

Elastic Rate Limiting for Spatially Biased Wireless Mesh Networks

Vincenzo Mancuso, Omer Gurewitz, Ahmed Khattab, Edward Knightly

► **To cite this version:**

Vincenzo Mancuso, Omer Gurewitz, Ahmed Khattab, Edward Knightly. Elastic Rate Limiting for Spatially Biased Wireless Mesh Networks. INFOCOM 2010, Mar 2010, San Diego, CA, United States. 2010. <inria-00506753>

HAL Id: inria-00506753

<https://hal.inria.fr/inria-00506753>

Submitted on 28 Jul 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Elastic Rate Limiting for Spatially Biased Wireless Mesh Networks

Vincenzo Mancuso, Omer Gurewitz, Ahmed Khattab and Edward W. Knightly

Abstract—IEEE 802.11-based mesh networks can yield a throughput distribution among nodes that is spatially biased, with traffic originating from nodes that directly communicate with the gateway obtaining higher throughput than all other upstream traffic. In particular, if single-hop nodes fully utilize the gateway’s resources, all other nodes communicating with the same gateway will attain very little (if any) throughput.

In this paper, we show that it is sufficient to rate limit the single-hop nodes in order to give transmission opportunities to all other nodes. Based on this observation, we develop a new rate limiting scheme for 802.11 mesh networks, which counters the spatial bias effect and does not require, in principle, any control overhead.

Our rate control mechanism is based on three key techniques. First, we exploit the system’s inherent priority nature and control the throughput of the spatially disadvantaged nodes by only controlling the transmission rate of the spatially advantaged nodes. Namely, the single-hop nodes collectively behave as a proxy controller for multi-hop nodes in order to achieve the desired bandwidth distribution. Second, we devise a rate limiting scheme that enforces a utilization threshold for advantaged single-hop traffic and guarantees a small portion of the gateway resources for the disadvantaged multi-hop traffic. We infer demand for multi-hop flow bandwidth whenever gateway resource usage exceeds this threshold, and subsequently reduce the rates of the spatially advantaged single-hop nodes. Third, since the more bandwidth the spatially disadvantaged nodes attain, the easier they can *signal* their demands, we allow the bandwidth unavailable for the advantaged nodes to be elastic, i.e., the more the disadvantaged flows use the gateway resources, the higher the utilization threshold is.

We develop an analytical model to study a system characterized by such priority, dynamic utilization thresholds, and control by proxy. Moreover, we use simulations to evaluate the proposed elastic rate limiting technique.

I. INTRODUCTION

A key challenge for rate control in IEEE 802.11 mesh networks is that the total traffic originating from any node that is directly connected to a gateway forms a high-priority aggregate flow when competing with other traffic. That is, for gateway-destined traffic, traffic originating from multi-hop nodes is disadvantaged and can only use resources not demanded by advantaged gateway-connected nodes.

In this paper, we develop a new elastic rate control method for mesh networks with three key techniques. First, we *exploit* and indeed require the distributed priority property that derives

from MAC layer behavior. Because the low priority traffic obtains only resources not utilized by the high priority traffic, a rate control system only needs to control the high priority traffic in order to control the throughput of both classes. Thus, we design a rate control method in which the single-hop nodes collectively behave as a proxy controller for multi-hop nodes: single-hop (advantaged) nodes allow multi-hop (disadvantaged) nodes to increase their rate by decreasing their own aggregate rate. Likewise, single-hop nodes can force multi-hop nodes to decrease their rate by increasing their own rate.

Second, we address a critical problem in the priority system: how can single-hop nodes allocate resources on behalf of multi-hop nodes without knowing their demands? That is, while a particular single-hop node is in range of a small number of multi-hop nodes, it is out of range of many others; without knowing if multi-hop nodes are even backlogged, how can a single-hop node make rate control decisions on their behalf? Moreover, multi-hop nodes cannot simply make use of *additive increase* when their demand increases, as they can only increase their throughput if single-hop nodes decrease their aggregate rate. To tackle these issues, we propose a distributed mechanism that operates a Gateway Airtime Partitioning (GAP), in which a small portion of the gateway airtime is reserved for multi-hop nodes to signal their demand (*signaling bandwidth*), whilst all nodes have open access to the remaining gateway airtime.

Third, the amount of signaling bandwidth determines the capability of multi-hop nodes to signal their demand and the time needed for the rate control mechanism to converge to a new stable state after the traffic demand changes. However, the higher the multi-hop traffic, the lower the additional demand that multi-hop nodes are expected/allowed to request, and, in turn, the less signaling bandwidth is needed. Hence, when multi-hop traffic is high, we allow the signaling bandwidth to be used for multi-hop data traffic, i.e., we elastically reduce the signaling bandwidth as a function of the multi-hop offered load.

Thus, we design our rate control algorithm as follows: (i) single-hop nodes are aware of the status of the gateway airtime utilization, i.e., they can estimate or get notified whether the gateway utilization is above a predefined threshold; (ii) single-hop nodes agree that the aggregate traffic originating from all single-hop nodes should never exceed a particular utilization threshold; (iii) thus, if the utilization threshold is exceeded, it is due to multi-hop traffic; (iv) we interpret crossing the utilization threshold as a request for a multi-hop flow to be in the *additive increase* mode.

With steps (i)-(iv), each (advantaged) single-hop node has a means to infer whether the (disadvantaged) multi-hop nodes

V. Mancuso is with Équipe-projet MAESTRO at INRIA Sophia Antipolis Méditerranée, France. O. Gurewitz is with Department of Communication Systems Engineering, Ben Gurion University of the Negev, Beer Sheva, Israel. A. Khattab is with Department of Electrical and Computer Engineering, University of Louisiana at Lafayette, LA. E. Knightly is with Department of Electrical and Computer Engineering at Rice University, Houston, TX. All authors were with Rice when the research for this work has been started.

This research was supported by NSF grants CNS-0751173 and CNS-08102501 and by the Cisco Collaborative Research Initiative.

want to increase their rates. With this view of the state of both classes of nodes, single-hop nodes can set their own rates according to rate control objectives, and, because of the priority nature of the system, they have also set the rates of the multi-hop flows.

We develop a simple fluid model to isolate and study the aforementioned system properties of priority, utilization thresholds, and control by proxy. We use the model both as a means to precisely define these properties and to show that our rate control scheme is able to converge to a desired operating point in this simplified priority system.

Finally, we present an extensive set of *ns-2* simulation experiments to evaluate the performance of the proposed scheme in more realistic settings than the fluid model. Our results show that our approach achieves targeted rates similar to those of a scheme that has perfect knowledge of all offered loads and rate limits flows according to their ideal rates. Moreover, we show that the very limited computational overhead of our approach yields significant gains compared to techniques that use rate control messages such as the IEEE 802.11s [1].

The rest of the paper is organized as follows. In Sections II and III we define the network model and the key system's properties, respectively. We propose our rate control framework, GAP, in Section IV. We model the system's properties and analyze our rate control scheme in Section V. In Section VI, we evaluate the performance of the proposed scheme. We overview related work in Section VII, and conclude the paper in Section VIII.

