Grid IC

Major Review Meeting
18th September 2008

J. M. Dana,
X. Grehant,
S. Jarp

Contents

- Resource allocation (X. Grehant)
- Tycoon (J.M. Dana)

Contents

- Resource allocation (X. Grehant)
- Tycoon (J.M. Dana)

A model to quantify performance

- From microarchitecture to grids
  - Quantify accuracy of job/server matching
  - What is the best microarchitecture for the job?
  - What is the best memory hierarchy for the job?
- Cache misses prediction
  - The ratio of cache misses determines the cost of stall conditions

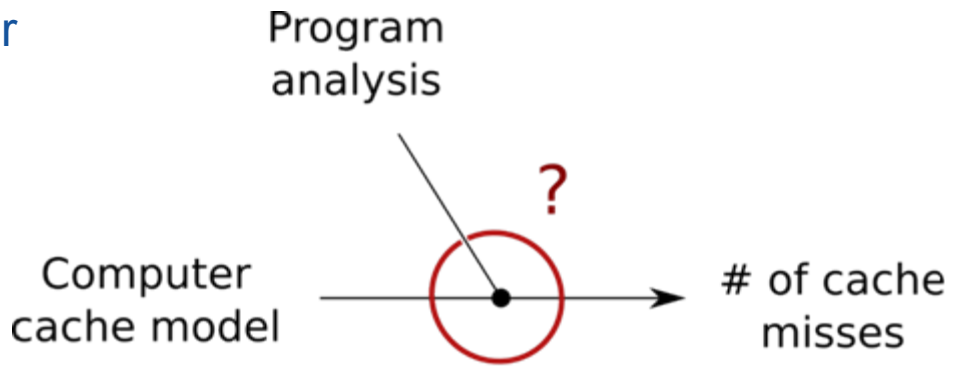

Cache misses prediction

- A metric and a method
  - New idea: statistical analysis of stack distance
  - New estimation algorithm: more accurate on high values
- Stack distance
  - Determines if a memory access is a cache miss
- Its probability distribution
  - Is estimated quickly
  - Let us do fast, cross-platform predictions

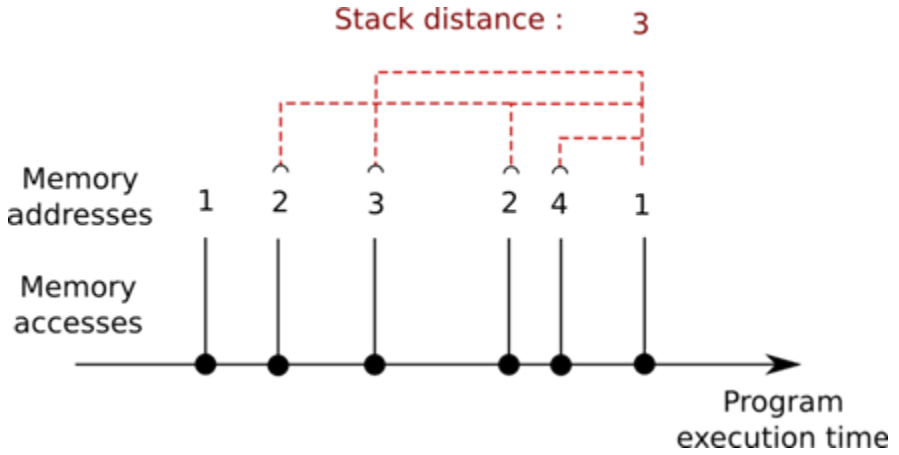

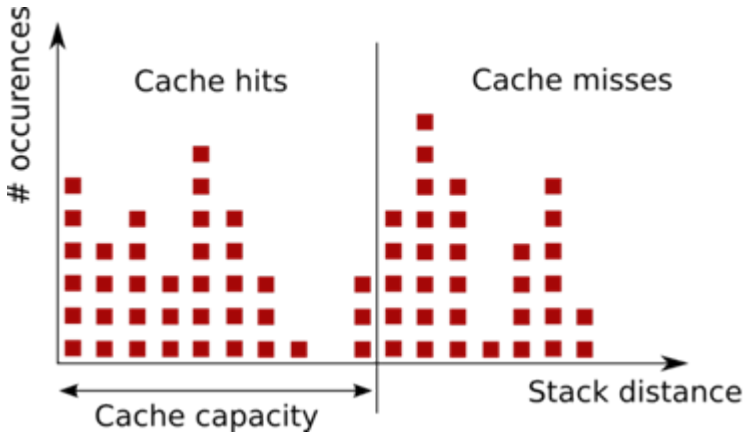

Cache misses prediction

- Results
  - Around 99% accuracy
  - Maximum 1 sample / 10 000 collected for program analysis
  - Program signature = distribution parameters: 2 floats
  - Instantaneous prediction: calculation of the cumulative density function of the distribution
- Status
  - Presented at PDPTA'08 in July
  - Code and data available at code.google.com/p/mtc-project
  - Now working on Cache thrashing prediction

