

Performance of wireless ad hoc networks under balanced fairness

Aleksi Penttinen and Jorma Virtamo

Networking laboratory, Helsinki University of Technology, P.O. Box 3000,
FIN-02015 HUT, Finland
{Aleksi.Penttinen, Jorma.Virtamo}@hut.fi

Abstract. Balanced fairness is a new resource sharing concept recently introduced by Bonald and Proutière. We extend the use of this notion to wireless networks where the link capacities at the flow level are not fixed but depend on the scheduling of transmission rights to interfering nodes on a faster time scale. The balance requirement together with the requirement of maximal use of the network's resources jointly determine both a unique state-dependent scheduling and bandwidth sharing between the contending flows. The flow level performance under the resulting scheme is insensitive to detailed traffic characteristics, e.g., flow size distribution. The theoretical and computational framework is formulated and illustrated by two examples for which the performance in terms of average flow throughputs in a dynamic system is explicitly worked out.

Keywords: bandwidth sharing, ad hoc networks, scheduling, balanced fairness, flow throughput.

1 Introduction

As in fixed networks, also in wireless ad hoc networks the performance perceived by the users sending elastic traffic mainly manifests itself on the flow level. A flow of elastic traffic typically comprises a transfer of a document, file or message such that the transmission can use all the bandwidth that is available but can also adapt the transmission speed to the congestion and share the bandwidth with other concurrent flows. The performance, such as the average duration of transfer of a document of a given size, clearly depends on dynamic behavior of the system and on how the bandwidth is shared between different flows. Thus it is necessary to study the system in a dynamic setting where new flows arrive at the network, are transferred across the network, and upon completion depart from the system. As far as we are aware, no analysis of this type has been done for ad hoc networks previously.

In order to facilitate the analysis, a certain degree of abstraction is necessary. In particular, we ignore the detailed packet level behavior although the actual communication in an ad hoc network consists of a sequence of packet level operations of channel access and data forwarding over a multihop route. From a conceptual point of view we distinguish two layers of operations. On the lowest level, which we call *scheduling*, one basically determines which transmitters are

allowed to send data at any given time (of course, because of the interference not all transmitters are allowed to operate simultaneously). As the set of permissions is switched on a fast time scale, the resulting network appears at flow level as a virtual network with links with capacities that depend on the schedule. So the schedule determines the capacity to send data on different links. On the upper level of operation, which is referred to as *bandwidth sharing* or *bandwidth allocation* one determines how the capacity of the virtual network is shared between the flows, i.e. viewed from a lower level, what data is sent when a sender has a permission to send.

The main difference of the analysis of flow level performance of ad hoc networks from that in fixed networks stems indeed from the fact that by scheduling one can, in certain limits, shift capacity of the network from one link to another, and that this degree of freedom can and should be used to improve the performance.

In the flow level abstraction we assume that the flow durations are long compared with the time scale of the operation of a schedule as well as with the time it takes for whatever protocol or flow control mechanism is used for the bandwidth sharing to find a steady state. Thus we assume that as soon as a flow arrives or departs, a new steady state, with a new virtual network and its resource sharing, is reached instantaneously.

Fairness of the bandwidth sharing has been recognized an important consideration for fixed networks and different fairness concepts have been introduced, cf. [14]. In ad hoc networks this issue has received attention only recently. In particular the application of the notion of max-min fairness in a static setting has been studied in [6, 12, 13]. The problem is again more complex than in fixed networks because it is entangled with the lower layer scheduling. Using closely related network models and scheduling constraints to those in the present paper, the authors of the mentioned papers have presented centralized and distributed methods that achieve max-min fair rates for given set of flows. However, the fact that an allocation is optimal in the sense of an utility function does not necessarily guarantee that the system converges to a steady state that is optimal [3]. Also, an analysis of the performance of max-min fair resource sharing in a dynamic setting would be prohibitively difficult.

A new concept of balanced fairness (BF) has recently been introduced by Bonald and Proutière [1, 2]. This is a very interesting notion on two accounts. First and foremost, it leads to a network performance which does not depend on the traffic characteristics except the traffic intensity on the different paths, in other words, the performance under BF is *insensitive*. Secondly, BF often allows an explicit analysis of the performance of simple systems in the dynamic setting.

