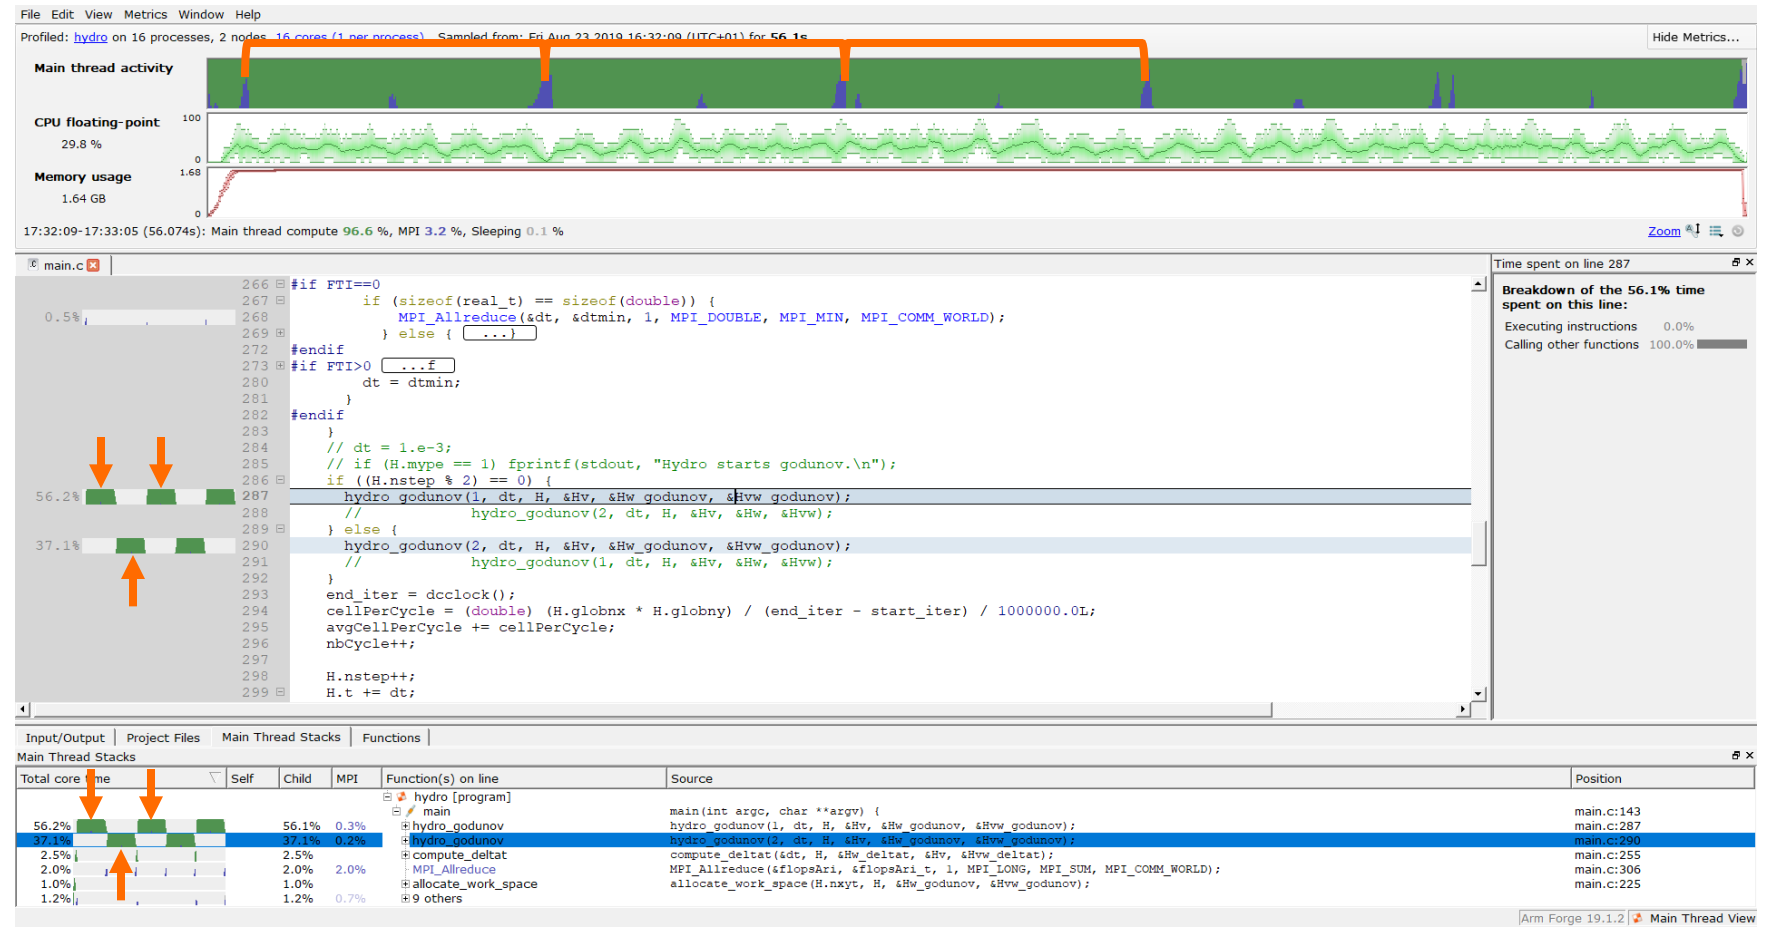
For C/C++, FORTRAN, Python HPC applications

- Easy to use
  - `$ map mpirun ./myapp`
- Lightweight
  - Results can be viewed offline
- Highly scalable
  - For GPGPU, OpenMP, MPI and hybrid applications
- Cross-platform
  - Arm v8-A, x86\_64 and POWER



