Reaction-Diffusion Based Transmission Patterns for Ad Hoc Networks

Mathilde Durvy

School of Computer and Communication Sciences
EPFL
CH-1015 Lausanne, Switzerland
mathilde.durvy@epfl.ch

Patrick Thiran

School of Computer and Communication Sciences
EPFL
CH-1015 Lausanne, Switzerland
patrick.thiran@epfl.ch

Abstract— We present a new scheme that mimics pattern formation in biological systems to create transmission patterns in multi-hop ad hoc networks. Our scheme is decentralized and relies exclusively on local interactions between the network nodes to create global transmission patterns. A transmission inhibits other transmissions in its immediate surrounding and encourages nodes located further away to transmit. The transmission patterns created by our medium access control scheme combine the efficiency of allocation-based schemes at high traffic loads and the flexibility of random access schemes. Moreover, we show that with appropriately chosen parameters our scheme converges to collision free transmission patterns that guarantee some degree of spatial reuse.

I. INTRODUCTION

Wireless ad hoc networks rely on a common transmission medium, called the channel. The role of the Medium Access Control (MAC) scheme is to coordinate the access of the network nodes to this channel to guarantee an efficient usage of this resource.

In local area networks, where all transmissions take a single hop, the only resource is time: a good MAC protocol schedules transmissions from different users in order to maximize the temporal usage of the channel, while maintaining fairness between the users. Wireless multi-hop networks pose a greater challenge, as now the resources are time and space: since a transmission consumes only a spatially restricted portion of the channel, a good MAC protocol should schedule transmissions from multiple users in order to maximize not only the temporal usage of the channel, but also its spatial usage. This latter property, called spatial reuse, is the central theme of this paper. How can a MAC protocol pack as many simultaneous transmissions as possible in a wireless multi-hop network?

Existing MAC protocols can be divided into allocation-based and random access schemes. In allocation-based schemes, a central authority shares the channel resources among the multiple users in a fixed manner. For example, in

Time-Division Multiple Access (TDMA), the network nodes are divided into disjoint subsets, which are then granted access to the channel in a round-robin manner. Figure 1(a) depicts a transmission pattern created by a TDMA scheme on a two-dimensional grid network. Active transmissions are highlighted by black disks and the gray area around an active transmission corresponds to its exclusion domain, which is the area where other nodes must remain silent to avoid a collision with the active transmission. In random access schemes, the network nodes contend for the medium in a decentralized manner. For example, in Carrier-Sense Multiple Access (CSMA), when a node has a packet to send, it first senses the medium. If the medium is idle, it transmits; whereas if the medium is busy, it backs-off for a random time. Carrier-Sense Multiple Access inspired numerous random access schemes, including the widely used IEEE 802.11.

Allocation-based MAC protocols perform well in environments where the traffic is predictable and the network topology is static. Random access schemes are more flexible but lack efficiency, especially when demand for bandwidth is high.

We propose a MAC scheme that combines the advantages of allocation-based and random access schemes to provide an efficient scheduling of the transmissions in ad hoc networks with high traffic load and contention. Like allocation-based schemes, it creates dense, collision-free transmission patterns. However, the transmission patterns are obtained dynamically in a distributed manner, like in random access schemes.

Our MAC scheme is based on the so-called reaction-diffusion mechanism. Reaction-diffusion was proposed by Turing in 1952 [1] to explain the formation of patterns in biological systems. Typical examples are mammalian coat patterns such as zebra stripes or felines' spots (Figure 1(b)). In Reaction-Diffusion MAC (RDMAC), we use the reaction-diffusion mechanism to create transmission patterns. Figure 1(c) illustrates the creation of a transmission pattern by RDMAC on a two-dimensional network where the nodes are placed on a grid. As in biological systems, the pattern formation is spontaneous and can adapt to changes in the environment.

(a) Transmission pattern created by a centralized TDMA strategy

(b) Mammalian coat pattern

(c) Transmission pattern created by the decentralized RDMAC scheme

Fig. 1. Illustration of different types of patterns: a crystalline pattern obtained by a centralized TDMA scheme (a), and a less regular pattern obtained by the decentralized RDMAC scheme (c), which is inspired by the mechanism governing pattern formation in biological systems such as mammalian coat patterns (b). In (a) and (c) we consider a two-dimensional grid network. Active transmissions are highlighted by black disks and the gray area around an active transmission corresponds to its exclusion domain.

The rest of the paper is organized as follows. In Section 2 we discuss related work. In Section 3 we introduce the system model and its assumptions. In Sections 4 and 5 we describe the RDMAC algorithm and discuss the choice of parameters. In Section 6 we present and analyze the simulation results. Finally, we conclude in Section 7.

II. RELATED WORK

Many MAC schemes have been designed to coordinate the access of network stations to the wireless medium. However, only decentralized MAC algorithms are suitable for ad hoc networks. The IEEE 802.11 Distributed Coordination Function (DCF) [2] is the most well-known decentralized MAC protocol. It is also one of the few that is actually implemented in industrial products. Yet, detailed experimental studies of the IEEE 802.11 DCF protocol reveal poor performance [3], especially when the traffic load is high. The quest for a decentralized, adaptive, efficient MAC protocol for ad hoc networks therefore remains, more than ever, an active area of research.

Recent work [4] gives some valuable insight on how the transmissions of the network stations should be organized in a one-dimensional multi-hop ad hoc network. It demonstrates that an ideal schedule consists of k equal time slots, where in each time slot the network nodes emit following a rotation of a single transmission pattern. These results are reminiscent of Time-Division Multiple Access schemes, such as Spatial TDMA [5], which proved very efficient in the context of networks with fixed topology and predictable, stable traffic loads. Spatial TDMA is a generalization of the TDMA protocol for multi-hop wireless networks. It

creates collision-free transmission patterns with a maximal number of simultaneous transmissions. However, the creation of the transmission pattern requires a global knowledge of the network topology and is thus not decentralized. In addition, the computation of a Spatial TDMA slot allocation is NP-Complete, which makes the algorithm unpractical for most network topologies. How to achieve such collision-free transmission patterns in a decentralized manner is still an open problem.

