SWINBURNE UNIVERSITY OF TECHNOLOGY

Predicting supercomputing workload using per user information

Tuan V. Dinh, Lachlan Andrew and Philip Branch

13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing.

Delft, 14th -16th May, 2013

Motivation

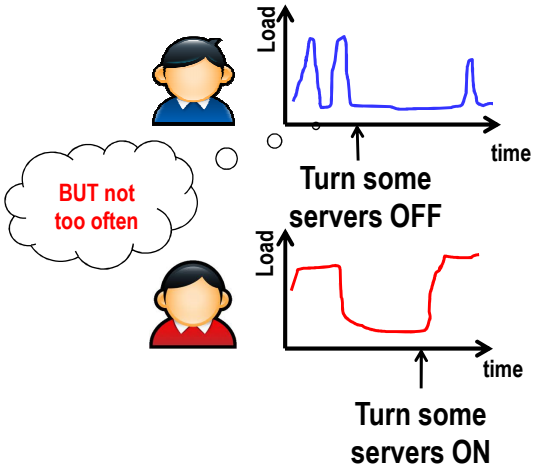

Another Application: valley filling

Outline

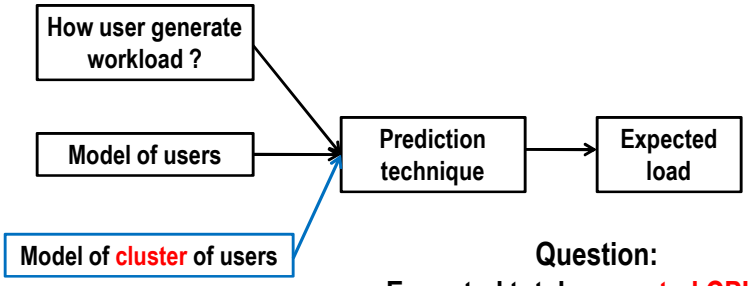

Question:
Expected total requested CPUs of incoming jobs ?

Related work

1. Characterising supercomputer job arrival process:

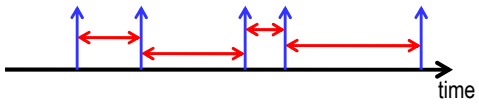

Feitelson, 1999: Poisson process

Exponential distribution — i.i.d

Lublin et al., 2003: Gamma distributions

Squilante et al., 1999: studied dependence between arrivals

Li et al., 2005: Heavy-tail distributions (Log-normal or Pareto)

primarily descriptive, developing synthetic workload models

Related work

2. Estimating future CPU utilisation:

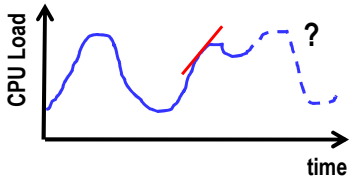

arrivals ~ derivative of the CPU load

Auto-regression (AR) models: use historical correlation

- **Wu et al., 2010:** successful for Grid computing
- **Liang et al., 2013:** use AR model, but employ data filters (Kalman).

Outline

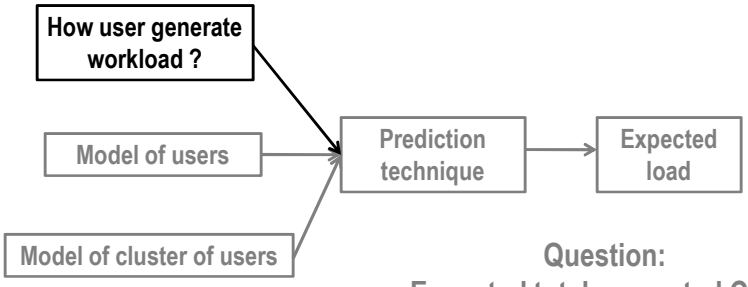

Question:
Expected total requested CPUs of incoming jobs ?

