Price-based Channel Time Allocation in Wireless LANs

Samarth H. Shah and Klara Nahrstedt *
Department of Computer Science
University of Illinois at Urbana-Champaign
Email: {shshah, klara}@cs.uiuc.edu

Abstract

In access-point based wireless networks employing the IEEE 802.11 protocol in DCF mode, without bandwidth management and rate control, users pump data into the network as fast as it is generated. This results in a loss of QoS for the user and performance degradation of the network. There is a need for bandwidth allocation and for users to co-operatively control their rates based on the allocated bandwidth. In this paper, we use price as a signal for bandwidth allocation in wireless hot spot networks. We allocate more bandwidth to users that pay more. At the same time, we also aim to maximize channel utilization and increase mean satisfaction across all the users.

1. Introduction and Motivation

Recent years have witnessed the emergence of the phenomenon of near-ubiquitous Internet access through wireless "hot spot" networks. A *hot spot* is a hotel, airport, cafe or restaurant providing access to the Internet via an access-point (AP) based wireless local-area network (WLAN) that usually employs the IEEE 802.11 protocol in DCF mode. One major challenge in hot spot networks is the distribution of the shared resource of *bandwidth* among the numerous users. Users can have different requirements, based on the applications they are running. Different users can experience different channel qualities. Channel quality can also vary with time for the same user. In addition, different users may also be willing to pay different sums of money for the same quantity of bandwidth. This is because the same amount of bandwidth is worth more to some users (e.g. a stock-broker who needs to make a buying/selling decision based on a stock-price) and less to others (e.g. a person who is just idly surfing). In the absence of a bandwidth arbitration mechanism, each user will pump data into the network as fast as it is generated and if the sum of the rates exceeds the channel capacity, then there will be severe performance degradation. Bandwidth management and co-ordinated rate control enhances the QoS of the individual applications, as shown in [12]. In this paper, we address this problem of how to distribute bandwidth in a wireless LAN between users with different requirements, different channel qualities, and willing to pay different sums of money. We attempt to maximize channel utilization as well as system revenue, while keeping user satisfaction also in mind.

Each user, once allocated a certain portion of the total network bandwidth, must co-operate with the network provider and control his/her rate to conform to the bandwidth fraction allocated. While we do not address this problem of co-operative and collective rate-control, it is an active area of research at various layers of the OSI protocol stack. Fair scheduling [14] addresses this problem of rate control at the MAC layer. We address the problem from the point of view of the *policy* in allocating a portion of the overall network bandwidth to each user. The *mechanism* of how each user conforms to its allotted share falls under the domain of fair scheduling at the MAC layer and other rate control schemes.

The rest of this paper is organized as follows. Section 2 formalizes the problem. Section 3 discusses the concept of channel time proportion. In Section 4, we present the overall architecture of our scheme and our channel time allocation algorithm. Section 5 contains a summary of our experimental setup and results. Section 6 discusses the related work in the area and Section 7 concludes the paper.

2. Network Model

We assume a single IEEE 802.11 subnet with a single AP and a set S of users with wireless hosts. The users only communicate with the AP, not directly with each other. In practice, IEEE 802.11 DCF is used even for such a communica-

* This work was supported by the DoD Multi-disciplinary University Research Initiative (MURI) program administered by the Office of Naval Research under Grant NAVY CU 37515-6281, and the NSF EIA 99-72884 grant. Any opinions, findings and conclusions are those of the authors and do not necessarily reflect the views of the above agencies.

tion model, not just for peer-to-peer wireless transmissions. In [12], bandwidth management over peer-to-peer wireless LANs is discussed, using the same system architecture described in this paper.

Each user $f \in S$ has a maximum bandwidth requirement $b(f, max)$ and a minimum bandwidth requirement $b(f, min)$, both in bits per second. The maximum and minimum bandwidth requirements include both uplink and downlink bandwidths. They can be estimated by the user based on the applications he/she intends to run and can be changed at any time during the user's activity. For best-effort flows, $b(f, min) = 0$. Such a utility function, defined by the minimum and maximum bandwidth requirements has also been used previously in literature [12, 10].

