Lightweight Cross–layer Control Algorithms for Fairness and Energy Efficiency in CDMA Ad–Hoc Networks

Anastasios Giannoulis
Dept. of Electrical and Computer Engineering
Rice University, Houston TX 77005, USA

Konstantinos P. Tsoukatos Leandros Tassiulas
Communications and Computer Engineering Dept.
University of Thessaly, Greece

***Abstract*— We consider a CDMA wireless ad–hoc network in the high SINR regime. We introduce a suite of cross–layer algorithms for joint flow control, routing, scheduling and power control. The algorithms guarantee forwarding of all incoming traffic, with an energy expenditure that can get arbitrarily close to the minimum possible. When traffic arrival rates lie outside the stable throughput region supported by the wireless network, the algorithms ensure fair allocation of resources. Compared to other algorithms that have been proposed in the past in a more general setting, our scheme is of considerably lower complexity. It relies on iterative methods for solving convex throughput optimization problems for CDMA networks in the high SNIR regime. The resulting cross–layer control algorithms are promising in practical implementations, for they operate in real–time, i.e., evolve in parallel with network dynamics, with limited computational complexity between successive control epochs.**

I. Introduction

In modern high–rate wireless data networks, improved performance depends on efficient use of the scarce resources (e.g. power, bandwidth, codes, antennas) in the presence of a volatile wireless channel. In contrast to traditional wireline networks, this calls for designs that jointly consider the physical together with higher layers in the networking stack. As a result, making the most out of the wireless network resources has generated sustained interest. A thread of research that goes back to at least [15] deals with identifying throughput maximizing policies in general time–varying networks, subject to server dependencies. Recent related work is reported, among others, in [3], [10], [11], [14], [17]. In addition to throughput maximization, these papers tackle the issue of fair allocation of network resources. The latter gains significance as wireless networks are expected to support a wide range of applications (from voice to file transfers, to video streaming) with diverse QoS requirements.

In this paper, we follow up on previous work [5], where two distributed power control algorithms were proposed. These algorithms ensure that the network achieves maximum possible throughput, given its topology and power budget. However, system operation for arrival rates beyond the throughput region was not addressed. Whenever demand for rates cannot be satisfied, not all participant devices require the same portion of the system resources. Sharing of resources must be carried out to achieve a certain notion of fairness, expressed in terms of utility functions. Starting with [9], this topic has been extensively studied in TCP–based wireline networks. A significant advance in that direction is reported in [11]. Another important issue is energy management. This is crucial in wireless ad–hoc and sensor networks, where battery drain may shorten network lifetime. An approach that minimizes energy consumption in addition to fairness and throughput objectives is reported in [10].

In this paper, focusing on CDMA wireless networks in the high SINR regime, we obtain algorithms for achieving the objectives above that are considerably simpler computationally. The approach can be viewed in three stages. 1) Introduction of distributed power control algorithms that are energy efficient; coupled with 2) Flow control techniques, which ensure fair sharing of the wireless network transmission rates, for any arrival rate scenario, be it inside or outside the stable throughput region supported by the network, and 3) Real–time joint dynamic operation in parallel with system evolution; the schemes adaptively track changes in traffic without requiring knowledge of arrival rates. In the resulting suite of algorithms different layers are coupled through queue length information, in an instance of cross–layer interaction: Local queue lengths are used for flow control at ingress nodes, whereas differential queue length and interference related information exchanged between links drives the power control function.

This paper is organized as follows. Section II presents the wireless network model. In Section III we recall the back–pressure power control algorithms. In Section IV we extend these power control algorithms so that they meet throughput demands of flows with the minimum possible energy consumption. Section V introduces flow control schemes running on top of power control; these provide fair resource allocation among competing flows. Section VI contains simulation results which illustrate the joint operation of the algorithms and affirm the desired network performance. Section VII concludes the paper.

II. System Model

We consider a wireless multihop network consisting of N nodes. Each node n can transmit information only to a set of neighbors N_n. Let there be a total of L possible transmitting links, i.e., pairs of neighbors. For each link l_i, the sender transmits data selecting a power level p_i and $\mathbf{p} = \{p_i,\ i =$

$1,\ldots,L\}$ is the system's power vector. By P_i^{max} we denote the maximum allowable instantaneous power for link l_i. At the receiver of link l_i let the signal-to-interference plus noise ratio (SINR) be

$$\gamma_i := \frac{p_i G_{ii}}{I_i + \eta_i},$$

η_i the noise power, and I_i the interference

$$I_i := \sum_{j=1, j\neq i}^{L} p_j \rho_{ji} G_{ji}.$$

In the expression above G_{ji} denotes the path loss between transmitter of link l_j and receiver of link l_i (abbreviated as $xmt(l_j)$ and $rcv(l_i)$ in the following), and ρ_{ji} is the corresponding coding gain. In CDMA wireless networks, the transmission rate over the channel may be modelled as a concave (or quasiconcave) curve with respect to the received SINR $\gamma_i(\mathbf{p})$. Here we assume that the transmission rate C_i has a functional dependence on γ_i similar to Shannon's capacity, i.e., $C_i = \log(1+\gamma_i)$. A further assumption instrumental in the results is that the system operates in the high SINR regime; this is often valid in CDMA systems. Whenever this is the case, the transmission rates can be closely approximated by $C_i = \log(\gamma_i)$.

