

Code Profiling Tools

Shirley Moore
Innovative Computing Laboratory
University of Tennessee
shirley@cs.utk.edu

9 April 2003

ICL/UTK

1

Definitions – Profiling

- **Profiling**
 - Recording of summary information during execution
 - inclusive, exclusive time, # calls, hardware statistics, ...
 - Reflects performance behavior of program entities
 - functions, loops, basic blocks
 - user-defined “semantic” entities
 - Very good for low-cost performance assessment
 - Helps to expose performance bottlenecks and hotspots
 - Implemented through
 - **sampling**: periodic OS interrupts or hardware counter traps
 - **instrumentation**: direct insertion of measurement code

9 April 2003

ICL/UTK

2

Definitions – Tracing

- **Tracing**
 - Recording of information about significant points (**events**) during program execution
 - entering/exiting code region (function, loop, block, ...)
 - thread/process interactions (e.g., send/receive message)
 - Save information in **event record**
 - timestamp
 - CPU identifier, thread identifier
 - Event type and event-specific information
 - **Event trace** is a time-sequenced stream of event records
 - Can be used to reconstruct dynamic program behavior
 - Typically requires code instrumentation

9 April 2003

ICL/UTK

3

Available Profiling Tools

- prof, gprof
- PAPI (profiling based on timers or on any PAPI hardware counter metric)
- Dynaprof (requires dyninst or DPCL)
- GuideView (OpenMP) (being phased out)
- Vampir (MPI)
- TAU (OpenMP, MPI, MPI/OpenMP)
- SvPablo
- HPCView
- Vprof
- Vendor specific tools (e.g., SGI IRIX perfx and srrun, IBM tprof and trace, Cray PAT)

9 April 2003

ICL/UTK

4

prof, gprof

- Available on many Unix platforms (e.g., IBM AIX, Sun Solaris, HP/Compaq Tru64)
- Produces “flat” profile (prof) or “call graph” profile (gprof)
- Compile with special flag or set environment variable or rewrite executable to enable instrumentation and produce profile data file
- Collect profile data by periodically sampling the program counter during execution and/or calling monitoring routines
- See
 - man prof
 - man gprof

9 April 2003

ICL/UTK

5

prof Profile of FSPX Benchmark

%time	cumulative	self	calls	ms/call	tot/call	name
21.71	18.93	18.93	6080	3.11	3.11	flux_
19.99	36.36	17.43	9124	1.91	3.91	proflux_
8.26	43.56	7.20	6080	1.18	1.18	pde_
8.11	50.63	7.07	6080	1.16	4.17	phase_
7.96	57.57	6.94	100061386	0.00	0.00	cpsintg_
7.46	64.08	6.51	100061388	0.00	0.00	cpsintg_
6.05	69.36	5.28	49807360	0.00	0.00	tsofx_
5.60	74.24	4.88	49807362	0.00	0.00	tlofx_
4.07	77.79	3.55	62202877	0.00	0.00	cpl_
2.44	79.92	2.13	37371906	0.00	0.00	cps_
1.67	81.38	1.46	37371904	0.00	0.00	hl_
1.43	82.63	1.25	37371904	0.00	0.00	hs_
1.07	83.56	0.93	24903681	0.00	0.00	elqds_
0.89	84.34	0.78	37371904	0.00	0.00	aks_

9 April 2003

ICL/UTK

6

IBM Profiling Utilities

- In addition to man pages, see the AIX 5 Performance Management Guide
- prof
- gprof
- tprof

9 April 2003

ICL/UTK

7

SGI Profiling Utilities

- sstrun
 - Collects Speedshop and Workshop performance data
 - Types of experiments: totaltime, usertime, pcsamp, ideal, hardware counter
- prof
 - Analyzes and displays SpeedShop performance data generated by sstrun
- perfex
 - Runs program and collects hardware counter data

