

DISTRIBUTED SOURCE RATE LIMITING MECHANISM FOR WIRELESS MESH NETWORK

Navdeep Saluja*¹ and Rajesh Shrivastava²

^{1,2}Department of C.S.E

Shriram Institute of Technology

Jabalpur, M.P., (India)

navdeep_saluja2000@yahoo.co.in¹, rajesh_5479@rediffmail.com²

Abstract- Wireless Mesh Networks (WMNs) built with commodity 802.11 radios are a cost-effective means of providing last mile broadband Internet access. Their multi-hop architecture allows for rapid deployment and organic growth of these networks. In this paper we focus on fair rate allocation requirements. Our approach does not require any changes to individual mesh routers. Further, it uses existing data traffic as capacity probes, thus incurring a zero control traffic overhead. We propose the mechanism based on this approach on aggregate rate control (ARC). ARC limits the aggregate capacity of a network to the sum of fair rates for a given set of flows. We show that the resulting rate allocation achieved approximately max-min fair. We show how it can be used to achieve weighted flow rate fairness. Our comparative analysis show that our mechanisms improve fairness indices when compared with networks without any rate limiting, and are approximately equivalent to results achieved with distributed source rate limiting mechanisms that require software modifications on all mesh routers.

Keywords- Wireless Mesh Network, Distributed Source Mechanism, Aggregate Rate Control.

INTRODUCTION

Wireless Mesh Networks (WMNs) are a type of multi-hop wireless network in which the mesh nodes act both as a host as well as a traffic relay for other nodes in the network. WMNs have two types of nodes: regular *mesh nodes* that can act both as data sources as well as routers, and *gateway nodes* that bridge traffic between the mesh network and a wired network, typically the Internet. In IEEE 802.11s standards terminology, these nodes are referred to as *Mesh Points (MP)* and *Mesh Point Portal (MPP)*, respectively. Client devices connect to their preferred mesh node either via wire or over a (possibly orthogonal) wireless channel, and use the multi-hop wireless relay to communicate with the gateway [1].

The IEEE 802.11 standard specifications were originally conceived for single-hop communication in a Wireless Local Area Network (WLAN). Studies have indicated that these radios exhibit suboptimal performance in multi-hop networks. The benefits of commodity 802.11 radios, however, seem to far outweigh these performance challenges. A large number of commercial WMN vendors as well as research test-beds use 802.11 radios, preferring to address any performance challenges through modifications in other layers of the network stack. Service model of a community wireless network with the last mile access provided through WMNs, we believe that wireless Internet Service Providers (ISPs) can build a business case for serving rural communities. ISPs only need to provide an Internet point-of-presence (PoP) by installing a gateway mesh router with always-on broadband Internet connectivity. In remote communities, this gateway connection to the Internet may also be a wireless link through satellite or WiMax networks. Community residents interested in subscribing to the ISP's Internet service can simply configure their commodity 802.11-based mesh routers to communicate with this gateway, either directly or through multi-hop wireless links. Our focus in this

dissertation is on understanding and addressing the performance challenges associated with enforcing a policy-driven resource management in 802.11 based WMNs. Our goal is to develop a set of mechanisms that enable an ISP to efficiently manage their network resources while conforming to their desired resource allocation criterion.

Resource allocation has been extensively studied in wired networks. It is often modeled as a constrained optimization problem. The set of constraints in a wireless network are fundamentally different from that of a wired network, and this necessitates a fresh perspective into the problem. The wireless channel is a broadcast medium with its spectral resource shared between all contending nodes. In a localized neighborhood, only a single node can transmit at a time, as concurrent transmissions will result in collisions and subsequent packet loss. These networks also face other sources of packet loss and interference, including scattering and multi-path fading from obstructions in the area [2] and [3].

Our research objectives can be summarized as follows:

1. We wish to understand the requirement for managing the allocation of network resources in a WMN. In particular, we are interested in exploring the behavior of 802.11 MAC in multi-hop networks, the response of transport-layer protocols, and the resulting interaction across these layers under varying traffic loads and network conditions.
2. Devise a framework of mechanisms that can enforce an efficient and a policy driven allocation of available network capacity. The spectral resource of a wireless network is susceptible to temporal variance in capacity due to unpredictable losses from collisions, interference or other physical-layer phenomenon specific to the radio channel. We are interested in developing solutions that can adapt to these vagaries of the wireless channel.
3. We wish to limit the scope of any proposed resource management framework to a set of mechanisms that can

be supported on the commodity 802.11 hardware. Further, the mechanisms need to be incrementally deployable for them to be of any practical utility to a network service provider. Most of the prior literature for enforcing a rate allocation in WMNs proposes some variant of distributed rate limiting protocols. These protocols require periodic flooding of time-varying state information to enable this distributed computation. Interpreting and reacting to this information requires software changes on all mesh routers [1] and [4].

