IP-FP6-015964

Algorithmic Principles for Building Efficient Overlay Computers

AEOLUS Workshop 2008

September 8-10, 2008

Sophia Antipolis, France

The 2008 AEOLUS Workshop is organized by CNRS/INRIA (David Coudert, Herve Rivano - Workshop Chairs) in Sophia-Antipolis, France on September 8-10, 2008, collocated with the 7th International Conference on Ad-Hoc Networks & Wireless (AdHoc-Now 08). The workshop website is at:
`http: //www-sop.inria.fr /mascotte /adhocnow /aeolus /#`
AEOLUS sites have contributed with presentations that are based either on published papers or on research work in progress related to AEOLUS. Furthermore, representatives from two other EU projects, ARRIVAL (FP6-021235-2: Algorithms for Robust and online Railway optimization Improving the Validity and reliAbility of Large scale systems) and NEON (IST-2005-027595: Lifecycle Support for Networked Ontologies), have presented the projects as well as research results achieved in their context.

Near-Optimal Replication in Unstructured P2P Networks

Mauro Sozio
MPII, Germany

Replicating data in distributed systems is often needed for availability and performance. In unstructured peer-to-peer networks, with epidemic messaging for query routing, replicating popular data items is also crucial to ensure high probability of finding the data within a bounded search distance from the requestor. This paper considers such networks and aims to maximize the probability of successful search. Prior work along these lines has analyzed the optimal degrees of replication for data items with non-uniform but global request rates, but did not address the issue of where replicas should be placed and was very very limited in the capabilities for handling heterogeneity and dynamics of network and workload.

This paper presents the integrated P2R2 algorithm for dynamic replication that addresses all these issues, and determines both the degrees of replication and the placement of the replicas in a provably near-optimal way. We prove that the P2R2 algorithm can guarantee a successful-search probability that is within a factor of 2 of the optimal solution. The algorithm is efficient and can handle workload evolution. We prove that, whenever the access patterns are in steady state, our algorithm converges to the desired near-optimal placement. We further show by simulations that the convergence rate is fast and that our algorithm outperforms prior methods.

Reference

Mauro Sozio, Thomas Neumann, Gerhard Weikum.
Near-Optimal Dynamic Replication in Unstructured Peer-to-Peer Networks.
ACM Symposium on Principles of Database Systems (PODS), 2008.

Towards a Practical and Fair Rate Allocation for Multihop Wireless Networks

Rémi Vannier
INRIA, France

IEEE 802.11 is often considered as the underlying wireless technology of multihop wireless networks. But the use of 802.11 in such networks raises issues, like efficiency and/or fairness issues. Different kinds of solutions have been proposed to overcome these problems. One approach is to design new MAC protocols that provide alternatives to the IEEE 802.11 MAC protocol. Although these solutions are of some interest, it should probably take some time before new wireless network interface cards based on one of these solutions are developed and released. Another approach is to consider that 802.11 will remain the underlying wireless technology and to design solutions above it. Several solutions based on rate allocation have been proposed so far. The main drawback of the proposed solutions is that they rely on a radio medium sharing model that is difficult to compute in a wireless, distributed and mobile environment. Indeed, very few of these solutions have been derived into a network protocol.

In this article, we propose a distributed and dynamic rate allocation solution that is based on a simple sharing model. Due to its simplicity, we can derive a network protocol that can be practically used in multihop wireless networks. This protocol provides a fair bandwidth sharing between end-to-end flows while maintaining an efficient overall throughput in the network. This solution has been implemented in NS2 and evaluated by simulations.

Reference
Rémi Vannier and Isabelle Guerin-Lassous.
Towards a Practical and Fair Rate Allocation for Multihop Wireless Networks based on a Simple Node Model.
INRIA Technical Report 6538, 2008.

Sharemind: a framework for fast privacy-preserving computations

Dan Bogdanov
CYB, Esthonia

Gathering and processing sensitive data is a difficult task. In fact, there is no common recipe for building the necessary information systems. In this paper, we present a provably secure and efficient general-purpose computation system to address this problem. Our solution - SHAREMIND - is a virtual machine for privacy-preserving data processing that relies on share computing techniques. This is a standard way for securely evaluating functions in a multi-party computation environment. The novelty of our solution is in the choice of the secret sharing scheme and the design of the protocol suite. We have made many practical decisions to make large-scale share computing feasible in practice. The protocols of SHAREMIND are information-theoretically secure in the honest-but-curious model with three computing participants. Although the honest-but-curious model does not tolerate malicious participants, it still provides significantly increased privacy preservation when compared to standard centralised databases.