The scheme we propose provides a solution to this challenging problem. The key property of RDMAC is the creation of transmission patterns similar to the TDMA transmission patterns, but contrary to TDMA in a decentralized manner. The RDMAC scheme explicitly assigns a Medium Access Probability (MAP) p_i to each connection. In this aspect, our protocol is similar to the SR-Aloha protocol [6] and to the SEEDEX protocol [7]. Nevertheless, the three protocols strongly differ in the way the Medium Access Probabilities (MAP) are computed. In the SR-Aloha protocol, which is one of the first decentralized MAC protocols to optimize spatial reuse, the MAP is time- and node- invariant. In the SEEDEX protocol, the MAP of a network station depends on the state of its 2-hop neighbors. Network stations can be in two states, the Possibly Transmit (PT) state and the Listening (L) state. A station s in the PT state transmits to a station t in the L state with probability $p = 1/(n + 1)$, where n is the number of neighbors of t which are in the PT state. Finally, in our RDMAC protocol, the choice of the medium access probabilities is based on the reaction-diffusion mechanism, described in more detail in Section IV.

Fig. 2. MAC layer connections in a grid network topology. The transmission of connection i is successful if there is no other active node in its exclusion domain (the light gray area in the figure).

III. ASSUMPTIONS AND SYSTEM MODEL

A. Model

We assume a simple physical model where all active nodes emit at a fixed power P and two nodes s and t can communicate if their distance is lower than their *connection range* d. We call a direct *connection* (or a link) two nodes that can communicate directly. A transmission between two connected nodes is successful if no other transmission takes place in its *exclusion domain*. Consider a connection i between node s_i and node t_i. The exclusion domain $\mathcal{D}_E^i$ of a connection i is thus the area around this connection where no other node can be active in order to guarantee a successful transmission between nodes s_i and t_i. We assume that $\mathcal{D}_E^i$ contains at least all nodes in the communication range of either s_i or t_i. Moreover, since the transmission power is identical for all nodes, we further assume that the shape of the exclusion domain is the same for all connections. Finally, we define the exclusion range R_E as the largest distance between a border point of an exclusion domain $\mathcal{D}_E^i$ and the closest end-node s_i or t_i of the corresponding connection i. Any interfering transmission in the exclusion domain of an already active connection will cause a *collision*. Figure 2 depicts a grid network topology, where nodes occupy all vertices of a square lattice, as well as the set of possible connections between nodes. The exclusion domain of connection i between node s_i and t_i is depicted in gray. In this setting, a transmission between s_i and t_i is successful if there is no other active node in their 1-hop neighborhoods.

We consider bidirectional exchanges between connected nodes. A typical exchange consists of a data packet followed by an acknowledgment. In some protocols, the exchange might also include a preliminary handshake (e.g. the Request-To-Send, Clear-To-Send handshake in the IEEE 802.11 protocol). Consequently, a transmission from s_i to t_i or from t_i to s_i will result in the same exclusion domain $\mathcal{D}_E^i$ since both nodes should be able to receive. Hence, for the purpose of MAC algorithm design, we do not distinguish between a

transmission from node s_i to t_i and a transmission from node t_i to s_i. Our approach follows the framework developed in [8] where connections, and not nodes, contend to access the channel.

We study probabilistic MAC schemes. At time t the MAC protocol explicitly assigns a *Medium Access Probability* (MAP) $p_i(t)$ to each connection i. If $e_i(t)$ is the random variable describing the state of connection i at time t, we have thus

$$e_i(t) = \begin{cases} 1 & \text{with probability } p_i(t) \\ 0 & \text{with probability } 1 - p_i(t). \end{cases} \quad (1)$$

Connection i is said to be *active* if $e_i(t) = 1$ and *idle* if $e_i(t) = 0$. It is said to be *almost surely active*, or a.s. active, if $p_i(t) = 1$, and *almost surely idle* (a.s. idle) if $p_i(t) = 0$. Denote by $S_i(t)$ the random variable that indicates that transmission i is successful at time t. We can then write $S_i(t)$ as

$$S_i(t) = e_i(t) \prod_{j \in \mathcal{D}_E^i} (1 - e_j(t)), \quad (2)$$

which means that transmission i is successful at time t ($S_i(t) = 1$) if no other transmission takes place in its exclusion domain $\mathcal{D}_E^i$ ($e_j(t) = 0$ for all $j \in \mathcal{D}_E^i$).

B. Spatial Reuse

The goal of this work is to study the ability of MAC algorithms to create dense transmission patterns in a distributed manner. We need a metric to evaluate what level of spatial reuse is achieved. We define the *density of spatial reuse* of a transmission pattern, which we denote by ρ, as the average number of successful transmissions (thus excluding transmissions subject to collisions) per unit area in this pattern. The throughput obtained at high traffic loads by a MAC algorithm is proportional to the density of spatial reuse it can achieve.

We do not address the issue of information sharing between the nodes, which is bound to arise when it comes to a real implementation of the algorithm. Rather, we assume that each node has a perfect knowledge of the information related to the other nodes in a constant size neighborhood around itself. We do not study other factors that clearly impact the MAC performance, such as the back-off mechanism when there are collisions, or the combination of routing with MAC, so that we can focus our study on the metric of spatial reuse.