In this paper we propose, design, and evaluate a set of mechanisms that can centrally manage the rate allocation process for these adaptive traffic streams using only the information locally available at the gateway. Through extensive experimental analysis, we establish that our proposed mechanisms can effectively limit these traffic flows to their allocated share of the network capacity.

BACKGROUND

MAC-layer Enhancements:-

By far the largest body of literature specifically devoted to wireless network fairness is that of MAC-layer solutions. Such approaches tend to assume that contending flows span a single-hop and fairness may be achieved by converging the MAC contention windows to a common value. Single-hop fairness, however, does not translate to end-to-end fairness in a multi-hop network.

Hidden and Exposed Terminals:-

The impact of hidden terminals in degrading network capacity and producing flow rate unfairness has been described in prior work. RTS/CTS handshake works well in a BSS architecture where any two stations are at most two hops apart. However, it does not resolve all hidden terminals when multiple BSS are co-located in a given space, and hidden nodes may now exist in the interference range of a receiver. Typically, such hidden nodes are mitigated by assigning non-overlapping channels to neighboring BSS. This does not work in single-radio WMNs where all radios operate on a common channel to ensure connectivity.

Haas and Deng have proposed Dual Busy Tone Multiple Access (DBTMA) that solves both the hidden and the exposed terminal problem. DBTMA builds on to the idea of Receiver-Initiated Busy-Tone Multiple Access (RI-BTMA) scheme in which the receiver broadcasts an out-of-band busy tone during the reception of data. This allows neighboring nodes to detect an ongoing communication at the receiver, thus preventing them from transmitting. In addition to this receive busy tone, DBTMA introduces an out-of-band transmit busy tone. A transmitter initiates this tone when sending out an RTS packet, thus protecting this packet and increasing the probability of its successful reception at the intended receiver. This design allows exposed terminals to initiate a new communication because they do not listen on the shared data channel for receiving acknowledgment from the intended receiver; this acknowledgment is instead sent as a receive busy tone.

Similarly, hidden terminals can respond to RTS requests by enabling the receive busy tone. While DBTMA MAC design

is simple, its practical implementation remains a challenge. First, it needs extra transmitting and sensing circuitry per node for generating and receiving the two busy tones. Second, there needs to be a considerable spectral separation between the data channel and the control channel for the tones to prevent any co-channel interference. However coordinating radios across widely separate frequency bands is hard because of frequency-dependent radio propagation characteristics. Finally, though both transmit and receive busy tones are narrowband tones, yet they still incur an overhead by consuming a finite amount of bandwidth that cannot instead be used by the data channel [1] and [2].

Prioritized MAC Access:-

CSMA/CA based MACs provide each node with an equal opportunity to transmit. However, in WMN, nodes closer to the gateway have to transmit their own traffic along with the aggregate traffic from other nodes in the network towards the gateway. As such, congested nodes higher in the connectivity graph should have prioritized access to the medium compared to its child nodes.

There are different mechanisms for prioritizing access to the wireless medium. If the fair share of a flow is already known, nodes along the multi-hop link can set their contention delays so as to achieve the desired rate. Hull et al. recommend using the scheme first outlined; a node higher in the connectivity graph has its randomized back off window set to $1/2^N$ the size of each of its N children, as on average this gives the parent node as many transmission opportunities as each of its children. In previous, the authors propose adjusting TXOP and CWmin values as described in 802.11e enhancements to restore fairness amongst multi-hop flows. We note that such schemes do not provide fairness amongst the same access category and can only scale to small networks [3] and [4].

Receiver-initiated/Hybrid MAC:-

The contention resolution mechanism in CSMA/CA is sender-initiated. We have seen that when the channel state at the sender is incomplete, its transmissions result in collisions at the receiver. To resolve this, receiver-initiated transmission protocols have been proposed. However, these schemes work well only when the intended receiver is aware of the exact traffic load information. Hybrid MAC protocols that combine both sender and receiver-initiated transmissions have also been proposed. Results show that such hybrid schemes can only improve fairness in some scenarios without significantly degrading the network capacity. Receiver-initiated extensions to traditional CSMA/CA protocol have also been proposed. MACAW uses a receiver-initiated Request-for-Request-To-Send (RRTS) control packet. It handles the use case when a receiver receives a RTS but cannot immediately respond with CTS till its NAV counter expires. At NAV expiration, the receiver transmits a RRTS frame to the original sender, requesting that RTS packet be retransmitted. Note that this works only when the initial RTS is decodable at the receiver. If this RTS is non-decodable due to a collision (lasting the duration of RTS frame), the receiver can decide to broadcast a RRTS frame based on a given probability distribution [74]. We note that while RRTS control frame can resolve some scenarios with missed transmission

opportunities, the overhead associated with the probabilistic use of an additional control frame can be a performance challenge for wireless networks [5].

