Selective Data Replication for Online Social Networks with Distributed Datacenters

Guoxin Liu [*], Haiying Shen [*], Harrison Chandler[*]

Presenter: Haiying (Helen) Shen

Associate professor

*Department of Electrical and Computer Engineering, Clemson University, Clemson, USA

Outline

- Introduction
- Related work
- Data analysis
- Selective data replication
- Evaluation
- Conclusion

Introduction

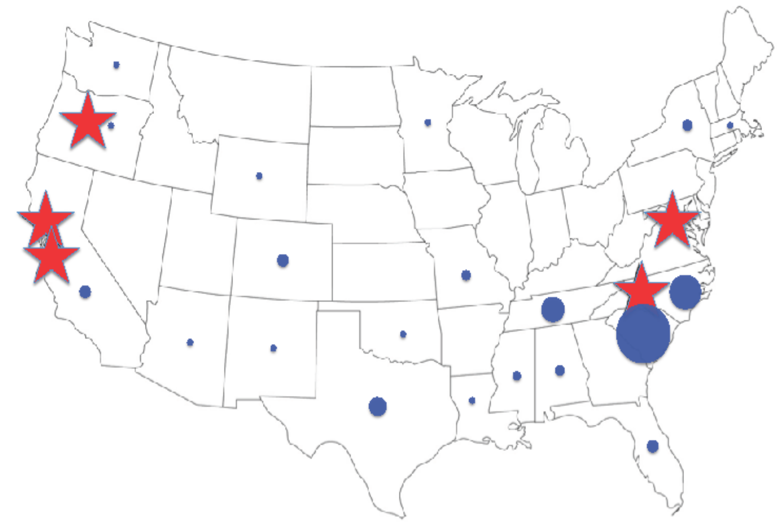

- Facebook's growth*
  - Monthly active users:
    - 700 millions in 2011
    - 800 millions in 2013
  - Users distribution:
    - 70% outside US and Canada in 2011
    - 80% outside US and Canada in 2013
  - Challenges for service scalability:
    - Global distribution: low service latency and costly service to distant users
    - Scaling problem: bottleneck of the limited local resources

*http://www.facebook.com/press/info.php?statistics.

