



Debugging and Profiling

Aaron Birkland

Cornell Center for Advanced Computing

Special thanks to the Texas Advanced Computing Center for
some slide content.



Introduction

Debugging

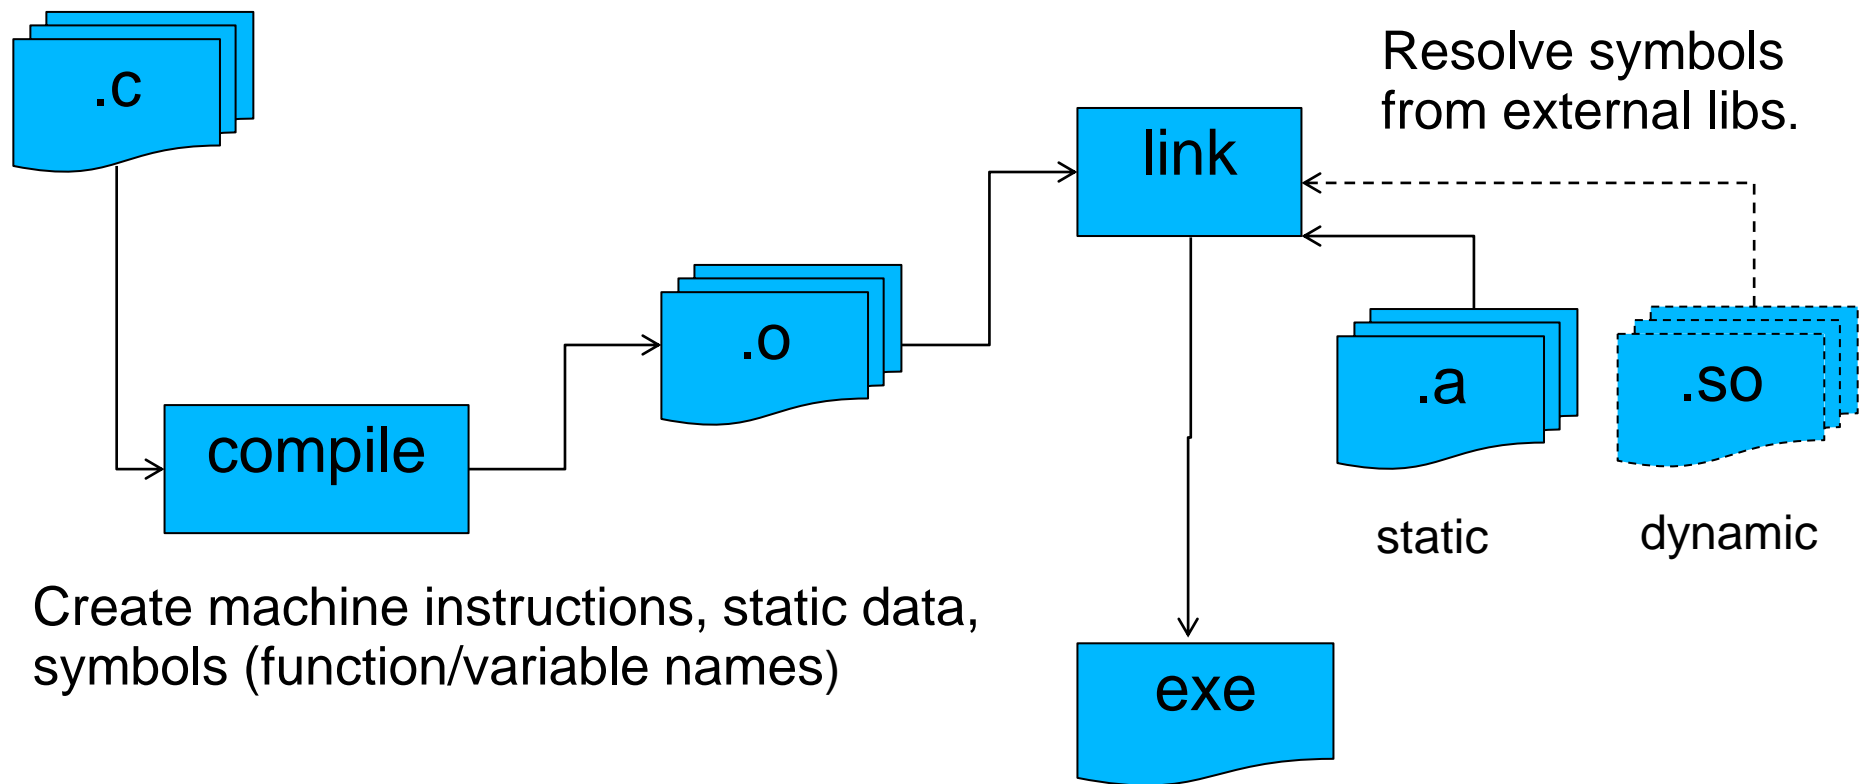
- Find defects, analyze failures, verify expected program flow.
- Debugger tools: Inspect or modify state of running program, post-mortem analysis of memory dumps.
- Harder in parallel!

Profiling

- Measure performance characteristics, Identify areas for improvement.
- Profiler tools: collect performance measurements of a running program, analyze afterward.
- Harder in parallel!

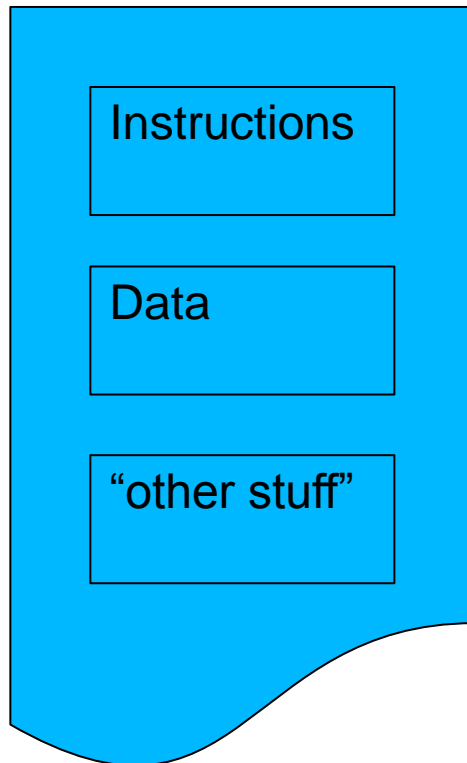


Background: Compiling/Linking





Background: Executable Files



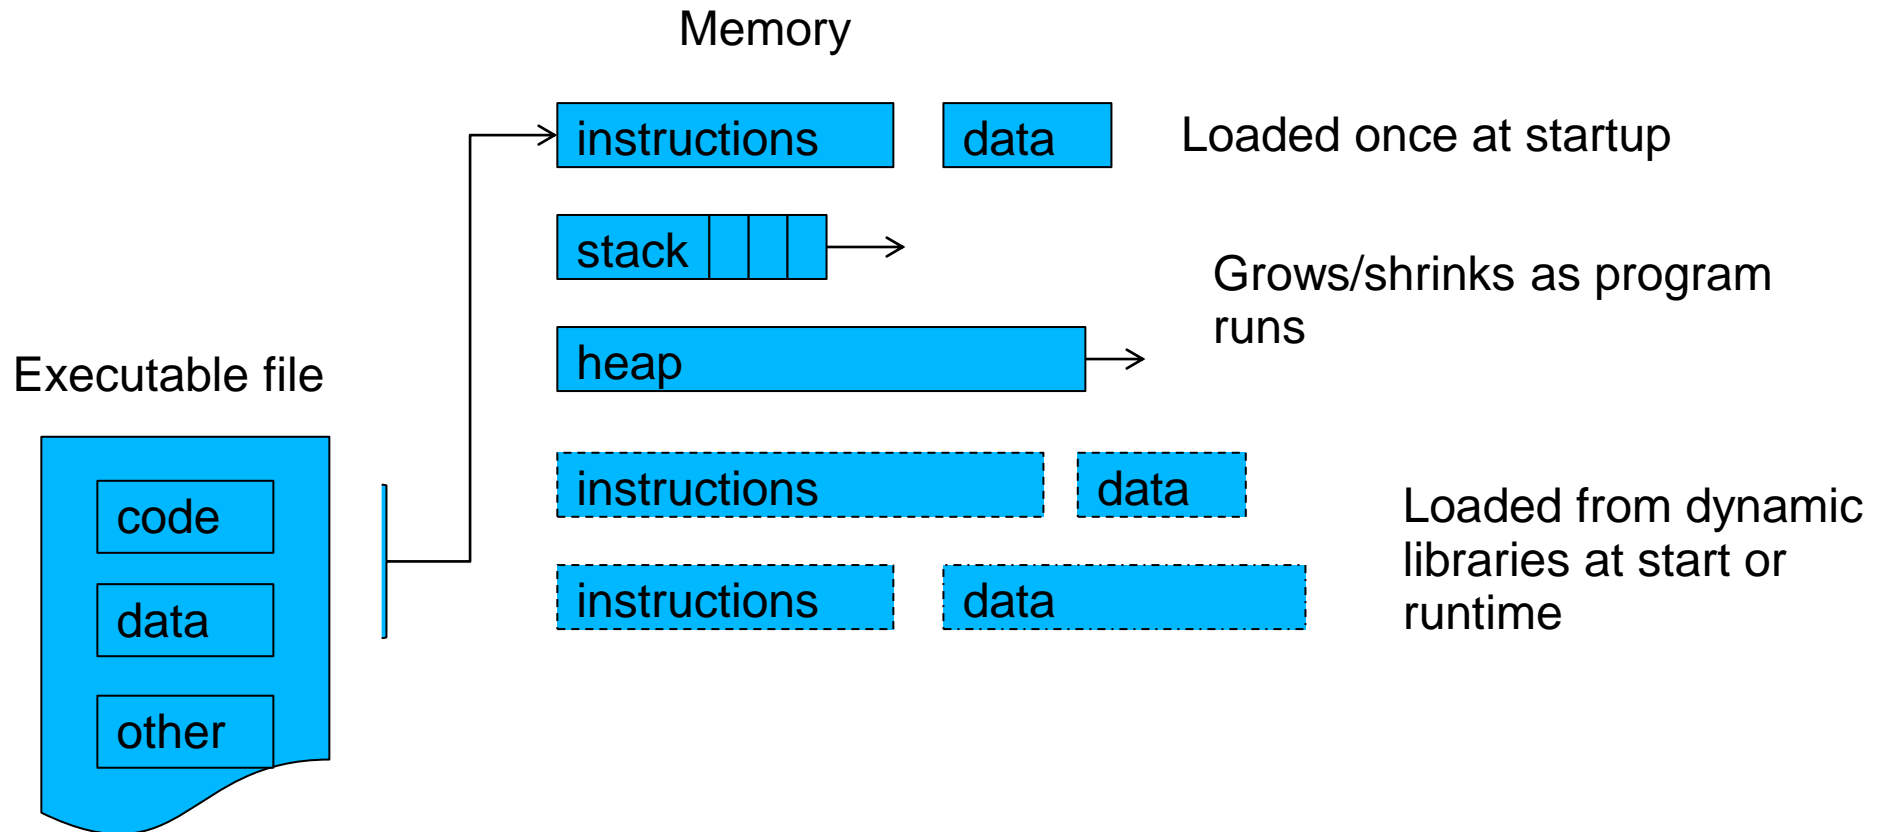
Machine instructions, memory addresses

Global and static variable data

Symbol table, linked library filenames,
compiler version, other metadata.