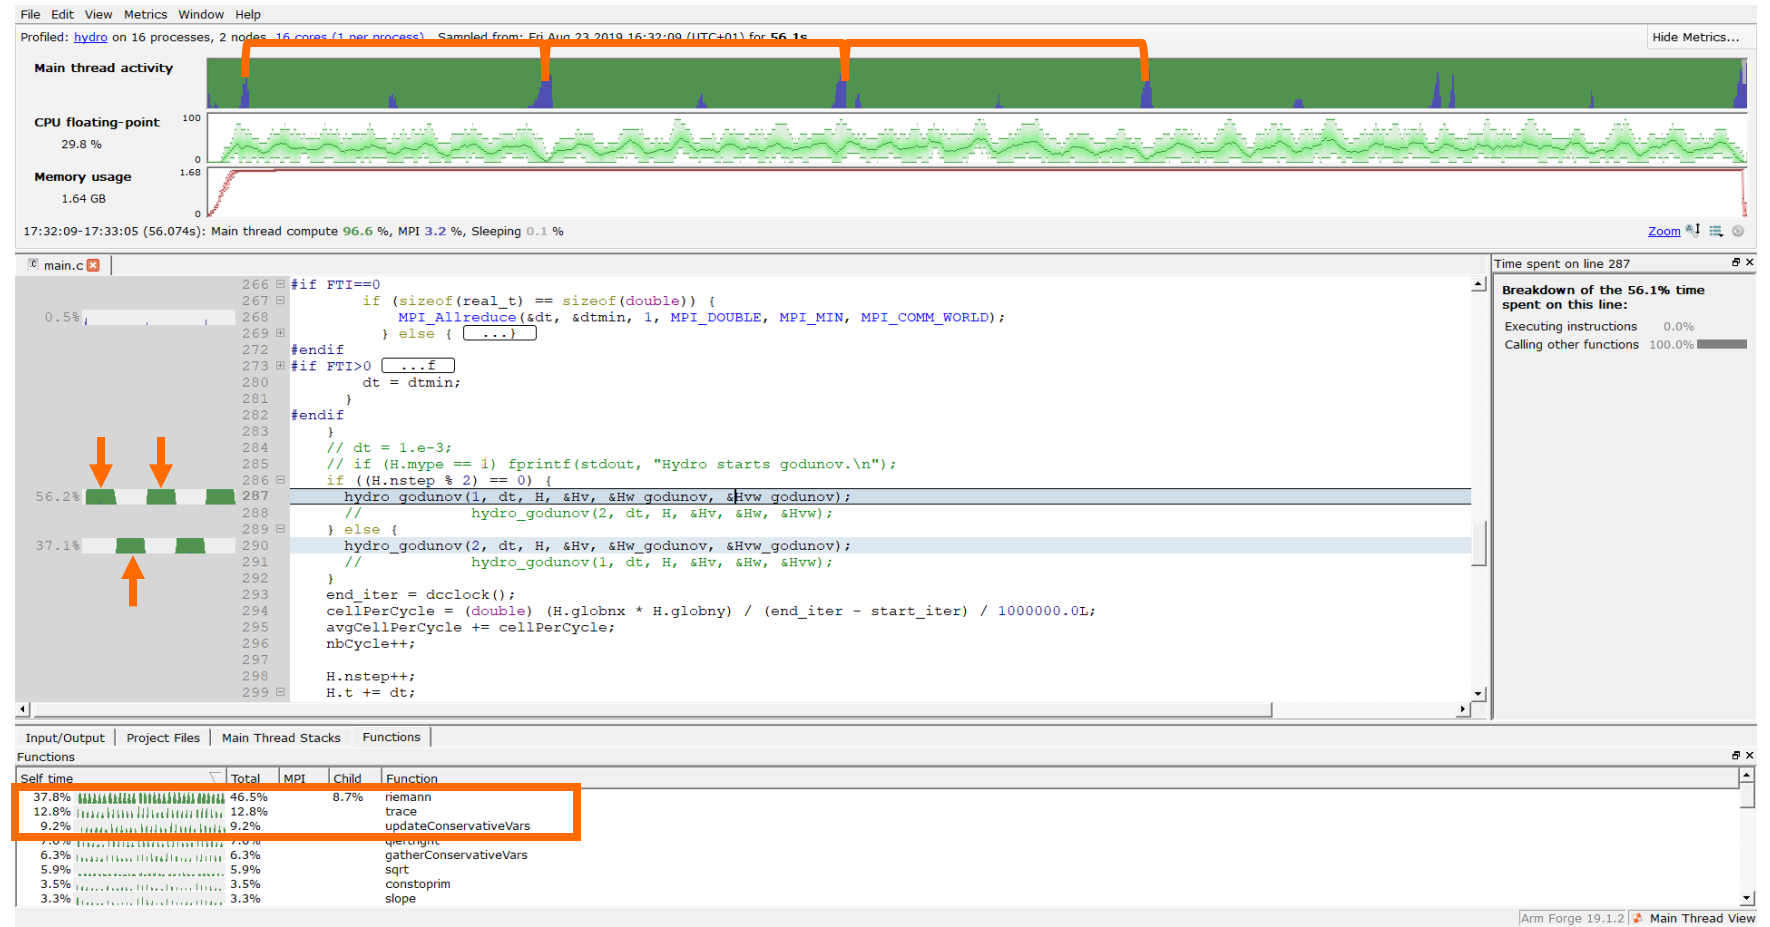
# Analyzing profiling results

- Example based on HydroC
  - <https://github.com/HydroBench/Hydro>
- 2 nodes, 16 processes
- Runtime: 56s
- Iterative pattern



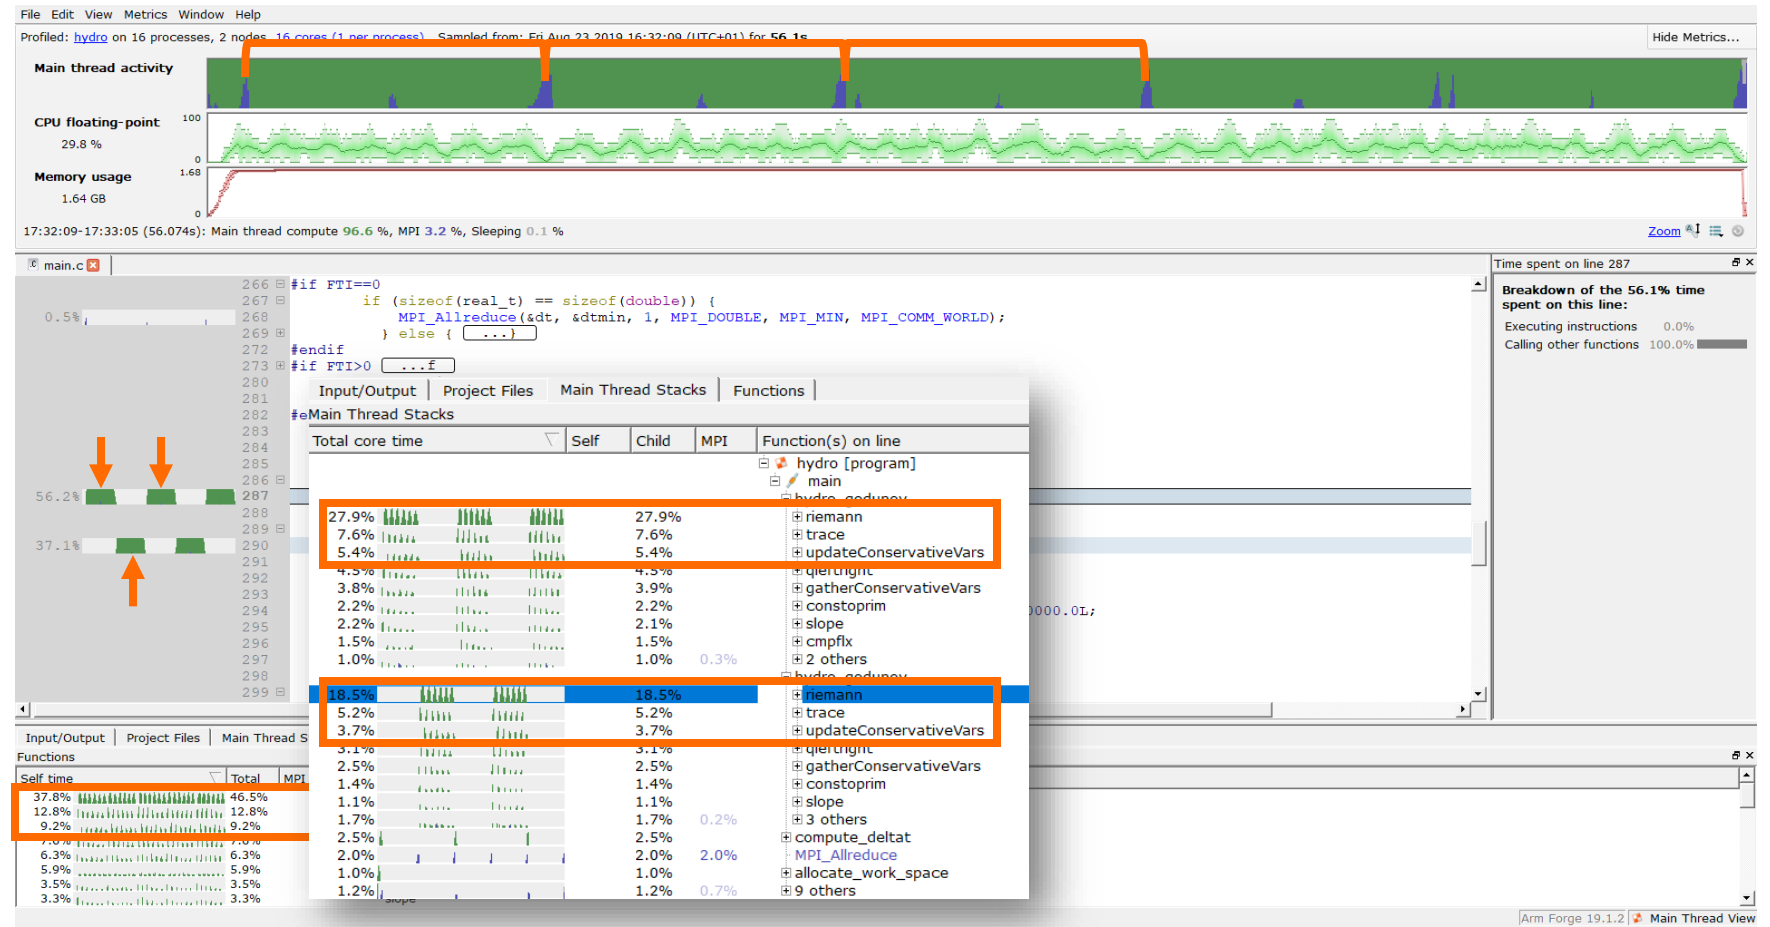
# Analyzing profiling results

- 2 nodes, 16 processes
- Runtime: 55s
- Iterative pattern
- Bottlenecks:
  - *riemann*
  - *trace*
  - *updateConservativeVars*