Current Facebook datacenters

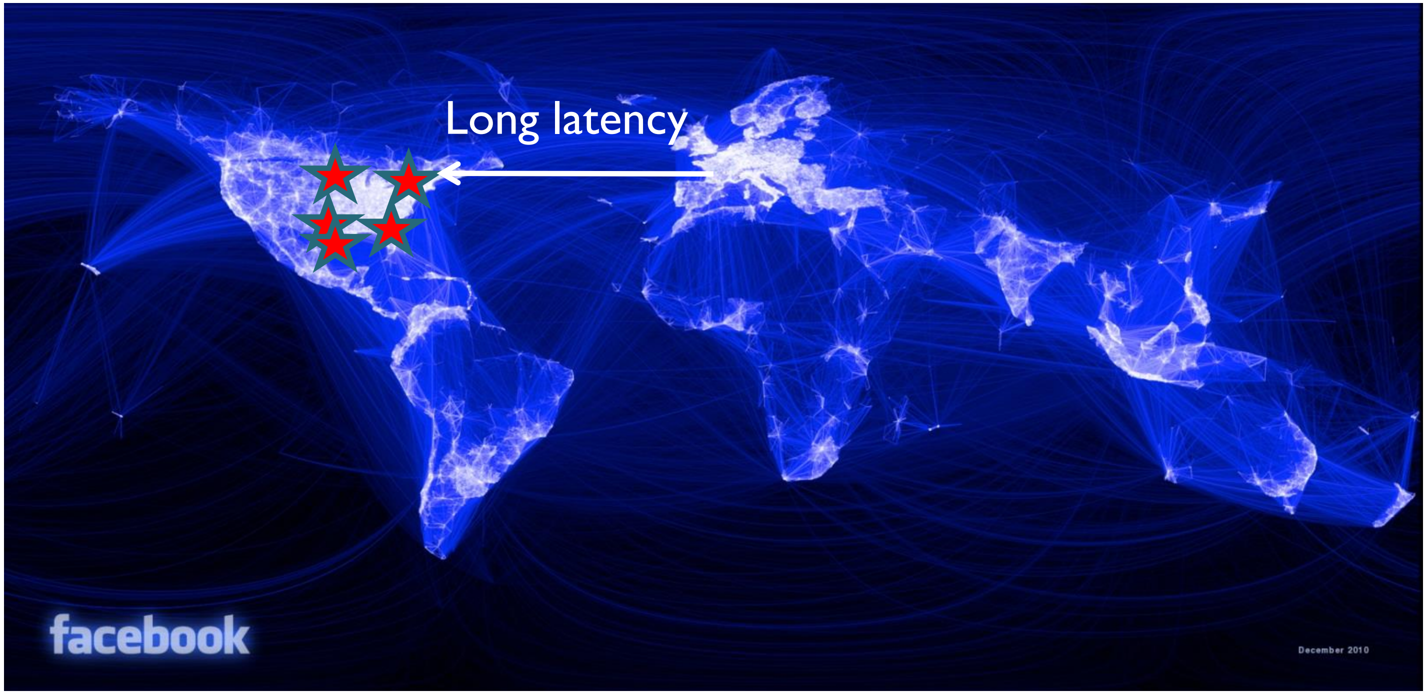

OSN distributed small datacenters

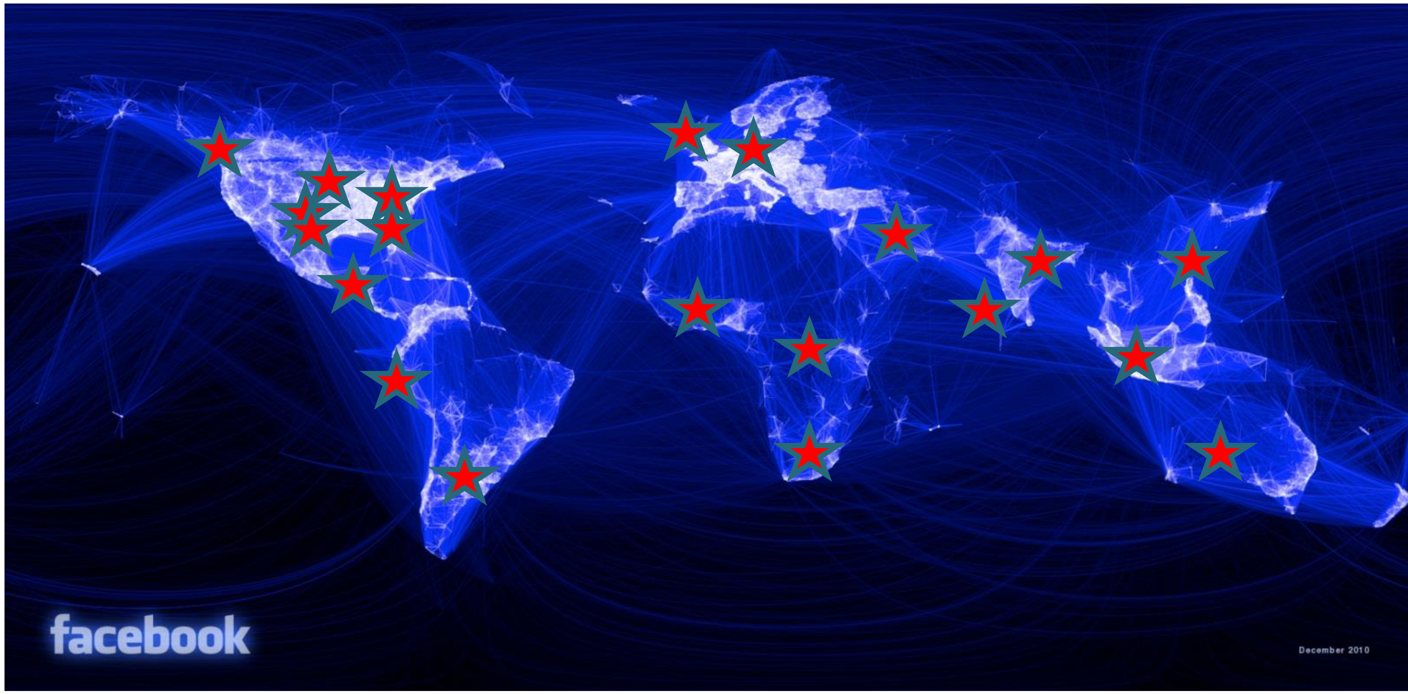

- New datacenter infrastructure
  - Globally distributed small datacenters
    - Luleå datacenter in Sweden: reducing the service latency of European users

