Chapter 14
Data Gathering in Wireless Networks

Vincenzo Bonifaci, Ralf Klasing, Peter Korteweg, Leen Stougie, and Alberto Marchetti-Spaccamela

Abstract In this chapter, we address the problem of gathering information in a specific node of a radio network when interference constraints are present. Nodes can communicate data using a radio device; we consider a synchronous time model, where time is divided into rounds. The interference constraints limit the possibility of simultaneous data communication of nodes to the same region of the network. The survey focuses on two interference models, the general interference model and the distance-2 interference model. We survey recent complexity results and approximation algorithms for several variants of the problem. We consider several interference scenarios, the uniform and non-uniform data models, different optimization parameters, and the off-line and online settings of the problem. The objective functions we consider are the minimization of maximum completion time, maximum flow time, and average flow time.

Vincenzo Bonifaci
Dipartimento di Ingegneria Elettrica e dell'Informazione, Università degli Studi dell'Aquila, Poggio di Roio, 67040 L'Aquila, Italy, and
Dipartimento di Informatica e Sistemistica, Sapienza Università di Roma, Italy, e-mail: bonifaci@dis.uniroma1.it

Ralf Klasing
CNRS - LaBRI - Université Bordeaux 1, France, e-mail: klasing@labri.fr

Peter Korteweg
Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands, e-mail: peterkorteweg@hotmail.com

Leen Stougie
Department of Economics and Business Administration, Free University, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, and
CWI, Amsterdam, The Netherlands, e-mail: lstougie@feweb.vu.nl

Alberto Marchetti-Spaccamela
Dipartimento di Informatica e Sistemistica, Sapienza Università di Roma, Italy, e-mail: alberto@dis.uniroma1.it

A.M.C.A. Koster, X. Muñoz (eds.), *Graphs and Algorithms in Communication Networks*, Texts in Theoretical Computer Science. An EATCS Series,
DOI 10.1007/978-3-642-02250-0_14,

Key words: data gathering, packet radio network, interference, completion time, flow time, approximation algorithm, online algorithm

14.1 Introduction

The wireless gathering problem was proposed by FRANCE TELECOM in the context of providing wireless Internet access to villages [12]. The houses of a village are equipped with a computer, and the computers are interconnected through a wireless local network. To provide Internet access to each of the houses, the computers have to send (and receive) information to a gateway, or *sink node*, which connects the village with the Internet. This creates a special many-to-one information flow demand in which access to the gateway must be provided through multi-hop communication. The radio transmissions which are necessary for data communication are subject to different interference constraints. We are interested in providing interference-free data gathering, minimizing a function of the time required to do so. The underlying problem of gathering data under interference constraints is a fundamental problem in wireless communication, and is also an important building block in more complex communication problems [1, 4]. We call this class of problems *wireless gathering problems* (WGPS). In this chapter we present an overview of recent models and results that are related to WGP.

First, we briefly describe several important features which influence models for wireless gathering in radio networks. These features are all related to the use of radio signals to communicate data, which distinguishes WGP from gathering problems in wired networks.

In radio networks nodes communicate with each other using a radio transmitter. A node broadcasts data to a region surrounding its radio transmitter; the radio signals are transmitted at a certain frequency or within a certain range of frequencies, called the broadcast channel. We restrict our discussion to WGP models with a single broadcast channel.

There are two types of transmitters, based on the antennas being either *unidirectional* or *omnidirectional*. In the omnidirectional case the signal is broadcast in every direction. In this case, under ideal circumstances, the broadcast region can be described as a ball centered at the sender node. In the unidirectional case, the antenna is pointed in a specific direction; hence, the broadcast region can be described as a narrow cone centered at the sender node.

There are two models for radio communication. There is the*half-duplex* model, in which at any instant a node can either send or receive data, and there is the *full-duplex* model, in which a node can both send and receive data simultaneously. We only consider WGP in the half-duplex model. When two nodes communicate, we assume that there is a *sender node*, which has data to send, and a *receiver node* which wishes to receive the data. Data is communicated from the sender node to the receiver node, but the receiver node may use acknowledgement packets (ACKs) to confirm the data reception.

Radio signals have two important properties: fading and interference. *Fading* is the effect of radio signal loss due to physical circumstances. These circumstances are the composition of the space between the sender and receiver nodes, e.g., free space or obstacles, and the distance between the sender and receiver nodes. The strength of a radio signal is a decreasing function of the distance d between the sender node and the receiver node, and the function is in the order of d^{-2} to d^{-6} [1, 18, 37]. For a transmitter to receive data, the radio signal should be of a certain strength. As a consequence of this minimum signal strength and fading, the reachable broadcast region can be described as a closed ball centered at the sender node, where the radius of the ball is called the *communication radius*.

Interference, also called *collision*, is the effect of radio signal loss due to the fact that multiple nodes communicate simultaneously on the same broadcast channel, within the same geographical region. As with data communication, interference occurs if the radio signal is strong enough. When a node broadcasts data, its radio signal is propagated to a region surrounding this node. This interference region is a closed ball centered at the sender node, similarly to the transmission region. We call the radius of this ball the *interference radius*. Note that if a node sends data at a certain power, the interference radius is at least as large as the communication radius, but may be larger, with typical factors between 2 and 3 [22, 37]. A typical assumption is that if a node receives signals from multiple nodes, all data is lost. In some scenarios it could be possible to detect that a collision has occurred, but in this chapter we will assume that no such a collision detection mechanism is available.

The properties of fading and interference highly influence the design of wireless networks and communication algorithms. On the one hand, fading makes communication over long distances costly, and interference limits the data throughput of the network. On the other hand, fading allows multiple nodes to use the same broadcast channel simultaneously, as long as the receiver nodes are sufficiently far apart. This is known as spatial frequency reuse [33, 34].

We assume that not all nodes can communicate directly with the sink, either due to physical constraints or because such communication is too costly in terms of energy usage. We assume nodes use multi-hop communication to communicate data to the sink node. We also assume that the routing network is given. Typically, this routing network is set up via some distributed algorithm [32].

Another feature of many problems is the distinction between a *uniform* and a *non-uniform* data model [7, 13, 22, 25]. In a uniform data model one assumes that each node has the same demand for data communication, and offers the same supply of data communication. In the case of gathering problems this translates into the assumption that each node, except the sink, has an equal number of packets to send; we focus on the case where each node has exactly one packet to communicate to the sink. A non-uniform data model does not impose any restrictions on data demand. Also, most of the models studied in the literature allow the buffering of packets at each node of the network. For a study of gathering protocols in a model where buffering is not allowed see [9].

We present an overview of recent advances in wireless gathering problems. As even recent literature on wireless networks is vast, we have to limit the scope of the models that we consider. In particular, we will mostly consider the case of omnidirectional antennas, which is where the interference constraints play a key role. When we consider unidirectional antennas we will explicitly say so.

WGP consists of finding an interference-free schedule, in which packets are sent to the sink as fast as possible. We use completion times and flow times as performance measures for the schedule. A completion time model is appropriate for wireless networks which partition data reception and data communication into two phases [22, 25], while a flow time model is appropriate for wireless networks where data reception and communication occur simultaneously.

We focus on theoretical results, which consist of complexity results and worst-case analyses of algorithms using approximation theory and competitive analysis. For complexity theory we refer to Garey and Johnson [23] and Papadimitriou [31]. For a background on approximation theory see the books of Ausiello et al. [2], Hromkovič [26], and Vazirani [39]. For a background on online algorithms and competitive analysis see Borodin and El-Yaniv [17]. We do not consider any empirical studies.