9 April 2003

ICL/UTK

8

HP/Compaq Profiling Utilities

- prof
 - Displays profiling data from -p or pixie
- pixie
 - Instruction-counting profiler
- gprof
 - Displays profiling data from -pg or hiprof
- hiprof
 - Creates instrumented version of program for call-graph profiling
- uprofile
 - Profiles a program with Alpha on-chip performance counters

9 April 2003

ICL/UTK

9

Overview of PAPI



- **Performance Application Programming Interface**
- The purpose of the PAPI project is to design, standardize and implement a portable and efficient API to access the hardware performance monitor counters found on most modern microprocessors.
- Parallel Tools Consortium project
- Being installed and supported at the DoD HPC Centers as part of PET CE004

9 April 2003

ICL/UTK

10

PAPI Counter Interfaces



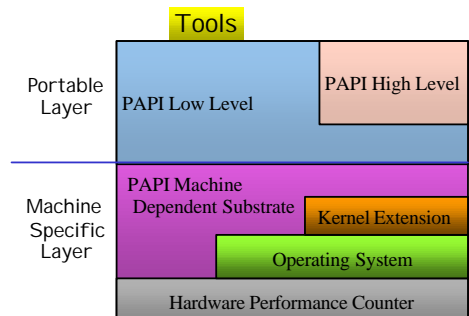
- PAPI provides three interfaces to the underlying counter hardware:
 1. The low level interface manages hardware events in user defined groups called *EventSets*.
 2. The high level interface simply provides the ability to start, stop and read the counters for a specified list of events.
 3. Graphical tools to visualize information.

9 April 2003

ICL/UTK

11

PAPI Implementation



9 April 2003

ICL/UTK

12

PAPI Preset Events

- Proposed standard set of events deemed most relevant for application performance tuning
- Defined in `papiStdEventDefs.h`
- Mapped to native events on a given platform
 - Run `tests/avail` to see list of PAPI preset events available on a platform

9 April 2003

ICL/UTK

13

High-level Interface

- Meant for application programmers wanting coarse-grained measurements
- Not thread safe
- Calls the lower level API
- Allows only PAPI preset events
- Easier to use and less setup (additional code) than low-level

9 April 2003

ICL/UTK

14

High-level API

- | | |
|----------------------------------|-----------------------------------|
| • C interface | • Fortran interface |
| <code>PAPI_start_counters</code> | <code>PAPIF_start_counters</code> |
| <code>PAPI_read_counters</code> | <code>PAPIF_read_counters</code> |
| <code>PAPI_stop_counters</code> | <code>PAPIF_stop_counters</code> |
| <code>PAPI_accum_counters</code> | <code>PAPIF_accum_counters</code> |
| <code>PAPI_num_counters</code> | <code>PAPIF_num_counters</code> |
| <code>PAPI_flops</code> | <code>PAPIF_flops</code> |

9 April 2003

ICL/UTK

15

PAPI_flops

- `int PAPI_flops(float *real_time, float *proc_time, long_long *flpins, float *mflops)`
 - Only two calls needed, `PAPI_flops` before and after the code you want to monitor
 - `real_time` is the wall-clocktime between the two calls
 - `proc_time` is the “virtual” time or time the process was actually executing between the two calls (not as fine grained as `real_time` but better for longer measurements)
 - `flpins` is the total floating point instructions executed between the two calls
 - `mflops` is the Mflop/s rating between the two calls

9 April 2003

ICL/UTK

16

Low-level Interface

- Increased efficiency and functionality over the high level PAPI interface
- About 40 functions
- Obtain information about the executable and the hardware
- Thread-safe
- Fully programmable
- Callbacks on counter overflow

9 April 2003

ICL/UTK

17

Callbacks on Counter Overflow

- PAPI provides the ability to call user-defined handlers when a specified event exceeds a specified threshold.
- For systems that do not support counter overflow at the OS level, PAPI sets up a high resolution interval timer and installs a timer interrupt handler.

