This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the ICC 2007 proceedings.

Fulfillment-based Fairness: A New Fairness Notion for Multi-AP Wireless Hotspots

Wei Zhou and Daji Qiao
Department of Electrical and Computer Engineering
Iowa State University, Ames, IA 50011
Email: {zhouwei, daji}@iastate.edu

Abstract— **Today, most wireless hotspots deploy multiple APs (Access Points) to improve the network performance and to provide fair and satisfactory services to clients. However, without a well-defined fairness objective function and the corresponding service provision schemes, the desired improvements may not be achieved. On the other hand, it has been realized that existing fairness notions proposed for single-AP wireless hotspots exhibit various performance anomalies in multi-AP wireless hotspots: *bandwidth anomaly* with Bandwidth-based Fairness (BbF) and *association anomaly* with Timeshare-based Fairness (TbF). To answer this challenge, we propose a new fairness notion, called the Fulfillment-based Fairness (FbF), for multi-AP wireless hotspots. It emphasizes allocation of bandwidth to clients in proportion to their respective maximum attainable bandwidth allocations. Extensive simulation shows that FbF outperforms BbF and TbF in terms of aggregate system throughput by up to 40% and 70%, respectively. It performs particularly well when the clients present a high degree of transmission rate diversity and/or in the presence of bottleneck clients that can only communicate with a single AP at low transmission rates.**

I. Introduction

The IEEE 802.11 WLAN (Wireless LAN) [1] has become the dominant technology for indoor broadband wireless networking. Known as *wireless hotspots* [2], public WLANs are springing up in conference venues, airport lounges, bookstores, cafes, and other public places to allow people to use their own portable devices such as laptops and PDAs to access the Internet. From January 2005 to January 2006, the number of wireless hotspots worldwide has grown 87% from 53,779 in 93 countries to 100,355 in 115 countries [3]. Most wireless hotspots deploy multiple APs (Access Points) to improve hotspot capacity and network performance, and to provide fair and satisfactory services to clients. However, without a well-defined fairness objective function and the corresponding service provision schemes, the desired capacity/performance/service improvements may not be achieved.

The popular Bandwidth-based Fairness (BbF) has been studied jointly with maximization of system throughput by many researchers [4], [5]. In fact, these two goals create inherent conflicts between them. For example, maximum system throughput may be achieved if each AP is assigned a non-interfering frequency channel[1] and serves a single client with the highest data rate (among all clients that are associated with this AP) while all other clients are starved. Clearly, this is unfair. In general, it is very difficult to achieve both goals at the same time. A more plausible objective is to provide network-wide fair bandwidth allocation to clients while maximizing the fair share of each client. This type of fairness is known as *max-min fairness*. In other words, bandwidth allocation among clients is called *max-min fair* if the bandwidth allocation of a client can not be increased without decreasing those of other clients who have already been allocated with smaller shares. The 802.11 DCF (Distributed Coordination Function) [1] was designed for this purpose, which guarantees long-term equal channel access probabilities among all clients in a single-AP wireless hotspot. In the rest of this paper, all fairness notions refer to max-min fairness. It has been realized in recent years that, with BbF as one of the network management goals, bandwidth allocated to each client is upper-bounded by the lowest transmission rate among all clients and, hence, serious system throughput degradation is inevitable in the presence of transmission rate diversity. Such *performance anomaly* in single-AP wireless hotspots was first discovered experimentally in [6] and later studied in depth via modeling and analysis in [7]. Similar problem can be observed in multi-AP wireless hotspots as well.

Since then, Timeshare-based Fairness (TbF) has been introduced [8] and deemed as a feasible remedy to the performance anomaly in single-AP wireless hotspots. *Timeshare* is defined as the fraction of time a client is able to access the channel to either transmit or receive packets from the AP. By reserving a fixed share of channel access time for each client, regardless of its transmission rate, TbF successfully prevents high-rate users from being "dragged down" by low-rate users in a single-AP wireless hotspot. However, in multi-AP wireless hotspots, we observe the the following phenomenon with TbF: since a client in a multi-AP wireless hotspot may have the option of associating with one of several available APs (and most likely communicating with them at different data rates), under certain circumstances, the client may opt to associate with an AP and communicate at a lower rate, in order to obtain a larger timeshare. Due to such *association anomaly*, TbF generally does not perform well in multi-AP wireless hotspots.

