Equalizing Throughputs in Random-Access Networks

P.M. van de Ven
Eindhoven University of Technology

S.C. Borst
Eindhoven University of Technology

D. Denteneer
Philips Research

A.J.E.M. Janssen
Eindhoven University of Technology

J.S.H. van Leeuwaarden
Eindhoven University of Technology

ABSTRACT

Random-access algorithms such as CSMA provide a popular mechanism for distributed medium access control in large-scale wireless networks. In recent years, tractable models have been shown to yield accurate throughput estimates for CSMA networks. We consider the saturated model on a general conflict graph, and prove that for each graph, there exists a vector of activation rates (or mean back-off times) that leads to equal throughputs for all users. We describe an algorithm for computing such activation rates, and discuss a few specific conflict graphs that allow for explicit characterization of these fair activation rates.

Keywords

CSMA, fixed point, global invertibility, loss networks, Markov processes, random access, throughput, wireless networks

1. INTRODUCTION

Carrier-Sense Multiple-Access (CSMA) is a popular mechanism for distributed medium access control in wireless networks. Under CMSA, a node attempts to activate after a certain random back-off time, but freezes its back-off timer whenever it senses activity of an interfering node, until the medium is again sensed idle. The local interactions lead to rather complex behavior on a macroscopic scale, which can be studied via models that assume that the interference constraints can be represented by a general conflict graph, and that the various nodes activate asynchronously at exponential rates whenever none of their neighbors are presently active. Such models were first pursued in the context of IEEE 802.11 systems by Wang & Kar [18], and further studied in [3, 4, 5]. These models in fact long pre-date the IEEE 802.11 standard and were already considered in the 1980's [1, 2, 10, 12]. The model has strong connections with Markov random fields and migration processes, and can under certain assumptions be interpreted as a special instance of a loss network [9, 11, 15, 19].

In the classical models the activation rates (mean back-off times) are assumed to be fixed, and the users are assumed to have saturated buffers. When the activation rates are allowed to be adapted and the buffer contents are driven by packet arrivals and departures, there are simple necessary and sufficient conditions for stability to be achievable. Several authors have proposed clever backlog-based algorithms for adapting activation rates that achieve stability whenever feasible to do so at all [6, 7, 8, 13, 14]. In contrast, we consider a scenario with saturated buffers, and aim at determining the activation rates that will lead to all nodes transmitting the same fraction of the time, thus ensuring a long-term fair sharing of the medium (equal throughputs). In this paper we address this problem for general conflict graphs, extending the results in [16] for linear networks.

2. MODEL

Consider a network of nodes sharing the wireless medium according to a CSMA-type protocol. The network is described by an undirected conflict graph (V, E) where the set of vertices $V = \{1, \ldots, N\}$ represents the various nodes of the network and the set of edges $E \subseteq V \times V$ indicates which pairs of nodes interfere. Nodes that are neighbors in the conflict graph are prevented from simultaneous activity, and the independent sets correspond to the feasible joint activity states. An inactive node is said to be blocked whenever any of its neighbors is active, and unblocked otherwise.

We consider a scenario where nodes are saturated, i.e., always have packets to transmit. The transmission times of node i are independent and generally distributed with mean 1. After each transmission, a node starts a back-off period. The back-off periods of node i are independent and generally distributed with mean $1/\lambda_i$. If a node becomes blocked, it freezes its back-off period until all neighboring nodes have become inactive.

Define $\Omega \subseteq \{0, 1\}^N$ as the set of all feasible joint activity states of the network, i.e., the incidence vectors of the independent sets of the conflict graph. Let $Y(t) \in \Omega$ represent the activity state of the network at time t, with $Y_i(t)$ indicating whether node i is active ($Y_i(t) = 1$) at time t or not ($Y_i(t) = 0$). Denote by $\pi(\omega) = \lim_{t\to\infty} P\{Y(t) = \omega\}$ the limiting probability of the joint activity state $\omega \in \Omega$. This probability distribution has a product-form solution (see e.g. [17])

$$\pi(\omega) = Z^{-1} \prod_{i=1}^{N} \lambda_i^{\omega_i}, \quad \omega \in \Omega, \tag{1}$$

where Z is the normalization constant.

Denote the total number of feasible states by $K + 1$, and write the state space as $\Omega = \{\Omega_0, \ldots, \Omega_K\}$, where we index the states such that state Ω_0 is the empty state, and Ω_i, $i = 1, \ldots, N$ represents the state in which only node i is active. Introduce the $N \times K$ incidence matrix X with $X_{ik} = 1$ when the ith element in the state Ω_k equals 1, $k = 1, \ldots, K$. All other elements of X are zero. So the columns of X jointly describe all feasible activity states, with the exception of the empty state.

