STEADFAST: A New Wireless Network at MIT

Itamar Belson, Madeleine Severance, Tzer Wong
{ibelson, madksev, tzer}@mit.edu

1. Introduction

MIT is seeking a new, well-designed wireless network that enables users to efficiently and effectively connect to the Internet through access points (APs) distributed throughout campus. Offering a wired connection to the rest of the web, these APs provide clients with a dependable interface through which they can provide individualized data, optimize personal connections, and transmit and receive information.

Utilizing the existing infrastructure, our proposed system design is able to scale to support the size of MIT's campus, consisting of up to 25,000 simultaneous clients attempting to connect through the 4,000 APs distributed throughout campus. The design utilizes the three main modules of the system –the client, the APs, and the IS&T server– while prioritizing user performance, network utilization, scalability, and simplicity. Through consideration of various use cases that may arise, the system is able to support a wide variety of MIT's needs; whether a single client connecting to an under-utilized portion of the network, hundreds of clients trying to simultaneously connect in a large lecture hall, or many mobile clients traveling through a hallway, the system is able to effectively handle the circumstance while abiding by the design goals.

We designed the system with the intent of maximizing user happiness and optimizing network utilization while adhering to the requirements imposed by IS&T and the scale of MIT's campus. This proposal outlines our design by specifying our decisions for each of the modules and how these individual decisions impact our overarching system.

2. System Overview

The system can be split into three main modules: the client, the APs, and the IS&T server. These modules work together to provide users wireless Internet connection. Detailed below are several specifications and assumptions for each module that are critical to our design.

2.1 Clients

The client is the user device that wishes to connect to the Internet. It has two main components: the monitor and the controller. The monitor knows the client's performance, specified by the required throughput, which we assume does not change over time, and actual achieved throughput, which may fluctuate. The monitor also receives input from the user if he or she is unhappy. The controller is the client's main module that communicates with the APs. The controller polls the monitor every second to determine the current level of performance that dictates the client's connection requirements. We assume that clients fall into two different categories: large and small. A large client has a large required throughput and achieves close to that amount consistently. A small client has a smaller required throughput that may fluctuate. Our system must be flexible in order to deal with these two types of users with their different requirements.

2.2 Access Points

APs belong to the MIT network and are placed throughout campus in predetermined locations. A single AP may support up to 128 clients at once and has a wireless connection range of 125 feet. APs broadcast heartbeats every 30 ms so that clients may detect which APs are in range. APs broadcast their data on one of 11 predetermined channels and have a maximum bandwidth of 54 Mbits/sec. Once a client connects to an AP that can handle its required load, it is free to transmit and receive data with the network. We assume that a reliable transport protocol already exists and that data delivery is guaranteed between an AP and its connected clients. In order to effectively connect clients to the rest of the network, APs provide a wired, reliable connection to the Internet and to the IS&T server.

2.3 IS&T Server

MIT IS&T has indicated that they would like to accumulate relevant data in order to monitor and assess network performance. As such, they have provided a server with 10 TB of memory to store the amassed data. APs communicate with the server through a wired connection, while clients can communicate with the server through the wireless connection of an AP. The server already contains a table that is prepopulated with the unique MAC addresses and locations of every AP in the network.

3. System Design

The general functions and assumptions of each of the existing modules provide a foundation upon which our system is designed. Through the consideration of their qualities and interactions and the analysis of tradeoffs inherent to each existing module, we were able to design a system that ensures user satisfaction and network optimization even in the most strenuous of network conditions.

3.1 Server Data

The first component of our system is the IS&T server. While the primary purpose of this server is to act as a storage center for data requested by IS&T, its ability to reliably communicate with all APs on the network through a wired protocol enables it to play a more central role in our design schema. As such, our design utilizes the server not only as an aggregator of historical network data but also as a monitor of the current status of the network APs.

3.1.1 Server for Historical Network Data

IS&T would like the server to store, for each AP, the number of bytes transferred and the estimated number of users connected every second. The server requests this information, which the AP already knows instantaneously, and can be transferred promptly. As this data will accumulate rapidly, our system sets aside a sizeable portion of the 10 TB server solely for the purpose of this data aggregation. Once records are received by the server, they are bundled together and stored in the appropriate row and column according to the AP's MAC address and current timestamp for easy monitoring by the IS&T administrators. Although it may seem like a significant load to query every AP each second, the server's wired connection to the APs provide a reliably efficient connection with adequate bandwidth for such a procedure.