Above observations on BbF and TbF motivate the need for defining a more fitting fairness notion for multi-AP wireless hotspots, which is the key contribution of this paper. More specifically, we propose a new fairness notion, called the

[1] In a multi-AP wireless hotspot, each AP operates on an administrator-assigned frequency channel and each client typically associates with an AP. All communications between an AP and its associated clients occur on the channel assigned to the AP.

Fulfillment-based Fairness (FbF), to emphasize fair bandwidth fulfillment among clients. A client's *bandwidth fulfillment level* is a new concept. It is not an absolute bandwidth measurement (in Mbps) but a percentage value (i.e. no unit). It is defined as the ratio of a client's bandwidth allocation to its maximum attainable bandwidth allocation – which is achieved with the client-AP association plan that favors this client the most. Simulation results show that the proposed FbF works well in multi-AP wireless hotspots by addressing both *performance anomaly* and *association anomaly* and leads to significantly improved system throughput. Hence, it seems to be a more reasonable fairness notion than both BbF and TbF when designing/managing multi-AP wireless hotspots.

In multi-AP wireless hotspots, a common approach to provide fair services to clients is via load balancing. With the default 802.11 setting, a client always associates with an AP with the strongest RSSI (Received Signal Strength Indicator). This clearly may lead to unevenly distributed loads among APs and consequently potential degradation in aggregate system throughput [9]–[11]. To address this problem, an effective solution is to consider more parameters in addition to RSSI when making the client-AP association decision, such as load information of APs, channel variation and interference [12]–[14]. Another type of approaches is to use the cell breathing technique [15], [16], which allows APs to adjust their coverage areas by varying the transmit power of beacon frames. In [17], the authors proposed an efficient client-based approach for frequency assignment and load balancing in 802.11 WLANs that leads to better usage of the wireless spectrum. The authors of [18] proposed a load balancing algorithm by carefully planning client-AP association to balance loads among APs. Although the aforementioned schemes were all designed with Bandwidth-based Fairness as the target fairness criterion, we expect that they would work with FbF as well with necessary modifications. How to modify existing schemes or design new schemes to provide fulfillment-based fair services to clients in multi-AP wireless hotspots is not the focus of this paper, and will be addressed as part of the future work.

The rest of this paper is organized as follows. In Section II, we investigate the limitations of existing fairness notions in multi-AP wireless hotspots. Section III describes the details of the proposed Fulfillment-based Fairness (FbF) and gives a brief discussion on the implementation issues. Through extensive simulation, the performance of FbF and existing fairness notions are evaluated and compared in Section IV and, finally, we conclude the paper in Section V.

II. Limitations of Existing Fairness Notions

In this section, we discuss the limitations of existing fairness notions in multi-AP wireless hotspots. Before proceeding to the details, we first give the formal definition of *max-min fairness* and all the fairness notions in this paper refer to max-min fairness. A resource allocation plan is said to be *max-min fair* if the allocated resource share of one user can not be increased without sacrificing those of others who have already been allocated with smaller resource shares. Formally, given a resource allocation plan S, and let s_i denote the resource share (e.g. bandwidth, timeshare, or bandwidth fulfillment level) allocated to user *i*. Let $\vec{S}$ denote the corresponding resource allocation vector of plan S sorted in the non-decreasing order. Given two vectors $\vec{S} = \{s_1, s_2, ..., s_n\}$ and $\vec{S'} = \{s'_1, s'_2, ..., s'_n\}$, we say that $\vec{S}$ has a higher lexicographic order than $\vec{S'}$ if $s_1 > s'_1$ or $\exists\ \ell \geqslant 2$ such that $s_\ell > s'_\ell$ and $s_k = s'_k$ for all $1 \leqslant k \leqslant \ell$. We call resource allocation plan S *max-min* fair if $\vec{S}$ has the same or higher lexicographic order than that of any other plan.

