Physics-Inspired Methods for Networking and Communications

David Saad[1], Chi Ho Yeung[1], Georgios Rodolakis[2], Dimitris Syrivelis[2], Iordanis
Koutsopoulos[2], Leandros Tassiulas[3], Rüdiger Urbanke[4], Paolo Giaccone[5], Emilio Leonardi[5]

[1] Aston University, UK

[2] Centre for Research and Technology Hellas (CERTH), Greece

[3] Department of Electrical Engineering & Yale Institute of Network Science, Yale University

[4] Ecole Fédérale Polytechnique de Lausanne (EPFL), Switzerland

[5] Politecnico di Torino, Italy

E-mail: d.saad@aston.ac.uk chbyeung@gmail.com {rodolakis, jsyr}@iti.gr {jordan, leandros}@uth.gr
ruediger.urbanke@epfl.ch {paolo.giaccone, emilio.leonardi}@polito.it

Abstract

Advances in statistical physics relating to our understanding of large-scale complex systems have
recently been successfully applied in the context of communication networks. Statistical mechanics
methods can be used to decompose global system behavior into simple local interactions. Thus, central-
ized problems can be solved or approximated in a distributed manner with iterative lightweight local
messaging. This survey discusses how statistical physics methodology can provide efficient solutions to
hard network problems that are intractable by classical methods. We highlight three typical examples in
the realm of networking and communications. In each case we show how a fundamental idea of statistical
physics helps in solving the problem in an efficient manner. In particular, we discuss how to perform
multicast scheduling with message passing methods, how to improve coding using the crystallization
process and how to compute optimal routing by thinking of routes as interacting polymers.

I. Introduction

Large communication networks, such as the Internet are very complex distributed systems, composed
of millions of architectural elements, interacting according to complex and unpredictable patterns. Despite
the fact that the Internet is man-made and its internal element architectures and protocols are well known
and standardized, its overall behavior is still largely unclear. On this regard we quote Eric Schmidt,

Google-CEO: "The Internet is the first thing that humanity has built that humanity doesn't understand, the largest experiment in anarchy that we have ever had."

For these reasons, analyzing, optimizing and predicting the performance of large scale communication networks, raises several challenges. Traditional approaches, such as Markov chains, control theory, queuing theory, Monte-Carlo simulations, which have been largely employed in network design and protocol analysis, permit a very detailed microscopic description of the dynamics of some network elements, but fail to provide scalable tools for the comprehension of emerging dominant macroscopic dynamics.

Recently, techniques borrowed from statistical physics have been successfully applied to represent emerging macroscopic phenomena in computer networks: Mean field approaches have been employed to analyze the dynamics of congestion control algorithms (such as TCP) in high capacity networks where thousands of users share the limited bandwidths [1]. In the context of wireless networks, percolation theory results have been applied for capacity analysis of large-scale networks [2]. A parallelism with interacting Fermionic particle systems has been invoked to explain phase transition phenomena observable in large scale dense WiFi networks [3]. A new optimal-performing MAC protocol, inspired by Glauber dynamics over graphs, has been proposed as a viable alternative to the current WiFi scheme [4]. Message passage algorithms have been employed to solve specific packet scheduling problems in a scalable and distributed fashion [5]. Other methods adopted from statistical physics to address problems in communication and networking have been reviewed in [6].

Despite the rich body of work on employing statistical physical concepts in networks, a systematic statistical mechanics-inspired networking theory is still missing. This survey is a first step, based on the EU-FP7 project STAMINA, to fill this gap in part, by providing a general introduction to statistical physics methods that can be broadly employed for performance analysis of communication networks as well as network design and optimization. The main advantage offered by statistical physics methods is in the range of tools developed to address large scale problems of non-linear nature, through the study of *typical* case behavior and relying on approximation techniques that are well established within the physics community. The disadvantage of these methods is that many of them are non-rigorous and that some tend to fail in small-scale systems.

The methods suggested throughout the paper all stem from the statistical physics of disordered systems and have been adapted to the various problems depending on the suitability of a particular tool to the problem at hand. In particular, we present a description of a suitable statistical physics framework (Section II) and, via indicative application examples, the potential of three classes of techniques: multicast

scheduling and message passing in Section III, reliable transmission and crystallization in Section IV, routing and interacting polymers in Section V.

II. STATISTICAL PHYSICS AND COMMUNICATION PROBLEMS

So what can physics contribute to the quest for improved and more principled solutions for increasingly complex and difficult communication problems?