Challenges: MAC-layer Modifications:-

Modifying MAC-layer protocols to support multi-hop flow fairness involves the following challenges:

1. Schemes requiring MAC-layer modifications have incremental deployment challenges, as the MAC-layer is fundamental for establishing link-level connectivity amongst the network nodes. This clearly limits its practical utility for the type of broadband wireless access networks based on heterogeneous node configurations that we consider in this dissertation.

2. Only a subset of the MAC-layer modifications described above can be practically implemented on 802.11 radio chipsets. This is because some functionality of these radios (such as carrier sensing) is inherent to the firmware which is often not exposed by the manufacturers. Switching to a different radio platform may often be infeasible because the cost dynamics of commodity 802.11 hardware is a significant factor in making WMNs an attractive last mile access technology.

TCP Enhancements:-

There is extensive literature on understanding and improving the performance of TCP in wireless networks. We provide a brief overview of broad categories of this research relevant to our work, referring the reader to a comprehensive survey. The performance of TCP in one-hop wireless networks has been studied extensively. The congestion control mechanisms of TCP have been optimized for wired networks. Wireless networks have fundamentally different characteristics in terms of bandwidth, propagation delay, and link reliability. As a result, packet loss can no longer be treated simply as an artifact of congestion in the network, disrupting the foundations of TCP's congestion control mechanisms.

Balakrishnan *et al.* classified the research work addressing TCP performance limitations in a one-hop wireless network into three major categories: end-to-end proposals, split-connection proposals, and link-layer proposals. The end-to-end protocols attempt to make TCP senders differentiate between losses from congestion or from other errors through the use of Explicit Loss Notification (ELN) mechanism. They also use variants of selective acknowledgments (SACKs) to enable the sender to recover from multiple packet losses in a window. The split-connection protocols hide the wireless link from a wired TCP sender by terminating its TCP connection at the wireless base station (BS), and using a separate reliable transport protocol on the wireless link between the BS and the wireless client.

Finally, the link-layer protocols attempt to shield the TCP sender from wireless losses by implementing local retransmissions and forward error correction [6] and [7].

Previous Mechanisms:-

Rate-control mechanisms that operate independently of the transport-layer protocols have also been proposed. We provide a summary of this work below.

A. Router-assisted Control:-

In general, router-assisted resource management encompasses a set of mechanisms including congestion signaling, packet scheduling, and queue management. Active Queue Management (AQM) protocols like Random Early Detection (RED) are examples of such router-assisted congestion control mechanisms. RED gateways are typically used at network boundaries where queue build-ups are expected when flows from high throughput networks are being aggregated across slower links. RED gateways provide congestion avoidance by detecting incipient congestion through active monitoring of average queue sizes at the gateway. When the queue size exceeds a certain threshold, the gateway can notify the connection through explicit feedback or by dropping packets. RED gateways require that the transport protocol managing those connections be responsive to congestion notification indicated either through marked packets or through packet loss. We observe that there has been little work exploring the applicability (or lack thereof) of AQM techniques in multihop wireless networks. One noticeable exception is Neighborhood RED (NRED) scheme that drops packets based on the size of a virtual distributed "neighborhood" queue comprising all nodes that contend for channel access. NRED only identifies a subset of these contending nodes, as it misses the flows that are outside the transmission range but still interfere. Additionally, this proposed mechanism is closely tied with a particular queue management discipline required on all mesh routers [8] and [9].

B. Rate-based Protocols:-

In contrast to window-based protocols, rate-based protocols require the receiver or the network to inform the sender of the rate at which it can support that connection. EXACT is an end-to-end rate-based flow control technique for ad hoc networks. It requires each router to periodically estimate its local capacity and use that to compute the fair share of all egress flows. A special IP header (called flow control header) is then populated with the lowest end-to-end rate along the path of a flow. This rate is communicated back to the source where it is enforced locally. However, maintaining per-flow state at intermediate routers is a challenge in scaling EXACT to large multihop networks.

Further, EXACT only defines a rate based flow control scheme; additional mechanisms like Selective Acknowledgments (SACK) need to be separately overlaid on EXACT for reliable data transmission. Stateless rate-based protocols have also been proposed in the literature. ATP (Ad hoc Transport Protocol) uses packet delay feedback from intermediate routers to allow a source to compute its rate. The routers maintain the average delay (sum of queuing and transmission delays) experienced by all egress packets.

These routers can then update the header of a passing packet such that the header always carries the largest delay encountered along the path of a packet. The receiver aggregates these delay values over a time interval, computes a new rate, and communicates it to the sender. Results show that while ATP improves the average rate allocation compared to TCP, the fairness index was no better than TCP and some flows still experienced starvation [8] and [10].

