Maximum Likelihood Estimation of Closed Queueing Network Demands from Queue Length Data

Weikun Wang *(Imperial College London, UK)*

Giuliano Casale *(Imperial College London, UK)*

Ajay Kattepur *(TCS Innovation Labs, India)*

Manoj Nambiar (TCS Innovation Labs, India)

ACM/SPEC ICPE 2016 - March14th, 2016

SPE & Model-Driven Engineering

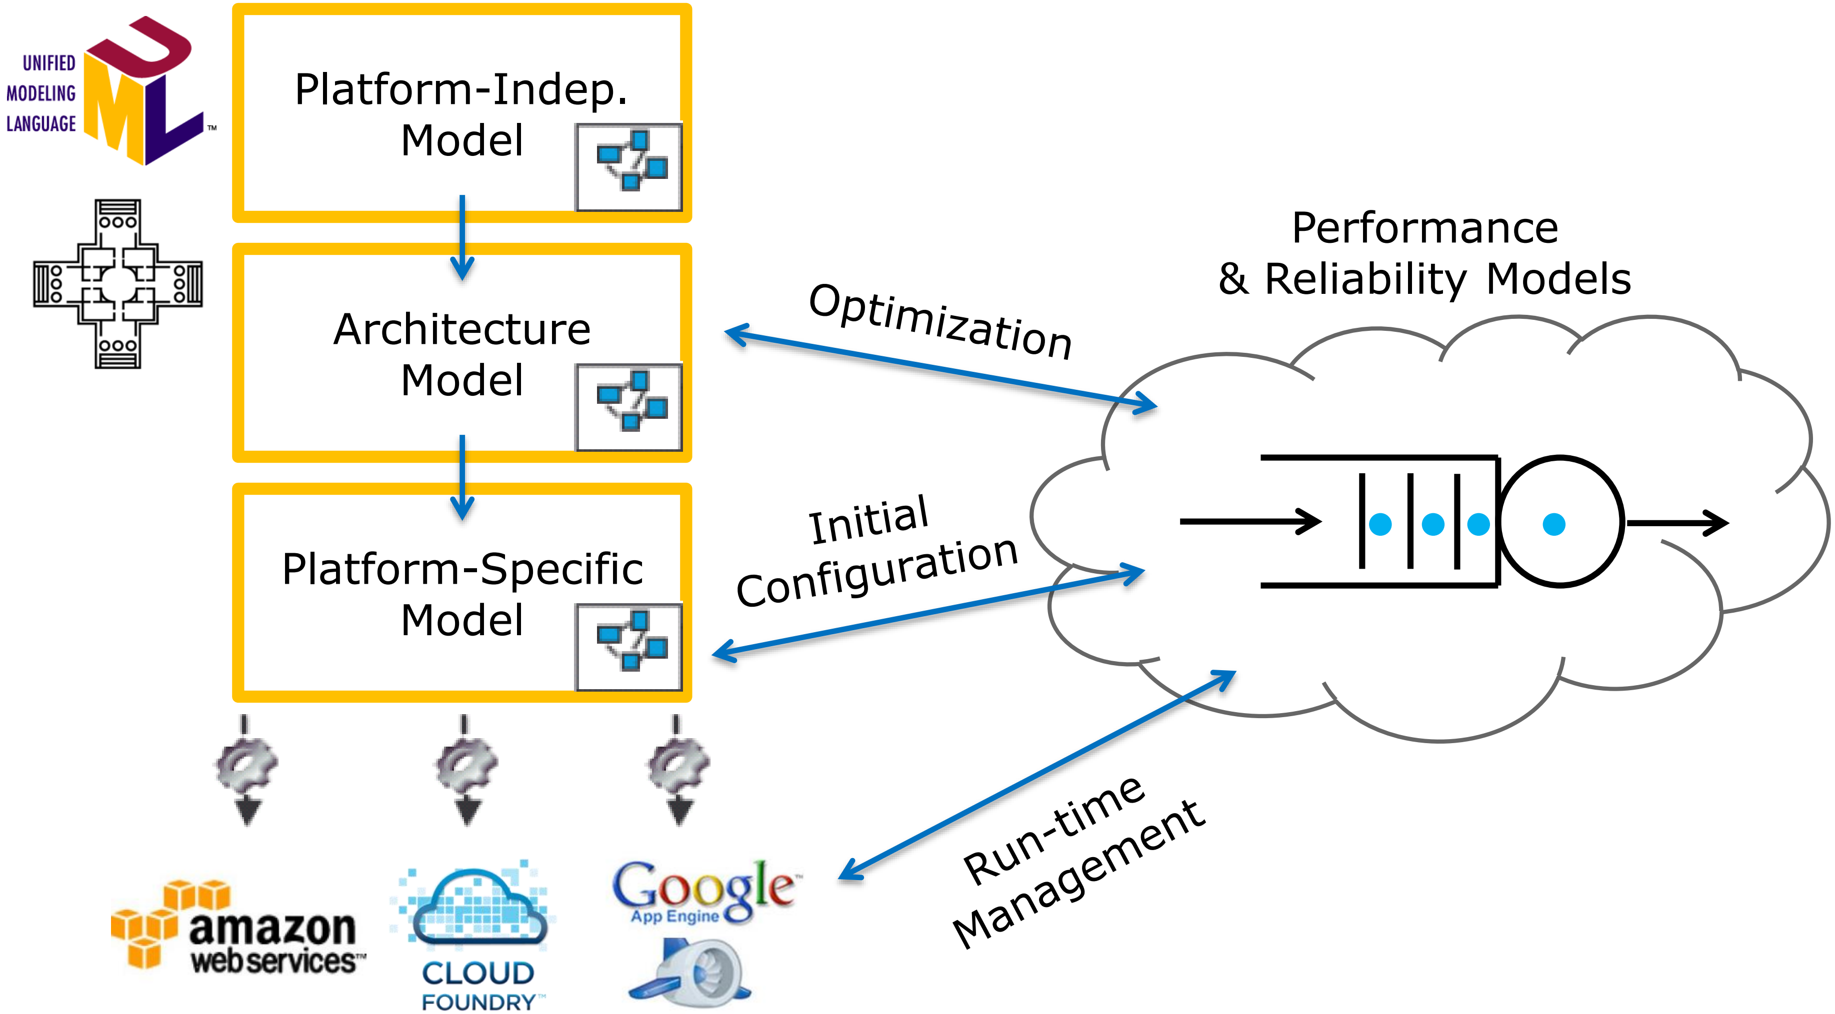

CQN Characterization

- Service demand of a request
  - CPU time, bandwidth consumed, …
- Multi-threaded software
  - *e.g.*, web servers