II. NETWORK MODEL

We consider non-mobile multi-hop wireless networks that forward traffic to and from gateway nodes. The physical topology of the mesh networks does not need to show any particular structure. In contrast, the *forwarding topology* yields a set of trees rooted in gateway nodes. Thus, we consider networks with unconstrained physical topology, which embed a forwarding tree of degree $N \geq 1$ and depth $D \geq 2$ per gateway. In particular, we focus on a single-gateway mesh, in which a routing protocol establishes a set of forwarding links yielding a tree structure. Note that since the gateway capacity is of the order of few tens of Mbps, due to the 802.11 maximum transmission rate, typical values of D are 2 to 3, and $N \leq 10$ to ensure sufficient per-node resources.

For a particular tree, mesh nodes other than the gateway node can be classified into two sets: a set \mathcal{S} of nodes that can directly connect to the gateway, and a set \mathcal{M} which includes the remaining nodes. Fig. 1 depicts an example of data forwarding tree network topology. Note that the physical interconnection is not a tree in the example.

For the sake of simplicity we consider wireless nodes equipped with a single half-duplex radio. Nodes in the two sets \mathcal{S} and \mathcal{M} compete for gateway access through a shared wireless interface. However, our approach can be easily extended to multiple-radio nodes. In fact, it would suffice to logically decompose the network into groups of wireless transceivers, each group using a single 802.11 channel, and apply our analysis to each group separately.

We assume that no peer-to-peer traffic is allowed within the mesh, i.e., all non-gateway nodes can open connections

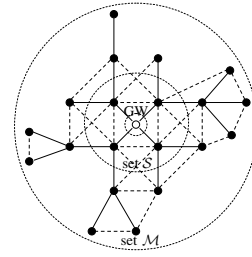


Fig. 1. A possible network topology. Solid and dotted lines represent data-forwarding and non-data-forwarding links, respectively. Data-forwarding links form a tree which is rooted in the gateway node.

with the gateway or with remote Internet nodes only. Since we are only interested in those performance factors that are originating in the mesh network, we do not consider the effects of the connection path between the gateway and any other remote Internet node. Hence, for the sake of simplicity, we model all downstream flows as originating at the gateway, and all upstream flows as terminating at the gateway node. However, in the rest of the paper we focus on the performance of upstream flows only. This assumption is motivated by the fact that previous work already showed that downstream traffic impairments do not arise in case of UDP, and TCP performs very similarly both upstream and downstream [2].

Hence we look at the upstream gateway capacity as a fixed bandwidth and define the upstream gateway airtime utilization, $U(t)$, as the amount of time the gateway transceiver is active or deferring due to other data transmissions within its carrier sense range. The time the gateway spends in idle state or counting down its backoff counter represents unutilized gateway airtime. All nodes are assumed to be using the same rate for data packet transmissions, thus gateway airtime utilization is equivalent to the consumed bandwidth.

Finally, we refer to flow i as the aggregate traffic originating from node i (i.e., flow i includes all upstream flows originated by node i , but does not include traffic forwarded on behalf of other nodes). We denote by $R_i(t)$ the rate at which node i throttles flow i at time t .

III. CHARACTERIZING 802.11 FORWARDING TREES

In this section, we identify three fundamental properties of 802.11 mesh networks that affect our rate control design.

A. Distributed Approximate Priority Property

A number of studies have described the uneven distribution of bandwidth among flows in multi-hop wireless networks (e.g., [3], [4], [5]). Specifically, due to spatial bias some flows can capture more resources than others. In wireless mesh networks, the spatial orientation of flows going to and from the gateway can further exacerbate severe throughput imbalances, giving higher bandwidth to nodes that are closer to the gateway. In particular, the most advantaged nodes are the ones directly connected to the gateway (one hop away from it), hence whenever nodes in the single-hop and multi-hop node sets compete for accessing the tagged gateway interface, the distributed traffic originating at all nodes in the single-hop set \mathcal{S} is of almost-strict priority in accessing the gateway

resources compared to the traffic originating at nodes in \mathcal{M} . Thus, the traffic originating at any member of the multi-hop set \mathcal{M} can only access the residual gateway resources not used by all nodes of the single-hop set \mathcal{S} . Hence, if the combined demands of all nodes in the single-hop node set \mathcal{S} consume all of the available gateway resources, the traffic originating at any multi-hop node in \mathcal{M} will not be able to access the gateway resources. We refer to such a system as a *distributed approximate priority system* wherein the distributed load of the sets \mathcal{S} and \mathcal{M} are the high priority and the low priority traffic entities, respectively.

In order to demonstrate the *distributed approximate priority system* we performed an experiment using mesh nodes running LocustWorld open-source mesh networking software. Each mesh node is equipped with a single SMC 2532-b IEEE 802.11b wireless interface operating at the 2.4 GHz ISM band with maximum transmission power of 24 dBm. The physical layer rate is set to 11 Mbps. We use such nodes to realize the network topology shown in Fig. 2(a) in an indoor environment. We use *iperf* to generate traffic at source nodes. The duration of each experiment run is 120 seconds. In the experiment we gradually increase the offered load of node S_2 from 128 Kbps up to 6 Mbps in 512 Kbps steps. The multi-hop node M_1 carries backlogged traffic. Node S_1 does not offer any traffic and just forwards node M_1 's traffic.

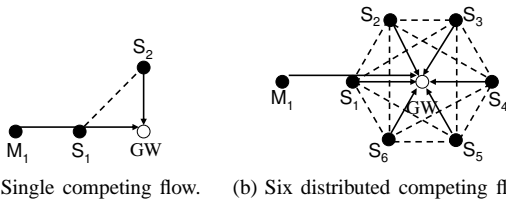


Fig. 2. Examples of topologies of a two-hop node competing with single-hop node(s). Dotted lines indicate that connected nodes are within sensing range.

Fig. 3 depicts the throughput of each flow as well as the gateway utilization versus the single-hop node's offered load, both for UDP and TCP traffic. The single-hop node's received throughput is almost equal to its demands regardless of whether the two-hop node's offered load is backlogged. Thus, the single-hop node's traffic has almost-strict priority in accessing the gateway resources (a perfect strict priority system would yield the $Y = X$ dashed line until saturation). Meanwhile, the two-hop node's traffic can only partially utilize the gateway airtime that is not utilized by the single-hop node's traffic.¹ When both flows are continuously backlogged, the single-hop traffic utilizes almost all of the gateway airtime, thereby starving the two-hop node. When we added a three-hop node to the two-hop branch (not shown here), we observed similar behavior. Note that the transmissions of the two-hop nodes in \mathcal{M} include both the locally generated and forwarded traffic. Thus, when the two-hop nodes are of approximate low priority in accessing the gateway, all the multi-hop traffic is of approximate low priority when competing with single-hop traffic. We repeat similar experiments using the *ns-2* simulator with the default 802.11 parameters on a broader set of topologies such as the topology depicted in Fig. 2(b). The

¹Gateway utilization drops when both flows share the gateway due to time spent in contention, collisions, and collision resolution.

