Twittering by Cuckoo – Decentralized and Socio-Aware Online Microblogging Services

Tianyin Xu*, Yang Chen*, Xiaoming Fu* and Pan Hui+

*Computer Networks Group, University of Goettingen, Goettingen, Germany

+Deutsche Telekom Laboratories/TU-Berlin, Berlin, Germany

INTRODUCTION

WHAT IS CUCKOO?

A decentralized and socio-aware microblogging system:

- Take advantage of the inherent social relations
- Leverage peer-to-peer (P2P) techniques

towards scalable and reliable microblogging services.

MOTIVATION

WHY CUCKOO? (TWITTER [1] ALREADY STANDS THERE)

Current microblogging systems depends on centralized architecture – Polling-style: blind, sticky, superfluous.

- Placing heavy bandwidth burdens[1] both on peers and servers;
- Performance bottleneck[2] and single point of failure[3];
- Vulnerable to service blocking[4] and malicious attacks[5].

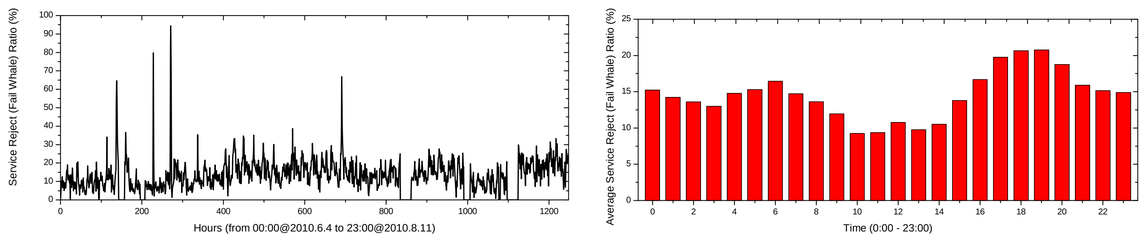

FIGURE 1: Service Reject (Fail Whale) Ratio

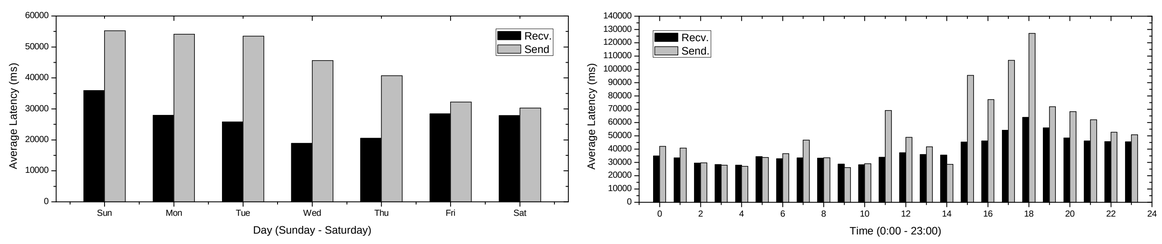

FIGURE 2: Response Latency (Send & Recv.)

CUCKOO IN A NUTSHELL

DECENTRALIZATION: build on a hybrid overlay structure.

- **Structured Overlay (DHT)**
  – Provide location service (find any online user in $O(log(N))$ hops);
  – Improving system availability;
- **Unstructured Overlay (Gossip Protocol)**
  – Propagate updates to all the subscribers;
  – Share micro-news with users having same interest.

SOCIO-AWARENESS: take advantage of inherent social relations.

- **Friend**
  – Reciprocate social link between two users;
  – Help each other to balance load and improve availability;
- **Partner**
  – Relationship between users with same interests;
  – Using gossip among partners to share micro-news;
- **Followee/Follower**
  – Most common one-way connection;
  – Followees push updates to followers instead of polling-style pull.

[1] In 2007, Twitter experienced 5 days, 23 hours of downtime in total due to server overloaded [3].
[2] Twitter currently has the rate limit that only allows 150 requests per hour for normal users [1].
[3] In June 2008, as many as 3% of page requests in Twitter yielded "over capacity" errors [4].
[4] Due to political reasons, Twitter is currently blocked by several countries.
[5] In August 2009, Twitter did be crippled by DDoS attack [2].