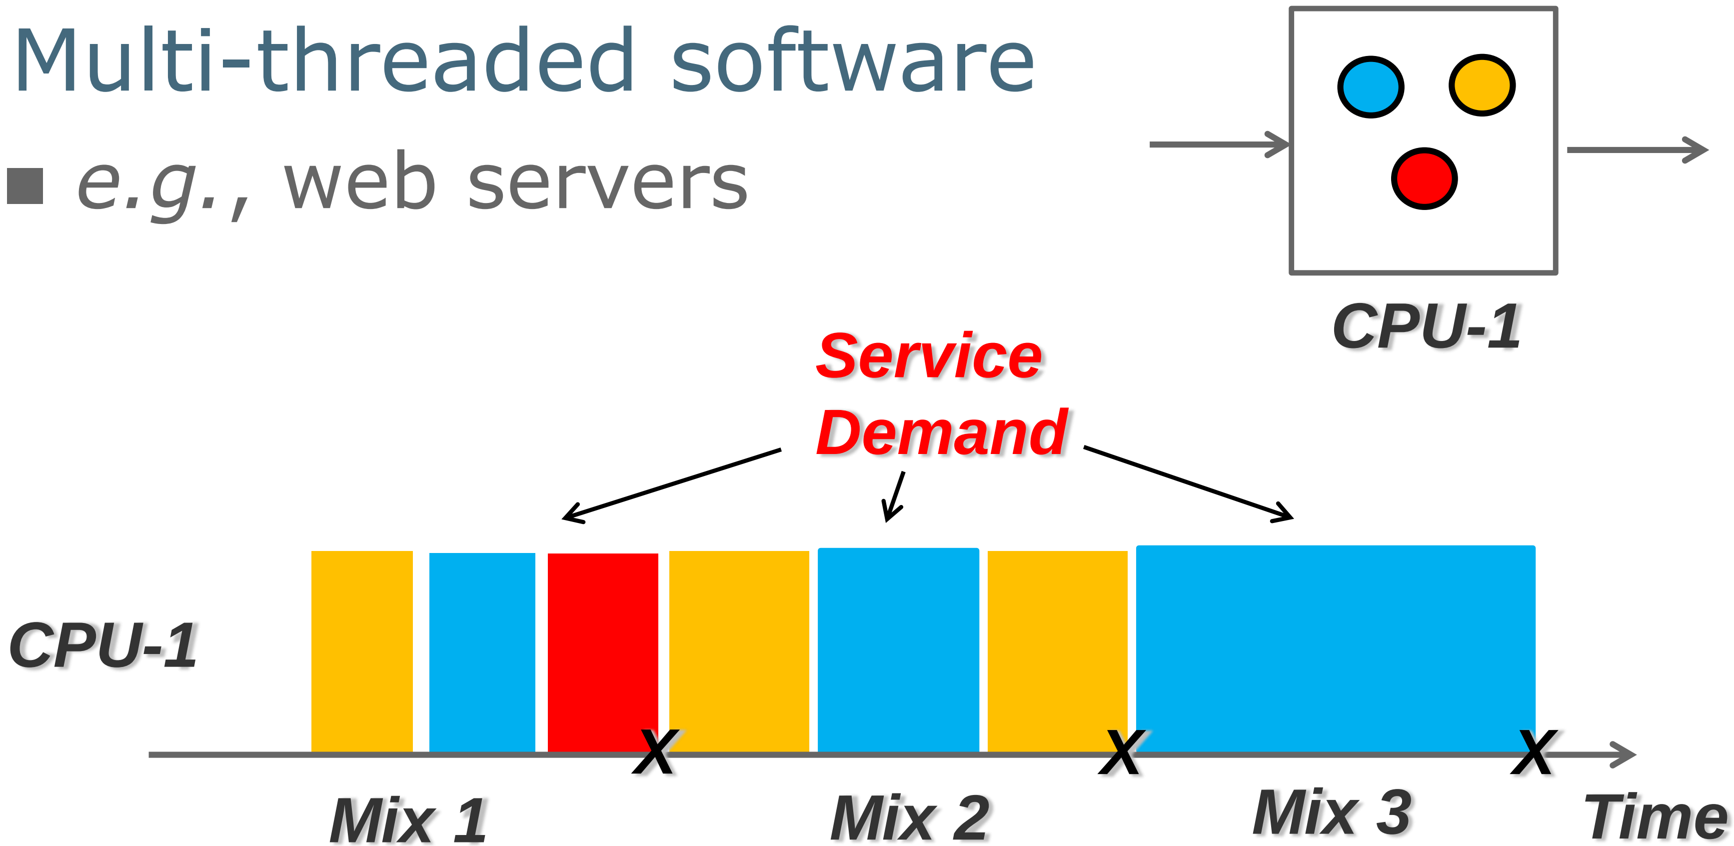

Drawback of Existing methods

Imperial College London
TATA CONSULTANCY SERVICES

- Utilization based approaches
  - Regression based on utilization and throughput
  - Issues: collinearities, load-dependence, outliers, utilization unreliable/unavailable, …

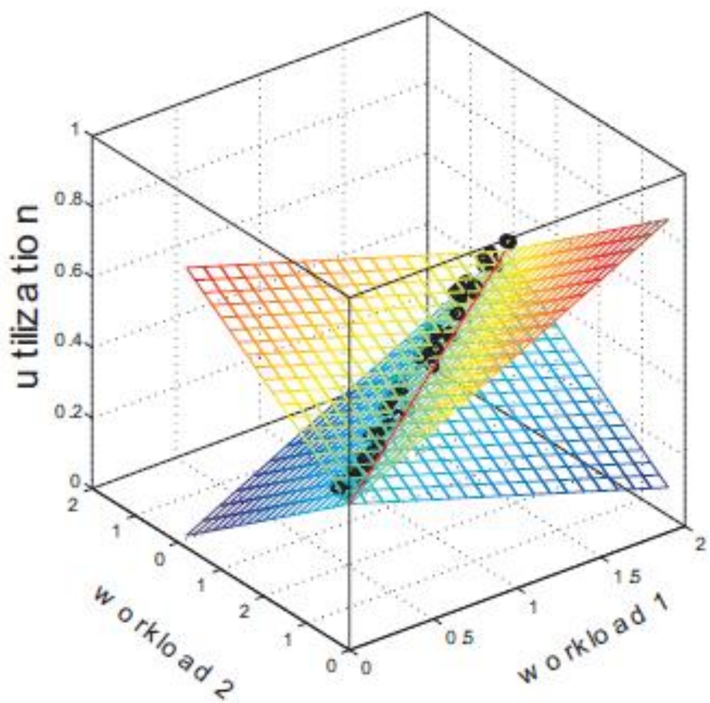

Queue length samples

- State observations
  - Dataset (L points): $\boldsymbol{n}^l \in \boldsymbol{D}$
  - CQN State: $\boldsymbol{n} = (n_{01}, \ldots, n_{0R}, n_{11}, \ldots, n_{1R}, \ldots, n_{MR})$

n =(0,1,3, 1,1,0, 0,0,0)

