

### Profile Data Mining with PerfExplorer



Sameer Shende Performance Reseaerch Lab, University of Oregon http://TAU.uoregon.edu



#### **TAU Analysis**





VI-HPS

#### **TAUdb: Performance Data Mgmt. Framework**



#### **Performance Analysis Programs TAU Performance System** cluster scalability profile ParaProf analysis metadata analysis 00 00 raw profiles **Query and Analysis Toolkit** Data Mining (Weka) 00 **Statistics** gprof (R / Omega) \* mpiP psrun \* HPMtoolkit Java PerfDMF API $^{*}$ ... SQL (PostgreSQL, MySQL, DB2, Oracle) XML document formatted ⊲--profile data

#### VI-HPS

#### Using TAUdb



• Configure TAUdb (Done by each user)

% taudb\_configure --create-default

- Choose derby, PostgreSQL, MySQL, Oracle or DB2
- Hostname
- Username
- Password
- Say yes to downloading required drivers (we are not allowed to distribute these)
- Stores parameters in your ~/.ParaProf/taudb.cfg file
- Configure PerfExplorer (Done by each user)

% perfexplorer\_configure

Execute PerfExplorer

% perfexplorer

#### **Using PerfExplorer**

```
VI-HPS
```

```
% wget <u>http://tau.uoregon.edu/data.tgz</u> (Contains CUBE profiles from Score-P)
% module load UNITE tau
% taudb configure --create-default
(Chooses derby, blank user/passwd, yes to save passwd, defaults)
% perfexplorer configure
(Yes to load schema, defaults)
% paraprof
(load each trial: DB -> Add Trial -> Type (Paraprof Packed Profile) -> OK) OR use
   taudb loadtrial -a "app" -x "experiment" -n "name" file.ppk
Then,
% perfexplorer
(Select experiment, Menu: Charts -> Speedup)
```



- Development of the TAU portal
  - Common repository for collaborative data sharing
  - Profile uploading, downloading, user management
  - Paraprof, PerfExplorer can be launched from the portal using Java Web Start (no TAU installation required)
- Portal URL

http://tau.nic.uoregon.edu



- Performance knowledge discovery framework
  - Data mining analysis applied to parallel performance data
    - comparative, clustering, correlation, dimension reduction, ...
  - Use the existing TAU infrastructure
    - TAU performance profiles, taudb
  - Client-server based system architecture
- Technology integration
  - Java API and toolkit for portability
  - taudb
  - R-project/Omegahat, Octave/Matlab statistical analysis
  - WEKA data mining package
  - JFreeChart for visualization, vector output (EPS, SVG)



- Performance data represented as vectors each dimension is the cumulative time for an event
- *k*-means: *k* random centers are selected and instances are grouped with the "closest" (Euclidean) center
- New centers are calculated and the process repeated until stabilization or max iterations
- Dimension reduction necessary for meaningful results
- Virtual topology, summaries constructed

#### **PerfExplorer - Cluster Analysis (sPPM)**







• Describes strength and direction of a linear relationship between two variables (events) in the data

€ ● ●			PerfExplorer	Client					
<u>File Analysis Views Charts Visualization H</u>	<u>H</u> elp								
Performance Data	Analysis Management     Cluster Results     Correlation Results								
🔻 河 Database Profiles									
► 📁 AVUS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.0
BigScience	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.7
► D CFDSHIP	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.5
FLASH_2.5_hydro_radiation	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.2
V 💭 LLNL_UBGL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0
▼ 💭 0064									
F 🕒 Time	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.0
▶ 📁 0124	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.7
► 📁 0256	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.5
▶ 📁 0512	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.2
V 💭 1024	0.00 0.5 1.0	0.00	0.00 0.5 1.0	0.00 0.5 1.0	0.00 0.5 1.0	0.00 0.5 1.0	0.00 0.5 1.0	0.00	0.0
🕨 🔍 Time									
IAMMPS (Large-scale Atomic Molecular N	1.00	1.00	1.00	1.00	1.00	1.9	1.00	1.00	1.0
🕨 📁 Miranda	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.7
► 📁 POP	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.5
► I SHAMRC	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.2
▶ [] SMG2000	0.00 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.00 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0
▶ 📁 SPhot									
▶ 📁 Uintah	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.0
► 📁 WRF	0.75	0.75	0.75	0.75	0.76	0.75	0.75	0.75	0.7
▶ 📁 bacc_app	0.50	0.50	0.50	0.50	0.50	0.50	0.00	0.50	0.5
▶	0.25	0.25	0.00	100	0.00	0.00	0.00	0.00	0.2
gyro.B1-std.HPM	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0
▶					100			1.00	
gyro.B2-cy.HPM	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	1.0
▶	0.60	0.50		0.60	0.50	0.50	0.50	0.50	0.5
gyro.B3-gtc.HPM	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.2
► 📁 sPPM	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0
socorro_Si256_input	0.0 0.5 1.0	0.0 0.8 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	
▶ 📁 Views	100 (	1.00	1.00	1.00 (	1.00	1.00 (	1 00 (	1 00 (	10
	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.7
	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.5
	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.2
	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0
	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	0.0 0.5 1.0	
	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.0
	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.7 🔻
									4 +
)++									

