

MATHEMATISCHES FORSCHUNGSINSTITUT OBERWOLFACH

Report No. 48/2010

DOI: 10.4171/OWR/2010/48

Mathematical Challenges in Stochastic Networks

Organised by
Serguei Foss, Edinburgh
Günter Last, Karlsruhe
Michel Mandjes, Amsterdam
Balaji Prabhakar, Stanford

October 17th – October 23rd, 2010

ABSTRACT. The workshop was devoted to the discussion of recent progress in modern stochastic network theory and to the exploration of open mathematical challenging problems in the field. The workshop covered a wide range of mathematical topics; while being centered around applied probability, it also included a substantial amount of graph theory and (combinatorial) optimization.

Mathematics Subject Classification (2000): 60xx.

Introduction by the Organisers

Stochastic Networks is a flourishing area in the very heart of modern Applied Probability. Its primary aim is to obtain a fundamental understanding of the mathematical properties of complex interacting random systems. Among the main goals are the design, analysis and evaluation of important benchmark systems, as well as the development of efficient tools for the optimization and simulation of networks.

Among the 52 participants there were many leading experts in the field as well as many very promising young scientists. For about a half of the participants the workshop was their first opportunity to visit Oberwolfach.

The programme included three 2-hour lectures, nineteen 40-minute talks, and two 90-minute open problems sessions. The lectures formed the organizational backbone of the workshop. They were delivered by Bruce Hajek on "Peer to peer communication in networks - issues, models, and analysis", by Alexandre Proutiere on "Short and long-term behaviors of Markov processes through mean

field asymptotics” and by Volker Schmidt on ”Distribution of cost functionals in spatial network models: Scaling limits and Monte- Carlo methods”. All the lectures have received a very warm welcome from the audience.

Bruce Hajek gave an impressive overview on the modelling and analysis of peer to peer networks and, in particular, on problems related to the so-called ”single-chunk syndrome” problem. These models are very new, they are influenced by the practice, and their analysis is very challenging and practically important. Alexandre Proutiere was speaking about modern problems in understanding and analysis of the dynamics of high dimensional Markov processes using the mean-field approximation. He also discussed two practical examples, of the epidemic diffusion of viruses and of large wireless networks with random access. Volker Schmidt lectured on a stochastic geometry approach to the modeling and performance analysis of spatial stochastic networks. The focus was on telecommunication networks involving road systems.

The contributed talks were typically linked to one or more issues covered by these lectures. In more detail, the topics discussed during the workshop were related to the following areas. One main theme was the probabilistic analysis of stochastic processes (uni- and multivariate Levy processes, random walks, Markov additive processes, etc.) relevant to stochastic networks (with talks by V. Anantharam, S. Asmussen, E. Baurdoux, O. Kella, L. Leskela, V. Wachtel, R. Szekli). Another direction was around various scaling techniques and related asymptotic analysis of stochastic networks. In addition to the lecture by A. Proutiere, there were several other talks (T. Dieker, D. Gamarnik, R. Johari, D. McDonald, A. Stolyar) on scaling in general Markov processes, parallel server systems, stochastic games, large deviations, and back-pressure algorithms. Another topic was the optimization, complexity, and scheduling problems in communication and spatial networks, and related asymptotic analysis (B. Hajek, V. Schmidt, I. Norros, D. Shah, P. Thiran). There were also several presentations (M. Lelarge, G. Hooghiemstra, J. Salez) on routing, matching and related problems in random graphs. Finally, C. Wichelhaus gave a talk on non-parametric inference for general stochastic networks.

An integral part of the programme was provided by two problem sessions organized by S. Foss and M. Mandjes. These sessions provided space for participants to bring up new ideas and discuss open problems in an informal manner. There were 8 presentations which covered a broad range of problems and, in particular, on classical stability analysis of Markov processes (P. Glynn), on greedy server and vacuum cleaner models (T. Konstantopoulos), and on sequential algorithms for solving linear equations (D. Wischik). Finally, Balaji Prabhakar gave a very informal presentation on how to use incentives to make ”Societal Networks” more efficient. For example, how to decongest the roads? How to increase the use of public transit? How to reduce energy waste? He described the results of actual deployments and illustrated the role played by computer communication networks and incentive algorithms in these deployments.