Let λ_{ij} be the exogenous traffic arrival rate at node i with destination node j, and $\boldsymbol{\lambda}$=$\{\lambda_{ij},\ i,j=1,\ldots,N\}$ the arrival rate vector of the system. Each node i maintains a separate queue per destination, containing W_i^j backlogged amount of data pending for transmission, with final destination node j. If $j \notin N_i$ then data contained in W_i^j can reach the destination j, after a multihop route. Let Λ be the network stability region, i.e., the set of all arrival rates for which there exists some policy stabilizing the network queues.

III. Back–Pressure Power Control

We briefly review the maximum throughput distributed power control algorithms presented in [5].

To ensure maximum throughput, scheduling and routing decisions take place according to a *back–pressure* routing and scheduling policy [15]. Central to such adaptive routing and scheduling are the maximum differential backlogs $X_i(t)$ during time slot t, over each link l_i: A traffic flow that is scheduled for transmission over link l_i during time slot t is one that attains the maximum differential backlog. We mention that in the high SINR regime the network feasible rate region is convex [12], therefore all links should be activated. Consequently, the back–pressure algorithm incurs little computational effort, as long as there is no need to examine combinations of transmissions. It is only required that each link searches for the flow with the maximum differential backlog.

The fact that transmission rates depend upon transmit powers, leads us in [5] to formulate the maximum throughput policy of [15] as an optimization problem over available system power. The objective is to find power updates that guarantee maximum throughput. It turns out that such updates can be derived from the solution to the following optimization problem:

MAXTHRU

$$\max \sum_{i=1}^{L} X_i(t) C_i(\gamma_i(\mathbf{p}))$$
$$\text{subject to} \quad 0 < p_i < P_i^{max},\ i=1,\ldots,L.$$

The problem above is amenable to distributed solutions. These entail the calculation of a cross–layer *interference price* $\pi_i(t)$ at each link l_i, which depends on the maximum differential backlog, as well the link sensitivity to interference from other links. This interference price is subsequently communicated to all other links, which use it to update their own power levels. In particular, with the help of the cross-layer price

$$\pi_j(t) := -X_j(t)\ \frac{\partial C_j(\gamma_j(\mathbf{p}(t)))}{\partial I_j(\mathbf{p}(t))}, \quad j=1,\ldots,L \qquad (1)$$

we introduced two distributed coordinated solutions to this optimization problem.

1) Back–Pressure Best Response (BPBR) Power Control, where at every time slot $t=1,\ldots$ each link $l_i, i = 1,\ldots L$ performs the updates

$$\pi_i(t) = \frac{X_i(t)}{I_i(\mathbf{p}(t)) + \eta_i}$$

and

$$p_i(t) = \min\left(\frac{X_i(t-1)}{\sum_{j=1, j\neq i}^{L} \pi_j(t-1)\rho_{ij}G_{ij}},\ P_i^{max}\right).$$

2) Back–Pressure Gradient Projection (BPGP) Power Control, based on the same price recursion

$$\pi_i(t) = \frac{X_i(t)}{I_i(\mathbf{p}(t)) + \eta_i}$$

and power recursion

$$p_i(t) = \left[p_i(t-1) + \kappa\,\frac{X_i(t-1)}{p_i(t-1)} - \kappa \sum_{j=1, j\neq i}^{L} \pi_j(t-1)\rho_{ij}G_{ij}\right]_0^{P_i^{max}}$$

Details on these schemes can be found in [5].

IV. Energy Control

The abovementioned distributed power control schemes, derived from the solution to *MAXTHRU*, ensure maximization of network throughput. However, there is no guarantee about the level of transmit powers. In fact, it was observed in [5] that relatively high arrival rates and long queues lead to power usage several times reaching the upper constraint P_i^{max}, irrespective of how high the maximum power constraints are. The reason for this is that the power update value was a function of the ratios of the previous power value, the queue lengths and interference prices. As the algorithm runs in parallel with system evolution, with random incoming traffic, these ratios often exceed the maximum power constraint P_i^{max}, in an effort to mitigate interference and reach the global optimum (as viewed at each time slot t). This effect was even more noticeable with the best–response scheme.

In a mobile ad–hoc network, power minimization is of primary interest. Operating at maximum power might be far from desirable, especially if we consider networks with limited energy resources. We are interested in finding ways to achieve throughput maximization, while simultaneously regulating energy consumption. Here we present power control schemes that maintain throughput optimality while also keeping transmit powers to the minimum possible level. Previous work provides such a framework. Following [10], [14] a power vector $\mathbf{p}^\star$ can be determined so that average power becomes arbitrarily close to the minimum average power P_{av}^* required for stability. Proximity to P_{av}^* is determined by a power cost parameter V, at the expense of larger backlogs, and hence larger delays. The power updates are obtained from the solution to the following optimization problem:

ENERGY

$$\max \sum_{i=1}^{L} \Big(X_i(t) C_i(\gamma_i(\mathbf{p})) - V p_i \Big)$$
$$\text{subject to} \quad 0 < p_i < P_i^{max}, \; i = 1, \ldots, L.$$

We mention that for $V = 0$ the problem above reduces to *MAXTHRU* [5], where no energy regulation is taken into account. As before, we adopt the Shannon capacity model for the wireless link, with the approximation $C_i = \log \gamma_i$ in the high SINR regime. If we consider frozen backlogs (used as capacity weights) and put $V = 0$, we get a problem similar to [2], [7]. The developments in [2] relied on the fact that in the high SINR regime the aggregate utility is a concave function of an exponential transform of the transmit powers, hence a unique global optimum exists. This means that iterative schemes can be sought that converge to the global optimum of this static optimization problem, namely maximization of a weighted sum of the channel capacities.

