A Closer Look at the Capacity of Wireless Mesh Networks

Bassam Aoun, Youssef Iraqi and Raouf Boutaba

***Abstract*— In this work, we study the *fair* capacity of Wireless Mesh Networks (WMN). We argue that the issue of fairness is fundamentally a network layer problem: *if every node transmits at a rate derived from its fair share of network resources, global fairness can be achieved*. In contrast to other works suggesting modifications of the MAC protocol, we tackle fairness by a network layer, centralized solution. We therefore propose an algorithm for *max-min* capacity calculation, formulated in term of collision domains.**

***Index Terms*— Max-min Capacity, Wireless Mesh Networks.**

I. Introduction

Wireless is well established for narrowband access systems, but its use for broadband access is relatively new. Wireless mesh architecture is a first step towards providing high-bandwidth network coverage. Mesh architecture sustains signal strength by breaking long distances into a series of shorter hops. Intermediate nodes not only boost the signal, but cooperatively make forwarding decisions based on their knowledge of the network. Such architecture provides high network coverage, spectral efficiency, and economic advantage.

The use of WMNs as backbone for large wireless access networks imposes strict bandwidth requirements. It is therefore necessary to study and quantify the capacity of such systems. In this work, we argue that the capacity of WMNs should be addressed in the context of fairness to ensure proper operation of the WMNs. Among the fairness schemes, *max-min* fairness allows fair and efficient use of network resources. We therefore propose a framework for *max-min* capacity calculation. Next, we leverage the properties of *max-min* fairness to propose an algorithm for capacity calculations in term of collision domains.

A. WMN Characteristics

WMNs have a relatively stable topology except for occasional nodes failure or addition. The traffic, being aggregated from a large number of end users, changes infrequently. Practically all the traffic is either forwarded to or from a gateway, while in ad hoc networks the traffic flows between arbitrary pairs of nodes. A tree-based proactive routing scheme is used for traffic forwarding, since it easily allow flows aggregation and would minimize overhead, ensuring an optimal utilization of bandwidth.

Fairness in ad hoc networks has been addressed in various studies; however the above specific characteristics of WMNs motivate the approach presented in this work.

This research is partially supported by NORTEL, NSERC, and CITO.

B. A. author is with the Computer Science Department, University of Waterloo, Waterloo, ON N2L3G1 CANADA (e-mail: baoun@uwaterloo.ca).

Y. I. author is with the Computer Science Department, Dhofar University, Salalah, Sultanate of Oman (e-mail: y_iraqi@du.edu.om)

R. B. author is with the Computer Science Department, University of Waterloo, Waterloo, ON CANADA (e-mail: rboutaba@bbcr.uwaterloo.ca).

II. Problem Motivation

In this section, we illustrate that fairness could be obtained by limiting the transmission rate at each node to its fair capacity. We consider the IEEE 802.11 as the MAC protocol.

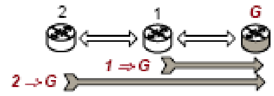

Figure 1: Two simultaneous flows.

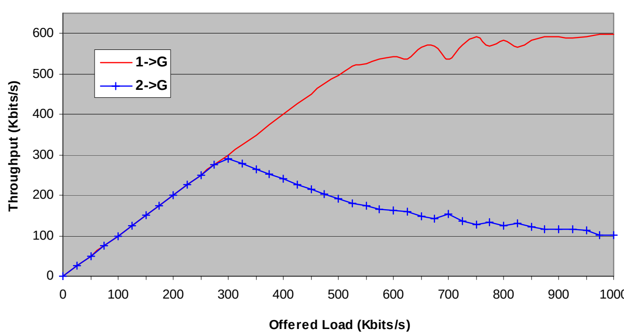

Figure 2: Throughputs of flows in Figure 1.

Figure 1 shows a simple chain topology consisting of two simultaneous flows. Their throughputs are plotted in Figure 2. We can see that if flow 1→G sends at a rate higher than 300 Kbits/s, it would starve flow 2→G. In this situation, 300 Kbits/s represents the *fair capacity* of both flows; beyond that value, the network exhibits unfairness, favoring flow 1→G at the expense of flow 2→G.

In addition to enhancing fairness between different flows, limiting the transmission rate of nodes to their fair capacity reduces packet losses (i.e. dropping) and prevents wasting bandwidth. For example, a node at the beginning of a chain could actually inject more packets than the subsequent nodes can forward. These packets are eventually dropped at later nodes. The time this node spends sending those packets decreases the delivered throughput since it prevents transmissions at subsequent nodes. Figure 4 shows that the throughput degrades as the offered load exceeds the available (called *fair*) capacity, 300 and 200 Kbits/s respectively.

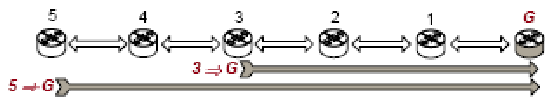

Figure 3: Flows transmitting individually.

