A Glimpse at Paul G. Spirakis

Ioannis Chatzigiannakis[1,2(✉)], Dimitris Fotakis[3], Spyros Kontogiannis[1,4], Othon Michail[1], Sotiris Nikoletseas[1,5], Grammati Pantziou[1,6], and Christos Zaroliagis[1,5]

[1] Computer Technology Institute and Press "Diophantus", Patras University Campus, N. Kazantzaki Str., 26504 Patras, Greece
{ichatz,michailo,nikole}@cti.gr
[2] Sapienza University of Rome, Rome, Italy
[3] Division of Computer Science, School of Electrical and Computer Engineering, National Technical University of Athens, 15780 Athens, Greece
fotakis@cs.ntua.gr
[4] Department of Computer Science and Engineering, University of Ioannina, 45110 Ioannina, Greece
kontog@cs.uoi.gr
[5] Department of Computer Engineering and Informatics, University of Patras, 26504 Patras, Greece
zaro@ceid.upatras.gr
[6] Department of Informatics, Technological Educational Institution of Athens, Egaleo, Greece
pantziou@teiath.gr

1 Introduction

Paul Spirakis is an eminent, talented, and influential researcher that contributed significantly to computer science. This article is a modest attempt of a biographical sketch of Paul, which we drafted with extreme love and honor.

2 Childhood, Education and Career

Paul G. Spirakis was born on 29 August 1955 in Didymoteicho, a city in the northeastern part of Greece, just 2Km from the Greek-Turkish border. His father George Spirakis originated from that city, while his mother Olga Avgoustinou originated from the island of Zakynthos (or Zante). She moved to Didymoteicho as a teacher in an elementary school of the area, where she met Paul's father. The family is complemented by a daughter (Eleni). Didymoteicho was the home of the Spirakis family and Paul's parents were among the prominent citizens of the city. George Spirakis was the City Mayor in the periods 1960–1966 and 1978–1982.

Paul finished the elementary school there as well as the five out of the six classes of the high-school. He was the top student in all classes with a passionate love for mathematics. This caused his mathematics teacher in high-school's fifth class (Mr. Ionas) to convince his parents to let Paul enroll in the advanced 1st

C. Zaroliagis et al. (Eds.): Spirakis Festschrift, LNCS 9295, pp. 3–24, 2015.
DOI: 10.1007/978-3-319-24024-4_1

(public) high-school in Thessaloniki for further training in mathematics. This indeed happened (to the full disappointment of Paul's literature teacher, who was trying to convince him and his parents to study literature) and Paul finished high-school in Thessaloniki in 1973, again as the top student of his class.

During the same academic year (1972–1973) he was also enrolled to a special school preparing students for the National Examinations for University Entrance. His mathematics teacher in that school (Mr. Mpallis), motivated by Paul's passionate love and talent for mathematics, soon developed a close relationship with him and devoted the majority of his time to Paul's further training in mathematics.

As a result, Paul received in 1973 one of the most prestigious prizes in Greece, the First Prize of the Greek Mathematics Society. This prize is awarded after a highly-competitive national examination among students all over Greece. In the same year, he also succeeded in the National Examinations for University Entrance, entering the School of Electrical Engineering at the National Technical University of Athens (NTUA). During his studies, he was constantly among the top 2 % of students, graduating with excellence in 1978 (5-year Diploma in Electrical Engineering from NTUA).

Subsequently, Paul was admitted (with scholarship) at Harvard University for pursuing postgraduate studies. He received his S.M. degree in Applied Mathematics & Computer Science in June 1979, and his Ph.D. in Applied Mathematics & Computer Science in early 1982. His Ph.D. Thesis "Probabilistic Algorithms, Algorithms with Random Inputs and Random Combinatorial Structures", under the supervision of Prof. John Reif, contributed significantly to the field of probabilistic and randomized algorithms.

After his Ph.D, Paul received a post-doctoral research fellowship at Harvard University and in September 1982 was appointed as Assistant Professor in the Courant Institute of Mathematical Sciences at New York University. In the summer of 1985 he was elected Associate Professor in the Department of Computer Engineering and Informatics at the University of Patras, and in October 1990 was promoted to a Full Professor (meanwhile, during the period October 1985–April 1987, he served his obligatory military service in the Greek Army). In 1996, Paul was appointed as the President (Chairman of the Board of Directors) of the Computer Technology Institute & Press "Diophantus". Among others, he is the Head of Research Unit 1 (Foundations of Computer Science, Relevant Technologies and Applications).

Since September 2013 he holds a chair professorship in the Department of Computer Science at the University of Liverpool, UK. He also leads the Networks Sciences and Technologies (NeST) Initiative of the University of Liverpool, and he is the chair of the research committee of the School of Electrical Engineering, Electronics and Computer Science at the same university.

3 Teaching, Mentoring, and Publications

During the 33 years of his academic career, Paul Spirakis taught a variety of classes at both undergraduate and graduate level, covering a broad spectrum of subjects within Computer Science, including *Algorithms and Combinatorial Optimization, Approximation algorithms to Hard Problems, Cryptography, Discrete Mathematics, Distributed Computing and Systems, Game Theory, Economics and Algorithms, Operating Systems, Parallel Algorithms, Probability Theory, Probabilistic Techniques in Algorithms,* and *Theory of Computation.*

Most of these classes were introduced and taught for the first time in a Greek University. All students of those classes surely remember Paul's inspiring lectures, and for most of them those lectures had a predominant influence on their careers.

A key ability of Paul Spirakis is the identification of core problems in technologies that are still in their inception, many years before they become apparent to the scientific community. This requires an almost unique inter-scientific ability to combine a broad set of techniques and methodologies to examine a given problem.

Paul Spirakis is among the most visionary thought leaders of our generation, with a great talent to inspire and guide new researchers. He has invested a huge amount of effort and time in mentoring students and young scientists, and introducing them to the mysteries of computation and its applications. He supervised numerous undergraduate and postgraduate students and he awarded more than 32 doctoral dissertations.

It did not then come as a surprise that, in a study conducted in 2004, he was reported among the *top 50 nurturers* in Computer Science Research [47].

Paul Spirakis has a tremendous publication record. He has published:

- More than 120 research papers in prestigious peer-reviewed journals.
- More than 290 research papers in refereed proceedings of top international computer science conferences.
- 12 books (3 in English and 9 in Greek).
- 17 chapters in books or edited volumes.
- More than 40 research papers in archived repositories (representing work in progress).

His publications have received more than 5300 citations (h-index: 37, g-index: 59).

4 Awards and Distinctions

For his research work and scientific achievements, Paul Spirakis has been honored with numerous awards and distinctions, which include:

- Fellow of the European Association for Theoretical Computer Science (EATCS), since 2014.
- Member of Academia Europaea, since 2010.
- Member of the ACM Europe Council, sice 2009.
- Recipient of a Technology Excellence Award, as a Technology Pioneer, in 2009 in the competition of Technology Excellence Awards 2009 organized by the PC and T3 magazines (Greek Editions).

- Vice President of the European Association for Theoretical Computer Science (elected unanimously), since 2002.
- Member of the Council Board of the European Association for Theoretical Computer Science (elected unanimously), since 1997.
- Acknowledged as among the *top 50 nurturers* in Computer Science Research [47].
- Distinguished Visiting Scientist of the Max-Planck Institute for Informatics, Fall 2001.

