

A Measurement-Based Algorithm to Maximize the Utility of Wireless Networks

Julien Herzen[†], Adel Aziz[†], Ruben Merz[‡], Seva Shneer^{*}, Patrick Thiran[†]

[†] School of Computer and Communication Sciences, EPFL

[‡] Deutsche Telekom Laboratories

^{*} Heriot-Watt University

{julien.herzen, adel.aziz, patrick.thiran}@epfl.ch, ruben.merz@telekom.de, v.shneer@hw.ac.uk

ABSTRACT

The goal of jointly providing fairness and efficiency in wireless networks can be seen as the problem of maximizing a given utility function. The main difficulty when solving this problem is that the capacity region of wireless networks is typically unknown and time-varying, which prevents the usage of traditional optimization tools. As a result, scheduling and congestion control algorithms are either too conservative because they under-estimate the capacity region, or suffer from congestion collapse because they over-estimate it.

We propose a new adaptive congestion control algorithm, called *Enhance & Explore* (E&E) [2]. It maximizes the utility of the network without requiring any explicit characterization of the capacity region. E&E works above the MAC layer and is decoupled from the underlying scheduling mechanism. It provably converges to a state of optimal utility.

We evaluate the performance of the algorithm in a WLAN setting, using both simulations and measurements on a real testbed composed of IEEE 802.11 wireless routers.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless Communication*

General Terms

Algorithms, Design, Measurement, Performance, Theory

1. INTRODUCTION

There are some situations where wireless networks fail to achieve both efficiency and fairness. This is the case for instance when the devices have unequal channel conditions and need to use different physical modulation rates. The nodes using slower modulation rates need more time to transmit a packet. In this case, the algorithms that provide some fairness on a per-packet basis (e.g., IEEE 802.11) are inefficient, because most of the airtime is consumed by the slow nodes, whereas the faster ones are unnecessarily throttled. This caveat can already be observed in a very basic scenario where

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

S3'11, September 19, 2011, Las Vegas, Nevada, USA.

Copyright 2011 ACM 978-1-4503-0868-7/11/09 ...\$10.00.

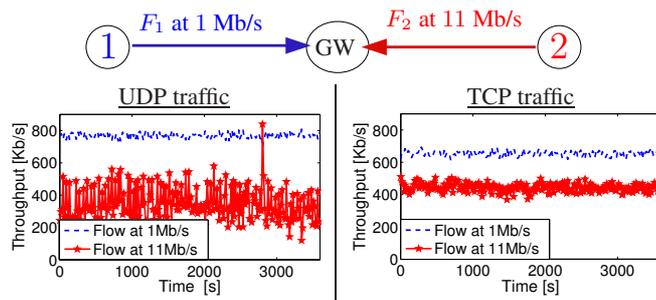


Figure 1: Illustration of the fairness problem occurring when different rates are used within a WLAN. Even though flow F_1 operates at 1 Mb/s and flow F_2 at 11 Mb/s, flow F_2 obtains the lowest throughput. Our measurements show that this rate anomaly problem occurs both with UDP (left) and TCP (right).

two nodes using modulation rates of 1 Mb/s and 11 Mb/s send UDP traffic to a common gateway (see Figure 1). In this case, most of the airtime is consumed by the flow at 1 Mb/s, which gets more throughput than the one at 11 Mb/s. This phenomenon is known as the *rate anomaly problem* [4]. Similar inefficiencies can happen in case of unequal channel conditions, such as when the capture effect occurs.

Depending on the notion of efficiency and fairness that one wants to achieve (for instance proportional fairness would equally share the airtime among the nodes), such a situation is far from optimal.

The problem of achieving both fairness and efficiency in networks can be cast to the problem of maximizing a given utility function [3]. While such a formulation has appealing properties, existing methods to find optimal operation points require the capacity region of the network to be explicitly characterized by a fixed and well defined interference graph. This requirement relies on a model of the capacity region that often requires assumptions that are unrealistic for wireless networks, such as perfect carrier sensing, symmetry of interference among the nodes and the absence of hidden terminal.

