



Poor Man's Social Network

Consistently Trade Freshness For Scalability

Zhiwu Xie, Jinyang Liu, Herbert Van de Sompel,
Johann van Reenen and Ramiro Jordan



Outline

- Scaling feed following
- Algorithm
- Experiment and results
- Conclusions



Feed Following Scalability

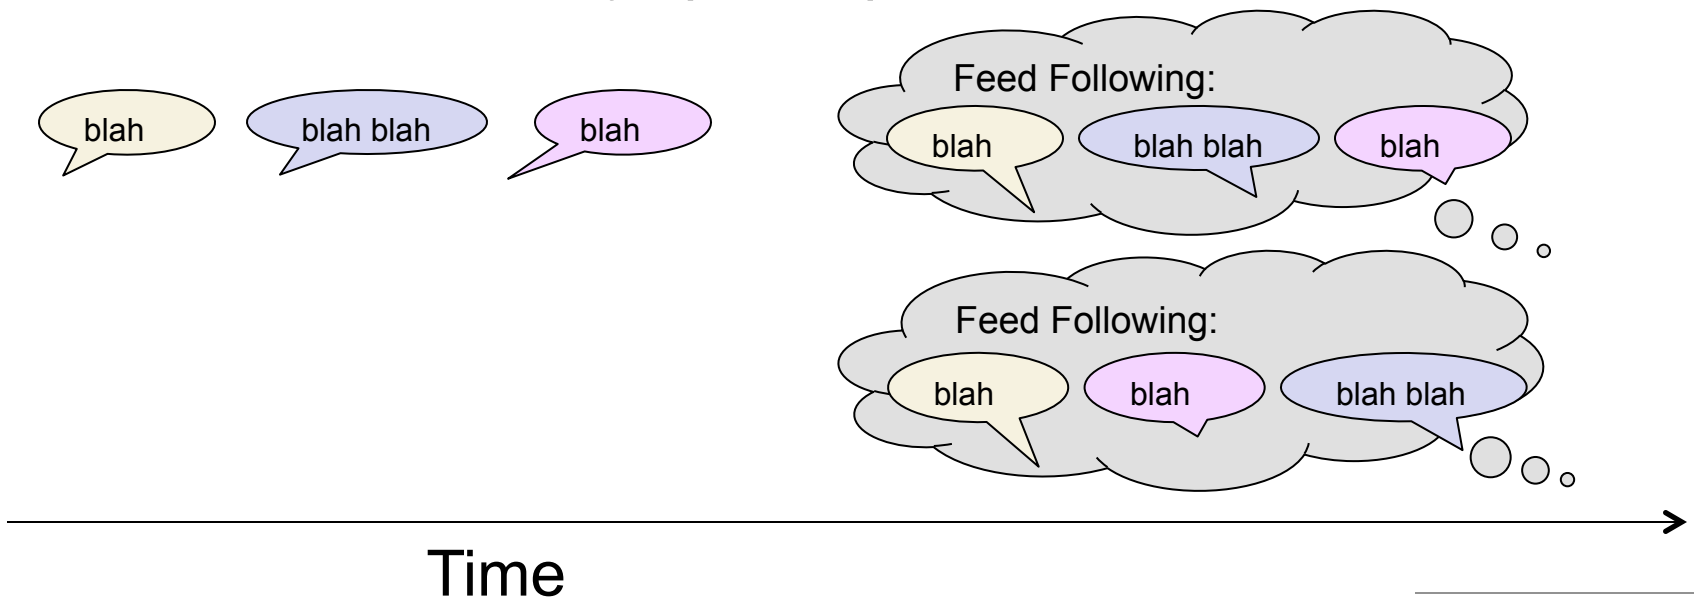
Give me the 20 **most recent** tweets sent by **all** the people **I** follow

- Individualized queries
- Fast changing global state
- Partitioning, replication, and caching
- NoSQL: trade consistency for scalability