5 Research

The research work of Paul Spirakis is so broad that it would be extremely difficult (if at all possible) to describe it in an succinct and unambiguous way. It encompasses two main streams that fruitfully interact with each other throughout his career, *Algorithms and Complexity*, and *Computer Systems and Networks*. The former stream includes his contributions to the foundations of computer science. The latter stream includes mainly applied work that it has heavily benefited from the former stream.

For this reason, we have chosen here to focus on the *Algorithms and Complexity* stream and attempt at providing a short overview of his contributions. In all cases, the vast majority of Paul's work is characterized by a set of recurring themes: probabilistic and randomized approaches, fundamental methods of parallel and distributed computing, graph-theoretic approaches, approximation methods, and (lately) algorithmic game theory and evolution methods & structures.

5.1 Probabilistic and Randomized Algorithms

From the very start of his academic formation, Paul Spirakis developed a fundamental relation to probabilistic methods and their applications in theoretical computer science. In his Ph.D. Thesis, Paul suggested some novel, major uses of discrete probability, in combinatorial structures as well as in the design and analysis of randomized algorithms.

In particular, in [72] he introduced randomness in the fundamental combinatorial structure of matroids, towards its powerful generalization to random matroids. For the independent set problem under this new model, nonconstructive existence proofs as well as efficient randomized algorithms were provided. Also, some nice applications of this structure to classic Erdős-Rényi random graphs were given. The second part of Paul's Ph.D. Thesis anticipated in a characteristic way his major future contributions to randomized methods in distributed computing throughout his career. In [73], he suggested the employment of probabilistic choice for interprocess communication (and symmetry breaking). It is worth noting that this has been one the first few uses of random choice in distributed computing.

Since this early stage, and throughout his career, random methods have been the main connecting thread in the research of Paul Spirakis. Several applications of randomness to diverse topics are also discussed elsewhere in this chapter. We highlight here two fundamental (per se) uses of randomness.

The first one concerns his constant investigation of random combinatorial models. Further to a deep exploration of crucial properties in $G_{n,p}$ random graphs, Paul liked to introduce nice extensions of such models, motivated both by mathematical curiosity and major trends in modern technology. Such extensions characteristically include the introduction of failures to random regular graphs models [65] and the exploration of interesting variations of random intersection graphs [67].

Relevant persistent features of his research include the heavy use of random processes to model and analyze interesting computational phenomena (see e.g., [17,51]) and the use of randomness in algorithms for graph-theoretic problems (see e.g., [19,29]).

In a second line, the exploration of major combinatorial properties of such models has been based on a sophisticated use and even extension of probabilistic techniques collectively referred to as the Probabilistic Method established by Paul Erdős. It is worth noting that Paul Spirakis maintained a deep devotion to such probabilistic techniques, and effectively passed on his dedication to several of his students and colleagues. He not only applied in a brilliant way randomized techniques but several times he also further developed and extended these techniques themselves, such as in [35], where a new series of tail bounds for occupancy problem has been provided.

5.2 Parallel Algorithms and Complexity

A few years after his Ph.D, Paul Spirakis started an intense research activity in the field of parallel algorithms and complexity. He studied various fundamental problems in parallel computing and his results contributed significantly in promoting the field. We mention here a few indicative cases.

One of his first attempts was to devise efficient algorithms that are sensitive to properties of the input which can be determined only at run-time. For instance, in the case of parallel addition in shared memory models, it is interesting to devise algorithms whose bounds depend on the number of non-zero elements. In [77], Paul designed such an algorithm for the fundamental problem of parallel addition. In an input of n numbers, m of which are non-zero entries, he devised a randomized parallel algorithm for a CRCW PRAM, which runs in $O(\log m)$ expected time with m processors using $O(m)$ shared space. He also applied this algorithm to the related problem of processor identification.

A fair part of his work dealt with the development of efficient parallel algorithms and the investigation of the parallel complexity of several problems on graphs within the context of shared memory parallel computing:

- Along the former, he derived (along with Grammati Pantziou and Christos Zaroliagis) efficient deterministic parallel algorithms for shortest paths and other problems (see e.g., [38,71]) as well as parallel algorithms which exhibited a remarkable average case performance. In particular, he studied with John Reif in [76] the parallel average case complexity of several problems on random instances of undirected and directed graphs. In that paper, a bulk of algorithms was developed, based on new results on the diameter of random (directed) graphs, that are able to solve a host of graph problems on an n-node random (directed) graph (connectivity, biconnectivity, transitive closure, minimum spanning trees, and all pairs minimum cost paths) in $O(\log \log n)$ expected parallel time on a CRCW PRAM. These are exponentially faster algorithms than their deterministic counterparts.

– Along the latter, he investigated the parallel complexity of the problem of testing whether a given graph G contains an induced subgraph of vertex (edge) connectivity at least k. Paul Spirakis, with Maria Serna and Lefteris Kirousis, proved in [41] that this problem is P-complete for any fixed $k \geq 3$. This result came as a surprise, since the related problem of computing the triconnected (or Tutte) components of G (maximal subgraphs of G such that for any four vertices in any of them, any two of these vertices can be connected by a path in G that avoids the other two) was known to be in NC. In addition, they provided interesting NC approximability and inapproximability results.

Another thread of his research concerned fault-tolerant parallel computing. Paul Spirakis together with KZvi Kedem and Krishna Palem investigated in [39] the problem of executing efficient robust parallel computations on a PRAM whose processing elements are prone to failure. In particular, they devised a general strategy for simulating an arbitrary step of an ideal CRCW PRAM on a PRAM with faulty processors at a small multiplicative time overhead and at a small (per processor) additive constant space overhead.

An equally important thread of Paul's research focused on more realistic models of parallel computing. In particular, he (along with co-authors) investigated in [4] a quantitative comparison of the BSP and LogP models of parallel computation. Both models are successful paradigms of the so-called bridging models of parallel computation, where one seeks for balancing simplicity (that eases software development), accuracy (to enable realistic performance predictions), and generality (to enable software portability across various architectures). In BSP the fundamental primitives are global barrier synchronization and the routing of arbitrary message sets. LogP lacks explicit synchronization and imposes a more constrained message-passing style which aims at keeping the load of the underlying communication network below a specified capacity limit. Intuitively, BSP offers a more convenient abstraction for algorithm design and programming, while LogP provides better control of machine resources. In [4], very efficient cross simulations between BSP and the stall-free LogP were derived, showing their substantial equivalence for algorithmic design guided by asymptotic analysis. It was also shown that the two models can be implemented with similar performance on most point-to-point networks.

5.3 Networks and Distributed Computing

The contribution of the research work of Paul Spirakis in the field of networks and distributed computing has been very important for the further development of the field. He has coauthored a significant number of research articles while he contributed as a program committee member of related scientific conferences and as editor of journals. In the following, we highlight some of the most important research works of Paul Spirakis and his coauthors in the field.

Paul Spirakis with John Reif addressed in [75] the fundamental problem of synchronizing communication between distributed processes whose speeds vary dynamically, and they showed how to implement a distributed local scheduler to find matching pairs of processes which are willing to communicate. In [73,74] they considered the probabilistic approach to synchronization of communication in a network of distributed asynchronous processes and presented probabilistic synchronization algorithms that have real time response (the establishment of communication is taking place within a specified time interval with high probability). The algorithms are applied to solve a large class of real time resource synchronization problems.

