ACAC 2016

11th Athens Colloquium on Algorithms and Complexity

August 25-26, University of Athens, Greece

– *Abstracts* –

ordered by presentation time

1. **Costas Iliopoulos**, King's College London, University of London
   Popping Supperbubbles and Identifying Clumps
   The information that can be inferred or predicted from knowing the genomic sequence of an organism is astonishing. String algorithms are critical to this process. This talk provides an overview of two particular problems - superbubbles and clumps - that arise during computational molecular biology research, and recent algorithmic developments in solving them. Although the motivation comes from computational biology, the proposed algorithms can be applied to various other domains such as musicology, web-mining, image processing etc.

2. **Loukas Georgiadis**, University of Ioannina
   2-Connectivity in Directed Graphs
   Graph connectivity is a fundamental concept in graph theory with numerous practical applications. In this talk we survey some recent results on 2-connectivity in directed graphs. For undirected graphs it has been known for over 40 years how to compute 2-connectivity information in linear time by using depth first search. The analogous problems in digraphs revealed to be much more challenging, and efficient algorithms for their solution have been developed only recently.

3. **Kostas Draziotis**, Aristotle University of Thessaloniki
   Improved Knapsack Algorithms
   In the present study we consider two variants of Schnorr-Shevchenko method (SS) for solving hard knapsack problems, which are in average faster than the SS method. Furthermore, we apply a variant of SS method to find small integer solutions on linear Diophantine equations and also find solutions in a specific interval. Finally, we provide an id-scheme based on the compact knapsack problem. *(This is joint work with A.Papadopoulou)

4. **Grigorios Koumoutsos**, Eindhoven University of Technology
   New Bounds for the (h,k)***-Server Problem***
   We study the k-server problem in the resource augmentation setting i.e., when the performance of the online algorithm with k servers is compared to the offline optimal solution with $h \leq k$ servers. The problem is very poorly understood beyond uniform metrics. For this special case, the classic k-server algorithms are roughly $(1+1/\epsilon)$-competitive when $k = (1+\epsilon)h$, for any $\epsilon > 0$. Surprisingly however, no $o(h)$-competitive algorithm is known even for HSTs of depth 2 and even when k/h is arbitrarily large. We obtain several new results for the problem. First we show that the known k-server algorithms do not work even on very simple metrics. In particular, the Double Coverage algorithm has competitive ratio $\Omega(h)$ irrespective of the value of k, even for depth-2 HSTs. Similarly the Work Function Algorithm, that is believed to be optimal for all metric spaces when $k = h$, has competitive ratio $\Omega(h)$ on depth-3 HSTs even if $k = 2h$. Our main result is a new algorithm that is $O(1)$-competitive for constant depth trees, whenever $k = (1+\epsilon)h$ for any $\epsilon > 0$. Finally, we give a general lower bound that any deterministic online algorithm has competitive ratio at least 2.4 even for depth-2 HSTs and when k/h is arbitrarily large. This gives a surprising qualitative separation between uniform metrics and depth-2 HSTs for the (h,k)-server problem, and gives the strongest known lower bound for the problem on general metrics.

5. **Angelina Vidali**, Teesside University
   TBA
   TBA

6. **Georgios Birmpas**, Athens University of Economics and Business
   On Truthful Mechanisms for Maximin Share Allocations
   We study a fair division problem with indivisible items, namely the computation of maximin share allocations. Given a set of n players, the maximin share of a single player is the best she can guarantee to herself, if she would partition the items in any way she prefers, into n bundles, and then receive her least desirable bundle. The objective then is to find an allocation, so that each player is guaranteed her maximin share. Previous works have studied this problem mostly algorithmically, providing constant factor approximation algorithms. In this work we embark on a mechanism design approach and investigate the existence of truthful mechanisms. We propose three models regarding the information that the mechanism attempts to elicit from the players, based on the cardinal and ordinal representation of preferences. We establish positive and negative (impossibility) results for each model and highlight the limitations imposed by truthfulness on the approximability of the problem. Finally, we pay particular attention to the case of two players, which already leads to challenging questions.

