Interference-Aware Time-Based Fairness for Multihop Wireless Networks

Douglas M. Blough
School of Electrical and Computer Engr.
Georgia Institute of Technology
Atlanta, GA 30332-0765
doug.blough@ece.gatech.edu

Giovanni Resta
Istituto di Informatica e Telematica
Consiglio Nazionale delle Ricerche
Pisa, Italy
g.resta@iit.cnr.it

Paolo Santi
IIT-CNR, Pisa – Italy, and
MIT Senseable City Lab
Cambridge, MA, USA
paolo.santi@iit.cnr.it

***Abstract*—We consider the problem of maximizing performance in multihop wireless networks while achieving fairness among flows. While time-based fairness has been widely recognized as the appropriate fairness mechanism in single-hop wireless networks, no analogous notion has been developed for multihop wireless networks. We define the first general notion of time-based fairness for multihop networks by abstracting a network into a virtual single-hop network and applying the single-hop time-based fairness notion. This produces rate shares for each flow in the network, and we develop a constructive method for achieving these rate shares through physical-interference-aware scheduling. When combined with an appropriate link transmission policy, this scheduling approach preserves the time-based-fair rate shares for flows even with spatial reuse and the resulting rate reductions that occur among concurrent links. To our best knowledge, this is the first constructive approach for achieving fair rate shares in multihop wireless networks with or without interference consideration. We also prove that, with an appropriate scheduling algorithm, this approach produces an aggregate rate that is within a constant factor of the maximum aggregate rate subject to time-based fairness. Finally, we perform extensive simulations, which show that our approach as much as doubles the aggregate rate of a solution that approximates max-min fairness, while achieving a more natural fairness property.**

I. Introduction

Network and protocol designers have long struggled to balance the competing notions of overall performance and treating individual users fairly. Typically, maximum overall performance is achieved by focusing resources on high-performing individual users at the expense of lower performers. Attempts to employ fairness constraints, i.e. ensuring some minimum level of performance to each individual user, inevitably lower overall network performance. Finding the right balance between these two objectives remains a challenging problem in many network settings. In this paper, we consider the problem of defining an appropriate fairness objective for multihop wireless networks that achieves good overall performance.

Time-based fairness, now widely accepted as the appropriate fairness concept for single-hop wireless networks, was developed in response to observations of the undesirable performance characteristics of rate-based fairness [5]. Rate-based fairness in single-hop networks means that different users are given an equal number of communication opportunities during each of which a fixed amount of data is transmitted. This results in the rates of all users degrading to the rates of the lowest-performing users, who require more air time to communicate a fixed amount of data [5]. Rate-based fairness also lowers the aggregate throughput of the system compared to alternative approaches. In time-based fairness, users are given equal amounts of *time* to access the wireless channel and they communicate whatever amount of data they can within their allocated time [14]. In this case, users' rates are proportional to the quality of their links.

D. Blough's research was supported in part by the National Science Foundation under Grants CNS-1319455 and CNS-1513884.

Most fairness approaches proposed in multihop wireless networks have not been time based. Most prior works have focused instead on max-min fairness. A flow rate vector is said to be max-min fair if a given flow's rate can be increased only by reducing the rate of another flow whose current rate is equal to or less than that of the given flow. It has been proven that, under fairly general assumptions on the multihop wireless network, the rates of all flows are equal in a max-min fair flow rate vector [11]. This result can be compared to rate-based fairness in single-hop networks, in that flows with all high-quality links are penalized by having to reduce their rates to equalize them with low-performing flows. Our goal, therefore, is to develop a time-based fairness approach for multihop wireless networks that is analogous to the widely adopted time-based approach for single-hop networks.

The approach we propose in this paper is to abstract the flows in a multihop network into *virtual single-hop links*. Virtual links are characterized by a *virtual rate*, defined as the maximum possible rate achievable on the flow in absence of interference from other flows. We then apply the notion of time-based fairness for single-hop networks to these virtual links and corresponding virtual rates, and show how this can be translated into a set of link-level constraints that enable this approach to be implemented efficiently in practical multihop wireless networks. We also show how this process can seamlessly handle spatial reuse and link rate variations caused by physical interference when generating the link constraints and scheduling transmissions. The result is a rigorous and general notion of time-based fairness for multihop wireless networks that achieves good overall performance, meets a natural and intuitively pleasing fairness objective, and is efficiently realizable in practical networks.

978-1-4673-9953-1/16/$31.00 ©2016 IEEE

II. Related work

Our work is inspired by research on fairness in multi-rate single-hop wireless networks [1], [5], [14]. Heusse, et al., were the first to show that the classic rate-based fairness approach severely penalizes users with high-rate links by equalizing their throughputs with the throughput of the lowest-rate link [5]. Tan and Guttag then defined the concept of time-based fairness, where different links were given equal times in the wireless channel instead of equal transmission opportunities [14]. Since the publication of [14], time-based fairness has been widely accepted as the best approach for single-hop wireless networks. Tinnirello and Choi showed how the TXOP mechanism in 802.11e could be used to achieve time-based fairness in WLANs [16]. Later, Blough, Resta, and Santi generalized the notion of time-based fairness for the multiple access point, but still single-hop, setting while accounting for physical interference between links [1].

There has been less clarity on the best approach to balancing performance and fairness in multihop wireless networks. Much of the work has considered max-min fairness or variations of it, e.g. [11], [12], [15], [17], [18]. Thulasiraman, Chen, and Shen considered the problem of multipath routing and bandwidth allocation to achieve max-min fairness [15]. To our knowledge, this is one of the only other works in this area that accounts for physical interference in its analyses. Raniwala, et al., proposed a congestion control mechanism at the transport layer that can provide max-min fairness on top of the 802.11 MAC [12]. Wang, et al. proposed a modified version of max-min fairness for multihop networks that accounts for intra-flow contention [17]. Zhang, et al. showed how max-min fairness can be achieved in multihop 802.11 wireless networks by fulfilling some simple local conditions at each station [18]. Finally, Radunovic and Le Boudec defined the notion of utility fairness, which generalizes both proportional fairness and max-min fairness [11]. This paper also demonstrated that, under fairly general conditions, max-min fairness results in equal rates being achieved by all flows.

