Most Balanced Overload Response in Sensor Networks

Leonidas Georgiadis
Aristotle University of Thessaloniki
Electrical Engineering Dept.
Thessaloniki 54124
E-mail: leonid@auth.gr

Leandros Tassiulas
Computer Engineering and Telecommunications Dept.
University of Thessaly
Volos, Greece
E-mail: leandros@uth.gr

***Abstract*— We consider the operation of a network in overload situations, that is, when the incoming traffic is outside the feasibility region determined by the network topology. In such a situation nodes will be overloaded and it is important to maintain a balanced network overload while ensuring that maximum amount of traffic reaches the sink nodes. We formulate the problem as lexicographic optimization of node overloads, study the properties of the solution and provide a distributed flow reallocation mechanism whose node overloads converge to the optimal solution.**

I. Introduction

Unpredictability in traffic load variations, link capacity fluctuations, topology modifications, node failures or various types of intentional misbehavior may lead a network to overload conditions. A smooth and balanced system response in those stressful situations is essential for effective crisis management in the network. This is more of an issue in wireless ad-hoc and sensor networks where due to the nature of the system and the likely scenarios of operation, anomalous behavior of that type is more likely to occur.

We consider a network consisting of an arbitrary spatial arrangement of nodes. Information may be generated at any node in the network and needs to be forwarded to a collection of hub (sink) nodes. The spatial distribution of traffic generation intensity is specified by the vector of traffic generation rates at each node. In overload condition this vector may lie outside the feasibility region of the system, that is there may be no feasible flow to transfer the information to the sinks, given the capacity of the system. This may occur for instance in a sensor network where traffic generation is event driven and activity scenarios may light up different portions of the system creating spatially localized, temporary overloads. In that case inevitably traffic backlogs will occur in the nodes. The distribution of the backlog build-up is an indication of the behavior of the system.

In this work we study the operation of the system in overload. A fluid model is considered where the information flow induced by the routing policy is represented by superflows. A superflow is a generalized notion of flow, where the aggregate incoming flow in a node may exceed the outgoing, i.e. flow conservation at the nodes need not necessarily hold. Superflows permit us to model more accurately overload situations. A superflow is a vector with a nonnegative element for each network link representing the information forwarding rate at that link. The difference of incoming minus the outgoing flow from a node is the backlog buildup rate at the node. We call this difference "node overload". The vector of node overloads under a certain routing policy is the quantitative performance objective that represents the overload response of the network to the routing policy.

We show that in the space of node overload vectors there is one that is lexicographically minimal and we characterize it. The overload corresponding to this vector also maximizes the information rate that reaches the sinks. Furthermore we show that this vector is the unique solution for a wide class of optimization problems where the optimization objective function is the sum of any nondecreasing convex function of node overloads. We call that vector "most balanced" overload vector and any superflow that induces the most balanced overload vector, "most balanced" superflow. Finally we present an adaptive superflow reallocation policy converging to a most balanced superflow.

The paper is organized as follows. In Section II we present some related work. In Section III we present the system model and the optimization problem under consideration. Properties and characterization of the optimal solution are provided in Section IV. In Section V we propose an adaptive superflow reallocation policy converging to a most balanced superflow.

II. Preliminaries and some related work

There are two viewpoints in studying information flow in networks like the one we described above, the microscopic and the macroscopic. At the microscopic level we keep track of the dynamic evolution of the system at the packet level modeling the instantaneous information backlog dynamics through appropriate stochastic queueing networks; the associated routing and flow control algorithms are viewed operating at the packet level. At the macroscopic level, under the assumption that the stochastic dynamic flows at the links and nodes of the network have long term averages, we focus on average flows; hence we have a fluid model of the system and we study different routing policies through the properties of their induced fluid flows.