IV. RDMAC: DESCRIPTION OF THE ALGORITHM

The novelty of the RDMAC scheme is embedded in the computation of the medium access probabilities. We already mentioned that the choice of the MAPs is inspired by the reaction-diffusion mechanism, which has been used in the past to explain the formation of patterns in biological systems. A reaction-diffusion system involves two substances, the *activator* and the *inhibitor*, which both diffuse within the system boundaries. The state of each point of the system depends on the relative concentration of the activator and inhibitor at

its location. Denote by $a(\mathbf{z}, t)$ and $h(\mathbf{z}, t)$ the concentration of the activator and inhibitor in the system at location $\mathbf{z}$ and time t. The reaction-diffusion mechanism is described in terms of second order partial differential equations ([1],[9]) of the form

$$\frac{\partial a}{\partial t} = f(a, h) + D_A \nabla^2 a$$
$$\frac{\partial h}{\partial t} = g(a, h) + D_H \nabla^2 h$$

PSfrag replacements

where ∇^2 is the Laplacian operator with respect to $\mathbf{z}$, where f and g are nonlinearities describing the reaction kinetics, and where D_A and D_H are the diffusion rate of respectively the activator and the inhibitor. Simpler cellular automaton models such as the *Cellular Neural Network* (CNN) can also capture the most salient features of the reaction-diffusion mechanism [10].

A CNN is an array of identical systems, the cells, which are only locally connected. Time is continuous or discrete. In the discrete time setting, the state x_i of cell i at time $t+1$ depends on its output y_i at time t and on the activatory and inhibitory inputs it gets from neighboring cells (Figure 3). Such an input can be positive, which induces activity, or negative, which inhibits it.

PSfrag replacements Input from other cells x_i $y_i = f(x_i)$

Fig. 3. A typical cell in a CNN. The state of the cell x_i depends on its previous output y_i and on the influence of the other cells around itself.

In our model, a cell corresponds to a direct connection at the MAC layer, and the output y_i of the cell is the probability p_i that the connection is active. Any transmission in the exclusion domain $\mathcal{D}_E^i$ of an already active connection i must decrease its probability p_i to send in order to avoid collision. Reciprocally, connection i inhibits other connections located in the area $\mathcal{D}_E^i$. However, in order to maximize spatial reuse, an active connection encourages network nodes just outside its exclusion domain to transmit. More specifically, the connections located in an area $\mathcal{D}_A^i$ surrounding $\mathcal{D}_E^i$ are activated by connection i. The area $\mathcal{D}_A^i$ is called the *activation domain* of connection i. Recall the grid topology introduced in Section III-A. The resulting activation and exclusion domains, $\mathcal{D}_A^i$ and $\mathcal{D}_E^i$, and the corresponding activation range and exclusion range, R_A and R_E, are represented in Figure 4. The size of the activation area is arbitrary. Yet in order to maximize spatial reuse it should neither be empty nor be too large, since we want a connection to activate nodes as close as possible to the border of its exclusion domain.


Fig. 4. The activation $\mathcal{D}_A^i$ (dark gray) and exclusion $\mathcal{D}_E^i$ (light gray) domains of connection i in a two-dimensional grid network.

These reaction-diffusion principles are translated in the update rule of the medium access probabilities. Whether we choose a continuous-time (differential equation) or discrete time (recurrence equation) for this update rule does not matter a lot, as the quantity of interest is the density of spatial reuse reached in steady state. Going for a discrete-time formulation, we embed the properties of the RDMAC algorithm in the MAP update equation

$$p_i(t+1) =$$
$$f\left(l p_i(t) - s \sum_{z \in \mathcal{D}_E^i} p_z(t) + r \sum_{z \in \mathcal{D}_A^i} p_z(t) \right), \quad (3)$$

where the function f is the piecewise linear function shown in Figure 5. The MAP update equation (3) admits three positive parameters: l, the self-activation parameter, s, the inhibition parameter, and r, the activation parameter. Transmissions in the spatial domain $\mathcal{D}_E^i$ have an inhibitory influence on connection i, whereas transmissions in $\mathcal{D}_A^i$ have an activating effect on connection i. The coupling parameters l, s, and r are time invariant (they do not depend on the time t), and space invariant (they do not depend on the connection number i).

In a real network, the nodes are usually not perfectly synchronized. Therefore, it is more realistic to assume an asynchronous update of the MAP, where all connections successively update their MAP. We assume that all connections update their MAP once in a time interval $[t, t+1]$ and we refer to the global update of the MAP as the *sequential iteration* at time t. A random permutation π of $\{1, \ldots, N\}$ determines the order in which the connections update their MAP.

Updating the probability to transmit according to Equation (3) creates spontaneous transmission patterns, like the one of Figure 1(c). Each active connection is surrounded by an exclusion domain where no other transmission is active. The transmission patterns created depend on the values of the parameters l, s and r in Equation (3). We will show in Section V that there exist values of l, s and r such that RDMAC can guarantee the establishment of collision free transmission patterns.

Fig. 5. The nonlinearity $f(\cdot)$ as in the MAP update equation (3)

V. WHAT IS A GOOD CHOICE OF PARAMETERS?

The RDMAC scheme has only the three parameters, l, s and r of Equation (3). We next discuss how to set these parameters to maximize the number of simultaneous successful transmissions in the network.

A. Characterization of Optimal Transmission Patterns

The problem of maximizing the number of simultaneous successful transmissions in a network is equivalent to the maximum independent set problem in an undirected graph $G = (V, L)$. The set $V = \{1, 2, \ldots, N\}$ of vertices corresponds to the direct connections of the network and the set L includes an edge for each pair of vertices i, j such that $i \in \mathcal{D}_E^j$ and $j \in \mathcal{D}_E^i$. The maximum independent set problem is NP-Complete [11] and can only be solved for small network topologies. Figure 1(a) illustrates an optimal transmission pattern, i.e., a solution of the maximum independent set problem, on a 20×20 grid topology. For such a network size it is already computationally unfeasible to solve the maximum independent set problem, but the regularity of the grid topology makes it possible to extend by symmetry the solution from a smaller instance of the problem. Once the transmission patterns corresponding to the solutions of the maximum independent set problem are known, they can be combined to form a TDMA schedule. In the remaining of the paper we denote by $\mathbf{e_{opt}}$ a state vector that corresponds to an optimal transmission pattern ($e_i = 1$ if connection i is active in the optimal transmission pattern and $e_i = 0$ otherwise).

