

Throughput-Delay Trade-off of CSMA Policies in Wireless Networks

Mahdi Lotfinezhad and Peter Marbach

Department of Computer Science, University of Toronto

E-mail: {mahdi,marbach}@cs.toronto.edu

Abstract—We consider CSMA policies for multihop wireless networks. CSMA policies are simple policies that can be easily implemented in a distributed manner. However, the delay performance of CSMA policies can be very poor as the delay can grow exponentially in the network size. As a result, CSMA policies are not practical for delay-sensitive traffic even for mid-sized networks. In this paper, we consider a slight variant of the classical CSMA policies and show that it leads to a much improved delay performance. In particular, we show that the delay does not depend on the network size. Using this result, we also characterize the delay-throughput trade-off of the proposed CSMA policy. At the heart of our analysis is a result that shows that CSMA policies quickly converge to a maximum schedule, i.e., converge to a maximum schedule at a rate that does not depend on the network size. Using this insight, we consider a CSMA policy that periodically “unlocks” the transmission pattern of a CSMA policy, and show that this unlocking mechanism can be used to obtain a much improved delay performance without significantly reducing the throughput. While our analysis has been carried out for the special case of an interference graph with a grid (lattice) topology, we provide numerical case studies for general network topologies and show that the intuition obtained from the analysis carries over to these general cases. We also illustrate the performance of the proposed CSMA policy when combined with a flow control mechanism.

I. INTRODUCTION

In this paper, we study the performance of CSMA policies [1]–[6]. CSMA policies are attractive as they are simple and can be easily implemented in a distributed manner. Moreover, these policies are known to be throughput optimal in that they can stabilize any packet arrival rate that can be stabilized by any other policy [4]–[6]. However, it is well-known that they can have a very poor delay performance where the delay can grow exponentially in the network size. As a result, CSMA policies are not practical for delay-sensitive traffic even for mid-sized networks.

In the following, we consider the question whether the fact that the delay can grow exponentially in the network size is a fundamental property of CSMA policies, or whether it is possible to obtain a much improved delay performance by slightly changing the classical CSMA policy. The main result of our analysis is that this is indeed the case: by slightly modifying the classical CSMA policy we obtain a new policy for which the delay is constant, and does not grow, as the size of the networks grows. We then also use this result to characterize the throughput delay trade-off of the proposed CSMA policy.

Our analysis is based on an idealized CSMA policy as it has been considered in [4], [7]: assuming a continuous distribution of the random back-off time of CSMA, we can avoid packet collisions. In addition, we focus in our analysis on wireless networks in which the interference can be characterized by a grid-like interference graph. The reason for assuming a grid (lattice) topology of the interference graph is that this topology allows a formal analysis.

For the classical CSMA policy, a grid (lattice) topology for the interference graph leads to the well-known hardcore model that has been extensively studied [8], [9], and has been used to study the delay performance of CSMA policies. In particular, it has been shown that under the grid interference graph, the delay of CSMA can grow exponentially in the network size. Therefore, even for simple interference topologies such as the grid topology, CSMA policies incur an exponentially large delay, which makes these policies less than attractive for large-scale implementation.

In the view of the above, our main contributions are as follows:

- 1) We formulate a model that allows to characterize the transient behavior of CSMA, i.e. that allows to study how the throughput of a CSMA policy changes over time if we start with a network where all nodes are silent (do not transmit). The model is based on two simplifying assumptions that allow a formal analysis of the transient behavior.
- 2) Using the model, we show that CSMA policies converge to a maximum schedule that achieves a high throughput at a rate that is independent of the network size. This insight is one of the key results of the paper.
- 3) Using the above insight, we propose a slight variant of the classical CSMA policy and show that its delay performance does not depend on the network size, i.e. the delay stays bounded even when the networks size approaches infinity. This result is rather surprising and at first counter-intuitive.
- 4) We also characterize the throughput-delay trade-off of the proposed CSMA policy and show that this trade-off is also independent of the network size.
- 5) While our analysis has been carried out for the special case of an interference graph with a grid (lattice) topology, we provide numerical case studies for general network topologies and show that the intuition obtained from the analysis carries over to these general cases. We

also illustrate the performance of the proposed CSMA policy when combined with a flow control mechanism. The above results suggest that CSMA policies are not only simple and can be implemented in a distributed manner, but that they can also achieve a very good performance both in terms of throughput and delay. The analysis shows that CSMA policies can obtain good performance even for very large networks, i.e. the delay performance and the delay-throughput trade-off are not dependent on the network size. To the best of our knowledge, this is the first time that such a result has been obtained. The result opens up the possibility of using CSMA policies for traffic that requires both a high-throughput and a small delay.

II. HIGHLIGHTS OF TECHNICAL RESULTS

To explain our main results, we first briefly review the CSMA delay behavior in the grid interference graph. In this graph, every wireless link is represented by a *node*. Each node itself is represented by coordinates (i, j) , $i, j \in \{0, 1, \dots, n - 1\}$, $n > 1$, so that all wireless links are nodes in a two dimensional grid (lattice), which can be interpreted as a subset of \mathbb{Z}^2 . Hereafter, we use the terms “wireless link” and “node” interchangeably. We use $L = n^2$ to denote the total number of nodes in the grid. We assume every node interferes with those nodes of the grid that are at unit distance from the node. Hence, every node has at most four interfering nodes. We assume that the packet arrival rate to all links is the same and equal to λ .