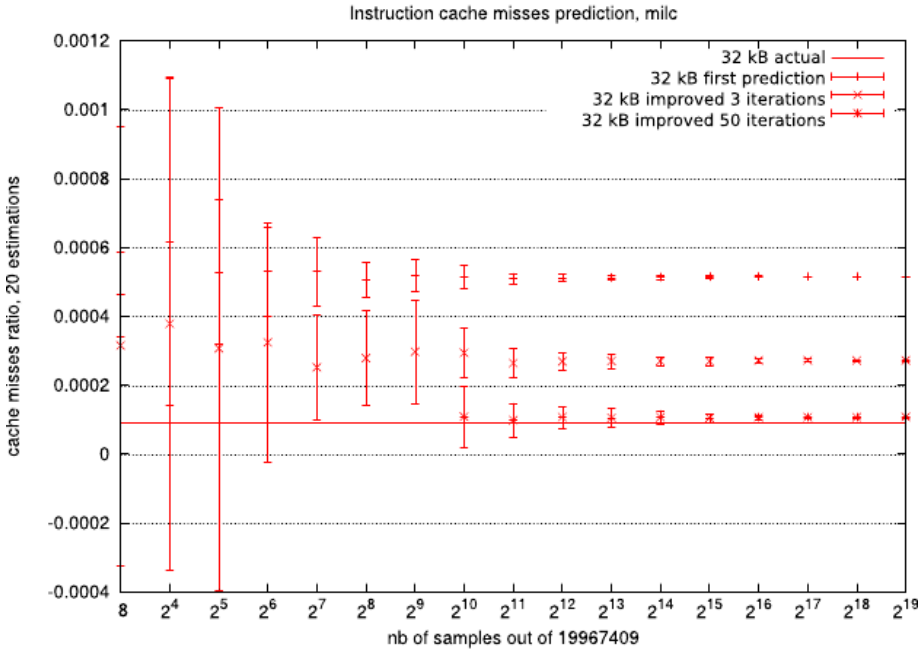

Contents

- Resource allocation (X. Grehant)
- Tycoon (J.M. Dana)

Distributed Computing Workshop

- Tycoon was presented in the Distributed Computing Workshop held in London (UK) the 21st May 2008

- One of the topics of the workshop was “cloud computing”

