

Throughput Optimization and Fair Bandwidth Allocation in Multi-Hop Wireless LANs

Qunfeng Dong

Department of Computer Sciences
University of Wisconsin-Madison
Madison, WI 53706, USA
Email: qunfeng@cs.wisc.edu

Suman Banerjee

Department of Computer Sciences
University of Wisconsin-Madison
Madison, WI 53706, USA
Email: suman@cs.wisc.edu

Benyuan Liu

Department of Computer Science
University of Massachusetts at Lowell
Lowell, MA 01854, USA
Email: bliu@cs.uml.edu

Abstract—There is an inherent well-known conflict between fairness and throughput that arises in many networking scenarios. A number of researchers have studied this problem in the context of (single-hop) wireless local area networks (WLANs), where clients directly exchange traffic with access points (APs). More recently, researchers have proposed multi-hop extensions to WLANs where client traffic is forwarded via a series of client-client links. In this paper, we show that the objective of improving throughput without sacrificing fairness can be much better met in multi-hop WLANs. We decouple this objective into two separate but related problems. First, we need an algorithm to organize clients into a multi-hop structure such that fair bandwidth allocation within this structure leads to improved throughput. Second, we need algorithms for performing fair bandwidth allocation within the determined multi-hop structure. In this paper, we first design optimal fair bandwidth allocation algorithms for both max-min throughput fairness and max-min time fairness in multi-hop WLANs. Subsequently, we design an efficient algorithm to find desirable multi-hop structures. With slight modifications, our results in this paper can be generalized to other multi-hop wireless networks as well. Our proposed solutions seamlessly integrate with legacy devices and hence are incrementally deployable. Simulation results demonstrate that our solutions can effectively improve throughput (by up to 114% or more) as well as network coverage while preserving fairness.

I. INTRODUCTION

Wireless local area networks (WLANs) have mushroomed at hotspots like office buildings, libraries, coffee shops, airports, hotels, etc. In a typical deployment, each *client* is equipped with an 802.11 interface and communicates over the air with an *access point* (AP) that is connected to Internet. This wireless communication is susceptible to signal quality degradation caused by fading, noise, interference, multi-path reflection, and user mobility, etc. When the average received signal strength is consistently below the threshold required for successful packet reception, the receiver experiences significant packet losses. To communicate more reliably, the sender can transmit at a lower *bit rate* (using a more resilient modulation scheme) so that the channel bit error rate is reduced. Many 802.11 vendors have implemented automatic schemes for such bit rate control [1]–[3]. Researchers [4], [5] have reported that rate diversity is prevalent in many in-door WLANs and exists even in small rooms, which are most appropriate for hotspots.

Previous studies of corporate WLANs [6] and campus-wide wireless networks [7] have shown that WLANs often carry sig-

nificant traffic and contain many APs that have frequent busy or congested periods. When multiple clients contend for access to the same wireless channel, a channel allocation scheme is needed to distribute channel access time among competing clients, according to some fairness policy. The Distributed Coordination Function (DCF) MAC protocol used by 802.11 tends to give equal long-term channel access opportunities to all competing clients [8], [9]. In particular, each node has (approximately) the same number of opportunities to transmit a data frame, regardless of its bit rate and hence the amount of channel access time needed. If clients transmit packets of similar sizes and experience similar loss rates, they achieve approximately the same throughput irrespective with their bit rates. This is referred to as *throughput-based fairness* [5]. Consequently, aggregate throughput and throughput of high bit rate clients may be dramatically brought down, because low bit rate clients will occupy more channel access time to transmit an equal amount of data. Such “performance anomaly” of 802.11 WLANs has been reported by Heusse *et al.* in [10].

This inefficiency leads to a number of potential problems. For example, in 54Mbps 802.11g networks that are deployed alongside 11Mbps 802.11b networks, 802.11g users may see far less performance improvement than expected and thus hesitate on upgrading to 802.11g. In order to address these inefficiencies, Tan and Guttag [5] proposed *time-based fairness*, where each client is assigned an equal amount of channel access time, regardless of its bit rate. Clearly, time-based fairness protects high bit rate clients from drastic throughput degradation by reserving a fixed share of channel access time for them. However, this disadvantages low bit rate clients. In many cases, this is considered undesirable, too.

A key problem with throughput-based fairness and time-based fairness is that so far they are both constrained to maneuver with client-AP links, some of which may have a low bit rate. We refer to such WLANs as single-hop WLANs. Furthermore, in currently deployed WLANs containing multiple APs, clients are typically restricted to use the link between themselves and the AP that has the strongest received signal strength indicator (RSSI). Bejerano *et al.* [11] propose to relax these constraints by allowing clients to switch to others APs in order to balance workload among APs. However, although such load balancing techniques can effectively improve fair-

ness, they are still restricted to use client-AP links and thus achievable throughput improvement is limited.

To improve throughput, some recent work (e.g. [12]) have made the observation that peer-to-peer links between nearby clients often possess high bit rates. By forwarding client traffic via high quality client-client links, multi-hop WLANs have the potential to significantly improve client throughput. However, ad hoc mechanisms to choose multi-hop path can often hurt performance. In this paper, we address the problem of finding appropriate multi-hop paths in multi-hop WLANs. We show that carefully choosing multi-hop paths can significantly improve throughput without sacrificing fairness. To achieve that, two relevant problems need to be addressed. First, we need an algorithm to organize clients into a multi-hop structure such that fair bandwidth allocation within this structure leads to improved throughput. Second, we need algorithms to perform fair bandwidth allocation within the determined structure.

Recently, Gambiroza *et al.* [13] attempt to address the second problem in the context of wireless backhaul networks, where the network topology is relatively static. In their work, the authors only provide an optimal solution for bandwidth assignment in the special case where all the links interfere with each other, i.e., the link contention graph is a clique. However, in many scenarios not all APs and clients are in direct interference range of each other. Finding an optimal solution in this general case is more difficult.

In this paper, we address both of the two relevant problems in the general setting, and we also examine how our solutions should adapt to dynamic changes of network topology. Such changes can be fairly common in WLAN scenarios. In particular, we make the following key contributions.

Key contributions

- We consider both max-min throughput fairness and max-min time fairness in multi-hop WLANs. For each fairness policy, we define and analyze an optimal algorithm for fair bandwidth allocation within a multi-hop WLAN.
- We design an efficient algorithm to smoothly improve the structure of WLANs. Simulation results demonstrate that our solutions can effectively improve throughput (by up to 114% or more) as well as network coverage while preserving fairness.
- Our solutions seamlessly integrate legacy client devices (that do not implement our algorithms) with smart client devices (that implement our algorithms) and hence allows a smooth transition through incremental deployment. Legacy client devices are directly associated with APs as usual and need not participate our algorithms to form a multi-hop WLAN. Such incremental deployment is not only more feasible than global upgrade, but also better motivated by our solutions, which reward individual clients investing to upgrade their device with much more perceivable performance improvement. These properties make our solutions practically interesting ones.
- With slight modifications, our results can be generalized to other multi-hop wireless networks as well.

Roadmap

The rest of the paper is organized as follows. We first present relevant models, definitions, notations, and formulations in Section II. Optimal bandwidth allocation algorithms for individual fairness policies are presented in Section III and IV. For practical purposes, we then propose a number of practical extensions to these algorithms in Section V. Our algorithm for improving network structure is presented in Section VI. After presenting simulation results in Section VII and reviewing related work in Section VIII, we conclude the paper in Section IX.

II. FORMULATION

In our network model, a WLAN is represented by a graph $G = (V, E)$, where V is the set of nodes representing APs and clients, and E is the set of edges representing communication links between nodes. For each node i , let b_i denote the bandwidth assigned to node i . We define the *bandwidth vector*, $\vec{B} = (b_1, b_2, \dots, b_n)$, as the clients' bandwidths sorted in non-decreasing order. For simplicity, we assume that clients are named according to this non-decreasing order of bandwidth.

The link layer model we use in this paper is based on the globally used DCF mode of IEEE 802.11 [14], where transmitting stations contend for channel access and the intended receiver is not aware of intended transmissions until it hears from the transmitter. Link layer details (such as bit rates, packet loss rates, channel access contention, etc) are implicitly incorporated in two link quality parameters called *effective bit rate (EBR)* and *actual bit rate (ABR)*, which are periodically measured by clients and APs. The EBR of a link represents the rate at which data can be successfully delivered over that link, if the wireless channel is exclusively occupied by that link and hence there is no contention. The ABR of a link represents the rate at which data can be successfully delivered over that link, if the sender has to contend for channel access.

