

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

(Extended Abstract)

Andrea E.F. Clementi¹, Angelo Monti², Francesco Pasquale^{1, **},
and Riccardo Silvestri²

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 frequency 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.

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

^{**} Corresponding author.

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]. 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 constructs 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 3.8) are a natural and broad generalization of *grid* networks. Due to lack of space, some proofs are omitted and they are available in the full version [CMPS07].

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 $O(\text{opt} \log n)$ 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 $O(\text{opt} + \log n)$ 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 $\mathcal{S} = \{S_1, \dots, S_k\}$ of V' , such that, for each $S \in \mathcal{S}$ and for each $x, y \in S$ with $x \neq y$, $N^{\text{out}}(x) \cap N^{\text{out}}(y) = \emptyset$.

In the sequel, we assume that, given any directed graph $G(V, E)$, we have at hand a collision-free family $\mathcal{S} = \{S_1, S_2, \dots, S_k\}$ 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 $\text{PRIO-SELECT}(\mathcal{S})$ 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 , $\text{PRIO-SELECT}(\mathcal{S})$ completes fault-tolerant gossiping in G within $O(nk)$ time slots and message complexity $O(n^2)$.*

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

Protocol $\text{MULTI-SELECT}(\mathcal{S})$. 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.

Theorem 2.3. *Given a collision-free family $\mathcal{S}$ of size k for a directed graph G , $\text{MULTI-SELECT}(\mathcal{S})$ 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 $\text{MULTI-SELECT}(\mathcal{S})$ 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

Centralized Construction. Given a set V of points (i.e. nodes) in $\mathbb{R}^2$ and a range assignment $r : V \rightarrow \mathbb{R}^+$, 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., $B(v) = \{x \in \mathbb{R}^2 : d(v, x) \leq r(v)\}$. 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 $V \subseteq \mathbb{R}^2$, a function $r : V \rightarrow \mathbb{R}^+$)

```

1  Let  $X := V' = \{v \in V : N^{out}(v) \neq \emptyset\}$ ;  $\mathcal{S} := \emptyset$ ;  $i := 0$ ;
2  while  $X \neq \emptyset$  do
3       $i := i + 1$ ;  $S_i := \emptyset$ ;  $U := \emptyset$ ;  $Y := X$ ;
4      while  $Y \neq \emptyset$  do
5          Choose  $v \in Y$  such that  $r(v)$  is maximum;
6          if  $U \cap B(v) = \emptyset$  then
7               $S_i := S_i \cup \{v\}$ ;  $U := U \cup B(v)$ ;
8               $Y := Y - \{v\}$ ;
9           $\mathcal{S} := \mathcal{S} \cup \{S_i\}$ ;  $X := X - S_i$ ;
10 return  $\mathcal{S}$ .
```

It is easy to see that, by using standard data structures, the algorithm works in $O(n^2)$ time. Moreover, family $\mathcal{S}$ returned by the algorithm is collision free by construction. 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., $I(v) = \{w \in V' : B(v) \cap B(w) \neq \emptyset \text{ and } r(w) \geq r(v)\}$.

Lemma 3.1. *Family $\mathcal{S}$ has size at most $\max_{v \in V'} |I(v)|$.*

Proof. At every iteration of the external loop (line 2), 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 $S_1, S_2, \dots, S_{i-1}$ constructed in the previous iterations. For every $j = 1, 2, \dots, i-1$, S_j must contain at least one node in $I(v)$. Indeed, assume by contradiction that there exists $j \leq i-1$ such that $S_j \cap I(v) = \emptyset$. Then, for every $w \in S_j$ with $r(w) \geq r(v)$, it holds that $B(w) \cap B(v) = \emptyset$. When the algorithm selects v in line 5, the condition at line 6 is true, so v should be inserted in S_j : a contradiction. Since the sets of $\mathcal{S}$ are pairwise disjoint, the number of iterations of the external loop does not exceed $\max_{v \in V'} |I(v)|$. $\square$

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 $I(v)$.