We are primarily interested in the long-term behavior of nodes, characterized by their throughputs. Let $\theta = (\theta_1, \ldots, \theta_N)$ denote the throughput vector, where θ_i represents the fraction of time node i is active. We have that

$$\theta = X \cdot \Pi, \tag{2}$$

with $\Pi = (\pi(\Omega_1), \ldots, \pi(\Omega_K))$.

3. MAIN RESULTS

We aim to solve for the vector $\lambda^* = (\lambda_1^*, \ldots, \lambda_N^*) \in [0, \infty)^N$ that leads to equal throughputs for all nodes. From (2) it is seen that this inversion problem involves solving a system of nonlinear equations. We may rewrite (2) as

$$\begin{aligned}\theta(\lambda) &= \Big(\sum_{k=1}^{K} \mathbf{1}(j \in \Omega_k) \pi(\Omega_k) \Big)_{j=1,\ldots,N} \\ &= \frac{1}{Z} \Big(\sum_{k=1}^{K} X_{jk} \prod_{i=1}^{N} \lambda_i^{X_{ik}} \Big)_{j=1,\ldots,N} = \frac{1}{Z} X \mathrm{e}^{X^t \ln \lambda}, \end{aligned} \tag{3}$$

with $\lambda = (\lambda_1, \ldots, \lambda_N)$. We first consider the mapping $\lambda \mapsto \eta(\lambda) = X\mathrm{e}^{X^t \ln \lambda}$, for now ignoring the normalization constant Z.

THEOREM 1. *The mapping η is globally invertible on $(0,\infty)^N$. Thus, given a vector $\bar{c} \in (0,\infty)^N$, there is a unique $\lambda^* = \lambda(\bar{c}) \in (0,\infty)^N$ such that $\eta(\lambda) = \bar{c}$.*

Sketch of the proof. For any $\lambda \in (0,\infty)^N$, the functional matrix

$$\left(\frac{\partial \eta_j}{\partial \lambda_i} \right)_{i,j=1,\ldots,N} = \left(\sum_{k=1}^{K} X_{jk} X_{ik} \frac{1}{\lambda_i} \prod_{l=1}^{N} \lambda_l^{X_{lk}} \right)_{i,j=1,\ldots,N} \tag{4}$$

is the product of a sum of positive semi-definite matrices and a diagonal matrix with positive diagonal elements. The functional matrix has full rank since $X_{ik} = \delta_{ik}$ for $k = 1, \ldots, N$ (Kronecker delta). Furthermore, it can be shown that

$$\max_j |\ln(\eta_j(\lambda))| \to \infty \tag{5}$$

as $\max_i |\ln \lambda_i| \to \infty$. Hence, the result follows from the global inversion theorem. □

We henceforth assume that $\bar{c} = c\tau$, with $c \in \mathbb{R}_+$ and τ the all-1 vector. To solve $\eta(\lambda) = \bar{c}$ we write

$$(X\mathrm{e}^{X^t \ln \lambda})_k = \lambda_k + \lambda_k \mathcal{G}_k(\bar{\lambda}_k) \tag{6}$$

with $\bar{\lambda}_k = (\lambda_1, \lambda_2, \ldots, \lambda_{k-1}, \lambda_{k+1}, \ldots, \lambda_N) \in \mathbb{R}^{N-1}$ and

$$\mathcal{G}_j(\bar{\lambda}_k) = \sum_{k=N+1}^{K} X_{jk} \prod_{\substack{i=1 \\ i \neq j}}^{N} \lambda_i^{X_{ik}}. \tag{7}$$

Define the mapping $\mathcal{H} : \mathbb{R}_+^N \mapsto \mathbb{R}_+^N$ as

$$\mathcal{H}(\lambda) = \Big(\frac{c}{1 + \mathcal{G}_k(\bar{\lambda}_k)} \Big)_{k=1,\ldots,N}. \tag{8}$$

