Optimizing Mobility Support in Large Switched LANs

Cristina Hristea and Fouad Tobagi

Stanford University

Abstract

With the recent explosion in wireless mobile communication, a great deal of work is taking place to solve the problem of handling mobile users in the Internet. This problem encompasses the design of schemes for tracking and routing to mobile users in various networks and subnetworks that make up the Internet infrastructure.

Switched LANs today employ the transparent learning protocol which allows switches to learn about the location of users by promiscuously listening to the traffic emitted by the users. Unfortunately, as this paper will show, fast moving users render the transparent learning protocol inefficient so that high bandwidth consumption and even packet loss may occur. To prevent packets form being mis-routed or even lost, one can employ a lightweight control-based mechanisms for user tracking, based on an existing LAN-protocol known as the Generic Attribute Registration Protocol (GARP). However, this solution can lead to large databases at the switches, which can in turn significantly increase the cost of the Internet infrastructure.

In this paper, we propose to combine and tune the transparent learning protocol and the GARP protocol in an efficient way as to jointly minimize packet loss, bandwidth consumption and the size of the tracking databases in a large switched network. Our optimization is shown to be effective for a wide range of mobility speeds and application characteristics, and consistent with the state of technology.

1 Introduction

Wireless networks such as the 802.11 have become widely deployed and used over the last decade. Today, 802.11 wireless cells populate offices, homes, schools, airports and even coffee shops. Given the convenience of untethered communication, the future will bring the deployment of 802.11 networks over large areas, such as across cities, anywhere from small to large and educational and industrial campuses and plants. As 802.11 networks become larger, supporting many fast moving users and increased application demands, a scalable design of the 802.11 network becomes paramount for ensuring that high quality of service is delivered to each individual user and their applications.

Unfortunately, current technologies significantly limit the number of individual addresses that can be tracked by a switch, as well as the bandwidth available for switching. For example, a single memory chip can store up to a few ten thousand entries and can be accessed at speeds of a few tens of MHz [4]. At an increased cost in the overall solution, multiple chips may be used for tracking mobile users at a given switch in a large network. Nonetheless, to minimize cost, it remains crucial that the tracking and routing schemes operate efficiently to reduce the number of tracking entries and the bandwidth consumption at the switches in the network.

Switched LANs today employ the transparent learning protocol which allows switches to temporarily track the location of users by promiscuously listening to the traffic emitted by users. Fundamentally, the transparent learning protocol reduces the need for packet flooding in the network, thus aims to minimize bandwidth consumption, while also maintaining small tracking databases at the switches. Unfortunately, as we will show in this paper, fast moving users render the transparent learning protocol inefficient, so that high bandwidth consumption and even packet loss may occur. To prevent packets form being mis-routed or even lost, one can employ a lightweight control-based mechanisms for user tracking, based on an existing LAN-protocol known as the GARP. However, this solution can lead to large databases at the switches if appropriate care is not taken.

In this paper, we propose to combine and tune the transparent learning protocol and the GARP protocol in an efficient way as to jointly minimize packet loss, bandwidth consumption and the size of the tracking databases in a large switched network. Our optimization is shown to be effective for a wide range of mobility speeds and application characteristics, and consistent with the state of technology. In the rest of the paper, we describe the state-of-the-art and present our proposed schemes. In Section 3 we describe the simulation environment and the results achieved for different mobility and application characteristics. In Section 4 we conclude the paper.

2 Schemes

2.1 LAN Switching (SNOOP)

In LAN networks, switches distributedly forward data packets to users. For this purpose, switches learn about the location of users via a transparent learning scheme (also known as snooping), which works as follows: switches listen to traffic originated from users and record the port via which the users can be reached. A lifetime is associated with each entry to ensure they do not become stale in face of changes in the network or user states. A packet addressed to a user that does not currently appear in the forwarding tables of a switch will be broadcast across all the ports of that switch (other than the port via which the packet was received). It is then assumed that this destination user will reply to the broadcast and that future packets addressed to it from the same source will not be broadcast.

Unfortunately snooping works on the assumption that information gathered in the past can be used for some period of time in the future. However, when users in the network are very mobile, one of two things may happen: 1) entries become stale (large lifetime values), thus packet loss may occur or 2) entries expire early (short lifetime values), resulting in a large of amount of data broadcasts. These problems can be alleviated if mobiles send data packets frequently (relative to the rate of movement), constantly updating the tracking databases, however this may not be always possible (e.g. portable devices where data transmissions can be expensive in terms of battery power).

A key question becomes at what values should the lifetime values be set so as to achieve low packet loss and still limit bandwidth consumption at the switches? We identify this value by running numerical simulations for various user mobility and application characteristics.

