Stability of a standard decentralised medium access

Seva Shneer[⋆] and Alexander Stolyar[†]

[⋆]Heriot-Watt University, Edinburgh, UK
[†]University of Illinois at Urbana-Champaign, USA

ABSTRACT

We consider a stochastic queueing system modelling the behaviour of a wireless network with nodes employing a discrete-time version of the standard decentralised medium access algorithm. The system is *unsaturated* – each node receives an exogenous flow of packets at the rate λ packets per time slot. Each packet takes one slot to transmit, but neighbouring nodes cannot transmit simultaneously. The algorithm we study is *standard* in that: a node with empty queue does *not* compete for medium access; the access procedure by a node does *not* depend on its queue length, as long as it is non-zero. Two system topologies are considered, with nodes arranged in a circle and in a line. We prove that, for either topology, the system is stochastically stable under condition $\lambda < 2/5$. This result is intuitive for the circle topology as the throughput each node receives in a saturated system (with infinite queues) is equal to the so-called *parking constant*, which is larger than $2/5$. (This fact, however, does not help to prove our result.) The result is not intuitive at all for the line topology as in a saturated system some nodes receive a throughput lower than $2/5$.

Keywords

Stochastic stability, queueing networks, non-monotone process, discrete parking process, wireless systems, medium access protocols, Carrier-Sense Multiple Access

1. INTRODUCTION

In this paper we consider a stochastic queueing model, primarily motivated by the MAC (Medium Access Control) protocols (algorithms) in wireless networks. The model in this paper corresponds to those MAC protocols that prevent neighbours from transmitting simultaneously, thus making any collision and loss of packets impossible. We focus on the *single-hop* wireless networks, where each generated packet only needs to be transmitted once, by its source node (where the packet was generated).

We are interested in (stochastic) stability of a network, i.e. the ability of all nodes in the network to transmit packets, without the packet queues building up to infinity. If an access algorithm guarantees stability as long as such is in principle possible, it is called maximally stable (or throughput-optimal). The celebrated MaxWeight (or, BackPressure) algorithms, originated in [9] and later much extended and generalised, are known to be maximally stable. However, MaxWeight algorithms are centralised, in that they need a central controller to know the states of the queues of all nodes and then to solve an (often hard) optimisation problem to make every access control decision. This makes a practical use of MaxWeight difficult or infeasible in large networks. There has been therefore a need for the decentralised algorithms where each node regulates its own access to the medium.

A well-known and widely used (most notably in IEEE 802.11) example of a decentralised MAC algorithm preventing collisions is the CSMA (Carrier Sense Multiple Access) protocol. It represents a type of multiple access which we will refer to as *standard CSMA*, where 'standard' refers to the following properties: (a) each node does not know (and does not try to explicitly learn) its neighbours, (b) the access procedure is same regardless of the node queue length, as long as it is non-zero, and (c) the node does not access the channel when its queue length is zero (i.e., no packets to transmit). Most results on networks governed by the standard CSMA protocol assume that the system is saturated, i.e. each node is assumed to always have packets to transmit, and the question of interest is the throughputs of the nodes. See e.g. [2] and references therein for this type of work.

In practice the systems are unsuturated, i.e. the nodes "receive" a random process of generated packets and thus do not always compete for medium. Such unsaturated systems may *not* be approximated by saturated ones, or even bounded by them. For example, if one exchanges an unsaturated node for a saturated one, this will be detrimental to the performance of the node's neighbours but beneficial for the neighbours' neighbours (with whom the original node does not interfere). In other words, the random process describing such a system does *not* have the *monotonicity* property. Monotonicity, informally speaking, means that two versions of the process, with initial state of one dominating that of the other, can be coupled so that this dominance relation persists at all times. The absence of monotonicity dramatically complicates the process analysis, including establishing stability conditions. It is also impractical in applications to keep a device accessing the network if it does not have packets to transmit. (Hence property (c) of a standard CSMA is important.) The analysis of general unsaturated networks is extremely difficult, because, primarily due to lack of monotonicity, the queue dynamics and transmission schedule of any node depends on the states of all other nodes, in a very complicated manner.