One of the major contributions of our scheme is that we convert bandwidth requirements into *channel time proportion (CTP)* requirements. The channel time proportion is the fraction of unit time a user uses the channel for sending and receiving his/her data. This is directly related to the bandwidth because, larger this fraction, the more data in unit time a user can send/receive, and thus the larger his/her throughput. The minimum bandwidth requirement of user f, $b(f, min)$ is converted into a minimum CTP requirement $c(f, min)$ and the maximum bandwidth requirement $b(f, max)$ is converted into a maximum CTP requirement $c(f, max)$. $0\% \leq c(f, min) \leq c(f, max) \leq 100\%$.

Each user f also has a maximum price $mp(f)$ in cents that he/she is willing to pay per minute for 1% of the channel time [1]. $B(f) = mp(f) \cdot c(f, max)$ is thus the *maximum bid* (in cents per minute) that the user f is willing to make for the service it is requesting. The values of $b(f, min)$, $b(f, max)$ and $B(f)$ are valid for the entire session of user f, unless changed by the user in between.

A centralized *Bandwidth Manager (BM)* situated at the AP takes as input the CTP requirements of all the users as well as their maximum bids. It outputs current price p (in cents per minute for 1% of the channel time) and the CTP $c(f, a)$ allocated to each user $f \in S$, which are calculated using the algorithm presented in Section 4. $c(f, min) \leq c(f, a) \leq c(f, max)$. The revenue from a particular user f, in cents per second, is $p \cdot c(f, a)$. If this is smaller than user f's bid $B(f)$, then the remainder is refunded to the user. The refund is very important because it encourages the user to bid high. If the bid is excessive, no harm is done, the remainder is refunded. The instantaneous revenue of the network provider, in cents/minute, is $R = p \cdot \sum_{f \in S} c(f, a)$. The total revenue is R aggregated over the entire time of operation of the network. If the set S changes (users arrive or leave), or if a user changes his/her parameters, then the price p and the CTPs of all the users f have to be recomputed, and the value of R in cents/minute also changes.

1 We have arbitrarily picked minutes and cents, respectively, as the units of time and money in this paper. Similarly, the minimum resolution of CTP used in this paper (1%) has been arbitrarily chosen. Network providers can set these parameters to suit their network.

3. Channel Time Proportion (CTP)

In the previous section, we mentioned that we converted bandwidth requirements of the users into channel time proportion (CTP) requirements. The reason for this was that each user's communication with the AP is affected by different and varying levels of medium contention, fading and interference. In [3, 12], we proposed a simple method to estimate maximum throughput over a wireless link, that considered these time- and location-dependent effects. Using this method, we obtain a different maximum throughput estimate $b_e(f)$ for the links between *each* user f and the AP, since these effects are different for different links. The estimation is performed for each user-AP link in the 802.11 device driver at the AP. The BM program at the AP periodically probes its MAC layer API for the $b_e(f)$ estimate for the link to each user f.

We use the maximum throughput estimate $b_e(f)$ on user f's link with the AP to convert user f's bandwidth requirements to its CTP requirements. Assume k bits per second can be transmitted between user f and AP over a wireless channel, given a certain level of contention and physical errors on the user-AP link. Assume the user requires a minimum throughput of l bits per second, then in effect the user requires a fraction $\frac{l}{k}$ of unit time on the channel. The CTP requirements of user f, $c(f, min)$ and $c(f, max)$, can thus be obtained by simply dividing its respective bandwidth requirements $b(f, min)$ and $b(f, max)$ by $b_e(f)$. $c(f, min) = \frac{b(f,min)}{b_e(f)}$ and $c(f, max) = \frac{b(f,max)}{b_e(f)}$. The value of $b_e(f)$ reflects the quality of the channel between user f and the AP. The concept of CTP brings users with different channel qualities on the same level plane for the channel time allocation algorithm.

