

IQU: Practical Queue-Based User Association Management for WLANs

Amit P. Jardosh, Kimaya Mittal, Krishna N. Ramachandran
Elizabeth M. Belding, and Kevin C. Almeroth
MOMENT Lab and NMSL, Department of Computer Science, UC Santa Barbara
{amitj,kimaya,krishna,ebelding,almeroth}@cs.ucsb.edu

¹WLANs are indispensable for providing Internet access to users at locations such as universities, corporate offices, conferences, airports, and coffee shops. Many of these environments often experience flash crowds, which we define to be a sudden surge in the number of users simultaneously attempting to access the WLAN. When flash crowds occur, WLANs are likely to suffer from destructive interference, excessive channel load, and unsustainable packet processing at access points (APs). These conditions lead to a plethora of problems, such as a deterioration in network throughput, heavy packet loss, intermittent connectivity, overwhelmed APs, and sometimes, a network collapse.

To verify these claims, we present two case studies of operational WLANs that experienced the aforementioned problems. The two WLANs each consisted of over 100 APs and more than 1000 simultaneous users, deployed at recently held 62nd and 64th Internet Engineering Task Force (IETF) meetings [7]. In the first case study, a high concentration of users in adjacent rooms led to frequent packet collisions and detrimental interference. As a result, users experienced unusably low throughputs. Figure 1 shows the per-user throughput of each user during a one-second interval versus the number of instantaneous users during the same one-second interval. We observe from the figure that as the number of active users increases from 1 per second to 80 per second, the per-user throughput decreases significantly. In the second case study, users failed to establish associations with any APs due to either frequent packet collisions or excessive, unsustainable packet processing at the APs. The repeated association attempts made by users resulted in high control packet overhead, compounding the problem. The channels and the APs could not sustain such heavy workloads. The result was sparse or no connectivity for users in the network and eventual network collapse.

The connectivity and usage problems experienced by users at these events are not unique. Similar problems often occur in other scenarios, particularly those that are prone to high user concentrations, such as conferences and conventions. We predict that, as the popularity of WLANs continues to increase, these problems will become even more frequent and widespread and WLANs will have a greater need to handle flash crowds and large user concentrations.

As a result, an effective solution to manage a large number of users in a WLAN is imperative. The solution should not only avoid network breakdown, but also ensure connectivity and high user throughput. Several approaches to manage heavily loaded WLANs have been presented and evaluated in previous work. These approaches can be classified into four categories: *over-provisioning*, *selective dropping* [3], *load balancing* [2, 4, 8] and *traffic shaping* [5, 10]. Each category has its benefits and can marginally im-




Figure 1: Per-user throughput.

prove performance during a flash crowd. However, they each have drawbacks as well. Over-provisioning is expensive, inefficient and limited by bandwidth availability, while selective dropping may lead to starvation of some users. Balancing load among neighboring APs is of limited help when the total load is high enough to overwhelm all APs in the vicinity. Traffic shaping limits individual throughput in order to accommodate all users, and therefore, when the number of simultaneous users is very high, traffic shaping alone may result in unacceptably poor performance for most users.

WLANs that need to support a large number of users are thus in critical need of a practical and effective system to handle heavy loads and flash crowds. In this paper, we propose *IQU*, a practical queue-based user association management system for heavily loaded WLANs. The premise of user association management is to control the frequency and duration of user associations with the network when the number of users trying to access the network is greater than what the network can support. *IQU* maintains a queue of users requesting network access. Only as many users as can be simultaneously accommodated are granted access to the network. Any remaining users wait for admission in a queue. Admitted users are assigned periods of access, called *work-periods*, within which they can execute any network-related tasks. If the network is underloaded, the user queue will be empty and users can continue to access the network even after their work-period expires. In a heavy load situation, the expiration of the work-period causes the user to be disassociated from the network and placed back into the queue. A different user from the head of the queue is then admitted into the network. Users with network access are updated with their remaining work period so that they can plan their network-related tasks accordingly². Similarly, users waiting in the queue are given esti-

¹A complete version of this research has been recently accepted as a conference publication [6].

²Applications designed for disconnected operation [9] can also leverage this information.

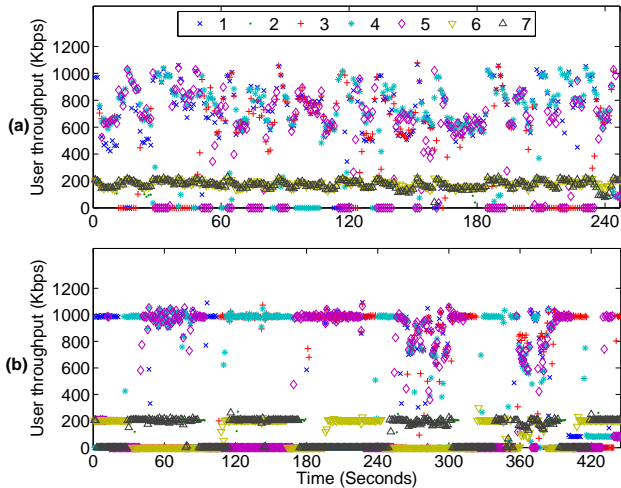


Figure 2: (a) Client throughput without IQU. (b) Client throughput with IQU.

mates of their wait time for network access and the duration of the work period they will be granted. Thus, unlike the scenarios presented in our case studies, there is no uncertainty about the availability of network access. This information prevents users from making repeated unsuccessful association attempts, thereby both reducing network overhead and considerably improving user experience.

