On Utilizing Experiment Data Repository for Performance Analysis of Parallel Applications

Hong-Linh Truong[1], Thomas Fahringer[2]

[1] Institute for Software Science,
University of Vienna, Austria
truong@par.univie.ac.at

[2] Institute of Computer Science,
University of Innsbruck, Austria
Thomas.Fahringer@uibk.ac.at

www.par.univie.ac.at/project/scalea

AURORA

APART

Outline

- Motivation
- SCALEA
- Experiment Data Repository
  - Experiment-related Data
  - Implementation Overview
  - Experiment-related APIs
- Achievements
  - Search and Filter Performance Data
  - Multi-Experiment Analysis
  - Tools Integration
- Conclusions and Future Work

Motivation

- The need to collect and archive performance data for
  - Multi-experiment analysis
  - Performance comparison
- Limitation on supporting code complexity:
  - Handle large-scale performance data
  - Basic search on performance data
- Lack of capabilities to export, share, and exchange performance data and tools interoperation

SCALEA (SC01, Euro-Par02, PVMMPI02)

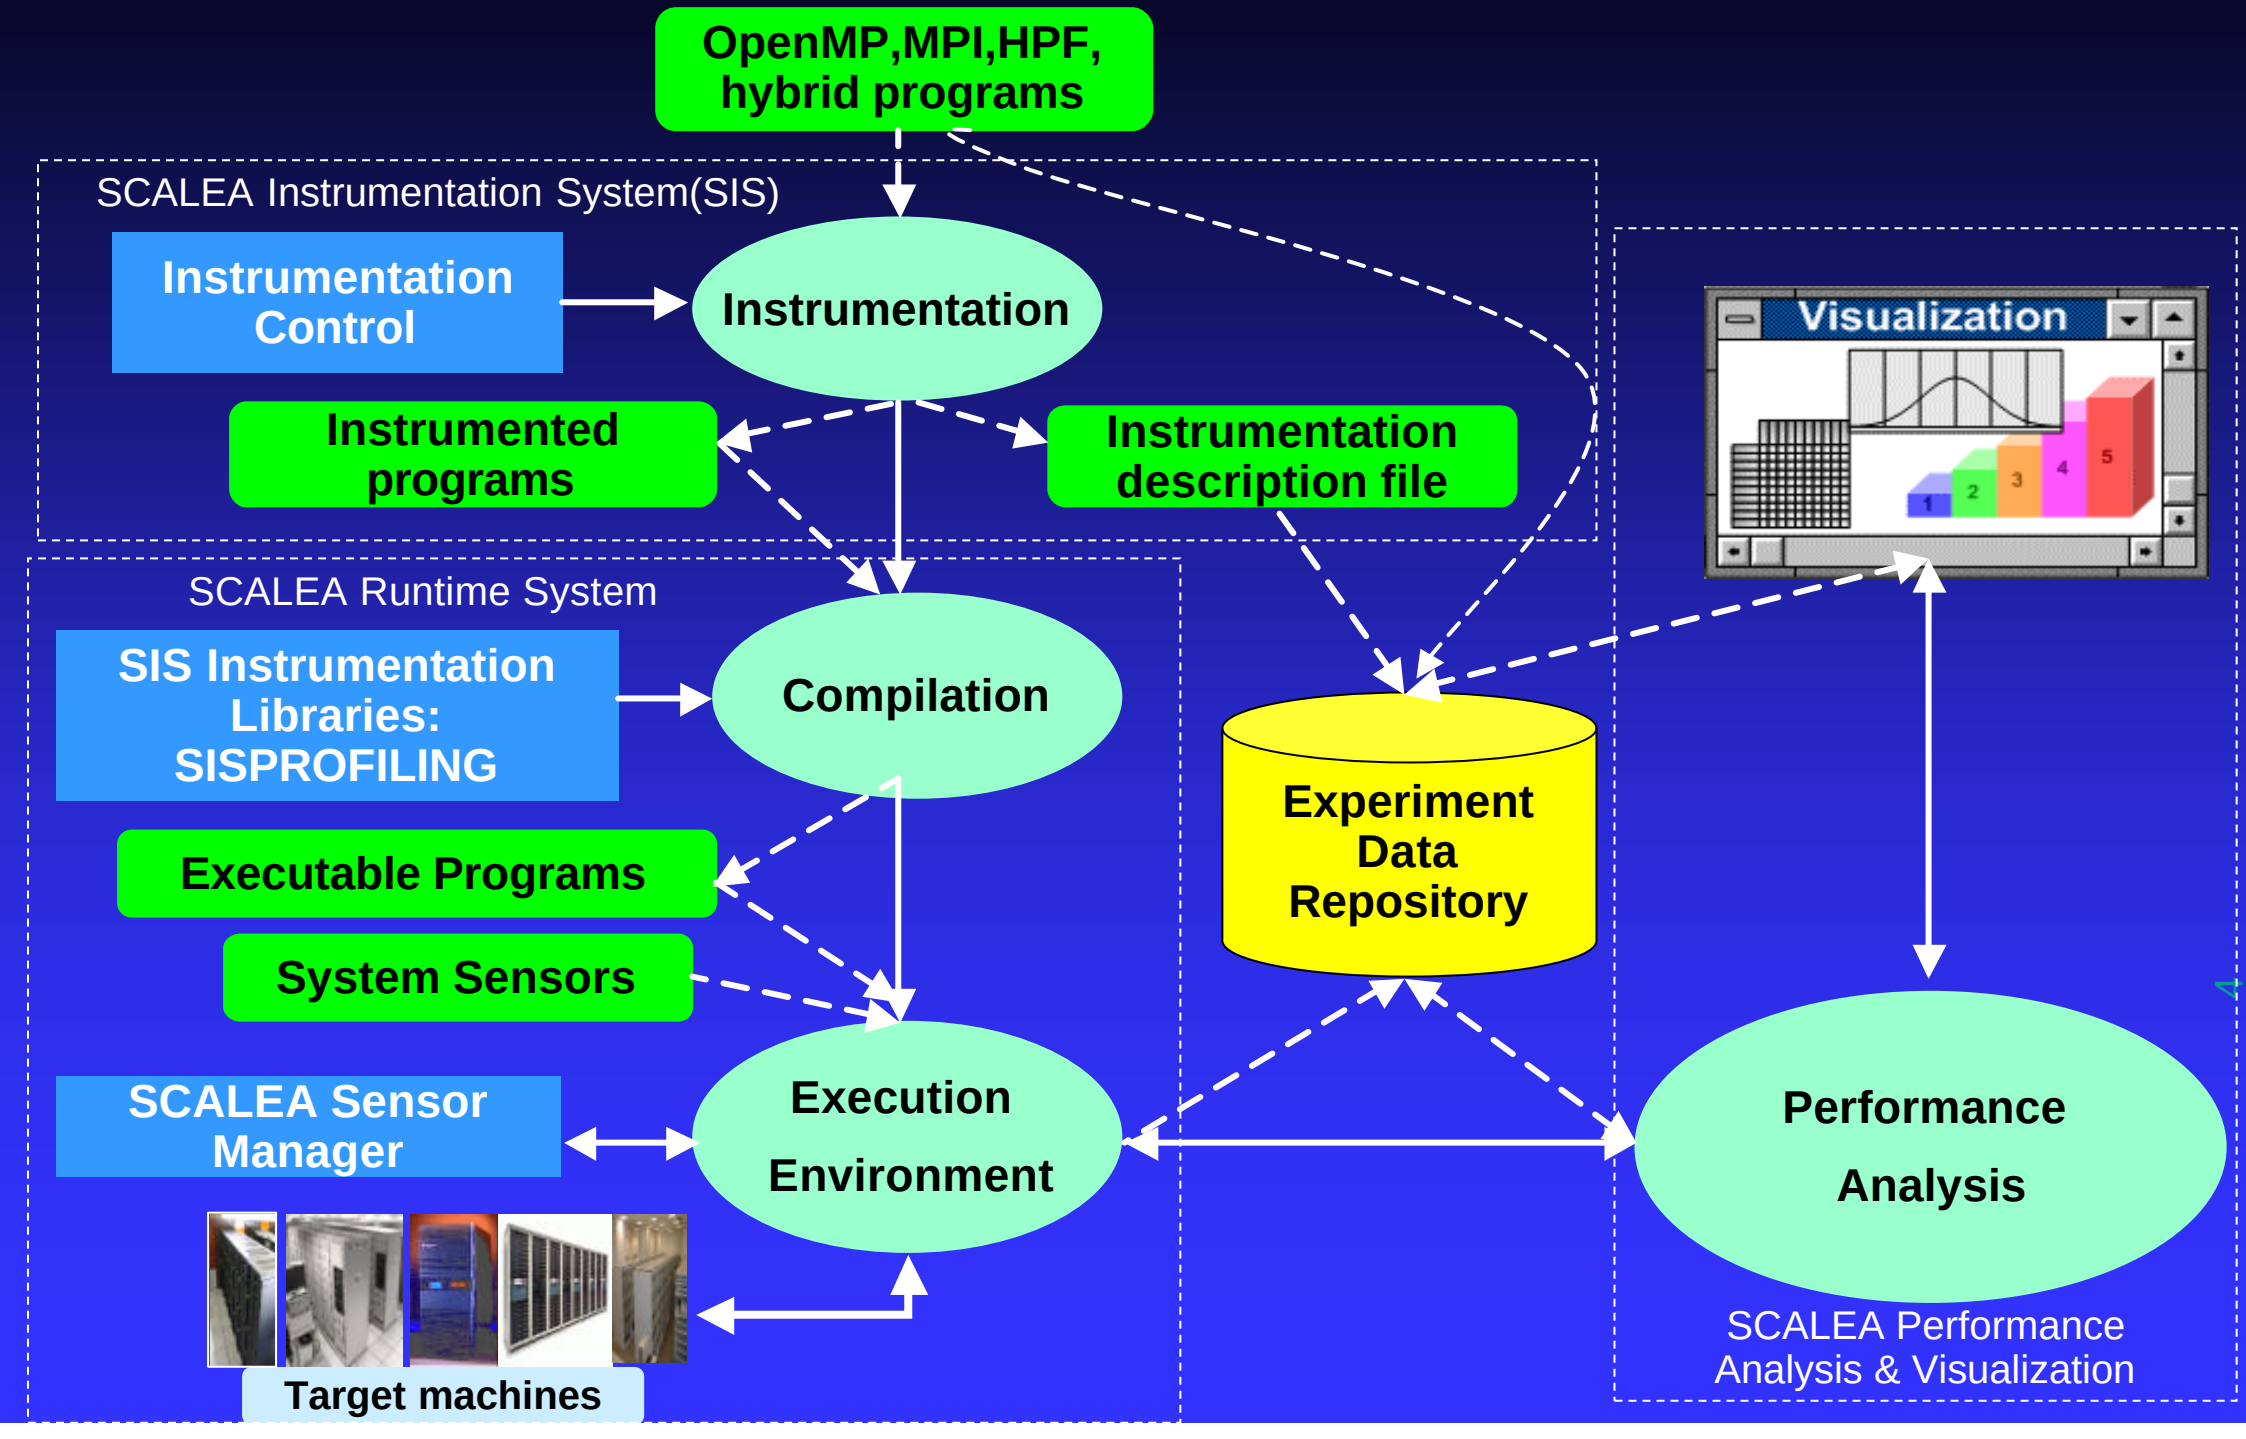

Experiment Data Repository

- For comparing performance across experiments with large amounts of data
  - Need to archive performance data
  - Require a very strong and flexible database system
- Data repository is first step required by mining, knowledge discovery in performance data
- Experiment Data Repository:
  - Stores information about application, source code, machine information, and performance data
  - Associates performance results with source code and machine information
  - Supports sophisticated analysis and easy integration with other tools