C. Alternative Distributed Protocol Designs:-

Distributed algorithms for enforcing fairness in multihop networks have been proposed, amongst others. In general, distributed computation of fair rates requires periodic signaling between contending senders. For example, a change in the status of a stream activity needs to be propagated to all nodes along the path of the stream and as well as their contending neighbors.

Further, contending nodes need to periodically synchronize their estimate of the network capacity. This may be done by interpreting the queue size as indicators of local contention, but when this contention is asymmetric, explicit signaling is required. All nodes in the network need to understand and correctly interpret these signaling messages. Jain et al. and Raniwala et al. have proposed distributed algorithms to achieve these requirements based on this conflict graph approach. Gambiroza et al. proposed a time-fairness reference model that removes the spatial bias for flows traversing multiple hops, and propose a distributed algorithm that allows them to achieve their fair-rate allocation.

TCP has a link-centric view of congestion that works well for wired networks. In wireless networks, however, congestion is a neighborhood phenomenon, i.e., it is not local to a single node but common to all nodes sharing the radio channel. IFRC is a distributed protocol for a set of sensor nodes to detect incipient congestion, communicate this to interfering nodes, and to follow an AIMD rate control mechanism for converging to the fair-rate. IFRC supports the many-to-one communication paradigm of sensor networks, but fails to identify and signal all set of interfering nodes in one-to-many traffic scenarios (typical downloads in WMNs). Using a clean-slate approach, the authors extend this work and propose WCP, a new rate-based protocol that shares congestion information with interfering nodes and uses synchronized AIMD of flow rates to achieve fairness.

With WCP the rate of a flow is inversely proportional to the number of congested neighborhoods it traverses. We note that this allocation criterion is not consistent with any of the commonly used fairness notions described. Further, WCP identifies the congestion neighborhood of a link based on transmission ranges, while it is known that nodes in the interference range (which is typically larger than transmission range) may interfere with a transmission.

Other recent work proposed rate control measures that also require modifications to the MAC layer. Zhang et al. transform the global max-min fairness objectives into a set of conditions locally enforceable at each node, and propose a rate adaptation algorithm based on these conditions. This work assumes a modified 802.11 MAC that transmits a packet only when a receiver has buffer to store it [6] and [11].

Hybrid Networks: WiMAX and 802.11 WMNs:-

While WiMAX MMR specifications have only recently been ratified, 802.11 radios have become a commodity platform with over 387 million chipset sales reported in 2008 alone. This economy of scale has resulted in 802.11 becoming the preferred radio platform for developing last

mile access networks for a large number of commercial entities. It is expected that as WiMAX matures, the two technologies will be used in a complementary manner, with WiMAX providing a long distance wireless backhaul for a last-mile distribution network which is primarily based on 802.11-based WMNs.

IEEE 802.11s:-

The IEEE 802.11 Task Group s is working on standardizing a set of amendments to the 802.11 MAC to create an Extended Service Set (ESS) of Mesh Points (MPs) that are connected via a multihop Wireless Distribution System (WDS).




Figure 1: Congestion Control Mode Identifier field in a Mesh Beacon announces the supported congestion control protocol.

Congestion Control in 802.11s:-

Intra-mesh congestion control is based on three mechanisms: congestion monitoring and detection, congestion notification, and congestion resolution via local rate control algorithms. The 802.11 standard only provides for a signaling framework for exchanging congestion notification messages; congestion monitoring and detection as well as subsequent rate-control algorithms are considered out of scope. This reflects the fact that while congestion control is necessary, there is no single solution that can be optimized for all the different mesh applications. The standard thus allows for an extensible congestion control framework.

Mesh Points (MPs) can support multiple congestion control protocols; though only a single protocol may be active at a given time in a network. This protocol is identified by the Congestion Control Mode Identifier (CCMI) field that is a part of the mesh beacon (Figure 1). A null value for this identifier means that the network does not support any congestion control scheme. An identifier may be reserved for an open, standardized congestion control protocol or may even be vendor specific. The draft 802.11s standard specifies a congestion control signaling protocol. When an MP detects local congestion, it may transmit a Congestion Control Notification (CCN) frame. This frame may either be sent to neighboring MPs or directed to the MP sourcing the traffic causing this congestion. The recipient may then choose to adjust their frame rate to the MP that sent the CCN frame. Each CCN frame contains Congestion Notification Element (CNE). The default CNE described by the standard contains the estimated time the congestion is expected to last for each of the four access categories in 802.11e (Figure 2); a value of zero specifies that there is no congestion detected for that specific access category.

However, as previously described, the framework is extensible to accommodate CNEs that contain additional information as required by the new congestion control protocols.

