Journal of Computer and System Sciences ••• (••••) •••–•••

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Deterministic multi-channel information exchange

Stephan Holzer [a,*,1], Thomas Locher [b], Yvonne Anne Pignolet [b], Roger Wattenhofer [c]

[a] MIT, Cambridge, USA
[b] ABB Corporate Research, Baden, Switzerland
[c] ETH Zurich, Switzerland

A R T I C L E I N F O

Article history:
Received 23 October 2014
Received in revised form 10 February 2017
Accepted 18 February 2017
Available online xxxx

Keywords:
Information exchange problem
k-selection
Many-to-all communication
Multi-message broadcast
Multiple access channel
Multi-channel
Wireless computing
Single-hop networks
Deterministic algorithms and lower bounds

A B S T R A C T

We study the information exchange problem on a set of multiple access channels: k arbitrary nodes have information they want to distribute to the entire network of n nodes via a shared medium partitioned into channels. We devise a deterministic algorithm running in asymptotically optimal time $\mathcal{O}(k)$ using $\mathcal{O}(n^{\log(k)/k})$ channels if $k \leq \frac{1}{6}\log n$ and $\mathcal{O}(\log^{1+\rho}(n/k))$ channels otherwise, where $\rho > 0$ is an arbitrarily small constant. This is a super-polynomial improvement over the best known bounds [20]. Additionally we show that our results are significantly closer to the optimal solution by proving that $\Omega(n^{\Omega(1/k)} + \log_k n)$ channels are necessary to achieve a time complexity of $\mathcal{O}(k)$.

1. Introduction

A fundamental problem of many communication systems that rely on a shared communication medium, e.g., wireless and bus networks, is (co-channel) interference, which occurs when more than one network entity tries to transmit a message over the same communication channel at the same time. Such simultaneous (or interleaved) transmissions of two or more messages over the same channel are commonly referred to as *collisions*. Often, a collision distorts all transmitted messages significantly, which entails that none of the messages can be decoded successfully at the receivers. Hence, there is a need for mechanisms that schedule the message transmissions appropriately in order to enable an efficient exchange of messages over the communication medium. There are various techniques to address or simplify this basic scheduling problem: By introducing a notion of time, the network entities can transmit in synchronized *time slots*, which reduces the potential for collisions. Another common trick is to use randomization, as in, e.g., the Aloha protocol [1]. If the network entities further have the ability to detect collisions, which allows the entities to learn that other entities strive to transmit as well, back-off mechanisms can be applied to ensure an eventual transmission of all messages.

---

[*] Corresponding author.
E-mail addresses: holzer@csail.mit.edu (S. Holzer), thomas.locher@ch.abb.com (T. Locher), yvonne-anne.pignolet@ch.abb.com (Y.A. Pignolet), wattenhofer@ethz.ch (R. Wattenhofer).

[1] Part of this work was done at ETH Zurich. At MIT the author was supported by the following grants: AFOSR Contract Number FA9550-13-1-0042, NSF Award 0939370-CCF, NSF Award CCF-1217506, NSF Award number CCF-AF-0937274.

http://dx.doi.org/10.1016/j.jcss.2017.02.006

Please cite this article in press as: S. Holzer et al., Deterministic multi-channel information exchange, J. Comput. Syst. Sci. (2017), http://dx.doi.org/10.1016/j.jcss.2017.02.006

Moreover, in various communication systems, such as Bluetooth or IEEE 802.15.4, several non-conflicting communication channels are available, which can be leveraged to disseminate information. While there is a large body of work on scheduling message transmissions for various models of a single communication channel, surprisingly little is known about the benefits and limits of using multiple channels for the purpose of information dissemination. This is the focus of this article, which addresses the question of how many communication channels are required in order to solve an information exchange problem as quickly as possible. More generally, we study the power of having additional channels at one's disposal when trying to disseminate information. We believe that this is an important missing piece in the study of communication over shared channels. Before giving a more formal definition of the considered information exchange problem, we present the communication model used throughout this article.