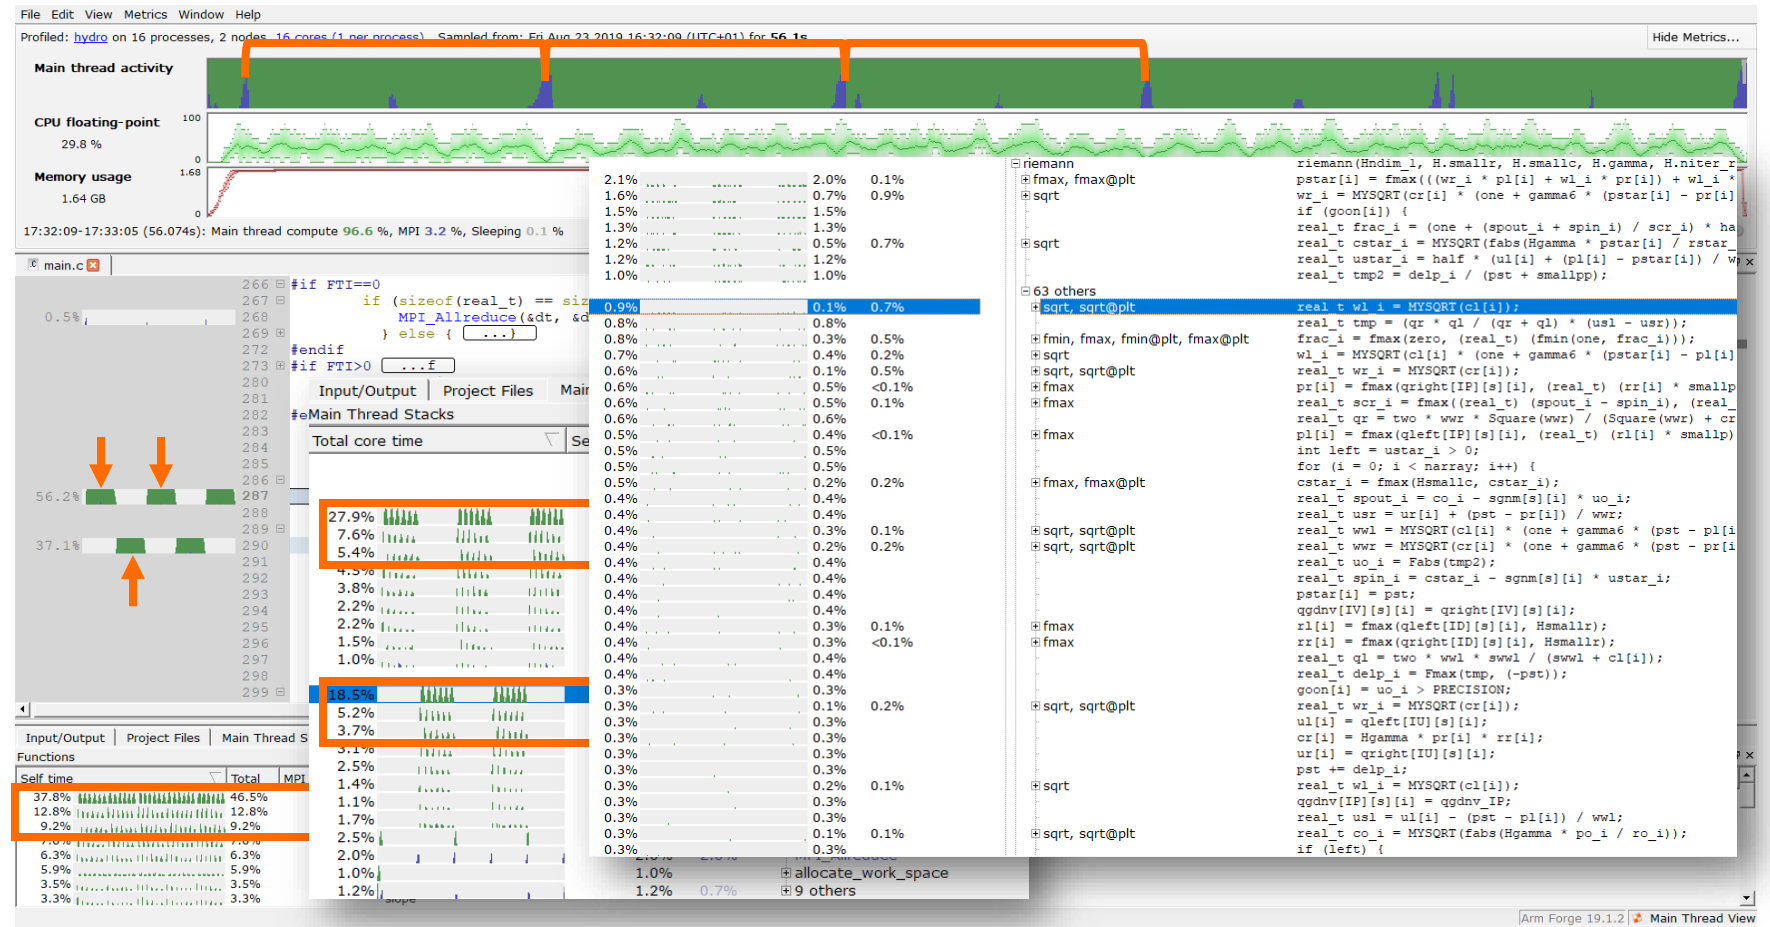
# Analyzing profiling results

- 2 nodes, 16 processes
- Runtime: 55s
- Iterative pattern
- Bottlenecks:
  - *riemann*
  - *trace*
  - *updateConservativeVars*
- Called in 2 code paths



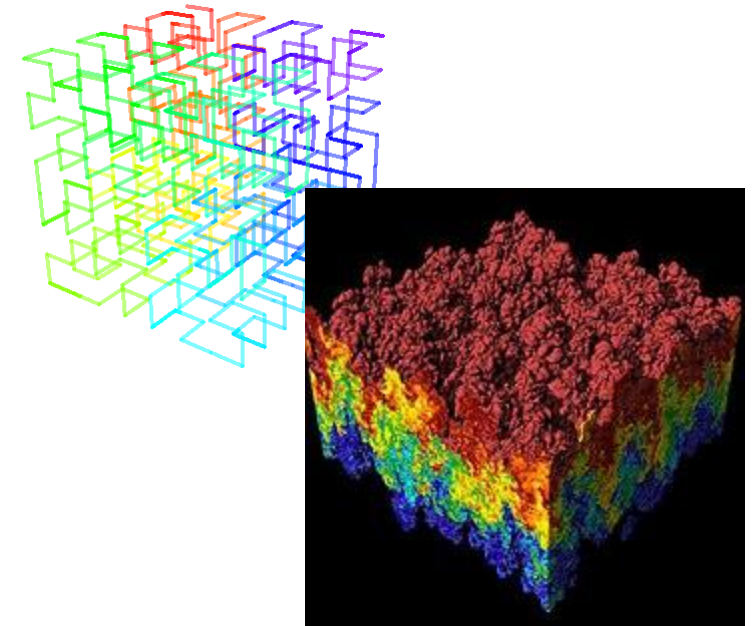
# Analyzing profiling results

- 2 nodes, 16 processes
- Runtime: 55s
- Iterative pattern
- Bottlenecks:
  - *riemann*
  - *trace*
  - *updateConservativeVars*
- Called in 2 code paths
- *updateConservativeVars*
  - 77% in 3 lines of code
  - Total of 70 lines of code
- *trace*
  - 50% in 11 lines of code
  - Total of 200 l.o.c.
- *riemann*: flat profile
  - 44% in 7 lines of code
  - Total of 340 l.o.c.



# Code region instrumentation with Caliper

- Performance data collection and analysis tool
  - <https://computing.llnl.gov/projects/caliper>
- Instrumentation-based for C, C++, FORTRAN codes
- Provides application and performance introspection
- Features:
  - Profiling, tracing of code regions,
  - MPI, PAPI stats
  - Annotations for third-party tools: TAU, Nvprof, now Arm MAP



# Instrument functions with Caliper

- Let's instrument the following functions:
  - *riemann*
  - *trace*
  - *updateConvervativeVars*
  - And a few others: *qlleftright*, *gatherConvervativeVar*, *constoprim*, *slope*
- Example: *riemann.c*

```
void riemann(int narray, const real_t Hsmallr, ... )
{
    int i, s, ii, iimx;
    ...

    CALI_MARK_FUNCTION_BEGIN;
    [...]
    CALI_MARK_FUNCTION_END;
}
```

# Instrument code sections with Caliper

- In *riemann*, the scientists left a few comments to describe the physics taking place
- Pseudo-code:

```
#compute pressure, density, velocity for each slice
for(s=0; s<=slices; s++)
{
    #precompute values for this slice
    for(i=0; i<=narray; i++)
    { [...] }
    #solve the Riemann problem on the interfaces of this slice
    for(iter=0; iter<=Hniter_riemann; iter++)
    { [...] }
    #main Riemann solver on arrays of this slice
    for(i=0; i<=narray; i++)
    { [...] }
}
```

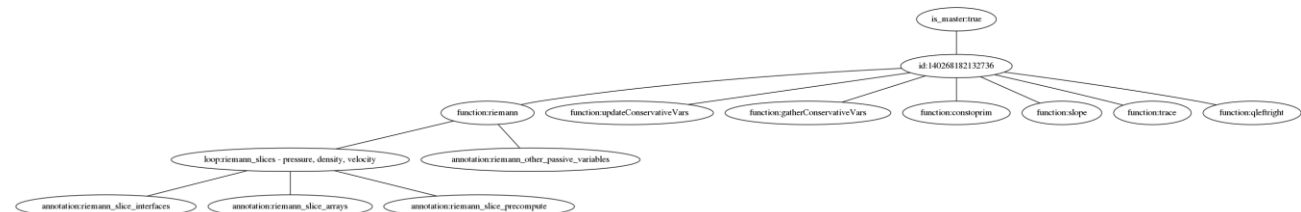
Caliper

```
CALI_MARK_LOOP_BEGIN(riemann_slice_id);
#compute pressure, density, velocity for each slice
for(s=0; s<=slices; s++)
{
    CALI_MARK_ITERATION_BEGIN(riemann_slice_id, s);
    CALI_MARK_BEGIN("riemann_slice_precompute");
    for(i=0; i<=narray; i++)
    { [...] }
    CALI_MARK_END("riemann_slice_precompute");
    CALI_MARK_BEGIN("riemann_slice_interfaces");
    for(iter=0; iter<=Hniter_riemann; iter++)
    { [...] }
    CALI_MARK_END("riemann_slice_interfaces");
    CALI_MARK_BEGIN("riemann_slice_arrays");
    for(i=0; i<=narray; i++)
    { [...] }
    CALI_MARK_END("riemann_slice_arrays");
}
CALI_MARK_LOOP_END(riemann_slice_id);
```

# Generating Caliper reports

- Modules or “services” provide measurement data
  - E.g. PAPI, CUPTI
- Can be configured through configuration files and environment variables
  - `$ CALI_CONFIG_PROFILE=runtime-report ./app`
- Output can be in stdout or in \*.cali files for post-processing

