

Bounded information dissemination in multi-channel wireless networks

**Yu Yan, Dongxiao Yu, Yuexuan Wang,
Jiguo Yu & Francis C. M. Lau**

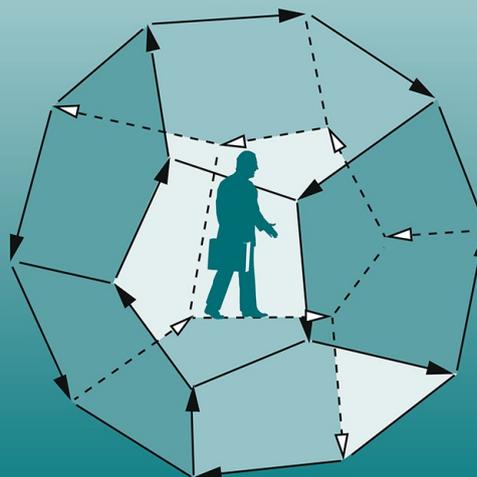
**Journal of Combinatorial
Optimization**

ISSN 1382-6905

J Comb Optim
DOI 10.1007/s10878-014-9804-3

Volume 28, No. 4, November 2014

**Journal of
Combinatorial
Optimization**



 Springer



 Springer

Your article is protected by copyright and all rights are held exclusively by Springer Science +Business Media New York. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Bounded information dissemination in multi-channel wireless networks

Yu Yan · Dongxiao Yu · Yuexuan Wang ·
Jiguo Yu · Francis C. M. Lau

© Springer Science+Business Media New York 2014

Abstract More and more wireless networks and devices now operate on multiple channels, which poses the question: How to use multiple channels to speed up communication? In this paper, we answer this question for the case of wireless ad-hoc networks where information dissemination is a primitive operation. Specifically, we propose a randomized distributed algorithm for information dissemination that is very near the optimal. The general information dissemination problem is to deliver k information packets, stored initially in k different nodes (the packet holders), to all the nodes in the network, and the objective is to minimize the time needed. With an eye toward the reality, we assume a model where the packet holders are determined by an adversary, and neither the number k nor the identities of packet holders are known to the nodes in advance. Not knowing the value of k sets this problem apart from broadcasting and all-to-all communication (gossiping). We study the information dis-

Y. Yan
Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084,
People's Republic of China
e-mail: yanyu02@gmail.com

D. Yu · Y. Wang (✉) · F. C. M. Lau
Department of Computer Science, The University of Hong Kong, Hong Kong,
People's Republic of China
e-mail: amywang@hku.hk

D. Yu
e-mail: dxyu@cs.hku.hk

F. C. M. Lau
e-mail: fcmlau@cs.hku.hk

J. Yu
School of Computer Science, Qufu Normal University, Rizhao 276800, People's Republic of China
e-mail: jiguoju@sina.com

semination problem in single-hop networks with bounded-size messages. We present a randomized algorithm which can disseminate all packets in $O(k(\frac{1}{\mathcal{F}} + \frac{1}{\mathcal{P}}) + \log^2 n)$ rounds with high probability, where \mathcal{F} is the number of available channels and \mathcal{P} is the bound on the number of packets a message can carry. Compared with the lower bound $\Omega(k(\frac{1}{\mathcal{F}} + \frac{1}{\mathcal{P}}))$, the given algorithm is very close to the asymptotical optimal except for an additive factor. Our result provides the first solid evidence that multiple channels can indeed substantially speed up information dissemination, which also breaks the $\Omega(k)$ lower bound that holds for single-channel networks (even if \mathcal{P} is infinity).

Keywords Information dissemination · Multi-channel wireless network · Distributed algorithm · Randomized algorithm

1 Introduction

More and more wireless devices, such as those using Wi-Fi ([IEEE 802.11 1999](#)) or Bluetooth ([Bluetooth 2007](#)), are going for multiple communication channels as opposed to relying on a single channel. This raises a fundamental question: How much can we gain by exploiting multiple channels in communications? Surprising, many answers to this question by researchers in computer science communities have only started to emerge recently. This paper tackles a case where this question is still largely unanswered: *information dissemination in a wireless ad-hoc network*. Wireless ad-hoc networks are among the most common ones in today's scenarios because of the proliferation of personal wireless devices and gadgets. This kind of ad-hoc networks have no pre-defined structure, and they are usually "one-hop" because of the small physical area or locality (e.g., a mall) within which they are formed. Due to the fact that in reality, wireless networks that are formed on the spot seldom have any centralized control, and nodes in such networks are too transient to serve as one, we pay our attention to distributed solutions.

Information dissemination is the most fundamental operation of all that support the smooth running of a network. More precisely, in a network with n nodes (wireless devices), where nodes have access to $\mathcal{F} > 1$ channels, the information dissemination problem (IDP) is to disseminate k information packets, initially stored at k unknown nodes (the packet holders), to the entire network. The objective is to minimize the time needed. Information dissemination is a basic building block for many upper-layer applications, such as routing ([Carzaniga et al. 2012](#)), network topology learning ([Gobjuka and Breitbart 2010](#)), and service/resource discovery ([Mian et al. 2009](#)). Two special cases of and well known IDPs are broadcasting ($k = 1$) and all-to-all communication (sometimes called gossiping) ($k = n$).

In single-channel networks, $\Omega(k)$ is an inherent lower bound for disseminating k packets. This is true even if a message can carry an unbounded number of information packets, since each packet holder needs to exclusively occupy the channel in a round in order to send its packet to at least one other node. Our question then is whether the availability of $\mathcal{F} > 1$ channels can significantly accelerate information dissemination, and if it can, whether \mathcal{F} times speedup is possible. To achieve the ideal speedup is easier hoped than done, for the simple reason that in each round a node can only operate

on one channel out of the \mathcal{F} channels; and to have a successful transmission, the sender must be the only one using the channel because of the possibility of collision. To avoid collisions, the nodes need to carefully coordinate their transmissions. This is difficult under the assumptions that there is no global information on the network topology and nodes have no knowledge about the other nodes, in particular who hold the packets initially. Much to our surprise, by using a novel information dissemination algorithm, these difficulties can actually be overcome and an \mathcal{F} times speedup is attainable. We show how this is done in the following.

1.1 Network model and problem definition

In a single-hop network, each node can directly communicate with any other node. There are \mathcal{F} channels in the network, denoted as $\{1, 2, \dots, \mathcal{F}\}$. Time is divided into synchronized rounds. Each node is equipped with a single half-duplex radio transceiver, such that in each round, each node can select only one of the \mathcal{F} channels to listen to or transmit on, but not both. A node operating on a channel in a given round learns nothing about events on the other channels. When a node v listens to a channel, it can receive a message if and only if there is only one node transmitting on the channel. If two or more nodes transmit on the same channel, a collision occurs and none of these transmissions would be successful. It is assumed that nodes can detect collisions, i.e., nodes can distinguish collision from silence.