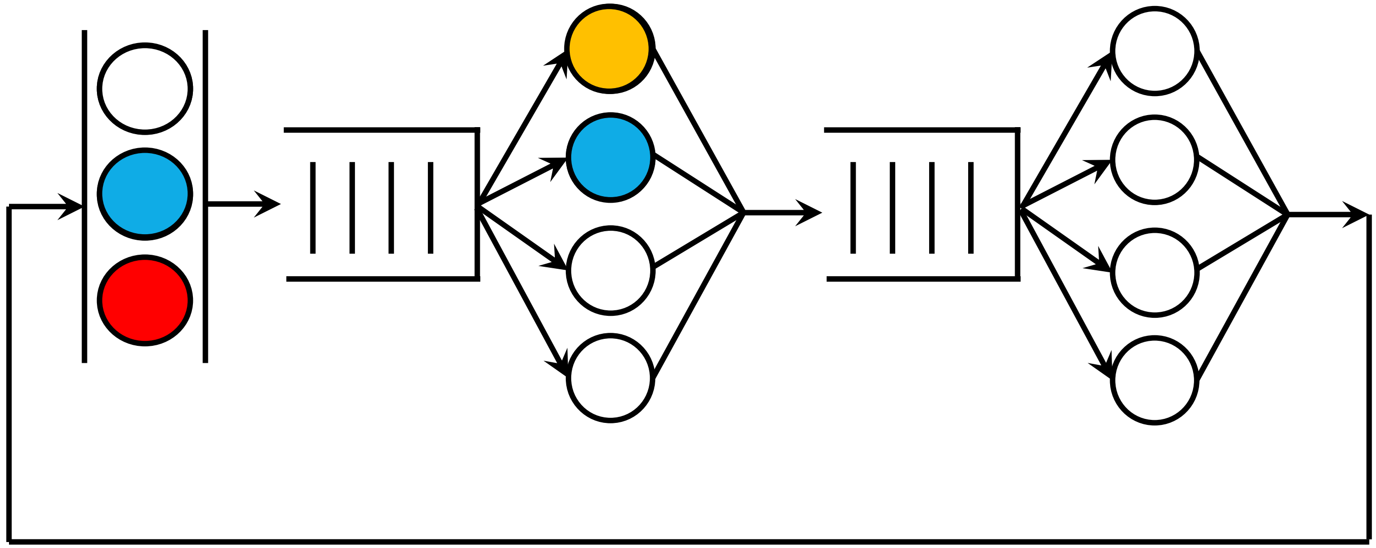

Queue length samples

- State observations
  - Dataset (L points): $\boldsymbol{n}^l \in \boldsymbol{D}$,
  - CQN State: $\boldsymbol{n} = (n_{01}, \ldots, n_{0R}, n_{11}, \ldots, n_{1R}, \ldots, n_{MR})$

n =(0,0,1, 1,1,1, 0,1,1)

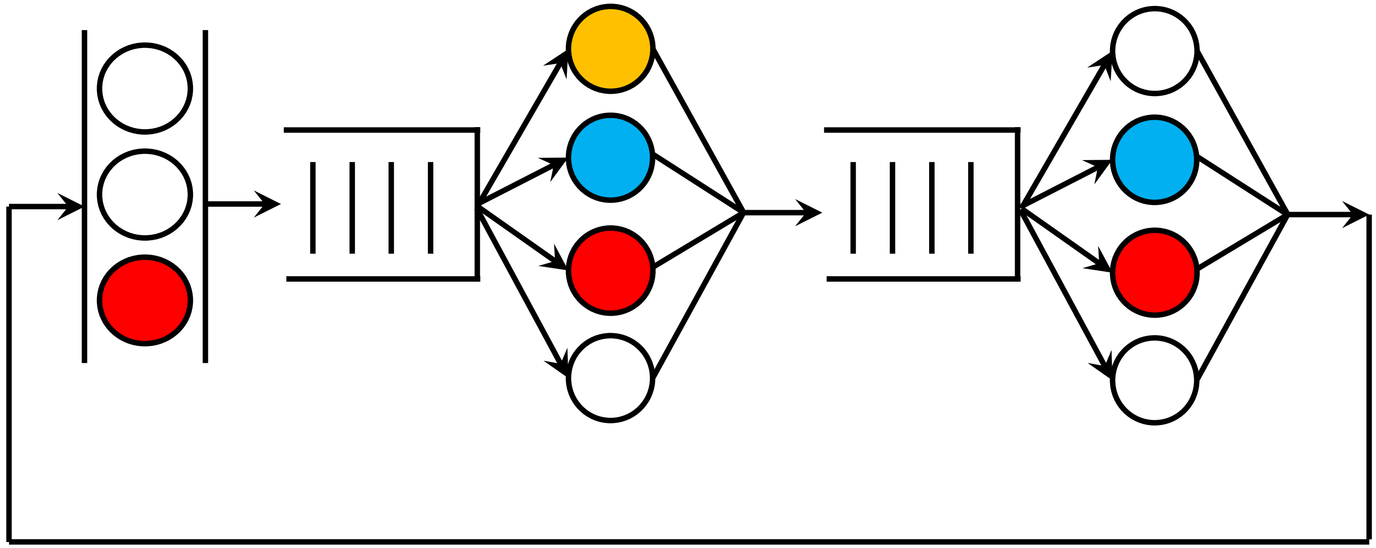

Queue length samples

- Assume product-form state probabilities

Service demand   Queue length

$$\mathbb{P}(\boldsymbol{n}|\boldsymbol{\theta}) = \prod_{j=1}^{R} \frac{\theta_{0j}^{n_{0j}}}{n_{0j}!} \prod_{i=1}^{M} n_i! \prod_{j=1}^{R} \frac{\theta_{ij}^{n_{ij}}}{n_{ij}!G(\boldsymbol{\theta})}$$

Normalizing constant

- Computationally challenging to evaluate $\mathbb{P}(\boldsymbol{n}|\boldsymbol{\theta})$
- Maximum likelihood estimation?
  - Infer demands with the probability

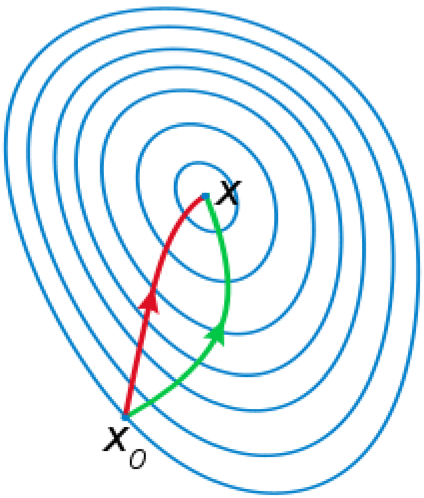

Maximum Likelihood Estimation

Imperial College London

TATA CONSULTANCY SERVICES

- Maximum likelihood estimation (MLE)

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \underbrace{\prod_{l=1}^{L} \mathbb{P}(\boldsymbol{n}^l | \boldsymbol{\theta})}$$

parameter space

Likelihood $\mathcal{L}(\boldsymbol{\theta})$

- Problem with direct computation
  - Evaluation of $\mathbb{P}(\boldsymbol{n}|\boldsymbol{\theta})$ for each observation
  - Slow due to the need for computing $G(\boldsymbol{\theta})$
  - Very small probabilities when L is large
- Any other solution?

- A necessary condition for a point $\hat{\boldsymbol{\theta}}$ inside $\boldsymbol{\Theta}$ to be a MLE is that

$$Q_{ij}(\hat{\boldsymbol{\theta}}) = \widetilde{Q}_{ij}(\boldsymbol{D}), \qquad \forall i, j,$$

theoretical mean queue length

observed mean queue length

Only ***mean*** queue length is required!