It is straightforward to show that for $V > 0$ each instance of *ENERGY* at each time slot t (with fixed queues) is a convex optimization problem too.

Proposition 1: At each time slot t the optimization problem *ENERGY* has a unique, global optimum.

Proof: This follows by a simple modification to the proof in [2]. We perform the change of variables $\tilde{p}_i := \log p_i$. The objective function becomes

$$\begin{aligned} J(\mathbf{X}(t), \tilde{\mathbf{p}}) \;&:=\; \sum_{i=1}^{L} \Big(X_i(t) \log(\gamma_i(\tilde{\mathbf{p}})) - V \exp(\tilde{p}_i) \Big) \\ &=\; \sum_{i=1}^{L} \Bigg[X_i(t) \Big(\log(G_{ii} \exp(\tilde{p}_i)) \\ &\quad - \log \Big(\eta_i + \sum_{\substack{j=1 \\ j \neq i}}^{L} \exp\left(\tilde{p}_j + \log(G_{ji}\rho_{ji})\right) \Big) \Big) \\ &\quad - V \exp(\tilde{p}_i) \Bigg]. \end{aligned}$$

The first term above is linear in $\tilde{\mathbf{p}}$ and the minus log term is concave in $\tilde{\mathbf{p}}$, because the logarithm of a sum of exponentials is a convex function [1]. The third term is concave in $\tilde{\mathbf{p}}$, and the sum of concave functions is concave. Therefore *ENERGY* is a convex optimization problem with a unique solution. ∎

Consequently, it is possible to obtain a unique solution to the static optimization problem that accounts for regulation of energy consumption. Note that as time elapses the optimization problem constantly changes, for queue lengths at each time slot also change, due to data transmissions and stochastic arrivals. However, we do not require convergence for each instance of the optimization problem. This might take several iterations per time slot and render the approach more complex and perhaps impractical. Instead, we propose algorithms that perform a single iteration per slot, towards the global optimum as perceived at each time slot t. Of course, convergence to this operating point never takes place, since the optimization problem changes at the next slot $t+1$, and so does the global optimum. Still, we prove in [6] that despite the stochastic fluctuations, running the algorithms in parallel with system evolution is sufficient to guarantee maximum throughput.

A. Best–Response Algorithm

To determine the unique solution to the optimization problem *ENERGY* with $C_i(\gamma_i(\mathbf{p})) = \log(\gamma_i(\mathbf{p}))$ we write the Karush-Kuhn-Tucker (KKT) conditions for the optimal power vector $\mathbf{p}^*(t)$, namely

$$X_i(t) \, \frac{\partial C_i(\gamma_i(\mathbf{p}(t)))}{\partial p_i(t)} + \sum_{\substack{j=1 \\ j \neq i}}^{L} X_j(t) \frac{\partial C_j(\gamma_j(\mathbf{p}(t)))}{\partial p_i(t)} \Bigg|_{\mathbf{p}(t) = \mathbf{p}^*(t)} - V = \nu_i - \mu_i$$

and

$$\nu_i(p_i^\star(t) - P_i^{max}) = 0, \quad \mu_i p_i^\star(t) = 0, \quad \nu_i, \mu_i \geq 0,$$

where ν_i, μ_i are the Lagrange multipliers associated with the power constraints. Consider the cross–layer pricing scheme $\boldsymbol{\pi} = \{\pi_i, \; i = 1 \ldots L\}$ of (1) introduced in [5], where each link charges other links for causing interference to its transmission. These prices convey both interference sensitivity (as in [7]), as well as backlog information, for they are scaled by the link's maximum differential backlog $X_j(t)$. When the inequality constrains are inactive the Lagrange multipliers are zero and the KKT conditions take the form

$$\frac{X_i(t)}{p_i(t)} - \sum_{\substack{j=1 \\ j \neq i}}^{L} \pi_j(t) \rho_{ij} G_{ij} - V = 0 \tag{2}$$

for each $i = 1, \ldots, L$. Assuming prices charged by other links are known, each link l_i may solve (2) to find its own power from

$$p_i(t) = \frac{X_i(t)}{V + \sum_{\substack{j=1 \\ j \neq i}}^{L} \pi_j(t) \rho_{ij} G_{ij}}, \quad i = 1, \ldots, L.$$

The equations above motivate the following energy efficient algorithm:

Back–Pressure Best–Response Energy Efficient Power Control

For every time slot $t = 1, 2, \ldots$, each link l_i, $i = 1, \ldots, L$:

1. Computes the differential backlog

$$X_i^m(t) := \begin{cases} W_{xmt(l_i)}^m(t) - W_{rcv(l_i)}^m(t), & rcv(l_i) \neq m \\ W_{xmt(l_i)}^m(t), & rcv(l_i) = m \end{cases}$$

   for each flow with destination $m = 1, \ldots, N$. Let the maximum differential backlog at link l_i be

$$X_i(t) := \max_{m=1,\ldots,N} X_i^m(t).$$

2. Schedules for transmission a flow $m^\star(i)$ achieving the maximum differential backlog, i.e., one for which $X_i^{m^\star(i)}(t) = X_i(t)$.
3. Computes a cross–layer interference price

$$\pi_i(t) = \frac{X_i(t)}{I_i(\mathbf{p}(t)) + \eta_i},$$

   The price $\pi_i(t)$ is subsequently communicated to all links.
4. Transmits with power given by

$$p_i(t+1) = \min\left(\frac{X_i(t)}{V + \sum_{\substack{j=1 \\ j \neq i}}^{L} \pi_j(t) \rho_{ij} G_{ij}}, \quad P_i^{max} \right)$$

   where the constant $V > 0$ is very large.

