

Realistic Request Arrival Generation In Storage Benchmarks

**Rekha Pitchumani,
Shayna Frank, and Ethan L. Miller**

Center for Research in Storage Systems
UC Santa Cruz



Key-Value Stores

- Vital component of cloud computing applications and high performance web scale databases
- A simple and versatile interface
 - Applicable to both the large distributed stores and the individual storage nodes
 - A device agnostic interface

Trace Based Evaluation

- No key-value workload trace is publicly available
- Only published study is that of an in-memory key-value caching layer, Facebook's Memcached deployment
 - Not publicly available
 - Not suitable for all systems

Yahoo! Cloud Serving Benchmark

- The standard benchmark of choice for evaluating key-value systems
 - Used both in the evaluation of large distributed key value stores, and individual key-value storage nodes
- Workload generator offers different mixes of operations and data sizes, and key selections based on different distributions
 - But no option to generate realistic request arrivals - only constant rate and throttled arrivals

Realistic Request Arrivals

- Needed for many systems
 - scrubbing, cache de-staging, data migration across tiers, automatic backups, garbage collection and data reorganization scheduling in Flash and SMR, tail latency reduction, load tolerance
- Temporal characteristics has been ignored by most benchmarks

YCSB - Real World Problem Evaluation


- YCSB used by simply adding and removing clients generating request at a constant rate
 - Elasticity controller designed to automatically respond to changes in workload
 - System for achieving datacenter-wide per-tenant performance isolation and fairness
 - Systems performing live database migration to tolerate load variations in multi-tenant databases

Arrival Process Models

- Existing studies show that disk, filesystem, network, and web traffic all exhibit some common temporal properties
 - Such as burstiness, self similarity, long range dependence, and diurnal activity
- Classify typically observed temporal patterns into three kinds of arrival processes
 - Poisson
 - Self similar
 - Envelope-guided

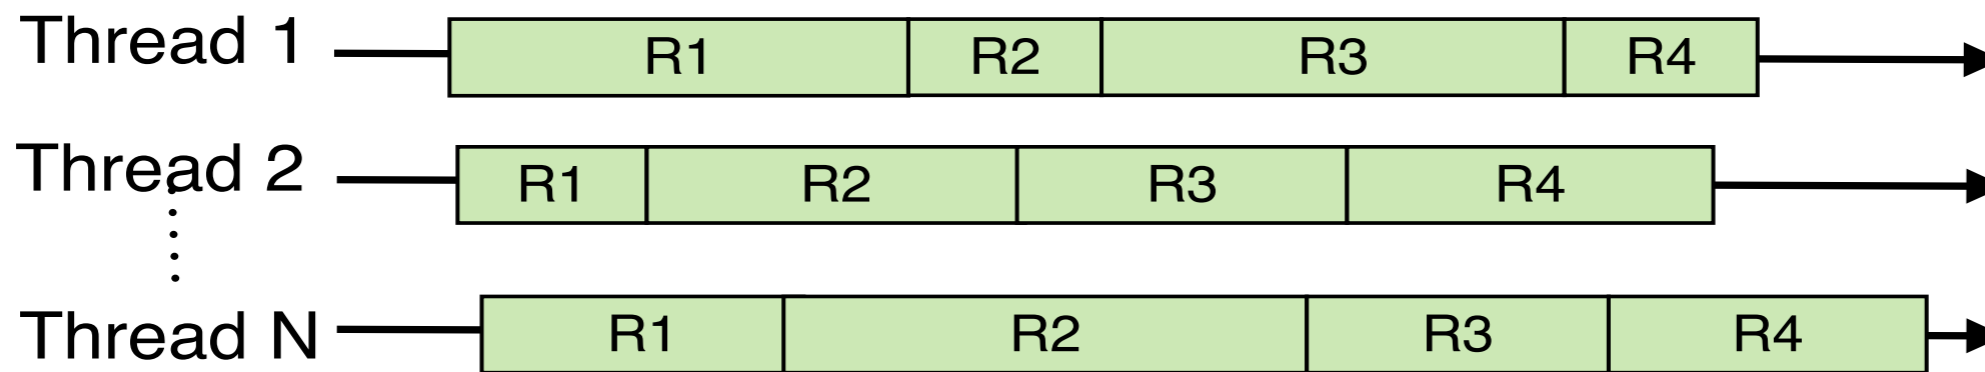
YCSB Client Architecture

- A thread-safe workload generator that generates requests according to user specifications
- A workload executor to execute the generated requests
- A database interface layer to connect with and pass requests to the database
- A separate thread to collect and report the status.