Obviously, $\mathcal{H} : [0,c]^N \mapsto [0,c]^N$ and since $\mathcal{H}$ is continuous, there exists a fixed point $\lambda^* \in [0,c]^N$ for which

$$\mathcal{H}(\lambda^*) = \lambda^*. \tag{9}$$

3.1 Fixed-point iteration

The fixed point can be determined by an iteration that starts with $\lambda^{(0)} = 0 \in [0,c]^N$, $\lambda^{(1)} = \mathcal{H}(\lambda^{(0)}) = \bar{c}$, and continues as

$$\lambda^{(i)} = \mathcal{H}(\lambda^{(i-1)}), \quad i = 1, 2, \ldots. \tag{10}$$

Under the condition that none of $\mathcal{G}_j \equiv 0$, it can be proved by induction that (componentwise inequalities)

$$\begin{aligned} 0 = \lambda^{(0)} < \lambda^{(2)} < \cdots < \lambda^{(2l)} < \lambda^* \\ < \lambda^{(2l-1)} < \lambda^{(2l-3)} < \cdots < \lambda^{(3)} < \lambda^{(1)} = \bar{c} \end{aligned} \tag{11}$$

for $l = 1, 2, \ldots$. The fixed-point iteration, however, may fail in the sense that $\lim_{l\to\infty} \lambda^{(2l)} < \lambda^* < \lim_{l\to\infty} \lambda^{(2l-1)}$ (strict inequality). In this case one can use a Newton method where the observation made in the proof of Theorem 1 on the functional matrix is instrumental.

3.2 Closed-form solutions

For certain specific conflict graphs we can determine the fixed point in closed form using the theory of Markov random fields. The first result in this direction is due to Kelly [10], who considered a tree with nearest-neighbor blocking. For a linear network of N nodes, in which a transmitting node blocks the first β nodes on both sides, the fixed point that renders equal throughput takes the form [16]

$$\lambda_i^* = \sigma(1+\sigma)^{\gamma(i)-\gamma(1)}, \tag{12}$$

with $\gamma(i)$ the number of nodes that can potentially block node i, and σ any positive constant. In this case the throughout is given by

$$\theta_i(\lambda^*) = \frac{\sigma}{1 + (1+\beta)\sigma}. \tag{13}$$

Another model for which a closed-form solution exists is the rectangular grid of size $2 \times L$ in which a transmitting node blocks its direct neighboring nodes (in horizontal and vertical direction). It turns out that for all $\sigma > 0$, the throughput is equalized by setting $\lambda_i^* = \sigma$ when i is a corner node, and

$$\lambda_i^* = \sigma \frac{1 + 2\sigma}{1 + \sigma} \tag{14}$$

otherwise. This choice of λ^* yields throughputs

$$\theta_i(\lambda^*) = \frac{\sigma(1+\sigma)}{1 + 2\sigma(2+\sigma)}. \tag{15}$$

4. DISCUSSION

We finally discuss some topics for further research and connections with other recent work.

4.1 Extended mapping

Recall from (3) that the mapping η only captures the non-normalized throughput. In order to control the actual throughput, we need to consider the extended mapping (with τ the all-1 vector)

$$\lambda \in (0,\infty)^N \mapsto \theta(\lambda) = \frac{X\mathrm{e}^{X^t \ln \lambda}}{1 + X\mathrm{e}^{X^t \ln \lambda}\tau} \in \mathrm{int}(\mathcal{A}). \tag{16}$$

Here $\mathcal{A}$ denotes the convex hull of all feasible states in Ω (see e.g. [13]). The mapping in (16) can, of course, not be inverted outside $\mathcal{A}$. The study of the mapping in (16) is work in progress.

An alternative approach to control the throughputs is to use the mapping η. The inverse $\lambda^* = \lambda(\bar{c})$ leads to a throughput vector $\theta = Z^{-1}\bar{c}$. Hence, Theorem 1 allows us to equalize throughputs, but we cannot control the actual value (since Z depends on λ). To achieve some specific throughput level θ^*, so that $\theta(\lambda) = \theta^*\tau \in \mathrm{int}(\mathcal{A})$, we can recursively apply the following two-step procedure:

(i) Take some $\alpha \in \mathbb{R}_+$, set $\bar{c} = \alpha\tau$ and find $\lambda(\bar{c})$ such that $\eta(\lambda) = \bar{c}$.

(ii) Calculate $\theta(\lambda(\bar{c})) = Z(\lambda(\bar{c}))^{-1}\bar{c}$, compare to $\theta^*\tau$, and adjust α accordingly.

4.2 Maximal schedules

It is well known that for large networks the calculation of the normalization constant Z presents numerical difficulties. Our inversion problem inherits these difficulties, as the construction of the incidence matrix X requires all possible states in Ω (as does the calculation of Z). An interesting approximation for the inversion problem, which might be more efficient, would be to restrict to the set of *maximal schedules*, defined as those states in which no additional node can become active.