4. System Architecture and Channel Time Allocation Algorithm

The components of the system and its overall architecture are very similar to that of the bandwidth management scheme described in [12]. The price-based channel time allocation algorithm, however, is completely different from the "fair" allocation algorithm in [12].

4.1. System Architecture

The overall architecture of the system is shown in Figure 1. The User Agent (UA), at the application layer, takes from the user f the values of $b(f, min)$ and $b(f, max)$. It obtains from the Bandwidth Manager (BM) at the AP the value of $b_e(f)$ for its link with the AP. The UA then computes the CTP requirements $c(f, min)$ and $c(f, max)$, as described in the previous section, and sends them, along with user f's maximum bid $B(f)$ for these CTP requirements, to the BM at the AP. The BM computes, using the algorithm in the next subsection, the system price p and CTP allocated to each user. It then conveys these values to all UAs of all the users through a subnet broadcast. The UA of each user passes the CTP allocated to it to the scheduler at the MAC level so that it can use it as the "weight" of this user in the fair scheduling algorithm. A performance evaluation of this basic scheme via ns-2 simulations as well as in a real 802.11 network testbed, coupled with a "fair" channel time allocation algorithm, is available in [12].

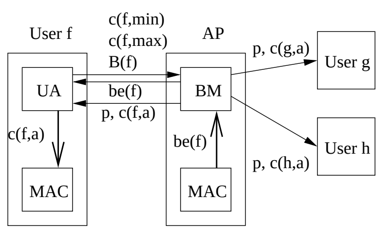

Figure 1. System architecture.

If at any stage, the user's requirements or bid change or the value of $b_e(f)$ varies significantly, the UA has to re-negotiate with the BM with these new parameters. The BM informs the UA of user f the $b_e(f)$ value it periodically retrieves from the AP's MAC layer API, and the UA can determine if it has changed significantly since the last $b_e(f)$ update. If it has, then new CTP requirements are computed using the new $b_e(f)$ value and these are used in re-negotiation with the BM. There is incentive for the user to re-negotiate even when requesting a smaller CTP, because the smaller the CTP requested, the lower the price he/she has to pay. If a user is within range of multiple APs, then his/her UA can just pick whichever one gives the user the best deal, in terms of CTP allocated and price charged. At any time, a user can query a BM for system information such as the current price p, reserve price $reservep$ (defined later), or channel utilization $\sum_{f \in S} c(f, a)$, and use this information in his/her bidding decisions. The BM may also provide a *new* user all the parameters of all the existing users, without actually naming them, so that the UA of the new user can *locally* execute the BM's channel time allocation algorithm, and obtain a resultant price and CTP allocation. Depending on this result, the new user may decide to change some of its parameters, or may decide to not bid at all, at this time.

4.2. Channel Time Allocation Algorithm

The pseudo-code of the channel time allocation algorithm is shown in Figure 2. This algorithm runs at the BM and is responsible for determining the current channel price p and the individual users' CTP allocation.

The network provider sets a fixed *reserve* price ($reservep$) which is derived from the cost of maintenance of the network. The price per second of 1% of CTP is never lower than the reserve price, i.e. $p \leq reservep$. In the trivial case where $\sum_{f \in S} c(f, max) \leq 100\%$, the price is just set to $max\{reservep, \min_{f \in S}\{mp(f)\}\}$. A reserve price ensures that users do not get away with setting their $mp(f)$ very small, and thus yielding very low revenue to the system. By setting the price to $\min_{f \in S}\{mp(f)\}$, when $\min_{f \in S}\{mp(f)\} > reservep$, our algorithm finds the largest price for which user satisfaction of all users is 100% and channel utilization is maximum. (See [13] for proof.) Alternatively, in the trivial case, the network provider may opt to just set p to $reservep$, irrespective of the $mp(f)$ values.