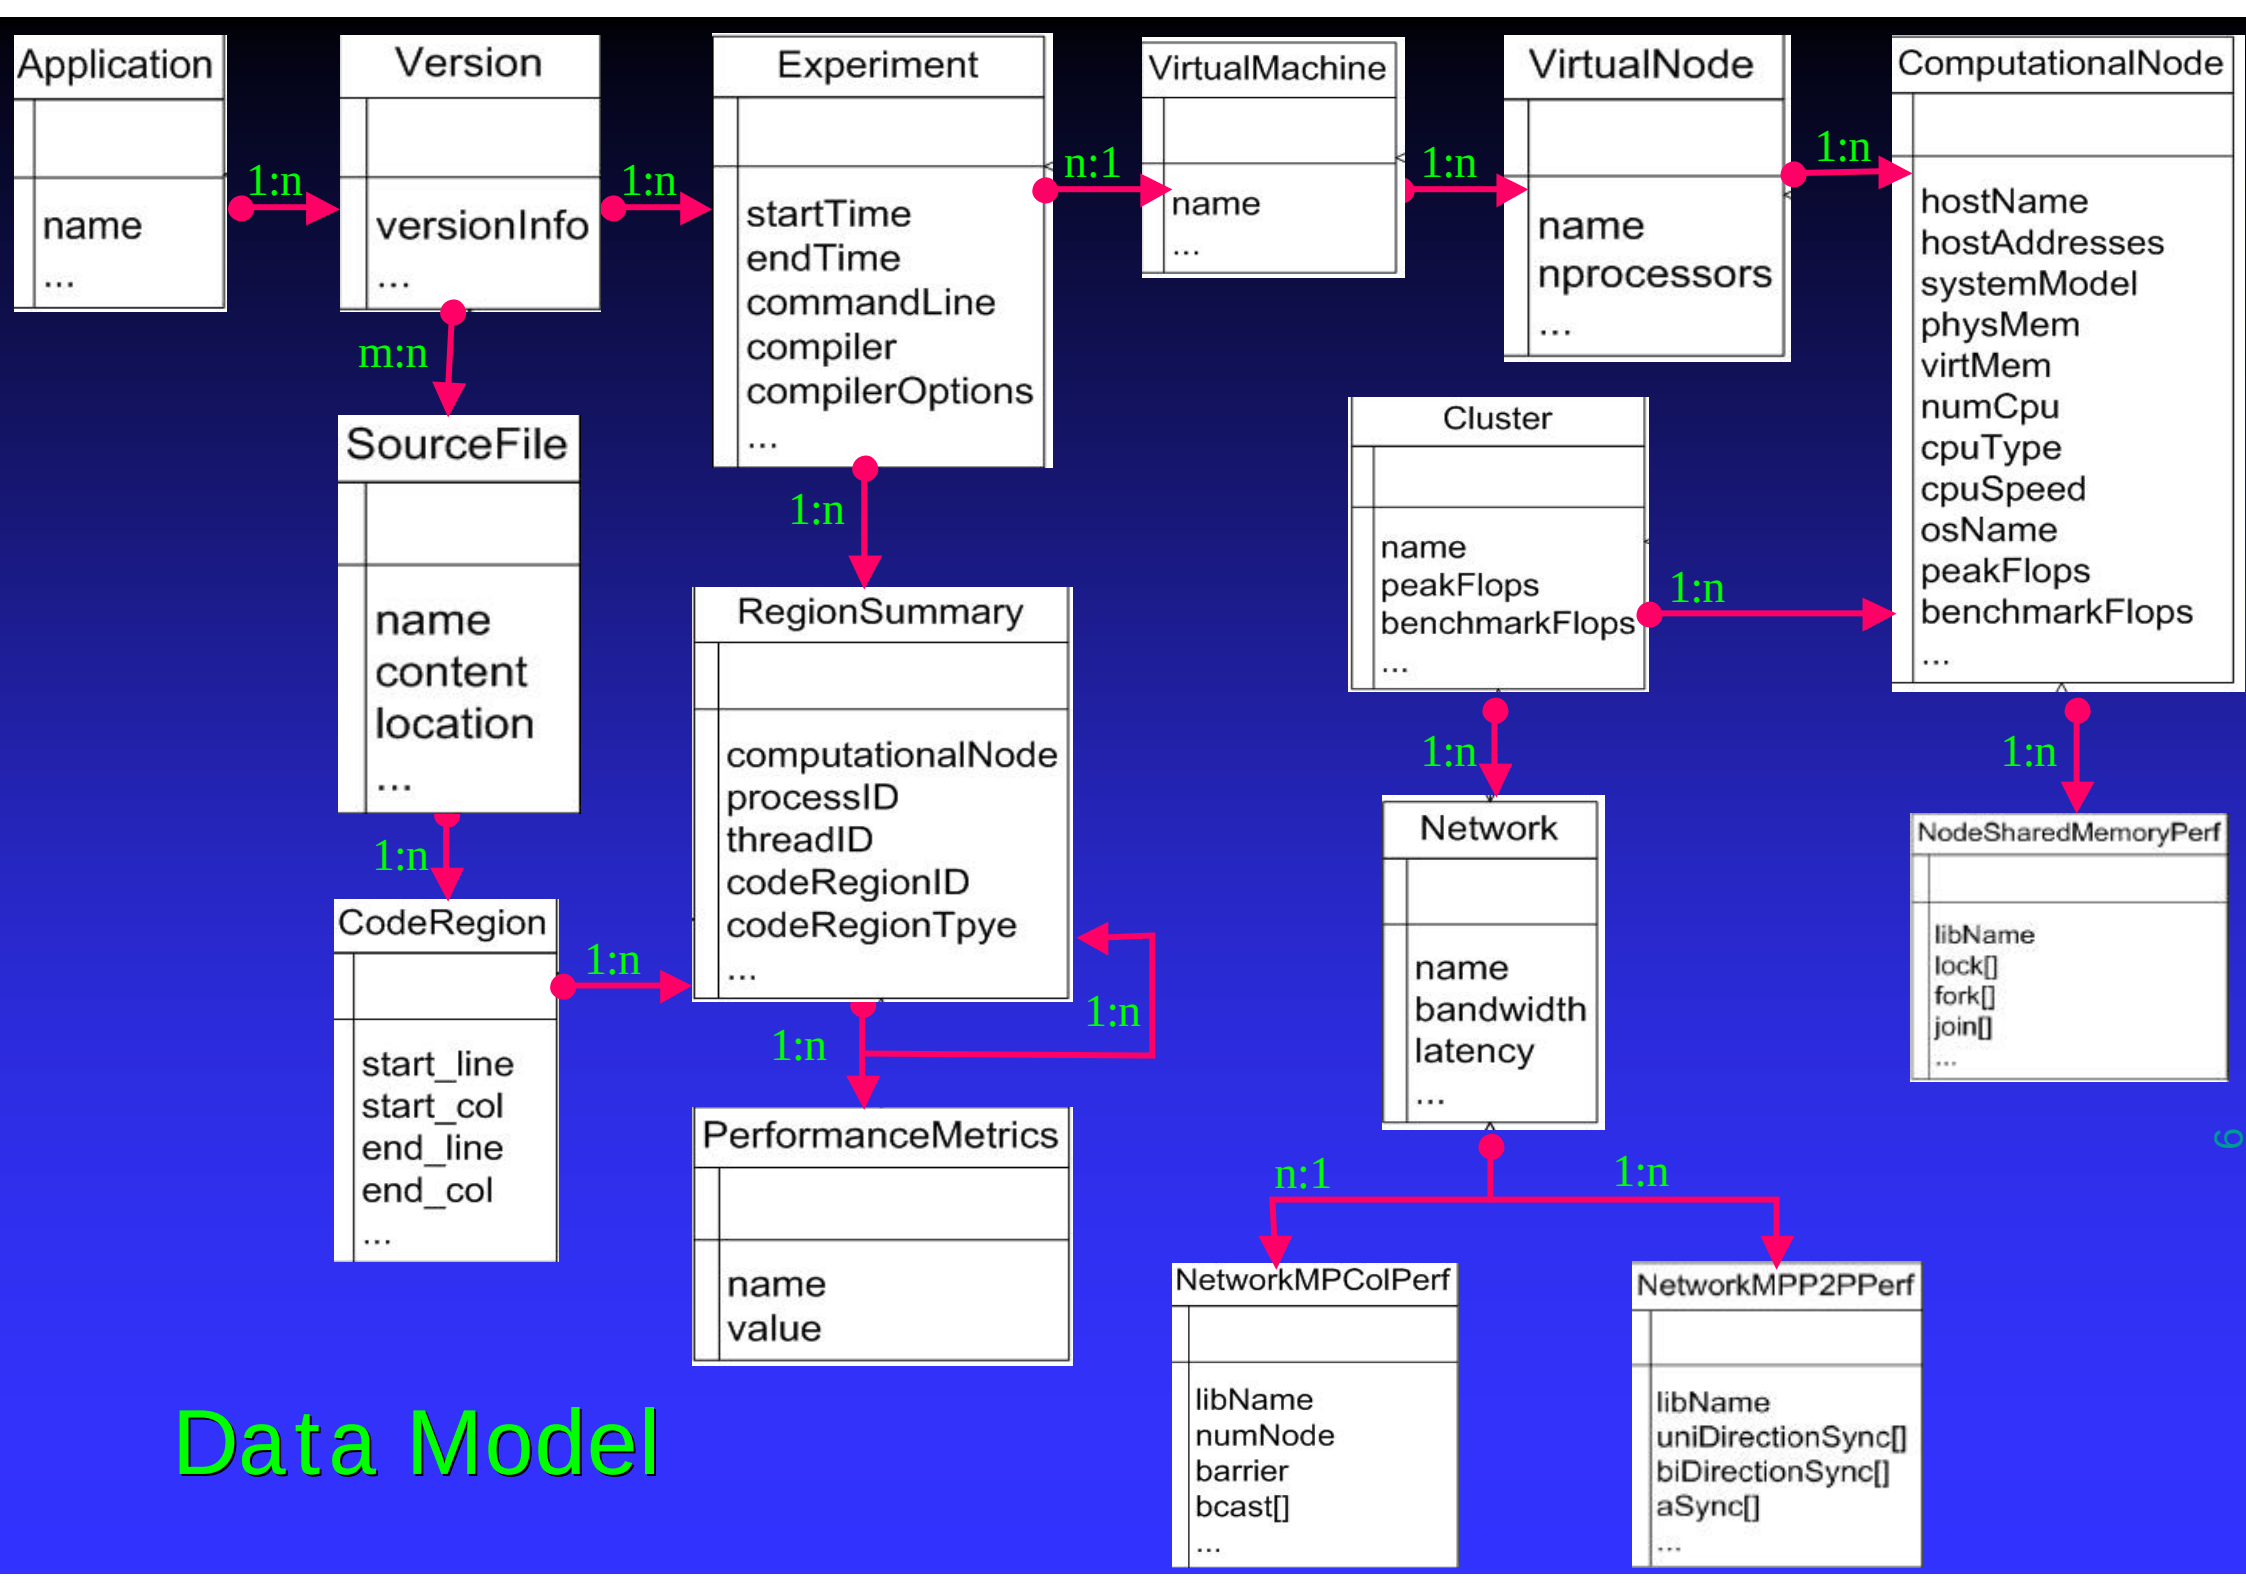
Application
name
...
1:n
Version
versionInfo
...
1:n
Experiment
startTime
endTime
commandLine
compiler
compilerOptions
...
n:1
VirtualMachine
name
...
1:n
VirtualNode
name
nprocessors
...
1:n
ComputationalNode
hostName
hostAddresses
systemModel
physMem
virtMem
numCpu
cpuType
cpuSpeed
osName
peakFlops
benchmarkFlops
...
m:n
SourceFile
name
content
location
...
1:n
CodeRegion
start_line
start_col
end_line
end_col
...
1:n
RegionSummary
computationalNode
processID
threadID
codeRegionID
codeRegionTpye
...
1:n
1:n
1:n
PerformanceMetrics
name
value
Cluster
name
peakFlops
benchmarkFlops
...
1:n
1:n
1:n
Network
name
bandwidth
latency
...
NodeSharedMemoryPerf
libName
lock[]
fork[]
join[]
...
n:1
1:n
NetworkMPColPerf
libName
numNode
barrier
bcast[]
...
NetworkMPP2PPerf
libName
uniDirectionSync[]
biDirectionSync[]
aSync[]
...
Data Model

Performance Metric Catalog

- Why performance metric catalog:
  - Documentation about metrics provided
  - Help tools/users understanding and interpreting performance data
- Performance metric has 4 properties
  - unique name is a value that distinguishes this metric from all others
  - data type describes data type of measurement which describes the value of the metric.
  - measurement unit is the unit of measurement
  - semantics (or well-defined meaning)
- All metrics are stored in a catalog in experiment data repository

Performance Metric Catalog

Metric Name	Data Type	Measurement Unit	Description
WTIME	double	us	Wallclock time
ODATA_L21	double	us	Overhead of data movement d...
ODATA_L32	double	us	Overhead of data movement d...
ODATA_SEND	double	us	Overhead of data movement d...
ODATA_RECV	double	us	Overhead of data movement d...
ODATA_PUT	double	us	Overhead of data movement d...
ODATA_GET	double	us	overhead of data movement d...
ODATA_P2P	double	us	overhead of data movement d...
ODATA_COL	double	us	Overhead of data movement d...
ODATA_FREAD	double	us	Overhead of data movement d...
ODATA_FWRITE	double	us	Overhead of data movement d...