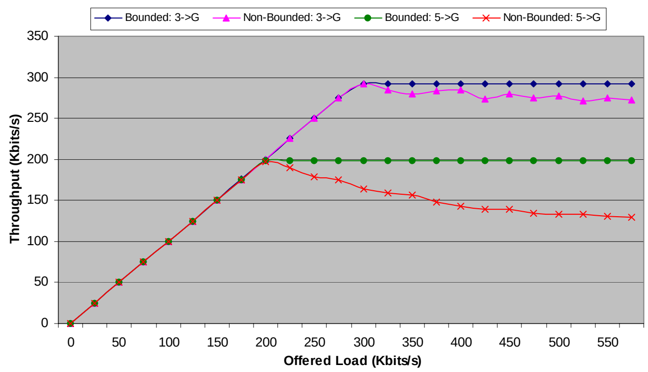

Figure 4: Throughputs of flows in figure 3.

In the following sections, we show how to calculate the *fair* capacity of each node.

III. Framework for WMN Capacity Calculation

A. Wireless Interference and Collision Domains

We recall that, in a wireless network, the resource of interest is not a link but a wireless channel. Neighboring wireless links contend and share the capacity of the local channel. The contention experienced by each wireless link is represented by the load of its collision domain, defined as follows:

Definition 1 (Collision Domain): The collision domain C_e of a wireless link $e \in E$ is the set of neighboring wireless links that share its local channel, and therefore interfere with its transmission. From a protocol point of view, the collision domain C_e of a wireless link is the set of neighboring wireless links that have to be inactive in order to the wireless link to transmit successfully.

B. Fairness Consideration

Capacity of a WMN should be addressed in the context of fairness to ensure proper operation of WMNs. The targeted granularity of fairness is an AP-aggregated flow f_v. In particular, each AP corresponds to a single residence, small business, or hot spot, and this AP traffic f_v should be treated as a single aggregate, independent of the number of TCP micro-flows or mobile devices supported by that AP.

Aggregate flows f_v for $\forall v \in V$ should be treated equally independent of the relative location of v with respect to the gateway; users should not be penalized for not having a nearby wireline Internet connection (i.e. a gateway). Such fairness mechanism is opposed to the *capacity-maximizing* allocation consisting of starving multi-hop flows and giving all the capacity to one-hop flows.

In addition to fairness, we would like to achieve efficient use of network resources. That is, network resources are to be reclaimed by other flows, when they are unused by flows bottlenecked elsewhere.

C. Fairness Reference Model

Among the fairness mechanisms, *max-min* fairness allows the fair and efficient use of network resources and is widely considered in wired and wireless networks. In our study, we assume that *max-min* fairness is enforced among APs' aggregate flows f_v.

We recall that vector $\vec{f}$ is defined such that the i^{th} coordinate is the rate allocation for AP $i \in V$ aggregate flow. We assume that all transmission rates are positive. We present the definition of *max-min* fairness as follows:

Definition 2 (Max-min Fairness): An allocation of rate $\vec{f}$ is *max-min* fair if and only if an increase of any rate i, within the domain of *feasible* allocations, must be at the cost of a decrease of some already smaller rate. Formally, for any other rate allocation $\vec{y}$, if $y_i > f_i$ then there must exist some $j \in V$ such that $f_j \le f_i$ and $y_j < f_j$.

The rate vector $\vec{f}$ is to be calculated, and the sum of its components $\sum_{i \in V} f_i$ is considered the total fair capacity of the WMN.

A *feasible* allocation of rates $\vec{f}$ is such that for $(\forall e \in E_R)$ $\sum_{w \in C_e} l_w \le W \Rightarrow \sum_{w \in C_e} \sum_{v \in V} A_{v,w} \cdot f_v \le W$, where W is the capacity of the wireless channel, l_w is the load on the wireless link w, and $A_{v,w}$ is a binary variable indicating whether the flow f_v crosses the wireless link w.

IV. Max-Min Capacity Calculation

A. Max-min Capacity Properties

Given the context of WMN, we present the following property of *max-min* fairness:

Theorem 1: A feasible allocation of rates $\vec{f}$ is *max-min* fair if and only if every flow has a bottleneck collision domain.

A *bottleneck* collision domain for flow f is a collision domain which is limiting, for the given allocation.

B. Algorithm for Capacity Calculation

In this section, we use the properties of *max-min* fairness to provide an algorithm for capacity calculation. We introduce *N(C,F)* to represent the network resource model of the WMN. F is the set of APs aggregate flows and C is the set of collision domains.

In order to ensure *max-min* fairness, we start by identifying the collision domain with the *smallest capacity available per flow*. Such collision domain is denoted as the *bottleneck collision domain of the network N(C,F).*

Definition 3: Collision domain $d \in C$ is called bottleneck with respect to the network *N(C,F)* if

$$\frac{Cap_d}{\sum_{e \in d} \sum_{f \in F} A_{f,e}} = \min_{s \in C} \frac{Cap_s}{\sum_{e \in s} \sum_{f \in F} A_{f,e}}$$

We now present the iterative *network bottleneck identification* method to allocate flows rate, achieving *max-min* fairness. First, we identify all bottleneck collision domains of *N(C,F)*. We share the capacity of each bottleneck collision domain equally among all flows traversing it. Then we remove these flows from the network, and reduce the capacity of every collision domain by the bandwidth consumed by the removed flows crossing it. The resulting network model is therefore reduced to $N^1(C^1,F^1)$, where F^1 is the remaining flows and C^1 is the updated capacity of collision domains. We now identify the *next* bottleneck collision domain of the reduced network $N^1(C^1,F^1)$ and repeat the procedure. We continue until all flows are assigned corresponding rates.