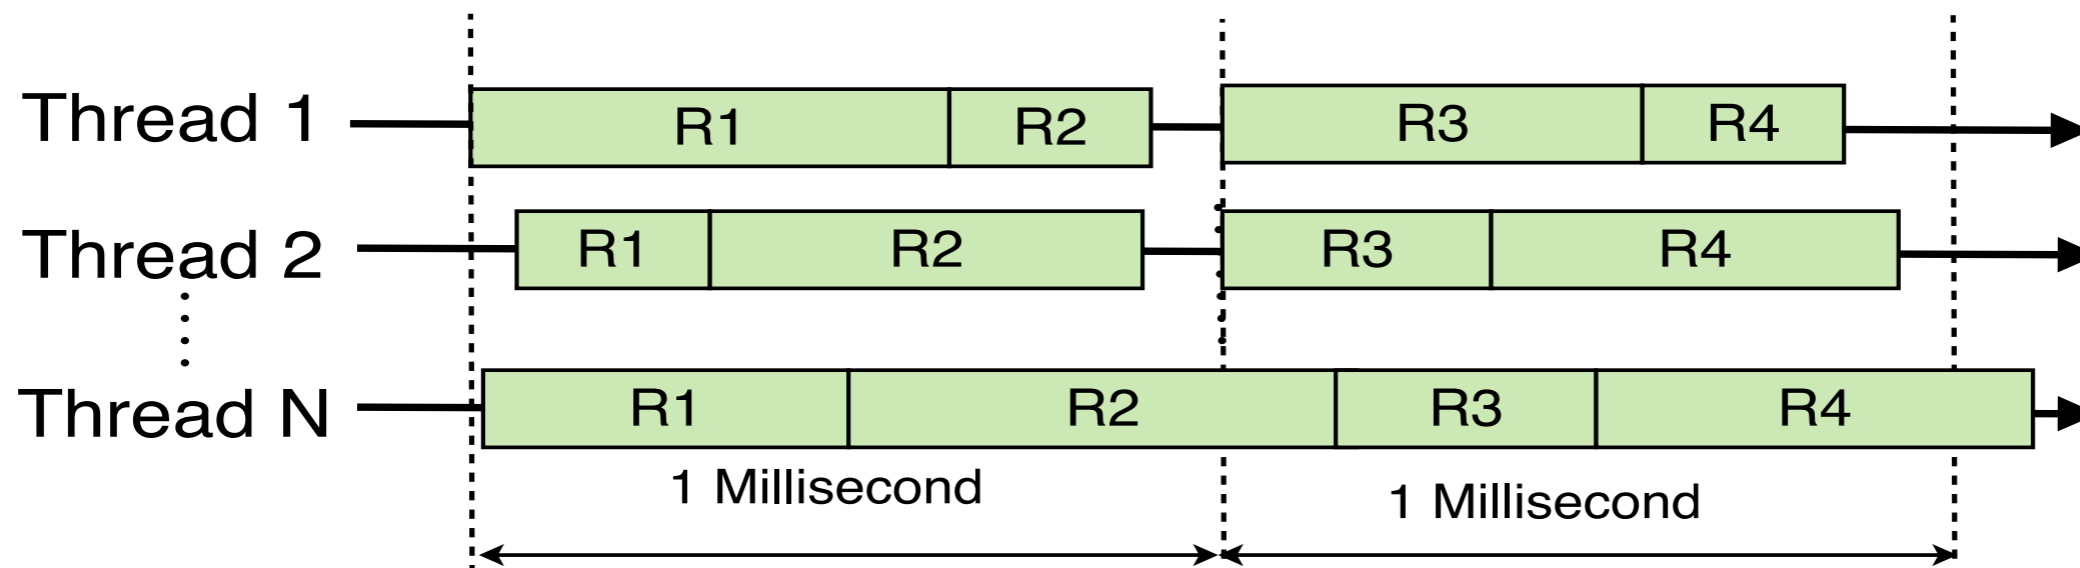


Request Execution