A few papers propose metrics other than max-min fairness that are tailored to the multihop network setting [4], [11]. Gao and Jiang define a custom fairness index that measures the amount of variation between flow rates, where an index of one means that all flows have the same rate and smaller values of the index indicate larger variations between individual rates [4]. As mentioned earlier, Radunovic and Le Boudec's utility fairness generalizes several existing notions of fairness [11].

To our knowledge, only a few prior papers have considered the notion of time-based fairness for multihop networks [3], [8]. Li, et al., propose using the TXOP parameter in 802.11e networks to equalize air time for flows across backlogged links in the network [8]. While this mechanism provides for time-based fairness at the link level and is certainly fairer than the basic 802.11 MAC, the authors did not demonstrate that it achieves any well-defined notion of time-based fairness for flows at the network level. Gambiroza, Sadeghi, and Knightly consider fairness in multihop wireless backhaul networks, primarily focusing on a "parking lot" scenario and under the assumption that every link substantially interferes with every other link (no spatial reuse) [3]. The authors define temporal fairness as giving equal total air time to each flow. Under their no spatial reuse assumption, this definition is almost equivalent to our notion of assigning equal times to the virtual links of each flow.

Major differences of our work from the prior work include: 1) providing a rigorous and quite general formal definition of time-based fairness for multihop networks, 2) demonstrating how both spatial reuse and physical interference, which can affect the link data rates in the network, are seamlessly handled by this definition, 3) formally showing that our notion of time-based fairness avoids the equalization of flow rates problem inherent to max-min fairness, and 4) providing a constructive and efficient method to achieve the specified fairness objective to a rigorously derived precision.

III. Network and Interference Models

A. Problem Setting

We consider a general multihop wireless network setting without mobility. The network consists of a set of nodes $V = \{v_1, v_2, \ldots, v_n\}$ and a set of (directed) links $L = \{l_{i,j} :$ v_j can receive from v_i at the minimum data rate$\}$. To simplify presentation and discussion, we assume an STDMA setting, where transmissions are scheduled within time slots and multiple transmissions can occur in each slot. We assume that, for a certain scheduling period, flows are given as a set of source, destination pairs, i.e. $F = \{f_k = (s_k, d_k) : s_k, d_k \in V$ and s_k has information to send to d_k in the current period$\}$. We assume single-path routing and that routes are known, which is the case for networks that employ deterministic routing algorithms. We further assume that flows are continuously backlogged, i.e., that there is always at least one packet to send along the respective flow when a link is scheduled to transmit. This assumption isolates our study from the effects of buffering, which is out of scope for this paper. We also do not consider communication latency as an objective.

Each link $l_{i,j}$ is characterized by an *intrinsic data rate* $dr_{i,j}$, which is the rate experienced on the link *in absence of interference*. The intrinsic data rate depends on several factors such as the distance between transmitter and receiver, the radio propagation environment, etc. Since we do not consider mobility in this paper, in what follows we assume that intrinsic data rates remain fixed throughout a scheduling period.

The rate of a flow is the end-to-end throughput of the flow. Given F, a flow rate vector, $R[F]$ is a vector containing rates for every flow. $R[f_k]$ denotes the flow rate of f_k.

B. Interference Model

We consider the physical interference model in which link performance is dependent on the signal-to-interference-plus-noise ratio (SINR) at the receiver. To be precise, we compute data rates according to the *graded SINR model*, originally proposed in [9], [10] and formally defined in [13]. According to this model, the *effective data rate* experienced by link $l_{i,j}$

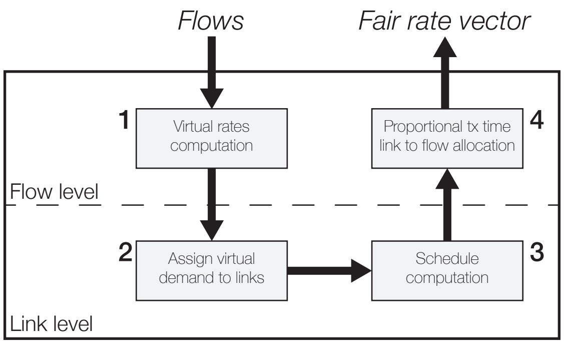

Fig. 1. Block diagram of the proposed time-based fairness framework.

when nodes in the set $\mathcal{T} = \{v_{t_1}, \ldots, v_{t_k}\}$ are transmitting simultaneously is given by

$$dr_{i,j}(\mathcal{T}) = f(\text{SINR}(i,j,\mathcal{T})) \ ,$$

where f is a non-decreasing function, and SINR$(i,j,\mathcal{T})$ is the SINR experienced at the receiver of $l_{i,j}$ when the nodes in $\mathcal{T}$ are transmitting.[1] When the set $\mathcal{T}$ contains only the transmitter of $l_{i,j}$, the data rate of $l_{i,j}$ is equal to its intrinsic data rate, i.e., $dr_{i,j}(\mathcal{T}) = dr_{i,j}$ in this case.

While the notion of multihop time-based fairness introduced in the next section does not depend on the specific rate function used, in this paper we consider two example functions for rates. A first example rate function is a staircase function, with data rates determined by SINR threshold values as shown in Table I [7]. The staircase function is based on 802.11a/g technology and we use it when studying concrete network scenarios through simulation in Section VI.

Data rate	Min SINR (dB)	Data rate	Min SINR (dB)
6 Mbps	6	24 Mbps	17
9 Mbps	8	36 Mbps	19
12 Mbps	9	48 Mbps	24
18 Mbps	11	54 Mbps	25

TABLE I
THRESHOLD SINR VALUES FOR 802.11A/G DATA RATES.

While the staircase function is based on a widely-deployed wireless technology, it is not well suited to mathematical analyses. To simplify our analyses, when deriving algorithmic performance bounds, we use another example function [1]. This rate function is defined in the following way: if the SINR (in dB) is below the value $\beta_0 = 0$, the rate is zero; if the SINR is above the value β_{max}, the rate is equal to a maximum value dr_{max}; and in the interval between β_0 and β_{max}, the rate increases linearly with SINR.

IV. Time-Based Fairness for Multihop Wireless Networks

A. Virtual Links

The concept of time-based fairness and its advantages over rate-based fairness were originally developed for an isolated single-hop network scenario [14]. Trying to define an appropriate version of time-based fairness for mulithop networks that is analogous to the one used in single-hop networks is complicated by several characteristics of multihop networks that are absent in the isolated single-hop case. Two characteristics, in particular, that impact the situation are intra-flow contention and inter-flow interference, neither of which is an issue in an individual single-hop network. Our approach is to abstract the multi-hop network so as to bring it closer to a single-hop network.