The outline of this chapter is the following. In Section 14.2 we formulate the basic wireless gathering problem mathematically. In Section 14.3 we analyze the complexity of several variants of the problem. In Section 14.4 we present online algorithms for several variants and we analyze their performance using approximation theory and competitive analysis. In Section 14.5 we summarize the models and results presented in this chapter, and outline some interesting open problems.

14.2 The Mathematical Model

The communication model for the wireless gathering problem is a generalization of the classic packet radio network model [3, 4, 7, 8]. Given are a graph $G = (V, E)$ with $|V| = n$, a sink $s \in V$, and a set of packets $J = \{1, 2, \ldots, m\}$. We assume that each edge has unit length. For each pair of nodes $u, v \in V$ we define the *distance* between u and v, denoted by $d(u, v)$, as the length of a shortest path from u to v in G. We have integers d_T and d_I, for the communication radius and interference radius respectively, where naturally we have $d_I \geq d_T$. Each packet $j \in J$ has a *release node* $v_j \in V$ and a *release date* $r_j \in \mathbb{Z}_+$ at which it enters the network. We consider the case where $r_j = 0$ for all j as a special case, which we refer to as WGP without release dates. If there is a single packet j released at each node $v \in V \setminus \{s\}$, the data is said to be *uniform*; otherwise, it is said to be *non-uniform*.

We assume that time is discrete; we call a time unit a *round*. The rounds are numbered $0, 1, \ldots$. During each round a node may be *sending* a packet, be *receiving* a packet, or be *inactive*. If $d(u, v) \leq d_T$ then u can send some packet j to v during a round. If node u sends a packet j to v in some round, then the pair (u, v) is called a *call* of packet j during that round.

We consider two interference models: the general interference model and the distance-2 interference model. In the general interference model two calls (u, v) and

(u', v') *interfere* if $d(u', v) \leq d_I$ or $d(u, v') \leq d_I$; otherwise, the calls are *compatible* [7, 8, 13]. The case $d_T = d_I = 1$ is a special case [3, 4]. In the distance-2 interference model [29] one assumes unit communication radius, and two calls (u, v) and (u', v') are compatible only if $\min_{x \in \{u,v\}, y \in \{u',v'\}} d(x, y) \geq 2$; that is, nodes involved in different calls should be apart at distance at least 2, so at any given round the set of calls forms a matching in the underlying graph. We observe the following relation between the distance-2 interference model and the general interference model: each feasible distance-2 interference schedule is a feasible general interference schedule for $d_T = 1$ and $d_I = 1$, and each feasible general interference schedule for $d_T = 1$ and $d_I = 2$ is a feasible distance-2 interference schedule.

A solution to WGP is a schedule of compatible calls such that all packets are sent to the sink. In principle, each radio transmission could broadcast the same packet to multiple nodes, but in the gathering problem, having more than one copy of each packet does not help – it suffices to keep the one that will arrive first at the sink. Thus, we assume that at any time there is a unique copy of each packet. Also, we assume that packets cannot be aggregated at nodes.

Given a schedule, let v_j^t be the unique node holding packet j at the beginning of round t. The *completion time* of a packet j is $C_j = \min\{t : v_j^t = s\}$. A packet j can be sent for the first time in round r_j. The *flow time* of a packet j is $F_j = C_j - r_j$. We consider the minimization of $\max_j C_j$, called the *makespan*, the minimization of $\max_j F_j$, and the minimization of $\sum_j F_j$. We refer to WGP minimizing the maximum completion time as CMAX-WGP, to WGP minimizing the maximum flow time as FMAX-WGP, and to WGP minimizing the total or average flow time as FSUM-WGP.

14.3 Complexity and Lower Bounds

We give an overview of complexity results and lower bounds on the competitive ratio for WGP.

14.3.1 Minimizing Makespan

The first NP-hardness proof for CMAX-WGP has been given by Bermond et al. [7, 8] by means of a reduction from a satisfiability problem. Here we give a proof that gives more insight into the graph-theoretical nature of the gathering problem. It is based on a reduction from the well-known NP-hard problem of determining the chromatic number of a graph [23] (a similar proof, but within a more general interference model, has been given by Coleri [20]). The chromatic number of a graph is the minimum number of colors needed to color all vertices of the graph so that no two adjacent vertices have the same color.

CHROMATIC NUMBER
Instance: a graph G and an integer k.
Question: does G have chromatic number at most k?

Theorem 14.1. CMAX-WGP *is* NP*-hard in the general interference model.*

Proof. Consider an instance of CHROMATIC NUMBER, that is, an integer k and a graph G, with vertex set $V(G) = \{v_1, \dots, v_n\}$. Let H be the graph consisting of n isolated vertices $\{u_1, \dots, u_n\}$. We construct a graph G' with vertex set $V(G') = V(H) \cup V(G) \cup \{s\}$ and edge set $E(G') = E(G) \cup \{[u_i, v_i] : i = 1, \dots, n\} \cup \{[v_i, s] : i = 1, \dots, n\}$ (see Figure 14.1). There is one packet in each vertex of H, the sink is the vertex s, $d_T = 1$, and $d_I = 2$.

We prove the theorem by showing that if G has chromatic number at most k, then there is a schedule for the CMAX-WGP instance on G' with makespan at most $k+n$, while if G has chromatic number at least $k+1$, then every schedule for the CMAX-WGP instance on G' has makespan at least $k+n+1$. The theorem then follows since CHROMATIC NUMBER is NP-hard.

Suppose G has chromatic number at least $k+1$. We claim that at any round, in any schedule, the vertices in $V(G)$ that are acting as receivers must form an independent set in G. To see this, notice that any useful transmission to vertex v_i must come from vertex u_i. But then, if (u_i, v_i) and (u_j, v_j) are compatible calls, it must be the case that $[v_i, v_j]$ is not an edge of G; otherwise, interference would occur. Thus, at least $k+1$ rounds are needed to transmit all the packets to vertices in $V(G)$. Additionally, when a vertex v_i transmits to s, no other vertex v_j can receive a packet, because v_j is at distance 2 from v_i. So, calls of the type (v_i, s) are not compatible with calls of the type (u_j, v_j), and a total of $k+1+n$ rounds is needed to gather all packets.

Assume now that G has chromatic number at most k. In a single round, we can forward from $V(H)$ to $V(G)$ any set of packets that corresponds to an independent set of G. Thus, in k rounds we can forward each packet to a vertex of $V(G)$. The remaining n rounds can be used to collect all packets at the sink, one by one. □

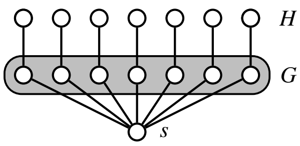

Fig. 14.1 The construction in the proof of Theorem 14.1

The complexity of uniform CMAX-WGP has been analyzed by Korteweg [28]. Kumar et al. [29] presented an inapproximability proof for packet routing in the distance-2 interference model; packet routing is a generalization of WGP in which each packet has to be sent from an arbitrary origin to an arbitrary destination. For both interference models, the NP-hardness of CMAX-WGP can be established by a reduction from the well-known problem of determining the chromatic number of a graph [23]. The following theorems can be shown to hold using proof ideas and techniques from Korteweg [28] and Kumar et al. [29].

