Exploiting Locality of Interest in Online Social Networks

Mike P. Wittie, Veljko Pejovic, Lara Deek,
Kevin C. Almeroth, Elizabeth M. Belding, Ben Y. Zhao
Department of Computer Science
University of California, Santa Barbara
Santa Barbara, CA 93106
{mwittie, veljko, laradeek, almeroth, ebelding, ravenben}@cs.ucsb.edu

1. INTRODUCTION

Online social networks (OSNs), such as Facebook, MySpace, Orkut, and many others, have expanded their membership rapidly over the last several years. These networks interconnect users through *friendship* relations and allow for asynchronous communications within thus defined social graph. While various OSNs support other types of interactions, including browsing of users' *profiles*, the bulk of traffic can be attributed to inter-user communications.

OSNs continue to expand, and as a result, an ever-increasing amount of computing power and bandwidth are needed to support the communications of the growing user base. At the center of an OSN is the social graph and user data, which are traditionally stored and operated on in a centralized data center. As the result, OSN services can appear unresponsive to users located far away from such data centers.

We focus our study on Facebook, the largest OSN. Facebook's highly centralized infrastructure is not well-suited to provide services in remote areas of the globe [1]. User-generated updates are routed to the master database in California, from which they propagate to a replica in Virginia [2]. Finally, content distribution networks (CDNs), such as Akamai, serve static content to users outside of the US. While the replication of Facebook's databases provides timely updates to users in US and Europe, latency measurements from other regions confirm a sluggish service elsewhere.

In this work we aim to evaluate Facebook infrastructure design choices. Few details of Facebook's infrastructure is publicly available, and so we reverse-engineer Facebook through analysis of the history of user interactions, interaction packet traces, and network performance measurement between discovered global infrastructure endpoints. Simulations of the observed interactions within the discovered infrastructure allow us to measure a number of key performance metrics, such as transaction delays as perceived by users. We also aim to quantify server load at each point of the infrastructure and relate that information to costs of network bandwidth and data center size.

Based on the discovered shortcomings of the current design of OSN infrastructure, we propose a number of changes. While the social graph is difficult to partition, we have observed that communication patterns within the social graph are highly localized. Based on locality of traffic, we aim to show that local handling of traffic improves service responsiveness for the global user base. We also aim to show that aggregation of traffic between regional and global infrastructure can reduce load and infrastructure cost.

2. REVERSE ENGINEERING FACEBOOK

Very little information about Facebook's infrastructure is available publicly. To understand the performance of Facebook's infrastructure and the quality of service offered to its global user base, we take a three-prong approach to reverse engineer Facebook. First, we analyze user interaction history within a regional social graph. Second, we characterize Facebook's traffic through packet traces. Finally, we measure Internet path characteristics between different points of Facebook's infrastructure the iPlane system.[1]

2.1 User Interaction Analysis

Crawls of Facebook state, performed by Wilson *et al.*, show a history of user interactions and the social graph within a number of regional networks [3]. In an initial analysis of user interactions we have noticed a certain *locality of interest* within user interactions. To quantify this phenomenon and understand how it could be leveraged to improve the OSN service, we analyze user interactions by their *directionality.*

Interactions can be categorized by the locality of the interacting users with respect to a regional network. Two categories of interactions, namely local-to-local ($L \rightarrow L$) and remote-to-local ($R \rightarrow L$), can be obtained from regional network crawls by scanning each user's update *feed.* The third category of interaction involving local users, local-to-remote ($L \rightarrow R$), is calculated by multiplying the number of each user's $R \rightarrow L$ interaction by their *reciprocity factor*, or their likelihood of replying to another local user's post.

To understand the locality of Facebook's traffic, we have analyzed the directionality of interactions in three regional networks India, South Africa, and Sweden. We have selected these regions based on geographic and socioeconomic diversity and the availability of Facebook crawls.

Figure 1 shows the relative volume of $L \rightarrow L$, $L \rightarrow R$, and $R \rightarrow L$ wall posts. We observe that most of the interaction, namely $L \rightarrow L$ and $L \rightarrow R$, are generated locally, or within each regional network. The prevalence of local traffic leads us to believe that processing and caching of local traffic within each region could improve user perceived performance and reduce bandwidth load on the global Facebook infrastructure.

To understand how updates are consumed, we assume that each wall post is read by friends of both the writer and owner of the wall, which is the default Facebook behavior. We then use the social graph to calculate the number of reads, or updates Facebook needs to send out to the feeds of interested users. Figure 2 shows the ratio of read to write events for wall posts categorized by the direction of the original write. In general, locally generated traffic, or $L \rightarrow L$ and $L \rightarrow R$ interactions, exhibit a high rate of consumption within each regional network.

The analysis of regional interaction patterns shows that traffic local to a region is produced and consumed in significant

[1]iplane.cs.washington.edu

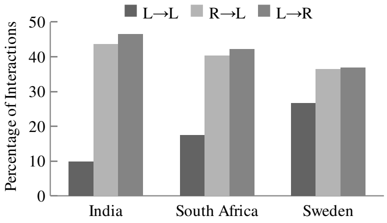

Figure 1: Wall posts.

