An analysis of distributed and asynchronous wireless group communication mechanisms

(Invited Paper)

Surendar Chandra (surendar@acm.org) and Xuwen Yu (xyu@vmware.com)

***Abstract*—Asynchronous group communication systems propagate updates from each group member to every other member. Contemporary users are wireless and operate from a variety of locations. Hence, we investigate the propagation performance for these users. An analysis of wireless users in academia, corporation and in a city-wide hotspot federation shows that the availability durations and the duration between sessions depends on the locale. A longer term analysis shows that the user session lengths are becoming smaller while the duration between sessions are becoming larger with significant user churn. We show that the propagation performance depends on the locale. New users incur a heavy cost in receiving updates from prior users; applications need to develop expiration mechanisms to reduce this overhead. Also, since users regularly leave the system, practical communication mechanisms need to choose their propagation periodicity in order to respond to scenarios when the user abruptly leaves the system without propagating their updates to other users.**

I. Introduction

Asynchronous group collaboration systems do not require the simultaneous availability of all the group members. They either use a centralized or a distributed group communication mechanism to implement their collaboration functionality. For example, Windows Live Mesh (mesh.live.com) and Apple iDisk (me.com) use a centralized mechanism while Windows Live Sync (mesh.live.com) uses a distributed approach.

Wireless laptops are ubiquitous. Global sales of laptops has exceeded that of desktops [1]. Laptops are resource rich. Distributed mechanisms among local workgroups can provide better capacity and reduce the wide area network cost. Hence, we focus on distributed approaches among wireless users.

Our analysis requires an understanding of when users are available for communication and when these users create updates that need to be propagated to other users. We use empirical wireless traces for the former and develop a metric that uses the duration that users are available for the latter.

Wireless users [2] work from a number of different locations: home, office as well as in public venues. Ideally, we need to simultaneously monitor all wireless locations (schools, offices, hotspots etc.) in a geographic locale in order to evaluate the system from the perspective of a mobile user who operates among these communities. However, such an analysis requires monitoring multiple administrative entities; currently unavailable to us and many researchers. An intrusive alternative is to install loggers in each monitored laptop. In this paper, we analyze the behavior of wireless users in different communities: a university, a corporate lab and in a city-wide hotspot federation at different times and locations (spanning 2001 through 2008 at several locations in North America). We used publicly available traces from the CRAWDAD [3] archive. We also collected application level traces from a university. Wireless access was *free* to its participants in all these traces; users can be expected to fully participate in a distributed collaboration scheme.

An analysis of some of these traces from a network capacity perspective had already been published [4], [5] by the original authors. We compared the different traces from the perspective of wireless group communication mechanisms. We show that the availability durations are longer in corporations followed by academia and then in hotspots. On the other hand, the duration between sessions are longer in hotspots followed by corporation and then in academia. Online sessions are becoming shorter while the duration between sessions is becoming longer. We observed a significant churn wherein nodes permanently left the system. We show that the amount of updates propagated increases tremendously for newer users. We answer important questions regarding the viability of group communications among wireless users.

We analyze wireless user availability in various locales and the performance of group communication mechanisms among these users in Sections II and III, respectively. We report related work in Section IV and conclude in Section V.

II. User availability analysis

We require information regarding when users became online and offline. We describe the locales where the wireless availability traces were collected.

A. Availability traces from various locales

We collected traces at a university and used traces from a corporate lab and a hotspot federation.

1) Academia: We collected availability data from Sep. 19, 2006 through Sep. 29, 2006 (800 APs, 2,036 users) as well as from Dec. 3, 2007 through Aug. 25, 2008 (1,300 APs, 4,063 users) using the Zeroconf *_workstation._tcp* service. We refer to these traces as *academia-2006* and *academia-2008*. The *academia-2008* trace included the end of the fall '07 semester, winter break, spring '08 semester, spring break as well as summer '08 sessions. We place particular importance on the two weeks starting at 12/04/2007 (when users were likely to be collaborating on course projects, though not necessarily using the mechanisms explored in this paper) as well as the 60 days starting from 12/04/2007 (includes the busy end of semester season, calm winter break as well as the beginning of a new

