ALGOSENSORS 2016

ALGO 2016, Aarhus, Denmark, August 2016

Marek Chrobak, Antonio Fernández Anta, Leszek Gąsieniec, and Ralf Klasing

August 25, 2016

ALGOSENSORS is an international symposium dedicated to the algorithmic aspects of wireless networks. Originally focused on sensor networks, it now covers algorithmic issues arising in wireless networks of all types of computational entities, static or mobile, including sensor networks, sensor-actuator networks, autonomous robots. The focus is on the design and analysis of algorithms, models of computation, and experimental analysis.

In year 2016 ALGOSENSORS has one invited talk and 9 regular presentations.

Programme Committee Chairs
Marek Chrobak, University of California Riverside, USA
Antonio Fernádez Anta, Institute MDEA Networks, Spain
Leszek Gąsieniec, University of Liverpool, UK (Program Chair)
Ralf Klasing, LaBRI - Université Bordeaux, France

Programme Committee
Fernando Boavida, University of Coimbra, Portugal
Costas Busch, Lousiana State University, USA
Bogdan Chlebus, University of Colorado Denver, USA
Marek Chrobak, University of California Riverside, USA
Jurek Czyzowicz, Universite du Quebec en Outaouais, Canada
Robert Elsässer, University of Salzburg, Austria
Antonio Fernádez Anta, Institute IMDEA Networks, Spain
Pierre Fraigniaud, CNRS and Université Paris Diderot, France
Leszek Gąsieniec, University of Liverpool, UK
Magnús M. Halldórsson, Reykjavik University, Iceland
Ralf Klasing, LaBRI - Université Bordeaux, France
Adrian Kosowski, INRIA and Université Paris Diderot, France
Evangelos Kranakis, Carleton University, Canada
Danny Krizanc, Wesleyan University, USA
Vincenzo Mancuso, Institute IMDEA Networks, Spain
Miguel Mosteiro, Kean University, USA
Tomasz Radzik, King's College London, UK
Gianluca Rizzo, HES SO Valais, Switzerland
Christian Scheideler, Universität Paderborn, Germany
Christopher Thraves, Universidad de Concepcion, Chile

Steering Committee

Josep Diaz, U.P. Catalunya, Spain
Magnús M. Halldórsson, Reykjavik U., Iceland
Bhaskar Krishnamachari, U. of Southern California, USA
P.R. Kumar, Texas A&M U., USA
Sotiris Nikoletseas, U. of Patras and CTI, Greece (Conference Chair)
Jose Rolim, U. of Geneva, Switzerland
Paul Spirakis, U. of Patras and CTI, Greece
Adam Wolisz, T.U. Berlin, Germany

Programme of Talks

Thursday 25 August 2016

WIRELESS (09:00-10:15) Chair: Annalisa De Bonis

– *Mohamad Ahmadi* and Fabian Kuhn*, Multi-Message Broadcast in Dynamic Radio Networks

– *Kokouvi Hounkanli, Avery Miller*, and Andrzej Pelc*, Global Synchronization and Consensus Using Beeps in a Fault-Prone MAC

– *Hicham Lakhlef, Michel Raynal and Francois Taiani**, Vertex Coloring with Communication and Local Memory Constraints in Synchronous Broadcast Networks

INVITED TALK (10:35-11:25) Chair: Leszek Gąsieniec

– *Fabian Kuhn**, Developing Robust Wireless Network Algorithms

WIRELESS & SEARCH (14:00-15:15) Chair: Dror Rawitz

– *Annalisa De Bonis* and Ugo Vaccaro*, A New Kind of Selectors and their Applications to Conflict Resolution in Wireless Multichannels Networks

– *Shouwei Li, Friedhelm Meyer Auf der Heide and Pavel Podlipyan**, The impact of the Gabriel subgraph of the visibility graph on the gathering of mobile autonomous robots

– *Konstantinos Georgiou, George Karakostas* and Evangelos Kranakis* Search-and-Fetch with One Robot on a Disk

APPROXIMATION (15:30-16:45) Chair: George Karakostas

– *Robert Benkoczi, Daya Gaur and Mark Thom**, A 2-approximation algorithm for barrier coverage by weighted non-uniform sensors on a line

– *Ariella Voloshin and Dror Rawitz**, Flexible Cell Selection in Cellular Networks

– *Manjanna B, Ramesh Jallu*, Gautam K Das and Subhas C Nandy*, The Euclidean k-Supplier Problem in $\mathbb{R}^2$

*The relevant speakers.