At the microscopic modeling regime let $a_i(t)$ be the total amount of information generated at node $i \in K$ (K is the set

of nodes where traffic is generated) in the time interval $[0, t]$ and $a_{ij}(t)$ the total amount of information transferred to node j from node i through link (i, j) in the same time interval. If $q_i(t)$ is the information backlog at node i at time t then,

$$q_i(t) = q_i(0) + a_i(t) + \sum_{j \in N_{in}(i)} a_{ji}(t) - \sum_{j \in N_{out}(i)} a_{ij}(t), \; i \in K, \tag{1}$$

where $N_{in}(i)$ and $N_{out}(i)$ are respectively the set of incoming and outgoing neighbors of node i.

Stability of the network means bounded backlogs over time, i.e.,

$$\lim_{t \to \infty} \sup E[q_i(t)] < \infty, \; i \in K.$$

Assuming that the long term averages of the stochastic flows $a_i(t), \; a_{ij}(t)$ exist,

$$\lim_{t \to \infty} a_i(t)/t = \lambda_i, \quad \text{a.s.},$$
$$\lim_{t \to \infty} a_{ij}(t)/t = f_{ij}, \quad \text{a.s.},$$

the stability of the network implies from (1) that,

$$\lambda_i = \sum_{j \in N_{out}(i)} f_{ij} - \sum_{j \in N_{in}(i)} f_{ji}, \; i \in K. \tag{2}$$

while the link capacity constraint implies that

$$0 \leq f_{ij} \leq c_{ij}, \; i \in K, \; j \in N_{out}(i). \tag{3}$$

Equations (2) and (3) are called flow conservation and link capacity conditions respectively and are necessary conditions for stability. While the arrival rate vector $\lambda = (\lambda_1, .., \lambda_K)$ is the average spatial statistical profile of the exogenous traffic and is not affected by the network control policy, the vector of flows $\mathbf{f} = (f_{ij} : i \in K, \; j \in N_{out}(i))$ is the result of the routing policy and we may say that it characterizes the routing policy as far as its long term behavior is concerned. A necessary condition for the feasibility of an arrival rate vector λ is that there is a flow vector $\mathbf{f}$ satisfying (2) and (3). Let F_λ be the collection of flow vectors satisfying (2) and (3) for arrival rate vector λ. Then each possible routing policy that guarantees stable operation of the network corresponds to a flow vector in F_λ.

The behavior of the system in the stability regime has been studied extensively in the past. A dynamic routing and flow control policy, the Adaptive Back Pressure policy has been proposed and analyzed, showing that it achieves maximum throughput in the network. ABP is a distributed control policy where node i controls the transmissions of its outgoing links based only on its own backlog as well as its outgoing neighbors, without any knowledge of the network topology or its statistics. In the context of the network under consideration ABP operates as follows at each node i.

- At time t node i compares its backlog $q_i(t)$ with the backlog of its one hop downstream neighbor $j \in N_{out}(i)$.
- If $q_j(t) \geq q_i(t)$ then link (i, j) idles and no packet is transmitted (flow control is performed).
- If $q_j(t) < q_i(t)$ then link (i, j) transmits full speed a packet from i to j.

The above policy was proposed initially in [1], [2], in the context of a multihop radio network and in combination with a max weight radio access control policy. It was shown that the combined routing-scheduling scheme achieves maximum throughput and stabilizes the network if it is possible to do so. More specifically it was shown that under the statistical assumption of i.i.d. arrivals and given that the arrival rate vector is such that F_λ is nonempty, the ABP policy achieves stability of the network. In other words its fluid flow profile behaves as one of the feasible flows of F_λ. A similar policy has been considered later in [3] in the context of a multiclass service network and its dynamic behavior was analyzed extending the results in [2] for deterministic (σ, ρ) traffic profiles, i.e. profiles that comply to the output streams of a (σ, ρ) regulator (see [4] for more details on (σ, ρ) regulated traffic). In [5] the policy was studied under general Markov modulated statistics and batch processing to account for synchronization deficiencies of different servers due to unequal service times. A generalization of the policy, incorporating power control for time-varying channels was presented in [6]. A policy similar to ABP was proposed and studied in [7], [8], in the context of adversarial queueing theory. That is, its performance was analyzed under arrival traffic patterns that might be the worst possible within a certain family of arrival patterns, for instance all possible arrival patterns at the output of a (σ, ρ) regulator. It was shown that the policy achieves maximum throughput in that context as well. Finally the ABP policy has been considered in the context of general service systems in manufacturing and transportation problems in [9], [10], and its maximum throughput properties were verified.