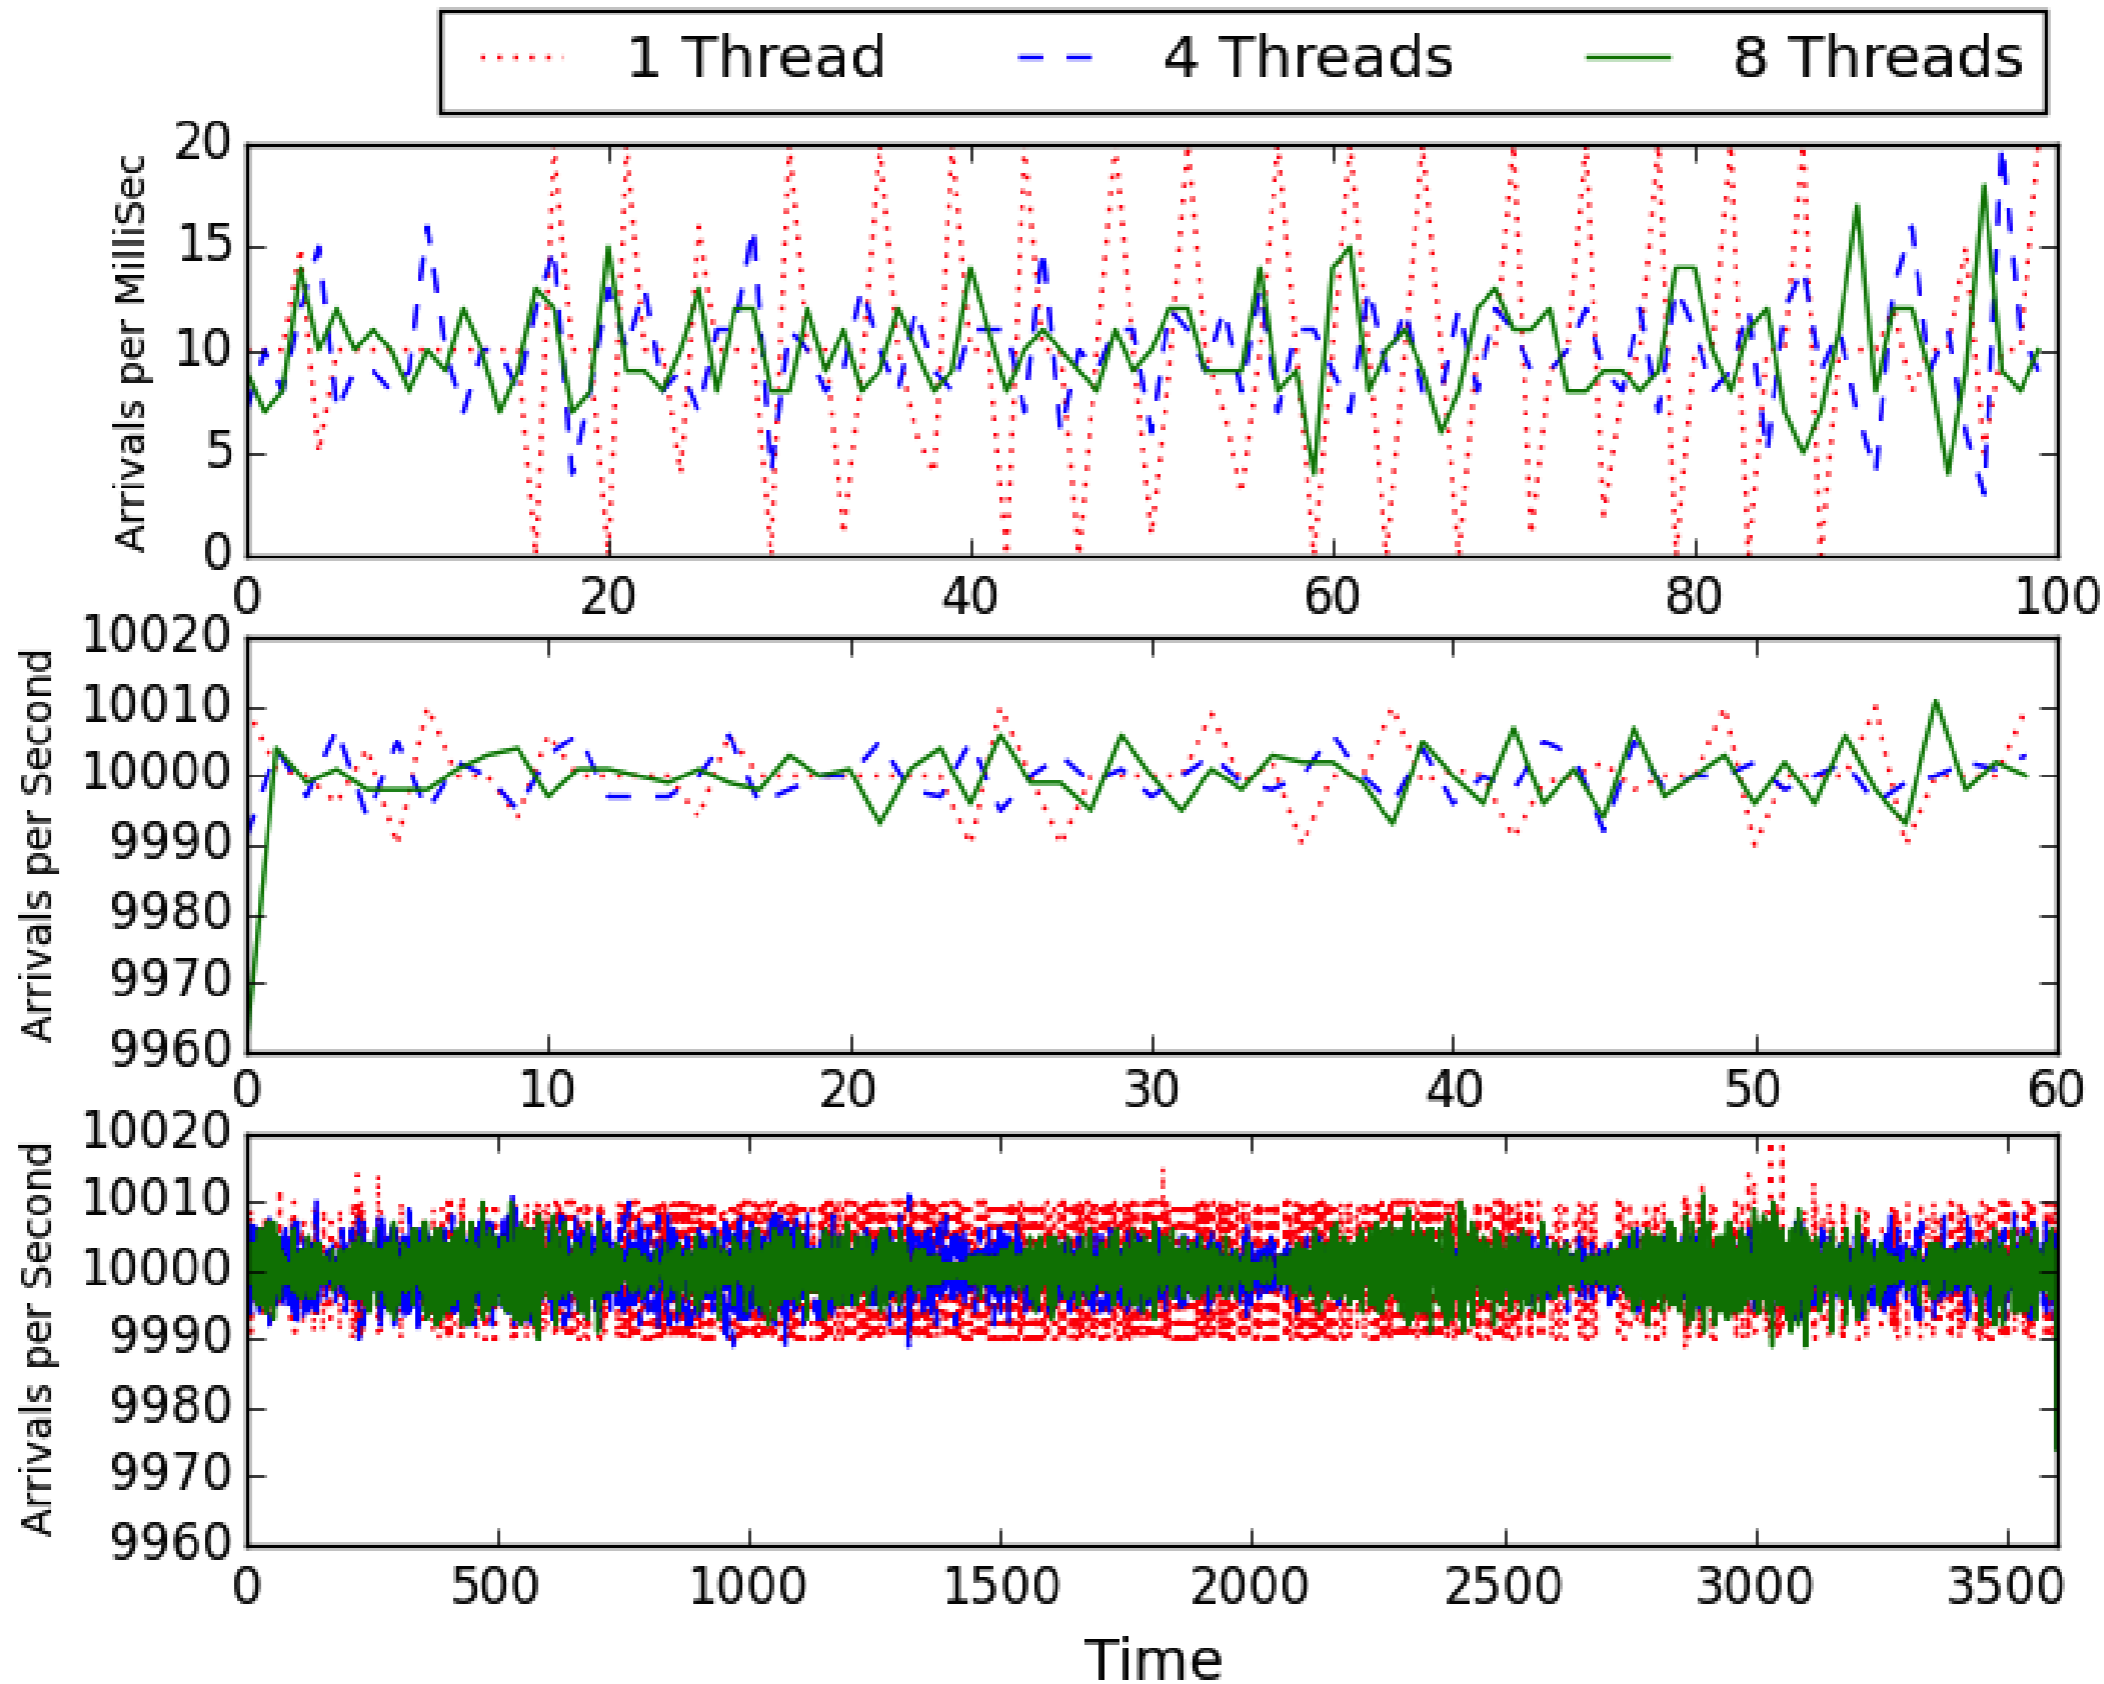
- By default, threads perform back-to-back synchronous IO



