

Track C

1. **Best paper in Track C**

Dariusz Dereniowski, Yann Disser, Adrian Kosowski, Dominik Pajak and Przemysław Uznański. *Fast Collaborative Graph Exploration*

Abstract: We study the following scenario of online graph exploration. A team of k agents is initially located at a distinguished vertex r of an undirected graph. In each time step, each agent can traverse an edge of the graph. All vertices have unique identifiers, and upon entering a vertex, an agent obtains a list of the identifiers of all its neighbors. We ask how many time steps are required to complete exploration, i.e., to make sure that every vertex has been visited by some agent.

We consider two communication models: one in which all agents have global knowledge of the state of the exploration, and one in which agents may only exchange information when simultaneously located at the same vertex. As our main result, we provide the first strategy which performs exploration of a graph with n vertices at a distance of at most D from r in time $O(D)$, using a team of agents of polynomial size $k = D n^{1 + \epsilon} < n^{2 + \epsilon}$, for any $\epsilon > 0$. Our strategy works in the local communication model, without knowledge of global parameters such as n or D .

We also obtain almost-tight bounds on the asymptotic relation between exploration time and team size, for large k . For any constant $c > 1$, we show that in the global communication model, a team of $k = D n^c$ agents can always complete exploration in $D(1 + \frac{1}{c-1} + o(1))$ time steps, whereas at least $D(1 + \frac{1}{c} - o(1))$ steps are sometimes required. In the local communication model, $D(1 + \frac{2}{c-1} + o(1))$ steps always suffice to complete exploration, and at least $D(1 + \frac{2}{c} - o(1))$ steps are sometimes required. This shows a clear separation between the global and local communication models.

2. **George Mertzios and Paul Spirakis.** *Strong Bounds for Evolution in Networks*

Abstract: This work extends what is known so far for a basic model of evolutionary antagonism in undirected networks (graphs). More specifically, this work studies the generalized Moran process, as introduced by Lieberman, Hauert, and Nowak [Nature, 433:312-316, 2005], where the individuals of a population reside on the vertices of an undirected connected graph. The initial population has a single *mutant* of a *fitness* value r (typically $r > 1$), residing at some vertex v of the graph, while every other vertex is initially occupied by an individual of fitness 1 . At every step of this process, an individual (i.e. vertex) is randomly chosen for reproduction with probability proportional to its fitness, and then it places a copy of itself on a random neighbor, thus replacing the individual that was residing there. The main quantity of interest is the *fixation probability*, i.e. the probability that eventually the whole graph is occupied by descendants of the mutant. In this work we concentrate on the fixation probability when the mutant is initially on a specific vertex v , thus refining the older notion of Lieberman et al. which studied the fixation probability when the initial mutant is placed at a random vertex. We then aim at finding graphs that have many "strong starts" (or many "weak starts") for the mutant. Thus we introduce a parameterized notion of *selective amplifiers* (resp. *selective suppressors*) of evolution, i.e. graphs with at least some $h(n)$ vertices

(starting points of the mutant), which fixate the graph with large (resp. with small) probability. We prove the existence of *strong* selective amplifiers (i.e. for $h(n)=\Theta(n)$ vertices v the fixation probability of v is at least $1-\frac{c(r)}{n}$ for a function $c(r)$ that depends only on r). We also prove the existence of quite strong selective suppressors. Regarding the traditional notion of fixation probability from a random start, we provide strong upper and lower bounds: first, we demonstrate the non-existence of "strong universal" amplifiers, i.e. we prove that for any undirected graph the fixation probability from a random start is at most $1-\frac{c(r)}{n^{3/4}}$. Finally we prove the *Thermal Theorem*, which states that for any undirected graph, when the mutant starts at vertex v , the fixation probability is at least $(r-1) / (r+\frac{\deg v}{\deg_{\min}})$. This implies a lower bound for the usual notion of fixation probability, which is almost tight. This theorem extends the "Isothermal Theorem" of Lieberman et al. for regular graphs. Our proof techniques are original and are based on new domination arguments which may be of general interest in Markov Processes that are of the general birth-death type.