978-1-4244-7265-9/10/$26.00

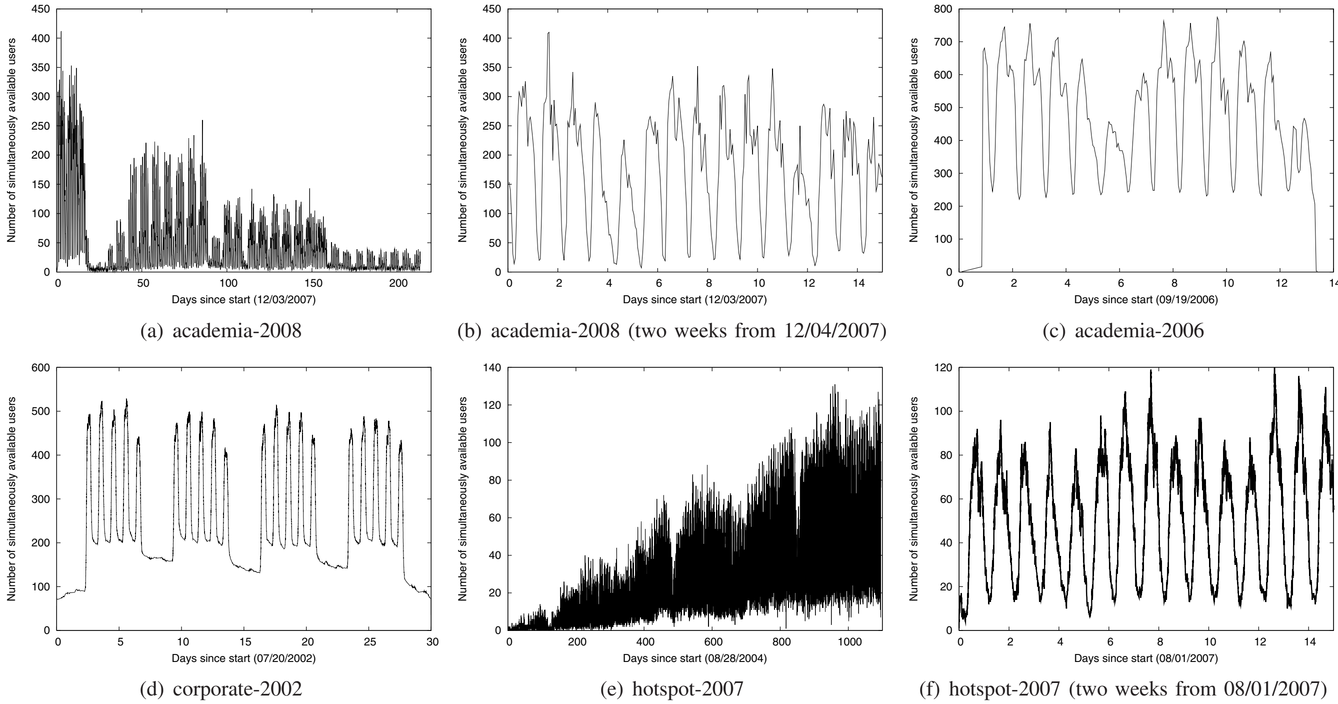

(a) academia-2008 (b) academia-2008 (two weeks from 12/04/2007) (c) academia-2006

(d) corporate-2002 (e) hotspot-2007 (f) hotspot-2007 (two weeks from 08/01/2007)

Fig. 1. Number of simultaneously available users

semester). Note that the network administrators progressively removed groups of APs into their own (unmonitored) VLANs during the beginning of the Spring '08 semester, Spring break '08 and Summer '08, reducing the number of simultaneously available users. Our analysis of each segment (not shown in this paper) showed that the user behavior was similar across all the times; albeit with fewer monitored users.

2) Corporation: Balazinska et al. [6] collected user availability traces (1,366 devices) at IBM Research from Jul. 22, 2002 through Aug. 17, 2002. They issued SNMP probes ranging from every 5 (55% of the traces) to 15 minute intervals. We refer to these access traces as *corporate-2002*.

3) Hotspot: $\hat{i}le\ Sans\ Fil$ (www.ilesansfil.org/welcome/) provides free coverage throughout Montreal using 206 hotspots. The access traces [7] last for three years from Aug. 28, 2004 through Aug. 28, 2007 (69,689 users). We place particular importance on the most recent two weeks (starting from 08/01/2007). We call these access traces as *hotspot-2007*.

B. Number of simultaneously available users

First, we analyze the number of simultaneously available users and plot the results by the time of the day for the various scenarios in Fig. 1. Similar to wired corporate desktops [8], users in all the traces exhibit a diurnal availability pattern.

