Rate-optimal decentralized broadcasting: the wireless case

Don Towsley
Dept. Computer Science
University of Massachusetts
Amherst, MA 01003, USA
towsley@cs.umass.edu

Andrew Twigg
Computer Laboratory
University of Cambridge
Cambridge, UK
andrew.twigg@cl.cam.ac.uk

***Abstract*—We consider the problem of broadcasting information from a source node to all nodes in a wireless network, using a local-control algorithm. For simplicity, we assume some oracle that provides the scheduling decisions required for interference. Under this assumption, we can show that a local control algorithm exists that achieves a stable system whenever the injection rate at the source is feasible for the network. We show results for the case of a single antenna and for multiple antennas operating independently.**

I. Introduction

Broadcasting poses an important challenge in military wireless networks, primarily due to the dynamic nature of the network. Moreover, efficient use of resources is extremely important in military networks as they are often severely energy constrained. Thus designing decentralized broadcasting algorithms that can minimize resource usage or, conversely maximize capacity given resource constraints is an important problem. In this paper we present and analyze a decentralized broadcast algorithm for a wireless network that achieves the *broadcast capacity* of the network. We prove its optimality by extending results in [6] developed for broadcast in a wireline network.

The broadcast algorithm is simple. Nodes exchange information with their neighbors regarding the identities of all data packets that they have received. Then, when granted access to the channel, a node randomly choose packets that are required by at least one of its neighbors. Last, the proof that this algorithm achives the broadcast capacity of the network, i.e., the maximum possible sending rate that can be supported by any algorithm, relies on studying the fluid limit of the algorithm and proving that the fluid model is stable by choosing an appropriate Lyapunov function.

The paper is organised as follows. In Section 2 we describe some of the related work in wireless broadcasting. In Section 3 we describe our setting, based on the so-called ‘scheduling-free model’ and single omni-directional antennas, and introduce the deterministic fluid approximations. In Section 4 we prove the main stability result. In Section 5 we consider extensions of the model to different antenna configurations.

II. Related work

The problem of performing multicast and broadcast in a wireless network has received increasing attention over the last several years. Closest to our work is that of Keshavarz-Haddad et al. [9], which considered the problem of broadcasting in a wireless network. They proved asymptotically tight bounds for the broadcast rate using a construction of minimum connected dominating sets (MCDS) that, although can be constructed in a distributed fashion, are not easy to maintain.

Less related are works that study the multicast and/or broadcast capacities of random spatial networks. Zheng studied the broadcast capacity of static wireless networks [19], where the nodes are distributed according to a 2-dimensional Poisson process sharing a Gaussian wireless channel whose capacity is given by the Shannon Signal to Interference Ratio (SINR). Jacquet and Rodolaski, [7], and Shakkottai and Srikant, [14] studied the multicast capacity of a dense wireless network as the number of nodes grows.

Last, considerable work on the development of specific algorithms for broadcast and multicast in wireless networks, [17], [13], [11], [16], [8]. However, these algorithms are heuristic in nature and make no claims about maximizing the broadcast capacity of the network.

III. Scheduling-free model

We represent a wireless network by a directed, edge-capacitated hypergraph $G = (N, E, \{c_v\}_{v \in V})$. Here $(u, V) \in E$ means that when $u \in N$ transmits and no other nodes transmit, the set $V \subset N$ receives the transmission at capacity c_u. We assume that each node has a single omni-directional antenna. Thus every node u has exactly one outgoing hyperedge. Thus V in the hyperedge (u, V) is the neighbour set of u and we will refer to it as $N(u)$. Last, we will let (u, v) denote an edge between u and v iff $v \in N(u)$.

We now consider the following broadcast problem. A source node s wants to send an infinite stream of data to all other nodes. Due to the local broadcast nature of wireless networks, a data transfer means replicating a given packet from some node u to *all* of its neighbours in G. In particular, different packets cannot be sent to two nodes v, v' with $v, v' \in V$ along edge (u, V) concurrently. The *injection rate* at the source is denoted by λ, and is the rate at which the source gets new packets.