In the non-trivial case $\sum_{f \in S} c(f, max) > 100\%$, the BM allocates the CTP so that revenue is maximized while keeping channel utilization at 100%. The users $f \in S$ are sorted in non-decreasing order of $mp(f)$. Let V be this sorted set. In the algorithm, set V represents users whose maximum CTP requirements are satisfied and their excess bid $B(f) - p \cdot c(f, max)$ cents per minute refunded. Set W represents users whose entire bid is consumed but their maximum CTP requirements can still not be fulfilled. Repeatedly, the user with the lowest value of $mp(f)$ is removed from V and added to set W until $\sum_{f \in V} c(f, max) < 100\%$. The price is now set to $p = \frac{\sum_{g \in W} B(g)}{100\% - \sum_{f \in V} c(f, max)}$. If $p \leq \min_{f \in V}\{mp(f)\}$, then the algorithm terminates here, else the current split of users between sets V and W is invalid. So, once again the user with the minimum $mp(f)$ is moved from V to W and the price p is recomputed for the new values of set V and set W. The procedure continues until a value of p is found that is $\leq \min_{f \in V}\{mp(f)\}$ (i.e., a valid split of users between sets V and W) or until set V is empty. The last value of p computed is the price for the set of users S. The portion of the algorithm described so far runs in $O(\|S\|)$ time.

The CTP allocated to each user $f \in S$ is simply $c(f, a) = min\{c(f, max), B(f)/p\}$. If $c(f, a) < c(f, min)$ then user f is deleted from S. He/she is *blocked* because he/she is not paying enough to even have his/her minimum CTP requirement be sat-

CTALLOC($V \leftarrow S$ sorted in non-decreasing order of $mp(f)$)

repeat
    $redo \leftarrow$ **false**
    if $\sum_{f \in V} c(f, max) \leq 100\%$
      then $p \leftarrow max\{reservep, \min_{f \in V}\{mp(f)\}\}$
          for each f **in** V
          do $c(f, a) \leftarrow min\{c(f, max), B(f)/p\}$
      else $W \leftarrow$Move users from top of V until $\sum_{f \in V} c(f, max) < 100\%$ or $V = \{\}$
          while $\left[p \leftarrow max\{reservep, \frac{\sum_{g \in W} B(g)}{100\% - \sum_{f \in V} c(f, max)}\}\right] > \min_{f \in V}\{mp(f)\}$
          do if $V = \{\}$
              then break
              else $W \leftarrow W +$First user $\in V$
          for each f **in** V
          do $c(f, a) \leftarrow c(f, max)$
          for each g **in** W
          do $c(g, a) \leftarrow B(g)/p$
    if $\exists f \in V \mid c(f, a) < c(f, min)$
      then Delete f from V
          $redo \leftarrow$ **true**
  until $redo =$ **false**

Figure 2. Channel time allocation algorithm.

isfied. The algorithm must run all over again for the new value of set S. The worst-case overall running time of the algorithm is thus $O(\|S\|^2)$.

Basically, in the non-trivial case when $\sum_{f \in S} c(f, max) > 100\%$, the algorithm maximizes revenue and channel utilization simultaneously. As mentioned in the next subsection, if the users were to enter into an auction, the price they will naturally settle upon at equilibrium is the same as the price our algorithm computes, when $\sum_{f \in S} c(f, max) > 100\%$. (See [13] for proof.) This price reflects the true worth of channel time among the users in the network.

When $\sum_{f \in S} c(f, max) > 100\%$, the CTP allocated to each user in set W is limited by its maximum bid, and the CTP allocated to each user in set V is limited by its maximum requirement. In other words, users in set V may have paid more than they require to satisfy their maximum CTP requirement and could hence receive a refund. Conversely, the users in set W have paid less than they require to satisfy their maximum CTP requirement and will have to settle for less CTP, while their entire bid is consumed by the system [2]. Incidentally, the algorithm ensures that users with the same value of $mp(f)$ are, ultimately, either placed all together in V or all together in W. Thus, when initially sorting the users in non-decreasing order of $mp(f)$ in V, ties can be broken arbitrarily.