9 April 2003

ICL/UTK

18

Statistical Profiling

- PAPI provides support for execution profiling based on any counter event.
- PAPI_profil() creates a histogram of overflow counts for a specified region of the application code.

9 April 2003

ICL/UTK

19

PAPI_profil

```
int PAPI_profil(unsigned short *buf, unsigned int
bufsiz, unsigned long offset, unsigned scale, int
EventSet, int EventCode, int threshold, int flags)
```

- buf – buffer of bufsiz bytes in which the histogram counts are stored
- offset – start address of the region to be profiled
- scale – contraction factor that indicates how much smaller the histogram buffer is than the region to be profiled

9 April 2003

ICL/UTK

20

What is DynaProf?

- A portable tool to instrument a running executable with *Probes* that monitor application performance.
- Simple command line interface.
- Open Source Software
- A work in progress...

No source code required!

9 April 2003

ICL/UTK

21

DynaProf Methodology

- Make collection of run-time performance data easy by:
 - Avoiding instrumentation and recompilation
 - Using the same tool with different probes
 - Providing useful and meaningful probe data
 - Providing different kinds of probes
 - Allowing custom probes

No source code required!

9 April 2003

ICL/UTK

22

Why the “Dyna”?

- Instrumentation is selectively inserted directly into the program’s address space.
- Why is this a better way?
 - No perturbation of compiler optimizations
 - Complete language independence
 - Multiple Insert/Remove instrumentation cycles

9 April 2003

ICL/UTK

23

DynaProf Design

- GUI, command line & script driven user interface
- Uses GNU readline for command line editing and command completion.
- Instrumentation is done using:
 - Dyninst on Linux, Solaris and IRIX
 - DPCL on AIX

9 April 2003

ICL/UTK

24

DynaProf Commands

```
load <executable>
list [module pattern]
use <probe> [probe args]
instr module <module> [probe args]
instr function <module> <function> [probe
args]
stop
continue
run [args]
Info
unload
```

9 April 2003

ICL/UTK

25

Dynaprof Probes

- papiprobe
- wallclockprobe
- perfometerprobe

9 April 2003

ICL/UTK

26

DynaProf Probe Design

- Can be written in any compiled language
- Probes export 3 functions with a standardized interface.
- Easy to roll your own (<1day)
- Supports separate probes for MPI/OpenMP/Threads

9 April 2003

ICL/UTK

27

Future development

- GUI development
- Additional probes
 - Perfex probe
 - Vprof probe
 - TAU probe
- Better support for parallel applications

9 April 2003

ICL/UTK

28

Exclusive Profile of Metric PAPI_TOP_INS.

Name	Percent	Total	Calls
TOTAL	100	3.031e+11	1
unknown	93.81	1.031e+11	1
profus_	27.75	8.418e+10	9324
phase_	15.44	4.48e+10	6080
flux_	2.907	7.398e+09	6080
pda_	0.4884	1.48e+09	6080

Inclusive Profile of Metric PAPI_TOP_INS.

Name	Percent	Total	SubCalls
TOTAL	100	3.031e+11	0
profus_	28.31	1.399e+11	2.242e+08
phase_	37.49	1.142e+11	1.247e+08
flux_	2.907	7.398e+09	0
pda_	0.4884	1.48e+09	0

1-Level Inclusive Call Tree of Metric PAPI_TOP_INS.

Parent/Child	Percent	Total	Calls
TOTAL	100	3.031e+11	1
profus_	100	1.399e+11	9324
- akl_	8.304	1.159e+10	3.737e+07
- ake_	8.4	1.161e+10	3.737e+07
- opl_	8.323	1.152e+10	3.737e+07
- ope_	8.325	1.152e+10	3.737e+07
- hl_	9.489	1.324e+10	3.737e+07
- ha_	9.564	1.313e+10	3.737e+07
flux_	100	7.398e+09	6080
pda_	100	1.48e+09	6080
phase_	100	1.142e+11	6080
- tsofa_	11.72	1.339e+10	2.49e+07
- tlofa_	11.49	1.319e+10	2.49e+07
- eslda_	12.88	1.473e+10	2.49e+07
- elofa_	12.49	1.449e+10	2.49e+07
- timo1_	4.999e-07	973	1
- timom0_	1.144	1.179e+08	7.271e+04
- ssofa_	0.101	1.103e+08	7.271e+04
- xlofa_	0.103	1.129e+08	7.271e+04
- opl_	8.913	1.038e+10	4.48e+07

fspx
Cycles
&
Instrs.