Paul Spirakis with Hermann Jung and Lefteris Kirousis presented in [34] an algorithm for scheduling a DAG of n nodes on a multiprocessor. The algorithm constructs an optimum schedule which uses at most n processors. They also gave lower bound results on the amount of recomputation needed, thus answering an open question posed by Papadimitriou and Yannakakis.

Paul Spirakis with Lefteris Kirousis and Philippas Tsigas addressed in [42] the problem of reading more than one variables in one atomic operation by only one reader while each of the variables is being written by a set of writers. They presented a deterministic protocol with linear in the number of processes space complexity, linear time complexity for a read operation and constant time complexity for a write operation, as well as a simple probabilistic algorithm with sublinear space complexity and time complexity for a read operation, thus improving significantly previous approaches which required at best, quadratic time and space complexity.

Paul Spirakis with Panagiota Fatourou studied in [21] the problem of efficiently scheduling strict multithreaded computations and presented the first fully distributed scheduling algorithm. The algorithm is asynchronous, on-line, and efficient not only in terms of its memory requirements and its execution time, but also in terms of its communication complexity. Their analysis applies to both shared and distributed memory machines.

Paul Spirakis with Josep Diaz, Dimitrios Koukopoulos, Sotiris Nikoletseas, Maria Serna, and Dimitrios Thilikos analyzed in [18] the stability properties of the FIFO protocol in the adversarial queueing model for packet routing. They presented an upper bound result for stability of any network under the FIFO protocol, answering partially an existing open question. In [45] Paul Spirakis with his co-authors Dimitrios Koukopoulos, Marios Mavronicolas, and Sotiris Nikoletseas, studied the problem of how network structure affects the stability properties of greedy contention-resolution protocols in the framework of the adversarial queueing theory. They came up with a comprehensive collection of structural results in the form of stability and instability bounds on injection rate of the adversary.

Another aspect of Paul's research concerned intrusion propagation in networks. In a joint work with Nikoletseas et al. [66], they studied the problem of intrusion propagation under the assumption of a rather limited in power intruder and how (under this assumption) intrusion can propagate in a perhaps highly secure network. To study this problem, they introduced a new general model for such an intrusion and its propagation in networks. As it turned out by analytic and experimental methods, even such an intruder can have a large penetration factor in the network. Moreover, it was also shown that it will not be easy for a detection mechanism to trace the origin of the intrusion, since it will have to trace a number of links proportional to the nodes captured.

Except for pure theoretical work, Paul Spirakis together with Christos Zaroliagis investigated in [78] implementation and experimentation aspects of distributed algorithms. When one engineers distributed algorithms, some special characteristics arise that are different from conventional (sequential or parallel) computing paradigms. These characteristics include: the need for either a scalable real network environment or a platform supporting a simulated distributed environment; the need to incorporate asynchrony, where arbitrary asynchrony is hard (if at all possible) to implement; and the generation of "difficult" input instances which is a particular challenge. In [78], the term *Distributed Algorithm Engineering* was coined to emerge the need for a systematic methodology to address the aforementioned characteristics as well as the considerable effort required to convert theoretically efficient and correct distributed algorithms to efficient, robust, and easily used software implementations on a simulated or real distributed environment. This conversion has to preserve the assumed properties and limitations of the distributed computing model. The study in [78] addresses several methodological issues in Distributed Algorithm Engineering and illustrates certain approaches to tackle them via case studies.

5.4 Internet, Mobile, and Evolution Networks

Paul was always interested in studying the theoretical foundations of networks using different techniques and testing alternative research direction. Yet, one of the most distinguishing aspects of his approach is to steer his theoretical curiosity into real problems emerging from newly introduced technologies. A typical example of this aspect of his research character is the case of the mobile networks.

In the 1990s, mobile telephony networks were attracting a lot of attention and were emerging as a new technological area. As a first attempt to examine this new technology, along with Grammati Pantziou and George Pentaris, he started by looking into the problem of call control using competitive analysis techniques [70]. Almost in parallel, he also looked into the problem of frequency assignment for fixed-infrastructure mobile networks (e.g., mobile telephony networks) along with Dimitris Fotakis, using graph theoretic techniques [30]. This lightning-fast (Blitzkrieg) examination gave him a good first understanding of this new technological area. Paul started to understand the intricacies and was ready to look beyond the first line of research problems and searched for the deep foundational questions.

In 1999 he forsees that apart from the mobile telephony networks, there is another, different kind of mobile networks where no fixed infrastructure exists. In cooperation with Kostas Hatzis, George Pentaris and Vasilis Tampakas they make one of the very first theoretical approaches into studying fundamental network control problems for the so-called "mobile ad-hoc networks" [32,33]. A key concept of this new approach is that nodes are free to move within the network area in any way they deem appropriate. Given this idea of free-mobility, they understand that all existing graph theoretic models where users are represented by the vertices of the graph and edges correspond to wireless communication channels for users that are within each-others communication range, are simply impractical for use in rigorous theoretical analysis. Among the main contribution of their work was the so-called "motion graph", where essentially the graph represents the area where the nodes move. Under this setting, they introduced efficient counting algorithms and proved (using markov chain theory) their correctness even under scenarios of extreme mobility.

The new model of "motion graph" was then used as the basis for studying other fundamental control problems in ad-hoc mobile network along with Ioannis Chatzigiannakis and Sotiris Nikoletseas. Having as a starting point the Markov chain theory, they looked into different network management protocols in large-scale networks, networks with faulty processors and where some nodes volunteer to support the operation of the network (e.g., see [9,11]). The key idea of the communication framework they introduced was to take advantage of the mobile nodes natural movement by exchanging information whenever mobile nodes meet incidentally. In some way, this idea resembled gossip like communication protocols where messages are spread among nodes like rumors.

Almost in parallel to the work of Paul and his team, similar gossip-like concepts were followed in [82] where an *epidemic algorithm* was designed. In their solution, messages are broadcast to all neighbors as long as there is enough storage space to hold the copies. When there is no room in the local storage of a node, oldest messages are evicted. However, in contrast to the work in [82] and other epidemic algorithms introduced later on in the relevant literature, the protocol framework of Paul and his collaborators did not rely on any assumption on the mobility patterns of the nodes. To deal with cases where nodes were spread in remote areas and would not move beyond these areas, they used these volunteer-nodes to reach them and establish connectivity by physically transporting messages between the sender and receiver nodes. Interestingly, these ideas were used to formulate the *Delay Tolerant Networking Architecture* in 2007 [5] and later on, in 2014, test them in the field as part of a project partially supported by NASA: a team of quadcopters was used to establish digital communication networks.

Another typical example of Paul's interest in emerging technologies, was the case of *Smart Dust* networks. These networks share many common points with the ad-hoc networks, however, the main difference is that nodes have extremely limited capabilities (i.e., in terms of computation, memory and energy). Once again his plan was to look into fundamental network management problems—and in particular the problem of event propagation. The first paper [10] in this new area appeared in 2002 in collaboration with Ioannis Chatzigiannakis and Sotiris Nikoletseas, and then another Blitzkrieg followed with a series of collaborators, to name a few: Peter Triantallou and Nikos Ntarmos [81], Tasos Dimitriou and Marios Mavronicolas [6,7], Ioannis Krontiris and Fotios Nikakis [20]. A few years later, the term Smart Dust networks was replaced by that of *Wireless Sensor Networks*, and a bit later with the term *Pervasive Computing*.