- Throttling for a target load of 2 requests per millisecond per thread




Original Arrivals



Modified Workload Executor


- Assume that the request sizes are independent of arrival times
- To facilitate generating bursts of IO activity at specified time intervals



Poisson Process

- A simple and widely used stochastic process for modeling arrival times
- Requests can be modeled as a poisson process if the request inter-arrival times are truly independent and exponentially distributed
- Most arrivals are correlated, and cannot be modeled accurately by a poisson process
 - But when many different kinds of independent workloads are run on a system, the resulting traffic could look like a poisson process.

Poisson Arrivals




Self Similar Process


- Self similarity means the series looks similar to itself at different time scales
 - bursts of increased activity and similar looking bursts at many different time scales
- Long range dependence means the series is correlated to not just its immediate past, but also its distant past
- Many real life observed workloads are both self similar and long range dependent
 - WWW traffic, Ethernet local area network (LAN) traffic, file-system traffic and also in disk-level I/O traffic

B-model

- A single characteristic parameter, bias b
- Split the work recursively into two portions in a proportion determined by the bias b
- b closer to 1 generates traffic with high local irregularity and $b = 0.5$ results in uniform traffic




Self Similar Arrivals



Envelope-Guided Process

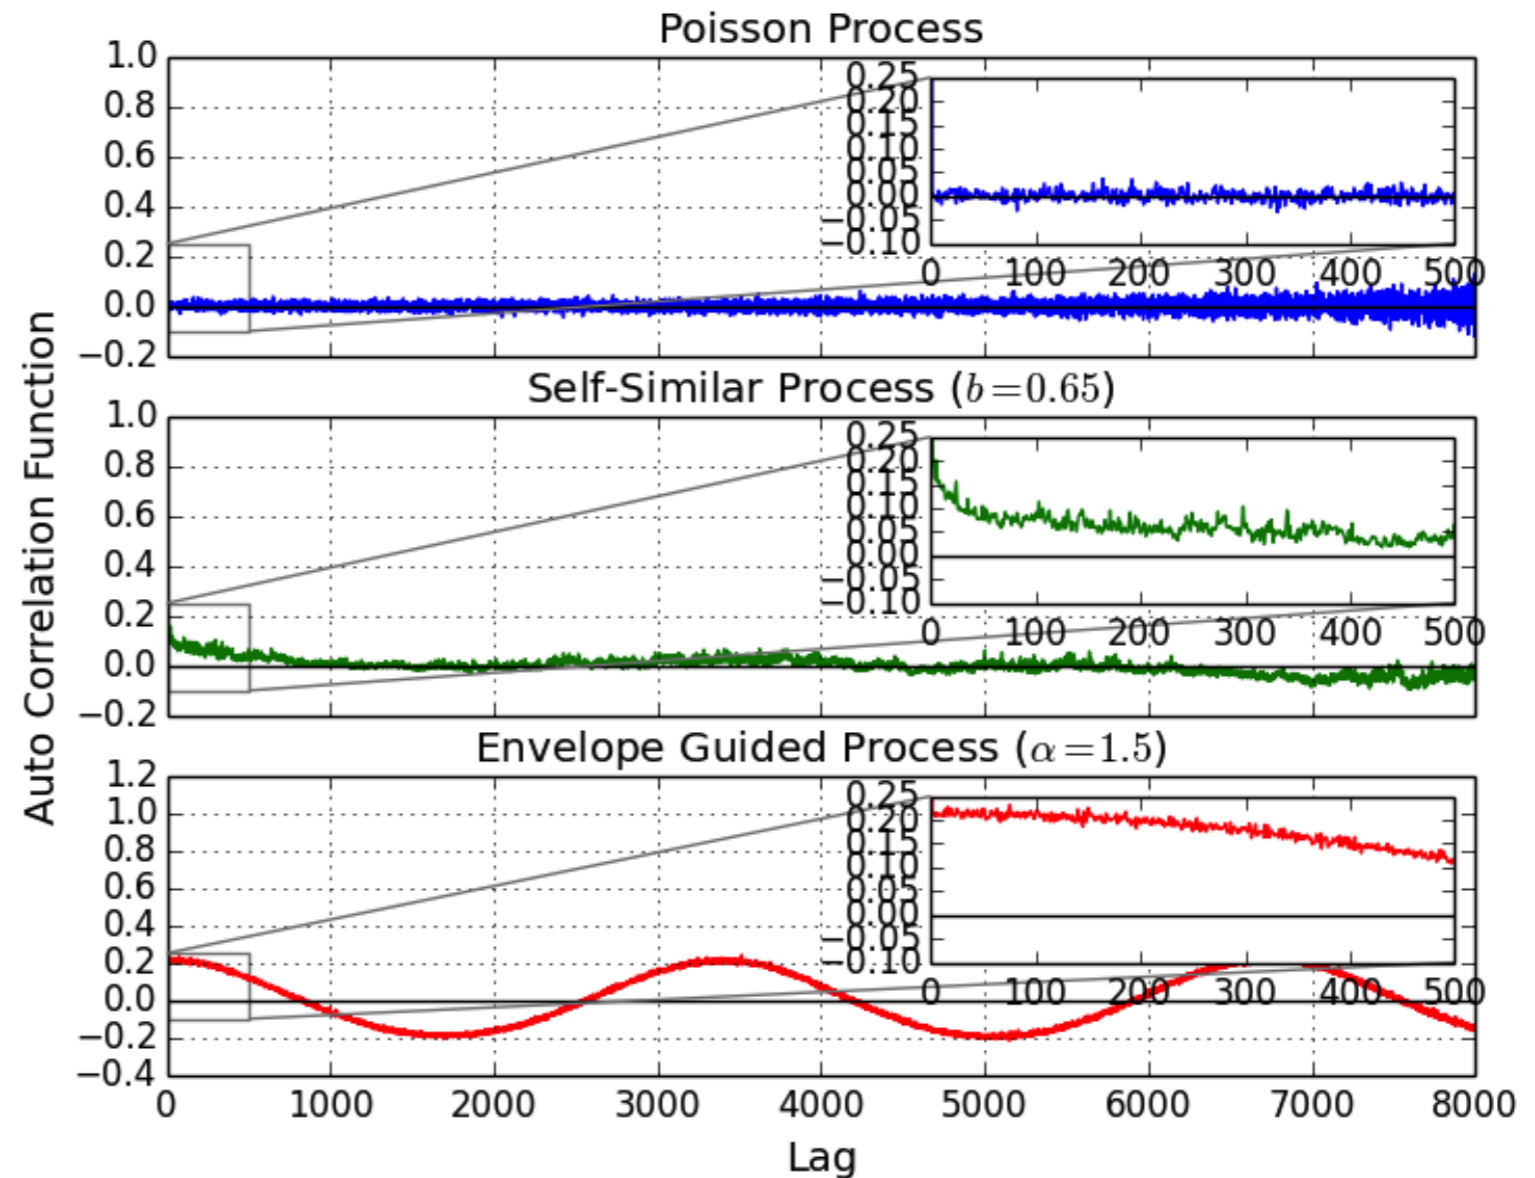
- At longer time scales, many workloads exhibit a clear diurnal pattern
 - But may exhibit self similar properties at smaller time scales
- Overall arrival rate determined by an envelope arrival rate function and burstiness determined by a secondary distribution selected
 - ‘Pseudo self similar’ processes using Pareto inter-arrivals