Abstracts

1.1 Session – Wireless

09:00-09:25 (25 minutes)

Multi-Message Broadcast in Dynamic Radio Networks

– *Mohamad Ahmadi and Fabian Kuhn*

Abstract We continue the recent line of research studying information dissemination problems in adversarial dynamic radio networks. We give two generic algorithms which allow to transform generalized version of single-message broadcast algorithms into multi-message broadcast algorithms. Based on these generic algorithms, we obtain multi-message broadcast algorithms for dynamic radio networks for a number of different dynamic network settings. For one of the modeling assumptions, our algorithms are complemented by a lower bound which shows that the upper bound is close to optimal.

09:25-09:50 (25 minutes)

Global Synchronization and Consensus Using Beeps in a Fault-Prone MAC

– *Kokouvi Hounkanli, Avery Miller, and Andrzej Pelc*

Abstract Global synchronization is an important prerequisite to many distributed tasks. Communication between processors proceeds in synchronous rounds. Processors are woken up in possibly different rounds. The clock of each processor starts in its wakeup round showing local round 0, and ticks once per round, incrementing the value of the local clock by one. The global round 0, unknown to processors, is the wakeup round of the earliest processor. Global synchronization (or establishing a global clock) means that each processor chooses a local clock round such that their chosen rounds all correspond to the same global round t.

We study the task of global synchronization in a *Multiple Access Channel* (MAC) prone to faults, under a very weak communication model called the *beeping model*. Some processors wake up spontaneously, in possibly different rounds decided by an adversary. In each round, an awake processor can either listen, i.e., stay silent, or beep, i.e., emit a signal. In each round, a fault can

occur in the channel independently with constant probability $0 < p < 1$. In a fault-free round, an awake processor hears a beep if it listens in this round and if one or more other processors beep in this round. A processor still dormant in a fault-free round in which some other processor beeps is woken up by this beep and hears it. In a faulty round nothing is heard, regardless of the behaviour of the processors. An algorithm working with error probability at most ϵ, for a given $\epsilon > 0$, is called *ϵ-safe*. Our main result is the design and analysis, for any constant $\epsilon > 0$, of a deterministic ϵ-safe global synchronization algorithm that works in constant time in any fault-prone MAC using beeps.

As an application, we solve the consensus problem in a fault-prone MAC using beeps. Processors have input values from some set V and they have to decide the same value from this set. If all processors have the same input value, then they must all decide this value. Using global synchronization, we give a deterministic ϵ-safe consensus algorithm which works in time $O(\log w)$ in a fault-prone MAC, where w is the smallest input value of all participating processors. We show that this time cannot be improved, even when the MAC is fault-free.

09:50-10:15 (25 minutes)

Vertex Coloring with Communication and Local Memory Constraints in Synchronous Broadcast Networks

– *Hicham Lakhlef, Michel Raynal and Francois Taiani*

Abstract The vertex coloring problem has been extensively studied in the context of synchronous round-based systems where, at each round, a process can send a message to all its neighbors, and receive a message from each of them. Hence, this communication model is particularly suited to point-to-point wired communication channels. Several vertex coloring algorithms suited to these systems have been proposed. They differ mainly in the number of rounds they require and the number of colors they use.

This paper considers a different broadcast/receive communication model in which message collisions and message conflicts can occur because processes share frequency bands. (A collision occurs when, during the same round, messages are sent to the same process by too many neighbors; a conflict occurs when a process and one of its neighbors broadcast during the same round.) This communication model is suited to systems where processes share communication bandwidths. More precisely, the paper considers the case where, during a round, a process may either broadcast a message to its neighbors or receive a message from at most m of them. This captures communication-related constraints or a local memory constraint stating that, whatever the number of neighbors of a process, its local memory allows it to receive and store at most

m messages during each round. The paper defines first the corresponding generic vertex multi-coloring problem (a vertex can have several colors). It focuses then on tree networks, for which it presents a lower bound on the number of colors K that are necessary (namely, $K = \lceil \frac{\Delta}{m} \rceil + 1$, where Δ is the maximal degree of the communication graph), and an associated coloring algorithm, which is optimal with respect to K.