Exclusive Profile of Metric PAPI_TOP_CTC.

Name	Percent	Total	Calls
TOTAL	100	5.017e+11	1
unknown	93.42	2.189e+11	1
profus_	27.75	1.399e+11	9324
phase_	14.3	7.179e+10	6080
flux_	0.6256	3.189e+09	6080

Inclusive Profile of Metric PAPI_TOP_CTC.

Name	Percent	Total	SubCalls
TOTAL	100	5.017e+11	0
profus_	27.42	2.078e+11	2.242e+08
phase_	38.92	1.959e+11	1.247e+08
flux_	0.6256	3.189e+09	0
pda_	0.6256	3.189e+09	0

1-Level Inclusive Call Tree of Metric PAPI_TOP_CTC.

Parent/Child	Percent	Total	Calls
TOTAL	100	5.017e+11	1
profus_	100	2.078e+11	9324
- akl_	7.945	2.281e+10	3.737e+07
- ake_	7.871	2.248e+10	3.737e+07
- opl_	8.441	2.346e+10	3.737e+07
- ope_	8.705	2.503e+10	3.737e+07
- hl_	9.292	2.646e+10	3.737e+07
- ha_	8.947	2.579e+10	3.737e+07
flux_	100	3.189e+09	6080
pda_	100	3.189e+09	6080
phase_	100	1.959e+11	6080
- tsofa_	13.42	2.425e+10	2.49e+07
- tlofa_	12.42	2.029e+10	2.49e+07
- eslda_	13.41	2.618e+10	2.49e+07
- elofa_	13.41	2.618e+10	2.49e+07
- timo1_	1.023e-06	1978	1
- timom0_	1.176	3.355e+08	1.271e+04
- ssofa_	0.1349	4.412e+08	7.271e+04
- xlofa_	0.133	4.294e+08	7.271e+04
- opl_	8.032	1.569e+10	2.483e+07

Perfometer

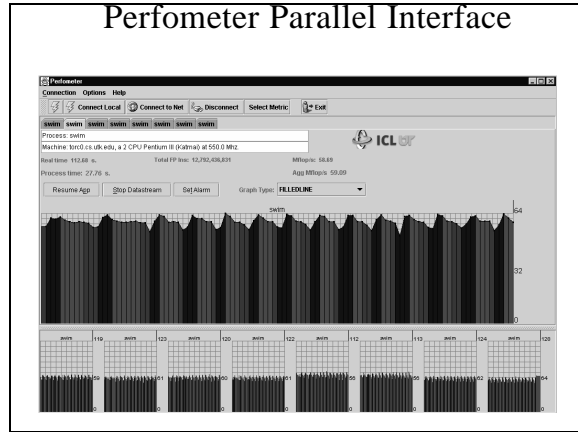

- Application is instrumented with PAPI
 - call perfometer()
 - call mark_perfometer(int color, char *label)
- Application is started. At the call to **perfometer**, signal handler and a timer are set up to collect and send the information to a Java applet containing the graphical view.
- Sections of code that are of interest can be designated with specific colors
- Real-time display or trace file



9 April 2003

ICL/UTK

30



KAP/Pro Toolset

- <http://www.kai.com/>
- Set of tools for developing parallel scientific software using OpenMP
- Components
 - Guide OpenMP compiler
 - Assure OpenMP debugger
 - GuideView performance analyzer
 - Utilities for translating older directives to OpenMP
- Licensed at ERDC MSRC and ARL MSRC
- Being phased out by Intel/KSL