A. Bandwidth-based Fairness (BbF)

The goal of BbF is to allocate fair bandwidth to clients regardless of their transmission rates. The 802.11 DCF was designed to provide BbF among clients in single-AP wireless hotspots. As a result, a high-rate client will be inevitably "dragged down" by low-rate clients, with its allocated bandwidth upper-bounded by the transmission rates of low-rate clients. Such *performance anomaly* [6] was originally discovered in single-AP wireless hotspots. Similar issues can be observed in multi-AP wireless hotspots, which is illustrated in the example below. Note that with max-min fairness, bandwidth allocations to high-rate clients are affected but not necessarily upper-bounded by the transmission rates of low-rate clients.

Example I. Consider a wireless hotspot shown in Fig. 1. Two IEEE 802.11a [19] APs (A_1 and A_2) operate on non-interfering frequency channels and two clients, C_1 and C_2, may associate with either AP. All stations are running the 802.11 DCF. Circles represent APs' coverage areas with radius of *r*. Each line represents a possible client-AP association and the number near the line represents the data rate (in Mbps) of the corresponding wireless link.

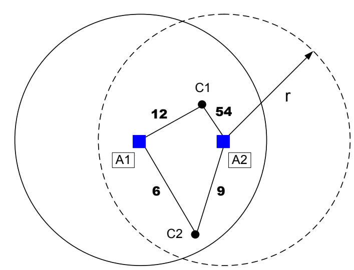

Fig. 1. Example I to illustrate the *performance anomaly* with BbF

Given a client-AP association plan, bandwidth allocations of clients are calculated using the simple load calculation model specified in [18], which ignores the transmission overheads such as contention window and backoff. Specifically, let C be the client of our interest and let A_C denote the AP that C is associated with. Furthermore, let $\{A_C\}$ denote the set of clients, each associated with A_C. Then the bandwidth allocated to C can be calculated by

$$B_C = \frac{1}{\sum_{z \in \{A_C\}} \frac{1}{R_{z,A_C}}}, \qquad (1)$$

where R_{z,A_C} is the transmission rate between client z and access point A_C. Possible client-AP association plans and the corresponding bandwidth allocations are compared in Table I. Clearly, the best association plan to achieve max-min BbF is to associate C_1 with A_1, and C_2 with A_2. Unfortunately, it results in a system throughput of 21 Mbps, which is only 35% of the maximum possible 60 Mbps. This example clearly shows the performance anomaly with BbF in multi-AP wireless hotspots, where bandwidth allocation to the high-rate client C_1 is affected by the low-rate client C_2.

TABLE I
BANDWIDTH COMPARISON WITH DIFFERENT CLIENT-AP ASSOCIATION PLANS IN EXAMPLE I

Client-AP Association Plan	B_{C_1}	B_{C_2}	B_{sys}	BbF Decision
$\{C_1 \leftrightarrow A_1, C_2 \leftrightarrow A_2\}$	12	**9**	21	√
$\{C_1 \leftrightarrow A_2, C_2 \leftrightarrow A_1\}$	54	**6**	60*	
$\{C_1 \leftrightarrow A_1, C_2 \leftrightarrow A_1\}$	**4**	**4**	8	
$\{C_1 \leftrightarrow A_2, C_2 \leftrightarrow A_2\}$	**7.7**	**7.7**	15.4	

B. *Timeshare-based Fairness (TbF)*

TbF was proposed to address the BbF-caused performance anomaly in single-AP wireless hotspots. Rather than allocating fair bandwidth to clients, the goal of TbF is to assign equal channel access time to all clients such that high-rate clients could transmit more data than low-rate clients during the same time period, thus yielding higher system throughput.

It may seem reasonable to apply TbF to multi-AP wireless hotspots to address performance anomaly. However, in a multi-AP wireless hotspot, a client may have the option of associating with several APs and most likely communicating with each of them at a different rate. Therefore, under certain circumstances, a client may opt to associate with a low-rate AP simply because such association allows the client to occupy the channel for longer time, hence increasing its timeshare. Such *association anomaly* is illustrated in the following example.