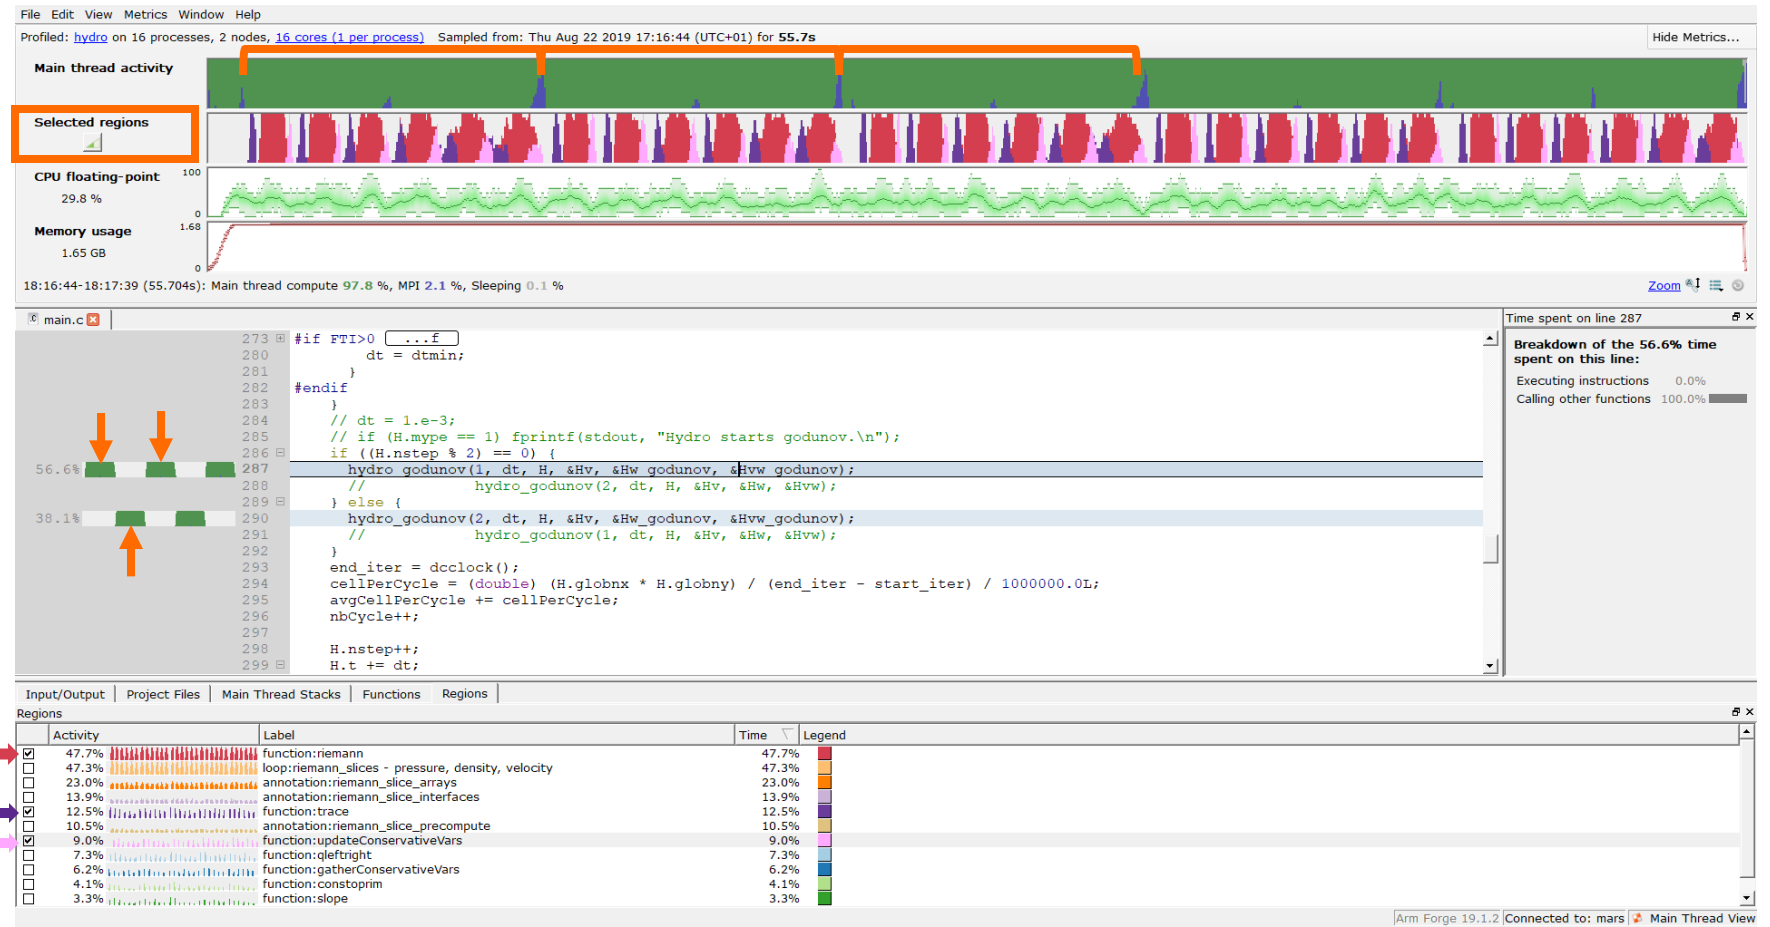
Path	Inclusive time (usec)	Exclusive time (usec)	Time (%)
updateConservativeVars	1637063.000000	1637063.000000	7.955088
riemann	8586175.000000	1317.000000	0.006400
riemann_slices - pressure	8584307.000000	1190957.000000	5.787295
riemann_slice_arrays	2907784.000000	2907784.000000	14.129986
riemann_slice_interfaces	2873562.000000	2873562.000000	13.963688
riemann_slice_precompute	1612004.000000	1612004.000000	7.833317
qlftright	1787885.000000	1787885.000000	8.687987
trace	2218274.000000	2218274.000000	10.779404
slope	1037166.000000	1037166.000000	5.039969
constoprim	1195203.000000	1195203.000000	5.807928
gatherConservativeVars	1559916.000000	1559916.000000	7.580202



```
$ cali-query -s "function=riemann" -a "sum(count)" -f "%[10]function% %[40]annotation% %[50]loop% %[10]count%" -o cali_selected.out
-T "FUNCTION ANNOTATIONS LOOP COUNT" hydro_0.cali
FUNCTION  ANNOTATIONS                                LOOP                                COUNT
riemann   riemann_other_passive_variables              30
riemann   riemann_slice_precompute        riemann_slices - pressure, density, velocity  30000
riemann   riemann_slice_interfaces        riemann_slices - pressure, density, velocity  30000
riemann   riemann_slice_arrays            riemann_slices - pressure, density, velocity  30000
```