B. Convergence of the RDMAC Algorithm to Equilibrium Points

The MAP update rule (3) defines a dynamical system. We show now that this system always converges to equilibrium points, and thus that it will not oscillate. The stable equilibrium points of the system correspond to the transmission patterns created by the RDMAC algorithm.

A vector $\mathbf{p} = [p_i]_{1 \leq i \leq N}$ is an equilibrium of the system defined by (3) if, for all $1 \leq i \leq N$,

$$p_i = f\left(l p_i - s \sum_{z \in \mathcal{D}_E^i} p_z + r \sum_{z \in \mathcal{D}_A^i} p_z \right). \qquad (4)$$

Theorem 1: All sample paths of the MAP update equation (3), with asynchronous updates, converge to equilibrium points.

Proof: Define the function $V[\mathbf{p}]$ by

$$V[\mathbf{p}] = -\frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} a_{ij} p_i p_j + \frac{1}{2} \sum_{i=1}^{N} p_i^2, \qquad (5)$$

where

$$a_{ij} = \begin{cases} l & \text{if } j = i \\ -s & \text{if } j \in \mathcal{D}_E^i \\ r & \text{if } j \in \mathcal{D}_A^i \\ 0 & \text{otherwise.} \end{cases} \qquad (6)$$

A sufficient condition for all trajectories of the system to converge to equilibrium points is (LaSalle's Theorem [12])

$$\mathbf{p}(t+1) \neq \mathbf{p}(t) \Rightarrow V[\mathbf{p}(t+1)] - V[\mathbf{p}(t)] < 0, \qquad (7)$$

which establishes that $V[\mathbf{p}]$ is a Lyapunov function. We can decompose the update of the medium access probabilities at time t into N successive steps, $t + k/N$, $1 \leq k \leq N$, where at each step only one of the connections updates its MAP, so that at time $t + 1$ all components of $\mathbf{p}$ have been updated exactly once. Therefore

$$V[\mathbf{p}(t+1)] - V[\mathbf{p}(t)] =$$
$$\sum_{k=1}^{N} \left(V\left[\mathbf{p}\left(t + \frac{k}{N} \right) \right] - V\left[\mathbf{p}\left(t + \frac{k-1}{N} \right) \right] \right).$$

Condition (7) is then satisfied if whenever

$$\mathbf{p}\left(t + \frac{k}{N} \right) \neq \mathbf{p}\left(t + \frac{k-1}{N} \right),$$
$$V\left[\mathbf{p}\left(t + \frac{k}{N} \right) \right] - V\left[\mathbf{p}\left(t + \frac{k-1}{N} \right) \right] < 0.$$

Consider, without loss of generality, that connection 1 is the first to update its MAP. Then

$$\mathbf{p}\left(t + \frac{1}{N} \right) = [p_1(t+1), p_2(t), \ldots, p_N(t)].$$

Rewrite (5) as

$$V[\mathbf{p}(t)] = -\frac{1}{2} \left(a_{11} p_1^2(t) + 2 p_1(t) \sum_{j=2}^{N} a_{1j} p_j(t) + \right.$$
$$\left. \sum_{i=2}^{N} \sum_{j=2}^{N} a_{ij} p_i(t) p_j(t) \right) + \frac{1}{2} \left(p_1^2(t) + \sum_{i=2}^{N} p_i^2(t) \right).$$

Using a similar expression for the function $V(\mathbf{p}(\cdot))$ at time

$t + 1/N$ we obtain

$$V[\mathbf{p}(t + \frac{1}{N})] - V[\mathbf{p}(t)] = -\frac{1}{2}\left(a_{11}p_1^2(t+1)+\right.$$

$$2p_1(t+1)\sum_{j=2}^{N} a_{1j}p_j(t) - p_1^2(t+1)$$

$$\left. -a_{11}p_1^2(t) - 2p_1(t)\sum_{j=2}^{N} a_{1j}p_j(t) + p_1^2(t)\right).$$

Adding and subtracting the terms $a_{11}p_1(t + 1)p_1(t)$ and $1/2a_{11}p_1^2(t)$ and rearranging the terms, we have

$$V[\mathbf{p}(t + \frac{1}{N})] - V[\mathbf{p}(t)] = -\frac{1}{2}(p_1(t+1) - p_1(t))$$

$$\left((a_{11} - 1)p_1(t+1) + 2\sum_{j=1}^{N} a_{1j}p_j(t) - (a_{11} + 1)p_1(t)\right).$$

With the notation (6), the MAP update equation (3) becomes

$$p_1(t+1) = f\left(\sum_{j=1}^{N} a_{1j}p_j(t)\right).$$

Consider the three regions defined by the three segments of the piecewise linear function f:

- If $\sum_{j=1}^{N} a_{1j}p_j(t) \geq 1$ then $p_1(t+1) = 1$ and

$$V[\mathbf{p}(t + \frac{1}{N})] - V[\mathbf{p}(t)] = -\frac{1}{2}(1 - p_1(t))$$

$$\left((a_{11} + (2\sum_{j=1}^{N} a_{1j}p_j(t) - 1)) - (a_{11} + 1)p_1(t)\right) < 0.$$

- If $\sum_{j=1}^{N} a_{1j}p_j(t) \leq 0$ then $p_1(t+1) = 0$ and

$$V[\mathbf{p}(t + \frac{1}{N})] - V[\mathbf{p}(t)] =$$

$$\frac{1}{2}p_1(t)\left(2\sum_{j=1}^{N} a_{1j}p_j(t) - (a_{11} + 1)p_1(t)\right) < 0.$$

- If $0 < \sum_{j=1}^{N} a_{1j}p_j(t) < 1$ then $p_1(t + 1) = \sum_{j=1}^{N} a_{1j}p_j(t)$ and

$$V[\mathbf{p}(t + \frac{1}{N})] - V[\mathbf{p}(t)] =$$

$$-\frac{1}{2}(a_{11} + 1)(p_1(t+1) - p_1(t))^2 < 0.$$

So for all three regions of the piecewise linear function f, $V[\mathbf{p}(t + 1/N)] - V[\mathbf{p}(t)] < 0$, which concludes the proof. ∎

C. Selection of the Stable Equilibrium Transmission Patterns

Can we choose the three parameters l, s and r such that all the trajectories of the non-linear system characterized by Equation (3) converge to equilibrium MAP vectors $\mathbf{p}$ close to $\mathbf{e_{opt}}$?