9 April 2003 ICL/UTK 33

GuideView

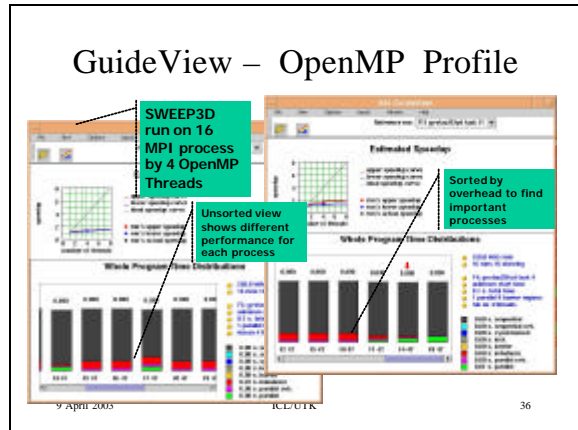
- Intuitive, color-coded display of parallel performance bottlenecks
- Regions of code are identified where improving local performance will have the greatest impact on overall performance.
- Configuration file Gvproperties.txt controls GUI features (fonts, colors, window sizes and locations, etc.)
- Online help system under Help menu
- Use `kmp_set_parallel_name` to name parallel regions so that name will be displayed by GuideView

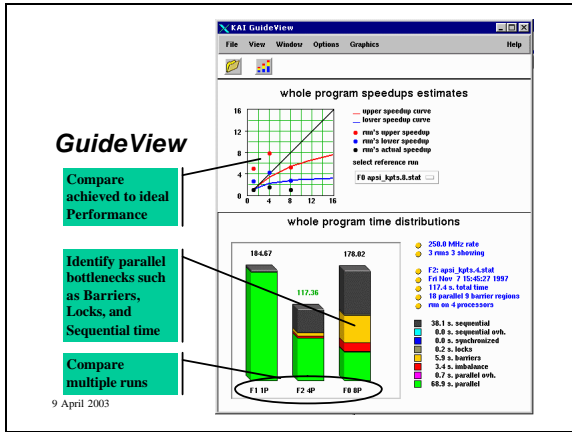
9 April 2003 ICL/UTK 34

Guide Instrumentation Options

- WGnoopenmp
 - Enable profiling but no OpenMP
- WGprof
 - Activates profiling for Vampir and GuideView; implies -WGstats
- WGprof_leafprune=<integer>
 - Sets the minimum size of procedures to retain in Vampir or GuideView profiles to <integer> lines
 - Use to reduce instrumentation overhead and tracefile size

9 April 2003 ICL/UTK 35

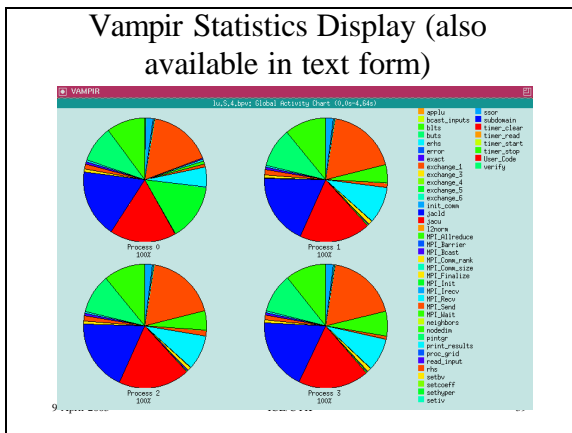




Vampir

- <http://www.nallas.com/e/products/vampir/index.htm>
- Primarily a tracing tool but also generates and displays profiling statistics
- Version 3 will support OpenMP and mixed MPI/OpenMP
- Version 3 will use PAPI to access hardware counter data
- Licensed at ERDC MSRC, ARL MSRC, and ARSC

