A Performance Evaluation Framework for Fair Solutions in Ad Hoc Networks

Manos Dramitinos, Rémi Vannier
INRIA, LIP, RESO Team
46, Allée d'Italie 69364 Lyon, France
Emmanouil.Dramitinos@ens-lyon.fr, Remi.Vannier@ens-lyon.fr

Isabelle Guérin Lassous
Univ. Claude Bernard - Lyon I
INRIA, LIP, RESO Team
46, Allée d'Italie 69364 Lyon, France
Isabelle.Guerin-Lassous@ens-lyon.fr

ABSTRACT

Fairness in multihop ad hoc networks has received considerable attention in the literature. Many schemes have been proposed, which attempt to compute the "optimal" bit rates of the transmitting mobile nodes so that a certain fairness criterion is met. As the related literature indicates, there is a trade-off between fairness and efficiency, since fairness schemes typically reduce the channel utilization. Also, it is questionable whether certain fairness schemes have a positive or negative impact on the QoS of certain user services. So far, there has been limited research on the impact of the varying short-term allocations of these protocols, due to their inherent features and also nodes mobility, on the user-perceived QoS (and social welfare) for services of long duration. In this paper, we introduce an assessment framework, based on *history-dependent utility functions* that can be used as a holistic performance evaluation tool of these fairness schemes. These functions quantify the satisfaction that the ad hoc users obtain from the way their long-lived service sessions are allocated bandwidth, due to the behavior of the MANETs fair schemes. This way we can unambiguously compare the performance of various fair solutions whose maximization goals are inherently different (max-min fairness, proportional fairness, etc.). Finally, we demonstrate the usefulness of this framework by applying it on different protocols.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communication* ; C.2.5 [**Computer Systems Organization**]: Computer Communication Networks—*Local and Wide-Area Networks* ; C.4 [**Performance of Systems**]: Measurement Techniques, Modeling Techniques, Performance Attributes; I.6.6 [**Computing Methodologies**]: Simulation and Modeling—*Simulation Output Analysis*

MSWiM'09, October 26–29, 2009, Tenerife, Canary Islands, Spain.

General Terms

Measurement, Performance

Keywords

Ad hoc, fairness, QoS, social welfare, utility function.

1. INTRODUCTION

Mobile ad hoc networks (MANETs) are self-configuring wireless networks of mobile nodes, the union of which form arbitrary topology. The nodes that also serve as routers are free to move arbitrarily; thus, the network's topology may change rapidly and unpredictably. MANETs can be used by their nodes to exchange content or acquire Internet access via the dynamic network topology. Due to its simplicity and its commercial availability, most of these networks are based on the underlying wireless technology IEEE 802.11. Some works have shown that this use raises issues in terms of efficiency and fairness [6]. Different protocols have been proposed to improve the unfair way that the various flows are allocated bandwidth in the standard IEEE 802.11 protocol (see for instance [3, 15, 18, 22]).

Most of these studies evaluate the performance of the proposed solutions on static scenarios: each scenario is evaluated under a given topology (that can be a random one) with fixed pairs communicating throughout the simulation time according to a given traffic. The studied parameters are very often the throughput per flow, the overall throughput, the cumulative distribution function of the per-flow throughput, the product of the flow throughputs as the Jain fairness index. These metrics are computed on average over the whole simulation or over several runs. Sometimes (but rarely), the short-term fairness is also evaluated via, for instance, the average Jain fairness index with a sliding window method or the distribution of the number of inter-transmissions.

But in a long time period, traffic patterns and/or the network configuration may vary over time, partly due to the nodes mobility and the traffic evolution. Thus, it is very likely that flows, penalized under a configuration during a given time period, do not encounter any fairness issues during other time periods. Moreover, average values (on throughput or on fairness index) give no clue on the volatility of flows throughput over time nor on the impact of the proposed schemes on different parameters that are also of some importance, like delay for instance, on the users services. Studies on short-term fairness give an idea on delays perceived by the stations for the channel access, but the way these delays are distributed over time cannot be deduced from these evaluations. These evaluation metrics indicate how the network is used but do not precisely reflect how users perceive the flows quality. But, users are very sensitive to the quality of service attained, while indifferent to the underlying network protocols and their optimization. Therefore, we argue that it is meaningful to perform a higher-level evaluation of the fair protocols in ad hoc networks, where a performance metric should reflect how "well" or "poorly" user flows were served by the network under a certain fairness scheme. This is the novelty of our approach.

The previous discussion motivates the use of *utility functions*, which can serve as a common ground of comparison of the performance of various fairness schemes whose maximization goals are inherently different (e.g. max-min fairness, proportional fairness, etc.). In this paper, we propose a utility-based framework for the *performance evaluation* of ad hoc fair protocols, based on the definition of history dependent utility functions pertaining to the various services. The performance metric of our approach is *social welfare*, i.e. the sum of all users' utilities. We argue that our methodology, as opposed to the existing evaluations of fair protocols[1] for ad hoc networks, accurately reflects how well users are served under a fair scheme. We also provide related simulation results thereof.

The remainder of this paper is organized as follows: In Section 2, we briefly overview related work and introduce our utility-based framework. Section 3 contains a classification of user services with respect to their sensitivity to various QoS parameters; we also define the utility functions for the services of these classes and discuss their properties and fitness to the assessment we wish to perform. In the penultimate section of the paper, we demonstrate the proposed framework's usefulness, by comparatively assessing the performance of the fairness scheme of [18] with that of standard IEEE 802.11 and providing some experimental results thereof. Section 5 contains directions of future research and some concluding remarks.