#### **PerfExplorer - Correlation Analysis (Flash)**



- -0.995 indicates strong, negative relationship
- As CALC\_CUT\_ BLOCK\_CONTRIBUTIO NS() increases in execution time, MPI\_Barrier() decreases



#### **PerfExplorer - Comparative Analysis**



- Relative speedup, efficiency
  - total runtime, by event, one event, by phase
- Breakdown of total runtime
- Group fraction of total runtime
- Correlating events to total runtime
- Timesteps per second

#### **PerfExplorer - Interface**

## VI-HPS



PerfExplorer Client





#### **PerfExplorer - Interface**



O O PerfExplorer Client								
<u>File Analysis</u> <u>Charts</u> <u>H</u> elp								
▼	Analysis Management	• Performance Explorer						
Set Event of Interest	Field	Value						
Set Total Number of Timesteps	Name	B1-std-nl2.cheetah.noaffnosng						
Timesteps Per Second	Experiment ID	16						
B1-sto	system_name							
Relative Efficiency by Event	system_mach system_os							
Relative Efficiency for One Event								
Relative Speedup	system_memory_size							
Relative Speedup by Event	system_processor_amt							
Relative Speedup for One Event	system_12_cache_size							
B1-stc Communication Time / Total Ru	ntime tem_userdata							
B1-stc. Runtime Breakdown								
B1-std.hockney	Select analysis							
B1-std.new.phoenix.0x002	complier_cc_version							
B1-std.phoenix.0x002	compiler_java_dirpath							
B1-std.phoenix.0x002scr	compiler_java_version							
B1-std.ram0x002.a	compiler_userdata							
🕨 📁 B1-std.ram0x002.b	configure_arch							
🔻 📁 B1-std.seaborg	configure_cpp							
B1-std.timing.seaborg.128	configure_cc							
B1-std.timing.seaborg.16	configure_jdk							
B1-std.timing.seaborg.256	configure_profile							
B1-std.timing.seaborg.32	userdata							
B1-std.timing.seaborg.512								
B1-std.timing.seaborg.64								
► □ B1-std.tg								
gyro.B2-cy	*							
F had all a constructions and the second sec								

#### **PerfExplorer - Relative Efficiency Plots**





#### **PerfExplorer - Relative Efficiency by Routine**





#### **PerfExplorer - Relative Speedup**





#### **PerfExplorer - Timesteps Per Second**





## Usage Scenarios: Evaluate Scalability HPS

- Goal: How does my application scale? What bottlenecks occur at what core counts?
- Load profiles in taudb database and examine with PerfExplorer



# Usage Scenarios: Evaluate Scalabity HPS



PerfExplorer





### PerfExplorer





### Performance Regression Testing VI-HPS



#### Support Acknowledgments

## VI-HPS

- •US Department of Energy (DOE)
  - Office of Science contracts
  - SciDAC, LBL contracts
  - LLNL-LANL-SNL ASC/NNSA contract
  - Battelle, PNNL contract
  - ANL, ORNL contract
- •Department of Defense (DoD)
  - PETTT, HPCMP
- •National Science Foundation (NSF)
  - Glassbox, SI-2
- •University of Tennessee, Knoxville
- •T.U. Dresden, GWT
- •Juelich Supercomputing Center



Pacific Northwest

















