On Sample-Path Staleness in Lazy Data Replication

Xiaoyong Li, Daren B.H. Cline and Dmitri Loguinov

Internet Research Lab
Department of Computer Science and Engineering
Texas A&M University

April 29, 2015

Agenda

- Introduction
- Staleness Formulation
- Staleness Cost
- Optimality
- Applications
- Conclusion

Introduction

- Highly-dynamic content
  - News
  - Weather
  - Road conditions
- An increasing number of applications need to maintain local copies of remote data sources
  - Search engines
  - Mash-up applications
  - Distributed caching
- Copies need to be synchronized constantly
  - To provide reliable services

Introduction

- Push-based policy
  - Sources send the update information to replicas
  - Requires the cooperation of sources
  - Hard to scale
- Pull-based policy
  - Replica retrievals the content explicitly
  - Scalable and less costly
  - Leads to staleness
- Need models and mechanisms for analyzing and controlling staleness
  - Previous works mainly consider Poisson updates

Agenda

- Introduction
- Staleness Formulation
- Staleness Cost
- Optimality
- Applications
- Conclusion

Staleness Formulation

- Information requests indicated by arrows:

- Model
  - Source experiences random updates via process N_U
  - Replica periodically downloads the content via process N_D

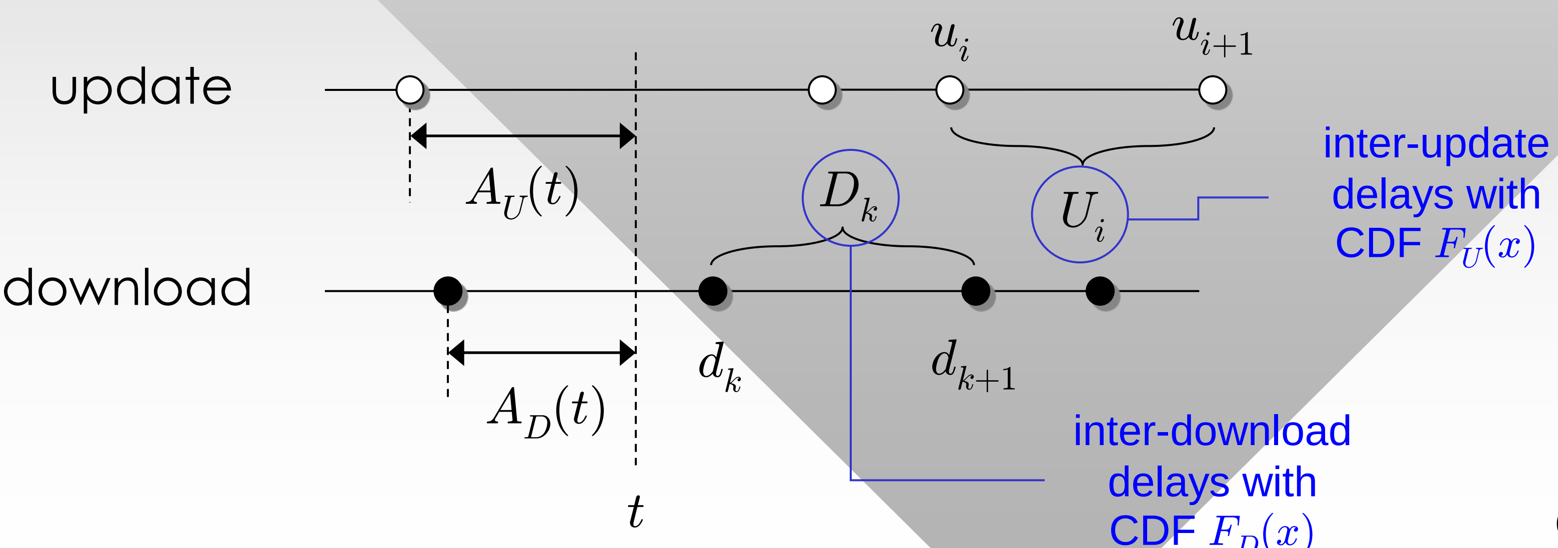

Metrics

- $M(t)$: the number of updates missing from the replica
  - $E[M(t)]$: the expected number of missing update