At any instant, the price for 1% CTP is the same for all users [3]. As users arrive and leave, the price varies depending on the competing demand. The user is billed the instantaneous price times his/her instantaneous allotment, aggregated over his/her entire session, when he/she leaves the network. We adopt an "acceptable or nothing" QoS philosophy. A user gets acceptable QoS (between minimum and maximum CTP requirements) or no service at all. When demand increases and CTP becomes dearer, users willing to pay too little for their minimum requirement are blocked. Best-effort users are always admitted due to their zero minimum CTP requirement. Users paying too little can either increase their maximum bid or decrease their minimum requirement and rejoin the network, after a time delay T. The algorithm works in such a way that, if $c(f, min) = 0$ and $c(f, max) = 100\%, \forall f \in S$, then the users will be allocated CTP in the ratio of their $mp(f)$ values. The price in this case will be $\sum_{f \in S} mp(f)$ (cents/minute)/%CTP.

2 It is not logical for a user with a larger $mp(f)$ to have its maximum requirement satisfied while one with a smaller $mp(f)$ value is not fully satisfied. This is the reason why users are moved from V to W only in non-decreasing order of $mp(f)$.

3 The price for 1 bps throughput is different for users with different channel qualities. This is also the case in [8].

4.3. Centralized Auction of Channel Time

For the case $\sum_{f \in S} c(f, max) > 100\%$, the algorithm described in the previous subsection is actually a centralized version of a distributed *auction* of channel time. An auction is a natural way of resource allotment when bidders value the resource differently. This is the case with wireless channel time in a hot spot network, as mentioned in Section 1.

Assume an ascending, multi-unit auction [6] of channel time between the users. Assume each user is allotted a CTP *proportional* to its bid, i.e. CTP allotted at any instant to a user is equal to the user's bid divided by the sum of all bids at that instant. Now, in the auction, users increase their bids in order to obtain more and more CTP. This bidding continues until all users either obtain their maximum CTP requirement $c(f, max)$ or reach the maximum they are willing to bid $B(f)$, i.e., the maximum sum of money they can afford. At this point, the system reaches equilibrium until a new user arrives, an existing user leaves, or some other parameter, such as a user's maximum bid, changes. The equilibrium price p is the true worth of CTP among the users. An example of such an auction can be found in [13].

Obviously, such a distributed auction is infeasible in our scenario. The repeated bids from every user, each resulting in a different CTP allotment for all users, constitute a very large overhead. The delay in attaining equilibrium can also be untenably long, and during this time, the users' CTP allotment will be continuously varying. We *centralize* this auction by having all the users provide the BM their limiting parameters: $c(f, max)$ and $B(f)$. The channel time allocation algorithm described earlier then *simulates* the ascending auction. It determines the same equilibrium price and CTP allotment as the distributed auction (see [13] for proof) in $O(\|S\|^2)$ time, while the distributed auction would have required several iterations of bidding from each user.

One possible problem with an auction-based channel time allocation policy is that of *collusion.* If bidders cooperate with each other and decide to make small bids, then the system's revenue can be adversely affected. As is well known in auction literature [9, 4], having a reserve price mitigates this problem to an extent. Furthermore, in a hot spot network at a public place such as an airport or cafe, it is impractical for *all* the users of the network to meet and agree to collude. The collusion fails if there is even a single user from outside the group of colluders present in the network. Thus, if all the colluders decide to have small maximum bids ($B(f)$) and only one user from outside the collusion group makes a normal maximum bid, that user will get most of the channel time and the colluders will all be adversely affected. The promise of more channel time for a larger bid, as well as the assurance that any extra money will be refunded, should encourage users to set their maximum bids $B(f)$ high, and thus increase system revenue. The refund helps to avoid the "winner's curse" effect [2].