Note that power updates above can be derived by interpreting the convex optimization problem as a power control game. Each link adjusts its power in a *best–response* fashion, in an effort to maximize its own net surplus [5], [7].

B. Gradient–Projection Algorithm

The second algorithm obtains by solving the *ENERGY* convex optimization problem with the *gradient–projection* method. This requires that the partial derivative of the objective function be calculated

$$\frac{\partial}{\partial p_i} J(\mathbf{X}(t), \mathbf{p}) = \frac{X_i(t)}{p_i(t)} - \sum_{\substack{j=1 \\ j \neq i}}^{L} \frac{X_j(t) \rho_{ij} G_{ij}}{I_j(\mathbf{p}(t)) + \eta_j} - V.$$

Convergence towards the maximum is provided by the updates

$$p_i(t+1) = p_i(t) + \kappa \frac{\partial}{\partial p_i} J(\mathbf{X}(t), \mathbf{p}), \quad i = 1, \ldots, L,$$

for sufficiently small step size $\kappa > 0$. Consequently, the second algorithm summarizes as follows:

Back–Pressure Gradient–Projection Energy Efficient Power Control

For every time slot $t = 1, 2, \ldots$, each link l_i, $i = 1, \ldots, L$:

1–3. Performs exactly the same steps 1–3 of the best–response algorithm.

4. Transmits with power given by

$$p_i(t) = \Bigg[p_i(t-1) \quad + \kappa \left(\frac{X_i(t-1)}{p_i(t-1)} - \sum_{\substack{j=1 \\ j \neq i}}^{L} \pi_j(t-1) \rho_{ij} G_{ij} - V \right) \Bigg]_0^{P_i^{max}},$$

   for sufficiently small stepsize $k > 0$, with the notation $[x]_a^b := \max(\min(x, b), a)$.

In both algorithms, the higher the power cost parameter V, the higher the queue lengths and the delays, the lesser the consumed energy. In the distributed algorithms we presented in [5], V was set to zero; this caused the minimum possible delays, however transmit powers were high. Still, for any fixed V convergence at each time slot t is turned towards a power vector $\mathbf{p}^\star(t)$ that results in throughput optimality.

We emphasize that both energy efficient algorithms employ back–pressure routing, and run in parallel with system operation. That is, the algorithms *do not* compute the solution $\mathbf{p}^\star(t)$ to *ENERGY* at every time slot t. Observe that node backlogs and flows transmitted over links vary over time. Therefore, a new optimization problem *ENERGY* arises at every time slot t, before convergence to the solution $\mathbf{p}^\star(t-1)$ for the previous time slot is achieved. Another observation is the following: The gradient projection algorithm, due to the small constant κ which is necessary for convergence, is more conservative in energy consumption in comparison with the best response scheme; this was also observed in [5] where $V = 0$. That is, for the same value of V, the gradient descent algorithm gives power levels closer to the minimum power required for stability.

The present setup can be compared with [10], where distributed implementation was presented for the special case of cell-partitioned networks, for which solution of *ENERGY* amounted to individual selection of the appropriate power at each link. In general, the transmission rate is a concave function of the SINR γ which, due to interference, depends on the entire power vector $\mathbf{p}$. As a result, the optimization problem is a complicated one, which requires global coordination. Our work provides distributed coordination schemes for solving *ENERGY*.

Finally, within the set of stabilizing network control policies, performance with respect to backlogs and delays may vary. In a single–hop network backlog ratios can be steered to desired targets by employing the so-called *exponential rule* of [13] in conjunction with power control, as done in [4]. In the current setup, each link l_i can be weighted differently, by means of its own power cost parameter V_i. Alternatively, each link l_i might use a flow related parameter V_{ij} (encapsulated in the header of the packet), so that the system could support applications with different QoS.

V. Flow Control

We now turn to providing performance guarantees for arrival rates beyond the achievable throughput region. In such an operating regime queues grow to infinity. It is clear that as long as a link's transmit power p_i is dependent upon the size of its differential backlog X_i then a relatively large arrival rate (beyond the system stable throughput region) might win exclusive use of the wireless channel. Other flows whose arrival rates are strictly inside the throughput region are lead to starvation. This necessitates the use of flow control.