OSN distributed small datacenters

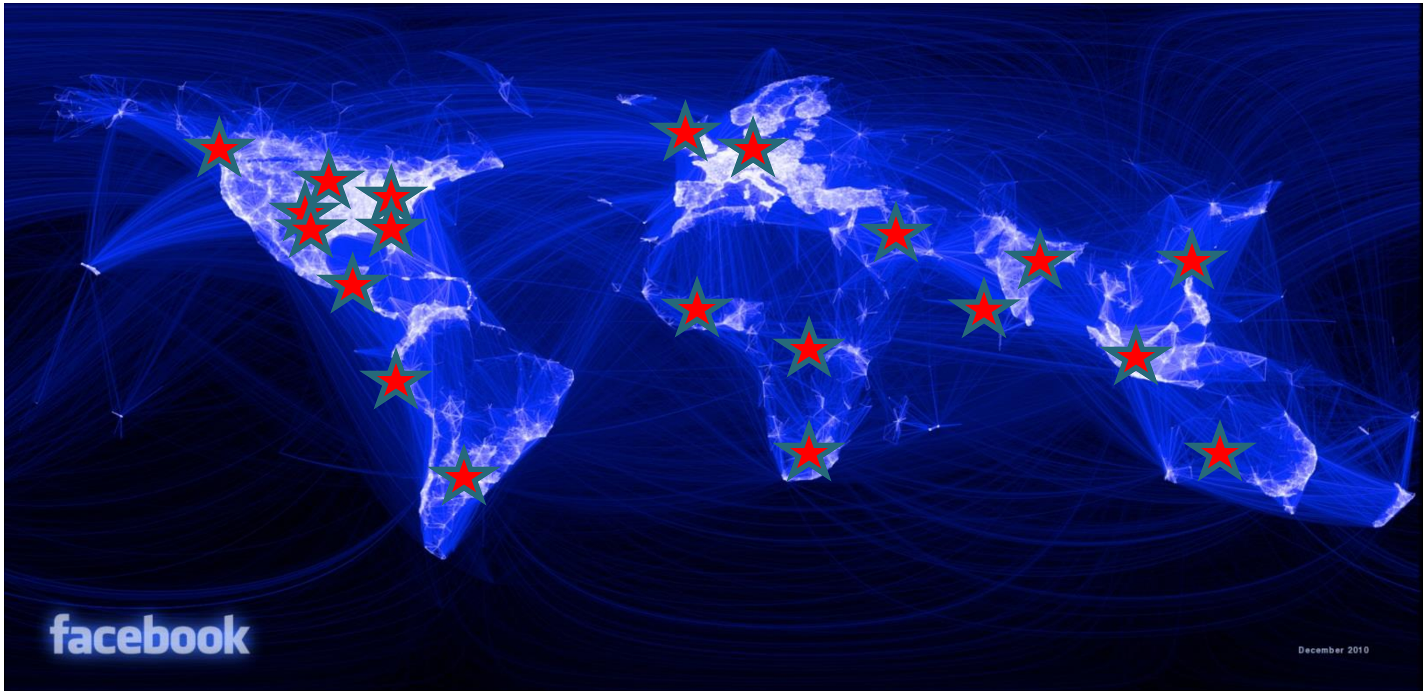

- New problems

Introduction

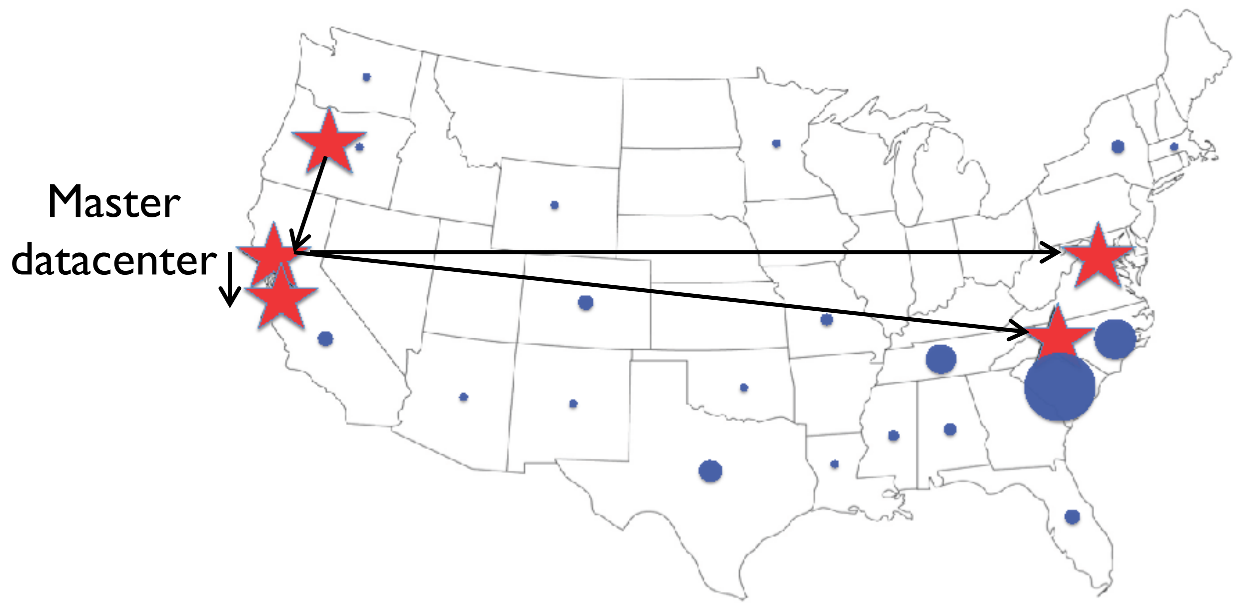

- Each datacenter has a full copy of all data
- Single-master replication protocol:
  - a slave datacenter forwards an update to the master datacenter, which then pushes the update to all datacenters

OSN distributed small datacenters

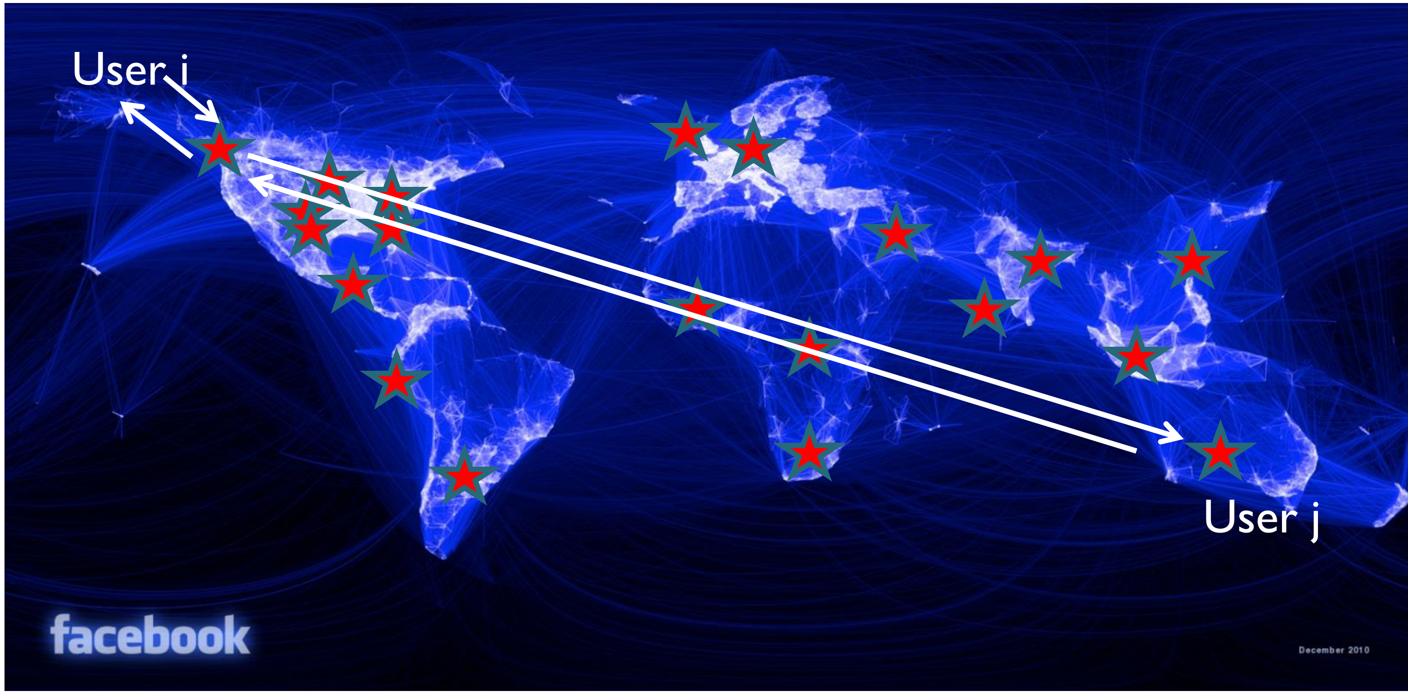

- New problems
  - Single-master replication protocol: tremendously high load
    - Ten million updates per second
  - Locality-aware mapping: stores a user's data to his/her geographically-closest datacenter
    - Frequent interactions between far-away users lead to frequent communication between datacenters

Introduction

- Key challenge:
  - How to replicate data in globally distributed datacenters to minimize the inter-datacenter communication load while still achieve low service latency

- Solution: Selective Data replication mechanism in Distributed Datacenters (SD^3)
  - Globally distributed small datacenters
    - Locality-aware mapping of users to master datacenters
  - Selective user data replication
  - Atomized user data replication