1.1. Model

In this article, we consider a simple network topology, the complete (single-hop) communication network in which every node can communicate with every other node. There are n nodes in total, each with a given, unique identifier in the range $[n] := \{1, \ldots, n\}$. We assume that multiple channels are available for communication and that local computations require zero time (since we focus on communication complexity). Additionally, we make the simplifying assumption that time is divided into synchronized time slots, i.e., we study slotted protocols: In any time slot, each node v may choose a channel i and perform exactly one of two operations, either `send`, which means that v *broadcasts* a message on channel i or `receive`, in which case v *listens* on channel i.[2] A transmission is *successful* if and only if exactly one node transmits its message on a given channel in a specific time slot. A node listening on a particular channel i only receives a message in a given time slot if there is a successful transmission on this channel. Messages are of bounded size, i.e., we assume that each message can only contain one information item (e.g., a node identifier). We further assume that there is no *collision detection*, i.e., if a node v does not receive a message when listening on a channel i, node v cannot determine whether there was a collision or no message was sent. This is a reasonable assumption as, e.g., simple wireless devices often do not have a reliable collision detection mechanism. Moreover, solutions in this model can be applied in settings with collision detection but not vice versa. We study the following problem.

Definition 1 *(Information Exchange Problem (a.k.a. Multi-Message Broadcast)).* There is an arbitrary subset of $k \leq n$ nodes (called *reporter nodes* or simply *reporters*) where each of the k nodes is given a distinct piece of information. This subset is determined by an adversary before the first time slot. The objective is to disseminate these k information items to every node in the network. The subset of reporters is not known initially. The number n of nodes and the number k of reporters may or may not be known.

This problem lies between two fundamental information dissemination problems: broadcasting (one-to-all communication) and gossiping (total information exchange). In other words, we generalize the *Information Exchange Problem* [18] (also known as *k-Selection* [22], *Multi-Message Broadcast* [21] and *Many-to-All Communication* [9]) for networks with several communication channels. In order to measure the quality of a solution to the Information Exchange Problem, we must define adequate complexity measures. Clearly, it takes a certain number of time slots to distribute all information items. As mentioned before, the goal is to disseminate all information items as quickly as possible. Therefore, the primary objective pursued in this article is to find an algorithm $\mathcal{A}$ with an optimal *time complexity*, which is defined as the maximum number of synchronous time slots that $\mathcal{A}$ requires to disseminate all k items for a worst-case selection of reporters. Since only one information item can be transmitted in any message, i.e., items cannot be bundled, and each node can only listen on one channel per time slot, it follows that the time complexity of any algorithm is at least $\Omega(k)$. The key question thus becomes how many channels does an algorithm for the Information Exchange Problem require in order to achieve an asymptotically optimal time complexity of $\Theta(k)$? Chlebus and Kowalski [8] prove that it is not possible to disseminate all information items in time $\mathcal{O}(k)$ with only one communication channel by giving a lower bound of $\Omega(k + \log n)$. If more channels are available, the lower bound $\Omega(k)$ can be matched using randomized algorithms [20]. However, these algorithms need a large number of channels and there is a (small) probability that these algorithms fail.

1.2. Contributions

We present deterministic algorithms for the Information Exchange Problem when n and k are known. In particular, we introduce two algorithms both exhibiting an asymptotically optimal time complexity of $\Theta(k)$, which are appropriate for different values of k, and give bounds on the number of channels that each algorithm requires for a given range of values k

---

[2] Naturally, a node may also choose not to perform any operation in a given time slot. For scenarios that take the energy efficiency into account this is a very important aspect.

Please cite this article in press as: S. Holzer et al., Deterministic multi-channel information exchange, J. Comput. Syst. Sci. (2017), http://dx.doi.org/10.1016/j.jcss.2017.02.006