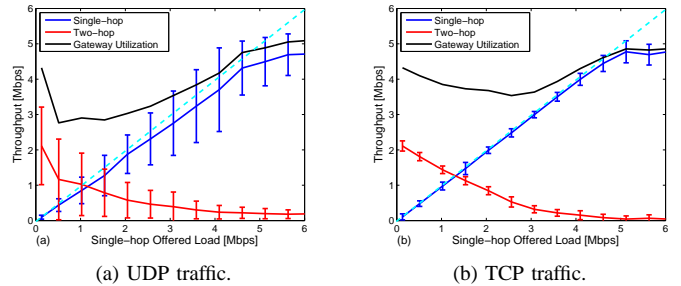


Fig. 3. Experimental demonstration of 802.11 approximate priority property.

results show similar trend as the one presented in Fig. 3, i.e., the existence of almost-strict priority in accessing the gateway airtime for single-hop traffic, whenever competing with a two-hop node. These results are presented in [6].

The two other properties which inspire our flexible rate control mechanism, relate to congestion around the gateway.

B. Gateway Airtime Saturation Property

The second property asserts that, assuming all links transmit at the same rate, a congested link in a single gateway topology (possibly a few hops away from the gateway) indicates congestion around the gateway (e.g., gateway airtime must be saturated). The reasoning is that since we assume that all traffic is going to and from the gateway, the congested link traffic also originates or terminates in the gateway. Hence, either the traffic traversing the congested link propagates all the way to the gateway, which implies congestion around the gateway, or the traffic cannot propagate all the way to the gateway due to some other congested region closer to the gateway. Extending the same argument will let us conclude that the congestion is localized in the vicinity of the gateway.

C. Proxy Control Property

The third property, which relates to the first two, asserts that enforcing free airtime in the extended gateway neighborhood will give nodes in set \mathcal{M} some transmission opportunities. Therefore, the single-hop nodes can act as a distributed proxy which controls the transmission opportunities of multi-hop nodes.

The reasoning for this property is based on the spatial reuse property in multi-hop topologies, i.e., nodes two and three hops away from the gateway can utilize the free gateway airtime to get some channel access, given that the single-hop nodes left some gateway airtime unused. Nodes four hops away from gateway can transmit simultaneously with single-hop nodes, which force two-hop nodes to backoff, and, in turn, provide four-hop nodes with transmission opportunities. This spatial reuse argument can be extended to nodes which are further away from the gateway. Obviously, since hidden nodes cannot coordinate their channel access, distance from the gateway is in inverse proportion to successful transmission opportunities.

IV. GAP: GATEWAY AIRTIME PARTITIONING THROUGH ELASTIC RATE LIMITING

We now introduce an elastic rate limiting strategy for single-hop nodes, that exploits the network characteristics presented

in Section III. The proposed strategy is named *Gateway Airtime Partitioning* (GAP) since it yields an airtime partition between gateway resources that can always be used by any node and resources that should be reserved for multi-hop demand indication.

A. GAP Framework

According to their unique system characteristics identified in Section III, we pose the rate control problem in 802.11 multi-hop trees as by which means to achieve an arbitrary bandwidth distribution between the two spatially biased node sets \mathcal{S} and \mathcal{M} , given their time-varying demand. GAP rate control objective is to ensure minimum rates that would be guaranteed under saturation load conditions, and to fairly share unused resources among all competing nodes. The distinct nature of our rate control problem comes from the approximate priority property of the spatially advantaged node set \mathcal{S} that allows its distributed demands to be almost fully served, irrespective of the demands of the spatially disadvantaged node set \mathcal{M} . The demands of the disadvantaged multi-hop nodes can only utilize the gateway airtime that is not used by the advantaged node set \mathcal{S} .

Our new rate control approach is to have the set of single-hop nodes collectively behave as a proxy controller for multi-hop nodes. Thus, the main challenge in GAP is how can single-hop nodes allocate resources on behalf of multi-hop nodes without knowing their demands. Nodes in \mathcal{S} should neither leave gateway resources more than the demands of nodes in \mathcal{M} nor less than \mathcal{M} 's minimum guaranteed rate. Moreover, nodes in \mathcal{S} can neither infer the demands nor observe the throughput of all nodes in \mathcal{M} due to the distributed nature of the network, and since we do not flood the network with per-node control messages.

Our GAP framework is based on two fundamental ideas. First, we have all nodes in the spatially advantaged node set agreeing that the combined gateway airtime utilization of their locally generated traffic, not including forwarded traffic, is limited to a particular threshold, rather than the entire gateway airtime. Consequently, spatially disadvantaged nodes' traffic can always use the residual gateway airtime for successful data transmissions (either from a two-hop node in \mathcal{M} to a node in \mathcal{S} , or a node in \mathcal{S} forwarding a multi-hop packet to the gateway). Since all nodes in \mathcal{S} are aware that their traffic's should not exceed a predefined threshold, nodes in \mathcal{S} interpret the excess gateway airtime utilization as only due to transmissions of multi-hop traffic. Hence, the only piece of information that nodes in \mathcal{S} need is whether the gateway airtime utilization exceeded the threshold. This information might be encoded into a 1-bit message that is *high* when the traffic threshold is exceeded, and otherwise is *low*. This message might be sent from the gateway to all nodes in \mathcal{S} , e.g., by encoding the bit into a currently unused *subtype value* in the Frame Control field of the 802.11 ACK, or by including a traffic indicator field in beacons regularly transmitted by the gateway, or also by allowing the gateway to transmit a new type of management frame, newly defined for traffic indication. Another way to obtain the same information on the gateway airtime utilization might consist in letting single-hop

nodes estimate the gateway activity by overhearing gateway's ACKs. This approach is in principle possible at single-hop nodes, but is particularly prone to estimation errors due to frame collisions and failures in decoding some ACKs, e.g., due to variations in the SNR, not all single-hop nodes might be able to decode gateway's ACKs transmitted at the highest modulation rate. In both ways, explicitly or implicitly, each single-hop node obtains a **gateway utilization indicator** I_U , whose binary value is *high* if the bandwidth reserved for multi-hop traffic is in use.

We define the disadvantaged-flow **Signaling Bandwidth**, $B_D = \gamma U^*$, $\gamma \ll 1$, as a small portion of the system's resources that nodes in \mathcal{S} collaboratively agree not to use for the transmission of their locally generated traffic to the gateway. Instead, this bandwidth will be used exclusively by the set of spatially disadvantaged flows. Any flow originating at a node in \mathcal{M} will use this bandwidth to transmit data, thereby, expressing that its current demands would like more bandwidth, if possible, to the spatially advantaged node set \mathcal{S} . Consequently, the distributed single-hop proxy controller will have all nodes in \mathcal{S} collaboratively adapting their rates to realize the rate control objective.