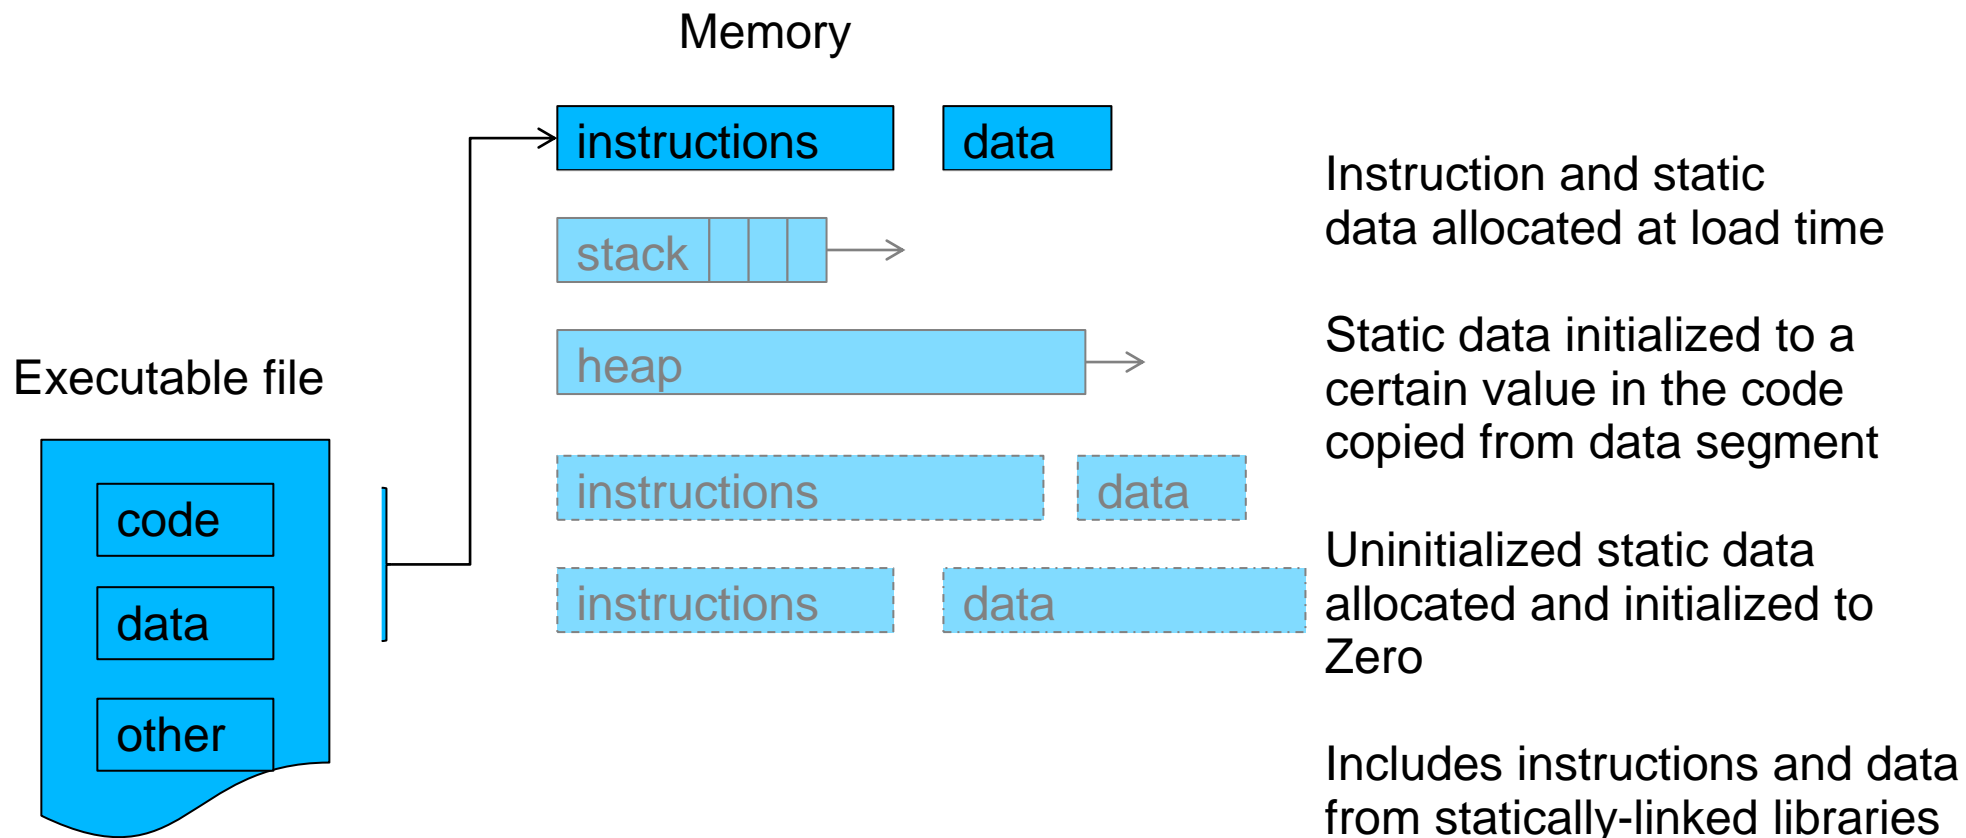


Background: Execution & Memory



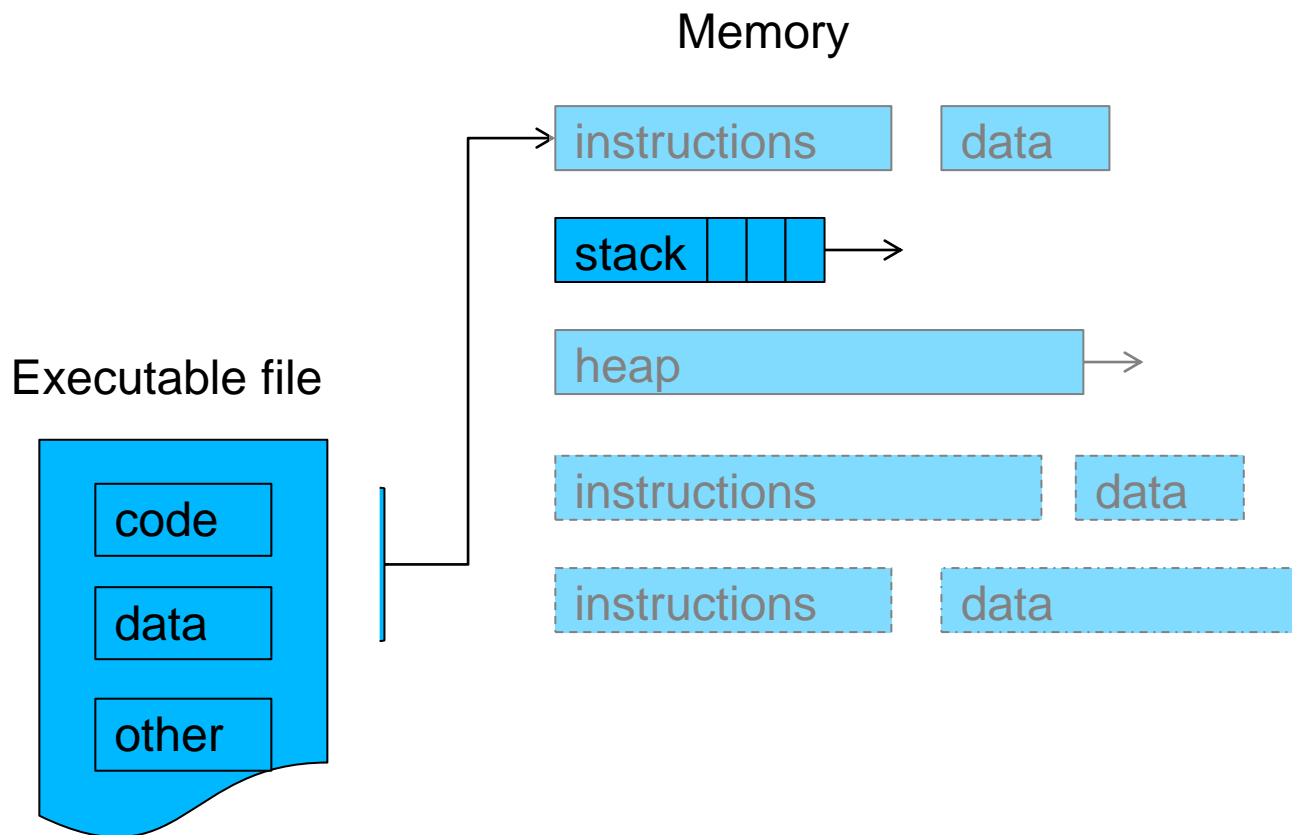


Background: Execution & Memory





Background: Execution & Memory



Stack composed of *frames* created/destroyed each time a function is called/exited

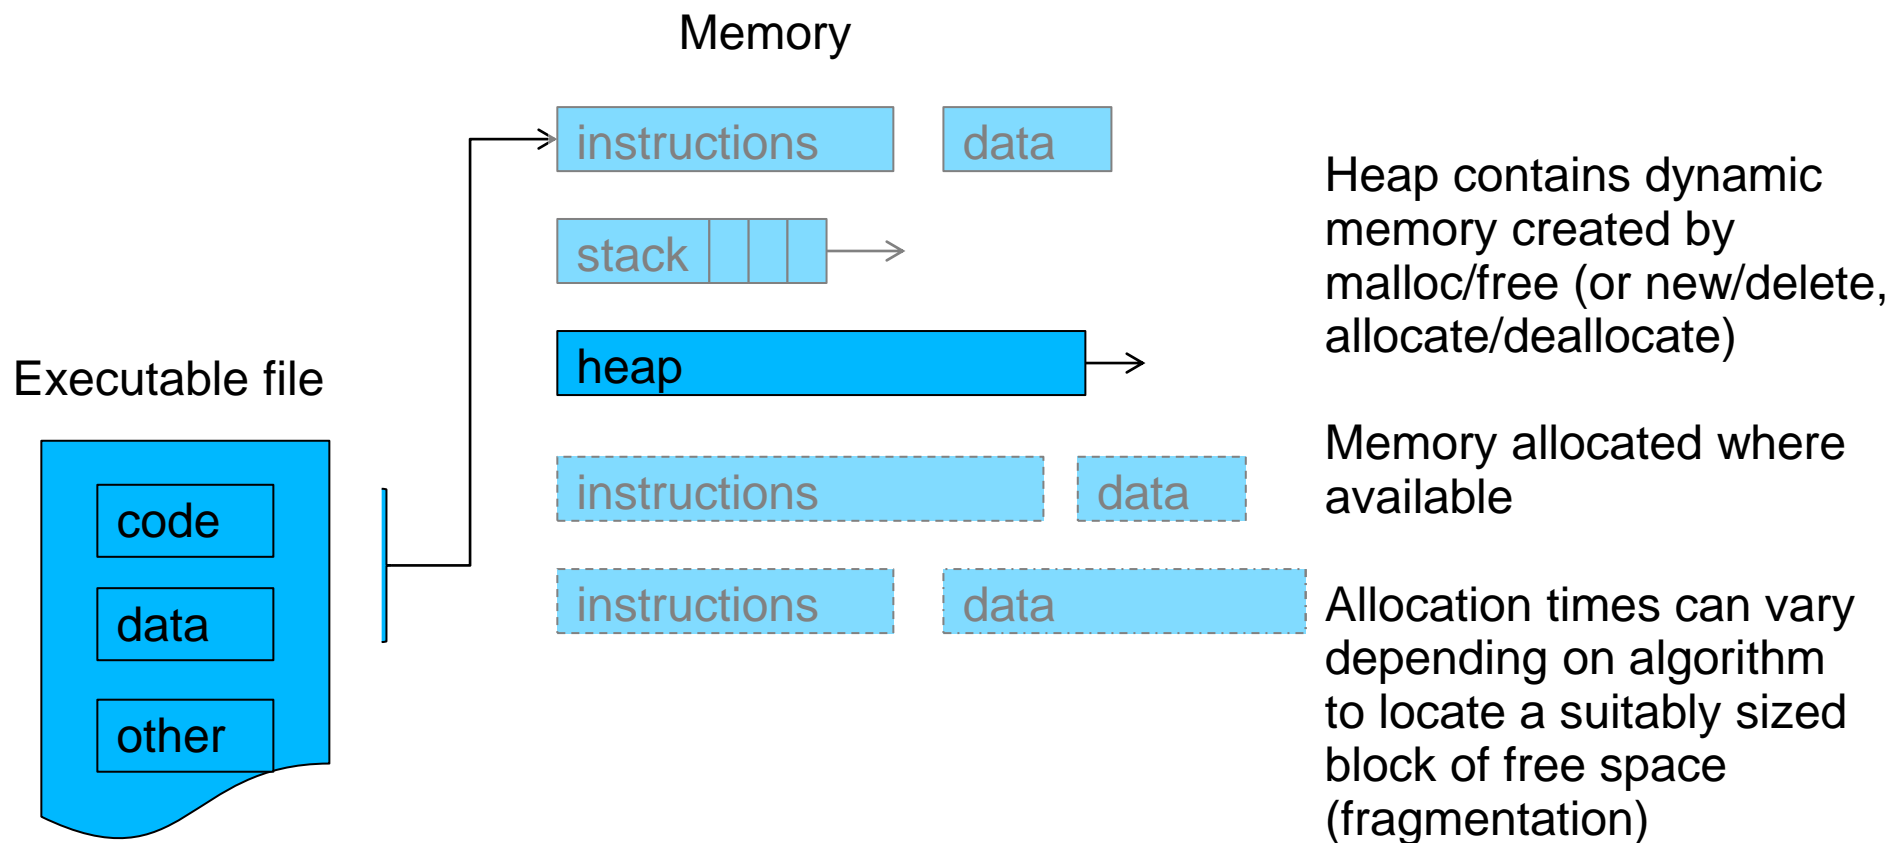
Each frame contains:

- All automatic local variables
- All arguments passed in
- Address of caller's frame

Frames are added/removed from *end* of stack only



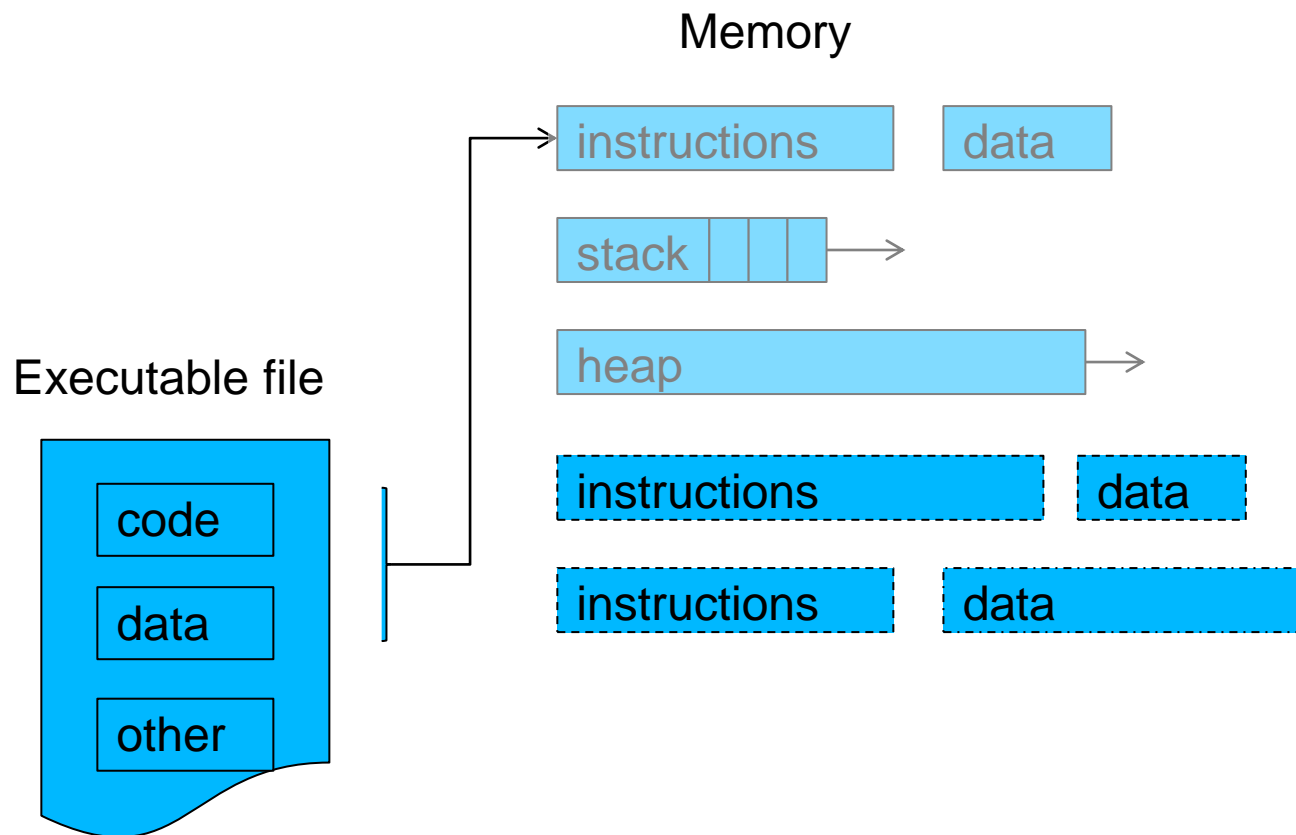
Background: Execution & Memory



Watch for memory leaks!



Background: Execution & Memory



“Other” segments in executable file contain names of dynamic libraries.

OS loads these libraries into Memory upon startup

Possible to load on demand (e.g. via `dlopen/dlsym`)

Won't run unless all referenced libraries are present! (undefined symbol)

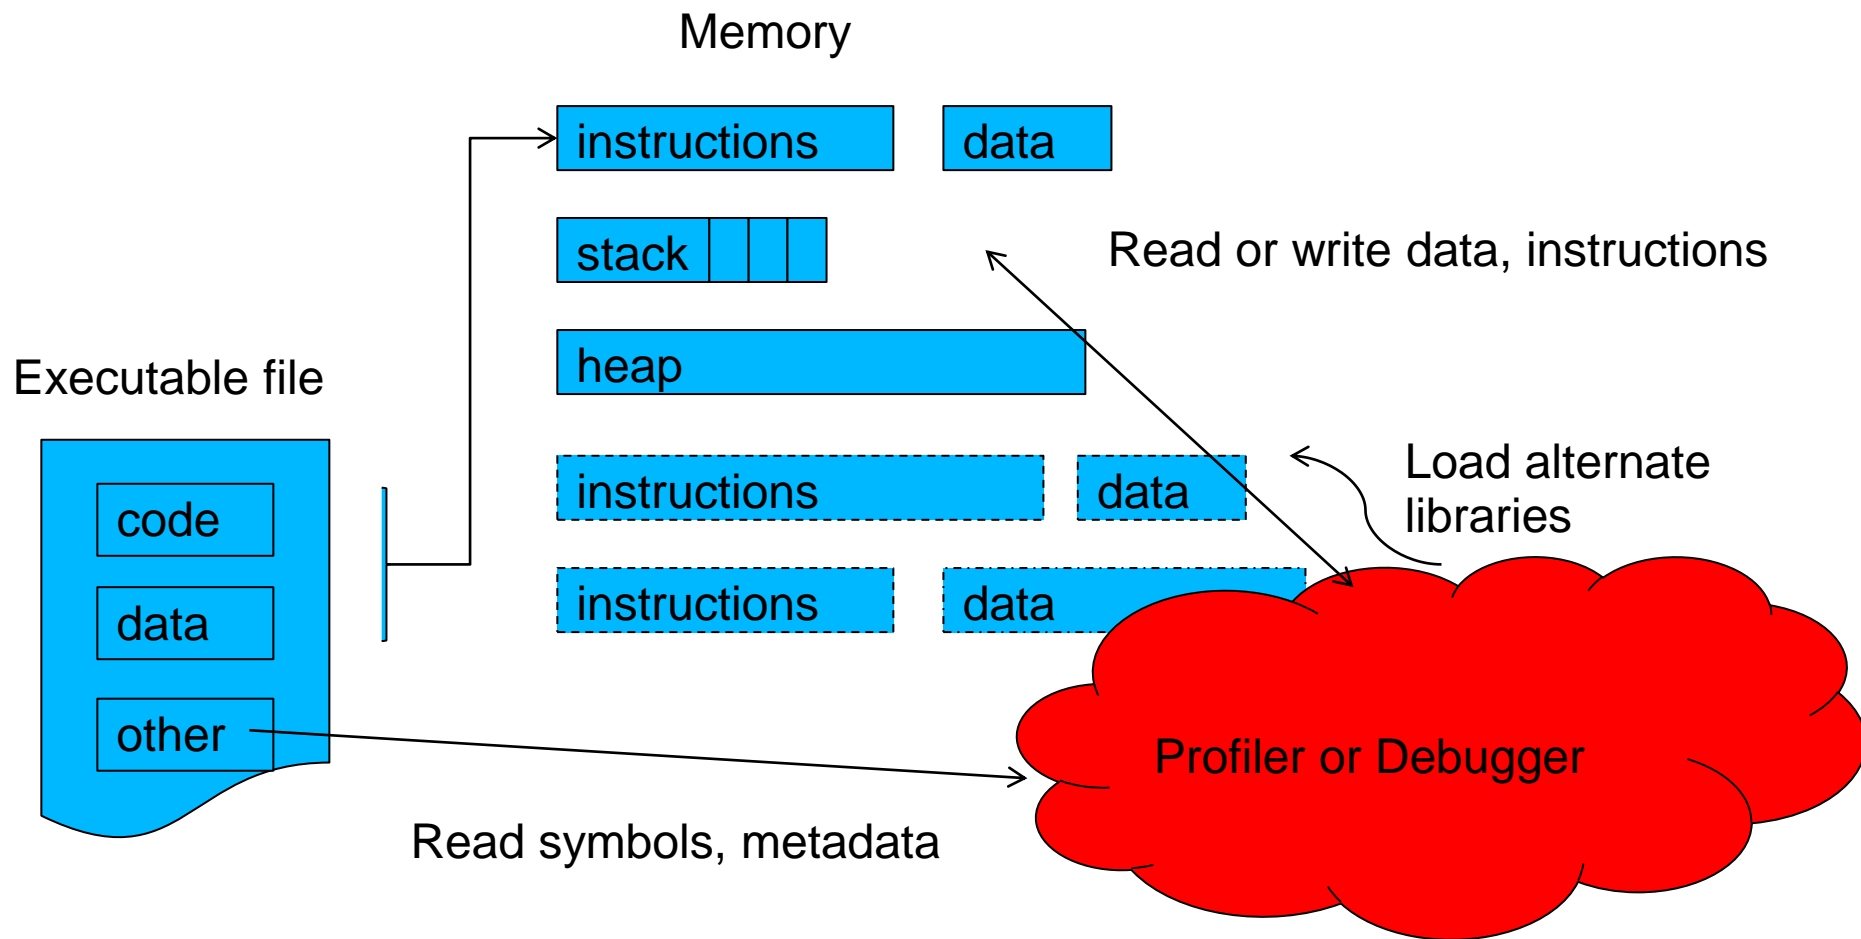


Background: OS and Hardware

- OS can provide API for inspecting and controlling process execution
- Wrap a program at startup or attach to running process
- Example: Linux `ptrace()`
 - Pause execution
 - Modify in-memory instructions
 - Inspect or modify data memory or registers
 - Catch signals and traps
- CPU can provide hardware counters
 - Cache hits/misses, TLB hits/misses, FLOPs, etc



Background: Profilers and Debuggers in control





Debugging

- Inspect program state, compare to one's own assumptions and expectations
 - Step through code line by line
 - Inspect variables/memory at specific points
 - Inspect memory and call stack after a crash
- For MPI, OpenMP 'state' gets more complex
 - Many remote processes with own memory
 - Message status and timing
 - Step through individual processes or thread independent of rest (while others may still be running!)



Debugging: printf and logging

```
int main (int argc, char** argv) {  
    printf("Starting main...");  
    int iterations = 5;  
    int val = 0, val2=0;  
    printf("Initialized val to %d and val2 to %d", val, val2);  
    while (iterations --) {  
        val = sometime();  
        print("Sometime() returned %d\n", val);  
        val2 = moretime();  
        printf("moretime() returned %d\n", val);  
    }  
    printf("Exiting main, iterations ==%s%d", iterations);  
}
```



Debugging: printf and logging

- Easy and intuitive
 - Target specific sections of code, under specific conditions
 - Simply analyze log(s) after execution, even for parallel or multithreaded jobs
 - Great for rare/transient or timing related bugs
- Invasive and messy
 - Need to re-compile when logging statement added/removed
 - Can slow down execution
 - Easy to forget statements are there
 - Can be hard to correlate output with statements.
 - Jumbled output with threads printing simultaneously



Debugging: printf and logging

- Logging frameworks an improvement over printf (e.g. Log4c)
 - Filter by log levels (WARN, INFO, DEBUG)
 - Timestamps, formatting, runtime configuration changes
 - Control over where/how log is written (console, large file, rolling file, remote server, database, etc)



Debugging: printf and logging

```
int main (int argc, char** argv) {  
    log4c_init();  
    mycat = log4c_category_get("sillyapp.main");  
    int iterations = 5;  
    log4c_category_log(mycat, LOG4C_PRIORITY_DEBUG, "Debugging app 1  
- loop %d", iterations);  
    int val = 0, val2=0;  
    log4c_category_log(mycat, LOG4C_PRIORITY_ERROR, "Some error"  
printf("Initialized val to %d and val2 to %d", val, val2);  
    ...  
}
```

[Header]

2009-05-13 15:21:14,315 [11] WARN Logger.Program Pretty sure I'm getting ready to die!