It should be noted that BF does not represent a solution to an utility optimization problem. Neither does it guarantee Pareto efficient use of the resources. However, studies of fixed networks have shown that in many cases the performance of a network under BF is similar to that under max-min fairness. BF provides therefore a useful approximation tool for evaluating network performance.

In this paper we introduce a natural extension of the notion of balanced fairness to take into account the fact that the capacities of the virtual network can be changed by the schedule. The resulting BF resource sharing problem is a joint problem of determining both the state-dependent schedule and bandwidth allocation to maximize the use of the network resources while at the same time retaining the balance and the related insensitivity properties. We demonstrate the analysis of the performance under BF by two examples illuminating different computational approaches. Furthermore, the paper contributes to general wireless network optimization problems by establishing a condition under which so called clique constraints provide sufficient conditions for the feasibility of a given set of link capacities of the virtual network.

The paper is organized as follows. Section 2 gives a formal description of the balanced fairness concept with an extension to variable link capacities. It also illustrates the main principles involved in this paper by carefully examining a simple example. These principles are then extended to general ad hoc networks in Section 3, while Section 4 brings forward an alternative approach to allow performance analysis without the need to explicitly solve the scheduling. Section 5 concludes the paper.

2 Extension of balanced fairness

Consider a network consisting of J unidirectional links and carrying N classes of flows using pre-defined routes. Let \mathcal{F}_j denote the set of flows using link j . The network state is represented by a vector $x = (x_1, \dots, x_N)$ in which x_i is the number of class- i flows in progress. Let $\phi_i(x)$ be the bandwidth allocation to class i in state x . This bandwidth is equally shared by all flows of class i . An allocation is said to be balanced if it holds that

$$\frac{\phi_i(x - e_j)}{\phi_i(x)} = \frac{\phi_j(x - e_i)}{\phi_j(x)} \quad \forall i, j, x_i > 0, x_j > 0$$

where e_i is an N -vector with 1 in the i th component and 0 elsewhere. It can be shown [1] that an allocation is balanced if and only if the allocations can be expressed in terms of a so-called balance function $\Phi(x)$ as

$$\phi_i(x) = k \Phi(x - e_i), \tag{1}$$

where the proportionality constant is $k = 1/\Phi(x)$. Conversely, any positive function $\Phi(x)$ defines a balanced allocation by (1).

Balanced fairness as defined by Bonald and Proutière [1, 2] refers to the *most efficient* balanced allocation in a fixed network in the sense that in each state of the system at least one of the links is saturated. This leads to a unique allocation: the balance function is uniquely defined by the recursion,

$$\Phi(x) = \max_j \frac{1}{C_j} \sum_{i \in \mathcal{F}_j} \Phi(x - e_i). \tag{2}$$

The seed of the recursion can be arbitrarily set, e.g. $\Phi(0) = 1$. The *extended* balanced fairness principle just says that, in each state x , the proportionality constant $k = 1/\Phi(x)$ is chosen as large as allowed by whatever constraints the system is subject to. In the case of fixed routes and fixed link capacities this leads to (2).

Now assume that we have fixed routes but the link capacities $C_j(p)$ depend on the schedule p , i.e. the capacities are related to a virtual network defined by schedule p . Then the constraints are

$$\sum_{i \in \mathcal{F}_j} \phi_i(x) \leq C_j(p), \quad \forall j, \quad \Rightarrow \quad k \leq \min_j \frac{C_j(p)}{\sum_{i \in \mathcal{F}_j} \Phi(x - e_i)},$$

and the maximizing k is

$$k = \max_p \min_j \frac{C_j(p)}{\sum_{i \in \mathcal{F}_j} \Phi(x - e_i)},$$

that is, one finds the most constraining link for any schedule p and then makes this constraint as loose as possible by changing p . Accordingly, the balance function $\Phi(x) = 1/k$ of the balanced fairness is now uniquely determined by the recursion

$$\Phi(x) = \min_p \max_j \frac{1}{C_j(p)} \sum_{i \in \mathcal{F}_j} \Phi(x - e_i). \quad (3)$$

For each state x , this recursion defines both the balance function $\Phi(x)$ and the schedule $p = p(x)$. Recursion (2) is, of course, a special case of this, when the capacities are fixed and no scheduling alternatives are available.