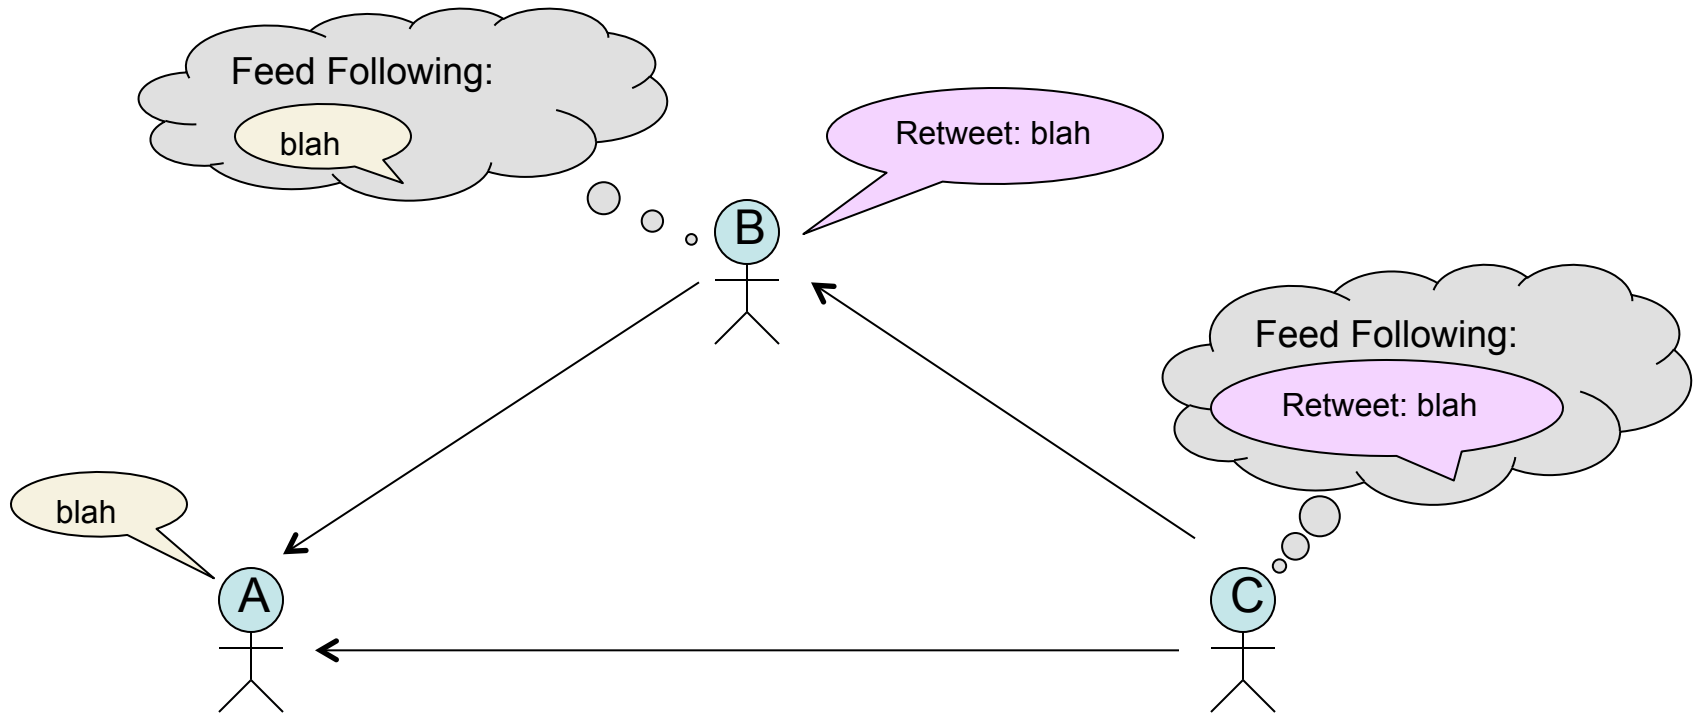


Consistency

- Atomicity, Linearizability, or One-copy Serializability (1SR)



Retweet Anomaly





New Approach: TimeMap Query

Who have created new tweets during the **past scheduled release** periods?

