Optimal Gossiping in Directed Geometric Radio Networks in Presence of Dynamical Faults*

Andrea E.F. Clementi[1], Angelo Monti[2], Francesco Pasquale[1]*, and Riccardo Silvestri[2]

[1] Dipartimento di Matematica, Università di Roma "Tor Vergata"
{clementi,pasquale}@mat.uniroma2.it
[2] Dipartimento di Informatica, Università di Roma "La Sapienza"
{monti,silvestri}@di.uniroma1.it

Abstract. We study deterministic fault-tolerant gossiping protocols in *directed Geometric Radio Networks* (in short, directed GRN). Unpredictable node and link faults may happen during every time slot of the protocol's execution.
We first consider the *single-message* model where every node can send at most one message per time slot. We provide a protocol that, in any directed GRN G of n nodes, completes gossiping in $O(n\Delta)$ time (where Δ is the maximal in-degree of G) and has message complexity $O(n^2)$. Both bounds are then shown to be optimal.
As for the *combined-message* model, we give a protocol working in optimal completion time $O(D\Delta)$ (where D is the maximal source eccentricity) and message complexity $O(Dn)$. Finally, our protocol performs the (single) broadcast operation within the same optimal time and optimal message complexity $O(n)$.

1 Introduction

In a *radio* network, every node (station) can directly transmit to some subset of the nodes depending on the power of its transmitter and on the topological characteristics of the surrounding region. When a node u can directly transmit to a node v, we say that there is a (wireless) directed link (u, v). The set of nodes together with the set of these links form a directed communication graph that represents the radio network. In the radio network model [BGI92,CGR02,CGGPR00,CR06], the communication is assumed to be synchronous: this allows to focus on the impact of the *interference* phenomenon on the network performance. When a node sends a message, the latter is sent in parallel on all outgoing links. However, since a single radio frequence is used (see [ABLP89,BGI92,CGGPR00]), when two or more neighbors of a node transmit at the same time slot, a *collision* occurs (due to interference) and the message is lost. So, a node can recover a message from one of its incoming links if and only if this link is the only one bringing in a message. The broadcast task consists of sending a *source message* from a given *source* node to all nodes of the network. The *completion time* of a broadcast protocol is the number of time slots required by the protocol to inform all (reachable) nodes. A node is *informed* if it has received the source message.
Another important task in radio networks is *gossiping*, i.e., n simultaneous and independent broadcast operations, each one from a different node [CGR02,CMS03,GPX05].

* Research partially supported by the EU under the EU/IST Project 15964 AEOLUS.

The completion time of a gossiping protocol is the number of time slots the protocol requires so that every source message m is received by all nodes reachable from the source of m. We will consider two transmission models: the *single-message* model [BII93] and the *combined-message* one [CGR02]: in the former every node can transmit and receive at most one source message per time-slot while, in the latter, source messages can be arbitrarily combined and sent/received in one time slot [CGR02,GPX05]. Broadcasting and gossiping are fundamental communication tasks in radio networks and they are the subject of several research works in both algorithmic and networking areas [BGI92,CGR02,CGGPR00,PR97,R96]. It is reasonable to claim that almost all major theoretical questions related to such tasks can be considered closed as far as *static* networks are considered: the network never changes during the entire protocol's execution (see Subsection 1.1).

However, radio networks are typically adopted in scenarios where *unpredictable* node and link faults happen very frequently. Node failures happen when some hardware or software component of a station does not work, while link failures are due to the presence of a new (artificial or natural) hurdle that does not allow the communication along that link. In ad-hoc networking, while it is sometimes reasonable to assume that nodes (thus the protocol) know the *initial* topology, nothing is known about the duration and the location of faults. Such faults may clearly happen *even during the execution of a protocol*. In the sequel, such kind of faults will be called *dynamical faults* or, simply, faults.

Theoretical results on broadcast and gossiping protocols in any scenario where the network topology may change during the protocol's execution are very few (see Subsection 1.1).

The model of faulty networks. We follow a high-level approach by considering *adversarial networks* [AS98,ABBS01,CMS04,P02,S01]. Arbitrary dynamical faults are decided by a deterministic adaptive *adversary*. We analyze the completion time and the message complexity (i.e. maximum number of transmitted messages) of broadcast and gossiping protocols with respect to *worst-case* adversary's strategies.

The (worst-case) completion time of a *Fault-tolerant Broadcast* (in short, *FB*) protocol on a network G is defined as the maximal number (with respect to any possible adversarial strategy) of time slots required to *inform* all nodes reachable from the source in the *unpredictable* fault-free part of the network. More precisely, according to the fault-tolerance model adopted in [KKP98,P02,CMS04], a *fault pattern* F is a function (managed by the adaptive adversary) that maps every time-slot t to the subset $F(t)$ of nodes and links that are faulty during time slot t. The residual subgraph G^F is the graph obtained from G by removing *all* those nodes and links that belong to $F(t)$, *for some* time-slot t during the protocol's execution. Then, a FB protocol for a graph G is a broadcast protocol that, for any source s, and for any *fault pattern* F, guarantees that every node, which is reachable from s in the *residual subgraph* G^F, will receive the source message. The *residual eccentricity* of a node v is its eccentricity in the residual graph. The eccentricity of v is the maximal oriented distance (i.e. number of hops) from v to a reachable node.

The above definitions can be easily extended to *Fault-tolerant Gossiping* (in short *FG*) protocols: For any source s, message m_s must be received by every node reachable from s in G^F, for any choice of fault pattern F.

It is important to remark that if a node v is not reachable from a source in the residual subgraph, then the arrival of m_s to v is not considered in the analysis of the completion time. This assumption might be considered too strong but it is *necessary*. Indeed, it is easy to see that any attempt to consider larger residual subgraphs makes the worst-case completion time of *any* deterministic FG protocol *infinite*. This is well-explained by the following simple game. Consider k informed nodes that are in the in-neighborhood of a non informed node w. It is easy to see that *any* deterministic protocol, trying to inform w, fails *forever* against the following simple adversary's strategy: if at least two of the k in-neighbors transmit then the adversary leaves all edges on, while if there is exactly one of them transmitting, then the adversary makes only this link faulty. Observe that w is always *connected* to the informed part of the network but it will never receive the message (w is indeed not in the *residual* graph).

On the other hand, broadcasting and gossiping (and their analysis) in the residual graph is much harder than the same operation in fault-free radio networks. This is mainly due to the presence of unknown collisions that the adversary can produce at any time-slot *on the residual graph too*. As a matter of fact, while the completion time of broadcast on general fault-free radio networks of source eccentricity D is $O(D + \log^3 n)$ [GPX05], it turns out that there is a class of radio networks of *constant* source eccentricity where the same operation, in the above fault model, requires $\Theta(n\sqrt{n})$ time slots [CMS04]. So, in general graphs of "small" source eccentricity, the completion time gap may be exponential. The lower bound $\Omega(n\sqrt{n})$ in [CMS04] provides also a strong evidence of the significant difference between *dynamical* faults (on the residual graph) and *permanent* faults: in the latter network scenario, worst-case broadcasting time is $O(n \log^2 n)$ [CGR02].