Workshop: Mathematical Challenges in Stochastic Networks**Table of Contents**

Venkat Anantharam (joint with Vivek Borkar)	
<i>Stochastic Approximation with Long-Range-Dependent and Heavy-Tailed Noise</i>	2749
Søren Asmussen (joint with Asger Hobolth)	
<i>Markov bridges, bisection and variance reduction</i>	2749
Erik J. Baurdoux	
<i>Scale functions for spectrally negative Lévy processes and some applications</i>	2751
Ton Dieker	
<i>Spectral gaps for large-scale Markov processes</i>	2754
David Gamarnik (joint with David Goldberg)	
<i>Steady state GI/GI/n queue in the Halfin-Whitt regime</i>	2756
Peter W. Glynn	
<i>Wide-sense Regeneration for Harris Recurrent Markov Processes: An Open Problem</i>	2758
Bruce Hajek (joint with Laurent Massoulié, Sujay Sanghavi, Kuang Xu, Ji Zhu)	
<i>Peer to Peer Communication in Networks – Issues, Models, and Analysis</i>	2760
Gerard Hooghiemstra (joint with Shankar Bhamidi, Remco van der Hofstad)	
<i>Shortest weight routing in complex networks</i>	2764
Ramesh Johari	
<i>Equilibrium analysis of large scale stochastic games</i>	2767
Offer Kella (joint with Onno Boxma, Marc Yor, Andreas Löpker and David Perry)	
<i>Some recent observations regarding growth collapse processes and their generalizations</i>	2768
Marc Lelarge (joint with Hamed Amini, Moez Draief)	
<i>Diameter in weighted random graphs</i>	2770
Lasse Leskelä	
<i>Markov couplings, stochastic orders, and stochastic relations</i>	2773
David McDonald (joint with Ivo Adan, Bob Foley)	
<i>Nonlinear large deviations: Exact Asymptotics</i>	2775

Ilkka Norros (joint with Hannu Reittu)	
<i>On the stability of population processes</i>	2782
Alexandre Proutiere (joint with Charles Bordenave, David McDonald)	
<i>Short and Long-term behavior of Stochastic Systems through Mean Field Asymptotics</i>	2783
Justin Salez (joint with Charles Bordenave, Marc Lelarge)	
<i>Counting and sampling matchings on infinite graphs</i>	2785
Volker Schmidt (joint with Frank Fleischer, Catherine Gloaguen and Florian Voss)	
<i>Distribution of cost functionals in spatial network models: Scaling limits and Monte-Carlo methods</i>	2788
Devavrat Shah (joint with David Tse, John Tsitsiklis, Damon Wischik, Yuan Zhong)	
<i>Three metrics for stochastic networks: capacity, queue-size and complexity</i>	2789
Alexander Stolyar	
<i>Scaling properties of queues under back-pressure algorithm</i>	2795
Ryszard Szekli (joint with Hans Daduna, Pawel Lorek)	
<i>Remarks on convergence to stationarity for unreliable queueing networks</i>	2797
Patrick Thiran (joint with Florence Bénézit, Vincent Blondel, Alexandros G. Dimakis, John Tsitsiklis, Martin Vetterli)	
<i>Gossip Algorithms Using (Weighted) Path Averaging</i>	2799
Vitali Wachtel (joint with Denis Denisov)	
<i>Ordered random walks</i>	2802
Neil Walton (joint with Peter Key)	
<i>Auction Networks a Stochastic Network Enviroment</i>	2804
Cornelia Wichelhaus (joint with Roland Langrock and Susan Pitts)	
<i>Nonparametric inference for networks of queues</i>	2806

Abstracts

Stochastic Approximation with Long-Range-Dependent and Heavy-Tailed Noise

VENKAT ANANTHARAM

(joint work with Vivek Borkar)

We analyze stability and convergence properties of stochastic approximation algorithms when the noise includes a long-range-dependent component, modeled by a fractional Brownian motion, and a heavy-tailed component, modeled by a symmetric stable process, in addition to the usual martingale noise. The proofs are based on comparing suitably interpolated iterates with a limiting ordinary differential equation. Key prior work providing the tail estimates that replace the traditional tail estimates for martingale noise in the proof scheme are Fernique's inequality, see the book of Berman [1] and Joulin's inequality, see [3]. The proof scheme is a variation of that found in the book of Borkar [2].

REFERENCES

- [1] S. M. Berman, *Sojourns and Extremes of Stochastic Processes*, Wadsworth and Brooks-Cole, Belmont, CA, (1992).
- [2] V. S. Borkar, *Stochastic Approximation: A Dynamical Systems Viewpoint*, Hindustan Publishing Agency, New Delhi, and Cambridge University Press, Cambridge, UK, (2008).
- [3] A. Joulin, *On maximal inequalities for stable stochastic integrals*, *Potential Analysis* **26** (2007), 57–78.