Definition 1 (*Information dissemination*) In a wireless network of n nodes, where there is a subset of $k \leq n$ nodes each having a distinct piece (packet) of information, the information dissemination problem (IDP) is to disseminate all these k packets to every node in the network in the fewest rounds.

In order that our algorithm can work in a wide range of scenarios including the worst-case ones, we assume that the initial packet holders are determined by an adversary. Nodes have no knowledge about the network except a polynomial upper bound on the network size n . As shown in the subsequent sections, a polynomial estimation only affects the time complexity of the proposed algorithm by a constant factor; in reality, a rough estimate of n is not difficult to obtain if it is really necessary. For simplicity, we use n itself to denote the estimate.

We study the IDP under the bounded-size message model which we think should be practical. It defines a limit \mathcal{P} on the number of packets in a transmitted message. The design of our algorithm does not assume any particular specific values for \mathcal{P} , although we believe for some special values, further improvements are possible.

In what follows, when we say an event occurs with high probability, it means that the event occurs with probability $1 - n^{-c}$ for a certain constant $c > 0$.

1.2 Our result

In multi-channel single-hop networks, we present a randomized algorithm that can deliver all k packets to the entire network in $O(k(\frac{1}{\mathcal{F}} + \frac{1}{\mathcal{P}}) + \log^2 n)$ rounds with high

probability. We also give an $\Omega(\frac{k}{\mathcal{F}} + \frac{k}{\mathcal{P}})$ lower bound. In the cases where $\mathcal{P} \geq \mathcal{F}$, our algorithm has running time $O(\frac{k}{\mathcal{F}})$, which achieves a \mathcal{F} times speedup over relying on just a single channel (for which $\Omega(k)$ is always a lower bound even when $\mathcal{P} = \infty$). This is the best possible speedup considering the $\Omega(\frac{k}{\mathcal{F}})$ lower bound.

In the unbounded-size message model, i.e., $\mathcal{P} = n$, the best known distributed algorithm for information dissemination in single-hop networks with \mathcal{F} channels is given in [Daum et al. \(2013\)](#), which can disseminate k packets in $O(k + \frac{\log^2 n}{\mathcal{F}} + \log n \log \log n)$ rounds. In contrast, our algorithm has running time $O(\frac{k}{\mathcal{F}} + \log^2 n)$ in this case, which is \mathcal{F} times faster when k is large.

Almost all previous algorithms use a direct way to accomplish information dissemination, i.e., they let each packet holder send its packet directly to all other nodes. This would yield an $\Omega(k)$ lower bound, even when multiple channels are used and each message can carry multiple packets, since each node can receive at most one message at any one time. Here we use an indirect approach to break this lower bound. The packets are first collected to a small number of elected disseminators by transmissions on multiple channels, and then these disseminators would broadcast the collected packets on a special dissemination channel. The key feature of the indirect approach is that each packet holder only needs to send its packet to a disseminator on one channel, rather than to every other node. Meanwhile, the number of disseminators is upper bounded by $O(\mathcal{F} \log n)$, which ensures that the contention level on the dissemination channel is not high and each disseminator can broadcast their packets very quickly. By properly selecting the transmission channels and adjusting the transmission probability, the packet collection procedure can be accomplished in $O(\frac{k}{\mathcal{F}})$ rounds with high probability, which is a key contributor in the efficient running time of the given algorithm.

1.3 Related work

The information dissemination in single-channel single-hop networks is a classical problem. Distributed algorithm research on this problem can be dated back to 1970s ([Hayes 1978](#)), and there is a long line of literature addressing this fundamental problem under different settings, e.g., [Chlebus et al. \(2006\)](#), [Clementi et al. \(2001\)](#), [Fernandez-Anta et al. \(2013\)](#), [Goldberg et al. \(2004\)](#), [Kowalski \(2005\)](#), [Martel \(1994\)](#), and [Yu et al. \(2012\)](#). Because in single-channel networks, the capacity of a message does not significantly impact the complexity of information dissemination, almost all previous work assumed unit-size messages each of which can only take one packet.

The study of distributed algorithms in multi-channel wireless networks is relatively recent. There are some recent papers focusing on information dissemination in multi-channel single-hop radio networks. Except [Daum et al. \(2013\)](#), the known results are all under the unit-size message model. The impact of multiple channels on information dissemination with unit-size messages is limited if nodes can only operate on one channel at a time. $\Omega(k)$ is a trivial lower bound for this case, since each node can receive at most one message in a round. In [Holzer et al. \(2011\)](#), [Holzer et al. \(2012\)](#), addressed the issue of how many channels are enough to achieve the

optimal $\Theta(k)$ bound. In a very recent paper [Daum et al. \(2013\)](#), by proposing a randomized algorithm with running time $O(k \log n + \frac{\log^2 n}{\mathcal{F}} + \log n \log \log n)$, Daum et al. showed that when k is small, the availability of multiple channels can give a result slightly better than that for single-channel networks. A different result is given in [Shi et al. \(2012\)](#), with the assumption that nodes can listen to and receive messages from multiple channels concurrently, [Shi et al. \(2012\)](#) gave an $O(\log k \log \log k)$ time randomized algorithm using $\Theta(n)$ channels. Regarding information dissemination in the unbounded-size message model, the only known result was given in [Daum et al. \(2013\)](#), where the presented randomized algorithm could accomplish information dissemination in $O(k + \frac{\log^2 n}{\mathcal{F}} + \log n \log \log n)$ rounds with high probability. This result does not break the $\Omega(k)$ lower bound for single-channel networks. So far, there are not known results on the general information dissemination problem for multi-channel networks with bounded-size messages.

There are also some recent work addressing other fundamental problems in multi-channel wireless networks, such as maximal independent set ([Daum et al. 2013](#)), leader election ([Daum et al. 2012a](#)), wake-up ([Daum et al. 2012b](#)) and broadcast ([Dolev et al. 2011](#)).

Bounded-size messages have been considered in some studies of the gossip problem. In particular, in [Bermond et al. \(1998\)](#), with the same bounded-size message model as assumed here, Bermond et al. studied the exact complexity of the gossiping in some special graphs. In [Bar-yehuda et al. \(1993\)](#), with a $\log n$ restriction on the message size, Bar-Yehuda et al. presented a randomized gossiping algorithm for unknown topologies. The $b(n)$ -gossip problem, where the number of packets in a single-message is bounded by an integer function $b(n)$, was studied in [Christersson et al. \(2002\)](#).