Lemma 3.2. *For every $v \in V'$, it holds that $|B(v) \cap I(v)| \leq \Delta$.*

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., $I(v) \cap B(v) \subseteq N^{in}(v)$. $\square$

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

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

Lemma 3.3. *Let $1 < \lambda < 2$ and let $k \in \mathbb{N}$ be large enough such that $\cos \frac{2\pi}{k} \geq \lambda/2$. 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 $]r(v), \lambda r(v)] \times [\vartheta_i, \vartheta_{i+1}[$ where $\vartheta_i = \frac{2\pi i}{k}$ for $i = 0, 1, \dots, k-1$. Then, since $\cos \frac{2\pi}{k} \geq \lambda/2$, it is easy to see that the square of the maximal distance between two points in the same region is $r(v)^2 + \lambda^2 r(v)^2 - 2\lambda r(v)^2 \cos \frac{2\pi}{k} \leq r(v)^2$. For any $w \in I(v)$, it holds that $r(w) \geq r(v)$, 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. $\square$

Consider the function $g(\lambda) = \frac{\lambda^2 + 2\lambda - 1}{2\lambda^2}$ and observe that $1/2 < g(\lambda) < 1$, for any $\lambda > 1$. It is possible to prove that, if k is such that $\cos \frac{2\pi}{k} \geq g(\lambda)$, then for any $a \geq b \geq \lambda$ it holds that

$$a^2 + b^2 - 2ab \cos \frac{2\pi}{k} \leq (a - 1)^2. \quad (1)$$

We will use this fact in proving next Lemma.

Lemma 3.4. *Let $\lambda > 1$ and let $k \in \mathbb{N}$ be large enough such that $\cos \frac{2\pi}{k} \geq g(\lambda)$. 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 $[\lambda r(v), +\infty[\times [\vartheta_i, \vartheta_{i+1}[$ where $\vartheta_i = \frac{2\pi i}{k}$ for $i = 0, 1, \dots, k-1$. 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 $\varrho_x \geq \varrho_y$. Then, two constants $a, b \in \mathbb{R}$ exist with $a \geq b \geq \lambda$ such that $\varrho_x = a \cdot r(v)$ and $\varrho_y = b \cdot r(v)$. We thus get

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

where in the inequality we used the fact that x and y lie in the same region. From (1), we get $d(x, y)^2 \leq r(v)^2 (a - 1)^2 = (a \cdot r(v) - r(v))^2 = (\varrho_x - r(v))^2$. Since $x \in I(v)$, it must hold that $B(x) \cap B(v) \neq \emptyset$, so $\varrho_x - r(v) \leq r(x)$, and $d(x, y)^2 \leq r(x)^2$. 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. $\square$

Lemma 3.5. *Let $1 < \lambda < 2$ and let $k \in \mathbb{N}$ be such that $\cos \frac{2\pi}{k} \geq \max\{\frac{\lambda}{2}, g(\lambda)\}$. Then, for any $v \in V'$, it holds that $|I(v)| \leq (1 + 2k)\Delta$.*

Proof. Consider the partition of $\mathbb{R}^2$ into the three sets: $B(v)$, $C_\lambda(v)$, and the complement of $B(v) \cup C_\lambda(v)$. By combining Lemmas 3.2, 3.3, and 3.4, we get $|I(v)| \leq (1 + k + k)\Delta$. $\square$

Theorem 3.6. *Algorithm CFF returns a collision-free family $\mathcal{S}$ for $G_r(V)$ of size at most $c\Delta$, where $c \leq 33$.*

Proof. Family $\mathcal{S}$ is collision-free for $G_r(V)$ by construction. Let λ be such that $1 < \lambda < 2$. From Lemmas 3.1 and 3.5, we obtain $|\mathcal{S}| \leq \max_{v \in V'} |I(v)| \leq (1 + 2k)\Delta$, with $k \in \mathbb{N}$ such that $\cos(2\pi/k) \geq \max\{\lambda/2, g(\lambda)\}$. Then, in order to minimize k ,

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

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

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 set of squares.