We follow a different approach for solving this problem. We take advantage of the fact that the gateway can naturally act as a centralized controller for the network, and we devise a *measurement-based* rate limiting algorithm that does not rely on any form of modeling of the capacity region. The algorithm provably converges to operation points of optimal utility. We validate the feasibility of the algorithm using both ns-3 simulations and a deployment on a real 802.11 testbed. The results show that E&E steadily increases the utility of the network, and converges to optimal rate allocations in practice.

2. THE E&E ALGORITHM

2.1 Overview

The algorithm itself runs at the gateway of the system and acts on the upstream traffic (i.e., from the nodes to the gateway). In addition, it uses a congestion control module that runs on all the nodes of the network, except the gateway, and throttles the rate at which packets are given to the scheduling layer.

If one wants to achieve proportional fairness in the example of Figure 1, the intuition is that the node at 1 Mb/s should send a little less packets in order for the node at 11 Mb/s to obtain a drastic throughput increase. The main challenge with this approach is that the optimal rate allocation is not known in advance, and varies with the network conditions. Indeed, the whole set of rate allocations that the system can sustain (which we denote the *rate region*) is also unknown, and difficult to characterize in practice.

We consider a usual utility maximization problem, where the objective function represents the desired tradeoff between performance and fairness. For instance, achieving proportional fairness amounts to find a rate allocation that maximizes the utility function

$$U_{prop} = \sum_i \ln x_i,$$

where x_i denotes the throughput (service) received by the i -th flow. Similarly, maximizing the sum of the throughputs received by the flows (disregarding fairness) trivially consists in maximizing the utility function

$$U_{max} = \sum_i x_i.$$

More generally, the algorithm applies to any form of utility functions that are expressible as a sum of continuously differentiable, strictly increasing and non-negative functions.

2.2 Enhance & Explore

We provide here the insight behind the E&E algorithm and we refer to [2] for a detailed description. The E&E algorithm uses the fact that the gateway has an accurate knowledge of the throughput received by each flow present in the network. Hence, it knows the utility achieved by the system at any point in time and it can decide new rate allocations that, hopefully, increase this utility. These decisions are broadcasted to the direct neighbors of the gateway and the flows are throttled accordingly by the wireless nodes. If a rate allocation decided by the gateway happens to be infeasible (i.e., it lies outside of the rate region), at least one of the flows does not get its assigned throughput, and the gateway detects it.

The main challenge is thus for the gateway to select rate allocations that maximize a given utility function, *without knowing the rate region*. The procedure itself consists in a succession of gradient ascents (*enhance* phases) and random searches (*explore* phases), and is illustrated in Figure 2 for two flows. The algorithm starts from any rate allocation that is achievable by the scheduling layer and that, by definition, belongs to the rate region (Figure 2 (a)). We denote μ_1 the utility of this allocation. Then, the algorithm attempts a new rate allocation that follows the direction of the gradient of the utility function, which defines a new targeted utility $\mu_2 > \mu_1$ (Figure 2 (b)). If the chosen allocation is equal to the throughputs achieved by each of the flows, the allocation is feasible and thus it is inside the rate region (by definition); in this case the algorithm makes another gradient ascent. If the attempted allocation is not feasible (outside of the rate region), then the algorithm tries another allocation that is randomly chosen but has the same utility μ_2 (Figure 2 (c)). After a few attempts, the algorithm halves