- Global time across partitions
- Schedule releasing
- Client-side processing and caching
- Consistently trade freshness for scalability



CAP Theorem

- Preconditioned on the asynchronous network model: the only way to coordinate the distributed nodes is to pass messages
- In the partially synchronous model, where global time is assumed to be available, CAP may indeed be simultaneously achievable most of the time

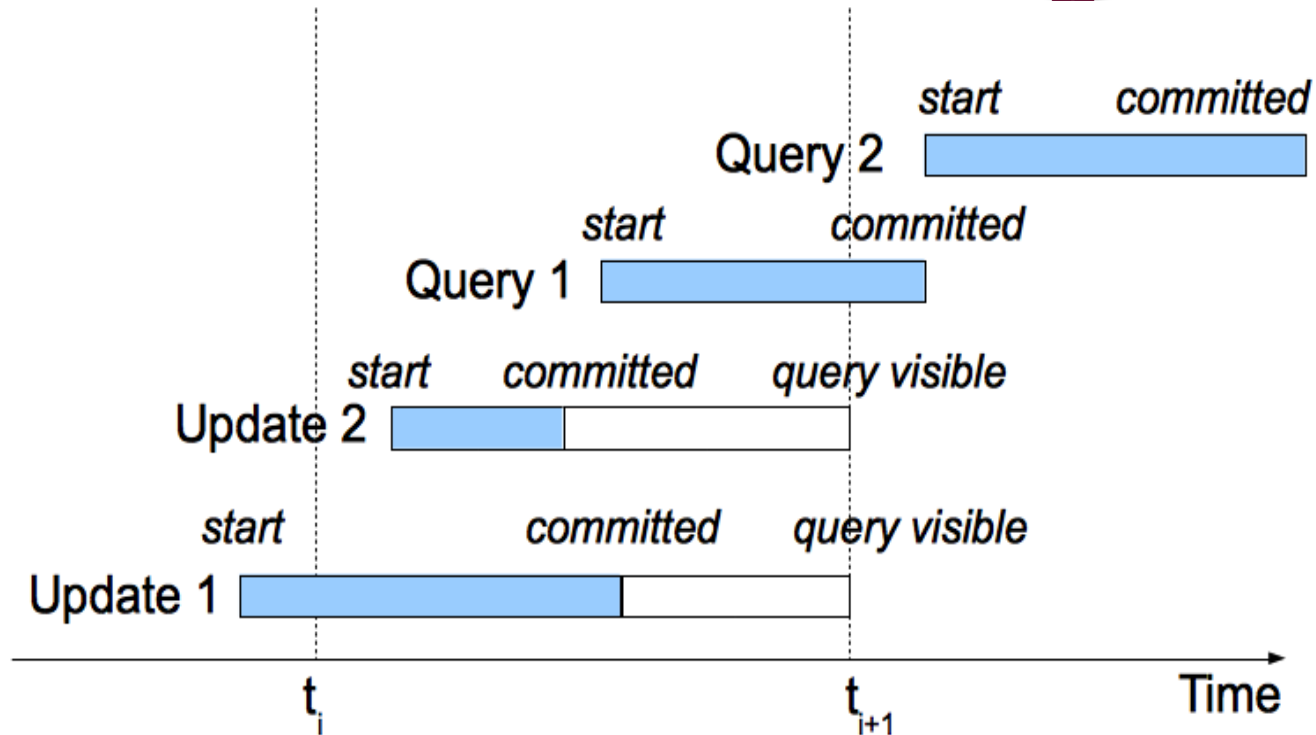


Global Time

- “One of the mysteries of the universe is that it is possible to construct a system of physical clocks which, running quite independently of one another, will satisfy the Strong Clock Condition.”
 - Time, Clocks and the Ordering of Events in a Distributed System, by Leslie Lamport