Eventually these technologies became broadly known as the *Internet of Things*. Regardless of the name and ephemeral keywords, the technological goal remained the same: to deploy low-cost network nodes and integrate them with the Internet with the ultimate goal to interconnect the digital and physical domains. Paul's innovative ideas of occasionally-connected networks, and the need to provide sets of protocol families are now interwined with the Internet of Things and uniformly considered by all developers and researchers as a basic concepts. Yet at the very beginning of the inception of the *Internet of Things*, this was not the case. Paul Spirakis was among the very few visionary to predict that in order to realize all these new types of applications we would have to address the problem of intermittent connectivity in networks with long delays between sending and receiving messages, or long periods of disconnection.

Moving forward from the *Smart Dust* networks and their evolution into the *Internet of Things*, Paul had in 2007 one more glimpse into the future. His vision on future systems included the orchestration of myriads of units/nodes, web services, business processes, people, companies and institutions that would be continuously integrated and connected, while preserving their individual properties, objectives and action. Paul's idea on the evolution of the *Web* and the *Internet of Things* was about tiny agents that operate using simple local rules and interact by exchanging short messages. Applications are continuous never-ending processes that eventually reach a state where the agents have developed a self-understanding of the global state.

The first attempt to deal with these systems was based on the Population Protocols model of computation that captures the way in which complex behavior of systems can emerge from the underlying local interactions of agents. Agents are usually anonymous and the local interaction rules are scalable (independent of the size n of the population). Such protocols can model the antagonism between members of several "species" and relate to evolutionary games. In the recent past, Paul was involved in joint research studying the discrete dynamics of cases of such protocols for finite populations. Such dynamics are, usually, probabilistic in nature, either due to the protocol itself or due to the stochastic nature of scheduling local interactions. Examples are (a) the generalized Moran process (where the protocol is evolutionary because a fitness parameter is crucially involved) [17] (b) the Discrete Lotka-Volterra Population Protocols (and associated Cyclic Games) [14] and (c) the Majority protocols for random interactions [52].

Such protocols are usually discrete time transient Markov chains. However the detailed states description of such chains is exponential in size and the state equations do not facilitate a rigorous approach. Instead, ideas related to filtering, stochastic domination and Potentials (leading to Martingales) help in understanding the dynamics of the protocols. Paul looked into the question of fast (in time polynomial in the population size) convergence (to an absorbing state). He addressed questions of "most probable" eventual state of the protocols (and the computation of the probability of such states). Several aspects of such discrete dynamics are wide open and it seems that the algorithmic thought can contribute to the understanding of this emerging subfield of science.

5.5 Algorithmic Game Theory

In the last 15 years, Paul has had a significant contribution to Algorithmic Game Theory (AGT) and has played a key role to the development of the field. He has been one of the pioneers in AGT and his ground-breaking work, mostly on algorithmic properties of load balancing and congestion games and on the efficient computation of exact or approximate Nash equilibria. He has significantly advanced our knowledge and has inspired many researchers to work on AGT.

Load-Balancing and Congestion Games. Paul had understood, from the very beginning, that load balancing and congestion games were bound to play a central role in AGT. His work has shaped the research agenda in the topic for many years and touches many important questions, from existence and complexity of Nash equilibria, to the inefficiency of equilibria, and to computationally efficient mechanisms for improving the Price of Anarchy. Notably, the examples below include four of Paul's most cited papers in Scopus.

In one of the first AGT papers, Marios Mavronicolas and Paul Spirakis [49] initiated the study of *fully mixed* Nash equilibria for load balancing games. They determined the fully mixed Nash equilibrium for identical links and weighted players. Building on it, they proved that the Price of Anarchy for n weighted players and m identical links is $O(\ln n/ \ln\ln n)$, for $n \geq m$. Shortly afterwards, Koutsoupias et al. [46] improved this bound to an almost tight bound of $\Theta(\ln m/ \ln\ln m)$.

Subsequently, Fotakis et al. [27] introduced the notion of *generalized fully mixed* Nash equilibria for uniformly related parallel links and proved that they almost maximize the Price of Anarchy for identical players. They also gave two proofs of the fact that load balancing games admit pure Nash equilibria. The proof techniques, one based on a greedy argument and the other on a lexicographic potential, have found numerous applications since then.

Investigating necessary and sufficient conditions for the existence of pure Nash equilibria in congestion games with weighted players, Fotakis et al. [26] proved that if the delay functions are linear, such games admit a weighted potential function. The construction is versatile and has been extended to several generalizations of congestion games. For instance, Paul and his coauthors extended it for linear congestion games with static coalitions of players [28] and for linear congestion games in a social network [24]. Moreover, Panagopoulou and Spirakis [68] presented a potential function for weighted congestion games with exponential delays. Today we know that there are no other classes of weighted congestion games with a potential function [31].

Investigating the inefficiency of equilibria in congestion games, Christodoulou et al. [13] proved general tight bounds on the Price of Anarchy and the Price of Stability for ε-approximate pure Nash equilibria of congestion games with linear delays.

Trying to eliminate the inefficiency of equilibria, Kaporis and Spirakis [37] introduced the *price of optimum*, namely the smallest fraction of coordinated players required to induce an optimal configuration in Stackelberg routing games, and showed how to efficiently compute it.

A most interesting and counterintuitive fact in selfish routing is the Braess paradox, namely that edge removal may improve the players' delay at equilibrium. Detecting and eliminating the Braess paradox is a notoriously hard computational problem. Fotakis et al. [25] proved that for non-atomic congestion games on single-commodity networks with linear delays, if the delays of almost all edges are strictly increasing, the most severe manifestations of the Braess paradox can be recognized in polynomial time.

Nash Equilibria in Bimatrix Games. One of the most appealing concepts in game theory is the notion of *Nash equilibrium*: A collection of strategies for the players, from which no player has an incentive to unilaterally deviate. The extremely nice thing about Nash equilibria is that they always exist in any finite normal-form game [64]. Nevertheless, the recent advances in the apparent hardness of the problem NASH(k) of computing an arbitrary Nash equilibrium in k-player games [15], even for 2-player (a.k.a. bimatrix) games [12], crucially question its appllicability as a realistic solution concept.

Therefore, a flurry of results have appeared in the literature during the last decade, concerning the polynomial-time construction of either an approximation of a Nash equilibrium for the general case, or an exact Nash equilibrium for certain important classes of bimatrix games. Spirakis has been one of the leading researchers in both these research trends.

With respect to the approximability of Nash equilibria, he essentially initiated the discussion on polynomial-time constructible approximations of NASH(2) [43], in parallel with the group of Papadimitriou [16]. Since then there has been a significant effort in the literature to reduce these upper bounds, for various notions of Nash equilibrium approximations. Spirakis and his colleagues have established novel techniques for constructing approximate equilibria in normalized bimatrix games, which essentially hold the record so far for the two most popular notions of Nash equilibrium approximations:

– For the most common notion of additive $\epsilon-$approximate Nash equilibria (ϵ-NE in short), in which no player may increase its payoff more than an additive term of ϵ given the strategies of the opponents, Spirakis and Tsaknakis [79]

suggested a gradient-based approach on the regret functions of the two players, and proved that it converges in polynomial time to a 0.3393-NE of an arbitrary normalized bimatrix game. For *symmetric* bimatrix games, Spirakis and Kontogiannis have provided a polynomial-time algorithm for constructing $(1/3+\delta)$-NE, for any constant $\delta > 0$ [44], which is also so far the best polynomial-time approximation guarantee. Although there is to date no theoretical evidence that one cannot go below these thresholds, all experimental evidence [23,44,80] demonstrate that it is indeed hard to break it with the existent approaches.
- For the more demanding notion of, again additive, $\epsilon-$well supported approximate Nash equilibria ($\epsilon-$WSNE in short), in which each player may assign positive probability mass only to actions that are $\epsilon-$approximate best responses to the opponent's strategy, Spirakis and Kontogiannis suggested a polynomial-time algorithm that is based on the tractability of linear programming and constructs $(2/3)-$WSNE. This result has only slightly been improved in [22] to $(2/3 - 0.005913759)$-WSNE. It is also known that going below the 2/3 threshold, one would need to check profiles of strategies with at least polylogarithmic support sizes [2], essentially ruling out enumerative approaches on the support sizes.