In all the works discussed above the system was studied in its stability region, i.e. when the load did not exceed capacity. In the following we study the system in the overload region resorting directly to the fluid model. The behavior of the system in overload has been considered recently by several researchers [11], [12], [13], [14], proposing flow control at the edge in combination with backlog balancing inside the network, to achieve desirable throughput. In the current paper we study the overload build-up and we characterize the behavior of the policies from that perspective.

Finally, we note that the model and the problem considered in this paper can address, in a limiting sense, the continuous assignment problem of Hajek in [15].

III. Definitions, Model and Assumptions

The topology of the network is represented by a directed graph $G = (N, L)$. The set N consists of network nodes that may generate and forward traffic, and a node d that represents collectively all gateway nodes to the infrastructure network. The set of links L includes a link between any two nodes that may communicate directly. It also includes a link to node d for any network node that may communicate directly with a gateway node. Link $l \in L$ has capacity c_l. Given a set of nodes $S \subset N$, let $L_{in}(S)$ be the set of links that start at some node out of S, in $S^c = N - S$ and end at some node in S. With $N_{in}(S)$ we denote the set of nodes in S^c that are starting

points of the links in $L_{in}(S)$. For simplicity of exposition define $L_{out}(S) = L_{in}(S^c)$ and $N_{out}(S)$ the set of nodes in S^c that are ending points of the links in $L_{out}(S)$. Without loss of generality we assume that $L_{out}(d) = \varnothing$. The set $L_{in}(S)$ will also be referred to as the set of "incoming links" of the cut (S, S^c). Similarly, the set $L_{out}(S)$ will be referred to as the set of "outgoing links" of the cut (S, S^c). For simplicity, if S consists of a single node i, $S = \{i\}$, we write simply $L_{in}(i)$ and similarly for the other notations.

Denote by K the set of nodes $N - \{d\}$. Information is generated at node $i \in K$ at rate $\lambda_i \geq 0$ and is destined to sink node d. A "superflow" $\mathbf{f} = \{f_l\}_{l \in L}$, is any nonnegative vector with one element for each link that satisfies the following constraints.

$$\lambda_i + \sum_{j \in N_{in}(i)} f_{ji} - \sum_{j \in N_{out}(i)} f_{ij} \geq 0, \ i \in K \tag{4a}$$

$$0 \leq f_l \leq c_l, \ l \in L. \tag{4b}$$

The inequality in (4a) may be strict since we allow for the possibility of overload at a node. In case equality holds for every node in K, the superflow reduces to the standard "flow" definition. The quantity

$$q_i = \lambda_i + \sum_{j \in N_{in}(i)} f_{ji} - \sum_{j \in N_{out}(i)} f_{ij},$$

is the rate at which traffic is accumulated at node $i \in K$ for the specific superflow vector. We refer to q_i as the "overload" at node i under superflow $\mathbf{f}$. We extend the definition of overload to sink node d by defining $q_d = 0$. We denote by $\widehat{F}_\lambda$ the set of superflows satisfying (4), and by $\widehat{Q}_\lambda$ the set of overload vectors induced by superflows in $\widehat{F}_\lambda$.

The throughput $T_{\mathbf{f}}$ of a superflow is defined as the sum of flow intensities at the links terminating at the sink node, i.e.,

$$T_{\mathbf{f}} = \sum_{j \in N_{in}(d)} f_{jd}. \tag{5}$$

The traffic load Λ of the network is the sum of exogenous arrivals intensities,

$$\Lambda = \sum_{i \in K} \lambda_i.$$

It can be derived from the definition of q_i that for any subset $S \subseteq K$ it holds

$$\sum_{i \in S} q_i = \sum_{i \in S} \lambda_i + \sum_{(i,j) \in L_{in}(S)} f_{ij} - \sum_{(i,j) \in L_{out}(S)} f_{ij} \geq 0. \tag{6}$$

From the above equation for $S = K$ we get

$$T_{\mathbf{f}} = \Lambda - \sum_{i \in K} q_i. \tag{7}$$

If the superflow is a flow then $q_i = 0$ for all $i \in K$ and $\Lambda = T_{\mathbf{f}}$, i.e., the throughput equals the traffic load.

