Performance Analysis of Wireless Multihop Data Networks

Nidhi Hegde[1] and Alexandre Proutiere[2]

[1] France Telecom R&D
38-40 rue du Général Leclerc, 92794 Issy-les-Moulineaux, France
[2] Royal Institute of Technology
Electrum 418, SE-164 40 Kista, Sweden
nidhi.hegde@rd.francetelecom.com, Alexandre.Proutiere@ens.fr

Abstract. We consider wireless multihop data networks with random multi-access mechanisms at the MAC layer. In general, our aim is to study the performance as perceived by users in a dynamic setting where data flows are generated randomly by users and cease upon completion. This task comprises two major difficulties: first, the behavior of random multi-access algorithms at slot-level in a multi-hop network is even more complex than in the case of a single hop hotspot. Second, in order to study user-level performance accounting for a dynamic population of flows, one has to first characterize the so-called rate region when the population is fixed. The rate region is defined by the set of rates at which the various active users can generate packets without inducing any instabilities in the network. Since links interact with each other through interference, characterizing the rate region is as difficult as studying the behavior of a set of interacting queues. In addition, the behavior of the congestion control algorithm must be taken into account since it impacts the set of active links and thus the interference. We propose a model, based on the so-called mean field approach, that circumvents both difficulties and allows the derivation of explicit expressions for the rate region.

1 Introduction

An emerging solution to expand network access to poorly-served or highly-loaded areas is to extend WLAN coverage. The additional WLAN access points are then connected to wired Internet gateways through wireless links to other access points or relays, creating wireless multi-hop networks with increasingly decentralized architectures. A fundamental problem is characterizing the performance of such networks, with an aim of creating decentralized algorithms that optimize performance.

We consider wireless multihop data networks with some random access mechanism at the MAC layer. Our focus is the performance as perceived by users, for example the time to transfer a data flow in a dynamic setting, where users begin finite-sized data transfers and cease transfers upon completion. A basic requirement is that the transfer times remain finite, meaning that the network should

J. García-Vidal and L. Cerdà-Alabern (Eds.): Wireless and Mobility, LNCS 4396, pp. 1–11, 2007.

be stable, that is, buffers should not overflow. We characterize the rate region, that is, the set of rates at which the active users can generate packets, such that the network is stable. In the case of wired networks or wireless networks with scheduling coordination, the performance has been extensively studied both at packet (i.e., when the number of active users is fixed) and at flow levels, see [1], [2] and references therein. Characterizing the performance of wireless networks with random multi-access algorithms presents additional challenges due to the interaction of several protocol layers. Slot-level contentions determine channel access, these transmission opportunities in turn determine packet-level rates, and this has an impact on the flow completion times.

At the slot-level, the set of active links, those with packets to transmit in the corresponding buffer, is fixed. The aim is to determine the instantaneous capacity of each link depending on the scheduling algorithm. There are two broad classes of algorithms: scheduling with coordination, and scheduling without coordination. With coordination, nodes are scheduled in a way that only non-interfering transmissions are allowed at any time. Without coordination, each node uses a random access mechanism, such as Aloha-type and IEEE 802.11 CSMA-based mechanism, to determine its transmission times. When all links interfere with each other, slot-level performance is well understood, see for example Bianchi's analysis [3] of the CSMA/CA algorithm. For the case where all links do not interfere with each other, there are few works on slot-level performance.

Previous literature on random-access multihop networks focus on slotted or unslotted Aloha and seek channel attempt probabilities that satisfy certain constraints. For Aloha-type schemes Kar et al [4] and Gupta and Stolyar [5] derive the attempt probabilities corresponding to proportional fairness, i.e., the probabilities such that the sum of the log of link throughputs is maximized. They also present distributed algorithms to achieve these probabilities. In [6] Wang and Kar identify attempt probabilities that realize max-min fairness. Furthermore, previous work considered saturated conditions where users always have a packet to transmit. Here we propose a general model for the analysis of the slot-level behavior for various types of random multi-access algorithms including those mentioned above. In addition, we do not consider saturated conditions and characterize the proportion of time a user may be active, which may depend on the congestion control mechanism implemented.