AP MAC ADDRESS	12:00:00	12:00:01
18:24:39:AC:BB:FC	(00000000, 00000000000000000000000000000000)	(00000000, 00000000000000000000000000000000)
18:25:28:AD:CL:EN	(00000010, 00000000000000000000000000110000)	(00000011, 00000000000000000000000001110000)
18:25:28:CA:EE:GA	(00001101, 00000000000000000000100100110000)	(00001101, 00000000000000000000100100111111)

Figure 1. Data Table for Historical Network Data. An example of two seconds worth of historical data for two network APs. The rows are organized according to AP's unique MAC addresses while the columns are organized and aggregated according to sequential seconds. Each record is represented as an 8 bit and 32 bit integer. The 8 bit integer represents the number of unique users while the 32 bit integer represents the number of transferred bytes by the AP since the last reset. To realize the number of bytes transferred in the one second time frame, subtract the value of the previous column from the column of interest.

3.1.2 Server for Current Network Status

In addition to the historical data accumulated by the server, the system design utilizes the server's storage space and centralized network position to collect current network status information. When a new client connects to a particular AP, it notifies the AP of its required throughput, which is passed to the server as it highlights this new connection. The server updates a data table that stores each AP's broadcast channel as well as its current required load. Our system uses the required throughput of clients as opposed to the achieved throughput as to minimize risk of user unhappiness upon successful connection to the network.

AP MAC ADDRESS	BROADCAST CHANNEL	REQUIRED THROUGHPUT
18:24:39:AC:BB:FC	1	00000000000000000000000000000000
18:25:28:AD:CL:EN	8	00000000000000000000000110000000
18:25:28:CA:EE:GA	11	00000000000000000000010100000101

Figure 2. Data Table for Current Network Status. An example of three APs with their broadcast channel and current required throughput. The rows are organized according to AP's unique MAC addresses while the columns are organized according to broadcast channel and required throughput. From this table, the server is able to identify the APs currently with the highest remaining bandwidth.

3.2 Client Connection

The second component of our system is the method by which clients are able to identify, locate, and connect to an optimal AP. The optimality of a client-AP connection depends on several factors including the distance between the client and AP, the current load on the particular AP, and the distribution of the overall load across all APs in the network. In addition to using quantifiable metrics, the method by which clients connect to the network takes into account the various client types as well as the various connection conditions that may arise.

3.2.1 Initial Contact

Upon attempting to connect to the system, a new client first establishes initial contact with the network through one of the available APs within range. In order to identify such an AP, the client listens for a heartbeat on a randomly chosen channel. If no heartbeat is discovered on the chosen channel, the client randomly picks another channel and listens for a heartbeat. This process continues until the first AP is detected, whereby the client connects to that AP, or until all channels have been exhausted and the client is notified of an unsuccessful connection. The randomness in channel selection reduces the instantaneous load from clients attempting to simultaneously connect to a single AP with a specific channel ordering. Once connected, the client will be able to communicate with other modules in the network via this preliminary connection.

3.2.2 Preliminary Connection

The preliminary connection serves as the client's first contact with the system. Through this connection, the client sends its required throughput to the server through the AP's wired connection, and using this value as a threshold, the server compiles and returns a prioritized list of potentially optimal APs for the client. This prioritization is calculated by taking into account several features such as the client's required throughput and the current load and bandwidth of the AP. The sequence of computations is summarized below:

1. The server uses the MAC address of the preliminary AP and the pre-populated table of MAC addresses to AP locations in order to identify all APs within a 125-foot radius of the preliminary AP. The 125-foot radius was chosen as no two APs in this radius can be broadcasting on the same channel.
2. The server uses the network status table described in [3.1.2] to identify the current load on each of the identified APs within range of the preliminary AP. These values are then sorted in order of decreasing available bandwidth such that the AP with the least load and most available bandwidth is first.
3. The server cuts the list where the remaining bandwidth can no longer support the required throughput of the client and removes any APs that are currently serving more than the maximum 128 clients. It then compiles the final prioritized list of available APs that can support the client's bandwidth with each AP's corresponding MAC address and broadcast channel.