Each node can be in one of two states: on or off. In the on state, it can transmit data. A node v can receive data from a neighbour u iff u is on and no other neighbour of v is on. Let $B(t) = (B_1(t), \ldots, B_n(t))$ be the configuration of the network where $B_u(t)$ =1 iff node u is on at the t-th time step. Given a configuration, we say that an edge $(u, v) \in E$ is active iff u is on and v can receive data from u. We call a configuration *interference-free* iff all edges $(u, v) \in E$ are active whenever u is on. This notion of interference between nodes complicates the scheduling decisions involved in the model. To simplify things, we first consider the following 'scheduling-free' model: we model $B(t)$ as a continuous-time Markov chain with the following properties: (1) at every time t, the configuration specified by $B(t)$ is interference-free and B is ergodic. We assume that such a chain is given by some scheduler that is not under our control. Thus, we must now only contend with the constraint that the same packet must be transmitted over all edges (u, v) for each u.

Given the transition rates of $B(t)$, we shall consider the steady state time $\pi(u)$ that u is active. Since B is ergodic the quantities $\pi(u)$ exist. We define the 'effective' capacity $c'_u = \pi(u)c_u$ of a node operating under this schedule. In what follows, we shall deal with capacities c_u, and leave the substitution of c_u by the effective capacity c'_u to the reader whenever these quantities exist. We do not claim right now that these capacities are indeed the actual capacities, but it will become clear in the construction of the fluid limits that these are a suitable approximation in the fluid scaling.

Following the development of [6], let X_S count the number of packets replicated exactly at nodes of $S \subseteq V$. In addition, let $A = \{G_1 = (W_1, W_1), \ldots, G_m = (W_m, F_m)\}$ is a set of subgraphs that describe the *active packets*: W_i is the set of nodes at which the i-th active packet is currently replicated; $F_i \subseteq W_i$ is the subset of nodes actively transmitting the i-th active packet to its neighbors.

A. Description of state space

We shall assume that at any given time, at most one packet is transferred along a given hyperedge. Thus, the total number of active packets is at most $|E|$. We shall further assume that the following constraints are met: for an active packet with description (W, F), for each $u \in F$, then $u \in W$, $N(u) \cap W \neq N(u)$ (there is at least one neighbour of u not in W), and there is no other $v \in W$ such that $N(u) \cap N(v) \neq \emptyset$.

We use the notation

$$X_{+u} = \sum_{S \in \mathcal{S}: u \in S, N(u) \cap \overline{S} \neq \emptyset} X_S$$

to count the number of packets replicated at u but not in at least one of its neighbours. In addition, we use the notation

$$X^a_{+u} = \sum_{(W,F) \in A} 1_{u \in W} 1_{N(u) \setminus W \neq \emptyset} 1_{\cup_{u' \in W \setminus F} N(u') \cap (N(u) \setminus W) \neq \emptyset}$$

to denote the number of active packets that could possibly be forwarded by u to one or more of its neighbors given the above constraints.

We enforce the following activity condition at all times: for any node u in the on state ($B_u(t) = 1$), there is active packet being transferred from u or else $X_{+u} = 0$ and $X^a_{+u} = 0$.

The system evolution is determined by the following transition mechanism.

B. Packet transmission algorithm

The algorithm is modelled by a markov chain having several types of transition.

a) Primary transitions: The first type of primary transition is due to a fresh packet arrival at the source. After such a transition, the state variable $X_{\{s\}}$ is updated to $X_{\{s\}} + 1$.

The second type of primary transition is due to completion of transfer of an active packet along some hyperedge. Let this packet be represented by (W, F), and let $u \in F$ be the node completing the transfer to its neighbours. If u is the only member of F then the packet is removed from the collection of active packets and the number of idle packets at $S = W \cup \{v\}$ is increased by 1. If $F \setminus \{u\}$ is not empty then W is updated to $W \cup N(u)$, and F is updated to $F \setminus \{u\}$.