Our results. We investigate *directed Geometric Radio Networks*, in short directed GRN [ENW00,KKKP00,CKOZ03,CCPRV01,DP07]. A directed GRN $G(V, E)$ is constructed by arbitrarily placing n nodes on the Euclidean plane; then, to each node v a transmission range $r_v \geq 0$ is assigned. These transmission ranges uniquely determine the set E of *directed* links: $(u, v) \in E$ iff $d(u, v) \leq r_u$, where $d(u, v)$ denotes the Euclidean distance between u and v. When all nodes have the same transmission range, the resulting graph is symmetric: this restriction is denoted as *symmetric* GRN.

We provide the first optimal bounds on the completion time and message complexity of FG protocols (and FB ones) in directed GRN for both single-message and combined-message models. More precisely, for the first model, given any directed GRN G of n nodes and maximal in-degree Δ, our FG protocol works in $O(n\Delta)$ time-slots and it has message complexity $O(n^2)$. Such bounds are then shown to be optimal.

Then, we consider the combined-message model and provide an FG protocol that works in optimal $O(D\Delta)$ time-slots (D denotes the maximal residual source eccentricity) and it has message complexity $O(n^2)$. We emphasize that this is the first FG protocol whose completion-time does not (explicitly) depend on n. Furthermore, the protocol can be easily analyzed for the (single) broadcast task: in this case, the completion time is still $O(D\Delta)$ while the message complexity reduces to $O(n)$. Both upper bounds are again optimal and, as for time complexity, it improves over the best (polynomial-time constructible) FB upper bound for general graphs by an $O(\log^3 n)$ factor ([CMS04] - see Subsection 1.1).

Adopted techniques. Since the fault pattern is unpredictable, an FG protocol must have the following "connectivity" property: it must consider all possible paths from a source to any node reachable from that source. To this aim, our protocols make an iterative use of collision-free families. A *collision-free family* is a set family (defined on the out-neighborhoods of the input graph - see Definition 2.1) that induces a transmission scheduling that somewhat guarantees the above connectivity property and yields *no* collision. So, when a node is scheduled as transmitter, its message is safely received by *all* its out-neighbors in the residual graph. This important fact is one of the key ingredients to get optimal message complexity (and thus energy efficiency) of our protocols. On the other hand, the size of the collision-free family is a linear factor in the completion time of our FG protocols. A crucial step in our protocol design is thus the efficient construction of a collision-free family for the input graphs. We indeed provide an algorithm that constructs an optimal-size collision-free family for any directed GRN working in time $O(n^2)$.
We observe that, given access to a collision-free family for the input graph, our protocols run in a fully-distributed fashion. However, in order to construct such optimal collision-free family it is necessary to know the *initial* graph topology. In Section 3, we also provide an efficient distributed construction of collision-free families under a much weaker knowledge condition: each node construct its own scheduling (so, "its" component of the collision-free family) by assuming that it only knows its position and a good approximation of the minimal distance among nodes. We then prove that if the (unknown) initial topology is *well spread* [CPS04], the returned collision-free family has optimal size, thus yielding the same protocol's performance given by the centralized construction. Well spread instances (see Definition 3.11) are a natural and broad generalization of *grid* networks.

1.1 Related works

Permanent faults. In [KKP98], the authors consider the broadcast operation in presence of permanent unknown node faults for two restricted classes of networks. They derive a $\Theta(D + \log \min\{\Delta, t\})$ bound where D is the source eccentricity in the residual graph and t is the number of faults. More recently, the issue of permanent-fault-tolerant broadcasting in general networks has been studied in [CGGPR00,CGR02,CMS03]. In these papers, several lower and upper bounds on the completion time of broadcasting are obtained in the *unknown* fault-free network model. We observe that the results obtained in unknown networks apply to general networks with permanent faults. In particular, in [CMS03], an $\Omega(n \log D)$ lower bound for the broadcast completion time is proved. The best general upper bound is $O(n \log^2 n)$ [CGR02]. In [CMS03], the authors provide a protocol having $O(D\Delta \log^2 n)$ completion time.
In [GL02], a gossiping protocol for unknown networks is given that works in $O(n^{1.5} \log^2 n)$ time. [CMS03] provides a permanent-fault tolerant gossiping protocol having $O(D\Delta^2 \log^2 n)$ completion time. The above results work for the combined-

message model. As for the single-message model, in [CMS03], a deterministic gossiping protocol is given that has $O(n\Delta^2 \log^3 n)$ completion time. We also mention the protocol for unknown directed GRN working in $O(n)$ time given in [DP07], even though it does not work for faulty networks.

Dynamical faults. We emphasize that *all* the above protocols *do not work* in presence of dynamical faults. As mentioned before, this is mainly due to the collisions yielded by any unpredictable wake-up of a faulty node/link during the protocol execution. Our dynamical fault model has been studied in [CMS04] where the *round robin* strategy is proved to be optimal for general graphs. Then, they show the existence of a deterministic FG protocol having $O(D\Delta \log n)$ completion time. The protocol is based on a probabilistic construction of *ad-hoc strongly-selective families* [CMS03,I02] for general graphs. Such families have a weaker property than collision-free ones: this weakness yields a not efficient message complexity. By adopting the efficient construction of such families in [I97], they can efficiently construct a FG protocol having $O(D\Delta \log^3 n)$ completion time. These protocols only hold for the combined-message model. In [PP05] an initial graph is given and, at each time slot, every node is faulty with probability p, where p is a fixed positive *constant* such that $0 < p < 1$. They prove an $O(\mathit{opt} \log n)$ bound for the broadcast completion time where *opt* is the optimal completion time in the fault-free case. They also prove that it is impossible to achieve $O(\mathit{opt} + \log n)$ completion time.
It is not hard to see that, when the graph is *symmetric*, any *distance-2 coloring* [C06] of size k yields a collision-free family of size k and viceversa. For some classes of undirected graphs, there are efficient constant-factor approximation algorithms that find a distance-2 coloring. In particular, for *unit disk graphs* [C06,CCJ90,SM97] a 7-approximation algorithm is presented in [SM97]. Since symmetric GRN in the plane are equivalent to unit disk graphs, the latter algorithm can be used to construct a collision-free family for this class of symmetric radio networks. However, this coloring algorithm *does not work* for *directed* GRN.

2 Collision-free families and fault-tolerant gossiping

In this section we introduce collision-free families and we show how to exploit them to design fault-tolerant gossiping protocols.