Implementation Overview

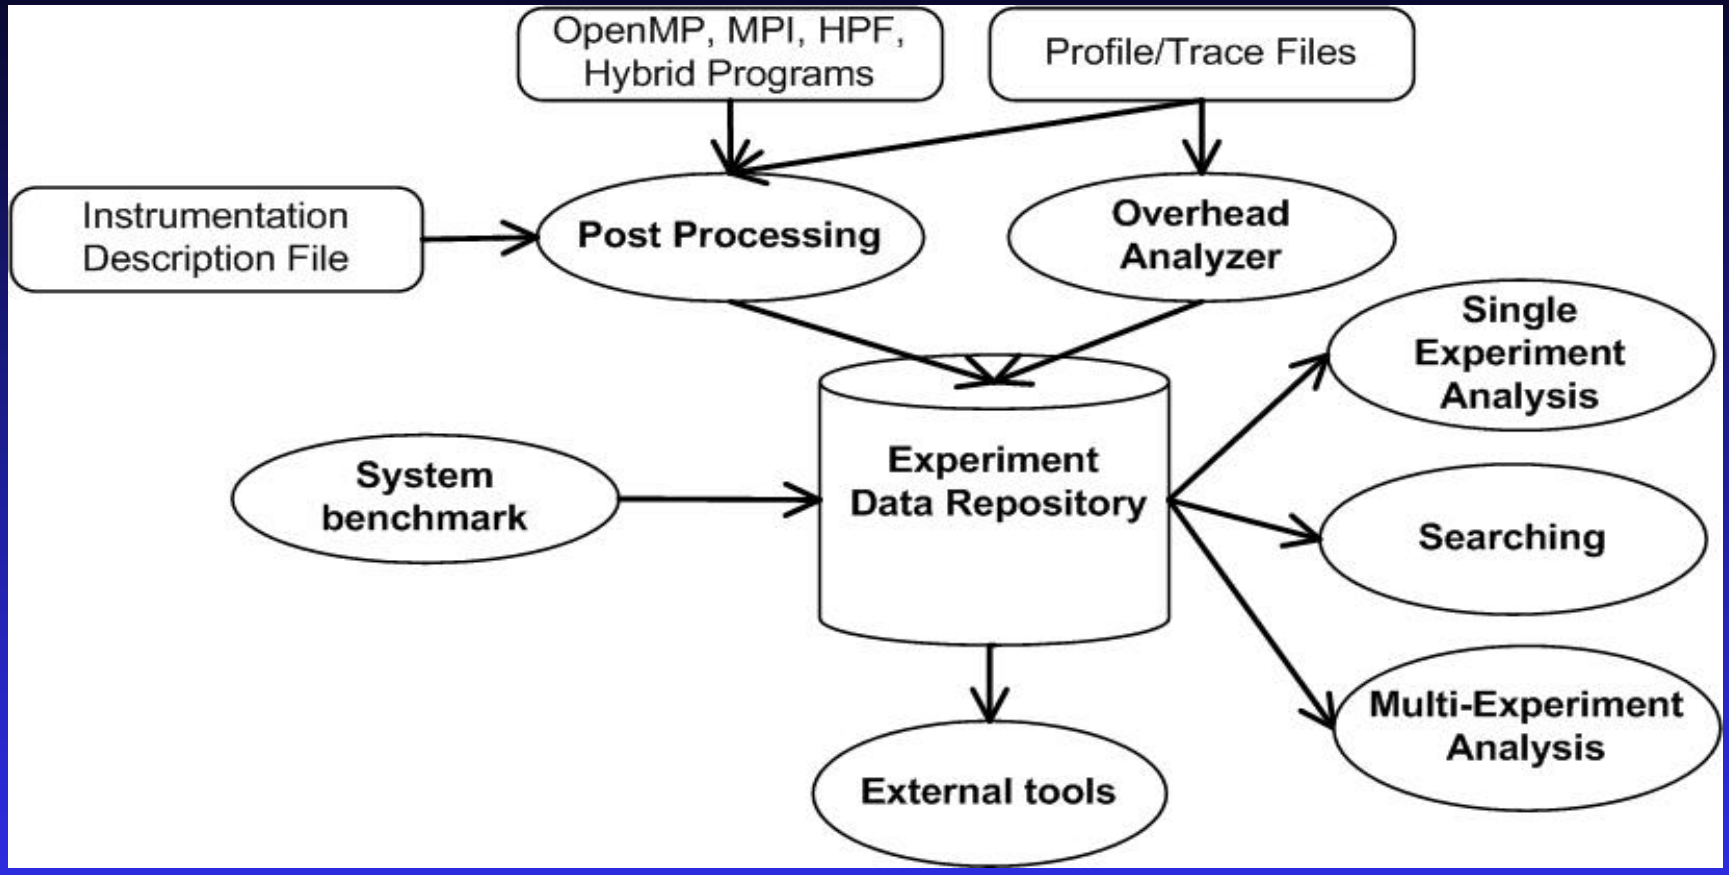

- Powering search and archive with great flexibility and robustness
  - Relational database based on PostgreSQL
- Taking advantages of portability and network capability
  - All components written in Java
  - Database connection powered by JDBC

Experiment-related Data APIs

- Well-defined APIs for accessing data in Experiment Data Repository
- Java implementation

```
public class PerformanceMetric {
   private String metricName;
   private Object metricValue;
   public String getMetricName(){…};
   public void   setMetricName(String metricName){…};
   public Object getMetricValue(){…};
   public void   setMetricValue(Object metricValue){…};
   ...
   }
```

```
public class ProcessingUnit {
   private String computationalNode;
   private int    processID ;
   private int    threadID ;
   …
   public ProcessingUnit(String node,int process, int thread) {...}
   …
   }
```

Experiment-related Data APIs (const.)

```
public class RegionSummary  {
  ...
  public PerformanceMetric[] getMetrics(){...}
  public PerformanceMetric getMetric(String metricName){...}
  ...
  }
```

```
public class ExperimentData {
  DatabaseConnection   connection;
  ...
  public ProcessingUnit[]  getProcessingUnits(Experiment e){...}
  public RegionSummary[]   getRegionSummaries(CodeRegion cr, Experiment
  e){...}
  public RegionSummary   getRegionSummary(CodeRegion cr, ProcessingUnit
  pu, Experiment e) {...}
  public RegionSummary   getRegionSummary(CodeRegion cr, Experiment
  e,ProcessingUnit pu, RegionSummary parent) {...}   public
  RegionSummary[]   getChildOfRegionSummary(RegionSummary rs){...}
  public RegionSummary    getParentOfRegionSummary(RegionSummary
  rs){...}
  ...
}
```

Example of Using APIs

```
CodeRegion cr = new CodeRegion(“IONIZE_MOVE”);
Experiment e = new Experiment(“9Nx4P,P4”);
ProcessingUnit
   pu =new ProcessingUnit(``gsr402.vcpc.ac.at'',2,0);

ExperimentData
   ed = new ExperimentData(new DatabaseConnection(...));
RegionSummary  rs = ed.getRegionSummary (cr,pu,e);


PerformanceMetric
         overhead =rs.getMetric(''oall_ident'');
PerformanceMetric
         wtime   =rs.getMetric(``wtime'');

double overheadRatio=
   ((Double)overhead.getMetricValue()).doubleValue()/
       ((Double)wtime.getMetricValue()).doubleValue();
```

XML Data

- Facilitate information exchange between tools
- Allow tools interoperation in a uniform way
- With/without source code
- Call-graph, selective metrics

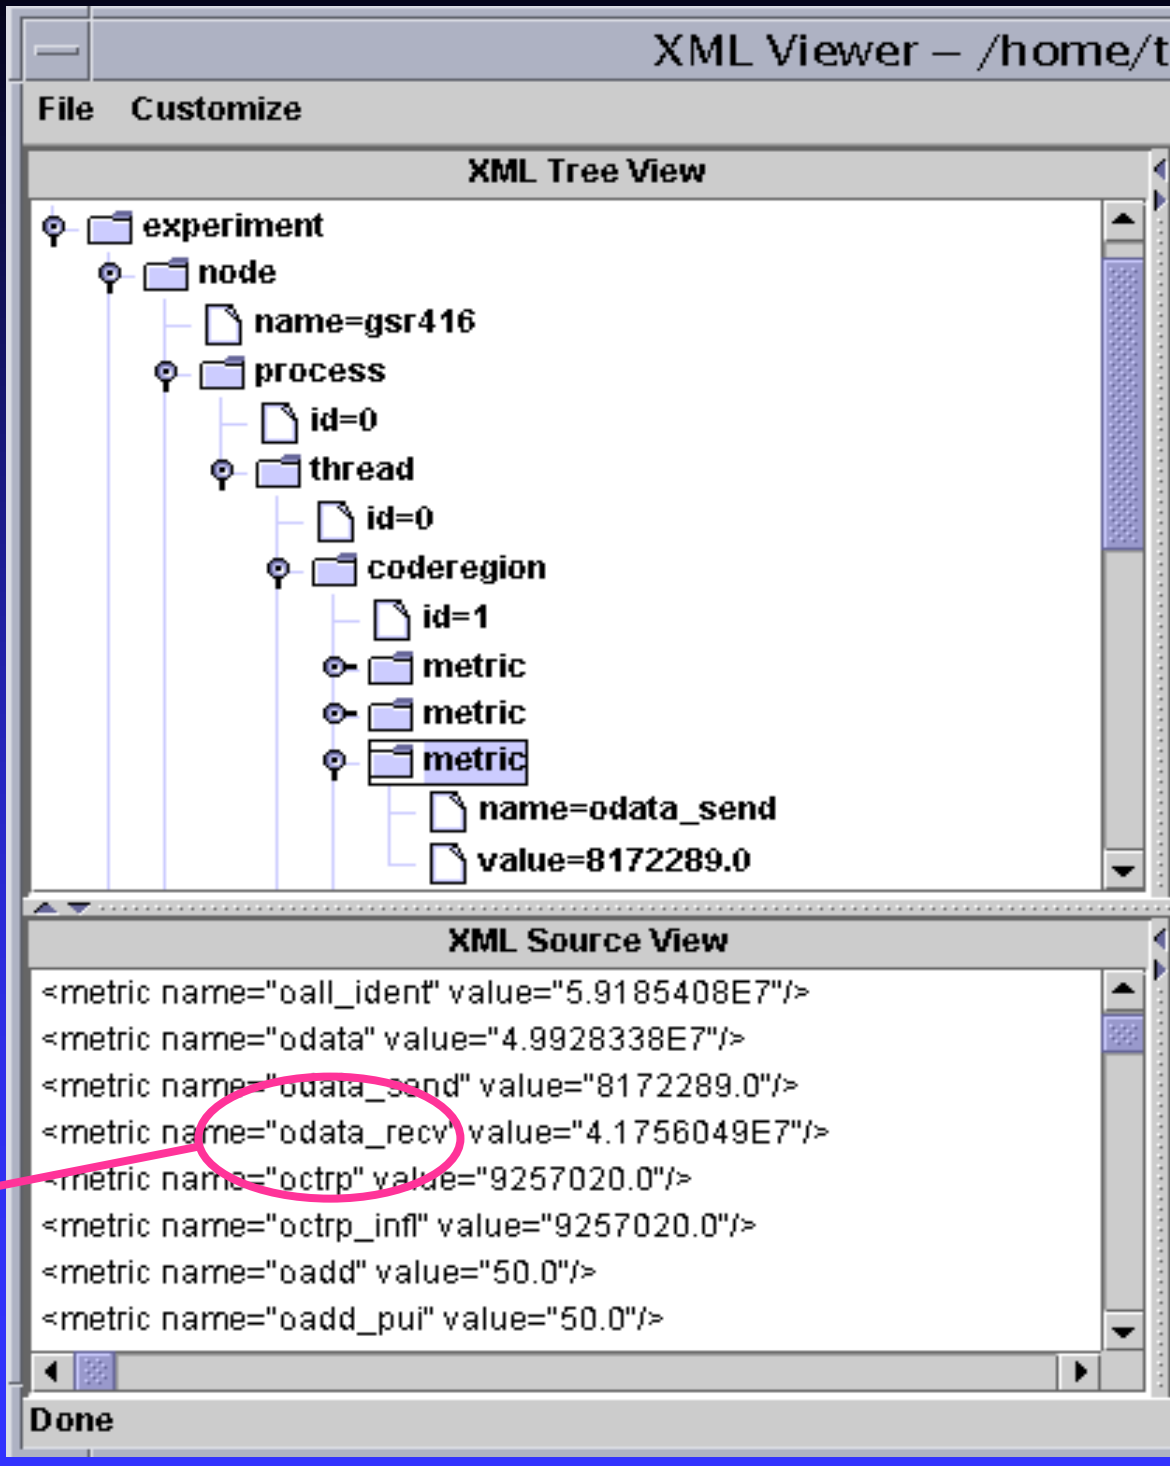

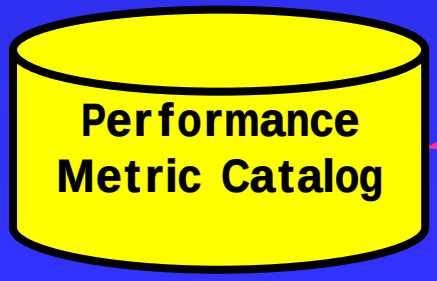

Search and Filter Capabilities

- Existing problems:
  - difficult to find interesting events via process time-lines, zooming, scrolling, call-graph
  - search and filter data based on files are not efficient, robust and fast

Picture taken from D. Kranzlmueller ´s talk

- Search and filter data based relational database and SQL: flexibility and robustness
- We support generic search and filter