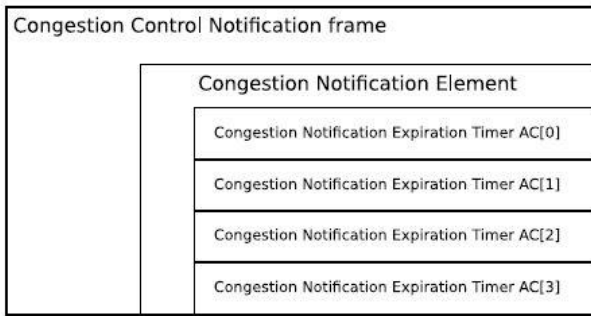


Figure 2: The default Congestion Notification Element in a Congestion Control

Notification frame contains the estimated congestion duration for each of the four Access Categories.

The congestion control framework in IEEE 802.11s draft specifications sacrifices robustness of the protocol in favor of supporting greater extensibility. Unfortunately, this leaves the framework with little functional specifications to ensure true interoperability across various vendors. It also raises the possibility of new security attacks, e.g., selfish nodes may not adjust their rates in response to CCN frames, or may even generate spurious CCN frames to slow down other MPs unnecessarily. There is ongoing work within the standard bodies to address congestion control as a part of 802.11s standards specifications. IEEE 802.11s draft standard only specifies a lightweight congestion control signaling framework. It does not specify as to which nodes should receive this notification or how the recipient should respond [4] and [7] and [10].

PROPOSED TECHNIQUES

Having established the efficacy of gateway nodes in enforcing rate control for adaptive traffic flows, our attention towards designing practical centralized rate controllers, i.e., we wish the computational model based approaches, and instead develop a framework of heuristics that are practically feasible for deployment operating a WMN. We propose Aggregate Rate Controller (ARC) that can enforce approximate max-min rate allocation using only the information locally available at the gateway. ARC manages the net amount of traffic allowed through the network, relying on the underlying max-min fairness characteristics to apportion this capacity fairly amongst all contending flows.

First, we show that distributed bottlenecks can exist even in a WMN where traffic patterns are skewed towards the gateway. Max-min rate allocation allows us to saturate these bottlenecks and efficiently utilize the available capacity. Second, we characterize the response of TCP flows in an 802.11 multi-hop network as a function of the aggregate capacity allowed through its traffic aggregation points. We show that it is possible to achieve approximate max-min flow rates simply by regulating the net amount of data traffic allowed through the gateway to its fair-aggregate capacity. Third, based on this behavior, we propose ARC, an aggregate rate-based scheduler that uses a measurement-based approach to determine this fair-aggregate capacity of a network, leading to approximate max-min allocation across contending flows.

Max-min Fairness in WMNs:-

WMNs used for Internet backhalls have a dominant traffic pattern in which flows are directed either towards or away from the gateway. This increases the spectrum utilization around the gateway, eventually becoming a bottleneck for flows that are not already bottlenecked elsewhere in the network. We know that the max-min algorithm described above assigns equal rate to all flows sharing a common bottleneck. Thus when the spectrum around the gateway constitutes the single bottleneck shared between all flows, max-min fairness results in equal rate allocations. Despite this skewed traffic pattern in a WMN, topological dependencies can create bottleneck collision domains other than those including the links around the gateway. We show how multi-rate links and node distributions create such distributed bottlenecks.

Multi-rate Links Figure 3 shows two variations of a simple chain topology with two backlogged flows f1 and f2 transmitting through a common gateway GW. The link rates are a mix of R and 2R, as shown. An equivalent wired network would yield a rate vector of (f1, f2) = (2R,R) in both of these topologies. With wireless links, however, the rate vector depends on the position of the slower link.

- Link f is the slower link in Figure 3a. Here f1 and f2 are bottlenecked by the collision domain of links c and d, respectively. The slower link is a part of d’s collision domain.




Figure 3: Max-min rate depends on the location of the slower link with capacity R. If this slower link is one of links a, b, c, d, or e, the max-min fair rate vector is (R/3 , R/3). If the slower link is at position f, the max-min rate vector is (f1, f2) = (2R/3 , R/3).

An equivalent wired network would always produce an allocation of (R, 2R), with the lower rate for the flow traversing the slower link.

The max-min rate allocation for this scenario is (f1, f2) = (2R/3 , R/3).

- If the slower link is one of the links b, c, d, or e, the max-min rate allocation is (f1, f2) = (R/3 , R/3). The collision domains of links c and d are fully saturated at this rate and form the bottleneck for flows f1 and f2 respectively.
- If the slower link is a, the max-min allocation is still (f1, f2) = (R/3 , R/3). However, in this case the two flows are both bottlenecked by a common collision domain of link c.