2. RELATED WORK AND OUR APPROACH

2.1 Related work

The works that propose solutions to the fairness issues in ad hoc networks usually base their performance evaluation on flows throughput (see for instance [3, 15]). Most of the tested scenarios are static: it means that the topology is given and the communicating source-destination nodes and respective flows are fixed and do not change during the simulation. For a general evaluation (and not restricted to one scenario), the flow throughputs are averaged over different runs. Even if a confidence interval is given, these averages give no indication on how the rates fluctuate *over time*. The short-term fairness studies, like the one in [4], give an indication on the delays perceived by the stations for the channel access via the average Jain fairness index with a sliding window method or the distribution of the number of inter-transmissions. But once more, average values or even a distribution do not capture the realization of the rates fluctuation, which greatly affects user satisfaction for many services. Finally, the Jain fairness index, often used in these studies, requires to know the targeted fair allocation. This last point is tricky because the most appropriate fair allocation for ad hoc networks is not known *a priori*, and it is difficult to know if it is better to compare the rates achieved with the proposed solution with an efficient allocation (that optimizes the overall throughput) or a max-min allocation or a proportional allocation or with an allocation in which all the flows have the same rate for instance. Therefore, this kind of evaluations do not measure the user satisfaction.

Since the seminal work of John von Neumann and Oscar Morgenstern [19], utility functions have been extensively used in economics, decision-making, game theory, grid and computing systems (see [7, 17, 20] and references therein). There have been some research efforts to apply utility functions in order to measure user satisfaction in wireless networks [11, 12, 14, 16]. The utility functions used in these works, depicted as Figure 1, quantify the benefit that users obtain from being allocated a certain bit rate at a given time. These benefit values are not correlated and if the rate fluctuates, the utility values will also fluctuate according to the utility function regardless of the rates achieved in the past. Thus, this approach neglects the severe impact of important QoS parameters, such as delay, on the user-perceived quality for services of long duration.

2.2 Our approach

We believe that assessing the impact of the varying over time rate allocations of the ad hoc fair schemes upon the resulting user-perceived quality and the social welfare attained is both of high importance and an open research issue. *History-dependent utility functions* are an extension of standard utility functions for expressing the value attained over a long time scale from receiving various levels of quality at short time scales. This short time scale is henceforth referred to as *slot*. The slot is defined as the time interval over which the number of packets delivered per service flow are counted so that the average rate received by each flow in this interval is computed and subsequently inputted to the respective service utility function, so that the user-perceived quality of this flow is quantified. The value of the slot depends on the service type, due to the difference tolerance of different services to quality degradation (e.g. high delay); hence it is different per service type.

The main merit of history-dependent utility functions is that *multiple* quality parameters such as the vector of instantaneous bit rates, delay and/or total quantity of resources allocated impact the values of the *correlated marginal utilities* and the overall expected level of users' satisfaction. We use the term "marginal utility" to denote the additional utility attained over each slot of the user's service session. Thus, these utility functions can accurately quantify the time-varying user-perceived quality. This kind of utility functions have originally been proposed in [8] for auction-based resource allocation in UMTS networks and subsequently used elsewhere [9, 13]. A detailed explanation of the differences between the utility functions we propose in this article and the ones of [8] is provided in Subsection 3.3.

In this paper, we propose a 3-tier framework, depicted as Figure 2. This framework uses history-dependent utility functions so as to quantify the satisfaction of the ad hoc users from the way their services are allocated bandwidth. In particular, we use a classification of user services, based on the QoS parameters that are of importance, in order to define a utility function and a slot value for each service

[1]The terms fair protocol and fairness protocol or scheme are used interchangeably throughout the paper.

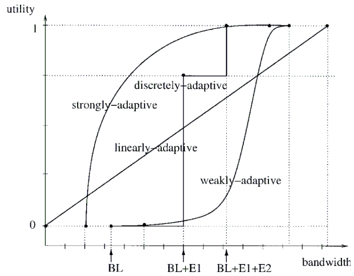

Figure 1: Typical utility functions. Source: [12].

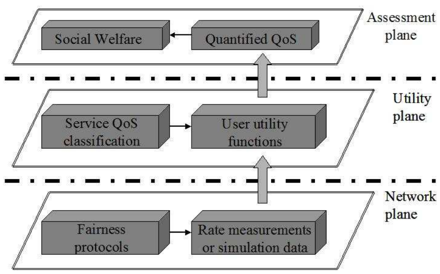

Figure 2: The proposed utility-based framework.

class. These utility functions are additive, i.e. defined as the sum of marginal utilities attained at each slot. The values of these marginal utilities are *correlated*, i.e. they are not independent and depending on the service type and the treatment of the user's flow by the ad hoc network, they vary so as to express the user's satisfaction for the quality of service experienced over time.

In order to comparatively assess the performance of ad hoc fair protocols, the vector of the rates of the user flows computed over slots (either by real measurements or simulations - Network plane) are inputted to the framework's utility functions (Utility plane). The latter quantifies and outputs both the per-flow user-perceived quality and the social welfare (Assessment plane), which are the performance values of the protocols under investigation. Social welfare is a widely used performance evaluation metric, also indicative of the acceptance of the proposed fairness schemes in practice: schemes resulting in low social welfare are expected to be of limited economic value and acceptance by the users. Note that this framework is a performance evaluation framework that has no impact on the protocols' running. It does not try to optimize the fair protocols, but it evaluates them with new metrics different from the ones currently used in ad hoc fairness studies.