Example II. Consider a wireless hotspot shown in Fig. 2. Two 802.11a APs (A_1 and A_2) operate on non-interfering channels and there are three clients in the network. C_2 may associate with either AP, but C_1 can only associate with A_1, and C_3 can only associate with A_2. We call clients such as C_1 and C_3 *1-AP clients*, because they are only able to communicate with a single AP. Similar to the calculation of bandwidth allocation discussed in the previous section, timeshare allocated to a client C can be calculated by

$$T_C = \frac{\frac{1}{R_{C,A_C}}}{\sum_{z \in \{A_C\}} \frac{1}{R_{z,A_C}}}. \quad (2)$$

Possible client-AP association plans and the corresponding bandwidth and timeshare allocations are compared in Table II. The best plan to achieve max-min TbF is to associate C_1 with A_1, while C_2 and C_3 with A_2. The resulting system throughput is 24 Mbps, which is significantly lower than the maximum possible 33 Mbps. In this example, association anomaly occurs to C_2 as it chooses to communicate with A_2 at the low 6 Mbps rather than with A_1 at the high 54 Mbps, which is caused by the low-rate 1-AP client C_3. In general, the presence of low-rate 1-AP clients is one of the key causes to association anomaly.

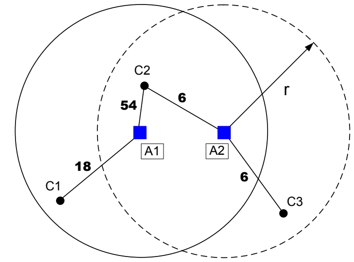

Fig. 2. Example II to illustrate the *association anomaly* with TbF

III. FULFILLMENT-BASED FAIRNESS (FBF)

Limitations of BbF and TbF motivate the need for defining a more fitting fairness notion for multi-AP wireless hotspots. In this section, we describe the details of our proposed new fairness notion for this purpose, called the Fulfillment-based Fairness (FbF).

A. *Definitions and Notations*

The goal of FbF is to address both performance anomaly and association anomaly in multi-AP wireless hotspots. It emphasizes fair bandwidth fulfillment among clients rather than fair allocation of the absolute bandwidth. A client's *bandwidth fulfillment level* is a new concept. Formally, it is defined as the ratio of a client's actual bandwidth allocation to its *maximum attainable bandwidth allocation*, which is achieved with the most favorable (with respect to this client) association plan that (i) reduces as much as possible the load of the AP this client is associated with, and (ii) still guarantees that each client is served by one of the APs.

A client's maximum attainable bandwidth allocation could be different from its maximum transmission rate, which will become clear when we revisit Example II in the next section.

B. *Examples Revisited*

Let's first revisit Example I with FbF. Since the most favorable association plans for C_1 and C_2 are $\{C_1 \leftrightarrow A_2, C_2 \leftrightarrow A_1\}$ and $\{C_1 \leftrightarrow A_1, C_2 \leftrightarrow A_2\}$, respectively, their maximum attainable bandwidth allocations are 54 Mbps and 9 Mbps, respectively, which happen to be the same as their maximum transmission rates. Possible client-AP association plans and the corresponding bandwidth allocations and fulfillment levels (denoted as F) are compared in Table III. Results show that the association plan to achieve max-min FbF indeed results in the highest system throughput, thanks to the fact that a client's bandwidth fulfillment level reflects not only its maximum transmission rate but also its available association options.

We now revisit Example II and from the comparison results shown in Table IV, we can see that max-min FbF and maximum system throughput are, again, achieved simultaneously. In this example, the most favorable association plan for C_1 is $\{C_1 \leftrightarrow A_1, C_2/C_3 \leftrightarrow A_2\}$, while $\{C_1/C_2 \leftrightarrow A_1, C_3 \leftrightarrow A_2\}$ is the most favorable association plan for both C_2 and

TABLE II

COMPARISON OF BANDWIDTH AND TIMESHARE ALLOCATIONS WITH DIFFERENT CLIENT-AP ASSOCIATION PLANS IN EXAMPLE II

Client-AP Association Plan	B_{C_1}	B_{C_2}	B_{C_3}	B_{sys}	T_{C_1}	T_{C_2}	T_{C_3}	TbF Decision
$\{C_1/C_2 \leftrightarrow A_1, C_3 \leftrightarrow A_2\}$	13.5	13.5	6	33*	0.75	**0.25**	1	
$\{C_1 \leftrightarrow A_1, C_2/C_3 \leftrightarrow A_2\}$	18	3	3	24	1	**0.5**	**0.5**	√