Theorem 14.2. *Uniform* CMAX-WGP *is* NP-*hard.*

Theorem 14.3. CMAX-WGP *is* NP-*hard in the distance-2 interference model.*

We now discuss lower bounds on the approximability and competitive ratio of CMAX-WGP. The following proposition, proved by Korteweg [28], provides a lower bound on the competitive ratio of any online algorithm for CMAX-WGP. Notice that, as is usual in lower bounds for online algorithms, the bound is independent of any hardness assumption.

Proposition 14.1. *No online algorithm for* CMAX-WGP *is better than 7/6-competitive, even if* $d_I = d_T$.

Proof. We give the proof for $d_I = d_T = 1$, the generalization to larger values being straightforward. Consider the graph depicted in Figure 14.2. The adversary releases packet 1 at u at time 0. Observe that for any algorithm that does not send packet 1 in the first round the lemma trivially holds. We assess both deterministic and randomized algorithms by applying Yao's minimax principle [40]. The adversary releases a second packet in round 1 either at u_3 or u_4, each with probability 1/2. Now, the expected number of rounds for any algorithm that sends a packet in the first round is $1/2 \cdot 4 + 1/2 \cdot 3$. In the optimal schedule packet 1 is sent to u_2 (u_1) in the first round if the adversary releases a packet at u_3 (u_4), which yields a makespan of 3. □

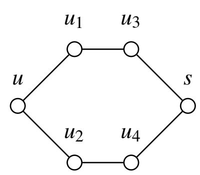

Fig. 14.2 No online algorithm for CMAX-WGP is better than 7/6-competitive ($d_T = d_I = 1$)

It is interesting to note that the example of Proposition 14.1 contains three packets, and there are no known constant lower bound results which hold for instances with an arbitrary number of packets.

Bonifaci et al. provided lower bounds on the approximability of shortest paths following algorithms [13]. A *shortest paths following algorithm* is an algorithm where each packet is sent over some shortest path towards the sink. This is a natural class of algorithms for routing problems, and in case of packet routing without interference it has been demonstrated that for some well-known greedy algorithms there is a gap between the completion times of routing over arbitrary paths and over shortest paths, in favor of routing over shortest paths [19]. The algorithms for WGP that we describe in the next section are shortest paths following. Following [13], for such algorithms we present a lower bound of $4 - 16/(m+4)$ on their approximation ratio for solving CMAX-WGP on m packets using a shortest paths following algorithm, in the case where $d_T = 1$ and $d_I = 2$.

Consider Figure 14.3. The nodes $u_1, \ldots, u_m$ have one packet each. Any shortest paths following algorithm sends all packets via u, yielding $\max_j C_j = 4m$. There is a solution with no packet passing u that implies $\max_j C_j^* \leq 4 + m$. The example can easily be extended for arbitrary $d_T, d_I = 2d_T$, such that no shortest path following algorithm is better than 4-approximate. In Section 14.4 we discuss a matching upper bound.

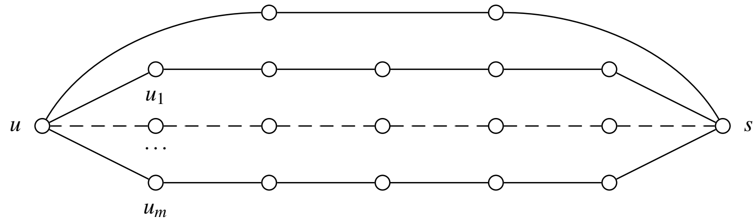

Fig. 14.3 No shortest paths following algorithm is better than 4-approximate for CMAX-WGP ($d_T = 1, d_I = 2$)

14.3.2 Minimizing Flow Times

Bonifaci et al. [15, 16] considered the problems FMAX-WGP and FSUM-WGP. For these versions even stronger results are possible than the one of Theorem 14.1. We present the result for FMAX-WGP.

The lower bound is based on the *induced matching* problem. A matching M in a graph G is an *induced matching* if no two edges in M are joined by an edge of G. The following rather straightforward relation between compatible calls in a bipartite graph and induced matchings will be crucial in the proof.

Proposition 14.2. *Let $G = (U, V, E)$ be a bipartite graph with node sets (U, V) and edge set E. Then, a set $M \subseteq E$ is an induced matching if and only if the calls corresponding to edges of M, directed from U to V, are all pairwise compatible, assuming $d_T = d_I = 1$.*

INDUCED BIPARTITE MATCHING (IBM)
Instance: a bipartite graph G and an integer k.
Question: does G have an induced matching of size at least k?

The proof by Bonifaci et al. uses the fact that the optimization version of IBM is hard to approximate: there exists an $\alpha > 1$ such that it is NP-hard to distinguish between graphs with induced matchings of size k and graphs in which all induced matchings are of size at most k/α. The current best bound for α is 6600/6599 [21].

Theorem 14.4. *Unless* $\mathsf{P} = \mathsf{NP}$*, no polynomial-time algorithm can have approximation ratio better than* $\Omega(m^{1/3})$ *for* FMAX-WGP *in the general interference model, even when* $d_T = d_I = 1$.

Proof. Let (G,k) be an instance of IBM, $G = (U,V,E)$. We construct a four-layer network with a unique source o (first layer), a clique on U and a clique on V (middle layers), and a sink s (last layer). Source o is adjacent to each node in U, and s to each node in V. The edges between U and V are the same as in G (see Figure 14.4). We set $d_T = d_I = 1$.

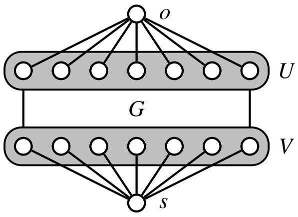

Fig. 14.4 The construction in the proof of Theorem 14.4

The FMAX-WGP instance consists of $m = (1-1/\alpha)^{-1}(1+k/\alpha)(2k+1)k = \Theta(k^3)$ packets with origin o. They are divided into m/k groups of size k. Each packet in the hth group has release date $(k+1)h$, $h = 0, \ldots, m/k-1$. Rounds $(k+1)h$ to $(k+1)(h+1)-1$ together are a *phase*.

We prove that if G has an induced matching of size k, there is a gathering schedule of cost $2k+1$, while if G has no induced matching of size more than k/α, every schedule has cost at least $(2k+1)k = (2k+1)\Theta(m^{1/3})$. The theorem then follows directly.

Assume G has an induced matching M of size k, say (u_i, v_i), $i = 0 \ldots k-1$. Then consider the following gathering schedule. In each phase, the k new packets at o are transmitted, necessarily one by one, to layer U while old packets at layer V (if any) are absorbed at the sink; then, in a *single* round, the k new packets move from U to V via the matching edges. More precisely, each phase can be scheduled in $k+1$ rounds as follows:

1. for $i = 0, \ldots, k-1$ execute in round i the two calls (o, u_i) and $(v_{i+1 \bmod k}, s)$;
2. in round k, execute simultaneously all the calls (u_i, v_i), $i = 0, \ldots, k-1$.

The maximum flow time of the schedule is $2k+1$, as a packet released in phase h reaches the sink before the end of phase $h+1$.

In the other direction, assume that each induced matching of G is of size at most k/α. By Proposition 14.2, at most k/α calls can be scheduled in any round from layer U to layer V. We ignore potential interference between calls from o to U and calls from V to s; doing so may decrease the cost of a schedule. As a consequence, we can assume that each packet follows a shortest path from o to s. Notice however that, due to the cliques on the layers U and V, no call from U to V is compatible with a call from o to U, or with a call from V to s.