- How to find the MLE?
  - Change the value of $\boldsymbol{\theta}$, until the mean queue length predicted with MVA match $\widetilde{Q}_{ij}(\boldsymbol{D})$
  - Fixed point iteration or an optimization program

Confidence Intervals

- Assume the MLE to be asymptotically normal
- Confidence intervals for the MLE demands

$$\hat{\theta}_{ij} \pm c\sqrt{(\boldsymbol{I}(\hat{\boldsymbol{\theta}})^{-1})_{ij,ij}}$$

  - $\boldsymbol{I}(\hat{\boldsymbol{\theta}}) = -\boldsymbol{H}(\hat{\boldsymbol{\theta}})$ is the Fisher Information matrix
  - $\boldsymbol{H}(\hat{\boldsymbol{\theta}})$ is the Hessian matrix $\boldsymbol{H}(\boldsymbol{\theta})_{ij,kh} = \frac{\partial^2 \log \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_{ij} \partial \theta_{kh}} \Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}}$
- $\boldsymbol{H}(\hat{\boldsymbol{\theta}})$ works with mean queue length only!
  - Obtained by using standard MVA, no probabilities!

QMLE Approximation

- Exact MLE can be found by direct search
  - Fixed-point iteration tends to be effective
- A simple approximation of the MLE:
  - Consider the demand vector $\boldsymbol{\theta}^{bs}$ where

$$\theta_{ij}^{bs} = \frac{\widetilde{Q}_{ij}(\boldsymbol{D})}{(N_j - \sum_{k=1}^{M} \widetilde{Q}_{kj})} \frac{\theta_{0,j}}{(1 + \sum_{h=1}^{R} \widetilde{Q}_{ih} - \widetilde{Q}_{ij}(\boldsymbol{D})/N_j)}$$

- Then it must be

$$Q_{ij}^{bs}(\boldsymbol{\theta}^{bs}) = \widetilde{Q}_{ij}(\boldsymbol{D})$$

observed mean queue length

Validation- Existing methods

- CI: Complete Information

  [J.F. Perez *et al.*, *IEEE Trans. Sw. Eng.'15]*

  - Full knowledge of sample path
  - Baseline approach

- ERPS: Extended Regression for Processor Sharing

  [J.F. Perez *et al.*, *IEEE Trans. Sw. Eng.'15]*

  - Based on mean response time and arrival queue

- GQL: Gibbs Sampling for Queue Lengths

  [W. Wang et al., Accepted to appear in ACM TOMACS]

  - Gibbs sampling based on queue length samples
  - Many iterations until convergence

Validation

- ≈20000 random models
  - Randomized number of stations, classes, jobs
  - Focus on QMLE instead of exact analysis
- Results
  - All the algorithms: below 10%
  - QMLE has less than 4% error
  - Confidence interval validated