Search and Filter Capabilities

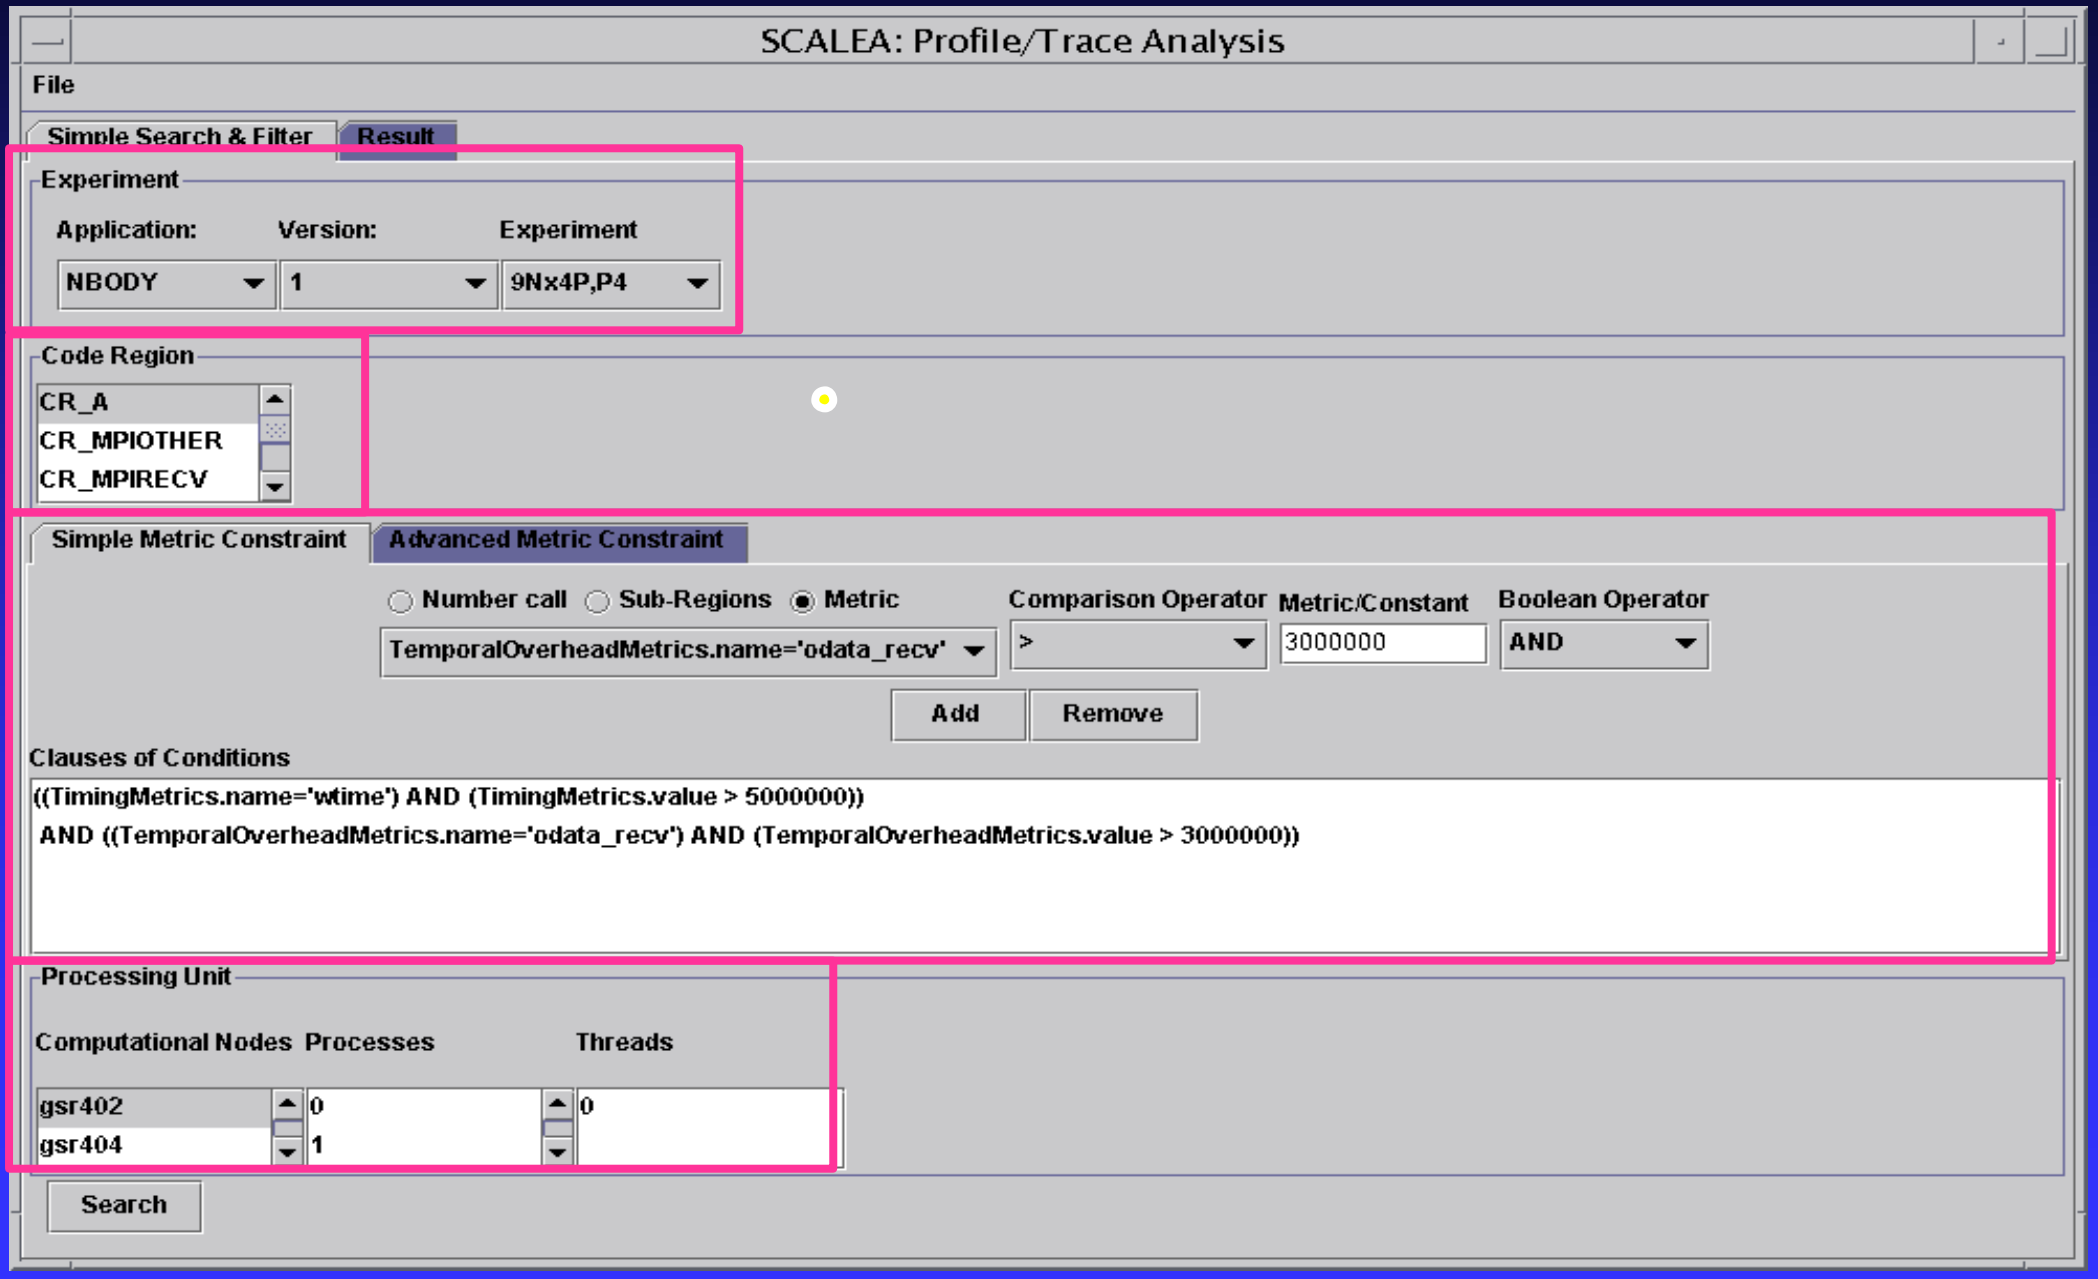

Search and Filter Capabilities (const.)

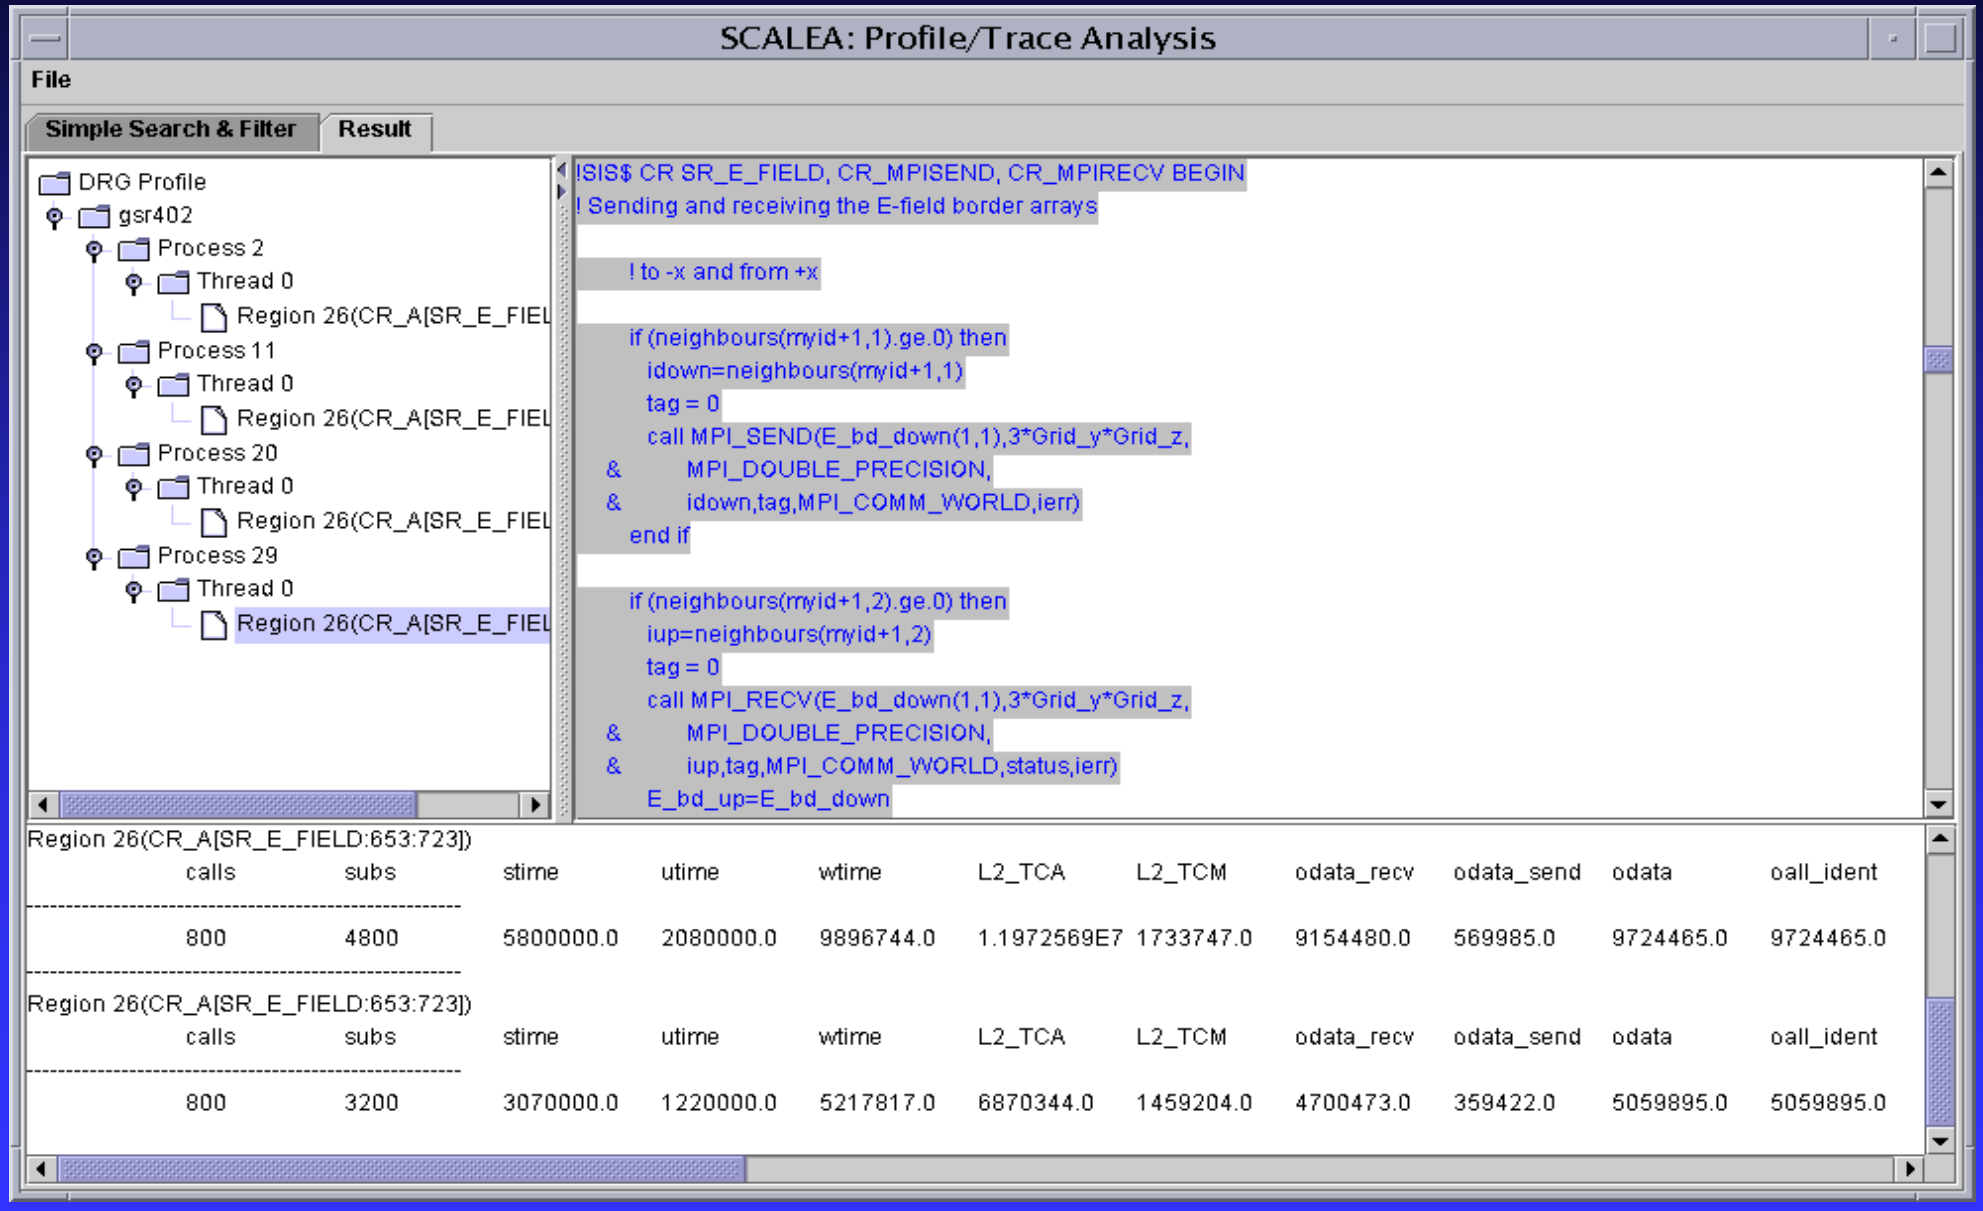

Search and Filter Capabilities (const.)

