

Usage-Based DHCP Lease Time Optimization

Manas Khadilkar, Nick Feamster, Matt Sanders, and Russ Clark
College of Computing, Georgia Tech

ABSTRACT

The Dynamic Host Configuration Protocol (DHCP) is used to dynamically allocate address space to hosts on a local area network. Despite its widespread usage, few studies exist on DHCP usage patterns, and even less is known about the importance of setting the lease time (the time that a client retains ownership over some IP address) to an appropriate value. Lease time can greatly affect the tradeoff between address space utilization and the number of both renewal messages and client session expirations. In this paper, using a DHCP trace for 5 weekdays from the Georgia Tech campus network, we present the largest known study of DHCP utilization. We also explore how various strategies for setting lease times can dramatically reduce the number of renewals and expirations without prohibitively increasing address space utilization.

Categories and Subject Descriptors

C.2.3 [Computer Communication Networks]: Network operations – network management

General Terms

Algorithms, Design, Management, Measurement

Keywords

DHCP, optimization, network management, usage

1. Introduction

The Dynamic Host Configuration Protocol (DHCP) [2] allows networks to automatically assign IP addresses to clients and set up various configuration parameters. Many networks use DHCP to reduce client configuration when allocating IP addresses, particularly when the network comprises many mobile or intermittently connected clients. DHCP is useful for managing IP address space allocation for networks where the total number of users outstrips the total number of concurrent users. For example, the Georgia Tech “LAWN” (Local Area Walkup and Wireless Network) [3] serves about 6,000 unique IP addresses per day, even though the maximum number of concurrent users never exceeds 2,700; this dynamism can preserve scarce address space if IP addresses are allocated and reclaimed properly.

DHCP’s effect on IP address space consumption depends on how it is configured to reclaim IP addresses. The DHCP server’s *lease*

time setting controls how often a DHCP server reclaims an allocated IP address: if a DHCP server does not receive a renew request within the lease time, it expires the client’s lease and reclaims its IP address for possible allocation to other clients. Ideally, a DHCP server could reclaim an IP address immediately after a client “leaves” the network. Unfortunately, DHCP does not mandate that clients explicitly notify the server when they leave, so servers commonly issue a single lease time for all clients¹; each client issues periodic “renew” requests to keep its lease active.

As a result, network operators must set the DHCP lease time judiciously. A lease time that is too large prevents a server from reclaiming IP addresses being used by clients that are no longer present on the network, unnecessarily wasting valuable IP addresses in a sometimes scarce address pool. On the other hand, setting a lease time that is too small can introduce substantial DHCP broadcast traffic from clients, which can overload switches and send unnecessary traffic to thousands of clients. Setting short lease times also unnecessarily expires client leases if the clients are mobile and leave for only short period of time. On many networks that use dynamic addressing, including the Georgia Tech campus network, DHCP leases are tied to authentication. Thus, a client whose lease expires must re-authenticate—a manual, inconvenient, and time-consuming process.

Today, configuring DHCP lease times is a black art: network operators typically choose a small fixed lease time and after carefully observing the network over some time period, cautiously increase the lease time if they think the network can sustain it. This paper presents three main contributions.

First, we present a measurement study of DHCP utilization on the Georgia Tech Local-Area Walkup/Wireless Network (LAWN), a dynamic campus-wide network with about 6,000 unique users per day (Section 4). We find that, for a typical day, the median client session time is about 75 minutes, and that more than 30% of clients return within 60 minutes after leaving the network.

Second, we present an emulation technique that can help operators evaluate the effects of longer DHCP lease times on address space consumption and DHCP renewals and expirations, given only an existing DHCP usage trace from an operational network (Section 5). We use this method to estimate utilization, expiry, and renewals on the LAWN (Section 6) and find that tripling the current lease time (to 90 minutes) reduces re-logins by almost 23%, while only increasing address space utilization by 14%.

Third, we explore two alternative strategies for setting lease times that further reduce the number of renewals and expirations for a given utilization of available address space: *single adjustment* and *exponential* (Section 7). In the single adjustment strategy, the DHCP server sets one lease time when a client first appears and a second lease time for all subsequent renewals from that client. The exponential strategy doubles the lease time for each renewal, up to a maximum possible lease time. We find that exponential backoff

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC '07 October 24–26, 2007, San Diego, CA.