Definition 2.1 (Collision-free families). *Let $G(V, E)$ be a directed graph and let V' be the set of nodes that have at least one out-neighbor. A* collision-free family $\mathcal{S}$ *for G is a partition $\mathcal{S} = \{S_1, \ldots, S_k\}$ of V', such that, for each $S \in \mathcal{S}$ and for each $x, y \in S$ with $x \neq y$, $N^{out}(x) \cap N^{out}(y) = \emptyset$.*

In the sequel, we assume that, given any directed graph $G(V, E)$, we have at hand a collision-free family $\mathcal{S} = \{S_1, S_2, \ldots, S_k\}$ for G. In Section 3 we will then show how to construct collision-free families of small size.

Single-message model. In this model every transmission can contain only one of the source messages. We assume that each message contains the unique ID number of its source so that different messages have different ID's. The following FG protocol makes use of message IDs to define a *priority* queue in every node.

Protocol PRIO-SELECT($\mathcal{S}$) consists of a sequence of consecutive *phases*. Each *phase* consists of $k = |\mathcal{S}|$ time-slots. At the very beginning, the priority queue of every node u contains only m_u. At the beginning of every phase, every node v extracts (if any) the message $\hat{m}$ of *highest priority* (i.e. the maximal ID number) from its priority queue. Then, at time-slot j of a phase, node v acts according to the following rules

– If $v \in S_j$ and $\hat{m}$ exists then v transmits $\hat{m}$.
– In all other cases, v acts as receiver. If v receives a message m for the *first time* then m is enqueued, otherwise it is discarded.

Theorem 2.2. *Given a collision-free family $\mathcal{S}$ of size k for a directed graph G,* PRIO-SELECT($\mathcal{S}$) *completes fault-tolerant gossiping in G within $O(nk)$ time slots and message complexity $O(n^2)$.*

Sketch of Proof. Let F be the fault-pattern. As a direct consequence of the collision-free property of $\mathcal{S}$, we have the following

> ***Claim.*** *Every node v, having a new message m to send at the beginning of a phase, will successfully send m to* all *its out-neighbors in the residual graph G^F by the end of the phase.*

Another important fact is that any message cannot be delayed "too much" from any other message. Indeed, consider any path $p(s, v)$ of length l, from a source s to a node v in G^F. If m_s does never meet any other message of higher priority in $p(s, v)$, then Claim 2 implies that v receives m_s within l phases. On the other hand, assume a message m_q of priority higher than m_s is in some priority queue of a node u in the path $p(s, v)$. It is possible to prove that m_q can delay m_s on $p(s, v)$ for at most one phase. Since at most n messages are contained in any priority queue in $p(s, v)$ and $l \leqslant n$, we get the $O(nk)$ upper bound.
The collision-free property of the family implies that every node transmits every source message at most once, so the message complexity is $O(n^2)$. □

Combined-message model. In this model, source messages can be arbitrarily combined and sent in one transmission.

Protocol MULTI-SELECT($\mathcal{S}$). Each node v keeps the set $M_{old}(v)$ of the messages already sent by node v and the set $M_{new}(v)$ of the messages that node v has to send. At the beginning of the protocol, $M_{new}(v)$ contains only the source message of node v and the set $M_{old}(v)$ is empty. The protocol consists of a sequence of consecutive *phases*. Each *phase* consists of $k = |\mathcal{S}|$ time-slots. All phases are identical. At time slot j of a phase, node v acts according to the following rules

– If $v \in S_j$ and $M_{new}(v)$ is not empty then v transmits all the messages in $M_{new}(v)$ and moves all these messages to the set $M_{old}(v)$;
– In all other cases, v acts as receiver. When v receives a message m, if it is not in $M_{old}(v)$ then it is added to $M_{new}(v)$. Otherwise m is discarded.

The proof of the following theorem is similar (even simpler) to that of Theorem 2.2.

Theorem 2.3. *Given a collision-free family $\mathcal{S}$ of size k for a directed graph G,* MULTI-SELECT($\mathcal{S}$) *completes fault-tolerant gossiping in G within $O(Dk)$ time-slots and message complexity $O(Dn)$, where D is the maximal residual source eccentricity. Moreover, an easy adaptation of* MULTI-SELECT($\mathcal{S}$) *for the broadcast operation works with the same completion time while the message complexity reduces to $O(n)$.*

3 Explicit constructions of collision-free families

In this section we first present a greedy algorithm that, given any directed GRN, constructs a collision-free family of optimal size. Then we show a distributed construction that, for *well spread* instances (see Definition 3.11), yields collision-free families of optimal size as well.

Centralized construction. Given a set V of points (i.e. nodes) in $\mathbb{R}^2$ and a range assignment $r : V \to \mathbb{R}^+$, the directed GRN is uniquely determined and it will be denoted as $G_r(V)$. Indeed, for each node $v \in V$, let $B(v)$ be the closed disk of center v and radius $r(v)$, i.e., $B(v) = \{x \in \mathbb{R}^2 \ : \ d(v,x) \leqslant r(v)\}$. We define the in-neighborhood of a node $v \in V$ as the set $N^{in}(v) = \{w \in V \ : \ v \in B(w)\}$. We define $\Delta(v) = |N^{in}(v)|$ and the maximal in-degree of $G_r(V)$ as $\Delta = \max_{v \in V} \Delta(v)$.
We will show that, given any directed GRN $G_r(V)$ as input, the following algorithm CFF returns a collision-free family $\mathcal{S}$ for $G_r(V)$ of size $O(\Delta)$. Since $\Omega(\Delta)$ is a trivial lower bound for such families, the one returned by CFF is asymptotically optimal.
The algorithm constructs every set of $\mathcal{S}$ by inserting nodes whose range disks are pairwise disjoint. Nodes are inserted in a non increasing order w.r.t. their ranges. This set construction is repeated until no node of V' is left outside $\mathcal{S}$.

```
Algorithm CFF (a finite set V ⊆ ℝ², a function r : V → ℝ⁺)
Let X := V' = {v ∈ V : N^out(v) ≠ ∅};
Let S := ∅;
Let i := 0;
while X ≠ ∅ do begin
        i := i + 1;
        S_i := ∅;
        U := ∅;
        Y := X;
        while Y ≠ ∅ do begin
                Choose v ∈ Y such that r(v) is maximum;
                if U ∩ B(v) = ∅ then begin
                        S_i := S_i ∪ {v};
                        U := U ∪ B(v);
                end;
                Y := Y − {v};
        end;
        S := S ∪ {S_i};
        X := X − S_i;
end;
return S.
```

It is easy to see that, by using standard data structures, the algorithm works in $O(n^2)$ time.