V. Simulation Results

NS-2 with CMU wireless extensions is used for simulations. The parameters are tuned to the commercially available 802.11-based WaveLan wireless cards. The effective transmission range is 250 meters and the sensing (interference) range is about 550 meters. The simulations involve nodes separated by 200 meters, which allows a node to connect only to neighboring nodes. The bandwidth is set to 1 Mbps and RTS/CTS exchange precedes all data packets. NS-2 is extended to support static routing, more adapted for WMN environments; we therefore eliminate overheads due to routes discovery and maintenance.

In this study we use UDP at the transport layer in order to isolate the effects of the 802.11 MAC protocol. TCP is not adequate in our analysis since it penalizes nodes located at farther hops due to longer round trip time and uses congestion control, starving flows which otherwise would only suffer unfairness when using UDP.

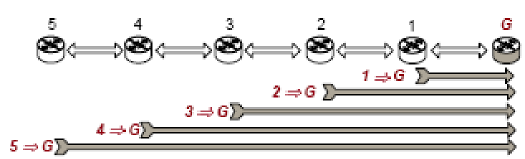

Figure 5: Chain Topology

In this section, we illustrate different concepts of our analysis. Figure 5 represents a chain topology consisting of 5 nodes in addition to the gateway.

Figure 6 plots the throughputs of each flow in Figure 5. Flows are transmitting individually. We note that the throughput scales as 1/*n*, where *n* is the number of nodes along the path to the gateway. The reason is that the load of the bottleneck increases as the number of neighboring transmission increases. However, the throughput stabilizes at *n* = 5, illustrating spatial reuse; transmissions farther away do not affect the bottleneck.

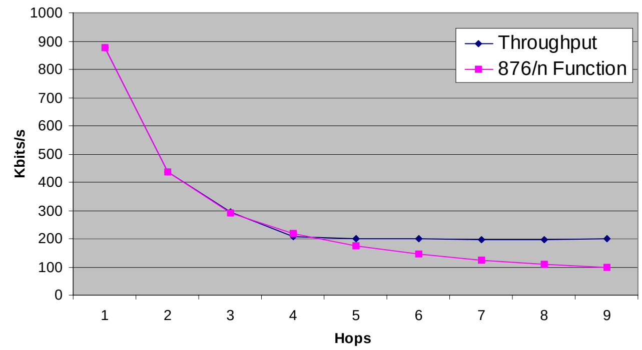

Figure 6: Throughputs of each flow in Figure 5, transmitting individually.

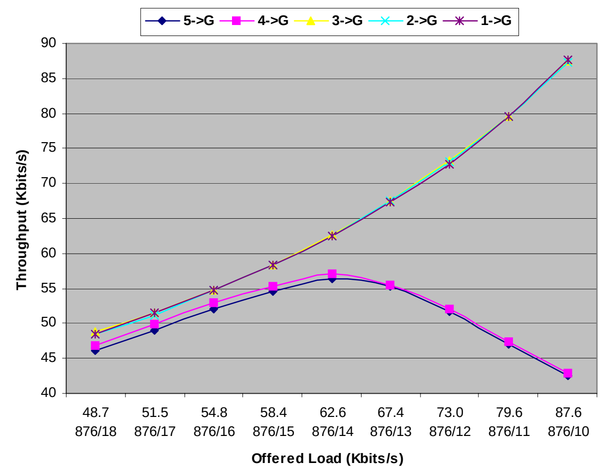

Figure 7: Throughputs of flows in Figure 5, transmitting simultaneously..

In Figure 7, we consider all 5 flows transmitting simultaneously. We note that the network exhibits unfairness when nodes transmit above 62.5 Kbits/s. That value can be obtained using the iterative *network bottleneck identification* method, presented in the previous section. Therefore, to impose fairness, we should limit flows throughputs to their *fair* capacity, derived from network congestion.

VI. Conclusion and Future Work

In this study, we argued that the issue of fairness is fundamentally a network layer problem: *if every node transmits at a rate derived from its fair share of network resources, global fairness can be achieved*. We established a framework for *max-min* capacity calculation and proposed an algorithm in term of collision domains.

The main challenge consists of computing the load of collision domains as it depends on the interference between adjacent nodes. In our study, we considered the IEEE 802.11 protocol. The load of collision domains is therefore derived by considering the coordinated channel access imposed by CTS/RTS. The followings are possible directions for future work. First, it would be interesting to study improvements on the capacity when considering directional antennas. Second, implementing a decentralized resource management algorithm to compute the *fair* capacity of APs is very important for admission control.