We claim that our approach is novel and can serve for the ad hoc research community as a performance evaluation tool of the various ad hoc fairness protocols. This way we can unambiguously compare the performance of various fair solutions whose maximization goals are inherently different (max-min fairness, proportional fairness, etc.). In this paper, we demonstrate this by utilizing our framework in order to perform a comparative assessment of the fair scheme of [18] and the standard IEEE 802.11.

3. THE UTILITY PLANE

This section presents the Utility plane, which is the core of the performance evaluation framework we propose.

3.1 Services and QoS

Prior to proceeding with the definition of the user utility functions, we focus on the various types of user services and their corresponding QoS requirements. There has been already substantial work in the context of UMTS networks by the 3GPP, in identifying the important QoS parameters of services and classifying them thereof. In particular, 3GPP report TS 23.107 [1] defines four QoS classes, namely *conversational*, *streaming*, *interactive* and *background*.

Conversational and streaming classes both pertain to delay-sensitive services such as voice and real-time audio/video streaming respectively. The prominent QoS factors for both these classes are the *guaranteed bit rate* and the *delay*. In particular, the maximum delay tolerance of voice services is $100msec$, while it is $250msec$ for streaming applications [1]. Interactive and background classes are well-suited to throughput-sensitive applications, such as Web browsing and email (or downloading) respectively, for which the quantity of accumulated data is of prominent importance.

In our work, we use these QoS classes and their respective dominant QoS attributes for the definition of the framework's utility functions, provided in the next subsection.

3.2 Utility functions definition

Without loss of generality, we assume that the utility $u_{s,i}$ that user i attains from service s is the sum of the marginal utilities attained at every slot t due to the vector of bit rates allocated $< x_i^{(1)}, ..., x_i^{(t)} >$ up to slot t. Thus,

$$u_{s,i}(x_i^{(1)}, ..., x_i^{(t_{s,i})}) = \sum_{t=1}^{t_{s,i}} v_{s,i}^{(t)}(x_i^{(1)}, ..., x_i^{(t)})$$

In this formula, $t_{s,i}$ is the duration of the user's service session and $x_i^{(t)}$ is the bit rate allocated to user i at slot t. Depending on the class of each service, we need to define an appropriate form of $v_{s,i}^{(t)}(x_i^{(1)}, ..., x_i^{(t)})$. For brevity reasons we henceforth denote $v_{s,i}^{(t)}(x_i^{(1)}, ..., x_i^{(t)})$ as $v_{s,i}^{(t)}$.

3.2.1 Streaming class

The services which comprise the streaming class are delay and rate sensitive, such as real-time streaming audio/video. For these services, we assume that there is a minimum bit rate r_0 under which the quality of the service is unacceptable for the user. This assumption has been verified by works on subjective QoS in wireless networks [21], [10], [5].

We define $v_{s,i}^{(t)}$ for this class as:

$$v_{s,i}^{(t)} = \begin{cases} [v_0(s,i,t) + \Delta V_s \cdot f(x_i^{(t)}, tp_i)] \cdot \alpha_s^{d_i} & \text{, if } x_i^{(t)} \geq r_0 \\ 0 & \text{, if } x_i^{(t)} < r_0 \end{cases}$$

We denote the corresponding marginal utility from being serviced constantly with a rate r_0 as v_0[2]. Therefore, if the average economic value of - and subsequently average user willingness to pay for - a streaming service of a rate r_0 and duration T is wtp, then $v_0 = \frac{wtp}{T}$. Any bit rate less than r_0 results in zero marginal utility for the user. Otherwise, the marginal utility consists of two parts: one part that reflects the user's satisfaction from being served with a rate of at

[2]For both r_0 and v_0 we omit the service subscript s for clarity reasons.

least r_0 and another part that measures the satisfaction of being served more than r_0.

The first part $v_0(s,i,t)$ is defined as:

$$v_0(s,i,t) = \begin{cases} v_0 & \text{, if } t = 0 \\ \frac{1}{\sqrt{d_i}} \cdot v_0 & \text{, if } d_i > 0 \\ \min\{v_0, v_0(s,i,t-1) + \beta \cdot v_0\} & \text{, if } d_i = 0 \end{cases}$$

where $\beta \in (0,1)$ and d_i is the distance between the current and the last previous slot during which a rate at least equal to r_0 was allocated to user i. $v_0(s,i,t)$ fluctuates over time rather than being constant. This way, the marginal utility has a "memory" of the quality degradation experienced over several slots and fluctuates accordingly.

Any additional quantity of bandwidth results in extra value $\Delta V_s \cdot f(x_i^{(t)}, tp_i)$ that depends on both the quantity of bandwidth $x_i^{(t)}$ allocated, as well as whether the user type tp_i is (a) discretely-adaptive, (b) linearly-adaptive, or (c) strongly-adaptive (see Figure 1). ΔV_s denotes the extra satisfaction of the user if she was awarded the entire network channel capacity C, while $f(.)$ is an increasing function of bandwidth whose form depends on the user's type and whose range is $[0,1]$. C is computed as being the theoretical maximum throughput that can be achieved by a saturating one-hop flow. Of course, this capacity would unlikely be achieved in practice in multihop networks. For the definition of the function f, we use some standard functions from the literature [12, 14]:

$$f(x_i^{(t)}, tp_i) = \begin{cases} \frac{x_i^{(t)} - r_0}{C - r_0} & \text{, if } tp_i = \text{linearly-adaptive} \\ \frac{log(x_i^{(t)} - r_0)}{log(C - r_0)} & \text{, if } tp_i = \text{strongly-adaptive} \end{cases}$$

Finally, for discretely-adaptive users, we define $f(x_i^{(t)}, tp_i)$ to be a step function whose steps are the points of the linearly-adaptive utility function, taken from $x_0 = r_0$ and for every *step* kb/s. Thus, the step utility function of this type of users is a discretization of the utility function used for linearly-adaptive users.