Copyright 2007 ACM 978-1-59593-908-1/07/0010 ...\$5.00.

¹Many DHCP servers allow lease times to be configured for specific clients, but large campus and enterprise networks typically use a single setting.

can reduce the number of client expirations by 27% while increasing peak address space utilization by only 17%.

Although our measurement results reflect usage patterns that are specific to the Georgia Tech network, both (1) the emulation technique we present for evaluating the effects of different lease time settings and (2) the strategies for determining the best lease time settings for a particular network are applicable to any network that uses DHCP to dynamically allocate addresses to clients. We believe that our techniques for studying various lease times will prove useful in any setting where operators must confront the tradeoffs imposed by DHCP lease time settings.

The rest of the paper is organized as follows. Section 2 presents background on DHCP and previous studies. Section 3 describes our measurement setup. Section 4 describes usage statistics for clients on the Georgia Tech campus network; these measurements provide insight into client activity patterns and can prove useful for setting DHCP lease times. Section 5 describes our algorithm for emulating the effects of longer lease times on address space utilization. Section 6 shows our analysis of the effects of various lease time settings, and Section 7 shows the effects of using several dynamic strategies for setting lease times. Section 8 concludes.

2. DHCP Background and Related Work

We briefly summarize DHCP operation and related work.

2.1 DHCP Background

The Dynamic Host Configuration Protocol (DHCP) [2] automates host configuration on a network (*e.g.*, IP address, subnet mask, default gateway, etc.) and helps network administrators dynamically allocate IP addresses to clients from a pool of IP addresses. Because it both automates client configurations and helps operators dynamically assign addresses from a relatively small pool of IP addresses, DHCP has gained widespread usage.

DHCP has four common message types. When a client first attempts to obtain configuration, it broadcasts a *discover* message to the subnet’s broadcast address (or to the broadcast destination 255.255.255.255). A DHCP server that is listening on that subnet then replies with an *offer* message, which contains the client’s MAC address, the offered IP address, the duration of the lease, and other specific configuration information. Upon receiving an offer, a client then replies to the offer message with a *request* for one of the offers it received from a DHCP server on the subnet; a DHCP server responds to a client request with an *acknowledgment*.

A DHCP server reclaims a client’s IP address when the client’s lease *expires* (we also say that a client’s “session” expires when its lease expires). An *active* client (*i.e.*, one that is present on the network) periodically “renews” its lease by sending DHCP REQUEST messages at an interval of half the lease time, at which point the client’s lease is extended for the duration of the full lease period. For example, if a client’s lease time is 30 minutes, it will issue DHCP requests every 15 minutes. We describe in Section 3.2 how we use these renewal messages to estimate when a client is active.

2.2 Related Work

Despite its widespread use, there is relatively little understanding about how to configure DHCP. In 1995, Perkins and Luo studied how to use DHCP in conjunction with mobile IP to support auto-configuration of mobile computers [5]; we measure the behavior of DHCP users on a large network comprised primarily of mobile hosts (*e.g.*, laptops). Previous work has studied wireless usage on campus networks [4] but focused on other issues (*e.g.*, mobility), rather than DHCP.

Birk *et al.* studied DHCP client behavior on two networks: one with a dynamic pool of 420 addresses and another with a pool of about 200 addresses [1]. They found diurnal and weekly trends in active usage patterns, which is consistent with our findings. Birk *et al.* acknowledge that choosing the size of the dynamic address pool and setting the default lease time allocate to clients can greatly affect DHCP “performance”. This paper offers the first study of the effects of various DHCP client lease time settings.

3. Measurement and Analysis Setup

We describe the campus network, our DHCP measurements, and our methods for estimating client activity.

3.1 Environment and Data

We measured DHCP behavior on the Georgia Tech Local Area Walkup/Wireless Network (LAWN), a campus-wide local-area network with a peak usage of more than 2,500 concurrent users and a usable address space of about dynamically allocatable 4,000 IP addresses. The LAWN comprises 1,000 access points and 2,800 “wired” network ports (most of which are not active at any point in time) across the campus. Because the LAWN is a single virtual local area network (VLAN), all client discover and request broadcast messages are served by a single DHCP server.