Among the main characteristics of statistical physics methods is that they deal with large-scale systems where interactions between system variables are non-linear and exhibit macroscopic emergent behavior, which is a characteristic of complex systems. Sophisticated techniques have been developed in this area to specifically address systems of this type and can be exploited to gain insight and develop optimization algorithms for real systems of similar characteristics.

The type of physical systems most suitable to represent optimization and constrained problems are termed disordered systems, where the disorder may refer to the underlying topology, i.e. the specific choice of edges in a graph, or strength of interaction between system constituents; these are assumed to be drawn from some distribution. Statistical mechanics methods for disordered systems in equilibrium revolve mostly around the concept of *free energy* that relates to the probability of the system to be in a given state. Calculating the averaged free energy over instances of the disorder which appears implicitly in the form of interaction between the various variables, is difficult and is the focus of the methods mentioned below.

Among the main methods adopted from the physics of disordered systems for the study of complex systems are the cavity and replica methods [7]. They facilitate the calculation of averages over all possible system instances in order to find the macroscopic system behavior, for instance over all networks of N vertices and given degree connectivity. While these methods focus on providing insight into the macroscopic behavior of the systems investigated, they usually also give rise to algorithms, such as belief propagation, that allow for the inference of variable values in specific instances.

III. MULTICAST SCHEDULING AND MESSAGE PASSING

A. Application: Scheduling

To keep up with the growing demand for communication resources in a cost-effective manner, optimal resource management becomes imperative, thus giving rise to hard optimization problems. One such example is scheduling, where a global objective, such as total throughput or energy-consumption, has to be

optimized while satisfying local constraints, such as conflicts between simultaneous wireless transmissions or simultaneous packet transmissions toward the same port interface in routers.

In particular, scheduling multicast traffic in Input Queue (IQ) switches requires solving a hard combinatorial optimization problem (the scheduling decision for each packet) in a very short time, comparable with the packet transmission time at the port interfaces. Consider the IQ switch illustrated in Figure 1(a). Based on the state of occupancy of the queues, the scheduler, must select a set of non-conflicting packets in order to maximize the total throughput. For multicast traffic, one *logical* queue is present for each possible input port and *fanout set* (i.e., a subset of output port destinations), thus making the scheduler's task highly complex.

This example constitutes an indicative test-case of the potential of the broader applicability of message passing techniques in communications. In this context, challenging optimization problems often require the design of algorithms that are both easy to implement, and efficient in terms of performance.

B. Physics: Cavity Method and Message-Passing

The optimization problem is mapped onto a statistical physics problem with a probability distribution over possible configurations, and the computation of the optimal solution is reduced to identifying the minimum-energy configurations at which the probability distribution concentrates. In order to find a configuration of low energy, one can use statistical physics techniques such as the cavity method. The latter is based on calculating the influence of neighboring nodes in the absence of the node in question, which leads to a set of coupled equations that can be solved iteratively, similar to density evolution, the macroscopic equivalent of belief propagation algorithms.

A convenient *graphical model* of the underlying system is a bipartite graph, with one set of nodes representing the system components (e.g, input ports of the IQ switch), while the other set consists of local interaction terms (such as conflicts in simultaneous packet transmissions). In Figure 1(b) we depict a toy example with 2 input ports and 3 interaction terms, modeling potential conflicts in queued packets due to intersecting multicast destination sets.

Each interaction node a sends a message $m(a \to i)$ to each neighbor node which contains a's belief about the state of node i. This belief is essentially a conditional probability estimated by a, based on all messages it receives from nodes other than i. Similarly, each node i computes its own belief and sends message $m(i \to a)$. These messages are computed separately by each node, they are iteratively updated and propagate through the system graph. This low-overhead iterative procedure usually leads to a low-energy configuration; in the case where the graph is a tree, it can be shown to converge to the

minimal-energy configuration. Thus, centralized problems can be solved in a distributed manner with iterative lightweight local messaging.

Recent success stories in discrete computational challenges attest to the strong potential of these techniques, including Shannon-capacity approaching codes and NP-hard problems like K-satisfiability [7], leading to more recent applications in wireless scheduling problems [5].

C. How Physics Helps in the Application

We present a specific application of the statistical physics methodology in the novel design and hardware implementation of a multicast switch.