For example, assume the link from client 4 to the AP in Figure 1(a) has a bit rate of 11Mbps and a packet loss rate of 20%. Its EBR is given by $11 \times (1 - 20\%) = 8.8\text{Mbps}$. Assume that due to channel contention with concurrent traffic between client 2 and client 3, even if client 4 is dedicated to the conversation with the AP, the link can only be active for 50% of the time. Without further considering other factors, the ABR of the link is thus given by $11 \times (1 - 20\%) \times 50\% = 4.4\text{Mbps}$. Some other factors may also affect the EBR and ABR of a link. For example, congestion window size has impact on the short term EBR of a link. In our model, EBR and ABR are measured by clients and APs to reflect a reasonably long term average. For each pair of node i and node j , we use $e_{i,j}$ and $a_{i,j}$ to denote the EBR and ABR of the link from i to j , respectively.

In this paper, we focus on the class of multi-hop WLAN structures where clients are organized into multi-hop trees each rooted at an AP. Such tree structures have been widely adopted by researchers for its simplicity and ease of management. Within a tree, each node i associates with a single parent node, denoted by P_i . Node i also has a set of child nodes, denoted

TABLE I
LIST OF NOTATIONS

b_i	bandwidth assigned to node i
P_i	parent node of node i in the multi-hop tree topology
Q_i	the set of child nodes of node i in the multi-hop tree topology
Q_i^+	$Q_i^+ = Q_i \cup \{i\}$
T_i	the subtree rooted at node i in the multi-hop tree topology
$ T_i $	the number of clients in the subtree rooted at node i
B_i	aggregate bandwidth assigned to nodes in T_i
$e_{i,j}$	EBR of the link from node i to node j
$a_{i,j}$	ABR of the link from node i to node j
t_j^i	the amount of node i 's time needed by B_j averaged over T_j
X_i	$X_i = 1$ if node i is a client; $X_i = 0$ if node i is an AP

by Q_i . For ease of presentation, we define $Q_i^+ = Q_i \cup \{i\}$ and denote the subtree rooted at node i by T_i . $|T_i|$ denotes the number of *client nodes* in T_i . The aggregate bandwidth assigned to nodes in T_i is denoted by B_i . The root node of a tree is an AP.


For ease of explanation, in the sequel we will focus on the uplink direction. The downlink direction similarly follows our discussion of the uplink direction. In the uplink direction of a tree rooted at an AP, every client needs to spend (wireless communication) time on receiving traffic from its children and on forwarding traffic for its children. The AP is a sink that does not generate its own traffic. Since it is the sink, the AP does not need to spend time on forwarding traffic that is received from its children. For simplicity, we define for each node i an indicator variable X_i such that $X_i = 0$ if node i is an AP and $X_i = 1$ if node i is a client. For each child node $j \in Q_i$ of node i , the fraction of node i 's time needed to receive traffic from node j at rate b_j is $b_j/e_{i,j}$, and the fraction of node i 's time needed to forward traffic from node j to node P_i (if any) at rate b_j is $X_i \cdot b_j/a_{i,P_i}$. The fraction of node i 's time needed to transmit its own traffic (if any) to node P_i at rate b_i is $X_i \cdot b_i/a_{i,P_i}$. Clearly, a bandwidth allocation (or bandwidth vector) \mathcal{B} is feasible at node i if and only if the workload on node i requires no more time than node i actually has. Namely, $\frac{X_i \cdot b_i}{a_{i,P_i}} + \sum_{j \in Q_i} \left(\frac{b_j}{e_{j,i}} + \frac{X_i \cdot b_j}{a_{i,P_i}} \right) \leq 1$. A bandwidth allocation \mathcal{B} is feasible if and only if it is feasible at every node.

For ease of reading, a list of notations is given in Table I.

A. Max-min throughput fairness

We now examine the fairness policies we study in this paper. As an extension of throughput-based fairness in single-hop WLANs that has been introduced in Section I, we consider the more general *max-min throughput fairness* [15] in multi-hop WLANs. Informally, a feasible bandwidth allocation is max-min throughput fair if and only if it is not possible to give any user more bandwidth without decreasing the bandwidth of some user with equal or already less bandwidth. Formally, max-min throughput fairness can be defined as follows.

Definition 1 (Max-Min Throughput Fairness): A feasible bandwidth allocation \mathcal{B} is max-min throughput fair if and only if its corresponding bandwidth vector $\vec{B} = (\bar{b}_1, \bar{b}_2, \dots, \bar{b}_n)$ has the same or higher lexicographical value than the bandwidth vector of any other feasible bandwidth allocation.



(a) max-min throughput fairness (b) max-min time fairness

Fig. 1. Example of max-min throughput fairness in a multi-hop WLAN. Numbers represent the bandwidth assigned to individual nodes.

In single-hop WLANs, clients associated with the same AP should receive equal bandwidth under max-min throughput fairness, which is consistent with throughput-based fairness. In multi-AP and multi-hop networks, max-min throughput fairness is better than throughput-based fairness in that it allows some clients to receive more bandwidth than other clients, if the latter are not able to consume more bandwidth. Thus, unnecessary waste of idle bandwidth can be avoided. This can be demonstrated by the example in Figure 1(a). In the example network, assume that the measured EBR and ABR of the link between the AP and client 4 are 2Mbps, and the measured EBR and ABR of the other links are 11Mbps. Under max-min throughput fairness, client 4 receives $\frac{4}{5}$ Mbps bandwidth and each of the other clients receives $\frac{11}{5}$ Mbps bandwidth. The aggregate throughput is $\frac{37}{5}$ Mbps.

B. Max-min time fairness

For each node i and each node $j \in Q_i$, we define node j 's *time share at node i* to be the amount of node i 's time needed to receive and forward traffic originating from nodes in T_j averaged over T_j , which is given by

$$t_j^i = \frac{\frac{B_j}{e_{j,i}} + \frac{X_i \cdot B_j}{a_{i,P_i}}}{|T_j|}.$$

For node i itself, its time share is its own time needed by transmitting its own traffic, which is $t_i^i = \frac{X_i \cdot b_i}{a_{i,P_i}}$. Given a bandwidth allocation \mathcal{B} whose bandwidth vector is $\vec{B} = (b_1, b_2, \dots, b_n)$, we define the *time share vector at node i* , $\vec{T}_i = (t_{j_1}^i, t_{j_2}^i, \dots, t_{j_k}^i)$, as the time shares of the $k = |Q_i^+|$ nodes in Q_i^+ sorted in non-decreasing order.

Ideally, time fairness in a multi-hop tree should ensure that for any node i , nodes in Q_i^+ receive equal time share at node i . However, if some node $j \in Q_i^+$ is not able to consume its time share at node i , its surplus time share at node i should be evenly distributed to other nodes in Q_i^+ . This ideal principle leads to our proposed notion of *max-min time fairness* formally defined as follows.

Definition 2 (Max-Min Time Fairness): A feasible bandwidth allocation \mathcal{B} is max-min time fair if and only if at each node i , its corresponding time share vector $\vec{T}_i = (t_{j_1}^i, t_{j_2}^i, \dots, t_{j_k}^i)$ has the same or higher lexicographical value than that of any other feasible bandwidth allocation.

The max-min time fair bandwidth allocation within the WLAN in Figure 1(a) is shown in Figure 1(b). Under max-min time fairness, client 1 and client 2 each receives $\frac{11}{6}$ Mbps

bandwidth, client 3 receives $\frac{11}{3}$ Mbps bandwidth, and client 4 receives $\frac{2}{3}$ Mbps bandwidth. The aggregate throughput is 8Mbps, which is more than the $\frac{37}{5}$ Mbps aggregate throughput achieved under max-min throughput fairness. Compared with max-min time fairness, max-min throughput fairness allocates more bandwidth to low EBR clients at the cost of high EBR clients and aggregate throughput, while max-min time fairness leads to a higher aggregate throughput by protecting high EBR clients at the cost of low EBR clients. This is consistent with the case of single-hop WLANs.

In single-hop WLANs, clients associated with the same AP should receive equal time of the AP under max-min time fairness, which is consistent with time-based fairness. In multi-hop WLANs, our proposed max-min time fairness turns out to be quite successful in two ways.

- As we will later see, compared with single-hop WLANs using time-based fairness, multi-hop WLANs using max-min time fairness universally improve client throughput.
- Compared with max-min throughput fairness, it generally leads to a higher aggregate throughput by protecting forwarding clients near the AP. This is appealing in many cases and more importantly, gives better motivation for clients to serve as a forwarding node near the AP, which means they will forward more traffic than their descendants in the tree.