We observed 2,036 new devices in *academia-2006* (Fig. 1(c)) and 2,729 new users in the first fifteen days of *academia-2008* (Fig. 1(b)). The number of simultaneously available wireless users decreased from 775 users (38% of the users) in Sep. 2006 to 410 users (15% of the users) in Dec. 2007.

We observed 1,366 unique devices in the *corporate-2002* traces (Fig. 1(d)). During the weekdays, the number of simultaneously available users were as high as 529 (38.7%). Even during late nights, we observe that over 200 users were available; likely devices without any users actively using them. Over the weekend, the number of devices dropped further to about 160; it is likely that about 40 users (who left their wireless devices at work during the weekday nights) took their wireless devices home for the weekend.

Finally, we plot the *hotspot-2007* traces in Fig. 1(e). These traces span over three years with 69,689 users observed during the entire duration. We observe a steady increase in the number of simultaneously available users. During the two week duration starting in 8/1/2007 (Fig. 1(f)), we observed 2,725 unique devices; about 120 of them (4.4%) were available together during the day times.

C. Online session duration and time between sessions

We plot the amount of time that a user was available and the time between sessions for the various traces in Fig. 2. The availability duration influences the amount of updates that can be created by the user while duration between sessions affects the amount of updates that need to be propagated to this user when they subsequently become online.

First, we plot the session behavior for the *academia-2008* and *academia-2006* traces in Fig. 2(a). We also plot the values for the two weeks starting in 12/04/2007. The session durations in *academia-2008* traces were short; 50% of the sessions were under 20 minutes and 95% of the sessions were less than 75 minutes. For the two weeks after 12/04/2007, we note that 50% of the sessions were under 20 minutes and 95% of the sessions were less than 70 minutes. Sessions from *academia-2006* were longer; 50% of the sessions were under one hour

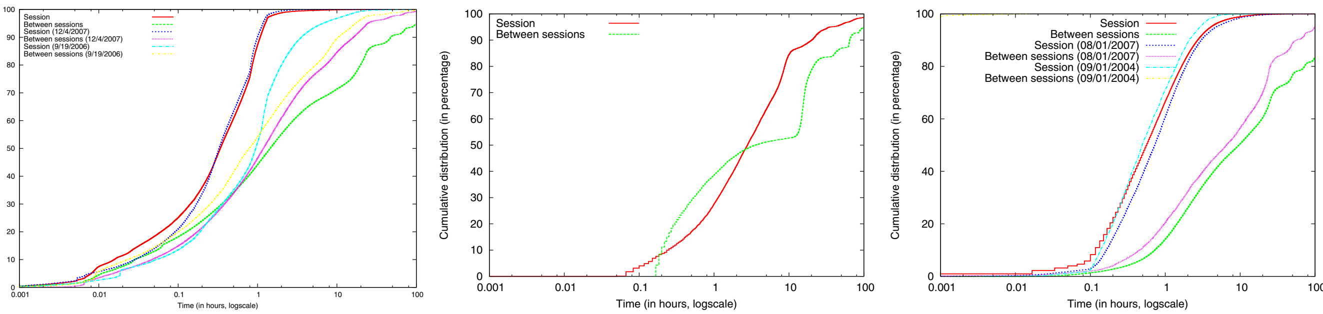

(a) academia-2008 and academia-2006 (b) corporate-2002 (c) hotspot-2007

Fig. 2. Online availability behavior

with 95% of the sessions were under 6.7 hours. The duration between sessions had increased between *academia-2006* and *academia-2008*. For the *academia-2008* trace, 50% of the duration between user sessions were less than 1.4 hours while 29% were longer than ten hours. Specifically focusing on the first two weeks in December 2007, these values were 1.2 hours for the median and 15% of the durations were longer than ten hours. On the other hand, *academia-2006* showed that 50% of the duration between sessions was less than 47 minutes, while 10.5% of the duration between sessions was longer than ten hours. Next, we analyze the duration between successive arrivals of a particular user in order to understand whether the shortening session durations equaled the increase in duration between sessions. For the first two weeks in December 2007, median values were 1.78 hours while 75% of users online every 5.5 hours. In 2006, the median values were 2.52 hours with 75% of users online every 6.9 hours. Even though the number of devices had increased (from 2,036 in *academia-2006* to 2,730 devices in the first two weeks of 2007), the session durations had decreased. Kotz et al. also observed a reduction in median session lengths from sixteen minutes in 2002 [9] to under ten minutes in 2004 [10]. This change might either be because users were using less computing time or because modern laptops offer reliable energy saving sleep modes and hence become unavailable for longer durations. Later, we show that this reduction in session durations reduces the performance of group communication systems.