Reference

Dan Bogdanov and Sven Laur and Jan Willemson.
Sharemind: a framework for fast privacy-preserving computations.
In Proc. of the 13th European Symposium on Research in Computer Security (ESORICS 08), 2008, to appear.

Long Running Trust Negotiations

Alberto Trombetta
Insubria, Italy

In general, trust negotiation is an authorization approach for establishing trust in an automated, distributed fashion. Specifically, a typical trust negotiation consists of an interactive process where two peers gradually establish trust by requesting and disclosing each others' properties, encoded through credentials. This approach differs from traditional authentication-based access control systems mainly because of the lack of pre-existing trust relationship between the involved peers. Moreover, within trust negotiations each peer can define access control policies to govern access to its sensitive credentials. Identity trust negotiations focus on the negotiation of users' identity attributes. This protocol share many similarities with conventional attributes, and presents unique challenges, mainly related to the distribution of identity attributes across different identity providers and locations. As a result, identity related trust negotiations may last a non-negligible timespan, because of the distributed nature of identity providers and possible location of identifiers. The longer the timespan, the higher is the probability that the negotiation has to be interrupted. Interruption in ongoing trust negotiations can be generated by external, unforeseeable events (e.g. peers' crashes, faulty transmission channel) or upon decision of the involved peers. A peer may in fact not be able to advance the negotiation for temporary lack of resources, such as network connection or memory capability. Or, the peer may not have readily available the credentials required by the counterpart, although eligible to it. In this work we address the latter case: the negotiation is gracefully suspended by the involved peers in a way that it can be resumed at a (possibly distant and agreed) future time.

Reference

Alberto Trombetta, Anna Squicciarini, Elisa Bertino, Stefano Braghin.
Long Running Trust Negotiations.
In Proc. of the 4th International ACM CCS Workshop on Digital Identity Management (ACM CCS2008 DIM), 2008, to appear.

BITPEER: continuous subspace skyline computation with distributed bitmap indexes

Evaggelia Pitoura
UOI, Greece

In this paper, we propose a bitmap approach for efficient subspace skyline computation in a distributed setting. Our approach computes extended skylines which have been shown to include all points necessary for computing the skyline at any subspace. We present an algorithm for computing extended skylines using a bitmap representation along with a storage efficient bucket-based variation of this representation. We provide a caching scheme so that subspace skyline queries can re-use the results of previously computed similar queries. We also introduce a method for grouping continuous subspace queries for supporting their efficient computation. Finally, we present preliminary experimental results of the performance of our approach.

Reference

Katerina Fotiadou and Evaggelia Pitoura.
BITPEER: continuous subspace skyline computation with distributed bitmap indexes.
In Proc. of the 2008 International Workshop on Data management in Peer-to-Peer Systems (DAMAP 08), pp. 35-42, 2008.

Network-Oblivious Algorithms

Francesco Silvestri
UNIPI, Italy

The design of algorithms that can run unchanged yet efficiently on a variety of machines characterized by different degrees of parallelism and communication capabilities is a highly desirable goal. We propose a framework for network-obliviousness based on a model of computation where the only parameter is the problem' s input size. Algorithms are then evaluated on a model with two parameters, capturing parallelism and granularity of communication. We show that, for a wide class of network-oblivious algorithms, optimality in the latter model implies optimality in a block-variant of the Decomposable BSP model, which effectively describes a wide and significant class of parallel platforms. We illustrate our framework by providing optimal network-oblivious algorithms for a few key problems, and also establish some negative results.

Reference

Gianfranco Bilardi and Andrea Pietracaprina and Geppino Pucci and Francesco Silvestri. Network-Oblivious Algorithms.
In Proc. of the 21st International Parallel and Distributed Processing Symposium (IPDPS 07), 2007.