2 Information dissemination algorithm

We now start presenting our algorithm. For easy of description and analysis, we denote $f = \mathcal{F} - 1$. The proposed algorithm needs an estimate of k as input. In this section, we assume that nodes possess the same estimate \hat{k} of k with $\hat{k} \in [k, 16k]$; we give an algorithm for deriving such an estimate in Sect. 3. We assume that $\mathcal{F} \leq \sqrt{\hat{k}/16 \log n}$.

Otherwise, the algorithm only uses the first $\sqrt{\hat{k}/16 \log n}$ channels, since the algorithm would not benefit from utilizing more channels.

2.1 Algorithm overview

The algorithm uses an indirect way to accomplish information dissemination. Specifically, the algorithm consists of two consecutive procedures: packet collection and packet dissemination. In the packet collection procedure, nodes send their packets to a small number of elected nodes called *disseminators*. The first f channels (of the \mathcal{F} channels) are used for collecting packets in this procedure. Then in the packet dissemination procedure, disseminators broadcast collected packets on the special dissemination channel (the \mathcal{F} -th channel).

At the start of the algorithm, a preprocessing stage is executed to elect f disseminators, each of which will listen on a different channel of the first f channels and collect packets transmitted by non-disseminators in the packet collection procedure. In the packet collection procedure, each non-disseminator selects a channel from the first f channels uniformly at random to transmit with well tuned probability in each round. Uniform selection of channels ensures that the contention on each channel is averaged out over multiple channels. The transmission probability of non-disseminators is adjusted based on the number of received messages by the disseminators, which reflects the contention level on the channels. With the adjustment strategy, in the analysis, it can be shown that the sum of the transmission probabilities of non-disseminators on each channel is always at the level of $\Theta(1)$. In other words, a constant number of non-disseminators transmit on each channel in expectation in each round, which ensures that each disseminator can receive $\Omega(1)$ messages on average. This is a key point in the derivation of the $O(\frac{k}{\mathcal{F}})$ bound on the running time. When the transmission probabilities of non-disseminators increase beyond a set threshold, the packet collection procedure ends and the packet dissemination procedure starts. In the packet dissemination procedure, disseminators broadcast stored packets on the special dissemination channel. Note that here the disseminators not only include those elected at the beginning, but also non-disseminators failing to send their packets to disseminators in the packet collection procedure whose number can be upper bounded by $O(\mathcal{F} \log n)$. So there are totally $O(\mathcal{F} \log n)$ disseminators operating in the packet dissemination procedure, which makes sure that the dissemination channel has a low contention level and each disseminator can broadcast their stored packets very quickly. The reason why we impose an ending condition instead of letting all packets to be collected to the elected disseminators is that when the number of non-disseminators is small, the utilization of multiple channels has no significant impact on packet collection.

There are four states that nodes may be in during the execution of the algorithm: the waiting state \mathbb{W} , the information submission state \mathbb{I} , the disseminator collection state \mathbb{C} and the disseminator broadcast state \mathbb{B} . States \mathbb{C} and \mathbb{I} are for the packet collection procedure, and state \mathbb{B} is for the packet dissemination procedure. Initially, all nodes without packets are in state \mathbb{W} in which nodes do nothing except listening on the \mathcal{F} -th channel.

2.2 Preprocessing stage

All packet holders first perform a preprocessing stage to elect the disseminators. The disseminators are elected using the leader election algorithm in [Daum et al. \(2012a\)](#) which can elect exactly one leader in $O(\frac{\log^2 n}{\mathcal{F}} + \log n)$ rounds with high probability. The preprocessing stage is divided into f phases, and each phase contains $\Theta(\frac{\log^2 n}{\mathcal{F}} + \log n)$ rounds which is an upper bound on the running time of the leader election algorithm. In each phase, packet holders that are not elected as leaders in previous phases execute the leader election algorithm. After the preprocessing stage, with high probability, f leaders are elected, which will work as disseminators in the subsequent procedures.

In the following description, we assume that exactly k leaders are elected. The error probability will be considered in the analysis. After the preprocessing stage, the elected disseminators enter state \mathbb{C} and other packet holders enter state \mathbb{I} .

2.3 Packet collection procedure

In the packet collection procedure, the disseminator elected in the i -th phase of the preprocessing stage listens on channel i and collects packets that are transmitted on this channel, while packet holders in state \mathbb{I} transmit on the first f channels in order to send their packets to the disseminators. At the beginning of the procedure, elected disseminators are in state \mathbb{C} and other packet holders are in state \mathbb{I} . The execution of this procedure is divided into phases, each of which contains $c_l \log n$ rounds, where c_l is a given constant. Between two consecutive phases, there is an adjusting round for adjusting the transmitting probability p_s of nodes in \mathbb{I} based on the number of received messages by the disseminators in the past phase. When p_s has increased to above a threshold, all remaining nodes in state \mathbb{I} and all disseminators in \mathbb{C} enter state \mathbb{B} , and start executing the packet dissemination procedure. In other words, these remaining nodes in \mathbb{I} work as disseminators in the subsequent packet dissemination procedure to directly broadcast their packets. We next give detailed descriptions for the operations in state \mathbb{C} and \mathbb{I} respectively, which is also given in Algorithm 1.

2.3.1 Information submission state \mathbb{I}

In each round of every phase, each node v in state \mathbb{I} selects one channel from the first f channels uniformly at random and transmits with a specific probability p_s on it. In the adjusting rounds, all nodes in state \mathbb{I} listen on the \mathcal{F} -th channel. The initial transmission probability p_s of v is set as $f/2\hat{k}$. After each phase, v adjusts p_s based on whether detected transmissions in the adjusting round. Once v has successfully transmitted a message on the selected channel, i.e., it transmits in a round and does not sense any other transmission on the selected channel, v enters state \mathbb{W} and listens on the \mathcal{F} -th channel from then on. When the transmission probability p_s is increased to at least $\frac{1}{c_l^2 \log n}$, v joins state \mathbb{B} .

2.3.2 Disseminator collection state \mathbb{C}

Denote by u_i the disseminator elected in the i -th phase of the preprocessing stage. In each round of every phase, u_i listens on channel i for receiving messages transmitted by nodes in \mathbb{I} . After each phase, if u_i has received at least $12 \log n$ messages in the past phase, it transmits on the \mathcal{F} -th channel in the subsequent adjusting round to adjust the transmission probability of nodes in state \mathbb{I} ; otherwise, u listens on the \mathcal{F} -th channel. Disseminators in \mathbb{C} keep updating parameters *count*, *phase* and p_s synchronously with nodes in \mathbb{I} . When the packet dissemination procedure begins, i.e., p_s has been increased to at least $\frac{1}{c_l^2 \log n}$, u_i joins state \mathbb{B} .