- Defining an expression of performance metrics
- Discretizing the expression into quantitative characteristics
- Search based on selected quantitative characteristics

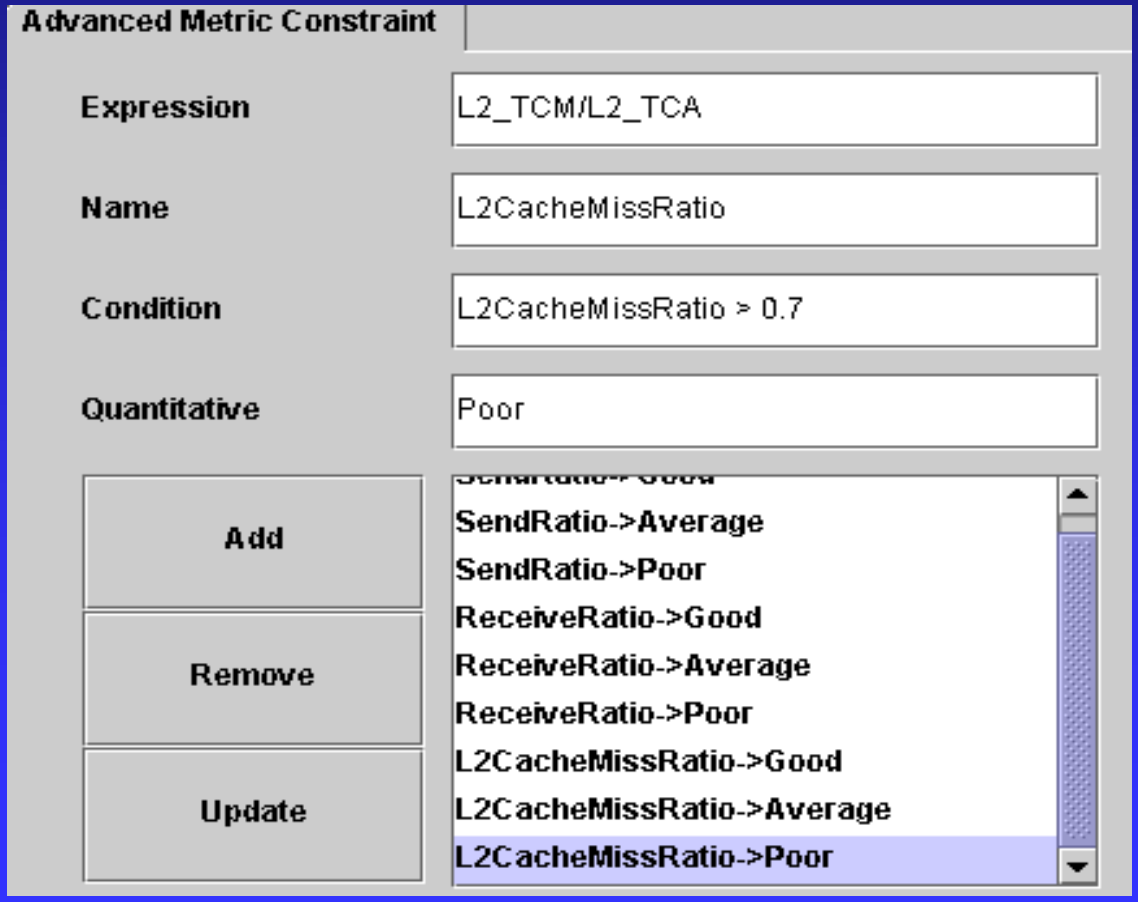

Multi-Experiment Analysis

- Many ways to specify what you want to analysis
  - Experiments
  - Code regions
  - Performance metrics
  - Statistic methods
- Various analyses
  - Performance comparison for different sets of experiments
  - Overhead analysis for multi-experiment
  - Performance speedup/improvement at both program and code region level