The second part of the GAP framework is to prevent the spatially disadvantaged flows from misusing B_D to get more than their minimum guaranteed rate if backlogged. In fact, if we would allow the single-hop nodes to unboundedly reduce their rates as long as B_D is partially or fully utilized, backlogged multi-hop nodes could use B_D to exclusively capture the system's resources, irrespective of the system's approximate priority behavior. Thus, we design GAP to not allow nodes in \mathcal{S} to reduce their rates if *their* gateway utilization does not exceed a minimum guaranteed rate.

The **Minimum Guaranteed Single-hop Rate**, U_S^* , is defined as the minimum bandwidth to be guaranteed for the node set \mathcal{S} under saturation load conditions achieving a specific distribution of the gateway airtime (e.g., max-min fairness, proportional fairness, etc.). U_S^* is the sum of the minimum guaranteed rates, F_i , of single-hop nodes. Likewise, U_M^* is defined as the minimum airtime guaranteed for \mathcal{M} under saturation loads. The sum of U_S^* and U_M^* is the saturation gateway bandwidth U^* . These values are apriori computed by means of analytical models (e.g., [7]) refined by measurements.

Note that adopting a minimum guaranteed rate for nodes in \mathcal{S} , not only prevents disadvantaged flows from misusing B_D , but also allows to tune the bandwidth effectively reserved for signaling according to the utilization of nodes in \mathcal{M} . In fact, as nodes in \mathcal{M} receive more bandwidth, less bandwidth is needed to indicate unserved demands. Thus, even though in a heavily utilized network the *maximum* unutilized bandwidth is B_D , the *actual* unutilized bandwidth can be reduced as soon as more disadvantaged demands are served. The U_S^* minimum guarantee is the tool that we use to allow the *signaling* bandwidth to be partially used by *data* flows originating at \mathcal{M} . For instance, at saturated load conditions, all of the system's resources are fully utilized and both traffic types will receive their guaranteed rates—i.e., nodes in \mathcal{S} will receive their minimum guaranteed rate U_S^* , while nodes in \mathcal{M} will take all the rest—and no bandwidth is left for additional signaling. In general, the single-hop nodes are allowed to exploit a

bandwidth that ranges from U_S^* to $(1 - \gamma)U^*$, the actual limit being dynamically and automatically determined by the demands of the multi-hop nodes, i.e., ideally:

$$R_S = \begin{cases} (1 - \gamma)U^* - \lambda_M, & \lambda_M \in [0, U_M^* - B_D]; \\ U_S^*, & \lambda_M > U_M^* - B_D; \end{cases} \quad (1)$$

where λ_M is the aggregate demand of nodes in \mathcal{M} , and R_S is the maximum aggregate rate allowed for nodes in \mathcal{S} , which is a piecewise linear function of λ_M . The corresponding unutilized part of the signaling bandwidth, under fully backlogged nodes in \mathcal{S} , ideally ranges between 0 and B_D :

$$B_D^{\text{free}} = \begin{cases} B_D, & \lambda_M \in [0, U_M^* - B_D]; \\ \max(0, U_M^* - \lambda_M), & \lambda_M > U_M^* - B_D. \end{cases} \quad (2)$$

In order to ensure that the single-hop set receive U_S^* under saturation loads, γ can be *at most* $\frac{U^* - U_S^*}{U^*}$. However, γ should be much less than $\frac{U^* - U_S^*}{U^*}$ in order to allow single-hop nodes to explore gateway resources more than U_S^* .

B. Distributed Algorithm

Here we present the algorithmic details of our distributed rate control scheme. When using GAP, all nodes in \mathcal{S} locally limit their transmission rates $R_i(t)$ based on the gateway utilization $U(t)$ and the local single-hop throughput. Algorithm 1 implements the proposed GAP framework at each single-hop node in \mathcal{S} for Additive Increase Multiplicative Decrease (AIMD) rate adaptation model. The algorithm is invoked every time a node in \mathcal{S} updates its gateway utilization indicator I_U , indicating whether or not $U(t) > (1 - \gamma)U^*$.

Algorithm 1

Input:

- I_U : Gateway Utilization Indicator
- $T_i(t)$: Throughput of Node $i \in \mathcal{S}$
- $\lambda_i(t)$: Instantaneous Offered Load of Node $i \in \mathcal{S}$
- F_i : Minimum Guaranteed Rate of Node i

Output:

- $R_i(t)$: The Rate Limit of Node $i \in \mathcal{S}$

Upon update of I_U Do

- 1: **if** $((T_i(t) < F_i(t) \parallel I_U == \text{low}) \ \&\& \ T_i(t) < \lambda_i(t))$
 - 2: $R_i(t) = R_i(t^-) + \alpha$
 - 3: **else if** $(T_i(t) > F_i(t) \ \&\& \ I_U == \text{high})$
 - 4: $R_i(t) = R_i(t^-) * \beta, \quad 0 < \beta < 1$
-

The GAP framework is implemented via two thresholds, namely, the maximum allowable single-hop utilization $(1 - \gamma)U^*$, and the minimum guaranteed single-hop rate $U_S^* = \sum_i F_i$. γ is a positive fraction that defines how much of the total gateway utilization U^* is allowed for disadvantaged flows' demand indication. If the maximum allowable single-hop utilization is not exceeded, all nodes in \mathcal{S} should have I_U set to *low*; otherwise, they all should have I_U set to *high*.

The distributed rate control process implemented at each node in \mathcal{S} is driven by the values of I_U and the achieved throughput: if the gateway airtime utilization is below the threshold $(1 - \gamma)U^*$ or the node throughput is below F_i , and




Fig. 4. Fluid system model.

the node demand is not satisfied, a single-hop node linearly increases its rate by α . On the other hand, when the gateway airtime utilization is above the $(1 - \gamma)U^*$ threshold, a single-hop node multiplicatively reduces its rate by a factor β *only* if its throughput exceeds its minimum guaranteed rate F_i . Thus, only nodes that are currently receiving more than their fair shares will reduce their rates upon traffic overflow indication. In Section V, we present a simple system model that isolates the main properties of 802.11 tree networks, and uses this model to analyze our rate control approach.

V. ANALYTICAL MODEL

In this section, we present a simple fluid model to isolate and study rate control in a system with priority service, utilization thresholds for inference of an *underserved state*, and control by proxy for high-priority nodes to control the rates of low-priority nodes. Our model omits all aspects of wireless networks and only focuses on these properties alone.