We proceed to characterize the stable equilibria of (3) in three steps. First, we show that at the system equilibrium,

all connections are either a.s. active ($p = 1$) or a.s. idle ($p = 0$). Second, we demonstrate that there is at least one a.s. active connection in the network. And finally, we show that a transmission pattern corresponding to an equilibrium of the system is collision free and guarantees some degree of spatial reuse.

Theorem 2: If $l > 1$, all connections are either a.s. active or a.s. idle at a stable equilibrium point of system (3). That is, $p_i \in \{0, 1\}$ for all connections i at a stable equilibrium.

Proof: Consider a group $\mathcal{M}$ of M connections at time t such that if $i \in \mathcal{M}$ then $0 < p_i(t) < 1$, and for all connections $j \in \mathcal{D}_E^i \cup \mathcal{D}_A^i$, either $j \in \mathcal{M}$, $p_j(t) = 0$ or $p_j(t) = 1$. We refer to the connections in group $\mathcal{M}$ as non-saturated. The MAP update equation (3) for this group of non-saturated connections are decoupled from the other equations. Consider the update of the non-saturated connection $i \in \mathcal{M}$, and assume without loss of generality that this connection is the first to be updated at iteration t, equation (3) can then be rewritten as

$$\mathbf{p}\left(t + \frac{1}{N}\right) = A_M\mathbf{p}(t) + \mathbf{b},$$

where A_M is the MxM matrix where rows are the same as those of the identity matrix except for row i that has coefficients a_{ij} equal to l, $-s$, r or 0 as defined by (6), and where b_i includes the contributions of the saturated connections on the update of connection i ($b_j = 0$ for all $j \neq i$). The sum of the eigenvalues of A_M is equal to its trace, which is $(M-1)+l > M$. Repeating the same reasoning with the next connection of $\mathcal{M}$ to be updated, we find that the group of non-saturated connections is described by a set of coupled linear recurrence equations, which is unstable because at each time at least one of the eigenvalues has a magnitude larger than 1. Therefore the system operating in the non-saturated region is unstable. ∎

Theorem 2 establishes that at the stable equilibrium of system (3), all connections are either a.s. active or a.s. idle. We now show that there is at least one a.s. active connection in the network.

Theorem 3: The all zero output, $\mathbf{p} = \mathbf{0}$, is an unstable equilibrium of the system.

Proof: Let us make an infinitesimal perturbation of the equilibrium $\mathbf{p} = \mathbf{0}$, as follows. Set $p_i(0) = \epsilon$ for one connection $1 \leq i \leq N$, where $\epsilon > 0$ is arbitrarily small, whereas $p_j(0) = 0$ for all $j \neq i$. As long as $p_j(t) = 0$ for $j \neq i$ and $p_i(t) \leq 1/l$,

$$p_i(t+1) \geq f(lp_i(t)) = lp_i(t) = l^{t+1}p_i(0) = \epsilon l^{t+1}.$$

Since $l > 1$, the latter equation shows that $p_i(t)$ grows exponentially away from ϵ, and therefore that the origin is an unstable equilibrium. ∎

Consequently, the network cannot stay in a state where no connection is active. Any small positive perturbation (due for example to the quantization noise in the computation process)

will push a connection to emit.

Suppose that we have a bound A_{max} on the maximum number of sender/receiver disjoint connections that are activated by a connection in the network, and define the distance between two connections as the distance between their closest end-nodes. Then we can compute parameters l, s and r of the MAP update equation (3) such that two neighboring a.s. active connections are at least at a distance equal to the exclusion range R_E from each other (and so do not collide) but at most at distance $2(R_E + d)$ (remember that d is the communication range), so that some level of spatial reuse is guaranteed. Consider the undirected graph $G' = (V', L')$ where the set V' of vertices corresponds to the network nodes and the set L' includes an edge for each direct connection in the network. We say that a network is connected if the corresponding graph G' is connected. More generally, the following theorem applies to any connected component of the network which covers an area large enough so that the exclusion domain of one of its connection cannot contain all the other connections of the component.

Theorem 4: If the network is connected, the parameter domain

$$l > 1 \quad s > l - 1 \quad 0 < r < \frac{1}{A_{max}}(1 - l + s) \qquad (8)$$

guarantees that, at a stable equilibrium point of system (3), (i) two a.s. active connections are at least at distance R_E and (ii) within a distance $2(R_E + d)$ of an a.s. active connection there is another a.s. active connection.

Proof: (i) By Theorem 2, $\mathbf{p}$ contains only 0 and 1 at the system stable equilibrium. In addition, the last inequality of (8) implies that

$$l - s + A_{max}r < 1,$$

which in turn implies that for any number $n_E \geq 1$ and for any number $0 \leq n_A \leq A_{max}$ of a.s. active connections in the exclusion and activation domains of an a.s. active connection

$$l - n_E s + n_A r < 1,$$

where l, s, $r > 0$. This proves that if there is an a.s. active connection j in the exclusion domain $\mathcal{D}_E^i$ of an a.s. active connection i, then the system cannot be at equilibrium.
(ii) We show by contradiction that within a distance $2(R_E + d)$ of an a.s. active connection there must be another a.s. active connection, otherwise the system is not at a stable equilibrium. Assume, on the contrary, that connection i is a.s. active at time t, that there is no other a.s. active connection at a distance less than or equal to $2(R_E + d)$ from connection i and that the system is at a stable equilibrium. Take the closest connection to connection i, which is not in its exclusion domain $\mathcal{D}_E^i$, and denote this connection by j. Connection j is at most at distance $R_E + d$ from connection i, otherwise connection i and j would belong to two disjoint components of the network and the network would not be connected. As the system is at