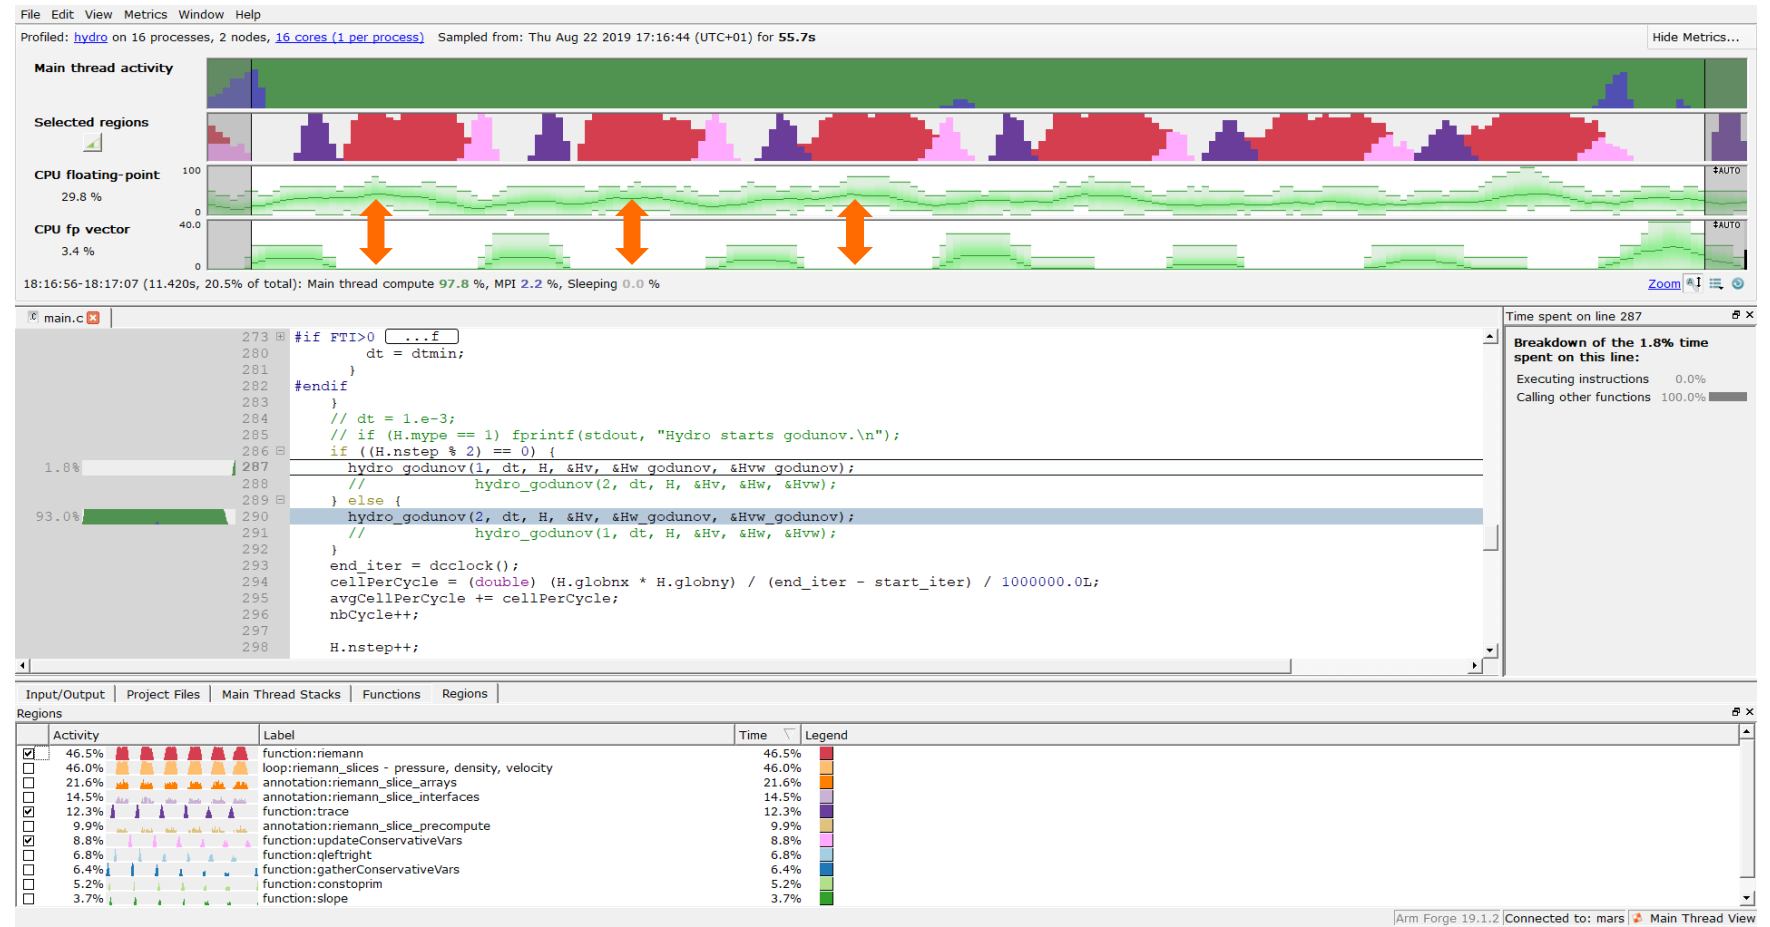
# Visualize Caliper data in Arm MAP

- A “Selected region” is displayed
- Caliper functions can be selected
- Iterative patterns appear more clearly
- The timeline shows how bottleneck functions are called
  - *riemann*
  - *trace*
  - *updateConservativeVars*



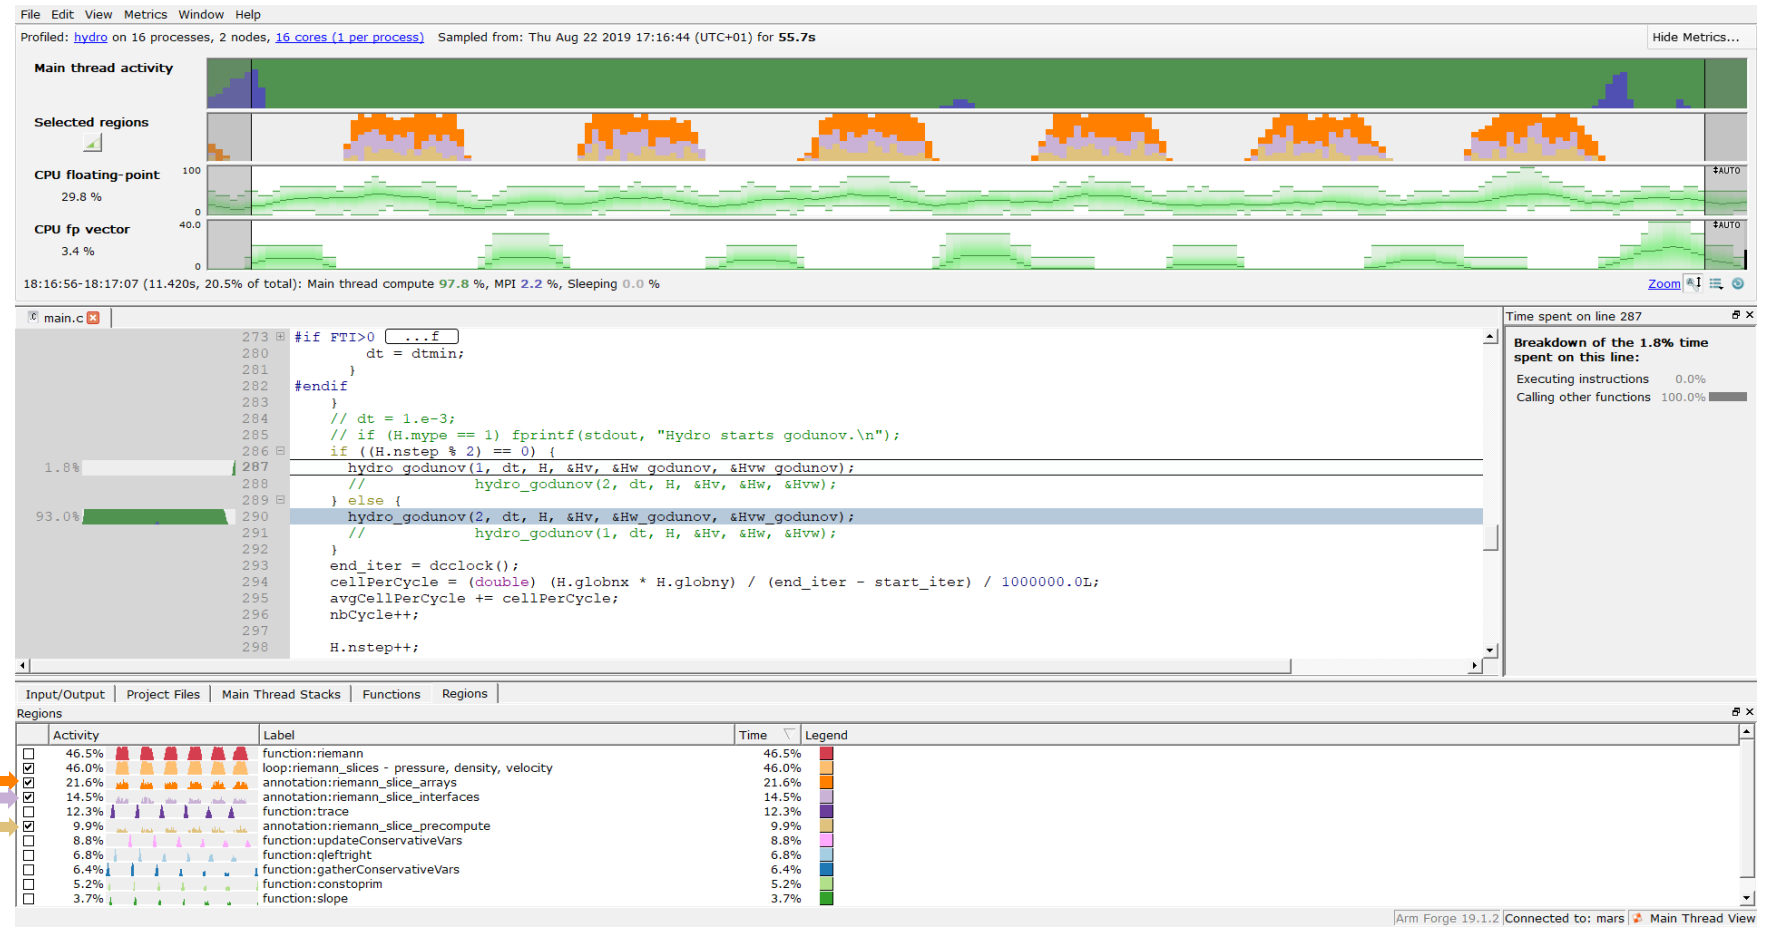
# Analyze Caliper data in Arm MAP

- Zooming in enables to correlate function calls with metrics
- e.g. the example show peak CPU floating point when the *riemann* function is executed
- However, no vectorization is performed