Lemma 3.1. *Family $\mathcal{S}$ is collision-free for $G_r(V)$.*

Proof. We first observe that $\mathcal{S}$ is a partition of V'. Let $S_i \in \mathcal{S}$ and $u, v \in S_i$ such that $u \neq v$. Thanks to line 11 of the algorithm, it holds that $B(u) \cap B(v) = \emptyset$; so $N^{out}(u) \cap N^{out}(v) = \emptyset$. □

We now provide a preliminary bound on the size of $\mathcal{S}$. For every $v \in V'$, we define the set $I(v)$ of all nodes of V' that could interfere with v and that have range not smaller than the range of v, i.e.,

$$I(v) = \{w \in V' \ : \ B(v) \cap B(w) \neq \emptyset \ \text{ and } \ r(w) \geqslant r(v)\}$$

Lemma 3.2. *Family $\mathcal{S}$ has size at most $\max_{v \in V'} |I(v)|$.*

Proof. At every iteration of the external loop (line 4), a new set of $\mathcal{S}$ is constructed. Consider the i-th iteration and let $v \in V'$ be any node not yet inserted in any of sets $S_1, S_2, \ldots, S_{i-1}$ constructed in the previous iterations. For every $j = 1, 2, \ldots, i-1$, S_j must contain at least one node in $I(v)$. Indeed, assume by contradiction that there exists $j \leqslant i-1$ such that $S_j \cap I(v) = \emptyset$. Then, for every $w \in S_j$ with $r(w) \geqslant r(v)$, it holds that $B(w) \cap B(v) = \emptyset$. When the algorithm selects v in line 10, the condition at line 11 is true, so v should be inserted in S_j: a contradiction. Since the sets of $\mathcal{S}$ are pairwise disjoints, the number of iterations of the external loop does not exceed $\max_{v \in V'} |I(v)|$. □

Our next goal is to prove that $\max_{v \in V'} |I(v)| \in O(\Delta)$. To this aim, we will show that, for every $v \in V'$, we can partition $\mathbb{R}^2$ into a constant number of *regions* so that each region contains at most Δ nodes of $I(v)$ (see Fig. 1).

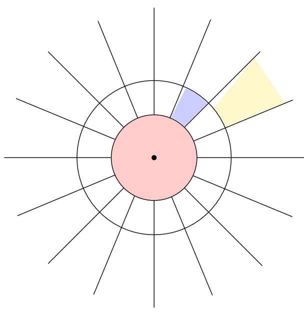

Fig. 1. Partition of $\mathbb{R}^2$ around node $v \in V$. In each region there are at most Δ points of $I(v)$

Lemma 3.3. *For every $v \in V'$, it holds that $|B(v) \cap I(v)| \leqslant \Delta$.*

Proof. Nodes in $I(v)$ have range at least $r(v)$. Hence, all nodes of $I(v)$ in $B(v)$ are points of $N^{in}(v)$, i.e., $I(v) \cap B(v) \subseteq N^{in}(v)$. □

We now consider the region outside disk $B(v)$ and define the circular crown

$$C_\lambda(v) = \{y \in \mathbb{R}^2 \ : \ r(v) < d(v,y) \leqslant \lambda r(v)\}, \text{ where } \lambda > 1$$

Lemma 3.4. *Let $1 < \lambda < 2$ and let $k \in \mathbb{N}$ be large enough such that $\cos \frac{2\pi}{k} \geqslant \lambda/2$. Then, for any $v \in V'$, $C_\lambda(v)$ contains at most $k\Delta$ nodes of $I(v)$.*

Proof. Consider a polar coordinate system centered in v and consider the partition of $C_\lambda(v)$ defined by the regions

$$]r(v), \lambda r(v)] \quad \times \quad [\vartheta_i, \vartheta_{i+1}[\text{ where } \vartheta_i = \frac{2\pi i}{k}, \text{ for every } i = 0, 1, \ldots, k-1$$

Then, since $\cos \frac{2\pi}{k} \geqslant \lambda/2$, it is easy to see that the square of the maximal distance between two points in the same region is

$$r(v)^2 + \lambda^2 r(v)^2 - 2\lambda r(v)^2 \cos \frac{2\pi}{k} \leqslant r(v)^2$$

For any $w \in I(v)$, it holds that $r(w) \geqslant r(v)$, so w is in the in-neighborhood of all points in the same region of w. So, in every region there are at most Δ points of $I(v)$ and, since there are k regions in $C_\lambda(v)$, the thesis follows. □

Lemma 3.5. *For each $\lambda > 1$, there exists an angle $\varphi > 0$ such that for each $a, b \in \mathbb{R}$ with $a \geqslant b \geqslant \lambda$, it holds that*

$$a^2 + b^2 - 2ab\cos\varphi \leqslant (a-1)^2$$

Proof. Let be $a \geqslant b \geqslant \lambda > 1$. It holds that

$$a^2 + b^2 - 2ab\cos\varphi \leqslant (a-1)^2 \Longleftrightarrow \cos\varphi \geqslant \frac{b^2 + 2a - 1}{2ab}$$

Consider the function

$$f(a,b) = \frac{b^2 + 2a - 1}{2ab} = \frac{b^2 - 1}{2b} \cdot \frac{1}{a} + \frac{1}{b}$$

For any fixed $b > 1$, function f is decreasing in variable a and, since it must be $a \geqslant b$, the maximum is obtained when $a = b$. Then, consider function

$$g(b) = f(b,b) = \frac{b^2 - 1 + 2b}{2b^2} = \frac{1}{2} + \frac{1}{b} - \frac{1}{2b^2}$$

its derivative is $g'(b) = \frac{1}{b^3} - \frac{1}{b^2}$. Hence $g'(b) < 0$ for each $b > 1$, so also $g(b)$ is decreasing and, since it must be $b \geqslant \lambda$, the maximum is obtained when $b = \lambda$. This maximum is

$$g(\lambda) = \frac{1}{2} + \frac{2\lambda - 1}{2\lambda^2}$$

By summarizing, for $\lambda > 1$ function g is decreasing, strictly less than 1 and tends to $1/2$ as $\lambda \to \infty$. So for each $\lambda > 1$, we have that $1/2 < g(\lambda) < 1$. Then for each $\lambda > 1$ we can choose an angle $\varphi > 0$ such that $\cos\varphi \geqslant g(\lambda)$. This implies that

$$g(\lambda) \geqslant \frac{b^2 + 2a - 1}{2ab}$$

for each $a \geqslant b \geqslant \lambda$. □

Consider the function $g(\lambda) = \frac{\lambda^2+2\lambda-1}{2\lambda^2}$ and observe that $1/2 < g(\lambda) < 1$, for any $\lambda > 1$.