Efficient Semi-Streaming Algorithms for Local Triangle Counting in Massive Graphs

Luca Becchetti
UDRLS, Italy

In this paper we study the problem of local triangle counting in large graphs. Namely, given a large graph $G = (V, E)$ we want to estimate as accurately as possible the number of triangles incident to every node $v \in V$ in the graph. The problem of computing the *global* number of triangles in a graph has been considered before, but to our knowledge this is the first paper that addresses the problem of *local* triangle counting with a focus on the efficiency issues arising in massive graphs. The distribution of the local number of triangles and the related local clustering coefficient can be used in many interesting applications. For example, we show that the measures we compute can help to detect the presence of spamming activity in large-scale Web graphs, as well as to provide useful features to assess content quality in social networks. For computing the local number of triangles we propose two approximation algorithms, which are based on the idea of min-wise independent permutations (Broder et al. 1998). Our algorithms operate in a semi-streaming fashion, using $O(|V|)$ space in main memory and performing $O(\log |V|)$ sequential scans over the edges of the graph. The first algorithm we describe in this paper also uses $O(|E|)$ space in external memory during computation, while the second algorithm uses only main memory. We present the theoretical analysis as well as experimental results in massive graphs demonstrating the practical efficiency of our approach.

Reference

Luca Becchetti, Paolo Boldi, Carlos Castillo, Aristides Gionis.
Efficient Semi-Streaming Algorithms for Local Triangle Counting in Massive Graphs.
In Proc. of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 08), 2008, to appear.

Approximation Algorithms for Scheduling Parallel Jobs: Breaking the Approximation Ratio of 2

Klaus Jansen
CAU, Germany

In this paper we study variants of the non-preemptive parallel job scheduling problem where the number of machines is polynomially bounded in the number of jobs. For this problem we show that a schedule with length at most $(1+\epsilon)OPT$ can be calculated in polynomial time, which is the best possible result (in the sense of approximation ratio), since the problem is strongly NP-hard. For the case when all jobs must be allotted to a subset of machines with consecutive indices a schedule with length at most $(1.5+\epsilon)OPT$ can be calculated in polynomial time. The previously best known results are algorithms with absolute approximation ratio 2. Furthermore, we extend both algorithms to the case of malleable jobs with the same approximation ratio.

Reference

Jansen, Klaus and Thöle, Ralf.
Approximation Algorithms for Scheduling Parallel Jobs: Breaking the Approximation Ratio of 2.
In Proc. of the 35th International Colloquium on Automata, Languages and Programming (ICALP 08), pp. 234-245, 2008.

Online scheduling on semi-related machines

Ulrich Schwarz
CAU, Germany

We consider the problem of preemptive makespan minimization on related machines, where machines are subject to breakdowns that are not known statically.

Two algorithms are presented: one guarantees optimal makespan if machine breakdown is announced a small amount of time in advance. This warning phase can be reduced at the cost of more preemptions.

The second one does not need a warning phase, but only achieves a makespan of $OPT + \epsilon$. Here, the number of preemptions increases with ϵ^{-1}, but also with $s_{\max}/s_{\min}$, where $s_{\max}$ and $s_{\min}$ are the speeds of the fastest and slowest machine involved.

Reference

Ulrich Schwarz.
Online scheduling on semi-related machines.
Information Processing Letters, Volume 108, Issue 1, pp. 38-40, 2008.

Early Obstacle Detection in Wireless Sensor Networks Routing

Luminita Moraru
CUI, Switzerland

Geographic routing scales well in sensor networks, mainly due to its stateless nature. Most of the algorithms in this area are concerned with guaranteeing a path toward the destination in the context of any network topology, while the optimality of the path is of little interest. We are presenting two novel geographic routing algorithms with obstacle avoidance properties. They aim at finding the optimal path from a source to a destination when some areas of the network are unavailable for routing due to low local density or obstacle presence. They locally and gradually with time evaluate and update the quality of the previously used paths and ignores non optimal paths for further routing. The first method is based on nodes' neighbourhood evaluation. The second is based on traffic analysis: on the identification of routing algorithms specific behaviours in the absence of relaying nodes. We are comparing its performance to existing state of the art protocols, showing that it performs better in terms of path length and hop count thus minimizing latency, overall traffic and energy consumption.

Reference

Luminita Moraru, Pierre Leone, Sotiris Nikoletseas, José D. P. Rolim.
Near optimal geographic routing with obstacle avoidance in wireless sensor networks by fast-converging trust-based algorithms.
In Proc. of the 3rd ACM Workshop on QoS and Security for Wireless and Mobile Networks (Q2SWiNet 07), pp. 31 - 38, 2007.