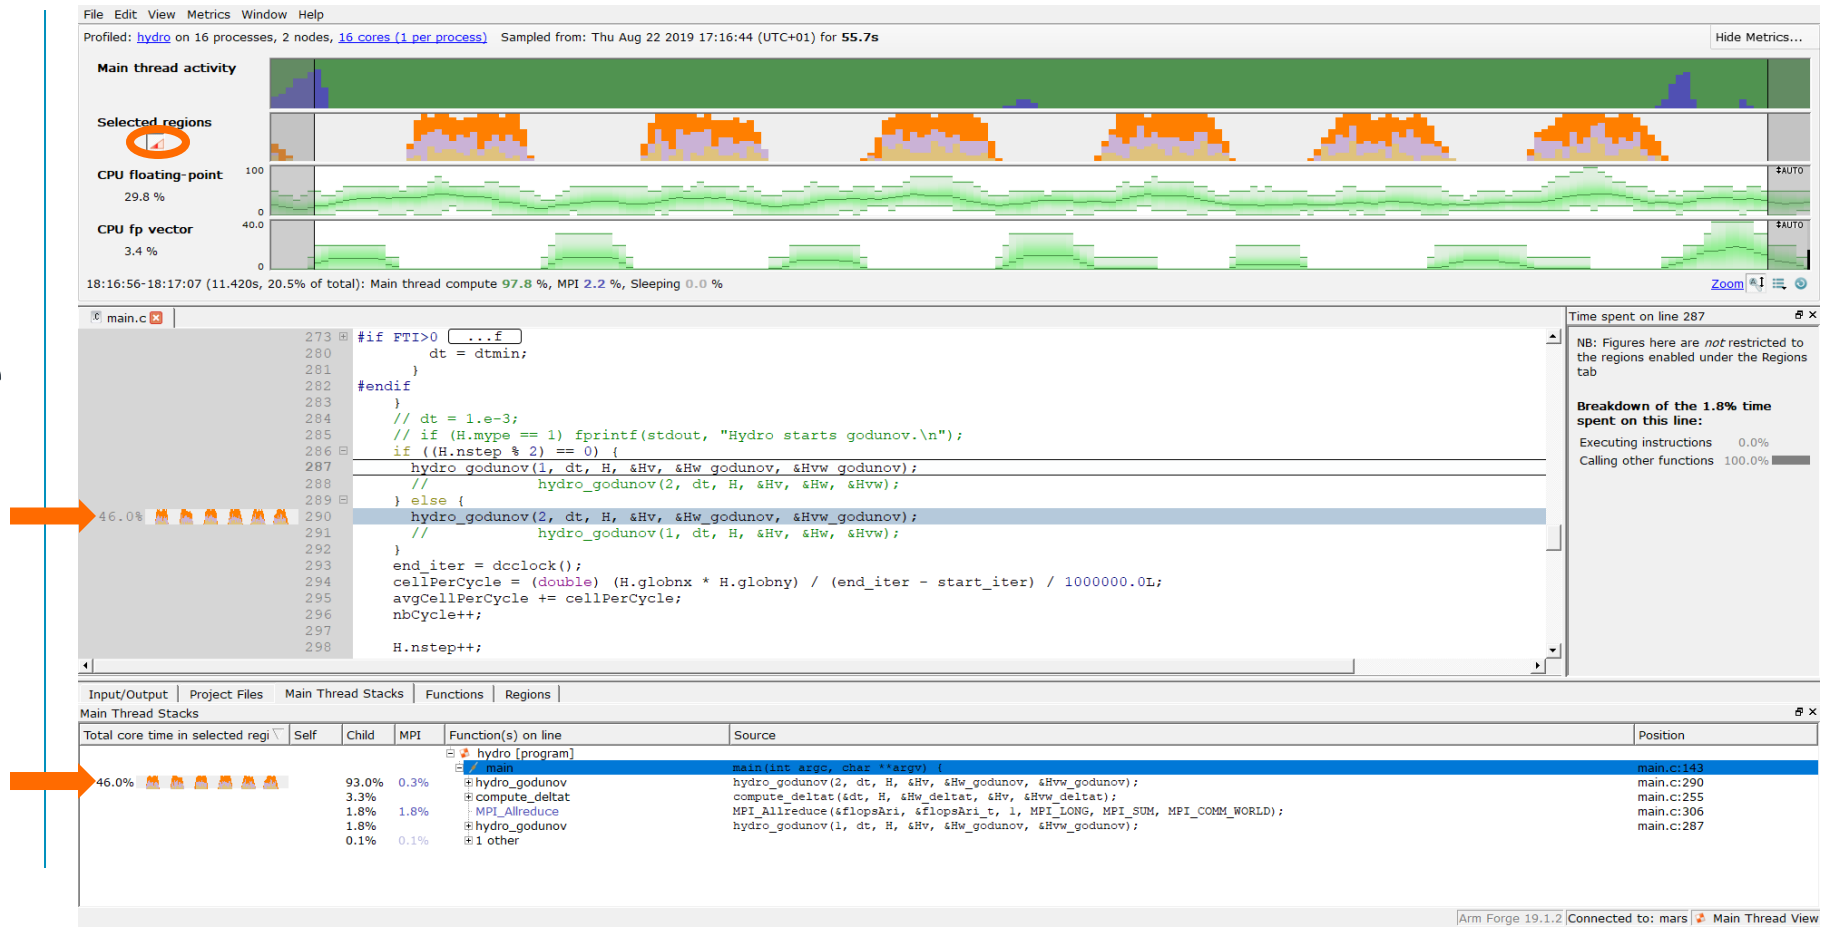
# Analyze Caliper data in Arm MAP

- Enabling *riemann* code regions give insights about how the time is spent inside the function
  - 22% in slice array solver
  - 15% in slice interface solver
  - 10% in slice precompute region



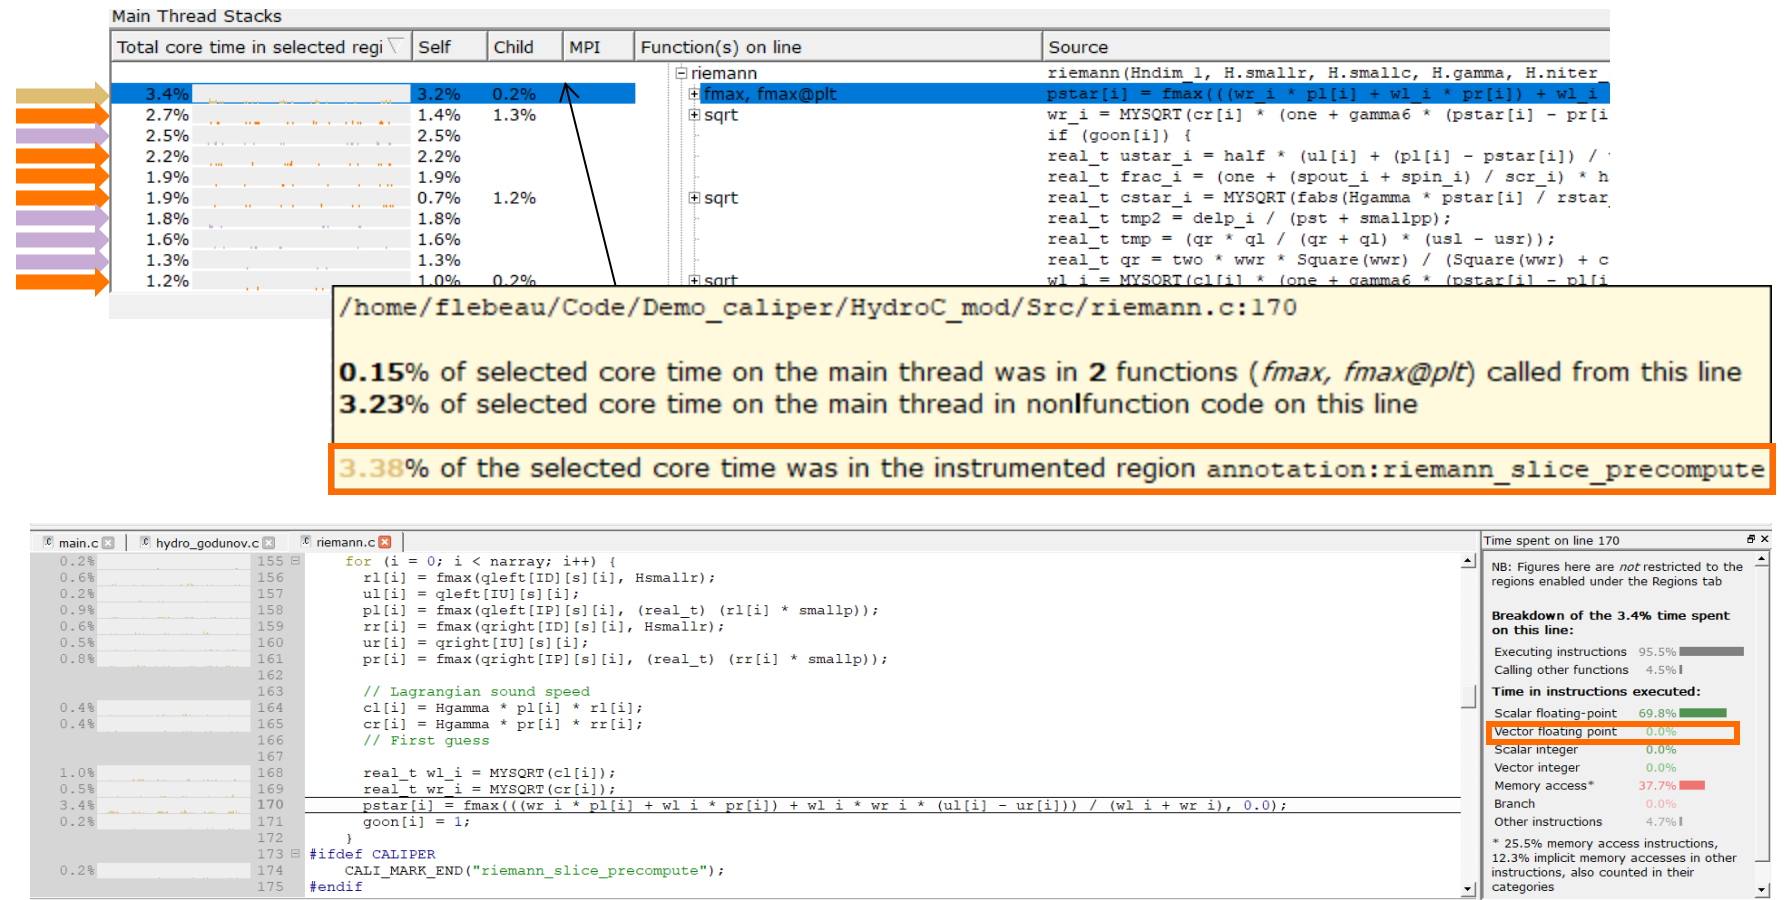
# Analyze Caliper data in Arm MAP

- The “Activate regions focused view” is available
- It displays Caliper information in the source code annotations and in the “Main thread stack”