A. GAP in a Strict Priority System Model

Strict Priority System Model. Our rate control approach is driven by the distributed approximate prioritization of 802.11 mesh networks. We present the simplest system model that formulates this property as a priority fluid server. In particular, we consider a work-conserving queuing system in which two traffic flows are served via a single fluid server with fixed capacity U^* . Since our rate control objective is to distribute the system's capacity between the two different aggregate flows (advantaged and disadvantaged), we consider a single infinite capacity buffer per traffic type as shown in the uncontrolled system depicted in Fig. 4(a). Traffic from the buffer for the spatially advantaged node set \mathcal{S} has strict priority access to the server: when this buffer is non-empty, it is drained at rate U^* . The other buffer enqueues the spatially disadvantaged multi-hop traffic and is served only when the high priority buffer is empty. We consider a fluid server utilization and arrival processes $U_S(t)$, $U_M(t)$ and $\lambda_S(t)$, $\lambda_M(t)$, respectively. If no rate control is performed in the system, $U_S(t)$ and $U_M(t)$ can be written as in [8]:

$$U_S(t) = U^* \mathbf{1}_{\{q_S(t) > 0\}}; \quad (3)$$

$$U_M(t) = [U^* - U_S(t)] \mathbf{1}_{\{q_M(t) > 0\}}; \quad (4)$$

where $\mathbf{1}_{\{expr\}}$ is an indicator function which is equal to 1 if the expression *expr* is *true* and otherwise is equal to zero. Using equations (3) and (4), the dynamic equations for the occupancy of the two system buffers are given by

$$\dot{q}_S(t) = \lambda_S(t) - U^* \mathbf{1}_{\{q_S(t) > 0\}}; \quad (5)$$

$$\dot{q}_M(t) = \lambda_M(t) - U^* (1 - \mathbf{1}_{\{q_S(t) > 0\}}) \mathbf{1}_{\{q_M(t) > 0\}}. \quad (6)$$

GAP Rate-Controlled System Model. GAP controls the server utilization of the low priority traffic by controlling the rate by which the high priority traffic accesses the server. Therefore, we add a buffered rate controller, $R_S(t)$, at the input of the high priority buffer. Since the GAP distributed rate control process is driven by the knowledge of the gateway utilization, the rate controller $R_S(t)$ is able to observe the strict priority server as depicted in Fig. 4(b). On the other hand, the low priority buffer has no knowledge of $U(t)$ due to the spatial disadvantage of the set of nodes it represents. Due to the strict priority nature of the system, the low priority traffic utilization $U_M(t)$ is the server capacity not used by the high priority traffic, and eq.(4) holds also for the case shown in Fig. 4(b). Hence, $U_M(t)$ is controlled by controlling $U_S(t)$.

Since GAP aims at preventing the high priority traffic from consuming all of the gateway capacity U^* , the service rate of the high priority traffic $U_S(t)$ will be always equal to the rate limit $R_S(t)$ in the fluid system, i.e.,

$$U_S(t) = R_S(t)\mathbf{1}_{\{q_c(t)>0\}} \quad (R_S(t) \geq U_S^*, \forall t \geq 0); \quad (7)$$

Consequently, the high priority system buffer q_S can be always considered empty (i.e., $q_S(t) = 0, \forall t \geq 0$) under the fluid approximation of the controlled system. Thus, the queue occupancy dynamic equations of the rate-controlled system are given by

$$\dot{q}_c(t) = \lambda_S(t) - R_S(t)\mathbf{1}_{\{q_c(t)>0\}}; \quad (8)$$

$$\dot{q}_M(t) = \lambda_M(t) - [(U^* - U_S(t))\mathbf{1}_{\{q_M(t)>0\}}]. \quad (9)$$

B. Algorithmic Operation

For an AIMD rate adaptation model, the dynamic equation of the high priority buffer controller is given by

$$\begin{aligned} \dot{R}_S(t) = & \alpha \mathbf{1}_{\{U(t) < (1-\gamma)U^*\}} \mathbf{1}_{\{R_S(t) < \lambda_S(t)\}} + \\ & -\beta R_S(t) \mathbf{1}_{\{U(t) > (1-\gamma)U^*\}} \mathbf{1}_{\{U_S(t) > U_S^*\}}. \end{aligned} \quad (10)$$

The first term in the right hand side of (10) corresponds to an additive increase in the rate whenever the feedback signal implies $U(t) < (1-\gamma)U^*$. Namely, the high priority traffic increases its rate and correspondingly decreases the allocated rate to low priority traffic whenever the server utilization is below the $(1-\gamma)U^*$ threshold. The constant α is the additive increase step size. The second term in the right hand side of (10) corresponds to multiplicative decrease. The constant β is the multiplicative decrease factor. The high priority buffer controller decreases its rate by β whenever the server utilization exceeds the $(1-\gamma)U^*$ threshold to allow low priority traffic demands to be served. The dynamic equation for the high priority traffic rate controller (10), along with the queue occupancy dynamics in (8) and (9), characterizes our rate rate-controlled system.

Based on the offered load, the three following conditions drive the system to a steady state point, i.e., $\dot{R}_S(t) = 0$.

(a) **Light high priority traffic demands.** If $\lambda_S(t)$ is below the minimum guaranteed single-hop utilization U_S^* , the strict priority behavior of the system allows the low priority demand λ_M to be served up to $U^* - \lambda_S$. In this case, both terms of (10) are equal to zero, $\lambda_S = U_S \leq U_S^* \leq R_S$, and $\lambda_M \leq U^* - \lambda_S$. A full gateway airtime utilization is hence possible.

(b) **Unsaturated traffic demands.** If the offered load of the high priority buffer exceeds the minimum guaranteed single-hop threshold U_S^* but it does not exceed the rate limiting value R_S , and the aggregate offered load is at most equal to $(1-\gamma)U^*$, there is no need for changes in the rate control, and both terms in (10) are zero ($U_S^* < U_S = \lambda_S \leq R_S$ and $U \leq (1-\gamma)U^*$). From this state, GAP allows the high priority buffer controller to increase its rate until the overall server utilization $U(t)$ reaches the threshold $(1-\gamma)U^*$, whilst the low priority traffic can express future increments in its demands using B_D plus the resources left unused by the high priority flows.

(c) **Saturated traffic demands.** Finally, consider the case wherein the high and low priority traffic demands are greater or equal to U_S^* and $(1-\gamma)U^* - U_S^*$, respectively. As explained earlier, the low priority traffic will keep using $B_D = \gamma U^*$ causing the high priority buffer controller to decrease its rate R_S until the high priority utilization U_S is equal to its minimum guaranteed value U_S^* . Thus, the system enters a steady state in which $U_S^* = R_S = U_S \leq \lambda_S$ and $U \geq (1-\gamma)U^*$. The low priority utilization is $\min(\lambda_M(t), U_M^*)$, and thereby a 100% utilization of the server capacity is possible if $\lambda_M(t) \geq U_M^*$, without any losses whatsoever due to any type of overhead.

C. Stability of the Equilibrium Points

Here, we study the stability of GAP described by the non-linear ordinary differential equation in (10). More specifically, we examine the stability of linearizations about the equilibrium points. Although such a technique does not lead to sufficient conditions for global stability (i.e., stability under arbitrarily large perturbations), it does lead to necessary conditions and provides useful insights [9].

Definition 1: An equilibrium point R_S^{eq} is said to be stable if small perturbations from that point decay with time. This condition is met only if the Taylor series expansion of the derivative of the right-hand side of equation (10) at R_S^{eq} is negative valued [9]. We refer to this condition as **the stability criterion**.

The following two theorems characterizes the stability of the equilibrium points of GAP.

Theorem 1: The equilibrium points of the proposed rate control framework are satisfying the stability criterion.

Theorem 2: Perturbations of the equilibrium points of the system are exponentially decaying with time with constant that is based on the equilibrium point.