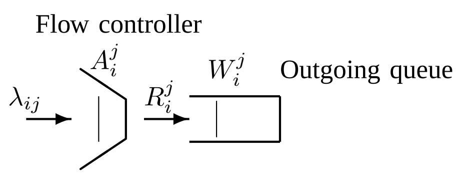

Fig. 1. Flow control based on local information

A. Problem Formulation

Stochastic arrivals entering the network at node i are initially stored at an intermediate reservoir whose total content is A_i, before becoming available for transmission. At each time slot t, flow control decisions determine the amount of traffic $R_i^j(t)$ with destination node j, to be removed from reservoir containing A_i^j (if enough available) and be placed at the outgoing queue with size W_i^j. Each node i perceives utility in his allocated flow rate for destination j, quantified by an increasing concave utility function $U_i^j(R_i^j)$. The objective is to maximize aggregate network utility:

FAIRNESS

$$\max \sum_{i=1}^{N}\sum_{j=1}^{N} U_i^j(R_i^j)$$

$$\text{subject to} \quad \mathbf{R} \in \Lambda.$$

The total outgoing traffic from node i to every possible direction at time t is denoted by $C_i^{out}(t)$. This depends on the power control updates of the previous section, involving the outgoing queues W_i^j for each node i. Introducing the Lagrange multipliers (congestion prices) $\{\xi_i(t),\ i = 1, \ldots N\}$ we consider the corresponding Lagrangian

$$\sum_{i=1}^{N}\sum_{j=1}^{N} U_i^j(R_i^j(t)) - \sum_{i=1}^{N} \xi_i(t)\left(\sum_{j=1}^{N} R_i^j(t) - C_i^{out}(t)\right).$$

The stationary points of the Lagrangian satisfy

$$\frac{\partial U_i^j(R_i^j(t))}{\partial R_i^j(t)} - \xi_i(t) = 0, \quad \text{i.e.,} \quad \xi_i(t) = (U_i^j)'(R_i^j(t))$$

and the optimal flow rates are related to congestion prices according to

$$R_i^j(t) = (U_i^j)'^{-1}(\xi_i(t)).$$

As in [9], a distributed solution to the dual problem can be sought using subgradient price updates of the form

$$\xi_i(t+1) = \left[\xi_i(t) + \frac{1}{K}\left(\sum_{j=1}^{N} R_{ij}(t) - C_i^{out}(t)\right)\right]^+ \qquad (3)$$

where $K > 0$ is a large constant, and setting

$$R_i^j(t+1) = (U_i^j)'^{-1}(\xi_i(t)).$$

A key observation in [11] is that (3) is a Lindley recursion, where ξ_i is the scaled queue length, i.e.,

$$\xi_i(t) = \frac{W_i(t)}{K}, \quad \text{where} \quad W_i(t) := \sum_{j=1}^{N} W_i^j(t)$$

is the *total local* queue length at node i, $\ i = 1, \ldots, N$. Thus the Lagrange multiplier computation is readily available from the network via local queue lengths. The parameter K regulates the behavior of the flow control mechanism. Higher K results in closer approximation of the desired fair operating point, however queues grow longer, hence delay is greater. In summary, we have:

Flow Control Based on Local Information

For every time slot $t = 1, 2, \ldots$ and for each destination $j = 1, \ldots, N$, each node $i = 1, \ldots, N$:

1. Removes from the exogenous traffic reservoir amount of traffic equal to

$$R_i^j(t) = \min\left((U_i^j)'^{-1}(W_i(t-1)/K), A_i^j(t)\right),$$

where $A_i^j(t)$ denotes the current amount of data in the exogenous traffic reservoir of node i destined for node j.

2. Subsequently places these data $R_i^j(t)$ in the respective outgoing queues with content $W_i^j(t)$.

The algorithm above is to be contrasted with traditional TCP–based flow control in wireline networks, where paths are fixed and the congestion price is the sum of node queue lengths along the path leading from source to destination [9]. Here, instead, flow control is based exclusively on local information: The aggregate (over all possible destinations) local queue length $W_i(t)$ at each node $i = 1, \ldots, N$, plays the role of congestion price, which is used to regulate traffic input to the network. Depending on the chosen flow utility functions U_i^j, flow control expresses different objectives. Routinely encountered objectives are proportional fairness, maximizing a weighted sum of throughputs or providing minimum rate guarantees; all these can be attained using the present approach.

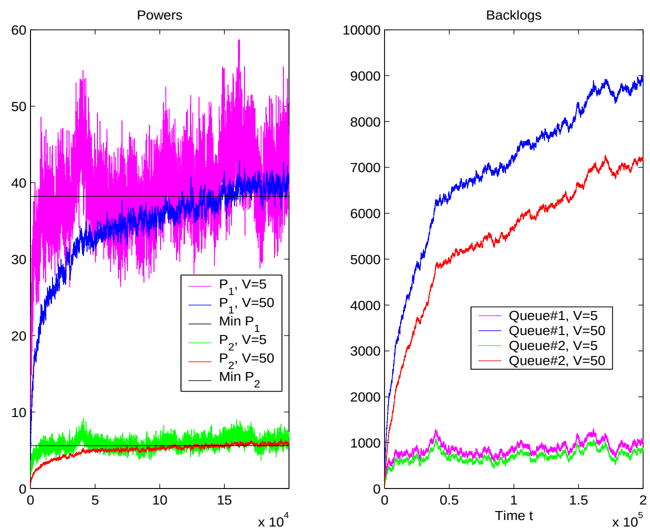

Fig. 2. Tradeoff between energy consumption and queueing delay

B. Proportional Fairness

Proportional fairness can be attained when the flow utility functions are of the form $U_i^j(R_i^j) = a_{ij}\log(R_i^j)$, where the weights $\{a_{ij},\ i,j = 1,\ldots,N\}$ determine fair proportion. In that case the rate updates solving *FAIRNESS* are given by

$$R_i^j(t+1) = \frac{a_{ij}}{\xi_i(t)} = \frac{a_{ij}K}{W_i(t)}, \quad i,j = 1,\ldots,N.$$