1.2 Session – Invited Talk

10:35-11:25 (50 minutes) Developing Robust Wireless Network Algorithms

– *Fabian Kuhn*

Abstract Over the last 30 years, we have seen a tremendous effort to develop distributed algorithms and abstract models to deal with the characteristic properties of wireless communication have been proposed. The models range from simple graph-based characterizations of interference to more accurate physical models such as the so-called signal-to-noise-and-interference (SINR) model.

As different as the typically considered models may be, most of them have one thing in common. Whether a node can successfully receive (and decode) a message is determined using some fixed, deterministic rule that depends on the structure of the network and some additional model parameters.

While in classical wired networks, assuming reliable communication might be a reasonable abstraction, this seems much more problematic in a wireless network setting. The propagation of a wireless signal depends on many diverse environmental factors and it does not seem to be realistic to explicitly model all of these factors or to exactly measure the properties of the wireless communication channels. In addition, the environmental factors might change over time and there can also be additional independent sources of signal interference that cannot be predicted or controlled by the network. Further, wireless devices might also be mobile so that we not only have unreliable communication channels, but potentially even almost arbitrary dynamically changing network topologies. Because the classic abstract wireless communication models do not capture such unpredictable behavior, many existing radio network algorithms might only work in the idealized formal setting for which they were developed.

In my talk, I describe ways to develop more robust wireless network algorithms. I will in particular show that complicated, unstable, and unreliable behavior of wireless communication can be modeled by adding a non-deterministic component to existing radio network models. As a result, any behavior which is too complex or impossible to predict is determined by an adversary.

Clearly, such models lead to less efficient algorithms. However, they also lead to more robust algorithms which tend to work under a much wider set of underlying assumptions. Very often, such models also lead to much simpler algorithms. I will discuss several existing results and I will sketch some general ideas and possible directions for dealing with adversarial uncertainty and more generally dynamic wireless networks.

1.3 Session – Wireless & Search

14:00-14:25 (25 minutes)

A New Kind of Selectors and their Applications to Conflict Resolution in Wireless Multichannels Networks

– *Annalisa De Bonis and Ugo Vaccaro*

Abstract We investigate the benefits of using multiple channels of communications in wireless networks under the full-duplex multi-packet reception model of communication. The main question we address is the following: Is a speedup linear in the number of channels achievable, for some interesting communication primitive? We provide a positive answer to this interrogative for the *Information Exchange Problem*, in which k arbitrary nodes have information they intend to share with the entire network. To achieve this goal, we devise and exploit a combinatorial structure that generalizes well known combinatorial tools widely used in the area of data-exchange in multiple access channels (i.e., strongly selective families, selectors, and related mathematical objects). For our new combinatorial structures we provide both existential results, based on the Lovász Local Lemma, and efficient constructions, leveraging on properties of error correcting codes. We also prove non existential results, showing that our constructions are not too far from being optimal. Our main results are the following:

Theorem 1 *Let k, q, and n be integers such that $1 \leq k \leq n$, and $2 \leq q \leq k$. There exists a conflict resolution algorithm for a multiple-access channel $\mathcal{C}$ without feedback (comprising of $q \geq 2$ individual channels) that schedules the transmissions of n stations in such a way that, for all possible subsets of up to k active stations, one has that all of them transmit successfully. The number of time slots t used by the conflict resolution algorithm is $O(\frac{k^2}{q} \log \frac{n}{k})$, and we can prove that for any conflict resolution algorithm for the Information Exchange Problem it holds that the number of time slots t is such that $\Omega(\frac{k^2}{q \log k} \log \frac{n}{k})$.*

Theorem 2 *Let k, q, and n be integers such that $2 \leq q \leq k \leq n$. There exists a conflict resolution algorithm for a multiple-access channel $\mathcal{C}$ with feedback (comprising of $q \geq 2$ individual channels) that schedules the transmissions of n stations in such a way that, for all possible subsets of k*

stations, one has that all active stations transmit successfully. The number t of time slots used by the conflict resolution algorithm is $O(\frac{k}{q}\log\frac{n}{k})$. Moreover, for any conflict resolution algorithm in this scenario it holds that the number of time slots t is $\Omega(\frac{k}{q}\log\frac{n}{k})$.