Swinburne supercomputer

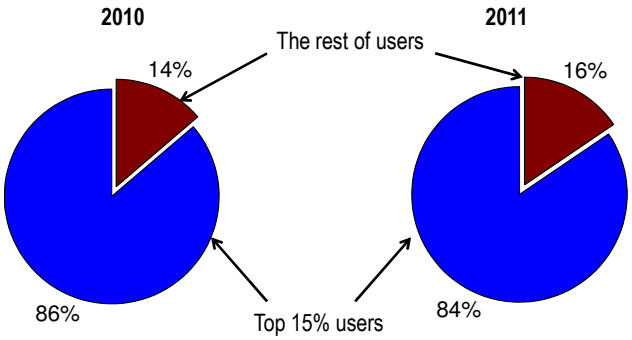

Year	Users
2010	86
2011	104
2010 & 2011	136

SWINBURNE UNIVERSITY OF TECHNOLOGY

User submission behaviors

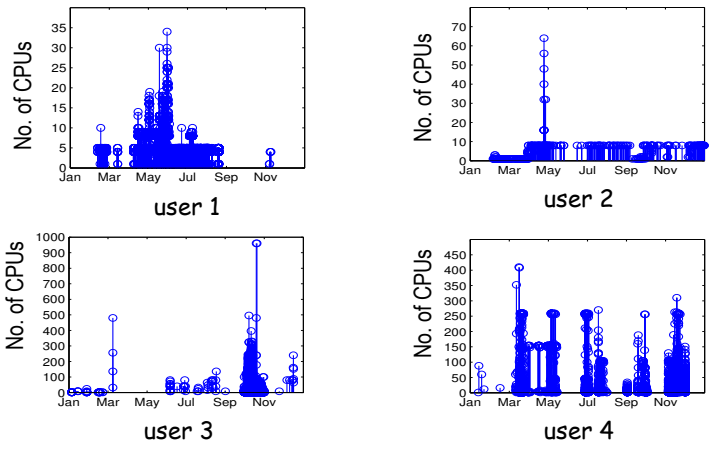

BIG USER: HETEROGENOUS

SWINBURNE UNIVERSITY OF TECHNOLOGY

Outline

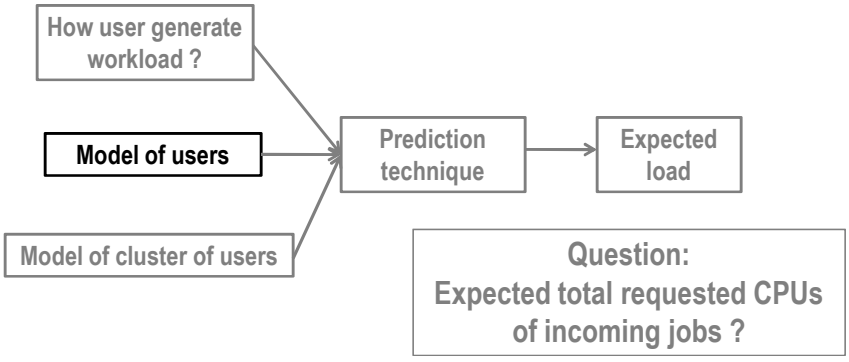

User's model

Use job inter-arrival time distribution of each user

$$\text{Hazard rate function} = \frac{\text{Probability of he/she about to submit one}}{\text{Given not submitted for } t \text{ unit of time:}} = \frac{\text{pdf(inter-arrival time)}}{1 - \Pr[\text{inter-arrival time} < t]}$$

Can be obtained from system

Measured hazard curves

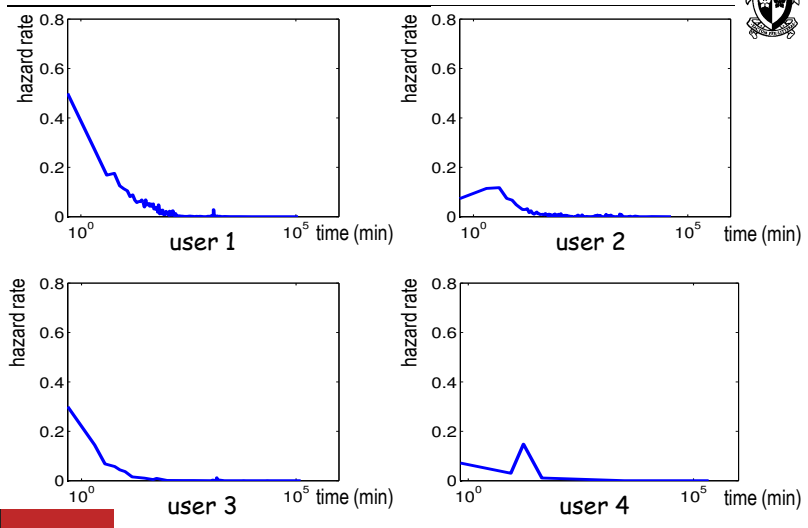

SWINBURNE UNIVERSITY OF TECHNOLOGY

Parametric fitting

Fitting family	Key properties
$h_W(t) \propto t^b$ (Weibull)	monotonic
$h_C(t) \propto t^b e^{-ct}$	