At the packet level the aim is to characterize the rate region when the number of active flows on each route is fixed. The rate region is defined by the set of feasible rates of the various flow classes - by feasible we mean that all buffers in the network must remain stable. The rate region strongly depends on the slot-level scheduling policy. When node transmissions are coordinated by a central scheduler, characterizing the rate region is fairly straightforward because the scheduler avoids interference. Scheduling coordination leads to the largest rate region possible, however is difficult to realize in a distributed way. An interesting approach is maximal scheduling, which has a distributed implementation. However, it may imply a large reduction of the rate region [7]. Without scheduling coordination, interfering nodes may attempt the channel simultaneously. In this case, characterizing the rate region is not trivial because the transmission rate of a link depends on the probabilities that the interfering links are active. At the packet level the network behaves as a set of interacting queues whose stability is largely unknown [8], [9], [10]. An additional difficulty in characterizing the rate region is that the behavior of the congestion control algorithm must be taken into consideration. Indeed, since the link capacities depend on the buffer contents of interfering links, the congestion control algorithm has a strong impact on the rate region. The greedy behavior of congestion control algorithms implies that along each route there is always one saturated buffer. We will show that this tends to reduce the rate region. The main contributions of this paper is to provide approximate expressions for the rate region, based on the mean field approach, and evaluate the impact of congestion control algorithms.

In the next section we present our model and discuss the impact of congestion control in wired networks. We describe the slot-level dynamics and the rate region at the packet level for centralized algorithms in Section 3 and for distributed algorithms in Section 4. In Section 5 we characterize the rate regions for some examples and we conclude in Section 6.

2 Model

2.1 Network Topology and Interference

We consider a network consisting of a set of nodes $\mathcal{N}$ and a set of links $\mathcal{L}$. We assume each node has a single outgoing link. This is a valid assumption for linear networks and networks based on a tree structure where traffic is aggregated at each point and there is a single path to a gateway node. A link may be present only between nodes that are within each other's transmission range. For a transmission on a given link l, another link k is an interfering link if link k's simultaneous transmission would 'collide' with link l's transmission and thus cancel it. Link k then interferes with link l if the source node of link k and the destination node of link l are within transmission range of each other. Furthermore a node cannot receive and transmit at the same time. We denote by $\mathcal{L}_l$ the set of links that interfere with a transmission on link l, more specifically, links that interfere with reception at the destination node of link l. This model of interference is often referred to as the exclusion model in the literature; it simplifies the analysis and the design of efficient distributed scheduling algorithms.

Remark 1. *The set of interfering links will in general depend not only on physical distances between nodes, but also to some extent on the random access mechanism in use. For example, the RTS/CTS virtual reservation scheme in IEEE 802.11 access mechanisms aims to avoid collisions due to so-called 'hidden nodes'. In doing so, it requires a larger set of nodes to be silent for a given transmission and thus increases the set of interfering links for a given link.*

2.2 Scheduling Algorithms

We assume that time is slotted, and for the sake of simplicity, that the duration of packet transmission is equal to the slot length. At each time slot, a scheduler decides which links are going to be activated. This decision can be done in a centralized manner, or in a distributed manner as in WLANs, mesh, and ad-hoc networks. Important examples of distributed scheduling schemes are random multi-access algorithms, such as Aloha or CSMA. We are particularly interested here in distributed algorithms since we believe that a distributed solution ensures the scalability and rapid growth of a network. However, we will present performance results with centralized scheduling for the sake of comparison.