It should be remembered that, since we are operating in the scheduling-free scenario, the packet is delivered to all neighbours $N(u)$ rather than some strict subset. The case where there is interaction between the set of 'on' nodes and the set of nodes that receive a given transmission from u is an open problem.

We wish to enforce the following activity condition on nodes which are on: for every node u, it is either transmitting a packet to its neighbors or in an off state.

b) Secondary transitions: These occur subsequently to primary transitions, to ensure that the activity condition is met. If, after a primary transition, there is a node u for which the activity condition is not met, this means that u is on and has $X_{+u} + X^a_{+u} > 0$ packets, one of which can be transferred from u to one or more of its neighbors. In this case, one of these packets will be selected uniformly at random, and start being transmitted to its neighbors.

More precisely, for each $S \in \mathcal{S}$ such that $u \in S, N(u) \cap \overline{S} \neq \emptyset$, with probability

$$\frac{X_S}{X_{+u} + X^a_{+u}},$$

the following state updates are made:

$$X_S \leftarrow X_S - 1,$$
$$A \leftarrow A \cup (S, (u, N(u))).$$

For each active packet (W, F) such that $u \in W$, with probability $1/(X_{+u} + X^a_{+u})$, the active set A is updated as follows:

$$A \leftarrow A \setminus (W, F) \cup (W, F \cup (u, N(u))).$$

Note that all these transition probabilities sum to 1, as required.

This secondary transition mechanism corresponds to what we shall call the "random useful" packet forwarding strategy: when a new useful packet transfer along from u can start, the packet that is actually transferred is selected uniformly at random from the total collection of packets present at u and not at some neighbour $v \in N(u)$, and not currently transferred towards any neighbour.

c) Scheduling transitions.: These occur whenever a node transits either from an on state to an off state, or an off state to an on state.

Consider the first case. Suppose that B_u changes from 1 to 0. We shall assume that the transmission is interrupted, and must be resent when the node is active again. Therefore, if there exists a $(W, F) \in A$ such that $u \in F$, then this transmission is stopped and we update

$$(W, F) \to (W, F \setminus \{u\})$$

(and if F is now empty, this packet is removed from the active set and we update the count of idle packets at W by $X_W \to X_W + 1$). On the other hand, if B_u changes from zero to one, i.e., u is allowed to transmit, then it executes the action as it would for a secondary transition.

We shall consider the Markovian case of transmission times: the interpacket arrival times at the source are exponentially distributed with mean λ^{-1}, and packet transfer times for (u, V) are exponential with mean c_u^{-1}. The memoryless property of the Poisson distribution will be important when considering the transmission interruptions arising from the scheduling transitions.

IV. Scheduling-free case: main result

In this section, we shall prove our main result. Theorem 1 says that, under a restriction on the injection rate of packets at the source, the Markov process of the previous section describes a stable system. Theorem 2 says that, under the scheduling-free assumption, this restriction is necessary, hence our result is optimal.

Our main result is the following:

Theorem 1: The Markov process $((X_S)_{S\in\mathcal{S}}, A)$ corresponding to random useful packet forwarding in the scheduling-free case with a single omni-directional antenna is ergodic under the condition

$$\lambda < \min_{S\in\mathcal{S}} \sum_{u\in S, N(u)\cap\overline{S}\neq\emptyset} c_u. \quad (1)$$

The theorem implies that the 'backlogs' at any node converge in probability, and so we say that the system is stable under the above restriction on λ. We will now show that this restriction is necessary, i.e. that the system is stable whenever the injection rate is strictly feasible, i.e. it is possible to broadcast at a rate of $\lambda + \epsilon$ for any $\epsilon > 0$. Thus, using our scheme we can broadcast at a rate arbitrarily close to the network capacity (remembering of course that we are operating in the scheduling-free scenario).

In the wired case we have independent edge transmissions, and the maximum broadcast rate is characterized in terms of packings of spanning arborescences. Edmonds [3] has shown that the maximum size of such a packing equals the *min-mincut*, denoted by

$$\mu = \min_{S\in\mathcal{S}} \sum_{u\in S, v\notin S} c_{uv}.$$

In our model, each node has to transmit the same packet over all its outgoing links. This corresponds to a packing of spanning arborescences where for every arborescence T in the packing, every internal node u of T has degree in T equal to its degree in G, i.e. $\deg_T(u) = \deg_G(u)$. We call such an arborescence a *restricted* arborescence.