2009-05-13 15:21:14,331 [11] ERROR Logger.Program uh-oh, no I wasn't!


2009-05-13 15:21:14,331 [11] FATAL Logger.Program blech. Out

[Footer]



Debugging: symbolic debugging

- Inspect process memory, correlate instructions & memory addresses with *symbols* from source code.
- Compiler option (`-g` for gcc, intel) tells compiler to store debugging symbols in the executable file

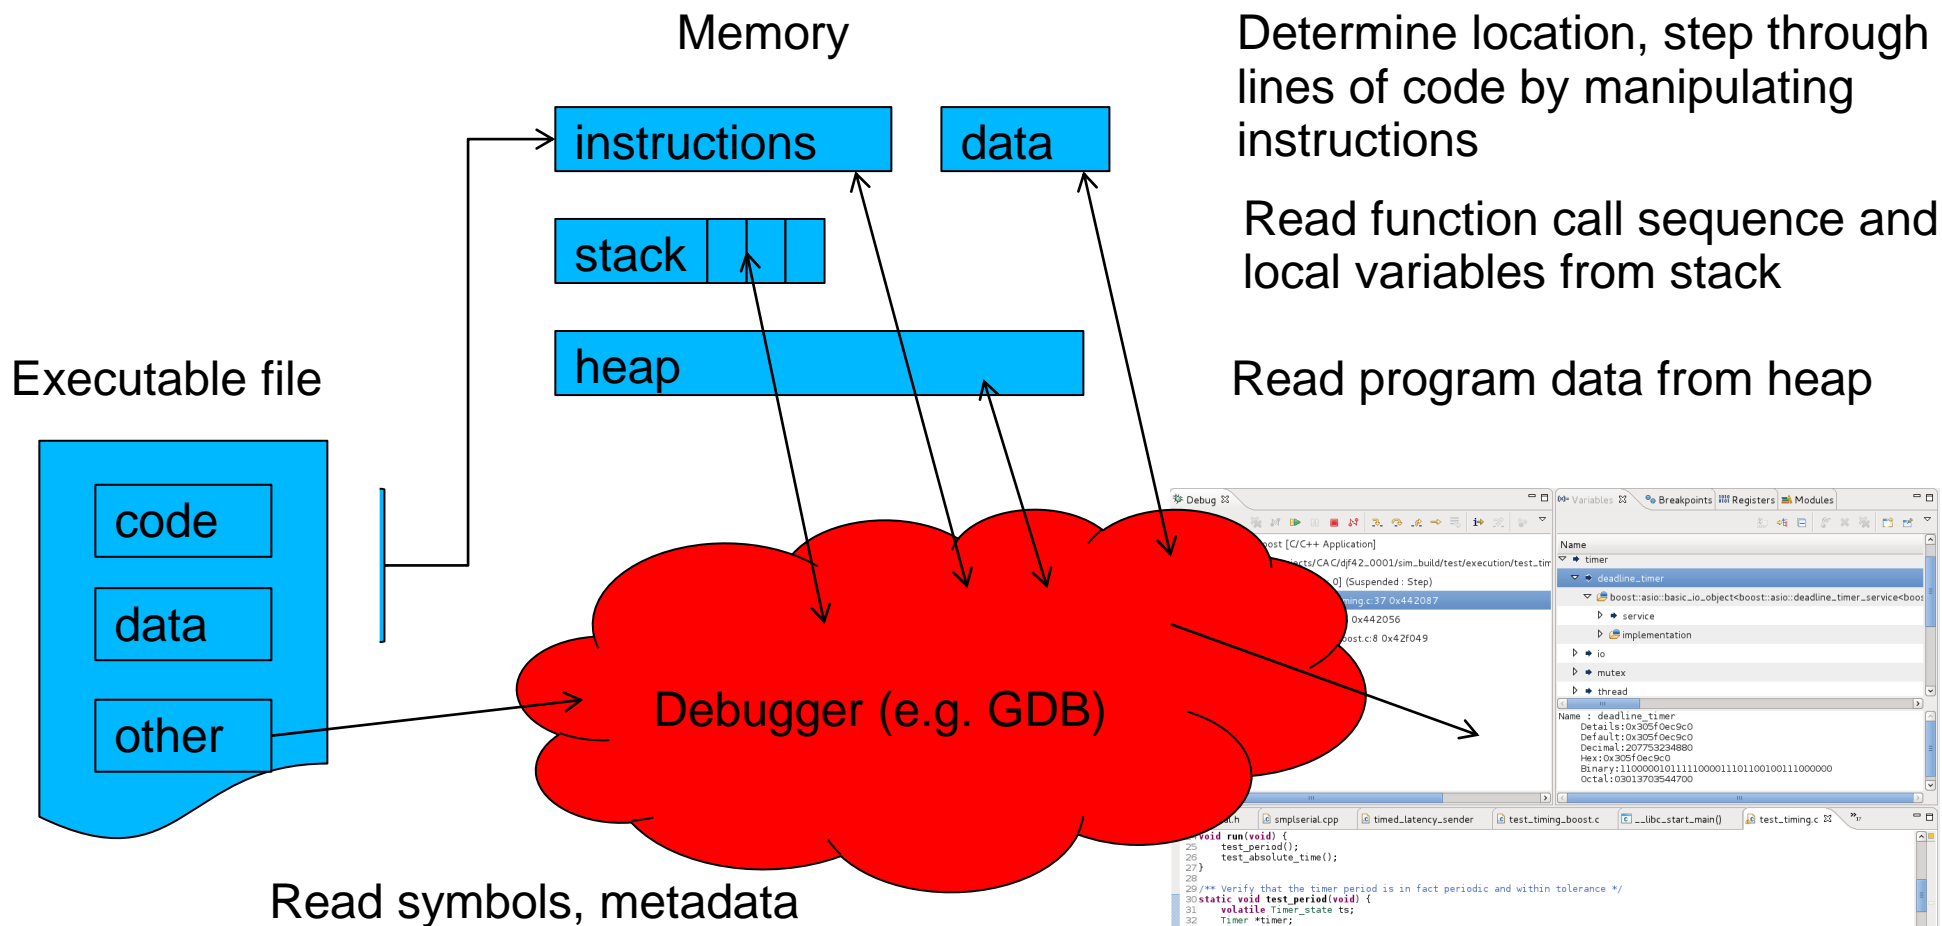


Human-readable symbols and correlation data stored in one of the “other” segments in an executable file.

- Not loaded into memory (no runtime overhead)
 - Some compilers MAY disable some optimizations
- Available for inspection by debugging tool
- Provides a very useful “map” for inspecting core dumps



Debugging: symbolic debugging: serial, threaded





Debugging: symbolic debugging: serial, threaded

- GDB (Gnu, almost ubiquitous), IDB (Intel)
 - Launch a program, analyze a dump, or attach to running process
 - Set conditional breakpoints, start/stop execution at will
 - Inspect and modify variables

Launch a process: `gdb <executable>`

Attach to process: `gdb <executable> 1234`

Analyze a dump: `gdb <executable> core.1234`
(check ulimit setting for max core file size!)



Debugging: symbolic debugging: GDB

- run – execute the program from beginning.
- backtrace – produce the backtrace from the last fault
- break <line number> or break <function-name> - break at the line number or at the use of the function
- step – step to next line of code (step into function if possible)
- next – step to next line of code (do not step into function)
- print <variable name> - print the value stored by the variable
- continue – run until next break point



Debugging: symbolic debugging

The screenshot shows a debugger window with three main panes. The top-left pane is the 'Stack view', showing the call stack for the current thread. The top-right pane is the 'Variables' pane, showing the current scope and its members. The bottom pane is the 'Source code and execution' pane, showing the source code of the current file.

Stack view: Shows the call stack for the current thread. The current frame is `test_period() at test_timing.c:37 0x442087`. Other frames include `run() at test_timing.c:25 0x442056` and `main() at test_timing_boost.c:8 0x42f049`.

Variables: Shows the current scope and its members. The current scope is `timer`. The members are `deadline_timer`, `service`, `implementation`, `io`, `mutex`, and `thread`. The `deadline_timer` member is expanded, showing its details: `Name : deadline_timer`, `Details: 0x305f0ec9c0`, `Default: 0x305f0ec9c0`, `Decimal: 207753234880`, `Hex: 0x305f0ec9c0`, `Binary: 11000001011111000011101100100111000000`, and `Octal: 03013703544700`.

Source code and execution: Shows the source code of the current file, `test_timing.c`. The current line of execution is `timer = timing_impl->start(0, DESIRED_INTERVAL, _track_timer, (Timer_state *) &ts);` at line 37.