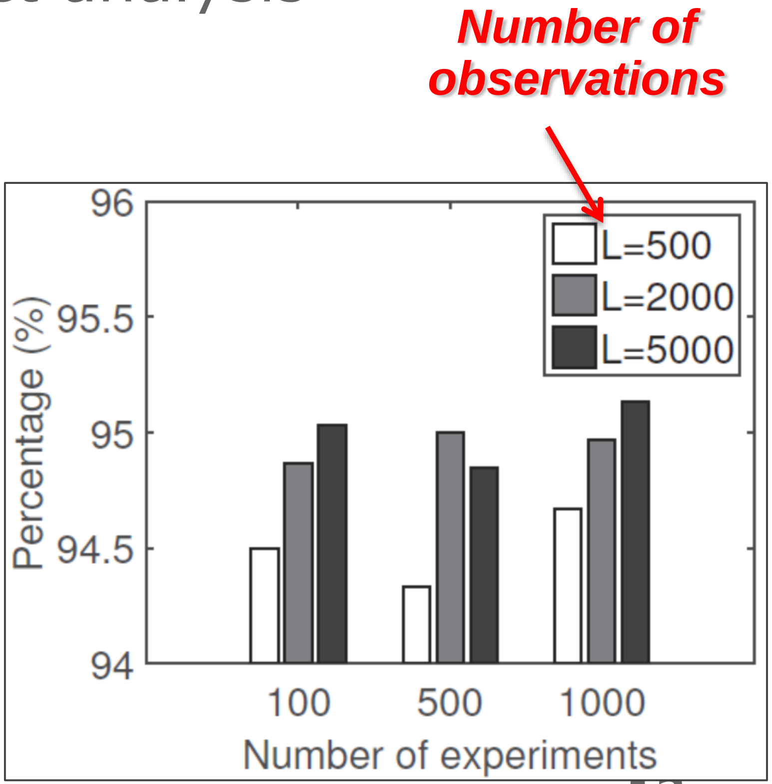

Execution time

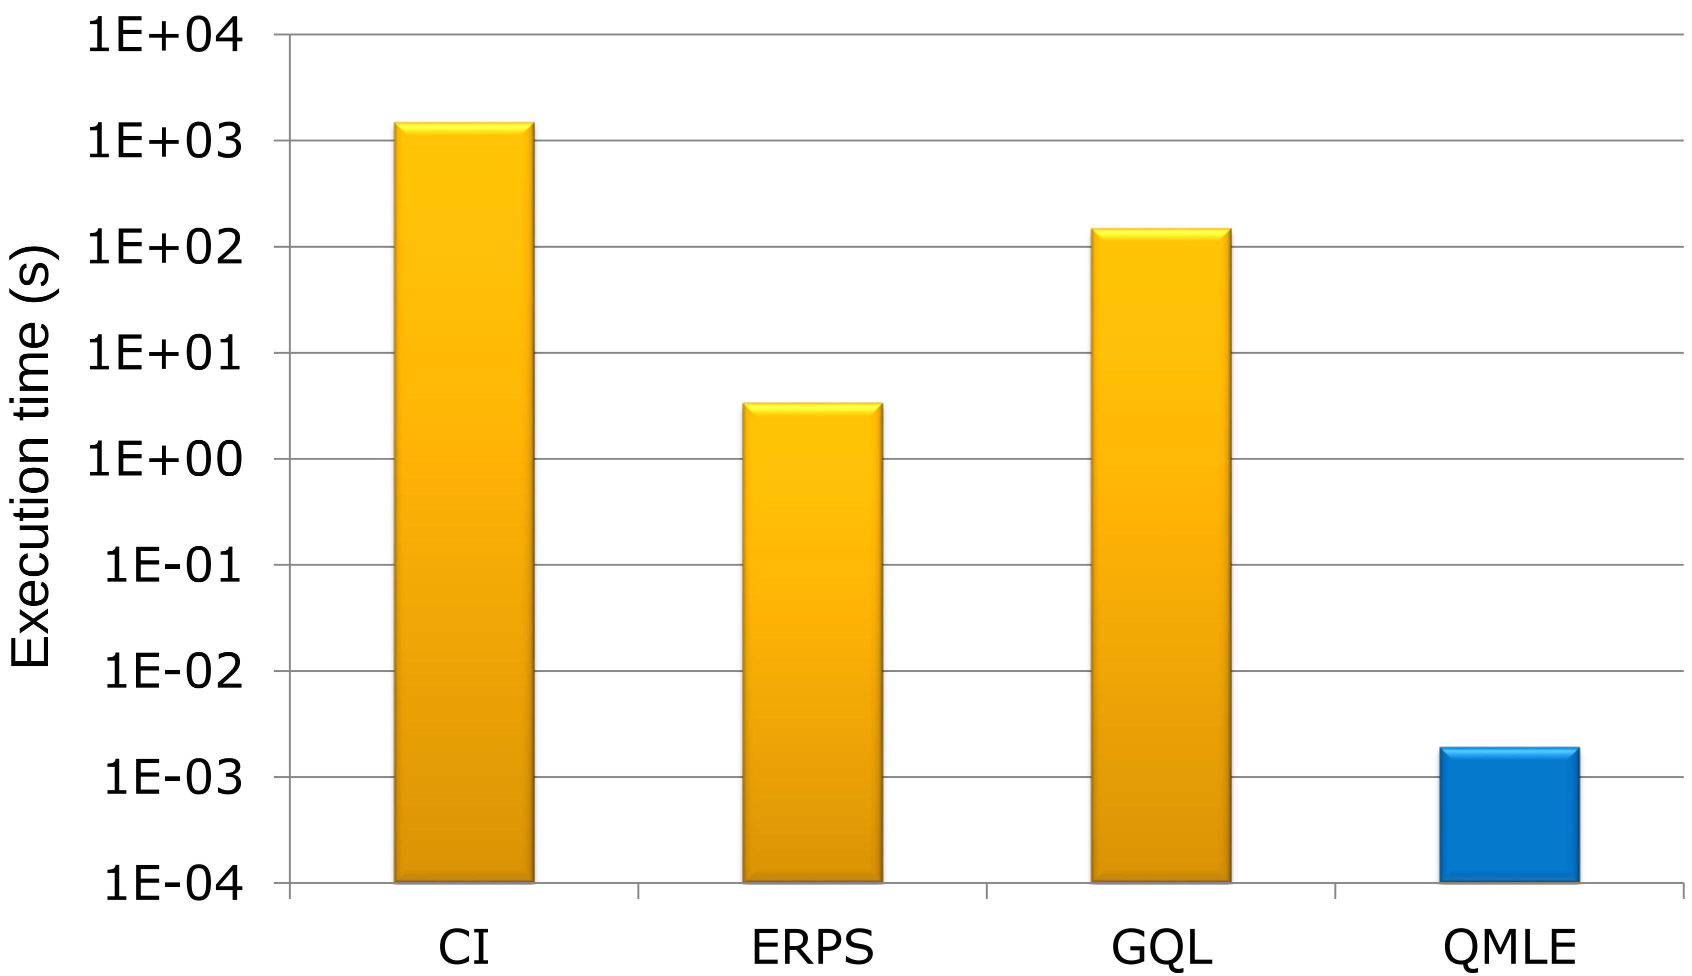

Load-dependent (LD) extension

- Mean demand varies under different load
- Real world system behavior
  - e.g. multi-core servers

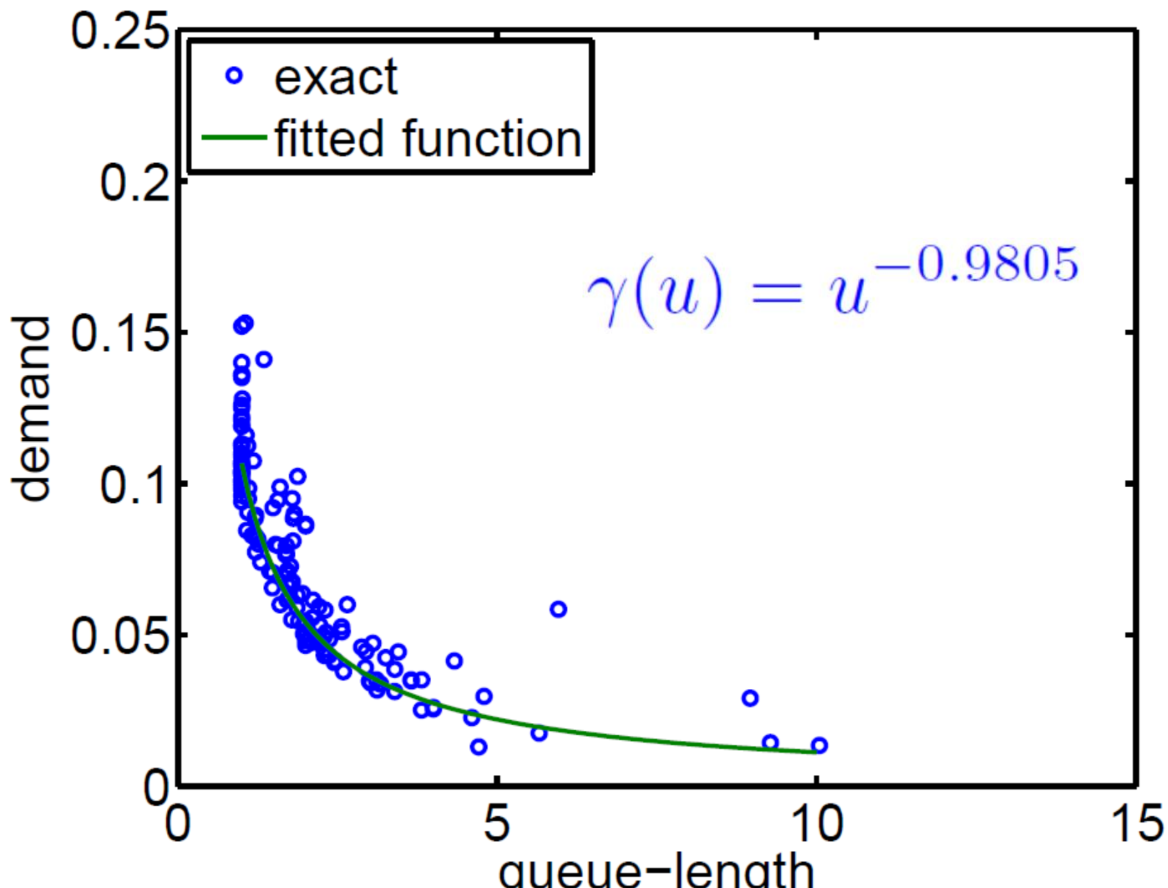

Load-dependent (LD) extension

- A scaling factor function $\gamma_i(u)$
  - load-independent : $\gamma_i(u) = 1, 1 \leq u \leq n_i$
- Product-form still holds

$$\mathbb{P}(\boldsymbol{n}|\boldsymbol{\theta},\boldsymbol{\gamma}) = \left(\prod_{j=1}^{R} \frac{\theta_{0j}^{n_{0j}}}{n_{0j}!}\right) \prod_{i=1}^{M} n_i! \prod_{j=1}^{R} \frac{\theta_{ij}^{n_{ij}}}{n_{ij}!G(\boldsymbol{\theta})} \underbrace{\prod_{u=1}^{n_i} \gamma_i(u)}_{\textit{new term}}$$

- MLE

$$(\hat{\boldsymbol{\theta}},\hat{\boldsymbol{\gamma}}) = \arg \max_{(\boldsymbol{\theta},\boldsymbol{\gamma})\in\boldsymbol{\Theta}} \prod_{l=1}^{L} \mathbb{P}(\boldsymbol{n}^l|\boldsymbol{\theta},\boldsymbol{\gamma})$$

MLE characterization

- Directly computation is infeasible
- A necessary condition for a point $(\hat{\boldsymbol{\theta}}, \hat{\gamma})$ inside $\boldsymbol{\Theta}$ to be a MLE is that

$$Q_{ij}(\hat{\boldsymbol{\theta}}, \hat{\gamma}) = \widetilde{Q}_{ij}(\boldsymbol{D}), \qquad \forall i, j$$

and

$$\mathbb{P}(n_k = v | \hat{\boldsymbol{\theta}}, \hat{\gamma}) = \mathbb{P}(\widetilde{n}_k = v | \boldsymbol{D}), \qquad \forall k, v$$

Theoretical marginal queue length probability

Empirical marginal queue length probability

- Works with marginal probability only!