In light of this interpretation, Theorem 1 says that there exists a packing of restricted arborescences of size $\min_{S\in\mathcal{S}} \sum_{u\in S, N(u)\cap\overline{S}\neq\emptyset} c_u$, i.e. λ is at least this value. To see this, take λ arbitrarily close to $\min_{S\in\mathcal{S}} \sum_{u\in S, N(u)\cap\overline{S}\neq\emptyset} c_u$ and note that each packet travels on exactly one restricted arborescence and consider the packing induced by the stationary distribution of arborescences chosen by the broadcasting scheme. The following argument shows that this bound is tight.

Theorem 2: For every edge-disjoint packing of restricted arborescences with value λ, we have

$$\lambda \leq \min_{S\subset V} \sum_{u\in S, N(u)\cap\overline{S}\neq\emptyset} c_u.$$

Proof: For a contradiction, assume that $\lambda > \sum_{u\in S, N(u)\cap\overline{S}\neq\emptyset} c_u$ for some set S. Then there necessarily exists some (possibly parallel) edges in some tree that do not exist in G, since each vertex u can contribute at most value c_u to the value of the packing. This is because either all its child edges are in a given tree, or none of them are.) ■

Together, this shows that any broadcasting scheme having the fluid limit defined in the next section is stable under injection rate λ, whenever rate $\lambda + \epsilon$ is feasible for $\epsilon > 0$, i.e. there exists a packing of restricted trees of value $(\lambda + \epsilon)$. A packing of restricted trees of value x corresponds to a distribution scheme having rate x under the "local broadcast" model.

It is interesting to ask how large can the gap be between the rates under local broadcast (say λ) and independent edge transmissions (say μ)? It is easy to see that $\mu/\lambda \leq \max_v \deg_G(v)$, and this is tight in the worst-case: consider a source connected to n nodes, which form a clique K_n where all capacities are unit. We can achieve a packing of n unrestricted trees, but any packing of restricted trees must have value at most 1 (for otherwise, a cycle in some tree would result). Note that this graph is a unit-disk graph, so trying to restrict oneself to such more realistic wireless topologies doesn't strengthen the result in the worst case.

A. Fluid limits

Definition 1: The real-valued non-negative functions $t \to y_S(t)$, $S \in \mathcal{S}$, are called fluid trajectories of the above Markov process if they satisfy the following conditions.

For all $S \in \mathcal{S}$, all $u \in S$, all $V \subseteq \overline{S} \cap N(u)$, there exist non-negative functions $t \to \phi_{S,(u,V)}(t)$ such that

$$\begin{aligned} y_{\{s\}}(t) &= y_s(0) + \lambda t - \sum_{V\subseteq N(u)} \phi_{\{s\},(s,V)}(t) \\ y_S(t) &= y_S(0) + \sum_{u\in S} \sum_{V\subseteq N(u)\cap S} \phi_{S\setminus V,(u,V)}(t) \quad (2) \end{aligned}$$

$$-\sum_{u\in S}\sum_{V\subseteq N(u)} \phi_{S,(u,V)}(t),$$
$$S \neq \{s\} \qquad (3)$$

and that are non-decreasing c_u-Lipschitz continuous, differentiable almost-everywhere.

For y_S, the additive terms count the total amount of fluid flowing into S due to transmissions from $u \in S$ along some hyperedges containing neighbours not currently in S.