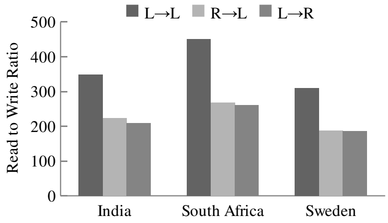

Figure 2: Wall post r/w ratio.

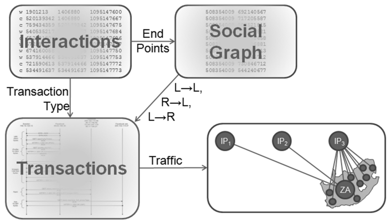

Figure 3: Simulator operation.

volumes. We believe this shows promise for a redesign of Facebook's infrastructure to cater to local interactions and, by doing so, improve user perceived responsiveness of the service and reduce infrastructure costs. To more fully understand the problem space, we continue our reverse engineering effort with the analysis of the network traffic involved in the studied interactions.

2.2 Transaction Analysis

To understand Facebook's traffic we also need to know the traffic involved in user interactions. By observing Facebook traffic of users in different regions, we can discover the addresses of servers they interact with, or infrastructure endpoints. Moreover, through sequencing and replaying of interactions we can discover caching behavior of the CDN network employed by Facebook. We capture IP packets involved and distill traffic flows between OSN server endpoints involved in each type of transaction. A transaction may contain the traffic for multiple interactions, such as the delivery of multiple wall posts, in which case we aggregate interaction data, but not transaction control traffic, within appropriate traffic flows.

2.3 Regional Network Measurement

Finally, to understand communication delays within Facebook's architecture, we measure network performance between server endpoints identified in packet traces. Packet traces are collected from a single Facebook account, but to realistically simulate connectivity between Facebook's and the many regional users, we assign these users to different subnets within each region. We identify these subnets from www.find-ip-address.org, which lists last mile subnets at a country level, and query iPlane to obtain median latency, loss, and bandwidth link characteristics between regional subnets and OSN servers.

3. EVALUATION METHODOLOGY

Our goal is to evaluate the current Facebook architecture on a network simulator driven by user interaction traces, transaction traffic, and network measurement. We would like to characterize the responsiveness of the Facebook service by measuring the delay of each transaction type. We would also like to measure server load and the amount of regionally cached data and relate these to infrastructure costs.

The high level operation of our simulator is illustrated in Figure 3. Timestamped interactions collected during crawls are precessed sequentially. Based on the interaction endpoints and the social graph, each interaction is classified as either $L \rightarrow L$, $L \rightarrow R$, or $R \rightarrow L$. Transaction type, such as 'wall post,' and its directionality allows for the identification of the corresponding traffic trace, which is then fed into a simulated network configured to represent the measured regional network. The simulator replays traffic traces of individual transactions and collects the metrics of transaction delay, and server and cache load.

The simulator will allow us to evaluate the reverse engineered Facebook infrastructure with real traffic. While this is still an ongoing effort, we propose a number of simple changes to Facebook's processing of local traffic, which we will also evaluate within our simulator framework.

4. DISTRIBUTED FACEBOOK STATE

Through analysis and simulations we have began to identify network latencies between globally distributed users and the US centralized Facebook infrastructure as the major source of user-observed delay. While Facebook mitigates some of this delay with regional CDNs, service responsiveness could be improved in two simple ways.

First, we propose that locally generated writes be handled first by a regional Facebook server before being propagated to the US for global consistency. Processing writes locally reduces transaction delay and makes the data available more quickly to other users in the regional network.

Second, we propose that communications between regional nodes and the US infrastructure be aggregated and compressed by the regional server before transmission to the US for global consistency. The aggregation of relatively contemporaneous requests will reduce network load, without introducing noticeable delays to users.

Additionally, because users now only need to communicate with the regional server, their TCP connections with US infrastructure is split by the server. Poor last-mile access links can achieve better throughput with a nearby server than with a distant server over a high-latency link and so greater interaction with regional OSN infrastructure can also improve users' effective bandwidth.

We aim to evaluate these proposed changes within the simulation framework configured through the reverse engineering effort. We hope to be able to show that the distribution of an OSN's state and traffic handling to the regional servers benefits both users and the OSN. We also hope to show that our findings generalize to other OSNs in follow up studies.

5. ACKNOWLEDGEMENTS

We would like to thanks Christo Wilson and Krishna Puttaswamy for useful discussions in the early stages of this project and for making available their crawls of Facebook.

6. REFERENCES

[1] N. Kennedy. Facebook's growing infrastructure spend. http://www.niallkennedy.com/blog/2009/03/facebook-infrastructure-financing.html, March 2009.

[2] J. Sobel. Engineering @ Facebook's notes: Scaling out. http://www.facebook.com/note.php?note_id=23844338919, August 2008.

[3] C. Wilson, B. Boe, A. Sala, K. P. N. Puttaswamy, and B. Y. Zhao. User interactions in social networks and their implications. In *EuroSys*, March 2009.