Lemma 3.6. *Let $\lambda > 1$ and let $k \in \mathbb{N}$ be large enough such that $\cos\frac{2\pi}{k} \geqslant g(\lambda)$. Then, for any $v \in V'$, there are at most $k\Delta$ nodes of $I(v)$ outside $B(v) \cup C_\lambda(v)$.*

Proof. Consider a polar coordinate system centered in v, and define a partition of the space outside $B(v) \cup C_\lambda(v)$ in the regions

$$[\lambda r(v), +\infty[\times [\vartheta_i, \vartheta_{i+1}[\text{ where } \vartheta_i = \frac{2\pi i}{k},\ i = 0, 1, \ldots, k-1$$

Let $x = (\varrho_x, \varphi_x)$ and $y = (\varrho_y, \varphi_y)$ two nodes of $I(v)$ that lie in the same region and suppose wlog that $\varrho_x \geqslant \varrho_y$. Then, two constants $a, b \in \mathbb{R}$ exist with $a \geqslant b \geqslant \lambda$ such that $\varrho_x = a \cdot r(v)$ and $\varrho_y = b \cdot r(v)$. We thus get

$$\begin{aligned} d(x,y)^2 &= \varrho_x^2 + \varrho_y^2 - 2\varrho_x\varrho_y \cos(\varphi_x - \varphi_y) \leqslant \varrho_x^2 + \varrho_y^2 - 2\varrho_x\varrho_y \cos\frac{2\pi}{k} \\ &= r(v)^2 \left(a^2 + b^2 - 2ab\cos\frac{2\pi}{k}\right) \end{aligned}$$

where in the first inequality we used the fact that x and y lie in the same region. From Lemma 3.5, we get

$$d(x,y)^2 \leqslant r(v)^2(a-1)^2 = (a \cdot r(v) - r(v))^2 = (\varrho_x - r(v))^2$$

Since $x \in I(v)$, it must hold that $B(x) \cap B(v) \neq \emptyset$, so $\varrho_x - r(v) \leqslant r(x)$, and $d(x,y)^2 \leqslant r(x)^2$. Therefore, y lies in $B(x)$ and, thus, $x \in N^{in}(y)$.
It follows that, for every region T, if $y \in T \cap I(v)$ is a node with minimum distance from v, i.e, a node with minimum ϱ_y, then $T \cap I(v) \subseteq N^{in}(y)$. This implies that in every region there are at most Δ points of $I(v)$: since the regions are k, the thesis follows. □

Lemma 3.7. *Let $1 < \lambda < 2$ and let $k \in \mathbb{N}$ be such that $\cos\frac{2\pi}{k} \geqslant \max\left\{\frac{\lambda}{2}, g(\lambda)\right\}$. Then, for any $v \in V'$, it holds that $|I(v)| \leqslant (1+2k)\Delta$.*

Proof. Consider the partition of $\mathbb{R}^2$ into the following three sets: 1. Disk $B(v)$; 2. Circular crown $C_\lambda(v)$; 3. The complement of $B(v) \cup C_\lambda(v)$. By combining Lemmas 3.3, 3.4, and 3.6, we get $|I(v)| \leqslant (1+k+k)\Delta$. □

Theorem 3.8. *Algorithm* CFF *returns a collision-free family $\mathcal{S}$ for $G_r(V)$ of size at most $c\Delta$, where $c \leqslant 33$.*

Proof. From Lemma 3.1, $\mathcal{S}$ is collision-free for $G_r(V)$. Let λ be such that $1 < \lambda < 2$. From Lemmas 3.2 and 3.7, we obtain $|\mathcal{S}| \leqslant \max_{v \in V'} |I(v)| \leqslant (1+2k)\Delta$, with $k \in \mathbb{N}$ such that $\cos \frac{2\pi}{k} \geqslant \max\{\lambda/2,\, g(\lambda)\}$. Then, in order to minimize k, we choose λ such that

$$\frac{\lambda}{2} = \frac{\lambda^2 + 2\lambda - 1}{2\lambda^2} \tag{1}$$

Consider the function $f(\lambda) = \lambda^3 - \lambda^2 - 2\lambda + 1$. Then $f(1) = -1$ and $f(2) = 1$, so (1) has a solution between 1 and 2. By numerical arguments, we can set $\lambda \approx 1.8$ and get

$$\cos \frac{2\pi}{k} \geqslant \max\left\{\frac{\lambda}{2},\, \frac{\lambda^2 + 2\lambda - 1}{2\lambda^2}\right\}, \quad \text{for any } k \geqslant 16$$

□

Distributed construction. Let us consider GRN $G_r(V)$ where $r(v) = R$ for each $v \in V$ (so $G_r(V)$ is symmetric). Directed GRN will be discussed at the end of this section. Our distributed construction of a collision-free family for $G_r(V)$ is based on the following idea. Consider a partition of $\mathbb{R}^2$ into squares small enough to guarantee that in each square there is at most one node of V. Then we partition the set of such small squares so that the distance between two squares in the same set of the partition is at least $2R$. Finally, consider the subsets of V obtained by collecting all nodes in the same squares' set (see Fig. 2).
Let $\gamma = \min\{d(u,v) \;:\; u, v \in V,\, u \neq v\}$. For any $x \in \mathbb{R}$ we define $[x]$ as

$$[x] = \begin{cases} \lfloor x \rfloor & \text{if } x - \lfloor x \rfloor \leqslant 1/2 \\ \lfloor x \rfloor + 1 & \text{if } x - \lfloor x \rfloor > 1/2 \end{cases}$$

We now assume that each node knows its own position, the transmission range R and the minimum distance γ. In the following algorithm, $\varepsilon > 0$ is an arbitrary small constant: we need it in order to have strict inequalities.

ALGORITHM FOR NODE u (position (x_u, y_u), transmission range R, min distance γ)

1 **Define** $\lambda = \gamma/\sqrt{2} - \varepsilon$;
2 **Define** $k = \lceil (2R + \varepsilon)/\lambda \rceil + 1$;
3 **Define** $\hat{x}_u = [x_u/\lambda]$ and $\hat{y}_u = [y_u/\lambda]$;
4 **Return** $f(u) = (\hat{x}_u \bmod k,\, \hat{y}_u \bmod k)$;

Let us consider the family $\mathcal{S} = \{S_{i,j}\}_{i,j=0,1,\ldots,k-1}$ where $S_{i,j} = \{u \in V \;:\; f(u) = (i,j)\}$. We now show that $\mathcal{S}$ is a collision-free family.