7. **Christos Zaroliagis**, University of Patras
   New Approaches for Time-Dependent Route Planning
   TBA

8. **Iannis Caragiannis**, University of Patras
   Discrete preference games in social networks
   It is often observed that users tend to imitate the behavior of their neighbors in a social network. This imitating behavior might lead to the strategic decision of adopting a public behavior that differs from what the agent believes is the right one and this can subvert the behavior of the population as a whole. In this talk, we consider the case in which agents express preferences over two alternatives and model social pressure with the majority dynamics: at each step an agent is selected and its preference is replaced by the majority of the preferences of her neighbors. We ask whether there are network topologies that are robust to social pressure. That is, we ask if there are graphs in which the majority of preferences in an initial profile always coincides with the majority of the preference in all stable profiles reachable from that profile. We completely characterize the graphs with this robustness property by showing that this is possible only if the graph has no edge or is a clique or very close to a clique. In other words, except for this handful of graphs, every graph admits at least one initial profile of preferences in which the majority dynamics can subvert the initial majority. We also show that deciding whether a graph admits a minority that becomes majority is NP-hard when the minority size is slightly smaller than half of the social network size.

   Joint work with Vincenzo Auletta, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano.

9. **George Krimpas**, University of Patras
   How Effective Can Simple Ordinal Peer Grading Be?
   Ordinal peer grading has been proposed as a simple and scalable solution for computing reliable information about student performance in massive open online courses. The idea is to outsource the grading task to the students themselves as follows. After the end of an exam, each student is asked to rank (in terms of quality) a bundle of exam papers by fellow students. An aggregation rule will then combine the individual rankings into a global one that contains all students. We define a broad class of simple aggregation rules and present a theoretical framework for assessing their effectiveness. When statistical information about the grading behaviour of students is available, the framework can be used to compute the optimal rule from this class with respect to a series of performance objectives. For example, a natural rule known as Borda is proved to be optimal when students grade correctly. In addition, we present extensive simulations and a field experiment that validate our theory and prove it to be extremely accurate in predicting the performance of aggregation rules even when only rough information about grading behaviour is available. Joint work with Ioannis Caragiannis and Alexandros Voudouris.

10. **Katerina Karanasiou**, University of Ioannina
On Low-High Orders of Directed Graphs: An Incremental Algorithm and Applications
A flow graph $G = (V, E, s)$ is a directed graph with a distinguished start vertex s. The dominator tree D of G is a tree rooted at s, such that a vertex v is an ancestor of a vertex w if and only if all paths from s to w include v. The dominator tree is a central tool in program optimization and code generation, and has many applications in other diverse areas including constraint programming, circuit testing, biology, and in algorithms for graph connectivity problems. A low-high order of G is a preorder δ of D that certifies the correctness of D, and has further applications in connectivity and path-determination problems. In this paper, we show how to maintain a low-high order of a flow graph G incrementally under edge insertions. For a sequence of m edge insertions in an initially empty flow graph with n vertices, our algorithm runs in $O(mn)$ total time. This immediately provides the first incremental certifying algorithm for maintaining a dominator tree in $O(mn)$ total time. Our result also implies incremental algorithms, with $O(mn)$ total running time, for maintaining fault tolerant reachability, two strongly divergent spanning trees, and a sparse certificate for 2-edge-connectivity of a directed graph. Hence, we provide a substantial improvement over the $O(m^2)$ simple-minded algorithms, which recompute the solution from scratch after each edge insertion. Furthermore, we show how to apply low-high orders to obtain a linear-time 2-approximation algorithm for the smallest 2-vertex-connected spanning subgraph problem.

11. **Anna Mpanti**, University of Ioannina
Two RPG Flow-graphs for Software Watermarking using Bitonic Sequences of Self-inverting Permutations
Software watermarking has received considerable attention and was adopted by the software development community as a technique to prevent or discourage software piracy and copyright infringement. A wide range of software watermarking techniques has been proposed among which the graph-based methods that encode watermarks as graph structures. Following up on our recently proposed methods for encoding watermark numbers w as reducible permutation flow-graphs $F[\pi^*]$ through the use of self-inverting permutations π^*, in this paper, we extend the types of flow-graphs available for software watermarking by proposing two different reducible permutation flow-graphs $F_1[\pi^*]$ and $F_2[\pi^*]$ incorporating important properties which are derived from the bitonic subsequences composing the self-inverting permutation π^*. We show that a self-inverting permutation π^* can be efficiently encoded into either $F_1[\pi^*]$ or $F_2[\pi^*]$ and also efficiently decoded from theses graph structures. The proposed flow-graphs $F_1[\pi^*]$ and $F_2[\pi^*]$ enrich the repository of graphs which can encode the same watermark number w and, thus, enable us to embed multiple copies of the same watermark w into an application program P. Moreover, the enrichment of that repository with new flow-graphs increases our ability to select a graph structure more similar to the structure of a given application program P thereby enhancing the resilience of our codec system to attacks.