5. Results

The channel time allocation algorithm described in the previous section is a *variable price* algorithm because the price changes according to the "richness" of the users in set S. In this section, we evaluate the performance of this algorithm and also compare it with two *fixed price* schemes:

1. In the *fixed price proportional* (FPP) scheme, each user f is initially allotted CTP $c_r(f, a) = B(f)/p_r$ where p_r is the fixed price. If $\sum_{f \in S} c_r(f, a) > 100\%$, then the allocations $c_r(f, a)$ are scaled down proportionately so that their sum equals 100%. If, for any user, $c_r(f, a) > c(f, max)$, then $c_r(f, a) = c(f, max)$. If, for any user, $c_r(f, a) < c(f, min)$, then this user is blocked on the grounds of insufficient budget.
2. The *fixed price greedy* (FPG) scheme greedily allocates CTP, using price p_g, to the users with the lowest maximum requirements. Users with larger maximum requirements get blocked, once 100% of the CTP has been allocated. Users who have too small a budget to even satisfy their minimum requirements, at the fixed price p_g, are also blocked.

Our simulation scenario consisted of 100 users with random arrival and departure times in a 5-hour time window. The minimum CTP requirements of the users ($c(f, min)$) were uniformly distributed in the range 0% to 2%, corresponding to 0 to 40kbps for a 2Mbps IEEE 802.11 network. The maximum CTP requirements were uniformly distributed in the range 2% to 10%, corresponding to 40kbps to 200kbps for a 2Mbps channel. The value of $mp(f)$, in (cents/minute)/%CTP, for each user f was randomly chosen from the set $\{0.1, 0.2, ..., 1.0\}$. The reserve price for the network was set at 0.1 (cents/minute)/%CTP.

Figure 3(b) is a plot of the number of users requesting service and the number of users admitted, using our variable price scheme. The remaining users are blocked. For the sake of simplicity, we assume that blocked users do not return and request service later. (Alternatively, we can assume that a user arriving later is actually a returning user, with different parameters.) Figure 3(a) shows the variation in price as users arrive and leave. The average price for the entire simulation run is 0.75 (cents/minute)/%CTP. Figure 3(c) illustrates the mean satisfaction of the admitted users. We measure satisfaction, as a percentage, for a user f as: $\frac{c(f,a)}{c(f,max)} \cdot 100$. Figure 3(d) is a plot of the channel utilization, in percentage, over the course of the simulation. As more users are admitted, channel utilization increases but the individual users are allocated less CTP, so their satisfaction falls. The mean user satisfaction and channel uti-