Network Response to Aggregate Rate Control:-

We analyze the flow goodput response as a function of the aggregate network capacity, *i.e.*, if the network is allowed to transport an aggregate of x units of traffic, how does the 802.11 MAC distribute these x units of capacity between multihop TCP flows all contending for channel access? In this work we experiment with managing this aggregate capacity via a single token bucket at the gateway router. We experiment with the two rate limiting mechanisms in Figure 4. The differences between the two architectures were described. To recap, in Figure 4a, all received data is stored in a shared buffer till there are enough tokens to send it out.

This simple architecture has known performance problems: it does not provide any isolation between flows and a shared FIFO queue can cause synchronization between TCP connections carrying bursty traffic. We address these issues by using per-node queuing at the gateway as shown in Figure 4b. This allows us to separate data from different subscribers irrespective of the number of TCP micro-flows a subscriber generates.

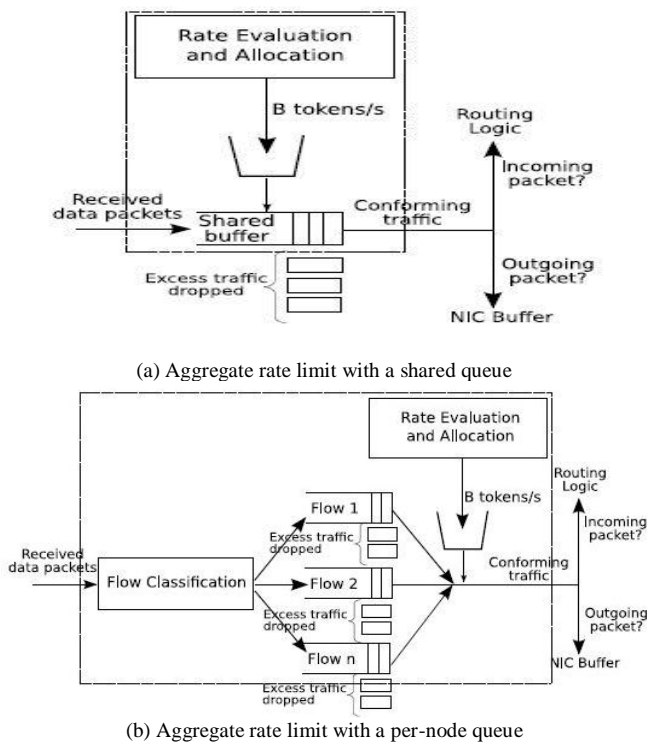


Figure 4: The main architectural components of ARC

Our proposed heuristic ARC uses a simple measurement-based adaptive rate allocation approach. It measures the rate obtained by all flows over a fixed interval called *epoch*. If all flow rates are equal, it assumes the network is underutilized and increases the aggregate capacity allocated at the gateway. If the rates are unequal, then the flows with lower rate are either bottlenecked locally or are experiencing unfairness. We differentiate between the two by maintaining state on historic flow rates. The aggregate capacity allocated at the gateway is decreased only when we suspect unfairness. Our heuristic thus mimics the behavior of adaptive protocols like TCP by probing the network for capacity information and adjusting its behavior in response. This closed-loop feedback system allows ARC to adapt to changing network and traffic conditions.

We now describe the ARC heuristic in detail. ARC is a system module on the gateway mesh router. It sits between the MAC layer and the network layer, operating transparently between them. Its three main components perform the following functions:

1. Flow classification
2. Rate evaluation and allocation
3. Flow rate enforcement

Flow Classification:-

In this first step, ARC performs flow classification for all data traffic (ingress and egress) through the gateway. Here flow refers to any suitable classification of traffic and its precise definition is left as a policy decision for the network operator. In this paper we have classified flows based on the source or destination mesh router. Thus a flow f_i represents the aggregate of all micro-flows originating from, or destined to, node in the network. In this context, we use nodes and flows interchangeably in our discussion. Our classification methodology requires a simple lookup of the packet header given a known offset, and can be performed efficiently. We note that such a classification is consistent with the common practices employed by ISPs on wired access networks, where capacity is managed on a per-subscriber basis.

Rate Evaluation and Allocation:-

Rate evaluation component measures the flow rate of all flows in a given epoch τt . These measured rates determine the aggregate capacity allocated at the gateway for the next epoch $\tau t + 1$. Let the duration of the epoch be δ . This value is configurable, though for stability it should operate at different timescales than the control action of TCP senders. For instance, δ can be set to multiples of roundtrip time so that TCP sources can react to changes in rate allocation and stabilize around their new values.