The performance of the policy in overload mode is quantified through the overload vector $\mathbf{q}$ of the corresponding superflow. Several physically important properties of a policy correspond to certain mathematical properties of the overload vector. From (7) we see that in order to maximize the network throughput it is enough to minimize the aggregate overload

$$\min_{\mathbf{q} \in \widehat{Q}_\lambda} \sum_{i \in K} q_i. \tag{8}$$

Also observe that if all the buffers at the nodes are equal, the time to buffer overflow of node i is the $(q_i)^{-1}$, and the time to first buffer overflow in the network is, $\min_{i \in K}\{q_i{}^{-1}\} = (\max_{i \in K} q_i)^{-1}$. Hence overload vectors that are solutions to the following problem,

$$\min_{\mathbf{q} \in \widehat{Q}_\lambda} \max_{i \in K} q_i, \tag{9}$$

maximize the time to first buffer overflow in the network. A stronger criterion than (9) is lexicographic optimization. This optimization, also known as min-max optimization [16], is based on the following order relation between vectors. Given a vector $\mathbf{v} = (v_1, ..., v_n)$, let $\overline{v}_i$, $i = 1, ..., n$ be the ith maximal coordinate of $\mathbf{v}$. We say that vector $\mathbf{v}$ is lexicographically smaller than vector $\mathbf{u}$, denoted by $\mathbf{v} \prec \mathbf{u}$, if either $\overline{v}_1 < \overline{u}_1$, or for some i, $1 \leq i < n$, $\overline{v}_j = \overline{u}_j$ for $1 \leq j \leq i$ and $\overline{v}_{i+1} < \overline{u}_{i+1}$. If in addition we allow for the possibility that $\overline{v}_i = \overline{u}_i$, for all $i = 1, 2, ..n$, we denote $\mathbf{v} \preccurlyeq \mathbf{u}$. Note that if for two vector $\mathbf{v}, \mathbf{u}$ we have $\mathbf{v} \prec \mathbf{u}$ then by definition $\max_i v_i \leq \max_j u_i$.

According to the previous discussion, attempting to maximize throughput and at the same time minimize the time to first buffer overflow amounts to solving simultaneously problems (8) and (9). It can be easily seen that solving each of these problems separately, does not guarantee a solution to the other one. However, as will be shown in the sequel, lexicographic optimization of node overloads does provide optimal solution to both of these problems. In fact, it turns out the an even stronger property than lexicographic optimization holds for the network under consideration. To this end, we introduce the following partial ordering. We say that vector $\mathbf{v}$ is more balanced than vector $\mathbf{u}$, denoted by $\mathbf{v} \vdash \mathbf{u}$, if the following inequalities hold

$$\sum_{l=1}^{i} \overline{v}_l \leq \sum_{l=1}^{i} \overline{u}_l, \ i = 1, ..., n. \tag{10}$$

Note that if for the overload vectors $\mathbf{q}^1$, $\mathbf{q}^2$ we have that $\mathbf{q}^1 \vdash \mathbf{q}^2$ then it follows that $\mathbf{q}^1 \preccurlyeq \mathbf{q}^2$and furthermore the throughput under $\mathbf{q}^1$ is larger than under $\mathbf{q}^2$. Hence a "most balanced" overload vector according to relation $\vdash$ is a very desirable property. A potential complication is that relation $\vdash$ is a partial ordering and not any two overload vectors are comparable with respect to that ordering, unlike the throughput or the lexicographic criterion that are total orderings. Hence while it is certain than an optimal throughput overload vector exist and the same holds for a lexicographically optimal ([17]), that is not clear for a most balanced overload vector. It is shown in the following that for the network under consideration a lexicographically optimal vector is also most balanced.