The throughput-optimal scheduling policy for multicast traffic allows for "fanout-splitting", i.e., a packet can be sent to just a subset of its destination ports, leaving some residual destinations for future transmissions. The main idea of the optimal policy is to serve at higher priority packets that are stored in large queues and that are possibly re-enqueued into smaller ones. In order to solve the resulting problem, one can resort to a Belief Propagation (BP) algorithm [8].

Messages are exchanged between each input and output port, and are updated iteratively by each port concurrently. The scheduler chooses the packets to be transferred, based on the final *beliefs*, i.e., local estimates at each input port of the throughput that can be achieved by choosing specific transmission fanout sets. Due to the "densely connected" constraints that prevent conflicting packets, there are several cases in which the message update phase does not converge. Such a difficulty can be overcome using BP with a fixed number of iterations, in conjunction with a centralized algorithm, which at each step chooses the transmission fanout set with the maximum belief.

The BP approach outperforms other greedy algorithms (such as longest queue first) in simulations, with a gain between, 6% and up to 48% under uniform traffic, and between 5% and 10% under worst-case concentrated traffic [8]. Interestingly, only a very small number of message update iterations is necessary to achieve this performance (as low as 1 or 2).

We have implemented the belief propagation scheduling algorithm as a hardware accelerator on the netFPGA platform [9]. The implementation consists of a software communication interface and a hardware scheduler state machine for a 4×4 switch, integrated in the emulation framework illustrated in Figure 2. The measured duration of the scheduling algorithm execution is $3.77\mu s$, $5.44\mu s$, 7.12μ, for $0, 1$ and 2 BP iterations, respectively, i.e., a 2.65x improvement over the performance if implemented only in software. In terms of energy efficiency, the power consumption of the hardware switch during the scheduling algorithm execution is 26.1W, versus the platform's idle power consumption of 23.8W.

This successful implementation demonstrates the feasibility and potential of the message passing approach in practical networking problems, with stringent requirements on efficiency and lightweight operation.

IV. Reliable transmission, crystallization and the nucleation phenomenon

A. Application: Reliable Transmission of Information

The reliable transmission of information is at the heart of any communication system. Whether it is noise due to thermal effects or packet losses due to buffer overflows, there are many physical phenomena and processes that lead to a loss or degradation of transmitted information. Error-correcting coding is the means of ensuring that, despite all these adverse effects, the end-to-end communication link is reliable.

Traditionally, error-correcting codes were based on algebraic notions of ever increasing sophistication, and codes were designed to maximize the Hamming distance between codewords, i.e., the minimum number of positions in which two distinct codewords differ. But in the last 20 years, codes based on sparse graphs and message-passing schemes have fundamentally changed the way codes are designed and have gradually replaced traditional schemes.

Sparse graph codes are based on bipartite graphs as shown in Figure 3 where the length of the code is $N = 7$.

Each round node on the left represents a bit, and each square node on the right represents a constraint. In particular, these constraints represent linear equations that have to be fulfilled. The code is the set of all binary sequences of length N which fulfill all these constraints. To transmit information we pick a codeword, transmit it, and use the redundancy which is inherent in the code to recover the transmitted word from the received information.

The term *sparse* graph code indicates that the blocklength N is typically in the thousands, but the degrees of both variables and checks are taken from a finite set and do not depend on N. Hence the number of edges is of order N and not N^2 and is therefore "sparse". The important point about such sparse graph codes is that their decoding is accomplished by the belief propagation algorithm; i.e., given a noisy version of the transmitted codeword, messages representing the current "beliefs" about the various bits are exchanged along the edges until these messages (hopefully) converge to the correct values. Such a decoder is inherently of low complexity and can conveniently be implemented in hardware.

We use such codes daily, since they are part of state-of-the art cell phones, WiFi modems, optical transmission schemes and hard drives. Despite intensive research on sparse graph codes and their wide deployment, there is still room for improvement. In particular, it is difficult to design a coding scheme

which allows for reliable transmission close to capacity and which has very low error probabilities as required, e.g., for storage applications (hard disks) or in the backbone of the Internet (optical communications). This is due to the fact that in order to achieve transmission close to capacity, typically, a large number of variable nodes of small degree are necessary but also cause relatively high "error floors". Moreover, such codes are usually not universal, i.e., codes which are designed and reliable for one channel might not allow reliable transmission over another, even if it has equal capacity. Quite recently, an interesting physical phenomenon has been shown to be useful in overcoming these two difficulties.