MLE characterization

- How to find the MLE?
  - Solve by optimization program
- Confidence intervals
  - Hessian matrix can still be derived
  - Computation requires marginal probabilities and mean queue length only
- Drawback
  - Computationally expensive because of LD-MVA

Validation

- Random models validation
  - 2 stations, 2 classes, 8 jobs, different think time
  - MATLAB *fmincon* solver
  - Compare the estimated $(\hat{\boldsymbol{\theta}}, \hat{\gamma})$ against exact ones

- Considered scaling factors
  - $\gamma_i(u) = 1/\min(u, C_i)$: resembles multi-core feature
    - C_i number of CPUs in queueing station i.

Estimation error on γ (scaling factors)

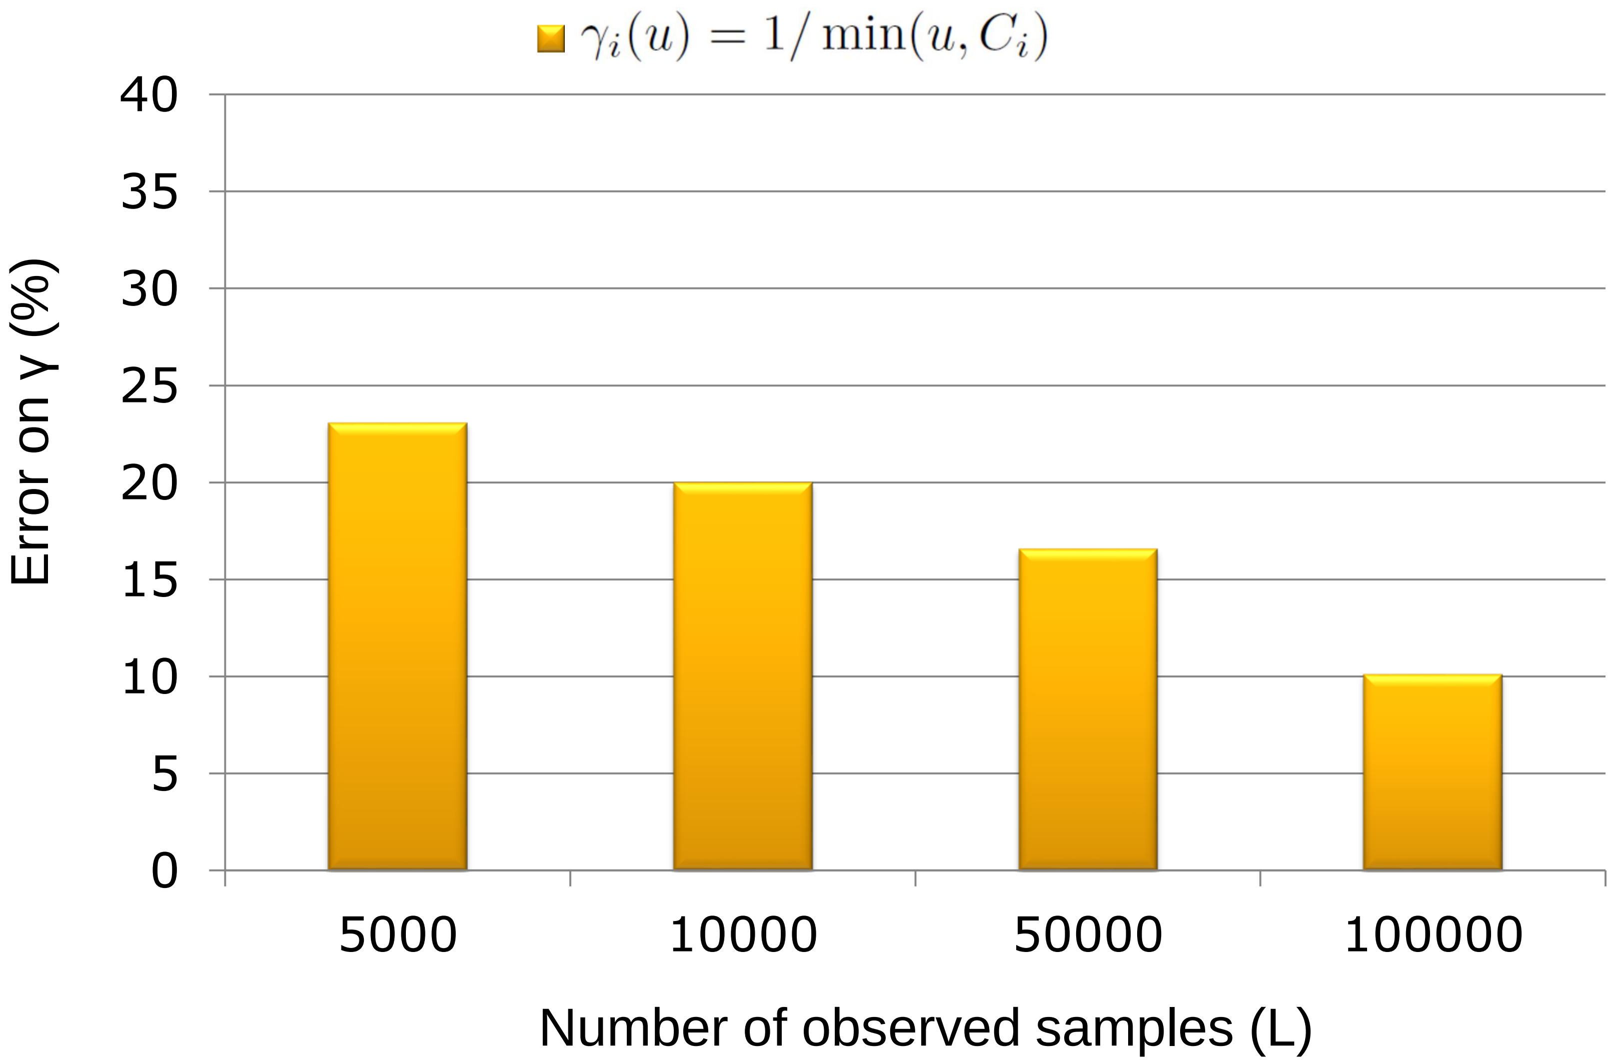

Execution time

Progress: We found a new approximation method for efficiently evaluating marginal probabilities, which reduces the execution time to < 20s on average!