Algorithm 1 BID: Packet Collection Procedure

Initially, $c_l = 192$; $l = c_l \log n$; $p_s = f/2k$; $count = 1$; $phase = 1$; $MsgNum = 0$

Initially

1: **if** $p_s \geq \frac{1}{c_l^2 \log n}$ **then** $state = \mathbb{B}$; **else** $state = \mathbb{I}$ **end if**

State I

2: **if** $count = l + 1$ **then**

3: listen the \mathcal{F} -th channel; $count = 1$; $phase = phase + 1$;

4: **if** there is not any transmission **then** $p_s = \min\{2p_s, \frac{1}{2}\}$ **end if**
*/** adjust the transmission probability after a phase*

5: **if** $p_s \geq \frac{1}{c_l^2 \log n}$ **then** $state = \mathbb{B}$; $phase = 1$ **end if**
*/** state transition from I to B*

6: **else**

7: uniformly at random select a channel from $\{1, 2, \dots, f\}$;

8: transmit with probability p_s on the selected channel
*/** select the operating channel*

9: **if** successfully transmitted a message **then** $state = \mathbb{W}$ **end if**
*/** state transition from I to W*

10: $count = count + 1$

11: **end if**

State C

12: **if** $count = l + 1$ **then**

13: **if** $MsgNum \geq 12 \log n$ **then** transmit on the \mathcal{F} -th channel; **else** listen on the \mathcal{F} -th channel **end if**
*/** adjust the transmission probability of non-disseminators*

14: update p_s according to the same rule as nodes in I; $phase = phase + 1$; $count = 1$; $MsgNum = 0$

15: **if** $p_s \geq \frac{1}{c_l^2 \log n}$ **then** $state = \mathbb{B}$; $phase = 1$ **end if**
*/** state transition from C to B*

16: **else**

17: listen on the specified channel

18: **if** receive a message **then** $MsgNum = MsgNum + 1$ **end if**

19: $count = count + 1$

20: **end if**

2.4 Packet dissemination procedure

At the beginning, disseminators are in state \mathbb{B} and all other nodes are in state \mathbb{W} . In this procedure, all nodes operate on the \mathcal{F} -th channel. The algorithm execution is also divided into phases, each of which consists of l rounds for $l = \Theta(\log n)$. Nodes in state \mathbb{W} do nothing expect listening. The operations of nodes in state \mathbb{B} are introduced next, which is also given in Algorithm 2.

2.4.1 Disseminator broadcast state \mathbb{B}

In each round of every phase, nodes in \mathbb{B} transmit with probability p_b and listen with probability $1 - p_b$. The initial transmission probability of nodes in \mathbb{B} is set as $\frac{1}{\beta f \log n}$, where β is a large enough constant for ensuring high probability results. After each phase, for a node v in \mathbb{B} , the transmission probability p_b is adjusted based on the number of received messages in the past phase. Because of the bound \mathcal{P} on the number of packets in a message, nodes coming from state \mathbb{C} may need to transmit multiple

messages for disseminating their collected packets. Once a node in \mathbb{B} has successfully transmitted all its stored packets, it enters state \mathbb{W} .

Algorithm 2 BID: Packet Dissemination Procedure

Initially, $p_b = \frac{1}{\beta f \log n}$; $MsgNum = 0$; $l = 192 \log n$

State \mathbb{B}

- 1: transmit on the \mathcal{F} -th channel with probability p_b ; $count = count + 1$;
 - 2: **if** received a message **then** $MsgNum = MsgNum + 1$ **end if**
 - 3: **if** $count = l$ **then**
 - 4: **if** $MsgNum < 16 \log n$ **then** $p_b = \min\{2p_b, \frac{1}{4}\}$ **end if**
*/** adjust the transmission probability after a phase*
 - 5: $count = 1$; $phase = phase + 1$; $MsgNum = 0$;
 - 6: **end if**
 - 7: **if** All stored packets has been successfully transmitted **then** $state = \mathbb{W}$ **end if**
*/** state transition from \mathbb{B} to \mathbb{W}*
-

2.5 Analysis

We next show that the proposed algorithm can accomplish information dissemination in $O(k(\frac{1}{\mathcal{F}} + \frac{1}{\mathcal{P}}) + \log^2 n)$ rounds with high probability. We also prove that any information dissemination algorithm needs $\Omega(\frac{k}{\mathcal{F}} + \frac{k}{\mathcal{P}})$ rounds even if collision detection is enabled. Thus the proposed algorithm is asymptotically optimal when k is large. Without confusion, we use \mathbb{I} , \mathbb{C} and \mathbb{B} to denote the set of nodes in their corresponding states, and denote $P_{\mathbb{I}}$ and $P_{\mathbb{B}}$ as the sum of transmission probabilities of nodes in state \mathbb{I} and state \mathbb{B} , respectively.

We first bound the time for the preprocessing stage. In Daum et al. (2012a), it is stated that the leader election algorithm can exactly elect one leader in $O(\frac{\log^2 n}{\mathcal{F}} + \log n)$ rounds with probability $1 - n^{-c}$ for a constant c . At the cost of adapting some involved constant parameters in the algorithm, c can be chosen arbitrarily. We give this result in the following Lemma 1.

Lemma 1 (Daum et al. 2012a) *With probability $1 - n^{-2}$, the leader election algorithm can elect exactly one leader in $O(\frac{\log^2 n}{\mathcal{F}} + \log n)$ rounds.*

From the above lemma it follows that f disseminators can be exactly elected with high probability in the preprocessing stage.

Lemma 2 *In the preprocessing stage, f disseminators are elected in $O(\log^2 n + \mathcal{F} \log n) \in O(\log^2 n + \frac{k}{\mathcal{F}})$ rounds with probability $1 - n^{-1}$.*

In the following, for simplicity, we assume that the leader election algorithm is correctly executed always, and the error probability will be considered at the end. Next we start bounding the time needed for the packet collection procedure and the packet dissemination procedure, respectively.

We first bound the time for the packet collection procedure. The basic idea is as follows: during the procedure, the adjustment strategy of the transmission probability

ensures that $P_{\mathbb{I}}$ are always maintained at a level of $\Theta(\mathcal{F})$. Because non-disseminators select the channels uniformly, the expected number of transmitters on each channel is $\Theta(1)$ in each round. Based on this claim, it can be shown that in each phase (consists of $\Theta(\log n)$ rounds), $\Omega(\mathcal{F} \log n)$ non-disseminators successfully send their packets to disseminators. Thus, the packet collection procedure lasts for at most $O(\frac{k}{\mathcal{F}})$ rounds.