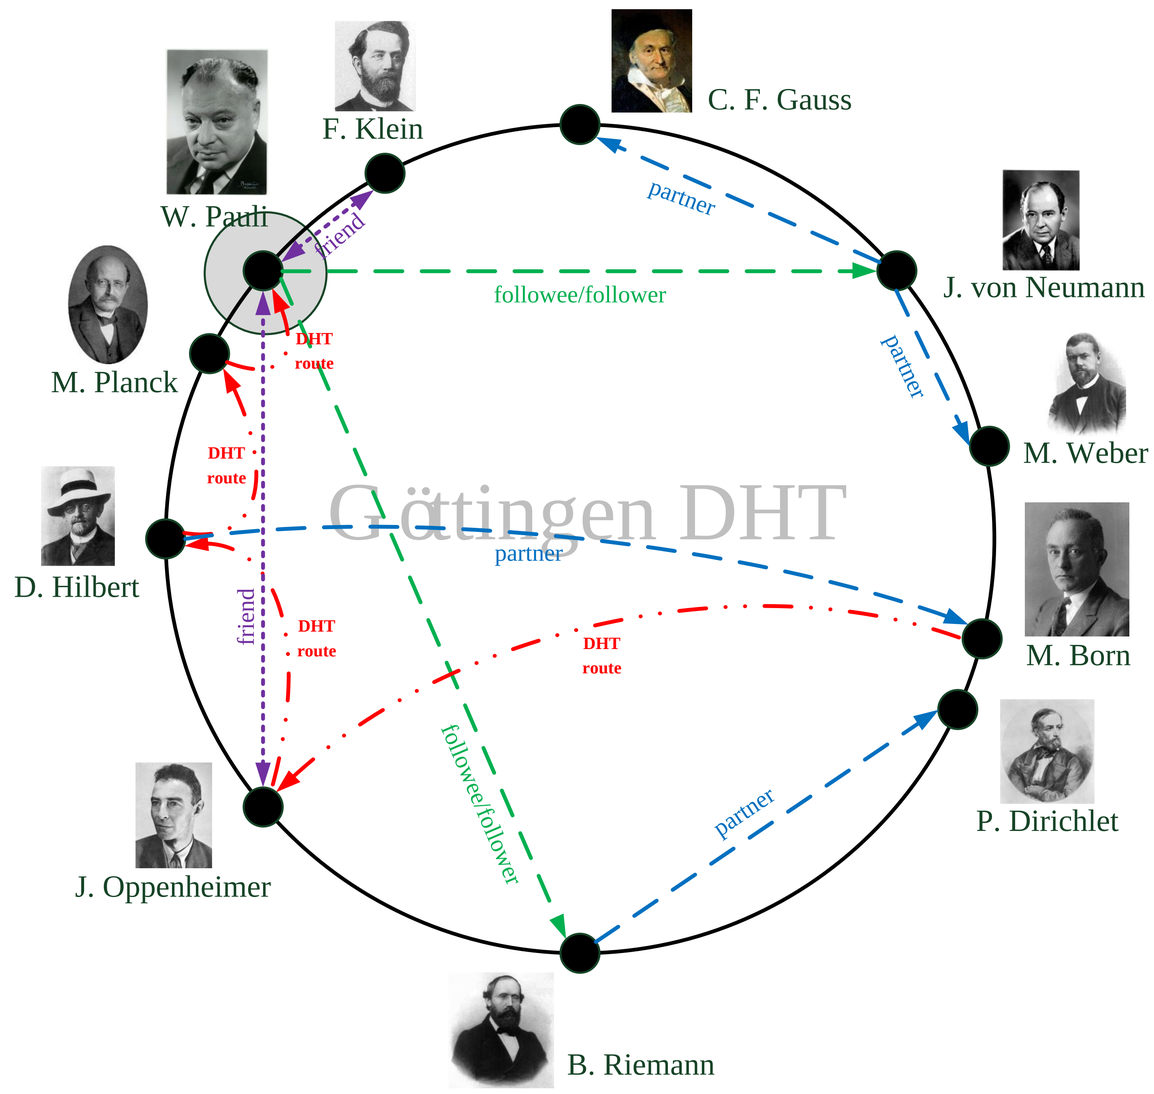

ROLE OF SERVICE PROVIDERS: CUCKOO IS PEER-ASSISTED INSTEAD OF FULLY DISTRIBUTED

- **Achieving better quality of service**
  – Support synchronization for peers with asynchronized access;
  – Guarantee high availability (always online);
- **Nothing to lose, nothing to change**
  – Fully compatible with current architecture;
  – Will not lose any functionalities nor user communities;
  – Keep all the precious resources (profiles & microblogs) as before;
- **Excellent platform for third party developers to enrich additional functions**
  – Simple functions on the server side and more colorful functions between peers.

CUCKOO USE CASES

WHAT CAN CUCKOO DO?

- **All Twitter's functions are supported.**
- **What's more?**
  – Allow to choose whether uploading to the central server;
  – Allow to choose dissemination style: direct or gossip (the former is enough for normal users & the latter is good for broadcasters);
  – Allow to share long micro-contents (larger than 140);
  – Can be used in regions where the service is blocked;
  – More functions are under development and can be envisioned;
  – Welcome to our website: http://mycuckoo.org!

References

[1] Twitter. http://twitter.com/
[2] CNET News. Twitter cripped by denail-of-service attack.
[3] Pingdom. Twitter growing pains cause lots of downtime in 2007.
[4] D. R. Sandler and D. S. Wallach. Birds of a FETHR: Open decentralized micropublishing. In Proc. of IPTPS, 2009.
[5] T. Xu, Y. Chen, J. Zhao, and X. Fu. Cuckoo: Towards Decentralized Socio-Aware Online Microblogging Services and Data Measurements. In Proc. of HotPlanet, 2010.