TABLE III

EXAMPLE I REVISITED WITH FBF

Association Plan	B_{C_1}	B_{C_2}	B_{sys}	F_{C_1}	F_{C_2}	FbF
$\{C_1 \leftrightarrow A_1, C_2 \leftrightarrow A_2\}$	12	9	21	**0.22**	1	
$\{C_1 \leftrightarrow A_2, C_2 \leftrightarrow A_1\}$	54	6	60*	1	**0.67**	√
$\{C_1 \leftrightarrow A_1, C_2 \leftrightarrow A_1\}$	4	4	8	**0.07**	0.44	
$\{C_1 \leftrightarrow A_2, C_2 \leftrightarrow A_2\}$	7.7	7.7	15.4	**0.14**	0.86	

C_3. Hence, the maximum attainable bandwidth allocations for C_1, C_2, and C_3 are 18 Mbps, 13.5 Mbps, and 6 Mbps, respectively. Notice the difference between C_2's maximum attainable bandwidth allocation of 13.5 Mbps and its maximum transmission rate of 54 Mbps. This is because, even with the most favorable association plan, C_2 still has to contend with C_1 to communicate with A_1. In fact, maximum attainable bandwidth allocations vary with the percentage of 1-AP clients in the network as well as their transmission rates. In general, the differences between maximum attainable bandwidth allocations and maximum transmission rates become less significant with smaller number of 1-AP clients present in the network. In the extreme case when there are no 1-AP clients in the network, i.e., each client can communicate with at least two APs, maximum attainable bandwidth allocations are the same as maximum transmission rates.

TABLE IV

EXAMPLE II REVISITED WITH FBF

Association Plan	B_{sys}	F_{C_1}	F_{C_2}	F_{C_3}	FbF
$\{C_1/C_2 \leftrightarrow A_1, C_3 \leftrightarrow A_2\}$	33*	**0.75**	1	1	√
$\{C_1 \leftrightarrow A_1, C_2/C_3 \leftrightarrow A_2\}$	24	1	**0.22**	0.5	

C. Implementation Options

The authors of [18] showed that it is an NP-hard problem to find the client-AP association plan that achieves max-min fairness in practical wireless hotspots, and several online intelligent association schemes [17], [18] have been proposed to achieve approximate max-min fairness. Although these schemes were designed with BbF as the target fairness criterion, they could work with FbF as well, only with necessary modifications.

In this paper, for simplicity, we evaluate the effectiveness of FbF against other fairness notions with the following simple association schemes: (i) the **Simple_maxmin** algorithm to find optimal client-AP association plans to achieve max-min fairness in small-scale networks via permutation test, whose pseudo code is shown in Fig. 3; and (ii) the **RS_maxmin** algorithm for large-scale networks via random shuffle. We assume there is a central administrator server (CAS), which can collect the network information from all APs and then execute **Simple_maxmin** or **RS_maxmin** to determine the client-AP association plan.

Algorithm 1 Simple_maxmin(A, C)

A : set of access points $\{a\}$
C : set of clients $\{c\}$
X : client-AP association plan $\{c \rightarrow a\}$
$\vec{S}$: resource allocation vector sorted in non-decreasing order

$X_{maxmin} = null$;
$\vec{S}_{maxmin} = \vec{0}$;
for $(\forall X)$ {
  $\vec{S}_X = \vec{0}$;
  Assign client-AP associations according to X;
  for $(\forall c \in C)$
    Calculate c's allocated resource s_c;
  if $(\vec{S}_X > \vec{S}_{maxmin})$ {
    $X_{maxmin} = X$;
    $\vec{S}_{maxmin} = \vec{S}_X$;
  }
}
return X_{maxmin}

Fig. 3. Pseudo-code of **Simple_maxmin**

Since **Simple_maxmin** checks all possible association plans by permutation, it is guaranteed to find the global optimal max-min fairness solution. The processing time depends on the number of clients and APs as well as the computation speed of CAS. The pseudo code of **RS_maxmin** is shown in Fig. 4. During each random shuffle, each client in the network is assigned with a random number and the client list is then sorted according to the random numbers. Starting from the client with the smallest number, each client determines its AP association that improves the resource allocation vector the most. This process continues and loops around the client list until the resource allocation vector stops improving. **RS_maxmin** approaches the optimal max-min fairness solution as the number of random shuffle increases, which is a tunable parameter.