IQU is a simple and powerful system for managing heavily loaded WLANs. However, IQU changes the basic access model to which today's WLAN users are accustomed. In heavily-loaded WLANs, IQU requires users to wait in a queue for access. Moreover, when access is granted, users must complete their network-related tasks within an allotted work period. This is a significant change from the current model of obtaining immediate access for unlimited durations of time. However, we believe that this change is inevitable in order to maintain good performance in a heavily-loaded WLAN. Moreover, we believe that the new model is intuitive and easy to understand. Users can be made aware of their assigned wait periods and work periods via a networking utility on the user's device. Note that the new access model may bring about unprecedented alterations in typical user behavior; for instance, users may generate traffic more quickly so that they can complete their tasks in the assigned work period.

To evaluate our system, we built a prototype of IQU and tested it on a testbed consisting of 8 Linux laptops (three IBM Thinkpads and five Toshiba Satellite laptops) equipped with Atheros chipset IEEE 802.11a/b/g wireless network cards. One laptop is configured to act as a wireless AP, while the remaining act as WLAN clients. The goal of this testbed is to convincingly demonstrate the practicality of an association management system such as IQU, as well as IQU's benefits in a real system. The AP and client laptops are placed within direct transmission range of each other. The wireless network cards are managed by the MADWiFi driver, which is a Linux kernel device driver module for Atheros-based WLAN devices [1]. We implement the IQU prototype by appropriately modifying this driver. For our experiments, we configure the wireless cards to use the IEEE 802.11b protocol and fix the data rate at 11 Mbps. We disable the RTS/CTS collision avoidance mechanism and MAC layer retransmissions. The impact of various IQU parameters on system performance is explored, and appropriate values for the parameters are identified.

We use our testbed to emulate a flash crowd and heavy load conditions. While three out of the seven clients maintain a UDP flow level of 200 Kbps in both directions combined, the remaining four clients simultaneously initiate UDP flow levels of 1 Mbps at the beginning of each experiment. Each UDP session lasts for four minutes. This traffic configuration emulates a flash crowd and heavy load conditions as observed in our case studies. Note that the four-minute durations do not include the time spent waiting in the user queue.

Figures 2(a) and 2(b) show the throughput achieved by each of the 7 clients without IQU enabled and with IQU enabled, respectively. In Figure 2(a) we observe significant losses and variations in the average individual throughput of all seven clients. This result clearly demonstrates the detrimental effect of a flash crowd and heavy load on network performance when IQU is not enabled. On the other hand, in Figure 2(b), we observe that IQU successfully controls the number of associated users such that, when admitted, each associated user's throughput close to the offered load.

Although IQU improves throughput during a flash crowd or high network load, the disadvantage is that it takes longer to service the clients in the network. This can be observed from the x-axis limits in Figures 2(a) and 2(b). The extent of increase in service time depends on the choice of parameter values and network traffic conditions. We argue that longer service times are an acceptable tradeoff for network administrators and users to avoid grossly unacceptable network performance or network collapse.

Due to its elegance and effectiveness, user association management has far-ranging implications as a tool for managing limited resources in sophisticated WLANs. It coalesces the benefits of over-provisioning, selective dropping, load balancing and traffic shaping, while avoiding their drawbacks. We believe that our work creates new directions for further research in this area. Different strategies can be explored for managing the user queue. Although we use a simple FIFO queue in this paper, priority-based queues may also be used to support different network access policies. Determination of the optimal number of users that may be permitted to simultaneously access the network and accurate estimation of user wait periods are other parts of this system that have potential for further exploration and research.

REFERENCES

- [1] Multiband Atheros Driver for Wireless Fidelity (MADWiFi). <http://madwifi.org>.
- [2] A. Balachandran, P. Bahl, and G. Voelker. Hot-Spot Congestion Relief in Public Area Wireless Networks. In *IEEE WMCSA*, Monterey, CA, Oct 2002.
- [3] A. Barbaresi, S. Barberis, and P. Gorla. Admission Control Policy for WLAN Systems based on the Capacity Region. In *IST Mobile Summit*, Dresden, Germany, Jun 2005.
- [4] Y. Bejerano, S. Han, and L. Li. Fairness and Load Balancing in Wireless LANs Using Association Control. In *ACM Mobicom*, Philadelphia, PA, Sep 2004.
- [5] C. Chiasserini and R. Rao. Performance of IEEE 802.11 WLANs in a Bluetooth Environment. In *IEEE WCNC*, Chicago, IL, Sep 2000.
- [6] A. P. Jardosh, K. Mittal, K. N. Ramachandran, E. M. Belding and K. C. Almeroth. IQU: Practical Queue-Based User Association Management for WLANs. To appear in *ACM Mobicom*, Los Angeles, CA, Sep 2006.
- [7] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding Congestion in IEEE 802.11b Wireless Networks. In *USENIX IMC*, Berkeley, CA, Oct 2005.
- [8] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh. A Client-driven Approach for Channel Management in Wireless LANs. In *IEEE Infocom*, Barcelona, Spain, Apr 2006.
- [9] L. Mummert, M. Ebling, and M. Satyanarayanan. Exploiting Weak Connectivity for Mobile File Access. In *ACM SOSP*, Copper Mountain, CO, Dec 1995.
- [10] M. Portoles, Z. Zhong, and S. Choi. IEEE 802.11 Downlink Traffic Shaping Scheme for Multi-User Service Enhancement. In *IEEE PIMRC*, Beijing, China, Sep 2003.