MAMA 2018 Workshop Irvine, California, USA

Very few results are known for unsaturated systems. For example, in [10] the authors consider a continuous-time model and study the question of stability. They demonstrate with an example that the condition that the arrival intensity for each node is smaller than the throughput of the node in the saturated system is not sufficient for stability. This means, in particular, that standard CSMA does *not* achieve maximum stability. The absence of maximal stability of the standard CSMA protocol led to the development, starting with [3], of queue-based algorithms where the access procedure of a node depends on its own queue length. In particular, an algorithm of this type was proposed in [5], which guarantees maximal stability for single-hop systems on any graph. However, the queue-based algorithms are known to lead to high delays. (Hence, property (b) of a standard CSMA is important).

In this paper, our goal is to characterise the performance of a standard CSMA algorithm. We would like to stress again that standard CSMA algorithms are important, because they are decentralised (and therefore easy to implement), because implementing a queue-based scheme leads to long delays, and because it is impractical in the unsaturated situation to keep all nodes active at all times.

2. FORMAL MODEL AND MAIN RESULT

Our formal model is as follows. We consider two versions of a simple single-hop system, consisting of N nodes, indexed $i = 1, \dots, N$, arranged in a circle or in a line, respectively. The systems operate in discrete time, with unit time slots. The exogenous arrival process (of packets) at each node is i.i.d. with mean λ. To model standard CSMA in discrete time, we will assume that at the beginning of each time slot the nodes are given access priorities, forming a permutation of numbers $1, \dots, N$, picked independently (across time slots), uniformly at random from all possible permutations. The node with the highest priority transmits in this slot if its queue is non-empty. The node with the second-highest priority transmits in the slot if its queue is non-empty and if none of its neighbour nodes is transmitting. And so on until all nodes are checked in their priority order. Each transmission takes one slot, so at the beginning of the next time slot no transmission is ongoing, and the medium access process is repeated independently. Note that this algorithm is easy to implement in a decentralised fashion, with an arbitrarily small loss in throughput (see [7]). Denote by $Q_i(n)$ the queue length at node i at time n. Then, $Q(n) = (Q_i(n)),\ n = 0, 1, 2, \dots$, is a countable Markov chain. The system (stochastic) stability is understood as positive recurrence of $Q(\cdot)$. We are interested in the stability conditions.

If we consider the circle system, if all nodes have non-empty queues, then by symmetry the average service rate (the expected number of transmissions per slot) is the same for all nodes and is equal to the so-called parking constant (sometimes referred to as jamming density), which we denote by $c_p(N)$. Therefore, if λ exceeds $c_p(N)$, the system is unstable. The parking constant $c_p(N)$ is equal to $1/2$ if $N = 4$, $2/5$ if $N = 5$ and it decreases over even values of N and increases over odd values of N to the same limit $1/2(1 - e^{-2}) \approx 0.4323$. On the other hand, each non-empty node always transmits if its priority is higher than the priorities of both its neighbours, which happens with probability $1/3$. Therefore $\lambda < 1/3$ trivially leads to stability. One of the goals of our work is to study how close the stability region of the standard CSMA algorithm on a circle is to $\lambda < c_p(N)$ (the best possibly achievable for this algorithm). The **main result** of the paper for the **circle topology** is that, under the standard CSMA, $\lambda < 2/5$ leads to stability for a system with $N \geq 4$. We conjecture that $\lambda < c_p(N)$ is sufficient for stability, but this is not proved.

It appears to be a lot more difficult to conjecture the stability condition for a system on a line. If all nodes have non-empty queues, then the process of transmission scheduling in a slot is again equivalent to the discrete parking problem, but in this situation the expected number of transmissions differs from node to node. For instance, the nodes on the edges will have transmission probabilities larger than $1/2$ and tending to $1 - e^{-1} \approx 0.6321$ as $N \to \infty$, but the nodes right next to them will have much lower transmission probabilities tending to $e^{-1} \approx 0.3679$ (see [6]). Our **main result** for the **line topology** is that, just like for the circle system, $\lambda < 2/5$ leads to stability for a system with $N \geq 4$. This result is not intuitive at all, given that in the saturated system the transmission probability of the second node from the edge is strictly less than $2/5$ (see above). The result also stresses once again that the saturated system does not provide a bound for the unsaturated one, that it is unreasonable to make all nodes (including empty) to compete for the transmission at all times, and that it is important to study unsaturated systems.