Outline

- Introduction
- Related work
- Data analysis
- Selective data replication
- Evaluation
- Conclusion

Related work

- Facebook community pattern:
  - Interaction communities exist
  - Interaction frequencies between friends vary
- Different atomized data types (e.g., wall/friend posts, personal info, photo/video comments) have different update/visit rates
- Facebook scalability
  - Inside datacenter
    - Collecting the data of users and their friends in the same server
  - Outside datacenter
    - Distributing region servers acting as Facebook service proxies
- Replication strategies in P2P and Cloud
  - Not suitable without considering the interactions among social friends

Outline

- Introduction
- Related work
- Data analysis
- Selective data replication
- Evaluation
- Conclusion

Data analysis

- Data crawling:
  - We used PlanetLab to evaluate an OSN's access latency and the benefits of globally distributed datacenters
  - We crawled status, friend posts, photo comments and video comments of 6,588 users from May 31-June 30, 2011
  - We crawled 22,897 friend pairs and their locations

Data analysis

- Basis of distributed datacenters
  - Service latency of the OSN
    - Typical latency budget 50-100 milliseconds
    - 20% of PlanetLab nodes experience service latency >102ms
  - Service latency with simulated globally distributed datacenters
    - more datacenters lead to lower service latency
  - Suggest distributing more small datacenters globally

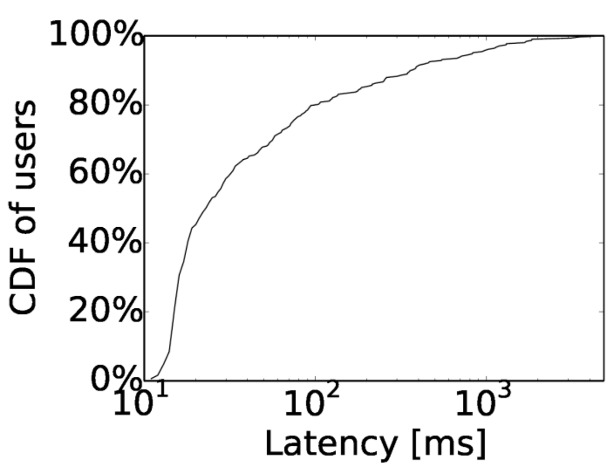

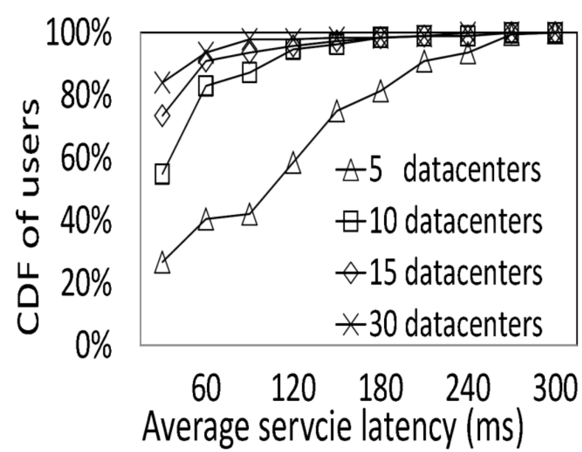

Data analysis

- Basis for selective data replication
  - Friend relationships do not necessarily mean high data visit/update rates
    - Interaction rate between some friends is not high
      - Replication based on static friend communities is not suitable
    - Interaction rate among friends vary over time
      - Visit/update rate of data replicas should be periodically checked

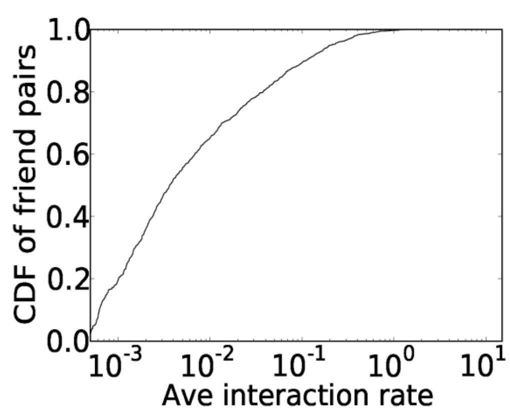

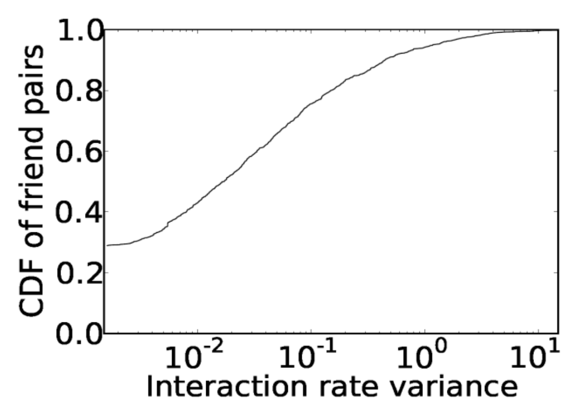

Data analysis

- Basis for atomized data replication
  - Different types of data have different update rates
  - The update rates of different types of data of a user vary
  - Exploiting the different visit/update rates of atomized data to make decision of replication separately
  - Avoid replicating infrequently visited and frequently updated atomized data to reduce inter-datacenter updates