Interacting Urns Processes and applications in sensor networks

Pierre Leone
CUI, Switzerland

We analyze a distributed variation on the Pólya urn process in which a network of tiny artifacts manages the individual urns. Neighboring urns interact by repeatedly adding the same colored ball based on previous random choices. We discover that the process rapidly converges to a definitive random ratio between the colors in every urn and that the rate of convergence of the process at a given node depends on the global topology of the network. In particular, the same ratio appears for the case of complete communication graphs. Surprisingly, this effortless random process supports useful applications, such as clustering and pseudocoordinate computation. We present preliminary numerical studies that validate our theoretical predictions.

Reference

Pierre Leone, Elad Michael Schiller.
Interacting urns processes: for clustering of large-scale networks of tiny artifacts.
In Proc. of the 23rd Annual ACM Symposium on Applied Computing (SAC 08), pp. 2046-2051, 2008.

Pseudonyms in cost-sharing games

Paolo Penna
UNISA, Italy

This work initiates the study of cost-sharing mechanisms that, in addition to the usual incentive compatibility conditions, make it disadvantageous for the users to employ pseudonyms. We show that this is possible only if all serviced users pay the *same* price, which in some cases means that such a mechanism does not exist. In practice, a user can increase her utility by lying in one way (misreport her willingness to pay) or another (misreport her identity). We also consider mechanisms that rely in some kind of "reputation" associated to the pseudonyms.

Reference

Paolo Penna, Riccardo Silvestri and Peter Widmayer.
Pseudonyms in cost-sharing games.
Technical Report, 2008.

Better bounds for online load balancing on unrelated machines

Ioannis Caragiannis
UOP, Greece

We study the problem of scheduling permanent jobs on unrelated machines when the objective is to minimize the L_p norm of the machine loads. The problem is known as load balancing under the L_p norm. We present an improved upper bound for the greedy algorithm through simple analysis; this bound is also shown to be best possible within the class of deterministic online algorithms for the problem. We also address the question whether randomization helps online load balancing under L_p norms on unrelated machines; this is a challenging question which is open for more than a decade even for the L_2 norm. We provide a positive answer to this question by presenting the first randomized online algorithms which outperform deterministic ones under any (integral) L_p norm for $p = 2, ..., 137$. Our algorithms essentially compute in an online manner a fractional solution to the problem and use the fractional values to make random choices. The local optimization criterion used at each step is novel and rather counterintuitive: the values of the fractional variables for each job correspond to flows at an approximate Wardrop equilibrium for an appropriately defined non-atomic congestion game. As corollaries of our analysis and by exploiting the relation between the L_p norm and the makespan of machine loads, we obtain new competitive algorithms for online makespan minimization, making progress in another longstanding open problem.

Reference

Ioannis Caragiannis.
Better bounds for online load balancing on unrelated machines.
In Proc. of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 08), pp. 972-981, 2008.

On the performance of approximate equilibria in congestion games

Elias Koutsoupias
UOA, Greece

We study the performance of approximate Nash equilibria for linear congestion games. We consider how much the price of anarchy worsens and how much the price of stability improves as a function of the approximation factor ϵ. We give (almost) tight upper and lower bounds for both the price of anarchy and the price of stability for atomic and non-atomic congestion games. Our results not only encompass and generalize the existing results of exact equilibria to ϵ-Nash equilibria, but they also provide a unified approach which reveals the common threads of the atomic and non-atomic price of anarchy results. By expanding the spectrum, we also cast the existing results in a new light. For example, the Pigou network, which gives tight results for exact Nash equilibria of selfish routing, remains tight for the price of stability of ϵ-Nash equilibria but not for the price of anarchy.

Reference

On the performance of approximate equilibria in congestion games.
G. Christodoulou and E. Koutsoupias and P. Spirakis.
In arXiv:cs/0804.3160, April 2008.