Let m_o and m_U be the number of packets at o and U, respectively, at the beginning of a given phase. Also, let $\beta = 1 + k/\alpha$. We associate with the phase a potential value $\psi = \beta m_o + m_U$, and we show that at the end of the phase the potential will have increased proportionally to k. Let c_o and c_U denote the number of calls from o to U and from U to V, respectively, during the phase. Since a phase consists of $k+1$ rounds, and in each round at most k/α calls are scheduled from U to V, we have $c_o + c_U/(k/\alpha) \leq k+1$, or, equivalently since $k/\alpha = \beta - 1$,

$$(\beta - 1)c_o + c_U \leq (\beta - 1)(k+1). \tag{14.1}$$

If m'_o, m'_U are the number of packets at o and U at the beginning of the next phase, and $\psi' = \beta m'_o + m'_U$ is the new potential, we have

$$\begin{aligned}
m'_o &= m_o + k - c_o \\
m'_U &= m_U + c_o - c_U \\
\psi' - \psi &= \beta(m'_o - m_o) + m'_U - m_U \\
&= \beta(k - c_o) + c_o - c_U \\
&= \beta k - (\beta - 1)c_o - c_U \\
&\geq \beta k - (\beta - 1)(k+1) \\
&= k - (\beta - 1) \\
&= (1 - 1/\alpha)k
\end{aligned}$$

where the inequality uses (14.1).

Thus, consider the situation after m/k phases. The potential has become at least $\Psi = (1 - 1/\alpha)m$. By definition of the potential, this implies that at least $\Psi/\beta = (1 - 1/\alpha)(1 + k/\alpha)^{-1}m = (2k+1)k$ packets reside at either o or U; in particular, they have been released but not yet absorbed at the sink. Since the sink cannot receive more than one packet per round, this clearly implies a maximum flow time of $(2k+1)k = (2k+1)\Theta(m^{1/3})$ for one of these packets. □

A similar construction shows that the same problem with minimization of total flow time FSUM-WGP cannot be approximated within a ratio better than $\Omega(m^{1-\varepsilon})$ for any $0 < \varepsilon < 1$ [15]. We also notice that a similar instance as that used in Section 14.3.1 constructed for proving inapproximability of shortest paths following algorithms for CMAX-WGP can be constructed here to prove that shortest paths following algorithms cannot approximate optimal solutions of FMAX-WGP and FSUM-WGP within a ratio better than $\Omega(m)$.

For the distributed model, Bonifaci et al. [14] provided lower bounds for FMAX-WGP which do not depend on the assumption $\mathsf{P} \neq \mathsf{NP}$. They consider a scenario in which the network is partitioned into layers based on distance to the sink. They assume that interference conflicts between transmissions from one layer of the tree to the next are resolved randomly: whenever several transmissions from a layer occur in the same round, only a uniformly chosen one succeeds; this is called the *random selection model*. This assumption seems natural for distributed algorithms, as they have no simple means of coordinating the transmitting nodes (or more precisely, coordinating the transmitting nodes is as hard as the original communication task). For distributed algorithms following a random selection model they present the following lower bound.

Theorem 14.5. *In the random selection model the approximation ratio of any algorithm for* FMAX-WGP *is at least* $\Omega(\log m)$.

In fact, Bonifaci et al. [14] argue that even resource augmentation using speed as a resource is not likely to improve this lower bound. The reason is that the lower bound is due to an adversarial selection of which packet to advance; and the probability of obtaining such a selection depends on the number of packets, and not on the speed of the algorithm.

14.4 Online Algorithms

14.4.1 Minimizing Makespan

14.4.1.1 Omnidirectional Antennas

Several authors have presented centralized online algorithms for omnidirectional WGP. The first algorithm, PIPELINE, was presented by Bermond et al. [7, 8]. The algorithm was analyzed in an off-line context, but can be implemented in an online setting. The idea of the algorithm is to pipeline packets towards the sink by partitioning the graph into intervals. The lengths of the intervals are chosen such that packets can advance in parallel without interfering with each other.

First, we introduce some notation to facilitate the exposition of the algorithm. An important concept used in this and other algorithms is that of critical radius. The *critical radius* R^* is the greatest integer R such that no two nodes at distance at most R from s can receive a packet in the same round. It is not hard to show that $R^* \geq \left\lfloor \frac{d_I - d_T}{2} \right\rfloor$ (see, for example, [7, 8]). The *critical region* is the ball centered at s of radius R^*. Thus, at any round at most one node in the critical region can receive a packet. We define $K^* = \left\lceil \frac{R^*+1}{d_T} \right\rceil \geq 1$ and $K = 1 + \left\lceil \frac{d_I+1}{d_T} \right\rceil$. Roughly stated, K gives an upper bound on the number of rounds during which a packet needs to be forwarded before a new packet can be safely forwarded from the same origin over the same path, while K^* gives a lower bound on the number of rounds during which a packet has to move inside the critical region, assuming it starts outside. We also let $K_0 = 1 + \left\lceil \frac{R^*}{d_T} \right\rceil$, $\mathrm{Rad} = \max_{u \in V} d(u,s)$, and $L = 1 + \left\lceil \frac{\mathrm{Rad} - K_0 d_T}{K d_T} \right\rceil$.

The algorithm partitions the set of distances to the sink $[1, \mathrm{Rad}]$ into L intervals $I_0, \ldots, I_{L-1}$. These are defined by $I_0 = [1, K_0 d_T]$ and, for $i = 1, \ldots, L-1$, $I_i = [K_0 d_T + 1 + (i-1)K d_T, K_0 d_T + iK d_T]$.

Additionally, each I_i is split into *areas* of length d_T, so I_0 is split into K_0 areas $I_0^j = [K_0 d_T + 1 - j d_T, K_0 d_T - (j-1)d_T]$, $j = 1, \ldots, K_0$; and I_i, $i = 1, \ldots, L-1$ is split into K areas $I_i^j = [K_0 d_T + 1 + iK d_T - j d_T, K_0 d_T + iK d_T - (j-1)d_T]$, $j = 1,$

…, K. We denote the set of vertices whose distance is in I_i (respectively I_i^j) by V_i (respectively V_i^j). Figure 14.5 shows a partition with $K = 4, K_0 = 3, d_T = 2$.

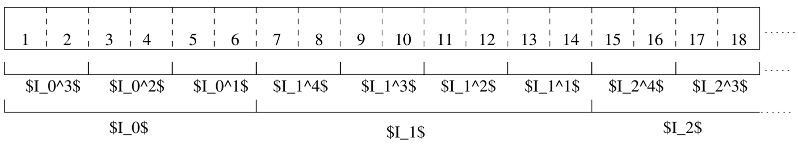

Fig. 14.5 Partitioning of distance intervals for $K = 4, K_0 = 3, d_T = 2$

We are now in position to describe the algorithm (Algorithm 14.1).