3. **Yoram Bachrach and Ely Porat.** *Sketching For Big Data Recommender Systems Using Fast Pseudo-Random Fingerprints*

Abstract: A key building block for collaborative filtering based recommender systems is identifying users with similar consumption patterns. Given access to the full data regarding the items consumed by each user, one can directly compute the similarity between any two users. However, for massive recommender systems such a naive approach requires a high running time and may be intractable in terms of the space required to store the full data. One way to overcome this is using sketching, a technique that represents massive datasets concisely, while still allowing calculating properties of these datasets. Sketching methods maintain very short fingerprints of the item sets of users, which allow approximately computing the similarity between sets of different users.

The state of the art sketch (Li and Koenig) has a very low space complexity, and a recent technique (Feigenblat, Shiftan, Porat) shows how to exponentially speed up the computation time involved in building the fingerprints. Unfortunately, these methods are incompatible, forcing a choice between low running time or a small sketch size. We propose an alternative sketching approach, which achieves both a low space complexity similar to that of (Li and Koenig) and a low time complexity similar to (Feigenblat, Shiftan, Porat).

We empirically evaluate our algorithm using the Netflix dataset. We analyze the running time and the sketch size using our approach and compare them to alternative approaches. Further, we show that in practice the accuracy achieved by our approach is even better than the accuracy guaranteed by the theoretical bounds, so it suffices to use even shorter fingerprints to obtain high quality results.

4. **Yoann Dieudonne and Andrzej Pelc.** *Deterministic Polynomial Approach in the Plane*

Abstract: Two mobile agents with range of vision 1 start at arbitrary points in the plane and have to accomplish the task of approach, which consists in getting at distance at most one from each other, i.e., in getting within each other's range of vision. An adversary chooses the initial positions of the agents, their possibly different starting times, and assigns a different positive integer label and a possibly different speed to each of them. Each agent is equipped with a compass showing the cardinal directions, with a measure of length and a clock. Each

agent knows its label and speed but not those of the other agent and it does not know the initial position of the other agent relative to its own. Agents do not have any global system of coordinates and they cannot communicate. Our main result is a deterministic algorithm to accomplish the task of approach, working in time polynomial in the unknown initial distance between the agents, in the length of the smaller label and in the inverse of the larger speed. The distance traveled by each agent until approach is polynomial in the first two parameters and does not depend on the third. The problem of approach in the plane reduces to a network problem: that of rendezvous in an infinite grid.

5. **Susanne Albers and Achim Passen.** *New Online Algorithms for Story Scheduling in Web Advertising*

Abstract: We study *storyboarding* where advertisers wish to present sequences of ads (stories) uninterruptedly on a major ad position of a web page. These jobs/stories arrive online and are triggered by the browsing history of a user who at any time continues surfing with probability β . The goal of an ad server is to construct a schedule maximizing the expected reward. The problem was introduced by Dasgupta, Ghosh, Nazerzadeh and Raghavan (SODA'09) who presented a 7-competitive online algorithm. They also showed that no deterministic online strategy can achieve a competitiveness smaller than 2, for general β .

We present improved algorithms for storyboarding. First we give a simple online strategy that achieves a competitive ratio of $4/(2-\beta)$, which is upper bounded by 4 for any β . The algorithm is also $1/(1-\beta)$ -competitive, which gives better bounds for small β . As the main result of this paper we devise a refined algorithm that attains a competitive ratio of $c=1+\phi$, where $\phi = (1+\sqrt{5})/2$ is the Golden Ratio. This performance guarantee of $c \approx 2.618$ is close to the lower bound of 2. Additionally, we study for the first time a problem extension where stories may be presented simultaneously on several ad positions of a web page. For this parallel setting we provide an algorithm whose competitive ratio is upper bounded by $1/(3-2\sqrt{2}) \approx 5.828$, for any β . All our algorithms work in phases and have to make scheduling decisions only every once in a while.