# Ensure vectorization of specific code regions

- The “Main thread stack” gives code semantic information
- In 1 iteration
  - 5 lines (~10%) in slice array solver
  - 4 lines (~7%) in slice interface solver
  - 1 line (~3%) in slice precompute region
- When selecting the source code lines, the code appear not to be vectorized

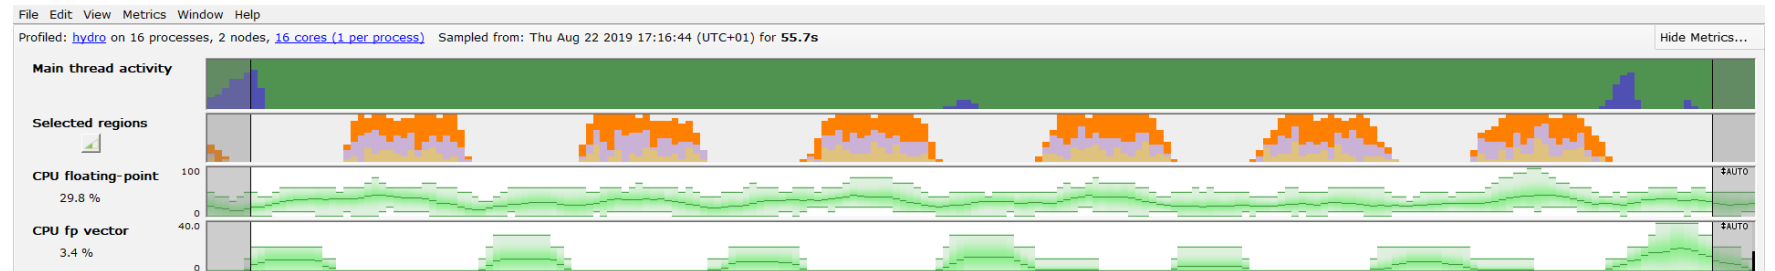


# Ensure vectorization of specific code regions

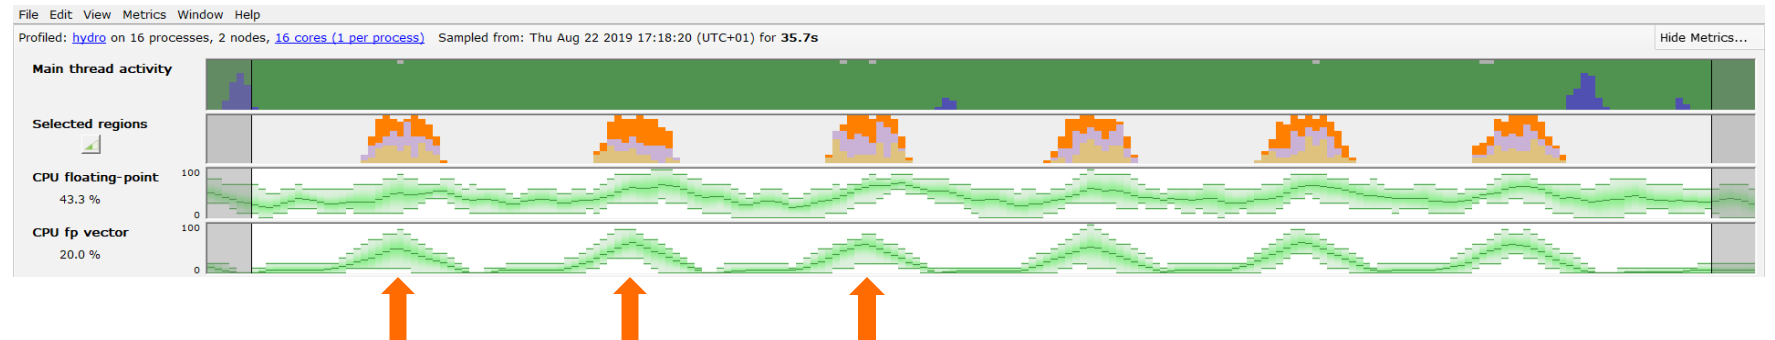
Use of `#pragma omp simd`

- 1.5 speedup
- Increased floating point performance
  - 43% vs. 30%
  - 20% vector vs. 3%

## Without pragma



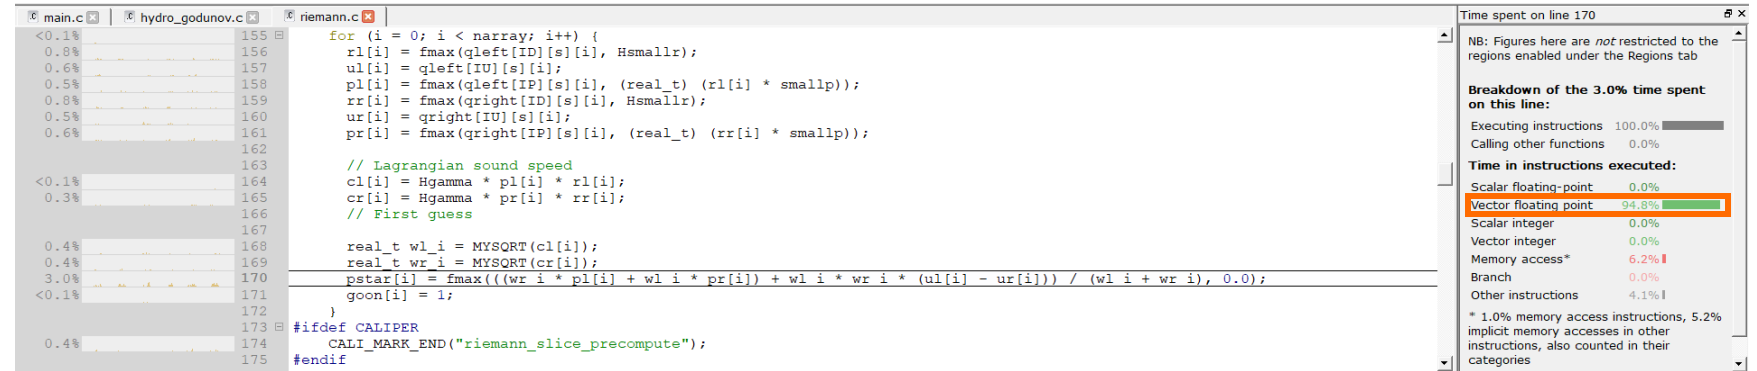
## With pragma



# Ensure vectorization of specific code regions

## Use of #pragma omp simd

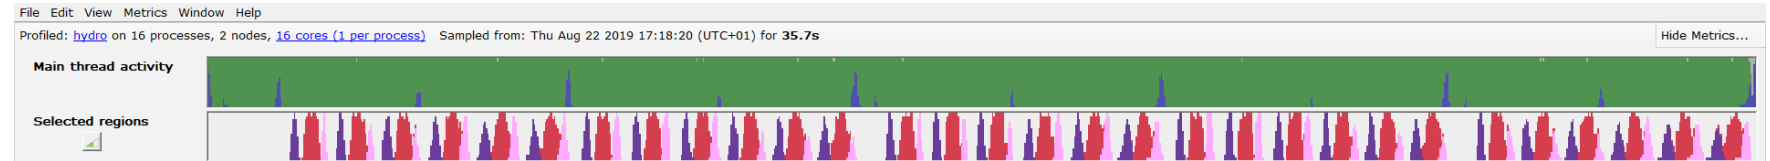
- 1.5 speedup
- Increased floating point performance
  - 43% vs. 30%
  - 20% vector vs. 3%
- Lines of code are well vectorized
- Riemann profile
  - **10%** in slice array solver
  - **8%** in slice interface solver
  - **9%** in slice precompute region