Algorithm 14.1 PIPELINE

The algorithm works in phases. Each phase, except possibly the last, consists of K rounds t_j, $j = 1, \ldots, K$. The algorithm uses the concepts of intervals and areas to construct a set of feasible calls in each round.

for each phase **do**
  for each round t_j, $j = 1, 2, \ldots, K$ **do**
    Select in each interval I_i a vertex u_i^j in V_i^j with an available packet to transmit (if such a vertex exists). Vertex u_i^j calls the closest vertex in the preceding area, i.e., if $d(u_i^j, s) = K_0 d_T + 1 + iKd_T - jd_T + \alpha$ for some $0 \leq \alpha < d_T$, then u_i^j calls a vertex v such that $d(v,s) = K_0 d_T + iKd_T - jd_T$. This means that if $i = 0$ and $j < K_0$ (or $i > 0$ and $j < K$) then $v \in V_i^{j+1}$, if $i > 0$ and $j = K$ then $v \in V_{i-1}^1$, and if $i = 0$ and $j = K_0$ then $v = s$.

We claim that PIPELINE creates a feasible schedule for WGP. First, let us show that for any round the calls scheduled by PIPELINE are all pairwise compatible. Indeed, consider two calls $(u,v) \neq (u',v')$ of the same round t_j. Then $d(u,s) = K_0 d_T + 1 + iKd_T - jd_T + \alpha$, for some $i \geq 0, 0 \leq \alpha < d_T$, and $d(v',s) = K_0 d_T + i'Kd_T - jd_T$ for some $i' \neq i$ (as $v \neq v'$). Therefore, $d(u,v') \geq |(i'-i)Kd_T - 1 - \alpha| \geq d_I + d_T - \alpha \geq d_I + 1$. Similarly one can show $d(u',v) \geq d_I + 1$, and the calls are compatible. To see why the algorithm delivers all the packets, observe that after a phase of K rounds, the protocol ensures that if a vertex of V_i contains a packet, then the last vertex of V_{i-1} has received a new packet.

To illustrate PIPELINE we show one phase of the algorithm in Figure 14.6.

Bonifaci et al. [13] presented a class of online centralized algorithms for CMAX-WGP, called PRIORITY GREEDY. In a PRIORITY GREEDY algorithm each packet is assigned a unique priority based on some algorithm-specific rules, and the priority ordering does not change over time. Then, in each round, packets are considered in order of decreasing priority and are sent towards the sink as far as possible while avoiding interference with higher priority packets. Thus, the schedule output by the algorithm is feasible by construction.

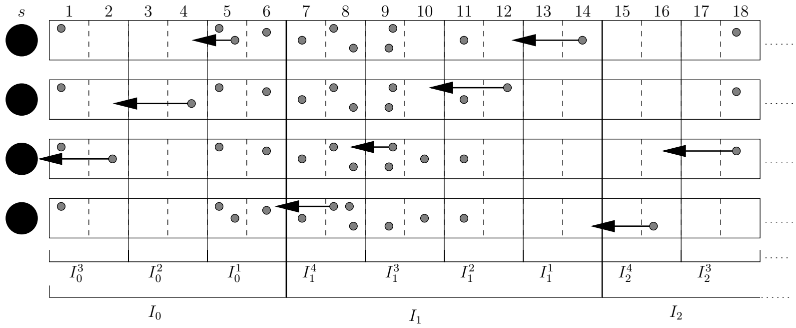

Fig. 14.6 A phase of PIPELINE, consisting of $K = 4$ rounds. Here, packets are represented as small balls. Notice that packets in the same cell are at the same distance from the sink, but they can be in different vertices

Algorithm 14.2 PRIORITY GREEDY

for each round $t = 0, 1, 2, \ldots$ **do**
Consider the available packets in order of decreasing priority, and send each packet as far as possible along a shortest path from its current node to the sink, without causing interference with any higher-priority packet.

Both Bermond et al. and Bonifaci et al. use similar concepts to derive upper bounds on their algorithms, as well as a lower bound on the makespan of an off-line optimal solution [7, 8, 13].

The lower bound on the completion time of any schedule is based on the observation that at most one packet can be sent from a node within the critical region. Let $\delta_j = \lceil \frac{d(v_j,s)}{d_T} \rceil$, the minimum number of calls required for packet j to reach s. Define also $\pi_j = \min\{\delta_j, K^*\}$ and $R_j = r_j + \delta_j - \pi_j$. Informally, π_j gives the number of rounds during which packet j has to move inside the critical region (irrespective of whether it originated inside or outside of it); R_j is the first possible time at which packet j can reach the border of the critical region. The following bound on the cost of an optimal solution can be proved by considering only the processing that has to be done inside the critical region [13].

Lemma 14.1. *Let $S \subseteq J$ be a nonempty set of packets, and let C_i^* denote the completion time of packet i in some feasible schedule. Then there is $k \in S$ such that $\max_{i \in S} C_i^* \geq R_k + \sum_{i \in S} \pi_i$.*

We sketch the idea behind the upper bound on the completion times; the sketch is based on the upper bound proof of the PRIORITY GREEDY algorithm. The idea is that if a packet is delayed, i.e., it is not sent as far as possible in each round until it reaches the sink, then this packet must be close to some other packet that is sent in that round. As a result we can provide a bound on the completion time of a packet by relating it to the completion time of another packet that delays this packet. This provides an upper bound on the makespan of a set of packets. Formally, we say that packet j is *blocked* in round t if $t \geq r_j$ but j is not sent over distance d_T in round

t. Note that in a PRIORITY GREEDY algorithm a packet can only be blocked due to interference with a higher priority packet. We define the following *blocking relation* on a PRIORITY GREEDY schedule: $k \prec j$ if in the last round in which j is blocked, k is the packet closest to j that is sent in that round and has a priority higher than j (ties broken arbitrarily). The blocking relation induces a directed graph $F = (J,A)$ on the packet set J with an arc (k,j) for each $k,j \in J$ such that $k \prec j$. Observe that, for any PRIORITY GREEDY schedule, F is a directed forest and the root of each tree of F is a packet which is never blocked. For each j, let $T(j) \subseteq F$ be the tree of F containing j, $b(j) \in J$ be the root of $T(j)$, and $P(j)$ be the set of packets along the path in F from $b(j)$ to j.

Lemma 14.2. *For each packet $j \in J$ in a* PRIORITY GREEDY *schedule, $C_j \leq R_{b(j)} + (K/K^*) \cdot \sum_{i \in P(j)} \pi_i$.*

Bonifaci et al. [13] considered a particular PRIORITY GREEDY algorithm called RPG in which packet j has higher priority than packet k if $R_j < R_k$ (ties broken arbitrarily). Combining Lemmas 14.1 and 14.2 they proved the following theorem.

Theorem 14.6. RPG *is K/K^*-competitive for* CMAX-WGP.

Similarly, Bermond et al. [7] presented the following theorem.

Theorem 14.7. PIPELINE *is K/K^*-competitive for* CMAX-WGP *without release dates.*

The exact ratio depends on d_T and d_I, but is always bounded: it is straightforward to verify that $2 \leq K/K^* \leq 4$ for all d_T and d_I, while $K/K^* \leq 3$ for $d_I = d_T$. Similarly, Korteweg [28] proved that PRIORITY GREEDY is $(K/K^* + 1)$-competitive for *any* fixed priority on the packets, using Lemmas 14.1 and 14.2. An interesting open problem is whether there exists a polynomial-time approximation scheme, or a $(1 + \varepsilon)$-approximation algorithm for general graphs for small values of $\varepsilon > 0$.

Notice that the algorithms PIPELINE and PRIORITY GREEDY can be implemented using only local information. Namely, it suffices that a node is informed about the state of nodes within distance $d_I + 1$.