The marginal utility (consisted of $v_0(s,i,t)$ and $V_s \cdot f(x_i^{(t)}, tp_i)$) is also multiplied with the $\alpha_s^{d_i}$ coefficient where $\alpha_s \in (0,1)$. Therefore, if the user's service is disrupted frequently, thus resulting in many positive d_i throughout its service session, a large number of marginal utilities will be reduced with the $\alpha_s^{d_i}$ coefficient so as to reflect the resulting degradation of QoS due to the incurred service interrupts (i.e. application content "freezes"). Indeed, if a user session is frequently interrupted, the value attained from watching video "snapshots" in non-consecutive slots is very low. Note that $\alpha_s^{d_i}$ captures the fact that instantaneous service after a big gap is of low value whereas $v_0(s,i,t)$ fluctuates on a bigger time scale so as to reflect the satisfaction of the user due to the average service provided on a bigger time window. To sum up, $v_0(s,i,t)$ captures better the "mean rate" allocated to the application while α_i^d captures the "burstiness" of the service interrupts.

As already explained, the slot value, i.e. the length of time over which the mean application rate is to be computed and then inputted to the utility function, is $250msec$. Therefore, for a service session of e.g. $60sec$, a vector of 240 rate values are computed, each computed over $250msec$ of the session and then inputted to the aforementioned utility so as to quantify the user-perceived QoS.

Concluding, this utility definition captures the effect of dissatisfaction stemming from both a) low bit transfer rates, since this results in both zero marginal utility and a reduction of the value of $v_0(s,i,t)$, and b) the frequency and duration of these service interrupts, which affect the values of the delay of the content delivery of the service, by multiplying the value of the marginal utility with the $\alpha_s^{d_i}$ coefficient whenever satisfactory service is resumed. It is also worth noting that the extra units of bandwidth allocated on top of r_0 result in different fractions of ΔV_s, depending on user's type, so as to cover all possible types of users.

3.2.2 Conversational class

As explained in Section 3.2, the services of the conversational class are affected by the same QoS parameters as the streaming class services. The difference is that the impact of delay results in faster and higher degradation of the user-perceived QoS. Therefore, it suffices to use the utility function definition of the streaming class modified only so that a steeper reduction of the marginal utility $v_0(s,i,t)$ occurs in cases of service interrupts. Therefore, $v_0(s,i,t)$ for the conversational class is defined as follows:

$$v_0(s,i,t) = \begin{cases} v_0 & \text{, if } t = 0 \\ \frac{1}{d_i} \cdot v_0 & \text{, if } d_i > 1 \\ \min\{v_0, v_0(s,i,t-1) + \beta \cdot v_0\} & \text{, if } d_i = 0 \end{cases}$$

In this class, we remind the reader that the slot value is $100msec$.

3.2.3 Background and Interactive class

Background and interactive class pertain to throughput-sensitive services, such as data downloading. We thus define the respective marginal utility as the product of the rate allocated at each slot times a constant utility normalization coefficient v_0: $v_{s,i}^{(t)} = v_0 \cdot x_i^{(t)}$. For a slot value of $1sec$, v_0 is the per second value of being allocated $1kb/s$.

3.3 Discussion

The utility functions definitions of this section are essentially not unique. This is is not important for the assessment our framework performs, as long as a) these definitions are rational, i.e. reflect the impact of the QoS parameters of importance on user utility and b) the *same utility definitions* are used for the comparative assessment of the protocols. Thus, the precise utility function definition and the resulting absolute values of the user-perceived QoS obtained by inputting the allocated rates under these protocols to the respective utility functions are not important, while their ordering is. This ordering is insensitive to the actual utility functions definitions, as long as the aforementioned conditions (conditions a) and b)) are met, and depicts which schemes perform better and for what kind of services, thus providing insight to the performance of the various ad hoc fairness schemes.

The proposed utility functions of the streaming and conversational classes are affected by a multitude of parameters, as opposed to those of the background and interactive classes which are only throughput-sensitive. Note that for the former service classes, the entire history of rate allocations of each service session are taken into account by affecting the values of both $v_{s,i}^{(t)}$ and $v_0(s,i,t)$. These marginal utilities are monotonically increasing in terms of the rate allocated and monotonically decreasing with respect to d_i. Furthermore, the relative reduction of the user utility due to service degradation is higher for the conversational class than for the streaming class, due to the latter's higher tolerance in delay. Thus, these utility functions provide an unambiguous ordering of users' preferences and subsequently of the value from being served under a certain fairness scheme, which is also in line with [1]. These facts, combined with the result of utility theory that any linear combination of utilities is order-preserving ensures that the ordering of the social welfare values is unambiguous even if multiple types of service flows, i.e. flows belonging to different service classes, are simultaneously served.