12. **Alkmini Sgouritsa**, University of Liverpool
Designing Cost-Sharing Methods for Bayesian Games
We study the design of cost-sharing protocols for two fundamental resource allocation problems, the Set Cover and the Steiner Tree Problem, under environments of incomplete information (Bayesian model). Our objective is to design protocols where the worst-case Bayesian Nash equilibria have low cost, i.e. the Bayesian Price of Anarchy (PoA) is minimized. Although budget balance is a very natural requirement, it puts considerable restrictions on the design space, resulting in high PoA. We propose an alternative, relaxed requirement called budget balance in the equilibrium (BBiE). We show an interesting connection between algorithms for Oblivious Stochastic optimization problems and cost-sharing design with low PoA. We exploit this connection for both problems and we enforce approximate solutions of the stochastic problem, as Bayesian Nash equilibria, with the same guarantees on the PoA. More interestingly, we show how to obtain the same bounds on the PoA, by using anonymous posted prices which are desirable because they are easy to implement and, as we show, induce dominant strategies for the players.

13. **Elias Tsigaridas**, INRIA Paris
A Randomized Algorithm to Decompose Binary Forms
Symmetric Tensor Decomposition is a major problem that arises in areas such as signal processing, statistics, data analysis and computational neuroscience. It is equivalent to write a homogeneous polynomial in n variables of degree D as a sum of D-th powers of linear forms, using the minimal

number of summands. This minimal number is called the rank of the polynomial/tensor. We consider the decomposition of binary forms; this corresponds to the decomposition of symmetric tensors of dimension 2 and order D and we present the the first quasi-optimal algorithm. Joint work with Matias Bender, Jean-Charles Faugere, and Ludovic Perret.

14. **Michael Lampis**, University Paris-Dauphine, LAMSADE
Trading Time for Approximation
The design of approximation algorithms is inherently a process with two competing objectives: we seek an algorithm with the maximum possible accuracy (i.e. good approximation ratio) and minimum possible running time. For the most important problems, the last 25 years of research have created a pretty clear picture of what can be done if we fix one of these objectives to the extreme. In other words, for most problems we understand how well we can approximate the optimal if we insist that the running time of our algorithm must be polynomial, and we understand how much time we need to devote to them if we insist that we need an exactly optimal solution. But what happens between these two extremes?

    In this talk we survey a number of recent results which try to explore and fully capture the Pareto curve between these two extremes. Interestingly, sometimes this curve turns out to be smooth and sometimes it turns out to contain huge jumps, in more or less convenient places. In order to understand this we will use a healthy mix of techniques from the theory of both approximation and exponential/parameterized algorithms.

15. **Giorgos Christodoulou**, University of Liverpool
Designing Games with Good Equilibria
We will discuss ways to cope with inefficiency of equilibria in some cases of interest like congestion games, scheduling and cost-sharing games.

    The main question we will address is: to what extent can prior knowledge of the underlying metric help in the design of good algorithms or protocols?

16. **Vangelis Markakis**, Athens University of Economics and Business
Inequity Aversion Pricing over Social Networks
We study a revenue maximization problem in the context of social networks. Namely, we consider a model introduced by Alon, Mansour, and Tennenholtz that captures *inequity aversion*, i.e., prices offered to neighboring vertices should not be significantly different. We first provide approximation algorithms for a natural class of instances, referred to as the class of single-value revenue functions. Our results improve on the current state of the art, especially when the number of distinct prices is small. This applies, for example, to settings where the seller will only consider a fixed number of discount types or special offers. We then resolve one of the open questions posed in the previous paper, by establishing APX-hardness for the problem. In doing so, we also establish APX-hardness for 3-Terminal Node Cut, which is interesting in its own right. Surprisingly, we further show that the problem is NP-complete even when the price differences are allowed to be relatively large. Finally, we also provide some extensions of the model, regarding either the allowed set of prices, or the demand type of the clients.