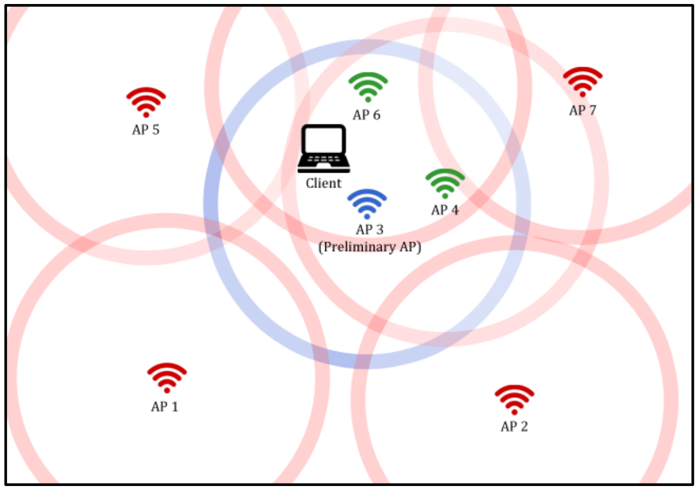

Figure 3. Diagram of Preliminary Connection. The client first identifies and connects to a preliminary AP within range (AP3). Through this AP, the server is able to identify other APs that may be within the 125-foot range of the client (AP4, AP6). These APs are sorted according to available bandwidth and returned to the client.

3.2.3 Optimized Connection

Following the above computations, the server returns the prioritized list of available APs to the client through the preliminary AP. Using this ordered list, the client disconnects from the preliminary AP and begins to listen for the heartbeat of each AP in sequence. Upon discovering an AP, the client first verifies the AP's MAC address. If it is indeed the same, then the client ensures that the AP can still support the required throughput before connecting. If the MAC address of the detected AP is different from that in the prioritized list, then the client checks if the available bandwidth of the AP is larger than that of the sought after AP, disconnecting if it is not. In the cases where no heartbeat is detected, a sought after AP can no longer support the client, or an incorrect AP is found, the client moves onto the next AP in the prioritized list. This process continues until either the client successfully connects to an AP or runs out of suggestions in the prioritized list. In the latter case, the client seeks a new, unexplored preliminary AP through the randomized process. The client then once again undergoes the sorting and prioritizing of nearby APs. If all of the nearby APs are explored without success, then the server uses the known locations of APs to recommend an acceptable AP within 500 feet of the client. This is accomplished through a similar process as described in [3.2.2], however, with a 375-foot radius.

3.2.4 Small Clients

While the described process is effective for large clients who require sizeable throughput, the method may not be as effective for small clients whose smaller throughput may fluctuate. As such, if the client self-identifies as a small client, it will forgo the prioritizing process and simply connect to the first AP that it detects with enough bandwidth to support its throughput. This is done through listening for random channel heartbeats similar to the method described in [3.2.1]. As such, in the worst-case scenario, the small client will listen on all eleven available channels before it either discovers an optimal AP or discovers that none of the APs in range are capable of supporting its required throughput.

3.2.5 Areas of High Mobility

While the process described so far is effective for areas of low mobility, it may not be as effective for regions of campus where mobility is typically high. These areas of high churn, such as hallways, may see many clients connecting to and disconnecting from APs at a rapid rate. As a result, the clients in these areas seek minimal connection time and cannot afford have the luxury of waiting for optimal AP selection. As such, our design implements APs in these designated areas with a different protocol. If a client connects to a preliminary AP in an area of high churn, the preliminary AP is promoted as the client's optimal AP. Although these APs may see times of unusually high load, the load will only ensue for a short period of time as clients will soon become out of range of the AP and reconnect to the next one in their path.

4. Analysis and Conclusion

Through focusing on the qualities of and primary interactions between the existing modules of the system, we were able to design a network schema that meets the desired specs and ensures user satisfaction and network optimization even in the most strenuous of conditions. First, IS&T is provided with accurate, up-to date information about the system's conditions so that they may effectively monitor and diagnose the network. In addition to historical data, we were able to leverage the centrality and power of the server to provide the system with a current network status that is used for effective client connection. Through the method of preliminary AP detection and connection, clients are able to rapidly connect to the system and provide it with the necessary details to ensure connection. Our system relied on the server to handle AP distribution among clients so that the network is able to make effective and intelligent decisions about client connection. With this centralized design focused on the server, our system is able to abstract complexity from the client and provide a simple interface for use. Due to the sorting algorithm and prioritized list creation, the network is able to balance the load of the network in strenuous circumstances. In connecting the clients, we focused on the predominant goal of user happiness and achieved this through treating clients of different size and mobility with different protocols. In addition, the system aims to maximize user happiness through connecting clients only to APs that are able to support their required throughput and through never disconnecting clients from APs once connected. While these two system design choices may affect overall system utilization, the happiness of the clients and the robustness of the system are of highest priority to our system. Overall, our design provides a well-designed system that is able to handle MIT's busy network effectively and efficiently.