Due to space limitations we only state the theorems. The detailed proofs are available in [6].

D. Applications of the Model

In [6] we present some applications of the model. Here, we overview such applications.

In our model we intentionally isolate the main 802.11 network properties driving our rate control approach. However, the model allows for the incorporation of other aspects of 802.11 networks for deeper understanding of the performance of our rate control approach. For example, in [6] we extend the model to replace the fluid service model with the 802.11 service model presented in [4].

We also show how to deduce the values of the additive increase and multiplicative decrease parameters that ensure the stability of the AIMD behavior at the steady state operation points of the algorithm.

VI. PERFORMANCE EVALUATION

In this section, we use simulations to demonstrate that by using elastic rate limiting, the GAP achieves the targeted gateway sharing with close accuracy compared to idealized 802.11s-compliant algorithms, yet GAP incurs very little or no overhead and thus achieves higher utilization.

A. Simulation Setup

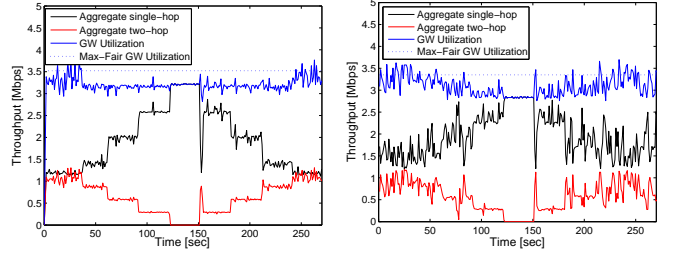
Simulations are conducted using the unmodified IEEE 802.11 MAC implementation in *ns-2* v.2.31 with default parameters. The carrier sensing range and the transmission range are set to 250m. The auto-rate mechanism is disabled and the data rate of all nodes is set to 11 Mbps. Each node is equipped with a single half duplex interface with an omnidirectional antenna. We disable the RTS/CTS exchange in order to avoid the associated overhead. We use the token bucket filter implementation in *ns-2* to implement the rate limiters at the single hop nodes. We do not simulate the gateway utilization estimation scheme, rather, the gateway utilization is made available to each single hop nodes via a centralized estimator employing EWMA filter. We denote x in $GAP(x)$ as the amount of signaling bandwidth B_D used in the experiment. At saturation load conditions, we target max-min fair rates for the two sets \mathcal{S} and \mathcal{M} .

Our benchmark is the Inter-TAP Fairness Algorithm (IFA) [7] that is an IEEE 802.11s-like link-layer rate control algorithm. In IFA, backhaul nodes periodically broadcast their demands and link capacities to adjacent nodes (every 300 ms) in control messages sent at the base rate (we simulate 50-byte packets). Using such information, all nodes individually calculate their fair shares, using the reference fairness model presented in [7], and communicate them to neighboring nodes in order for the fair share computation in a contention region to be coordinated. We consider an idealized variant of IFA that uses a global observer that measures the individual link capacities needed for the fair rate calculations.

We present results for TCP traffic upload traffic. Other traffic matrices including mixes of UDP and download traffic are available in [6]. The reported results, except for the instantaneous behavior results, are the average of 10 simulations.

B. Tracking Changes in the Multi-hop Traffic's Demands

In this experiment, we demonstrate the ability of the single-hop node set \mathcal{S} to track changes in the offered load of the multi-hop node set \mathcal{M} . We consider both a gateway of degree 1 and 4 with two-hop branches. Single-hop nodes carry backlogged traffic. The offered load of the two-hop node set \mathcal{M} is initially 1.2 Mbps (higher than the saturation share). Every 30 seconds, we decrease the offered load of the two-hop node set \mathcal{M} in 300 Kbps steps until it reaches zero, then increase it again starting at 150 seconds with the same step (in the 4-branch topology, we turn off/on a two-hop node every 30 seconds). As shown in Fig. 5, as long as the



(a) Single two-hop branch.

(b) Four two-hop branches.

Fig. 5. The instantaneous throughput of the two node sets \mathcal{S} and \mathcal{M} and the gateway utilization while varying the multi-hop set load.

multi-hop node set \mathcal{M} offered load is below its saturation fair share, the single-hop node set \mathcal{S} increases its rate for efficient gateway utilization. However, the single-hop traffic cannot utilize the signaling bandwidth (which is 256 Kbps and 512 Kbps in Fig. 5(a) and 5(b), respectively). Consequently, when the offered load of the two-hop nodes increases, nodes in \mathcal{M} transmit using B_D such that nodes in \mathcal{S} infer the change in \mathcal{M} 's offered load and reduce their rate accordingly. Fig. 5(b) shows that tracking demand changes is doable, even though more challenging, also when multiple single-hop nodes are competing, and the distributed AIMD mechanism generates higher throughput oscillations while keeping almost unchanged the average and the aggregate utilization.

C. Topological Factors

We evaluate the performance of the GAP mechanism and the 802.11s-like algorithm under the impact of topological factors such as the gateway degree, branch depth, and hidden terminals. We use three performance metrics: (i) the accuracy of the received throughput T with respect to the targeted throughput T^* , defined through the error

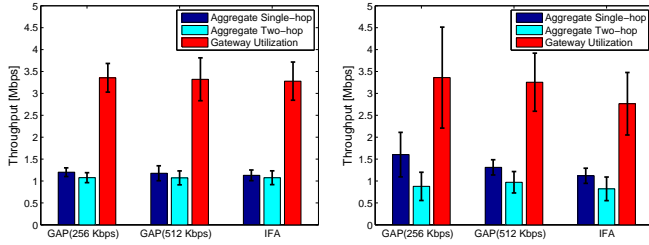
$$E = \frac{1}{2} \left[\left| \frac{T_S - T_S^*}{T_S^*} \right| + \left| \frac{T_M - T_M^*}{T_M^*} \right| \right]; \quad (11)$$

(ii) the time to rise t_{rise} , defined as the time needed, for fully backlogged single-hop nodes to reach the maximum allowable gateway utilization upon switching off all the multi-hop load;

(iii) the time to decay t_{decay} , defined as the time needed for fully backlogged single-hop nodes to decrease their aggregate rate to the minimum guaranteed single-hop bandwidth, after switching on fully backlogged multi-hop nodes. t_{rise} and t_{decay} define the *convergence time* of the selected mechanism.

Impact of gateway degree. First, we study the impact of the distributed nature of the network on the performance of both GAP and IFA. To address this issue, we vary the gateway degree N while having all nodes in \mathcal{S} in range of each other (in order to isolate hidden terminals effect and utilization estimation accuracy). Fig. 6 depicts the aggregate throughput of the two node sets and the total gateway utilization for gateway of degree of 1 and 8 and two-hop branches. Table I lists the error with respect to the targeted rates and the rise and decay times for these two topologies.

For a single branch topology, GAP closely converges to the targeted rates (with 6.7% error) with a signaling bandwidth as low as $B_D = 256 \text{ Kbps}$ (7.6% of U^*), as shown in the



(a) Single two-hop branch.