2.3 The Rate Region

In the present paper we restrict the analysis to packet-level performance. We assume that packets are categorized into K classes according to their route in the network. The routes are further assumed to be fixed. We denote by r_k the route of packets of class k, i.e., the set of links traversed by class-k flows. In general we are interested in deriving the rate region of the network, i.e., the set of vectors representing the rates at which the packets of the various classes can be generated without inducing instability in the network. In the following we denote by ϕ_k the rate in bit/s generated by packets of class k.

2.4 The Impact of Congestion Control

In addition to scheduling decisions, the transfers of data packets are operated under the control of TCP that defines the packet rates of sources, i.e., it defines the rates ϕ_k. This protocol adapts the source rates only through the experienced packet losses in the network. In wired networks where it can be assumed that packet losses occur only because of buffer overflow, TCP exhibits what we call a *greedy behavior*. This means that there is always at least one saturated link along each route r_k (because otherwise TCP would increase the corresponding source rate). In the wireless networks we consider, we also assume that packet losses are mainly due to buffer overflow. The greedy behavior of TCP must then be taken into account in deriving the rate region.

We now present a simple network example to illustrate the impact of the greedy behavior of TCP on the rate region. Consider the wired network depicted in Figure 1. When we do not account for the greedy behavior of TCP the rate region is the set of $\phi = (\phi_1, \phi_2)$ such that $\phi_1 + \phi_2 \leq C_1$ and $\phi_2 \leq C_2$. Note that we should take strict inequalities to ensure buffer stability. However as explained below, it is more convenient to define the rate region as the closure of the set of achievable vectors ϕ. When accounting for the greedy behavior of TCP, there must be at least one saturated buffer along each route, which means in our example that both buffers are saturated. In this case, the rate region reduces to the set of vectors ϕ such that $\phi_1 + \phi_2 = C_1$ and $\phi_2 = C_2$.

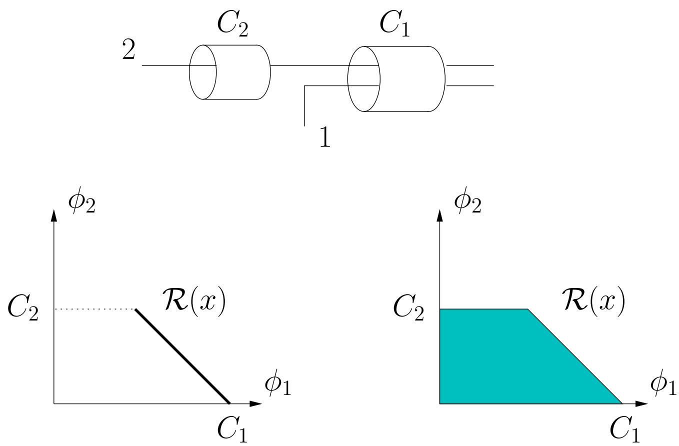

Fig. 1. A wired tree network and the corresponding rate regions accounting (left) or not (right) for the greedy behavior of TCP

In wired networks in general, without the greedy behavior of TCP, the rate region is the smallest coordinate convex[1] set containing the rate region obtained with the greedy behavior of TCP. As a consequence, this behavior does not really impact the user-level performance as explained in [2]. On the contrary, in the case of wireless networks with distributed scheduling, we will show that this may have a strong impact on the performance as perceived by users.

To summarize we give the definition of the rate region: it is the set of achievable rate vectors $\phi = (\phi_1, \ldots, \phi_K)$, where achievable means that ϕ is such that under the considered scheduling algorithm:

(i) All buffers are stable or at the stability limit;
(ii) ϕ is greedy in the sense that at least one buffer on each route is saturated.