C. Weighted Sum of Throughputs

The objective is to reach the operating point where the weighted sum of throughputs is maximized (also see [17]), where $\{a_{ij},\ i,j = 1,\ldots,N\}$ are now the throughput weights. This means that service should ensure that $\sum_{i=1}^N \sum_{j=1}^N a_{ij} r_i^j$ is maximized, where r_i^j is the throughput of exogenous traffic generated at node i, with destination node j. Hence $U_i^j(R_i^j) = a_{ij}R_i^j$, and the solution to *FAIRNESS* is

$$R_i^j(t+1) = \begin{cases} A_i^j(t), & W_i(t) \le a_{ij}K \\ 0, & W_i(t) > a_{ij}K \end{cases}, \quad i,j = 1,\ldots N.$$

D. Minimum Rates

Flow control tries to distribute some minimum requested rates $r_{i,min}^j$ for each flow with origin node i and destination j. Assume that such rates are indeed feasible, i.e., lie inside the stable throughput region of the network. These can be approximately provisioned by considering the concave function

$$U_i^j(R_i^j) = \delta\log(R_i^j - r_{i,min}^j)$$

where $\delta > 0$ is sufficiently small. With this utility function the source perceives little gain in receiving rates that are substantially larger than the minimum ones. The resulting rate updates are given by

$$R_i^j(t+1) = r_{i,min}^j + \frac{\delta K}{W_i(t)}, \quad i,j = 1,\ldots N.$$

E. Joint Real–Time Operation

The joint network design, involving flow control, routing scheduling and power control, is obtained by concatenating the flow control algorithm with any one of the energy efficient algorithms listed in Section IV. Real–time energy and flow control operations take place simultaneously. Note that average delay is affected by both control parameters V and K. Also note that different layers are coupled through queue length information, in an instance of cross–layer interaction: Local queue lengths W_i are used as congestion prices for flow control at network nodes, whereas differential queue length (along with interference related) information exchanged between links drives the physical layer power control function. Whenever arrival rates lie outside the network stable throughput region, power control still operates as in the stable scenario. In that case, the (unstable) flow arrival rates are essentially substituted by the arrival rates of the desired fair operating point. Also, as mentioned before, flow control decisions do not depend upon control information sent to the source across the network. The rate of each flow can be determined locally at each source, since only local information is used to determine allowable rates: 1) local total queue length and 2) flow specific utility function. Hence, flow control does not incur additional communication overhead.

Similar flow control techniques have been proposed and used before [8], [11]. However, the present work proposes real–time interaction of components across different layers: Flow control, depending on local queue lengths W_i, interacts in real–time (through the queue lengths) with the adaptive power control algorithms, that run in parallel with system evolution, to determine the instantaneous service rates $C_i^{out}(t)$ and hence affect queue lengths and flow control decision at the next time slot. Consequently, correct real–time operations of the resulting suite of algorithms under dynamic traffic patterns is not a priori guaranteed. In Section VI we provide extensive simulation results that illustrate convergence of the joint algorithms; these results verify that the desired real–time network control is indeed achieved.

Lastly, we comment that the present model did not account for time variations in the wireless channel. This important issue is a subject for further work.

VI. Simulation Results

We start by demonstrating that, with the proposed power updates, energy expenditure, regulated by the parameter V, can get arbitrarily close to the minimum possible. A simple scenario with two single–hop flows, i.e., two pairs of transmitters and receivers is considered. In Figure 2, we show the evolution of the power levels and the queue lengths over time, for two different values of the regulating parameter $V = 5$ and $V = 50$. We run the best–response algorithm, but similar results hold for the gradient–projection algorithm as well. The arrival rates are close to the boundary of (but inside) the system stable throughput region, with the one flow having almost double rate as the other. The horizontal black lines are the transmit powers that give target SINRs whose corresponding

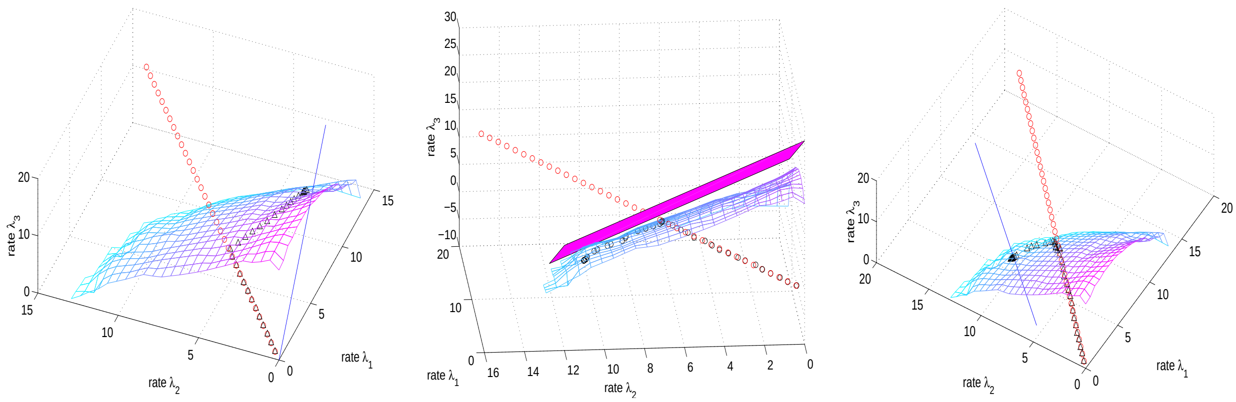

Fig. 4. Multihop network: Proportional fairness (left), maximum weighted sum of throughputs (center) and minimum rates (right)