IV. PERFORMANCE EVALUATION

In this section, we evaluate the performance of the proposed Fulfillment-based Fairness using the Qualnet simulator [20].

A. Simulation Setup

In the simulation, we assume that

- All clients and APs are static;
- Each station is equipped with an IEEE 802.11a interface that may transmit at one of the eight available rates: 6, 9, 12, 18, 24, 36, 48, and 54 Mbps;
- MAC protocol is the 802.11 DCF;

```
Algorithm 2 RS_maxmin(A, C)

A, C, X, S⃗ : same as those in Simple_maxmin

X_maxmin = null;
S⃗_maxmin = 0⃗;
while (n < num_shuffle){
    X'_maxmin = X' = null;
    S⃗'_maxmin = S⃗' = 0⃗;
    Let C' be a random permutation of C;
    while TRUE{
        for (∀c ∈ C) {
            for (∀a ∈ A) {
                if (c associating with a improves S⃗')
                    Associate c with a in X';
            }
        }
        if (S⃗' > S⃗'_maxmin) {
            X'_maxmin = X';
            S⃗'_maxmin = S⃗';
        }
        else break;
        if (S⃗'_maxmin > S⃗_maxmin) {
            X_maxmin = X'_maxmin;
            S⃗_maxmin = S⃗'_maxmin;
        }
    n++;
}
return X_maxmin
```

Fig. 4. Pseudo-code of **RS_maxmin**

- The reachability and the maximum transmission rate between a client and an AP is determined by the distance between them;
- Rate adaptation is disabled;
- All APs operate on non-interfering frequency channels.

We simulate two types of network scenarios: (i) small-scale networks with 3 APs and 10 clients, where optimal client-AP association plans to achieve max-min fairness are determined by the **Simple_maxmin** algorithm; (ii) large-scale networks with 10 APs and 40 clients, where client-AP association plans are determined by the **RS_maxmin** algorithm.

We compare the performances of client-AP association plans corresponding to Fulfillment-based Fairness (FbF), Bandwidth-based Fairness (BbF), and Timeshare-based Fairness (TbF), as well as the naive Strongest Signal First (SsF). Performance metric is the aggregate system throughput. In each simulation run, clients send CBR flow to their associated APs, and the CBR rates are set high enough to saturate the channel. Each point in the figures is averaged over 100 simulation runs.

B. Simulation Results

1) Small-Scale Networks: We first compare the performances of testing schemes when there are no 1-AP clients in the network. Simulation results are shown in Fig. 5. Note that Y-axis is not the absolute bandwidth measurement but the throughput improvement (in percentage) of FbF over other testing schemes. X-axis represents the standard deviation of the maximum transmission rates among all clients, denoted

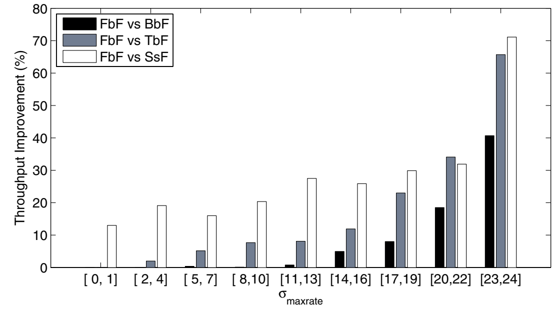

Fig. 5. Small-scale networks: comparison of all testing schemes under various rate diversity among clients

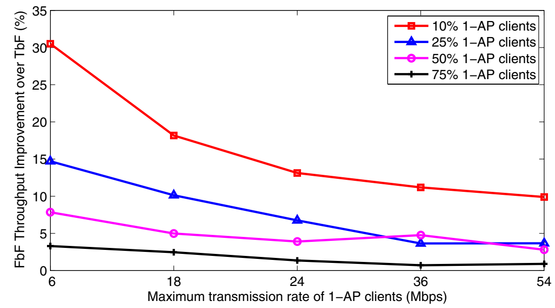

Fig. 6. Small-scale networks: comparison of FbF and TbF under various percentage of 1-AP clients and various maximum transmission rates of such 1-AP clients

by $\sigma_{maxrate}$, which we use to characterize the transmission rate diversity. In 802.11a networks, $\sigma_{maxrate}$ can be as large as 24 Mbps since the maximum difference between available transmission rates is 48 Mbps.