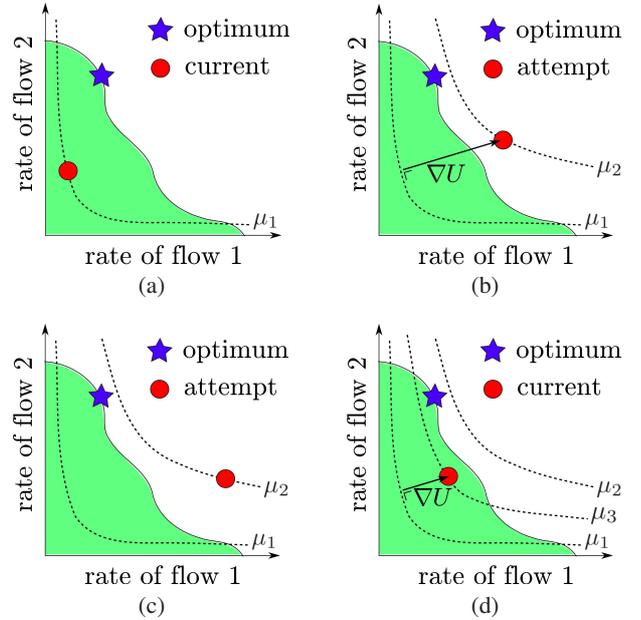


Figure 2: Illustration of the operation of the E&E algorithm for two flows. The shaded area represents the rate region. The dashed lines represent rate allocations that have the same utility (for the utility function of proportional fairness U_{prop}). (a) The algorithm starts from a point having utility μ_1 . (b) The algorithm attempts a gradient ascent to a point having utility μ_2 but this point is outside of the rate region. (c) The algorithm tries another, randomly chosen allocation having utility μ_2 . Again, the corresponding point is outside of the rate region. (d) The algorithm halves the size of the gradient, yielding a feasible allocation of utility μ_3 . Then the procedure repeats. Note that $\mu_1 < \mu_3 < \mu_2$.

the size of the gradient ascent and the procedure repeats (Figure 2 (d)).

2.3 Convergence

An interesting feature of the E&E algorithm is that it does not require the rate region to be convex, but only coordinate-convex¹ (as it is the case for the rate region shown in Figure 2). This is an extremely reasonable assumption for a rate region: given a feasible allocation, decreasing one or several rates still yields a feasible allocation.

The algorithm starts from the throughput achieved by the MAC layer (without congestion control) and gradually increases the utility. One can show that, for a fixed but unknown rate region, the algorithm eventually converges to a rate allocation having the best possible utility (see [2] for a formal proof). In real settings, where the rate region is time-varying, E&E keeps adapting the rate vector in the direction of the best allocation currently sustainable by the system.

¹Let $n \in \mathbb{N}$. A set $S \in \mathbb{R}_+^n$ is coordinate-convex when the following is true: if $\vec{b} \in S$, then for all $\vec{a} : \vec{0} \leq \vec{a} \leq \vec{b}$, $\vec{a} \in S$, with \leq denoting the component-wise comparison.

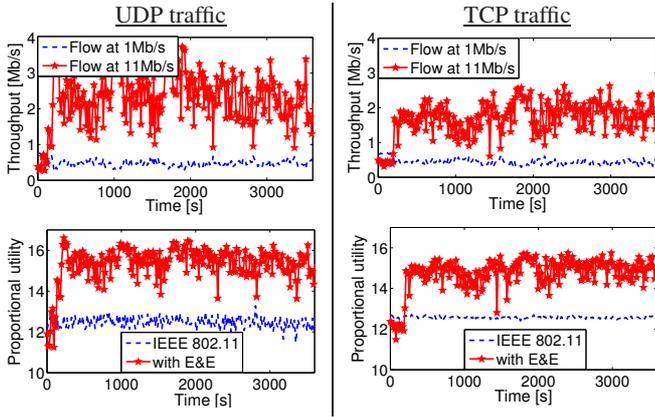


Figure 3: Illustration of the gain in utility achieved by the E&E algorithm in the scenario of Figure 1. The top figures show the new throughputs achieved by the flows with E&E using the utility function U_{prop} . The bottom figures show the achieved utility either with E&E or without E&E (i.e., IEEE 802.11 without rate limitation).

3. RESULTS

We confirm the feasibility and the good performance of our algorithm by using ns-3 simulations and measurements on a real 802.11 testbed deployed on EPFL campus [1]. We use the *Click* modular router [5], which allows us to use the exact same code for the real deployment and the simulations. In both cases, we replayed the scenario of Figure 1, where two nodes use physical modulation rates of 1 Mb/s and 11 Mb/s, respectively.

3.1 Testbed Results

Figure 3 shows the throughput achieved by the system when the E&E algorithm is used, and the corresponding gain in utility. We use the utility function of proportional fairness (U_{prop}) in this scenario.