In the following Lemma 3, we first present an upper bound on $P_{\mathbb{I}}$.

Lemma 3 *In the first $O(n^2)$ rounds of the packet collection procedure, with probability $1 - O(n^{-1})$, $P_{\mathbb{I}} \leq f/2$.*

Proof We prove the lemma by contradiction. Assume that t is the first round such that the lemma does not hold. Since the transmission probability may only be increased in adjusting rounds, t must be the first round of a certain phase. Assume this phase be phase i . Note that initially, by the value range of \hat{k} , $p_s = \frac{f}{2\hat{k}} \cdot k \leq \frac{f}{2}$. So we have i is not the first phase of the packet collection procedure.

In the adjusting round before phase i , the transmission probability of nodes can be at most doubled. We can get $P_{\mathbb{I}} \in (\frac{f}{4}, \frac{f}{2}]$ during phase $i - 1$. Let v be a disseminator in \mathbb{C} during phase $i - 1$. We next show that v receives at least $12 \log n$ messages from nodes in \mathbb{I} after phase $i - 1$ with probability $1 - n^{-3}$.

Assume that v is the disseminator on channel k with $1 \leq k \leq f$. Denote P_r as the probability that v receives a message in a round of phase $i - 1$. Then

$$\begin{aligned}
 P_r &= \sum_{u \in \mathbb{I}} p_s \cdot \frac{1}{f} \cdot \prod_{w \in \mathbb{I} \setminus \{u\}} \left(1 - p_s \cdot \frac{1}{f}\right) \\
 &\geq \sum_{u \in \mathbb{I}} p_s \cdot \frac{1}{f} \cdot \left(\frac{1}{4}\right)^{\sum_{w \in \mathbb{I}} p_s \cdot \frac{1}{f}} \\
 &\geq \sum_{u \in \mathbb{I}} p_s \cdot \frac{1}{f} \cdot \left(\frac{1}{4}\right)^{\frac{f}{2} \cdot \frac{1}{f}} \\
 &\geq \frac{1}{8} \tag{1}
 \end{aligned}$$

In other words, v will receive a message in each round of phase $i - 1$ with constant probability. And v can receive at least $24 \log n$ messages in phase $i - 1$ in expectation. By the Chernoff bound, during phase $i - 1$, v can receive at least $12 \log n$ messages with probability $1 - n^{-3}$. In the adjusting round after phase $i - 1$, v transmits on the \mathcal{F} -th channel. After detecting the transmissions, all nodes in \mathbb{I} do not change the value of p_s in the subsequent phase by the algorithm. In round t , $p_s \leq \frac{f}{2}$ with probability $1 - n^{-3}$, which contradicts the assumption on t . Then none of the first $O(n^2)$ rounds are the first violating one with probability $1 - O(n^{-1})$. \square

We next present a lower bound of $P_{\mathbb{I}}$. We call a phase in the packet collection procedure an increasing one if the transmission probability p_s is doubled after the phase; otherwise an unchanging one. We first present a sufficient condition for a phase to be increasing in the following lemma.

Lemma 4 *If in a phase of the packet collection procedure, $P_{\mathbb{I}} \leq \frac{f}{32}$, the phase is increasing with probability at least $1 - n^{-2}$.*

Proof Assume that phase i satisfies the condition. By the algorithm, we only need to show that all disseminators in \mathbb{C} receive less than $12 \log n$ messages during phase i . Then in the subsequent adjusting round, none of the disseminators in \mathbb{C} transmits on the \mathcal{F} -th channel. Consequently, after phase i , all nodes double the value of p_s , which will complete the proof.

In each round of phase i , the probability that a disseminator v in state \mathbb{C} receives a message is

$$\sum_{u \in \mathbb{I}} p_s \cdot \frac{1}{f} \cdot \prod_{w \in \mathbb{I} \setminus \{u\}} \left(1 - p_s \cdot \frac{1}{f}\right) \leq \sum_{u \in \mathbb{I}} p_s \cdot \frac{1}{f} \leq \frac{1}{32}. \tag{2}$$

Disseminator v can receive at most $6 \log n$ messages in expectation in phase i . The number of messages v received is less than $12 \log n$ with probability $1 - n^{-3}$ by the Chernoff bound. Each disseminator receives less than $12 \log n$ messages during phase i with probability $1 - n^{-2}$, which completes the proof. \square

Lemma 5 *In the first $O(n)$ phases of the packet collection procedure, with probability $1 - O(n^{-1})$, $P_{\mathbb{I}} \geq \frac{f}{128}$.*

Proof Assume that phase j is the first one during which there is a round such that $P_{\mathbb{I}} < \frac{f}{128}$. By the initial setting of p_s , $P_{\mathbb{I}} \geq \frac{f}{32}$ at the beginning of the algorithm. Furthermore, note that $p_s \leq \frac{1}{c_l^2 \log n}$ if there are nodes joining state \mathbb{I} . Thus, there are at least $2c_l f \log n$ nodes in state \mathbb{I} . During the first phase of stage 1, at most $c_l f \log n$ nodes in \mathbb{I} join state \mathbb{W} after successfully transmitting their messages. Consequently, after this phase, $P_{\mathbb{I}} \geq \frac{1}{2} \cdot \frac{f}{32} = \frac{f}{64}$. So phase j is not the first phase of the packet collection procedure.

In phase $j - 1$, we still have $P_{\mathbb{I}} \geq \frac{f}{128}$ with the assumption that phase j is the first violating one. Similar to the above, if there are still nodes in state \mathbb{I} , $P_{\mathbb{I}}$ can be decreased by at most a factor $\frac{1}{2}$ in any phase of the packet collection procedure with probability $1 - O(n^{-1})$. So it is easy to get $P_{\mathbb{I}} \leq \frac{f}{32}$ in phase $j - 1$ with the assumption on j . By Lemma 4, phase $j - 1$ is an increasing phase with probability $1 - n^{-2}$, which means $p_s \geq \frac{f}{128} \cdot 2 \cdot \frac{1}{2} = \frac{f}{128}$ during phase j . This contradiction shows that j is not the first violating phase with probability $1 - n^{-2}$. With probability $1 - O(n^{-1})$, none of the first $O(n)$ phases in the packet collection procedure is the first violating one. \square

With Lemma 5 in the above, we can lower bound the number of nodes leaving state \mathbb{I} in each phase of the packet collection procedure.