Concerning the second research trend, of polynomial-time algorithms for efficiently solving exactly NASH(2) in certain classes of bimatrix games, the most typical class is that of *constant-sum* games. Spirakis and Kontogiannis have determined another class of polynomial-time solvable bimatrix games, called *mutually-concave* games [44]. The main idea was based on the polynomial-time tractability of convex quadratic programs, along with a parameterization of the quadratic-program formulation of Mangasarian and Stone for Nash equilibria in bimatrix games [48]. It was proved in [44] that this class of games is essentially equivalent to the class of *strategically-zero-sum* games proposed by Moulin [63], but the approach in [44] provides a faster algorithm for computing an equillibrium than the one suggested by Moulin. The class of mutually-concave games is strictly larger than that of constant-sum games, and is incomparable to the other widely studied class of *constant-rank* games, for which it is known that a FPTAS exists [36]. Spirakis and Panagopoulou have also proved that randomly constructed bimatrix games are easy to solve [69].

5.6 Population Protocols and Temporal Graphs

In the last 7 years, Paul Spirakis has been studying in a systematic way formal models and problems inspired from Dynamic Distributed Computing Systems and Networks. His work on this modern subject can be broken down into two main strands: *Population Protocols* and *Temporal Graphs*.

Population Protocols (PPs) of Angluin et al. [3] is a distributed computing model of highly restricted computational entities (e.g., tiny sensor nodes or nanorobots) that move and interact *passively*, following the dynamicity of their environment. In terms of modeling, the entities are automata that interact in pairs according to a fair (or random) scheduler. The main result known when Spirakis started working on the subject was that PPs are very limited computationally, being able to compute only *semilinear predicates* on input assignments. The work carried out by Spirakis, together with Othon Michail, Ioannis Chatzigiannakis, and other collaborators, is prominent between those that established PPs as a very active and rapidly growing sub-area of *Distributed Computing*. Working on their *Mediated Population Protocols* (MPPs) [55], an extension that allows the automata to additionally establish (physical or virtual) links with each other, they proved a quite surprising (at that time) result: *the n automata can now be programmed to simulate any nondeterministic Turing machine of space*$O(n^2)$. Recently, building upon the MPP model, Othon Michail and Paul Spirakis initiated the study of protocols that can algorithmically construct stable networks [59], trying at the same time to answer computability questions and questions related to the potential of their model to represent chemical and biological self-assembly processes. Between the several other important contributions of Spirakis and his collaborators on PPs, one can distinguish the parameterized study of PPs according to the size of their local memories [8], papers that have promoted our understanding on the *Termination* problem [61,62], his work with George Mertzios, Sotiris Nikoletseas, and Christoforos Raptopoulos on the *Majority* problem [52], and with Czyzowicz et al. on the dynamics of discrete Lotka-Volterra PPs [14]. His involvement in PPs has led to a remarkable research record since 2008: 1 research monograph [54], 1 book chapter, around 20 publications in prestigious journals and conferences, and numerous invited talks and lectures.

Though very recent, and with its full story remaining to be written, the work of Spirakis on PPs has already been recognized at the highest level, being one of the four research areas for which he received his 2014 EATCS Fellowship (see Sect. 4).

Inspired by the theoretical developments in dynamic distributed computing systems, together with Othon Michail and other collaborators, Paul Spirakis started envisioning a *Temporal* extension of *Graph Theory* and he is currently working intensively to put down its founding stone. Informally, a *temporal graph* is a graph that changes with time. It can be thought of as a special case of labeled graphs, where the (usually discrete) labels represent some measure of time and constrain connectivity to respect the additional time-dimension, captured in the notion of *time-respecting* paths. The work of Spirakis on the subject begins with [50]. That paper has received attention, as it gave the first temporal analogue of Menger's theorem (in contrast to a famous negative result proving that the classical formulation of Menger's theorem fails on temporal graphs [40]) and introduced the intriguing problem of designing a cost-optimal temporal graph under various combinations of connectivity constraints and cost parameters. Together with Othon Michail, he defined and initiated the study of temporal versions of several well-known combinatorial optimization graph-problems, such as the *Temporal Exploration* problem, the *Temporal Traveling Salesman Problem with Costs One and Two*, and the *Temporal Matching* problem, and gave a first set of approximation algorithms and hardness results for them [60]. Moreover, Spirakis may be considered as one of the very few people that have initiated the formal study of *Random Temporal Graphs*, on which he is currently working with Eleni Akrida and other collaborators [1]. For a more detailed introduction to the work of Spirakis and others on the subject of *Temporal Graphs*, the reader is referred to [53], which is included in this volume.

Finally, we should mention that Paul Spirakis, together with Othon Michail and Ioannis Chatzigiannakis, has worked on the modern area of distributed computation in highly-dynamic networks, where protocols must perform their task against a *worst-case adversary* that controls and alters the communication topology. His contribution to the subject can be found in [56–58].

6 Other Professional Activities

Paul Spirakis has been extremely successful in receiving project grants—a condition *sine qua non* for establishing a research group and descent conditions for pursuing research.

He devoted a considerable amount of his time in submitting proposals and getting highly-competitive project grants, most of which came from the European Union (other funding organizations included NSF and the Greek State). Overall, he has been involved in about 50 projects—a fair portion of which he acted as coordinator.

One of Paul's major efforts, since 1996, was to run (as Chairman of the Board of Directors) the Computer Technology Institute & Press "Diophantus" (CTI), one of the largest Research & Development Institutes in Greece. He carried out this challenging job with great success, managing to run CTI without governmental funding. CTI is a self-funded organization employing more than 300 experienced and specialized scientists (including faculty members of the Department of Computer Engineering and Informatics at the University of Patras), engineers, PhD students, and supporting staff. In the last 5 years, CTI has successfully undertaken more than 130 R&D projects, 95 of which were funded by the EC. In the last 20 years, CTI has exhibited substantial basic and applied research activity in areas such as algorithms and complexity, optimization, wired as well as wireless and sensor networks, computer and network security, ubiquitous and distributed computing, e-learning, complex information systems design and development, production systems, embedded systems, integration and sustainable development.

One of the memorable great successes of Paul Spirakis was the development of close ties between CTI and the Greek State, making CTI the official technical consultant of the Ministry of Education on IT issues continuously in the past 20 years. CTI places particular emphasis on education, by developing and deploying conventional and digital media in education and lifelong learning, publishing printed and electronic educational materials, administrating and managing the Greek School Network, and supporting the organization and operation of the electronic infrastructure of the Greek Ministry of Education and all educational units. In addition, CTI undertakes the lifelong training of school teachers in the new education technology and material.