Stack view

Variables

Evaluation, expressions

Source code and execution

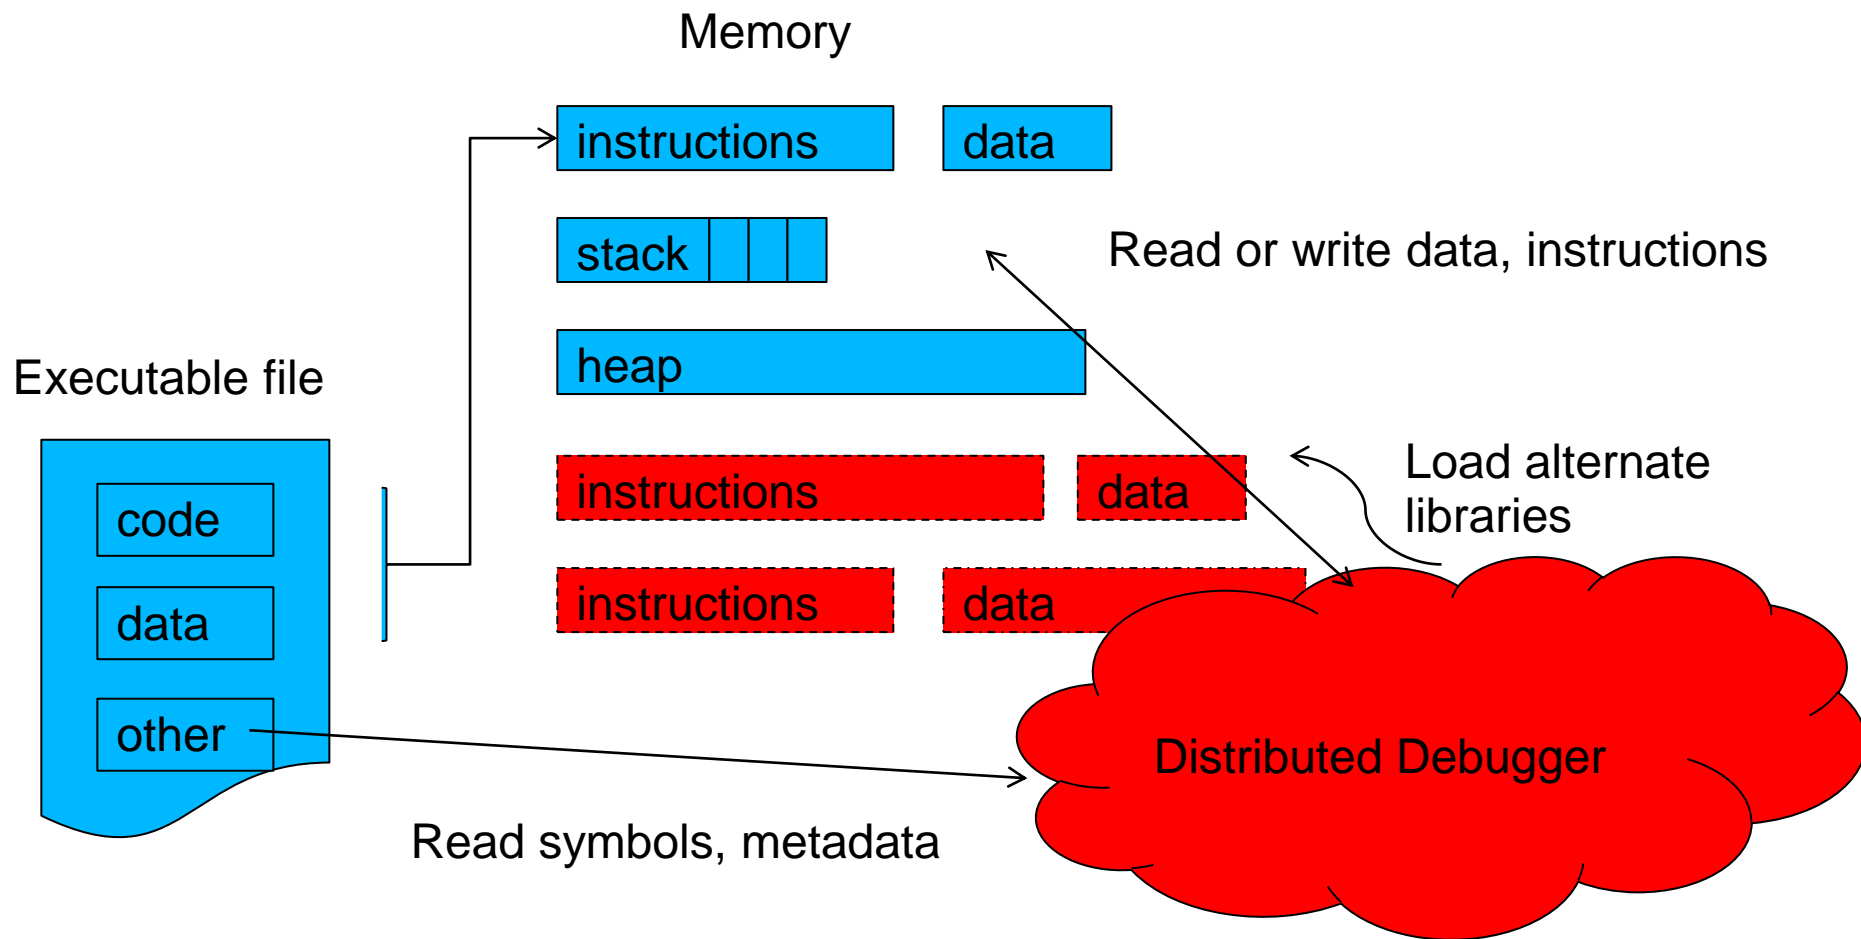


Debugging: symbolic debugging: Optimized code

- Aggressive optimizations (e.g. `-O3`) cause machine instructions to diverge from machine code!
 - Loop unrolling, function inlining, instruction re-ordering, optimizing out variables, etc
- Effects: debugger much less predictable
 - Setting some breakpoints are impossible (instructions optimized out or moved)
 - Variables are optimized out, or appear to change unexpectedly
 - Stepping through code follows arbitrary execution order
- Easiest to debug with NO optimizations (`-O0`)

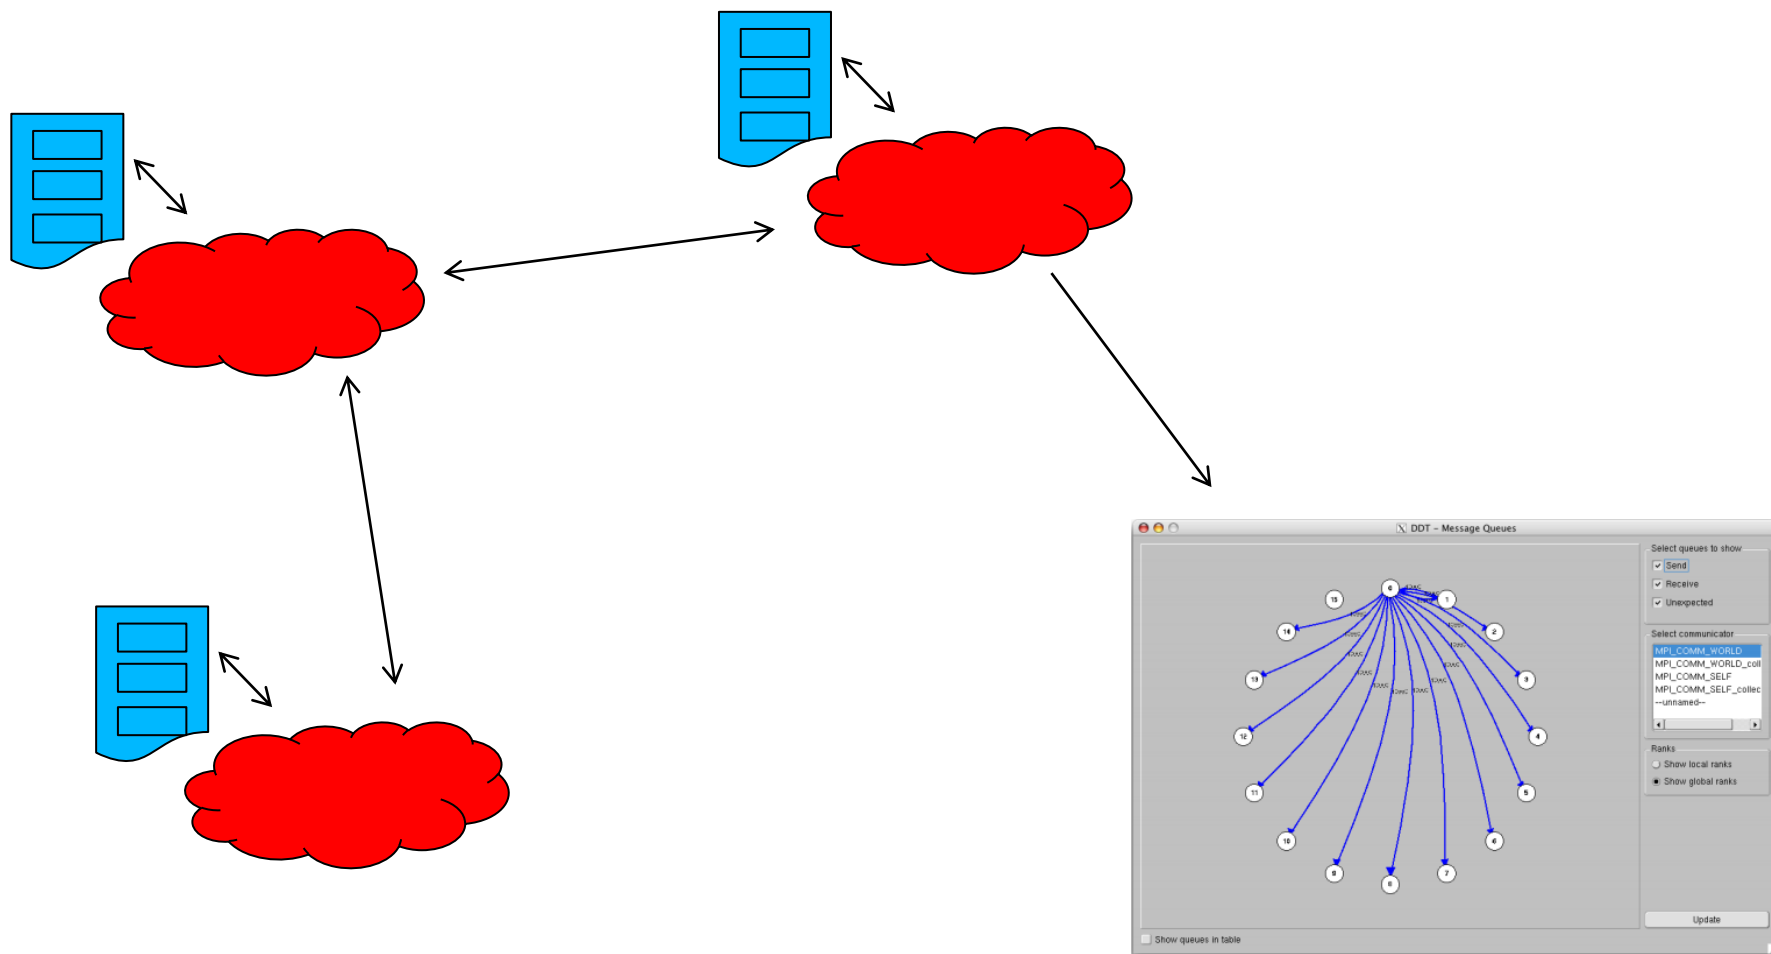


Debugging: symbolic debugging: Distributed





Debugging: symbolic debugging: Distributed





Debugging: symbolic debugging: distributed: DDT

- DDT (Allinea Distributed Debugger Tool)
- Proprietary, GUI-oriented
- Large-scale OpenMP, MPI debugging
 - MPI message tracking
 - View queues and communication patterns for running procs
 - Supports all MPI distributions on Ranger
- Jobs submitted through DDT
 - Remember, it needs to “wrap” and control each task
- Usage: Compile with `-g`, then `module load ddt`, then `ddt <executable>` and go from there.



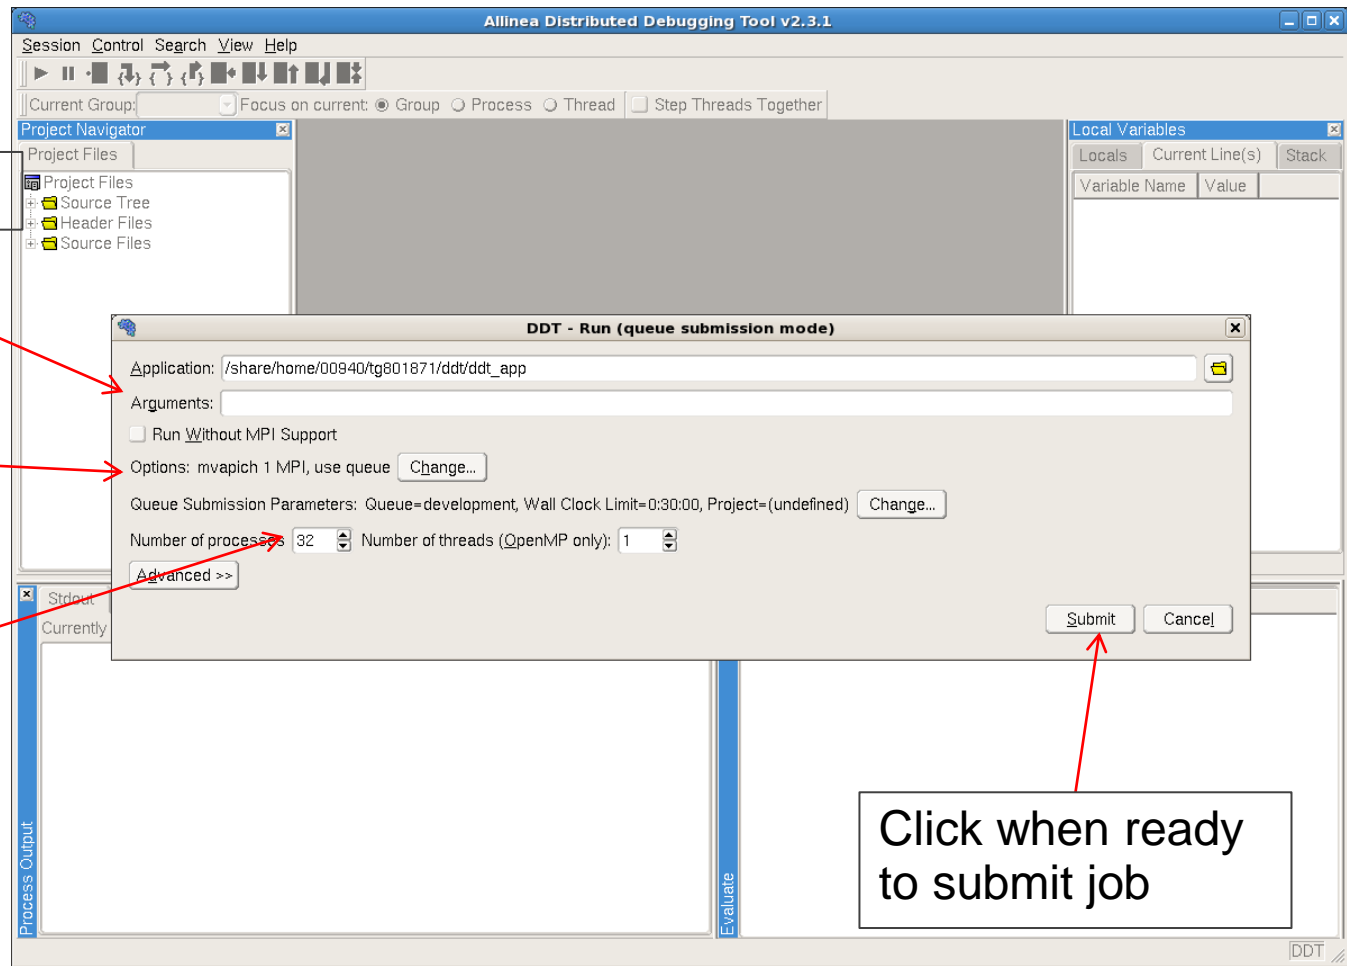
Debugging: symbolic debugging: distributed: DDT

Add any arguments

Ranger default

Sets number of nodes

Click when ready to submit job





Debugging: symbolic debugging: distributed: DDT

The screenshot shows the Allinea Distributed Debugging Tool (DDT) v2.3.1 interface. A dialog box titled "DDT - Job Submitted" is open, displaying the following text:

Your debugging job has been submitted to the queue. DDT will continue automatically once the job has been started.
You may cancel the job by closing DDT or clicking on the button below.