For the *corporate-2002* traces (Fig. 2(b)), 50% of the sessions were longer than 2.8 hours with 95% of the sessions less than 36 hours. Analyzing the duration between sessions, we note that 50% of them were over 3.5 hours with 48% over ten hours. Analyzing the *hotspot-2007* traces from Fig. 2(c), we note that 50% of the sessions were smaller than 35 minutes with 95% of the sessions less than four hours. Analyzing the Verizon hot spot network in Manhattan, Blinn et al. [11] also observed that 45.74% of the sessions lasted more than one hour. The duration between sessions for our hotspot traces was over 9.6 hours for 50% of the cases with 49% of the duration between sessions longer than ten hours.

Next, we analyze the long term average daily session lengths for *academia-2008* and *hotspot-2007*. For *academia-2008* (Fig. 3(a)), the session lengths are longer during breaks

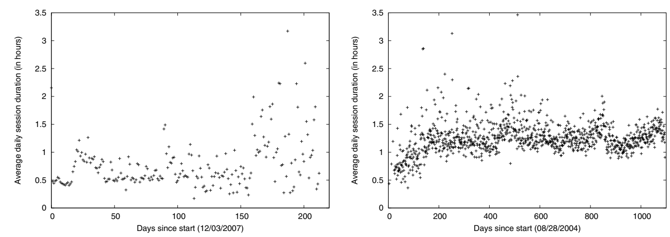

(a) academia-2008 (b) hotspot-2007

Fig. 3. Evolution of session duration

(weekend and semester breaks) than when the classes were in session. For the *hotspot-2007* traces (Fig. 3(b)), we note a slight increase in the session duration with less variability, especially when compared to the first six months. The *hotspot-2007* traces were collected among a federation of hotspots throughout Montreal. On the other hand, wireless hotspots in locations such as cafes and airports (where users are more mobile) are likely to show even smaller session durations.

D. User churn behavior

Finally, we investigate the appearance of new nodes as well as attrition of nodes for the various traces in Fig. 4. A new node will need to receive all the group communications from prior users. In the academic scenario, we expected new users to arrive at the beginning of a semester while generally remaining stable throughout the semester. Analyzing the *academia-2008* traces in Fig. 4(a), we note that over 2,750 users (of the total 4,000 users) were seen within the first few weeks. However, after the winter break (about 50 days since the start of the trace), new users were added steadily. On the other hand, about 500 users were never seen after the Fall 2007 semester (first few weeks). During the Spring 2008 semesters, users constantly left the system till about 150 days into the trace. Note that APs were removed from our monitoring abruptly during the spring break (and not gradually as indicated by Fig. 4(a)). We also observed a similar trend in the *corporate-2002* traces (Fig. 4(c)) where about 50% of the users appeared at the beginning of the traces (delayed because some users were likely offline during the start of the trace collection). However, the rest of the users continued to trickle in during

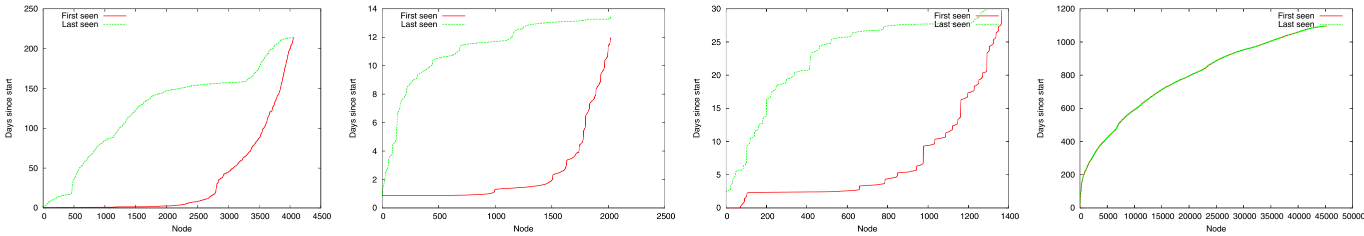

(a) academia-2008 (b) academia-2006 (c) corporate-2002 (d) hotspot-2007 (both the value almost overlap)