Both for UDP and TCP traffic, the utility increases from roughly 12 to 15. This difference is substantial as the utility function for proportional fairness is logarithmic: The increase is obtained by dividing the throughput of flow F_1 by a factor 2 and by multiplying the throughput of flow F_2 by a factor 4–7 (hence multiplying their product by a factor more than two).

3.2 Simulation Results

Using simulations allows us to experiment in a more controlled environment. In this case, the rate region is significantly less time-varying and it is possible to empirically sample points on its boundary, by using several combinations of rate allocations for the two flows. From the rate region, we find the rate allocations that provide the optimal utility for the two functions that we consider, namely proportional fairness (U_{prop}) and sum of the throughputs (U_{max}). Figure 4 depicts the measured rate region and the computed optima. It also shows the average throughputs achieved when E&E is running (for 200 s, with UDP traffic), for the two utility functions. We see that E&E adapts extremely well to the chosen utility and that it manages to obtain average throughputs close to the optima.

4. CONCLUSION

The Enhance & Explore algorithm consists in a congestion control module that runs at each node of the network and limits the rate

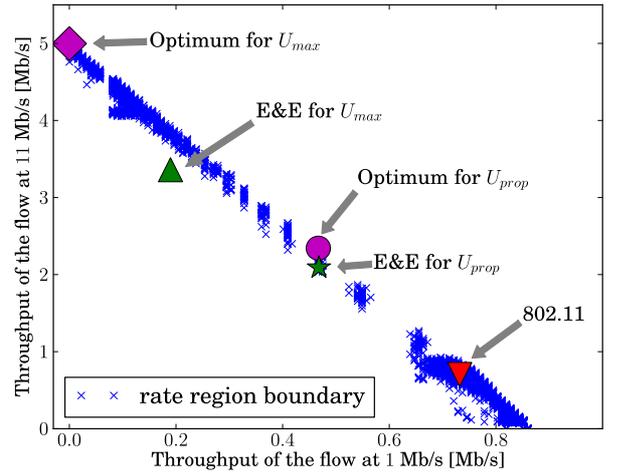


Figure 4: Simulation-based measurements of the rate region, along with the optima for the two considered utility functions (U_{prop} and U_{max}). Also shown are the average throughputs obtained by E&E and 802.11, respectively.

at which the scheduling layer receives traffic. The rate allocations for the different flows are decided by the gateway of the network and periodically broadcasted to its neighbors. This architecture renders possible a measurement-based method that adapts to unknown channel conditions. It does not rely on the usual but unrealistic assumptions that the rate region is convex or that it can be explicitly determined as a function of a well-defined interference graph. Yet, the algorithm provably converges to an optimal rate allocation. A notable feature of E&E is that it starts from the allocation given by the underlying MAC layer, and continues to improve the utility from this point. This ensures that, even if an optimal allocation might be reached only for long-lived flows whose duration exceeds the transient phase of the adaptive algorithm, it is not at the expense of short-lived flows. Its feasibility is demonstrated in practice on a real testbed composed of 802.11 nodes and ns-3 simulations. The evaluations confirm the good performance and the practicality of the algorithm.

5. REFERENCES

- [1] A. Aziz, A. E. Fawal, J.-Y. L. Boudec, and P. Thiran. Aziala-net: Deploying a scalable multi-hop wireless testbed platform for research purposes. In *Mobihoc S'3 2009*, <http://icawww1.epfl.ch/aziala/>.
- [2] A. Aziz, J. Herzen, R. Merz, S. Shneer, and P. Thiran. Enhance & explore: an adaptive algorithm to maximize the utility of wireless networks. In *Proc. of MobiCom*, Las Vegas, NV, Sept. 2011.
- [3] M. Chiang, S. Low, A. Calderbank, and J. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. *Proceedings of the IEEE*, 95(1):255–312, jan. 2007.
- [4] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance anomaly of 802.11b. In *Proc. of Infocom*, 2003.
- [5] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The click modular router. *ACM Trans. Comput. Syst.*, 18(3):263–297, Aug 2000.