2.2 REG

A different method for forwarding in switched LANs, which is currently used only for multicast applications and the formation of virtual LANs (VLANs), is to explicitly register users with the appropriate switches via a registration scheme designed especially for that purpose, known as GARP [1]. Users spread across the switched LAN can join a multicast group and can receive packets that are sent to that group. A message sent to a multicast group does not embed the address of the destination users. Instead, a multicast address is used to represent a group of member users, and switches assume the burden of mapping that address to all the group members and delivering the group packets to them. Switches learn about the arrival and departure of members via explicit registration, deregistration and acknowledgment messages sent by the (soon to become) members. Like in the transparent learning scheme, lifetimes are associated with each forwarding entry to ensure that the databases do not become stale in the event that member users are no longer alive or the network configuration has changed. Unlike in the transparent learning scheme however, tracking here happens pro-actively, even in the absence of any data communication. Also, all users are tracked at all the switches in the active topology of a LAN, including switches that do not belong to the path that connects users with their communicating counter-parties. These features are necessitated by the nature of the application which supports the open host group concept [1]. One drawback of this design is the large size of the tracking databases when many populous multicast groups or VLANs are simultaneously active.

It is interesting to ask whether a registration-based scheme like GARP can be used to track mobile users in a given network. Unlike in GARP however, registering all the users with all the switches in the active topology may not be desirable, since it can result in large tracking databases at the switches and registrations that take a long time to complete. Instead, only a subset of the switches may be used to track a given user, which we refer to as the write set of that user. Switches outside the write set for a given user simply broadcast packets sent for that user.

In our design, we consider the write set of a user to be the subset of switches in the neighborhood of the user that are located at a distance d from the user, such that d is smaller than some radius R. When a user turns on or when it arrives at some switch S, it registers with S and all the switches in the network that are separated from S by no more than R hops. When a user moves from some switch S1 to another switch S2, it first deregisters itself from the write set at S1. An acknowledgment is sent in the opposite direction of the deregistration message. Upon the receipt of the acknowledgment, the user registers itself with the write set at S2. To simplify the implementation, we opt not to use timers to invalidate entries at the tracking databases.

Evidently, the explicit control messages add overhead in the network which increases with R and with the user speed. The number of data broadcasts decreases with R while the size of the tracking databases and the time to complete a location update increases with R. To identify the optimal value of R that balances all these factors, we simulate REG for various application and mobility characteristics.

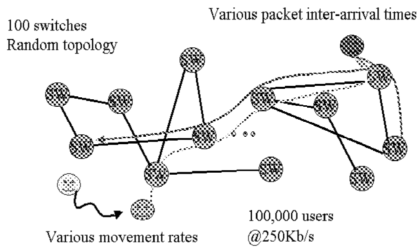

Figure 1: The simulation model.

2.3 REG_SNOOP

Since they are not mutually exclusive, it becomes interesting to ask whether the REG and SNOOP schemes can be used in combination to create an optimal scheme for tracking and routing to mobile users. We notice that REG can result in a large number of data broadcasts when the size of the write set is small. By using snooping outside the write set of a user, we can limit the number of data broadcasts that take place outside the write set of a user when routing packets to that user. Furthermore, if snooping takes place far away from the user, then the lifetimes values can be higher (e.g. mostly determined by topological network changes), hence fewer data broadcasts may occur. However, by tracking users both inside and outside some write set, the average size of the tracking databases may increase. Furthermore, packet loss may also increase (compared to the REG scheme) for a given lifetime value. Evidently, the optimal lifetime value S will depend on R, and the mobility and application characteristics of the user. We identify S and R via simulation, for various application and mobility characteristics.

3 Simulations

Our simulation environment is based on the Scalable Simulation Framework (SSF) package distributed by researchers at the Dartmouth University. For generality purposes, we use a random network topology, as generated by the Waxman scheme [3]. The simulation model is depicted in Figure 1.

Using SSF abstraction, we implement a mobile user model, with various mobility and application usage characteristics. There are a total of 800 users, initially distributed uniformly across the network. We use a mobility model inspired by existing literature [2] that works as follows: A user is stationary at a certain node for some period of time that is randomly drawn from an exponential stochastic process. When this period elapses, the user moves to one of its network neighbors (movement to any neighbor is equally likely) where it again resides for some random period of time, given by the exponential stochastic process. In our simulations, we vary the rate of the exponential process, thus affecting the rate of user movement. We consider three such rates to be of interest: once every second (high), once every 10 seconds (medium) and once every 100 seconds (slow). To understand the choice of these rates, consider a geographical area populated with 802.11 wireless cells which form an 802 extended LAN. The high mobility rate corresponds to the rate at which square cells of size 30mX30m would be crossed by a car moving at about 110 km/hr. The medium mobility rate corresponds to the rate of crossing 300mX300m cells or via pedestrians moving at about 11 km/hr. Finally, the last rate corresponds for example, to pedestrians crossing 300mX300m cells.