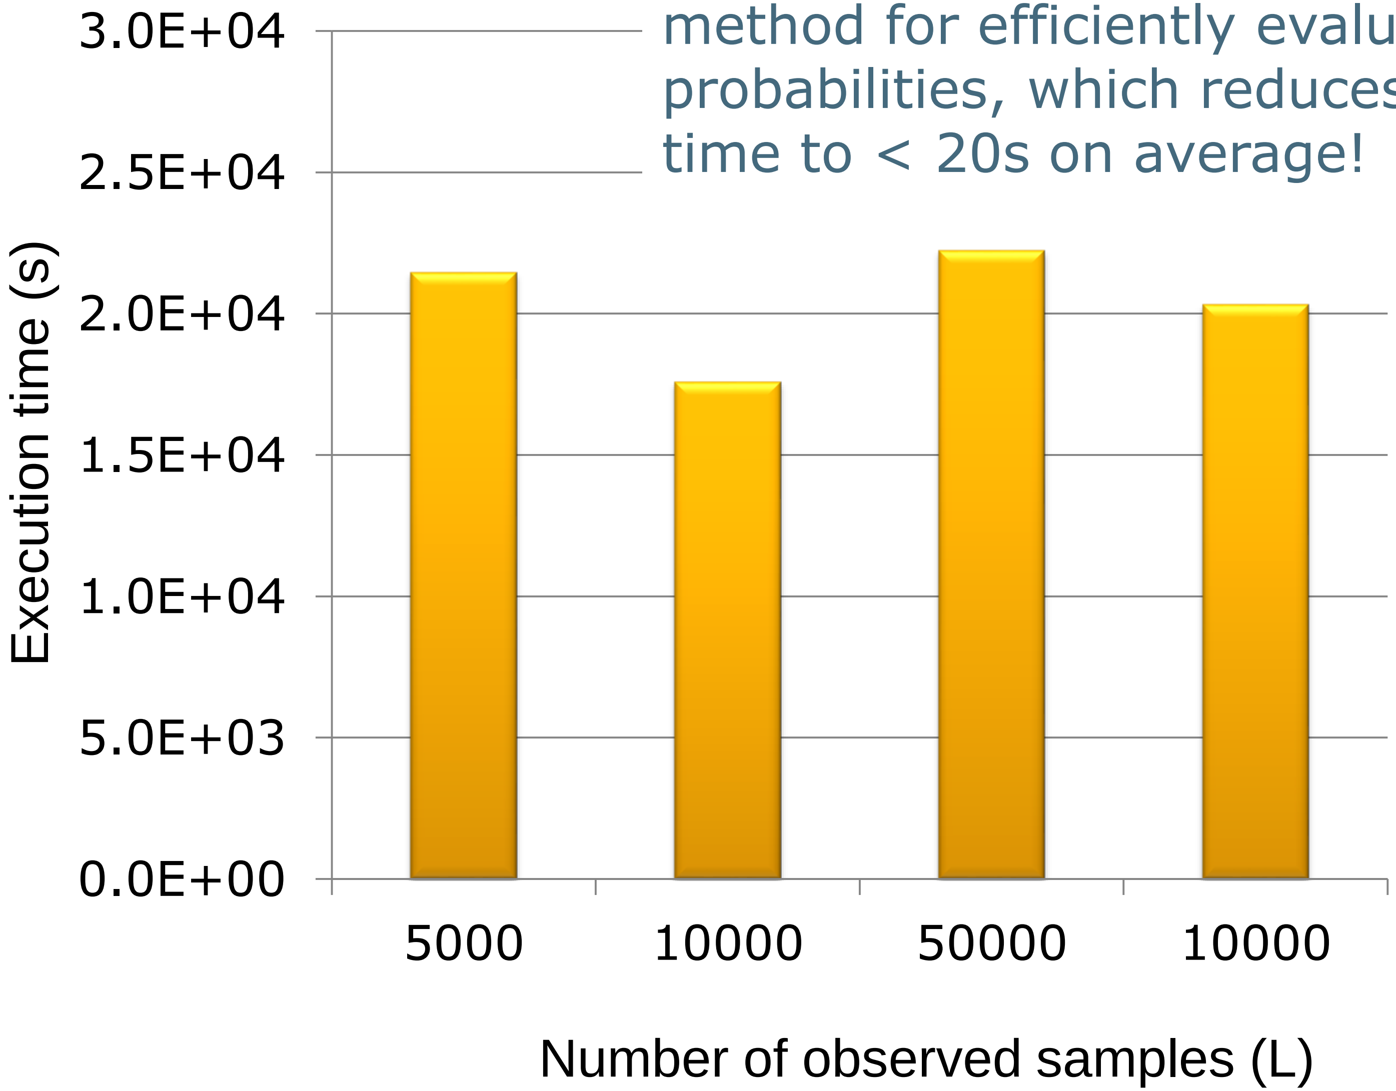

Case Study (MyBatis JPetStore)

Imperial College London
TATA CONSULTANCY SERVICES

- 3-tier commercial application
  - Transactions grouped in R=1 class
  - 5 GB user data