Select what to be analyzed

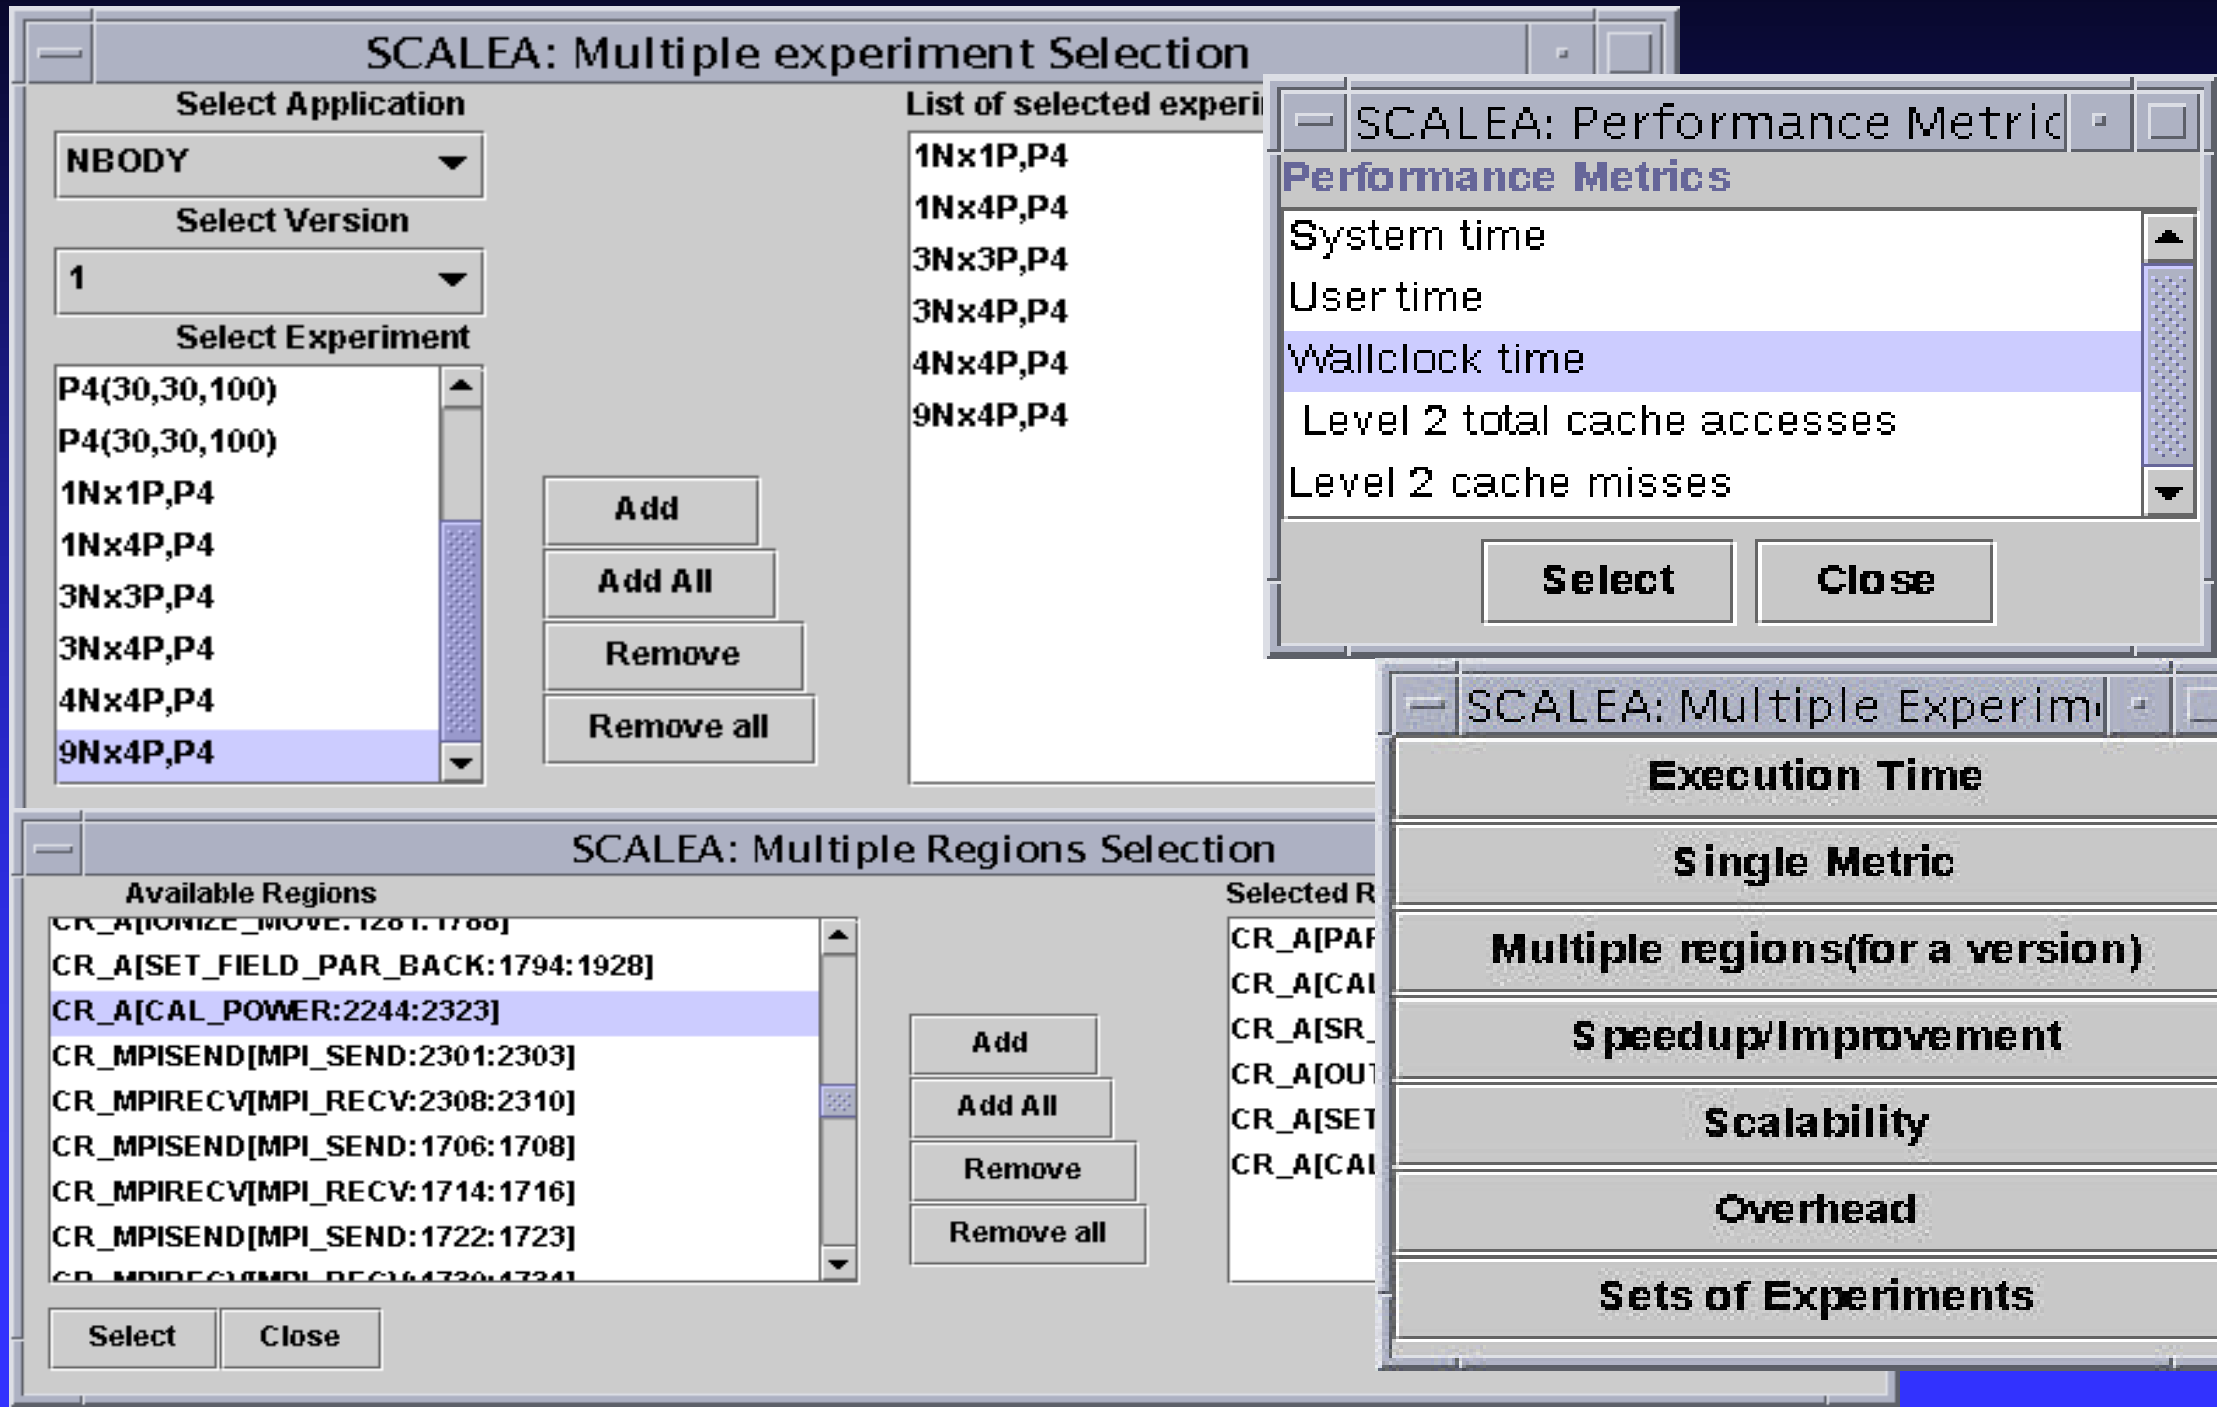

Different Sets of Experiments

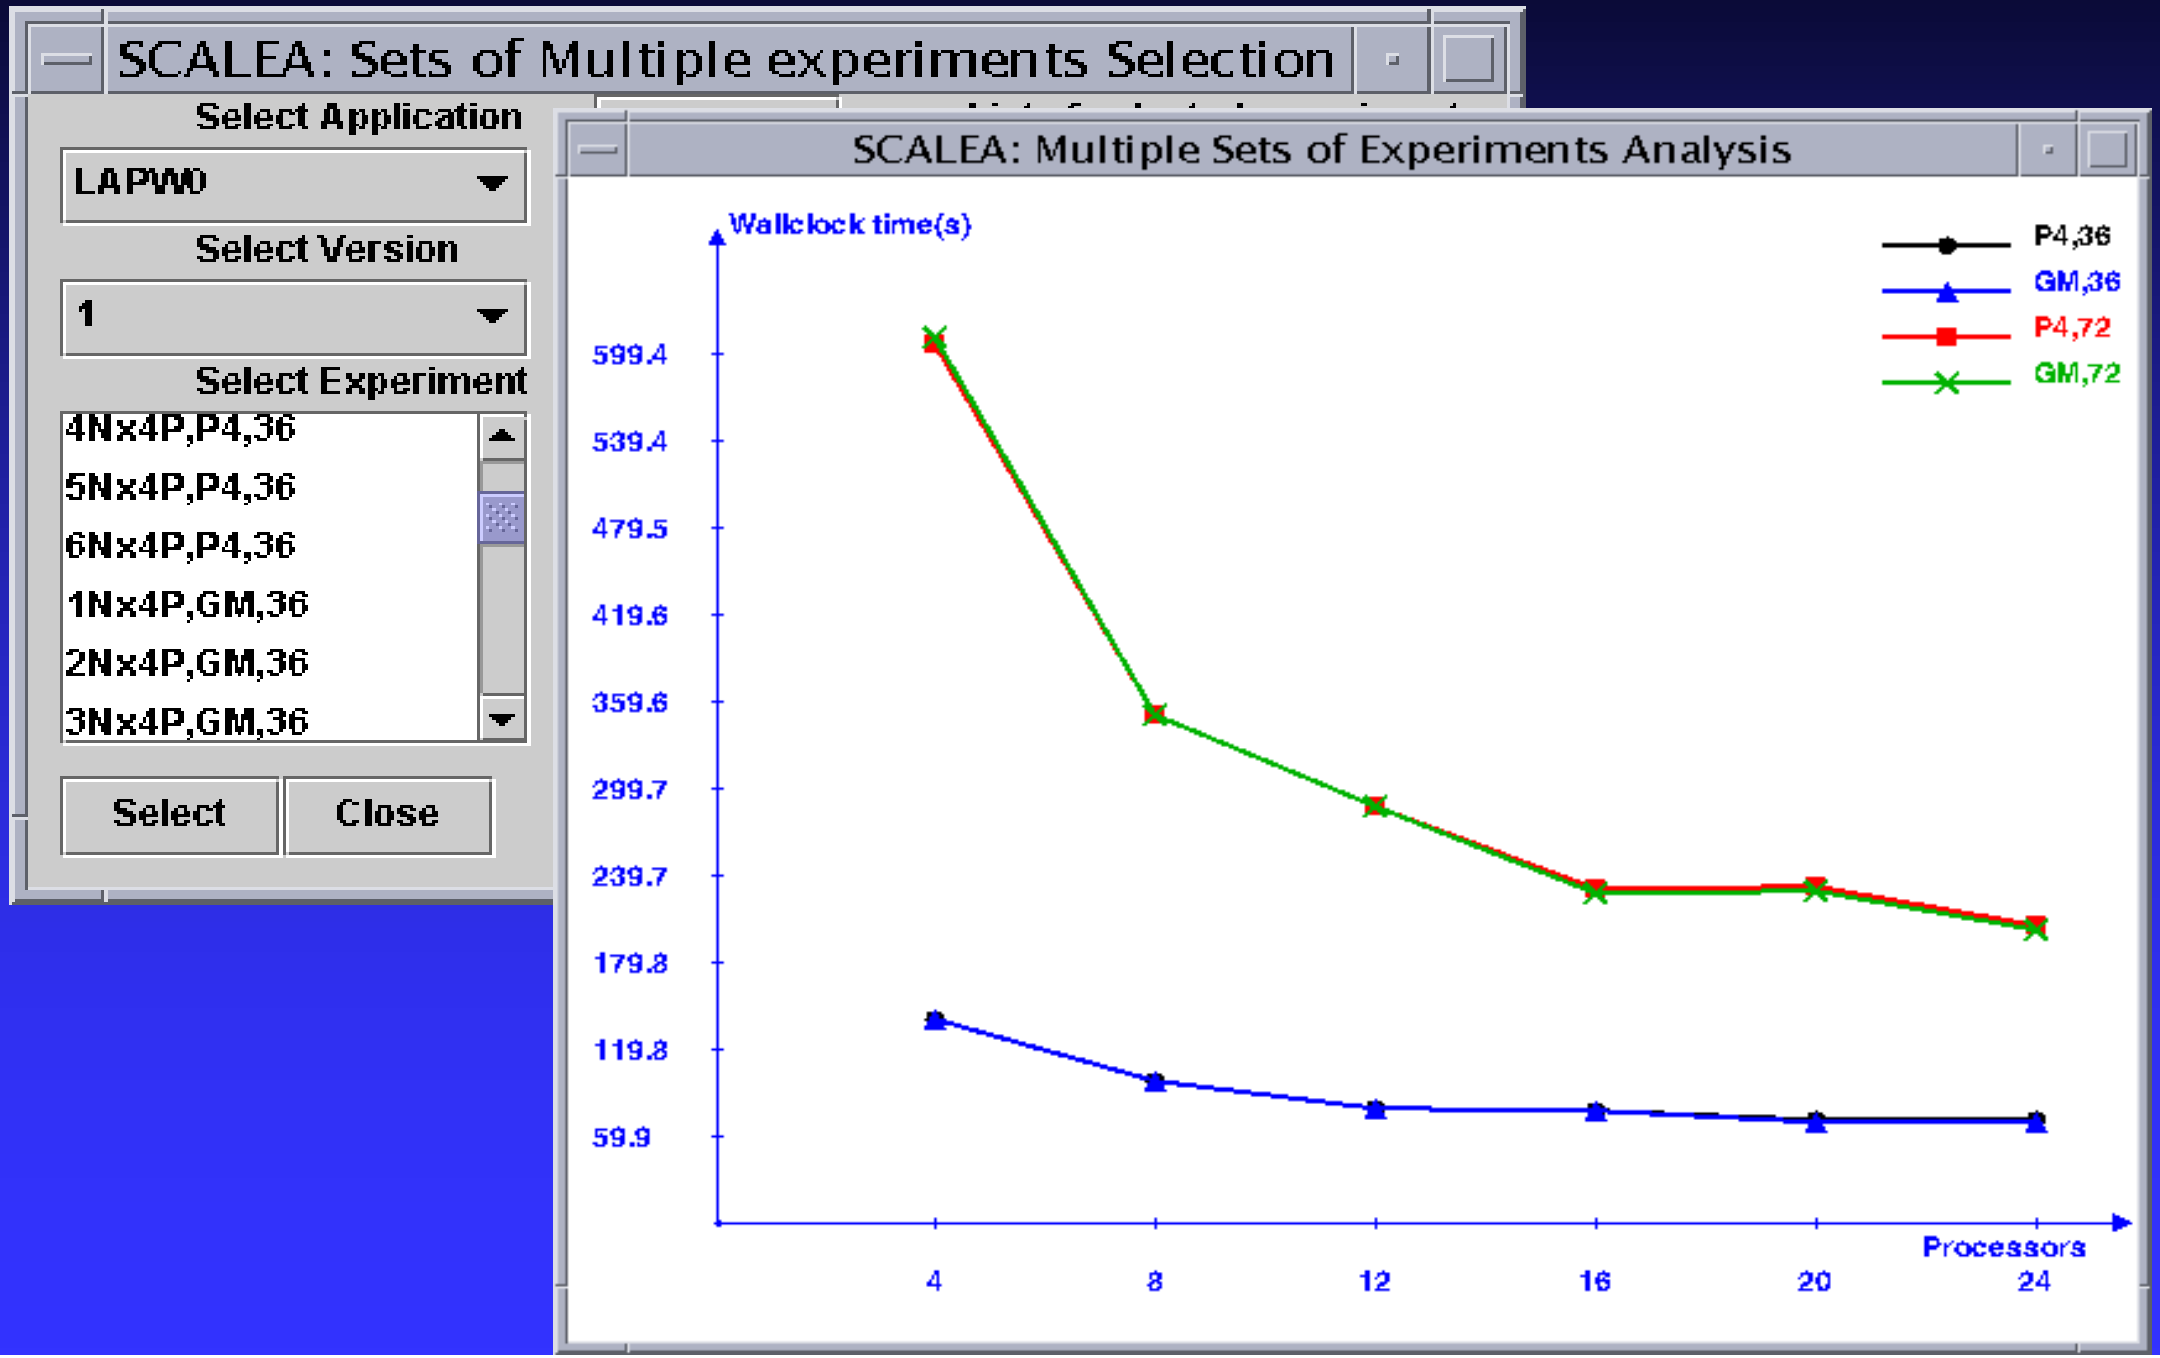

Improvement/Speedup, Efficiency

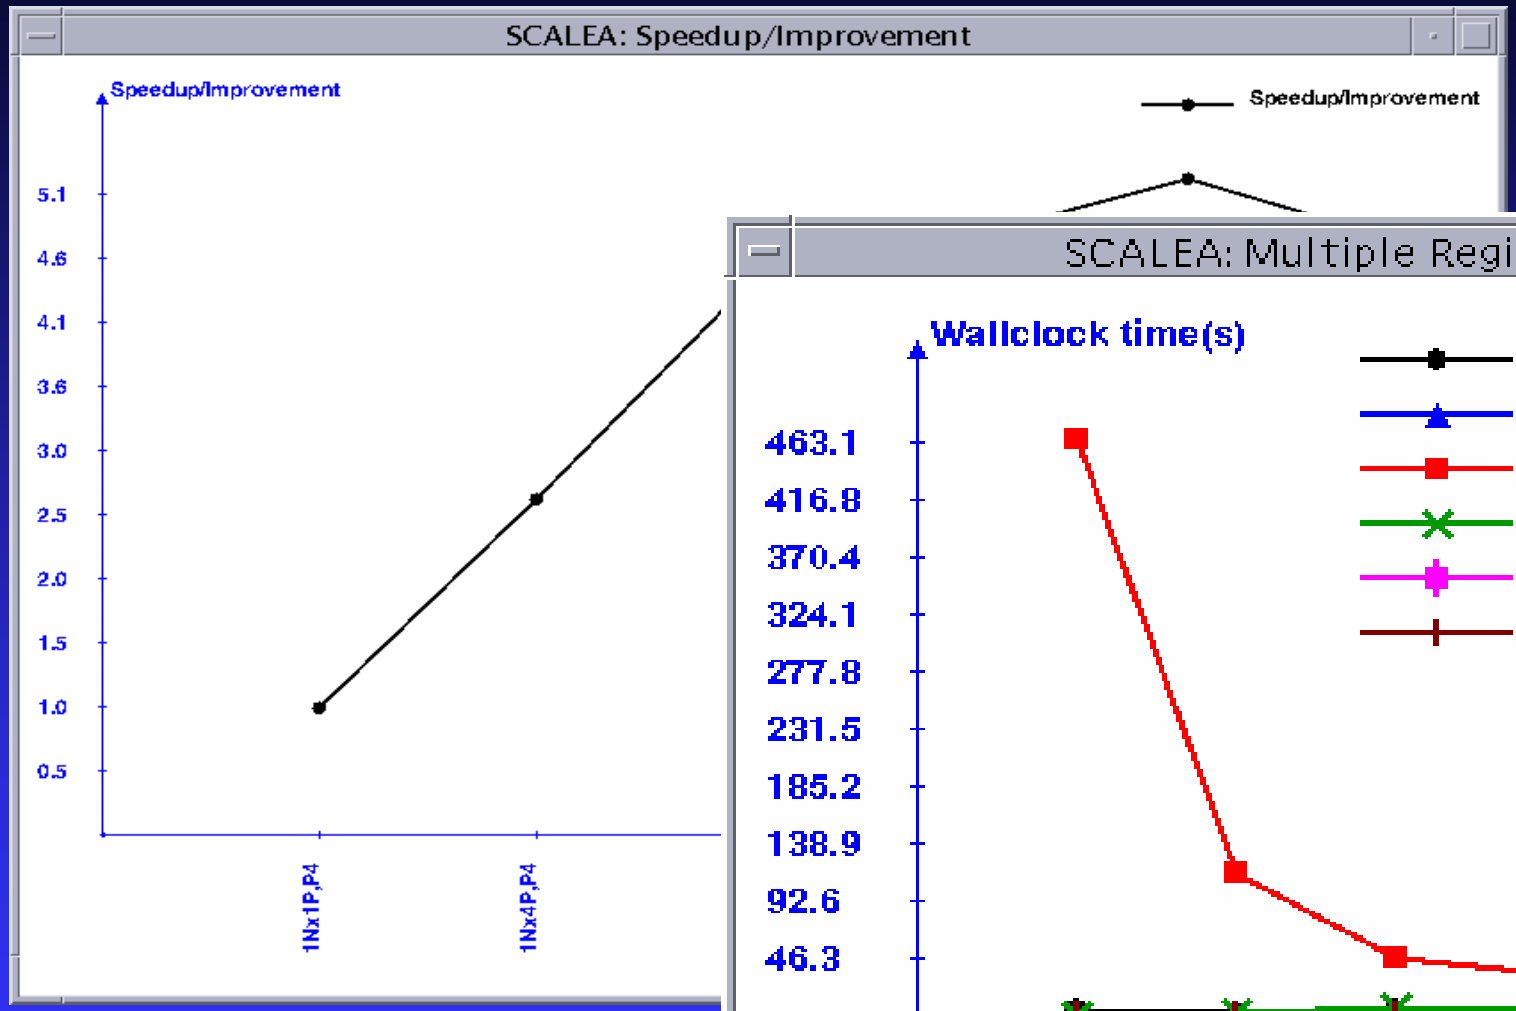

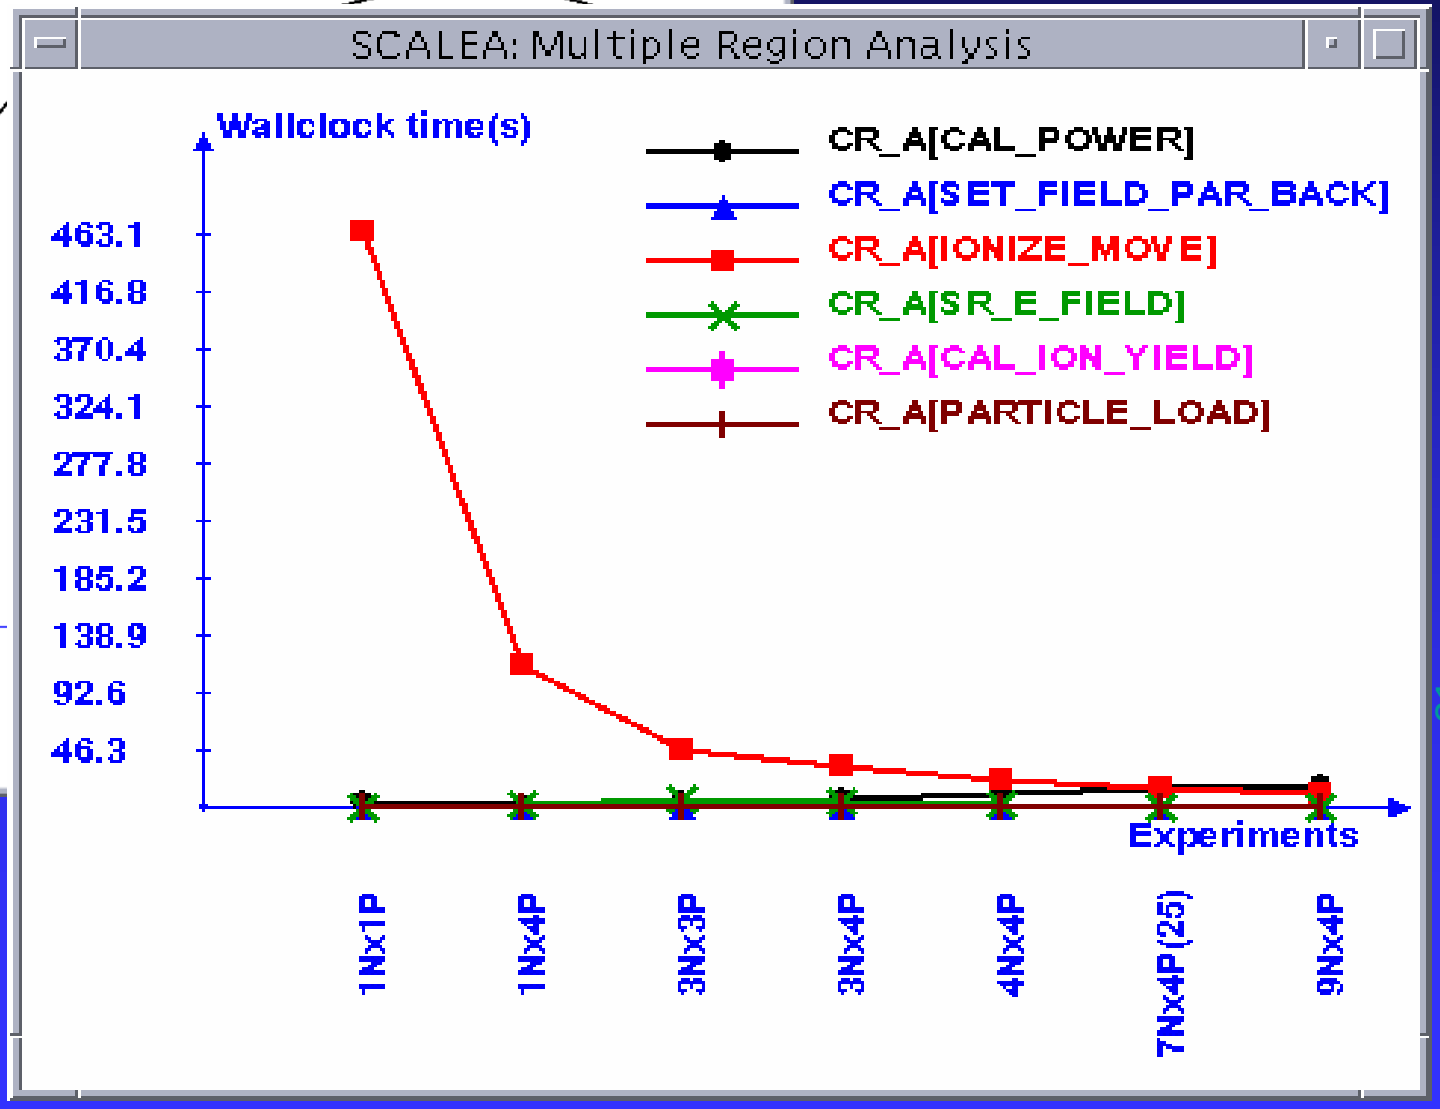

System Information

SCALEA: Network Selection

Cluster

gescher

Network

FastEthernet

Select | Close

SCALEA: Network Analysis

MPI Barrier Synchronization

MPI Broadcast

MPI P2P Blocking Send/Receive

