

STAT

VI-HPS Tuning Workshop

October 10, 2013 | Alexandre Otto Strube

Outline

- Overview
- How?
- Working

Overview

- **Stack Trace Analysis Tool**

Overview

- **Stack Trace Analysis Tool**
- Shows merged call tree of whole program

Overview

- **Stack Trace Analysis Tool**
- Shows merged call tree of whole program
- Useful to detect deadlocks

Overview

- **Stack Trace Analysis Tool**
- Shows merged call tree of whole program
- Useful to detect deadlocks
- Very lightweight helper tool

Overview

- **Stack Trace Analysis Tool**
- Shows merged call tree of whole program
- Useful to detect deadlocks
- Very lightweight helper tool
- **Scales to millions of processes**

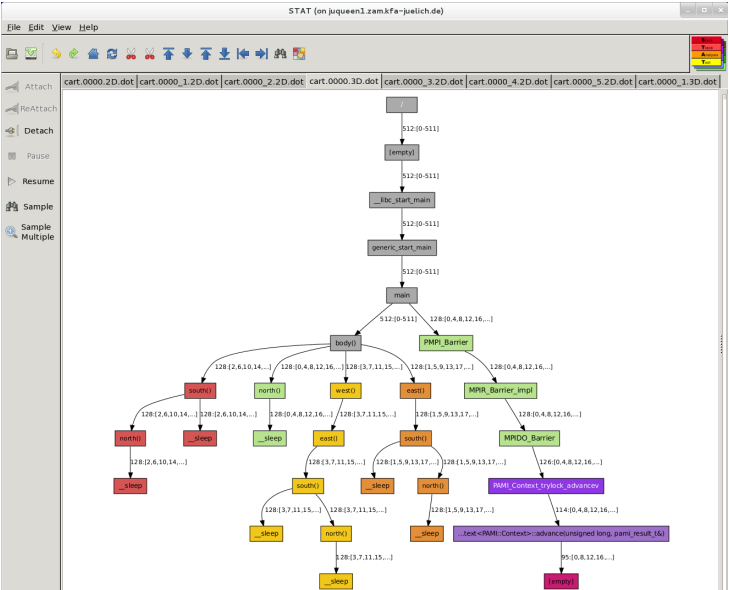
http://www.hpcwire.com/hpcwire/2012-12-03/bug_repellent_for_supercomputers_proves_effective.html

Overview

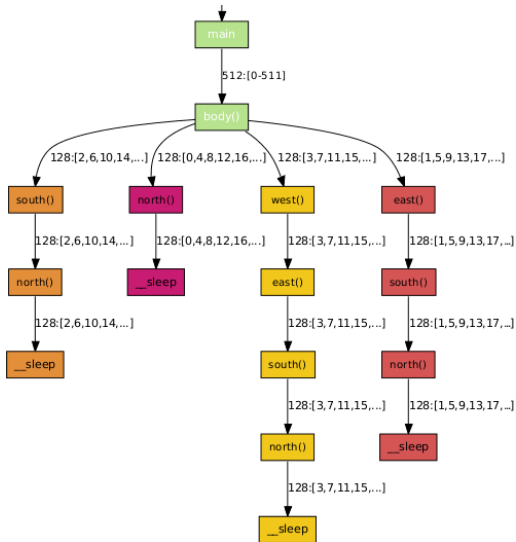
- **Stack Trace Analysis Tool**
- Shows merged call tree of whole program
- Useful to detect deadlocks
- Very lightweight helper tool
- Scales to millions of processes
http://www.hpcwire.com/hpcwire/2012-12-03/bug_repellent_for_supercomputers_proves_effective.html
- Pinpoint individual problems

Overview

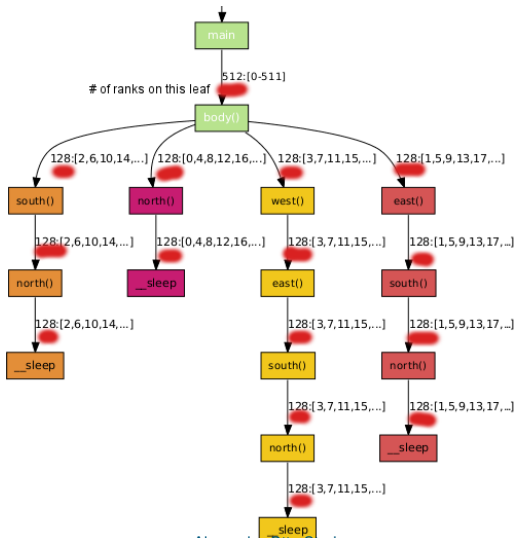
- **Stack Trace Analysis Tool**
- Shows merged call tree of whole program
- Useful to detect deadlocks
- Very lightweight helper tool
- Scales to millions of processes
http://www.hpcwire.com/hpcwire/2012-12-03/bug_repellent_for_supercomputers_proves_effective.html
- Pinpoint individual problems
- NOT a debugger