Lemma 6 *In each phase of the first $O(n)$ ones in the packet collection procedure, if there are still nodes in \mathbb{I} , at least $\Omega(f \log n)$ of these nodes enter state \mathbb{W} with probability $1 - O(n^{-1})$.*

Proof Assume phase j satisfies the given condition. In each round of phase j , the expected number N_s of nodes in \mathbb{I} that can successfully transmit messages is

$$\begin{aligned}
 N_s &= \sum_{v \in \mathbb{I}} p_s \prod_{u \in \mathbb{I} \setminus \{v\}} \left(1 - p_s \cdot \frac{1}{f}\right) \\
 &\geq \sum_{v \in \mathbb{I}} p_s \left(\frac{1}{4}\right)^{\sum_{u \in \mathbb{I}} p_s / f} \\
 &\geq \frac{f}{256}
 \end{aligned}
 \tag{3}$$

The last inequality is by Lemma 3 and Lemma 5. Using a standard Chernoff bound argument, it is easy to show that during phase j , with probability $1 - n^{-2}$, $\Omega(f \log n)$ nodes will transit to state \mathbb{W} . By the union bound and taking the error probability of Lemma 3 and Lemma 5 into account, the lemma is proved. \square

Now we are ready to bound the time needed for the packet collection procedure.

Lemma 7 *With probability $1 - O(n^{-1})$, the packet collection procedure takes at most $O(\frac{k}{f})$ rounds.*

Proof To prove the lemma, we need to bound the number of rounds from the beginning to the time when p_s exceeds the threshold $\frac{1}{c_f^2 \log n}$.

By Lemma 6, in each of the first $O(n)$ phases of the packet collection procedure, if there are still nodes in \mathbb{I} , with probability $1 - O(n^{-1})$, $\Omega(f \log n)$ of those nodes switch state. Then after the algorithm has started for $O(\frac{k}{f \log n})$ phases, there will be at most $\frac{c_f^2 f \log n}{128}$ nodes staying in \mathbb{I} . By Lemma 5, $P_{\mathbb{I}} \geq \frac{f}{128}$. Thus, at this time, $p_s \geq \frac{1}{c_f^2 \log n}$. This means that the packet collection procedure has ended at or before this time. Finally, note that each phase has $\Theta(\log n)$ rounds. Then combining the error probability of Lemma 6 and Lemma 5, the lemma is proved. \square

We still need to bound the length of the packet dissemination procedure, i.e., the time for nodes in state \mathbb{B} . Before that, we first bound the number of nodes that transit to state \mathbb{B} from state \mathbb{I} . The bound is obtained based on the transition condition and the upper bound on $P_{\mathbb{I}}$ given in Lemma 3.

Lemma 8 *With probability $1 - O(n^{-1})$, the number of nodes that transit to state \mathbb{B} from \mathbb{I} is at most $O(f \log n)$.*

Proof When the transition condition is satisfied, i.e., $p_s \geq \frac{1}{c_f^2 \log n}$, by Lemma 3, with probability $1 - n^{-1}$, the number of nodes is at most $\frac{c_f^2}{2} f \log n$, which completes the proof. \square

Note that the number of disseminators transiting to state \mathbb{B} from \mathbb{C} is f . Combining with Lemma 8, for a large enough constant β , the number of nodes that join state \mathbb{B} can be bounded by $\frac{\beta}{2} f \log n$. Before bounding the time for the packet dissemination procedure, we first give an upper bound of $P_{\mathbb{B}}$.

Lemma 9 *In the first $O(n^2)$ rounds in stage 2, with probability $1 - O(n^{-1})$, $P_{\mathbb{B}} \leq \frac{1}{2}$.*

Proof Assume that round i in phase j is the first round in which the upper bound is broken. Obviously, i is the first round of phase j , since the probability can be increased only at the end of a phase. Furthermore, by Lemma 8 and the initial probability setting, it is easy to show that phase j is not the first phase of the packet dissemination procedure with probability $1 - O(n^{-1})$. Denote \mathbb{B}_j as the set of nodes in state \mathbb{B} at the beginning of phase j .

Note that $P_{\mathbb{B}}$ can be at most doubled after a phase and j is the first violating phase. Thus, in any round of phase $j - 1$, $P_{\mathbb{B}} \in (\frac{1}{4}, \frac{1}{2}]$. Then using a similar argument to that in Lemma 3, we can show that all nodes in \mathbb{B}_j can receive at least $16 \log n$ messages with probability $1 - n^{-4}$ during phase $j - 1$. This is true for all nodes in \mathbb{B}_j with probability $1 - n^{-3}$. Thus, all nodes in B_j will not change the transmission probability after phase $j - 1$, i.e., during phase j , $P_{\mathbb{B}} \leq P_{\mathbb{B}_j} \leq \frac{1}{2}$. This contradiction shows that round i is not the first violating round with probability $1 - n^{-3}$. None of the first $O(n^2)$ rounds are the first violating round with probability $1 - O(n^{-1})$, which completes the proof. \square

Now we are ready to bound the time for the packet dissemination procedure.

Lemma 10 *With probability $1 - O(n^{-1})$, after the packet dissemination procedure has started for $O(\frac{k}{\mathcal{P}} + f \log n + \log^2 n) \in O(\frac{k}{\mathcal{P}} + \frac{k}{\mathcal{F}} + \log^2 n)$ rounds, all nodes in state \mathbb{B} would have successfully broadcast the stored packets.*

Proof We first bound the total number of successful transmissions needed to disseminate all packets. Denote B_I as the set of nodes that transit to state \mathbb{B} from \mathbb{I} , and denote B_C as the set of nodes that transit to state \mathbb{B} from \mathbb{C} . By Lemma 8, $|B_I| \in O(f \log n)$. Each node in B_I only needs to successfully transmit once as it only needs to disseminate its own packet. For disseminators in B_C , there are at most $O(k)$ packets stored in these nodes as each packet is only stored in one node by the algorithm. As each node in B_C transmits at most one message that contains fewer than \mathcal{P} packets, there are $O(k/\mathcal{P} + f)$ successful transmissions by these nodes. Combining everything together, the total number of successful transmissions by nodes in \mathbb{B} is $O(\frac{k}{\mathcal{P}} + f \log n)$.