The first step in dealing with the multi-hop network is to abstract away inter-flow interference by considering operation of a single flow at a time, as if flows performed time sharing of the complete network. Next, we deal with intra-flow contention by optimally scheduling the links of each flow according to the given interference model. The result of this is to obtain a rate for each flow, which is the rate of the flow in the absence of inter-flow interference but considering intra-flow contention, which we call the *virtual rate*. This is analogous to the instrinsic rate of a link in single-hop networks [1], since they both capture the performance of a user in the absence of other users. To complete the abstraction, we treat each flow as if it were a single link with an instrinsic rate equal to the calculated rate of the flow in the absence of inter-flow interference. We refer to these abstracted links as *virtual links*. We denote the virtual rate of flow f_i (flow rate in absence of inter-flow interference) by vr_i. We have now abstracted the multihop network into a set of virtual single-hop flows, which is much closer to the scenario of a single-hop network, where time-based fairness is well defined and well understood. This corresponds to Step 1 of our methodology, as depicted in the block diagram of Figure 1.

We note that computing an optimal schedule, even for links in a single flow, is NP-hard. However, due to a flow's performance dropping off fairly rapidly with the number of links in the flow, practical deployments typically limit maximum flow length to at most 3 or 4 by design. In these situations, it is feasible to compute, using exhaustive search, the optimal schedule for the links in a single flow in the absence of inter-flow interference. This result is verified in the simulations reported in Section VI, where we compute optimal schedules for flows over a variety of network scenarios. In the unusual case of a network with long flow lengths, one could use a good heuristic scheduler to compute virtual rates. In this case, the virtual rates would be approximations instead of exact values.

B. Time-based Fairness with Virtual Links

An abstracted network made up of single-hop flows across virtual links can be used to define time-based fairness similarly to what is done in single-hop networks. To be specific, we *allocate equal times to each virtual link*. Thus, in a (virtual) sense, flows are allocated equal times. Of course, in the real multi-hop network, flow paths cross multiple links of varying intrinsic rates and, therefore, the actual times that the flows are occupying the wireless channel will be different.

Allocating resources to flows based on virtual links allows us to determine the demand on each link in the flows' paths. Ultimately, this also allows us to calculate the actual rates achieved by the flows in the network. Assume that the time

[1] We do not consider transmission power control and, therefore, a particular transmitting node always uses the same power and the SINR is determined solely by the set of simultaneous transmitters.

allocated to each virtual link is tv. Then, a flow f_k will deliver $tv{\cdot}vr_k$ bits during one scheduling period and this demand must be met on each link $l_{i,j}$ on flow f_k's path. Since flow paths can overlap, a single link $l_{i,j}$ can appear in multiple flow paths; denote them as $\mathcal{F}_{i,j}$. For each link in the network, we simply sum the demands on the link across all flows and we obtain a set of link demands across all links. Formally, the demand $d_{i,j}$ on link $l_{i,j}$ is computed as:

$$d_{i,j} = \sum_{f_k \in \mathcal{F}_{i,j}} tv \cdot vr_k = tv \sum_{f_k \in \mathcal{F}_{i,j}} vr_k \ . \tag{1}$$

Computing virtual demands for each link in the network corresponds to Step 2 of our proposed methodology in Figure 1.

Once a set of link demands is known, the links can be scheduled efficiently, using known good heuristic schedulers for the given interference model, in order to obtain a virtual-time-fair schedule. This is Step 3 of our approach in Figure 1. Note that there is no feasibility issue in this scheduling problem, because we are not scheduling based on rates but based on demands, which are given in bits. Thus, the shorter the schedule is, the higher is the overall performance that will be achieved but there is always a long-enough schedule that will allow for all demands to be satisfied.

Once the schedule is computed, links must make appropriate use of their transmission times. This is defined by the *proportional transmission time (PTT) allocation policy*, which is defined below (Figure 1 – Step 4).

Definition 1: Let $l_{i,j}$ be a link in L. The proportional transmission time policy dictates that $l_{i,j}$ must send packets of flow $f_k \in \mathcal{F}_{i,j}$ for a portion of its allocated airtime that is proportional to $\frac{vr_k}{\sum_{f_h \in \mathcal{F}_{i,j}} vr_h}$, i.e., in proportion to the virtual rate vr_k of flow f_k.

Under the above policy, a flow f_k with virtual rate vr_k delivers $tv \cdot vr_k$ bits over the course of one scheduling period. If the length of the schedule is T, then the achieved rate of flow f_k is $\frac{tv \cdot vr_k}{T}$. This rate is *virtual time fair*, in the sense that the rate achieved by f_k is proportional to its virtual rate. In other words, the share of bandwidth allocated to f_k equals:

$$bw_{k,fair} = \frac{\frac{tv \cdot vr_k}{T}}{\sum_{f_h \in \mathcal{F}} \frac{tv \cdot vr_h}{T}} = \frac{vr_k}{\sum_{f_h \in \mathcal{F}} vr_h} \ , \tag{2}$$

where $\mathcal{F}$ is the set of all flows in the network.

C. Summary of Approach

Figure 2 shows the steps of the procedure to construct a virtual-time-fair schedule for our approach.

First, the virtual rates are calculated for all flows. Then, the link demands for all links are calculated from the virtual rates. Finally, a link scheduling algorithm is used to compute a link schedule that meets all of the link demands. When the links apply the PTT allocation policy to their transmissions over the execution of the schedule, virtual-time-fair operation is achieved.

The procedure for calculating virtual rates for a flow requires efficiently scheduling the links of the flow to determine

Input: a link set L
a flow set F
a virtual link allocation time tv
Output: a virtual-time-fair schedule S and rate allocation vector R[F]

for each flow f_k
  compute vr_k using *virtual rate calculation procedure*
for each link $l_{i,j}$
  $d_{i,j} = 0$
  for each flow f_k
    if f_k travels across $l_{i,j}$ then $d_{i,j} = d_{i,j} + tv \cdot vr_k$
compute link schedule S that satisfies $d_{i,j}$ for every link $l_{i,j}$

apply PTT allocation policy on every link

Fig. 2. Steps to Compute a Virtual-Time-Fair Schedule

its achievable rate. As noted earlier, this can be done through an optimal link scheduling procedure when the number of hops in each flow is reasonably small, which is the typical case for multihop wireless networks. The simulation results in Section VI are generated using an optimal procedure that exhaustively searches for the optimal schedule. In the uncommon case where flow lengths become long, approximate virtual rates can be calculated using heuristic schedulers.