The results of this paper can be easily extended to the case where it is of interest to provide a weighted most balanced overload response in the sense of achieving lexicographic optimization of q_i/α_i for given constants $\alpha_i > 0$, $i \in K$. This may be of interest in situations where the buffer sizes of network nodes may differ. For simplicity in the discussion we avoid introducing the weights in the current presentation.

IV. Properties and Characterization of Most Balanced Superflows

Consider a superflow characterized by the following inequalities that hold for any link $(i,j) \in L$.

$$\text{If } q_i < q_j, \text{ then } f_{ij} = 0, \tag{11a}$$

$$\text{If } q_i > q_j, \text{ then } f_{ij} = c_{ij}. \tag{11b}$$

A superflow satisfying inequalities (11) is called "Superflow of Adaptive Back Pressure policy", SABP in short. The reason for this terminology is that, as will be seen in Section V, a SABP superflow can be obtained as limit of superflows induced by an adaptive flow update policy that is similar to the ABP policy described in the introduction. The SABP superflow can also be thought of as the equilibrium point of "selfish routing" in cases where the only information available to users is the backlog change rate at the node where they are located and at the outgoing neighbors of that node. Hence the node (or agents located at the node) directs its traffic only to nodes with smaller overloads in the hope that this way the traffic will encounter smaller congestion. Problems related to selfish routing have been the subject of several studies, see [18] and the references therein.

The main result of the paper is summarized in the following theorem.

Theorem 1: A superflow induces a most balanced overload vector if and only if it is SABP. The most balanced overload vector is unique - however there may be more than one superflow inducing the most balanced overload vector.

The proof of the theorem is outlined in the following. For more details the reader is referred to

http://users.auth.gr/~leonid/public/OverloadRespExt.pdf.

We start with the following lemma which shows that a superflow inducing a lexicographically optimal node overload vector is SABP.

Lemma 2: Let $\mathbf{f}^*$ be a superflow inducing a lexicographically optimal vector. Then $\mathbf{f}^*$ is SABP.

To proceed we need some further properties that are satisfied by the overload vector induced by a superflow. Because of (6) and (4b) we have for any $S \subseteq K$,

$$\sum_{i\in S} q_i = \sum_{i\in S}\lambda_i + \sum_{(i,j)\in L_{in}(S)} f_{ij} - \sum_{(i,j)\in L_{out}(S)} f_{ij} \geq \left(\sum_{i\in S}\lambda_i - \sum_{(i,j)\in L_{out}(S)} c_{ij}\right)^+. \tag{12}$$

Let us define for any $S \subseteq K,\ S \neq \varnothing$,

$$B(S) \triangleq \left(\sum_{i\in S}\lambda_i - \sum_{(i,j)\in L_{out}(S)} c_{ij}\right)^+. \tag{13}$$

For $S = \varnothing$ we use the convention $B(S) = 0$. The following lemma provides a lower bound on the maximum overload values on subsets of K. For a set X, $|X|$ denotes the number of its elements.

Lemma 3: Under any superflow $\mathbf{f}$ inducing overload vector $\mathbf{q}$, for any $S \subseteq K$,

$$\max_{i\in K} q_i \geq \max_{S\subseteq K} B(S)/\,|S|\,. \tag{14}$$

Let

$$\widehat{R}_1 = \max_{S\subseteq K,\ S\neq\varnothing} \frac{B(S)}{|S|},$$
$$\mathcal{S}_1 = \left\{S : \frac{B(S)}{|S|} = \widehat{R}_1,\ S \subseteq K,\ S \neq \varnothing\right\},$$
$$\widehat{S}_1 = \cup_{S\in\mathcal{S}_1} S.$$

The next lemma shows the basic property of SABP superflows related to min-max optimization.

Lemma 4: Let $\mathbf{f}$ be a SABP superflow inducing overload vector $\mathbf{q}$. Then (14) is achieved with equality and $\widehat{S}_1$ is the set of nodes with maximal overloads under $\mathbf{f}$, i.e.,

$$q_i = \max_{j\in K}\{q_j\} = \widehat{R}_1,\ i \in \widehat{S}_1.$$

Hence, $\widehat{S}_1$ is the set of nodes with maximal overload under a lexicographically optimal superflow.