Assume that the flows are generated by sessions, each session being composed of a random number of flows separated by think times. Flow sizes and think time durations can be arbitrarily distributed and need not to be independent. If balanced resource allocation is used and the *sessions* arrive as a Poisson process then, as shown in [2, 3], the steady state distribution of the network state is given by

$$\pi(x_1, \dots, x_N) = \frac{1}{G(\rho)} \Phi(x_1, \dots, x_N) \rho_1^{x_1} \dots \rho_N^{x_N}, \quad (4)$$

and depends on the traffic characteristics only through the traffic loads ρ_i of different routes. Load ρ_i is the product of the flow arrival rate and mean flow size on route i .

In equation (4), $G(\rho)$ is the normalization constant

$$G(\rho) = \sum_{x_1=0}^{\infty} \dots \sum_{x_N=0}^{\infty} \Phi(x_1, \dots, x_N) \rho_1^{x_1} \dots \rho_N^{x_N}, \quad (5)$$

which depends on the traffic load vector $\rho = (\rho_1 \dots \rho_N)$. The normalization constant $G(\rho)$ is an important quantity as the performance measures can be derived from it. Under a specific condition detailed in [4] the constant can be calculated

recursively directly without even solving $\Phi(x)$. In the examples presented in this paper the condition is satisfied and the normalization constant can indeed be obtained in a simple way.

A key performance measure for class- i flows is the throughput γ_i , defined as the ratio of the mean flow size to the mean flow duration. By Little's result this is equal to $\rho_i/E[x_i]$. The denominator can be obtained by derivation yielding

$$\gamma_i = \frac{G(\rho)}{\frac{\partial}{\partial \rho_i} G(\rho)}. \quad (6)$$

Balanced fairness has the very desirable property that the performance of the network is insensitive to traffic details. Moreover, it allows us to evaluate the performance by (3), (5) and (6). Next we illustrate the concept of extended balanced fairness by a very simple example, where the scheduling problem is trivial. A more detailed discussion of the scheduling under interference constraints is deferred to Section 3.

2.1 Example 1

The system consists of three nodes (A,B,C) and two radio links (1,2), each of nominal capacity 1. Node A cannot reach node C directly. There are two flow classes: class-1 flows from B to C utilize link 1 only while class-2 flows from A to C use both link 1 and link 2. Figure 1 illustrates the system.




Fig. 1. Example 1

Assume that the two links interfere and cannot be used at the same time. The schedule is now defined by a single parameter p defining which portion of time is scheduled for link 1. The effective link capacities are thus,

$$C_1(p) = p, \quad C_2(p) = 1 - p,$$

and the recursion (3) reads

$$\Phi(x) = \min_p \max \left\{ \frac{\Phi(x - e_1) + \Phi(x - e_2)}{p}, \frac{\Phi(x - e_2)}{1 - p} \right\}.$$

The minimum with respect to p is obtained when the two expressions are equal, implying

$$p = \frac{\Phi(x - e_1) + \Phi(x - e_2)}{\Phi(x - e_1) + 2\Phi(x - e_2)}, \quad 1 - p = \frac{\Phi(x - e_2)}{\Phi(x - e_1) + 2\Phi(x - e_2)}, \quad (7)$$

and we have the recursion

$$\Phi(x) = \Phi(x - e_1) + 2\Phi(x - e_2).$$

This recursion is solved by the balance function

$$\Phi(x) = \binom{x_1 + x_2}{x_1} 2^{x_2}. \quad (8)$$

The normalization constant can be easily calculated,

$$G(\rho) = \frac{1}{1 - \rho_1 - 2\rho_2}.$$

This is recognized to be the same as the normalization constant of a single fixed link of unit capacity shared by three flow classes with loads ρ_1 , ρ_2 and ρ_2 (or two flow classes with loads ρ_1 and $2\rho_2$). Indeed, we have a single resource, the time slot, contended by three transmissions (a class-1 transmission from B to C and class-2 transmissions from A to B and from B to C). From the normalization constant one readily obtains the throughputs using (6),

$$\gamma_1 = 1 - \rho_1 - 2\rho_2, \quad \gamma_2 = \frac{1}{2}(1 - \rho_1 - 2\rho_2). \quad (9)$$

The throughputs are greater than those resulting from the basic BF using any fixed scheduling. Note that for negligible loads the throughputs are 1 and $\frac{1}{2}$, as they should, and also that for any loads the class-2 throughput is always half of that of class 1.