Markov bridges, bisection and variance reduction

SØREN ASMUSSEN

(joint work with Asger Hobolth)

Let X be a continuous-time Markov process with state space E . A Markov bridge (T, a, b) is then defined as the distribution of X in $[0, T]$ conditioned on $X(0) = a$, $X(T) = b$. Such bridges are important in many areas, in particular statistical inference for Markov processes and a number of biological applications like coalescence problems, and there is a considerable literature on simulation algorithms, see the references. For example, in statistical inference X may only be observed at discrete time points $0, h, 2h, \dots, Nh$ and a Markov bridge with $T = h$, $a = X((k-1)h)$, $b = X(kh)$ may then be useful to estimate sufficient statistics like the total time spent in a given state, which is a key step in the stochastic EM algorithm. In coalescence, a simplified problem is to obtain information on the mutations at a single site of a DNA string from a common ancestor to two species, say mouse and human. Modeling mutations as a reversible Markov process, the evolution from mouse via ancestor to human is then a Markov process where the endpoints are observable, i.e. a Markov bridge.

We consider here simulation algorithms for the case of a finite state space. Some previous algorithms are reviewed and we present a new one using bisection. Preliminary examples are presented that the algorithm has good potential for variance reduction.

For processes such as diffusions or Lévy processes ($E = R$ or $E = R^d$), it is not obvious how to simulate a Markov bridge and existing algorithms tend to deal with special cases and be difficult to implement. The situation of a finite E is typically much easier, since simple rejection sampling applies: just simulate forward from $X(0) = a$ using the description in terms of exponential holding times and Markovian changes of states, and accept the first sample path with $X(T) = b$. However, if the event $X(T) = b$ is rare, rejection sampling is inefficient, and therefore alternatives have been suggested (see [15] for a survey). We consider here yet another, bisection based on subsequently dividing the interval into two. The main idea is to classify intervals as having 0, 1 or at least 2 jumps, and noting that the sample path in intervals with 0 or 1 jumps can easily be generated. The algorithm then uses a recursion to proceed with intervals with at least 2 jumps.

Bisection is sometimes faster than existing algorithms and sometimes not. However, our main argument in its favor is the potential for variance reduction. As for Brownian motion (e.g. [10], [11]), one then identifies the ‘most important variables’ with $X(T), X(T/2), (X(T/4), X(3T/4)), \dots$ (in that order), selects a finite number of them and applies variance reduction techniques like stratification, importance sampling or quasi-Monte Carlo to the selected variable. Examples from ongoing experimentation were presented in the talk and show excellent potential for variance reduction.

The final version of this research is expected to appear in [2].

REFERENCES

- [1] S. Asmussen and A. Hobolth, Bisection ideas in end-point conditioned Markov process simulation. In *Proceedings of the 7th International Workshop on Rare Event Simulation* (G. Rubino & B. Tuffin, eds.) (2008), 121–130. Rennes, France.
 - [2] S. Asmussen and A. Hobolth, Markov bridges, bisection and variance reduction. *Proceedings of MCQMC2010*, Warszawa, Springer-Verlag (2011).
 - [3] S. Asmussen and P.W. Glynn, *Stochastic Simulation. Algorithms and Analysis*, Springer-Verlag (2007).
 - [4] A.N. Avramidis, P. L’Ecuyer and P.-A. Tremblay Efficient simulation of gamma and variance-gamma processes. In *Proceedings of the 2003 Winter Simulation Conference* (S. Chick et al., eds.) (2003).
 - [5] A. Beskos, O. Papaspiliopoulos and G.O. Roberts, *Retrospective exact simulation of diffusion sample paths with applications*, Bernoulli **12** (2006), 1077–1098.
 - [6] A. Beskos, O. Papaspiliopoulos and G.O. Roberts, A new factorization of diffusion measure with view towards simulation, To appear (2007).
 - [7] A. Beskos, G.O. Roberts, A. Stuart and J. Voss, *MCMC methods for diffusion bridges*, Stoch. Dyn. **8** (2008), 319–350.
 - [8] M. Bladt, A. Gonzales and S.L. Lauritzen, *The estimation of phase-type related functionals using Markov chain Monte Carlo*, Scand. Act. J. **4** (2003), 280–300.
 - [9] M. Bladt and M. Sørensen, Simple simulation of diffusion bridges with application to likelihood inference for diffusions (2009).
- <http://www.math.ku.dk/~michael/diffusionbridge0809.pdf>