Fig. 4. User churn behavior

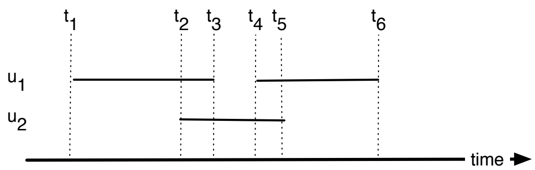

Fig. 5. Asynchronous update propagation between users u_1 and u_2

the remainder of the trace duration (with a two day jump, likely because of lull over the weekend). Finally, we observed constant churn for the *hotspot-2007* traces (Fig. 4(d)). This behavior is likely to be typical among hotspots.

III. Analysis of group communication mechanism

Next, we investigate the implications of user availability for group communication mechanisms.

A. System architecture

In distributed asynchronous group collaboration systems, each user maintains a local copy of the shared contents. When online, each user creates updates on these shared contents. Updates are first applied locally and then propagated to other group members where they are applied to their copies of the shared contents. In systems such as Bayou [12], out-of-order messages create database style update roll-backs and hence the propagation order is important. However, in our system, the update order in different user replicas is not important. The goal is to propagate updates to all the group members. Note that our system differs from delay tolerant routing [13], [14] protocols that propagate updates between a pair of nodes.

Gossip protocols [15] can propagate updates among all the group members; users who are online periodically exchange theirs as well as updates from other users that they had previously encountered. The system eventually reaches a consistent state without strict bounds.Given the network capacity of wireless LANs, the time to transfer updates is also negligible.

Next, we needed data about when users were available for propagation and when they created updates. All users in a trace are considered to be part of the collaboration group. Because of the lack of empirical data regarding update creation rates, we assume that all nodes create updates at a constant rate.

B. Metrics for asynchronous group communications

We evaluated the system by measuring the *amount of updates propagated* with time as well as the number of contributing nodes for a particular propagation. The amount of updates that are created in a particular session is proportional to the duration that a user is available. We measure the amount of updates that need to be propagated to each user as the amount of time that other group members were available since the last time that a particular user was available. We illustrate our metric for users u_1 and u_2 in Fig. 5. u_1 was available from t_1 to t_3 and again from t_4 to t_6. Similarly, u_2 was available from t_2 to t_5. The total amount of updates created by u_1 is directly proportional to the duration $(t_3 - t_1) + (t_6 - t_4)$. We measure the *amount of updates propagated* for u_2 at t_2 as $(t_2 - t_1)$ (updates propagated from u_1 to u_2), from t_2 to t_3 for both u_1 and u_2 as one (both were simultaneously available and so the updates can be continuously propagated. Note that practical policies only periodically propagate updates). At t_4, for the user u_1, the *amount of updates propagated* is $(t_4 - t_3)$, which is the amount of updates accumulated since user u_1 was unavailable. The number of contributing nodes in this example is one. An user u_3 coming online between t_5 and t_6 will receive updates from u_1 with the number of contributing nodes of two (u_1 and u_2).

C. Analysis

We compare the maximum and average *amount of updates propagated* as well as the number of users that contributed to the average and plot the results in Fig. 6. Note that the maximum value reflects updates propagated to new nodes joining the system (Fig. 4); these nodes require updates from all the prior nodes. For example, the thousandth user arriving on the fiftieth day needs updates from the prior 999 users that were accumulated over the initial fifty day duration. Also, note that this metric reflects all the updates accumulated by prior users and does not consider mechanisms described by Demers et al. [16] that remove older updates from further propagation. These removal policies are application specific.

First, we plot the results for the academic setting (Fig. 6). During the first fifteen days of *academia-2008* (which corresponds to the duration immediately prior to the final exams of Fall '07 term), the maximum values for the *amount of updates propagated* reaches over 2,000 days (Fig. 6(a)).