3 Rate Region: Centralized Scheduling

Assume here that there is a central scheduler choosing at each slot the set of links that should transmit. The chosen set should always be optimal in the sense that it is not possible to add a new link that has packets in its buffer to this set without adding interference to one of the links already in the set. Denote by $\mathcal{P}$ the set of optimal link sets and by P its cardinality. A link l may belong to different optimal link sets, we denote by $\mathcal{P}_l$ this set of optimal sets. Further denote by $\mathcal{T}$ the set of vectors $\tau = (\tau_1, \ldots, \tau_P)$ such that $\tau_p \geq 0$, and $\sum_{p=1}^{P} \tau_p = 1$. Here τ_p may be interpreted as the fraction of time the link set p is scheduled. Note finally that in this setting, a link may be scheduled even if the corresponding

[1] A set $\mathcal{Y} \subset \mathbb{R}_+^K$ is coordinate convex iff $\forall y = (y_1, \ldots, y_K) \in \mathcal{Y}$ then $(z_1, \ldots, z_K) \in \mathcal{Y}$ for all $z_i \leq y_i$.

buffer is empty. Then the scheduling $\tau \in \mathcal{T}$ is compatible with the packet-level traffic $\phi = (\phi_1, \dots, \phi_K)$ only if for any link l, we have:

$$\sum_{k:l \in r_k} \phi_k \leq \sum_{p \in \mathcal{P}_l} \tau_p.$$

Now the rate regions considering the greedy behavior of TCP or not are given by the following:

$$\mathcal{R}^{TCP} = \{\phi : \exists \tau \in \mathcal{T} : \forall l, \sum_{k:l \in r_k} \phi_k = \sum_{p \in \mathcal{P}_l} \tau_p\},$$

$$\mathcal{R} = \{\phi : \exists \tau \in \mathcal{T} : \forall l, \sum_{k:l \in r_k} \phi_k \leq \sum_{p \in \mathcal{P}_l} \tau_p\}.$$

The above result was first proved by Tassiulas and Ephremides [11] in a much more general context. Note that the rate regions are in general convex, and that $\mathcal{R}$ is obtained as the smallest convex and coordinate convex set containing $\mathcal{R}^{TCP}$. In case of centralized scheduling, the situation is very close to that of wired networks.

4 Rate Region: Distributed Scheduling

In the absence of scheduling coordination, it is impossible in general to exactly characterize the packet-level stability condition, and then to derive the rate region. As mentioned above, this is due to the fact that the capacity of a particular link depends on buffer contents of interfering links. We now present a heuristic based on mean field asymptotics that allow us to circumvent this problem.

4.1 Slot-Level Dynamics

Consider a random multi-access algorithm such as Aloha or CSMA. In order to determine the traffic that link l can handle, we must infer the stationary probability p_l that link l attempts to use the channel. To do so, we apply the mean field approach, first implicitly introduced by Bianchi [3] and recently theoretically justified [12]. This approach assumes that in order to compute the probability p_l, we can consider that the behavior of link l depends on that of the interfering links only through a constant stationary collision probability c_l. This assumption then provides a formula relating p_l and c_l:

$$p_l = F_l(c_l). \tag{1}$$

For Aloha-type multi-access algorithms, the functions F_l are constants. For IEEE 802.11 algorithms, Bianchi identified this function for all links l as follows:

$$F_l(c) = \frac{2(1-2c)}{(CW+1)(1-2c) + CW \times c(1-(2c)^m)},$$

where CW and m are parameters of the 802.11 backoff mechanism.

Remark 2. *Note that in the latter case, the assumption that the time to transmit a packet is one slot, is crucial. Indeed, Bianchi's analysis consists in modelling the behavior of one link via a Markov chain representing the corresponding backoff at the instants of a point process. These instants correspond to the end of empty slots, packet transmissions, and collisions. For the analysis to be valid, all links must share the same point process. This implies that either all links interfere with each other, which is not necessarily true here, or that the duration of a packet transmission or of a collision are equal to that of an empty slot.*

For any link m, let a_m denote the proportion of time link m is active, that is, has a packet to transmit. Then the collision probability for link l is given by:

$$c_l = 1 - \prod_{m \in \mathcal{L}(x) \cap \mathcal{L}_l} (1 - a_m p_m). \tag{2}$$