- [10] R.E. Caflish and B. Moskowitz, *Modified Monte Carlo methods using quasi-random sequences*, in Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing (H. Niederreiter & P.J.-S. Shine, eds.), Lecture Notes in Statistics **106** (1995), 1–16, Springer-Verlag.
- [11] R.E. Caflish, W. Morokoff and A.B. Owen, *Valuation of mortgage-backed securities using Brownian bridges to reduce effective dimension*, J. Comput. Finance **1** (1997), 27–46.
- [12] P. Fearnhead and C. Sherlock, *An exact Gibbs sampler for the Markov-modulated Poisson process*, J.R. Statist. Soc. B, **68** (2006), 767–784.
- [13] P. Fitzsimmons, J. Pitman and M. Yor, *Markovian bridges: construction, Palm interpretation and splicing*, In Seminar on Stochastic Processes. Progress in Probability **32** (1992), 101–134.
- [14] A. Hobolth, *A Markov chain Monte Carlo expectation maximization algorithm for statistical analysis of DNA sequence evolution with neighbour-dependent substitution rates*, Journal of Computational and Graphical Statistics **17** (2008), 138–164.
- [15] A. Hobolth and E.A. Stone, *Efficient simulation from finite-state, continuous-time Markov chains with incomplete observations*, Ann. Appl. Statist. (2009).
- [16] G. Leobacher, *Stratified sampling and quasi-Monte Carlo simulation of Lévy processes*, Monte Carlo Methods and Applications **12** (2006), 231–238.
- [17] C. Moler and C. Van Loan, *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later* SIAM Review **45** (2003), 3–49.
- [18] B. Moskowitz and R.E. Caflish, *Smoothness and dimension reduction in quasi-Monte Carlo methods*. Journal of Mathematical and Computer Modeling **23** (1996), 37–54.
- [19] R. Nielsen, *Mapping mutations on phylogenies*, Syst. Biol., **51** (2002), 729–739.
- [20] C. Ribeiro and N. Webber (2003) *Valuing path-dependent options in the variance-gamma model by Monte Carlo with a gamma bridge*, J. Comput. Finance **7**, 81–100.
- [21] C. Ribeiro and N. Webber, *Correction for simulation bias in Monte Carlo methods to value exotic options in models driven by Lévy processes*, Appl. Math. Finance, **13** (2006), 333–352.
- [22] G.O. Roberts and O. Stramer, *On inference for partially observed nonlinear diffusion processes using Metropolis-Hastings algorithms*, Biometrika **88** (2001), 603–621.

Scale functions for spectrally negative Lévy processes and some applications

ERIK J. BAURDOUX

Spectrally negative Lévy processes (i.e. those without positive jumps and which are not the negative of a subordinator) appear naturally in queueing theory. Indeed, the workload of an M/G/1 queue is given in terms of a compound Poisson process with negative drift, reflected at its supremum. In many queueing applications the compound Poisson process with negative drift is replaced by a spectrally negative Lévy process, see for example [7, 10, 11, 12, 13, 14, 18] and references therein. For this class of processes the so-called scale functions play a vital role in exit problems and excursion calculations.

Since a spectrally negative Lévy process X does not have positive jumps it holds for any $\lambda \geq 0$ that

$$\mathbb{E}[e^{\lambda X_t}] = e^{\psi(\lambda)t},$$

where ψ is the Laplace exponent which is of the form

$$\psi(\lambda) = a\lambda + \frac{\sigma^2}{2}\lambda^2 + \int_{(-\infty, 0)} (1 - e^{\lambda x} - \lambda x 1_{\{x > -1\}}) \Pi(dx).$$

Here $a \in \mathbb{R}$, $\sigma^2 \geq 0$ and Π is called the Lévy measure of X which satisfies $\int_{(-\infty,0)}(1 \wedge x^2)\Pi(dx) < \infty$. It is easily shown that ψ is a strictly convex function on $[0, \infty)$ and that $\psi(\infty) = \infty$. This allows us to introduce the right inverse of ψ which we denote by Φ .

Denote by $\overline{X}_t := \sup_{0 \leq s \leq t} X_s$ the supremum process and for $q \geq 0$ let \mathbf{e}_q be an exponentially distributed random variable with parameter q , independent of X (for $q = 0$ we interpret \mathbf{e}_q as being infinite with probability 1). For spectrally negative Lévy processes the Wiener–Hopf factorisation simplifies, since in this case $\overline{X}_{\mathbf{e}_q}$ can be shown to be exponentially distributed with parameter $\Phi(q)$. If \underline{X} denotes the infimum process, this leads to

$$(1) \quad \mathbb{E}[e^{\lambda \underline{X}_\infty}] = \psi'(0+) \frac{\lambda}{\psi(\lambda)}$$

whenever $\psi'(0+) > 0$, which corresponds to the situation that X drifts to $+\infty$. Further details can be found in Chapter VII of [3] and Chapter 8 of [15] and references therein.

For $a, b \in \mathbb{R}$, introduce the first passage times

$$\tau_a^+ = \inf\{t \geq 0 : X_t > a\} \quad \text{and} \quad \tau_b^- = \inf\{t \geq 0 : X_t < b\}$$

and for $x \in \mathbb{R}$ denote by \mathbb{P}_x the law of X given that $X_0 = x$. The following definition of a scale function is essentially due to Zolotarev [23] and Takács [22] (further references can be found in the books [3] and [15]).