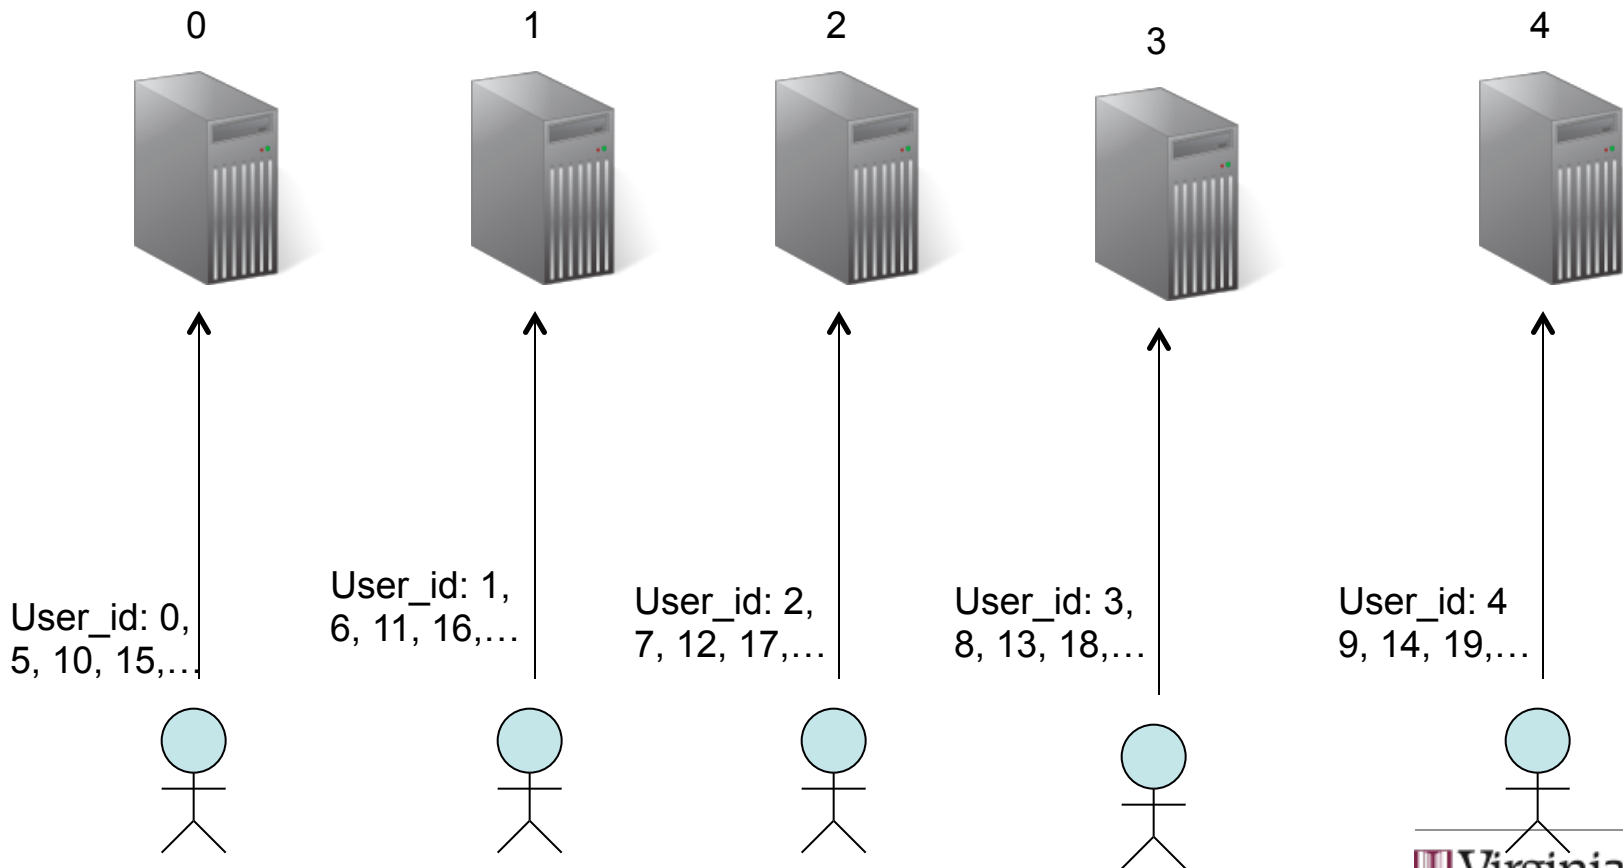
Scheduled Release Algorithm



Who have created new tweets during the **past scheduled release** periods?