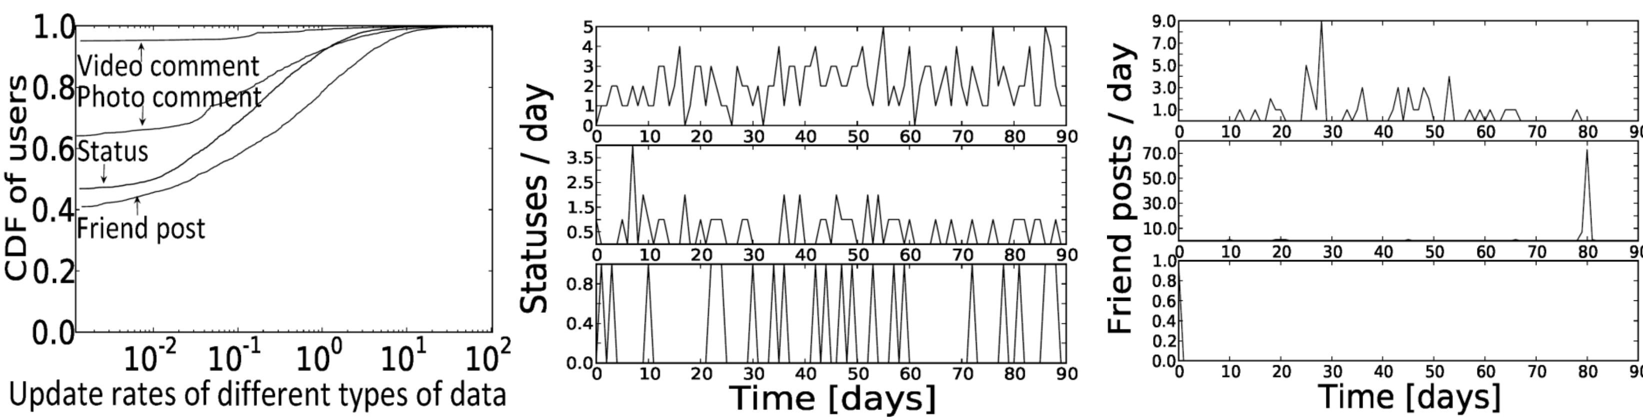

Outline

- Introduction
- Related work
- Data analysis
- Selective data replication
- Evaluation
- Conclusion

Selective data replication

- An overview of SD^3
  - Deploy worldwide distributed smaller datacenters
    - Map users to their geographically closest datacenters as their master datacenters
  - Replicate data only when the replica saves network load
  - Atomize a user's data based on different types

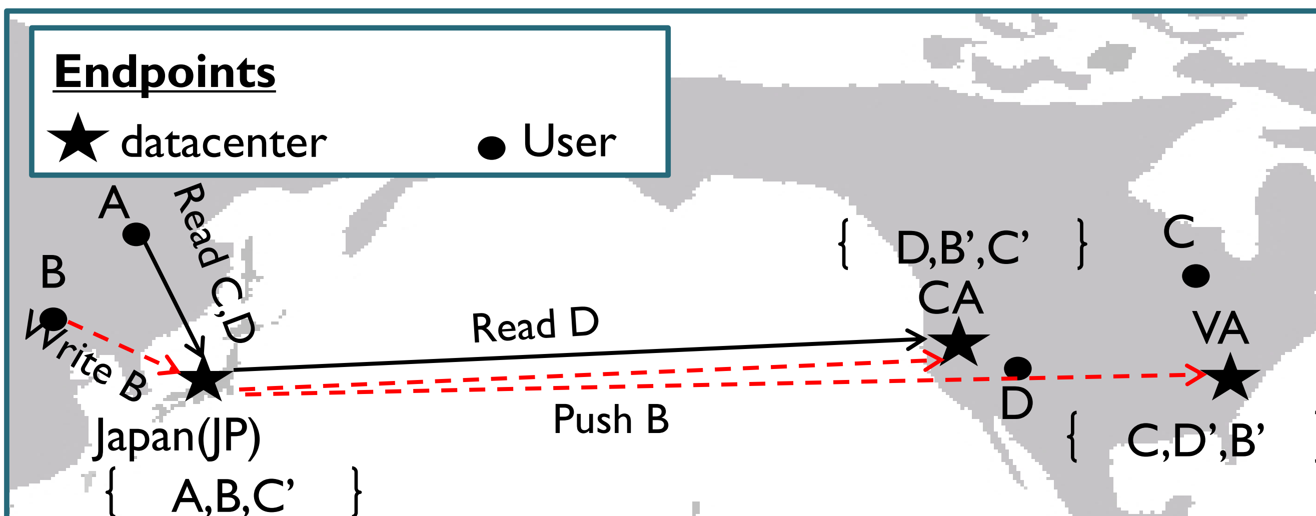

Selective data replication

- Local replicas of friends' data
  - Reduce service latency (related to visit rate)
  - Generate data update load (related to update rate)
- Selective data replication (SD^3): minimize network load while maintain low service latency
  - Consider both visit rate and update rate of a user's data to decide replication
  - Adopt a simple measurement for network load:
    - Package size × traffic distance

Selective data replication

- For a specific replica set of all datacenters:
  - Network load benefits:
    - $B_{total} = O_s - O_u$
  - O_s: saved network load
    - The total differences of visit network load between with and without all replicas
  - O_u: update network consumption
    - The total update network load with all replicas
  - Goal: maximizing B_{total}
  - Solution:
    - For each datacenter's non-master user data
      - $B_{c,j} = O_{s,j} - O_{u,j} = (V_{c,j}S_j - U_jS_u)D_{c,cj}$
    - Maximize the benefits of each user data replica