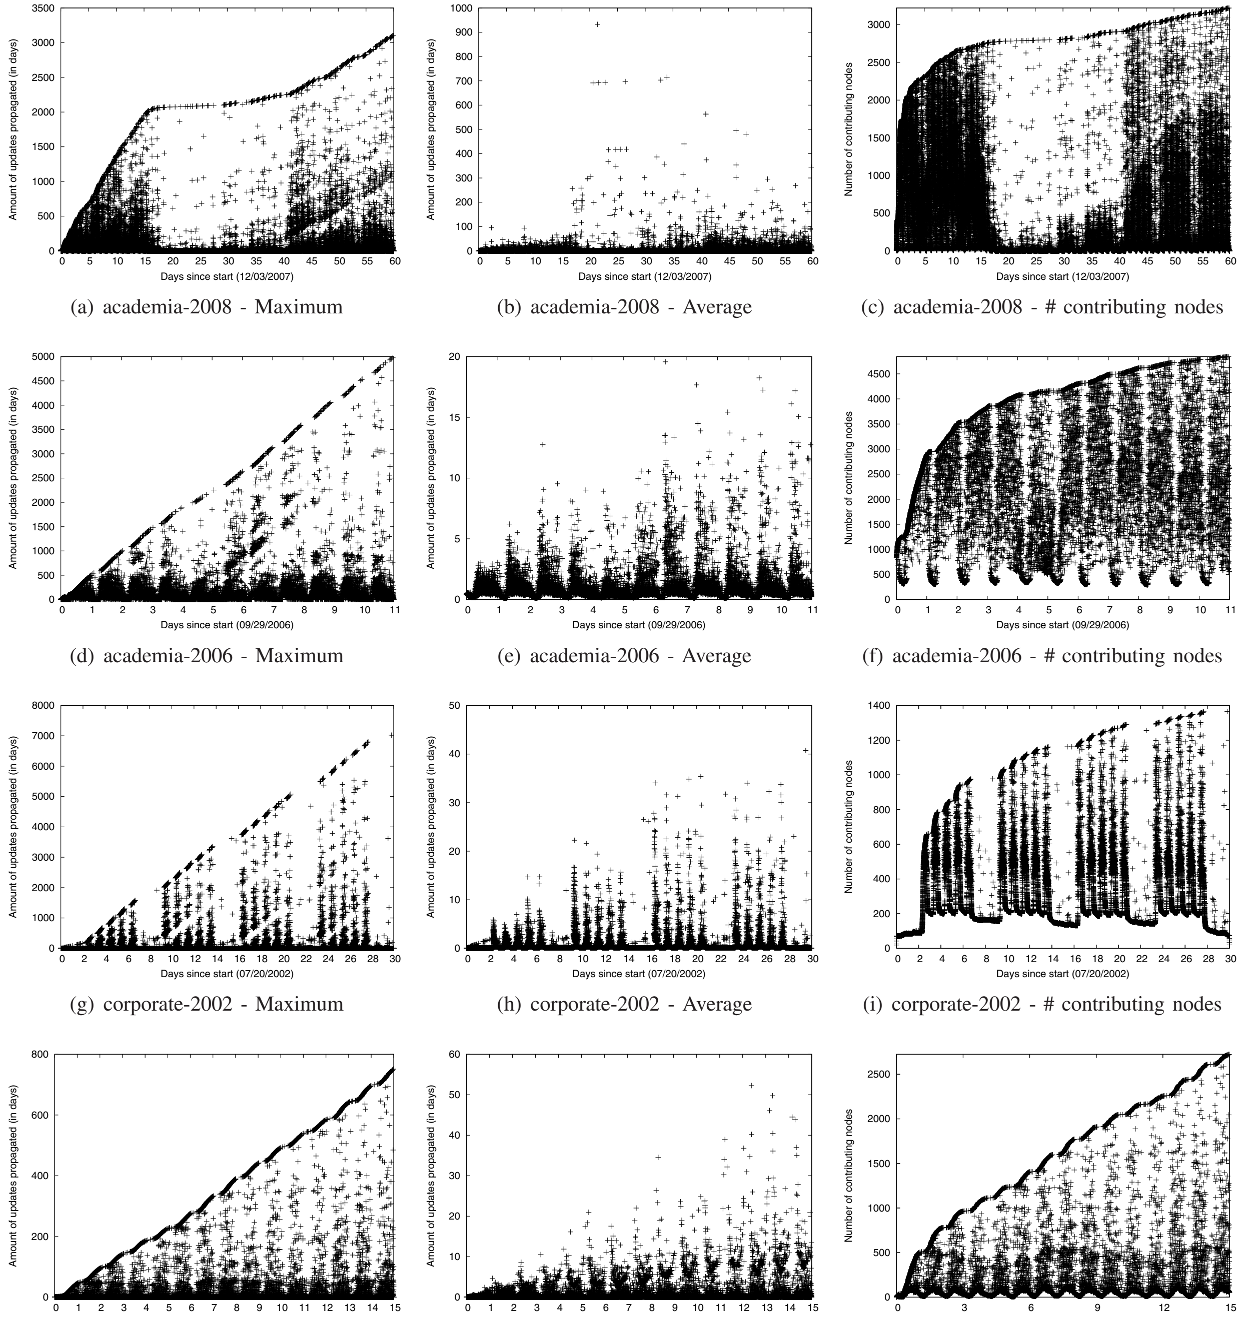

(a) academia-2008 - Maximum (b) academia-2008 - Average (c) academia-2008 - # contributing nodes

(d) academia-2006 - Maximum (e) academia-2006 - Average (f) academia-2006 - # contributing nodes

(g) corporate-2002 - Maximum (h) corporate-2002 - Average (i) corporate-2002 - # contributing nodes

(j) hotspot-2007 - Maximum (k) hotspot-2007 - Average (l) hotspot-2007 - # contributing nodes

Fig. 6. Performance of wireless and distributed asynchronous group communication systems

During this duration, 2,716 new nodes joined the system. A node joining on the 15th day will require all the prior updates for these fifteen days from the prior 2,715 nodes which translates to about 2,000 days worth of updates for a single node. Depending on the application dictates of the rate of update creation as well as the average size of the updates, asynchronous propagation requires large amounts of updates for newly arriving nodes. On the other hand, the average *amount of updates* only reached about 50 days (Fig. 6(b)), confirming the viability of distributed approaches. The number of contributing nodes (Fig. 6(c)) varies up to 2,716; these values lead to the low average values. Over the next three weeks, the campus was on winter break with few wireless users. The few new users coming online during this duration helped maintain the maximum values at about 2,000 days during the break. However, the average values during this duration was also as high as 950 days (i.e., average was much closer to the maximum). Even though newer updates were not created at the same pace as during the end of the term (fewer users - Fig. 6(c)), newer nodes still require relatively large amounts of updates. After the beginning of the Spring '08 semester, the average and maximum values reverted to similar behavior as in Fall '07. For *academia-2006*, we note that the maximum values for *updates propagated* which leads to more updates being created for nodes which enter the system while the average values were lower because more of the nodes were simultaneously available. Maximum values from Fig. 6(d)) can be as high as 5,000 