Since each node has a single omni-directional antenna, we can write the following expression:

$$y_{+u}(t) > 0 \Rightarrow \frac{d}{dt}\phi_{S,(u,N(u)\cap\overline{S})}(t) = c_u \frac{y_S(t)}{y_{+u}(t)}, \qquad (4)$$

where we have used the notation

$$y_{+u}(t) := \sum_{S: u\in S, N(u)\cap\overline{S}\neq\emptyset} y_S. \qquad (5)$$

Also, define

$$\frac{d}{dt}\phi_{S,(u,V)} = 0 \quad \text{for} \quad V \neq N(u)\cap\overline{S}. \qquad (6)$$

Together, these conditions enforce that fluid can flow only from $u \in S$ to all of its neighbours not in S, rather than some selective subset of neighbours (as in the independent edge transmission case).

We now argue that the fluid limits above do indeed describe the Markov process after some suitable rescaling. The proof is along the same lines as given in [6], except that we must consider the completion of packet transfers along hyperedges, which are treated just as simple edges, and we must consider the effect of the scheduling transmissions. A full proof will be given in the full version of the paper.

Proof outline. As in [6], given a sequence $(z_N)_N$ of positive reals and initial conditions $(X^N(0), A^N(0))_N$ such that the limit $\lim_{N\to\infty}\frac{1}{z_N}X^N(0) = x(0)$ exists, let us define the rescaled Markov process

$$Y_S^N(t) z_N X_S^N(z_N t)$$

for all $S \in \mathcal{S}$. We wish to show that for all $T > 0$, the process Y^N restricted to the interval $[0, T]$ converges in probability to the fluid trajectories defined above.

To this end, consider the total number of packets present in S, given by $\tilde{X}_S = X_S + \sum_{(W,F)\in A} \mathbf{1}_{W=A}$ and similarly define the rescaled quantity

$$\tilde{Y}_S^N(t) z_N \tilde{X}_S^N(z_N t) = Y_S^N(t) + 11 z_N \sum_{W\in A^N} \mathbf{1}_{W=S}.$$

We now consider the process counting the number of packets sent over a given hyperedge $(u, N(u))$. This is an interrupted Poisson process (IPP), a special case of the 2-state Markov-modulated Poisson process where one of the states has zero intensity. Assume a unit-rate Poisson process P_u modulated by a stationary continuous time Markov process $B(t)$ with jump times for each node u independent and exponentially-distributed intervals of mean 'on' length $\frac{1}{\pi_u}$ and 'off' length $1 - \frac{1}{\pi_u}$ where π_u is the stationary probability of u in $B(t)$. Katsinis et al. [5] show that this modulated process has mean intensity π_u and Alwakeel [1] shows that in the limit it behaves as a scaled Poisson process with intensity π_u (intuitively, the memoryless property guarantees that on a jump from on to off, the expected time until the next arrival of P_u is the same when P_u is restarted following the next jump off to on). Roughly, since packet transmissions are interrupted and must be restarted (rather than continuing on the next jump), and since a transmission begins immediately after the on to off jump, the count of the number of arrivals of the process P_u in the on period will, in expectation, count the number of completed transmissions during this interval.

Following [6] and scaling P_u in time we can write

$$\Phi_{S,(u,N(u))}^N(t) = P_u\left(\pi_u c_u \int_t^t \sum_{(W,F)\in A^N(s)\mathbf{1}_{W=S}\mathbf{1}_{(u\in F}} \mathrm{d}s\right)$$

to approximate the total number of completions of transfers of packets previously replicated at S, along hyperedge $(u, N(u))$. The remainder follows [6], noting that we approximate the IPP by a rescaled Poisson process and thus the probability of a large deviation $(1+\delta)$ remains exponentially small in δ.

The functions $\phi_{S,(u,N(u)}$ are handled in a similar way: it is not difficult to see that they are c_u-Lipschitz, and indeed the probability that a packet selected for transmission along $(u, N(u))$ is an idle packet replicated at S is, in the limit, $y_S(t)/y_{+u}(t) + O(h)$. This ensures the claimed derivatives of the functions. It remains to understand how the scheduling transitions affect things. Since packet transfer times are exponential and memoryless, and the schedule $S(t)$ is ergodic, we have that the 'on' periods for u are exponential with mean $1/\pi(u)$. ∎

B. *Fluid dynamics: stability*

We now establish stability of the fluid trajectories under the condition on the injection rate mentioned earlier. Define the linear Lyapunov function

$$L(y(t)) = \sup_{S\subset V} \beta_{|S|} y_{\subseteq S}.$$

We want to show that there exist positive constants $\beta_1, \ldots, \beta_{n-1}$ so that L is continuous and strictly decreasing. As in [6], the ergodicity of the process will follow by applying a version of Foster's theorem, which we shall omit here for clarity (the reader can refer to the proof in [6] for technical details).