stable equilibrium, connection j is either a.s. active ($p_j(t) = 1$) or a.s. idle ($p_j(t) = 0$). The first situation ($p_j(t) = 1$) contradicts our initial assumption since there is an a.s. active connection (connection j) at a distance at most $R_E + d$ from connection i. We show that in the second situation ($p_j(t) = 0$) the system cannot be at a stable equilibrium. Set $p_j(t) = \epsilon$, where $\epsilon > 0$ is arbitrarily small. Since connection j is at most at distance $R_E + d$ from connection i, all connections in $\mathcal{D}_E^j$ are at most at distance $2(R_E + d)$ from connection i and are thus by assumption a.s. idle. Therefore, as long as $p_j(t) \leq 1/l$,

$$p_j(t + 1) \geq f\left(lp_j(t)\right) = lp_j(t).$$

Since $l > 1$, this equation shows that connection j must become a.s. active.
Consequently, our initial assumption is contradicted since there is either an a.s. active connection at a distance less than or equal to $2(R_E + d)$ or the system is not at a stable equilibrium. ∎

In the grid topology $A_{max} = 4$. In more random topologies, A_{max} is not known. However, we find that taking A_{max} equal to the average number of nodes in the activation area of a connection works well in practice as demonstrated by the numerical results in the next section.

VI. NUMERICAL RESULTS AND ANALYSIS

In this section, we provide quantitative results on the spatial reuse achieved by the RDMAC algorithm and validate the fact that the RDMAC algorithm performs well on irregular or mobile topologies.

A. Candidate Algorithms

We compare the performance of the RDMAC algorithm with four other decentralized MAC algorithms. These algorithms were selected because they provide idealized models for well-known MAC protocols or reflect some innovative ideas in the literature.

- *The Hard-core algorithm.* A hard-core point process is a point process in which the constituent points are forbidden to lie closer than a certain minimum distance [13]. In our case, we want active connections to be sufficiently far apart in order to avoid collisions. A simple way to do so is to independently mark each direct connection by a random number uniformly distributed in a fixed interval. A connection is active if its mark is larger than the marks of all the other connections in its exclusion area. The hard-core process models well the HiPERLAN type 1 access scheme [14]. In this scheme a node seeking to access the channel transmits a burst of random length (i.e., its mark), the node with the longest burst gains access to the channel.

- *The Random Pick algorithm.* Connections are picked successively in a random order. A newly selected connection is added to the set of active connections if it does not lie in the exclusion area of an already active connection. The

(a) Node Poisson distributed in a 20x20 square and the corresponding direct connections.

(b) Active set of connections at a given time, as computed by RDMAC algorithm.

Fig. 6. Illustration of the transmission patterns created on a Poisson topology (active transmissions are highlighted by black disks while the gray areas correspond to the exclusion domains). Contrary to the grid topology, the direct connections have now of variable lengths.

Random Pick algorithm can be seen as an idealization of the IEEE 802.11 protocol, where RTS-CTS exchanges would be performed sequentially throughout the entire network, thus avoiding collisions and the need for a backoff mechanism.

- *The SEEDEX algorithm [7].* In the SEEDEX algorithm, nodes can be either in the Possibly Transmit (PT) state or in the Listening (L) state. A node is in the PT state with probability p_s and in the L state with probability $1 - p_s$. A node t in the PT state transmits to a station l in the L state with probability $p = 1/(n + 1)$, where n is the number of neighbors of l which are in the PT state.

- *The SR-Aloha algorithm [6].* In the SR-Aloha scheme all nodes emit with the same probability p. p is chosen to maximize the mean number of connections that can successfully transmit per unit area. [6] derives the optimal value of p for Poisson topologies and exponentially distributed sending powers. The analytical results of [6] cannot be applied directly here since we do not assume exponentially distributed sending powers but a fixed sending power P. Therefore, we compute the optimal value of p experimentally.

The SEEDEX algorithm and the SR-Aloha algorithm assume unidirectional exchanges and thus only attempt to silence the nodes around the receiver. Such a model can potentially schedule more simultaneous transmissions than the bidirectional model assumed in this paper but does not provide a framework for reliable acknowledgments. In order not to penalize these two algorithms, we keep them in their original model; although this might be slightly unfair to the other algorithms.

B. Simulation Setting

Nodes are located in a 20×20 square. We consider two different types of topology, the grid topology and the Poisson topology. In the grid topology 400 nodes are placed on a grid at a distance of 1 of each other. In the Poisson topology the nodes are uniformly distributed on the 20x20 square and the number of nodes in the square is a Poisson process of intensity 1. The average number of nodes is thus the same as in the grid topology. Figure 6(a) shows an example of a Poisson topology. We assume that the communication range $d = \sqrt{4/\pi}$. The radius of the exclusion area is $R_E = \sqrt{4/\pi} = d$ and the radius of the activation area $R_A = \sqrt{8/\pi}$. With these values, each node has on average 4 nodes in its exclusion area and 4 nodes in its activation area for both types of topology.

For completeness we now mention the parameters used by each algorithm. The Reaction-Diffusion algorithm uses the parameter values $l = 1.01$, $s = 1.01$ and $r = 0.25$. These values fall in the parameter domain (8) and thus guarantee no collision at the system equilibrium for the grid topology. Each connection is assigned an initial mark (i.e., $p(0)$) uniformly distributed between $[0, 1/100]$; the same mark is used for the Hard-core algorithm. The update order is a random permutation π of $\{1, \ldots, N\}$; the same order is used for the Random pick algorithm. In the SEEDEX algorithm $p_s = 0.26$ as specified in [7]. Finally, we find by simulation that $p = 0.30$ yields the best performance for the SR-Aloha algorithm.