9 April 2003 ICL/UTK 38

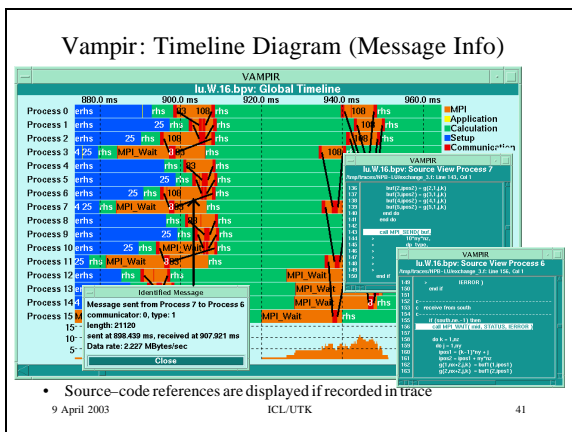


Vampir: Timeline Diagram

- Functions organized into groups
- Coloring by group
- Message lines can be colored by tag or size

Information about states, messages, collective, and I/O operations available by clicking on the representation

9 April 2003 ICL/UTK 40



Vampir: Profile Statistics Displays

- Aggregated profiling information: execution time, # calls, inclusive/exclusive
- Available for all/any group (activity)
- Available for all routines (symbols)
- Available for any trace part (select in timeline diagram)

9 April 2003 ICL/UTK 42

Vampir: Communication Statistics Displays

- Bytes sent/received for collective operations
- Byte and message count, min/max/avg message length and min/max/avg bandwidth for each process pair
- Message length statistics
- Available for any trace part

9 April 2003 ICL/UTK 43

Vampir: Other Features

- Dynamic global call graph tree
- Parallelism display
- Powerful filtering and trace comparison features
- All diagrams highly customizable (through context menus)

9 April 2003 ICL/UTK 44

Vampir: Process Displays

- Activity chart
- Call tree
- Timeline
- For all selected processes in the global displays

9 April 2003 ICL/UTK 45

Vampir (NAS Parallel Benchmark – LU)

- Timeline display
- Callgraph display
- Parallelism display
- Communications display

9 April 2003 ICL/UTK 46

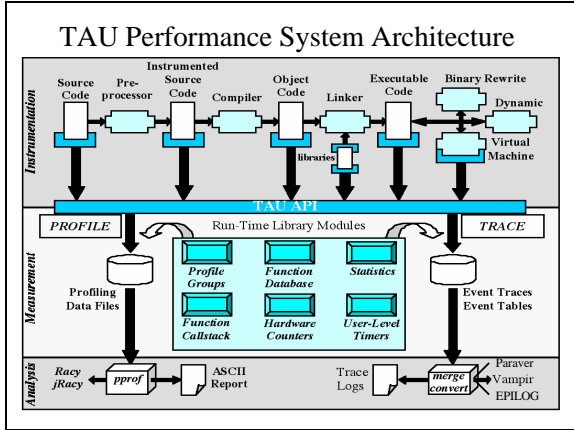
Vampir v3.x: Hardware Counter Data

9 April 2003 ICL/UTK 47

TAU

- Tuning and Analysis Utilities
- <http://www.cs.uoregon.edu/research/paracomp/tau/>
- Portable profiling and tracing toolkit for performance analysis of parallel programs
 - Fortran 77/90, C, C++, Java
 - OpenMP, Pthreads, MPI, mixed mode
- In use at DOE ASCI Labs and at NCSA
- Being installed and supported at DoD HPC Centers as part of PET CE002

9 April 2003 ICL/UTK 48



TAU Instrumentation

- Manually using TAU instrumentation API
- Automatically using
 - Program Database Toolkit (PDT)
 - MPI profiling library
 - Opari OpenMP rewriting tool
- Uses PAPI to access hardware counter data

9 April 2003 ICL/UTK 50

Program Database Toolkit (PDT)

