

# RESOURCE USE PATTERN ANALYSIS FOR OPPORTUNISTIC GRIDS

Marcelo Finger   Germano C. Bezerra   Danilo R. Conde

Department of Computer Science (IME-USP)  
University of São Paulo

Supported by CNPq/Brazil project 550895/2007-8.

# TOPICS

- 1 OPPORTUNISTIC GRIDS AND SCHEDULING
- 2 THE UPA METHOD
- 3 SIMULATION
- 4 IMPLEMENTATION
- 5 EXPERIMENTS
- 6 RELATED WORK
- 7 CONCLUSION

# TOPICS

- 1 OPPORTUNISTIC GRIDS AND SCHEDULING
- 2 The UPA Method
- 3 Simulation
- 4 Implementation
- 5 Experiments
- 6 Related Work
- 7 Conclusion

# GRIDS AND OPPORTUNISM

- Applications for Grid Computing
  - computationally intensive
  - distributed
  - heterogeneous environments

# GRIDS AND OPPORTUNISM

- Applications for Grid Computing
  - computationally intensive
  - distributed
  - heterogeneous environments
- Opportunistic Grid Computing
  - Idle time of machines

# GRIDS AND OPPORTUNISM

- Applications for Grid Computing
  - computationally intensive
  - distributed
  - heterogeneous environments
- Opportunistic Grid Computing
  - Idle time of machines
  - High-performance computation

# GRIDS AND OPPORTUNISM

- Applications for Grid Computing
  - computationally intensive
  - distributed
  - heterogeneous environments
- Opportunistic Grid Computing
  - Idle time of machines
  - High-performance computation
  - Resource-owners have to give permission

# GRIDS AND OPPORTUNISM

- Applications for Grid Computing
  - computationally intensive
  - distributed
  - heterogeneous environments
- Opportunistic Grid Computing
  - Idle time of machines
  - High-performance computation
  - Resource-owners have to give permission
  - QoS must remain high



# GRIDS AND OPPORTUNISM

- Applications for Grid Computing
  - computationally intensive
  - distributed
  - heterogeneous environments
- Opportunistic Grid Computing
  - Idle time of machines
  - High-performance computation
  - Resource-owners have to give permission
  - QoS must remain high
- InteGrade: opportunistic grid infrastructure

# EFFECTIVE OPPORTUNISTIC COMPUTING

- Desirable: prediction of resource availability
- Prediction available for grid scheduler
- The better the prediction, the lower impact on QoS

# EFFECTIVE OPPORTUNISTIC COMPUTING

- Desirable: prediction of resource availability
- Prediction available for grid scheduler
- The better the prediction, the lower impact on QoS
- Proposed Solution:

Resource Use Pattern Analysis

# UPA: (RESOURCE) USE PATTERN ANALYSIS

- Consists of:
  - Detecting the local **use pattern** of each **resource** at each **machine** in the grid

# UPA: (RESOURCE) USE PATTERN ANALYSIS

- Consists of:
  - Detecting the local **use pattern** of each **resource** at each **machine** in the grid
- Basic Hypothesis:
  - Each **resource** has a (temporal) **“pattern”** of use

# GOAL

- To develop a **method**

# GOAL

- To develop a **method**
- that **automatically** performs resource use pattern analysis

# GOAL

- To develop a **method**
- that **automatically** performs resource use pattern analysis
- for machines belonging to **opportunistic grids**



# STRATEGY


- Discover the prototypical patterns of use (off-line)
  - Unsupervised machine learning
  - Clustering analysis
- Runtime prediction
  - Comparing prototypical with “current” pattern of use
- Development method:
  - Simulation: parameter setting
  - Implementation: LUPA module for InteGrade grid

# TOPICS

- 1 Opportunistic Grids and Scheduling
- 2 THE UPA METHOD**
- 3 Simulation
- 4 Implementation
- 5 Experiments
- 6 Related Work
- 7 Conclusion

# RESOURCE USE OBJECTS


- A resource use object:



- Sampled at every 5 min
- Span of 48h used for prediction
- Resources: **CPU use, available RAM**, disk space, swap space, network and disk I/O

# RESOURCE USE OBJECTS

- A resource use object:



- Sampled at every 5 min
- Span of 48h used for prediction
- Resources: **CPU use**, **available RAM**, disk space, swap space, network and disk I/O
- Objects represent **availability** (not only use)

# THE UPA METHOD

- Resource Use Pattern Analysis
  - Unsupervised machine learning
  - Obtain fixed number of **use classes**

# THE UPA METHOD

- Resource Use Pattern Analysis
  - Unsupervised machine learning
  - Obtain fixed number of **use classes**
    - A class represents a frequent use pattern
    - E.g. busy work day, light work day, holiday, etc
  - Each class represented by a **prototypical** object

# THE UPA METHOD

- Resource Use Pattern Analysis
  - Unsupervised machine learning
  - Obtain fixed number of **use classes**
    - A class represents a frequent use pattern
    - E.g. busy work day, light work day, holiday, etc
  - Each class represented by a **prototypical** object
- Two phases:
  - training/learning phase: off-line
  - execution/prediction phase: runtime

# OFF-LINE LEARNING

- Inputs a large amount of objects
- Collected from regular operation
- Clustering is applied to training data
- Reliability depends on amount of data
  - At least 60 objects, or 2 months of data.