Theorem 1. *For $q \geq 0$ there exists a unique continuous function (called scale function) $W^{(q)} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ satisfying*

$$\int_0^\infty e^{-\lambda x} W^{(q)}(x) dx = \frac{1}{\psi(\lambda) - q} \quad \forall \lambda > \Phi(q).$$

and for $0 < x < a$

$$\mathbb{E}_x[e^{-q\tau_a^+} 1_{\{\tau_a^+ < \tau_0^-\}}] = \frac{W^{(q)}(x)}{W^{(q)}(a)}.$$

One way of proving this result is to first assume that $q = 0$ and $\psi'(0+) > 0$ and by taking $W^{(0)}(x) = \mathbb{P}(\underline{X}_\infty \geq -x)/\psi'(0+)$. From (1) it follows that for any $\lambda \geq 0$

$$\int_0^\infty e^{-\lambda x} W^{(0)}(x) dx = \frac{1}{\psi(\lambda)}.$$

Furthermore, for $0 < x < a$ an application of the strong Markov property yields

$$\begin{aligned} \mathbb{P}(\underline{X}_\infty \geq -x) &= \mathbb{E}_x[\mathbb{P}_x(\underline{X}_\infty \geq 0 | \mathcal{F}_{\tau_a^+})] \\ &= \mathbb{E}_x[1_{\{\tau_a^+ < \tau_0^-\}} \mathbb{P}_a(\underline{X}_\infty \geq 0)] + \mathbb{E}_x[1_{\{\tau_a^+ > \tau_0^-\}} \mathbb{P}_{X_{\tau_0^-}}(\underline{X}_\infty \geq 0)]. \end{aligned}$$

It does not take much work to show that the second quantity in the last line is equal to zero, and hence $\mathbb{P}_x(\tau_a^+ > \tau_0^-) = W^{(0)}(x)/W^{(0)}(a)$. The other cases ($q > 0$ or $\psi'(0+) \leq 0$) can be deduced by using a change of measure based on the martingale $e^{\Phi(q)X_t - qt}$ which leads to a probability measure under which the process drifts to $+\infty$.

Explicit examples of scale functions are given in [4, 16, 17] and numerical schemes for scale functions can be found in [20] and [21]

To indicate the usefulness of scale functions in the fluctuation theory of spectrally negative Lévy processes we mention two applications. Firstly, the last passage time below zero of a spectrally negative Lévy process, which was studied in [6] with an application to insurance mathematics in mind. In [2] this is extended as the Laplace transform is given for various last passage times before an exponential independent, exponentially distributed time. As a corollary, we deduce an extension to a result from [8] where it was shown that for a spectrally negative stable process X of index α

$$\mathbb{P}(\exists t > 0 : X_t = \bar{X}_t = t) = \frac{1}{\alpha}.$$

Secondly, in [19], using martingale techniques, the resolvent measure was found for a reflected spectrally negative Lévy process. In [9] a new proof was given based on excursion calculation making use of the relationship between the scale function and the process conditioned to stay positive. Using results from [5] about a general Lévy process conditioned to stay positive this method was extended in [1] and, as an application, the resolvent measure was found for a reflected symmetric stable process.

REFERENCES

- [1] E.J. Baurdoux, *Some excursion calculations for reflected Lévy processes*, ALEA Lat. Am. J. Probab. Math. Stat. **6** (2009), 149–162.
- [2] E.J. Baurdoux, *Last exit before an exponential time for spectrally negative Lévy processes*, J. Appl. Probab. **46** (2009), 542–558.
- [3] J. Bertoin, *Lévy processes*, Cambridge University Press (1996).
- [4] J. Bertoin, *On the first exit time of a completely asymmetric stable process from a finite interval*, Bull. London Math. Soc. **28** (1996), 156–169.
- [5] L. Chaumont and R.A. Doney, *On Lévy processes conditioned to stay positive*, Electron. J. Probab. **10** (2005), 948–961.
- [6] S.N. Chiu and C.C. Yin, *Passage times for a spectrally negative Lévy process with applications to risk theory*, Bernoulli **11** (2005), no. 3, 511–522.
- [7] K. Dębicki, T. Dieker and T. Rolski, *Quasi-product forms for Lévy driven fluid networks*, Math. Oper. Res. **32** (2007), 629–647.
- [8] R.A. Doney, *Hitting probabilities for spectrally positive Lévy processes*, J. London Math. Soc. **44** (1991), 566–576.
- [9] R.A. Doney, *Some excursion calculations for spectrally one-sided Lévy processes*, Sémin. de Prob. XXXVIII (2005), 5–15.
- [10] P. Glynn and M. Mandjes, *Simulation-based computation of the workload correlation function in a Lévy-driven queue*, to appear in J. Appl. Probab. (2010).
- [11] O. Kella, *Parallel and tandem fluid networks with dependent Lévy inputs*, Ann. Appl. Probab. **3** (1993), 682–695.
- [12] O. Kella, O.J. Boxma and M. Mandjes, *On Lévy-driven vacation models with correlated busy periods and service interruptions*, Queueing Syst. **64** (2010), 359–382.
- [13] O. Kella and W. Whitt, *A tandem fluid network with Lévy input*. In: Queueing and Related Models, Oxford University Press (1992).
- [14] T. Konstantopoulos, G. Last and S. Lin, *On a class of Lévy stochastic networks*, Queueing Syst. **46** (2004), 409–437.