B. Physics: Crystallization

Crystallization is the process which describes how solid crystals form from a liquid solution. This solution is typically a *meta-stable* state, which does not correspond to the lowest-energy configuration and which therefore eventually goes into the stable crystalline form. Nevertheless, the meta-stable state can persist on long time scales and a *nucleus*, i.e., a seed, is needed to get the crystallization process started.

Nucleation and the crystallization process can take on several forms, and we are all familiar with them in several disguises: (i) Reusable heat packs typically contain sodium acetate enclosed in a suitable container. When heated, sodium acetate takes on a liquid form which is a meta-stable state that is stable over long periods of time. The nucleation process takes on the form of bending a small metal disc which is contained in the heat pack. This starts the crystallization process and releases thereby heat. (ii) Water can be brought into a supercooled state in which it is still in liquid form even though it is considerably below the freezing point, as long as it is cooled in a clean container. The nucleation process starts by, e.g. shaking the container violently or by adding a little seed. A quick search on YouTube for "supercooled water" shows many instances how the crystallization process starts suddenly through a suitable nucleation process. Other examples are cloud or hail formation.

C. How Physics Helps in the Application

The nucleation phenomenon can be exploited to build codes that are provably capacity-achieving under message-passing decoders and universal for large classes of channels. Figure 4 shows an incarnation of the basic idea.

Rather than using an unstructured graphical model, we "spatially couple" a number of such graphical models along a chain in such a way that neighboring models interact, but that the local degree structure

of each model stays unchanged. In addition, we properly "terminate" the chain in such a way that the problem is made easier at the boundary.

Applying the basic belief-propagation algorithm on such a code, an interesting phenomenon occurs. The code can be successfully decoded up to a higher noise value than what is possible for the underlying "component" code. In fact, codes constructed in such a way can be decoded up to the maximum a-posteriori threshold of the underlying code. This is the highest threshold achievable under *any* decoding algorithm, no matter how complex. The decoding happens along a "wave." At the boundary, due to the special termination, decoding is easier and bits are decoded first. Once the bits at the boundary have been decoded, the "interface" between decoded and uncoded part moves inwards, and this decoding wave advances at a constant speed.

Mathematically and physically there is a direct analogy to nucleation and the crystallization process. The special termination at the boundary acts as the nucleus and gets the decoding started. Without the boundary the belief-propagation decoder, which is in general suboptimal, is not strong enough to drive the system into the lowest-energy state, which corresponds to the correct, i.e., transmitted codeword, but instead ends up in a meta-stable state. But once decoding has started at the boundary, the process continues like in the crystallization process, and the whole system moves at constant speed towards the lowest-energy configuration. Codes constructed based on this principle can be designed to work arbitrarily close to the capacity of the channel with high reliability and are inherently universal, i.e., they work well over whole classes of channels and do not have to be tuned to a particular application [10], [11].

This phenomenon can also be exploited in other areas. E.g., it has been used in compressive sensing [12] and to analyze constraint satisfaction problems [13].

V. ROUTING AND POLYMERS

A. Application: Routing

Finding optimal routes, given some measure of optimality, is a difficult task with implications on a large number of application domains, from water distribution networks and VLSI design to journey planners. Clearly, it has many applications in the area of communication networks, ranging from sensor and optical networks to peer-to-peer and wireless communication.

Optimal routing is also a very hard computational problem, being non-localized with non-linear interactions at vertices (routers) and/or edges (communication lines); therefore, most existing routing algorithms are based on localized selfish decisions and rely on (mostly non-adaptive) routing tables to identify the shortest weighted path to destination regardless of the individual decisions made, e.g., the

celebrated Dijkstra algorithm. Dynamic routing protocols do exist, but they are mostly heuristic and insensitive to other individual routing decisions which dynamically constitute the traffic.

To optimize the use of resource, a more global approach is required that takes into account all individual routing decisions and makes efficient use of the over-stretched network infrastructure. The cost to be minimized is defined according to the task at hand. For instance, in many cases one would like to suppress congestion in order to avoid bottle-necks by minimizing overlaps between routes, either at vertices or at edges, possibly attributing weights to vertices and/or edges to reflect preference, capacity or delays. At the other end of the spectrum, one may employ an objective function that aims to decrease the number of active vertices by consolidating paths to reduce infrastructure demands or energy consumption. This is particularly relevant in the context of communication networks, as the Internet can consume up to 4% of the electricity generated at peak times. These objective functions are typically non-linear and represent interactions between non-localized objects - multiple routes.