# OFF-LINE LEARNING

- Inputs a large amount of objects
- Collected from regular operation
- Clustering is applied to training data
- Reliability depends on amount of data
  - At least 60 objects, or 2 months of data.
- Several **parameters** have to be set

# LEARNING PARAMETERS

- Number of clusters: 5, 10

# LEARNING PARAMETERS

- Number of clusters: 5, 10
- Data normalisation: no normalisation, variational

# LEARNING PARAMETERS

- Number of clusters: 5, 10
- Data normalisation: no normalisation, variational
- Computation of prototypical element: centroid, centre

# LEARNING PARAMETERS

- Number of clusters: 5, 10
- Data normalisation: no normalisation, variational
- Computation of prototypical element: centroid, centre
- Similarity measurement: Euclidean distance between points.
  - Between clusters: single linkage, complete linkage, centroid method, Ward's method

# LEARNING PARAMETERS

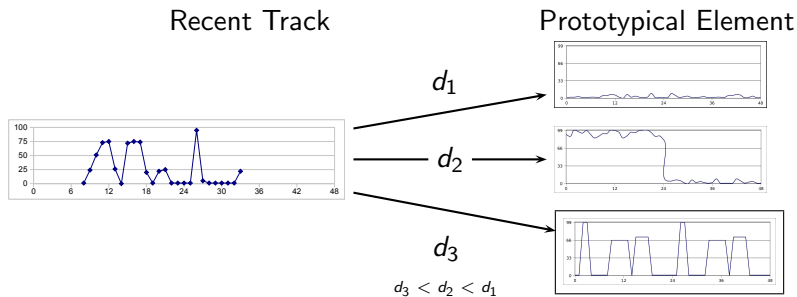
- Number of clusters: 5, 10
- Data normalisation: no normalisation, variational
- Computation of prototypical element: centroid, centre
- Similarity measurement: Euclidean distance between points.
  - Between clusters: single linkage, complete linkage, centroid method, Ward's method
- Clustering algorithms: hierarchical, sequential,  $k$ -means, etc.

# LEARNING PARAMETERS

- Number of clusters: 5, 10
- Data normalisation: no normalisation, variational
- Computation of prototypical element: centroid, centre
- Similarity measurement: Euclidean distance between points.
  - Between clusters: single linkage, complete linkage, centroid method, Ward's method
- Clustering algorithms: hierarchical, sequential,  $k$ -means, etc.

Simulation was performed to choose parameters


# RUNTIME PREDICTION






# RUNTIME PREDICTION

Recent Track



Prototypical Element



# PREDICTIONS

- Recent 24h record is compared against 48h use classes
- The closest class is the **current use class**.
- Predicted future in current use class: **[now-48h]**
- Request at 6am can predict 18 hours
- Request at 6pm can predict 6 hours (only)

# PREDICTIONS

- Recent 24h record is compared against 48h use classes
- The closest class is the **current use class**.
- Predicted future in current use class: **[now-48h]**
- Request at 6am can predict 18 hours
- Request at 6pm can predict 6 hours (only)
- Options for longer predictions:
  - Use less than 24 of current record.

# PREDICTIONS

- Recent 24h record is compared against 48h use classes
- The closest class is the **current use class**.
- Predicted future in current use class: **[now-48h]**
- Request at 6am can predict 18 hours
- Request at 6pm can predict 6 hours (only)
- Options for longer predictions:
  - Use less than 24 of current record. Predicts  $< 48h$
  - Chain of predictions (not implemented yet)

# TOPICS


- 1 Opportunistic Grids and Scheduling
- 2 The UPA Method
- 3 SIMULATION**
- 4 Implementation
- 5 Experiments
- 6 Related Work
- 7 Conclusion

# TRACE-DRIVEN SIMULATION

- Data collected for CPU and RAM use, 120 days
- Linux machines with very different types of users:
  - 1 a general purpose machine with more than 30 users;
  - 2 a single user machine;
  - 3 a general purpose machine with 6 users;
  - 4 a multi-user machine employed for testing heavy computational linguistics programs.
- Several parameters compared

# SIMULATION RESULTS

Correct predictions above 75% in all parameter combinations



# PARAMETERS LEARNED

- Number of clusters: 5, 10
- Data normalisation: no normalisation, variational
- Computation of prototypical element: centroid, centre
- Similarity measurement: Distance between clusters: points.
  - single linkage, complete linkage, centroid method, Ward's method
- Clustering algorithms: hierarchical, sequential,  $k$ -means, etc.