Note that, for the conversational and streaming class services, we strictly bound the maximum marginal utility of a user by $2 \cdot v_0$. Thus, a high utility v_0 is obtained for each of these services if the minimum rate r_0 is allocated, while the allocation of additional units (on top of r_0) results in lower value than v_0. Therefore, the social welfare attained over a slot if an additional user gets served even with the minimum rate r_0 is higher than the one attained if the same quantity of bandwidth was given to a set of users already being served. This is clearly in line with the fairness rationale. Moreover, the proposed utility functions by construction take into account both the volume and the variance over time of the rates allocated to each service throughout its session, so as to express the impact of both channel utilization and variance of rate allocations on the user-perceived QoS. A precise setting of the v_0 and α_s can be performed by means of applying standard methods of experimental economics, such as the Becker-DeGroot-Marschak mechanism [2] or the Mean Opinion Score (MOS); this is beyond the scope of this paper and does not affect the comparison we perform.

The value of the social welfare under different traffic mixes and conditions indicates the sensitivity of the applications served to the underlying protocols behavior. The impact of admitting more flows or altering the nodes topology and mobility is depicted in the difference of values of the social welfare attained in these cases. Also note that under different loads, the degree of fairness of a scheme may be beneficial or not, depending on whether r_0 can be provided or not. Hence, the proposed methodology is capable of providing better qualitative information and insight regarding the actual QoS experienced from the user services, as opposed to the average network metrics or the Jain fairness index.

Finally, we state the main differences of this paper's utility functions with those of [8]. In particular, the utility functions of [8] are tailored for auction-based resource allocation of a 3G network where the base station has full control on the service specification, service rates and allocations, as opposed to ad hoc networks where no central coordination point is present. Thus, the utility functions of [8] are defined for discrete rates, corresponding to the encoded rates of applications, and are not continuous, as opposed to this paper's definitions where also multiple types of users are considered. Also, in [8], the only feasible allocations for a service within a slot are the values 0 and r_0, which is clearly not the case in the present paper. Also, due to the fact that in [8] an auction for resource allocation is proposed, this mandates a common slot duration for all service types in [8], as opposed to this paper. Finally, the utility functions of [8] are more simplified, due to the fact that they are also used as bidding functions, so that the auction can be repeated in fast time scales, while more expressive - but also more complicated - utilities are used in this paper to assess the performance of ad hoc fairness schemes.

4. USEFULNESS OF FRAMEWORK

We demonstrate the usefulness of our framework by using it to comparatively assess the performance of the scheme of [18] and standard IEEE 802.11. Note however that it can be used for any scheme applying fairness in any network layer. The scheme of [18] was designed to solve some of the fairness issues of the IEEE 802.11 standard. By exchanging rate information at the routing layer, this scheme computes the rates achievable by the MAC layer so that all flows are granted a part of the capacity according to a proportional fairness rationale. An extensive comparative assessment was conducted. The ns2 simulator version 2.33, the protocol implementation code of [18] and a set of Perl scripts implementing the Utility and Assessment planes of the framework (see Figure 2) comprise the software used for this assessment.

Each experiment is conducted for T seconds over a square $S \times S$ terrain. For all protocols assessed, each set of experiments regards the same set of flows crossing the network, whose origin and destination are randomly selected. The first second of transmission of each flow is randomly selected from a uniform distribution having support in $[1, T]$. This way, we are able to capture the effect of admitting more flows within each experiment, so as to study the protocols robustness and performance under varying network conditions. Streaming flows are simulated as UDP/CBR flows, while FTP flows use TCP as the transport protocol. The parameters for ns2 are depicted as Table 1, while the utility functions parameters are $v_0 = 10$, $\alpha_s = 0.97$, $\beta = 0.1$, $\Delta V_s = v_0$, $r_0 = 128kb/s$, $C = 5Mbps$.

Table 1: The simulation parameters.

NS version	2.33
Physical rate	11 Mbps
Real throughput	5 Mbps
Transmission range	200 m
Carrier sensing range	397 m
Capture threshold	10 dB
Radio propagation model	TwoRayGround
Routing protocol	AODV
Packet size	1000 bytes
RTS/CTS	disabled

The first set of experiments regards the transmission of 10 video flows transmitting with $384kb/s$ over a MANET of 25 nodes randomly deployed on a $500m \times 500m$ terrain. Each experiment is run for $T = 100sec$. The average values of the social welfare computed over 10 simulation runs, as well as the confidence intervals, are depicted as Table 2; the large confidence interval values are mostly due to the varying nodes placement and number of the multi-hop flows of the MANET among different runs of the simulation set.

The scheme of [18] outperforms IEEE 802.11. More importantly, this ordering in the social welfare values was observed not only in the averages provided in Table 2, but also for *all* individual runs comprising this set of simula-

Table 2: Average social welfare for 10 video flows.

User Type	*Average social welfare*	
	IEEE 802.11	Scheme of [18]
Linearly-adaptive:	14788 [10843,19559]	15101 [11757,20300]
Discretely-adaptive:	14086 [10344,18574]	14458 [11375,19366]
Strongly-adaptive:	23041 [16829,30697]	23252 [16987,31598]