Let $\gamma = \min\{d(u, v) : u, v \in V, u \neq v\}$. For any $x \in \mathbb{R}$ we define $[x]$ as the nearest integer to x . We now assume that each node knows its own position, the transmission range R and the minimum distance γ . In the following protocol, $\varepsilon > 0$ is an arbitrary small constant: we need it in order to have strict inequalities.

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

Theorem 3.7. *Family $\mathcal{S}$ is collision-free family of size $O(R^2/\gamma^2)$ for $G_r(V)$.*

Sketch of the Proof. By definition of k (line 2 of Protocol) 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 $d(u, w) \leq R$ and $d(v, w) \leq R$. By triangular inequality we get $d(u, v) \leq 2R$, but this is a contradiction, because since u, v are two different nodes in the same set of the family, by construction it must be $d(u, v) > 2R$. $\square$

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

Definition 3.8 (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 $\Gamma/\gamma \leq c\sqrt{n}$.

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

Theorem 3.9. *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 $R^2/\gamma^2 \in O(c^2 \Delta)$. $\square$

Our distributed construction also works for directed GRN where parameter R is replaced by the maximal node range $R_{\max}$. Theorem 3.7 holds with $R_{\max}$ in place of R and Theorem 3.9 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. *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 $G_r(V)$.*

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 .

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 $\Omega(n^2)$.*

Combined-Message Model

Corollary 4.3. *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.*

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 $O(n)$ that is optimal.

Theorem 4.4. *For any n , Δ and D such that $D\Delta \leq n$, 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 $\Omega(n)$.*

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 Open Questions

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] Adler, M., Scheideler, C.: Efficient communication strategies for ad hoc wireless networks. Proc. of 10th ACM SPAA, pp. 259–268 (1998)
- [ABLP89] Alon, N., Bar-Noy, A., Linial, N., Peleg, D.: A lower bound for radio broadcast. Journal of Computer and System Science 43, 290–298 (1991)
- [ABBS01] Awerbuch, B., Berenbrink, P., Brinkmann, A., Scheideler, C.: Simple routing strategies for adversarial systems. In: Proc. of 42th IEEE FOCS, pp. 158–167 (2001)
- [BG192] Bar-Yehuda, R., Goldreich, O., Itai, A.: 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 (1992)
- [BII93] Bar-Yehuda, R., Israeli, A., Itai, A.: Multiple communication in multi-hop radio networks. SICOMP 22(4), 875–887 (1993)
- [C06] Calamoneri, T.: The $L(h, k)$ -Labeling problem: Survey and Annotated Bibliography. The Computer Journal 49, 585–608 (2006)
- [CKOZ03] Calinescu, G., Kapoor, S., Olshevsky, A., Zelikovsky, A.: Network Life Time and Power Assignment in ad hoc Wireless Networks. In: Di Battista, G., Zwick, U. (eds.) ESA 2003. LNCS, vol. 2832, pp. 114–126. Springer, Heidelberg (2003)
- [CGGPR00] Chlebus, B., Gasieniec, L., Gibbons, A., Pelc, A., Rytter, W.: Deterministic broadcasting in unknown radio networks. In: Proc. of 11th ACM-SIAM SODA, pp. 861–870 (2000)
- [CGR02] Chrobak, M., Gasieniec, L., Rytter, W.: Fast Broadcasting and Gossiping in Radio Networks. Journal of Algorithms 43(2), 177–189 (2002)
- [CCJ90] Clark, B.N., Colbourn, C.J., Johnson, D.S.: Unit disk graphs. Discrete Mathematics 86, 165–177 (1990)
- [CCPRV01] Clementi, A.E.F., Crescenzi, P., Penna, P., Vocca, P.: On the complexity of minimum energy consumption broadcast subgraphs. In: Proc of 11th STACS, pp. 121–131 (2001)
- [CMS03] Clementi, A.E.F., Monti, A., Silvestri, R.: Distributed broadcast in radio networks of unknown topology. Theoretical Computer Science 302, 337–364 (2003)
- [CMS04] Clementi, A.E.F., Monti, A., Silvestri, R.: Round Robin is optimal for fault-tolerant broadcasting on wireless networks. Journal of Parallel and Distributed Computing 64(1), 89–96 (2004)