Kumar et al. [29] presented a decentralized algorithm for packet routing under the distance-2 interference model. The authors presented an $\mathcal{O}(\Delta \log^2 n)$-approximation algorithm, where Δ is the maximum graph degree and n is the number of nodes. Their algorithm assumes that each node knows upper bounds on the maximum number of packets per node and the network diameter. The first assumption seems restrictive from a practical point of view, where packets arrive online over time. The algorithm proceeds in phases, and at the start of each phase nodes communicate with nodes up to a distance 3 to determine interference-free schedules for the round in the next phase. As such the algorithm, like those discussed above, is decentralized, but not distributed in the sense that nodes use information about neighboring nodes.

Bar-Yehuda et al. [4] considered a distributed algorithm for CMAX-WGP in the special case where $d_T = d_I = 1$ and there are no release dates. We refer to their algorithm as DISTRIBUTED GREEDY (DG). The idea behind DG is the following.

To reduce interference between nodes, DG partitions nodes into layers, and assigns a label to nodes in a layer. A *layer* is a set of all nodes at the same distance from the sink. A node at distance d from the sink is assigned *label* $d \mod 3$. Each node can be either *active* during a round or *inactive*; only active nodes will transmit a packet. A node will not be active if its packet buffer is empty.

DG uses a procedure to establish communication from a set of active nodes. The procedure, first introduced and studied by Bar-Yehuda, Goldreich, and Itai [3], is called DECAY and requires $2\log\Delta$ rounds; the time needed for a single execution of the procedure is called a *phase*.

Algorithm 14.3 DECAY(u,v)

for $j = 1,2,\ldots,2\log\Delta$ **do**
  u sends to v the oldest packet from its buffer;
  u deactivates itself for the rest of the phase with probability 1/2.

In fact, the original description does not describe which packet v to advance from the buffer, because for the analysis of completion times the choice of this packet is not relevant. For flow times the choice can be relevant; hence, we choose to advance the oldest packet. We now present the description of the DISTRIBUTED GREEDY algorithm (Algorithm 14.4).

Algorithm 14.4 DISTRIBUTED GREEDY (DG)

for each phase $k = 1,2,\ldots$ **do**
  Activate each node with label $k \mod 3$ that has a nonempty packet buffer;
  Execute DECAY$(u,\text{parent}(u))$ in parallel for each active node u.

Although the algorithm does not model acknowledgement of packets explicitly, it is easy to include them, e.g., by doubling the number of rounds, having communication in odd rounds and acknowledgements in even rounds, as observed by Bar-Yehuda et al. Using this, we can assume that successful receipt of a packet (by the parent of the sending node in the communication tree) is acknowledged immediately. Only at that time does it get removed from the sender's buffer.

By the transmission protocol in DG, where in phase k only nodes of layer $k \mod 3$ transmit, if two nodes transmit, then either they are at the same layer or they are at least distance 3 apart. Hence, in DG two nodes can only interfere if both sender nodes are in the same layer.

A *superphase* consists of three consecutive phases. Another important ingredient in the analysis of DG is the following, proved by Bar-Yehuda et al. [4].

Theorem 14.8. *Let i be any layer of the tree containing some packet at the beginning of a superphase. There is probability at least* $= e^{-1}(1-e^{-1})$ *that during this superphase* DG *sends successfully a packet from at least one node u in layer i to the parent node of u in the communication tree.*

This theorem shows that, during a superphase, each nonempty layer forwards a packet with probability   to the following layer. Notice however that there is no guarantee on which particular packet is advanced. The use of superphases and labels, i.e., a synchronous model, is essential to the proof of Theorem 14.8. If the DECAY procedure is applied in an asynchronous model, it is not clear whether a similar constant probability   is attainable.

Theorem 14.8 suffices to bound the completion time of packets in a schedule constructed by DG.

Theorem 14.9. DISTRIBUTED GREEDY *is in expectation* $\mathscr{O}(\log \Delta)$*-competitive for* CMAX-WGP *without release dates, and* $d_I = d_T = 1$.

14.4.1.2 Special Topologies

The hardness results for CMAX-WGP on general graphs of Section 14.3 motivate the study of specific topologies, such as the path, balanced stars, and the two-dimensional grid. With the additional assumption that the data is uniform (every node holds exactly one packet) it is possible to provide algorithms whose performance differs only by an additive constant from the theoretical minimum. We are not aware of studies for specific topologies in the case of non-uniform data (in particular, it is unknown whether optimal polynomial-time algorithms are possible for path or tree topologies), although it is certainly possible to at least improve the approximation guarantees in this setting.

As an example of the uniform data model on a specific topology, Figure 14.7 shows an optimal gathering schedule using 18 rounds for a path of seven vertices (each having one data packet), with $d_T = 1$, $d_I = 2$, and $s = 0$. The schedule has a regular structure and this regularity can be exploited to give a general algorithm for paths whose cost differs by the optimal one only by an additive constant term (though the constant may depend on d_T and d_I).

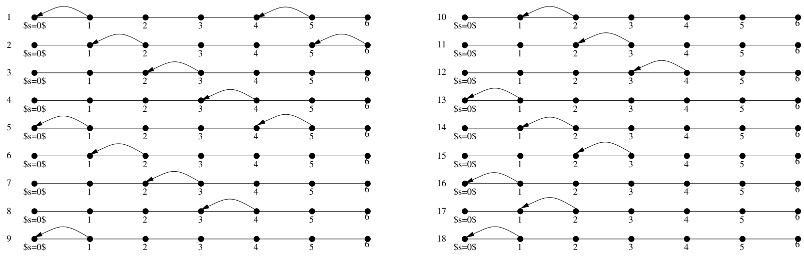

Fig. 14.7 A gathering schedule in the path when $d_T = 1$, $d_I = 2$, and every vertex has one packet to send to the sink $s = 0$

In fact, the uniform model has first been studied in the case of specific graph topologies for specific values of d_I and d_T. In particular, the case $d_T = 1$ was studied in [5] for the case where the graph is a path, and in [10] for the case where the graph is the two-dimensional square grid. An optimal algorithm for trees when $d_T = d_I = 1$ is given by Bermond and Yu [11].

Bermond et al. [6, 30] consider the uniform model for paths and grids, for any value of d_T. Even though their algorithms do not match the lower bounds, they are again larger by an additive constant that depends only on d_I and d_T. The authors also study the case of stars and show that the general lower bound of [7, 8] is tight up to a constant that does not depend on the size of the network. Table 14.1 shows the main results of Bermond et al. [6, 30]. The notation is the following:

- LB (UB) is a lower bound (upper bound) on the number of rounds for gathering in the corresponding topology.
- P_n is the path with n vertices $0, 1, \ldots, n-1$. Vertex i is adjacent to vertex $i+1$ for any $i = 1, \ldots, n-2$. Therefore, the sink s is simply an integer such that $0 \leq s \leq n-1$.
- $S_{K,l}$ is the balanced spider graph with K branches. $S_{K,l}$ consists of K copies of P_l (called branches) sharing a common extreme, the sink s.
- $G^2(p,q)$ is the two-dimensional grid, i.e., the graph $G = (V,E)$ where $V = \{(i,j) : -p \leq i \leq p, -q \leq j \leq q\}$. So $n = (2p+1)(2q+1)$, and (x,y) and (x',y') are connected when $|x-x'| + |y-y'| = 1$. We assume that $p,q \geq d_I + d_T + 1$ and $s = (0,0)$.