Scheduling to maximize participation

Evi Papaioannou
UOP, Greece

We study a problem of scheduling client requests to servers. Each client has a particular latency requirement at each server and may choose either to be assigned to some server in order to get serviced provided that her latency requirement is met, or not to participate in the assignment at all. From a global perspective, in order to optimize the performance of such a system, one would aim to maximize the number of clients that participate in the assignment. However, clients may behave selfishly in the sense that, each of them simply aims to participate in an assignment and get serviced by some server where her latency requirement is met with no regard to the overall system performance. We model this selfish behavior as a strategic game, show how to compute pure Nash equilibria efficiently, and assess the impact of selfishness on system performance. We also show that the problem of optimizing performance is computationally hard to solve, even in a coordinated way, and present efficient approximation and online algorithms.

Reference

Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Evi Papaioannou.
Scheduling to maximize participation.
Theoretical Computer Science, Vol. 402(2-3), pp. 142-155, 2008.

On-going Work on Analysis and Simulation of a Reliable Architecture for Data Storage

Frédéric Giroire
CNRS & INRIA, France

We present our on-going work on analysis and simulations of the SPREADS system, which is a P2P-based reliable architecture for data storage.

Since peers are continuously failing, the SPREADS system must be self-repairing and resilient to faults. Instead of storing replicas of data among peers, the resilience may be achieved using erasure correcting codes which can be used to create a large scale redundant array of disks (RAID) with high reliability.

To maintain the redundancy level such a system needs to recompute the erasure codes and reinsert data into peers continuously. This leads to the following problems: given constraints on the disk usage efficiency and bandwidth usage, how to manage the system to ensure a desired data perenity and given a parametrized control algorithm, which parameters of the control algorithm ensure a given level of perenity.

We present a formal stochastic model for the SPREADS system and its control algorithm. Two methods are proposed to analyze the system behavior in its stationary phase. The comparison with simulations show that our second analytical model captures the system behavior with reasonable accuracy.

Reference

Olivier Dalle, Frederic Giroire, Julian Monteiro and Stephane Perennes.
On-going Work on Analysis and Simulation of a Reliable Architecture for Data Storage.
http://golgoth.inria.fr/wiki/Contrats/SPREADS

TSP: 27 Edges to PLS-Completeness

Dominic Dumrauf
UPB, Germany

In this paper, we investigate the complexity of computing *locally optimal* solutions for the Traveling Salesman Problem in the framework of PLS, as defined in Johnson et al. [1]. The neighborhood structure in our setting is superimposed by the well-known *k-OPT algorithm* (up to k edges may be exchanged in two neighboring solutions). We call this problem $TSP-(k)$ and show that $TSP-(k)$ *is PLS-complete for* $k \geq 27$. This significantly extends a result from Krentel [2], who shows that $TSP-(k)$ is PLS-complete for *some constant* $k \in N$ with $1,000 \leq k \leq 10,000$.

Similar to Krentel, [2], we use a reduction from Circuit/Flip to $(p,q,r,)-GSat$. The problem $(p,q,r,)-GSat$ is an extension of weighted Generalized Satisfiability to higher valued variables. Here, p is the maximum number of variables occuring in a weighted predicate, q is the maximum number of appearances of a variable, and r is the valuedness of the variables. Krentel, [2], uses that $(p,q,r,)-GSat$ is reducable to Satisfiability over binary variables and on to $TSP-(k)$, applying standard reductions at the expense of a large neighborhood in the TSP-problem. We instead perform a *direct reduction* from Circuit/Flip to $TSP-(k)$ and use the completeness of $(3,6,3,)-GSat$ as an improved starting point for our reduction. TSP-gadgets are defined that mimic the *weighted predicates* over *ternary variables* from the PLS-completeness proof of $(3,6,3,)-GSat$.

References

[1] David S. Johnson, Christos H. Papadimitriou, Mihalis Yannakakis. How Easy is Local Search? Journal of Computer and System Sciences, 37(1), pp. 79-100, 1988.

[2] Mark W. Krentel. Structure in Locally Optimal Solutions. In Proc. of the 30th Annual Symposium on Foundations of Computer Science (FOCS 89), pp. 216-221, 1989.

Reference
Dominic Dumrauf and Burkhard Monien.
TSP: 27 Edges to PLS-Completeness.
Technical Report, 2008.

Information Spreading in Dynamic Networks

Francesco Pasquale
UDRTV, Italy

Graphs that evolve over time are currently a very hot topic in computer science. The analysis of communication primitives in dynamic scenarios is a challenging task. In this talk we give an overview of two papers in which we studied the speed of information spreading in different models of dynamic networks.

