



Performance of Resource Allocation Policies in Grid Systems

Stylianos Zikos
PhD Student
Department of Informatics
Aristotle University of Thessaloniki




Outline

- Grid systems
- A system model
- Scheduling jobs
- Service demand variability
- Performance metrics
- Current work
- Future work
- References



Grid systems



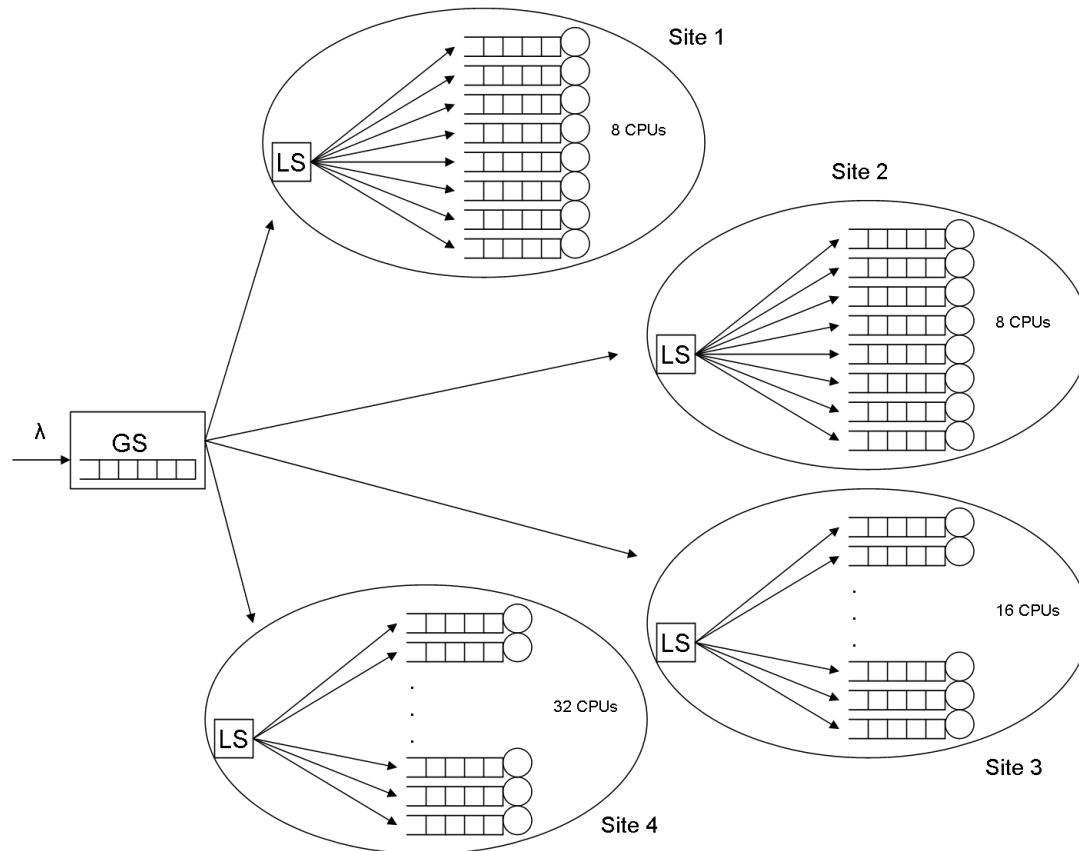
- **Grid characteristics:**
 - 1) distributed system
 - 2) heterogeneous resources
 - 3) no centralized control
- Main types: Computational Grids, Data Grids
We focus on computational grids

Grid systems




- A hierarchical architecture is very common for these systems
- For example, it may include a grid scheduler, various local schedulers and many resources
- All the above entities belong to different levels → job scheduling occurs at multiple levels
- The simulation models we created are 2-level models

A system model




- Open queueing network model of a 2-level grid
- Jobs arrive at the Grid Scheduler
- Heterogeneous sites (different number of processors)

A system model



- Discrete event simulation
- We implemented simulation applications in C.
- Other tools for discrete event simulation of grid systems: GridSim, GangSim.