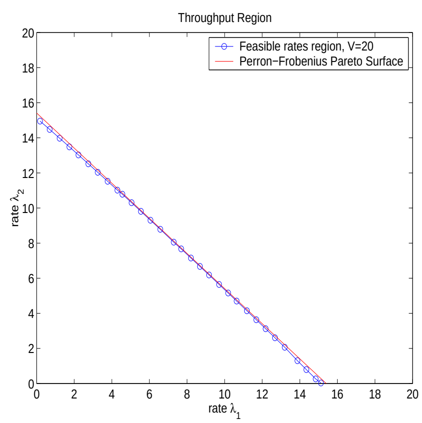

Fig. 3. Throughput region for energy regulating algorithms

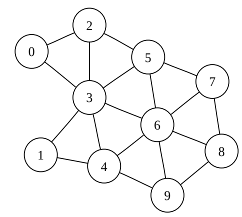

Fig. 5. Multihop network topology

constant channel transmission rates are equal to the (unknown to the algorithm) arrival rates. For a small value $V = 5$ the powers of the two flows fluctuate heavily around the minimum powers (black horizontal lines). For a larger value $V = 50$ fluctuations around the minimum powers are much lower. However, the corresponding (see color code) queue lengths and delays for flows are significantly larger when $V = 50$. The tradeoff between energy consumption and queueing delay is clear.

In general, energy regulation may not imply throughput optimality. We show that the proposed schemes address both objectives simultaneously. Figure 3 depicts the feasible rate region for the single–hop, two receivers/transmitters pair network. In [5], the rate region achieved in the simulation experiments indeed approximated the Pareto surface, determined in terms of the Perron–Frobenius eigenvalue [12]; this provides the theoretical throughput region of the system in the high SNIR regime. Here, in Figure 3, we verify that the theoretical maximum throughput region is again achieved while also regulating energy consumption with $V = 20$. The results shown concern the best–response scheme, however similar results hold for the gradient–projection algorithm too.

Next, we illustrate the operation of the proposed joint suite of algorithms, including flow control, routing, scheduling and power control. We consider the multihop network of Figure 5. Three flows originating from nodes 0, 1 and 5 have destination nodes 8, 7 and 4 respectively. The throughput region in Figure 4 is convex, as expected. A sequence of increasing arrival rates (red color) is considered, and as long as these remain inside the stable throughput region, they coincide with the service rates (black color) provided by the network. Since we intend to maximize network throughput, the desired operating point is given by the intersection of the fair regions and the Pareto surface. For the case of proportionally fair utilities (left), fair points lie on the line

$$\frac{\lambda_1}{5} = \lambda_2 = \frac{\lambda_3}{3}.$$

For maximizing a weighted sum of throughputs (center), the operating point is given by the intersection of the Pareto surface with the plane $5\lambda_1 + 6\lambda_2 + 4\lambda_3 = c$, with the maximum possible value of $c = 106$. Finally, for providing minimum rates (right), fair points lie on the line

$$r_1 - 3 = r_2 - 8 = r_3 - 5.$$

We confirm that when arrival rates (red color) are beyond the throughput region of the network, service rates (black color) move along the boundary in the right direction, so as to *approach the fair point*. Thus, the proposed energy regulating algorithms successfully interact with the flow control scheme to jointly determine resource allocation in real–time.

A final, cleaner example on the joint algorithm operation is provided by a single–hop network with two flows. In all cases of Figure 6 we generate a sequence of Poisson arrival processes, and show the arrival rates and corresponding service

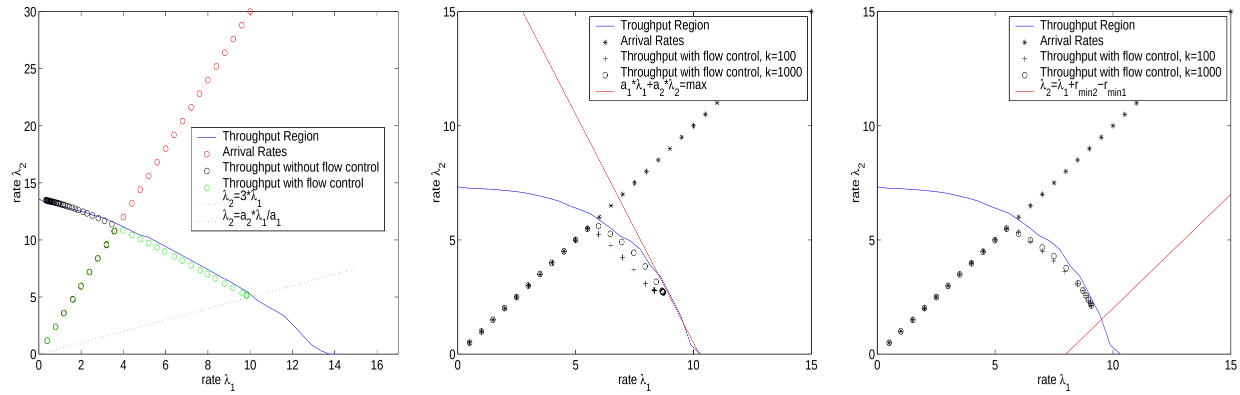

Fig. 6. Two traffic flows: Proportional fairness (left), maximum weighted sum of throughputs (center) and minimum rates (right)