An important question for the above network is to determine the maximal packet arrival rate λ that can be stabilized by a CSMA policy. By now it is well-known that for any packet arrival rate $\lambda < 0.5$, there exists a CSMA policy that can stabilize the network [7], [9]. Roughly, this can be achieved by choosing a CSMA policy in which all nodes use the same transmission attempt-rate z , and by letting z become larger and larger. In the case where the attempt rate z becomes large, the network state will mainly alternate between two transmission patterns. The first (respectively second) pattern consists of transmitting nodes that are mostly even (respectively odd) nodes, i.e., those for which $i + j$ is even (respectively odd). However, transitions between these two transmission patterns occur very infrequently, i.e., the CSMA policy tends to “lock into” one of the two transmission patterns for a very long time before it switches to the other pattern.

It is well-known that this “locking-in” behavior of the CSMA policy is necessary in order to achieve a high throughput, i.e., in order to support an arrival rate λ close to 0.5. However, while locking into a transmission pattern benefits the throughput, it dramatically hurts the delay performance. An exact delay analysis of this case is difficult; the currently best characterization of the delay performance (see, e.g., [5]) provide an upper-bound for the delay of the form z^L , i.e. a delay bound which increases exponentially in the network size L . To support an arrival rate of $\lambda = 0.5 - \epsilon$, $\epsilon > 0$, it has been shown that an attempt rate z roughly of the order $z = 1/\epsilon$ is needed [8]. As a result, in order to support an arrival rate of $\lambda = 0.5 - \epsilon$, $\epsilon > 0$, the upper-bound on the delay can grow as

fast as $(1/\epsilon)^L$. Furthermore, for the case where ϵ is sufficiently small, a lower-bound on the delay has been derived that grows exponentially with exponent $\sqrt{L}/(\log L)^2$ [10]. These bounds imply that the delay performance of CSMA becomes very poor for large networks, i.e. as L becomes large.

To better understand the delay performance of CSMA policies, we propose a model to characterize the transient behavior of these policies, i.e., how fast these policies converge to maximum schedules. For this model, we show that, quite surprisingly, the rate of convergence of CSMA policies to maximum schedules is independent of the network-size. We then use this insight to design a new CSMA policy. We show that the delay of the new CSMA policy is $O(\frac{1}{\epsilon^3})$, which is independent of the network-size. In other words, the delay does not increase as the network-size increases to infinity. Using this result, we obtain that the delay-throughput trade-off is also independent of the network-size.

While our analysis was for the special case of the grid interference graph with uniform attempt-rates, numerical case studies that we performed suggest that the intuition obtained for the grid-topology carries over to more general interference topologies. In addition, our numerical results show that when the new CSMA policy is combined with a flow control mechanism, then the delay is also of the order $O(\frac{1}{\epsilon^3})$, where ϵ is defined as the “distance” to the optimal network utility. These results suggest that the proposed CSMA mechanism performs well in practical settings.

REFERENCES

- [1] M. Garetto, T. Salonidis, and E. W. Knightly, “Modeling per-flow throughput and capturing starvation in csma multi-hop wireless networks,” *IEEE/ACM Trans. Netw.*, vol. 16, no. 4, pp. 864–877, 2008.
- [2] J.-W. Lee, M. Chiang, and A. Calderbank, “Utility-optimal random-access control,” *IEEE Trans. Wireless Commun.*, vol. 6, no. 7, pp. 2741–2751, July 2007.
- [3] C. Bordenave, D. McDonald, and A. Proutiere, “Performance of random medium access control, an asymptotic approach,” in *SIGMETRICS '08: Proceedings of the 2008 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, 2008, pp. 1–12.
- [4] L. Jiang and J. Walrand, “A distributed algorithm for maximal throughput and optimal fairness in wireless networks with a general interference model,” EECS Department, University of California, Berkeley, Tech. Rep., Apr 2008. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-38.html>
- [5] S. Rajagopalan, D. Shah, and J. Shin, “Network adiabatic theorem: an efficient randomized protocol for contention resolution,” in *SIGMETRICS '09: Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems*, 2009, pp. 133–144.
- [6] J. Ni and R. Srikant, “Distributed CSMA/CA algorithms for achieving maximum throughput in wireless networks,” 2009. [Online]. Available: <http://www.citebase.org/abstract?id=oai:arXiv.org:0901.2333>
- [7] M. Durvy, O. Dousse, and P. Thiran, “Border effects, fairness, and phase transition in large wireless networks,” in *IEEE INFOCOM'08*, April 2008, pp. 601–609.
- [8] D. S. Gaunt and M. E. Fisher, “Hard-Sphere Lattice Gases. I. Plane-Square Lattice,” *Journal of Chemical Physics*, vol. 43, pp. 2840–2863, Oct. 1965.
- [9] D. Aldous, G. Grimmett, C. Howard, H. Kesten, F. Martinelli, L. Saloff-coste, and J. Steele, *Probability on Discrete Structures*. Springer, 2004.
- [10] C. Borgs, J. Chayes, A. Frieze, J. H. Kim, P. Tetali, E. Vigoda, and V. H. Vu, “Torpid mixing of some monte carlo markov chain algorithms in statistical physics,” in *40th Annual Symposium on Foundations of Computer Science*, 1999, pp. 218–229.