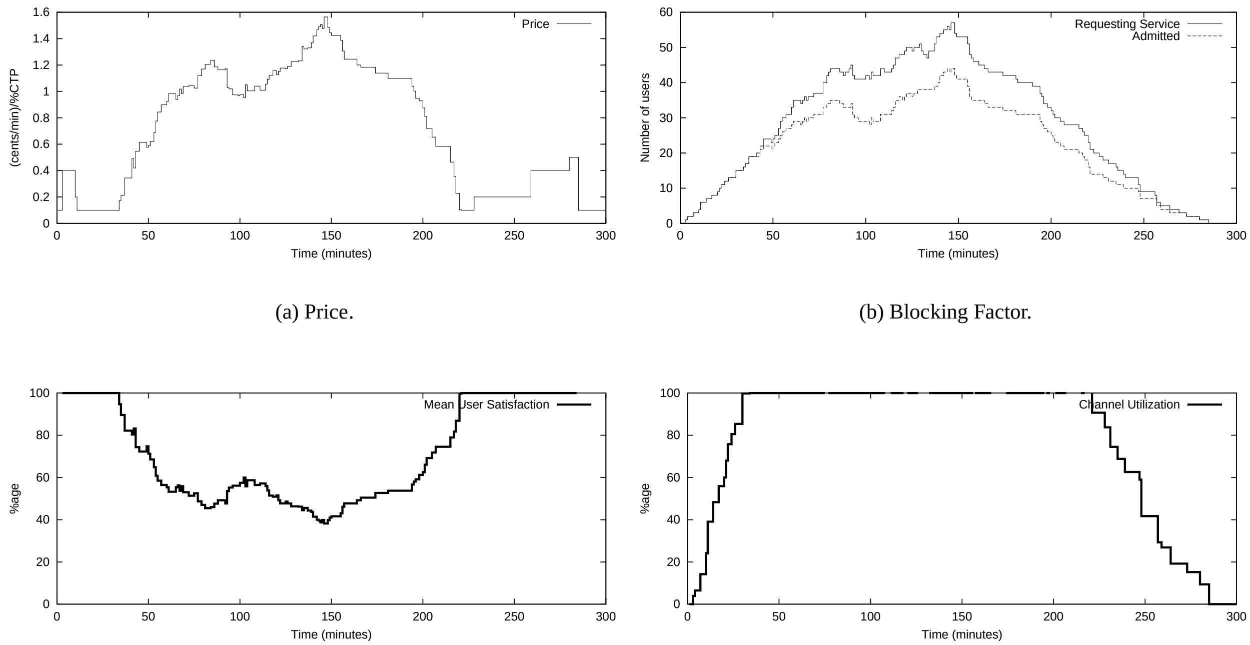

(a) Price.

(b) Blocking Factor.

(c) Mean User Satisfaction.

(d) Channel Utilization.

Figure 3. Performance evaluation of our variable price channel time allocation algorithm.

lization, averaged over the entire 5-hour simulation run, are shown under the "Variable Price" row of Table 1. The "Revenue" column of Table 1 shows the total revenue earned at the end of the 5-hour simulation run.

Table 1 presents the results of the comparison between our variable price channel time allocation algorithm and the fixed-price algorithms. For lower fixed prices, the revenue for the system is lower, but the mean user satisfaction and channel utilization are higher. This is because at lower prices, the users' budgets can buy them large CTPs. At higher prices, the revenue is higher, but mean user satisfaction, channel utilization and the blocking factor are worse. Our variable price scheme attempts to simultaneously optimize all the performance parameters. While there may be prices for which the fixed price schemes slightly outperform our scheme under one of the performance metrics, the corresponding penalty paid by the fixed price scheme in terms of the other metrics is large. For example, the FPP scheme, with a price of 1.5 (cents/minute)/%CTP yields a 10% higher revenue than our scheme, but the mean user satisfaction is nearly halved, channel utilization falls from 83% to 51%, and blocking factor is also worse.

It should be noted that, even at the low fixed price of 0.2 (cents/minute)/%CTP, the fixed price schemes do not necessarily perform better than our variable price algorithm in terms of channel utilization and mean user satisfaction. This is because our reserve price is lower than 0.2 (cents/minute)/%CTP. Thus, when our scheme defaults to the reserve price, it results in much better utilization and mean user satisfaction at those instants, which affects the overall mean utilization and mean user satisfaction comparisons.

Also note that the fixed price schemes cannot know the average worth of the channel time to the users before hand. But even if this knowledge *were* available, and the fixed price were set to this average worth (i.e., the average price of the variable price scheme 0.75 (cents/minute)/%CTP), the performance still does not match the performance of the variable price scheme.

6. Related Work

In this section, we survey the literature in the field of price-based bandwidth allocation in different types of wireless networks. We point out the differences from our work, and highlight our contributions.

An early work on pricing and bandwidth adaptation in wireless networks was the TIMELY project [1]. Users benefitting from adaptation were charged and those suffering from adaptation were compensated. However, how the exact charges and credits were calculated, was not specified. In [8], the authors discuss price-based resource allocation

Algorithm	Price ((c/min)/%CTP)	Revenue (c)	Avg. Satisfaction	Channel Utilization	Users blocked
Variable Price	Variable	19617	71%	83%	24
FPP	Fixed @ 0.2	4496	68%	80%	24
	Fixed @ 0.75	15718	59%	74%	24
	Fixed @ 1.5	21410	38%	51%	26
FPG	Fixed @ 0.2	3750	94%	69%	50
	Fixed @ 0.75	13245	66%	64%	40
	Fixed @ 1.5	20766	38%	50%	27

Table 1. Comparison of variable price and fixed price channel time allocation algorithms.

on the downlink of a time-slotted or CDMA-based wireless LAN. They assume that users do not know each others' utility functions. In [7], the authors divide the network bandwidth into stable (low bandwidth) and instantaneous (unstable, high bandwidth) classes, and broadcast a price-service menu for these classes periodically. The users are given an incentive to truthfully declare their required bandwidth and service class. As far as we know, the concept of *channel time proportion* is unique to our IEEE 802.11-based scheme. Our scheme also directly auctions channel time so that users who value it more obtain more of it.

