Fine-scale Burstiness Spells Bad News for Elasticity!

Sadeka Islam[1], Srikumar Venugopal[1,2], Anna Liu[1,3] (1. UNSW 2. IBM Research - Ireland 3. NICTA)

sadeka.islam@student.unsw.edu.au, srikumarv@ie.ibm.com, anna.liu@nicta.com.au

1. Problem Statement

• **Fine-scale burstiness: severe fluctuations** at small timescales (e.g., seconds) in the request arrival rates.
  • existing workload models **don't preserve** the empirical statistical regularities of fine-scale burstiness.
  • impact of fine-scale burstiness on elasticity is **still unexplored**.

• Overlooking these issues leads to **imperfect elasticity**, thereby affecting operational expenses and overall revenue.

Wikipedia workload snippet (Oct 1, 2007)

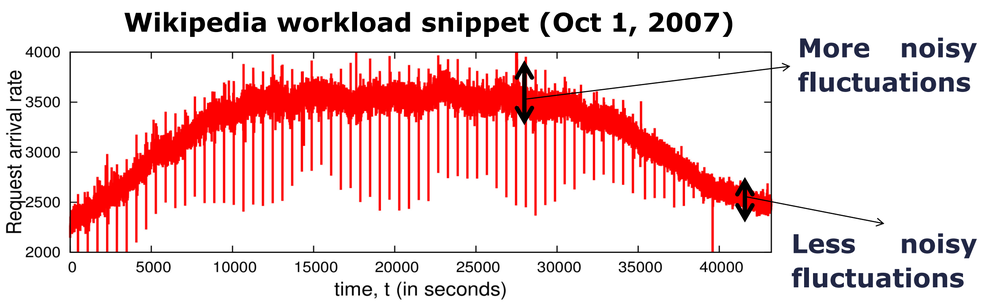

2. Goals

• **Realistic modeling of fine-scale bursty workloads** so that the statistical regularities are resembled in the generated prototype.
• **Investigating the impact** of fine-scale burstiness on elasticity.

3. Regularity Behavior

• **Regularity behavior**: **how randomness evolves** over time.
• Holder exponent, h: determines regularity at a point t with respect to its neighborhood.

Regularity Behavior

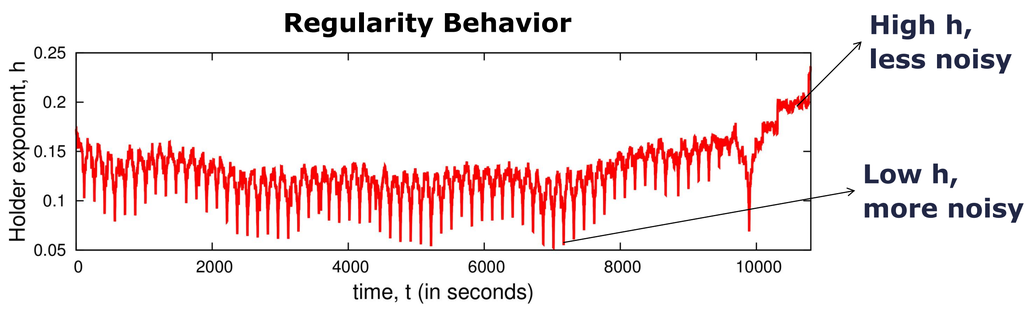

4. Realistic Workload Model

• Preserves **actual workload's regularity behavior** in prototype**.**
• Preserves amplitude-regularity correlation: existing approach (-0.52), our approach(-0.617) , actual workload(-0.659).

How to model fine-scale bursty workloads?