4.3 Critical load

As stated in the introduction, there are simple backlog-based algorithms for adapting activation rates that achieve stability whenever feasible at all [6, 7, 8, 13, 14]. In these algorithms, the activation rate of a node i is some function $F_i(Q) = F_i(Q_1, \ldots, Q_N)$ of the queue lengths at the various nodes. When the load vector $\rho = (\rho_1, \ldots, \rho_N)$ is 'near' the boundary $\partial\mathcal{A}$ of the feasible rate region, there will be a 'nearly unique' throughput vector $\theta \in \mathrm{int}(\mathcal{A})$ such that $\rho < \theta$. This suggests that in critical load the queue lengths will behave in such a manner that the functions $F_i(Q_1, \ldots, Q_n)$ correspond to the activation rates λ_i that lead to the throughput vector θ, providing an interesting connection between weighted fair activation rates and heavy-traffic limits.

5. REFERENCES

[1] R.R. Boorstyn, A. Kershenbaum (1980). Throughput analysis of multihop packet radio. In: *Proc. ICC*, 1361–1366.

[2] R.R. Boorstyn, A. Kershenbaum, B. Maglaris, V. Sahin (1987). Throughput analysis in multihop CSMA packet radio networks. *IEEE Trans. Commun.* **35**, 267–274.

[3] M. Durvy, O. Dousse, P. Thiran (2007). Modeling the 802.11 protocol under different capture and sensing capabilities. In: *Proc. Infocom 2007*, Anchorage AK, May 6–12, 2356–2360.

[4] M. Durvy, O. Dousse, P. Thiran (2009). On the fairness of large CSMA networks. *IEEE J. Sel. Areas Commun.* **27 (7)**, 1093–1104.

[5] M. Durvy, P. Thiran (2006). A packing approach to compare slotted and non-slotted medium access control. In: *Proc. Infocom 2006*, Barcelona, Spain.

[6] L. Jiang, D. Shah, J. Shin, J. Walrand (2010). Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inform. Theory*, to appear.

[7] L. Jiang, J. Walrand (2010). A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Trans. Netw.*, **18(3)**, 960-972.

[8] L. Jiang, J. Walrand (2008). A distributed algorithm for optimal throughput and fairness in wireless networks with a general interference model. Preprint.

[9] F.P. Kelly (1979). *Reversibility and Stochastic Networks*. Wiley, Chichester.

[10] F.P. Kelly (1985). Stochastic models of computer communication systems. *J. Roy. Stat. Soc. B* **47 (3)**, 379–395.

[11] F.P. Kelly (1991). Loss networks. *Ann. Appl. Prob.* **1 (3)**, 319–378.

[12] A. Kershenbaum, R.R. Boorstyn, M. Chen (1987). An algorithm for evaluation of throughput in multihop packet radio networks with complex topologies. *IEEE J. Sel. Areas Commun.* **5 (6)**, 1003–1012.

[13] J. Liu, Y. Yi, A. Proutière, M. Chiang, H.V. Poor (2008). Maximizing utility via random access without message passing. Technical report MSR-TR-2008-128, Microsoft Research.

[14] S. Rajagopalan, D. Shah, J. Shin (2009). Network adiabatic theorem: an efficient randomized protocol for content resolution. In: *Proc. ACM SIGMETRICS/Performance 2009*, Seattle WA, June 15–19, 133–144.

[15] Y. Suhov, U.A. Rozikov (2004). A hard-core model on a Cayley tree: An example of a loss network. *Queueing Systems* **46 (1-2)**, 197–212.

[16] P.M. van de Ven, J.S.H. van Leeuwaarden, D. Denteneer, A.J.E.M. Janssen (2009). Spatial fairness in wireless multi-access networks. In: *Proc. ValueTools 2009*, Pisa, October 20–22.

[17] P.M. van de Ven, J.S.H. van Leeuwaarden, S.C. Borst, A. Proutière (2010). Insensitivity and stability of random-access networks. To appear in: *Proc. Performance 2010*, Namur, November 15–19.

[18] X. Wang, K. Kar (2005). Throughput modeling and fairness issues in CSMA/CA based ad-hoc networks. In: *Proc. Infocom 2005*, Miami FL, March 13–17.

[19] S. Zachary, I. Ziedins (1999). Loss networks and Markov random fields. *J. Appl. Prob.* **36 (2)**, 403–414.