We would like to have, for S^* achieving the supremum, the following property, for any fixed $\alpha > 0$: for $u \in S^*, N(u) \cap \overline{S^*} \neq \emptyset$, then for all $S' \not\subseteq S^*$ with $u \in S', N(u) \cap \overline{S}' \neq \emptyset$, we have $y_{S'} \leq \alpha y_{+u}$. Then we would have the following evaluation:

$$
\begin{aligned}
&\frac{d}{dt} y_{\subseteq S^*} \\
&= \sum_{S \subseteq S^*} \frac{d}{dt} y_S \\
&= \lambda - \sum_{u \in S^*, N(u) \cap \overline{S^*} \neq \emptyset} \sum_{S' \subseteq S^*, u \in S'} \sum_{V \subseteq N(u) \cap \overline{S^*}} \frac{d}{dt} \phi_{S',(u,V)}
\end{aligned}
$$

Using (3-5), this is equal to

$$
\begin{aligned}
& \lambda - \sum_{u \in S^*, N(u) \cap \overline{S^*} \neq \emptyset} \sum_{S' \subseteq S^*, u \in S'} \frac{d}{dt} \phi_{S',(u,N(u) \cap \overline{S^*})} \\
= \; & \lambda - \sum_{u \in S^*, N(u) \cap \overline{S^*} \neq \emptyset} c_u \left(1 - \sum_{\substack{S' \not\subseteq S^*, u \in S' \\ N(u) \cap \overline{S'} \neq \emptyset}} y'_S / y_{+u} \right) \\
\leq \; & \lambda - \sum_{u \in S^*, N(u) \cap \overline{S^*} \neq \emptyset} c_u + n 2^n \alpha \max_u c_u
\end{aligned}
$$

where we used the following equality, for any vertex u and set of vertices S with $N(u) \cap \overline{S} \neq \emptyset$:

$$
\sum_{\substack{S' \subseteq S, u \in S' \\ N(u) \cap \overline{S} \neq \emptyset}} y_{S'} / y_{+u} + \sum_{\substack{S' \not\subseteq S, u \in S' \\ N(u) \cap \overline{S'} \neq \emptyset}} y_{S'} / y_{+u} = 1.
$$

Proposition 1: For any fixed $\alpha > 0$, there exists constants $\beta_1, \ldots, \beta_{n-1}$ such that $S^* = \arg\sup_{S \subset V} \beta_{|S|} y_{\subseteq S}$ satisfies the property described above.

Proof. We want to show that for all $S \subset V, u \in S, N(u) \cap \overline{S} \neq \emptyset$,

$$
y_{+u} < \epsilon_{|S|} y_{\subseteq S} \Rightarrow \beta_{|S|-1} y_{\subseteq (S-u)} > \beta_{|S|} y_{\subseteq S}.
$$

So assume $\exists S \subset V$ with $u \in S, N(u) \cap \overline{S} \neq \emptyset$ with $y_{+u} < \epsilon_{|S|} y_{\subseteq S}$. Then write

$$
\begin{aligned}
y_{\subseteq S} & \leq y_{\subseteq (S-u)} + \sum_{S' : u \in S', S' \subseteq S, N(u) \cap \overline{S'} \neq \emptyset} y_{S'} \\
& \leq y_{\subseteq (S-u)} + y_{+u} \\
& < y_{\subseteq (S-u)} + \epsilon_{|S|} y_{\subseteq S}.
\end{aligned}
$$

Hence $y_{\subseteq S}(1 - \epsilon_{|S|}) < y_{\subseteq (S-u)}$ and the claim will be satisfied if $(1 - \epsilon_i)\beta_{i-1} \geq \beta_i$ for $i = 2, \ldots, n-1$.