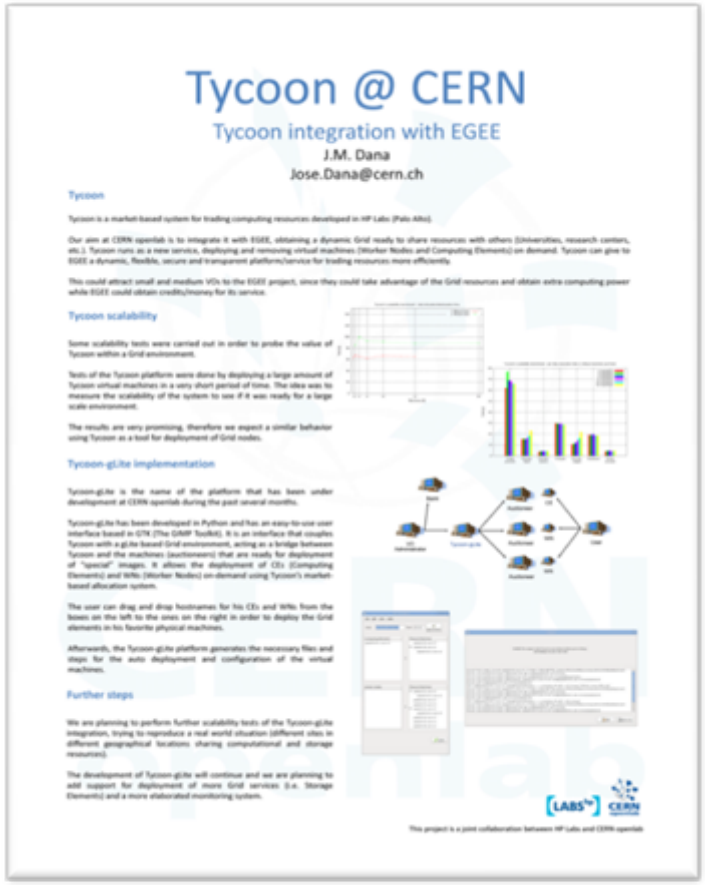

Tycoon @ CERN

Tycoon integration with EGEE

J.M. Dana

Jose.Dana@cern.ch

Tycoon

Tycoon scalability

Tycoon-gLite implementation

Further steps

The Open Cloud Computing initiative

- *Cloud Computing" is a style of computing where IT-related capabilities are provided "as a service", allowing users to access technology-enabled services "in the cloud" without knowledge of, expertise with, or control over the technology infrastructure that supports them*

- Cloud computing research testbed with 3 initial sponsors:
  - HP
  - Intel
  - Yahoo!

Tycoon and the Cloud Computing

- Tycoon is being used as a resource allocation system in “the cloud”
- We believe that our collaboration with HP Labs has helped Tycoon to be part of this new research initiative

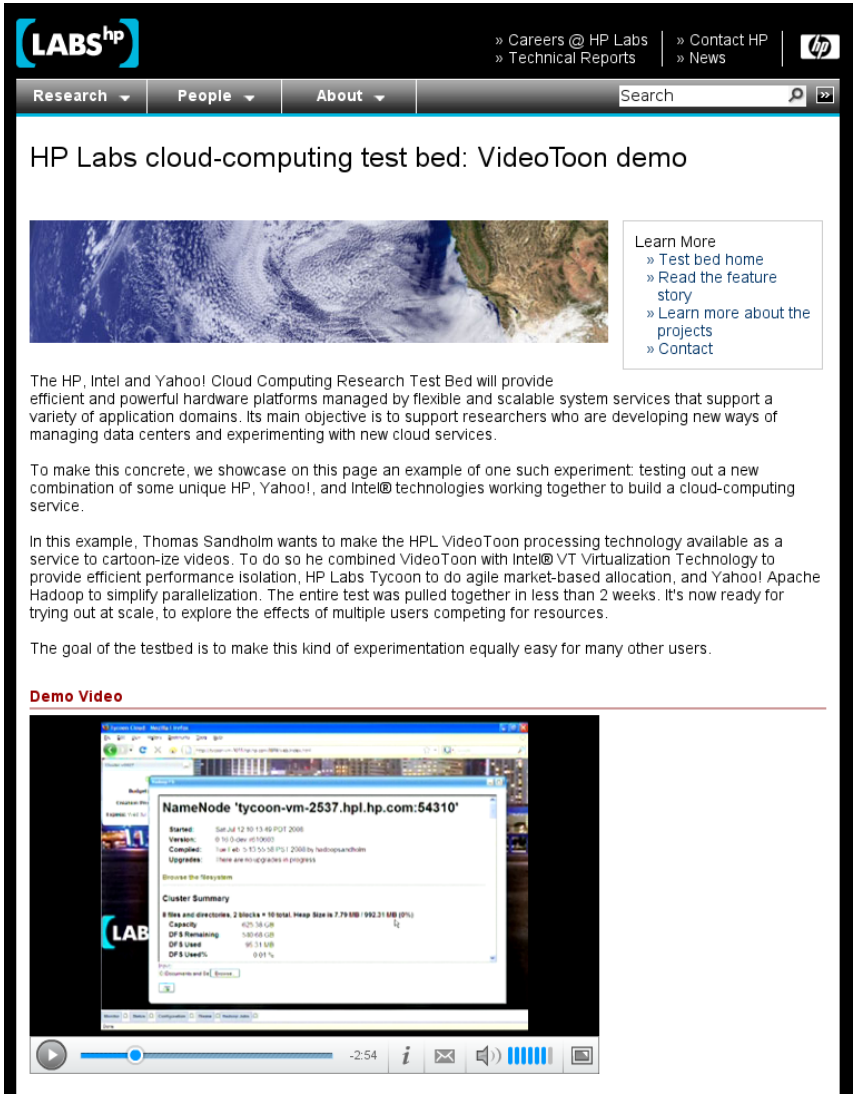

HP Labs' 2008 Request for Proposals

- CERN submitted 6 different proposals in June
  - From FIO, DM, GS groups
- Also visit to Palo Alto by A. Pierson/S. Jarp
  - Discussions with R. Friedrich (Director of Open Innovation Project) and several Directors of relevant Labs
- Unfortunately none of the CERN proposals were accepted!

CERN openlab

Q & A