We studied DHCP behavior using DHCP server logs from the LAWN DHCP server collected from February 12, 2007 to February 16, 2007. We present the results from a five-day trace because it represents a typical week on campus; other weeks exhibit similar behavior. Each entry in these logs represents a single message and indicates the type of DHCP message that was sent or received, the time of the message, and the client’s MAC address and an anonymized IP address. We used client MAC addresses to uniquely identify individual clients. The LAWN is an order of magnitude larger than other public DHCP studies [1], in terms of both the size of the address pool and the number of active clients.

3.2 Estimating Duration of Client Activity

To study both basic usage patterns and the effects of changing DHCP lease times on traffic patterns and address space utilization, we must first understand the dynamics of client activity. Because clients regularly issue DHCP messages, we can use the DHCP messages themselves to infer client activity. Clients issue “renew” messages at time intervals that are roughly half the lease time. The DHCP protocol expects the clients to renew their lease when half of the issued lease time has passed. The periodicity at which the client renews its lease is different for different lease periods; on the LAWN, the lease time is 30 minutes. We found that over 84% of the client sessions renewed their leases in the 15th minute (for a lease period of 30 minutes); more than 15% of the sessions renewed sooner than 15 minutes, and only 1% renewed after 15 minutes.

DHCP renew requests allow us to estimate when the client is active. The time window of the estimated client departure can be determined to within the expected renewal time: a client could have left the network anywhere between its last seen request and the expected renewal time after that request. We assume that clients leave within this period according to a uniform distribution and thus estimate that a client is active for a period of 7.5 minutes since its last request (*i.e.*, half of the renewal period in the case of the LAWN).

4. Usage Statistics

This section presents statistics of client activity on the LAWN. We first examine five days of DHCP logs to determine the network’s address space consumption under the current 30-minute




Figure 1: Address space utilization for five days on the Georgia Tech LAWN campus network, with 30-minute lease times.

lease time setting. We also examine the distributions of *on time* (i.e., the time that a client is active) and *off time* (i.e., the time between periods of client activity). These statistics serve not only as a useful point of comparison to previous work [1], but also help us better understand how different lease times affect client renewals and expires.

4.1 Address Space Utilization

Figure 1 shows activity for the weekdays February 12–16, 2007. The vertical dotted lines indicate midnight for each day of the week. As expected, client activity on the campus network shows diurnal and weekly trends (we note that Friday shows lower client activity, likely due to the fact that fewer classes are held on Friday). Address space utilization peaks at around 2 p.m. each day (at approximately 2,700 concurrent users for each weekday except for Friday) and is lowest at around 6 a.m. (at about 600 concurrent users). We expect that many of the 600 concurrent DHCP clients around 6 a.m. are persistently connected via either a wireless access point or by a wired port.

Although the usage patterns on the LAWN are similar to those seen in previous work [1], we observe two notable differences (other than the magnitude of the activity). First, the usage patterns shown in Figure 1 show a drop in activity after the peak at roughly 2 p.m., but this drop in utilization slows at around 9 p.m. on Monday through Thursday (though not on Friday). Birk *et al.* likely did not observe this characteristic because they observed DHCP utilization in an office building and research laboratory, where users are less likely to become inactive and later return to the same network.

Second, usage patterns are slightly more erratic. We attribute this characteristic to two underlying causes: (1) the existence of more “high frequency” activity (e.g., clients disconnecting and re-connecting as students move between classes); and (2) considerably shorter client lease times (30 minutes, as opposed to 48 hours), which preserves such high-frequency activity.

4.2 Individual Client Dynamics

In this section, we present findings about *individual client* activity. We studied two characteristics of individual client sessions:

- *On-Time* is the duration of time for which a client remains active.
- *Off-Time* is the duration of time between when a client initiates a new session and the time that the client’s previous lease expires.




Figure 2: Cumulative distribution of client on-times.




Figure 3: Cumulative distribution of client off-times.

The rest of this section discusses the motivation for studying each of these statistics, explains the difficulties in estimating individual on-time and off-time values, and presents the distribution of these statistics for the LAWN.

On-time Statistics. Computing the on-time for each client helps us determine how a particular lease time might reduce the number of lease renewals, since increasing lease times reduces the frequency at which a client must issue requests for lease renewal. Unfortunately, computing the exact on-time for a particular session is difficult because the DHCP specification does not mandate that clients issue a RELEASE when they leave the network [2], and almost 93% of clients in our trace do *not* issue a RELEASE when they leave. As a result, a client could have left the network any time between its last seen REQUEST and the time when its lease expires. In the case of the LAWN, a client could have “left” anytime between within 30 minutes after its last observed request.