6. **Emanuele Guido Fusco, Andrzej Pelc and Rossella Petreschi.** *Learning a ring cheaply and fast*

Abstract: We consider the task of learning a ring in a distributed way: each node of an unknown ring has to construct a labeled map of it. Nodes are equipped with unique labels. Communication proceeds in synchronous rounds. In every round every node can send arbitrary messages to its neighbors and perform arbitrary local computations. We study tradeoffs between the time (number of rounds) and the cost (number of messages) of completing this task in a deterministic way: for a given time T we seek bounds on the smallest number of messages needed for learning the ring in time T . Our bounds depend on the diameter D of the ring and on the delay $\theta = T - D$ above the least possible time D in which this task can be performed. We prove a lower bound $\Omega(D^2/\theta)$ on the number of messages used by any algorithm with delay θ , and we design a class of algorithms that give an almost matching upper bound: for any positive constant $0 < \epsilon < 1$ there is an algorithm working with delay $\theta \leq D$ and using $O(D^2 (\log^* D)/\theta^{1-\epsilon})$ messages.

7. **Luca Becchetti, Vincenzo Bonifaci, Michael Dirnberger, Andreas Karrenbauer and Kurt Mehlhorn.** *Physarum Can Compute Shortest Paths: Convergence Proofs and Complexity Bounds*

Abstract: Physarum Polycephalum is a slime mold that is apparently able to solve shortest path problems. A mathematical model for the slime's behavior in the form of a coupled system of differential equations was proposed in~\cite{Tero:2007}. We prove that a discretization of the model (Euler integration) computes a $(1 + \epsilon)$ -approximation of the shortest path within $O(m L (\log n + \log L) / \epsilon^3)$ iterations, with arithmetic on numbers of $O(\log(nL/\epsilon))$ bits. We also obtain two results for the directed Physarum model of~\cite{Ito:2011}: convergence of a nonuniform version of the model and convergence of and complexity bounds for the discretization of the uniform model.

8. **Martin Hoefer and Lisa Wagner.** *Locally Stable Marriage with Strict Preferences*

Abstract: We study two-sided matching markets with locality of information and control. Each male (female) agent has an arbitrary strict preference list over all female (male) agents. In addition, each agent is a node in a fixed network. Agents learn about possible partners dynamically based on their current network neighborhood. We consider convergence of dynamics to locally stable matchings that are stable with respect to their imposed information structure in the network. While existence of such states is guaranteed, we show that reachability becomes NP-hard to decide. This holds even when the network exists only among one side. In contrast, if only one side has no network and agents remember a previous match every round, reachability is guaranteed and random dynamics converge with probability 1. We characterize this positive result in various ways. For instance, it holds for random memory and for memory with the most recent partner, but not for memory with the best partner. Also, it is crucial which partition of the agents has memory. Finally, we conclude with results on approximating maximum locally stable matchings.

9. **Sebastian Faust, Carmit Hazay and Daniele Venturi.** *Outsourced Pattern Matching*

Abstract: In secure delegatable computation, computationally weak devices (or clients) wish to outsource their computation and data to an untrusted server in the cloud. While most earlier work considers the general question of how to securely outsource any computation to the cloud server, we focus on concrete and important functionalities and give the first protocol for the pattern matching problem in the cloud. Loosely speaking, this problem considers a text T that is outsourced to the cloud S by a client C_T . In a query phase, clients C_1, \dots, C_l run an efficient protocol with the server S and the client C_T in order to learn the positions at which a pattern of length m matches the text (and nothing beyond that). This is called the outsourced pattern matching problem and is highly motivated in the context of delegatable computing since it offers storage alternatives for massive databases that contain confidential data (e.g., health related data about patient history). Our constructions offer simulation-based security in the presence of semi-honest and malicious adversaries (in the random oracle model) and limit the communication in the query phase to $O(m)$ bits plus the number of occurrences -- which is optimal. In contrast to generic solutions for delegatable computation, our schemes do not rely on fully homomorphic encryption but instead uses novel ideas for solving pattern matching, based on a reduction to the subset sum problem. Interestingly, we do not rely on the hardness of the problem, but rather we exploit instances that are solvable in polynomial-time.

