

Delay tolerant collaborations among campus-wide wireless users

Xuwen Yu and Surendar Chandra

Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556, USA
{xyu2,surendar}@nd.edu

Abstract—The ubiquitous deployment of wireless LAN networks are allowing students to embrace laptops as their preferred computing platform. We investigated the viability of building collaborative applications to share contents amongst student groups. In our application scenario, the university will provide wireless infrastructure throughout the campus but not the storage infrastructure required to store the shared contents. Laptops will likely exhibit weak availability. Hence, these collaborative applications need to tolerate long delays in propagating updates amongst the participants. In this paper, we presented a preliminary analysis of message forwarding behavior under realistically resource constrained node scenarios. Our experiments were based on the observed wireless user behavior at the University of Notre Dame. Our experiments showed the inherent limits of epidemic propagation in real campus wireless network scenarios.

I. MOTIVATION

Wireless laptops are gradually replacing desktops as the primary computing platform for many users. USA today [1] described the emergence of about 30 million (American) mobile laptop users. At our university, there were 12,252 (as of Nov 1, 2007) wireless devices. We expect similar trends among user communities in large enterprises as well.

Newer laptops are matching the resources available in a desktop device; boasting resources as high as 2.6 GHz dual core processors, 300 GB of hard drive storage and 4 GB of main memory. Traditionally laptops were dependent on storage services provided by a wired infrastructure. Mobile laptops may be offline for long durations making any dependence on infrastructure storage particularly problematic. Besides, it is harder to provide the collaborative storage infrastructure that can match the aggregate contents in these resource rich laptops. For example, even if each laptop in our university volunteered as little as 10 GB of its local storage, the aggregate capacity can easily reach over 100 TB. Our goal was to understand the viability of collaboration among wireless users without depending on the wired storage infrastructure.

Laptops go through periods of disconnection. Hence, collaborative applications among laptops need to operate using local copies of the data. Since all the group members are not simultaneously available, epidemic algorithms [2] can be used to asynchronously propagate local updates to other group members when they are available. Depending on the node availability, updates among the nodes can be delayed for long durations. The local data may not be consistent across all the participating nodes. Such a collaborative system critically depends on the rates at which updates are created, update size

(updates are propagated to all the group participants), as well as by the node availability. It is important to analyze the rates at which updates can be propagated in a real deployment. Earlier research efforts did not have the benefit of the critical mass of available wireless devices for their analysis. We used wireless LAN users at the University of Notre Dame, as the basis for our study. Unlike Hui et al. [3], we do not require that the collaborators be in close proximity to each other. We assumed that any two devices that were simultaneously online can communicate with each other; either using the university wireless infrastructure or through ad hoc networks. We also assumed that the infrastructure itself did not provide any storage resources to assist in message propagation.

In a companion paper, we investigated the availability characteristics of wireless devices at Notre Dame. We observed that wireless devices tended to exhibit short durations during which they were available followed by extended durations when they were not available. The churn frequency itself was not high. The node availability exhibited diurnal distribution with far fewer nodes available early in the morning. However, we showed that the temporal consistency values were high: both for analyzing the same users availability behavior or for any two pairs of users. Users who were part of the high consistency set can provide better collaborative services.

In this paper, we analyzed the epidemic propagation rates among these users for varying group sizes. Our preliminary analysis showed that the propagation can exhibit large delays. On average, a single update can reach about 60% of the collaborators in about 24 hours while reaching 90% of the members in over ten days. Buffer constraints on the intermediate nodes severely affected the propagation durations. Unlike Vahdat et al. [4] who used a random node mobility model, our realistic analysis could not achieve 100% propagation rates even after ten days. Further work is needed to improve realistic routing mechanisms for delay tolerant applications.

The rest of the paper was organized as follows: we describe related work in Section II. We describe the system architecture in Section III with results from our experimental analysis in Section IV. We conclude in Section V.

II. RELATED WORK

We build on advances in understanding of mobile disconnected access and in delay tolerant networks (DTN).

Developed over a decade back, applications such as Bayou [5] operated without centralized storage infrastructure and

used epidemic algorithms [2] to propagate updates. However, the behavior of these systems fundamentally depends on the mobility patterns of typical users. Wireless laptops are recently far more ubiquitous than when these systems were developed. Our focus in this paper was to analyze the availability patterns of modern wireless users in order to understand the expected propagation rates.