In [1] we studied the completion time of distributed broadcast in dynamic radio networks. The dynamic network was modeled by means of adversaries. We considered two of them: A "deterministic worst-case adversary", that at any time step can make any network change, and an "oblivious, memoryless random adversary": At each time step t a graph G_t is selected according to the Erdos-Renyi $G_{n,p}$ model. In both cases we proved tight bounds on the completion time of the broadcast operation.

In [2] we introduced stochastic time-dependency in dynamic random graphs: Starting from an arbitrary initial graph, at every time step, every edge changes its state (existing or not) according to a two-state Markov chain. Such evolving graph model is a wide generalization of time-independent dynamic random graphs of [1]. We studied the speed of information dissemination in such dynamic graphs providing nearly tight bounds on the completion time of the flooding mechanism.

References

[1] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Communication in Dynamic Radio Networks. In Proc. of the 26th Annual ACM SIGAT-SIGOPS Symposium on Principles of Distributed Computing (PODC 07), ACM Press, pp. 205-214, 2007.

[2] A. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding Time in edge-Markovian Dynamic Graphs. In Proc. of the 27th Annual ACM SIGAT-SIGOPS Symposium on Principles of Distributed Computing (PODC 08), 2008, to appear.

Reference
Andrea Clementi, Claudio Macci, Angelo Monti, Francesco Pasquale and Riccardo Silvestri.
Information Spreading in Dynamic Networks.
Technical Report, 2008.

Adaptive Algorithms for Sensor Networks

Sotiris Nikoletseas
RA CTI, Greece

Motivated by emerging applications, we consider sensor networks where the sensors themselves (not just the sinks) are mobile. We focus on mobility scenarios characterized by heterogeneous, highly changing mobility roles in the network. To capture these high dynamics we propose a novel network parameter, the mobility level, which, although simple and local, quite accurately takes into account both the spatial and speed characteristics of motion. We then propose adaptive data dissemination protocols that use the mobility level estimation to improve performance. By basically exploiting high mobility (redundant message ferrying) as a cost-effective replacement of flooding, e.g., the sensors tend to dynamically propagate less data in the presence of high mobility, while nodes of high mobility are favored for moving data around. These dissemination schemes are enhanced by a distance-sensitive probabilistic message flooding inhibition mechanism that further reduces communication cost, especially for fast nodes of high mobility level, and as distance to data destination decreases. Our findings demonstrate significant performance gains of our protocols compared to non-adaptive protocols.

Reference

Athanassios Kinalis and Sotiris Nikoletseas.
Adaptive Algorithms for Sensor Networks.
In Proc. of the 11th ACM International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 08), 2008, to appear.

Overview of optimization of Wireless Mesh Networks

Herve Rivano
CNRS & INRIA, France

Wireless mesh network performance issues have been modeled by the Joint Routing and Scheduling Problem (JRSP) in which a maximum per-flow throughput is computed. A classical relaxation of JRSP, denoted as the Round Weighting Problem (RWP), consists in assigning enough weight to sets of compatible simultaneous transmissions (rounds), while minimizing the sum of them, thus maximizing the relative weight of each round, which model the throughput. In this talk, we present a new linear formulation of RWP focused on the transport capacity over the network cuts, thus eliminating the routing. We show its equivalence with existing formulations with flows and formalize a primal-dual algorithm that quickly solves this problem using a cross line and column generations process. An asset of this formulation is to point out a bounded region, a "bottleneck" of the network, that is enough to optimize in order to get the optimal RWP of the whole network. The size and location of this area is experimentally made through simulations, highlighting a few hop distant neighborhood of the mesh gateways. One would then apply approximated methods outside this zone to route the traffic without degrading the achieved capacity.

References

C. Molle, F. Peix, S. Pérennes and H. Rivano.
Formulation en Coupe/Rounds pour le Routage dans les réseaux radio maillés.
In Proc. of the 10èmes Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel 08), 2008.

C. Molle, F. Peix and H. Rivano.
Génération de colonnes pour le routage et lordonnancement dans les réseaux radio maillés.
In Proc. Colloque francophone sur lingénierie des protocoles (CFIP 08), 2008.