10. **George Mertzios, Othon Michail, Ioannis Chatzigiannakis and Paul Spirakis.** *Temporal Network Optimization Subject to Connectivity Constraints*

Abstract: In this work we consider temporal networks, i.e. networks defined by a labeling λ assigning to each edge of an underlying graph G a set of discrete time-labels. The labels of an edge, which are natural numbers, indicate the discrete time moments at which the edge is available. We focus on path problems of temporal networks. In particular, we consider time-respecting paths, i.e. paths whose edges are assigned by λ a strictly increasing sequence of labels. We begin by giving two efficient algorithms for computing shortest time-respecting paths on a temporal network. We then prove that there is a natural analogue of Menger's theorem holding for arbitrary temporal networks. Finally, we propose two cost minimization parameters for temporal network design. One is the temporality of G , in which the goal is to minimize the maximum number of labels of an edge, and the other is the temporal cost of G , in which the goal is to minimize the total number of labels used. Optimization of these parameters is performed subject to some connectivity constraint. We prove several lower and upper bounds for the temporality and the temporal cost of some very basic graph families such as rings, directed acyclic graphs, and trees.

11. **Nikolay Gravin and Pinyan Lu.** *Competitive Auctions for Markets with Positive Externalities*

Abstract: In the digital good auctions the auctioneer sells an item in unlimited supply to a set of potential buyers. The objective is to design a truthful auction that maximizes the total profit of the auctioneer. Motivated from the observation that the valuations for the good of the buyers might be interconnected through a social networks, we study digital goods auctions with positive externalities among the buyers. This defines a multi-parameter auction design problem where the private valuation of every buyer is a function of the set of other winning buyers. The main contribution of this paper is a truthful competitive mechanism for subadditive valuations. Our competitive result is with respect to a new solution benchmark $\mathcal{F}^{(3)}$. On the other hand, we show a surprising impossibility result if comparing to the stronger benchmark $\mathcal{F}^{(2)}$, where the latter has been used quite successfully in digital goods auctions without externalities~\cite{Goldberg2006}.

12. **Sepp Hartung, André Nichterlein, Rolf Niedermeier and Ondrej Suchy.** *A Refined Complexity Analysis of Identity Anonymization on Graphs*

Abstract: Motivated by a strongly growing interest in graph anonymization in the data mining and databases communities over the last five years, we study the NP-hard problem to make a graph k -anonymous by adding as few edges as possible. Herein, a graph is k -anonymous if for every vertex in the graph there are at least $k-1$ other vertices of same degree. Our main result is to show that there is a polynomial-size problem kernel for the problem parameterized by the maximum vertex degree. This result is in a sense tight since we also show that the problem is already NP-hard for H-index three, implying NP-hardness for "stronger" parameters such as degeneracy and average degree. Moreover, our algorithmic results shed light on the performance quality of a popular heuristic due to Liu and Terzi [ACM SIGMOD 2008]; in particular, we show that the heuristic provides optimal solutions in case that many edges need to be added.

13. **George Christodoulou and Martin Gairing.** *Price of Stability in Polynomial Congestion Games*

Abstract: The Price of Anarchy in congestion games has attracted a lot of research over the last decade. This resulted in a thorough understanding of this concept. In contrast the Price of Stability, which is an equally interesting concept, is much less understood. In this paper, we consider congestion games with polynomial cost functions with nonnegative coefficients and maximum degree d . We give matching bounds for the Price of Stability in such games, i.e., our technique provides the exact value for any degree d . For linear congestion games, tight bounds were previously known. Those bounds hold even for the more restricted case of dominant equilibria, which may not exist. We give a separation result showing that already for congestion games with quadratic cost functions this is not possible; that is, the Price of Anarchy for the subclass of games that admit a dominant strategy equilibrium is strictly smaller than the Price of Stability for the general class.

14. **Jurek Czyzowicz, Evangelos Kranakis and Eduardo Pacheco.** *Localization For a System of Colliding Robots*

Abstract: We study the *localization problem* in the ring: a collection of n anonymous mobile robots are deployed in a continuous ring of perimeter one. All robots start moving at the same time along the ring with arbitrary velocity, starting in clockwise or counterclockwise direction around the ring. The robots bounce against each other when they meet. The task of each robot is to find out, in finite time, the initial position and the initial velocity of every deployed robot. The only way that robots perceive the information about the environment is by colliding with their neighbors; any type of communication among robots is not possible.