17. **Yorgos Amanatidis**, Athens University of Economics and Business
Coverage, Matching, and Beyond: New Results on Budgeted Mechanism Design
We study a type of reverse (procurement) auction problems in the presence of budget constraints. The general algorithmic problem is to purchase a set of resources, which come at a cost, so as not to exceed a given budget and at the same time maximize a given valuation function. This framework captures the budgeted version of several well known optimization problems, and when the resources are owned by strategic agents the goal is to design truthful and budget feasible mechanisms, i.e. elicit the true cost of the resources and ensure the payments of the mechanism do not exceed the budget. Budget feasibility introduces more challenges in mechanism design, and we study instantiations of this problem for certain classes of submodular and XOS valuation functions. We first obtain mechanisms with an improved approximation ratio for weighted coverage valuations, a special class of submodular functions that has already attracted attention in previous works. We then provide a general scheme for designing randomized and deterministic polynomial time mechanisms for a class of XOS problems. This class contains problems whose feasible set forms an independence

system (a more general structure than matroids), and some representative problems include, among others, finding maximum weighted matchings, maximum weighted matroid members, and maximum weighted 3D-matchings. For most of these problems, only randomized mechanisms with very high approximation ratios were known prior to our results. Joint work with Georgios Birmpas and Evangelos Markakis.

18. **John Livieratos**, University of Athens
Aggregation of Votes with Multiple Positions on Each Issue
We consider the problem of aggregating votes cast by a society on a fixed set of issues, where each member of the society may vote for one of several positions on each issue, but the combination of votes on the various issues is restricted to a feasible set of voting patterns. We explore the connection between sets of feasible voting patterns where non-dictatorial aggregation is possible, with the complexity of the conservative multi-sorted satisfiability problem defined over such sets. In order to achieve that, we provide a characterization for possibility domains, that is, sets of feasible voting patterns where non-dictatorial aggregation is possible and for the stronger notion of possibility domains in each issue, where non-dictatorial aggregation is not just required for the whole set, but also issue-wise.

19. **Orestis Telelis**, University of Piraeus
Envy-Free Revenue Approximation for Asymmetric Buyers with Budgets
We study the computation of revenue-maximizing envy-free outcomes in a monopoly market with budgeted buyers. Departing from previous works, we focus on buyers with asymmetric combinatorial valuation functions over subsets of items. We first establish a hardness result showing that even with two identical additive buyers, the problem is inapproximable. As an attempt to identify tractable families of instances, we introduce the notion of budget compatible buyers, placing a restriction on the budget of each buyer in terms of his valuation function. Under this assumption, we establish approximation upper bounds for buyers with submodular valuations over preference subsets as well as for identical subadditive valuation functions. Finally, we also analyze an algorithm for arbitrary additive valuation functions, which yields a constant approximation for a constant number of buyers. We conclude with intriguing open questions regarding budgeted buyers with asymmetric valuation functions. This is a joint work with Vangelis Markakis.

20. **Michael Sioutis**, University of Artois
Efficiently Reasoning about Qualitative Constraints through Variable Elimination
We introduce, study, and evaluate a novel algorithm in the context of qualitative constraint-based spatial and temporal reasoning, that is based on the idea of variable elimination, a simple and general exact inference approach in probabilistic graphical models. Given a qualitative constraint network N, our algorithm enforces a particular directional local consistency on N, namely, directional weak path consistency. Our discussion is restricted to distributive subclasses of relations, i.e., sets of relations closed under converse, intersection, and weak composition and for which weak composition distributes over non-empty intersections for all of their relations. We demonstrate that enforcing directional weak path consistency on a given qualitative constraint network defined over a distributive subclass of relations allows us to decide its satisfiability. The experimentation that we have conducted with random and real-world qualitative constraint networks defined over a distributive subclass of relations of the Region Connection Calculus, shows that our approach exhibits unparalleled performance against competing state-of-the-art approaches for checking the satisfiability of such constraint networks.