The link scheduling done in the last step is where the interference-awareness is critical. We achieve this by using physical-interference-aware scheduling to meet the specified demands. If scheduling is done using simple interference models, achieved rates can vary dramatically from actual rates [1]. This could result in a large variation from the target demands, which would cause the virtual-time-fairness condition to be violated.

Since the final link scheduling step is done over an entire multihop network, it requires a non-optimal algorithm due to the NP-hard nature of the scheduling problem. We consider two different physical-interference-aware scheduling algorithms for this step. In Section V-B, we use a physical-interference-aware link scheduler with a known approximation bound on performance to prove that we can approximately satisfy virtual time fairness while being within a constant factor of the optimal throughput. In the results of Section VI, which are intended to be reflective of typical multihop wireless network scenarios, we use the GreedyPhysical scheduling algorithm [2], which is a well-known algorithm for link scheduling under the physical interference model with good performance on realistic networks. Those results show that virtual time fairness can again be approximately satisfied while achieving much higher throughput than with max-min fairness.

D. An Example

Consider the network in Figure 3 with intrinsic link rates, in Mbps, as shown for each link. For simplicity, we assume binary interference in the example. By binary interference, we mean that interference between two links is either total, meaning the links cannot successfully transmit even at the lowest data rate, or non-existent, meaning the links can transmit without any reduction in their rates. Binary interference can be represented by an interference graph defining the pairs

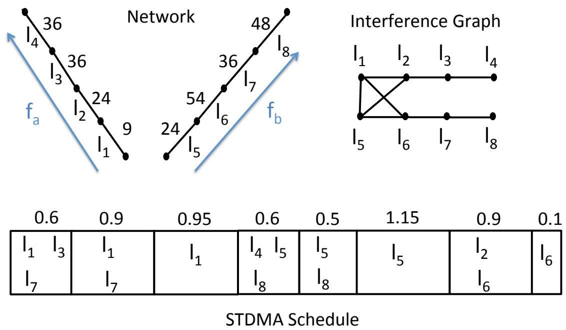

Fig. 3. Example of multihop network, corresponding interference graph, and schedule for a virtual-time-fair solution.

of links that interfere. An interference graph that could result from this network under the binary interference assumption is shown in the figure.

The virtual rates in the network of Figure 3 are 6.55 Mbps for f_a and 16 Mbps for f_b, which are obtained by optimally scheduling each flow by itself. For reference purposes, in the max-min fair solution for this network, each of the two flows achieves 5.14 Mbps for an aggregate throughput of 10.28 Mbps. One can clearly observe that the performance of the high-performing flow f_b is significantly degraded in the max-min fair solution while the low-performing flow is hardly penalized at all. The throughput-maximizing solution in this network is to let f_b have all of the resources, achieving an aggregate throughput of 16 Mbps, which is more than three times that of the max-min fair solution. Clearly, however, the throughput-maximizing solution is highly unfair, since it completely starves f_a.

To construct our virtual-time-fair solution, we begin by allocating equal time tv to the virtual links of the two flows. Thus, f_a delivers $6.54\,tv$ Mb in its allocated time and f_b delivers $16\,tv$ Mb in its time. These demands are then placed on each link in the flows' paths. Since there are no links that are present in both flows, the total demand on l_1 to l_4 is $6.54\,tv$ Mb and the total demand on l_5 to l_8 is $16\,tv$ Mb. Using these demands with $tv = 3.375$ (so that the time on l_6 is normalized to 1), we schedule the links as efficiently as possible subject to the interference graph of Figure 3. The result is the virtual-time-fair schedule shown at the bottom of the same figure. Since no link is used by more than one flow, there is no need for the PTT allocation policy in this case, so the link schedule is all that is needed.

The rate achieved by f_a in this solution is $\frac{22.1}{5.7} = 3.87$ Mbps and the rate of f_b is $\frac{54}{5.7} = 9.47$. Note that the actual rates have the same ratio (1 to 2.45) as the ratio of the virtual rates of the flows. Comparing this to the max-min fair allocation, f_a's rate is reduced by 25%, while f_b's rate is almost doubled. The aggregate throughput with the virtual-time-fair solution is 13.34 Mbps, which is 84% of the max-throughput solution but this is achieved while maintaining a fair balance between the flows.

We also use the example to point out that accounting for spatial reuse has a significant effect on the time-fair solution. According to the definition of time fairness in [3], which does not account for spatial reuse, each flow should get the same total air time. However, in this example, f_b gets more air time (4.75) than f_a (3.95) but the ratio of the two flows' final rates is the same as it would be in the definition of [3] without spatial reuse. The solutions are time fair in the sense that the flows are assigned equal air times in a solution in which they time share the network and then this solution is translated into a more efficient one with spatial reuse but maintaining the same rate ratio. Assigning equal air times in the final schedule with spatial reuse would make the ratio of the rates dependent on how the scheduling is done and, in this example, would tend to penalize f_b in favor of f_a violating the spirit of time fairness and producing lower overall throughput.

V. Analysis of Virtual Time Fairness

A. ϵ-Approximate Virtual Time Fairness

We now define a notion of ϵ-approximate virtual time fairness, where the fairness constraints are weakened to account for discretization effects due to discrete data rate values and time slot duration.

Definition 2: Let $R[F]$ be a rate vector for the flows in F as achieved by a certain schedule S of total length T combined with a link transmission time allocation policy P. We say that $R[F]$ is *ϵ-approximate virtual time fair* if and only if, for each $f_k \in F$, we have:

$$(1-\epsilon)bw_{k,fair} \leq bw_k \leq (1+\epsilon)bw_{k,fair} \ ,$$

where $bw_{k,fair}$ is the virtual time based fair bandwidth allocation for f_k as defined in (2), and bw_k is the bandwidth allocated to f_k according to $R[F]$, i.e.,

$$bw_k = \frac{\frac{tv \cdot R[f_k]}{T}}{\sum_h \frac{tv \cdot R[f_h]}{T}} \ .$$

Equivalently, we say that a (schedule, policy) pair (S, P) is ϵ-approximate virtual time fair if the resulting rate vector $R[F]$ is ϵ-approximate virtual time fair.