- **Program code analysis framework** for developing source-based tools
- **High-level interface** to source code information
- **Integrated toolkit** for source code parsing, database creation, and database query
 - commercial grade front end parsers
 - portable IL analyzer, database format, and access API
 - open software approach for tool development
- Target and integrate **multiple source languages**
- Use in TAU to build **automated performance instrumentation tools**

9 April 2003 ICL/UTK 51

PDT Components

- **Language front end**
 - Edison Design Group (EDG): C, C++
 - Mutek Solutions Ltd.: F77, F90
 - creates an intermediate-language (IL) tree
- **IL Analyzer**
 - processes the intermediate language (IL) tree
 - creates “program database” (PDB) formatted file
- **DUCTAPE** (Bernd Mohr, ZAM, Germany)
 - C++ program **D**atabase **U**tilities and **C**onversion **T**ools **A**pplication **E**nvironment
 - processes and merges PDB files
 - C++ library to access the PDB for PDT applications

9 April 2003 ICL/UTK 52

PDT Status

- **Program Database Toolkit (Version 2.2, web download)**
 - EDG C++ front end (Version 2.45.2)
 - Mutek Fortran 90 front end (Version 2.4.1)
 - C++ and Fortran 90 IL Analyzer
 - DUCTAPE library
 - Standard C++ system header files (KCC Version 4.0f)
- **PDT-constructed tools**
 - TAU instrumentor (C/C++/F90)
 - Program analysis support for SILOON and CHASM
- **Platforms**
 - SGI, IBM, Compaq, SUN, HP, Linux (IA32/IA64), Apple, Windows, Cray T3E, Hitachi

9 April 2003 ICL/UTK 53

OPARI: Basic Usage (f90)

- Reset **OPARI** state information
 - `rm -f opari.rc`
- Call **OPARI** for each input source file
 - `opari file1.f90`
 - ...
 - `opari fileN.f90`
- Generate **OPARI** runtime table, compile it with ANSIC
 - `opari -table opari.tab.c`
 - `cc -c opari.tab.c`
- Compile modified files `*.mod.f90` using OpenMP
- Link the resulting object files, the **OPARI** runtime table `opari.tab.o` and the TAU **POMP** RTL `libpomp.a`

9 April 2003 ICL/UTK 54

TAU Analysis

- Profile analysis
 - pprof**
 - parallel profiler with text-based display
 - racy**
 - graphical interface to pprof (Tcl/Tk)
 - jracy**
 - Java implementation of Racy
- Trace analysis and visualization
 - Trace merging and clock adjustment (if necessary)
 - Trace format conversion (ALOG, SDDF, Vampir)
 - Vampir** (Pallas) trace visualization
 - Paraver (CEPBA) trace visualization