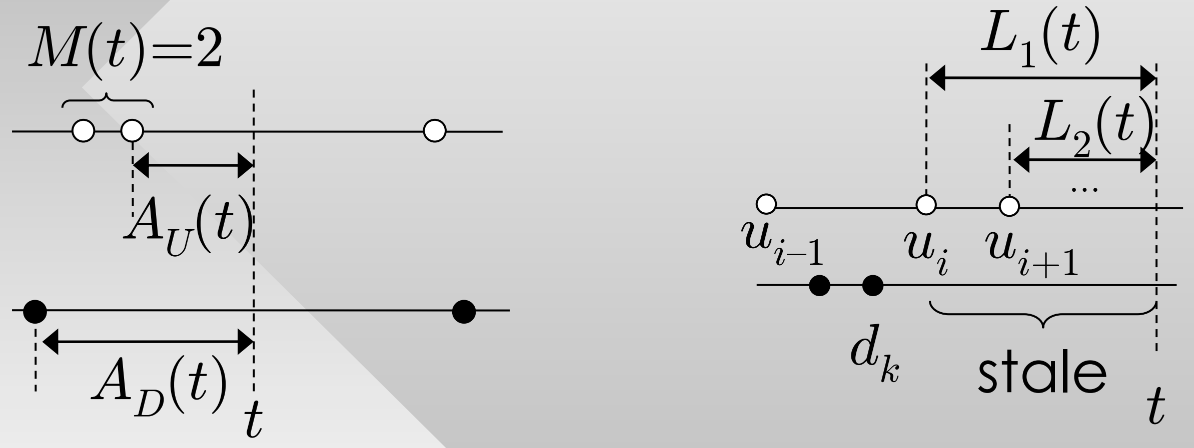

- Backward delays to each unseen update

$$L_1(t) > L_2(t) > \ldots > L_{M(t)}(t)$$

- Apply weight function $w(x)$ to each lag
  - Maps staleness lags to actual cost

Computer Science, Texas A&M University

Metrics

- Two different cost metrics:

$$\eta(t) = \begin{cases} w(L_1(t)) & M(t) > 0 \\ 0 & \text{otherwise} \end{cases} \qquad \rho(t) = \begin{cases} \sum_{i=1}^{M(t)} w(L_i(t)) & M(t) > 0 \\ 0 & \text{otherwise} \end{cases}$$

- Both $\rho(t)$ and $\eta(t)$ generalize metrics in previous work
  - $\eta(t)$ becomes *staleness* and *age* [Cho 2000] with $w(x) = 1$ and $w(x) = x$
  - $\rho(t)$ with $w(x) = 1$ and $w(x) = x$ lead to *blur* [Denev 2009] and *addictive age* [Ling 2004]
- Both $\rho(t)$ and $\eta(t)$ are random variables
  - Using expectation requires multiple sample-paths
  - One single sequence available in practice (sample-path)

$$\bar{\eta} = \lim_{T\to\infty} \frac{1}{T}\int_0^T \eta(t)dt \quad \text{and} \quad \bar{\rho} = \lim_{T\to\infty} \frac{1}{T}\int_0^T \rho(t)dt.$$

Staleness Cost

- Phase-lock problem prevents us from getting the solution of the above two metrics
  - Constant update interval with interval 1
  - Constant download interval with interval 2
  - The staleness metrics depends on their initial states
- To avoid phase-lock cases, we propose age independence assumption:
  - Random query time of consumers Q_T : uniform in $[0,T]$
  - Two points processes N_U and N_U are called age-independent if $\forall x, y \geq 0$

$$\lim_{T\to\infty} P(A_D(Q_T) < x | A_U(Q_T) = y) = G_D(x)$$

sample path equilibrium distribution of A_D

Staleness Cost

- Theorem 1: Source penalty:

update rate

$$\bar{\eta} = \lambda\mu \int_0^\infty \bar{F}_U(y) \int_0^\infty w(x)\bar{F}_D(x+y)dxdy$$

download rate

  - Matches previous results with exponential update
  - When $\lambda = \mu$, left scenario gives age 1.5 hours; right 19 hours