Each simulation is repeated 200 times. The results presented consist of values averaged over the 200 experiments, and of 95% confidence intervals.
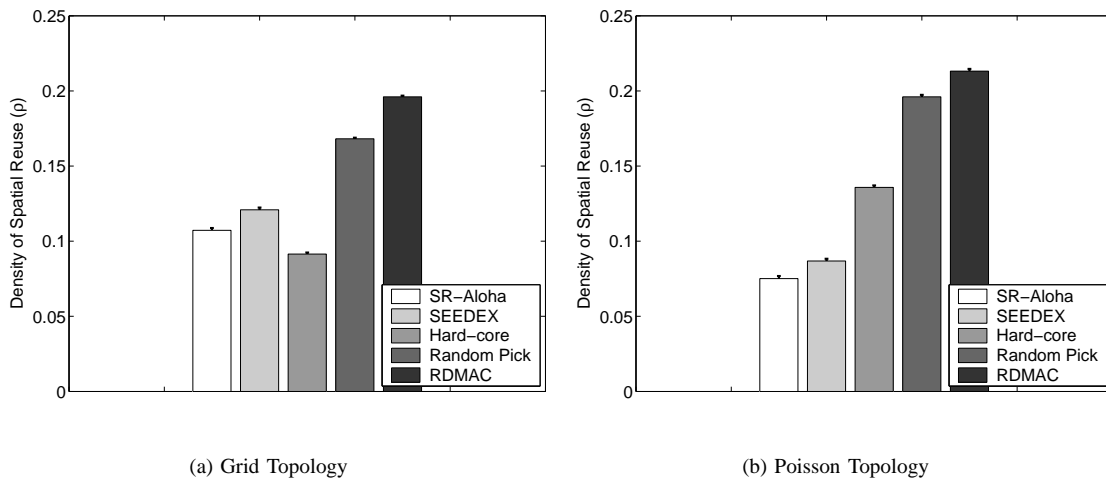
(a) Grid Topology
(b) Poisson Topology

Fig. 7. Density of Spatial reuse ρ of the transmission patterns created by the different algorithms.

C. Spatial Reuse under Scrutiny

How dense are the transmission patterns created by the different algorithms?

1) Definitions: Recall that the *density of spatial reuse* ρ of a transmission pattern is the average number of successful transmissions per unit area in this pattern. Transmissions subject to collisions are not included in the computation of the density of spatial reuse since they cannot be used for data transmission. The density of spatial reuse achieved by a MAC algorithm is proportional to its throughput at high traffic loads.

2) Centralized upper bound on the density of spatial reuse: Due to the regularity of the grid topology it is possible to solve the corresponding instance of the maximum independent set problem (Section V-A) and infer the optimal transmission patterns. The density of spatial reuse of an optimal transmission pattern on the grid topology is $1/4$, this value is thus an upper bound on the density of spatial reuse that can be achieved by any MAC algorithms. Unfortunately, in the case of a Poisson topology, solving the maximum independent set problem is computationally intractable (due to the NP-Complete nature of the problem). Consequently, the maximal density of spatial reuse achievable in a Poisson topology is not known.

3) Numerical results: Figure 7 shows the performance of the four candidate algorithms and the RDMAC algorithm in term of spatial reuse.

The RDMAC algorithm performs the best, it achieves a density of spatial reuse $\rho \simeq 0.20$ in the grid topology and $\rho \simeq 0.21$ in the Poisson topology. It is thus only 20% away from the centralized upper bound on the density of spatial reuse for the grid topology. Figure 6(b) gives a visual illustration of a typical transmission pattern created by the

RDMAC algorithm on a Poisson topology. The transmission pattern is very dense, all connections are either active or inhibited by other active connections. Observe also that the transmission pattern is free of collisions, despite the fact that for Poisson topologies Theorem 4 (Section V) does not always hold, since we replace a deterministic bound on the value A_{max} by an average value. The RDMAC algorithm performs slightly better in the Poisson topology than in the grid topology. The reason lies in the distribution of the connection lengths. In the grid topology all direct connections have length 1 while in the Poisson topology the average connection length is $2d/3 = 4/3\sqrt{\pi} < 1$. Since smaller connection lengths result in smaller exclusion domains, it is indeed reasonable to expect that a MAC algorithm achieves a higher spatial reuse in the Poisson topology.

The Random Pick algorithm performs surprisingly well, with a density of spatial reuse only 7% to 9% lower than the RDMAC algorithm. This result is interesting since it clearly points out that the weakness of the IEEE 802.11 protocol does not lie in its pattern formation ability but most likely in its inability to perform the RTS-CTS exchanges efficiently. Previous works, for example [15], have already shown that the performance of the IEEE 802.11 protocol depends greatly on the parameters of its backoff mechanism. Our results tend to confirm that the backoff mechanism of the IEEE 802.11 protocol is indeed a key factor affecting the performance of the protocol.

The Hard-core, SR-Aloha and SEEDEX algorithms have a significantly lower densities of spatial reuse. The Hard-core algorithm is collision free, but is too conservative in its way of allocating the channel and tends to silence too many connections. This suggests that the HiPERLAN type 1 access scheme does not provide a spatial usage of the channel as good as the IEEE 802.11 protocol. The transmission patterns

created by the SEEDEX and the SR-Aloha algorithms suffer from collisions, which reduces their density of spatial reuse. They are also more sensitive to irregularities in the topology as they are the only ones to perform less well on the Poisson topology than on the grid topology.

D. Spatial Reuse and Mobility

How does mobility affect the spatial reuse of the RDMAC algorithm?