Envelope-Guided Arrivals



AutoCorrelation Function

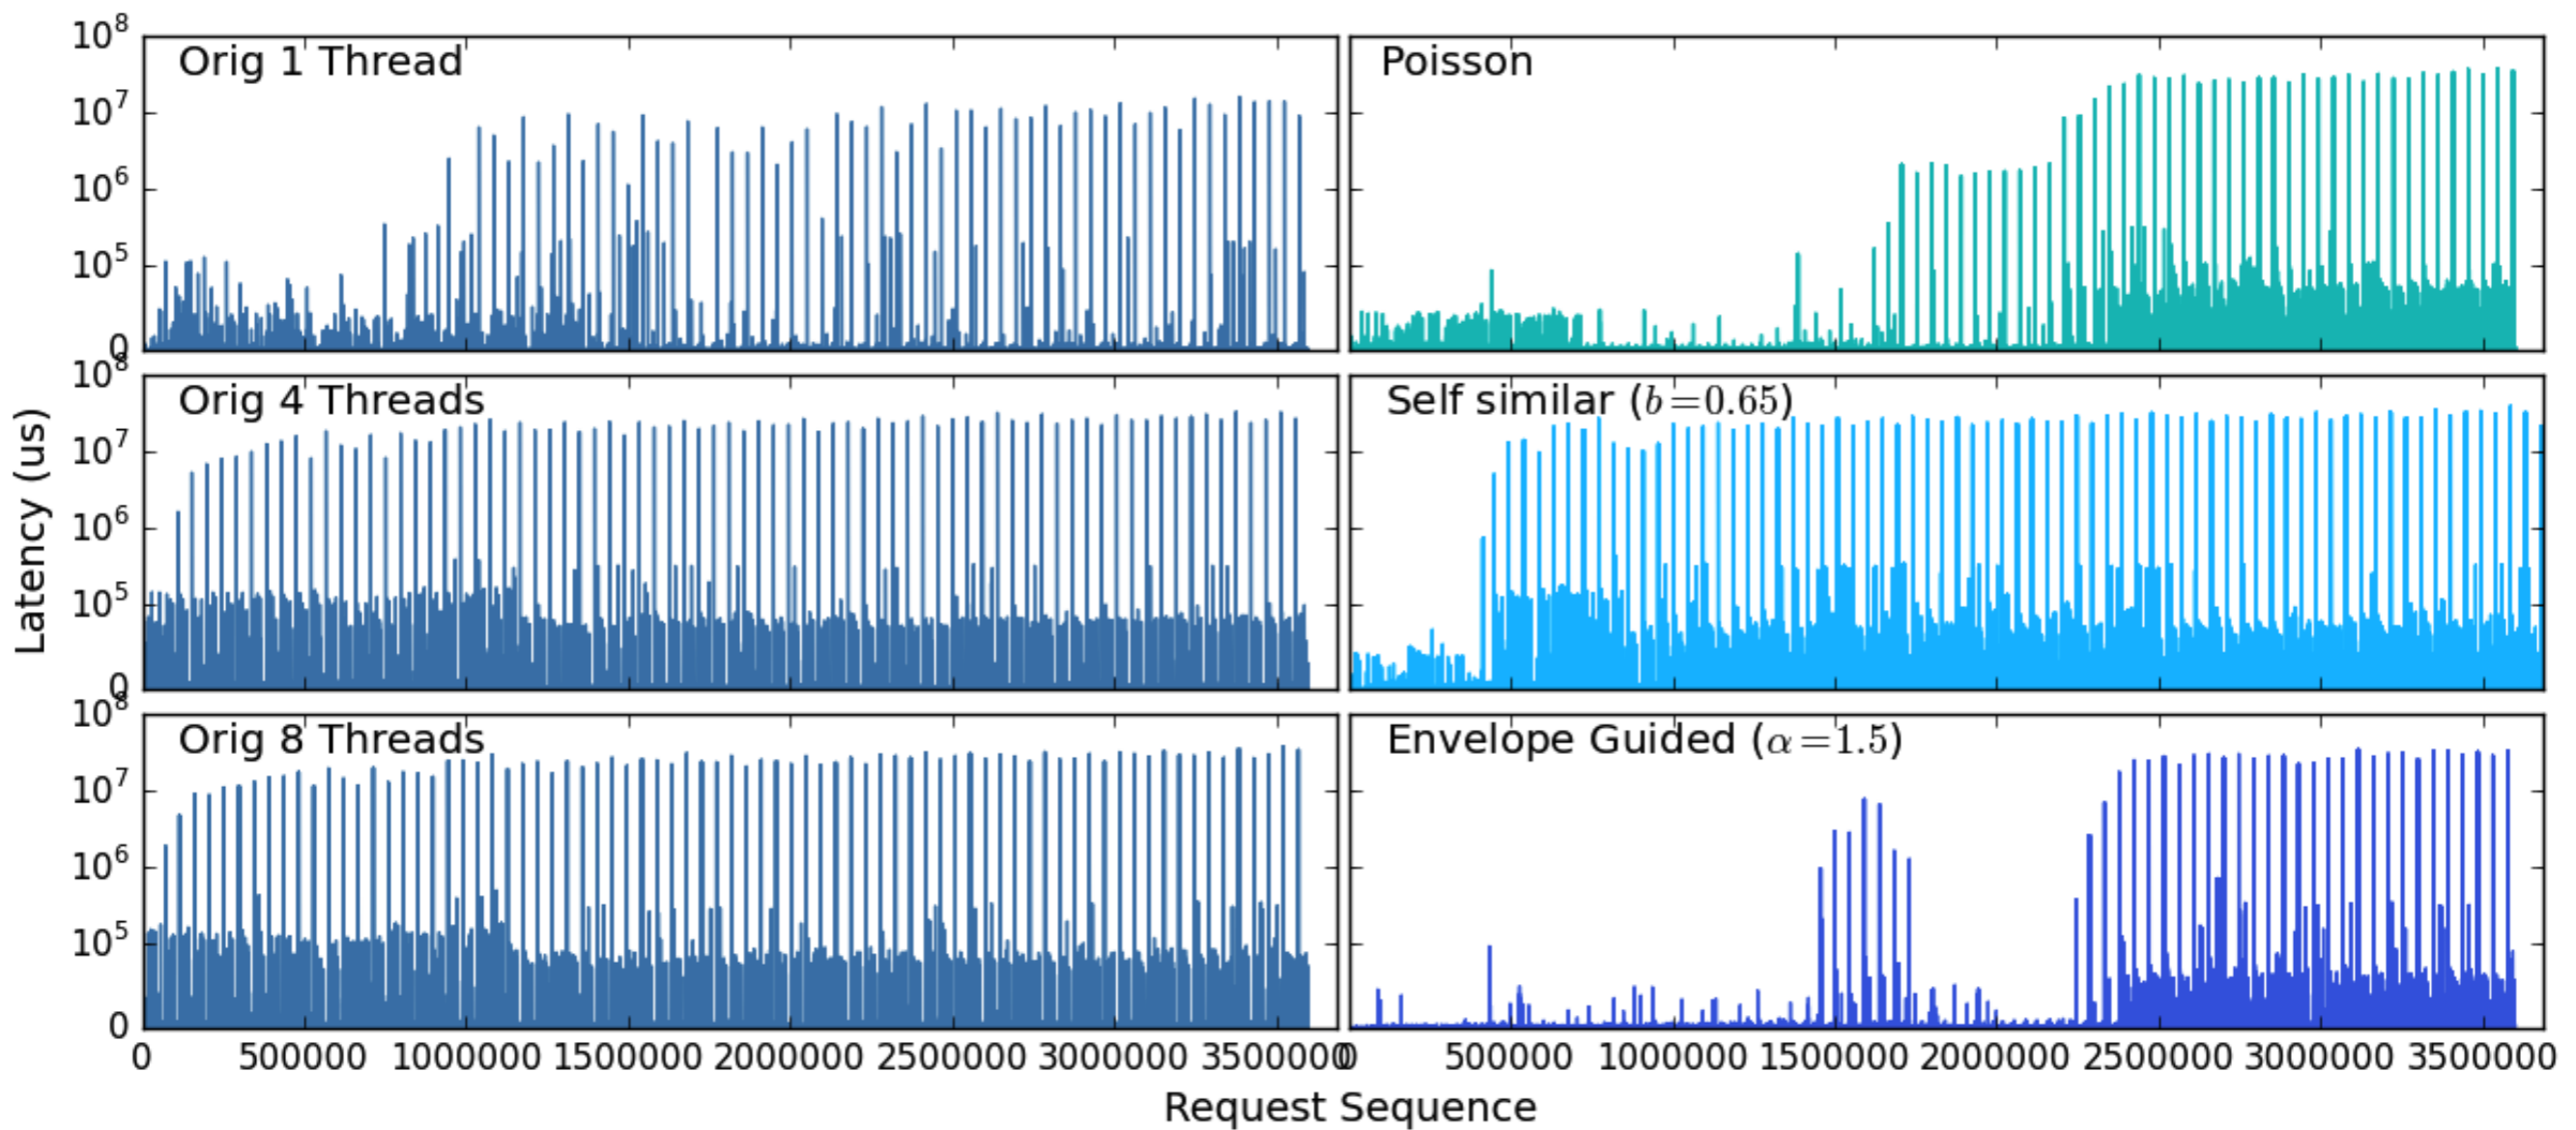
- Autocorrelation is the cross-correlation of a time series with itself
- ACF of a poisson process is usually low and close to zero even at lag 1
- ACF of a long range dependent process decays slowly
- If the traffic is periodic, ACF oscillates corresponding to the periodicity of the original time series