9 April 2003

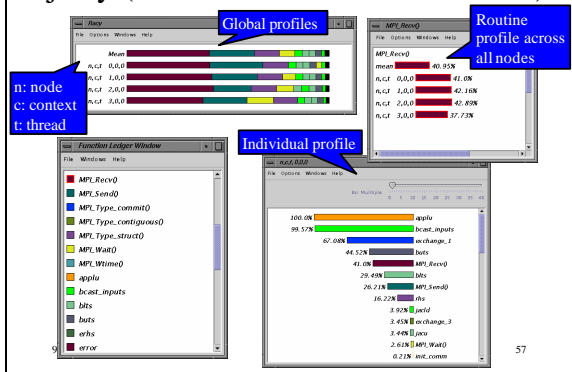
ICL/UTK

55

TAU Pprof Display

Time	Exclusive msec	Inclusive Total msec	#Call	#Subrs	Inclusive Name
100.0	1	3,111,293	1	15	191293269 applu
99.4	3,667	2,110,463	3	37517	63487928 bcact_inputs
67.1	491	2,051,326	37200	17200	3450 exchange_1
44.5	6,461	1,275,159	9300	18600	9127 bits
41.0	1,118,436	1,115,436	18600	0	4217 MPI_Recv()
29.5	5,778	56,407	9300	0	6068 bits
26.2	50,142	50,142	17504	0	1561 MPI_Send()
16.2	24,451	31,031	301	602	10396 rns
3.3	7,201	7,201	9300	0	807 jactd
3.4	538	6,594	504	1812	10918 exchange_3
3.4	6,590	6,590	9300	0	703 jactv
2.6	4,989	4,989	608	0	8206 MPI_Wait()
0.2	0.44	400	1	4	400081 init_com
0.1	398	399	1	0	399634 MPI_Init()
0.1	140	247	1	47616	247086 setiv
0.1	131	131	57282	0	2 erract
0.1	89	103	1	2	103168 errns
0.1	0.966	96	1	0	96458 read_input
0.0	1	24	9	0	10603 MPI_Bcast()
0.0	26	44	1	7937	44878 error
0.0	24	12	1	0	40 MPI_Recv()
0.0	15	15	1	5	15630 MPI_Finalize()
0.0	4	12	1	1700	12336 setiv
0.0	7	8	3	3	2893 L2norm
0.0	1	3	1	0	491 MPI_Allreduce()
0.0	1	3	1	1	3874 pintgr
0.0	0.116	0.837	1	4	837 exchange_4
0.0	0.512	0.512	1	0	512 MPI_Keyval_create()
0.0	0.121	0.353	1	2	353 exchange_5
0.0	0.024	0.151	1	2	151 exchange_6
0.0	0.103	0.103	6	0	17 MPI_Type_contiguous()

jracy (NAS Parallel Benchmark – LU)



1-Level Callpath Implementation in TAU

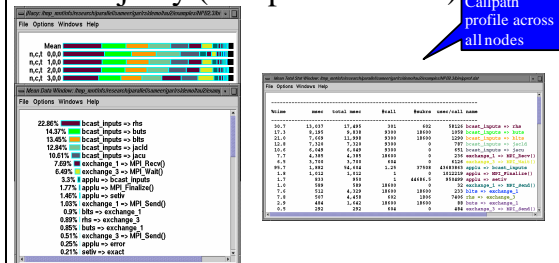
- TAU maintains a performance event (routine) callstack
- Profiled routine (child) looks in callstack for parent
 - Previous profiled performance event is the parent
 - A *callpath profile structure* created first time parent calls
 - TAU records parent in a *callgraph map* for child
 - String representing 1-level callpath used as its key
 - "a()->b()": name for time spent in "b" when called by "a"
- Map returns pointer to callpath profile structure
 - 1-level callpath is profiled using this profiling data
- Build upon TAU's performance mapping technology
- Measurement is independent of instrumentation
- Use `-PROFILECALLPATH` to configure TAU

9 April 2003

ICL/UTK

58

jracy (Callpath Profiles)



9 April 2003

ICL/UTK

59

SvPablo

- Sourceview Pablo
- <http://www.pablo.cs.muc.edu/Project/SVPablo/SVPabloOverview.htm>
- Collects profile data and maps it to constructs in the original source code
- Automatic and interactive instrumentation of Fortran 77/90, C, MPI, OpenMP
- Uses PAPI to access hardware counter data
- In use at DOE sites and at NCSA

9 April 2003

ICL/UTK

60

SvPablo

9 April 2003 61

HPCView

- <http://www.cs.rice.edu/~dsystem/hpview/>
- Combines multiple set of program profile data
- Correlates profile data with source code
- Computes derived performance metrics
- Generates browsable database
- Uses srrun, uprofile, and PAPI to collect profile data
- In use at DOE ASCI Labs and at NCSA

9 April 2003 62

VProf

- Visual Profiler
- <http://aros.ca.sandia.gov/~cljanss/perf/vprof/>
- Generate profiling data sorted by file, function, line
- Display using vprof GUI or cprof command-line utility
- Ported to Linux/IA-32/IA-64, AIX/RS6000, Linux/Alpha
- In use at DOE ASCI sites and at NCSA

9 April 2003 63

vprof

9 April 2003 64