OSN distributed small datacenters

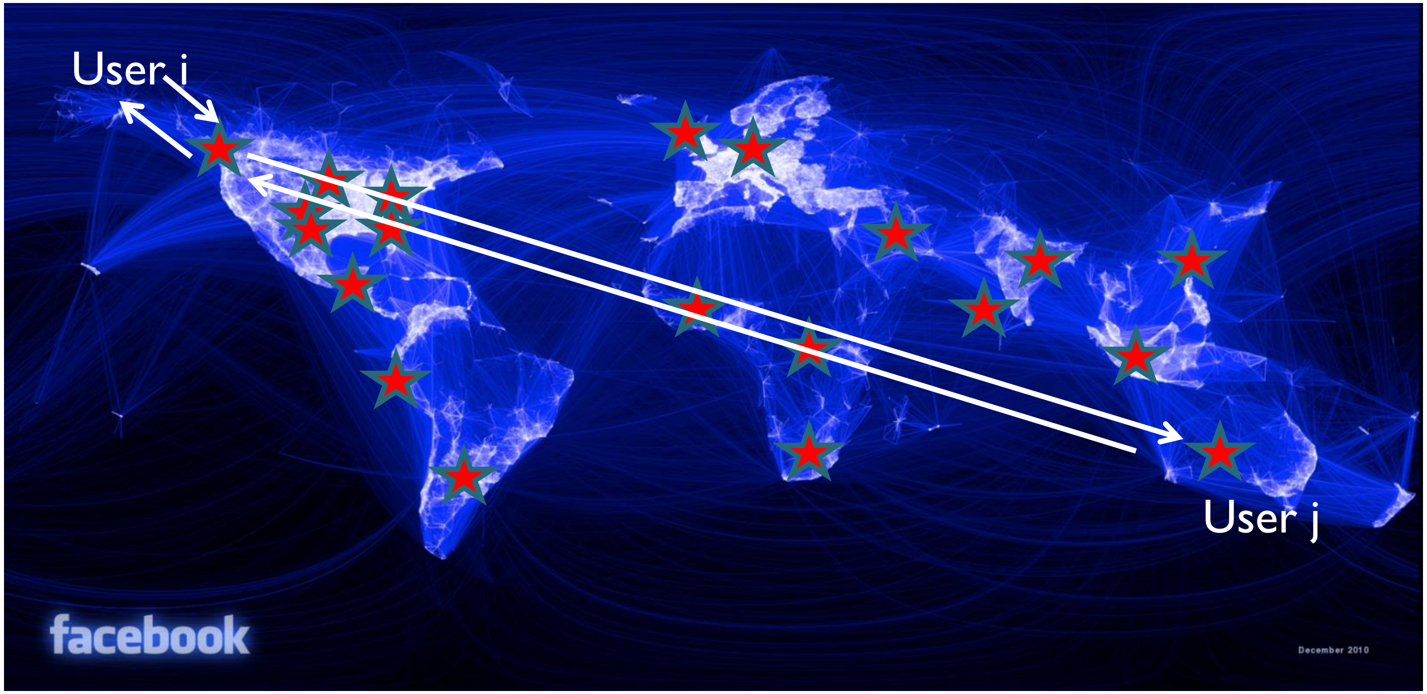

Selective data replication

- Decision of replication based on prediction
  - Constant visit rate and update rate
    - All user data j that $B_{c,j}$>0
  - Large variance of visit and update rates
    - Introduce two thresholds: T_{Max} and T_{Min}
      - $B_{c,j} > T_{Max}$, create a new replica of user data j
      - $B_{c,j} < T_{Min}$, remove the replica of user data j
    - Decision of thresholds:
      - Based on user service latency constraint, saved network load, replica management overhead and so on

Selective data replication

- Algorithm analysis of SD^3
  - Performance
    - SPAR: replicating all friends data
    - RS: replicating all visited data
    - SD^3: selective replication
  - Time complexity of SD^3:
    - $O(n)$ (n: num. of users)
- Enhancement:
  - Atomized user data replication
    - Handle different types of user data separately to decide replication

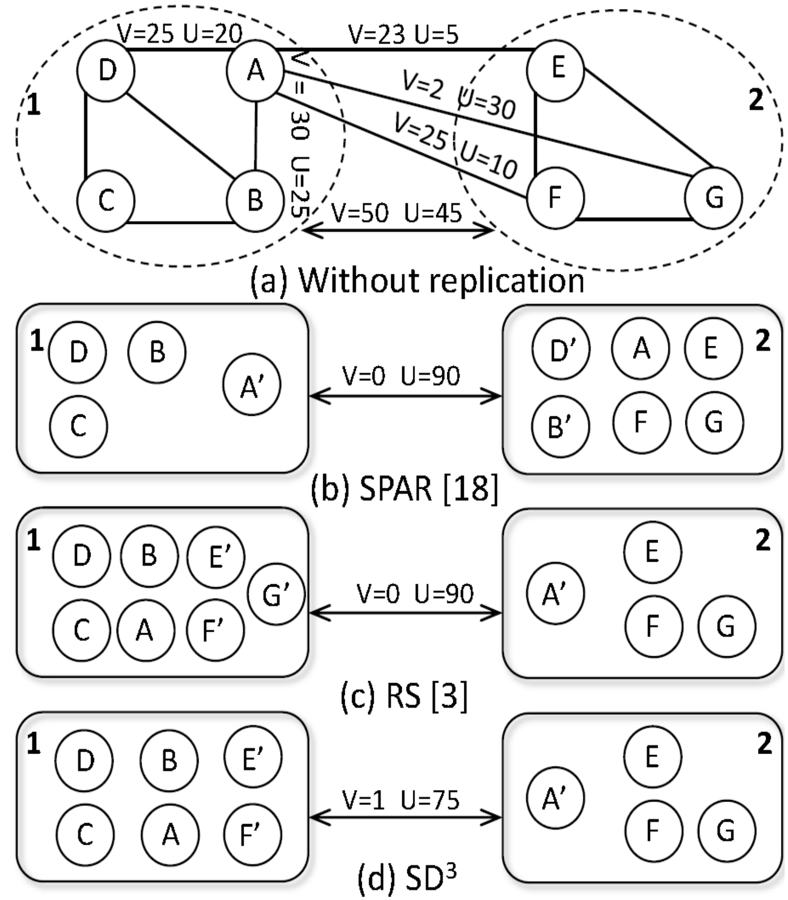

[3] M. P. Wittie, V. Pejovic, L. B. Deek, K. C. Almeroth, and B. Y. Zhao. Exploiting locality of interest in online social networks. In Proc. of ACM CoNEXT, 2010.

[18] J. M. Pujol, V. Erramilli, G. Siganos, X. Yang, N. Laoutaris, P. Chhabra, and P. Rodriguez. The little engine(s) that could: scaling online social networks. In Proc. of SIGCOMM, 2010.

Outline

- Introduction
- Related work
- Data analysis
- Selective data replication
- Evaluation
- Conclusion

Evaluation

- Used crawled the OSN data for
  - Update rate of each user data type
    - Derived visit rate according to [11]
  - Number of friends and friend distribution
  - Visit rate distribution of a user data type among friends
- 13 simulated datacenters
- 36,000 simulated users
- Comparison:
  - SPAR [18]: replicating all friends data
  - RS [3]: replicating all visited data and keep within a certain time
    - RS_L and RS_S
  - LocMap: without replication

[3] M. P. Wittie, V. Pejovic, L. B. Deek, K. C. Almeroth, and B. Y. Zhao. Exploiting locality of interest in online social networks. In Proc. of ACM CoNEXT, 2010.
[11] F. Benevenuto, T. Rodrigues, M. Cha, and V. Almeida. Characterizing user behavior in online social networks. In Proc. of ACM IMC, 2009.
[18] J M. Pujol, V. Erramilli, G. Siganos, X. Yang, N. Laoutaris, P. Chhabra, and P. Rodriguez. The little engine(s) that could: scaling online social networks. In Proc. of SIGCOMM, 2010.