SWINBURNE UNIVERSITY OF TECHNOLOGY

Fitting quality

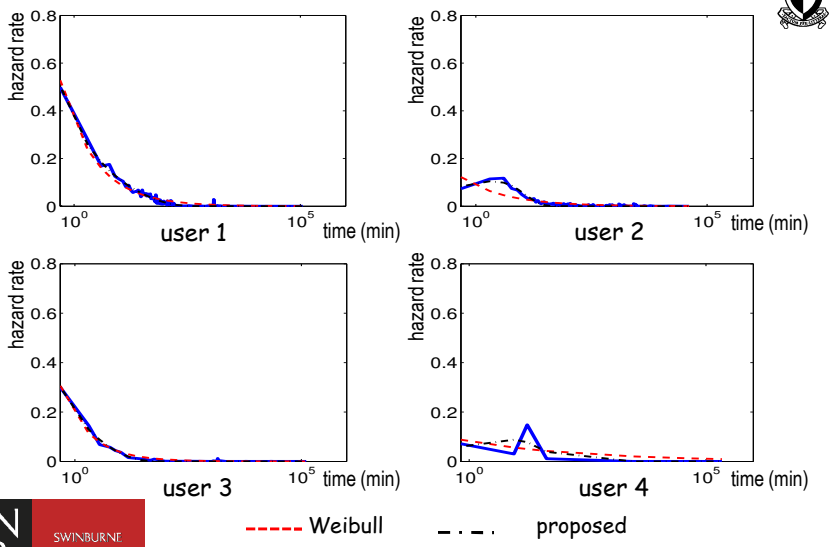

SWINBURNE UNIVERSITY OF TECHNOLOGY

Outline

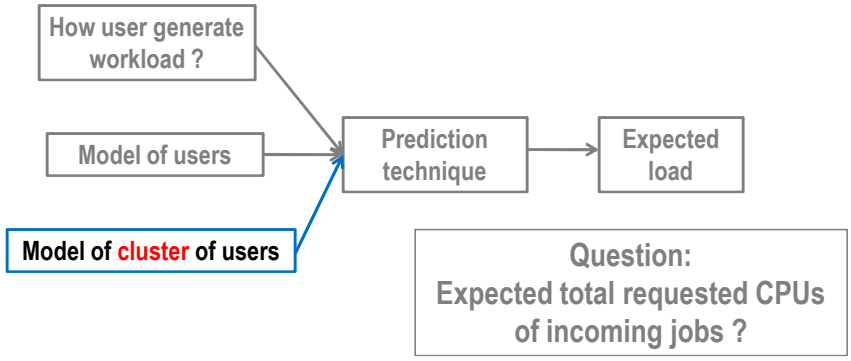

SWINBURNE UNIVERSITY OF TECHNOLOGY

User clustering

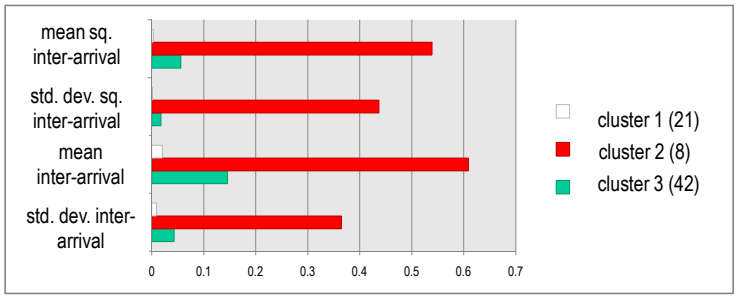

User clustering for 2010 (except for the top 15 users). C = 3.
Algorithm: k-means algorithm. Software: WEKA [*]