There has also been research in the area of price-based resource allocation for wireless ad-hoc networks [15, 11]. In [15], the authors argue that the shared resource is not a link, as in the case of a wireline network, but a wireless neighborhood clique. On the basis of this, they adapt the concepts of Kelly et al's seminal paper [5] on price-based resource allocation in wireline networks to mobile ad hoc networks.

7. Conclusion

In this paper, we presented a price-based channel time allocation scheme for a wireless hot spot network, in order to alleviate the problem of congestion, and provide admitted users acceptable QoS. We made two main contributions in this work. First, we introduced the concept of *channel time proportion* (CTP), and converted users' throughput requirements to channel time requirements. Secondly, we presented a centralized auction algorithm, of quadratic time-complexity, to determine the instantaneous price of the channel and the users' respective CTP allocations. We compared the performance of our algorithm with two fixed-price channel allocation schemes and found that our algorithm does the best job of simultaneously maximizing revenue, user satisfaction and channel utilization.

References

[1] V. Bharghavan, K. Lee, S. Lu, S. Ha, J. Li, and D. Dwyer. The timely adaptive resource management architecture. *IEEE Personal Communication Magazine*, 5(4), August 1999.

[2] J. Bulow and P. Klemperer. Prices and the winner's curse. *RAND Journal of Economics*, 2002.

[3] K. Chen and K. Nahrstedt. Exact: An explicit rate-based flow control framework in manet. Tech. Report UIUCDCS-R-2002-2286/UILU-ENG-2001-1750, Dept. of Computer Science, UIUC, July 2002.

[4] K. Hendricks and R. Porter. Collusion in auctions. *Annales D'Economie et de Statistique*, 15/16, 1989.

[5] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: Shadow prices, proportional fairness and stability. *Journal of Operational Research Society*, 49, 1998.

[6] P. Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 1999.

[7] R. Liao, R. Wouhaybi, and A. Campbell. Incentive engineering in wireless lan based access networks. In *IEEE ICNP*, Paris, France, November 2002.

[8] P. Marbach and R. Berry. Downlink resource allocation and pricing for wireless networks. In *IEEE InfoCom*, New York, NY, June 2002.

[9] R. McAfee and J. McMillan. Bidding rings. *American Economic Review*, 82, 1992.

[10] M. Mirhakkak, N. Schult, and D. Thomson. Dynamic bandwidth management and adaptive applications for a variable bandwidth wireless environment. *IEEE JSAC*, 19(10), October 2001.

[11] Y. Qiu and P. Marbach. Bandwidth allocation in ad-hoc networks: A price-based approach. In *IEEE InfoCom*, San Francisco, CA, March-April 2003.

[12] S. Shah, K. Chen, and K. Nahrstedt. Dynamic bandwidth management in single-hop ad hoc wireless networks. In *IEEE PerCom*, Fort Worth, TX, March 2003.

[13] S. Shah and K. Nahrstedt. Price-based channel time allocation in wireless lans. Tech. Report UIUCDCS-R-2003-2390/UILU-ENG-2003-1776, Dept. of Computer Science, UIUC, December 2003.

[14] N. Vaidya, P. Bahl, and S. Gupta. Distributed fair scheduling in a wireless lan. In *ACM MobiCom*, Boston, MA, August 2000.

[15] Y. Xue, B. Li, and K. Nahrstedt. Price-based resource allocation in wireless ad hoc networks. In *IEEE IWQoS*, Monterey, CA, June 2003.