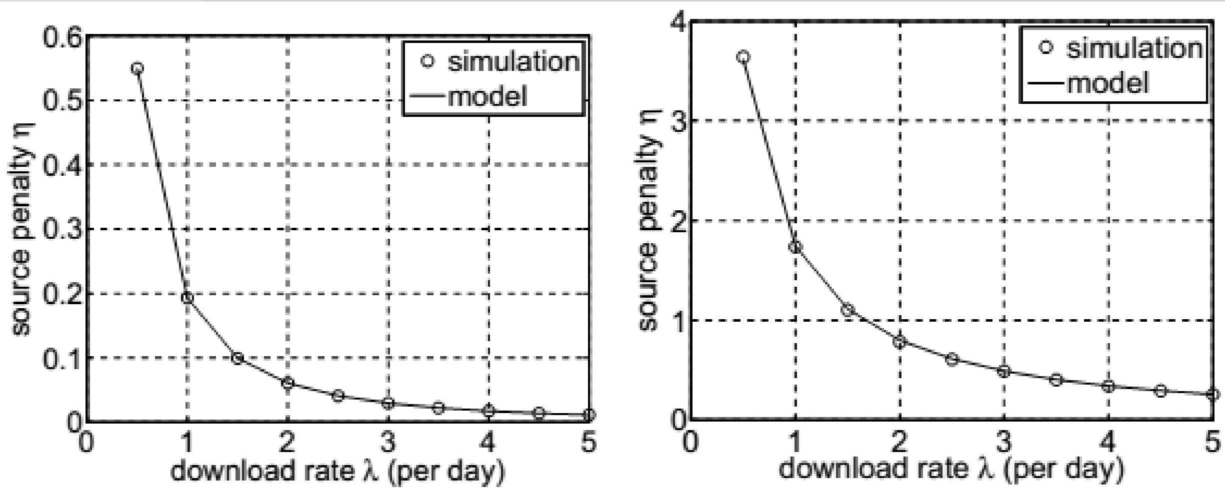

(a) Pareto U, constant D (b) constant U, Pareto D

$$w(x) = x, \mu = 2$$

Computer Science, Texas A&M University

Staleness Cost

- Theorem 2: Update penalty:

$$\bar{\rho} = \mu E[w_2(A_D)] = \lambda\mu E[w_3(D)] \qquad w_n(x) = \int_0^x w_{n-1}(y)dy$$

  - Matches previous results with exponential update
  - Allows decaying function $w(x) = 1/(1+x)$

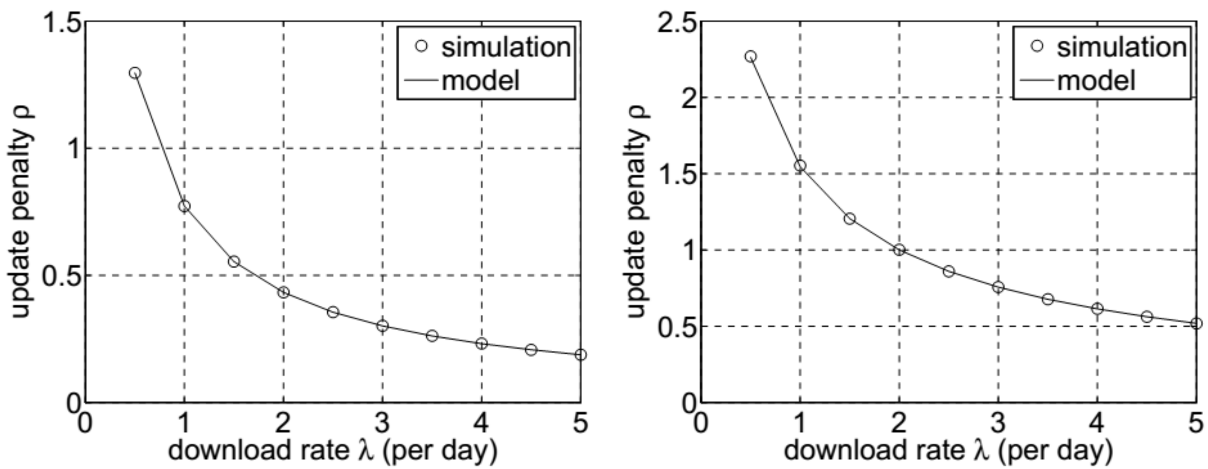

(a) Pareto U, constant D (b) constant U, Pareto D

Agenda

- Motivation
- Staleness Formulation
- Staleness Cost
- Optimality
- Applications
- Conclusion