Unicast communication takes place between users in the network, modeled using stochastic processes. The packet inter-arrival time is drawn from an exponential random process with various means.

To measure network performance, we use the following metrics: average per-switch bandwidth consumption, packet loss, and the average size of the databases used for routing at the switches. The packet latency is also measured, however it is very small for all the schemes studied, hence it is not discussed in this paper. When measuring bandwidth consumption, we include data packet, and the control overhead, weighted by 0.01 (as a control packet is roughly 100 times smaller than a data packet). The packet loss is defined as the percentage of packets that do not arrive at their intended destinations relative to the total number of packet emitted. The average database size is the average of the snapshot of the database size at all the switches in the network, taken every 10 ms. By using these metrics and simple scaling techniques, we analyze a network consisting of 100,000 users, each generating 250 Kb/s in data traffic.

3.1 Results

In Figure 2, 3 and 4 we plot respectively the average per-switch bandwidth consumption, average packet loss and average database size at the switches in the network, for the SNOOP, REG and REG_SNOOP protocols. The mean mobility rate is 1 moves/second and the mean packet inter-arrival time is 1 second. On the y-axis we show the lifetime value and on the x-axis the radius of the registration set. The y and x axis roughly correspond to using the SNOOP and REG protocols respectively.

As seen in Figure 2, the largest amount of bandwidth consumption is present at the origin, where every switch

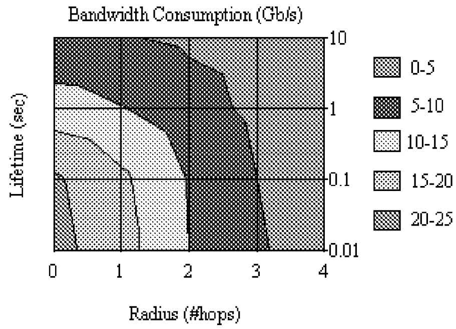

Figure 2: The average bandwidth consumption at every switch for different radii and lifetime values.

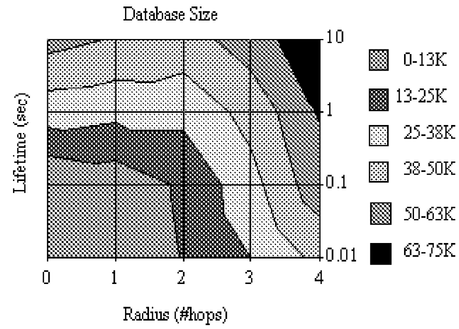

Figure 4: The average database size at every switch for different radii and lifetime values.

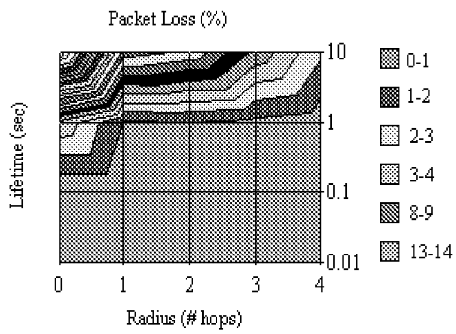

Figure 3: The average packet loss for different radii and lifetime values.

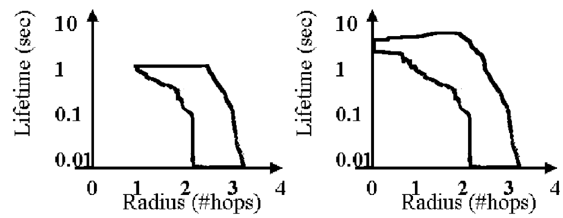

(a) The average packet inter-arrival time is 1 sec. and the average mobility rate is 1 move/sec.

(b) The average packet inter-arrival time is 1 sec. and the mobility rate is 0.01 moves/sec.

Figure 5: The set of datapoints satisfying the design constraints.

sees every packet that a user sends (every packet is basically flooded). As the lifetime value increases, up to five times less traffic is averaged at the switches. No points along the y-axis intersect the region of minimum bandwidth consumption (between 0-5 Gb/s), hence the SNOOP protocol does not attain minimal bandwidth consumption. To achieve this minimum, the REG protocol must be employed either in isolation, or in combination with the SNOOP protocol.

In Figure 3 we see that SNOOP can result in packet loss in the double digits for large lifetime values. By exploiting also the REG protocol, packet loss can be reduced for the same lifetime value.

In Figure 4 we see that as the radius increases, the average database size increases many times. Similarly, as the lifetime value increases, the average database size increases as snooped entries expire more slowly at the switches.

We can see how tradeoffs can be made in the selection of the appropriate lifetime value and radius size for the operation of the network. A large lifetime value results in less bandwidth consumption but it can lead to increased packet loss and database sizes. A large radius size results in lower bandwidth consumption at the switches and lower packet loss, but it can lead to larger databases. The optimal tuning of a network must take these tradeoffs into consideration. This is discussed at great length in the next section.