Evaluation

- Effect of Selective User Data Replication
  - Avoid replicating rarely visited and frequently updated user data
    - SD^3 generates a small number of replicas

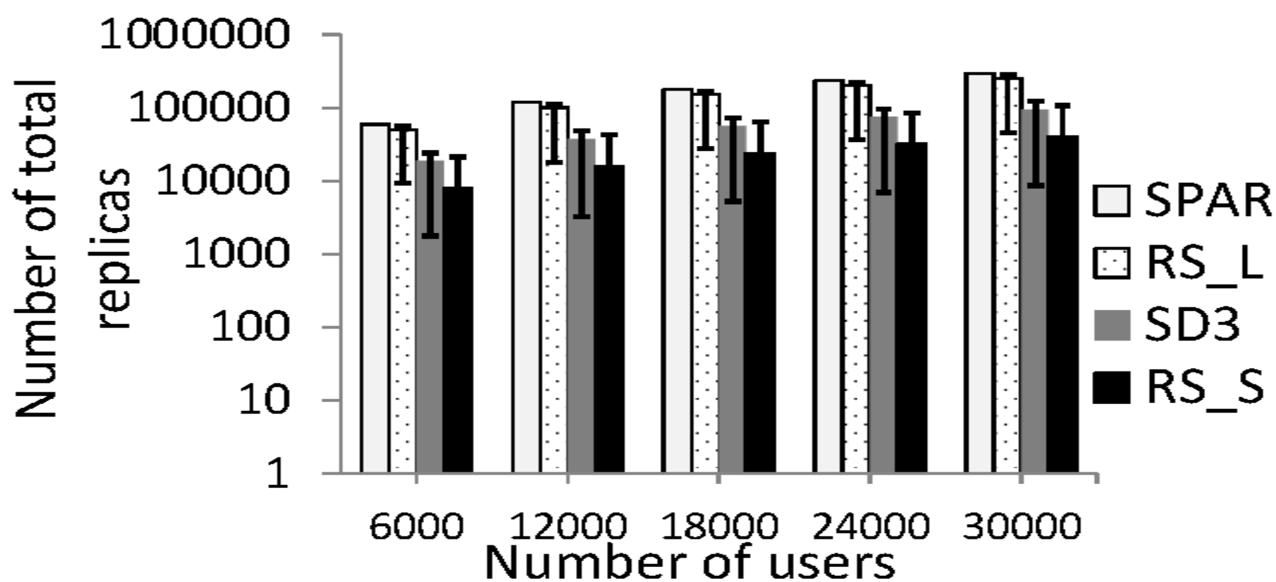

Evaluation

- Effect of Selective User Data Replication
  - Avoid replicating rarely visited and frequently updated user data
    - SD^3 saves the highest network load

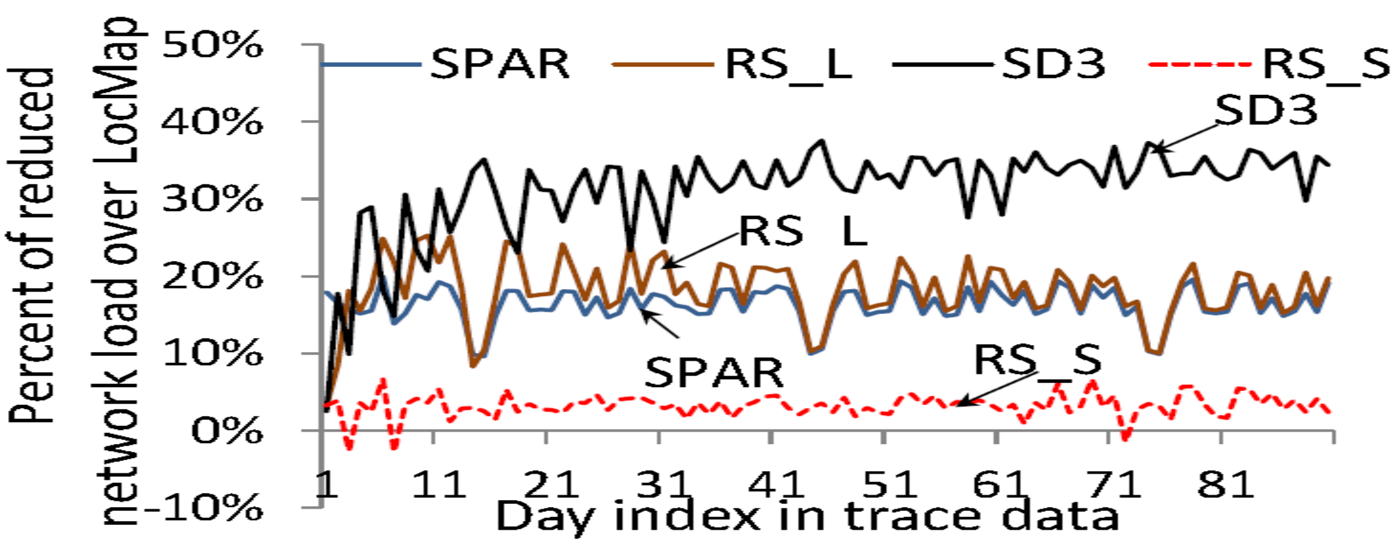

Evaluation

- Effect of Selective User Data Replication
  - Avoid replicating rarely visited and frequently updated user data
    - SD^3 achieves a small service latency