Definition 3: Let $R[F]$ be a rate vector for the flows in F as achieved by a certain schedule S of total length T, combined with a link transmission time allocation policy P. We say that the pair (S, P) is *ϵ-approximate virtual time throughput optimal* if and only if S provides the highest aggregate throughput among all the schedules S' such that $\exists P' : (S', P')$ is ϵ-approximate time fair.

We now show that ϵ-approximate virtual time fairness, a *flow-level* property, can be achieved if a related *link-level* property holds. We start defining the link-level property, which we call *η-approximate demand satisfaction*.

Definition 4: Let the demand $d_{i,j}$ of link $l_{i,j}$ be calculated as defined in (1). Given a schedule $S = \{T_1, T_2, \ldots, T_t\}$, consisting of t slots of equal time duration τ, let $d^h_{i,j}$ be the demand of $l_{i,j}$ satisfied in slot h ($d^h_{i,j} = 0$ if $l_{i,j}$ is not scheduled in slot h). We say that S satisfies the *η-approximate demand satisfaction* property if, for every $l_{i,j} \in L$, the demand is η-*approximately satisfied* at the end of the scheduling period $T = t \cdot \tau$. Formally:

$$\forall l_{i,j}, \quad \sum_{h=1}^{t} d_{i,j}^h \in [d_{i,j}(1-\eta), d_{i,j}(1+\eta)] \ .$$

To calculate $d_{i,j}^h$ for a given schedule, we must determine the effective rate achieved by $l_{i,j}$ in slot h, which depends on the set of transmitters scheduled in the slot. To be precise, the demand $d_{i,j}^h$ is computed as:

$$d_{i,j}^h = \tau \cdot dr_{i,j}(\mathcal{T}_h) \ ,$$

where $\mathcal{T}_h$ is the set of transmitter nodes in slot T_h.

B. Approximation Bounds

The following theorem is the main result of this section: it shows how ϵ-approximate time based fairness can be achieved by η-approximate demand satisfaction combined with an adequate link-level time allocation policy.

Theorem 1: If a schedule S satisfies the η-approximate demand satisfaction property and link transmission time in S is allocated according to the PTT allocation policy, then the flow rate vector $R[F]$ resulting from (S, PTT) is ϵ-approximate time based fair, where $\epsilon = \frac{2\eta}{1-\eta}$.

Proof: We first observe that the demand on each link is η-approximately satisfied on each link by hypothesis, and that, following the PTT policy, each link allocates its airtime according to the virtual time fair share of each flow that traverses it. Hence, denoting by vr'_k the virtual rate achieved by flow f_k in the schedule S, we have that, for any k,

$$vr'_k \in [vr_k(1-\eta), vr_k(1+\eta)] \ ,$$

where vr_k is the virtual rate of flow k. We can then upper bound the share of bandwidth achieved by f_k in S as follows:

$$\begin{aligned} bw_k &= \frac{vr'_k}{\sum_{f_h \in \mathcal{F}} vr'_h} \\ &\leq \frac{vr_k(1+\eta)}{\sum_{f_h \in \mathcal{F}} vr_h(1-\eta)} \\ &= bw_{k,fair} \cdot \left(\frac{1+\eta}{1-\eta}\right) \ . \end{aligned}$$

A lower bound on bw_k can be obtained similarly:

$$bw_{k,fair}\left(\frac{1-\eta}{1+\eta}\right) \leq bw_k$$

The proof follows by observing that $\frac{1-\eta}{1+\eta} = 1 - \frac{2\eta}{1+\eta} = 1 - \epsilon_1$, that $\frac{1+\eta}{1-\eta} = 1 + \frac{2\eta}{1-\eta} = 1 + \epsilon_2$, and that $\epsilon_2 > \epsilon_1$. ■

Theorem 1 assumes that proportional transmission time allocation as defined by the PTT policy can be implemented in a "perfect" way. In practice, making a perfectly proportional allocation of the transmission times assigned to each flow sharing a link might not be possible, due to discretization effects caused by finite duration of a transmission slot and discrete packet sizes. However, if schedule S is repeated several times during a relatively long scheduling period, such discretization effects can be substantially reduced by keeping trace of the actual amount of data transmitted by all the flows sharing the link during the scheduling period.

Theorem 2: Assume the schedule S is computed according to the algorithm INTTIMEFAIR of [1], where demands on each link are defined according to equation (1). Further, assume η is an arbitrary constant with $\frac{1}{7} \leq \eta < 1$, and that link transmission times are allocated according to the PTT policy. Then, (S, PTT) is ϵ-approximate virtual time fair with $\epsilon = \frac{2\eta}{1-\eta}$, and the aggregate throughput provided by S is within a constant factor from the optimal virtual time fair throughput.

Proof: By Theorem 1 of [1], the schedule computed by algorithm INTTIMEFAIR satisfies η-approximate demand satisfaction property, for any η with $\frac{1}{7} \leq \eta < 1$. We can then apply Theorem 1 and conclude that, by combining S with the PTT policy, we obtain a rate allocation vector that is ϵ-approximate virtual time fair, with $\epsilon = \frac{2\eta}{1-\eta}$. By Theorem 4 of [1], the aggregate throughput provided by INTTIMEFAIR is within a constant factor from the optimal time-fair throughput. Notice that the one stated in Theorem 4 of [1] is a *link-level* fairness property. However, the PTT policy ensures that link-level fairness can be translated in a flow-level fairness property, from which the theorem follows. ■

VI. SIMULATION RESULTS

To evaluate our approach, we have also performed extensive simulation experiments. In each simulation experiment, we have distributed uniformly at random a number $\sharp AP$ of Access Points (APs) and a number $\sharp Users$ of users in a square area with a minimum side length of 1.5 Km. We also constrain APs to be at distance at least 200 m from each other. To compute SNR and SINR values, we have used the log-distance radio propagation model with path-loss exponent $\alpha = 3.8$, transmission power $P = 20dBm$, and noise $N = -80dBm$. We have also constrained link length to be at most 180 m, to prevent the use of very weak wireless links which could compromise overall performance.

Once users and APs have been distributed, each user is connected to the closest AP in terms of number of hops, with ties broken randomly. For each user, a downlink connection to the assigned AP is then established with probability $p = 0.9$, and an uplink connection otherwise (with probability $1 - p$).