B. Physics: Polymers

Although the techniques developed for the study of disordered systems, such as the cavity method (message passing), have the potential to model interactions between simple system constituents such as network nodes, they may be difficult to apply when interactions involve more complicated objects such as routes. In the case of routing, additional techniques should be employed in order to verify that routes are contiguous, leading from sources to destinations. Techniques developed in the study of polymers are ideally suited for these tasks.

A polymer can be viewed as a chain of molecules connected to one another in a manner where the end of one segment is the beginning of another. Routes on a network can be viewed as polymers placed on a graph such that segments correspond to edges and interact at vertices, as shown in Figure 5. The aim is to choose the shortest routes (polymers) while incorporating an interaction between them on vertices or edges: *repulsion*, in case one wants to minimize congestion and make traffic as uniform as possible; and *attraction*, when one wants to consolidate routes and reduce the number of active vertices and/or edges.

C. How Physics Helps in the Application

We studied two routing scenarios using statistical physics methods:

1) The case in which randomly selected nodes on a given graph communicate with specific routers while minimizing congestion or maximizing route consolidation [14]. This study revealed the

macroscopic behavior in terms of cost and average path length as the number of communication sources increases and provided a distributed algorithm to find optimal solutions in specific instances.

2) A scenario whereby we analyze macroscopic properties of generic path-optimization problems between arbitrarily selected communicating pairs; we also derive a simple, principled, generic and distributive routing algorithm, capable of considering simultaneously all individual path choices.

While the method is suitable for many networking and communication problems, we choose a more graphical and real-life problem to demonstrate the efficacy of the new algorithm - routing passengers on the London underground network based on real Oyster-card data. Figure 6(a) shows how congestion is reduced by the algorithm when the objective function chosen aims to repel routes, as reflected in the fairly uniform traffic distribution even in the central region; the cost obtained by our algorithm is 20.5% smaller than that of the shortest path configuration obtained by the Dijkstra algorithm and slightly better than other state-of-the-art algorithms [15], with only a slight increase in average path length by 5.8%. In contrast, Figure 6(b) shows how paths for the same passenger set are consolidated at major routes and stations, when a cost aimed at consolidating routes is chosen. This scenario may be relevant at times when the service is reduced for some reason, for instance a strike or at off-peak hours to decrease costs. Due to the concave nature of the cost, there are no efficient competitive algorithms for carrying out these tasks.

These methods have direct relevance to various communication problems, such as node-disjoint routing, a hard-computational problem that is essential to prevent blocking in optical networks [16]. The insight gained by employing methods of statistical mechanics to routing problems and the efficient algorithms derived are of great potential and will help to provide more efficient and scalable individualized routing algorithms.

The main challenges one needs to address for making these routing techniques more applicable are: a) to accommodate temporal interaction between routes, i.e., taking into account the time it takes to arrive at nodes/edges and limiting the interaction between routes to concurrent traffic; b) devising approximation techniques in the case of node/edge-disjoint multi-colored routing in the presence of a large number of different colors, such as in the case of realistic multi-wavelength optical networks.

VI. Conclusion and Future Perspectives

Physics-inspired methods can lead to breakthroughs when applied to communications and networking tasks. In this paper we listed some instances of such problems. The presentation spans a range of examples,

from aspects of theoretical modeling using physics-inspired methods up to how such methods find their way into systems that are amenable to implementation.

The range of possible applications of physics-inspired methods is in no way limited to the three applications we discussed above. In concluding, we discuss three challenging research directions.

Non-equilibrium methods: Most existing physics-inspired methods have been developed for understanding homogeneous equilibrium systems, while most real communication systems are heterogeneous, hierarchical and inherently *not* in equilibrium. Adapting non-equilibrium methods from statistical physics to manage and optimize dynamical and hierarchical communication systems is a significant challenge. This new approach may also reveal hidden network properties and shed light on system design choices.

Big Data: The amount of data generated in recent years is unprecedented and, more often than not, the rate of generation exceeds the rate at which they can be absorbed and understood, by humans or machines. Often data streams should be processed online and low-complexity schemes are needed, to create meaningful information out of the deluge of data. Being equilibrium-based methods, current message passing methods are limited in their ability to deal with data streams and novel dynamical inference methods should be devised to address this challenge.