Now assume we have an $S \subset V$ with $u \in S, N(u) \cap \overline{S} \neq \emptyset$ and $\epsilon_{|S|} y_{\subseteq|} \leq y_{+u}$. We want to show that if $\exists S' \not\subseteq S$ with $u \in S', N(u) \cap \overline{S'} \neq \emptyset$ and $y_{S'} > \alpha y_{+u}$, then $\beta_{|S \cup S'|} y_{\subseteq (S \cup S')} > \beta_{|S|} y_{\subseteq S}$. To show this, write

$$
\begin{aligned}
\beta_{|S \cup S'|} y_{\subseteq (S \cup S')} & \geq \beta_{|S \cup S'|} (y_{S'} + y_{\subseteq S}) \\
& > y_{\subseteq S} (\alpha \epsilon_{|S|} + 1).
\end{aligned}
$$

The claim will be satisfied if $\beta_{|S \cup S'|}(1 + \alpha \epsilon_{|S|}) \geq \beta_{|S|}$, which is satisfied by $\beta_{n-1}(1 + \alpha \epsilon_i) \geq \beta_i$ for $i = 1 \ldots n-2$.

Hence we need to find β_i's that satisfy these two inequalities, and these are exactly the same inequalities considered in Massoulié et al. [6]. Using their choice of ϵ_i and β_i will establish the proposition. ∎

V. Different antenna models

Our results can extend to other antenna models. In particular, we have the following

A. Multiple omnidirectional antennas

Assume that each node u has k antennas, each with capacity c_u. It is not difficult to see that this is equivalent to the single antenna case with capacities scaled by a factor of k.

B. Fixed directional antennas

In this case, we allow transmission to selected subsets of neighbours. For example, in a planar embedding of the network (say), we would have some θ and draw a cone of width θ and fixed radius, that defines the set of neighbours reachable by a given antenna pointing in the central direction of the cone. Since we do not consider graphs necessarily embeddable into surfaces of bounded genus, we shall simply specify for each node u a collection of hyperedges $E(u) = (u, \{V_i\})$ that u may selectively transmit to. If it has k independent antennas, it may select at most k of these edges (they may overlap).

We shall assume (for technical reasons explained below) that antennas are fixed wrt a particular hyperedge. In reality, this means that transmission is fixed between u and a collection of different neighbours (so the antenna may physically move but transmission is always to the same set of receivers). Hence we have a collection of k possibly overlapping hyperedges $(u, V_1(u)), \ldots, (u, V_k(u))$. When u is clear from the context, we shall simply write V_i instead of $V_i(u)$. We shall allow different antennas to have different capacities (for example, because of radius of transmission, different noise and signal characteristics, etc.). This is specified by assigning the hyperedge capacities c_{u,V_i}. Furthermore, we assume that all antennas at u can operate independently.

For a node u and $V \subseteq N(u)$ define

$$
y_{+u-V} = \sum_{S : u \in S, V \cap \overline{S} \neq \emptyset} y_S
$$

and for antenna i,

$$
y_{+u-V_i}(t) > 0 \Rightarrow \frac{d}{dt} \phi_{S,(u,V_i \cap \overline{S})}(t) = c_{u,V_i} \frac{y_S}{y_{+u-V_i}}.
$$

Also define

$$
\frac{d}{dt} \phi_{S,(u,V)}(t) = 0 \quad \text{for} \quad V \notin \{V_1, \ldots, V_k\}.
$$

To characterize the maximum rate in terms of arborescence packings, consider the collection of restricted arborescences where in each tree T, each node u is either a leaf or appears with child set the union of some $V_i(u)$'s. As usual, we require that every pair of nodes (u, v) appears in the graph with cardinality at most $k|\{i : v \in V_i(u)\}|$.

We consider the following natural extension of the random forwarding scheme: each antenna (u, V_i) transmits at maximum capacity a packet chosen at random from the set of packets at u but not at some node in V_i. The notion of "scheduling-free" remains the same as before. We can prove the following result (the proof will appear in a full version).