Vahdat et al. [4] used epidemic routing to propagate updates in an ad hoc networking scenario. They simulated a random node mobility pattern and analyzed the propagation behavior by varying the radio range. They showed that epidemic routing achieved eventual delivery of 100% of messages. Similarly, Davis et al. [6] investigated propagation among wearable computers using simulated human mobility. They investigated the effects of message duplication and buffer overhead. Recently, delay tolerant network technologies (DTN) are used to asynchronously propagate updates among a set of clients. Fall [7] introduced a network architecture that operated without continuous network connectivity among the participating nodes and investigated [8] the routing behavior across a DTN. They used simulations and progressively increased the amounts of network topology information available to the routing mechanism. They showed that the systems performed better with the addition of more topology information. Our primary goal was to validate these propagation rates using realistic node mobility behavior rather than by using simulated behavior.

Chaintreau et al. investigated [9] opportunistic forwarding algorithms using the captured contact information of volunteer conference attendees at INFOCOM 2005 [3]. Similarly, Song et al. [10] used the access point records to collect contact patterns among wireless users. They observed that the epidemic propagation can be unacceptably long, especially among casual users (some users might never meet each other in the future). We consider that any two nodes that are online are also accessible to each other either by using the campus network infrastructure or by using ad hoc routing mechanisms. Hence, our node availability was expected to be far better than that was observed using user contact measurements. However, even with this improved availability, we observed poor update propagation performance. Our work places serious doubts on the viability of many prior systems.

III. SYSTEM ARCHITECTURE

We target scenarios where collaborations were effected among a group of users whom belonged to a larger community of users. Students collaborating on a project with students from the same class is an example of such a system. Project members in a corporate setting could be expected to exhibit similar application requirements. The members of this group communicate with other online users using the campus network infrastructure or by using wireless ad hoc networks. We assume that the infrastructure itself does not provide any storage facilities to hold the updates for propagation to other users. The viability of this system depends on the expected mobility pattern of the individual users. For example, users in

a corporate setting could be expected to be available during the 9 AM - 5 PM duration while students in a campus setting may be available throughout the day, especially if the wireless was also available in the dormitories.

A. Experiment setup

We analyzed the behavior of WLAN users at Notre Dame. Notre Dame university used over 800 access points to provide coverage in residence halls, class rooms and the laboratories. The WLAN network was widely used by undergraduate students, graduates, faculty, staff and guests. For our study, we used the Zeroconf [4] service discovery protocol to collect the usage statistics of Notre Dame campus WLAN users. The discovery protocol itself pushed the service availability information to the monitoring client using link local multicast. Since these multicast packets are not routed, we require the monitoring station to be co-located inside the monitored wireless VLAN. The data collection lasted for eleven days from September 19, 2006 through September 29, 2006. During this duration, the entire campus wireless LAN infrastructure was configured to route all Zeroconf service discovery packets to the monitoring station. This allowed us the flexibility of not installing a monitoring station inside each of the campus WLANs. We used the *dns-sd* tool to monitor the *workstation.tcp* file sharing service used by Mac OSX and Linux clients. This service was discovered whenever the machine booted to an useful state. We captured over 2,000 unique hosts. This number was a significant portion of the entire university population.

In a companion paper, we describe the availability characteristics of wireless campus users. We analyzed the online as well as offline durations. Excellent update propagation rates require that all users be simultaneously online. We observed that the users tended to be available for relatively short durations followed by long durations when they were unavailable. The churn frequency was not high.

IV. RESULTS

We performed experiments to answer questions such as: Given the user mobility characteristics, what were the expected propagation rates for using epidemic algorithms? The propagation rates likely depend on the group size.

A. Message propagation behavior

First, we investigated the suitability of these wireless users for asynchronous collaborations. Epidemic algorithms [2] asynchronously propagated individual updates to other users who were simultaneously accessible. Any such updates were applied against the local replica. The participating nodes maintained information about the current updates that were successfully applied locally. The amount of time it took for an update to reach all the members of a collaborative group depended on the user availability. We investigated the epidemic propagation rates for our user population. We analyze the system for propagating a single message as well as in considering the buffer constraints at the intermediate nodes.




Fig. 1. Epidemic propagation rate for various group size

In order to evaluate the message propagation characteristics, first we need to understand the number of students in a course. Students registered to a particular class can be expected to share contents with other students. Large classes are likely to provide more replicas. We analyzed the number of students registered in each class at Notre Dame for the Fall 2006 semester. We noted that 75% of the classes had six or more students, while about 25% of the classes had thirty or more students. For the rest of the paper, we analyzed the system for group sizes of six and thirty.