[*]: http://www.cs.waikato.ac.nz/ml/weka/

Outline

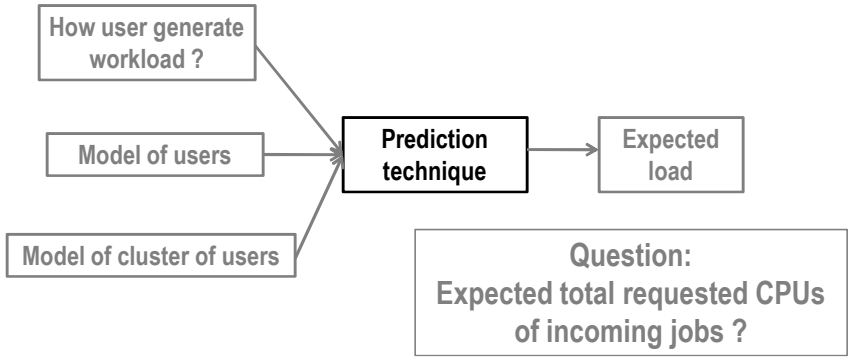

Schemes

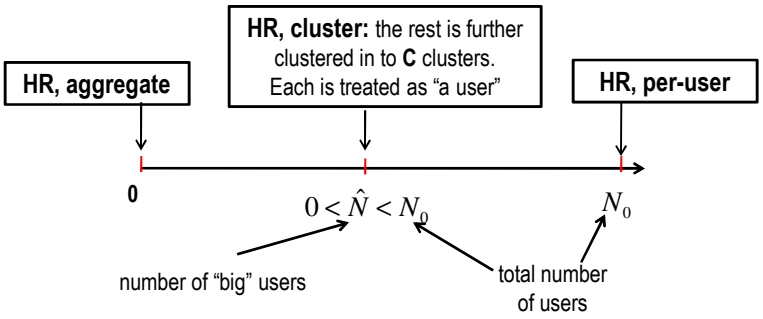

SWINBURNE UNIVERSITY OF TECHNOLOGY

Performance evaluation

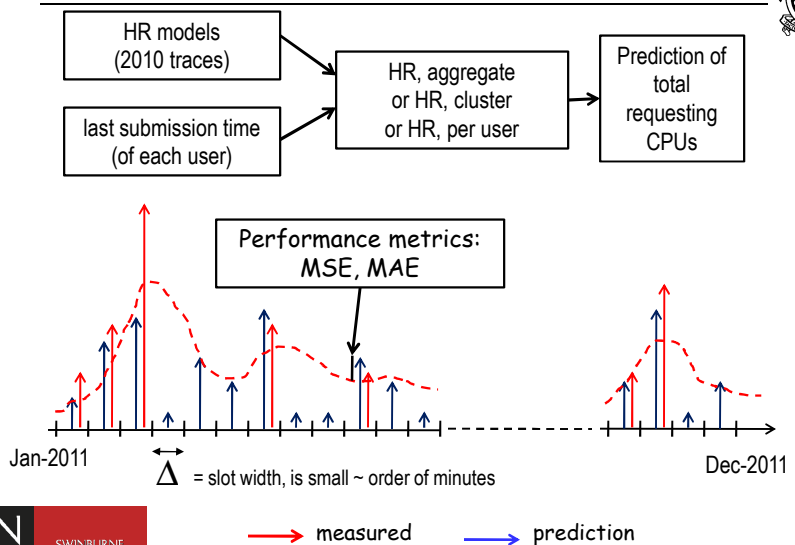

SWINBURNE UNIVERSITY OF TECHNOLOGY

Performance evaluation

Schemes	no smoothing		smoothing	
	MSE	MAE	MSE	MAE
AR(p=35)	4,300	11.67	2,310	9.29
HR, per-user	2,749	10.68	381	7.46
HR, aggregate	2,738	8.60	377	6.10
HR, clustered	2,733	8.44	371	5.86

- No. of individual users: **15**
- Slot time = 2 minutes
- C = 3, clustering algorithm: k-means algorithm. Attributes: mean and std. dev. inter-arrival time; mean and std. dev. *squared* inter-arrival time

SWINBURNE UNIVERSITY OF TECHNOLOGY

Degree of clustering ?

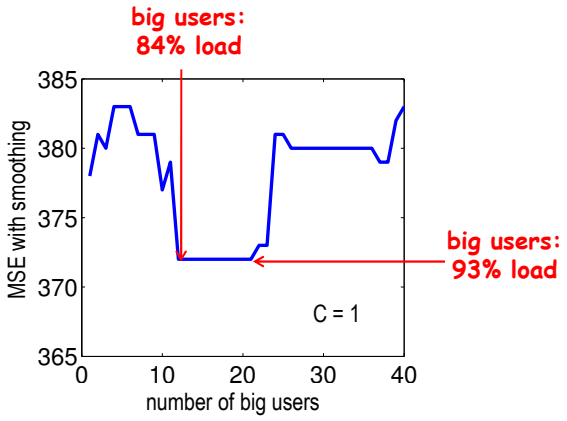

SWINBURNE UNIVERSITY OF TECHNOLOGY

Remarks and future work

Limitations:

- experiment with only ONE workload
- assumptions on user correlations

Future work:

- conduct experiment with more workload
- optimising user clustering process

SWINBURNE UNIVERSITY OF TECHNOLOGY

Conclusions

User: heterogeneous

- Big users dominate
- User's submission behaviours are different
- Cluster small users

"HR, cluster" works best

- Modelling **all** users may give worse results
- Choosing right number of big users

Next (or remaining) questions ?

- Other traces ?
- Optimise the clustering process

SWINBURNE UNIVERSITY OF TECHNOLOGY