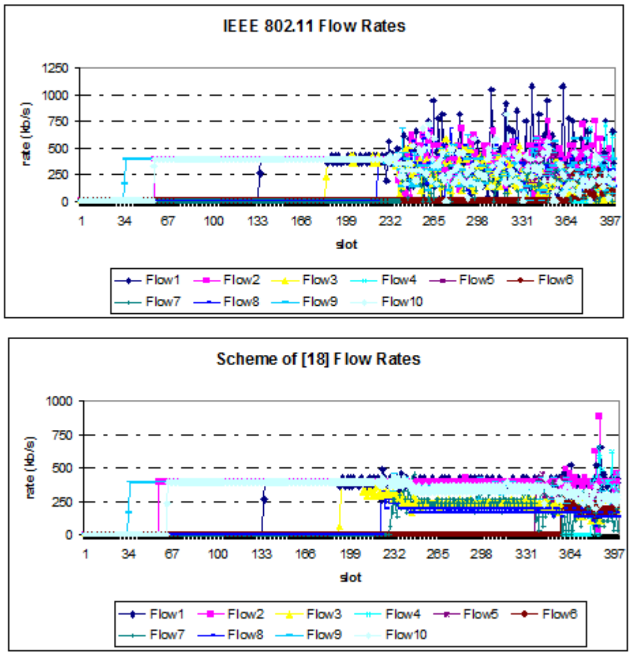

Figure 3: The flows rate allocations for a MANET of 30 nodes and 10 video flows.

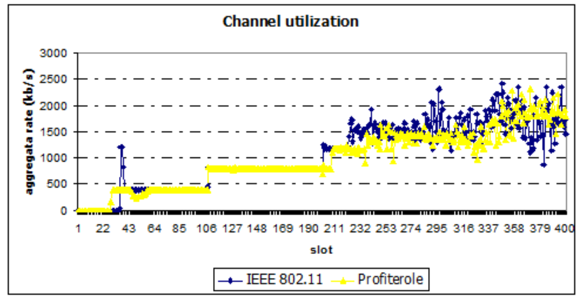

Figure 4: The (aggregate) channel utilization for a MANET of 30 nodes and 10 video flows.

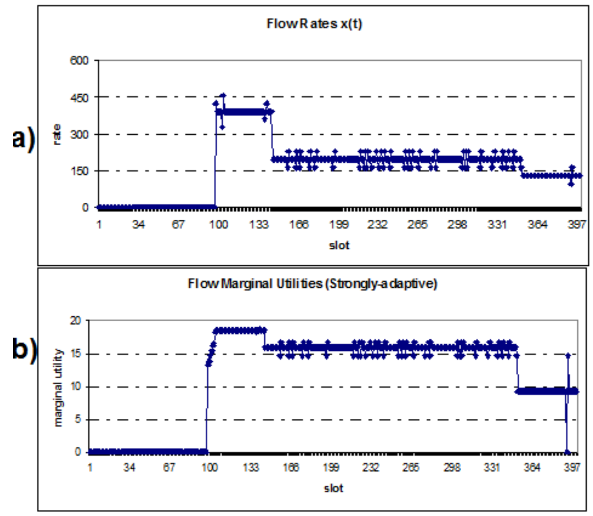

Figure 5: The rates and marginal utilities of a flow.

tions. This is due to the fact that the scheme of [18] combines high channel utilization and fairness. IEEE 802.11 has slightly higher channel utilization, i.e. it delivers on the average 17024 packets per simulation, which exceeds the 15998 packets of the scheme of [18]. On the other hand, it exhibits higher variance in the rates allocated to the competing flows, thus more frequently failing to meet the r_0 constraint of the video flows and limiting users' satisfaction; the rates allocated to the 10 competing flows for one simulation are depicted as Figure 3. This, combined with the inherent bad performance of IEEE 802.11 in terms of fairness results in lower user satisfaction and respective social welfare value. In fact, the higher the number of competing flows, the higher the variance of the IEEE 802.11. This higher variance is evident both in the individual flow allocations of Figure 3, and the aggregate channel utilization, depicted as Figure 4.

The Jain index is a metric used in the literature to depict fairness. For N flows with throughputs x_i, $i \in N$, it is defined as $\frac{(\sum_i x_i)^2}{N \cdot \sum_i x_i^2}$. The closer its value is to 1, the more fair the scheme. The average value of the Jain index computed over all simulations of this set is 0.83 for IEEE 802.11 and 0.88 for the scheme of [18]. Indeed, it is often observed that IEEE 802.11 does better than the scheme of [18] in terms of channel utilization but worse in terms of fairness. None of the two metrics suffices to deduce which of the protocols is overall better. Our framework resolves this ambiguity by means of the social welfare metric, which quantifies the impact of the protocols performance on users' flows.

Note that the ordering of the social welfare values of Table 2 is the same for every possible user type; this comprises additional evidence of the "insensitivity" of the ranking of the performance of the schemes we compare to the utility function definition. As expected, the value of the social welfare assuming that users are discretely-adaptive is less than that attained under the linearly-adaptive type, since the former is a step discretization of the latter. Also, due to the fact that for strongly-adaptive users additional units of bandwidth (on top of r_0) result in higher value than both linearly-adaptive and discretely-adaptive users, the respective value of the social welfare is the greatest.