Remark 3. *Our model differs from that of Bianchi in how the interference set $\mathcal{L}_l$ is defined. We have defined this set in Section 2.1 from the perspective of the destination node of link l. Thus, this set includes all links k such that the destination node of link l is in the transmission range of the source node of link k. The set then includes so-called* hidden nodes *that are sensed by the receiver but not by the transmitter. Therefore, while the form of (2) is similar to Bianchi's analysis, the collision probability here is calculated over a set of links defined appropriately for a mesh network.*

4.2 Packet-Level Dynamics

To complete the model, we now must infer the proportion of time each link is active, i.e., we compute the a_l's. Assume here that packets of class k are emitted at rate ϕ_k. Then when the rate vector ϕ is achievable, the quantities (a_l, p_l, c_l) are determined using (1), (2), and following conservation laws, for all $l \in \mathcal{L}(x)$ we have the following:

$$\sum_{k: l \in r_k} \phi_k = a_l p_l (1 - c_l). \tag{3}$$

In the above model, under the packet-level traffic $\phi = (\phi_1, \ldots, \phi_K)$ the network is stable only if the system of equations (1)-(2)-(3) has a unique solution such that for all links l, $a_l, p_l, c_l \in [0, 1]$. Note that the mean field approach alleviates the two main difficulties in evaluating the performance of mesh networks, namely the interaction between the backoff processes at the various links and the impact of buffer contents of interfering links on the capacity of a given link.

Finally, using the mean field heuristic, we are able to characterize the rate region. The conditions (i) and (ii) are (see Section 2):

$$\forall l \in \mathcal{L}(x), \exists p_l, c_l, a_l \in [0:1] \text{ satisfying (1), (2), and (3),} \tag{4}$$

$$\forall k, \exists l \in r_k : a_l = 1. \tag{5}$$

Now the rate regions with or without the greedy behavior of TCP are given by:

$$\mathcal{R}^{TCP} = \{\phi : (4),(5) \text{ are satisfied}\}. \tag{6}$$

$$\mathcal{R} = \{\phi : (4) \text{ is satisfied}\}. \tag{7}$$

5 Examples

5.1 Impact of the Greedy Behavior of the Congestion Control Algorithm

In wireless networks with no scheduling coordination, the greedy behavior of TCP can result in a significant reduction in the rate region. To illustrate this, consider a simple network composed of two interfering links depicted in Figure 2. There are 2 classes of flows, class-1 (resp. class-2) flows use link 1 (resp. 2) only.

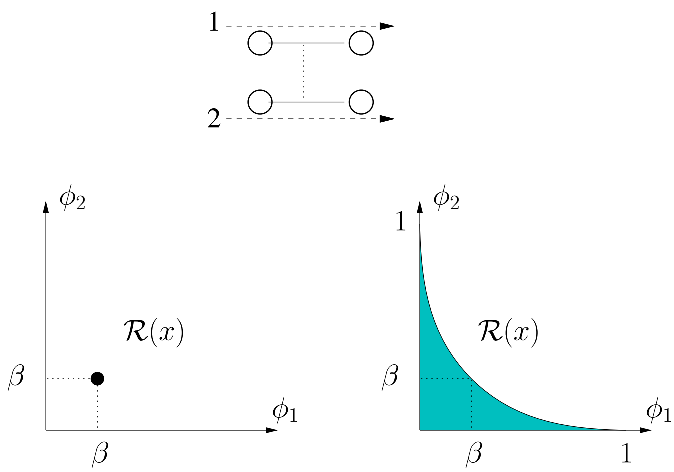

Fig. 2. The two-link network and the corresponding rate region when $x_1, x_2 > 0$ accounting (left) or not (right) for the greedy behavior of TCP - $\beta \approx 0.24$. When representing networks, links are solid lines, a dotted line between two links indicates interfering links.