- [15] A.E. Kyprianou, *Introductory Lectures on Fluctuations of Lévy Processes with Applications*, Springer (2006).
- [16] A.E. Kyprianou and V. Rivero, *Special, conjugate and complete scale functions for spectrally negative Lévy processes*, *Elect. J. Probab.* **13** (2008), 1672–1701.
- [17] A.E. Kyprianou and F. Hubalek, *Old and new examples of scale functions for spectrally negative Lévy processes* To appear in *Sixth Seminar on Stochastic Analysis, Random Fields and Applications*, Birkhäuser (2010).
- [18] P. Lieshout and M. Mandjes, *Asymptotic analysis of Lévy-driven tandem queues*, *Queueing Syst.* **60** (2008), 203–226.
- [19] M.R. Pistorius, *On exit and ergodicity of the spectrally one-sided Lévy process reflected at its infimum*, *J. Theoret. Probab.* **17** (2004), 183–220.
- [20] L.C.G. Rogers, *Evaluating first-passage probabilities for spectrally one-sided Levy processes*, *J. Appl. Probab.* **37** (2000), 1173–1180.
- [21] B.A. Surya, *Evaluating scale functions of spectrally negative Lévy processes*, *J. Appl. Probab.* **45** (2008), 135–149.
- [22] L. Takács, *Combinatorial methods in the theory of stochastic processes*. Wiley (1967).
- [23] V.M. Zolotarev, *The first-passage time of a level and the behavior at infinity for a class of processes with independent increments*, *Theory Probab. Appl.* **9** (1964) 653–662.

Spectral gaps for large-scale Markov processes

TON DIEKER

The research presented in the talk aims to assess *dominant behavior* of large-scale stochastic processes through quantities that can be provably *calculated numerically* by small-scale algorithms. The particular focus lies on Markov chains on finite but large state spaces, such as those arising in the context of graphs. Here small-scale algorithms must be understood as algorithms which operate on objects of size many orders of magnitude less than the size of the state space. Many key properties of such large-scale processes can be deduced from the eigenvalues and eigenspaces of their transition matrices. For instance, as a consequence of the spectral theorem, the so-called spectral gap determines the ‘speed’ at which the process traverses its state space (the spectral gap is defined as the difference between the trivial eigenvalue 1 and the second largest eigenvalue of the transition matrix).

Calculating eigenvalues numerically becomes prohibitive for Markov chains on large state spaces, which poses significant challenges for determining spectral characteristics such as spectral gaps. This motivates the search for theory which, under structural assumptions on the Markov chain, allows one to connect the dominant large-scale eigenvalues and eigenvectors to corresponding small-scale quantities which can be computed quickly (e.g., numerically).

My talk describes a first result of this kind. Given a permutation σ of $1, \dots, n$, let the neighborhood of σ be all assignments σ' which are different from σ at exactly two places. (This ‘transposition’ neighborhood is a standard neighborhood structure for local search in the quadratic assignment problem [1], as the change in the objective function from σ to σ' can be calculated relatively efficiently.) Suppose that $\sigma(i)$ and $\sigma(j)$ are interchanged with probability α_{ij} , where $\sum_{1 \leq i < j \leq n} \alpha_{ij} = 1$. This leads to a Markov chain $\{\sigma_m : m \geq 0\}$ on the set of all permutations, to be referred to as the ‘large Markov chain’ since the state space is of size $n!$. A much

smaller Markov chain, of size n , arises when looking at the position of a fixed number in the permutation (say number 1); this is the process $\{\sigma_m^{-1}(1) : m \geq 0\}$. The dynamics of this smaller process are readily described in terms of the α_{ij} . Indeed, if the state of the small process is currently i , then it becomes j in the next step with probability α_{ij} . We call the process the ‘small Markov chain’. Note that, under natural irreducibility assumptions, the small chain remains in its current state with positive probability.

The following theorem has originally been conjectured by David Aldous and Persi Diaconis in 1992 for a special case, and has received much attention over the years [3, 6, 7, 8]. The proof relies on several novel insights, and I intend to explore the extent of which these insights can be further utilized and exploited. The theorem is the main result from my article [5] as well as the simultaneous work by Caputo *et al.* [2]. The paper by Caputo *et al.* contains the only known proof of a key technical lemma needed in the proof, which was conjectured to hold in my paper [5]. Recall that the spectral gap of a Markov chain is defined as the difference between the largest and the second largest eigenvalue of the transition matrix.