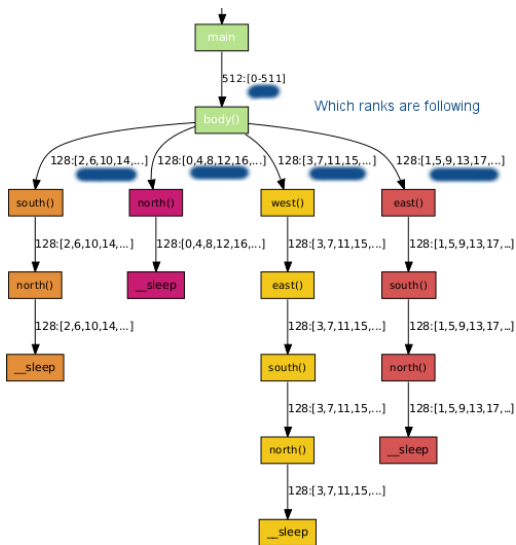
Zoom



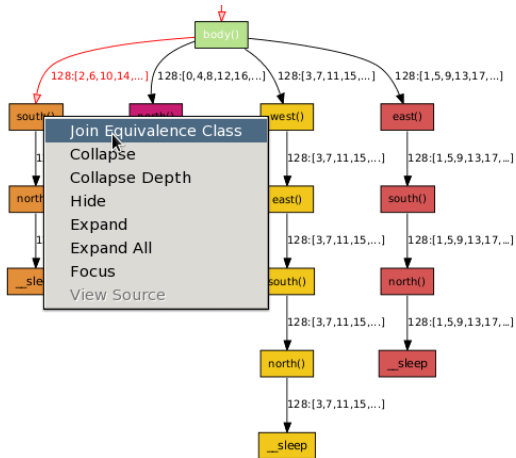
Zoom




Zoom



Equivalence



Equivalence



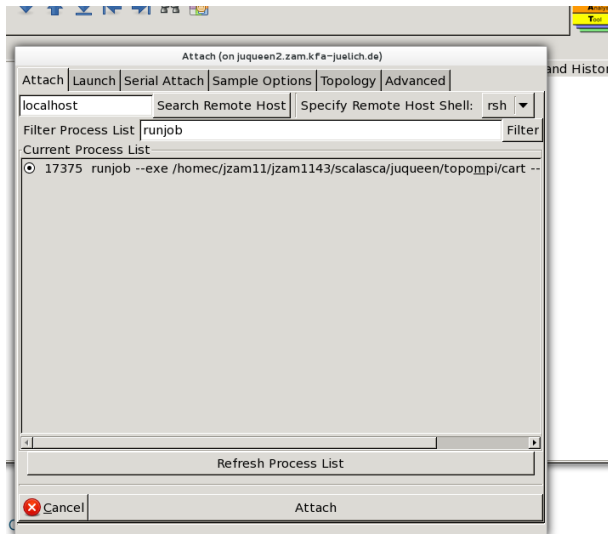
Launching STAT

```
$ module load UNITE stat  
$ STATGUI
```


Caveat

If launching from juqueen1, add this to the submission script
@ requirements = (Machine == "juqueen1")

Attaching to a process



Collecting multiple samples



Collecting multiple samples

Stack Sample Preferences (on juqueen2.zam.kfa-juelich.de)

Per Sample Options

With Threads

Gather Python Traces

Stack Frame (node) Sample Options

function only

function and pc

function and line

Process Set (edge) Sample Options

full list

count and representative

Run Time Before Sample (sec)

▶ Advanced

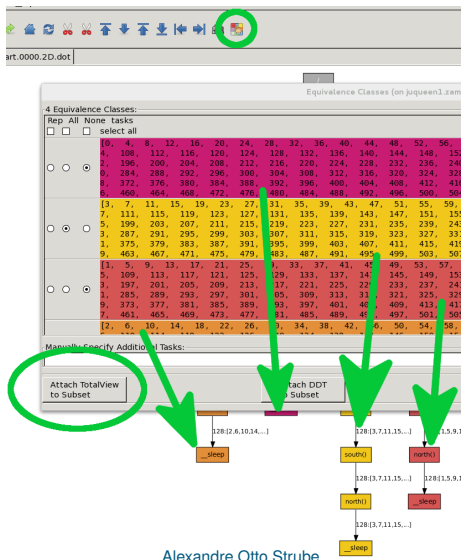
Multiple Sample Options

Num Traces Trace Frequency (ms)

Gather Individual Samples

Clear On Sample

Attach to equivalent



The screenshot shows a software interface with a table of equivalence classes and a task dependency graph below it.