Theorem 3: The process described above is stable (under the scheduling-free assumption) whenever the injection rate satisfies

$$\lambda < \min_{S \in \mathcal{S}} \sum_{u \in S} \sum_{i: V_i(u) \cap \overline{S} \neq \emptyset} c_{u,V_i}. \qquad (7)$$

VI. Summary

We considered the problem of broadcasting a stream of data in a wireless network, using local decisions only at each node. We proved that, under some fairly restrictive scheduling assumptions, a rate-optimal protocol exists based on the ideas of Massoulié et al. [6]. A natural next step would be to try to remove the so-called scheduling-free assumption; this appears to be a difficult open problem because of the interaction it would introduce between scheduling decisions and packet transmissions that only reach a strict subset of neighbours.

VII. Acknowledgments

This work was sponsored by US Army Research laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] M. Alwakeel, Equivalent Poisson Process for Interrupted Poisson Process with on-off Periods, *J. Sci. Med. Eng.*, Vol 19, No. 1, 103–112, 2007.

[2] J.G. Dai, On positive Harris recurrence of multiclass queuing networks: a unified approach via fluid limit models, *Ann. of Applied Probability* 5, 49–77, 1995.

[3] J. Edmonds, Edge-disjoint branchings, *Combinatorial Algorithms*, 21–31, 1972.

[4] B. Korte and J. Vygen, *Combinatorial Optimization Theory and Algorithms*, Springer, 2005.

[5] C. Katsinis, A. Volz, A network traffic shaping technique based on waiting time, *Int. J. of Computers and Applications*, Vol 21, 44–49, 1996.

[6] L. Massoulié, A. Twigg *Rate-optimal decentralized broadcast algorithms*. In Journal of Performance Evaluation, 2008.

[7] P. Jacquet, G. Rodolakis, "Multicast scaling properties in massively dense ad hoc networks," *Proc. ICPADS05*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 93–99.

[8] S.-J. Lee, M. Gerla, C.-C. Chiang, "On-Demand Multicast Routing Protocol," *Proc. of WCNC*, September 1999.

[9] A. Keshavarz-Haddad, V. Ribeiro and R. Riedi, "Broadcast capacity in multihop wireless networks," *Proc. MobiCom 2006*.

[10] P. Robert, *Stochastic Networks and Queues*, Springer, 2003.

[11] E. M. Royer, C. E. Perkins, "Multicast Operation of the Ad Hoc On-Demand Distance Vector Routing Protocol," *Mobicom 1999*, Seattle, WA, Aug. 1999.

[12] A.N. Rybko and A.L. Stolyar, Ergodicity of stochastic processes describing the operations of open queueing networks. *Problemy Pederachi Informatsii,* **28**, 2–26, 1992.

[13] P. Chaporkar, S. Sarkar. "Wireless Multicast: Theory and Approaches," *IEEE Trans. on Information Theory,* **51**(6), pp. 1954–1972, June 2005.

[14] S. Shakkottai, X. Liu, R. Srikant. "The multicast capacity of ad hoc networks," *Proc. Mobihoc 2007.*

[15] A. Shrijver, *Advanced graph theory and combinatorial optimization*, Lecture Notes available at http://homepages.cwi.nl/∼lex/, 2001.

[16] P. Sinha, R. Sivakumar, V. Bharghavan. "MCEDAR: Multicast Core-Extraction Distributed Ad Hoc Routing," *Proc. of WCNC 1999*, New Orleans, LA, September 1999.

[17] J. E. Wieselthier, G. D. Nguyen, A. Ephremides. "On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks," *Proceedings of IEEE INFOCOM 2000,* Tel-Aviv, Israel, March 2000.

[18] H.Q. Ye, J.H. Ou and X.M. Yuan, Stability of Data Networks: Stationary and Bursty Models. *Operations Research,* **53**(1), 107-125, 2005.

[19] R. Zheng, "Information dissemination in power-constrained wireless network," *Proc. INFOCOM 2006.*