Because of his extensive involvement in the theory and applications of Computer Science (CS) as well as of the Information & Communication Technologies (ICT) in general, Paul Spirakis had a prominent role in promoting the importance and in determining the research agenda of CS and ICT within Europe. He also served as a consultant in numerous bodies of European Union and the Greek State.

Paul writes occasionally articles in Greek newspapers, mostly in special columns devoted to science and technology, as well as articles expressing his liberal opinions about several aspects of our political and socio-economic environment.

7 Contributions to the Scientific Community

Paul Spirakis delivered more than 125 lectures as invited speaker of various scientific conferences, research institutes, and universities around the globe.

He serves in the Editorial Boards of prestigious computer science journals, including *Acta Informatica, Algorithmica, Computer Science Review, International Journal of Computing Theory, Journal of Parallel and Distributed Computing, Parallel Processing Letters, Theoretical Computer Science*, and *Theory of Computing Systems*. He has also been editor of *Computational Geometry Journal—Theory and Applications*, and area editor of the *Encyclopedia of Algorithms*, published by Springer.

Paul Spirakis was the chairman of more than 25 and a member in more than 130 Program Committees of major international conferences in computer science (incuding STOC, ESA, IPDPS, ICALP, PODC, SPAA, and STACS). He has also been editor in several special issues of prestigious journals dedicated to selected papers from some of these major conferences.

8 Personal

Paul Spirakis is married to (the mathematician) Asimina Chrisofaki. They have two daughters, Olga, an actress and choreographer, and Zeta (Georgia), a computer scientist who currently pursues postgraduates studies (MA) in *The London Film School*. Paul has also a grandson from Olga. He greatly believes in the importance of family, having said among others that "investing time on family is a very good investment".

A first acquaintance with Paul makes you immediately realizing that you are confronted with a perspicacious and open-minded man. Paul is warm, generous, and cultured. Despite being an eminent scientist, he remains simple and approachable, and his door is always open to students and young aspiring researchers. Paul greatly believes in the new generation and in the proper education of young people. He is extremely patient with students believing in their potential and helping each one of them self-realize his/her own special talents.

This makes him at the same time a teacher, a good friend, and a father of his students.

Another immediate observation from a first acquaintance with Paul is that he is a heavy smoker and a passionate drinker of Greek Frappe (cold instant) coffee.

He is characterized by an extremely good sense of humor and high optimism, which make him a pleasant company. He laughs out loudly, especially when he listens to or tells a funny joke (something that he enjoys a lot).

Paul tries every day to learn or discover something new in mathematics. If not, he (claims that he) feels guilty. He also studies books and papers from other scientific disciplines (physics, biology, etc.), because he truly believes that multidisciplinarity is the key to new ideas and approaches. Time seems to have almost no effect on him. He is full of energy and always keen and ready to tackle the next open problem.

In his free time, he likes to read. Among the great variety of subjects that he reads, he is a passionate reader of crime and spy novels, science fiction, and history. Paul is also a very talented chess player.

Paul has a unique talent in accomplishing heavy managerial and administrative duties and at the same time high-level research and teaching. His passion for research was evident at Friday nights, when, after a whole day (and preceding week) of administrative meetings, he remained in the building of CTI (usually until after midnight) in order to have some fruitful research meetings with his students and collaborators.

Paul deeply believes in the ideals of democracy, social justice, and the promotion of individual talents and initiatives.

9 Epilogue

This article offers a glimpse at the life and career of Paul Spirakis. It is an effort (hopefully successful) to present main aspects of his personal and professional life and to emphasize on his visions and achievements. In any case, the goal was to highlight Paul's excessive offering as a teacher, researcher and scientist, and to express our gratitude and appreciation to him.

Acknowledgements. We are indebted to Eleni Spiraki (Paul's sister) for her great help in the preparation of this article.

References

1. Akrida, E., Gasieniec, L., Mertzios, G., Spirakis, P.: Ephemeral networks with random availability of links: diameter and connectivity. In: Proceedings of 26th ACM Symposium on Parallelism in Algorithms and Architectures—SPAA 2014, pp. 267–276 (2014)

2. Anbalagan, Y., Norin, S., Savani, R., Vetta, A.: Polylogarithmic supports are required for approximate well-supported Nash equilibria below 2/3. In: Chen, Y., Immorlica, N. (eds.) WINE 2013. LNCS, vol. 8289, pp. 15–23. Springer, Heidelberg (2013)
3. Angluin, D., Aspnes, J., Diamadi, Z., Fischer, M.J., Peralta, R.: Computation in networks of passively mobile finite-state sensors. Distrib. Comput. **18**(4), 235–253 (2006)
4. Bilardi, G., Herley, K., Pietracaprina, A., Pucci, G., Spirakis, P.: BSP versus LogP. Algorithmica **24**, 405–422 (1999). Preliminary version in SPAA 1996
5. Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, B., Scott, K., Fall, K., Weiss, H.: Delay-tolerant networking architecture. Technical report, The IETF Trust (2007)
6. Chatzigiannakis, I., Dimitriou, T.D., Mavronicolas, M., Nikoletseas, S.E., Spirakis, P.G.: A comparative study of protocols for efficient data propagation in smart dust networks. In: Kosch, H., Böszörményi, L., Hellwagner, H. (eds.) Euro-Par 2003. LNCS, vol. 2790, pp. 1003–1016. Springer, Heidelberg (2003)
7. Chatzigiannakis, I., Dimitriou, T., Mavronicolas, M., Nikoletseas, S., Spirakis, P.: A comparative study of protocols for efficient data propagation in smart dust networks. Parallel Process. Lett. **13**(4), 615–627 (2003)
8. Chatzigiannakis, I., Michail, O., Nikolaou, S., Pavlogiannis, A., Spirakis, P.: Passively mobile communicating machines that use restricted space. Theor. Comput. Sci. **412**(46), 6469–6483 (2011)
9. Chatzigiannakis, I., Nikoletseas, S.E., Paspallis, N., Spirakis, P.G., Zaroliagis, C.D.: An experimental study of basic communication protocols in ad-hoc mobile networks. In: Brodal, G.S., Frigioni, D., Marchetti-Spaccamela, A. (eds.) WAE 2001. LNCS, vol. 2141, pp. 159–171. Springer, Heidelberg (2001)
10. Chatzigiannakis, I., Nikoletseas, S., Spirakis, P.: Smart dust protocols for local detection and propagation. In: Proceedings of 2nd ACM International Workshop on Principles of Mobile Computing—POMC 2002, pp. 9–16 (2002)
11. Chatzigiannakis, I., Nikoletseas, S., Spirakis, P.: Distributed communication algorithms for ad hoc mobile networks. J. Parallel Distrib. Comput. **63**(1), 58–74 (2003)
12. Chen, X., Deng, X., Teng, S.-H.: Settling the complexity of computing two-player Nash equilibria. J. ACM **56**, 3, Art. No. 14 (2009)
13. Christodoulou, G., Koutsoupias, E., Spirakis, P.: On the performance of approximate equilibria in congestion games. Algorithmica **61**(1), 116–140 (2011)
14. Czyzowicz, J., Gąsieniec, L., Kosowski, A., Kranakis, E., Spirakis, P.G., Uznański, P.: On convergence and threshold properties of discrete Lotka-Volterra population protocols. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015. LNCS, vol. 9134, pp. 393–405. Springer, Heidelberg (2015)
15. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. SIAM J. Comput. **39**(1), 195–259 (2009). Preliminary version in ACM STOC 2006
16. Daskalakis, C., Mehta, A., Papadimitriou, C.: A note on approximate Nash equilibria. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 297–306. Springer, Heidelberg (2006)
17. Díaz, J., Goldberg, L.A., Mertzios, G.B., Richerby, D., Serna, M.J., Spirakis, P.: Approximating fixation probabilities in the generalized Moran process. Algorithmica **69**(1), 78–91 (2014). Preliminary version in SODA 2012
18. Diaz, J., Koukopoulos, D., Nikoletseas, S., Spirakis, P., Serna, M., Thilikos, D.: Stability and instability of the FIFO protocol. In: Proceedings of 13th ACM Symposium on Parallel Algorithms and Architectures—SPAA 2001, pp. 48–52 (2001)