MPI P2P Nonblocking Send/Receive

MPI P2P Syncronous Ssend/Irecv

SCALEA: MPI Broadcast

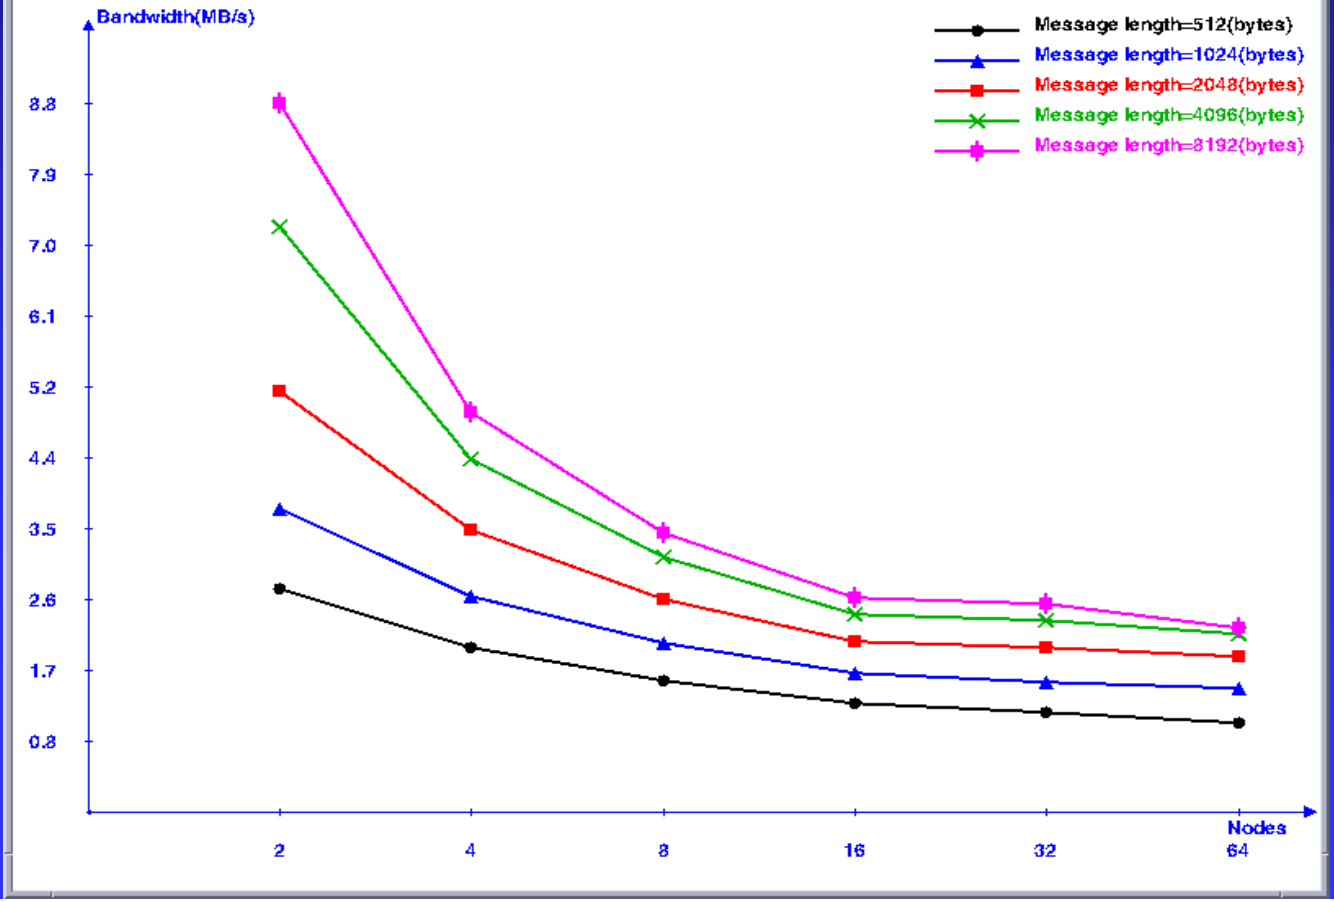

SCALEA: Host Selection

Cluster

gescher

Host

gsr402

Select | Close

SCALEA: Computational Node

Hostname	gsr402
Host Aliases	gsr402
Host Addresses	192.168.184.2
System Model	Intel 698 MHz Pentium III
Physical Memory	896 MB
Virtual Memory	4.0 GB
Number of CPUs	4
Cpu Type	Pentium III
Cpu Speed	698 MHz
Os Name	Linux
Os Version	2.4.19-PMC-SMP