MyBatis

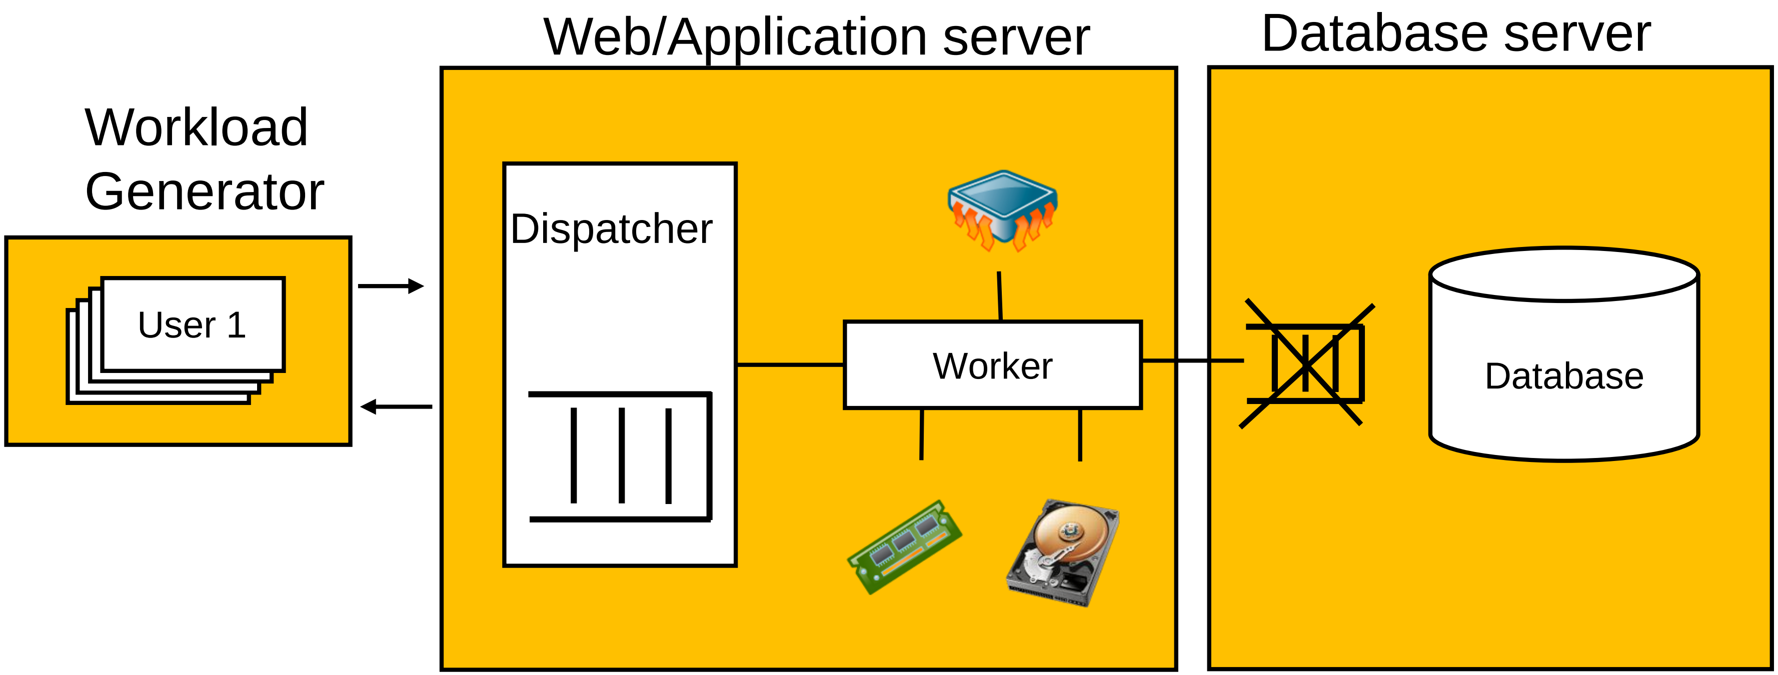

Observed performance matching

Imperial College London

TATA CONSULTANCY SERVICES

- Exact demand unknown
  - Estimated demands using QMLE
  - Validate observed throughput with estimated demands

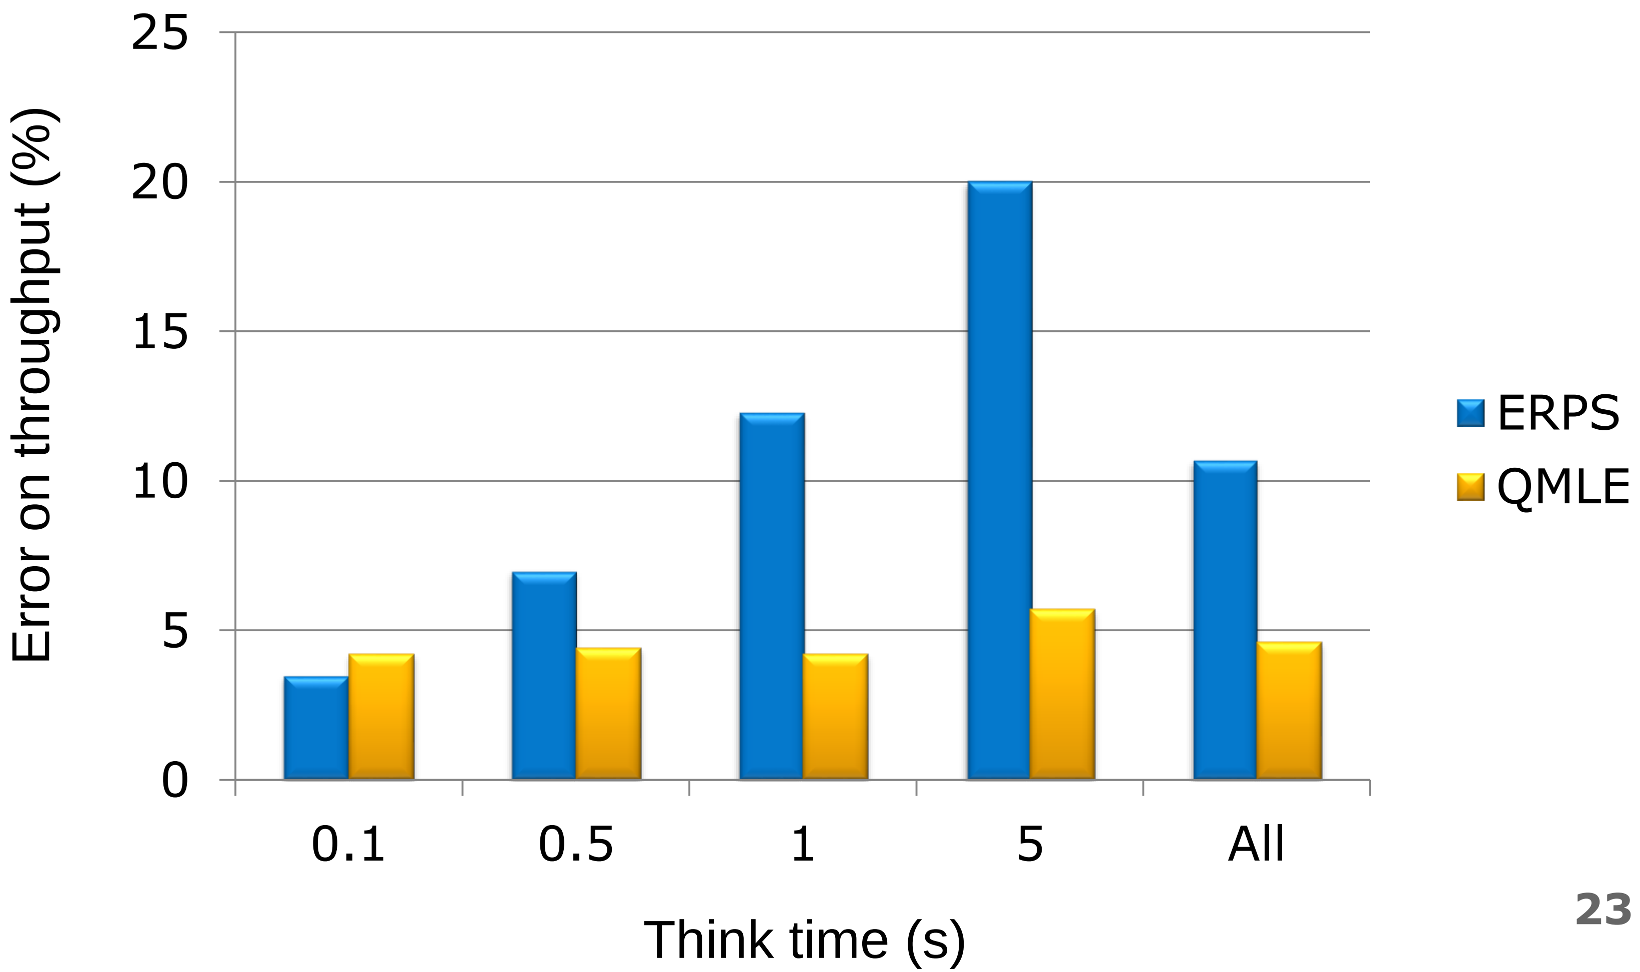

Conclusion

- Demand estimation from queue length
  - Efficient
  - Confidence interval characterization
  - Load-dependent extension
- Ongoing work
  - Accelerate the load-dependent estimation
  - More experimental evaluations

Funded by FP7 MODAClouds, H2020 DICE, EPSRC OptiMAM

Thanks!

weikun.wang11@imperial.ac.uk