For each setting of the simulation parameters, 100 *valid* random network deployments are generated and used to produce statistically significant simulation results, where a deployment is valid if and only if all users can be connected to an AP.

We compare our approach to two alternate approaches: $TDMA$, corresponding to a solution in which a single flow is active in the network at one time, but links within a flow can still exploit spatial reuse; and *equal rate*, corresponding to a solution in which the virtual rate of each flow is set to the same, fixed value. By scheduling equal virtual rates efficiently, we can approximate the maximum rate achievable with equal rate allocation, which approximates max-min fairness, because under quite general conditions, as discussed in [11], max-min fairness produces an equal rate allocation with maximum rate.

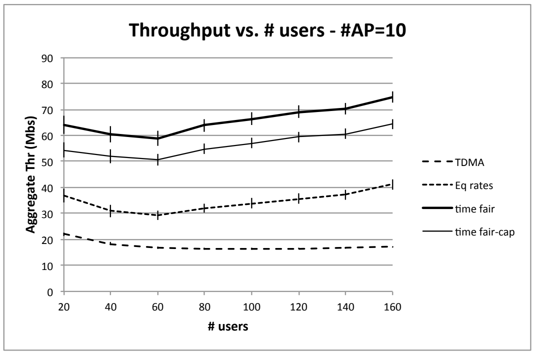

Fig. 4. Aggregate throughput as a function of the number of users. The number of APs is set to 10. Markers are shown as 95% confidence intervals.

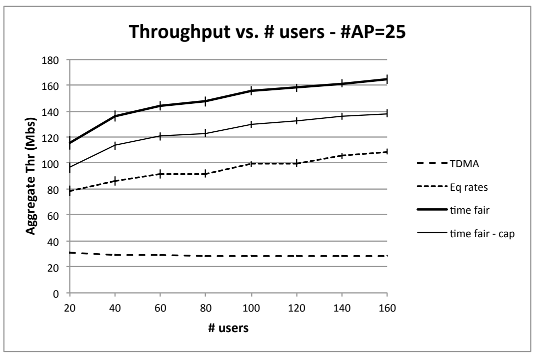

Fig. 5. Aggregate throughput as a function of the number of users. The number of APs is set to 25. Markers are shown as 95% confidence intervals.

We want to stress that our approach is very efficient in terms of running time, since computation of virtual rates and subsequent scheduling step can be executed in fraction of a second even for networks composed of 200 users.

A. Increasing density

In a first set of experiments, we have kept the simulation area fixed to 1.5×1.5 Km, and increased the number of users from 20 to 160. The results of the experiments are reported in Figure 4 when the number of APs is set to 10, and in Figure 5 when the number of APs is set to 25. As seen from the figures, our approach substantially increases throughput with respect to the TDMA solution, as well as to the equal rate allocation solution. Notably, the total aggregate throughput of our virtual time fair approach is increased by from 50% to 100% with respect to equal rate allocation.

B. Fixed density

In a second set of experiments, we have increased the number of users or APs, while leaving the user or AP density fixed. This is achieved by properly scaling the network area. In a first scenario we have fixed the number of users to 200, set the side of the simulation area to $375 \cdot \sqrt{\sharp AP}\ m$, and varied $\sharp AP$ from 10 to 30. The results are reported in Figure 6. It is interesting to observe that our virtual time-based fair approach is able to take full advantage of the increased number of APs, showing a linear increase of aggregate throughput as $\sharp AP$ increases. Also the equal rate solution shows an increasing trend of throughput vs $\sharp AP$, but with a much slower slope with respect to the virtual time fair approach.

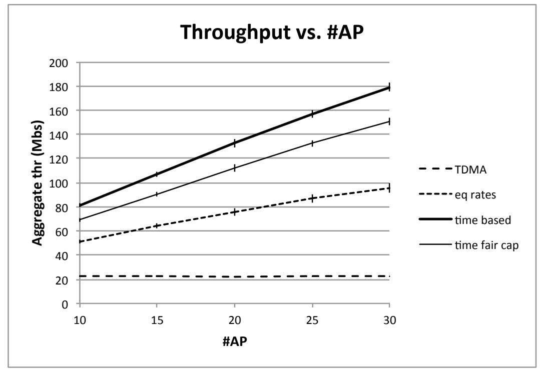

Fig. 6. Aggregate throughput as a function of the number of APs. The number of users is set to 200. Markers are shown as 95% confidence intervals.

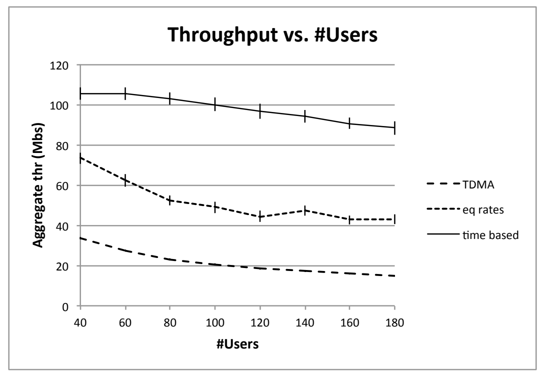

Fig. 7. Aggregate throughput as a function of the number of users. The number of APs is set to 15. Markers are shown as 95% confidence intervals.

In a second scenario, we have set the number of APs to 15, set the side of the simulation to $150 \cdot \sqrt{\sharp Users}\ m$, and varied $\sharp Users$ from 40 to 180. The results are reported in Figure 7. All the approaches show a decreasing trend of the aggregate throughput vs. $\sharp Users$. This is due to the fact that, given a fixed number of APs, a larger number of users have to compete for the same resources (connection to an AP), and this comes at the expense of overall performance. In case of increased user density per unit area (Figures 4 and 5), this negative effect is more than compensated by the fact that average link length decreases with $\sharp Users$, which implies higher data rates. In the case of fixed density shown in Figure 7, average link length does not change, and the negative effect of increased competition cannot be compensated. It is notable that in this case the virtual time-based solution still shows the best results, displaying a relatively smaller decrease of aggregate throughput for increasing number of users as compared to the other approaches.