1) *Epidemic propagation rate for a single message:* First, we analyzed the epidemic propagation rate for sending a single update to other members. We defined epidemic propagation rate as the percentage of members who had received the update. We expected the epidemic propagation rate to depend on the group size and the user availability durations. Larger groups increased the likelihood of finding other group members who already carried the update. For simplicity, we assumed that the updates were transferred instantaneously; if two users were simultaneously online we assumed that both of them received the updates from each other. Realistically, this duration depended on the network used and the update size. With the increasing prevalence of IEEE 802.11n networks, this was not an unreasonable assumption. For our experiments, we initiated the epidemic propagation process from a random node and computed the average epidemic propagation rates to reach all group members. We repeated this process for a thousand times and plotted the average cumulative distributions of the time it took to propagate the update to all nodes for group sizes of six, thirty and all users in the system in Fig. 1. From Fig. 1, we note that the epidemic propagation rate was reasonably constant across the various group sizes. In general, the propagation rates were poor, reaching about 20% of the group members in a few hours. However, even after 24 hours, the updates had not propagated to over 40% of the users. Surprisingly, including all the users only marginally improved the epidemic propagation rate; reaching 65% of the members in 24 hours (as opposed to 50% for a group of size six). The

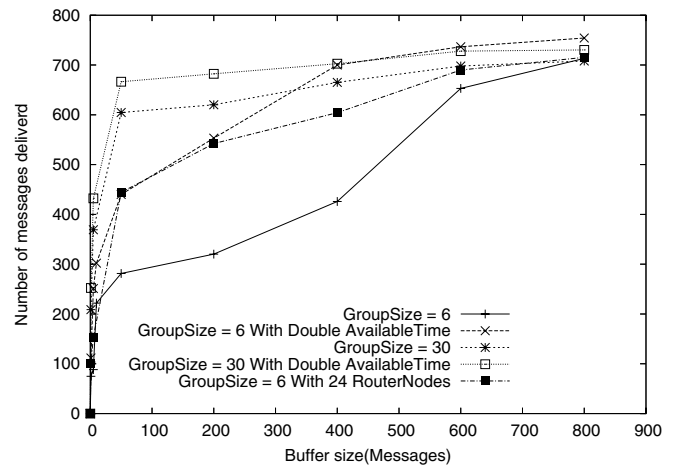


Fig. 2. Epidemic propagation rate for constrained local buffers

epidemic propagation rates were influenced by group members who were unavailable for large durations.

2) *Epidemic propagation with constrained buffer size:* In general, epidemic algorithms were restricted by the availability rates of nodes that were unavailable for long durations. The epidemic propagation rates can also be expected to be constrained by the amounts of storage available to hold the updates at the intermediate nodes. In order to understand this effect, we conducted experiments by varying the amount of buffer space available at each node (buffer space expressed in message count, similar to Vahdat et al. [4]). For this experiment, we considered the propagation of 800 messages that were randomly generated by each participant to the various group members for group sizes of six and thirty. Each of the group participants were restricted by the amount of messages that they can hoard. Once this buffer limit was reached, the local node will preempt the oldest message, effectively making this update unavailable for propagation to other group members from this node. We ran the experiments for a hundred different groups and reported the average epidemic propagation rate at the end of the eleven day data collection interval. We plotted the results in Fig. 2. We observed that for a group of size six, with a buffer size limit of fifty messages, the epidemic propagation rate was about 275 messages in eleven days. Using a buffer size of 600 messages, 600 messages were delivered by the end of eleven days. For a group size of thirty, the propagation rates were higher even with smaller local buffers because the increased group size allowed for retrieving the updates from another node. In order to understand the effects of the node availability, we repeated the experiments by artificially doubling the observed node availability durations. Increasing the availability rates increased the epidemic propagation rates. For a group size of six, with a buffer size of fifty messages, 450 messages were delivered to all members in eleven days.

We also introduced routing nodes where a group of six nodes utilized the services of 24 random nodes to ferry updates to other members of the group. We observed little improve-

ment in the message delivery durations for this scenario; for a limit of a fifty message buffer, 450 messages were delivered in eleven days (as opposed to 275 messages for a group size of six and similar to the behavior for group size of thirty).

Unlike the observations by Vahdat et al. [4], we never achieved 100% delivery rates. Vahdat simulated a random node mobility model. Our results show importance of experimentally validating the propagation rates with real user mobility patterns. Since node propagation did not improve significantly by either (artificially) increasing the node availability or by using random router nodes, we investigated choosing router nodes specifically in the next section.