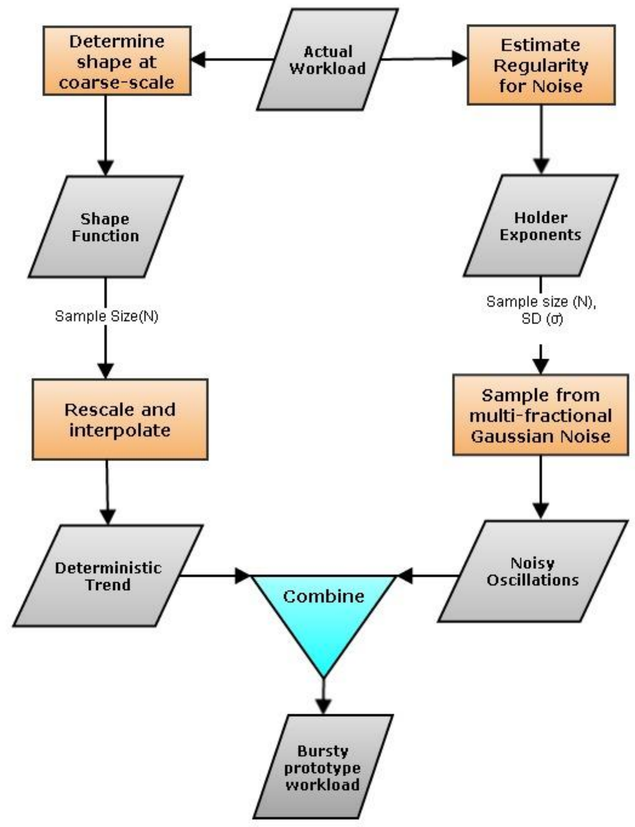

Fine-scale bursty prototype (Our approach)

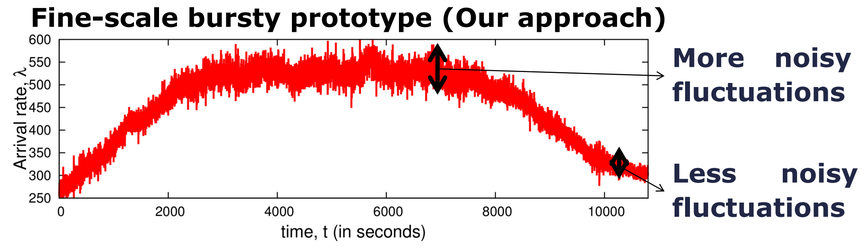

Fine-scale bursty prototype (Existing approach)

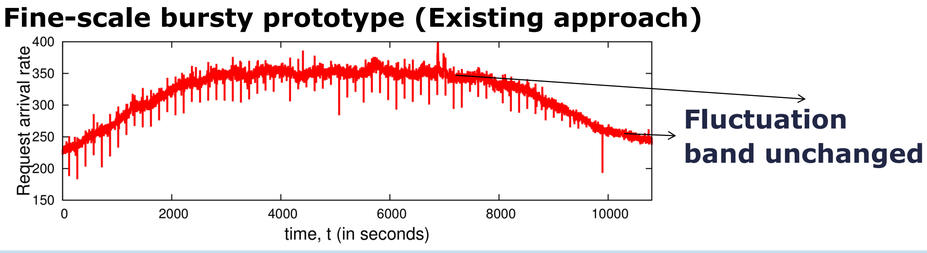

5. Experimental Setup

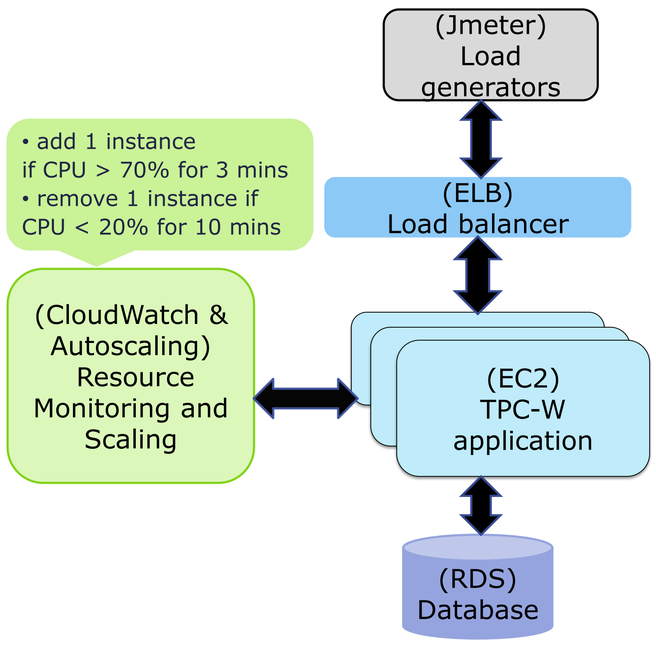

6. Detrimental Effect on Elasticity

1. High request rates balanced by low request rates → Reduced CPU utilization → 2nd Scale-out delayed
2. Highly variable arrival rates → Scarce resources, only a fraction of requests served → Retried requests built up overwhelming demand surge