C. Fairness evaluation

We have also evaluated the fairness of our approach vs. the exact virtual time based fair allocation, using the fairness index suggested in [1],

$$FI = \frac{1}{e^{\frac{1}{n} \cdot \sum_k \left| \ln \frac{bw_{k,fair}}{bw_i} \right|}} \ ,$$

where n is the total number of flows, $bw_{k,fair}$ is the virtual time-fair bandwidth share of flow i, and bw_i is the actual bandwidth share of the flow as resulting from the scheduling algorithm and link transmission policy.

The index takes values in $(0, 1]$, with 0 representing complete lack of fairness and 1 exact virtual time based fairness. In all simulation experiments, our solution showed consistently high fairness index values of 0.94 and above. This indicates that our proposed approach does indeed provide a close approximation to virtual time fairness. As an example, the values of the fairness index corresponding to the scenario with 10 APs and increasing density are reported in Table II.

To further quantify fairness, we have also evaluated how much the rate of the flow f_{min} with the lowest virtual rate in the time-based fair solution is reduced with respect to the equal rate allocation. The results (averaged over 100 random deployments) for a sample configuration with 100 users and 10 APs show that the rate of f_{min} can be as low as 20% and as high as 70% of the equal rate, with an average of about 40%.

$\sharp Users$	FI	$\sharp Users$	FI
20	0.96	100	0.964
40	0.961	120	0.964
60	0.963	140	0.964
80	0.964	160	0.964

TABLE II
FAIRNESS INDEX OF OUR TIME-BASED FAIR APPROACH IN THE SCENARIO WITH 10APS, AND INCREASING NUMBER OF USERS.

VII. SINGLE-HOP VS. MULTI-HOP FLOWS

Prior work, e.g. [3], [6], has pointed out that, in some settings, multi-hop flows can experience very poor performance as compared to single-hop flows. This is usually attributed to multi-hop flows having to contend for the channel multiple times, which is an issue that is specific to CSMA/CA settings. To address this issue, several approaches have been proposed to boost performance of multi-hop flows by allocating more resources to them. In this section, we evaluate the performances of single-hop and multi-hop flows within our virtual-time-fair framework.

A. Single-hop and Multi-Hop Flows with Basic Virtual Time Fairness

Since we consider STDMA settings, the contention issue is not a factor. Nevertheless, it is still important to understand the relative performances of single-hop and multi-hop flows in these settings. We begin by pointing out that boosting performance of multi-hop flows, and therefore reducing performance of single-hop flows as a consequence, is counter to the basic concept of time-based fairness we are proposing. In our concept of time-based fairness, flows have some intrinsic performance, i.e. their maximum achievable rate with no interfering flows. Allocating equal (virtual) times to each flow means that the performance they ultimately attain is in the same ratio to other flows' performances as the ratios of their intrinsic rates. Artificially boosting low-rate flows would penalize higher-performing ones and tends toward max-min fairness and equalizing performance. This would be similar to rate-based fairness in the single-hop wireless network case, where the performance of a user with a high-rate link is reduced by the presence of a low-rate user. Nevertheless, we would like to compare single-hop vs. multi-hop performance to see if there is something that could be perceived as bias against multi-hop flows in our approach and, if so, whether the issue can be addressed without deviating too far from our basic conceptual approach.

As a starting point, we consider some simple examples to illustrate that, with multirate links, the performance of flows in our approach is determined not only by the number of hops, but also by the quality of the individual links. Compare, for example, a single-hop flow across a 6 Mbps link and a two-hop flow across two 54 Mbps links. The virtual rate of the single hop flow remains 6 Mbps, while the two-hop flow's virtual rate is 27 Mbps. A three-hop flow with all 54 Mbps links could also have a virtual rate of close to 27 Mbps depending on the level of interference between the first and third hops. The lowest its virtual rate could be is 18 Mbps, which is still 3 times the rate of a single-hop flow with a 6 Mbps link. Thus, in our setting, multi-hop flows with high-quality links will outperform single-hop flows with low-quality links and hence, there is no perfect correlation between the number of hops and performance. On the other hand, the best-performing flows in our approach will be single-hop flows with high-quality links, which could have virtual rates as high as 54 Mbps, which is twice the maximum value that can be achieved by a multi-hop flow.

B. Analysis of Single-hop and Multi-hop Flow Performance

To better understand the relationship of performance vs. number of hops, we look at data from the simulations reported in Section VI. We have considered two reference scenarios, both with 10 APs, but with different number of users: 100 (Scenario 1) and 160 (Scenario 2), respectively. For each scenario, we have generated 100 random topologies as described in Section VI, and collected statistics of the resulting flows. Since we have a flow associated with each user in the network, the total number of considered flows in the two scenarios was 10,000 and 16,000, respectively. We have analyzed 1-, 2-, and 3-hops flows, which are the overwhelming majority of flows in both scenarios: 82% and 87% of the total flows in Scenarios 1 and 2, respectively.

The results of the evaluation are summarized in Table III. As seen from the table, in both scenarios it is possible to have 2-hop flows with higher virtual rates than 1-hop flows, and 3-hop flows with higher virtual rates than 2-hop flows. This actually occurs for about 2% of the 2-hop flows, and for about 15% of the 3-hop flows.

Based on the results of our evaluation, we can conclude that, although possible, the event of a longer flow having a higher virtual rate than a shorter flow is fairly unlikely. In particular, 1-hop flows are quite likely to have the highest virtual rates in the network. If this is perceived as an unjustified bias towards single-hop flows by the network designer, it is possible to lessen the problem by applying a simple modification to our approach. The idea is to introduce an upper bound u to the

$\sharp$ users	$\sharp$1-hop flows	(max, min) rate	$\sharp$2-hop flows	(max, min) rate	$>$ 1-hop	$\sharp$ 3-hop flow	(max, min) rate	$>$ 2-hop
100	3,669	(54,18)	2,668	(21.6,9)	59	1713	(13.5,6)	267
160	6,109	(54,18)	4,916	(21.6,9)	95	2802	(13.5,6)	390

TABLE III

EVALUATION OF PERFORMANCE VS. NUMBER OF HOPS. COLUMN "$>$ h-HOP" REPORTS THE NUMBER OF $(h+1)$-HOP FLOWS WITH A VIRTUAL RATE HIGHER THAN A h-HOP FLOW.