C. Molle, F. Peix, and H. Rivano.
An optimization framework for the joint routing and scheduling in wireless mesh networks.
In Proc. of the 19th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 08), 2008.

How Many Attackers Can Selfish Defenders Catch?

Vicky Papadopoulou
UCY, Cyprus

In a distributed system with *attacks* and *defenses,* an economic investment in defense mechanisms aims at increasing the degree of system protection against the attacks. We study such investments in the *selfish* setting, where both *attackers* and *defenders* are self-interested entities. In particular, we assume a *reward-sharing* scheme among *interdependent* defenders; each defender wishes to maximize its own *fair share* of the attackers caught due to him (and possibly due to the involvement of others).

Addressed in this work is the fundamental question of determining the *maximum* amount of protection achievable by a number of such defenders against a number of attackers if the system is in a *Nash equilibrium*. As a measure of system protection, we adapt the *Defense-Ratio* [1], which describes the expected proportion of attackers caught by defenders. In a *Defense-Optimal* Nash equilibrium, the Defense-Ratio is optimized.

We discover that the answer to this question depends in a quantitatively subtle way on the invested number of defenders. We identify graph-theoretic *thresholds* for the number of defenders that determine the possibility of optimizing a Defense-Ratio. In this vein, we obtain, through an extensive combinatorial analysis of Nash equilibria, a comprehensive collection of *trade-off* results.

References

[1] M. Mavronicolas, V. Papadopoulou, A. Philippou and P. Spirakis. A Network Game with Attacker and Protector Entities. In Proc. of the 16th Annual International Symposium on Algorithms and Computation (ISAAC 05), LNCS 3827, Springer, pp. 288297, 2005.

References
Marios Mavronicolas, Burkhard Monien and Vicky Papadopoulou.
How Many Attackers Can Selfish Defenders Catch?
In Proc. of the 41st Hawaii International International Conference on Systems Science (HICSS 08), pp. 470-478, 2008.

Interference Games

Luca Moscardelli
UNISA, Italy

We present a game-theoretic approach to the study of scheduling communications in wireless networks and introduce and study a class of games that we call *Interference Games*. In our setting, a player can successfully transmit if it "*shouts strongly enough*"; that is, if its transmission power is sufficiently higher than all other (simultaneous) transmissions plus the environmental noise. This physical phenomenon is commonly known as the Signal-to-Interference-plus-Noise-Ratio (SINR). Specifically, we show that the game has no *pure Nash equilibrium* and *mixed Nash equilibria* are either unfair (in the sense that one player has a much larger expected utility than the other players) or have social welfare 0. This state of things calls for the investigation of other notions of equiibrium. We show that there are *Sink* and *Correlated* equilibria with value of social welfare close to optimal. We then move to the study of Repeated Interference Games in which two players play possibly infinitely many instances of the Interference Game. We show that there exists a fair *Subgame Perfect Equilibrium* with maximal social welfare. Notice that such an equilibrium exists in the *partial monitoring*, that is, when each player only knows whether its previous transmission attempts have been successful. This result cannot be obtained with the well known Folk Theorem which requires *full monitoring*, that is each player has complete knowledge of the past actions of the other players.

Reference

Vincenzo Auletta, Luca Moscardelli, Paolo Penna and Giuseppe Persiano.
Interference Games.
Technical Report, 2008.

Gathering in radio networks

Jean-Claude Bermond
CNRS & INRIA, France
Invited talk

We study the problem of gathering information from the nodes of a multi-hop radio network into a predefined destination node under reachability and interference constraints. This problem has been asked by France Telecom in the context of how to bring Internet in villages; it appears also in collecting data in sensor networks. In such a network, a node is able to send messages to other nodes within reception distance, but doing so it might create interference with other communications. Thus, a message can only be properly received if the receiver is reachable from the sender and there is no interference from another message being simultaneously transmitted. The network is modeled as a graph, where the vertices represent the nodes of the network and the edges, the possible communications. The interference constraint is modeled by a fixed integer d, which implies that nodes within distance d in the graph from one sender cannot receive messages from another node. We mainly consider the case where each node has one (or more) unit-length message to transmit and, furthermore, we suppose that it takes one unit of time (slot) to transmit a unit-length message and during such a slot we can have only calls which do not interfere (called compatible calls). A set of compatible calls is referred to as a round. We give protocols and lower bounds on the minimum number of rounds for the gathering problem for various networks and under various hypothesis (like the possibility or not to store messages in transit nodes). We also consider the dynamic version of the problem where we want to give to each node some bandwidth.