Capacity crunch of communication networks: The ever-increasing demand for information, both due to the increasing number of mobile networks and the powerful and information-hungry networked devices, bring communication networks close to capacity. Smartphones, tablets and watches with multiple embedded sensors capture data from their surroundings and transmit it; incoming mobile video traffic is also increasing with exponential rate. Clearly, this introduces unprecedented requirements to communication network operators and challenges the provisioning of end-user Quality of Experience. Novel schemes are needed, in multi-dimensional resource configurations, to optimize content transmission to the end-user. For example, intelligent caching of content at various locations or even at neighboring devices is needed in conjunction with joint optimization across resource domains like energy, computational capacity and control, the latter being reflected in the choice of transmission rate and power, access point association etc. Physics-inspired techniques could be crucial in optimizing performance with multiple objectives and constraints.

REFERENCES

[1] F. Baccelli, A. Chaintreau, D. De Vleeschauwer, and D. McDonald. A mean-field analysis of short lived interacting TCP flows. In *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '04/Performance '04, pages 343–354, New York, NY, USA, 2004. ACM.

[2] M. Franceschetti, O. Dousse, D.N.C. Tse, and P. Thiran. Closing the gap in the capacity of wireless networks via percolation theory. *Information Theory, IEEE Transactions on*, 53(3):1009–1018, March 2007.

[3] M. Durvy, O. Dousse, and P. Thiran. On the fairness of large CSMA networks. *Selected Areas in Communications, IEEE Journal on*, 27(7):1093–1104, September 2009.

[4] L. Jiang, M. Leconte, Jian Ni, R. Srikant, and J. Walrand. Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling. In *INFOCOM, 2011 Proceedings IEEE*, pages 371–375, April 2011.

[5] S. Sanghavi, D. Shah, and A.S. Willsky. Message passing for maximum weight independent set. *Information Theory, IEEE Transactions on*, 55(11):4822–4834, 2009.

[6] C. H. Yeung and D. Saad. Networking - a statistical physics perspective. *Journal of physics A: mathematical and theoretical*, 46(10):103001, 2013.

[7] M. Mezard and A. Montanari. *Information, physics, and computation*. Oxford University Press, USA, 2009.

[8] P. Giaccone and M. Pretti. A belief-propagation approach for multicast scheduling in input-queued switches. In *Workshop on networking accross disciplines: Communication networks, complex systems and statistical physics (NETSTAT)*, pages 1403–1408, 2013.

[9] D. Syrivelis, P. Giaccone, I. Koutsopoulos, M. Pretti, and L. Tassiulas. On emulating hardware/software co-designed control algorithms for packet switches. In *Emutools*, 2014.

[10] A. J. Felström and K. S. Zigangirov. Time-varying periodic convolutional codes with low-density parity-check matrix. 45(5):2181–2190, September 1999.

[11] S. Kudekar, T. Richardson, and R.L. Urbanke. Spatially coupled ensembles universally achieve capacity under belief propagation. *Information Theory, IEEE Transactions on*, 59(12):7761–7813, Dec 2013.

[12] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun, and L. Zdeborová. Statistical-physics-based reconstruction in compressed sensing. *Phys. Rev. X*, 2:021005, May 2012.

[13] S. Hamed Hassani, N. Macris, and R. Urbanke. Thresholds saturation in spatially coupled constraint satisfaction problems. E-print: http://arxiv.org/abs/1112.6320, December 2011.

[14] C. H. Yeung and D. Saad. Competition for shortest paths on sparse graphs. *Phys. Rev. Lett.*, 2012.

[15] C. H. Yeung, D. Saad, and K. Y. M. Wong. From the physics of interacting polymers to optimizing routes on the London underground. *Proc Natl. Academy of Sci.*, 110:13717–13722, 2013.

[16] C. De. Bacco, S. Franz, D. Saad, and C. H. Yeung. Shortest node-disjoint paths on random graphs. *Jour. Stat. Mech.*, in press, 2014.

David Saad is Professor of Information Mathematics in Aston University, UK. He received a BA in Physics and BSc in Electrical Engineering from Technion, an MSc in Physics and PhD in Electrical Engineering from Tel-Aviv University. He joined Edinburgh

University in 1992 and Aston in 1995. His research focuses on the application of statistical physics methods to several fields, including neural networks, error-correcting codes, multi-node communication, network optimization, routing, noisy computation and advanced inference methods.

Chi Ho Yeung received his B. Sc., M. Phil. and Ph. D. degrees in Physics from the Hong Kong University of Science and Technology in 2004, 2006 and 2009 respectively. He is currently a lecturer at the Hong Kong Institute of Education. His major research interests include statistical physics, disordered systems, optimisation, routing, recommendation algorithms, complex and social networks.