We approximate on-time for each session by using a client’s REQUEST messages as an indication that the client is still active, as described in Section 3.2: specifically, we subtract the last seen REQUEST for a client from the first one seen since an expiry, and add 7.5 minutes to this value.

Algorithm 1 Generating Logs with Increased Lease Times

```
Client is expired
if Previous Message from Client is a RELEASE then
  Log a REQUEST at the current lease time
else
  Could be a live or expired client
  if Client is timed out based on original lease time then
    if Current time is greater than the next replay renewal time then
      Log a REQUEST at the current lease time
    end if
  else
    Client is live
    if The expected "live" time of the client is greater than the next
      replay renewal time then
      Log a REQUEST at the next replay renewal time
    end if
  end if
end if
```

Figure 2 shows the cumulative distribution of on-times for the client sessions for February 12, 2007, which had more than 14,000 unique sessions. About 20 percent of sessions lasted 30 minutes or less (*i.e.*, they issued two renew requests or fewer), and 59% of all sessions lasted 90 minutes or less. This distribution suggests that, subject to address utilization constraints, a significant fraction of DHCP renew traffic could be saved by increasing lease time to 90 minutes. In fact, we see in Section 6 that increasing the lease time to 90 minutes for all clients saves about 70% of the renew messages seen with a lease time of 30 minutes without prohibitively increasing peak address space utilization.

Off-time Statistics. Setting longer lease times not only reduces renewal traffic for clients that are active, but it also prevents temporarily inactive clients from being prematurely logged out. Computing the off-time for each client allows us to determine how lease time settings affect the number of times that a client that leaves for a short period of time and then returns might experience lease-time expiry, and, as a result, need to obtain a new lease (and, in the case of the LAWN, need to re-authenticate).

We approximate off-time with a method that is similar to our on-time estimation: We compute a client’s session off-time by subtracting the time that we last saw a renewal message from the time that we see another client renewal message and subtracting 30 minutes from this time interval. Figure 3 shows the distribution of off-times for the client sessions on February 12, 2007; the DHCP logs indicated about 7,000 distinct instances of DHCP off-times. About 70% of off-time instances are less than 210 minutes; in other words, when a client leaves, it is likely to return within 210 minutes about 70% of the time. Similarly, more than 30% of off-time instances are less than 60 minutes (*i.e.*, clients renew their leases within 90 minutes of their last seen REQUEST message about 30% of the time). Thus, increasing lease times to 90 minutes could significantly reduce the number of sessions where clients are forced to re-authenticate.

5. Emulating Longer Lease Times

Network operators typically want to set the DHCP lease time to the longest possible interval that does not put undue pressure on the available address space. Unfortunately, to date, operators have not had the ability to evaluate how larger lease times might affect peak address utilization without actually increasing the lease time on a running network. In this section, we describe a method for helping operators emulate the effects of longer DHCP lease times using existing DHCP usage logs, and the underlying assumptions of the analysis. We have implemented this method in an emulation tool,




Figure 4: Address space utilization for 4 days for 4 lease times.

which we used in our evaluation of various lease time optimization strategies in later sections.

The algorithm uses the original log to generate a “replay log”, which approximates the times when the client would have sent renew REQUEST messages for longer lease times. This technique generates logs which can be used to determine the effects of longer lease times on renewal traffic, premature expires, and address space utilization.

We first introduce the following terms:

- *current time*: the timestamp associated with DHCP message in the original logs.
- *replay renewal time*: the time when the client is expected to renew its lease for a given lease time; equal to the time of the last replay message seen from the client plus half of the given lease time.

Algorithm 1 summarizes how a replay log is generated from the original DHCP request logs. Of course, a client’s initial REQUEST, and any REQUEST following a RELEASE will occur at the same time for any lease time. Lease renewals are more subtle. If a client’s lease appears to have expired between two requests (*i.e.*, if the original logs show a DISCOVER before the client’s next REQUEST or if the time difference between the client’s two consecutive REQUESTS is greater than the original lease time) the algorithm first determines whether the current time is past the time at which the client was due for the next replay renew. If so, the algorithm logs a REQUEST message in the replay logs. For each renew sent by the client in the original logs, the algorithm logs a renew REQUEST whenever it determines that the client is likely to be active” past the time of when a renew would be scheduled in the case of the longer lease time.