Partitioning: Send A New Tweet






Partitioning: TimeMap






Client Side Processing





If the current time is 1:05:37PM, please tell me who (no matter if I follow any of them or not) have sent new tweets from 1:05:30PM to 1:05:35PM. I'll figure out by myself if any of these new tweets are relevant to me, and if so, I'll retrieve these tweets separately by myself.



If the current time is 1:05:37PM, please tell me who (no matter if I follow any of them or not) have sent new tweets from 1:05:30PM to 1:05:35PM. I'll figure out by myself if any of these new tweets are relevant to me, and if so, I'll retrieve these tweets separately by myself.



Staleness vs. Latency





Trade Freshness For Scalability

- Mass transit system vs. private car
- Lose flexibility, but gain overall efficiency by sharing resources
- Stale up to the length of the schedule release period, e.g., 5 seconds.



Experiment

- Implemented on AWS
- A Twitter like feed following application
- Server side: Python/Django, PostgreSQL, PL/pgSQL
- Client side: emulated browser, implemented in Python/Django and PostgreSQL



Experiment: Configurations

- Used ~ 100 cloud instances from Amazon
- Most are used for emulated browsers
- 3 to 6 c1.medium as servers
- Use memcached to simulate caches

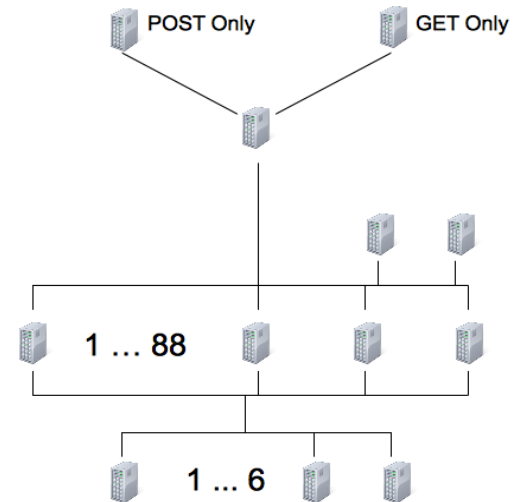
Load Generator
c1.xlarge
Httpperf + autobench