# PARAMETERS LEARNED

- Number of clusters: **5**, 10
- Data normalisation: no normalisation, **variational**
- Computation of prototypical element: **centroid**, centre
- Similarity measurement: Distance between clusters: points.
  - single linkage, complete linkage, **centroid method**, Ward's method
- Clustering algorithms: **hierarchical**, sequential, *k*-means, etc.

# PARAMETERS LEARNED

- Number of clusters: **5**, 10
- Data normalisation: no normalisation, **variational**
- Computation of prototypical element: **centroid**, centre
- Similarity measurement: Distance between clusters: points.
  - single linkage, complete linkage, **centroid method**, Ward's method
- Clustering algorithms: **hierarchical**, sequential, *k*-means, etc.

**Simulation validated UPA hypotheses** on the use of past patterns to predict future behaviours

# TOPICS

- 1 Opportunistic Grids and Scheduling
- 2 The UPA Method
- 3 Simulation
- 4 IMPLEMENTATION**
- 5 Experiments
- 6 Related Work
- 7 Conclusion

# IMPLEMENTATION OBJECTIVES

- Explore ways the scheduling of grid applications using the UPA method

# IMPLEMENTATION OBJECTIVES

- Explore ways the scheduling of grid applications using the UPA method
- Compare UPA scheduling with other methods

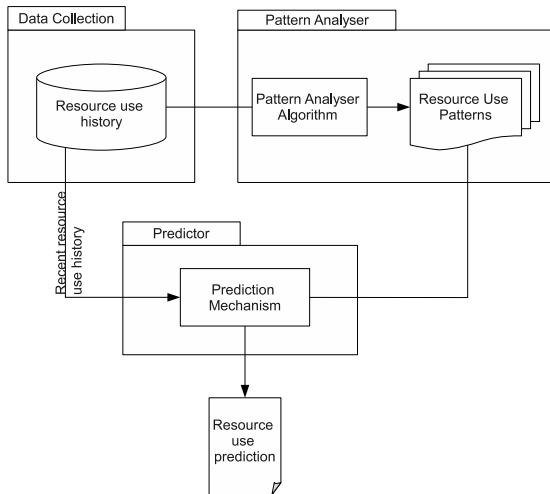
# IMPLEMENTATION OBJECTIVES

- Explore ways the scheduling of grid applications using the UPA method
- Compare UPA scheduling with other methods
- Identify initialisation strategies
  - because UPA needs a few months of resource use data collection to be reliable

# DESIGN DECISIONS

- UPA data analysed locally, not centrally
  - Advantages: privacy, no network traffic due to UPA processing
- Local resource Use Pattern Analyser (LUPA) module installed in all grid machines
- LUPA module is part of InteGrade middleware
- C++, on several Linuxes

# LUPA ARCHITECTURE





# LUPA SUBMODULES

- **Data Collection**: resource use sampling.
- **Pattern Analyser**: off-line clustering.  
Patterns are recomputed once a day.
- **Predictor**: runtime predictions. Interface

```
double[] getPrediction(resource r, int hours);
```

returns a vector of values representing the `r`-use prediction for the next `hours`, in 5-minute intervals.

LUPA is **not** yet integrated with the scheduler.

# TOPICS

- 1 Opportunistic Grids and Scheduling
- 2 The UPA Method
- 3 Simulation
- 4 Implementation
- 5 EXPERIMENTS**
- 6 Related Work
- 7 Conclusion

# EXPERIMENT GOALS

Compare UPA scheduling with other methods

# EXPERIMENT GOALS

Compare UPA scheduling with other methods

Scheduling methods for choosing  $n$  machines for  $h$  hours:

# EXPERIMENT GOALS

Compare UPA scheduling with other methods

Scheduling methods for choosing  $n$  machines for  $h$  hours:

- **RR**: round robin. No prediction

# EXPERIMENT GOALS

Compare UPA scheduling with other methods

Scheduling methods for choosing  $n$  machines for  $h$  hours:

- **RR**: round robin. No prediction
- **last4**: Predicts to the future average of last 4h.