Finally, note that from the balance function (8) the explicit expressions for the state-dependent allocations and schedules can be obtained by (1) and (7), respectively, yielding

$$\left\{ \begin{array}{l} \phi_1(x) = \frac{x_1}{x_1 + x_2}, \\ \phi_2(x) = \frac{\frac{1}{2}x_2}{x_1 + x_2}, \end{array} \right. \quad \left\{ \begin{array}{l} p(x) = \frac{x_1 + \frac{1}{2}x_2}{x_1 + x_2}, \\ 1 - p(x) = \frac{\frac{1}{2}x_2}{x_1 + x_2}. \end{array} \right.$$

These equations can be interpreted at the scheduling level as follows. The time slots are equally shared by all active flows, e.g. on a rotational basis. A class 2 flow needs two time slots for an end-to-end transmission; every second time slot assigned to it is used for transmission from A to B and every second time slot for transmission from B to C.

We reiterate that when the network is operated under this scheme, the performance (9) is insensitive to any detailed traffic characteristics. In this example, the scheme is also Pareto efficient in the sense that no resources are wasted; $p(x) = \phi_1(x) + \phi_2(x)$ and $1 - p(x) = \phi_2(x)$.

3 Scheduling under interference constraints

In the previous example, the two links *interfered* with each other and could not be used at the same time, resulting in effective capacities of the links that depend on scheduling.

The same principle extends to general wireless networks as follows. One can define a set of permissible simultaneous transmissions, a transmission mode, consisting of the directed links that can be used at the same time. It suffices to consider only the links used by some flow class and the maximal transmission modes, i.e. those which are not contained in another mode. Denote the set of directed links in a maximal transmission mode by τ and the set of all maximal transmission modes by \mathcal{T} .

In this general setting, we define a schedule as the vector $p = \{p_\tau, \tau \in \mathcal{T}\}$, with the meaning that in each time slot one of the τ 's is used in some order such that, on average, transmission mode τ is used the portion of time p_τ .

As in Example 1, the effective capacities of links can be defined. We assume that the flow and scheduling time scales are well separated, i.e. the duration of a typical flow is much longer than the time slot. Then, given the schedule p , we have on the flow level a virtual network with ‘fixed pipes’. Link j of this fixed network has the effective capacity

$$C_j(p) = C_j \sum_{\tau \in \mathcal{T}: j \in \tau} p_\tau, \quad (10)$$

where C_j is the nominal capacity (bandwidth) of the radio channel on link j . With these capacities one can solve the recursion (3) numerically or, as in Example 1, analytically. In general, we call a set of link capacities d_j , $j = 1, \dots, J$, *feasible* if a schedule p exists such that $d_j \leq C_j(p)$ for all j .

3.1 Modeling the interference

Interference determines which links can transmit simultaneously. Here we go briefly through different alternatives for modeling interference in ad hoc networks and define the concept of link graph.

Elementary interference models set the following constraints to the links in the network. A node may not transmit and receive simultaneously and it cannot transmit or receive more than one packet at a time. In other words, all the links connected to a given node belong to different transmission modes. Such models are often justified by the assumption that other transmissions in the vicinity of the node can be in progress without conflict using locally distinct frequencies [5–7].

A more detailed model would entail that no two links can be simultaneously active if either of the receiving ends is interfered by the other transmission. In the model presented in [8] a transmission can prevent reception everywhere within the transmission range, whereas in the widely applied *protocol model* [9] the interference depends on the locations of the transmitting node so that the closest (with a selected margin) transmission can be successfully received.

Note that the above models define interference as a pairwise property of the links. Generally, the successful reception depends on the signal to interference ratio calculated at each receiver, where the interference depends on all other links in use. This, however, complicates the search of transmission modes and in this paper we restrict ourselves to the pairwise models. No further restrictions are made on model selection.

Pairwise interference can be described using a link graph. Given a network and a set of flows with their routes, the corresponding link graph is constructed as follows: each *active* directed link in the network is mapped to a vertex and an edge connects two link graph vertices if the corresponding links interfere with each other. By definition, each τ corresponds to a maximal independent set in the link graph, i.e. a maximal subset of the vertices such that no two vertices in the subset represent an edge of the link graph. Therefore, in principle, one can find \mathcal{T} by enumerating all independent sets of the link graph.