Theorem 2. *For any n and any choice of $\{\alpha_{ij}\}$, the spectral gap of the large Markov chain equals the spectral gap of the small Markov chain. In particular, the spectral gap of the large-scale chain can be determined by a small-scale algorithm.*

The intuition behind this theorem relies on the connection, briefly mentioned earlier, between the spectral gap and the speed at which a Markov chain traverses its state space. The theorem intuitively says that the large Markov chain cannot reach equilibrium before the small Markov chain, and vice versa it is impossible for the small chain to reach equilibrium with the large chain being ‘away’ from equilibrium. Thus, the speed at which the large-scale process converges to equilibrium is completely determined by a small-scale quantity. Interestingly, there exist numerical examples for which a result like Theorem 2 does not hold even though there is no (apparent) reason for the intuition to fail.

It is my hope to establish further results in the spirit of Theorem 2, specifically focusing on the following two directions:

- (1) Prove results in the spirit of Theorem 2 for neighborhood structures different from the one arising by swapping two numbers. Although the proof of Theorem 2 has recently been found, the key inequality proved by Caputo *et al.* [2] arguably relies on ‘brute force’ techniques which cannot be expected to hold much value when building a general theory for results of this type. Thus, novel insights will be needed.
- (2) Establish theory outside of the setting of permutations. The spectral gap is determined by the largest nontrivial eigenvalue, but other eigenvalues of the large Markov chain are of interest from the point of view of stochastic processes on other discrete structures than permutations; examples include perfect matchings, Hamiltonian cycles, or colorings of a graph [2, 4, 5]. The objective of this research task is to find analogs of Theorem 2 in these

cases. In general one cannot expect that the small Markov chain has a state space of size n , but perhaps of size polynomial in n .

REFERENCES

- [1] R. E. Burkard, E. Cela, P. M. Pardalos, and L. S. Pitsoulis, *The quadratic assignment problem*, Handbook in Combinatorial Optimization (D.-Z. Du and P. M. Pardalos, eds.), vol. 3, Kluwer, 241–337 (1998).
- [2] P. Caputo, T. M. Liggett, and T. Richthammer, *Proof of Aldous’ spectral gap conjecture*, J. Amer. Math. Soc. **23** (2010), 831–851.
- [3] F. Cesi, *On the eigenvalues of Cayley graphs on the symmetric group generated by a complete multipartite set of transpositions*, arxiv.org/abs/0902.0727 (2009).
- [4] P. Diaconis and S. P. Holmes, *Random walks on trees and matchings*, Electron. J. Probab. **7** (2002), no. 6.
- [5] A. B. Dieker, *Interlacings for random walks on weighted graphs and the interchange process*, SIAM J. Discrete Math. **24** (2010), 191–206.
- [6] S. Handjani and D. Jungreis, *Rate of convergence for shuffling cards by transpositions*, J. Theoret. Probab. **9** (1996), 983–993.
- [7] B. Morris, *Spectral gap for the interchange process in a box*, Electron. Commun. Probab. **13** (2008), 311–318.
- [8] S. Starr and M. Conomos, *Asymptotics of the spectral gap for the interchange process on large hypercubes*, arxiv.org/abs/0802.01368 (2009).

Steady state GI/GI/n queue in the Halfin-Whitt regime

DAVID GAMARNIK

(joint work with David Goldberg)

Parallel server $GI/GI/n$ queueing systems can operate in a variety of regimes that balance between efficiency and quality of offered service. This is captured by the Halfin-Whitt (H-W) regime, which can be described as critical with respect to the probability that an arriving customer has to wait for service. Namely, in this regime the stationary probability of wait is bounded away from both 0 and 1 as the number of servers grows. Although studied originally by Erlang [1] and Jagerman [4], the regime was formally introduced in [3] by Halfin and Whitt, who established the criticality of the probability that a customer waits in terms of the square-root spare capacity rule for large multi-server (FCFS $GI/GI/n$) queues with exponentially distributed processing times. They also showed that as the size of the multi-server queue grows, the queueing process (properly scaled) converges weakly to a non-trivial positive recurrent diffusion. In addition, they showed that the sequence of normalized steady-state queue length distributions $\{(Q^n(\infty) - n)^+ n^{-\frac{1}{2}}\}$ is tight, and explicitly computed the associated weak limit, where $Q^n(\infty)$ denotes the steady state number of jobs in the $GI/M/n$ queueing system.