Prior to proceeding with the presentation of additional results on the performance of the protocols under evaluation, we provide some indicative comparative plots of the rates allocated to a video session under the scheme of [18] and the corresponding marginal utility assuming that the user is of the strongly-adaptive type (similar behavior is exhibited for all the flows/user types). This flow is active between the 24^{th} and 100^{th} second. Note that since the attained rates of the flow at every $250msec$ slot is in general higher than the minimum acceptable rate of $128kb/s$, the marginal utility,

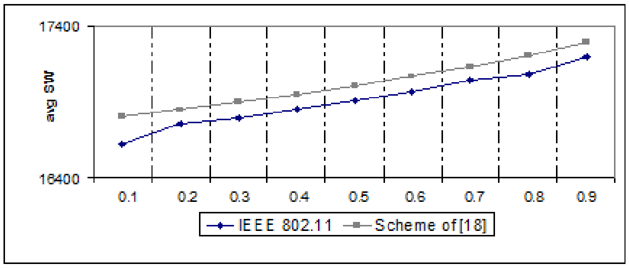

Figure 6: Average social welfare for various α_s.

depicted as Figure 5(b), follows closely the trend of the rate allocations, depicted as Figure 5(a). As more flows emit over time, the fair rate of the flows is reduced, hence the different "steps" in the respective rates plot, and this is clearly depicted to the marginal utilities as well. Instantaneous spikes in the rates allocated also affect the marginal utility. Finally, when the minimum acceptable rate of $128kb/s$ is not met - for slot 393 - the marginal utility drops to zero. Overall, the marginal utility accurately reflects the quality of service experienced by the user over time.

We have also performed a sensitivity analysis of the results for various values of α_s, which indicates that the ordering of the results remains unchanged, for all values of α_s. The values of the aggregate social welfare computed over all simulation runs and for all three types of users for the aforementioned simulations set are provided as Figure 6. We refrain from providing a different plot for every type of users so as to make the figure more readable.

For lower loads of the network, both protocols have similar performance, with the scheme of [18] being slightly better in terms of average values of social welfare. Table 3 depicts the average values of the social welfare computed over 10 simulation runs, as well as the confidence intervals; again it is $T = 100sec$, $S = 500m$. Once more, IEEE 802.11 delivers (slightly) more packets on the average, i.e. 10679 compared to the 10374 of the scheme of [18], but still exhibits more variance in the rates allocated. The performance gains of the scheme of [18] are less than those in the simulation set-up with 10 competing UDP flows, due to the limited amount of competition over the network.

Table 3: Average social welfare for 5 video flows.

User Type	*Average social welfare*	
	IEEE 802.11	Scheme of [18]
Linearly-adaptive:	9054 [7542, 12483]	9199 [7446, 14628]
Discretely-adaptive:	8592 [7160,11845]	8756 [7124,13624]
Strongly-adaptive:	14231 [11838,19639]	14410 [11540,22330]

The following set of experiments regards the transmission of 15 video flows transmitting with $384kb/s$ over a MANET of 25 nodes randomly deployed on a $500m \times 500m$ terrain. Each experiment is run for $T = 100sec$. The average values of the social welfare computed over 10 simulation runs, as well as the confidence intervals, are depicted as Table 4.

Once more, the scheme of [18] outperforms IEEE 802.11, with the difference in the social welfare values attained being larger compared to the simulations of 5 and 10 flows.

Table 4: Average social welfare for 15 video flows.

User Type	*Average social welfare*	
	IEEE 802.11	Scheme of [18]
Linearly-adaptive	19646 [15944,26326]	22549 [17961,28173]
Discretely-adaptive	18819 [15291,25275]	21751 [17294,27272]
Strongly-adaptive	30238 [24539,40224]	33894 [27012,41000]

Moreover, the average number of packets delivered by the scheme of [18] is 22573, exceeding the 21571 packets of IEEE 802.11. Therefore, employing a fairness scheme when multiple flows compete over a multi-hop ad hoc network can be beneficial not also in terms of fairness but also in terms of channel utilization.

Also, it is interesting to compare the values of the social welfare attained with the simulations run for 10 video/UDP flows. It is clear that the value of the social welfare for all schemes has increased due to the admission of more flows. This is an indication that the network is not saturated and the admission of more flows increases the overall benefit. On the contrary, had an excessive number of flows been admitted, very low values of social welfare would have been observed. It is worth emphasizing that such qualitative information is not captured by the other performance evaluations metrics used in the literature.

Finally, simulations indicate that the scheme of [18] outperforms IEEE 802.11, even if a mix of UDP and TCP applications utilize the network. The following set of experiments is indicative of the schemes' performance and regards the transmission of 10 video flows transmitting with $384kb/s$ and 5 FTP/TCP applications over a MANET of 30 nodes randomly deployed on a $500m \times 500m$ terrain. For the computation of social welfare, the per second value of being allocated $1kb/s$ for the FTP applications is set to 0.01. Each experiment is run for $T = 100sec$. The average values of the social welfare computed over 10 simulation runs, as well as the confidence intervals, are depicted as Table 5.

Table 5: Average social welfare for 10 video/UDP and 5 FTP/TCP flows.

User Type	*Average social welfare*	
	IEEE 802.11	Scheme of [18]
Linearly-adaptive	12392 [10443,15069]	13164 [10053,15993]
Discretely-adaptive	11860 [10016,14450]	12758 [9743,15407]
Strongly-adaptive	18931 [15830,22823]	19307 [14656,23396]

Since the traffic mix comprises of both UDP and TCP traffic and applications with inherently different properties and sensitivity to QoS, the social welfare value comprises a metric which can be used to comparatively assess the performance of the two schemes. On the contrary, it would not be possible to apply a metric such as the Jain index in this context, due to the heterogeneous traffic mix.

Overall, the utility-based framework captures the inherent properties of the schemes under investigation and depicts the varying QoS that flows receive over time in the respective utility values. Fairness schemes, as [18], allocate smoother rates over time; this is very desirable for multi-hop QoS-sensitive flows, thus resulting in higher values of the social welfare attained. This comes for a relatively small expense of channel utilization if the network is not congested, while this is not always the case under high competition where the ranking of the two protocols performance in terms of channel utilization is ambiguous.