Georgios Rodolakis graduated in Electrical and Computer Engineering from Aristotle University, Greece, and obtained his D.E.A. and Ph.D. from Ecole Polytechnique and INRIA, France, in 2006. After a research fellowship at Macquarie University, Sydney, Australia, he joined CERTH, Greece, in 2011. His main research interests are in the areas of mobile networks, information theory, design and analysis of algorithms.

Dimitris Syrivelis received his bachelor from Technical University of Crete in 2003 and Phd from University of Thessaly in 2009 on h/w-s/w codesigned systems for massively parallel architectures. He has developed several networking systems for wired and wireless applications, most notably for wireless network coding, information-centric networking and software-defined networking. He is a regular contributor to open source networking projects (Click Modular Router, Linux Kernel) and he currently has 18 publications in well-known conferences.

Iordanis Koutsopoulos is Assistant Professor at the Department of Computer Science of Athens University of Economics and Business (AUEB). He obtained the M.S and Ph.D. degrees in Electrical and Computer Engineering from the University of Maryland, College Park (UMCP) in 1999 and 2002 respectively. His research interests are in the broad area of network control and optimization, with applications to wireless networks, social networks, smart grid control, sensor networks and cloud computing systems.

Leandros Tassiulas is Professor in the Department of Electrical Engineering and the Yale Institute of Network Science at Yale University. His research interests are in the field of computer and communication networks with emphasis on fundamental mathematical models and algorithms of complex networks, architectures and protocols of wireless systems, sensor networks, novel internet architectures and experimental platforms for network research.

Ruediger L. Urbanke (EPFL) is principally interested in the analysis and design of iterative coding schemes, which allow reliable transmission close to theoretical limits at low complexities. He is a co-author of the book "Modern Coding Theory" (Cambridge University Press), a co-recipient of the 2002 and the 2013 IEEE Information Theory Society Paper Award, the 2011 IEEE Koji Kobayashi Award, as well as the 2014 IEEE Hamming Medal.

Emilio Leonardi is an Associate Professor at Politecnico di Torino. He has co-authored over 150 papers in the area of telecommunication networks. He participated in several national and European projects He has also been involved in several consulting and research projects with private industries. He has been the coordinator of the European 7-th FP-STREP "NAPA-WINE" on P2P streaming applications. His research interests are in the field of: social networks, network science, caching systems.