Non-bursty workload (reference) vs. Fine-scale bursty workload

Application server behavior

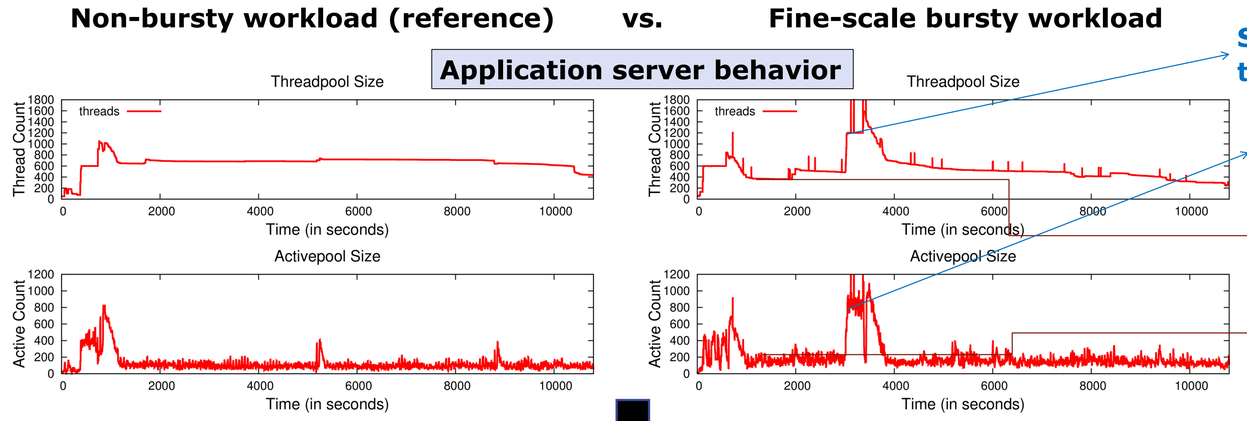

Impact on Elasticity

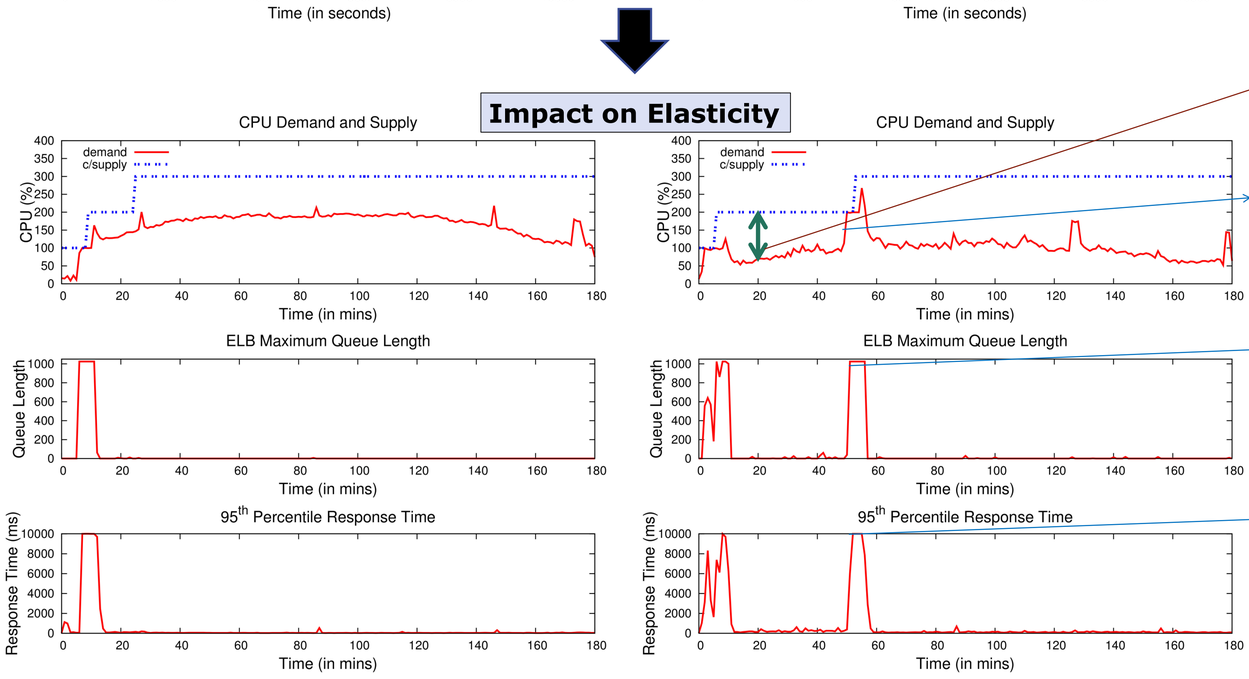

Key Insights

• Fine-scale burstiness
  • causes **highly volatile demand** pattern.
  • increased over- and under-provisioning.
  • **negative impact** on elasticity.

• Requirements for good elasticity
  • fine-grained monitoring metrics.
  • more agile adaptive scaling.

http://sisl023.web.cse.unsw.edu.au/burstiness