3.2 LP-formulation for the recursion step

In some simple cases, such as the one in Example 1, one can find the optimal schedule p and bandwidth allocation analytically. In general, however, one has to resort to numerical analysis. In this case it is useful to formulate the recursion step and the corresponding schedule optimization as an LP-problem. We outline the approach, though it is not utilized in this paper.

Let $y = (p, k)^T$ be the decision vector and let $a = (0, \dots, 0, 1)$. If the $\Phi(x - e_i)$ are known for all i , we obtain $\Phi(x) = 1/k$ and the corresponding schedule $p(x)$ from y by finding the maximum in the LP-problem:

$$\max ay$$

subject to

$$My \leq (\bar{0}, 1)^T, \quad y \geq 0,$$

where $\bar{0}$ is a J -vector of zeros. M is the constraint matrix given by

$$M = \begin{pmatrix} -A & B(x) \\ \bar{1} & 0 \end{pmatrix},$$

where $\bar{1}$ is a $|\mathcal{T}|$ -vector of ones, A is the $J \times |\mathcal{T}|$ matrix $A_{j,\tau} = 1_{j \in \tau}$ and $B(x)$ is the column vector $B_j(x) = \sum_{i \in \mathcal{F}_j} \Phi(x - e_i)$. Note that only the element $B(x)$ in the matrix M needs to be updated in the recursion.

4 Maximal clique constraints

In Example 1, the bandwidth shares after schedule optimization could be interpreted to be limited solely by the fact that the three transmissions contend for the same time slot.

In a general scheduling and bandwidth sharing problem we can similarly identify one or several sets of transmissions such that transmissions in a given set contend for a common time slot. Such transmissions constitute a *clique* in the corresponding link graph. Each clique q imposes a necessary condition on the bandwidth allocation. The most stringent set of conditions is set by the maximal cliques, i.e. cliques that are not a subset of another clique. Thus, we have the necessary conditions for a feasible bandwidth allocation

$$\sum_{j \in q} \frac{1}{C_j} \sum_{i \in \mathcal{F}_j} \phi_i(x) \leq 1, \quad \forall q \in \mathcal{Q}, \quad (11)$$

where \mathcal{Q} denotes the set of all maximal cliques. The maximal cliques can be enumerated, e.g. by an algorithm from [11].

An interesting question is whether the maximal clique constraints (11) also give sufficient conditions for a feasible bandwidth allocation, i.e. whether there exists a schedule that allows attaining the link capacities required by an allocation. By counter examples one can easily see that this is not generally true. In Appendix A, however, we prove a lemma stating that the maximal clique constraints do provide necessary and sufficient conditions for the feasibility of an allocation if the link graph is a perfect graph¹.

When this is the case the bandwidth allocation is only limited by conditions (11) which are of the same type as the link capacity constraints in a fixed network. Note, however, that now there is one constraint for each clique as opposed to one constraint for each link in a fixed network, and also that a bandwidth allocation ϕ_i can appear in a condition several times as the flows in class i can traverse several links in the same clique. The balance function can then be calculated using the recursion

$$\Phi(x) = \max_{q \in \mathcal{Q}} \sum_{j \in q} \frac{1}{C_j} \sum_{i \in \mathcal{F}_j} \Phi(x - e_i). \quad (12)$$

While this greatly simplifies the task of evaluating the flow level performance as there is no need to explicitly consider the scheduling, it indeed leaves the schedule indetermined. Therefore, when needed, the scheduling must be worked out separately.

4.1 Example 2

Consider now a slightly more complicated example consisting of six nodes, three flow classes and five active unidirectional links of unit capacity, as shown in Figure 2. The interferences shown in the link graph result from the protocol model (see Section 3.1) and the link graph has four maximal independent sets and two maximal cliques. Note that the link graph is triangulated (i.e. it contains

¹ A graph G is called perfect if the chromatic number $\chi(H)$ every induced subgraph $H \subseteq G$ equals the maximum clique size $\omega(H)$ of the subgraph. A conjecture by Berge says that a graph G is perfect if and only if neither G nor its complement \bar{G} contains an odd cycle of length at least 5 as an induced subgraph [10].