(b) Eight two-hop branches.

Fig. 6. Throughput of the two node sets \mathcal{S} and \mathcal{M} , and aggregate gateway utilization under saturated demands.

leftmost set of bars in Fig. 6(a). As N increases, the amount of signaling bandwidth needed for the GAP algorithm to accurately converge to the targeted rates increases as depicted in Fig. 6(b). Note that the accuracy and stability of both GAP and IFA algorithms decrease with N . However, only IFA causes the gateway utilization to decrease with N . Such performance degradation of IFA is attributed to the increased overhead of the periodically exchanged control messages. Moreover, GAP converges to the target rate faster than IFA as more nodes cooperate in rate control (Table I). In IFA, nodes need more time to convey the changes in their loads to the gateway to compute the fair rates and communicate the rates back to different nodes.

Impact of branch depth and hidden nodes. Our results presented in [6] extend the validity of the results presented here to network with more than two hops, and with different degree of coordination between nodes. Our experiments show that the required signaling bandwidth increases with the branch depth, and decreases when nodes in \mathcal{S} are not coordinated (i.e., *hidden*). In fact, the presence of hidden nodes slows down the transmission rates of single-hop nodes, thereby, more transmission opportunities occur for multi-hop nodes. In contrast, the throughput of the IFA system further decreases with both the branch depth and the presence of hidden nodes.

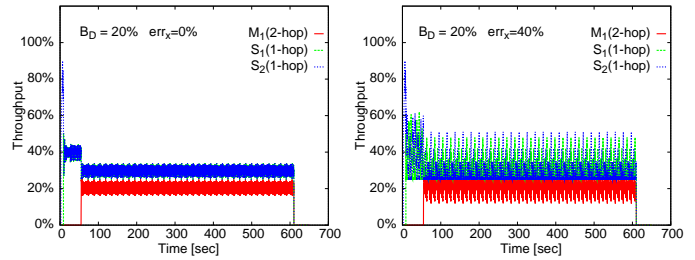
D. Robustness of GAP

The gateway utilization indication can be affected by errors due to the utilization estimation process performed either by the gateway or by each single-hop node independently. In this set of experiment we want to observe the behavior of two single-hop nodes (S_1 and S_2) and one two-hop node (M_1) (as in Fig. 2(a)), when the gateway's resource estimation process is affected by a random error.

We numerically evaluate Algorithm 1 in a scenario in which all nodes are fully backlogged and single-hop nodes are controlled by the GAP. In particular we compare the results obtained under error-free estimation and under a $\pm 40\%$ uncertainty in each estimation. The parameters for the simulation are

TABLE I
CONVERGENCE TIME AND ACCURACY OF DIFFERENT ALGORITHMS.

Parameter	Single Branch			Eight Branches		
	GAP 256K	GAP 512K	IFA	GAP 256K	GAP 512K	IFA
t_{rise} [s]	0.9857	1.0341	0.3040	0.0170	0.0224	0.3564
t_{decay} [s]	0.0645	0.0484	0.3867	0.0332	0.0024	0.4045
E [%]	6.7035	5.7925	5.63	32.226	15.075	13.39



(a) No estimation error.

(b) $\pm 40\%$ estimation error.

Fig. 7. System behavior due to the GAP mechanism with and without estimation uncertainty.

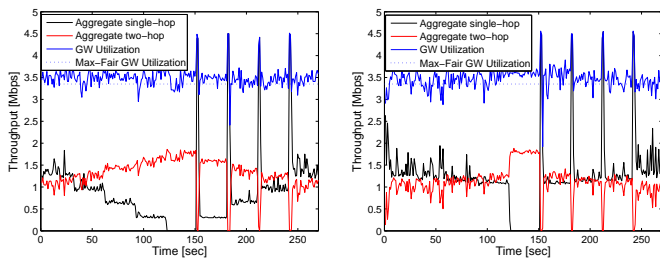
scaled to the gateway capacity U^* as follows: $B_D = 0.2U^*$, $U_S = 0.5U^*$, $\alpha = 0.01U^*$, and $\beta = 0.9$.

The error-free case depicted in Fig. 7(a) shows the throughput oscillations typical of an AIMD rate control algorithm. After the warm up phase, i.e., soon after the last flow enters the network, single-hop nodes never go below their guaranteed rate (25% of U^*), and the two-hop node stays close to its fair share (that is 25% after considering that two-hop node throughput consumes about two times the gateway airtime of single-hop throughput). In the error-prone scenario of Fig. 7(b), we can observe larger oscillations in the throughput, as expected due to the large estimation errors. However, the peaks in each throughput curve are quickly recovered by the AIMD mechanism adopted by the GAP. Note that the average throughput of each node remains closer to the selected share rate, i.e., 25% for each node. Therefore, the GAP is shown to be robust to very high estimation errors.

E. Per-node vs. Aggregate Minimum Guarantees

Finally, we investigate the performance of the GAP when only using a per-node minimum guaranteed rate (Fig. 8(a), representing results obtained with Algorithm 1), versus a modified GAP version using an aggregate single-hop minimum guaranteed rate (Fig. 8(b), representing the outcome of an algorithm in which the gateway utilization indication I_U is evaluated in conjunction with an aggregate single-hop utilization indication I_S , i.e., the binary value I_U is replaced by $(I_U \&\& I_S)$). As well as for I_U (see Sec. IV-A), the value of the indication variable I_S can be either estimated by each node or simply notified by the gateway with an extra signaling bit, or also encoded together with I_U . In our experiment, we consider a 4-branch topology in which all nodes are initially fully backlogged; then, every 30 seconds, one single-hop node is switched off after the other, until only multi-hop nodes remain active.

As can be seen in Fig. 8, when the offered multi-hop traffic is saturated and the load of the single-hop set is asymmetrically distributed among different nodes in \mathcal{S} , the adoption of per-node utilization guarantees can cause the algorithm based only on the maximum single-hop allowable utilization to fail realizing the targeted gateway sharing objective between the two sets \mathcal{S} and \mathcal{M} , but it conservatively advantages the multi-hop node set whilst not starving the single-hop node set. This is because the backlogged demand of \mathcal{M} keeps the signaling bandwidth fully utilized, and thereby, nodes in \mathcal{S} cannot increase their rates above their individual minimum guaranteed rate.



(a) Per node guarantee.

(b) Aggregate guarantee.

Fig. 8. The instantaneous throughput of the two competing sets S and M , and the gateway utilization with different available information.

VII. RELATED WORK

There have been significant research efforts investigating rate control in wireless multi-hop networks. This research was motivated by TCP congestion control throughput degradation [10] and severe unfairness [11] in such networks. Early work on wireless rate control aimed at either improving the performance of a multi-hop TCP flow using techniques that reduce flow self-contention and packet losses (e.g., reducing TCP transmission rate or window size) [12]–[16], or radically redesigning the transport protocol for multi-hop wireless environments as in [17] and [18]. In addition, [2] and [19] proposed link-layer techniques that slow down nodes along the multi-hop path to counter TCP unfairness.