Load Balancer
c1.xlarge
haproxy

Cache
m2.2xlarge
memcached

Emulated Browser
c1.xlarge
Django + httpd +
PostgreSQL

Server
c1.medium
Django + httpd +
PostgreSQL + memcached





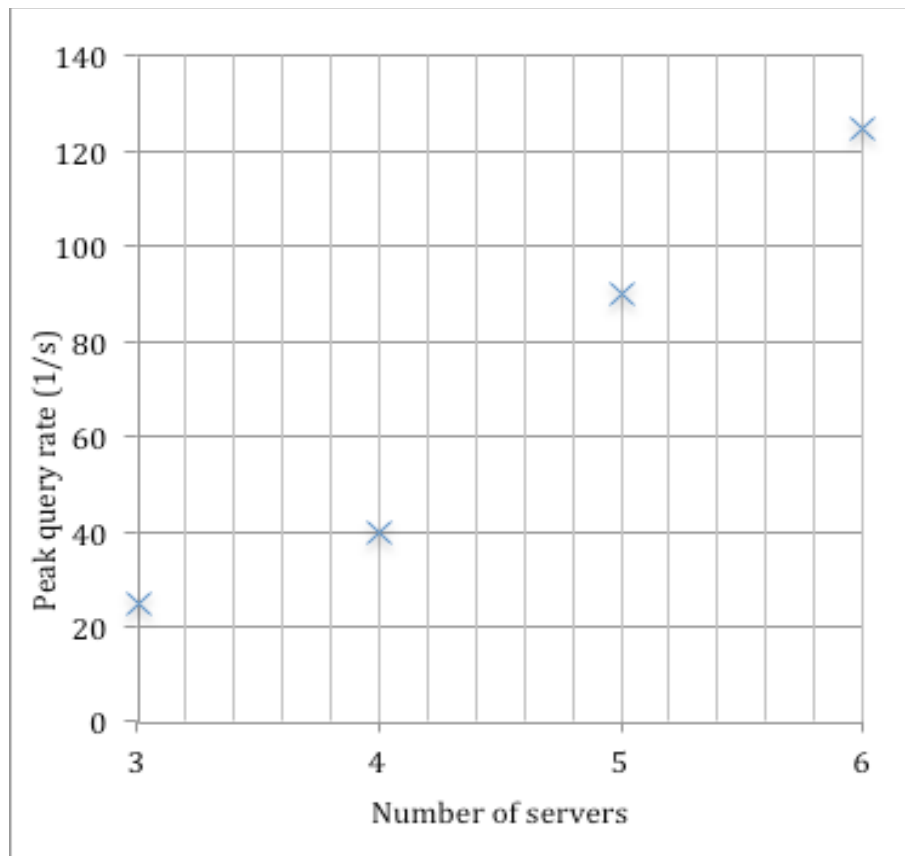
Experiment: Workload

- Work load similar to the Yahoo! PNUTS experiment
- A following network of ~ 200,000 users
- Synthetic workload generated by Yahoo! Cloud Serving Benchmark

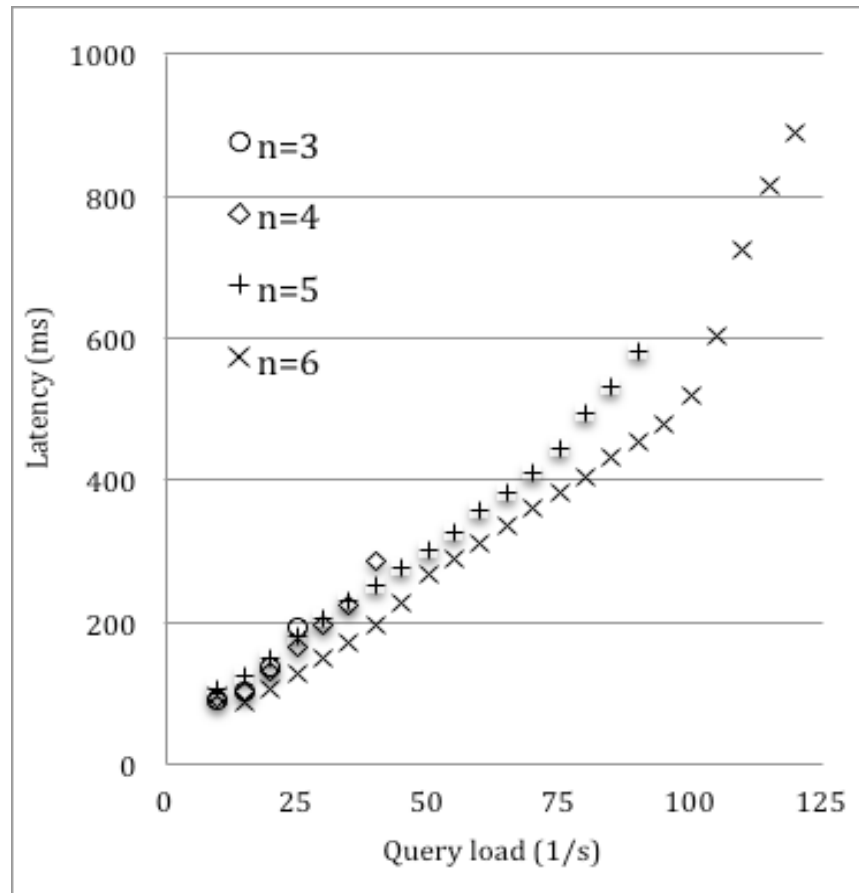
		PNUTS	This
Number of producers		67,921	67882
Number of consumers		200,000	196,283
Consumers per producer	Average	15.0	13.38
	Zipf parameter	0.39	0.39
Producers per consumer	Average	5.1	4.63
	Zipf parameter	0.62	0.62
Per-producer rate	Average	1/hour	1/hour
	Zipf parameter	0.57	0.57
Per-consumer rate	Average	5.8/hour	varied
	Zipf parameter	0.62	0.62