C. Objective and design

The objective of this paper is to design solutions for improving throughput without sacrificing fairness in multi-hop WLANs. To achieve this objective, two problems need to be addressed. First, we need to design a tree construction algorithm to organize clients into a multi-hop structure such that fair bandwidth allocation within this structure leads to improved throughput. Subsequently, we need to design algorithms taking the determined structure and link EBRs as input to perform fair bandwidth allocation within the structure. As the tree construction algorithm relies on the fair bandwidth allocation algorithms to evaluate the quality of a structure, we first present our fair bandwidth allocation algorithms for max-min throughput fairness and max-min time fairness in Section III and Section IV, respectively.

III. MAX-MIN THROUGHPUT FAIR ALLOCATION

In this section, we present the idea, design, and analysis of *Max-Min Fair Allocation (MMFA)*, an optimal algorithm for max-min throughput fair bandwidth allocation in multi-hop WLANs. For ease of understanding, we start with the simple case where the WLAN is organized into a tree rooted at the only AP within the WLAN. We will investigate the case of multi-AP WLANs as well as other extensions in Section V.

A. General idea

Given the tree structure of a multi-hop WLAN, MMFA takes a bottom-up approach. At each node i in the tree, MMFA first recursively conducts max-min throughput fair bandwidth allocation within the individual subtrees rooted at node i 's

child nodes (if any), and then conducts max-min throughput fair bandwidth allocation within \mathcal{T}_i by performing *Pump-and-Drain* at node i . The Pump-and-Drain operation is as follows.

- *Pump*: If node i is a client, MMFA assigns a certain amount of bandwidth to node i such that node i receives the highest bandwidth among nodes in \mathcal{T}_i and node i 's time is completely used. There is no need to perform Pump at APs since APs should receive 0 bandwidth.
- *Drain*: The bandwidth allocation resulting from Pump may not be feasible, because node i may be overloaded after being assigned the highest bandwidth among nodes in \mathcal{T}_i . In that case, we need to decrease the bandwidth assigned to nodes in \mathcal{T}_i to ensure that the resulting bandwidth allocation is feasible and max-min throughput fair within \mathcal{T}_i .

Based on this general idea, we next present the detailed design and correctness proof of MMFA.

B. Detailed design

To perform Pump-and-Drain, each node i maintains and reports to its parent the following information, which can be locally determined by aggregating the information reported by its children (if any).

- The bandwidth assigned to node i , namely b_i .
- The total bandwidth assigned to nodes in \mathcal{T}_i , which is given by $B_i = b_i + \sum_{j \in \mathcal{Q}_i} B_j$.
- The distinct amounts of bandwidth assigned to clients in \mathcal{T}_i , which are stored in array \mathcal{L}_i in non-decreasing order. For simplicity, we assume that \mathcal{L}_i is automatically compacted so that $|\mathcal{L}_i|$ is always equal to the current number of distinct amounts.
- In array \mathcal{N}_i , the k th item $\mathcal{N}_i[k]$ is the set of clients in \mathcal{T}_i whose assigned bandwidth is $\mathcal{L}_i[k]$.

MMFA is defined as a recursive procedure. In particular, the execution of MMFA at node i consists of two steps.

- MMFA recursively calls MMFA for each child node $j \in \mathcal{Q}_i$ to conduct max-min throughput fair bandwidth allocation within \mathcal{T}_j .
- After recursive MMFA executions at nodes in \mathcal{Q}_i have returned, the MMFA execution at node i concludes by performing Pump-and-Drain at node i to achieve a max-min throughput fair bandwidth allocation within \mathcal{T}_i . After that, the MMFA execution at node i returns with its locally maintained information and reports the information to the parent of node i (if any).

MMFA runs in a distributive way. For each tree rooted at an AP in the WLAN, MMFA is called for the AP, which then recursively calls MMFA for its descendants in the tree. The sequence of recursive MMFA executions propagate in a top-down fashion and return in a bottom-up fashion (reporting local information to their calling MMFA execution at their parent). Finally, the AP determines the max-min throughput fair bandwidth allocation within the whole tree and spreads the allocation to clients in a top-down fashion. The EBR and ABR of links are periodically measured by clients and APs,

and are reported in a bottom-up fashion along the tree so that the root node of each subtree has complete information to correctly perform bandwidth allocation within the subtree.

Pump-and-Drain is described in details as follows.

Pump: Once we have determined the value of b_i (which is currently initialized to 0), it will be straightforward to determine the value of B_i , \mathcal{L}_i , and \mathcal{N}_i according to their foregoing description and the data structures reported by nodes in \mathcal{Q}_i . For simplicity, we assume that these data structures are implicitly updated each time a bandwidth allocation adjustment is made. If node i is an AP, $b_i = 0$ and there is no need to perform Pump. If node i is a client, the resource constraint at node i dictates that $W_i = \frac{b_i}{a_{i,P_i}} + \sum_{j \in \mathcal{Q}_i} \left(\frac{B_j}{e_{j,i}} + \frac{B_j}{a_{i,P_i}} \right) \leq 1$, where W_i represents the fraction of node i 's time needed to support the bandwidth B_i assigned to nodes in \mathcal{T}_i . We refer to W_i as the *workload* on node i . If $W_i \geq 1$, node i is considered *saturated*. To make $W_i = 1$, the amount of bandwidth that should be assigned to node i is

$$\Delta = \left(1 - \sum_{j \in \mathcal{Q}_i} \left(\frac{B_j}{e_{j,i}} + \frac{B_j}{a_{i,P_i}} \right) \right) \cdot a_{i,P_i}. \quad (1)$$

- If Δ turns out to be the highest bandwidth assigned to nodes in \mathcal{T}_i , we assign Δ bandwidth to node i . Since $W_i = 1$, there is no need to perform Drain. Pump-and-Drain is thus done.
- If Δ is not the highest bandwidth, we assign the current highest bandwidth, $\mathcal{L}_i[|\mathcal{L}_i|]$, to node i . Consequently, $W_i > 1$ and Drain needs to be performed to decrease the bandwidth assigned to nodes in \mathcal{T}_i so that $W_i = 1$.

Drain: Drain is performed in an iterative fashion. During each iteration, only those nodes with the highest bandwidth in \mathcal{T}_i (i.e., nodes in $\mathcal{N}_i[|\mathcal{L}_i|]$) are decreased, each by an appropriate amount δ . Let n_j denote the number of nodes in \mathcal{T}_j that are also in $\mathcal{N}_i[|\mathcal{L}_i|]$, namely $n_j = |\mathcal{T}_j \cap \mathcal{N}_i[|\mathcal{L}_i|]|$. To make $W_i = 1$, the amount of bandwidth to be decreased at each node in $\mathcal{N}_i[|\mathcal{L}_i|]$, δ , is given by

$$W_i - \left[\frac{X_i \cdot \delta}{a_{i,P_i}} + \sum_{j \in \mathcal{Q}_i} \left(\frac{n_j \cdot \delta}{e_{j,i}} + \frac{X_i \cdot n_j \cdot \delta}{a_{i,P_i}} \right) \right] = 1$$

$$\Rightarrow \delta = \frac{W_i - 1}{\frac{X_i}{a_{i,P_i}} + \sum_{j \in \mathcal{Q}_i} \left(\frac{n_j}{e_{j,i}} + \frac{X_i \cdot n_j}{a_{i,P_i}} \right)}.$$

- If $\delta > \mathcal{L}_i[|\mathcal{L}_i|] - \mathcal{L}_i[|\mathcal{L}_i| - 1]$, we decrease the assigned bandwidth of each node in $\mathcal{N}_i[|\mathcal{L}_i|]$ by $\mathcal{L}_i[|\mathcal{L}_i|] - \mathcal{L}_i[|\mathcal{L}_i| - 1]$ and repeat this iterative adjustment again.
- Otherwise, we decrease the assigned bandwidth of each node in $\mathcal{N}_i[|\mathcal{L}_i|]$ by δ and it is now the case that $W_i = 1$. Drain is thus done.

C. Correctness proof

Theorem 1: MMFA achieves max-min throughput fairness.

Proof: Without loss of generality, we ignore the trivial case where the tree is a singleton of an AP. In a bottom-up

order, we prove by induction on the depth of nodes in the tree that, after MMFA is executed at any node i , the bandwidth allocation within \mathcal{T}_i (denoted by \mathcal{B}_i) is feasible and max-min throughput fair and that $W_i = 1$.