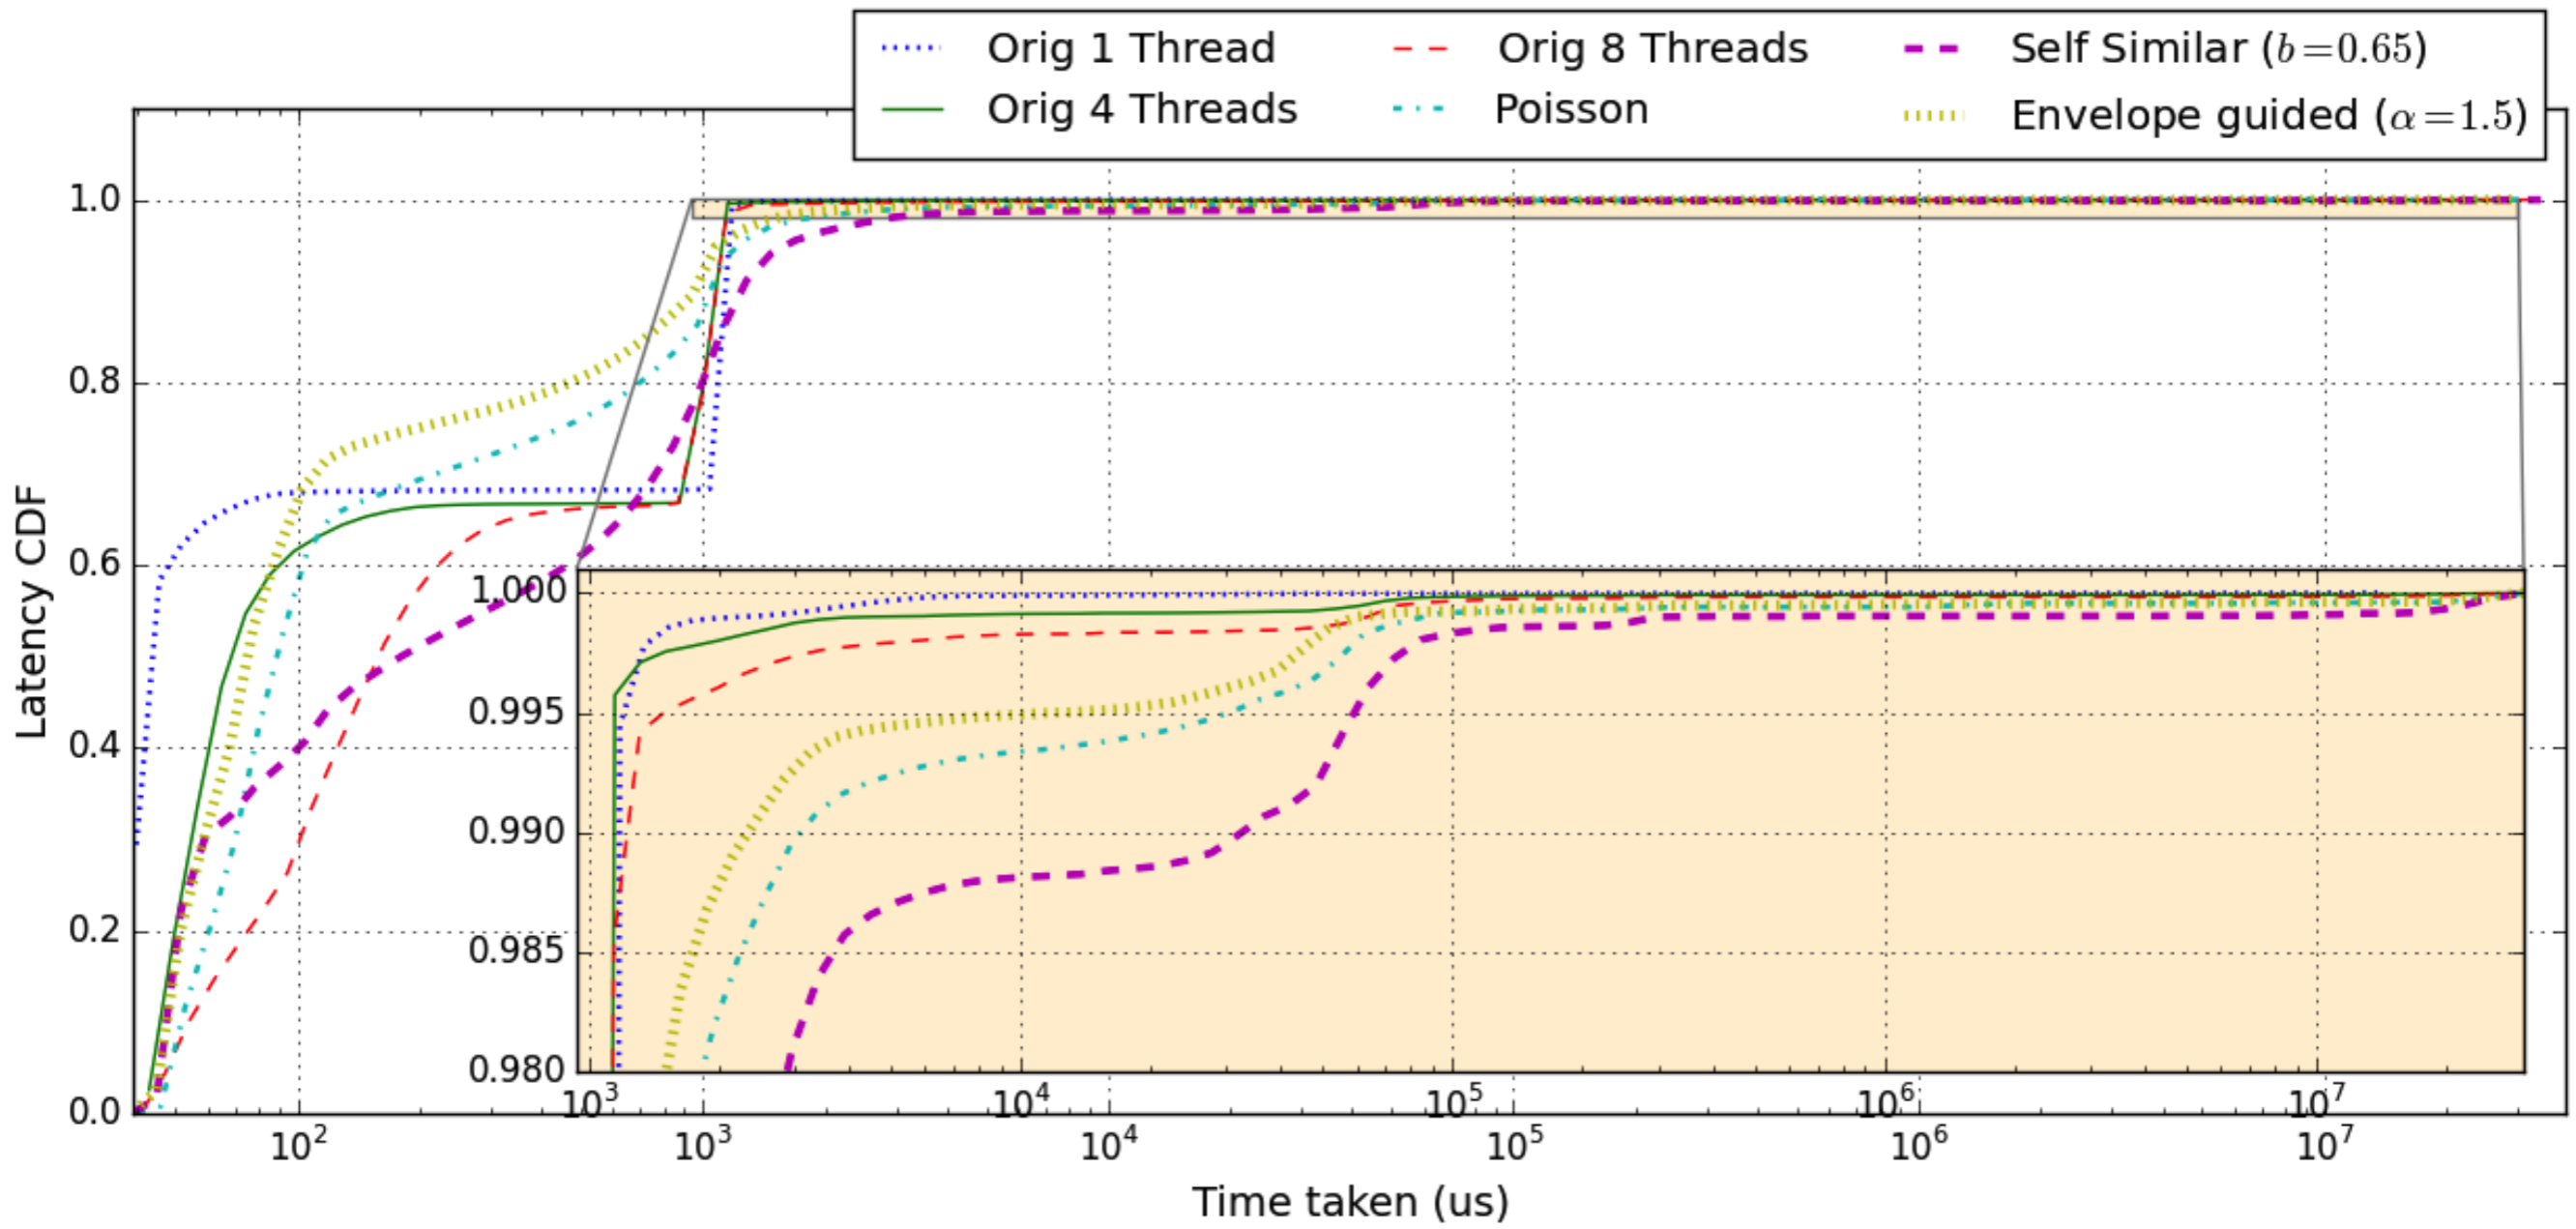
Demonstration

- LevelDB
 - Log-Structured Merge Tree-based embeddable KV database engine
 - Periodic background compaction process both cleans and reorganizes data
- Run on a dedicated 160GB single platter Seagate disk
 - MapKeeper server for communication

LevelDB



LevelDB



Future Work

- Scaled-up Realistic Request Arrivals
 - User generative model - each client chooses a model representative of a distinctive user of the shared storage system
 - Multi-client co-ordinated request arrivals would be better
- Content Generation
 - Realistic content with configurable levels of duplicity and compressibility
 - Realistic variable length key content
- Correlation In Request Sizes
 - Study if truly uncorrelated to arrival rate

Conclusion

- Workload's temporal characteristics are a big influencer on the system behavior
- We presented three categories of arrival process models that all benchmarks should provide
 - Poisson
 - Self similar
 - Envelope-guided
- We have implemented the models in YCSB and demonstrate the effect using database evaluation

Thank You

Rekha Pitchumani

rekhap@soe.ucsc.edu

<http://www.ssrc.ucsc.edu/person/rpitchumani.html>