3.2 Network Tuning

We apply a number of constraints on the design of the network, and highlight the set of radius and lifetime values that satisfy these constraints. The constraints on the network design are as follows: the average database size at a switch is limited to 38,000 entries, which is realistic given the state of technology. The average per-switch bandwidth available for switching is limited to 10 Gb/s, also a realistic constraint, and the average packet loss is target at 1%, which is consistent with the needs of most applications. Given these constraints, the set of lifetime and radius values that meet the requirements is reduced to the datapoints in-

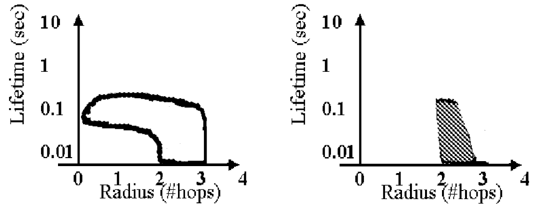

(a) The packet inter-arrival time is 10ms. and the mobility rate is 1 move/sec.

(b) The datapoints satisfying all the different application and mobility scenarios.

Figure 6: The set of datapoints satisfying the design constraints for various scenarios or combinations thereof.

side the polygons depicted in Figures 5 and 6. The average packet inter-arrival time is 1 second while the mobility rate is 1 move/sec in Figure 5a) and 0.01 moves/sec in Figure 5b); the average packet inter-arrival time is 10 msec and the average mobility rate is 1 move/sec in Figure 6a).

When the mobility rate is high, we observe that the solution does not intersect the y-axis, thus the SNOOP protocol can not be used in isolation without violating the design constraints. The use of explicit registration is required for this purpose. When the mobility rate is lower, additional datapoints can be used to satisfy the design constraints, as less packet loss is observed when the lifetime value nears 10 seconds.

Let us look at the impact of the application characteristics on these results, shown in 6a). The first observation is that the set of satisfying datapoints extends towards the y-axis, which means that the SNOOP protocol becomes a better choice for meeting our design goals. The reason for this is that, as the time spacing between packets is reduced, switches have the opportunity to listen on a more frequent basis, thus the bandwidth consumption can be reduced. The second observation is that fewer datapoints in the space where x>0.1 sec satisfy the constraints. One reason for this is that, as more snooping occurs, the average database size increases beyond the acceptable limits.

The filled polygon in Figure 6b) represents the intersection of the results achieved using the various mobility rates and packet inter-arrival times. We see that the set of satisfying datapoints is significantly reduced when all the different scenarios are taken into account.

Note that one may consider constraining the design using average lookup rates at the switches instead of the average bandwidth consumption. Let us assume that the target lookup rate is 4 million packets/sec. In this case, the worst design scenario becomes the application rate of 100 packets/sec, and the design set is given by the points inside the unfilled polygon shown in Figure 6a). Note that the control overhead may seem important here since every control packet contributes the same overhead as a data packet. However, given the difference in rates between the application and the user mobility, the control overhead is again negligible. This is not the case when the application rate is 1 packet/sec and the mobility rate is 1 move/sec, however the overall lookup rate is small for such a scenario, thus it can be easily accommodated without affecting the design.

In summary, we have seen that the SNOOP protocol, currently used in switched LANs, cannot be used efficiently to track fast moving users. Instead, the joint use of the SNOOP and REG protocols is to be desired. The selection of the radius and lifetime values must be done carefully as to take into account the design constraints, as well as the application and mobility characteristics.

4 Conclusions

With the recent explosion in wireless mobile communication, solving the problem of tracking mobile users in large switched LANs has become paramount. While LAN technology is a natural candidate for supporting mobile user tracking (the flat address space and the transparent learning protocol), it can lead to packet loss and high bandwidth consumption in a large network.

In this paper we show that the transparent learning protocol must be combined with the GARP protocol, currently used only for supporting multicast applications and VLAN formation, when tracking mobile users in a large switched network. Furthermore, we show that, when implementing these protocols, the selection of parameters can drastically affect performance, and thus must be done carefully as to take into account the design constraints, as well as the application and user mobility characteristics.

References

[1] IEEE 802.1d MAC Layer Bridging Standard.

[2] Akyildiz, I.F., Lin, Y.-B., Lai, W.-R., Chen, R.J. A New Random Walk Model for PCS Networks. IEEE JSAC , 18(7): 1254-1260, 2000.

[3] B. Waxman. Routing of Multipoint Connections. IEEE J. Select. Areas Commun., December 1988.

[4] V. Lines et al, "66MHz 2.3M Ternary Dynamic Content Addressable Memory", IEEE 2000.