Base case: In the base case, node i is a leaf client node. Pump will assign a_{i,P_i} bandwidth to node i , which is clearly feasible and max-min throughput fair for the singleton \mathcal{T}_i , and $W_i = 1$.

Inductive case: If $\Delta \geq \mathcal{L}_i[|\mathcal{L}_i|]$, it is clear that \mathcal{B}_i is feasible within \mathcal{T}_i and that $W_i = 1$. \mathcal{B}_i is max-min throughput fair, because there is no way to increase the bandwidth of any node in \mathcal{T}_i without decreasing the bandwidth of another node in \mathcal{T}_i that has equal or already less bandwidth.

- On one hand, there is no way to increase the bandwidth of node i without decreasing the bandwidth of another node in \mathcal{T}_i (which must have the same or less bandwidth), since node i is saturated.
- On the other hand, for any node k in \mathcal{T}_j rooted at some node $j \in \mathcal{Q}_i$, there is no way to increase node k 's bandwidth without decreasing the bandwidth of another node in \mathcal{T}_j with the same or less bandwidth. Because by our inductive assumption, bandwidth allocation within \mathcal{T}_j (denoted by \mathcal{B}_j) has been max-min throughput fair.

We next examine the case where $\Delta < \mathcal{L}_i[|\mathcal{L}_i|]$.

- On one hand, it is clear from the description of Drain that \mathcal{B}_i is feasible after Drain, and $W_i = 1$. Since node i has been saturated, it is not possible to increase the bandwidth of any node with the highest bandwidth in \mathcal{T}_i without decreasing the bandwidth of another node in \mathcal{T}_i , which must have the same or less bandwidth.
- On the other hand, consider any node k in \mathcal{T}_i whose bandwidth is not the highest. It is clear that $k \neq i$ and thus node k must reside in \mathcal{T}_j rooted at some node $j \in \mathcal{Q}_i$. Meanwhile, b_k is not decreased by Drain, since only nodes in $\mathcal{N}_i[|\mathcal{L}_i|]$ are ever decreased. To prove by contradiction, assume that we can increase b_k to b'_k without decreasing the bandwidth of another node in \mathcal{T}_i with b_k or less bandwidth. Let \mathcal{B}'_j denote the resulting bandwidth allocation within \mathcal{T}_j , which is clearly feasible. Compared with the original \mathcal{B}_j prior to Pump-and-Drain, the nodes in \mathcal{T}_j with b_k or less bandwidth have never been decreased, since they are not in $\mathcal{N}_i[|\mathcal{L}_i|]$. However, bandwidth of node k can be increased. This contradicts the inductive assumption that \mathcal{B}_j has been max-min throughput fair. ■

We illustrate the dynamics of MMFA and the effectiveness of multi-hop WLANs using the example in Figure 3. In the example network, solid lines represent direct associations between clients and the AP. Dashed lines represent unused links between nodes. Client-clients links have an EBR and ABR of 11Mbps. Links connecting client 8,9 to the AP also have an EBR and ABR of 11Mbps. Links connecting client 5,6,7 to the AP have an EBR and ABR of 5.5Mbps. Links connecting client 1,2,3,4 to the AP have an EBR and ABR of 0

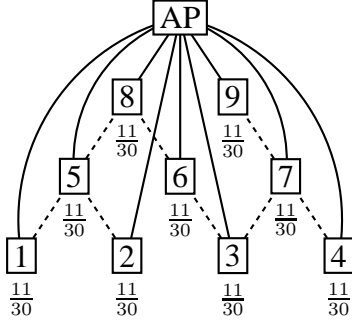


Fig. 3. Max-min throughput fair bandwidth allocation in a single-hop WLAN. Numbers represent bandwidth assigned to individual nodes.

2Mbps. Under max-min throughput fairness, each client will receive $\frac{11}{30}$ Mbps bandwidth.

In Figure 2, a multi-hop tree organization of the same network as well as a level-by-level illustration of MMFA applied on the tree are presented. By utilizing high quality client-client links, MMFA significantly improves throughput for every client in the network. In particular, some nodes receive 1Mbps bandwidth and the others receive $\frac{5}{3}$ Mbps bandwidth, which are almost $3 \sim 5$ times as much as the $\frac{11}{30}$ Mbps assigned to each client in the single-hop organization.

IV. MAX-MIN TIME FAIR ALLOCATION

Following the general idea of Pump-and-Drain in a bottom-up fashion, we here present the design and analysis of *Time-Based Fair Allocation (TBFA)*, an optimal algorithm for max-min time fair bandwidth allocation in multi-hop WLANs. At this point, we also consider the simple case of single-AP WLANs where clients are organized into a tree rooted at the AP. Multi-AP WLANs as well as other extensions will be investigated in Section V.

A. General idea

Given the multi-hop tree of a WLAN, TBFA also takes a bottom-up approach. At each node i in the tree, TBFA first recursively conducts max-min time fair bandwidth allocation within the subtrees rooted at node i 's child nodes (if any), and then conducts max-min time fair bandwidth allocation within \mathcal{T}_i by performing a similar but different Pump-and-Drain operation at node i . The Pump-and-Drain operation of TBFA is as follows.

- *Pump*: We divide node i 's time within \mathcal{Q}_i^+ to ensure that: (1) If node i is an AP, it receives 0 time; otherwise, it receives the highest time share among nodes in \mathcal{Q}_i^+ . (2) Each node $j \in \mathcal{Q}_i$ either receives the highest time share at node i , or receives the amount of node i 's time that is required to support the aggregate throughput B_j of nodes in \mathcal{T}_j . Let B_j^p denote the aggregate throughput of nodes in \mathcal{T}_j that can be supported with node i 's time that is allocated to nodes in \mathcal{T}_j by Pump.
- *Drain*: For each node $j \in \mathcal{Q}_i$, if $B_j > B_j^p$, we need to decrease the bandwidth of nodes in \mathcal{T}_j appropriately to ensure that $B_j = B_j^p$ and that the resulting bandwidth

allocation within \mathcal{T}_j (denoted by \mathcal{B}_j) is feasible and max-min time fair. Consequently, the resulting bandwidth allocation within \mathcal{T}_i (denoted by \mathcal{B}_i) is max-min time fair as well.

B. Detailed design

TBFA is a recursive procedure. In particular, the execution of TBFA at each node i consists of two steps.

- TBFA recursively calls TBFA for each child node $j \in \mathcal{Q}_i$ to conduct max-min time fair bandwidth allocation within \mathcal{T}_j .
- After these recursive TBFA executions at nodes in \mathcal{Q}_i have returned, the TBFA execution at node i concludes by performing Pump-and-Drain at node i to achieve a max-min time fair bandwidth allocation within \mathcal{T}_i . After that, the TBFA execution at node i returns with its locally maintained information and reports the information to the parent of node i (if there is one). To implement Pump-and-Drain, node i needs to maintain and report to its parent b_i , B_i , and $|\mathcal{T}_i|$, which can be locally determined by aggregating the information reported by its children.

TBFA runs in a distributive way. For each tree rooted at an AP in the WLAN, TBFA is called for the AP, which then recursively calls TBFA for its descendants in the tree. The sequence of recursive TBFA executions expand in the top-down order and return in the bottom-up order (reporting local information to their calling TBFA execution at their parent). Finally, the AP determines the max-min time fair bandwidth allocation within the whole tree and spreads the allocation to clients in a top-down fashion. The EBR and ABR of links are periodically measured by clients and APs. Unlike MMFA, there is no need to spread link EBRs and ABRs.

Now we describe the Pump-and-Drain of TBFA in details.

Pump: Pump is done in an iterative fashion. Initially, each node in \mathcal{Q}_i^+ is assigned 0 time share at node i . We refer to nodes in \mathcal{Q}_i^+ that have been assigned time of node i as *solved* nodes, and refer to the other nodes in \mathcal{Q}_i^+ as *unsolved* nodes. Let \mathcal{U}_i denote the set of unsolved nodes in \mathcal{Q}_i . If a child node j in \mathcal{Q}_i is assigned time of node i , we consider all nodes in \mathcal{T}_j to have been assigned time of node i as well.