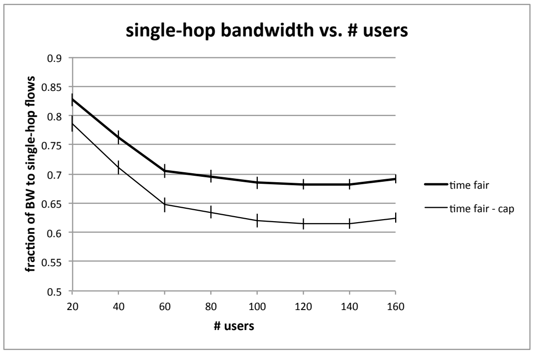

Fig. 8. Fraction of total bandwidth allocated to single-hop flows. The number of APs is set to 10. Markers are shown as 95% confidence intervals.

virtual rate of single-hop flows, and to use u instead of the actual virtual rate vr of a single-hop flow whenever $vr > u$.

The results of this modification to our basic strategy, with u set to 27 $Mbps$[2], are reported in Figures 4–7 with the curve designated by "time fair - cap". As expected, the total aggregate throughput is reduced with respect to the virtual time fair solution by about 10%–15%, but the throughput remains substantially higher than the equal rate solution. While the aggregate throughput is reduced compared to the uncapped solution, the amount of bandwidth assigned to multihop flows is substantially increased – see Figure 8.

VIII. CONCLUSION

In this paper, we have presented a near-optimal, constructive approach to achieve a general notion of time-based fairness for multi-hop wireless networks. Results of simulation experiments have shown that time-based fairness achieves a substantially higher aggregate throughput than that of an approach approximating max-min fairness, highlighting for the first time the significant benefits of time-based fairness in the context of multihop wireless networks. The notions of virtual link and virtual rate presented herein allow the network designer to implement a variety of fairness policies by acting on virtual demands. For example, equal virtual demands produce a rate allocation approximating max-min fairness, and virtual demands proportional to cumulative virtual rates on the link produce a time-based fair allocation. We believe the herein presented approach presented of abstracting flow performance by means of virtual links/rates might be generalized to other resource allocation problems for multihop wireless networks.

REFERENCES

[1] D. Blough, G. Resta, P. Santi, "Interference-aware Proportional Fairness for Multi-Rate Wireless Networks", *Proceedings of IEEE Infocom,* pp. 2733–2741, 2014.

[2] G. Brar, D. Blough, and P. Santi, "Computationally Efficient Scheduling with the Physical Interference Model for Throughput Improvement in Wireless Mesh Networks," *Proc. ACM Mobicom* pp. 2–13, 2006.

[3] V. Gambiroza, B. Sadeghi, E.W. Knightly, "End-to-End Performance and Fairness in Multihop Wireless Backhaul Networks", *Proc. ACM Mobicom*, pp. 287–201, 2004.

[4] M. Gao and H. Jiang, "Throughput Performance and Fairness in Multihop Wireless Mesh Networks," *Proc. Symp. on Wireless Networking and Multimedia,* pp. 867–872, 2014.

[5] M. Heusse, F. Rousseau, G. Berger-Sabbatel, A. Duda, "Performance anomaly of 802.11b", *Proc. IEEE INFOCOM*, 2003.

[6] J. Hoblos, "Fairness Enhancement using Transmission Scheduling in Multi-hop Wireless Mesh Networks", *Proc. IEEE Int. Conference on Communications and Information Technology*, pp. 251–255, 2013.

[7] T.-S. Kim, H. Lim, J.C. Hou, "Improving Spatial Reuse Through Tuning Transmit Power, Carrier Sense Threshold, and Data Rate in Multi-Hop Wireless Networks", *Proc. ACM Mobicom*, pp. 366-377, 2006.

[8] T. Li, D.J. Leith, V. Badarla, D. Malone, Q. Cao, "Achieving End-to-end Fairness in 802.11e Based Wireless Multi-hop Mesh Networks Without Coordination", *Mobile Networks and Applications*, Vol. 16, pp. 17–34, 2011.

[9] R. Maheshwari, S. Jain, S.R. Das, "A Measurement Study of Interference Modeling and Scheduling in Low-Power Wireless Networks", *Proc. 6th ACM SenSys*, pp. 141–154, 2008.

[10] R. Maheshwari, J. Cao, S.R. Das, "Physical Interference Modeling for Transmission Scheduling on Commodity WiFi Hardware", *Proc. IEEE Infocom*, pp. 2661–2665, 2009.

[11] B. Radunovic and J. Le Boudec, "Rate performance objectives of multi-hop wireless networks", *IEEE Transactions on Mobile Computing*, Vol. 3, No. 4, pp. 334–349, 2004.

[12] A. Raniwala, P. De, S. Sharma, R. Krishnan, and Tzi-cker Chiueh, End-to-End Flow Fairness over IEEE 802.11-based Wireless Mesh Networks, *Proc. of IEEE INFOCOM Mini-Symposium,* pp. 2361–2365, 2007.

[13] P. Santi, R. Maheshwari, G. Resta, S. Das, D. Blough, "Wireless Link Scheduling under a Graded SINR Interference Model", *Proc. ACM FOWANC*, pp. 3–12, 2009.

[14] G. Tan, J. Guttag, "Time-based Fairness Improves Performance in Multi-Rate WLANs", *Proc. USENIX Conf.*, pp. 23–36, 2004.

[15] P. Thulasiraman, J. Chen, and X. Shen, "Multipath Routing and Max-Min Fair QoS Provisioning under Interference Constraints in Wireless Multihop Networks," *IEEE Transactions on Parallel and Distributed Systems,* Vol. 22, pp. 716–728, May 2011.

[16] I. Tinnirello and S. Choi, Temporal Fairness Provisioning in Multi-Rate Contention-Based 802.11e WLANs, *Proc. of IEEE WOWMOM,* pp. 220–230, 2005.

[17] P. Wang, H. Jiang, W. Zhuang, and V. Poor, "Redefinition of Max-Min Fairness in Multi-Hop Wireless Networks," *IEEE Transactions on Wireless Communications,* Vol. 7, pp. 4786–4791, Dec. 2008.

[18] L. Zhang, W. Luo, S. Chen, Y. Jian, "End-to-end Maxmin Fairness in Multihop Wireless Networks: Theory and Protocol", *Journal of Parallel and Distributed Computing*, Vol. 72, n. 3, pp. 462–474, 2012.

[2]This caps the virtual rate of single-hop flows to the maximum possible virtual rate of a multi-hop flow.