Experiment Result: Query Rate

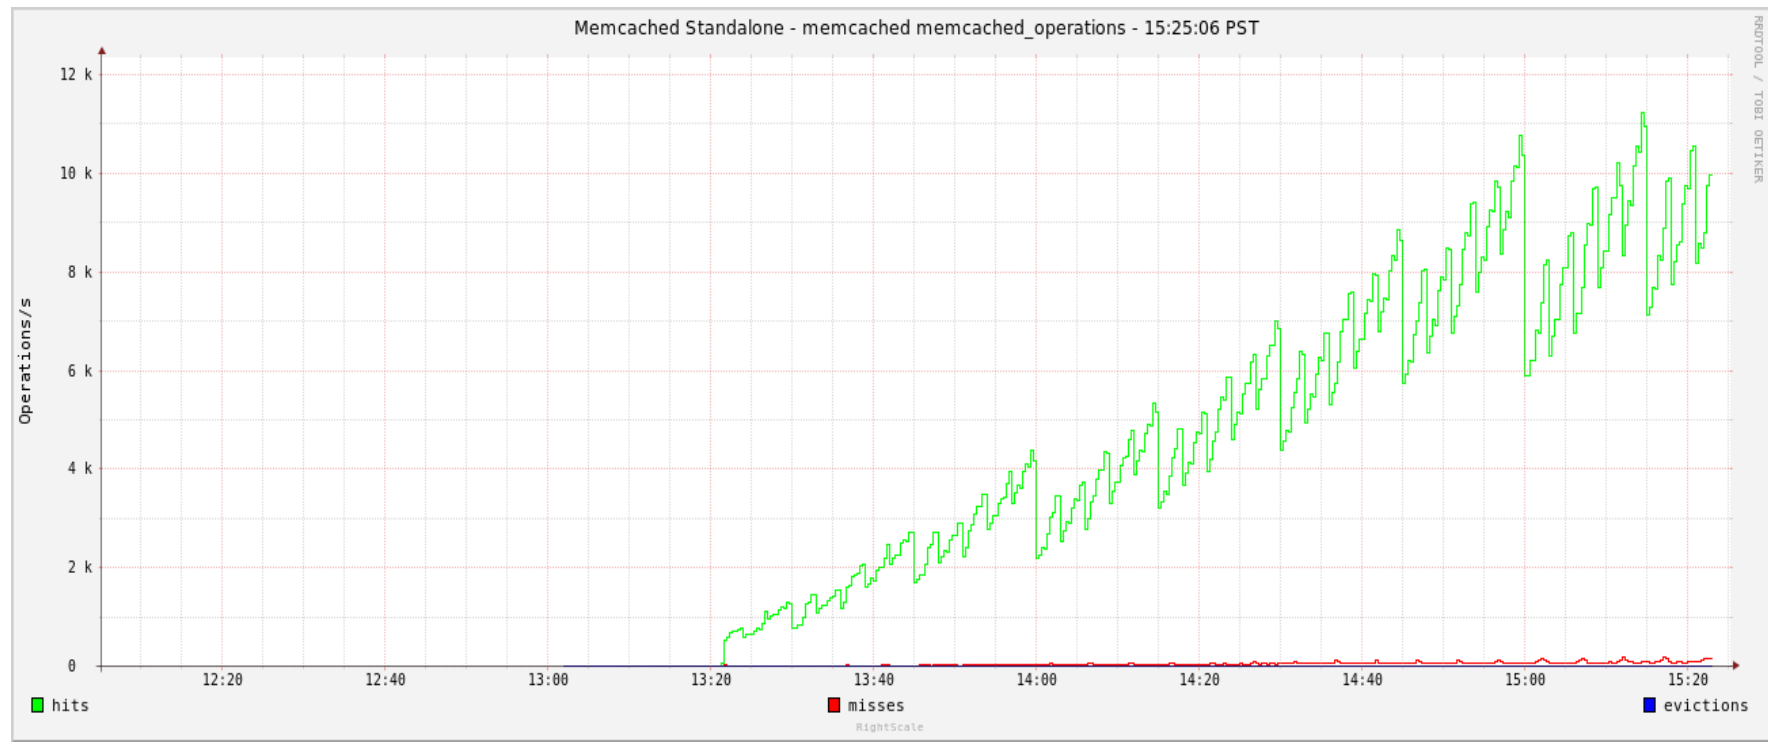


Experiment Result: Latency





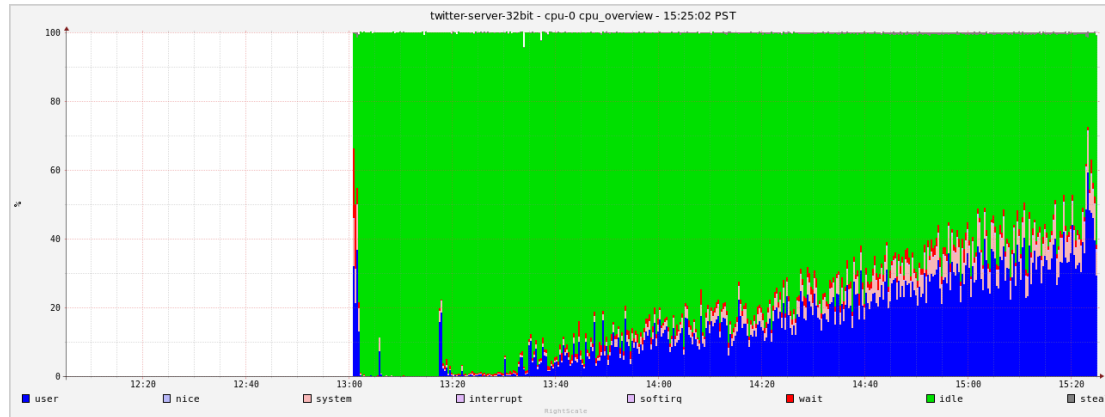
Experiment Results: Caching