During each iteration, based on the fraction of time available at node i , denoted by C_i ($C_i \leq 1$), Pump calculates the average amount of node i 's time that should be assigned to each unsolved node, which is given by

$$\frac{C_i}{X_i + \sum_{k \in \mathcal{U}_i} |\mathcal{T}_k|}.$$

Then, the fraction of node i 's time that should be assigned to an unsolved child node $j \in \mathcal{Q}_i$ is clearly

$$\frac{C_i \cdot |\mathcal{T}_j|}{X_i + \sum_{k \in \mathcal{U}_i} |\mathcal{T}_k|}.$$

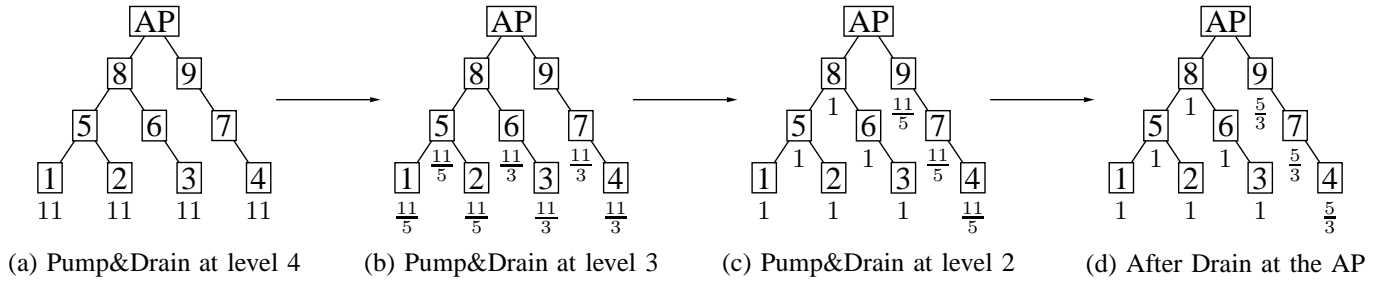


Fig. 2. Level-by-level illustration of MMFA in a multi-hop WLAN. Numbers represent bandwidth assigned to individual nodes.

The aggregate throughput B_j^p of nodes in \mathcal{T}_j that is allowed by node i 's time allocated to nodes in \mathcal{T}_j is thus given by

$$\frac{B_j^p}{e_{j,i}} + \frac{X_i \cdot B_j^p}{a_{i,P_i}} = \frac{C_i \cdot |\mathcal{T}_j|}{X_i + \sum_{k \in \mathcal{U}_i} |\mathcal{T}_k|} \implies B_j^p = \frac{\frac{C_i \cdot |\mathcal{T}_j|}{X_i + \sum_{k \in \mathcal{U}_i} |\mathcal{T}_k|}}{\frac{1}{e_{j,i}} + \frac{X_i}{a_{i,P_i}}}.$$

If node i is a client, its allowed bandwidth is given by

$$\frac{b_i}{a_{i,P_i}} = \frac{C_i}{1 + \sum_{k \in \mathcal{U}_i} |\mathcal{T}_k|} \implies b_i = \frac{C_i \cdot a_{i,P_i}}{1 + \sum_{k \in \mathcal{U}_i} |\mathcal{T}_k|}.$$

If node i is an AP, b_i is always 0.

If for some unsolved child node $j \in \mathcal{U}_i$, $B_j < B_j^p$, the assumption that bandwidth allocation within \mathcal{T}_j is max-min time fair implies that node j is not able to consume all of its allocated time of node i . Let ΔC denote the amount of node i 's time required to support B_j at node i , which is given by

$$\Delta C = \frac{B_j}{e_{j,i}} + \frac{X_i \cdot B_j}{a_{i,P_i}}.$$

We solve such node j by allocating ΔC time of node i to node j and removing node j from \mathcal{U}_i . After all such unsolved child nodes are solved, the current iteration terminates and we repeat the iterative procedure again.

If for every node $j \in \mathcal{U}_i$ it is the case that $B_j \geq B_j^p$, we node j its calculated time share at node i , and Pump is done. Node i is always solved during the last iteration.

Drain: Drain is also defined as a recursive procedure. In order to carry out the Drain operation, each node i locally maintains the following information.

- \mathcal{L}_i is a complete list of distinct amounts of time share at node i assigned to nodes in \mathcal{Q}_i^+ . For ease of presentation, we assume that \mathcal{L}_i is an array organized in non-decreasing order. Namely, $\mathcal{L}_i[1] < \mathcal{L}_i[2] < \dots < \mathcal{L}_i[|\mathcal{L}_i|]$. Moreover, \mathcal{L}_i is automatically compacted so that the length of \mathcal{L}_i is always equal to the number of distinct amounts.
- The k th element of array \mathcal{N}_i , $\mathcal{N}_i[k]$, is the set of nodes in \mathcal{Q}_i^+ whose time share at node i is $\mathcal{L}_i[k]$.

Assume that we need to perform Drain at node j because $B_j > B_j^p$. We decrease the time share at node j of nodes in \mathcal{Q}_j^+ in an iterative procedure that is reverse to the iterative procedure of Pump. During each iteration, only those nodes with the highest time share at node j (i.e., nodes in $\mathcal{N}_j[|\mathcal{L}_j|]$)

are decreased, each by an appropriate amount $\delta > 0$ such that $B_j = B_j^p$. δ is given by

$$\delta \cdot X_j \cdot a_{j,i} + \sum_{k \in \mathcal{Q}_j \cap \mathcal{N}_j[|\mathcal{L}_j|]} \frac{\delta \cdot |\mathcal{T}_k|}{\frac{1}{e_{k,j}} + \frac{X_j}{a_{j,i}}} = B_j - B_j^p,$$

which gives us

$$\delta = \frac{B_j - B_j^p}{X_j \cdot a_{j,i} + \sum_{k \in \mathcal{Q}_j \cap \mathcal{N}_j[|\mathcal{L}_j|]} \frac{|\mathcal{T}_k|}{\frac{1}{e_{k,j}} + \frac{X_j}{a_{j,i}}}}.$$

If $\delta > \mathcal{L}_j[|\mathcal{L}_j|] - \mathcal{L}_j[|\mathcal{L}_j| - 1]$, we decrease the time share at node j of nodes in $\mathcal{N}_j[|\mathcal{L}_j|]$ by $\mathcal{L}_j[|\mathcal{L}_j|] - \mathcal{L}_j[|\mathcal{L}_j| - 1]$ and repeat this iterative adjustment again. Otherwise, we decrease their time share at node j by δ and it is now the case that $B_j = B_j^p$. Then for each child node $k \in \mathcal{Q}_j$ whose time share at node j has been decreased during the iterative adjustment, we recursively perform Drain at that node, too. After these recursive Drain executions have returned, the Drain operation at node j returns as well.

C. Correctness proof

Theorem 2: TBFA achieves max-min time fairness.

Proof: In a bottom-up order, we prove by induction on the depth of nodes in the tree that after TBFA has finished executing at any node i : (1) Node i is saturated. (2) The computed bandwidth allocation within \mathcal{T}_i is feasible and max-min time fair. Without loss of generality, we here ignore the trivial case where the tree is an AP singleton.

Base case: If TBFA is executed at a leaf node i , node i 's time is totally allocated to its own traffic. It is straightforward that node i is saturated and bandwidth allocation at singleton i is feasible and max-min time fair.

Inductive case: For a non-leaf node i , it is clear from the description of Pump and Drain that resource constraints at individual nodes are always obeyed. Therefore, TBFA achieves a feasible bandwidth allocation. Moreover, node i is saturated after Pump-and-Drain at node i .

From the description of Pump, it is clear that for each node i , any node $j \in \mathcal{Q}_i^+$ either receives t_j^i time of node i or receives the highest time share at node i among nodes in \mathcal{Q}_i^+ . Drain operation reduces the time share at node i of nodes in \mathcal{Q}_i^+ in decreasing order of their time share at node i , and nodes with the highest time share at node i always remain among

the nodes with the highest time share at node i . Therefore, it remains to be the case that any node $j \in Q_i^+$ either receives t_j^i time of node i or receives the highest time share at node i among nodes in Q_i^+ .

Since node i is saturated, there is no way to increase the time share at node i of any node with the highest time share among nodes in Q_i^+ , without decreasing the time share at node i of some node in Q_i^+ that has equal or already less time share at node i . Any node $j \in Q_i^+$ that receives t_j^i time of node i can not receive more time share at node i , since node j is saturated (by inductive assumption). The conclusion is that bandwidth allocation within T_i is max-min time fair. ■

A level-by-level bottom-up illustration of TBFA applied on the example network in Figure 3 is given in Figure 4. The multi-hop tree structure is the same as Figure 2. Compared with the client bandwidth vector computed by MMFA in Figure 2, TBFA generally leads to a higher aggregate throughput by favoring forwarding clients near the AP. This property is appealing in many cases and more importantly, gives better motivation for clients to serve as a forwarding node near the AP, which means they will forward more traffic than their descendants in the tree.