Algorithm 1: Rate evaluation algorithm:-

```

Input: Epoch duration  $\delta$ , inter-flow unfairness threshold  $\gamma$ , intra-flow rate cutoff threshold  $\theta$ , measured rate  $r_i$  and historical  $r_{i_{max}}$ ,  $\forall i \in N$ ,
Output: Rate increase or decrease decision,
while once every  $\delta$  time units do
if  $r_i/r_j < \gamma$  and  $r_i/r_{i_{max}} < \theta$  for  $\forall i, j \in N$  then
decAggRate;
else
 $r_{i_{max}} = r_i$ ;
incAggRate;
end;
end;
    
```

The rate allocation component may use any number of heuristics to determine the new Cest. It has to search through the space of feasible allocations for the new capacity estimates. A simple algorithm using exponential increase/decrease in aggregate capacity is shown in Algorithm 2.

Algorithm 2

```

Input :  $C_{means} = \sum_{i=1}^N r_i$ 
Output: New aggregate rate limit Cest.
incAggRate begin
Cest =  $\alpha \times C_{means}$ ; where  $\alpha > 1$ 
End;
decAggRate begin
Cest =  $\beta \times C_{means}$ ; where  $\beta < 1$ 
End;
    
```

Flow Rate Enforcement

ARC uses a single token bucket at the gateway to control the aggregate network capacity. The token generation rate B is controlled by the rate allocation mechanism. Note that ARC can be used with either of the two enforcement mechanisms. Our evaluation uses ARC, it provides better isolation between flows leading to improved fairness characteristics.

WMN despite its dominant traffic pattern consisting of flows directed towards or away from the gateway. Max-min rate fairness with its Pareto optimality allows us to efficiently utilize these bottlenecks while maximizing the minimum allocation. We proposed ARC, a measurement-based rate controller that can be implemented at the gateway router and manages traffic as a single aggregate bundle instead of distinct flows. Also, we proposed heuristics for achieving approximate max-min rate allocation through gateway-enforced rate control in a WMN.

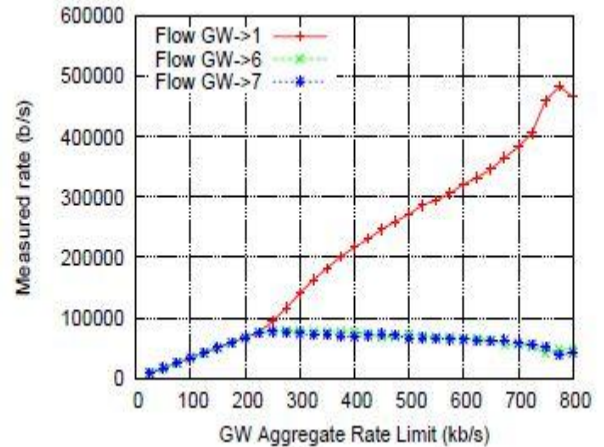
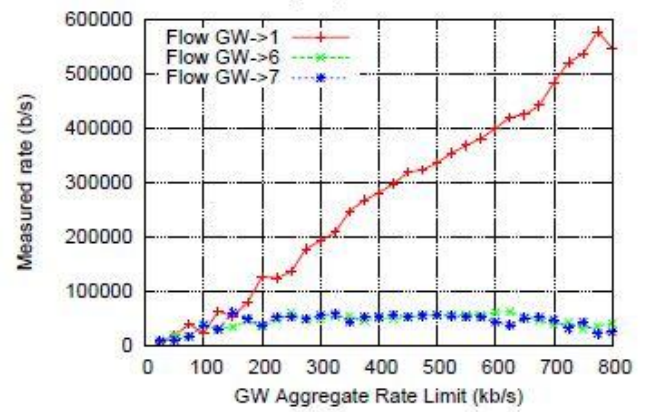
RESULTS

We simulate a number of network topologies with gateway rate limiting the aggregate TCP traffic it bridges between the wired and the wireless network. The measured flow throughput as a function of the aggregate capacity for the network in Figure 3 with 1 Mb/s 802.11 links. Each data point for a given rate limit represents the average of 5 experimental runs. Recall that the optimal max-min fair rate computed for this topology is $(f_1, f_2, f_3) = (R/5, R/10, R/10)$. The nominal capacity of a 1 Mb/s 802.11 link is approximately 800 Kb/s, assuming a perfect channel and no collisions. This translates to $(f_1, f_2, f_3) = (160 \text{ Kb/s}, 80 \text{ Kb/s}, 80 \text{ Kb/s})$ and a fair-aggregate capacity of 320 Kb/s.

We make the following observations from Figure 5. First, both plots show an initial increase in rate of *all* flows with increasing rate limit. These rate increases are prolonged (up to an aggregate rate limit of 225 Kb/s) and more uniform (all rates within 2% of each other up to this rate limit) in Figure 5b where we use a predestination node queue at the gateway. In contrast, Figure 5a shows a deviation of up to about 50% between the maximum and the minimum rate flows between 0–150 Kb/s rate limit. This deviation increases with increasing aggregate rate limit.