B. Improving message propagation using node correlation

In the previous section, we observed that propagation rates were low. Next, we investigate mechanisms that can improve these rates by selectively choosing the router nodes. We analyze the user temporal behavior and the correlation among different users to choose good nodes.

1) *User temporal behavior*: First, we analyzed the system to see if users exhibited predictable behavior by being consistently available at the same time every day (for the eleven days that we investigated). Suppose a significant percentage of the users were available at (say) 10 AM. One can then imagine achieving good propagations rates during this time on every day (when other users were also simultaneously online). On the other hand, if we noticed that a significant number of users were not consistently available at (say) 10 AM, then one can safely ignore this time period. Even though one cannot build a collaborative system at this particular time, it may not matter much because there were no other users who could generate or consume updates.

For this analysis, we chose three particular times of the day; 3 AM (late night), 3 PM (work time) and 9 PM (evening). We analyzed the users who were available at this time in all of the eleven days (that we monitored the system). Note that we would not count an user who become unavailable at say 2:59 AM. We tabulated the results for analyzing both as counts, as well as percentages in Table I. The table should be read as follows: 59.5% of the users (1,211) were never seen at 3 AM, 11.2% (227) were at least seen once, while 3.59% (73) were seen at 3 AM on all eleven days. From Table 1, we note that a significant number of users (59.5% (1,211)) are never seen at late night. As a percentage of the total user population, there is little temporal available consistency wherein the users were always available at the same time in all the eleven days.

In order to further understand the user behavior, we computed the rate of available time consistency as follows: we define consistency at a specific hour by the metric that the user will either be consistently available or unavailable on all the eleven days. For example, if either the user was not available on all eleven days or the user was unavailable on all eleven days at a specific time, we compute the consistency at that time as one. If the user was available for half the time and unavailable for the other half, then the consistency is zero. We compute the consistency values for the user in




Fig. 3. rate of available time consistency

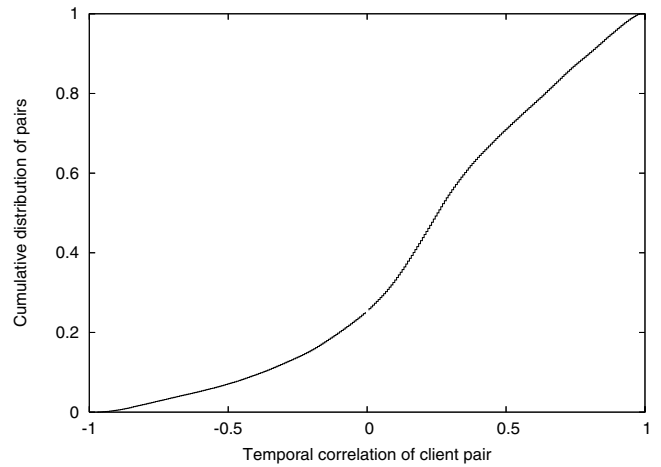


Fig. 4. Distribution of temporal correlation value for all machine pairs

steps of one hour for the entire 24 hour day and normalize it by dividing by 24. We plot a cumulative distribution of the rate of available time consistency in Fig. 3. A value of one indicates that all users were consistent (always available at all times or unavailable at all times) whereas a value of zero implies that users are equally likely to be either available or unavailable with no consistent way to predict their behavior. A collaborative system prefers these values to be one so that updates can be propagated to other users in a timely fashion. From Fig. 3, we note that half the users had consistency values of over 0.7. Only 5% of the users had values of 0.4 or lower. These values suggest that, even though there may not be many users who are available at all times, a large number of users were predictable in terms of times that they are available (or unavailable) and a small number of users consistently form a recurring group. This behavior has significant implications for developing targeted router nodes.

2) *Temporal correlation of user pairs*: Next, we plot the cumulative distribution of the temporal correlation value for all the user pairs in Fig. 4. This function was previously

Time of day	Percentage of users seen in n days in all the eleven days											
	$n = 0$	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$	$n = 6$	$n = 7$	$n = 8$	$n = 9$	$n = 10$	$n = 11$
3 AM	59.5	11.2	7.17	4.28	3.00	2.26	1.96	1.57	2.36	1.57	1.62	3.59
3 PM	16.2	11.7	9.04	9.19	6.83	7.67	6.49	7.27	7.91	6.54	5.36	5.85
9 PM	34.2	11.3	5.31	5.06	5.11	5.35	5.65	5.99	5.90	7.08	5.01	4.12