days while the average values (Fig. 6(e)) are lower at around 20: the sessions lengths were longer and the duration between sessions were shorter (Fig. 2(a). We also observed the diurnal nature of the system: the average update counts are larger in early mornings when nodes become available and need the updates from users who were available during the prior night.

Larger session durations in *Corporate-2002* traces lead to maximum *update amounts* (Fig. 6(g)) as high as 7,000 days with an average value (Fig. 6(h)) of 25 during work days. Similarly, for the two weeks since 08/01/2007 in the *hotspot-2007* traces, the maximum and average values were 800 days (Fig. 6(j)) and 20 days (Fig. 6(k)), respectively. The number of contributing nodes reached 2,000 (Fig. 6(l)).

IV. Related work

A number of prior research efforts investigated the behavior of typical users in order to build group communications systems. For example, Bolosky et al. [8] analyzed the long term behavior of wired desktops in a corporate setting using ping messages. Their analysis [17] showed that the availability was sufficient to support up to four nines availability using three replicas in the Farsite [18] storage system.

Song et al. [14] used the AP SNMP records to synthesize contact patterns among wireless users. They notes that asynchronous update propagation can be unacceptably long, especially among casual users (some users might never meet each other in the future). We assumed the availability of a wireless infrastructure; either via access points or by using ad hoc/mesh networks. Given the prevalence of wireless access points, such an assumption is not unrealistic. This assumption allowed us to ignore the spatial mobility patterns of the users and focus on asynchronous collaboration; we consider any two nodes that were online anywhere on the wireless trace to be available. We expected our node availability to be far better than that was observed using user vicinity contact measurements. However, we observed limited system performance. This places serious doubts on the viability of mechanisms that require that collaborating group be co-located.