21. **Kyriakos Axiotis**, National Technical University of Athens
Connectivity in Temporal Graphs
A lot of real world applications (eg transportation networks, social networks) are inherently dynamic and cannot be modeled effectively by traditional graphs. The more general notion of Temporal Graphs can capture this kind of time-changing behavior. We consider temporal graphs with discrete time labels and investigate connectivity questions that revolve around Minimum Temporally Connected Subgraphs (MTCS) - a generalization of spanning trees in temporal graphs. We study the worst case size of an MTCS, as well as the approximability of two related optimization problems. For the former, we establish a construction that achieves asymptotically tight size, thus closing the previous large gap. For the latter we provide several upper and lower bounds on the approximation

ratio for the general case, as well as for several special cases. Our reductions show that Temporal Connectivity problems are closely linked to Directed Steiner problems. Joint work with Dimitris Fotakis.

22. **Ioannis Kontoyiannis**, Athens University of Economics and Business
Model Selection Algorithms for Discrete Time Series
We introduce a class of novel methodological and algorithmic tools for effective Bayesian inference and model selection on arbitrary discrete time-series data. Our approach is based on the use of a very rich class of hierarchical models, and the observation that the so-called "context tree weighting" algorithm (developed by Willems and co-authors in the early 1990s) admits broad extensions that provide effective computational tools for inference in very general settings.

    The statistical model and algorithmic aspects will be described in detail. Applications in a very broad variety of areas of machine learning and statistics ranging from finance to neuroscience will also be briefly discussed.

    Based on joint work with Athina Panotopoulou and Maria Skoularidou.

23. **Katia Papakonstantinopoulou**, University of Athens
Memory-Optimized Distributed Graph Processing through Novel Compression Techniques
A multitude of contemporary applications now involve graph data whose size continuously grows and this trend shows no signs of subsiding. This has caused the emergence of many distributed graph processing systems including Pregel and Apache Giraph. However, the unprecedented scale now reached by real-world graphs hardens the task of graph processing even in distributed environments and the current memory usage patterns rapidly become a primary concern for such contemporary graph processing systems. We seek to address this challenge by exploiting empirically-observed properties demonstrated by graphs that are generated by human activity. In this paper, we propose three space-efficient adjacency list representations that can be applied to any distributed graph processing system. Our suggested compact representations reduce respective memory requirements for accommodating the graph elements up to 5 times if compared with state-of-the-art methods. At the same time, our memory-optimized methods retain the efficiency of uncompressed structures and enable the execution of algorithms for large scale graphs in settings where contemporary alternative structures fail due to memory errors. Last but not least, we propose a tree-based space-efficient out-edge representation that favors mutations as well.

24. **Anna Karasoulou**, University of Athens
Approximating Multidimensional Subset Sum and Minkowski Decomposition of polytopes
We consider the approximation of two NP-hard problems: Minkowski Decomposition (MinkDecomp) of lattice polygons in the plane and the closely related problem of Multidimensional Subset Sum (kD-SS) in arbitrary dimension. In kD-SS we are given an input set S of k-dimensional vectors, a target vector t and we ask if there exists a subset of S that sums to t. We prove, through a gap-preserving reduction, that, for general dimension k, kD-SS does not have a PTAS although the classic $1D$-SS does. On the positive side, we present an $O(n^7/\epsilon^4)$ approximation algorithm for $2D$-SS, where n is the cardinality of the set and ϵ bounds the difference of some measure of the input polygon and the sum of the output polygons. Applying this algorithm, and a transformation from MinkDecomp to $2D$-SS, we can approximate MinkDecomp. For an input polygon Q and parameter ϵ, we return two summands A and B such that $A + B = Q'$ with Q' being bounded in relation to Q in terms of volume, perimeter, or number of internal lattice points and an additive error linear in ϵ and up to quadratic in the diameter of Q. A similar function bounds the Hausdorff distance between Q and Q'.