For nodes in \mathbb{B} , we call a phase an increasing one if the transmission probability is doubled and unchanging one otherwise. If a phase is unchanging, there are at least $16 \log n$ successful transmissions occurring during the phase. Thus, by the above given bound on the number of successful transmissions needed, there are at most $O(\frac{\frac{k}{\mathcal{P}} + f \log n}{\log n}) = O(\frac{k}{\mathcal{P} \log n} + f)$ unchanging phases. For a node v in \mathbb{B} , after at most $O(\frac{k}{\mathcal{P} \log n} + f + \log n)$ phases, there will be enough increasing phases such that the transmission probability of v increases to at least $\frac{1}{4}$. After that, in each round, by Lemma 9, the probability that v can successfully send a message on the \mathcal{F} -th channel is $\frac{1}{4} \cdot \prod_{u \in \mathbb{B} \setminus \{v\}} (1 - p_b^u) \geq \frac{1}{4} \cdot (\frac{1}{4})^{\sum_{u \in \mathbb{B} \setminus \{v\}} p_b^u} \geq \frac{1}{8}$. Since v needs to successfully transmit at most $O(\frac{k}{\mathcal{P}})$ messages for disseminating packets, using a standard Chernoff bound argument, it is easy to show that from then on, v can successfully transmit all stored packets with probability $1 - O(n^{-2})$ after at most $O(\frac{k}{\mathcal{P}} + \log n)$ rounds. Combining all these together and taking the error probability of Lemma 8 and Lemma 9 into account, we get that after $O(\frac{k}{\mathcal{P} \log n} + f + \log n) \cdot O(\log n) + O(\frac{k}{\mathcal{P}} + \log n) =$

$O(\frac{k}{\mathcal{P}} + f \log n + \log^2 n)$ rounds, with probability $1 - O(n^{-1})$, all nodes in \mathbb{B} can successfully broadcast their stored packets to all other nodes. \square

Theorem 1 *With probability $1 - O(n^{-1})$, the algorithm can disseminate all packets in $O(k(\frac{1}{\mathcal{F}} + \frac{1}{\mathcal{P}}) + \log^2 n)$ rounds. Any information dissemination algorithm has running time of $\Omega(\frac{k}{\mathcal{F}} + \frac{k}{\mathcal{P}})$.*

Proof The running time of the proposed algorithm can be obtained by Lemma 2, Lemma 7 and Lemma 10.

For each node v possessing a packet, if it joins state \mathbb{C} , it will finally send all stored packets on the \mathcal{F} -th channel after joining state \mathbb{B} by Lemma 10. If v stays in state \mathbb{I} , it either successfully transmits to a disseminator in state \mathbb{C} or joins state \mathbb{B} after the packet collection procedure. Thus, its packet will also get broadcast by Lemma 10. Furthermore, note that when a successful transmission occurs on the \mathcal{F} -th channel in the packet dissemination procedure, all nodes except the transmitting one listen on the \mathcal{F} -th channel by the algorithm, which ensures that all nodes will receive the transmitted packets. All the above analysis is based on the fact that disseminators are elected correctly. Taking into account the error probability in Lemma 2 and combining everything together, the correctness of the algorithm is proved.

To accomplish information dissemination, each node needs to transmit its packet to at least one other node, i.e., each node needs to occupy a channel exclusively. In each round, at most \mathcal{F} nodes can successfully transmit messages. Thus, $\Omega(\frac{k}{\mathcal{F}})$ rounds are needed for accomplishing information dissemination. In addition, each transmitted message can contain at most \mathcal{P} packets. Each node needs to receive $\frac{k}{\mathcal{P}}$ messages for acquiring all the packets. Thus, $\frac{k}{\mathcal{P}}$ rounds are needed. Combining all things together, we get the claimed lower bound. \square

3 Estimating k

In this section, we give an algorithm for deriving an estimate of k as required in the information dissemination algorithm in Sect. 2. The detailed algorithm is given in Algorithm 3. The packet holders execute the algorithm on the \mathcal{F} -th channel. The algorithm execution is divided into phases, each phase consisting of $\Theta(\log n)$ rounds. In each round, each packet holder transmits with a specific transmission probability which exponentially decreases after each phase. When there are enough successful transmissions in a phase, the algorithm then halts and nodes output the estimation based on the transmission probability in the particular phase.

We next show that after executing Algorithm 3 for at most $O(\log k \log n)$ rounds, nodes get an estimate \hat{k} which is a constant approximation to k .

Lemma 11 *With probability $1 - n^{-1}$, Algorithm 3 will not halt when $p_e > \frac{8}{k}$. Consequently, the estimate \hat{k} does not fall in the range $(0, k)$.*

Proof We only need to show that in each phase j with $p_e > \frac{8}{k}$, the number of successfully transmitting nodes is less than $8 \log n$ with high probability.

In each round of phase j , the probability that there is a successful transmission is $\sum_{i=1}^k p_e(1 - p_e)^{k-1} \leq \frac{k p_e}{1 - p_e} e^{-k p_e} < 16e^{-8}$. Then during phase j , the expected

Algorithm 3 Estimate k

Initially, $p_e = \frac{1}{2}$; $count = 1$; $phase = 1$; $MsgNum = 0$; $l = 128 \log n$
 1: **if** $count = l + 1$ **then**
 2: **if** $MsgNum \geq 8 \log n$ **then** output $\hat{k} = \frac{8}{p_e}$;
 3: **else** $p_e = p_e/2$; $count = 1$; $MsgNum = 0$; $phase = phase + 1$
 4: **end if**
 5: **else**
 6: transmit with probability p_e on the \mathcal{F} -th channel
 7: **if** received a message **then** $MsgNum = MsgNum + 1$ **end if**
 8: $count = count + 1$
 9: **end if**

number of successfully transmitting nodes is less than $4 \log n$. Then by the Chernoff bound, with probability $1 - n^{-2}$, there are less than $8 \log n$ successful transmissions during phase j , which means the algorithm will not halt after phase j . Meanwhile, note that after such a phase j , the transmission probability of nodes is halved. Thus, there are at most $\log k$ such phases. By the union bound, the algorithm will not halt after any of these phases with probability $1 - n^{-1}$. \square

Lemma 12 *With probability $1 - n^{-1}$, Algorithm 3 will halt before p_e decreases to less than $\frac{1}{2k}$. In other words, the estimation generated is at most $16k$.*

Proof We only need to show that with probability $1 - n^{-1}$, the algorithm will halt when $p_e \in [\frac{1}{2k}, \frac{1}{k})$. Assume in phase j , p_e falls into the particular interval. In each round of phase j , the probability that there is a successful transmission is $\sum_{i=1}^k p_e(1 - p_e)^{k-1} \geq \frac{1}{8}$. In expectation, there are at least $16 \log n$ successful transmissions in phase j . Using the Chernoff bound, there are at least $8 \log n$ successful transmissions in phase j with probability $1 - n^{-1}$, which completes the proof. \square

The following Theorem 2 is a direct corollary of the following Lemma 11 and Lemma 12.