6. Effects of Increased Lease Times

Longer lease periods improve usability (especially for intermittently connected clients) and reduce DHCP request traffic, but they also increase address space utilization. Accordingly, operators want to set the longest possible lease time that still leaves sufficient spare address space. We use the algorithm from Section 5 to determine the effects of longer lease times on address space utilization, DHCP renews, and premature session expires.

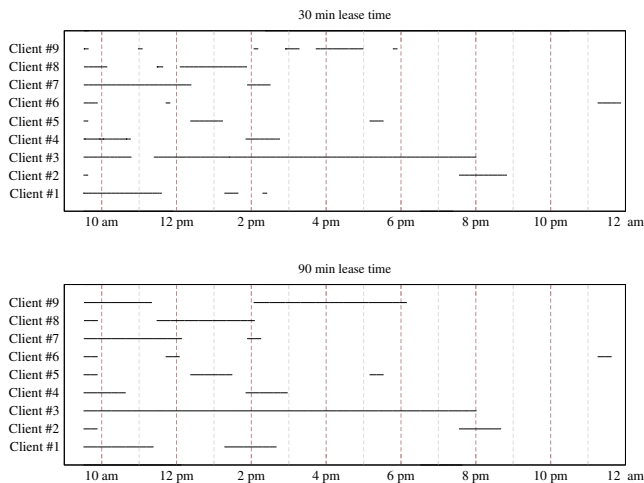


Figure 5: Sessions for 9 clients for 30 minute and 90-minute lease times.

6.1 Address Space Utilization

We analyzed address space utilization for different lease times. Operators prefer that peak address space utilization not exceed 80%, to accommodate unexpected bursts in utilization and the gradual growth of DHCP clients.²

Figure 4 shows the concurrent leases that would exist on the Georgia Tech campus network for various DHCP lease times. The existing lease time, 30 minutes, uses only 67.14% of the available address space. As expected, longer lease times increase address space utilization, particularly at peak times when many new clients enter the network. A lease time of 240 minutes exhausts the available address space (a /20) at peak utilization; lease times of 90 and 150 minutes use about 80% and 90% of the address space, respectively. A 90-minute lease period incurs a peak utilization of 81.54% over a week and an average utilization of 79.47%.

6.2 Reducing DHCP Renewals and Expires

We measured the number of client expires saved as a result of increasing the lease period from 30 minutes to 90 minutes over 5 days. Table 1 summarizes the savings in renewals and expires. We found that an average of 8,398 expires occur per day with 30-minute lease times; about 6,483 expires occur with 90-minute leases. On the Georgia Tech campus network, this implies that each day there are about 1,915 instances where the clients could avoid re-authentication if the lease period were 90 minutes, which would improve for clients and decrease load on the authentication server.

Figure 5 shows example DHCP sessions for 9 clients on February 12, 2007 to help provide intuition for why longer lease times can save re-authentications. The solid lines indicate the periods when the client was active (as determined from DHCP renew requests at 15-minute intervals). Clients 1,3,8 and 9 stand to benefit as a result of an increase in the lease period. In a network like Georgia Tech's where each expired client has to re-authenticate itself after its DHCP lease expires, six of Client 9's sessions are coalesced into just two sessions; as a result, Client 9 must re-authenticate only once, instead of five times.³

²This target utilization figure is a rule of thumb, but our results and analysis techniques would apply generally for other target utilization figures as well.

³Figure 5 also shows an artifact of our estimation where clients appear to have longer sessions with 30-minute lease times than with 90-minute lease times. Consider Client 7. In the 30-minute case, this client renews at the 15th minute and 30th minute after its start of session. From this, we estimate the client's on-time to be 37.5 minutes. This on-time will not generate any renewals in the replay log for the 90-minute case because the first expected