We assume the principle of momentum conservation as well as the conservation of energy, so robots exchange velocities when they collide. The capabilities of each robot are limited to: observing the times of its collisions, being aware of its velocity at any time, and processing the collected information. Robots have no control of their walks or velocities. Robots' walks depend on their initial positions, velocities, and the sequence of collisions. They are not equipped with any visibility mechanism.

The localization problem for bouncing robots has been studied previously in [\cite{us,batonPaper}](#) in which robots are assumed to move at the same velocity. The configuration of initial positions of robots and their speeds is considered *feasible*, if there is a finite time, after which every robot starting at this configuration knows initial positions and velocities of all other robots. Authors of [\cite{us}](#) conjectured that if robots have arbitrary velocities, the problem might be solvable, if the momentum conservation and energy preservation principles are assumed.

In this paper we prove that the conjecture in [\cite{us}](#) is false. We show that the feasibility of any configuration and the required time for solving it under such stronger constraints depend only on the collection of velocities of the robots. More specifically, if $\{v_i\}_{i=1}^n$ are the velocities of a given robot configuration s , we prove that s is feasible if and only if $v_i \neq \text{avg}$ for all $0 \leq i \leq n-1$, where $\text{avg} = \frac{v_0 + \dots + v_{n-1}}{n}$. To figure out the initial positions of all robots no more than $\frac{2}{\min_{0 \leq i \leq n-1} |v_i - \text{avg}|}$ time is required.

15. **Monika Henzinger, Sebastian Krinninger and Danupon Nanongkai.** *Sublinear-Time Maintenance of Breadth-First Spanning Tree in Partially Dynamic Networks*

Abstract: We study the problem of maintaining a breadth-first spanning tree (BFS tree) in partially dynamic distributed networks modeling a sequence of either failures or additions of communication links (but not both). We show $(1+\epsilon)$ -approximation algorithms whose amortized time (over some number of changes) is sublinear in D , the maximum diameter of the network. This breaks the $\Theta(D)$ time bound of recomputing "from scratch".

Our technique also leads to a $(1+\epsilon)$ -approximate incremental algorithm for single-source shortest path (SSSP) in the sequential (usual RAM) model. Prior to our work, the state of the art was the classical exact algorithm of (Even-Shiloach 1981), whose optimality is proven likely, under some assumptions (Roditty-Zwick 2004). Our result is the first to show that, in the incremental setting, this bound can be beaten in certain cases, if a small approximation is allowed.

16. **Helger Lipmaa and Tomas Toft.** *Secure Equality and Greater-Than Tests with Sublinear Online Complexity*

Abstract: Secure multiparty computation (MPC) allows multiple parties to evaluate functions without disclosing the private inputs. Secure comparisons (testing equality and greater-than) are important primitives required by many MPC applications. We propose two equality tests for ℓ -bit values with $O(1)$ online communication that require $O(\ell)$ respectively $O(\kappa)$ total work, where κ is a correctness parameter.

Combining these with ideas of Toft~\cite{PKC2011:Toft}, we obtain (i) a greater-than protocol with sublinear online complexity in the arithmetic black-box model ($O(c)$ rounds and $O(c \cdot \ell^{1/c})$ work online, with $c = \log \ell$ resulting in logarithmic online work). In difference to Toft, we do not assume two mutually incorruptible parties, but $O(\ell)$ offline work is required, and (ii) two greater-than protocols with the same online complexity as the above, but with overall complexity reduced to $O(\log \ell (\kappa + \log \log \ell))$ and $O(c \cdot \ell^{1/c} (\kappa + \log \ell))$; these require two mutually incorruptible parties, but are highly competitive with respect to online complexity when compared to existing protocols.

17. **David G. Harris, Ehab Morsy, Gopal Pandurangan, Peter Robinson and Aravind Srinivasan.** *Efficient Computation of Balanced Structures*

Abstract: Basic graph structures such as maximal independent sets (MIS's) have spurred much theoretical research in distributed algorithms, and have several applications in networking and distributed computing as well. However, the extant (distributed) algorithms for these problems do not necessarily guarantee fault-tolerance or load-balance properties: For example, in a star-graph, the central vertex, as well as the set of leaves, are both MIS's, with the latter being much more fault-tolerant and balanced --- existing distributed algorithms do not handle this distinction. We propose and study "low-average degree" or "balanced" versions of such structures. Interestingly, in sharp contrast to, say, MIS's, it can be shown that checking whether a structure is balanced, will take substantial time. Nevertheless, we are able to develop good sequential and distributed algorithms for such "balanced" versions. We also complement our algorithms with several lower bounds.