Let us investigate the rate region when $x_1, x_2 > 0$. If the congestion control algorithm is greedy, then we know that both buffers will be saturated and as a consequence, the rate region reduces to a single point $\phi = (p_1(1-p_2), p_2(1-p_1))$ where p_1, p_2 satisfies $p_1 = F_1(p_2)$ and $p_2 = F_2(p_1)$. Now if the congestion control is not greedy, the buffers may not always be saturated, and the set of rates at which the network can operate is then given by:

$$\mathcal{R}(x) = \{\phi : \exists p_1, p_2, a_1, a_2 \in [0:1], \\ \phi_1 = a_1 p_1 (1 - a_2 p_2), \phi_2 = a_2 p_2 (1 - a_1 p_1), \\ p_1 = F_1(a_2 p_2), p_2 = F_2(a_1 p_1)\}.$$

The rate region for both cases is shown in Figure 2 for the MAC algorithm DCF used in IEEE 802.11. The reduction of the rate region due to the greedy behavior of TCP is much more important here than in the case of wired networks. In this example, the rate region is reduced to a single point.

5.2 Further Examples

We now present numerical examples of rate regions, accounting for greedy congestion control, for simple networks running the IEEE 802.11 DCF algorithm with $CW = 4$. If the network reduces to a single link and a single flow class, the achievable rate is then $r_0 = 0.4$. When the network has two links with a single class going

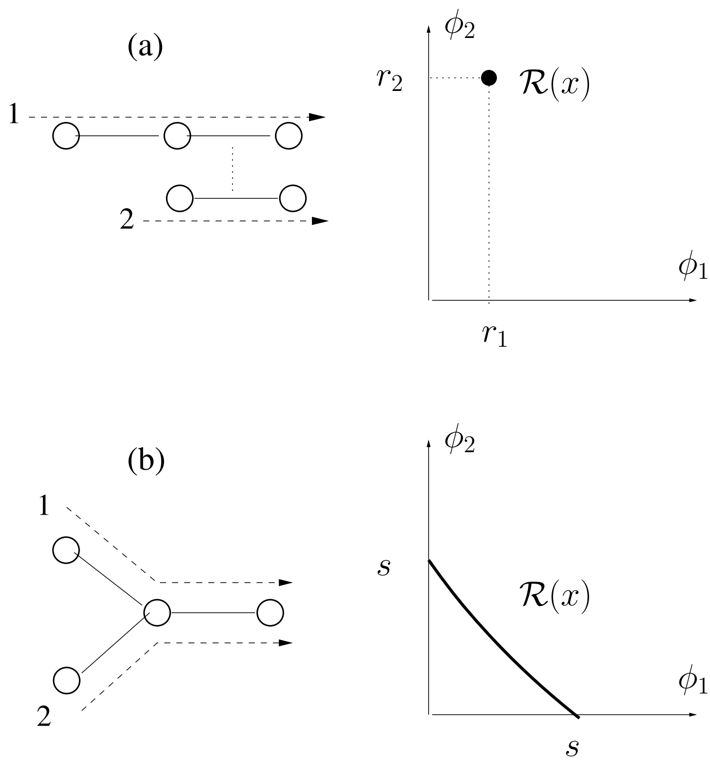

Fig. 3. Left: network topologies. Right: the corresponding rate regions when $x_1, x_2 > 0$; $r_1 \approx 0.08$, $r_2 \approx 0.29$, and $s \approx 0.2$.

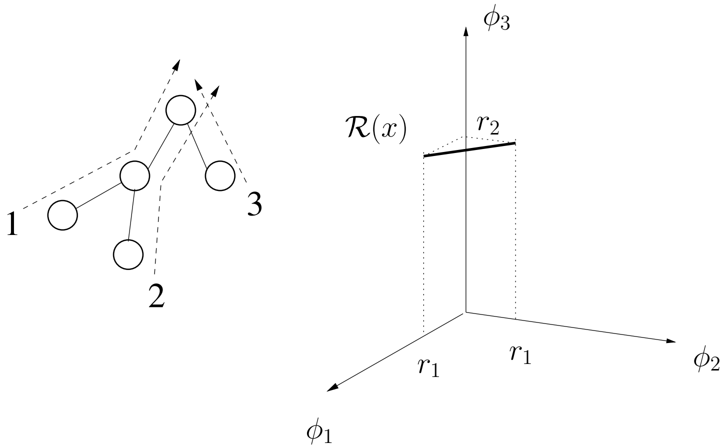

Fig. 4. A tree network and the corresponding rate region when $x_1, x_2, x_3 > 0$