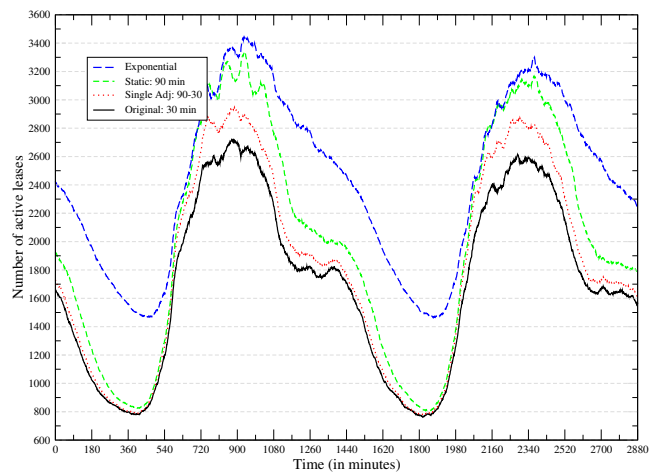


Figure 6: Address space utilization for two days for different strategies.

Strategy	Consumed	Expires Saved	Renews Saved
<i>Orig. (30 min)</i>	2,750 (67.14%)	8,398 (————)	175,088 (————)
<i>1-Adj. (90-30)</i>	2,950 (72.02%)	7,469 (11.06%)	124,291 (29.01%)
<i>Static (90 min)</i>	3,340 (81.54%)	6,483 (22.80%)	50,920 (70.92%)
<i>1-Adj. (150-30)</i>	3,350 (81.78%)	6,415 (23.61%)	112,157 (35.94%)
<i>Exponential</i>	3,450 (84.22%)	6,163 (26.62%)	38,615 (77.94%)

Table 1: Total addresses consumed (and percent utilization of the /20 address space), the number of DHCP lease expires saved, and the number of renewals saved, for various lease setting strategies.

Longer lease times reduce renewal traffic both on the network and on the DHCP server. We found that in the 90-minute case, an average of only 50,920 renewal messages are seen per day, versus 175,088 per day with 30-minute lease times; a 90-minute lease period thus saves up to 70.92% of DHCP renewal messages without significantly increasing DHCP utilization above 80%.

Although it might seem natural to set the DHCP lease time to about the time of a typical class period; however, in practice, usage patterns are highly variable. In particular, class periods are variable, wireless networks are deployed in many other parts of the campus (e.g., common areas), and the campus network must support an increasingly wider range of devices that use dynamic addressing (e.g., gaming consoles, phones). As these devices place increasing pressure on available address space, setting the appropriate lease time for each client becomes increasingly important. In the next section, we explore some strategies for setting the DHCP lease time dynamically to account for this variability in usage patterns.

7. Dynamic Lease Time Optimization

In this section, we explore two strategies for dynamically adjusting DHCP lease times for each client: *single adaptation* (whereby a client receives a different lease time after the first REQUEST) and *exponential* (whereby the lease time a client receives doubles each time it renews a lease, up to a maximum possible lease time).

7.1 Single Adaptation

Our results from Section 4 show that client on-times are concentrated within certain time ranges: specifically, Figure 2 shows that half of the client sessions are less than 75 minutes. Human behavior in the replay is at 45 minutes. As such, we estimate the client's on-time in the 90-minute replay log to be 22.5 minutes. This artifact affects our estimates of client session times but does not greatly affect our estimates for the number of renewals and premature session expires.

havior can explain this statistic: Georgia Tech has class periods of 60-minute and 90-minute durations, and the majority of the DHCP users are students with laptops. This aspect suggests that increasing the lease time to 60 or 90 minutes might reduce renewals.

Accordingly, we propose a *single adaptation* strategy, which sets a long initial lease and reduces the lease times for subsequent client renewals. This strategy reduces lease times for the common case but keeps the number of renewals and expires the same for the remaining cases. We attempted several single adaptation settings, two of which are summarized in Table 1. The best strategy we found was to set a lease period of 90 minutes for the initial REQUEST and to reduce subsequent lease times to 30 minutes. The initial lease period of 90 minutes also allows a user to have an off-time of anywhere from 0 to 90 minutes, thus preventing expires for clients that temporarily disconnect and re-attach to the network within a short time period (e.g., a single class period).

Single adaptation has two main advantages. First, it requires only minimal changes to the DHCP server, which need only distinguish DHCP Init-Reboot REQUEST messages from Renew-Reinit REQUEST messages. This differentiation is permitted by the DHCP specification and requires no additional state at the client or server [2]. Second, this scheme offers a tradeoff between increasing a lease time from a fixed lower value to a fixed upper value (30 and 90 in our case), as shown in Figure 6.

