

# *Service Level Agreement Based Scheduling Heuristics*

**GGF9**

**Chicago, IL**

**October, 2003**



# Why?

**We (Manchester) are interested in scheduling: building sophisticated brokers and superschedulers for the Grid**

**We want to be able to schedule workflows of complex jobs onto compute resources**

**To do this, we need to honour dependences between parts of the workflow**

**So we need to know when jobs are going to run – sometimes precisely when, but mostly we're only interested in the bounds (soonest start, latest end time)**

**How do we achieve this, given the current state-of-the-art of compute resource schedulers, i.e. batch queue systems?**



# State of the Art

**Current schedulers offer two basic levels of service:**

Run this when it reaches the head of the queue



Run this at a precise time (advance reservation)

# What's the problem?

Any Grid job, e.g. a sequence of tasks in a workflow, can only be implemented well with advance reservation

Exceptions are restricted cases, e.g. where all tasks in the job can be sent to the same Condor pool

Often, this is a sledgehammer being used to crack a nut

The other problem is that advance reservation doesn't fit very well into the batch processing model

It causes expensive gaps in the schedule that can't be plugged

- Checkpointing takes time. Up to an hour for large machines
- Suspend/Resume quicker, but impacts available resources/performance of the AR job

Experience suggests utilisation will decrease rapidly (worse than linear) as the percentage of AR jobs in the job mix increases



# What's the solution?

**Clients (users, brokers, superschedulers) often know or could easily define constraints for soonest start, latest end time – which could form part of an SLA**

**Additional constraints might be based on performance when running (e.g. time per iteration) or cost.**

**But all of this information is being ignored, i.e. it is available, but not captured!!!**

***So the solution is to capture and use this information!!!***



# A new approach to Scheduling?

So whenever a job is submitted to our scheduler, there is a negotiation, where constraints are agreed and an SLA is formed.

The scheduler bases what to run next on these SLAs

Also, the current set of SLAs will determine what new SLAs the scheduler can commit to

But of course, some jobs will fall over, or be withdrawn so some overbooking might come in handy

And of course, if someone came along with a huge amount of cash, we might want to break a couple of smaller agreements...

The problem looks different now –  
**WHAT'S GOING ON?**



# Or an old approach to Scheduling?

The problem doesn't look like a batch scheduling problem any more.

It looks like a “traditional scheduling” optimisation problem, like timetable optimisation, or optimising workflow in a factory.

There is a community with well developed techniques for solving such problems.

It seems that at some point in the past “scheduling” split into these two separate disciplines.

Now it's time to re-unite them.

The Grid needs this community, and their techniques.



# So what next?

**We've joined forces with the ASAP group from Nottingham, one of the best "traditional scheduling" groups in the world.**

**They bring their fuzzy logic and heuristic techniques to the table; we bring our Grid scheduling expertise.**

**These people have a lot to teach us – let's get them on board!**

**So, taking a longer view, there is a workshop at NeSC later this month which aims to bring together people doing scheduling on the Grid, and Traditional Scheduling people.**





# And finally...

We're looking for a researcher at Manchester with Grid experience, and Nottingham are looking for a researcher in traditional scheduling.

Job ad: <http://www.man.ac.uk/news/vacancies/research.html#761>

Trad sched job: <http://www.asap.cs.nott.ac.uk/>

Collectively, NeSC and the Interdisciplinary Scheduling Network (ISN) are looking for people to attend the cross-community scheduling workshop on the 21<sup>st</sup> and 22<sup>nd</sup> of October, in Edinburgh, title “Open Issues in Grid Scheduling”

More info: <http://www.nesc.ac.uk/esi/events/309/>