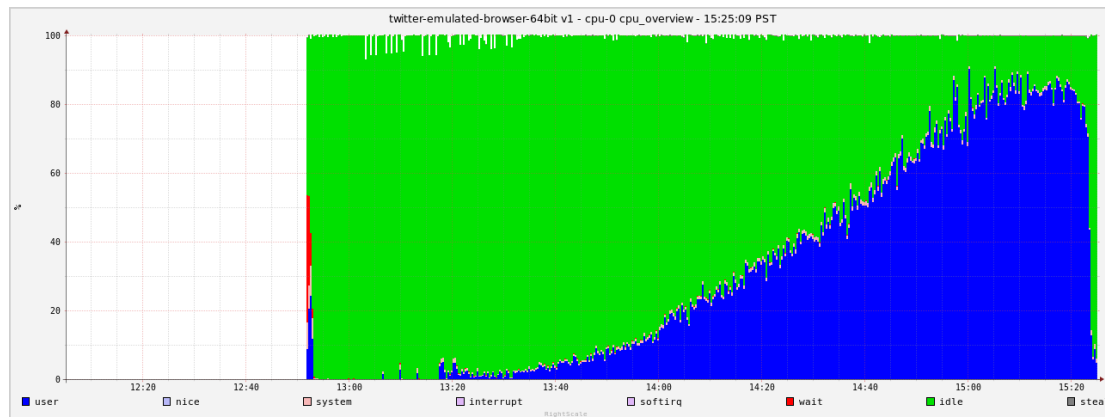


Experiment Results: CPU Load

Server



Client





Conclusions

- Consistently scale feed following
- Linear scalability
- Practical low cost solution



Thank You

- Questions?