As shown in Table 1, single adaptation with an initial lease time of 90 minutes and subsequent lease times of 30 minutes increased address space consumption by only 5% (as opposed to a 14% increase with a static 90-minute lease) while still saving on 11% (929 instances) expires and 29% renewals each day. Unlike previous work [1], our DHCP clients also included was collected from all over the Georgia Tech network including areas where the expected on-times of 60 minutes and 90 minutes may not be the common case. We expect that single adaptation might perform even better if applied only to clients using access points near classrooms.

7.2 Exponential Adaptation

The *exponential adaptation* strategy issues small leases to clients when they first arrive (which serves clients with small on-times) and doubles the lease every time the client does a renew (which reduces the number of times that persistent clients need to renew their leases). Because the scheme assumes a short lease time but gracefully adapts for clients that remain active for long periods of time, we expect that it will work well for networks whose usage characteristics are unpredictable or otherwise not well known. We experimented with an exponential strategy starting off with a lease time of 30 minutes and doubling all the way up to 960 minutes. We determined this upper bound by observing that for a single day, there were only a few clients (about 1%) who had off-times of more than 930 minutes. This strategy is motivated by the observation that if a client has been active long enough to do a renew, then it could be expected to stay on longer and deserves to be given a lease period of double the previous value. This scheme does require that the servers maintain state of the clients or that the clients report their current or expected lease time in the options field of their REQUEST, which is supported, though not required, by the protocol specification.

Figure 6 shows that exponential adaptation increases address space utilization during periods of low utilization but does not increase peak address space utilization substantially over fixed 90-minute lease times (presumably because it can also quickly reclaim leases for clients that are more transient. Exponential adaptation saves the most expires (26.62%) and renewals (77.94%), with the address space utilization reaching a maximum of 84.22%. As op-

posed to static lease time settings, exponential adaptation reduces client renewals by not forcing persistent clients to frequently renew. Figure 6 shows that exponential adaptation can save renewals for these persistent clients: The minimum address space consumption never drops below 1,600 concurrent addresses, compared to about 800 more active leases than the minimum of any other scheme. With exponential adaptation, any client with an on-time of more than 465 minutes (about 8 hours) receives a lease of 16 hours, thus allowing the client to return the next day without having to re-authenticate.

8. Conclusion

Despite DHCP's widespread usage and the importance of properly configuring DHCP client lease times, today's network operators have little understanding of how to best optimize lease times given an address pool and a group of dynamic users. This paper takes a first step towards demystifying this process with three main contributions: First, we study DHCP usage patterns on a campus-wide network with a peak usage of more than 2,500 users and 1,000 access points. Second, we present a tool that uses existing DHCP traces to analyze the effects of increasing lease times on DHCP traffic, lease expiration, and address space utilization. This tool can help operators evaluate the effects of alternate lease time settings and strategies *without having to experiment with these changes on a running network*. Finally, we use this tool to evaluate both the effects of longer lease times and alternate strategies for adjusting client-specific lease times. We find that dynamic lease time adjustment strategies can significantly reduce the amount of DHCP traffic and premature client session expirations without prohibitively increasing peak address space utilization.

The techniques we have presented could apply to any DHCP network. In future work, we would like to explore how alternative lease time strategies could apply on other networks with dynamic addressing that might have much different characteristics (e.g., cable modem users, users in developing regions).

Acknowledgments

We thank David Andersen for the suggesting the exponential lease time strategy. This research was funded by NSF CAREER Award CNS-0643974 and NSF grants CNS-0626950 and CNS-0721581.

REFERENCES

- [1] V. Birk, J. Stroik, and S. Banerjee. Debugging DHCP Performance. In *Proc. Internet Measurement Conference*, Taormina, Italy, Oct. 2004.
- [2] R. Droms. *Dynamic Host Configuration Protocol*, Mar. 1997. RFC 2131.
- [3] Georgia Tech Local Area Walkup/Wireless Network (LAWN). <http://www.lawn.gatech.edu/>, 2007.
- [4] R. Hutchins and E. W. Zegura. Measurements from a Campus Wireless Network. In *IEEE International Conference on Communications*, volume 5, pages 3161–3167, 2002.
- [5] C. Perkins and K. Luo. Using DHCP With Computers that Move. *Wireless Networks*, 1(3):341–353, Sept. 1995.