V. EXTENSIONS

So far we have been dealing with a simplified WLAN model. In real applications, there are a number of practically important factors that need to be considered. Here we extend MMFA and TBFA to address the following issues.

Legacy clients

It is important for our proposed bandwidth allocation schemes to be incrementally deployable. Namely, they should seamlessly adapt to the case where a considerable portion of clients use legacy interfaces that are not able to participate and respond to our schemes. We point out that our algorithms do not assume or rely on the cooperation of such legacy interfaces. For example, a legacy 802.11 client interface is still directly associated with the AP that presents it with the strongest received signal strength indicator, and it is not required to forward traffic for any other node. Since such direct associations form a part of the tree topology and bandwidth allocation can be done by APs and forwarding nodes, our bandwidth allocation schemes apply as usual.

Limited capacity backhaul

By now, we have been assuming that the AP is connected to the Internet infrastructure through a backhaul link with sufficiently high bandwidth, so that the aggregate throughput of the WLAN is always fully supported. This is the case in many office buildings with 100Mbps or Gigabit LAN infrastructure. However, there are also many cases where the WLAN is connected to the Internet using a limited capacity backhaul such as a 768Kbps DSL link. We point out that a slight extension of our bandwidth allocation algorithms suffices to handle such cases. In particular, we can create a




Fig. 5. Effectiveness of a good multi-hop structure. Numbers represent assigned bandwidth of individual nodes.

virtual switch node representing the backhaul link and connect it to the AP using a channel with infinite EBR and ABR. The virtual switch is also connected to the infrastructure using an uplink whose EBR and ABR are equal to the backhaul capacity. Like the AP, the virtual switch always receives 0 bandwidth, too. Our bandwidth allocation schemes apply on this new virtual WLAN as usual.

Multi-AP WLANs

In many locations such as a large office building, a number of APs may be deployed to provide improved coverage and throughput. We point out that it is also straightforward to extend our bandwidth allocation schemes to handle such cases. If the shared backhaul link has sufficient bandwidth, there is no need to do anything. Because there is no contention between APs for limited backhaul capacity. We just run our algorithms within the trees rooted at individual APs separately. In the presence of a shared backhaul link with limited capacity, we can extend our schemes in a similar way. In particular, we create a virtual switch node representing the backhaul link and connect it to the APs using channels with infinite EBR and ABR. The virtual switch is also connected to the infrastructure using an uplink whose EBR and ABR are equal to the capacity of the backhaul. The virtual switch always receives 0 bandwidth. Our bandwidth allocation schemes apply on this new virtual WLAN as usual.

VI. TREE CONSTRUCTION ALGORITHM

Our fair bandwidth allocation algorithms in combination with appropriate multi-hop structures can significantly improve client throughput without sacrificing fairness. For the example WLAN in Figure 3, the multi-hop structure in Figure 5 demonstrates the effectiveness of a good tree structure. Using the multi-hop structure in Figure 2(d), the max-min throughput fair bandwidth allocation assigns 1Mbps bandwidth to some clients and $\frac{5}{3}$ Mbps bandwidth to other clients. However, using the multi-hop structure in Figure 5 leads to a more throughput fair bandwidth allocation where each client receives $\frac{11}{9}$ Mbps bandwidth and the aggregate throughput remains the same. In this section, we study the problem of finding good tree structures and present our *Tree Construction Algorithm (TCA)* for that purpose.

Before we can proceed to find a “good structure”, it still remains to define the notion of “good structure”. Consider two structures, α and β , whose resulted client bandwidth vectors

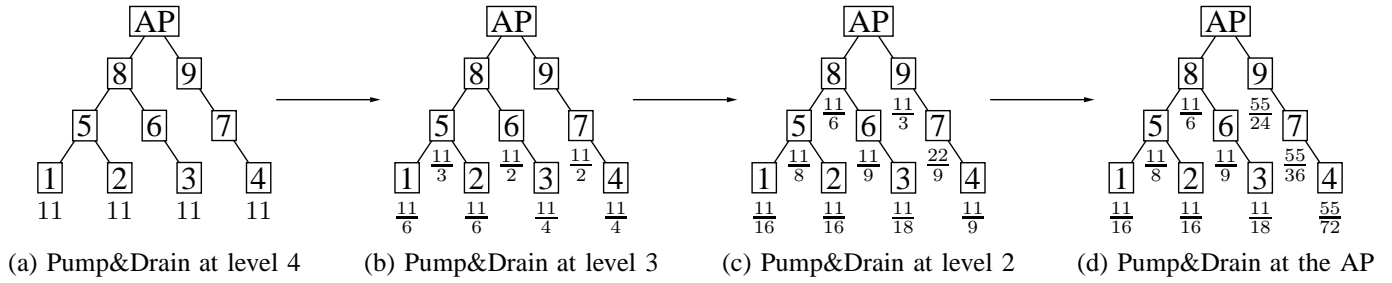


Fig. 4. Level-by-level illustration of TBFA in a multi-hop WLAN. Numbers represent bandwidth assigned to individual nodes.

under some certain fairness policy are \mathcal{B}_α and \mathcal{B}_β , respectively. We define \mathcal{T}_α to be “better” than \mathcal{T}_β if and only if \mathcal{B}_α has a higher lexicographical value than \mathcal{B}_β . We adopt this definition of “good structure” because our algorithms significantly improve throughput in multi-hop WLANs. Therefore, in defining “good structure” we should be more focused on fairness to balance between throughput and fairness. However, we point out that the reader is free to use any other notion of “good structure” in our TCA to find good structures of interest.

Finding the optimal network topology is a non-trivial task. In the case of max-min throughput fairness, this problem is actually provably intractable. Bejerano *et al.* [11] have shown that it is NP-hard to find the optimal association between APs and clients such that each client is associated with one single AP. The NP-hardness of finding the optimal multi-hop structure directly follows, since this is a more generalized problem.

Given that finding the optimal structure is a non-trivial and even intractable task, we here present our TCA, a heuristic to incrementally improve network structure in a smooth way. TCA runs in a (possibly periodic) iterative fashion. During each iteration, for each smart client node i (that implements our schemes), it probes to *migrate* node i with its subtree from its current location to become a child of another node j , which may be one of the APs or a smart client that is not in \mathcal{T}_i . The resulting client bandwidth vector is calculated using the corresponding bandwidth allocation algorithm. All such candidate (i, j) pairs are tested. If the bandwidth vector of the best (i, j) pair is better than the current bandwidth vector, TCA migrates client node i with its subtree to be a child of node j . This resulting new structure is the locally optimal structure that we can find at this point, and is accepted as the new structure. After that, TCA calls MMFA or TBFA for the new structure to perform fair bandwidth allocation, according to the adopted fairness policy. The migration decision and the corresponding bandwidth allocation within the WLAN are reported to the clients involved. The involved clients perform the migration and rate control accordingly. This iterative procedure halts at the point where such a locally better structure can not be found.

For example, assume the current structure of the WLAN in Figure 3 is the one in Figure 2(d). MMFA assigns 1Mbps bandwidth to some clients and $\frac{5}{3}$ Mbps bandwidth to other clients. During the following iteration, TCA decides that

migrating client 3 to become a leaf child of client 7 leads to the better structure in Figure 5, where MMFA assigns $\frac{11}{9}$ Mbps bandwidth to every client, which is considered better according to max-min throughput fairness.

Node arrival and departure can be smoothly handled in a similar way. Each time a new client joins the WLAN, TCA tests all (smart) accepting nodes and picks the one that will lead to the locally best new structure after accepting the new client as its child. The new client is then attached as its child, and bandwidth allocation within the new structure is conducted using the appropriate fair bandwidth allocation algorithm. If fast association is preferred, TCA may pick one of the APs based on some quick evaluation, and try to improve the resulting structure in successive iterations as usual. Such quick evaluation may be strongest received signal strength indicator (RSSI) or least-loaded-first (LLF), etc. When a node leaves the WLAN, its children (if any) can be directly attached to some AP (based on similar quick evaluations) in order to minimize the communication disruption perceived by clients which are descendants of the leaving node. After adjusting bandwidth allocation, TCA will try to improve the resulting structure as usual.