Consider now a SABP superflow. If $\widehat{S}_1 = K$ then according to the previous discussion any SABP superflow has overload vector $\mathbf{q}^*$ such that,

$$q_i^* = \frac{B(K)}{|K|},\ i \in K.$$

Assume next that $\widehat{S}_1 \subset K$. Since $\max_{i\in\widehat{S}_1^c} q_i < q$, we conclude from the definition of SABP superflow that,

$$f_{ij} = 0, \text{ for all } (i,j) \in L_{in}\left(\widehat{S}_1\right), \tag{15}$$

$$f_{ij} = c_{ij}, \text{ for all } (i,j) \in L_{out}\left(\widehat{S}_1\right). \tag{16}$$

Consider the reduced network where the subgraph that consists of the nodes in $\widehat{S}_1$ is removed and for each link in $L_{out}\left(\widehat{S}_1\right)$ we put an exogenous arrival source to the node in $\widehat{S}_1^c$ where the link terminates with intensity equal to the link capacity. We can apply the same argument to the reduced graph in order to determine the set $\widehat{S}_2$ of nodes on which the second largest overload for any SABP superflow is achieved. In this manner we end-up getting node sets $\widehat{S}_1, \widehat{S}_2, ..., \widehat{S}_L$, where on set $\widehat{S}_l$ the lth maximal overload for any SABP superflow is achieved. Hence any SABP superflow determines uniquely the node overloads. Since by Lemma 2 any most balanced superflow is also a SABP superflow, we conclude.

Lemma 5: A SABP superflow induces a lexicographically optimal overload vector if and only if it is SABP. The lexicographically optimal overload vector is unique.

The next lemma shows that a lexicographically optimal overload vector is also most balanced.

Lemma 6: Let $\mathbf{q}^*$ be the overload vector induced by a SABP superflow. Then

$$\mathbf{q}^* \vdash \mathbf{q} \text{ for all } \mathbf{q} \in \widehat{Q}_\lambda. \tag{17}$$

Combining Lemmas 5 and 6 we obtain Theorem 1.

The next Theorem shows that (17) is equivalent to minimizing the sum of any convex nondecreasing function of node overloads.

Theorem 7: It holds

$$\sum_{l=1}^{i} \overline{q}_l^* \leq \sum_{l=1}^{i} \overline{q}_l, \text{for all } i = 1, ..., |K|,$$

if and only if

$$\sum_{i=1}^{|K|} g\left(q_i^*\right) \leq \sum_{i=1}^{|K|} g\left(q_i\right), \tag{18}$$

for any convex nondecreasing function $g(q)$, $q \geq 0$.

We note that property (18) is the defining property of "most balanced" assignment in [15].

V. Distributed asynchronous SABP computation Policies

In this section we present an asynchronous distributed method for computing SABP superflows that relies on the following local adjustment of a flow, done on a per link basis.

For link (i, j) do the following flow update:

- If $q_i > q_j$ and $f_{ij} < c_{ij}$ then increase f_{ij}until either $f_{ij} = c_{ij}$ or $q_i = q_j$
- If $q_i < q_j$ and $f_{ij} > 0$ then decrease f_{ij}until either $f_{ij} = 0$ or $q_i = q_j$

If the above iteration is performed infinitely often by each link, without any need for coordination of successive iterations among links we have convergence to an SABP. More specifically, let $t_n, n = 0, ..$ be a sequence of flow adjustment times, $t_n > t_{n-1},\ n = 1, ..$ and $l_n,\ n = 0, ..$ be a sequence of links such that the origin node of link l_n performs the flow adjustment operation described above at time t_n on link l_n. Assume that the update operation on each link is performed infinitely often, i.e. for any time T and for any link $l = (i, j)$ there is an update instant $t_k > T$ at which node i performs the update on link l. We have the following theorem

Theorem 8: Let $\mathbf{f}_n$ be the superflow vector at time t_n and $\mathbf{q}_n$ the associated overload vector. If $\mathbf{q}^*$ is the unique overload vector associated with all SABP superflows, then

$$\lim_{n\to\infty} \mathbf{q}_n = \mathbf{q}^*$$
$$\lim_{n\to\infty} \mathbf{f}_n = \mathbf{f}^*.$$

where $\mathbf{f}^*$ is some SABP superflow.