18. **Seth Pettie and Hsin-Hao Su.** *Fast Distributed Coloring Algorithms for Triangle-Free Graphs*

Abstract: Vertex coloring is a central concept in graph theory and an important symmetry-breaking primitive in distributed computing. Whereas degree- Δ graphs may require palettes of $\Delta+1$ colors in the worst case, it is well known that the chromatic number of many natural graph classes can be much smaller. In this paper we give new distributed algorithms to find (Δ/k) -coloring in graphs of girth 4 (triangle-free graphs), girth 5, and trees, where k is at most $(\frac{1}{4}-o(1))\ln\Delta$ in triangle-free graphs and at most $(1-o(1))\ln\Delta$ in girth-5 graphs and trees. Specifically, for Δ sufficiently large we can find such a coloring in $O(k + \log^* n)$ time. Moreover, for any Δ we can compute such colorings in roughly logarithmic time for triangle-free and girth-5 graphs, and in $O(\log \Delta + \log_{\Delta} \log n)$ time on trees. As a byproduct, our algorithm shows that the chromatic number of triangle-free graphs is $(4+o(1))\frac{\Delta}{\ln\Delta}$, which improves on Jamall's recent bound of $(67+o(1))\frac{\Delta}{\ln\Delta}$.

19. **L. Elisa Celis, Dimitrios C. Gklezacos and Anna R. Karlin.** *On Revenue Maximization for Agents with Costly Information Acquisition*

Abstract: A prevalent assumption in traditional mechanism design is that buyers know their precise value for an item; however, this assumption is rarely true in practice. In most settings, buyers can "deliberate", i.e., spend money or time, in order to improve their estimate of the item's value. It is known that the deliberative setting is fundamentally different than the classical one, and the desirable properties of a mechanism such as equilibria, revenue maximization, or truthfulness, may no longer hold.

In this paper we introduce a new general deliberative setting in which users have independent private values that are a-priori unknown, but can be learned. We consider the design of dominant-strategy revenue-optimal auctions in this setting. Surprisingly, for a wide class of environments, we show the optimal revenue is attained with a sequential posted price mechanism (SPP). While this result is not constructive, we show how to construct approximately optimal SPPs in polynomial time. We also take first steps towards the design of revenue-optimal Bayes-Nash incentive compatible auctions for a deliberative setting.

20. **Tomasz Jurdzinski, Dariusz Kowalski and Grzegorz Stachowiak.** *Distributed Deterministic Broadcasting in Wireless Networks of Weak Devices*

Abstract: Many futuristic technologies, such as Internet of Things or nano-communication, assume that a large number of simple devices of very limited energy and computational power will be able to communicate efficiently via wireless medium. Motivated by this, we study broadcasting in the model of ad-hoc wireless networks of weak devices with uniform transmission powers. We compare two settings: with and without local knowledge about immediate neighborhood. In the latter setting, we prove $\Omega(n \log n)$ lower bound and develop an algorithm matching this formula. This result could be made more accurate with respect to network density, or more precisely, the maximum node degree Δ in the communication graph. If Δ is known to the nodes, it is possible to broadcast in $O(D \Delta \log^2 n)$ rounds, which is almost optimal in the class of networks parametrized by D and Δ due to the lower bound $\Omega(D \Delta)$. In the setting with local knowledge, we design a scalable and almost optimal algorithm accomplishing broadcast in

$O(D \log^2 n)$ communication rounds, where n is the number of nodes and D is the eccentricity of a network. This can be improved to $O(D \log g)$ if network granularity g is known to the nodes. Our results imply that the cost of "local communication" is a dominating component in the complexity of wireless broadcasting by weak devices, unlike in traditional models with non-weak devices in which well-scalable solutions can be obtained even without local knowledge.