ACTIVE JOBS-----
JOBID JOBNAME USERNAME STATE CORE HOST QUEUE REMAINING STARTTIME

0 active jobs : 0 of 3888 hosts (0.00 %)

WAITING JOBS-----
JOBID JOBNAME USERNAME STATE CORE HOST QUEUE WCLIMIT QUEUETIME

571888 DDTJOB tg801871 Waiting 32 2 development 00:30:00 Tue Mar 3 12:49:51

WAITING JOBS WITH JOB DEPENDENCIES---
JOBID JOBNAME USERNAME STATE CORE HOST QUEUE WCLIMIT QUEUETIME

UNSCHEDULED JOBS-----
JOBID JOBNAME USERNAME STATE CORE HOST QUEUE WCLIMIT QUEUETIME

Total jobs: 1 Active Jobs: 0 Waiting Jobs: 1 Dep/Unsched Jobs: 0

Cancel Job



Debugging: symbolic debugging: distributed: DDT

Processes and groups

The screenshot displays the Allinea Distributed Debugging Tool (DDT) v2.3.1 interface. At the top, a process tree shows 'All' (0-31), 'Root' (0), and 'Workers' (1-31). The main window shows Fortran source code for 'debug_code.f'. A 'Local Variables' panel on the right shows a variable 'i' with value 0. The bottom panel shows 'Stdout', 'Stderr', and 'Stdin' tabs, with 'Stdout' selected. An 'Evaluate' panel is also visible at the bottom right.

Source code and execution

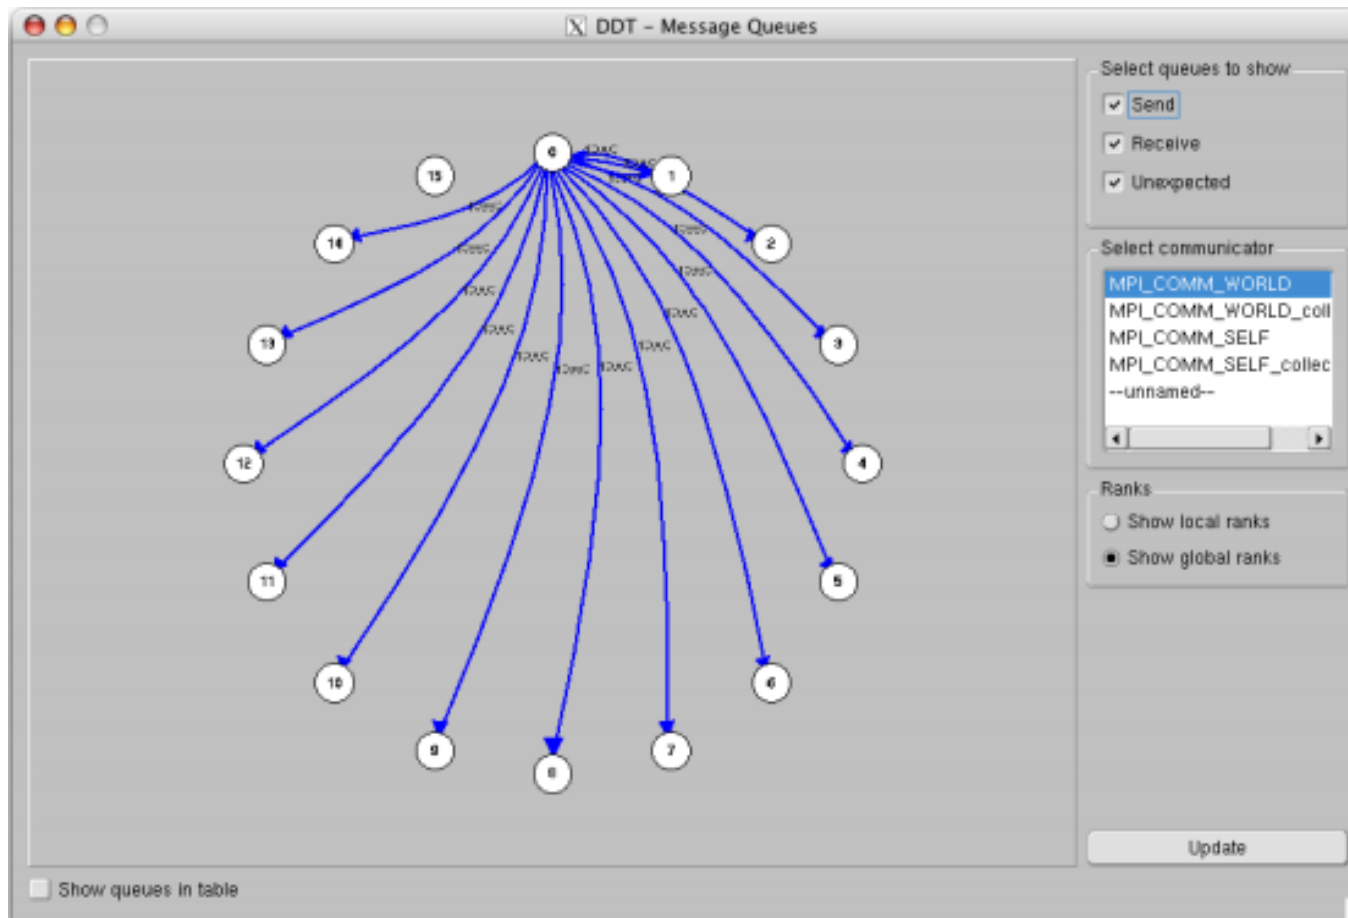
Variables

Stack view and stdout

Evaluation, expressions



Debugging: symbolic debugging: distributed: DDT



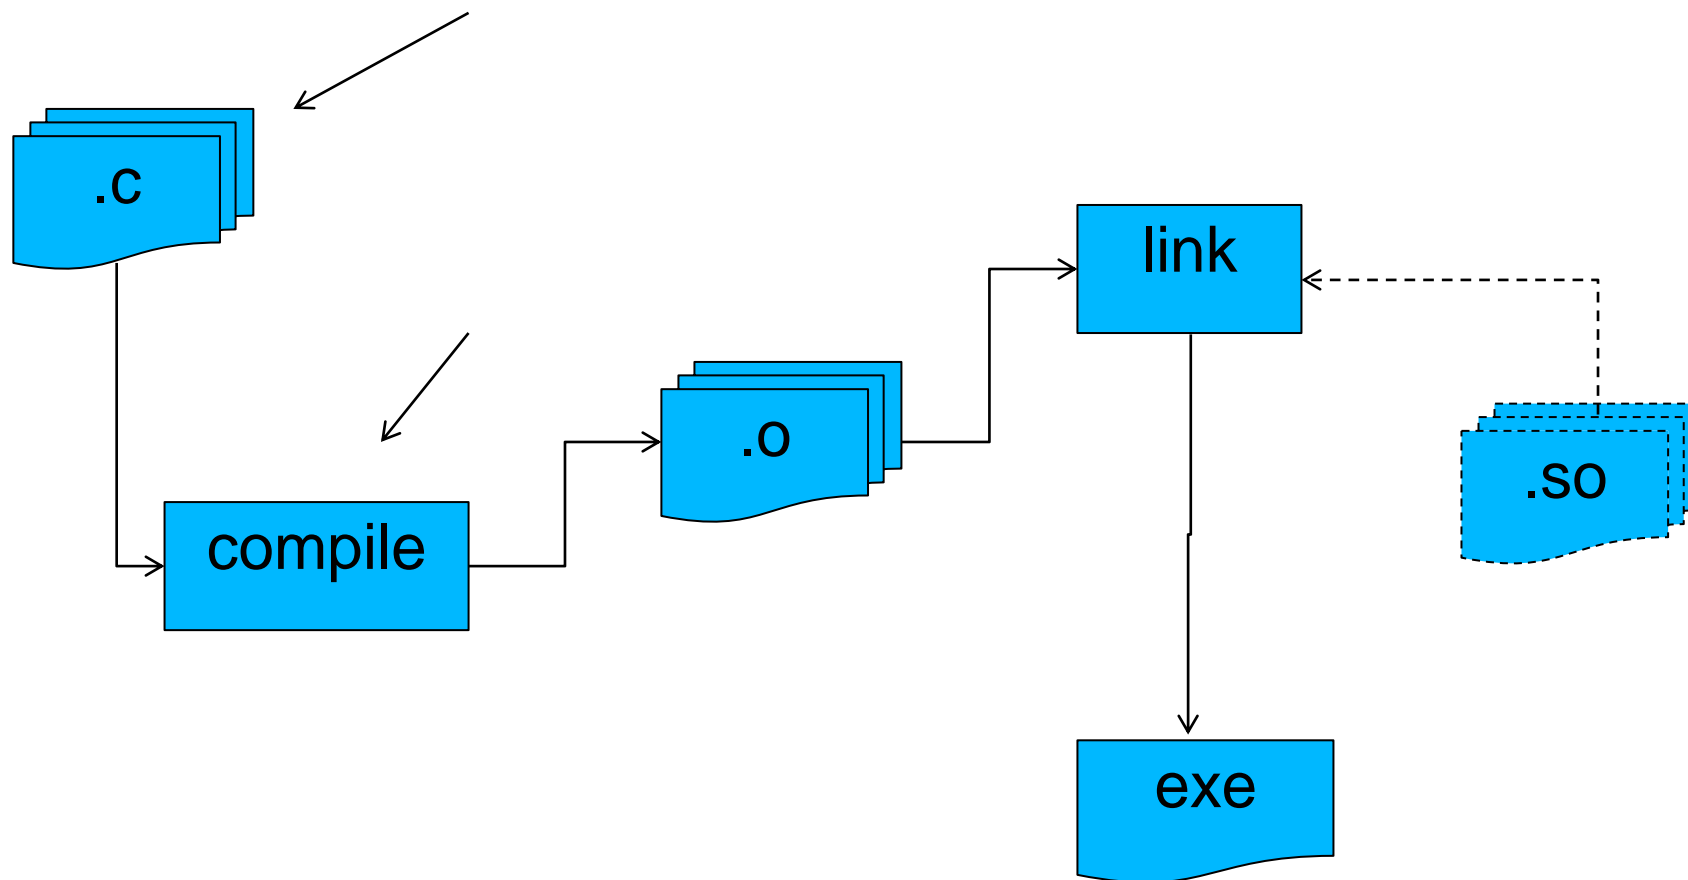


Profiling

- Measure performance characteristics, identify compute-intensive areas (e.g. “hot spots”) that *may* be worth improving
- Can suffer from “observer effect” – collecting performance data significantly degrades performance
- Two main approaches: instrumentation and statistical sampling
 - Instrumentation: add instructions to collect information (function call duration, number of invocations, etc)
 - Sampling: Query state of unmodified executable at regular intervals



Profiling: Instrumentation





Profiling: Instrumentation: printf and timers

- Check system time and printf at appropriate points
 - `SYSTEM_CLOCK` or `clock()` for fortran, C
- Very simple, great for targeting a specific area.
- Problem: printf statements are expensive, especially if there are many
- Problem: Timer precision and accuracy is system/implementation dependent.