Paolo Giaccone received the Dr.Ing. and Ph.D. degrees in telecommunications engineering from the Politecnico di Torino, Italy, in 1998 and 2001, respectively. He is currently an Assistant Professor in the Department of Electronics and Telecommunications at Politecnico di Torino. His main area of interest is the design of network algorithms, the theory of interconnection networks, and the performance evaluation of telecommunication networks through simulative and analytical methods.
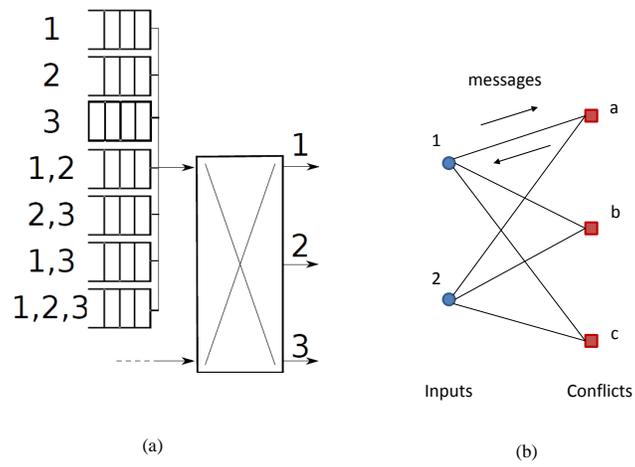
Fig. 1. (a) An IQ switch with 2 inputs and 3 outputs. Each logical queue is tagged with the corresponding fanout set, and (b) a toy example of a bipartite graphical model of the switch with blue nodes corresponding to inputs and red nodes to interaction terms (conflicting multicast packet transmissions).
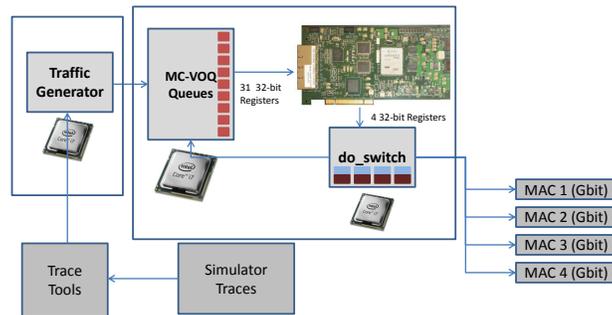
Fig. 2.   A hardware/software co-design schematic of BP scheduling implemented on NetFPGA.
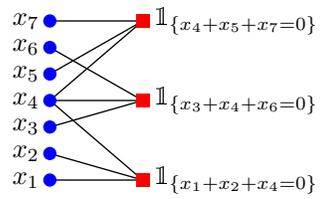
Fig. 3. Bipartite graph representing a sparse graph code of length $N = 7$. The code consists of the set of binary words of length 7 that fulfill the three linear constraints.
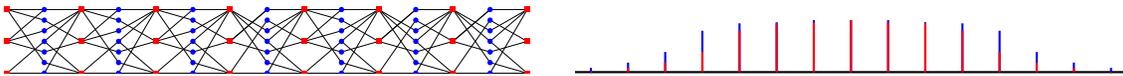
Fig. 4. Left: A spatially coupled code. Several "component" codes are placed along a line are neighboring components are "coupled" i.e., they share edges. At the boundary there are locally more check nodes than in the interior. This makes decoding easier and acts as a seed. Right: The "decoding wave." The horizontal axis corresponds to the spatial dimension of the code. The vertical dimension shows that local error probability at a certain point in the decoding process. As the number of iterations increases, the error probability decreases taking the shape of a "wave" that moves at constant speed from the boundary towards the middle.
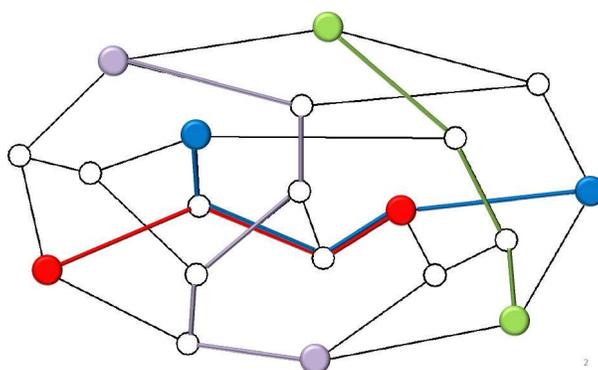
Fig. 5. Routes as polymers on a regular random network of degree 3. Full circles represent start and end points of communication paths, each one presented in a different color, while intermediate empty nodes may constitute part of the path(s). Interaction between polymers may take place on vertices and/or on edges.
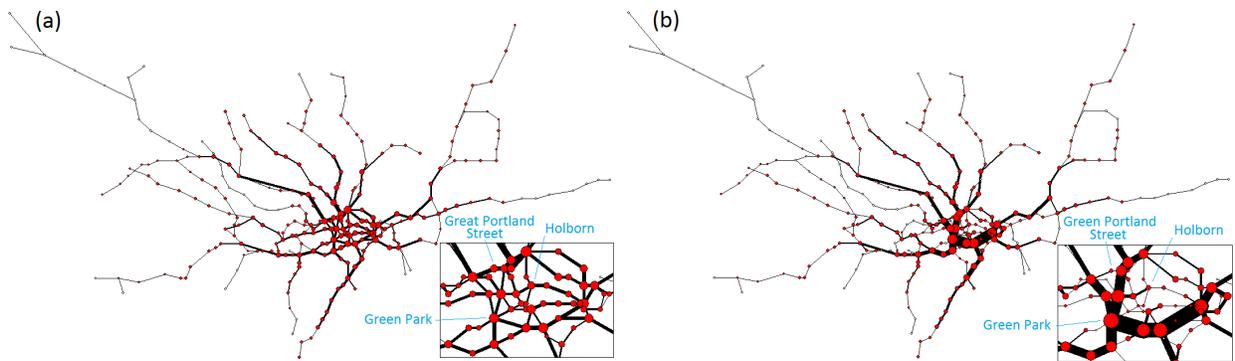
Fig. 6. Optimized traffic on the London subway network. A total of 218 real passenger source-destination pairs are optimized, corresponding to 5% of data recorded by the Oyster card system between 8:30am - 8:31am on one Wednesday in November 2009. The network consists of 275 stations. The corresponding costs per node are chosen to (a) repel and (b) consolidate routes. Red nodes correspond to stations. The size of each node and the thickness of each edge are proportional to traffic through them.