Equivalence Classes Table:

Rep	All	None	tasks
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	select all
0	4	8	12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 56, 60
4	108	112	116, 120, 124, 128, 132, 136, 140, 144, 148, 152
2	196	200	204, 208, 212, 216, 220, 224, 228, 232, 236, 240
0	284	288	292, 296, 300, 304, 308, 312, 316, 320, 324, 328
8	372	376	380, 384, 388, 392, 396, 400, 404, 408, 412, 416
6	460	464	468, 472, 476, 480, 484, 488, 492, 496, 500, 504
3	7	11	15, 19, 23, 27, 31, 35, 39, 43, 47, 51, 55, 59
7	111	115	119, 123, 127, 131, 135, 139, 143, 147, 151, 155
5	199	203	207, 211, 215, 219, 223, 227, 231, 235, 239, 243
3	287	291	295, 299, 303, 307, 311, 315, 319, 323, 327, 331
1	375	379	383, 387, 391, 395, 399, 403, 407, 411, 415, 419
9	463	467	471, 475, 479, 483, 487, 491, 495, 499, 503, 507
3	5	9	13, 17, 21, 25, 29, 33, 37, 41, 45, 49, 53, 57
5	109	113	117, 121, 125, 129, 133, 137, 141, 145, 149, 153
3	197	201	205, 209, 213, 217, 221, 225, 229, 233, 237, 241
1	285	289	293, 297, 301, 305, 309, 313, 317, 321, 325, 329
9	373	377	381, 385, 389, 393, 397, 401, 405, 409, 413, 417
7	461	465	469, 473, 477, 481, 485, 489, 493, 497, 501, 505
2	6	10	14, 18, 22, 26, 30, 34, 38, 42, 46, 50, 54, 58

Manually Specify Additional Tasks:

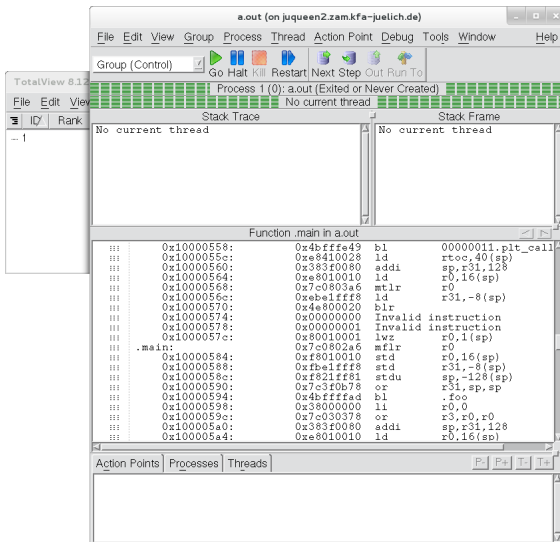
- Attach TotalView to Subset
- Attach DDT to Subset

Task Dependency Graph:

```

graph TD
    A["128{2,6,10,14,...}"] --> B[sleep]
    C["128{3,7,11,15,...}"] --> D[south]
    C --> E[north]
    D --> F["128{3,7,11,15,...}"]
    E --> G["128{15,9,1}"]
    F --> H[north]
    G --> I[sleep]
    H --> J["128{3,7,11,15,...}"]
    J --> K[sleep]
  
```

Totalview



The screenshot shows the TotalView 8.1.2 interface with the assembly view of the `.main` function in `a.out`. The assembly code is as follows:

```

Function .main in a.out
0x10000558: 0x4bfff649 bl 00000011.plt_call
0x1000055e: 0xe8410028 ld rtoc, 40(sp)
0x10000560: 0x383f0080 addi sp, r31, 128
0x10000564: 0xe8010010 ld r0, 16(sp)
0x10000568: 0x7c0803a6 mtlr r0
0x1000056c: 0xebef1ff8 ld r31, -8(sp)
0x10000570: 0x4e800020 blr
0x10000574: 0x00000000 Invalid instruction
0x10000578: 0x90000001 Invalid instruction
0x1000057c: 0x80010001 lwr r0, 1(sp)
0x10000584: 0x7c0802a6 mflr r0
0x10000588: 0xf8010010 std r0, 16(sp)
0x1000058c: 0xfbe1ff8 std r31, -8(sp)
0x10000590: 0xf821ff81 stdu sp, -128(sp)
0x10000594: 0xc3f0b78 or r31, sp, sp
0x10000598: 0x4bfff649 bl .foo
0x1000059c: 0x38000000 li r0, 0
0x100005a0: 0x7c030378 or r3, r0, r0
0x100005a4: 0x383f0080 addi sp, r31, 128
0x100005a8: 0xe8010010 ld r0, 16(sp)

```