In Table 14.1, $\mathscr{O}(1)$ is used to denote a constant that may depend on d_I and d_T but not on the size n of the network.

Table 14.1 Approximation results for gathering in specific topologies

Topology	LB	UB
P_n	$\frac{d_I+d_T+1}{d_T}\max[s,n-s] - \mathscr{O}(1)$	$\frac{d_I+d_T+1}{d_T}\max[s,n-s] + \mathscr{O}(1)$
$S_{K,l}$, $\lfloor d_I/d_T \rfloor$ odd	$\frac{1}{2}(1+\lfloor d_I/d_T \rfloor)n - \mathscr{O}(1)$	$\frac{1}{2}(1+\lfloor d_I/d_T \rfloor)n + \mathscr{O}(1)$
$S_{K,l}$, $\lfloor d_I/d_T \rfloor$ even	$\frac{1}{2}\lfloor d_I/d_T \rfloor n + \frac{n}{K} - \mathscr{O}(1)$	$\frac{1}{2}\lfloor d_I/d_T \rfloor n + \frac{n}{K} + \mathscr{O}(1)$
$G^2(p,q)$, $\lfloor d_I/d_T \rfloor$ odd	$\frac{1}{2}(1+\lfloor d_I/d_T \rfloor)n - \mathscr{O}(1)$	$\frac{1}{2}(1+\lfloor d_I/d_T \rfloor)n + \mathscr{O}(1)$
$G^2(p,q)$, $\lfloor d_I/d_T \rfloor$ even	$\frac{1}{2}\lfloor d_I/d_T \rfloor n + \frac{n}{4} - \mathscr{O}(1)$	$\frac{1}{2}\lfloor d_I/d_T \rfloor n + \frac{n}{4} + \mathscr{O}(1)$

14.4.1.3 Unidirectional Antennas

We discuss here some results for gathering problems with unidirectional antennas. Florens et al. [22] study makespan minimization in a model where each each node is equipped with directional antennas and has no buffering capacity. Furthermore, it is assumed that a node cannot receive and send simultaneously, that the communication radius is 1, and that there is no interference but each node can only receive one message at a time. Under these assumptions, Florens et al. give optimal (polynomial-time) gathering algorithms for path and tree networks. Their work has been extended to general graphs in the uniform case by Gargano and Rescigno [25]. Other results for specific topologies are discussed by Revah and Segal [35, 36] and Segal and Yedidsion [38].

A discussion of some algorithmic and graph-theoretic problems related to wireless data gathering with minimum makespan is contained in [24]. Finally, another related model can be found in Klasing et al. [27], where the authors study the case in which steady-state flow demands between each pair of nodes have to be satisfied.

14.4.2 Minimizing Flow Times

Most literature on gathering problems focuses on minimizing completion times. In this subsection we highlight some results on minimizing flow times. First, we consider the centralized model. Bonifaci et al. [16] analyzed FMAX-WGP in the general interference model. For this version they analyzed the performance of a particular PRIORITY GREEDY algorithm. Because it follows from Theorem 14.4 that there is no constant approximation algorithm for this problem, unless $\mathsf{P} = \mathsf{NP}$, they used resource augmentation to analyze the quality of the algorithm. They study a variant of PRIORITY GREEDY which orders packets based on release dates, i.e., packet j precedes k if $r_j \leq r_k$; ties ($r_j = r_k$) are broken arbitrarily. They call this variant FIFO after the well known first-in-first-out algorithm in scheduling theory, although in this case the term FIFO refers to the priority ordering; observe that the first packet in the system does not have to arrive earliest at the sink using FIFO. They use FIFO as a sub-routine in an algorithm which can be used in a resource augmentation setting based on speed. The algorithm is the so-called σ-speed algorithm, where the parameter σ denotes the ratio between the clock speed of the algorithm and the clock speed of the optimal solution to which the algorithm is compared. The algorithm is the following (Algorithm 14.5).

Algorithm 14.5 σ-FIFO

1. Create a new instance $\mathscr{I}'$ by multiplying release dates: $r'_j = \sigma r_j$;
2. Run FIFO on $\mathscr{I}'$;
3. Speed up the schedule thus obtained by a factor of σ.

The schedule constructed by σ-FIFO is a feasible σ-speed solution to the original problem because of step 1. Bonifaci et al. [16] prove that this algorithm is optimal for some σ which depends on K and K^*, but is never larger than 5.

Theorem 14.10. *For* $\sigma \geq K/K^* + 1$, σ-FIFO *is a* σ*-speed optimal algorithm for* FMAX-WGP *in the general interference model.*

In [15] complementary and indeed similar results have been obtained for the problem with the average completion time as objective, FSUM-WGP.

For the distributed model, for WGP minimizing flow times in the general interference model, the performance of algorithm DG is studied by Bonifaci et al. [14]. Again, the performance of the algorithm was studied in a resource augmentation setting with an increase in speed of factor σ, similarly to the centralized model. We refer to this version as σ-speed DG, although the algorithm is identical to DISTRIBUTED GREEDY. Also, they focused on minimizing average flow times instead of minimizing maximum flow times. The motivation for this objective over the objective minimizing maximum flow times is based on the proof of the lower bound of Theorem 14.5. As described above the proof indicates that for a general class of distributed algorithms, i.e., algorithms which base decisions on random selection, it is rather unlikely that there exists a constant competitive algorithm for this problem, even if one allows resource augmentation using extra speed. The same authors presented the following positive result.

Theorem 14.11. *Let* $0 < \varepsilon \leq 1$ *and* $\sigma = 6\ {}^{-1} \cdot \log \Delta \cdot \ln(\delta/\varepsilon)$. *Then* σ*-speed* DG *is in expectation* $(1 + 3\varepsilon)$*-competitive when minimizing the average flow time.*

14.5 Conclusion

The chapter surveys recent complexity results and approximation algorithms for several variants of the wireless gathering problem. It considers different interference models, the uniform and non-uniform data models, different optimization parameters, and the off-line and online settings of the problem.

Many interesting directions of future work arise from the considered problems. These include the attempt to close the existing gaps between the upper and lower bounds for the specific problems. Where good solutions on general graphs are not possible or not available, the focus on graph classes that are of interest from a practical point of view is of high importance. In the non-uniform data model many important questions are still to be resolved. Also, more work remains to be done on unidirectional antennas with or without buffering capabilities at the nodes. Finally, especially from a practical perspective, the study of distributed algorithms needs to be further intensified.

Acknowledgements This work was supported by EU COST action 293 – Graphs and Algorithms in Communication Networks (GRAAL). Research of the authors was partially supported by the Future and Emerging Technologies Unit of EC (IST priority – 6th FP) under contract no. FP6-021235-2 (project ARRIVAL), by the Dutch BSIK-BRICKS project, by EU ICT-FET 215270 FRONTS, by MIUR-FIRB Italy-Israel project RBIN047MH9, by ANR-project ALADDIN (France), by the project CEPAGE of INRIA (France), and by European project COST Action 295 Dynamic Communication Networks (DYNAMO).