Lemma 3.9. *Let u, v be two distinct nodes in the same set $S_{i,j}$, then $d(u,v) > 2R$.*

Proof. Let $u, v \in S_{i,j}$ then by line 4 in the algorithm we have

$$\begin{aligned} \hat{x}_u &= i + a_u k \quad & \hat{x}_v &= i + a_v k \\ \hat{y}_u &= j + b_u k \quad & \hat{y}_v &= j + b_v k \end{aligned}$$

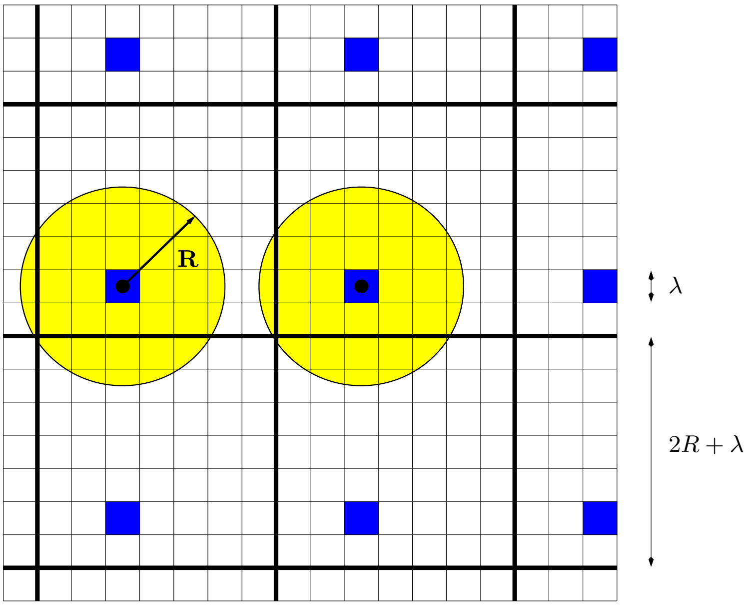

Fig. 2. In each small square there is at most one point. In each big square there are $\Theta(R^2/\gamma^2)$ small squares. Two nodes in the same set of the family do not create collisions.

with $a_u, a_v, b_u, b_v \in \mathbb{N}$. Hence

$$|\hat{x}_u - \hat{x}_v| = |a_u - a_v|k$$
$$|\hat{y}_u - \hat{y}_v| = |b_u - b_v|k$$

Observe that, by line 3 in the algorithm, it follows that

$$\frac{x}{\lambda} - \frac{1}{2} \leqslant \hat{x} \leqslant \frac{x}{\lambda} + \frac{1}{2} \tag{2}$$

and

$$\lambda\hat{x} - \frac{\lambda}{2} \leqslant x \leqslant \lambda\hat{x} + \frac{\lambda}{2} \tag{3}$$

If it were $a_u = a_v$ and $b_u = b_v$ then we would have $|\hat{x}_u - \hat{x}_v| = |\hat{y}_u - \hat{y}_v| = 0$. By (3), we get

$$|x_u - x_v| \leqslant |\lambda\hat{x}_u + \frac{\lambda}{2} - (\lambda\hat{x}_v - \frac{\lambda}{2})| = \lambda$$
$$|y_u - y_v| \leqslant |\lambda\hat{y}_u + \frac{\lambda}{2} - (\lambda\hat{y}_v - \frac{\lambda}{2})| = \lambda$$

So it would be $d(u, v)^2 \leqslant 2\lambda^2 < \gamma^2$. And this is a contradiction because γ is the minimal distance between two nodes. Now suppose wlog that $a_u \neq a_v$, then we have

$|\hat{x}_u - \hat{x}_v| = |a_u - a_v|k \geqslant k$. From Equation (2) it holds that

$$|\hat{x}_u - \hat{x}_v| \leqslant \left| \frac{x_u}{\lambda} + \frac{1}{2} - \frac{x_v}{\lambda} + \frac{1}{2} \right| \leqslant \frac{|x_u - x_v|}{\lambda} + 1$$

So $d(u, v) \geqslant |x_u - x_v| \geqslant (k-1)\lambda > 2R$. □

Theorem 3.10. *Family $\mathcal{S}$ is collision-free for $G_r(V)$ of size $O(R^2/\gamma^2)$.*

Proof. By definition of k (line 2 of the algorithm) we have that $|\mathcal{S}| = k^2 \in O(R^2/\gamma^2)$. Let $u, v \in S_{i,j}$ with $u \neq v$. Assume, by contradiction, that $N^{out}(u) \cap N^{out}(v) \neq \emptyset$ and let $w \in N^{out}(u) \cap N^{out}(v)$. So $d(u, w) \leqslant R$ and $d(v, w) \leqslant R$. By triangular inequality we get $d(u, v) \leqslant 2R$ thus contradicting Lemma 3.9. □

We now show that when nodes are *well spread*, the size of the family is asymptotically optimal.

Definition 3.11 (Well spread instances). *Let $V \subseteq \mathbb{R}^2$ be a set of n points in the Euclidean plane. Let γ and Γ be respectively the minimal and the maximal distance between two points in V. Let c be any positive constant, set V is said c-*well spread *if $\Gamma/\gamma \leqslant c\sqrt{n}$.*

Observe that square-grid networks are the most regular case of c-well spread instances where $c = \sqrt{2}$ [CPS04].

Theorem 3.12. *If $V \subseteq \mathbb{R}^2$ is a c-well spread instance, then $R^2/\gamma^2 \in O(c^2\Delta)$, where Δ is the maximal degree of $G_r(V)$.*

Proof. There exists a disk of radius Γ that contains all the n nodes. That disk can be covered with $O(\Gamma^2/R^2)$ disks of radius R. Then there exists a disk U with radius R such that it contains $\Omega\left(\frac{nR^2}{\Gamma^2}\right)$ nodes. Since V is c-well spread, $n/\Gamma^2 \in \Omega(1/c^2\gamma^2)$ and so disk U contains $\Omega(R^2/c^2\gamma^2)$ nodes. It follows that $R^2/\gamma^2 \in O(c^2\Delta)$. □

Our distributed construction also works for directed GRN where parameter R is replaced by the maximal node range $R_{\max}$. Theorem 3.10 holds with $R_{\max}$ in place of R and Theorem 3.12 holds with $R_{\min}$ in place of R, where $R_{\min}$ is the minimal node range. Thus if $R_{\max} \in O(R_{\min})$ the construction is still optimal.

4 Optimal Bounds

The results obtained in the previous two sections allow us to get optimal bounds for fault-tolerant protocols.

Single-message model.

Corollary 4.1.

(a) *Given a directed GRN $G_r(V)$, there exists an explicit FG protocol having completion time $O(n\Delta)$ and message complexity $O(n^2)$, where Δ is the maximal in-degree of $G_r(V)$.*

(b) *There exists a distributed FG protocol that, on any c-well spread symmetric GRN $G_r(V)$, completes gossiping in $O(nc^2\Delta)$ time slots and has message complexity $O(n^2)$. The protocol requires the knowledge of the minimal distance γ.*

Proof. Part (a) is a direct consequence of Theorem 2.2 and Theorem 3.8. Part (b) is obtained by combining Theorems 2.2, 3.10 and 3.12. □

Next theorem shows that these bounds are tight.