TABLE I
NUMBER OF TIMES THAT USERS ARE REPEATEDLY SEEN DURING THE ENTIRE ELEVEN DAY TRACE PERIOD

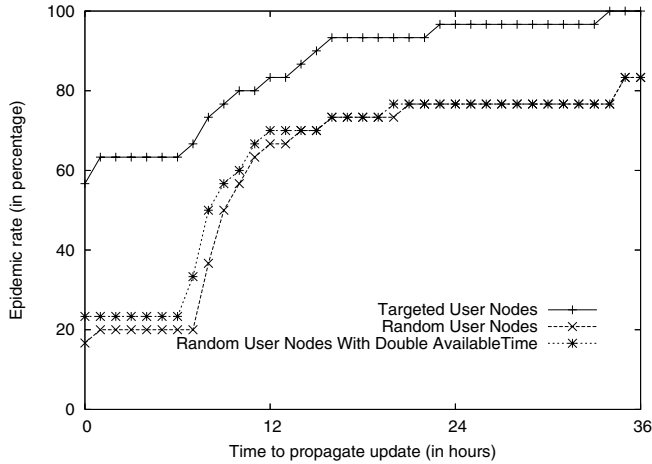


Fig. 5. Epidemic propagation rate for targeted group of size thirty

described by Bolosky et al. [11]. We computed the temporal correlation value for all user pairs as follows: add one when two users were either simultaneously available or unavailable and subtract one when only one of the two users were simultaneously available. We sample the system every hour. The results were normalized by dividing by the total sample count. As we described in Section IV-B1, we prefer values of one as it suggests that the pairs of users are either both available or unavailable. On the other hand, a value of -1 suggests that the behavior of the pair of users are unpredictable. From Fig. 4, we note that over 50% of the users have a temporal correlation of user pair values of 0.25. As compared to the observations on a corporate desktop [11], which observed values of 0.5 for 50% of the users, our scenario showed a lower correlation.

3) *Applying node temporal behavior to improve epidemic propagation rates:* Next, we applied the observed temporal behavior of the nodes to improve the propagation rates. Epidemic propagation rates were affected by the availability of the router nodes. Propagating to a node that was always available greatly improved the overall propagation rates. Users who exhibited predictable availability were also expected to have good propagation rates. We chose a random group of thirty users during the work-week as follows: throughout the five days (Mon-Fri), at least 50% of the users from this group were simultaneously available. We plotted the propagation rates amongst these users in Fig. 5. Note that not all users are simultaneously available. Still, the propagation characteristics of this set was significantly better than for choosing random set of users (in Fig. 2); reaching 66% of the group almost

instantaneously and reaching 95% in about 16 hours. Similar improvements were also observed against artificially doubling the availability duration among random users.

V. CONCLUSION

This paper analyzed the viability of developing delay tolerant collaborative applications among wireless users in a university campus setting. Earlier systems that used random node mobility models had shown that updates in such systems can be propagated quickly using epidemic algorithms. We showed that the achieved epidemic propagation rates were far worse than was reported in literature. Further work should focus on developing robust mechanisms that can improve propagation rates in realistic deployments.

ACKNOWLEDGMENTS

This work was supported in part by the U.S. National Science Foundation (CNS-0447671).

REFERENCES

- [1] M. R. della Cava, "Working out of a 'third place,'" USA Today, Oct. 2006, 10/5/2006.
- [2] 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 Symposium on Principles of distributed computing*, Aug. 1987, pp. 1–12.
- [3] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *2005 ACM SIGCOMM workshop on Delay-tolerant networking*, Aug. 2005, pp. 244–251.
- [4] A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," Duke University, Tech. Rep. CS-2000-06, Jul. 2000.
- [5] A. Demers, K. Petersen, M. J. Spreitzer, D. Terry, M. Theimer, and B. Welch, "The bayou architecture: support for data sharing among mobile users," in *Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, Dec. 1994, pp. 2–7.
- [6] J. A. Davis, A. H. Fagg, and B. N. Levine, "Wearable computers as packet transport mechanisms in highly-partitioned ad-hoc networks," in *IEEE Intl. Symp. on Wearable Computers*, 2001, pp. 141–148.
- [7] K. Fall, "A delay-tolerant network architecture for challenged internets," in *SIGCOMM '03*, 2003, pp. 27–34.
- [8] S. Jain, K. Fall, and R. Patra, "Routing in a delay tolerant network," in *Sigcomm '04*, 2004, pp. 145–158.
- [9] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on the design of opportunistic forwarding algorithms," in *INFOCOM 2006*, Barcelona, Spain, Apr. 2006, pp. 1–13.
- [10] 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.
- [11] 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 *SIGMETRICS '00: Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, 2000, pp. 34–43.