References

1. Akyildiz, I. F., Su, W., Sankarasubramaniam, Y., Cayirci, E.: Wireless sensor networks: a survey. Computer Networks **38**(4), 393–422 (2002)
2. Ausiello, G., Protasi, M., Marchetti-Spaccamela, A., Gambosi, G., Crescenzi, P., Kann, V.: Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties. Springer (1999)
3. 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 Sciences **45**(1), 104–126 (1992)
4. Bar-Yehuda, R., Israeli, A., Itai, A.: Multiple communication in multihop radio networks. SIAM Journal on Computing **22**(4), 875–887 (1993)
5. Bermond, J. C., Corrêa, R. C., Yu, M. L.: Gathering algorithms on paths under interference constraints. In: Proceedings of the 6th Italian Conference Algorithms and Complexity, *Lecture Notes in Computer Science*, vol. 3998, pp. 115–126. Springer (2006). Full version to appear in *Discrete Mathematics*
6. Bermond, J. C., Galtier, J., Klasing, R., Morales, N., Pérennes, S.: Gathering in specific radio networks. In: 8èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications, pp. 85–88. Trégastel, France (2006)
7. Bermond, J. C., Galtier, J., Klasing, R., Morales, N., Pérennes, S.: Hardness and approximation of gathering in static radio networks. Parallel Processing Letters **16**(2), 165–183 (2006)
8. Bermond, J. C., Galtier, J., Klasing, R., Morales, N., Pérennes, S.: Hardness and approximation of gathering in static radio networks. In: Proc. Foundation and Algorithms for Wireless Networking, pp. 75–79. Pisa, Italy (2006)
9. Bermond, J. C., Gargano, L., Rescigno, A.: Gathering with minimum delay in tree sensor networks. In: Proceedings of the 15th International Colloquium on Structural Information and Communication Complexity, *Lecture Notes in Computer Science*, vol. 5058, pp. 262–276. Springer (2008)
10. Bermond, J. C., Peters, J.: Efficient gathering in radio grids with interference. In: Septiémes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications, pp. 103–106. Presqu'île de Giens, France (2005)
11. Bermond, J. C., Yu, M. L.: Optimal gathering algorithms in multi-hop radio tree networks with interferences. In: Proc. of the Int. Conf. on Ad-Hoc, Mobile, and Wireless Networks, pp. 204–217 (2008)
12. Bertin, P., Bresse, J. F., Sage, B.: Accès haut débit en zone rurale: une solution "ad hoc". France Telecom R & D **22**, 16–18 (2005)
13. Bonifaci, V., Korteweg, P., Marchetti-Spaccamela, A., Stougie, L.: An approximation algorithm for the wireless gathering problem. Operations Research Letters **36**(5), 605–608 (2008)
14. Bonifaci, V., Korteweg, P., Marchetti-Spaccamela, A., Stougie, L.: The distributed wireless gathering problem. In: Proc. Int. Conf. on Algorithmic Aspects in Information and Management, *Lecture Notes in Computer Science*, vol. 5034, pp. 72–83. Springer (2008)
15. Bonifaci, V., Korteweg, P., Marchetti-Spaccamela, A., Stougie, L.: Minimizing average flow time in sensor data gathering. In: Proc. 4th Workshop on Algorithmic Aspects of Wireless Sensor Networks, *Lecture Notes in Computer Science*, vol. 5389, pp. 18–29. Springer (2008)
16. Bonifaci, V., Korteweg, P., Marchetti-Spaccamela, A., Stougie, L.: Minimizing flow time in the wireless gathering problem. In: Proceedings of the 25th International Symposium on Theoretical Aspects of Computer Science, pp. 109–120 (2008)
17. Borodin, A., El-Yaniv, R.: Online computation and competitive analysis. Cambridge University Press (1998)
18. Boukerche, A. (ed.): Handbook of Algorithms for Wireless Networking and Mobile Computing. Chapman & Hall (2005)
19. Cidon, I., Kutten, S., Mansour, Y., Peleg, D.: Greedy packet scheduling. SIAM Journal on Computing **24**(1), 148–157 (1995)

20. Coleri, S.: PEDAMACS: Power Efficient and Delay Aware Medium Access Protocol for Sensor Networks. Master's thesis, University of California, Berkeley (2002)
21. Duckworth, W., Manlove, D., Zito, M.: On the approximability of the maximum induced matching problem. Journal of Discrete Algorithms **3**(1), 79–91 (2005)
22. Florens, C., Franceschetti, M., McEliece, R. J.: Lower bounds on data collection time in sensory networks. IEEE Journal on Selected Areas in Communications **22**, 1110– 1120 (2004)
23. Garey, M. R., Johnson, D. S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. Freeman (1979)
24. Gargano, L.: Time optimal gathering in sensor networks. In: Proceedings of the 14th International Colloquium on Structural Information and Communication Complexity, *Lecture Notes in Computer Science*, vol. 4474, pp. 7–10. Springer (2007)
25. Gargano, L., Rescigno, A. A.: Optimally fast data gathering in sensor networks. In: Proceedings of the 31st Symposium on Mathematical Foundations of Computer Science, *Lecture Notes in Computer Science*, vol. 4162, pp. 399–411. Springer (2006)
26. Hromkovič, J.: Algorithmics for Hard Problems — Introduction to Combinatorial Optimization, Randomization, Approximation, and Heuristics. Springer (2001)
27. Klasing, R., Morales, N., Pérennes, S.: On the complexity of bandwidth allocation in radio networks. Theoretical Computer Science **406**, 225–239 (2008)
28. Korteweg, P.: Online Gathering Algorithms for Wireless Networks. Ph.D. thesis, Eindhoven Technical University (2008)
29. Kumar, V. S. A., Marathe, M. V., Parthasarathy, S., Srinivasan, A.: End-to-end packet-scheduling in wireless ad-hoc networks. In: J. I. Munro (ed.) Proceedings of the 15th Symposium on Discrete Algorithms, pp. 1021–1030 (2004)
30. Morales, N.: Algorithmique de réseaux de communication radio modelisés par de graphes. PhD thesis, Université de Nice-Sophia Antipolis (2007)
31. Papadimitriou, C. H.: Computational Complexity. Addison-Wesley (1994)
32. Perkins, C. E. (ed.): Ad Hoc Networking. Addison-Wesley Professional (2001)
33. Pottie, G. J., Kaiser, W. J.: Wireless integrated network sensors. Communications of the ACM **43**(5), 51–58 (2000)
34. Raghavendra, C. S., Sivalingam, K. M., Znati, T. (eds.): Wireless Sensor Networks. Springer (2004)
35. Revah, Y., Segal, M.: Improved algorithms for data-gathering time in sensor networks II: Ring, tree and grid topologies. In: Proc. of the 3rd IEEE Int. Conf. on Networking and Services (2007)
36. Revah, Y., Segal, M.: Improved lower bounds for data-gathering time in sensor networks. In: Proc. of the 3rd IEEE Int. Conf. on Networking and Services (2007)
37. Schmid, S., Wattenhofer, R.: Algorithmic models for sensor networks. In: Proceedings of the 20th International Parallel and Distributed Processing Symposium (2006)
38. Segal, M., Yedidsion, L.: On real time data-gathering in sensor networks. In: Proc. of the 3rd IEEE Int. Conf. on Mobile, Adhoc and Sensor Systems (2007)
39. Vazirani, V. V.: Approximation Algorithms. Springer, Berlin (2001)
40. Yao, A. C. C.: Probabilistic computations: Towards a unified measure of complexity. In: Proc. of the 18th Symp. on the Foundations of Computer Science, pp. 222–227 (1977)