through both links, the achievable rate is s around 0.2. In Figure 3, two networks with two flow classes and the corresponding rate regions are presented.

In Figure 4 a tree network with 3 flow classes and its rate region when $x_1, x_2, x_3 > 0$ is given.

Having characterized the rate regions of such networks, we are able to study the flow-level performance. Note that unlike wired networks where the rate regions are convex, the rate regions for wireless networks with distributed scheduling may be arbitrary. This has important implications on the flow-level performance, as shown in [13]. We leave the study of flow-level performance for the networks considered here for future work, but note that for the case of two flow-classes, this is characterized explicitly in [13].

6 Conclusion

We have presented an integrated slot and packet-level model based on the mean field approach that allows explicit approximate expressions for the rate region of wireless multi-hop data networks. These expressions may then used to evaluate the flow-level performance of the network. The analysis demonstrates that for networks with random access algorithms, the rate region can in general have an arbitrary shape and that may in particular be non-convex. Furthermore, this region may be significantly reduced due to the greedy behavior of the congestion control algorithm.

References

1. Georgiadis, L., Neely, M., Tassiulas, L.: Resource allocation and cross-layer control in wireless networks. Foundations and Trends in Networking **1**(1) (2006) 1–144
2. Bonald, T., Massoulie, L., Proutiere, A., Virtamo, J.: A queueing analysis of max-min fairness, proportional fairness and balanced fairness. Queueing Systems (2006)

3. Bianchi, G.: Performance analysis of the IEEE 802.11 distributed coordination function. IEEE Journal on Selected Areas in Communications **18**(3) (2000) 535–547
4. Kar, K., Sarkar, S., Tassiulas, L.: Achieving proportionally fair rates using local information in aloha networks. IEEE Transactions on Automatic Control **49**(10) (2004) 1858–1862
5. Gupta, P., Stolyar, A.: Optimal throughput allocation in general random-access networks. In: Proceedings of the Conference on Information Systems and Sciences (CISS). (2006)
6. Wang, X., Kar, K.: Distributed algorithms for max-min fair rate allocation in aloha networks. In: Proceedings of the 42nd Annual Allerton Conference. (2004)
7. Chaporkar, P., Kar, K., Sarkar, S.: Throughput guarantees through maximal scheduling in wireless networks. In: proceedings of the 43nd Annual Allerton Conference. (2005)
8. Luo, W., Ephremides, A.: Stability of n interacting queues in random-access systems. IEEE trans. on Information Theory **45**(5) (1999) 1579–1587
9. Szpankowski, W.: Stability conditions for some multiqueue distributed systems: Buffered random access systems. Advances in Applied Probability **26** (1994) 498–515
10. Bonald, T., Borst, S., Hegde, N., Proutiere, A.: Wireless data performance in multi-cell scenarios. In: Proceedings of ACM Sigmetrics. (2004)
11. Tassiulas, L., Ephremides, A.: Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multi-hop radio networks. IEEE Transactions on Automatic Control (1992) 1936–1948
12. Bordenave, C., McDonald, D., Proutière, A.: Random multi-access algorithms: A mean field analysis. In: Proceedings of Allerton conference. (2005)
13. Bonald, T., Proutiere, A.: Flow-level stability of utility-based allocations on non-convex rate region. In: Proceedings of the Conference on Information Systems and Sciences (CISS). (2006)