Theorem 4.2. *For any sufficiently large n and Δ, such that $n - \Delta \in \Omega(n)$, there exists a GRN $G_r(V)$ of n nodes and maximal in-degree Δ such that, for any FG protocol for $G_r(V)$, an adversary's fault-pattern F exists such that the protocol is forced to execute $\Omega(n\Delta)$ time-slots and to have message complexity $\Omega(n^2)$.*

Proof. For any sufficiently large n and for any Δ such that $n - \Delta \in \Omega(n)$, we consider a directed GRN $G(V, E)$ (see Fig. 4) consisting in a directed path C of $n - \Delta - 1$ nodes whose last node is connected to a set L of Δ nodes. These latter nodes are connected to one sink node w. Observe that it is always possible to set the relative distance among such nodes so that all the above described links are directed from left to right (no backwards edge exists). Consider any fixed FG protocol $\mathcal{P}$ working on G. The adversary first makes faulty all the edges among the nodes in L for the entire execution of $\mathcal{P}$. Such permanent-fault pattern will be denoted as $\mathcal{F}$. In this way, the nodes in L cannot exchange any information.

> ***Claim.*** *For any $u \in C$ and for any $v \in L$, there must exist a time slot τ where node v transmits m_u and all the other nodes in L do not transmit.*
>
> *Proof.* We first observe that, for any dynamical-fault pattern of the links from L to w that is combined with $\mathcal{F}$, the actions of $\mathcal{P}$ do not change. Assume by way of contradiction that the claim does not hold for some $\hat{u}$ and $\hat{v}$. Two cases may arise. If there is no time slot in which $\hat{v}$ transmit $m_{\hat{u}}$, then the adversary makes all the links from L to w faulty but link $(\hat{v}, w)$ which is always fault-free. In the second case, whenever $\hat{v}$ transmits $m_{\hat{u}}$, there is another node in L that transmits. Then, in this case, link $(\hat{v}, w)$ is always fault-free and all the other links from L to w are not faulty only when $\hat{v}$ transmits. It is then easy to see that, in both cases, w will never receive $m_{\hat{u}}$, but in the residual graph there is a path from $\hat{u}$ to w. □

From the above claim, it follows that FG Protocol $\mathcal{P}$ must execute at least $(n - \Delta - 1)\Delta$ time slots.
As for message complexity, we observe that every node in the rightmost half of path C must transmit all the $\Omega(n)$ source messages of the nodes in the leftmost half of path C. This clearly yields the $\Omega(n^2)$ lower bound. □

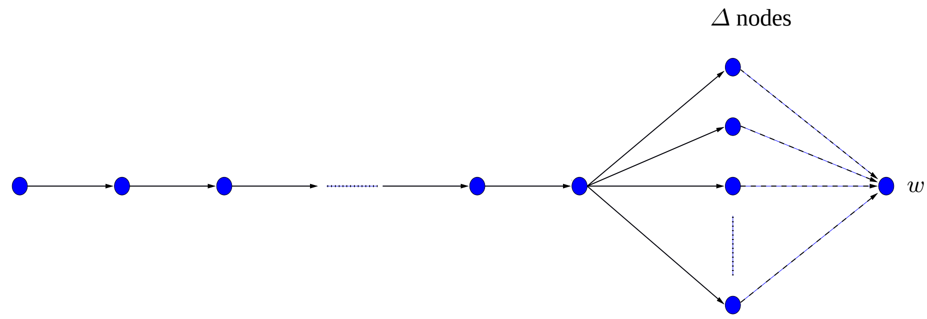

Fig. 3. Strategy of faults for lower bound $\Omega(\Delta n)$

Combined-message model.

Corollary 4.3.

(a) Given a directed GRN $G_r(V)$, there exists an explicit FG protocol having completion time $O(D\Delta)$ and message complexity $O(Dn)$, where D is the maximal residual source eccentricity.

(b) There exists a distributed FG protocol that, on any c-well spread symmetric GRN $G_r(V)$, completes gossiping in $O(Dc^2\Delta)$ time slots and has message complexity $O(Dn)$. The protocol requires the knowledge of the minimal distance γ.

As for the (single) broadcast operation, the same protocols work in the same completion time while the message complexity reduces to $O(n)$.

Proof. Part (a) is a direct consequence of Theorem 2.3 and Theorem 3.8. Part (b) is obtained by combining Theorems 2.3, 3.10 and 3.12. □

The above completion time bound is now shown to be tight.

Theorem 4.4. *For any n, Δ and D such that $D\Delta \leqslant n$, there exists a GRN $G_r(V)$ of n nodes and maximal in-degree Δ such that, for any FB protocol for $G_r(V)$, there are a source $s \in V$ and an adversary's fault-pattern F, yielding source eccentricity D, such that the protocol is forced to execute $\Omega(D\Delta)$ time-slots and to have message complexity $\Omega(n)$.*

Proof. Let PROT be any deterministic FB protocol. The adversary considers the following directed GRN G. Nodes are located on the plane and organized in *Layers*. All nodes of the l-th layer are (initially) connected to all nodes of the $l+1$-th layer, but not viceversa. This can be obtained by setting the distance between the l-th layer and the $l+1$-th one to be larger than the distance between the $l+1$-th layer and the $l+2$-th one, and so on. Nodes of the same Layer are connected to each other due to simple geometrical reasons. The first layer L_0 contains only the source node. The other layers are of two kinds and they alternate each other from left to right. Layer L_i consists in

$\Delta \geq 2$ nodes and layer L'_i consists in only one node (see Fig. 4). Clearly the number of layers is $O(n/\Delta)$.
The first adversary's action is to make all links between nodes of the same layer faulty, for all time slots, while links from L'_i to L_{i+1} are never faulty. Then, it allows the source to inform L_1. Now, the line of reasoning works by induction and it is based on the following fact.

Claim. *Let us assume level L_i is all informed at time slot t. Then, for any strategy of* PROT *the adversary can keep (at least) one link from L_i to L'_i never faulty (so it belongs to the residual graph) and L'_i is not informed before time-slot $t + \Delta$.*

Proof. Nodes in L_i are denoted as z_j, $j = 1, \ldots, \Delta$ and the node in L'_i is denoted as y. Then, the adversary reads the protocol's deterministic schedule and acts as follows. Consider any time slot t' such that $t \leq t' < t + \Delta$: If there is exactly one node $\tilde{z}_j$ in L_i that transmits, then the adversary makes (only) edge $(\tilde{z}_j, y)$ faulty, otherwise the adversary keeps all the edges on. It should be clear that, as far as $t' \leq \Delta - 1$, this adversary's strategy does not allow node y to be informed and it can keep at least one link from L_i to L'_i (always) non faulty. □

The above Claim implies that the adversary can define a fault pattern so that the source has eccentricity D and its message requires at least Δ time slots to go from any Level L_i to L'_i. Then, for any source eccentricity D and max in-degree Δ such that $D\Delta \in O(n)$, the directed GRN G together with the fault pattern described above yields the $\Omega(D\Delta)$ lower bound. □