Profiling: Instrumentation: GPROF

- GPROF (GNU profiler)
- Compile option `-pg` adds debugging symbols and additional data collection symbols
 - Slows program down, sometimes significantly
- Each time program is run, output file `gmon.out` is created containing profiling data
 - This data is then analyzed by `gprof` in a separate step, e.g. `gprof <executable> gmon.out > profile.txt`

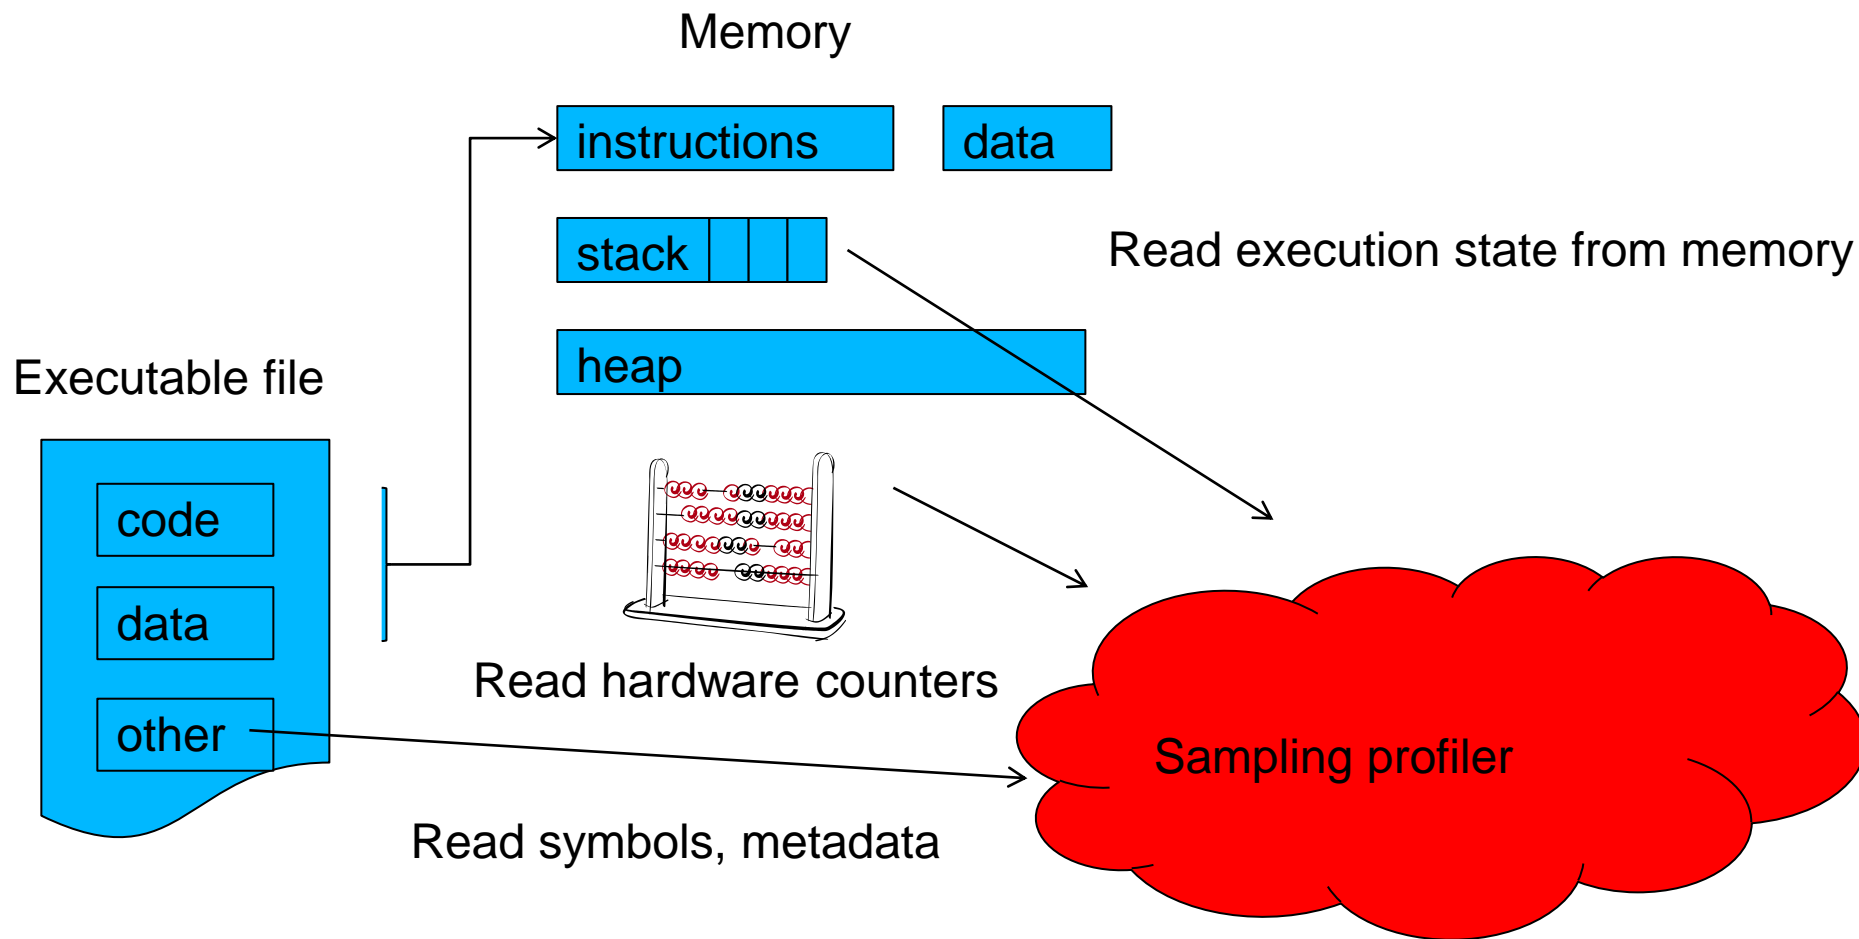


Profiling: Instrumentation: GPROF

- Flat profile
 - Lists each function with associated statistics
 - CPU time spend, number of times called, etc
 - Useful to identify expensive routines
- Call Graph
 - Number of times function was called by another, called others
 - Gives a sense of relationship between functions
- Annotated Source
 - Number of times a line was executed



Profiling: sampling



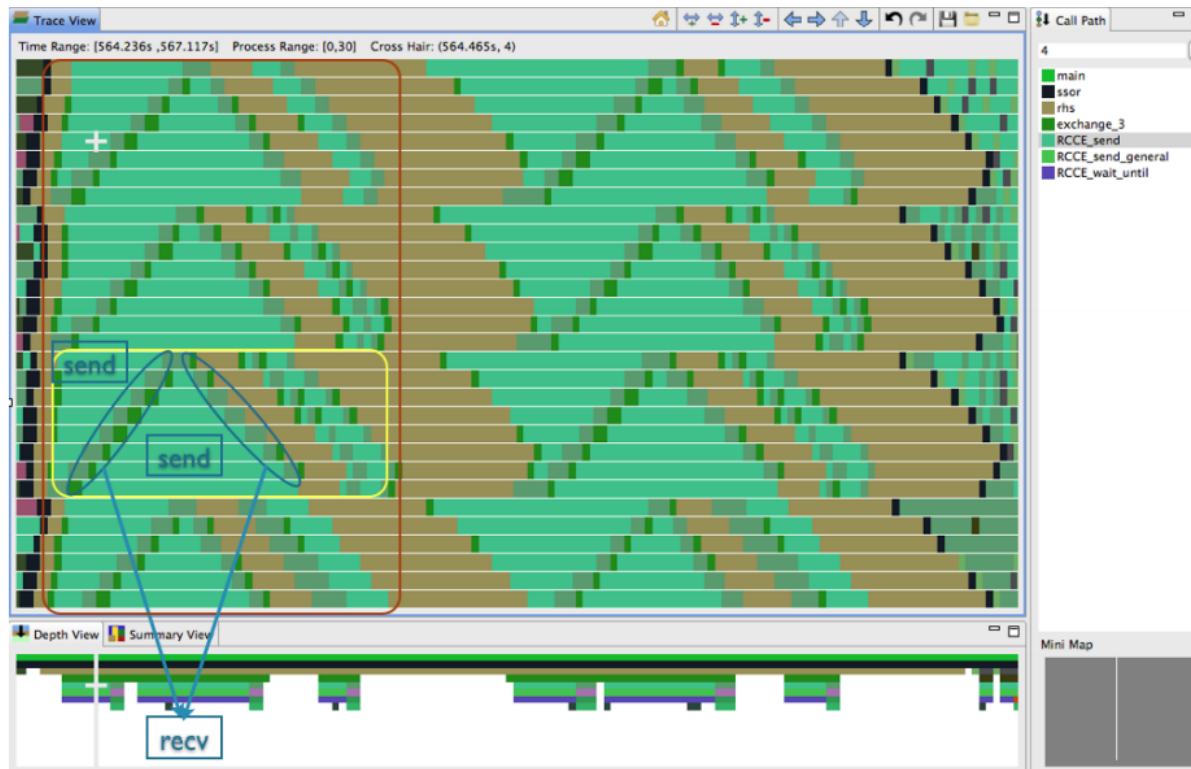


Profiling: sampling: HPCToolkit, PAPI

- PAPI: Provides access to hardware counters
 - API hides gory details of hardware/OS platform
 - Cache accesses, hits, misses
 - FLOPS
 - The kinds of data available depend very much on hardware
- HPCToolkit
 - Asynchronous sampling of running processes
 - Supports OpenMP, MPI, and hybrid
 - Supports running against optimized code
 - <http://hpctoolkit.org>



Profiling: sampling: HPCToolkit



From Xu Liu, John Mellor-Crummey, and Nathan R. Tallen (2012),
Analyzing Application Performance Bottlenecks on Intel's SCC.
Presented at TACC-Intel Highly Parallel Computing Symposium, Austin, TX



Profiling: sampling: PerfExpert

- Developed at TACC
- Easy to use interface over data collected via HPCToolkit and PAPI
- Provides suggestions and “what to fix”
- Runs against fully optimized code with debugging symbols
- <http://www.tacc.utexas.edu/perfexpert>



Profiling: sampling: PerfExpert

Code Section: Loop in function main() at Integrator.c:81 (98.9% of the total runtime)

=====
change the order of loops

```
loop i { loop j {...} } → loop j { loop i {...} }
```

employ loop blocking

```
loop i {loop k {loop j {c[i][j] = c[i][j] + a[i][k] * b[k][j];}}} →  
loop k step s {loop j step s {loop i {  
  for (kk = k; kk < k + s; kk++) {  
    for (jj = j; jj < j + s; jj++) {  
      c[i][jj] = c[i][jj] + a[i][kk] * b[kk][jj];}}}}}
```

apply loop fission so every loop accesses just a couple of different arrays

```
loop i {a[i] = a[i] * b[i] - c[i];} →  
loop i {a[i] = a[i] * b[i];} loop i {a[i] = a[i] - c[i];}
```