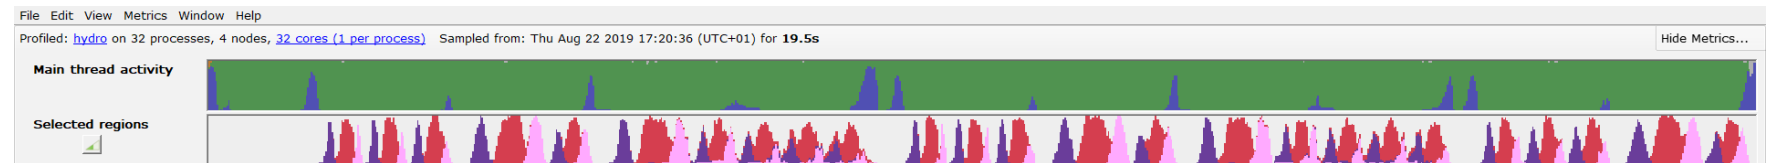
Regions		
	Activity	Label
<input type="checkbox"/>	27.6%	function:riemann
<input checked="" type="checkbox"/>	27.2%	loop:riemann_slices - pressure, density, velocity
<input type="checkbox"/>	12.7%	function:trace
<input type="checkbox"/>	11.2%	function:qlleftright
<input checked="" type="checkbox"/>	10.4%	annotation:riemann_slice_arrays
<input type="checkbox"/>	9.5%	function:updateConservativeVars
<input type="checkbox"/>	8.6%	function:gatherConservativeVars
<input checked="" type="checkbox"/>	8.5%	annotation:riemann_slice_precompute
<input checked="" type="checkbox"/>	8.4%	annotation:riemann_slice_interfaces
<input type="checkbox"/>	7.9%	function:constoprim
<input type="checkbox"/>	6.2%	function:slope

# Scale up

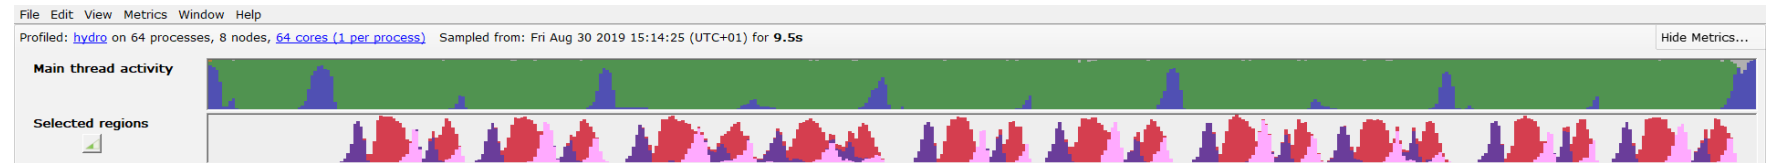
- 2 nodes



- 4 nodes

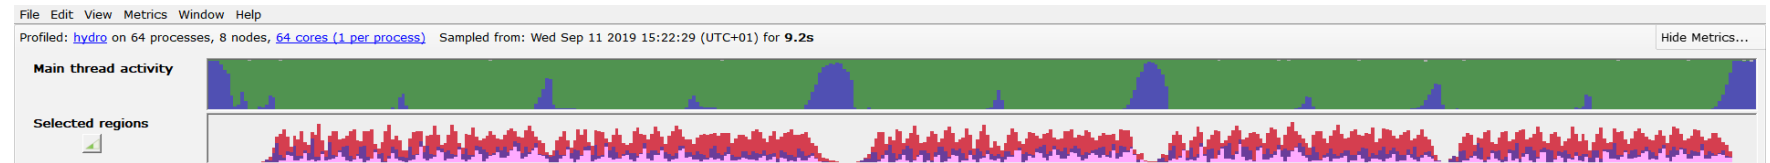
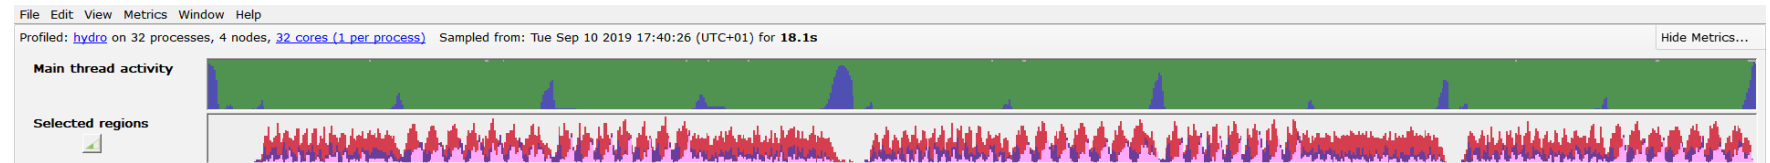
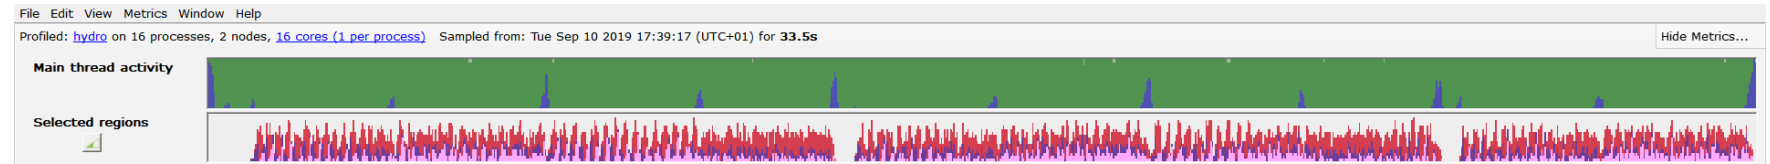


- 8 nodes



# Scale up – with different input data

- 2 nodes
- 4 nodes
- 8 nodes

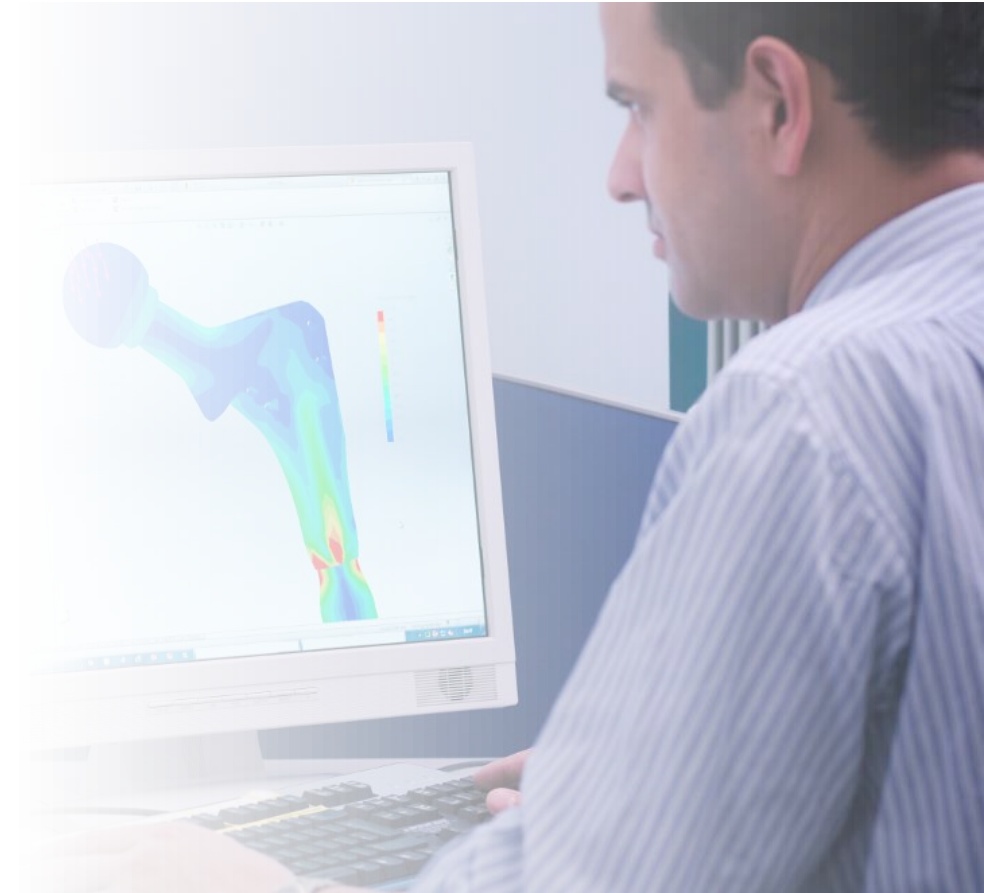


# Conclusion: topics covered

- Profiling at scale with Arm MAP
- Use Caliper annotations to define code regions
  - Supported by MAP version 19.1 onwards
  - Add application insight to profiling results
- Visualize results
  - Understand application behavior faster
  - Optimize code more efficiently

# Request a trial license

- If you would like to profile your own Caliper-instrumented applications, visit our [website](#)
- A trial licence of Arm Forge enables:
  - Multi-node and multi-threaded profiling with Arm MAP
  - On all HPC architectures
  - On C/C++, FORTRAN and Python codes
  - Debugging with Arm DDT
  - Access to our support channel
- Plenty of resources are available online for your trial: [developer.arm.com/hpc](https://developer.arm.com/hpc)



arm

Thank You

Danke

Merci

谢谢

ありがとう

Gracias

Kiitos

감사합니다

धन्यवाद

شكراً

תודה

# Questions

The ARM logo is displayed in a white, lowercase, sans-serif font. The letters are bold and closely spaced. The background is a dark blue with a subtle grid of small white plus signs.

The Arm trademarks featured in this presentation are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. All other marks featured may be trademarks of their respective owners.

[www.arm.com/company/policies/trademarks](http://www.arm.com/company/policies/trademarks)