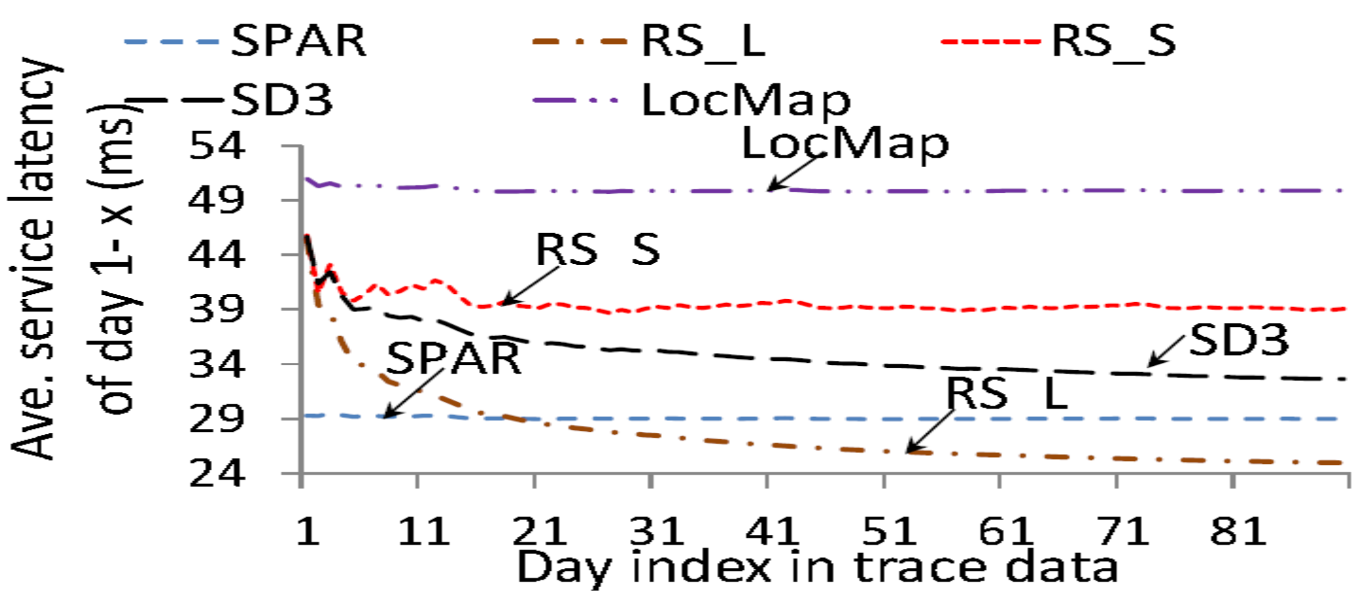

Evaluation

- Effect of Atomized User Data Replication
  - Separately handle different user data types
    - SD[3] with atomized user data replication saves at least 42% network load

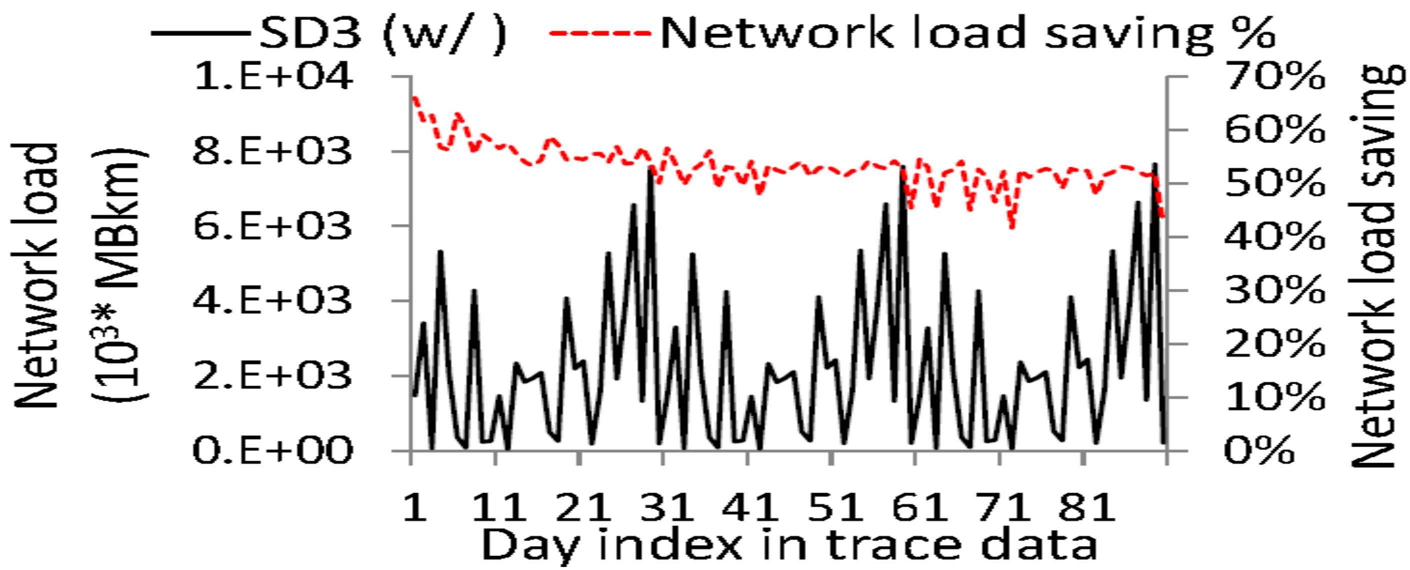

Outline

- Introduction
- Related work
- Data analysis
- Selective data replication
- Evaluation
- Conclusion

Conclusion

- Goal:
  - Low inter-datacenter network load and low service latency
- Selective data replication mechanism in Distributed Datacenters (SD^3)
  - Design supports:
    - Crawled trace data
  - Design principles:
    - Jointly consider both visit rate and update rate of a user data's to decide the replication in order to minimizing the network load
  - Enhancement:
    - Atomized data (each data type) handled separately
- Future wok:
    - Investigate the determination of all parameters to meet different requirements on service latency and network load

Thank you!
Questions & Comments?

Haiying (Helen) Shen, Associate Professor

shenh@clemson.edu

Pervasive Communication Laboratory

Clemson University

Clemson is hiring Postdoc in Connected Vehicles