Fig. 2. Flows, link graph, maximal independent sets and cliques of Example 3.

no induced cycles other than triangles) and thus a perfect graph (cf. [10]) and the maximal clique constraints apply. These can be written as

$$\begin{cases} 3\phi_1(x) + 2\phi_2(x) & \leq 1, \\ 2\phi_1(x) + \phi_2(x) + \phi_3(x) & \leq 1, \end{cases}$$

which results in the recursion

$$\begin{aligned} \Phi(x) &= \max\{3\Phi(x - e_1) + 2\Phi(x - e_2), 2\Phi(x - e_1) + \Phi(x - e_2) + \Phi(x - e_3)\} \\ &= 2\Phi(x - e_1) + \Phi(x - e_2) + \max\{\Phi(x - e_1) + \Phi(x - e_2), \Phi(x - e_3)\}. \end{aligned}$$

It can be shown by induction that when $x_3 > 0$ the latter expression realizes the maximum. Thus,

$$\Phi(x) = \begin{cases} 3\Phi(x - e_1) + 2\Phi(x - e_2), & x_3 = 0, \\ 2\Phi(x - e_1) + \Phi(x - e_2) + \Phi(x - e_3), & x_3 > 0. \end{cases}$$

As the recursion has a given form throughout a coordinate plane as well as in a positive part of the state space, the method of [4] can be applied to the calculation of the normalization constant in parts. For brevity we omit the details and only give the result

$$G(\rho) = \frac{1 - 2\rho_1 - \rho_2}{(1 - 3\rho_1 - 2\rho_2)(1 - 2\rho_1 - \rho_2 - \rho_3)}.$$

In the same way as in the first example, throughputs of different classes can be calculated by (6). The results are shown in Figure 3 for two different scenarios. In the first, ρ_2 and ρ_3 are kept fixed while ρ_1 is varied. In the second, ρ_3 is varied while the two others are fixed.

5 Conclusions

This paper studied flow-level dynamics of ad hoc networks and illustrated how the concept of balanced fairness can be extended to allow the derivation of performance measures of a wireless multihop network. The resource allocation is defined by a balance function which can be computed recursively using, e.g. one of the three methods presented in this paper: One can either explicitly write down the capacities of the links as a function of schedule or solve the integrated




Fig. 3. Throughputs of the flow classes in Example 3. Left; $\rho_2 = 0.2$, $\rho_3 = 0.2$. Right; $\rho_1 = 0.1$, $\rho_2 = 0.2$

recursion step and scheduling as an LP-problem. The third approach, applicable under special circumstances, is to use the maximal clique constraints that often yield a simple recursion for the balance function. When the link graph is perfect this approach is feasible and especially suitable for performance analysis since the actual schedule need not be worked out.

A noteworthy feature of the scheme is that the performance is insensitive to traffic details. This potentially allows one to develop simple and robust provisioning rules that depend only on traffic intensities.

A Necessity and sufficiency of maximal clique constraints

Denote the set of maximal cliques of a link graph G by \mathcal{Q} . Let d_j denote the capacity of link j in the virtual network. As discussed in Section 3, the capacities $\{d_j\}$ are feasible if there exists a conflict-free schedule such that each link j is scheduled for transmission the fraction of time d_j/C_j .

Lemma: *If the link graph G of an ad hoc network is perfect then both a necessary and a sufficient condition for the feasibility of capacities $\{d_j\}$ is*

$$\sum_{j \in q} \frac{d_j}{C_j} \leq 1, \quad \forall q \in \mathcal{Q}, \quad (13)$$

Proof: The condition is necessary since d_j/C_j is the fraction of the time that has to be scheduled for transmission on link j and for all $j \in q$ the transmissions must be non-simultaneous.

To prove sufficiency, we assume that the d_j/C_j can be written in the form

$$\frac{d_j}{C_j} = \frac{n_j}{N}, \quad \forall j,$$

where the n_j and N are integers (by choosing N large enough these relations can be satisfied to any desired accuracy). Our task now is to show that if

$$\sum_{j \in q} n_j \leq N, \quad \forall q \in \mathcal{Q}, \quad (14)$$

then there exists a conflict-free schedule in which each link j is given the fraction of time n_j/N . To this end, consider a frame of arbitrary duration and divide it into N time slots. Now, in order to realize the capacities $\{d_j\}$, one should be able to assign n_j time slots in the frame to each link j so that no conflicts occur.