We have the following observations: (i) In general, FbF outperforms other testing schemes, and when the rate diversity is high, the performance improvements of FbF over other schemes become more significant; (ii) When the rate diversity is low, BbF and FbF show comparable performances; this is because the inherent performance anomaly with BbF is less likely to occur when most stations transmit at similar data rates; on the other hand, when the rate diversity is high, most likely BbF will yield performance anomaly; as shown in the figure, when $\sigma_{maxrate}$ is between 23 and 24 Mbps, FbF outperforms BbF by more than 40%; (iii) TbF doesn't perform well because of its inherent association anomaly while the poor performance of SsF is most likely due to the resulting unbalanced loads among APs.

Fig. 6 shows the throughput improvement of FbF over TbF with various percentage of 1-AP clients in the network and various maximum transmission rates of such 1-AP clients. As shown in the figure, with a fixed percentage of 1-AP clients in the network, the improvement of FbF over TbF increases as the maximum transmission rate of 1-AP clients decreases. This confirms our earlier discussion in Section II that the presence of low-rate 1-AP clients is one of the key causes to association anomaly. On the other hand, with a fixed maximum transmission rate of 1-AP clients, the improvement of FbF over TbF decreases with more 1-AP clients in the network. This makes sense because, with more 1-AP clients, fewer clients in the network can adjust their associations, and consequently, the benefit of applying intelligent association control becomes less significant. In fact, when the percentage of 1-AP clients reaches 100%, each client can only communication with a single AP, i.e. the client-AP associations have already been determined. In this situation, all fairness notions are equivalent.

2) Large-Scale Networks: We repeat the above simulation for large-scale networks and the number of shuffles in the **RS_maxmin** algorithm is set to 1000. Simulation results are plotted in Figs. 7 and 8. Similar trends are observed in large-scale networks as those in small-scale networks, while the throughput improvement of FbF over other testing schemes becomes even more significant. This is because larger-scale networks offer more options for client-AP associations and hence more room for performance enhancement.

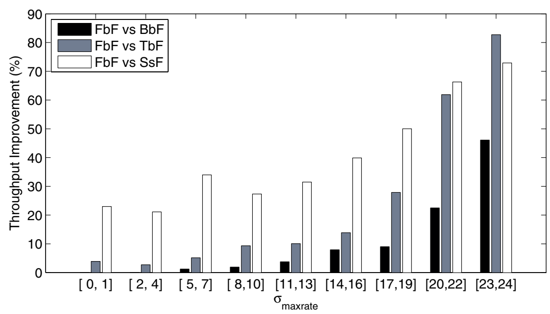

Fig. 7. Large-scale networks: comparison of all testing schemes under various rate diversity among clients

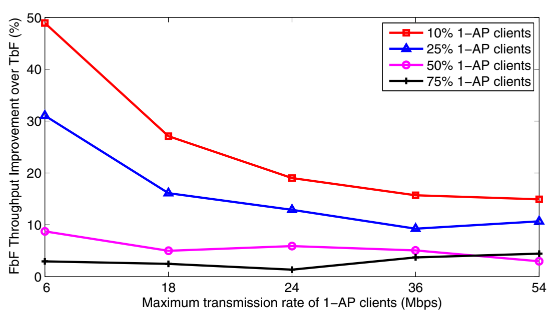

Fig. 8. Large-scale networks: comparison of FbF and TbF under various percentage of 1-AP clients and various maximum transmission rates of such 1-AP clients

V. Conclusion and Future Work

In this paper, we present Fulfillment-based Fairness (FbF) for multi-AP wireless hotspots. Fairness is, of course, a subjective notion, and we don't claim that the proposed FbF is "fairer" than others. In comparison to existing fairness notions, such as Bandwith-based Fairness (BbF) and Timeshare-based Fairness (TbF), the key idea of FbF is to allocate bandwidth to clients in proportion to their respective maximum attainable bandwidth allocations, which take into consideration not only clients' maximum transmission rates but also their association options to available access points. As a result, FbF does not suffer performance anomalies inherent with existing fairness notions in multi-AP wireless hotspots: performance anomaly with BbF or association anomaly with TbF.