Theorem 2 *Algorithm 3 can get an estimate $\hat{k} \in [k, 16k]$ in $O(\log n \log k)$ rounds with probability $1 - O(n^{-1})$.*

4 Conclusion

We proposed a fast distributed algorithm for information dissemination in multi-channel single-hop networks under the bounded-size message model (each message can carry at most \mathcal{P} packets). When \mathcal{F} channels are available in the network, our algorithm can disseminate k packets in $O(k(\frac{1}{\mathcal{F}} + \frac{1}{\mathcal{P}}) + \log^2 n)$ rounds. The result demonstrates that the availability of multiple channels can greatly accelerate information dissemination and lead to \mathcal{F} times speedup in some cases. Our result is also asymptotically optimal for cases with large k and \mathcal{P} when compared with the given $\Omega(\frac{k}{\mathcal{F}} + \frac{k}{\mathcal{P}})$ lower bound. In the future, it will be interesting to devise efficient information dissemination algorithms in multi-channel multi-hop networks by adapting the techniques developed in this work. In multi-hop networks, where the hops may

actually depend on individual nodes' transmission powers, the issue of interference and how it should be modeled would come into the picture, and add to the difficulty and challenge of the problem. Another future research direction would be to consider the minimum number of channels needed for attaining the bound $O(\frac{k}{C} + \frac{k}{P})$ for a given integer C and the message bound P .

Acknowledgments This work was supported in part by the National Natural Science Foundation of China Grants 61073174 and 61373027, Hong Kong RGC GRF Grant 714311, Shu Shengman Special Fund and Natural Science Foundation of Shandong Province Grant ZR2012FM023.

References

- Bar-yehuda R, Israeli A, Itai A (1993) Multiple communication in multi-hop radio networks. *SIAM J Comput* 22:875–887
- Bermond J, Gargano L, Rescigno AA, Vaccaro U (1998) Fast gossiping by short messages. *SIAM J Comput* 27(4):917–941
- Bluetooth Consortium (2007) Bluetooth Specification Version 2.1
- Carzaniga A, Khazaei K, Kuhn F (2012) Oblivious low-congestion multicast routing in wireless networks. In: Proceedings of the thirteenth ACM international symposium on mobile ad hoc networking and computing (MobiHoc '12). ACM, New York, NY, p 155–164
- Chlebus BS, Kowalski DR, Rokicki MA (2006) Adversarial queuing on the multiple-access channel. In: Proceedings of the twenty-fifth annual ACM symposium on principles of distributed computing (PODC '06). ACM, New York, NY, p 92–101
- Christersson M, Gasieniec L, Lingas A (2002) Gossiping with bounded size messages in ad hoc radio networks. In: Widmayer P, Eidenbenz S, Triguero F, Morales R, Conejo R, Hennessy M (eds) Proceedings of the 29th colloquium on automata, languages and programming (ICALP'02). Springer-Verlag, Berlin, Heidelberg, p 377–389
- Clementi AEF, Monti A, Silvestri Riccardo (2001) Selective families, superimposed codes, and broadcasting on unknown radio networks. In: Proceedings of the twelfth annual ACM-SIAM symposium on discrete algorithms (SODA '01). Society for Industrial and Applied Mathematics, Philadelphia, PA, p 709–718
- Daum S, Gilbert S, Kuhn F, Newport C (2012a) Leader election in shared spectrum radio networks. In: Proceedings of the 2012 ACM symposium on principles of distributed computing (PODC '12). ACM, New York, NY, p 215–224
- Daum S, Kuhn F, Newport C (2012b) Efficient symmetry breaking in multi-channel radio networks. In: Aguilera MK (ed) Proceedings of the 26th international conference on distributed computing (DISC'12). Springer-Verlag, Berlin, Heidelberg, p 238–252
- Daum S, Ghaffari M, Gilbert S, Kuhn F, Newport C (2013) Maximal independent sets in multi-channel radio networks. In: Proceedings of the 2013 ACM symposium on principles of distributed computing (PODC '13). ACM, New York, NY, p 335–344
- Dolev S, Gilbert S, Khabbazi M, Newport C (2011) Leveraging channel diversity to gain efficiency and robustness for wireless broadcast. In: David Peleg (ed) Proceedings of the 25th international conference on distributed computing (DISC'11). Springer-Verlag, Berlin, Heidelberg, p 252–267
- Fernandez-Anta A, Mosteiro MA, Ramon Munoz J (2013) Unbounded contention resolution in multiple-access channels. *Algorithmica* 67(3):295–314
- Gobjuka H, Breitbart YJ (2010) Ethernet topology discovery for networks with incomplete information. *IEEE/ACM Trans Netw* 18(4):1220–1233
- Goldberg LA, Jerrum M, Kannan S, Paterson M (2004) A bound on the capacity of backoff and acknowledgment-based protocols. *SIAM J Comput* 33(2):313–331
- Hayes JF (1978) An adaptive technique for local distribution. *IEEE Trans Commun* COM 26:1178–1186
- Holzer S, Pignolet Y-A, Smula J, Wattenhofer R (2011) Time-optimal information exchange on multiple channels. In: Proceedings of the 7th ACM SIGACT/SIGMOBILE international workshop on foundations of mobile computing (FOMC '11). ACM, New York, NY, p 69–76

- Holzer S, Locher T, Pignolet Y, Wattenhofer R (2012) Deterministic multi-channel information exchange. In: Proceedings of the twenty-fourth annual ACM symposium on parallelism in algorithms and architectures (SPAA '12). ACM, New York, NY, p 109–120
- IEEE 802.11 (1999) Wireless LAN MAC and physical layer specifications
- Kowalski DR (2005) On selection problem in radio networks. In: Proceedings of the twenty-fourth annual ACM symposium on principles of distributed computing (PODC '05). ACM, New York, NY, p 158–166
- Martel CU (1994) Maximum finding on a multiple access broadcast network. *Inf Process Lett* 52(1):713
- Mian AN, Baldoni R, Beraldi R (2009) A survey of service discovery protocols in multihop mobile ad hoc networks. *Pervasive Comput* 8(1):66–74
- Shi W, Hua Q-S, Yu D, Wang Y, Lau FCM (2012) Efficient information dissemination in single-hop multi-channel radio networks. In: Wang X, Zheng R, Jing T, Xing K (eds) Proceedings of the 7th international conference on wireless algorithms, systems, and applications (WASA'12). Springer-Verlag, Berlin, Heidelberg, p 438–449
- Yu D, Hua Q-S, Dai W, Wang Y, Lau FCM (2012) Dynamic contention resolution in multiple-access channels. In: Koucheryavy Y, Mamatas L, Matta I, Tsaoussidis V (eds) Proceedings of the 10th international conference on wired/wireless internet communication (WWIC'12). Springer-Verlag, Berlin, Heidelberg, p 232–243