If $n_j = 1$ for all j then finding a conflict free schedule is equivalent to the graph coloring problem on G in the following way. Let each vertex j represent a one-slot transmission on link j , which has to be scheduled to one of the N time slots in the frame, the position of which we refer to as the “color” of the one-slot transmission. That no two interfering transmissions can use the same time slot translates to that two vertices of G cannot have the same color if they are connected with an edge. As we assumed G to be perfect, the number of colors required equals the maximum clique size, which by (14), with $n_j = 1$ for all j , is no greater than N the number of available colors. So a conflict-free schedule for the case where each link has one one-slot transmission does exist.

We can extend this consideration to any numbers n_j satisfying (14). Make n_j copies of each vertex $j \in G$, each representing one of the n_j one-slot transmissions on link j . All replicas of vertex j have to be connected by an edge to each other and to all other nodes and their replicas in the same clique, as the corresponding transmissions interfere with each other. Such a graph can be constructed by expanding vertices one-by-one until each link j is represented by n_j vertices. A result of graph theory (Lemma 5.5.4 in [10]) states that a graph obtained from a perfect graph by expanding a vertex is again perfect. Thus, the final graph G^* , resulting from the $(n_j - 1)$ -fold expansion of each link j , is perfect. Consequently, the number of colors needed for coloring G^* equals its maximum clique size, $\max_{q \in \mathcal{Q}} \sum_{j \in q} n_j$, which again by (14) is no greater than the number of available colors N . \square

Corollary: A flow allocation $\{\phi_i(x)\}$ satisfying

$$\sum_{j \in q} \frac{1}{C_j} \sum_{i \in \mathcal{F}_j} \phi_i(x) \leq 1 \quad \forall q \in \mathcal{Q}, \quad (15)$$

is feasible if the corresponding link graph G is perfect.

Proof: Setting $d_j = \sum_{i \in \mathcal{F}_j} \phi_i(x)$ the result follows directly from the lemma.

References

1. T. Bonald and A. Proutière, *Insensitivity in processor-sharing networks*, Performance Evaluation, vol. 49, 2002, pp. 193-209.
2. T. Bonald and A. Proutière, *On performance bounds for balanced fairness*, to appear in Performance Evaluation, 2003.
3. T. Bonald and A. Proutière, *Insensitive bandwidth sharing in data networks*, Queuing Systems, vol. 44, 2003, pp. 69-100.
4. T. Bonald, A. Proutière, J. Roberts and J. Virtamo, *Computational aspects of balanced fairness*, In Charzinski et al. (eds.); Providing Quality of Service in Heterogeneous Environments, Elsevier Science, 2003, pp. 801-810.

5. B. Hajec and G. Sasaki, *Link scheduling in polynomial time*, IEEE Transactions on Information Theory, vol. 34, no. 5, Sep. 1988, pp. 910-917.
6. L. Tassiulas and S. Sarkar, *Maxmin fair scheduling in wireless networks*, In Proc. IEEE INFOCOM 2002, June 2002, pp. 764-772.
7. S. Sarkar and L. Tassiulas, *End-to-end bandwidth guarantees through fair local spectrum share in wireless ad hoc networks*, to appear in Proc. Control and Decision Conference (CDC) 2003.
8. E. Arikan, *Some complexity results about Packet Radio Networks*, IEEE Transactions on Information Theory, vol. IT-30, no. 4, July 1984, pp. 681-685.
9. P. Gupta and P. R. Kumar, *The capacity of wireless networks*, IEEE Transactions on Information Theory, vol. 46, no. 2, Mar. 2000, pp. 388-404.
10. R. Diestel, *Graph Theory* Electronic Edition 2000, <http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/>
11. C. Bron, J. Kerbosch *Algorithm 457: Finding all cliques of an undirected graph*, Communications of the ACM, vol. 16, no. 9, Sep. 1973, pp. 575-578.
12. X.L. Huang and B. Bensaou, *On max-min fairness and scheduling in wireless ad hoc networks: analytical framework and implementation* In Proc. ACM MobiHOC 2001, Oct. 2001, pp. 221-231.
13. H. Luo, S. Lu and V. Bharghavan, *A new model for packet scheduling in multihop wireless networks* In Proc. ACM MOBICOM 2000, Aug. 2000, pp. 76-86.
14. J. Mo and J. Walrand, *Fair end-to-end window-based congestion control*, IEEE/ACM Transactions on Networking, vol. 8, no. 5, Oct. 2000, pp. 556-567.