25. **Ioannis Psarros**, University of Athens
High-dimensional approximate r-nets
The construction of r-nets offers a powerful tool in computational and metric geometry. We focus on high-dimensional spaces and present a new randomized algorithm which efficiently computes approximate r-nets with respect to Euclidean distance. For any fixed $\epsilon > 0$, the approximation factor is $1 + \epsilon$ and the complexity is polynomial in the dimension and subquadratic in the number

of points. The algorithm succeeds with high probability. More specifically, the best previously known LSH-based construction of Eppstein et al. [EHS15] is improved in terms of complexity by reducing the dependence on ϵ, provided that ϵ is sufficiently small. Our method does not require LSH but, instead, follows Valiant's [Val15] approach in designing a sequence of reductions of our problem to other problems in different spaces, under Euclidean distance or inner product, for which r-nets are computed efficiently and the error can be controlled. Our result immediately implies efficient solutions to a number of geometric problems in high dimension, such as finding the $(1+\epsilon)$-approximate kth nearest neighbor distance in time subquadratic in the size of the input.

References:

[EHS15] D. Eppstein, S. Har-Peled, and A. Sidiropoulos. Approximate greedy clustering and distance selection for graph metrics. CoRR, abs/1507.01555, 2015

[Val15] G. Valiant. Finding correlations in subquadratic time, with applications to learning parities and the closest pair problem. J. ACM, 62(2):13, 2015.

26. **Valia Mitsou**, Hungarian Academy of Sciences
Double-exponential and triple-exponential lower bounds of choosability problems parameterized by treewidth
Choosability, introduced by Erdös, Rubin, and Taylor [Congr. Number. 1979], is a well-studied concept in graph theory: we say that a graph is c-choosable if for any assignment of a list of c colors to each vertex, there is a proper coloring where each vertex uses a color from its list. We study the complexity of deciding choosability on graphs of bounded treewidth. It follows from earlier work that 3-choosability can be decided in time $2^{2^{O(w)}}n^{O(1)}$ on graphs of treewidth w. We complement this result by a matching lower bound giving evidence that double-exponential dependence on treewidth may be necessary for the problem: we show that an algorithm with running time $2^{2^{o(w)}}n^{O(1)}$ would violate the Exponential-Time Hypothesis (ETH). We consider also the optimization problem where the task is to delete the minimum number of vertices to make the graph 4-choosable and demonstrate that dependence on treewidth becomes triple-exponential for this problem: it can be solved in time $2^{2^{2^{O(w)}}}n^{O(1)}$ on graphs of treewidth w, but an algorithm with running time $2^{2^{2^{o(w)}}}n^{O(1)}$ would violate ETH. The significance of the results is that these problems are apparently the first fairly natural graph-theoretic problems that require double-exponential or triple-exponential dependence on treewidth. Joint work with Dániel Marx.

27. **Panagiotis Liakos**, University of Athens
Local Dispersion-aware Link Communities
Real-life systems entailing interacting objects have been routinely modeled as graphs or networks. Revealing the community structure of such systems is crucial in helping us better understand their complex nature. Networks and the relationships they portray are exploited by community detection techniques that seek to facilitate the discovery of separate, overlapping, nested or fully hierarchical communities. Nevertheless, our perception of what a community is in a network of interacting objects has evolved over the years. In this respect, the detection of the possibly overlapping organizational groups of a node and our enhanced understanding of their intricate interactions remain open challenges. In this paper, we propose a local community detection approach focusing in a single target node in the network. We use agglomerative clustering to group pairs of links in the corresponding egonet and provide a richer hierarchical structure compared to previous attempts. To this end, we exploit the dispersion of established relationships among objects in the egonet. Our algorithm (termed Local Dispersion-aware Link Communities or LDLC) measures the similarity of such links as well as the extent of their participation in multiple contexts, to determine the order in which pairs of links should be grouped. Our experimental results showcase that our approach overcomes issues earlier techniques stumbled upon. Moreover, we evaluate the performance of our algorithm against ground-truth communities for a wide range of networks and show that LDLC significantly outperforms state-of-the-art methods.

28. **Christos Vlasopoulos**, NCSR Demokritos
Grand Challenge: Dynamic Graph Management for Streaming Social Media Analytics
We present a system for analytics on streaming social media that computes the most active posts, based on the age and the amount of comments for each post, and tracks the largest communities

that comprise friends that are fond of the same content. To deal with high velocity data streams, we implemented an algorithm for incrementally updating graphs expressing social networks. The evaluation of our system is based on the datasets of the DEBS 2016 challenge.

29. **Vangelis Paschos**, University Paris-Dauphine, LAMSADE
On the approximation of maximum k-vertex cover in bipartite graphs
TBA