Optimality

- With the same download rate λ, what distribution of synchronization intervals $F_D(x)$ is best?
- Definition 7: Variable X is stochastically larger than Y in second order, i.e., $X \geq_{st}^{2} Y$, if

$$\int_0^x \bar{F}_X(y)dy \geq \int_0^x \bar{F}_Y(y)dy \text{ for all } x \geq 0$$

- Theorem 7: For a given download rate λ and fixed update process N_U, both $\bar{\eta}$ and $\bar{\rho}$ decreases if download delays become stochastically larger in second order

Best Download Strategy

- Similarly, for a given update rate μ and fixed download process N_D, freshness increases if inter-update delays become stochastically smaller in second order
- Lemma 3: For a given mean, a constant stochastically dominates all other random variables in second order
- Corollary 1: Constant inter-synchronization delays are optimal under both $\bar{\eta}$ and $\bar{\rho}$, all suitable weights $w(x)$, and all update processes N_U

Agenda

- Motivation
- Staleness Formulation
- Staleness Cost
- Optimality
- Applications
- Conclusion

Real-Life Update Process

- The most frequently modified article in Wikipedia
  - "George W. Bush" with 44,296 updates in 10 years
  - (a) Pareto tail $(1+x/\beta)^{-\alpha}$ with $\alpha = 1.4, \beta = 0.93$
  - (b) Long-range dependence with Hurst parameter 0.81
  - (c) Non stationary
  - (d) Comparison between Poisson and real updates

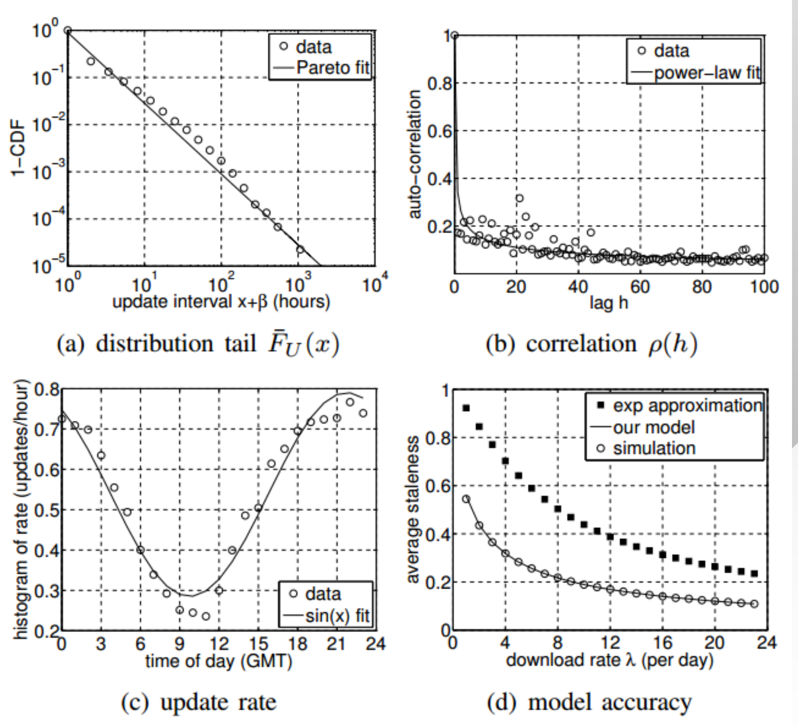

Fig. 7. George W. Bush page dynamics.