To summarise, the main result of this paper is

THEOREM 1. *The system with either the circle or line topology, and $N \geq 4$, is stable if $\lambda < 2/5$.*

REMARK 2. *Consider the unsaturated circle system, but the access protocol is modified so that* all nodes compete for transmission at all times, even when they are empty. *Such a system is then trivially stable under the condition $\lambda < c_p(N)$, because the average service rate received by each node is exactly $c_p(N)$ as long as the node is occupied. (This fact also does* not *imply our Theorem 1 for the circle system and does* not *help to prove it.) Therefore, the modified protocol is maximally stable when all nodes receive arrivals at the same rate λ. One may ask: Isn't the modified protocol then "better" than the standard CSMA that we study? The answer is: Yes, it might be,* if *it is known in advance that the system topology is a circle and* if *the arrival rates are equal. Our goal, however, is to study properties of a protocol, under which each node is completely unaware of the rest of the system – in this case making each node to compete for the channel at all times (even when empty) is clearly a "bad" approach, because often only a small fraction of network nodes have any traffic to transmit.*

Theorem 1 is proved using the fluid limit technique. The main difficulty in the proofs consists in dealing with situations when some of the queues of a fluid limit are at zero. One has to study rather complicated structure of the occupancy and activation processes of a neighbourhood of non-zero queues. We believe that the technique we developed may be used to study the behaviour of CSMA and similar decentralised algorithms on more complicated topologies.

3. DISCUSSION OF THE PROOF

The proof of Theorem 1 is quite technical. Due to space limitation, here we only present a discussion, highlighting the key difficulties. The complete proof is in [7].

Let us discuss the proof for the circle topology, to be specific. It uses the fluid limit technique [4, 1, 8]. Recall that the process fluid limit is, informally speaking, a process $q(t) = (q_i(t)),\ t \geq 0$ arising as an appropriate limit of a sequence of rescaled processes $Q^r(rt)/t,\ t \geq 0$, with the initial norm $\|Q^r(0)\| = \sum_i Q_i^r(0) = r$ increasing to infinity. To prove stability it is sufficient to show that every trajectory $q(\cdot)$ of the fluid limit, with initial norm $\|q(0)\| = \sum_i q_i(0) = 1$, reaches 0 within a finite time.

To prove the latter, we show that there exists $\epsilon > 0$ such that for any fluid limit trajectory, at any regular point t such that $\sum_{i=1}^N q_i(t) > 0$, $\sum_{i=1}^N q_i'(t) \leq -\epsilon$. (A time point t is regular if the derivatives $q_i'(t)$ for all i exist. Almost all points t are regular.)

The case when all $q_i(t) > 0$ is relatively easy. Indeed, this case corresponds to the situation when all queues of the original process are continuously occupied and, therefore, each of them receives the service at the average rate equal to the parking constant $c_p(N)$. This implies $q_i'(t) = \lambda - c_p(N)$, and it is not hard to show that $c_p(N) \geq 2/5$ for $N \geq 4$.

The main difficulty is to deal with states $q(t)$ in which $q_i(t) = 0$ for some queues i. (And this, in turn, stems primarily from the non-monotonicity of the process. For example, we *cannot* claim that the derivatives of positive $q_i(t)$ are dominated by the derivatives in the state where all $q_i > 0$.) In this case, a group of nodes $(k+1, k+2, \ldots, k+l)$ such that $q_k(t) = q_{k+l+1}(t) = 0$ and $q_{k+i}(t) > 0$ for all $i = 1, \ldots, l$, we will call a *positive group* of size l. (Node indices i and $N+i$ correspond to the same node i, when nodes arranged in a circle.) We prove that

$$\sum_{i=1}^{l} q_{k+i}'(t) < -\epsilon_1, \qquad (1)$$

for any positive group $(k+1, k+2, \ldots, k+l)$, for some fixed constant $\epsilon_1 > 0$.