Legacy clients still associate with APs as they usually do, and TCA does not interfere with them at all. They are neither required to migrate nor required to accept migrating clients as their children. As legacy clients do not have any child node, they do not need to participate bandwidth allocation algorithm, either. Because their own bandwidth can be controlled by the AP that they are associated with.

In summary, our solution combining our TCA and bandwidth allocation algorithms provides a practically interesting smooth transition from legacy single-hop WLANs to smart multi-hop WLANs by two means.

- Our solution seamless integrates legacy clients with smart clients and hence allows a smooth transition from legacy technology to smart technology through incremental deployment. Such incremental deployment is not only more feasible than global upgrade, but also better motivated by our proposed algorithms, which reward individual clients investing to upgrade their device with much more perceivable throughput improvement.
- TCA constantly improves network throughput by improving network structure. Arrival and departure of nodes can be handled as structure changes in the same smooth way.

VII. EVALUATION

In our simulated scenarios, we compare the performance of our proposed solutions with the strongest received signal strength indicator (RSSI) method as well as the *integral load balancing algorithm (ILBA)* proposed by Bejerano *et al.* [11]. Note that RSSI and ILBA only decide the association between clients and APs. Different fairness policies can be enforced on any given association. For throughput fairness, we compare our “TCA+MMFA” with both RSSI and ILBA. For time fairness, we compare our “TCA+TBFA” with RSSI. Since ILBA is designed for optimizing max-min throughput fairness, it is not compared under time fairness.

Similar to [11], we choose link EBRs and ABRs according to the bit rates commonly advertised by 802.11b vendors. In particular, we assume that the EBR and ABR are 11Mbps for links no longer than 50 meters, 5.5Mbps for links no longer than 80 meters, 2Mbps for links no longer than 120 meters, and 1Mbps for links no longer than 150 meters, respectively. The transmission range of an AP is 150 meters. The backhaul connecting APs to the infrastructure has a typical LAN capacity of 100Mbps, which is more than enough to support the aggregate throughput of the WLAN.

Performance metrics

The performance of our solutions has two key aspects.

- *Throughput*: We study the client throughput distribution provided by individual schemes. In particular, we examine the lowest client throughput, median client throughput, highest client throughput, and aggregate throughput.
- *Fairness*: We use Jain’s Fairness Index [16] to evaluate the fairness provided by individual schemes. The Jain’s Fairness Index of a bandwidth vector $\vec{B} = (b_1, b_2, \dots, b_n)$ is given by

$$\frac{(\sum_{i=1}^n b_i)^2}{n \times \sum_{i=1}^n b_i^2}.$$

Intuitively, a bandwidth vector’s Jain’s Fairness Index is 1 if it is perfectly fair (i.e., clients receive equal bandwidth), and is $\frac{1}{n}$ if it is completely unfair (i.e., only one client is assigned bandwidth and all other clients are not).

Additionally, as TCA tries to iteratively improve the multi-hop tree structure, we also study the converging speed and adaptation ability of our proposed solutions.

- *Convergence*: We examine the converging process of our proposed solutions, starting from the single-hop association based on RSSI until TCA converges to some locally optimal multi-hop tree structure.
- *Adaptation*: We examine the ability of our solutions to quickly adapt to network topology changes such as node join and node leave. Both regular topology changes and random topology changes are simulated.

Throughput and fairness

We first compare the client throughput provided by the schemes in comparison. In our network setting, 30 clients

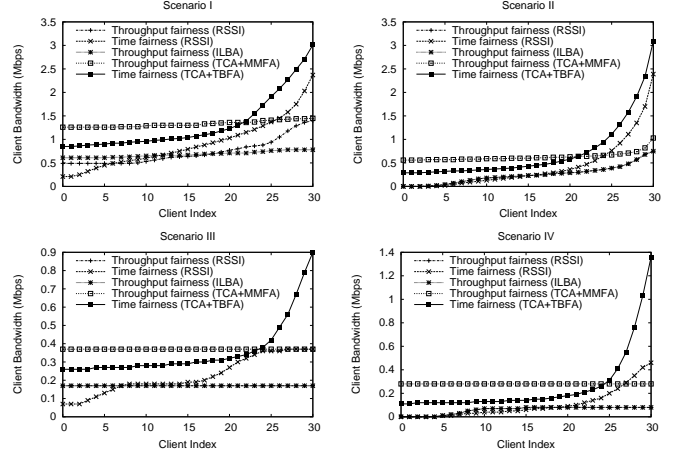


Fig. 6. Bandwidth allocation by different schemes in different scenarios.

are distributed in a square area uniformly at random. We believe this represents a moderately loaded network. Four representative scenarios are examined in our simulations.

- *Scenario I*: $150m \times 150m$ area, one AP at each corner.
- *Scenario II*: $300m \times 300m$ area, one AP at each corner.
- *Scenario III*: $150m \times 150m$ area, one AP at the center.
- *Scenario IV*: $300m \times 300m$ area, one AP at the center.

For each scenario, we conduct 1000 simulations and present the results in Figure 6. For each individual scheme in comparison, we present client throughput averaged over 1000 simulations in non-decreasing order. Note that in Scenario II, clients barely have access to more than one AP, so ILBA behaves essentially the same as RSSI. In Scenario III and IV, ILBA does not make difference at all since there is only one AP. Intuitively, the reader can think of the network setting as that of an indoor environment. Although in a typical indoor environment nodes may not be as far away from each other as in our network setting, various indoor signal degradation effects are similar to the effect of longer distance.

From the simulation results presented in Figure 6, the following conclusions clearly stand out. (1) Our proposed “TCA+MMFA” solution provides universally better client throughput than ILBA and RSSI for throughput fairness. (2) Although our proposed “TCA+TBFA” solution is not as fair as “TCA+MMFA”, it provides universally better client throughput than ILBA for throughput fairness and RSSI for both throughput fairness and time fairness. This property of “TCA+TBFA” is quite appealing. (3) Our multi-hop WLAN solutions not only improves client throughput, but also improves coverage. In Scenario II, clients located around the center can not directly associate with any AP. Techniques based on single-hop direct association such as RSSI and ILBA will allocate 0 bandwidth to these clients. In contrast, our multi-hop WLAN solutions can effectively provide good throughput for these clients.

Using the same client throughput data, we also examine the fairness and aggregate throughput achieved by schemes in comparison. To evaluate their fairness properties, we present




Fig. 7. Jain's Fairness Index of schemes in comparison.




Fig. 8. Aggregate throughput (Mbps) of schemes in comparison.

the Jain's Fairness Index [16] computed from the bandwidth vector of each scheme in Figure 7. The aggregate throughput of individual schemes is presented in Figure 8. From the figures, two conclusions can be drawn by comparison. (1) Our multi-hop schemes significantly improve throughput without sacrificing fairness. If clients can directly associate with APs (Scenario I and III), the aggregate throughput is improved by up to 114%. If some clients can not directly associate with any AP (Scenario II and IV), even much more throughput improvement can be achieved. (2) In cases where some clients can not directly associate with any AP, our multi-hop schemes significantly improve fairness as well, by providing good throughput for those clients.

Dynamic topologies, convergence, and adaptation

As TCA works in an iterative fashion, we also examine the convergence properties and adaptation properties of "TCA+MMFA" and "TCA+TBFA" using Scenario I and 40 clients. Simulation results of the the two examined solutions are presented in Figure 9 and Figure 10, respectively. For each solution, we first examine its convergence process starting from the single-hop configuration where each of the 40 clients directly associates with an AP, until after TCA converges to a locally optimal multi-hop tree structure. The lowest client bandwidth, the median client bandwidth, and the highest client bandwidth at each iteration of TCA are presented. Note that TCA aims to improve max-min throughput fairness, so the

lowest client bandwidth is only increasing.

After TCA has converged, the adaptation properties of each solution are investigated under two different scenarios: regular topology changes and random topology changes. The lowest client bandwidth at each iteration of TCA is presented. In the scenario of regular topology changes, topology changes occur every 20 iterations. In the scenario of random topology changes, a topology change occurs with probability 0.25 at each iteration. Each topology change involves a random node departure followed by a random node arrival. Topology changes are handled by TCA using the proposed quick evaluation based on RSSI. In the figures, each topology change is represented by a dashed vertical line. From the figures, the conclusion is that our solutions converge quickly and are quite responsive to topology changes.

VIII. RELATED WORK

Fairness in wire-line networks has been extensively studied in the literature, but resource allocation constraints significantly differ in wireless networks. Therefore, fair bandwidth allocation problems require a fresh investigation in the context of wireless networks. We briefly review the most relevant work.