19. Díaz, J., Serna, M.J., Spirakis, P.: On the random generation and counting of matchings in dense graphs. Theor. Comput. Sci. **201**(1–2), 281–290 (1998)
20. Dimitriou, T.D., Krontiris, I., Nikakis, F., Spirakis, P.G.: SPEED: Scalable Protocols for Efficient Event Delivery in sensor networks. In: Mitrou, N.M., Kontovasilis, K., Rouskas, G.N., Iliadis, I., Merakos, L. (eds.) NETWORKING 2004. LNCS, vol. 3042, pp. 1300–1305. Springer, Heidelberg (2004)
21. Fatourou, P., Spirakis, P.: Efficient scheduling of strict multithreaded computations. Theory Comput. Syst. **33**(3), 173232 (2000)
22. Fearnley, J., Goldberg, P.W., Savani, R., Sørensen, T.B.: Approximate well-supported Nash equilibria below two-thirds. In: Serna, M. (ed.) SAGT 2012. LNCS, vol. 7615, pp. 108–119. Springer, Heidelberg (2012)
23. Fearnley, J., Igwe, T.P., Savani, R.: An empirical study of finding approximate equilibria in bimatrix games. ArXiv/CoRR, abs/1502.04980 (2015)
24. Fotakis, D., Gkatzelis, V., Kaporis, A., Spirakis, P.: The impact of social ignorance on weighted congestion games. Theory Comput, Syst. **50**(3), 559–578 (2012)
25. Fotakis, D., Kaporis, A., Spirakis, P.: Efficient methods for selfish network design. Theor. Comput. Sci. **448**, 9–20 (2012)
26. Fotakis, D., Kontogiannis, S., Spirakis, P.: Selfish unsplittable flows. Theor. Comput. Sci. **348**, 226–239 (2005)
27. Fotakis, D., Kontogiannis, S., Koutsoupias, E., Mavronicolas, M., Spirakis, P.: The structure and complexity of Nash equilibria for a selfish routing game. Theor. Comput. Sci. **410**(36), 3305–3326 (2009)
28. Fotakis, D., Kontogiannis, S., Spirakis, P.: Atomic congestion games among coalitions. ACM Trans. Algorithms **4**, 4 (2008)
29. Fotakis, D., Nikoletseas, S., Papadopoulou, V., Spirakis, P.: Radiocoloring in planar graphs: complexity and approximations. Theor. Comput. Sci. **340**(3), 514–538 (2005)
30. Fotakis, D.A., Spirakis, P.G.: A Hamiltonian approach to the assignment of non-reusable frequencies. In: Arvind, V., Sarukkai, S. (eds.) FST TCS 1998. LNCS, vol. 1530, pp. 18–30. Springer, Heidelberg (1998)
31. Harks, T., Klimm, M., Möhring, R.H.: Characterizing the existence of potential functions in weighted congestion games. Theory Comput. Syst. **49**(1), 46–70 (2011)
32. Hatzis, K., Pentaris, G., Spirakis, P.G., Tampakas, B.: Counting in mobile networks: theory and experimentation. In: Vitter, J.S., Zaroliagis, C.D. (eds.) WAE 1999. LNCS, vol. 1668, pp. 95–109. Springer, Heidelberg (1999)
33. Hatzis, K., Pentaris, G., Spirakis, P., Tampakas, V., Tan, R.B.: Fundamental control algorithms in mobile networks. In: Proceedings of ACM SPAA 1999, pp. 251–260 (1999)
34. Jung, J., Kirousis, L.M., Spirakis, P.: Lower bounds and efficient algorithms for multiprocessor scheduling of directed acyclic graphs with communication delays. Inf. Comput. **105**, 94–104 (1993)
35. Kamath, A., Motwani, R., Palem, K., Spirakis, P.: Tail bounds for occupancy and the satisfiability threshold conjecture. In: Proceedings of 35th IEEE Symposium on Foundations of Computer Science—FOCS 1994, pp. 592–603 (1994)
36. Kannan, R., Theobald, T.: Games of fixed rank: a hierarchy of bimatrix games. Econ. Theory **42**, 157–173 (2010). Preliminary version in ACM-SIAM SODA 2007
37. Kaporis, A., Spirakis, P.: The price of optimum in Stackelberg games on arbitrary single commodity networks and latency functions. Theor. Comput. Sci. **410**(8–10), 745–755 (2009)

38. Kavvadias, D., Pantziou, G., Spirakis, P., Zaroliagis, C.: Hammock-on-Ears decomposition: a technique for the efficient parallel solution of shortest paths and other problems. Theor. Comput. Sci. **168**(1), 121–154 (1996). Preliminary version in MFCS 1994
39. Kedem, Z., Palem, K., Spirakis, P.G.: Efficient robust parallel computations. In: Proceedings of 22nd Annual ACM Symposium on Theory of Computing—STOC 1990, pp. 138–148 (1990)
40. Kempe, D., Kleinberg, J., Kumar, A.: Connectivity and inference problems for temporal networks. In: Proceedings of 32nd ACM Symposium on Theory of Computing—STOC 2000, pp. 504–513 (2000)
41. Kirousis, L., Serna, M., Spirakis, P.: The parallel complexity of the connected subgraph problem. SIAM J. Comput. **22**(3), 573–586 (1993). Preliminary version in FOCS 1989
42. Kirousis, L., Spirakis, P., Tsigas, P.: Reading many variables in one atomic operation solutions with linear or sublinear complexity. IEEE Trans. Parallel Distrib. Comput. **5**(7), 688–696 (1994)
43. Kontogiannis, S.C., Panagopoulou, P.N., Spirakis, P.G.: Polynomial algorithms for approximating nash equilibria of bimatrix games. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 286–296. Springer, Heidelberg (2006)
44. Kontogiannis, S., Spirakis, P.: Approximability of symmetric bimatrix games and related experiments. In: Pardalos, P.M., Rebennack, S. (eds.) SEA 2011. LNCS, vol. 6630, pp. 1–20. Springer, Heidelberg (2011)
45. Koukopoulos, D., Mavronicolas, M., Nikoletseas, S., Spirakis, P.: The impact of network structure on the stability of greedy protocols. Theory Comput. Syst. **38**(4), 425–460 (2005)
46. Koutsoupias, E., Mavronicolas, M., Spirakis, P.: Approximate equilibria and ball fusion. Theory Comput. Syst. **36**, 683–693 (2003)
47. Kumar, B., Srikant, Y.N.: The best nurturers in computer science research. In: Proceedings of 2005 SIAM International Conference on Data Mining, pp. 566–570; also Technical Report IISc-CSA-TR-2004-10, Computer Science and Automation, Indian Institute of Science, India, October 2004
48. Mangasarian, O.L., Stone, H.: Two-person nonzero-sum games and quadratic programming. J. Math. Anal. Appl. **9**(3), 348–355 (1964)
49. Mavronicolas, M., Spirakis, P.: The price of selfish routing. Algorithmica **48**(1), 91–126 (2007)
50. Mertzios, G.B., Michail, O., Chatzigiannakis, I., Spirakis, P.G.: Temporal network optimization subject to connectivity constraints. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M., Peleg, D. (eds.) ICALP 2013, Part II. LNCS, vol. 7966, pp. 657–668. Springer, Heidelberg (2013)
51. Mertzios, G.B., Nikoletseas, S., Raptopoulos, C., Spirakis, P.: Natural models for evolution on networks. Theor. Comput. Sci. **477**, 76–95 (2013)
52. Mertzios, G.B., Nikoletseas, S.E., Raptopoulos, C.L., Spirakis, P.G.: Determining majority in networks with local interactions and very small local memory. In: Esparza, J., Fraigniaud, P., Husfeldt, T., Koutsoupias, E. (eds.) ICALP 2014. LNCS, vol. 8572, pp. 871–882. Springer, Heidelberg (2014)
53. Michail, O.: An introduction to temporal graphs—an algorithmic perspective. In: Algorithms, Probability, Networks, and Games. LNCS, vol. 9295. Springer, New York (2015)