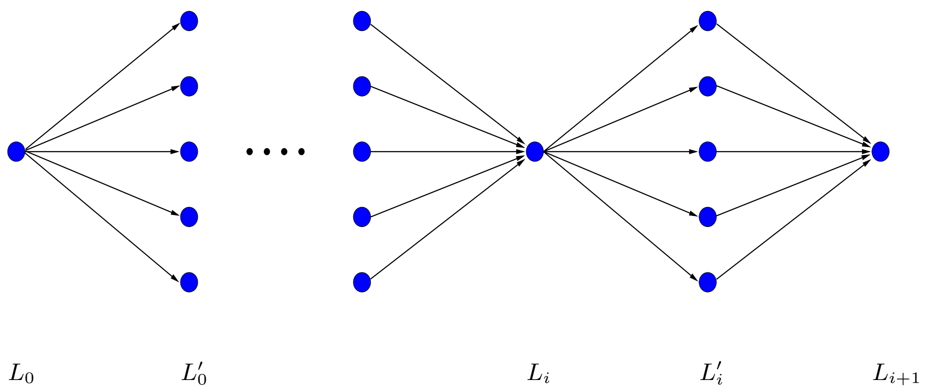

Fig. 4. Layered network

As for the case $D \cdot \Delta > n$, we observe that a lower bound $\Omega(n\sqrt{n})$ holds for FB protocols on directed GRN of unbounded maximal in-degree and residual source eccentricity

$D = \Theta(\sqrt{n})$. This result is an easy consequence of the lower bound for arbitrary graphs proved in [CMS04]: The graph yielding such lower bound is indeed a GRN of maximal in-degree $\Delta = \Theta(n)$.

5 Future work

In the combined-message model, we observe that our FB protocol has also optimal message complexity while it is an open question whether the $O(Dn)$ bound for the FG message complexity is optimal.
Another future work is that of extending our distributed construction of collision-free families to other important classes of radio networks.
Finally, an interesting issue is that of designing *randomized* FG protocols. Such protocols may yield a much better completion time on the residual graph and, more importantly, they might have good performances *outside* the residual graph too.

References

[AS98] M. Adler and C. Scheideler (1998), Efficient communication strategies for ad hoc wireless networks. Proc. of *10th ACM SPAA*, 259-268.

[ABLP89] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg (1991), A lower bound for radio broadcast, *Journal of Computer and System Science*, 43, 290-298.

[ABBS01] B. Awerbuch, P. Berenbrink, A. Brinkmann, and C. Scheideler (2001), Simple routing strategies for adversarial systems. Proc. of *42th IEEE FOCS*, 158-167.

[BGI92] R. Bar-Yehuda, O. Goldreich, and A. Itai (1992), On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization, *Journal of Computer and System Science*, 45, 104-126.

[BII93] R. Bar-Yehuda, A. Israeli, and A. Itai (1993), Multiple communication in multi-hop radio networks, *SICOMP*, 22 (4), 875-887.

[C06] T. Calamoneri (2006), The $L(h,k)$-Labeling problem: Survey and Annotated Bibliography, *The Computer Journal*, 49, 585-608.

[CKOZ03] G. Calinescu, S. Kapoor, A. Olshevsky, and A. Zelikovsky (2003), Network Life Time and Power Assignment in ad hoc Wireless Networks, Proc. of *11th ESA*, LNCS 2832, 114-126.

[CGGPR00] B. Chlebus, L. Gasieniec, A. Gibbons, A. Pelc, and W. Rytter (2000), Deterministic broadcasting in unknown radio networks. Proc. of *11th ACM-SIAM SODA*, 861-870.

[CGR02] M. Chrobak, L. Gasieniec, and W. Rytter (2002), Fast Broadcasting and Gossiping in Radio Networks. *Journal of Algorithms*, 43 (2), 177-189.

[CCJ90] B.N. Clark, C.J. Colbourn, and D.S. Johnson (1990), Unit disk graphs, *Discrete Mathematics*, 86, 165-177.

[CCPRV01] A.E.F. Clementi, P. Crescenzi, P. Penna, and P. Vocca (2001), On the complexity of minimum energy consumption broadcast subgraphs, Proc of *11th STACS*, 121-131.

[CMS03] A.E.F. Clementi, A. Monti, and R. Silvestri (2003), Distributed broadcast in radio networks of unknown topology. *Theoretical Computer Science*, 302, 337-364.

[CMS04] A.E.F. Clementi, A. Monti, and R. Silvestri (2004), Round Robin is optimal for fault-tolerant broadcasting on wireless networks, *Journal of Parallel and Distributed Computing*, 64(1), 89-96.

[CPS04] A.E.F. Clementi, P. Penna, R. Silvestri (2004), On the power assignment problem in radio networks, *MONET* (9), 125-140.

[CR06] A. Czumaj and W. Rytter (2006), Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60 (2), 115-143.

[DP07] A. Dessmark and A. Pelc (2007), Broadcasting in geometric radio networks, *Journal of Discrete Algorithms*, 5, 187-201.

[ENW00] A. Ephremides, G. Nguyen, and J. Wieselthier (2000), On the construction of energy-efficient broadcast and multi-cast trees in wireless networks, Proc. of *19th IEEE INFOCOM*, 585-594.

[GL02] L. Gasieniec and A. Lingas (2002), On adaptive deterministic gossiping in ad-hoc radio networks, Information Processing Letters, 2, 89-93.

[GPX05] L. Gasieniec, D. Peleg, and Q. Xin (2005), Faster Communication in Known Topology Radio Networks, Proc. of *24th ACM PODC*, 129-137.

[I97] P. Indyk (1997), Deterministic Superimposed Coding with Application to Pattern Matching, Proc. of *38th IEEE FOCS*, 127-136.

[I02] P. Indyk (2002), Explicit constructions of selectors and related combinatorial structures, with applications. Proc of *13th ACM-SIAM SODA*, 697-704.

[KKKP00] L. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc (2000), Power consumption in packet radio networks, *Theoretical Computer Science*, 243, 289-305.

[KKP98] E. Kranakis, D. Krizanc, and A. Pelc (1998), Fault-Tolerant Broadcasting in Radio Networks, Proc. of *6th ESA*, LNCS 1461, 283-294.

[P02] A. Pelc (2002), Broadcasting in radio networks, *Handbook of Wireless Networks and Mobile Computing*, John Wiley and Sons, Inc., 509-528.

[PP05] A. Pelc and D. Peleg (2005), Feasibility and complexity of broadcasting with random transmission failures. Proc. of *24th ACM PODC*, 334-341.

[PR97] E. Pagani and G. Rossi (1997), Reliable Broadcast in Mobile Multihop Packet Networks, Proc. of *3rd ACM-IEEE MOBICOM*, 34-42.

[R96] T. S. Rappaport (1996), *Wireless Communications: Principles and Practice*, Prentice Hall.

[SM97] A. Sen and E. Malesinska (1997), On Approximation algorithms for packet radio scheduling, Proc. of *35th Allerton Conference on Communication, Control and Computing*, Allerton, 573-582.

[S01] C. Scheideler (2001), Models and Techniques for Communication in Dynamic Networks. Proc. of *18th STACS*, 27-49.