Bandwidth Estimation

- Comparison of our results with previous work
  - Apply both models to George W. Bush
  - Poisson: 20% staleness requires 95 downloads/day
  - Real: 20% staleness requires 8 downloads/day

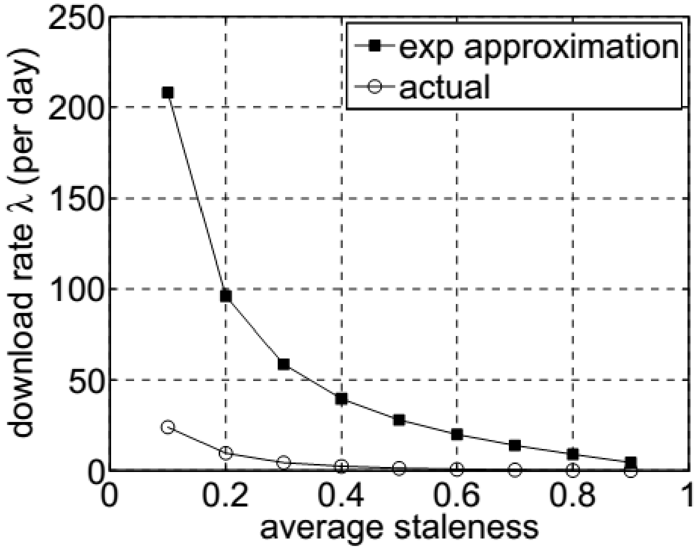

(a) crawl rate

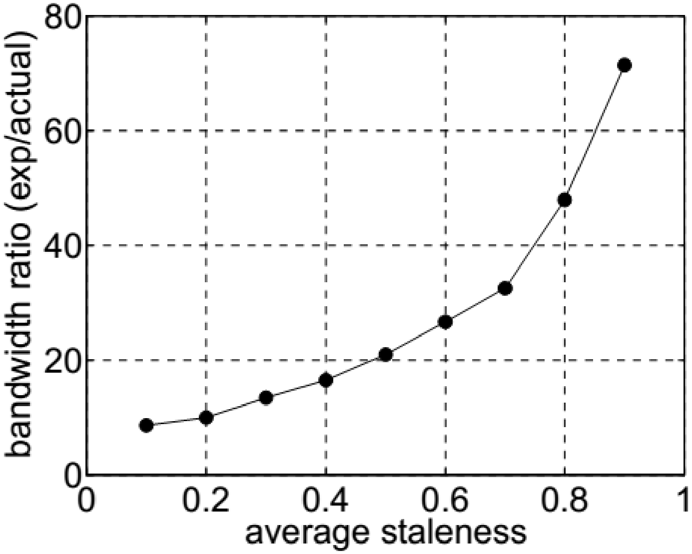

(b) ratio

Aggregation (Many-to-One)

- Resource allocation

$$\min \sum_{i=1}^{M} q_i C(\lambda_i, \mu_i) \quad \text{subject to} \quad \sum_{i=1}^{M} \lambda_i \leq \Lambda$$

q_i: probability that an incoming query quest i

μ_i: update rate of source i

λ_i: download rate of source i

- Using source penalty $\bar{\eta}$:
  - Page starvation exists: a source i will never be synchronized when total bandwidth Λ is relatively small compared to the source update rate μ_i
- Theorem 4: Assume $q_i\mu_i > q_j\mu_j > 0$ and constant download delay, optimal solution using $\bar{\rho}$ guarantees that $\lambda_i > \lambda_j > 0$
  - No page starves $\lambda_i = \Lambda \frac{\sqrt{q_i \mu_i}}{\sum_{j=1}^{M} \sqrt{q_j \mu_j}}.$

Load-Balancing (One-to-Many)

- Single source multiple replicas
  - The goal is to deduce the expected penalty afforded by the freshest member of all m replicas
  - Each replica has rate λ/m
  - Compare with single replica with rate λ
- The staleness at different replicas is no longer independent
  - Updates at the source make all copies outdated
  - The entire collection of replicas can be replaced by a single replica that that refresh pattern $N_D{}^*$, which is the superposition of all point processes $\{N_D^i\}_{i=1}^m$

Conclusion

- We proposed a novel framework for modeling staleness metrics under general update/download processes
- We established that constant inter-refresh intervals were optimal for all considered cases
- Finally, we consider a family of related problem stemming from $1 \times m$ and $M \times 1$ replication, showing that they can be solved from the preceding results

Questions?