We remark that the asymptotic upper bound of Theorem 2 holds also in the case when there is no a priori knowledge of the number k of active stations. In this work we also highlight a few interesting connections between the combinatorial structures introduced in this paper and the well known Frameproof Codes. Moreover, we also show the remarkable fact that, for an infinite set of the relevant parameters n and k, one can construct our combinatorial structures in polynomial time and of optimal (minimum) length.

14:25-14:50 (25 minutes)

The impact of the Gabriel subgraph of the visibility graph on the gathering of mobile autonomous robots

– *Shouwei Li, Friedhelm Meyer Auf der Heide and Pavel Podlipyan*

Abstract In this paper, we reconsider the well-known Go-To-The-Center algorithm for gathering in the plane n autonomous mobile robots with limited viewing range. This is a discrete, round-based algorithm that gathers group pf robots in $\Theta(n^2)$ rounds. Remarkably, this algorithm exploits the fact, that during its execution, many collisions of robots occur. Such collisions are interpreted as a success because it is assumed that such collided robots behave the same from now on. This is o.k. under the assumption, those robots have no extent. Otherwise, collisions should be avoided.

In this paper, we consider a continuous Go-To-The-Center (GTC) strategy in which the robots continuously observe the positions of their neighbors and adapt their speed (assuming a speed limit) and direction. Our first results are time bounds of $O(n^2)$ for gathering in two-dimensional Euclidean space, and $\Theta(n)$ for the one-dimensional case.

Our main contribution is the introduction and evaluation of a continuous algorithm which performs Go-To-The-Center considering only the neighbors of a robots w.r.t. the Gabriel subgraph of the visibility graph (GTGC). We show that this modification still correctly executes gathering in one and two dimensions, with the same time bounds as above. Simulations exhibit a severe difference of the behavior of the GTC and the GTGC strategy: Whereas lots of collisions occur during a run of the GTC strategy, typically only one, namely the final collision occurs during a run of the GTGC strategy. We can prove this "collisionless property" of the GTGC algorithm for the one-dimensional case. In the case of the two-dimensional Euclidean space, we conjecture that the "collisionless property" of the GTGC holds for almost every initial configuration.

14:50-15:15 (25 minutes)

Search-and-Fetch with One Robot on a Disk

– *Konstantinos Georgiou, George Karakostas and Evangelos Kranakis*

Abstract A robot is located at a point in the plane. A treasure and an exit, both stationary, are located at unknown (to the robot) positions both at distance one from the robot. Starting from its initial position, the robot aims to fetch the treasure to the exit. At any time the robot can move anywhere on the disk with constant speed. The robot detects an interesting point (treasure or exit) only if it passes over the exact location of that point. Given that an adversary controls the locations of both the treasure and the exit on the perimeter, we are interested in designing algorithms that minimize the treasure-evacuation time, i.e. the time it takes for the treasure to be discovered and brought to the exit by the robot. In this paper we differentiate how the robot's knowledge of the distance between the two interesting points affects the overall evacuation time. We demonstrate the difference between knowing the exact value of that distance versus knowing only a lower bound and provide search algorithms for both cases. In the former case we give an algorithm which is off from the optimal algorithm (that does not know the locations of the treasure and the exit) by no more than $\frac{4\sqrt{2}+3\pi+2}{6\sqrt{2}+2\pi+2} < 1.019$ multiplicatively, or $\frac{\pi}{2} - \sqrt{2} \leq 0.157$ additively. In the latter case we provide an algorithm which is shown to be optimal.

1.4 Session – Approximation

15:30-15:55 (25 minutes)

A 2-approximation algorithm for barrier coverage by weighted non-uniform sensors on a line

– *Robert Benkoczi, Daya Gaur and Mark Thom*

Abstract Barrier coverage is an approach to the intruder detection problem that relies on monitoring a perimeter, or barrier, of an area of interest using sensors placed around it. Czyzowicz *et al.* [2], [3] were among the first to consider barrier coverage problems on line segments using relocatable sensors. They cast the problem under several different objective functions, among them the MinSum problem, which seeks to minimize the total distance travelled by all sensors to their positions in the coverage. They showed that a constant factor approximation algorithm for the MinSum problem cannot exist unless $P = NP$.

In [1], the authors restrict the problem to that in which the initial position and covering range of each sensor is disjoint from the line segment to be covered. There, they develop an FPTAS

for the disjoint version of the problem, which they term DisjointMinSum.