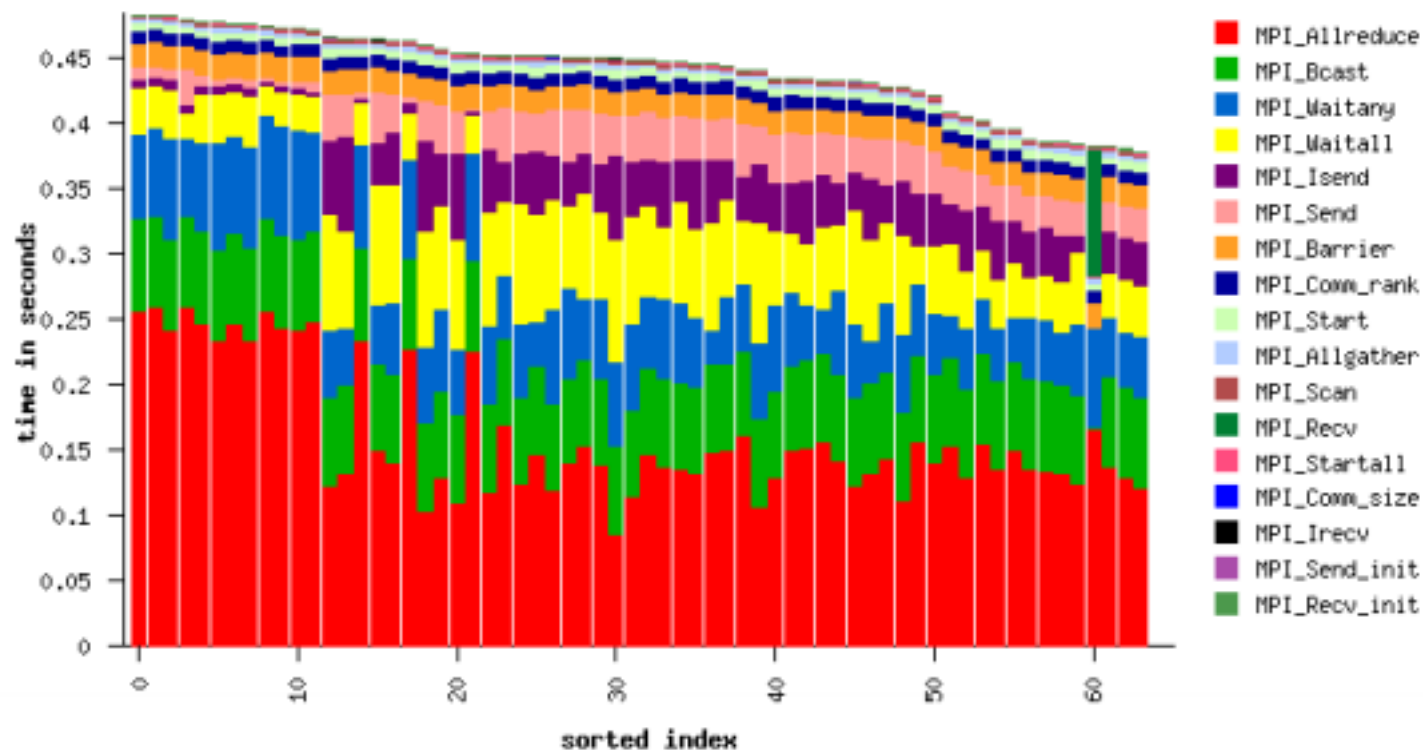

Profiling: IPM

- Integrated Performance Monitoring
- Run against fully optimized code with debugging symbols (-g)
- You need to explicitly pre-load ipm library:
 - `module load ipm`
 - `export LD_PRELOAD=$TACC_IPM_LIB/libipm.so`
 - `export IPM_REPORT=full`
 - `ibrun <my executable> <my arguments>`
- Produces text, html, xml reports of processing and communication statistics
- Very good for quick snapshot of MPI behaviour
- <http://ipm-hpc.sourceforge.net/>



Profiling: IPM

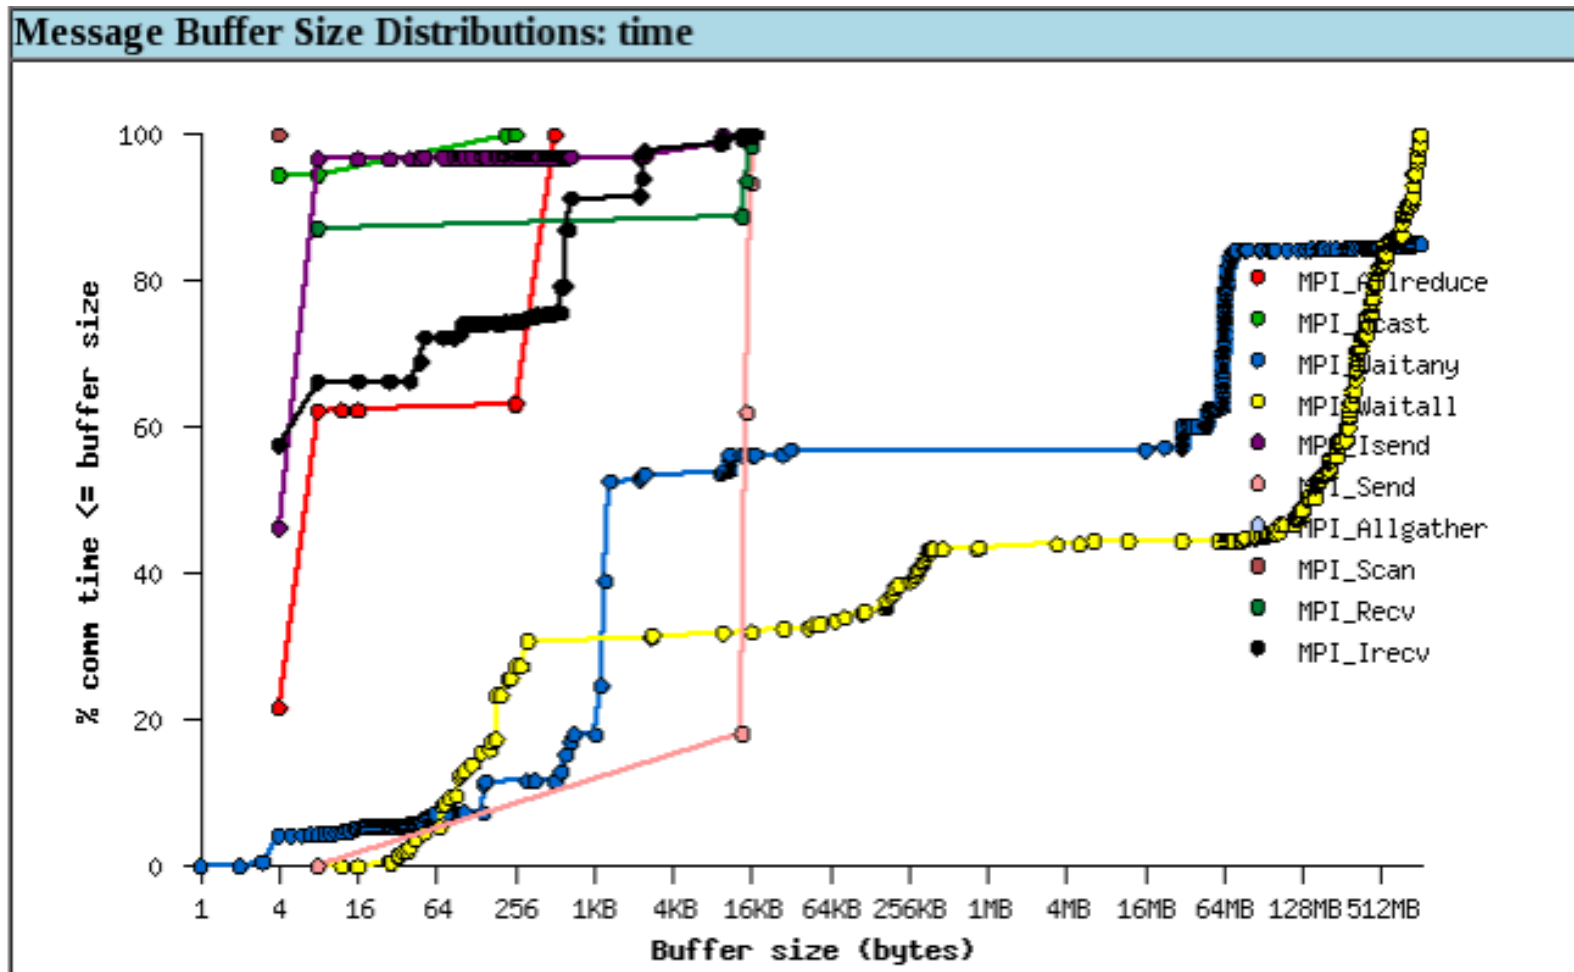
Communication balance by task (sorted by MPI time)



[by MPI rank](#), [time detail by MPI time](#), [time detail by rank](#), [call list](#)

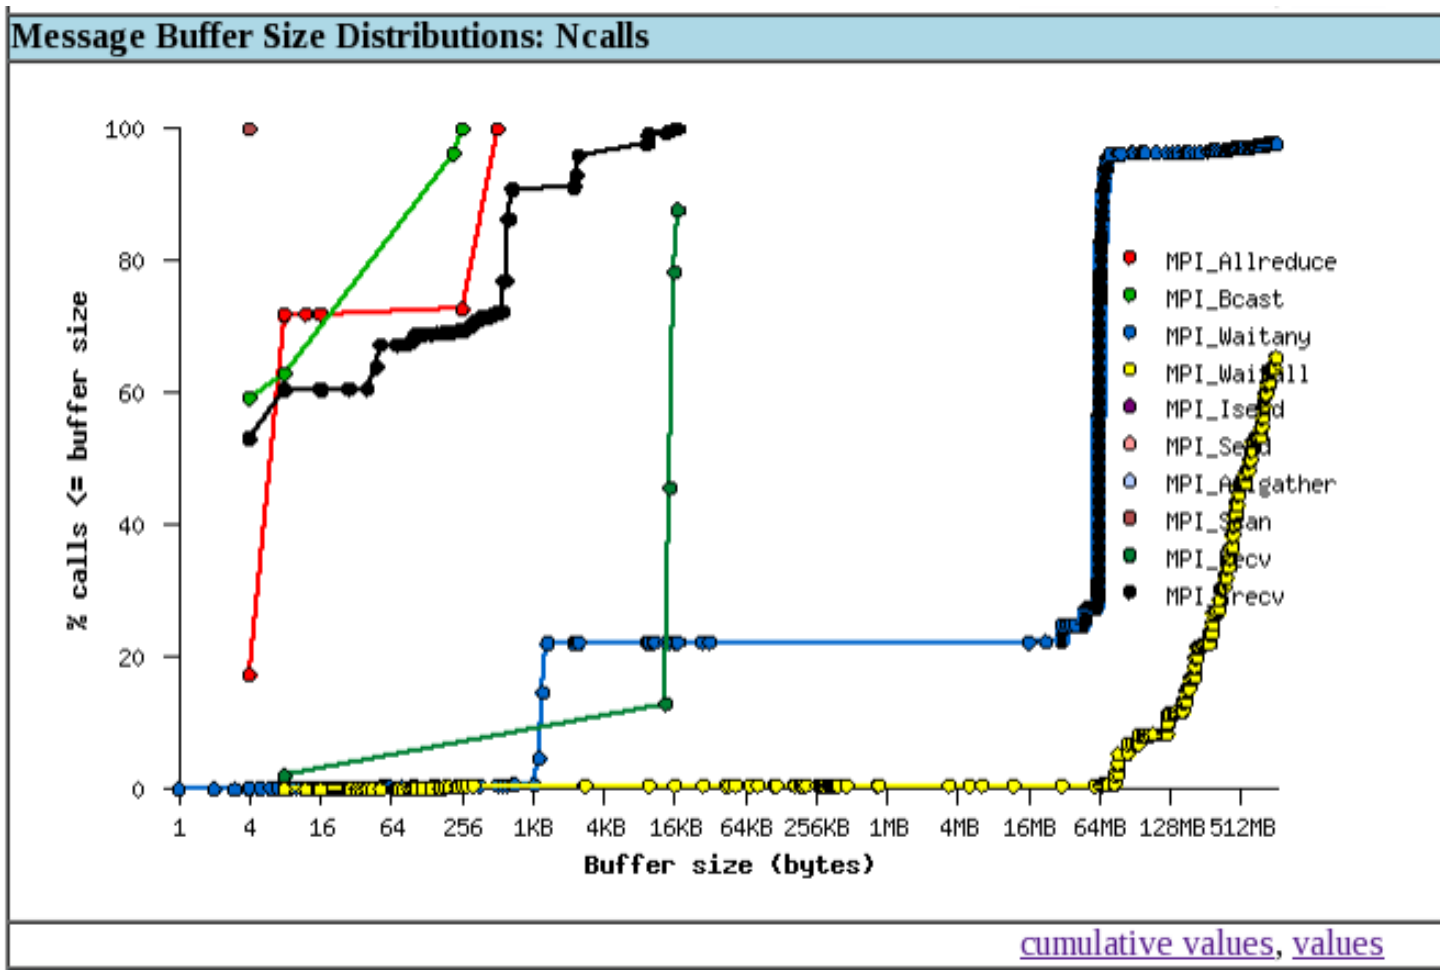


Profiling: IPM





Profiling:IPM





Profiling:IPM