References

J-C. Bermond, L. Gargano, A. Rescigno.
Gathering with Minimum Delay in Tree Sensor Networks.
In Proc. of the 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO 08), LNCS 5058, Springer, pp. 262-276, 2008.

J-C. Bermond, R. Corrêa, M-L. Yu.
Gathering Algorithms on Paths Under Interference Constraints.
In Proc. of the 6th Italian Conference on Algorithms and Complexity (CIAC 06), LNCS 3998, Springer, pp. 115-126, 2006.

J-C. Bermond, N. Morales, S. Pérennes, J. Galtier, R. Klasing.
Hardness and approximation of Gathering in static radio networks.
In Proc. of the 4th IEEE Conference on Pervasive Computing and Communications Workshops (PerCom 2006 Workshops), IEEE Computer Society, pp. 75-79, 2006.

J-C. Bermond, J. Galtier, R. Klasing, N. Morales, S. Pérennes.
Hardness and Approximation of Gathering in Static Radio Networks.
Parallel Processing Letters, 16(2), pp. 165-184, 2006.

NeOn: Lifecycle Support for Networked Ontologies

Mathieu d'Aquin
The Open University, United Kingdom
(Project NEON)
Invited talk

The growing availability of information has shifted the focus from closed, relatively data-poor applications, to mechanisms and applications for searching, integrating and making use of the vast amounts of information that are now available. Ontologies provide the semantic underpinning enabling intelligent access, integration, sharing and use of data and indeed this technology has now become so strategic to companies, that the Gartner market research firm has ranked taxonomies/ontologies third in their list of the top 10 technologies for 2005.

As ontologies are produced in larger numbers and exhibit greater complexity and scale, we now have an opportunity to build a new generation of complex systems, which can make the most of the unprecedented availability of both large volumes of data and large, reusable semantic resources. These systems will provide new functionalities in the emerging semantic web, in the automation of business to business relationships, and also in company intranets.

At the same time, we face a challenge: current methodologies and technologies, inherited from the days of closed, data-poor systems, are simply not adequate to support the whole application development lifecycle for this new class of semantic applications.

The aim of NeOn is to create the first ever service-oriented, open infrastructure, and associated methodology, to support the development life-cycle of such a new generation of semantic applications, with the overall goal of extending the state of the art with economically viable solutions. These applications will rely on a network of contextualized ontologies, exhibiting local but not necessarily global consistency.

Reference

http://www.neon-project.org/

ARRIVAL: Algorithms for Robust and online Railway optimization: Improving the Validity and reliAbility of Large scale systems

Spyros Kontogiannis
UOI & RA CTI, Greece
(Project ARRIVAL)
Invited talk

Algorithmic methods have reached a state of maturity as a consequence of decades of research where real-world problems were posed to the algorithms community triggering important developments in the field. Despite this success, the current state of algorithmic research still faces severe difficulties, or cannot cope at all, with highly-complex and data-intensive applications as those dealing with optimization issues in large-scale communication and transportation networks. The complexity and size of such optimization problems pose new challenges for algorithmic research, and their efficient solution requires a radically new foundational paradigm.

In this project, we are interested in establishing such a new paradigm and considerably advance the current state of algorithmic research by attacking optimization questions in perhaps the most complex and largest in scale (transportation) setting: that of railway systems. Railway optimization deals with planning and scheduling problems over several time horizons. We focus on two important and actually unexplored facets of planning that pose even harder optimization questions: robust planning and online (real-time) planning. These two, tightly coupled, facets constitute a proactive and a reactive approach, respectively, to deal with disruptions to the normal operation.

Our main goal is to develop the necessary foundational algorithmic research in order to provide ingenious and sound answers to the fundamental efficiency and quality issues encapsulated in robust and online planning of complex, large-scale systems as those in railways. We will endeavor to develop a thorough understanding of the fundamental issues that make robust and online railway optimization problems hard and to subsequently develop new algorithmic and complexity principles to deal with hardness through an integrated, interdisciplinary approach drawn from algorithmic, operations research, and constraint programming techniques.

Reference

http://arrival.cti.gr/