Here, we consider a weighted version of the problem, where the movement of each sensor is multiplied by a positive weight assigned to it. We develop a 2-approximation algorithm for the weighted disjoint barrier coverage problem, and use it to extend the unweighted FPTAS of [1] to the weighted problem. We show this can be done with a quadratic factor improvement in time complexity over the unweighted version.

References

[1] Benkoczi, R., Friggstad, Z., Gaur, D., and Thom, M. Minimizing total sensor movement for barrier coverage by non-uniform sensors on a line. In *Algorithms for Sensor Systems*. Springer, 2015, pp. 98–111.

[2] Czyzowicz, J., Kranakis, E., Krizanc, D., Lambadaris, I., Narayanan, L., Opatrny, J., Stacho, L., Urrutia, J., and Yazdani, M. On minimizing the sum of sensor movements for barrier coverage of a line segment. In sl Ad-Hoc, Mobile and Wireless Networks. Springer, 2010, pp. 29–42.

[3] Czyzowicz, J., Kranakis, E., Krizanc, D., Lambadaris, I., Narayanan, L., Stacho, L., Urrutia, J., and Yazdani, M. On minimizing the maximum sensor movement for barrier coverage of a line segment. In sl Proceedings of 8th International Conference on Ad Hoc Networks and Wireless (2002), pp. 22–25.

15:55-16:20 (25 minutes)

Flexible Cell Selection in Cellular Networks

– *Ariella Voloshin and Dror Rawitz*

Abstract We introduce the problem of Flexible Scheduling on Related Machines with Assignment Restrictions (FSRM). In this problem the input consists of a set of machines and a set of jobs. Each machine has a finite capacity, and each job has a resource requirement interval, a profit per allocated unit of resource, and a set of machines that can potentially supply the requirement. A feasible solution is an allocation of machine resources to jobs such that: (i) a machine resource can be allocated to a job only if it is a potential supplier of this job, (ii) the amount of machine resources allocated by a machine is bounded by its capacity, and (iii) the amount of resources that are allocated to a job is either in its requirement interval or zero. Notice that a job can be serviced by multiple machines. The goal is to find a feasible allocation that maximizes the overall profit. We focus on r-FSRM in which the required resource of a job is at most an r-fraction of (or r times) the capacity of each potential machine. FSRM is

motivated by resource allocation problems arising in cellular networks and in cloud computing. Specifically, FSRM models the problem of assigning clients to base stations in 4G cellular networks.We present a 2-approximation algorithm for 1-FSRM and a $\frac{1}{1-r}$-approximation algorithm for r-FSRM, for any $r \in (0, 1)$. Both are based on the local ratio technique and on maximum flow computations. We also present an LP-rounding 2-approximation algorithm for a flexible version of the GENERALIZED ASSIGNMENT PROBLEM that also applies to 1-FSRM. Finally, we give an $\Omega(\frac{r}{\log r})$ lower bound on the approximation ratio for r-FSRM (assuming $P \neq NP$).

16:20-16:45 (25 minutes)

The Euclidean k-Supplier Problem in $I\!R^2$

– *Manjanna B, Ramesh Jallu, Gautam K Das and Subhas C Nandy*

Abstract In this paper, we consider k-supplier problem in $I\!R^2$. Here, two sets of points P and Q are given. The objective is to choose a subset $Q_{opt} \subseteq Q$ of size at most k such that congruent disks of minimum radius centered at the points in Q_{opt} cover all the points of P. We propose a fixed-parameter tractable (FPT) algorithm for the k-supplier problem that produces a 2-factor approximation result. For $|P| = n$ and $|Q| = m$, the worst case running time of the algorithm is $O(6^k(n + m) \log(mn))$, which is an exponential function of the parameter k. We also propose a heuristic algorithm based on Voronoi diagram for the k-supplier problem, and experimentally compare the result produced by this algorithm with the best known approximation algorithm available in the literature [Nagarajan, V., Schieber, B., Shachnai, H.: The Euclidean k-supplier problem, In Proc. of 16th Int. Conf. on Integ. Prog. and Comb. Optim., 290–301 (2013)]. The experimental results show that our heuristic algorithm is slower than Nagarajan et al.'s $(1+\sqrt{3})$-approximation algorithm, but the results produced by our algorithm significantly outperforms that of Nagarajan et al.'s algorithm.

Contents