Close

Tools Integration

http://www.par.univie.ac.at/project/askalon/

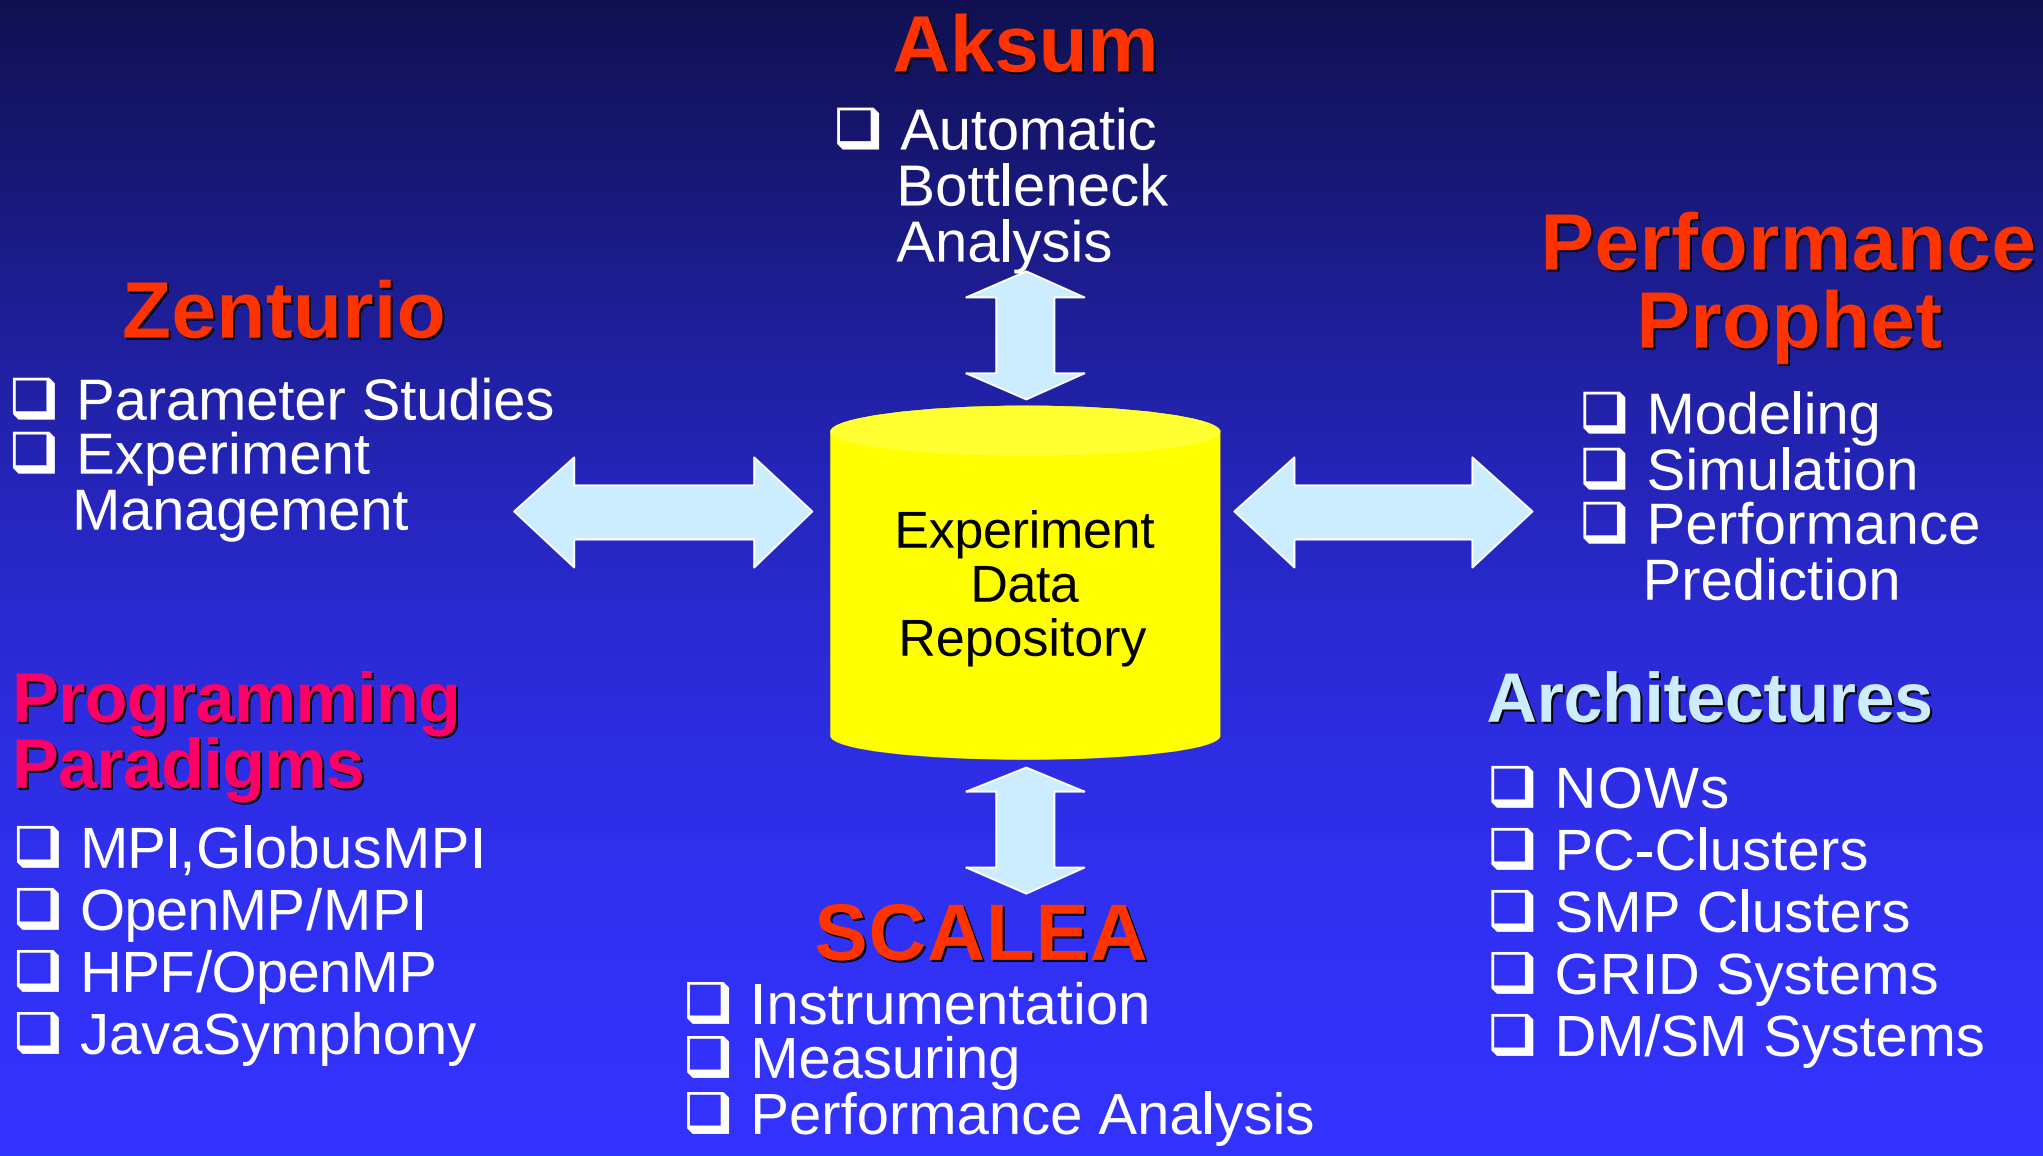

Case 1: AKSUM

- URL: http://www.par.univie.ac.at/project/aksum/
- Fahringer and Seragiotto, Jr.

AKSUM employs SCALEA to:

• instrument files given an arbitrary code region

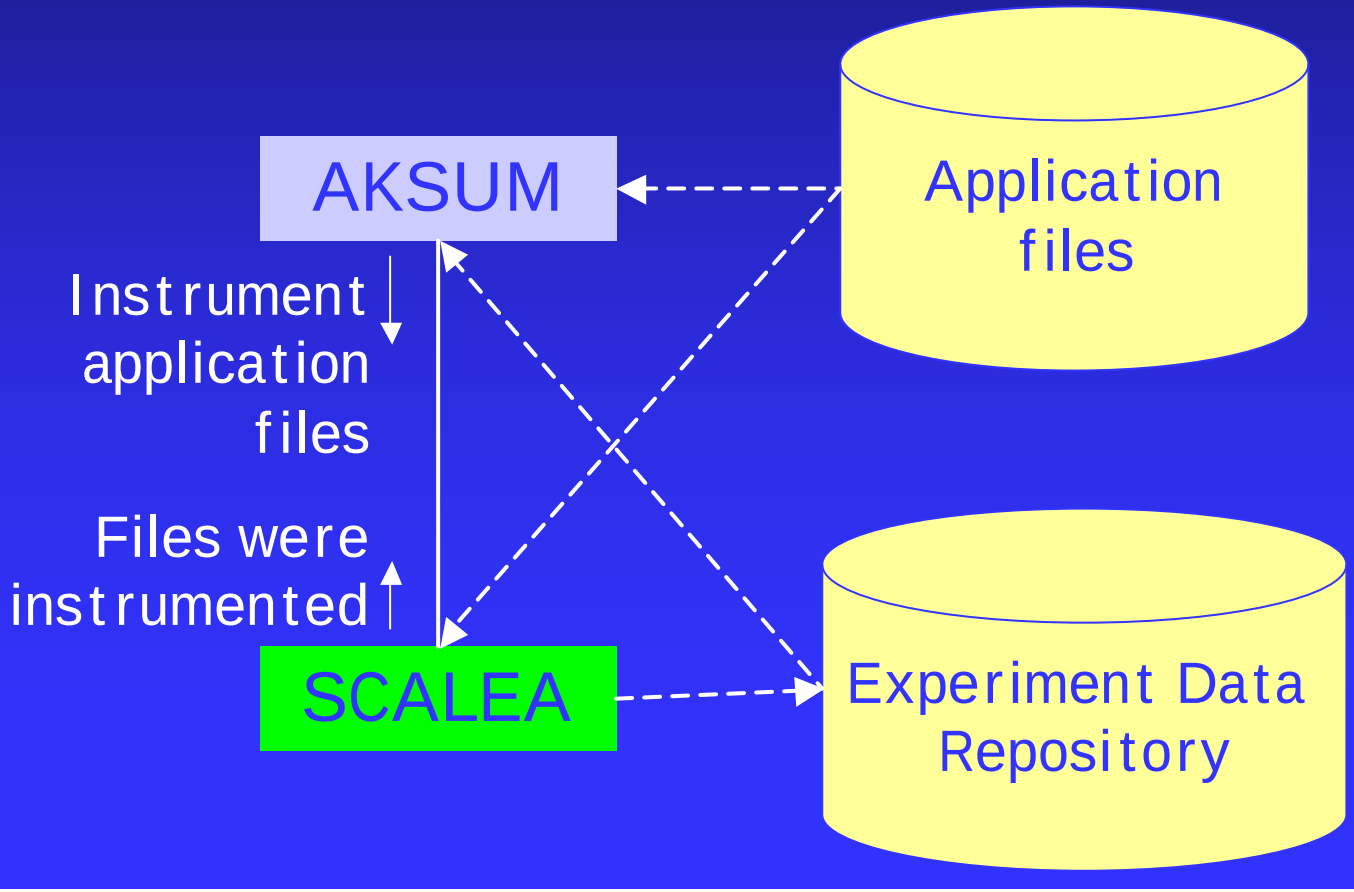

• transfer application files to the repository

• transfer the data generated by the monitoring to the repository

•use data provided by SCALEA for property analysis

Case 2: Performance Prophet

- URL: http://www.par.univie.ac.at/project/prophet/
- Fahringer, Pllana, and Testori
- PP is a performance modeling and prediction system
- PP utilizes SCALEA to obtain timing parameters for application.
  - PP uses performance data in XML format exported by SCALEA for automatic building of cost functions
  - Cost functions are used to develop a hybrid analytical and simulation model of the application

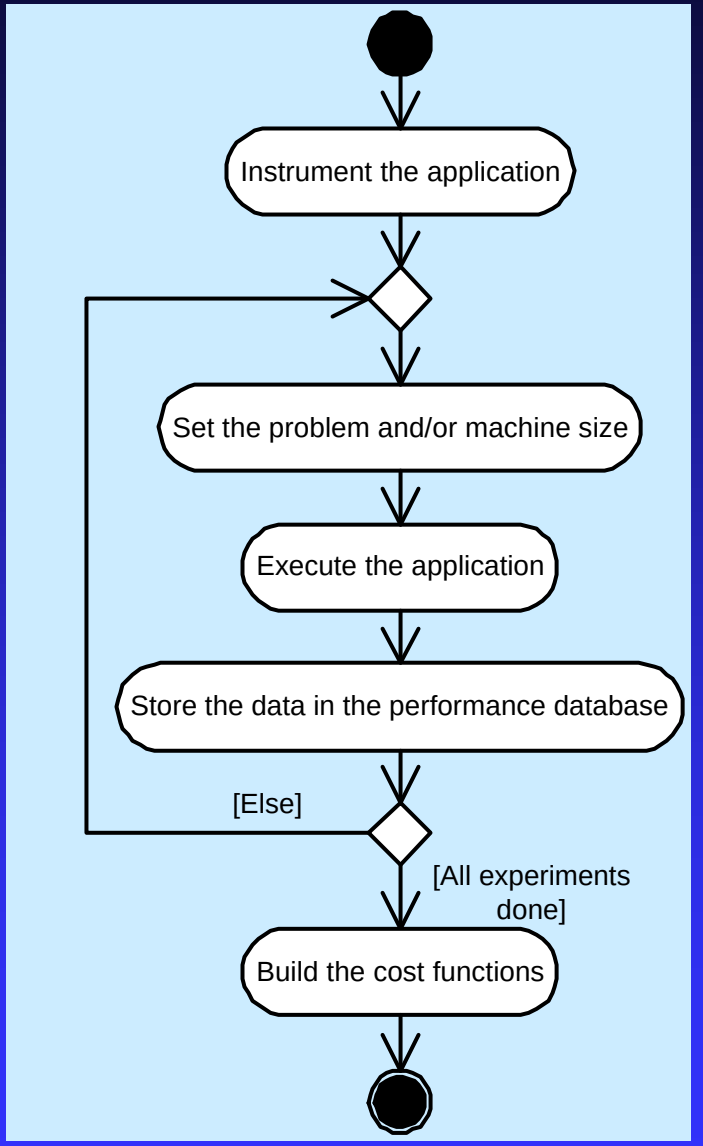

Conclusions and Future Work

- Conclusions
  - Design of experiment data repository for performance analysis tool
  - Demonstration of achievements gaining when employing experiment data repository

  → Data repository has increasingly supported the automation of performance analysis and optimization process

- Ongoing work
  - Working on simple and efficient way to search on performance data
  - Applying automatic scalable analysis techniques
  - Semantic representation of performance data

www.par.univie.ac.at/project/scalea