1) Mobility model: We use the well-known random waypoint mobility model to evaluate the performance of the RDMAC algorithm when nodes are mobile. In this model a node randomly chooses a destination point and a speed. It then moves at constant speed toward the selected destination. Once the destination is reached, it pauses for a certain time, picks a new destination and speed and starts moving again. Under the random waypoint model nodes tend to concentrate in the middle of the simulation area, to avoid this effect we use periodic boundary conditions (i.e., a torus abstraction). The initial distribution of the nodes is Poisson and a distance of 1 (see simulation setting) corresponds to 10m. For complexity reasons we reduce the area of simulation from a 20x20 to a 10x10 square. The simulations run for 10s. We consider three scenarios which correspond to different levels of mobility. In the first scenario, nodes move at 2ms^{-1} which is approximately the running speed of a human being. In the second scenario, nodes move at 14ms^{-1}, the driving speed in a city and in the last scenario nodes move at 25ms^{-1} which is about the driving speed on a regular country road. In all scenarios the pause time is zero.

2) Numerical results: Contrary to the other algorithms, RDMAC is a dynamic algorithm where the transmission patterns self-organize over time. Under mobility, the RDMAC algorithm progressively incorporates new connections in the existing transmission pattern. The price to pay for this soft transition between transmission patterns is the convergence time. Figure 8 demonstrates that despite this handicap the RDMAC algorithm can guarantee a high degree of spatial reuse even at high levels of mobility. Indeed, in the 'running speed' scenario, the RDMAC algorithm still achieves a level of spatial reuse higher than all the other algorithms. Moreover, it performs very close to the random pick algorithm in the two other mobility scenarios.

VII. CONCLUSION

Reaction-Diffusion Medium Access Control is a decentralized scheme that is remarkable by its ability to build global transmission patterns using only local interactions between network nodes. This first property of the RDMAC scheme is especially useful in ad hoc networks where decisions based only on local knowledge make it easier to adapt to failures or to mobility. Moreover, the RDMAC scheme creates


Fig. 8. Spatial reuse achieved by the RDMAC algorithm under three mobility scenarios. In the 'running speed' scenario nodes move at a speed of 2ms^{-1}, in the 'city car' scenario nodes move at 14ms^{-1} and finally in the 'country car' scenario nodes move at 25ms^{-1}.

collision-free transmission patterns with a high density of spatial reuse. This second property of the RDMAC scheme allows for an efficient scheduling of the transmissions in multi-hop wireless networks with high traffic loads where current random-access schemes often show poor performance.

To conclude, we propose some possible ways to exploit these key properties of the RDMAC algorithm in future implementations. First, the RDMAC algorithm could be used to precompute suboptimum TDMA schedules off line at low complexity. This approach is especially convenient for large networks where the computation of optimum TDMA schedules is impossible. Feeding different initial conditions to the RDMAC algorithm triggers the creation of different transmission patterns that can then be arranged into a TDMA schedule. Simulation results on the grid topology show that the RDMAC algorithm can indeed achieve almost 80% of the density of spatial reuse of an optimum TDMA transmission pattern in an average of 35 iterations of equation (3). A second idea is to use the RDMAC algorithm to create on-the-fly transmission patterns. Periodically, each connection reset its MAP, for example to the level of occupancy of its queue. As a result, a new transmission pattern emerges, which favors nodes with a high number of packets to send. The transmission pattern then remains active until the next periodic reset.

The reaction-diffusion mechanism was initially proposed to explain and reproduce the formation of patterns in biological systems. In the present work, we have applied the reaction-diffusion mechanism to the problem of medium access control in multi-hop wireless networks. We believe that the self-organization properties of the reaction diffusion mechanism will prove to be useful for other applications in the dynamic setting of ad hoc networking.

REFERENCES

[1] A. Turing, "The chemical basis of morphogenesis," in *Philosophical Transactions of the Royal Society B (London), 237,37-72*, 1952.

[2] *IEEE802.11, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std., Aug 1999.

[3] P. Gupta, R. Gray, and P. Kumar, "An experimental scaling law for ad hoc networks," Univ. of Illinois at Urbana Champaign, May 2001. [Online]. Available: citeseer.nj.nec.com/gupta01experimental.html

[4] B. Radunovic and J.-Y. L. Boudec, "Joint scheduling, power control and routing in symmetric, one-dimensional, multi-hop wireless networks," in *WiOpt Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, March 2003.

[5] R. Nelson and L. Kleinrock, "Spatial-TDMA: A collision-free multihop channel access control," in *IEEE Transactions on Communications*, vol. 33, 1985, pp. 934–944.

[6] F. Baccelli, B. Blaszczyszyn, and P. Muhlethaler, "A Spatial Reuse Aloha MAC Protocol for Multihop Wireless Mobile Networks," in *Allerton conference on communication, control, and computing*, 2003.

[7] R. Rozovsky and P. Kumar, "SEEDEX: A MAC protocol for ad hoc networks," in *Proceedings of MobiHoc*, 2001. [Online]. Available: citeseer.nj.nec.com/rozovsky01seedex.html

[8] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, "Achieving MAC layer fairness in wireless packet networks," in *Mobile Computing and Networking*, 2000, pp. 87–98. [Online]. Available: citeseer.nj.nec.com/nandagopal00achieving.html

[9] J. Murray, *Mathematical Biology. II: Spatial Models and Biomedical Applications (3rd ed)*. Springler-Verlag Berlin Heidelberg, 2003.

[10] P. Thiran, "Dynamics and self-organization of locally-coupled neural networks," in *PPUR, collection META*, 1997.

[11] D. Garey, M.R.and Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: Freeman, 1979.

[12] M. Vidyasagar, *Nonlinear Systems Analysis (2nd ed.)*. Prentice Hall, 1993.

[13] D. Stoyan, W. Kendall, and J. Mecke, *Stochastic geometry and its applications (2nd ed.)*. John Wiley & Sons Ltd, 1995, pp. 162–166.

[14] *HIgh PErformance Radio Local Area Network (HIPERLAN) Type 1; Functional specification, ETS 300 652*, Std., October 1996.

[15] M. C. F. Cali and E. Gregori, "Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit," in *IEEE/ACM Trans. on Networking, 8(6):785 799*, December 2000.