5. CONCLUSIONS

In this paper, we have presented a utility-based framework, based on *QoS-aware history-dependent* utility functions. These functions quantify the satisfaction that the users of the MANETs obtain from the way their long-lived service sessions are allocated bandwidth, due to mobility and the behavior of the fairness protocols proposed for ad hoc networks. We have argued that our approach is novel and can serve as an economic-aware performance evaluation tool of the various ad hoc fairness protocols. Finally, we have demonstrated the framework's usefulness, by performing a comparative assessment of the fairness scheme of [18] with IEEE 802.11. Using our framework so as to perform an in-depth study of the relationship of fairness and QoS in ad hoc networks, as well as a detailed evaluation of more fairness schemes, comprise interesting topics of future research.

6. REFERENCES

[1] 3GPP. Specification 23.107 quality of service (qos) concept and architecture. In *http://www.3gpp.org.*

[2] G. M. Becker, M. H. DeGroot, and J. Marschak. Measuring utility by a single-response sequential method. *Behavioral Science*, 9:226–232, 1964.

[3] B. Bensaou, Y. Wang, and C. C. Ko. Fair Medium Access in 802.11 based Wireless Ad-Hoc Networks. In *ACM MobiHoc*, pages 99–106, Boston, Massachusetts, 2000.

[4] G. Berger-Sabbatel, A. Duda, M. Heusse, and F. Rousseau. Short-term fairness of 802.11 networks with several hosts. In *Proceedings of the Sixth IFIP TC6/WG6.8 Conference on Mobile and Wireless Communication Networks (MWCN 2004)*, volume 162 of *IFIP International Federation for Information Processing*, pages 263–274, Paris, France, Oct. 2004. Springer Boston.

[5] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain. Watching television over an ip network. In *Proc. of the 8th ACM SIGCOMM conference on Internet measurement*, pages 71–84, Vouliagmeni, Greece, Oct. 2008.

[6] C. Chaudet, D. Dhoutaut, and I. Guérin Lassous. Performance Issues with IEEE 802.11 in Ad Hoc Networking. *IEEE Communication Magazine*, 43(7):110–116, July 2005.

[7] C. Courcoubetis and R. Weber. *Pricing Communication Networks: Economics, Technology and Modelling*. Wiley and Sons, 2003.

[8] M. Dramitinos, G. D. Stamoulis, and C. Courcoubetis. Auction-based resource reservation in 2.5/3g networks. *Kluwer/ACM Mobile Networks and Applications Special Issue on Mobile and Pervasive Commerce*, 9(6):557–566, Dec. 2004.

[9] M. Dramitinos, G. D. Stamoulis, and C. Courcoubetis. Auction-based resource allocation in umts high speed downlink packet access. In *Proc. of the First EuroNGI Conference Conference on Traffic Engineering*, pages 434–441, Rome, Italy, Apr. 2005.

[10] H. Knoche and J. D. McCarthy. Design requirements for mobile tv. In *Proc. of the 7th international conference on Human computer interaction with mobile devices and services*, pages 69–76, Salzburg, Austria, 2005.

[11] K. Lee. Adaptive network support for mobile multimedia. In *Proc. of the 1st annual international conference on Mobile computing and networking*, pages 62–74, 1995.

[12] R. F. Liao and A. T. Campbell. A utility-based approach for quantitative adaptation in wireless packet networks. *Kluwer Wireless Networks*, 7:541–557, 2001.

[13] P. Maillé and B. Tuffin. An auction-based pricing scheme for bandwidth sharing with history-dependent utility functions. In *Proc. of the First International Workshop on Incentive Based Computing (IBC'05)*, pages 30–49, Compiegne, France, Sept. 2005.

[14] D. Miras, R. J. Jacobs, and V. Hardman. Content-aware quality adaptation for ip sessions with multiple streams. In *Proc. of the 8th International Workshop on Interactive Distributed Multimedia Systems*, pages 168–180, London, UK, 2001. Springer-Verlag.

[15] T. Razafindralambo and I. Guérin Lassous. Increasing Fairness and Efficiency using the MadMac Protocol in Ad Hoc Networks. *Ad Hoc Networks Journal, Elsevier Ed.*, 6(3):408–423, May 2008.

[16] J. Sachs, M. Prytz, and J. Gebert. Multi-access management in heterogeneous networks. *Wireless Personal Communications*, 48(1):7–32, Jan. 2009.

[17] S. Shenker. Fundamental design issues for the future internet. *IEEE Selected Areas in Communications*, 13(7):1176–1188, Sept. 1995.

[18] R. Vannier and I. Guérin Lassous. Towards a Practical and Fair Rate Allocation for Multihop Wireless Networks based on a Simple Node Model. In *11th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, pages 23–27, Vancouver, Canada, Oct. 2008.

[19] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, second edition, 1947.

[20] W. Walsh, G. Tesauro, J. Kephart, and R. Das. Utility functions in autonomic systems. In *Proc. of the International Conference on Autonomic Computing*, pages 70–77, 2004.

[21] G. Wikstrand and J. Sun. Determining utility functions for streaming low bit rate soccer video. In *Proc. of IASTED International Conference on Internet and Multimedia Systems and Applications (IMSA)*, page Aug., Kauai, Hawaii, USA, 2004.

[22] Y. Xue, B. Li, and K. Nahrstedt. Optimal Resource Allocation in Wireless Ad Hoc Networks: A Price-based Approach. *IEEE Transactions on Mobile Computing*, 5(4):347– 364, April 2006.