Simulation results clearly show that FbF leads to vastly improved system throughput in the presence of high transmission rate diversity among clients and/or low-rate 1-AP clients that can only communicate with a single AP at low transmission rates, which can be often observed in practical wireless hotspots. Hence, we conclude that FbF seems to be a more reasonable fairness notion when designing/managing multi-AP wireless hotspots.

The work presented in this paper assumes that all clients have the same weight, and can be easily extended to scenarios with heterogenous weights among clients. Future work includes considering FbF jointly with channel assignment to APs and studying relevant problems, e.g., how to determine clients' bandwidth allocations and fulfillment levels when APs operate on overlapping or partially-overlapping frequency channels.

References

[1] IEEE 802.11, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, Aug. 1999.
[2] "http://www.wi-fihotspotlist.com/."
[3] "http://www.jiwire.com/press-100k-hotspots.htm."
[4] T. Nandagopal, T.-E. Kim, X. Guo, and V. Bharghavan, "Achieving mac layer fairness in wireless packet networks," in *Proc. ACM MobiCom'00*, 2000.
[5] L. Tassiulas and S. Sarkar, "Maxmin fair scheduling in wireless networks," in *Proc. IEEE INFOCOM'02*, 2002.
[6] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE INFOCOM'03*, 2003.
[7] D. Yang, K. Jang, and J. B. Chang, "Performance enhancement of multirate ieee 802.11 wlans with geographically scattered stations," in *IEEE Transactions on Mobile Computing*, 2006.
[8] G. Tan and J. Guttag, "Time-based fairness improves performance in multi-rate wireless lans," in *Proc. USENIX Technical Conference*, 2004.
[9] Cisco System Inc, Cisco aironet: Data sheet for cisco aironet 1200 series.
[10] M. Balazinska and P. Castro, "Characterizing mobility and network usage in a corporate wireless local-area network," in *Proc. ACM MOBISYS'03*, 2003.
[11] A. Vasan, R. Ramjee, and T. Woo, "Echos: Enhanced capacity 802.11 hotspots," in *Proc. IEEE INFOCOM'05*, 2005.
[12] A. Balachandran, P. Bahl, and G. M. Voelker, "Hot-spot congestion relief in public-area wireless networks," in *WMCSA '02: Proceedings of the Fourth IEEE Workshop on Mobile Computing Systems and Applications*, 2002.
[13] I. Papanikos and M. Logothetis, "A study on dynamic load balance for ieee 802.11b wireless lan." in *Proc. IEEE COMCON'01*, 2001.
[14] Y. Fukuda, T. Abe, and Y. Oie, "Decentralized access point selection architecture for wireless lans," in *Proc. IEEE VTC'04*, 2004.
[15] Y. Bejerano and S. Han, "Cell breathing techniques for load balancing in wireless lans," in *Proc. IEEE INFOCOM'06*, 2006.
[16] P. Bahl, M. Hajiaghayi, K. Jain, V. Mirrokni, L. Qiu, and A. Saberi, "Cell breathing in wireless lans: Algorithms and evaluation," in *IEEE Transactions on Mobile Computing*, 2006.
[17] A. Mishra, V. Brik, S. Banerjee, A.Srinivasan, and W. Arbaugh, "A client-driven approach for channel management in wireless lans," in *Proc. IEEE INFOCOM'06*, 2006.
[18] Y. Bejerano, S. J. Han, and L. E. Li, "Fairness and load balancing in wireless lans using association control," in *Proc. ACM MobiCom'04*, 2004.
[19] IEEE 802.11a, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-speed Physical Layer in the 5 GHz Band, Supplement to IEEE 802.11 Standard*, Aug. 1999.
[20] QualNet Simulator, http://www.scalable-networks.com/, Online Link.