Alternatively, [5], [7], [20]–[22] viewed wireless congestion as a neighborhood phenomenon rather a per flow problem. Thus, nodes in a congested neighborhood collaborate in detecting and resolving the congestion situation. Irrespective of the particular mechanism of such rate control schemes, nodes in these schemes explicitly exchange information that reflect their participation in the congestion state (e.g. queue length [5], [20], [21], or outgoing links’ capacities and offered loads [7]) in order to allow different nodes to compute their appropriate transmission rates. Then, nodes communicate their calculated rates in order to insure that all flows in the neighborhood are fairly sharing the available resources.

However, none of the prior work formulates the rate control problem in multi-hop mesh networks based on the system properties identified in this paper. Unlike TCP-specific schemes [12]–[16], GAP is a link-layer solution that is independent of the transport protocol. Moreover, our rate control approach does not deteriorate the network throughput with unnecessary control message exchanges which can cause severe throughput degradation over the wireless medium [23].

VIII. CONCLUSIONS

In this work, we exploit the specific properties of spatially biased 802.11 mesh networks that drive the rate control problem in such systems. We present GAP, an elastic rate control mechanism in which the set of nodes with spatial advantages, in terms of gateway access, locally and collaboratively control their transmission rates to share the gateway resources with other spatially disadvantaged nodes. We present a simple analytical model that isolates the system’s properties driving our rate control scheme and use it to analyze the proposed

GAP framework. Finally, we use simulations to show that GAP achieves comparable performance with overhead-prone algorithms without deteriorating the system’s throughput via explicit rate control messaging.

REFERENCES

- [1] J. Camp and E. Knightly, “The IEEE 802.11s extended service set mesh networking standard,” *IEEE Communications Magazine*, vol. 46, no. 8, pp. 120–126, Aug. 2008.
- [2] J. Shi, O. Gurewitz, V. Mancuso, J. Camp, and E. W. Knightly, “Measurement and Modeling of the Origins of Starvation in Congestion Controlled Mesh Networks,” in *Proceedings of IEEE INFOCOM*, Phoenix, AZ, USA, Apr. 2008.
- [3] M. Garetto, J. Shi, and E. Knightly, “Modeling media access in embedded two-flow topologies of multi-hop wireless networks,” in *Proceedings of ACM MOBICOM*, Cologne, Germany, Aug. 2005.
- [4] M. Garetto, T. Salonidis, and E. Knightly, “Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks,” *ACM/IEEE Transactions on Networking*, vol. 16, no. 4, pp. 864–877, Aug. 2008.
- [5] L. Qiu, Y. Zhang, F. Wang, M. Han, and R. Mahajan, “A general model of wireless interference,” in *Proceedings of ACM MOBICOM*, Montréal, QC, Canada, 2007.
- [6] A. Khattab, O. Gurewitz, V. Mancuso, and E. Knightly, “Overhead-free congestion control framework for IEEE 802.11 multi-hop mesh networks,” in *Rice University Technical Report TREE0806*, Aug. 2008.
- [7] V. Gambiroza, B. Sadeghi, and E. W. Knightly, “End-to-End Performance and Fairness in Multihop Wireless Backhaul Networks,” in *Proceedings of ACM MOBICOM*, Philadelphia, PA, USA, Sept. 2004.
- [8] Y. Liu and W. Gong, “On fluid queueing systems with strict priority,” *IEEE Transactions on Automatic Control*, vol. 48, no. 12, pp. 2079–2088, Dec. 2003.
- [9] A. Kumar, D. Manjunath, and J. Kuri, *Communication Networking: An Analytical Approach*, 2004.
- [10] M. Gerla, R. Bagrodia, L. Zhang, K. Tang, and L. Wang, “TCP over Wireless Multi-hop Protocols: Simulation and Experiments,” in *Proceedings of IEEE ICC*, Vancouver, BC, Canada, June 1999.
- [11] S. Xu and T. Saadawi, “Revealing TCP Unfairness behavior in 802.11 based Wireless Multi-hop Networks,” in *Proceedings of IEEE PIMRC*, San Diego, CA, USA, Oct. 2001.
- [12] E. Altman and T. Jimnez, “Novel delayed ACK techniques for improving TCP performance in multihop wireless networks,” in *Proceedings of IFIP PWC*, Venice, Italy, Sept. 2003.
- [13] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla, “The Impact of Multihop Wireless Channel on TCP Throughput and Loss,” in *Proceedings of IEEE INFOCOM*, San Francisco, CA, USA, Mar. 2003.
- [14] J. Chen, M. Gerla, Y. Z. Lee, and M. Sanadidi, “TCP with Variance Control for Multihop IEEE 802.11 Wireless Networks,” in *Proceedings of MILCOM 2006*, Washington, D.C., USA, Oct. 2006.
- [15] S. M. ElRakabawy, A. Klemm, and C. Lindemann, “TCP with Adaptive Pacing for Multihop Wireless Networks,” in *Proceedings of ACM MOBIHOC*, Urbana-Champaign, IL, USA, May 2005.
- [16] —, “Gateway Adaptive Pacing for TCP across Multihop Wireless Networks and the Internet,” in *Proceedings of ACM/IEEE MSWiM*, Malaga, Spain, Oct. 2006.
- [17] K. S. V. Anantharaman, H.-Y. Hsieh, and R. Sivakumar, “ATP: A Reliable Transport Protocol for Ad-hoc Networks,” in *Proceedings of ACM MOBIHOC*, Annapolis, MD, USA, June 2003.
- [18] P. Sinha, N. Venkitaraman, R. Sivakumar, and V. Bharghavan, “WTCP: A reliable transport protocol for wireless wide-area networks,” in *Proceedings of ACM MOBICOM*, Seattle, WA, USA, Aug. 1999.
- [19] L. Yang, W. K. Seah, and Q. Yin, “Improving fairness among TCP flows crossing wireless ad hoc and wired networks,” in *Proceedings of ACM MOBIHOC*, Annapolis, MD, USA, June 2003.
- [20] K. Xu, M. Gerla, L. Qi, and Y. Shu, “Enhancing TCP Fairness in Ad Hoc Wireless Networks Using Neighborhood RED,” in *Proceedings of ACM MOBICOM*, Seattle, WA, USA, Sept. 2003.
- [21] S. Rangwala, R. Gummadi, R. Govindan, and K. Psounis, “Interference-aware fair rate control in wireless sensor networks,” in *Proceedings of ACM SIGCOMM*, Pisa, Italy, Sept. 2006.
- [22] A. Jindal and K. Psounis, “Characterizing the Achievable Rate Region of Wireless Multi-hop Networks with 802.11 Scheduling,” *ACM/IEEE Transaction on Networking*, vol. 17, no. 4, pp. 1119–1131, Aug. 2009.
- [23] J. Camp, V. Mancuso, O. Gurewitz, and E. W. Knightly, “A Measurement Study of Multiplicative Overhead Effects in Wireless Networks,” in *Proceedings of IEEE INFOCOM*, Phoenix, AZ, USA, Apr. 2008.