V. Discussion

Modern wireless devices are resource rich and ubiquitous. In this work, we analyzed the behavior of distributed and asynchronous group communications mechanisms at a university, corporate lab and a hotspot federation. Our analysis is agnostic to the requirements of specific applications, both in terms of the frequency of updates to the shared contents as well as the size of these updates. We show that the availability behavior was better in corporate settings where the users were available for longer durations with a larger percentage of the wireless users available on weekdays. The session durations were smaller in academia. We show that the session durations appear to become shorter which affects the performance of active and passive sharing systems. The systems exhibited constant node churn which places heavy load on asynchronous group communication systems that need to transmit prior updates to these newer nodes. Our work highlights the need for robust expiration mechanisms for older updates. The amount of updates created is small enough to make distributed approaches viable. We also offer some cautionary notes for using wireless group communications.

Acknowledgments: Kevin Smyth helped in trace collection. Supported in part by the U.S. NSF (CNS-0447671).

References

[1] iSuppli Corp., "Notebook pc shipments exceed desktops for first time in Q3," isuppli.com/NewsDetail.aspx?ID=19823, Dec. 2008.

[2] M. S. Rosenwald, "Digital nomads choose their tribes," *Washington Post*, Jul. 2009.

[3] D. Kotz and T. Henderson, "Crawdad: A community resource for archiving wireless data at dartmouth," *IEEE Pervasive Computing*, vol. 4, no. 4, pp. 12–14, Oct. 2005.

[4] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan, "Characterizing user behavior and network performance in a public wireless lan," in *ACM SIGMETRICS '02*, 2002, pp. 195–205.

[5] M. Balazinska and P. Castro, "Characterizing mobility and network usage in a corporate wireless local-area network," in *MobiSys '03*, 2003, pp. 303–316.

[6] ——, "CRAWDAD data set ibm/watson (v. 2003-02-19)," http://crawdad.cs.dartmouth.edu/ibm/watson, Feb. 2003.

[7] M. Lenczner, B. Gregoire, and F. Proulx, "CRAWDAD trace ilesansfil/wifidog/session/04_07 (v. 2007-08-27)," http://crawdad.cs.dartmouth.edu/ilesansfil, Aug. 2007.

[8] W. J. Bolosky, J. R. Douceur, D. Ely, and M. Theimer, "Feasibility of a serverless distributed file system deployed on an existing set of desktop pcs," in *ACM SIGMETRICS*, 2000, pp. 34–43.

[9] D. Kotz and K. Essien, "Analysis of a campus-wide wireless network," in *ACM MobiCom '02*, 2002, pp. 107–118.

[10] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," in *MobiCom '04*, 2004, pp. 187–201.

[11] D. P. Blinn, T. Henderson, and D. Kotz, "Analysis of a Wi-Fi hotspot network," in *International Workshop on Wireless Traffic Measurements and Modeling (WiTMeMo '05)*, June 2005, pp. 1–6. [Online]. Available: http://www.cs.dartmouth.edu/~dfk/papers/blinn:hotspot.pdf

[12] K. Petersen, M. J. Spreitzer, D. B. Terry, M. M. Theimer, and A. J. Demers, "Flexible update propagation for weakly consistent replication," in *SOSP '97*. Saint Malo, France: ACM, 1997, pp. 288–301.

[13] S. Jain, K. Fall, and R. Patra, "Routing in a delay tolerant network," in *Sigcomm '04*, 2004, pp. 145–158.

[14] L. Song and D. Kotz, "Evaluating opportunistic routing protocols with large realistic contact traces," in *ACM MobiCom workshop on Challenged Networks (CHANTS 2007)*, September 2007.

[15] K. Birman, "The promise, and limitations, of gossip protocols," *SIGOPS Oper. Syst. Rev.*, vol. 41, no. 5, pp. 8–13, 2007.

[16] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry, "Epidemic algorithms for replicated database maintenance," in *ACM PODC*, Aug. 1987, pp. 1–12.

[17] J. R. Douceur and R. P. Wattenhofer, "Optimizing file availability in a secure serverless distributed file system," in *20th IEEE Symposium on Reliable Distributed Systems*, New Orleans, LA, Oct. 2001, pp. 4–13.

[18] A. Adya, W. J. Bolosky, M. Castro, G. Cermak, R. Chaiken, J. R. Douceur, J. Howell, J. R. Lorch, M. Theimer, and R. P. Wattenhofer, "Farsite: federated, available, and reliable storage for an incompletely trusted environment," *SIGOPS Oper. Syst. Rev.*, vol. 36, no. SI, pp. 1–14, 2002.