In [17], Nandagopal *et al.* propose scheduling schemes for maximizing the sum of user utility in wireless networks, and point out that max-min fairness can be achieved as a special case using a certain choice of utility function. Tassiulas and Sarkar [18] argue that the optimization scheme becomes inefficient in such special cases and that max-min fair bandwidth allocation should be addressed separately. To derive a solution for max-min fair bandwidth allocation, the authors use a network model with a number of simplifying assumptions. For example, it is assumed that links that do not share nodes will never contend for channel access. Moreover, only single-hop flows are considered. [19] and [20] study arbitrary link contention graphs, but stick to the formulation where only single-hop flows are considered. Recently, Gambiroza *et al.* [13] take efforts to formulate the case of multi-hop flows and arbitrary link contention graphs. However, rigid analysis and bandwidth assignment are only presented for the special case where only one link within the whole network can be active at any point of time. In other words, the link contention graph is a clique containing all the links.

This paper differs from previous work in the following ways. First, we focus on the case of multi-hop WLANs. Second, we focus on allocating achievable fair shares of bandwidth to individual flows instead of scheduling and queuing schemes. Third, we not only address multi-hop flows, but also use a more general formulation that (implicitly) accommodates arbitrary link contention graphs. Fourth, we also address the problem of choosing an appropriate multi-hop structure.

In the context of single-hop WLANs, a number of researchers have investigated fairness and throughput issues. To name a few, Heusse *et al.* [10] report the "performance anomaly" of IEEE 802.11 WLANs due to the throughput-based fairness they implement. Tan and Guttg [5] propose

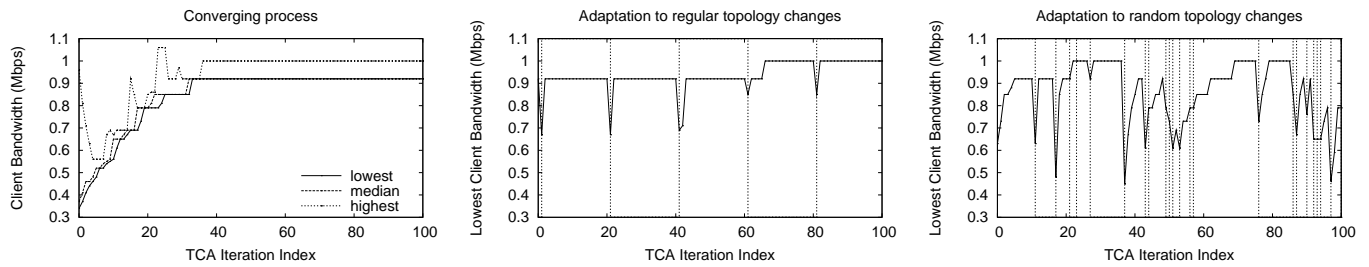


Fig. 9. Convergence/adaptation properties of TCA+MMFA under different scenarios. Vertical lines mark the time a topology change occurs.

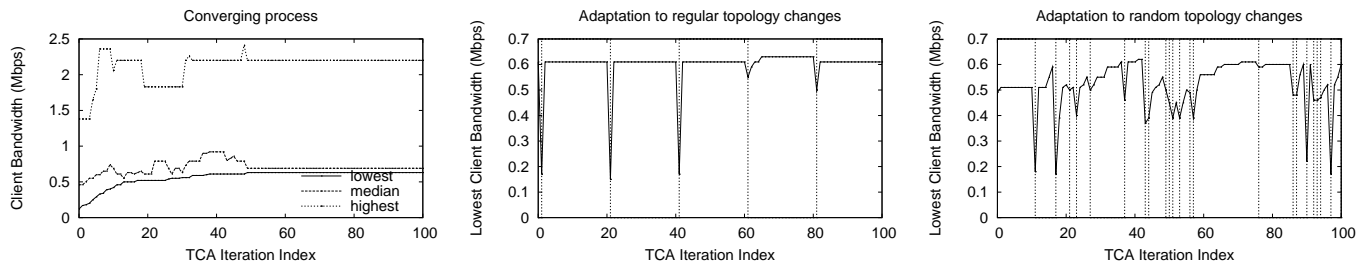


Fig. 10. Convergence/adaptation properties of TCA+TBFA under different scenarios. Vertical lines mark the time a topology change occurs.

time-based fairness to improve aggregate throughput by reserving a fixed share of channel access time for high bit rate clients. For multi-AP WLANs, Bejerano *et al.* [11] propose to improve throughput-based max-min fairness by balancing load among APs using association control.

There has also been some previous work on multi-hop WLANs (e.g. [12]). However, formal analysis and precise solutions for throughput optimization and fair bandwidth allocation have been lacking. To the best of our knowledge, this paper is the first of this type.

IX. CONCLUSION

In this paper, we study the problem of improving throughput while preserving fairness in multi-hop WLANs. To ensure fairness, we design optimal algorithms for max-min throughput fair bandwidth allocation and max-min time fair bandwidth allocation in multi-hop WLANs. To improve throughput, we design a heuristic to search for better network structures so that fair bandwidth allocation within the new structure results in a better client throughput profile. With slight modifications, our results can be generalized to other multi-hop wireless networks as well. The proposed solutions seamlessly integrate with legacy devices and hence are incrementally deployable. Simulation results demonstrate that our solutions can significantly improve both throughput (by up to 114% or more) and coverage while preserving fairness.

X. ACKNOWLEDGMENTS

Qunfeng Dong and Suman Banerjee were partially supported in this work by NSF award CNS-0520152.

REFERENCES

- [1] A. Kamerman and L. Monteban, "Wavelan II: A high-performance wireless lan for the unlicensed band," *Bell Labs Technical Journal*, pp. 118–133, Summer 1997.
- [2] Data Sheet of Cisco Aironet 350 Series Access Points.
- [3] ORiNOCO AS-2000 System Release Note.
- [4] D. Kotz, C. Newport, and C. Elliott, "The mistaken axioms of wireless network research," Department of Computer Science, Dartmouth College, Technical Report TR2003-467, July 2003.
- [5] G. Tan and J. Guttag, "Time-based fairness improves performance in multi-rate wireless lans," in *USENIX Annual Technical Conference*, 2004.
- [6] M. Balazinska and P. Castro, "Characterizing mobility and network usage in a corporate wireless local area network," in *MobiSys*, 2003.
- [7] D. Kotz and K. Essien, "Analysis of a campus-wide wireless network," in *ACM MobiCom*, September 2002.
- [8] C. E. Koksal, H. I. Kassab, and H. Balakrishnan, "An analysis of short-term fairness in wireless media access protocols," in *ACM SIGMETRICS*, 2000.
- [9] Y. Tay and K. Chua, "A capacity analysis for the IEEE 802.11 MAC protocol," *ACM/Baltzer Wireless Networks*, vol. 7, no. 2, pp. 159–171, 2001.
- [10] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *IEEE INFOCOM*, 2003.
- [11] Y. Bejerano, S.-J. Han, and L. E. Li, "Fairness and load balancing in wireless lans using association control," in *ACM MobiCom*, 2004.
- [12] S. Lee, S. Banerjee, and B. Bhattacharjee, "The case for a multi-hop wireless local area network," in *IEEE INFOCOM*, 2004.
- [13] V. Gambiroza, B. Sadeghi, and E. W. Knightly, "End-to-end performance and fairness in multihop wireless backhaul networks," in *ACM MobiCom*, 2004.
- [14] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification*, IEEE Std 802.11, LAN MAN Standards Committee of the IEEE Computer Society, 1999.
- [15] D. Bertsekas and R. Gallager, *Data Networks*. Prentice-Hall, 1987.
- [16] R. Jain, Ed., *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. Wiley-Interscience, 1991.
- [17] T. Nandagopal, T.-E. Kim, X. Guo, and V. Bharghavan, "Achieving mac layer fairness in wireless packet networks," in *ACM MobiCom*, 2000.
- [18] L. Tassiulas and S. Sarkar, "Maxmin fair scheduling in wireless networks," in *IEEE INFOCOM*, 2002.
- [19] X. L. Huang and B. Bensaou, "On max-min fairness and scheduling in wireless ad-hoc networks: Analytical framework and implementation," in *ACM MobiHoc*, 2001.
- [20] H. Luo, J. Cheng, and S. Lu, "Self-coordinating localized fair queuing in wireless ad hoc networks," *IEEE Transactions on Mobile Computing*, vol. 3, no. 1, 2004.