54. Michail, O., Chatzigiannakis, I., Spirakis, P.: New Models for Population Protocols. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, San Rafael (2011)
55. Michail, O., Chatzigiannakis, I., Spirakis, P.: Mediated population protocols. Theor. Comput. Sci. **412**(22), 2434–2450 (2011)
56. Michail, O., Chatzigiannakis, I., Spirakis, P.G.: Naming and counting in anonymous unknown dynamic networks. In: Higashino, T., Katayama, Y., Masuzawa, T., Potop-Butucaru, M., Yamashita, M. (eds.) SSS 2013. LNCS, vol. 8255, pp. 281–295. Springer, Heidelberg (2013)
57. Michail, O., Chatzigiannakis, I., Spirakis, P.: Causality, influence, and computation in possibly disconnected synchronous dynamic networks. J. Parallel Distrib. Comput. **74**(1), 2016–2026 (2014)
58. Michail, O., Chatzigiannakis, I., Spirakis, P.: Computing in dynamic networks. In: Dehmer, M., Emmert-Streib, F., Pickl, S. (eds.) Chapter 6 in Computational Network Theory: Theoretical Foundations and Applications, 1st edn, pp. 173–218, Wiley-VCH Verlag GmbH & Co. KGaA (2015)
59. Michail, O., Spirakis, P.: Simple and efficient local codes for distributed stable network construction. In: Proceedings of 33rd ACM Symposium on Principles of Distributed Computing—PODC 2014, pp. 76–85 (2014)
60. Michail, O., Spirakis, P.G.: Traveling salesman problems in temporal graphs. In: Csuhaj-Varjú, E., Dietzfelbinger, M., Ésik, Z. (eds.) MFCS 2014, Part II. LNCS, vol. 8635, pp. 553–564. Springer, Heidelberg (2014)
61. Michail, O., Spirakis, P.: Terminating population protocols via some minimal global knowledge assumptions. J. Parallel Distrib. Comput. **81**, 1–10 (2015)
62. Michail, O., Spirakis, P.: Distributed computation by connectivity preserving network transformers (2015) (to appear)
63. Moulin, H., Vial, J.P.: Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. Int. J. Game Theory **7**(3–4), 201–221 (1978)
64. Nash, J.: Noncooperative games. Ann. Math. **54**, 289–295 (1951)
65. Nikoletseas, S., Palem, K.V., Spirakis, P., Yung, M.: Short vertex disjoint paths and multiconnectivity in random graphs: reliable network computing. In: Proceedings of 21st International Colloquium on Automata, Languages and Programming—ICALP 1994, pp. 508–519 (1994)
66. Nikoletseas, S., Prasinos, G., Spirakis, P., Zaroliagis, C.: Attack propagation in networks. Theory Comput. Syst. **36**, 553–574 (2003). Preliminary version in ACM SPAA 2001
67. Nikoletseas, S.E., Raptopoulos, C., Spirakis, P.G.: The existence and efficient construction of large independent sets in general random intersection graphs. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 1029–1040. Springer, Heidelberg (2004)
68. Panagopoulou, P., Spirakis, P.: Algorithms for pure Nash equilibria in weighted congestion games. ACM J. Exp. Algorithmics **11**, 1–19 (2006)
69. Panagopoulou, P., Spirakis, P.: Random bimatrix games are asymptotically easy to solve (A simple proof). Theor. Comput. Sci. **54**(3), 479–490 (2014)
70. Pantziou, G., Pentaris, G., Spirakis, P.: Competitive call control in mobile networks. In: Algorithms and Computation—ISAAC 1997, pp. 404–413 (1997)
71. Pantziou, G., Spirakis, P., Zaroliagis, C.: Efficient parallel algorithms for shortest paths in planar digraphs. BIT **32**(2), 215–236 (1992). Preliminary vesrion in SWAT 1990

72. Reif, J., Spirakis, P.: Random matroids. In: Proceedings 12th ACM Symposium on Theory of Computing—STOC 1980, pp. 385–397 (1980)
73. Reif, J., Spirakis, P.: Distributed algorithms for synchronizing interprocess communication within real time. In: Proceedings of 13th ACM Symposium on Theory of Computing—STOC 1981, pp. 133–145 (1981)
74. Reif, J., Spirakis, P.: Real time synchronization of interprocess communications. ACM Trans. Programm. Lang. Syst. **6**(2), 215–238 (1984)
75. Reif, J., Spirakis, P.: Unbounded speed variability in distributed systems. SIAM J. Comput. **14**(1), 7592 (1985)
76. Reif, J., Spirakis, P.: Expected parallel time and sequential space complexity of graph and digraph problems. Algorithmica **7**, 597–630 (1992)
77. Spirakis, P.: Optimal parallel randomized algorithms for addition sparse addition and identification. Inf. Comput. **76**, 1–12 (1988)
78. Spirakis, P., Zaroliagis, C.: Distributed algorithm engineering. In: Fleischer, R., Moret, B.M.E., Schmidt, E.M. (eds.) Experimental Algorithmics. LNCS, vol. 2547, pp. 197–228. Springer, Heidelberg (2002)
79. Tsaknakis, H., Spirakis, P.G.: An optimization approach for approximate Nash equilibria. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 42–56. Springer, Heidelberg (2007)
80. Tsaknakis, H., Spirakis, P.G., Kanoulas, D.: Performance evaluation of a descent algorithm for bi-matrix games. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 222–230. Springer, Heidelberg (2008)
81. Triantafillou, P., Ntarmos, N., Nikoletseas, S., Spirakis, P.: Nanopeer networks and P2P worlds. In: Proceedings of 3rd International Conference on Peer-to-Peer Computing—P2P 2003, pp. 40–46 (2003)
82. Vahdat, A., Becker, D.: Epidemic routing for partially connected ad hoc networks. Technical report CS-200006, Duke University (2000)