To prove (1) we consider four cases separately, $l = 1, 2, 3$ and $l \geq 4$. The case $l = 1$, when the positive group consists of the single node $k+1$, is the most difficult. Indeed, $q_{k+1}'(t)$ depends on the joint occupancy distribution of all other nodes $i \neq k$, in the process "local steady-state" defined by the condition that node $k+1$ is continuously occupied. Such a steady-state is very hard to characterise exactly, and we derive quite intricate estimates based on considering a neighbourhood of node $k+1$ of up to 3 nodes on each side. In the case $l = 2$, we obtain a required estimate of $q_{k+1}' + q_{k+2}'$, which is based on considering nodes $k, k+1, k+2$; this estimate in fact applies to any pair of nodes at either edge of a positive group. In the case $l = 3$, we obtain an estimate of $q_{k+1}' + q_{k+2}' + q_{k+3}'$, based on considering nodes $k+1, k+2, k+3$. In the case $l \geq 4$ we apply the estimates obtained for $l = 2$ to the two pairs of nodes, $(k+1, k+2)$ and $(k+l-1, k+l)$, which are at the edges of the positive group; then we additionally derive a universal negative upper bound on $q_i'(t)$ which applies for any i at the distance at least 3 (including itself) from the edge of a positive group.

4. FUTURE WORK

We are currently also working on more difficult multi-hop networks where each message has a source and a destination and may therefore need to be transmitted by several nodes in the network. For a line we consider traffic arriving at node 1 and requiring to reach node N. Each message thus needs to be transmitted by every node in turn. We are interested in stability and end-to-end throughput of the system.

For a circle one can define a multi-hop network in, for instance, the following way. Fix a constant $m \geq 1$ and assume that each node gets on average λ/m new packets per time slot. Assume that medium-access competition is done in the same way as for the single-hop system but, upon a successful transmission from node i, a message leaves the system with probability $1/m$ and goes to the queue of node $i+1$ with probability $1 - 1/m$. It is easy to show that the average nominal traffic load for any node is λ and we conjecture that $\lambda < c_p(N)$ is sufficient for stability for any $m \geq 1$. As $m \to \infty$ and $N \to \infty$ the model on a circle may serve as an approximation for the behaviour of a large line segment far away from source.

We can in fact prove the above multi-hop stability conjecture for $N = 4$, but this proof heavily relies on the 4-node structure and does not generalise to larger values of N. Stability for all N and all $m \geq 1$ is a challenging and exciting question for further study.

5. REFERENCES

[1] J.G. Dai (1995). On the Positive Harris Recurrence for Open Multiclass Queueing Networks: A Unified Approach Via Fluid Limit Models. *Annals of Applied Probability*, **5**, 49–77.

[2] M. Durvy, O. Dousse, P. Thiran (2009). Self-organization properties of CSMA/CA systems and their consequences on fairness. *IEEE Transactions on Information Theory*, **55 (3)**, 931?-943.

[3] L. Jiang, J. Walrand (2010). A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Transactions on Networking (TON)*, **18(3)**, 960–972.

[4] A.N. Rybko and A.L. Stolyar (1992). Ergodicity of Stochastic Processes Describing the Operation of Open Queueing Networks. *Problems of Information Transmission*, **28**, 199–220.

[5] D. Shah, J. Shin (2012). Randomized scheduling algorithm for queueing networks. *Annals of Applied Probability*, **22(1)**, 128–171.

[6] S. Shnner, P. van de Ven (2017). Per-site occupancy in the discrete parking problem. *Statistics and Probability Letters*, **120**, 141–146.

[7] S. Shneer, A. L. Stolyar. Stability conditions for a discrete-time decentralised medium access algorithm. Submitted. https://arxiv.org/abs/1707.01548

[8] A.L. Stolyar (1995). On the Stability of Multiclass Queueing Networks: A Relaxed Sufficient Condition via Limiting Fluid Processes. *Markov Processes and Related Fields*, **1(4)**, 491–512.

[9] L. Tassiulas, A. Ephremides (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, **37(12)**, 1936–1948.

[10] P.M. van de Ven, S.C. Borst, J.S.H. van Leeuwaarden, A. Proutière (2010). Insensitivity and stability of random-access networks. *Performance Evaluation*, **67 (11)**, 1230–1242.