Thus separating flows while rate limiting at the gateway shows improved fairness characteristics. Second, increasing the rate limit beyond 225 Kb/s in Figure 5b only increases the rate of flow f_1 while the rate of flows f_2 and f_3 taper off. Between a rate limit of 250–350 Kb/s, for example, f_1 approximately doubles its rate while registering a drop of only 4% in the average rate of f_2 and f_3 . This drop in rate increases with increasing aggregate capacity, *e.g.*, at 400 Kb/s, the drop in average throughput of f_2 and f_3 is approximately 7% compared to the highs seen at a rate limit of 250 Kb/s. Third, at the computed fair-aggregate capacity of 320 Kb/s, the measured rates of flows f_1 , f_2 , and f_3 are all within 10% of their optimal max-min rates computed assuming no collisions.

(a) Shared queue with GW rate limiting



(b) Per-destination queue with GW rate limiting

Figure. 5a and 5b use the enforcement mechanism in Figs. 4a and 4b respectively at the gateway router.

CONCLUSION

Our main contributions in showed that a broad category of router-assisted traffic allocation mechanisms that are designed for *fair* allocation of bandwidth across different flows in a wired network are ineffective in the wireless environment. We identified that this is due to key differences in the abstraction of wired and wireless links. First, work-conserving packet scheduling mechanisms are based on the assumption that links can be scheduled independently. This is not possible in a wireless network where transmission on a given link prohibits any concurrent transmissions on all interfering links. Second, packet-loss in wired networks occurs primarily as queue drops at the congested router. In WMN, there are limited queue-associated drops even at the gateway mesh router that converges traffic across the entire network. Most packet drops are due to collisions that are distributed across the network. Cross-layer interaction between different layers of the protocol stack means that some of these collisions can lead to long-term unfairness and possible flow starvation. Having established the efficacy of centralized rate control, we then proposed heuristics for estimating the *fair* rate allocation for a given set of flows strictly using only the local information available at the gateway. Understanding the challenges associated with extending this work to larger deployments with real world traffic workloads, perhaps on a commercial-grade WMN, is a necessary and a logical next step.

REFERENCES

- [1]. Nessrine Chakchouk and Bechir Hamdaoui, "Traffic and Interference Aware Scheduling for Multiradio Multichannel Wireless Mesh Networks", IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY, VOL. 60, NO. 2, FEBRUARY 2011, pp 555-565.
- [2]. Erwu Liu, Qinqing Zhang and Kin K. Leung, "Clique-Based Utility Maximization in Wireless Mesh Networks", IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS 2011, pp 1-10.
- [3]. Yahya Al-Hazmi and Hermann de Meer, "Virtualization of 802.11 Interfaces for Wireless Mesh Networks", IEEE 2011 Eighth International Conference on Wireless On-Demand Network Systems and Services, pp- 44-51.
- [4]. Ronghui Hou, King-Shan Lui, Fred Baker, and Jiandong Li, "Hop-by-hop Routing in Wireless Mesh Networks with Bandwidth Guarantees", IEEE TRANSACTIONS ON MOBILE COMPUTING-2011, pp-1-15.
- [5]. IEEE LAN/MAN Standards Committee. IEEE 802.16 - Air Interface for Fixed Broadband Wireless Access Systems. IEEE, New York, 2004.
- [6]. IEEE LAN/MAN Standards Committee. IEEE 802.11 - Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. IEEE, New York, June 2007.
- [7]. Vasken Genc, Sean Murphy, Yang Yu, and John Murphy. IEEE 802.16j Relay-based Wireless Access Networks: An Overview. IEEE Wireless Communications, pp56-63, October 2008.
- [8]. Panos Gevros, Jon Crowcroft, Peter Kirstein, and Saleem Bhatti. Congestion Control Mechanisms and the Best Effort Service Model. IEEE Network Magazine, pp16-26, May/June 2001.
- [9]. Kamran Jamshaid and Paul A.S.Ward. Experiences using Gateway-Enforced Rate-Limiting Techniques in Wireless Mesh Networks. In Proc. of the IEEE WCNC '07, pp 3725-3730, March 2007.
- [10]. Ying Jian, Shigang Chen, Liang Zhang, and Yuguang Fang. New Adaptive Protocols for Fine-Level End-to-End Rate Control in Wireless Networks. In Proc. of the IEEE ICNP, pp 22-32, October 2008.
- [11]. Jangeun Jun, Pushkin Peddabachagari, and Mihail L. Sichitiu. Theoretical Maximum Throughput of IEEE 802.11 and its Applications. In Proc. of IEEE International Symposium on Network Computing and Applications, pp249-256, April 2003.