rates, as well as the fair lines. When considering proportional fairness, we also show the performance of flow control in comparison to the uncontrolled case. It is seen that in the uncontrolled case one flow is lead to starvation. As far as other notions of fairness are concerned, we illustrate the difference in the precision of approximating the fair operating point, for different values of the control parameter K. Recall that better tracking takes places for higher K at the expense of larger backlogs. As expected, allocated service rates are equal to the arrival rates when the latter lie inside the network's throughput region, otherwise they are steered towards the targeted fair operating point.

VII. Conclusion

In this paper we considered a CDMA wireless ad–hoc network in the high SINR regime, and extended work initiated in [5], where distributed maximum throughput power control schemes were introduced. We presented algorithms that jointly solve the flow control, routing, scheduling and power control problems. Main ingredients of the approach are back–pressure routing, and flow control at each node, based on local queue length information. The proposed algorithms leverage on iterative distributed solutions to convex optimization problems, and can be tuned to provide fairness and minimize energy consumption by trading off for queueing delay. A prime feature of the proposed schemes is that they operate in real–time, i.e., in parallel with system evolution, without knowledge of traffic statistics. Simulation results confirming desired operations abound.

Acknowledgement

This work was supported by the European Commission through the EuroNGI Network of Excellence specific research project RAWQOS.

References

[1] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, Cambridge (UK) 2004.
[2] M. Chiang, "Balancing Transport and Physical Layers in Wireless Multihop Networks: Jointly Optimal Congestion Control and Power Control," IEEE J. Sel. Areas Comm., vol. 23, no. 1, pp. 104-116, Jan. 2005
[3] A. Eryilmaz and R. Srikant, "Fair Resource Allocation in Wireless Networks using Queue-length-based Scheduling and Congestion Control," Proceedings of IEEE INFOCOM, Miami, FL, March 2005.
[4] A. Giannoulis, K. P. Tsoukatos and L. Tassiulas, "Cross–layer Power Control in Wireless Networks", to appear in special issue of ACM SIGMETRICS Performance Evaluation Review.
[5] A. Giannoulis, K. P. Tsoukatos and L. Tassiulas, "Maximum Throughput Power Control in CDMA Wireless Networks", accepted for presentation at IEEE ICC' 06, Instanbul, Turkey 2006.
[6] A. Giannoulis, K. P. Tsoukatos and L. Tassiulas, "Lightweight Cross–layer Optimal Control of CDMA Ad–Hoc Networks", in preparation.
[7] J. Huang, R.A. Berry, and M.L. Honig, "A Game Theoretic Analysis of Distributed Power Control for Spread Spectrum Ad Hoc Networks", in IEEE International Symposium on Information Theory (ISIT'05), Adelaide, Australia, September 2005.
[8] X. Lin and N. B. Shroff, "Joint Rate Control and Scheduling in Multihop Wireless Networks", 43rd IEEE Conference on Decision and Control, in Atlantis, Paradise Island, Bahamas.
[9] S.H. Low and D.E. Lapsley, "Optimization Flow Control, I: Basic Algorithm and Convergence", IEEE/ACM Transactions on Networking, 7(6), pp. 861–874, Dec. 1999.
[10] M. J. Neely, "Energy Optimal Control for Time Varying Wireless Networks", IEEE INFOCOM Miami, FL, March 2005.
[11] M. J. Neely, E. Modiano, and C-P. Li, "Fairness and Optimal Stochastic Control for Heterogeneous Networks", IEEE INFOCOM Miami, FL, March 2005.
[12] D. O' Neill, D. Julian, and S. Boyd, "Seeking Foschini's Genie: Optimal Rates and Powers in Wireless Networks", IEEE Transactions on Vehicular Technology, 2004.
[13] S. Shakkottai and A.L. Stolyar, "Scheduling Algorithms for a Mixture of Real-Time and Non-Real-Time Data in HDR", in Proceedings of the 17th International Teletraffic Congress - ITC-17, Salvador da Bahia, Brazil, Sep. 2001, pp. 793-804.
[14] A.L. Stolyar, "Maximizing Queueing Network Utility subject to Stability: Greedy Primal-Dual Algorithm" , Queueing Systems, 2005, Vol. 50, No.4, pp. 401-457.
[15] L. Tassiulas, A. Ephremides, "Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks", IEEE Transactions on Automatic Control, Vol. 37, No. 12, pp. 1936-1949, December 1992.
[16] L. Tassiulas, "Linear complexity algorithms for maximum throughput in radio networks and input queuing switches", Proceedings of INFOCOM 98, San Francisco, California, 1998.
[17] V. Tsibonis, L. Georgiadis, L. Tassiulas, "Exploiting Wireless Channel State Information for Throughput Maximization", IEEE Transactions on Information Theory, vol. 50, no. 11, November 2004, pp 2566-2582.