# EXPERIMENT GOALS

Compare UPA scheduling with other methods

Scheduling methods for choosing  $n$  machines for  $h$  hours:

- **RR**: round robin. No prediction
- **last4**: Predicts to the future average of last 4h.
- **UPA**

# EXPERIMENT GOALS

Compare UPA scheduling with other methods

Scheduling methods for choosing  $n$  machines for  $h$  hours:

- **RR**: round robin. No prediction
- **last4**: Predicts to the future average of last 4h.
- **UPA**

Predicting methods choose  $n$  machines with highest prediction of CPU availability for the next  $h$  hours




## EXPERIMENT DESCRIPTION

- 15 machines, logs of 41–120 days
- Sequence of tests for a fixed set of parameters:  $n$ ,  $h$
- Instant tests (1 experiment  $\approx$  527 tests)
  - Choose a day with  $m$  valid machines
  - 24 tests are executed for 3 methods, each for an hour of the day

$$performance_s(t) = 1 - \frac{\sum_{i=1}^n use(m_i, t, h)}{n}$$

- $s$ : scheduling algorithm
- $n$ : number of chosen machines
- $m_i$ : machine chose to run for  $h$  hours starting at  $t$
- $use(m_i, t, h)$ : average CPU at  $m_i$  for  $[t, t + h]$

# UPA *versus* LAST4



# COMPARISON OF SCHEDULING METHODS

	<b>UPA</b>	<b>last4</b>
Performance	Similar to <b>last4</b> , but better	Similar to <b>UPA</b> , but worse
Impact of $\uparrow m/n$	$\uparrow$ performance	no effect
Impact of $\uparrow h$	no effect	no effect

Performance of **RR** consistently below both **last4** and **UPA**

# ADAPTATIVE UPA METHOD

- Use **UPA** if more than 21 days of data collection is available;  
else
- Use **last4** if more than 4 hours of data collection is available;  
else
- Predict that resource use at request time persists.

# OVERHEAD ANALYSIS

- Running time for pattern analysis was always below 1s
- Running time for prediction calls was always below 3ms

# OVERHEAD ANALYSIS

- Running time for pattern analysis was always below 1s
- Running time for prediction calls was always below 3ms
- Measurements made on notebook [AMD Turion 64 1.8GHz CPU, 1GB RAM running Kubuntu 7.10 (32 bits) Linux]
- Running pattern analysis once a day poses negligible overhead

# TOPICS

- 1 Opportunistic Grids and Scheduling
- 2 The UPA Method
- 3 Simulation
- 4 Implementation
- 5 Experiments
- 6 **RELATED WORK**
- 7 Conclusion

## CORRELATED AVAILABILITY

- BOINC: Volunteer computing
- Correlated Availability in Internet-Distributed Systems (Kondo, Andrzejak & Anderson, 2008)
- Computes patterns of simultaneous **CPU availability** via clustering



## CORRELATED AVAILABILITY

- BOINC: Volunteer computing
- Correlated Availability in Internet-Distributed Systems (Kondo, Andrzejak & Anderson, 2008)
- Computes patterns of simultaneous **CPU availability** via clustering
- Monitors 112,268 hosts
- Global monitoring and pattern discovery
- Several interesting clusters of machines discovered, with potential applications.
- Not originally intended for scheduling

# PREEMPTIVE RESUME SCHEDULING

[ROY AND LIVNY, 2003]

- Condor:
  - Combination of dedicated and opportunistic scheduling
  - Opportunistic scheduling based on checkpointing
  - Preemptive resume scheduling
  - No prediction

# PREEMPTIVE RESUME SCHEDULING

[ROY AND LIVNY, 2003]

- Condor:
  - Combination of dedicated and opportunistic scheduling
  - Opportunistic scheduling based on checkpointing
  - Preemptive resume scheduling
  - No prediction
  - **No need to predict job duration**

# PREDICTIVE SCHEDULING

[YANG, SCHOPF AND FOSTER, 2003]

- One-step-ahead CPU load prediction
  - Load Tendency Prediction
  - Increase/decrease adaptation processes
- Interval Load Prediction
  - Average over aggregated CPU load time series
- Load Variance Prediction
  - Standard deviation time series
  - Combined with one-step-ahead prediction
- Conservative scheduling: combines interval and variance load prediction

# TOPICS

- 1 Opportunistic Grids and Scheduling
- 2 The UPA Method
- 3 Simulation
- 4 Implementation
- 5 Experiments
- 6 Related Work
- 7 **CONCLUSION**

# CONCLUSIONS

- Some form of prediction is always preferable to no prediction
- UPA method compares favourably w.r.t. other methods
- Small overheads involved
- Method can be used for practical grid scheduling

# CONCLUSIONS

- Some form of prediction is always preferable to no prediction
- UPA method compares favourably w.r.t. other methods
- Small overheads involved
- Method can be used for practical grid scheduling
  - but needs task duration prediction

# ONGOING AND FUTURE WORKS

- Scheduler using LUPA information
- Automated learning of task duration prediction
- Long term predictions
- Preemptive task migration using UPA
- Automated “booking” of machine resources for future executions using UPA



[www.integrade.org.br](http://www.integrade.org.br)

[mfinger@ime.usp.br](mailto:mfinger@ime.usp.br)