Scheduling jobs

- Job scheduling is applied at two levels:

Grid level


Local level

Grid level: Grid Scheduler (GS) or Broker selects the appropriate sites for jobs

Local level: Local Schedulers (LS) allocate jobs to machines inside sites

Goals:

- 1) Jobs must be served as soon as possible
- 2) All the available grid resources must be utilized



Scheduling jobs

Site allocation policies at Grid Scheduler (nonclairvoyant)

- **Random policy:** selection probability proportional to a site's computational capacity
- **Deferred policy**
 - Based on dynamic site load information that the GS receives from the LSs, every specified time periods
 - The GS stores incoming jobs in the queue and dispatches them when information from sites is available
 - For each job, the site with the minimum load is selected
 - Load information: number of jobs, number of idle processors

Scheduling jobs



- **Hybrid policy**
 - Combines Random and Deferred
 - Jobs arriving shortly after an allocation event, are dispatched according to Random policy (their delay in GS queue is avoided)
 - In our experiments Hybrid outperformed both Random and Deferred
- We also evaluated the performance of two novel GS policies (PAD, FZF-PAD), similar to Hybrid
 - Reduce site load information traffic
 - Provide utilization fairness



Scheduling jobs


Resource allocation policies at Local Schedulers (nonclairvoyant)

- **Random**
 - Easy to implement
 - Poor performance
- **Shortest Queue**
 - LS monitors queue lengths
 - Great performance

Service demand variability




- Users with different demands → jobs with different characteristics
- In a case where most jobs demand similar times but there are some which demand much more time: the demanding jobs can cause delays to the rest
- A local policy which we called ‘Selective Random’ can be used when job service demands are not known
 - Upon an arrival, the processors serving a job for time over a threshold, are excluded from random selection
 - Tries to detect demanding jobs
 - Outperforms Random




Performance metrics

- **Average response time of jobs**
- **Maximum response time**
- **Average slowdown**
 - Slowdown of a job: The job's response time divided by its service time
- **Load information traffic:** Counts the load information events from sites to Grid Scheduler (lower means lower communication overhead)
- **Resource utilization fairness:** Number of jobs serviced by a site divided by the theoretical number of jobs that it should serve according to its computational capability
 - A value close to 1 for all sites indicates that job distribution is fair



Current work

- Clairvoyant scheduling
 - Service demands of jobs are known to schedulers
- High variability of job service demands
 - Service demand times follow the Pareto distribution
- Goal: To find policies which minimize the average slowdown and balance the load



Future work


- Additional criterion for scheduling jobs, not related with performance:
Energy efficiency
- Power consumers: Processors, disks, networks
- Challenge: Heterogeneous resources have different power management features
- Goal: To find out grid and local policies that decrease the consumed power, without affecting performance considerably

References



1. S. Zikos and H.D. Karatza, “**Resource Allocation Strategies in a 2-level Hierarchical Grid System**”. In the Proceedings of the 41th Annual Simulation Symposium (ANSS), IEEE Computer Society Press, SCS, April 13-16, 2008, Ottawa, Canada, pp. 157-174.
2. S. Zikos and H.D. Karatza, “**The Impact of Service Demand Variability on Resource Allocation Strategies in a Grid System**”. ACM Transactions on Modeling and Computer Simulation (TOMACS), to appear.
3. S. Zikos and H.D. Karatza, “**Communication Cost Effective Scheduling Policies of Nonclairvoyant Jobs with Load Balancing in a Grid**”. The Journal of Systems and Software, Elsevier, to appear.

Our team



- Coordinator
Prof. Helen Karatza <karatza@csd.auth.gr>
- Members
Stylianos Zikos <szikos@csd.auth.gr>
Zafeirios Papazachos <zpapazac@csd.auth.gr>
Fotis Loukos <fotisl@csd.auth.gr>
Konstantinos Karaoglanoglou <kkaraogl@csd.auth.gr>
Georgios Stavriniades <gstavrini@csd.auth.gr>