Notice that while $\mathbf{q}^*$ is unique, $\mathbf{f}^*$is not. That is, there may be more than one SABP superflows. The particular SABP superflow to which convergence is obtained in Theorem 8 depends on the sequence $l_n,\ n = 0, ..$ according to which link flows are adjusted.

VI. Conclusions

Polynomial time algorithms for determining a SABP superflow exist, but they do not lead to distributed adaptive policies which are of main concern when ad-hoc and sensor networks are considered.

In this work we concentrated on the routing and forwarding aspects of information transmission. A topic of further investigation is to consider cross-layer issues where wireless node interactions, rate adaptations and power control are also taken into account. Another topic of interest is the study in overload conditions of policies operating at the packet level, and the relation of these policies to flow level policies considered in the current paper.

References

[1] L. Tassiulas, Dynamic link activation scheduling in multihop radio networks with fixed of changing topology, Ph.D. thesis, University of Maryland, College Park (1991).

[2] L. Tassiulas, A. Ephremides, Stability properties of constrained queueing systems and scheduling policies for maximum throughtput in multihop radio networks, IEEE Transactions on Automatic Control 37 (12) (1992) 1936–1949.

[3] L. Tassiulas, Adaptive back-pressure congestion control based on local information, IEEE Transactions on Automatic Control 40 (2) (1995) 236–250.

[4] R. Cruz, A calculus for network delay, part i: Network elements in isolation, IEEE Transactions on Information Theory 37 (1991) 114–131.

[5] L. Tassiulas, P. Bhattacharya, Allocation of interdendent resources for maximum throughput, Stochastic Models 16 (1).

[6] M. J. Neely, E. Modiano, C. E. Rohrs, Dynamic power allocation and routing for time varying wireless networks, IEEE Journal on Selected Areas in Communications, Special Issue on Wireless Ad-hoc Networks 23 (1) (2005) 89–103.

[7] W. Aiello, E. Kushilevitz, R. Ostrovsky, A. Rosen, Adaptive packet routing for bursty adversarial traffic, in: Proc. ACM Symposium on Theory of Computing, 1998.

[8] B. Awerbuch, P. Berenbrink, A. Brinkmann, C. Scheideler, Simple routing strategies for adversarial systems, in: Proc. IEEE Symposium on Foundations of Computer Science, 2001.

[9] G. Dai, W. Lin, Maximum pressure policies in stochastic processing networks, Operations ResearchTo appear.

[10] G. Dai, W. L. R. Moorthy, C. P. Teo, Berth allocation planning optimization in container terminals, preprint.

[11] V. Tsibonis, L. Georgiadis, L. Tassiulas, Exploiting wireless channel state information for throughput maximization, IEEE Transactions on Information Theory 50 (11) (2004) 2566–2582.

[12] M. J. Neely, E. Modiano, C.-P. Li, Fairness and optimal stochastic control for heterogenous networks, in: IEEE Infocom 2005, Miami, Florida, USA, 2005.

[13] A. Eryilmaz, R. Srikant, Fair resource allocation in wireless networks using queue-length based scheduling and congestion control, in: IEEE Infocom, Maimi, Florida, USA, 2005.

[14] A. L. Stolyar, Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm, Queueing SystemsTo appear.

[15] B. Hajek, Performance of global load balancing by local adjustment, IEEE Transactions on Information Theory 36 (6) (1990) 1398–1414.

[16] D. Bertsekas, R. Gallager, Data Networks, Prentice Hall, Englewood Cliffs, New Jersey 07632, 1992.

[17] L. Georgiadis, P. Georgatsos, K. Floros, S. Sartzetakis, Lexicographically optimal balanced networks, IEEE/ACM Transactions on Networking 10 (6) (2002) 818–829.

[18] T. Roughgarden, E. Tardos, How bad is selfish routing?, ACM 49(2) (2002) 236–259.