- [CMPS07] Clementi, A.E.F., Monti, A., Pasquale, F., Silvestri, R.: Optimal gossiping in directed geometric radio networks in presence of dynamical faults (2007), available at www.mat.uniroma2.it/~pasquale
- [CPS04] Clementi, A.E.F., Penna, P., Silvestri, R.: On the power assignment problem in radio networks. *MONET* (9), 125–140 (2004)
- [CR06] Czumaj, A., Rytter, W.: Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms* 60(2), 115–143 (2006)
- [DP07] Dessmark, A., Pelc, A.: Broadcasting in geometric radio networks. *Journal of Discrete Algorithms* 5, 187–201 (2007)
- [ENW00] Ephremides, A., Nguyen, G., Wieselthier, J.: On the construction of energy-efficient broadcast and multi-cast trees in wireless networks. In: *Proc. of 19th IEEE INFOCOM*, pp. 585–594 (2000)
- [GL02] Gasieniec, L., Lingas, A.: On adaptive deterministic gossiping in ad-hoc radio networks. *Information Processing Letters* 2, 89–93 (2002)
- [GPX05] Gasieniec, L., Peleg, D., Xin, Q.: Faster Communication in Known Topology Radio Networks. In: *Proc. of 24th ACM PODC*, pp. 129–137 (2005)
- [I97] Indyk, P.: Deterministic Superimposed Coding with Application to Pattern Matching. In: *Proc. of 38th IEEE FOCS*, pp. 127–136 (1997)
- [I02] Indyk, P.: Explicit constructions of selectors and related combinatorial structures, with applications. In: *Proc of 13th ACM-SIAM SODA*, pp. 697–704 (2002)
- [KKKP00] Kirousis, L., Kranakis, E., Krizanc, D., Pelc, A.: Power consumption in packet radio networks. *Theoretical Computer Science* 243, 289–305 (2000)
- [KKP98] Kranakis, E., Krizanc, D., Pelc, A.: Fault-Tolerant Broadcasting in Radio Networks. In: Bilardi, G., Pietracaprina, A., Italiano, G.F., Pucci, G. (eds.) *ESA 1998. LNCS*, vol. 1461, pp. 283–294. Springer, Heidelberg (1998)
- [P02] Pelc, A.: Broadcasting in radio networks. In: *Handbook of Wireless Networks and Mobile Computing*, pp. 509–528. John Wiley and Sons, Inc., Chichester (2002)
- [PP05] Pelc, A., Peleg, D.: Feasibility and complexity of broadcasting with random transmission failures. In: *Proc. of 24th ACM PODC*, pp. 334–341 (2005)
- [PR97] Pagani, E., Rossi, G.: Reliable Broadcast in Mobile Multihop Packet Networks. In: *Proc. of 3rd ACM-IEEE MOBICOM*, pp. 34–42 (1997)
- [R96] Rappaport, T.S.: *Wireless Communications: Principles and Practice*. Prentice-Hall, Englewood Cliffs (1996)
- [SM97] Sen, A., Malesinska, E.: On Approximation algorithms for packet radio scheduling. In: *Proc. of 35th Allerton Conference on Communication, Control and Computing*, Allerton, pp. 573–582 (1997)
- [S01] Scheideler, C.: Models and Techniques for Communication in Dynamic Networks. In: *Proc. of 18th STACS*, pp. 27–49 (2001)