Similar weak convergence results under the H-W scaling were subsequently obtained for $GI/GI/n$ systems with non-exponential service time distributions [8], [5], [6], [2], with the most general results appearing in [10] (and follow-up papers

[9],[7]). As the theory of weak convergence generally relies heavily on the assumption of compact time intervals, the most general of these results hold only in the transient regime. Indeed, with the exception of [3] (which treats exponential processing times), [5] (which treats deterministic processing times), and [2] (which treats processing times with finite support), all of the aforementioned results are for the associated sequence of normalized *transient* queue length distributions only, leaving many open questions about the associated *steady-state* queue length distributions. In [2], the authors show (for processing times with finite support) that the sequence of steady-state queue length distributions (normalized by $n^{\frac{1}{2}}$) is tight, and has a limit whose tail decays exponentially fast. In particular it is shown that $Q(\infty)$ which is the unique weak limit of the sequence $\{(Q^n(\infty) - n)^+ n^{-\frac{1}{2}}\}$, satisfies the property

$$(1) \quad \mathbb{E}[\exp(\theta Q(\infty))] < \infty, \quad \forall \theta < \frac{2\beta}{c_a^2 + c_s^2},$$

$$(2) \quad \mathbb{E}[\exp(\theta Q(\infty))] = \infty, \quad \forall \theta > \frac{2\beta}{c_a^2 + c_s^2},$$

where β is the spare capacity parameter in the $GI/GI/n$ queueing system in the Halfin-Whitt regime, and c_a^2 and c_s^2 are the coefficients of variations of the interarrival and service times, respectively. It was conjectured in [2] that these results should hold for more general processing time distributions, but prior to this work no further progress had been made along those lines.

In this paper we establish that the sequence of queue length of the $GI/GI/n$ system in the Halfin-Whitt regime, normalized by n is tight, under very minor assumptions on the interarrival and service times. Namely, we show that the sequence $\{(Q^n(\infty) - n)^+ n^{-\frac{1}{2}}\}$ is tight. Furthermore, we partially confirm the conjecture introduced in [2] regarding the decay rate. In particular, for every weak limit $Q(\infty)$ of the sequence $\{(Q^n(\infty) - n)^+ n^{-\frac{1}{2}}\}$ we show that the bound (1) holds. We also establish the matching lower bound (2) for the case when the arrival process is Poisson.

Our method of proof is based on a clever bounding the queueing process by another process which loosely speaking can be defined as $GI/GI/n$ queueing system which "never" idles. This is a system in which each server upon starting an idling phase creates an artificial job and works on it. The resulting queueing process has a cleanly defined weak limit which is a Gaussian process with negative drift. Then using the process level weak convergence technique, the tightness result is established, and the bounds on the decay rate for the limiting queue length are obtained from the known results on the supremum of a Gaussian process with a negative drift.

REFERENCES

- [1] A.K. Erlang, *On the rational determination of the number of circuits*, The Copenhagen Telephone Company, Copenhagen (1948).

- [2] D. Gamarnik and P. Momcilovic, *Steady-state analysis of a multi-server queue in the Halfin-Whitt regime*, Advances in Applied Probability **40** (2008), 548–577.
- [3] S. Halfin and W. Whitt, *Heavy-traffic limits for queues with many exponential servers*, Operations Research **29** (1981), no. 3, 567–588.
- [4] D. Jagerman, *Some properties of the Erlang loss function*, Bell System Techn. J. **53** (1974), no. 3, 525–551.
- [5] P. Jelenkovic, A. Mandelbaum, and P. Momcilovic, *Heavy traffic limits for queues with many deterministic servers*, Queueing Systems: Theory and Applications **47** (2004), no. 1–2, 53–69.
- [6] A. Mandelbaum and P. Momcilovic, *Queues with many servers: The virtual waiting-time process in the QED regime*, Math. Oper. Res. **33** (2008), no. 3, 561–586.
- [7] A.A. Puhalskii and J. Reed, *On many-server queues in heavy traffic*, Ann. Appl. Prob. **20** (2010), 129–195.
- [8] A. Puhalski and M. Reiman, *The multiclass GI/PH/N queue in the Halfin-Whitt regime*, Advances in Applied Probability **32** (2000), 564–595.
- [9] J. Reed, *The G/GI/N queue in the Halfin-Whitt regime II: Idle time system equations*, Preprint (2007).
- [10] J. Reed, *The G/GI/n queue in the Halfin-Whitt regime*, Annals of Applied Probability **19** (2009), 2211–2269.

Wide-sense Regeneration for Harris Recurrent Markov Processes: An Open Problem

PETER W. GLYNN

Harris recurrence is a widely used concept in the analysis of stochastic systems.

Definition 1 A (strong) Markov process $X = (X(t) : t \geq 0)$ taking values in a separable metric space S is said to be *Harris recurrent* if there exists a non-trivial (reference) measure η such that whenever $\eta(A) > 0$, then

$$\int_0^\infty I(X(t) \in A) dt = \infty \quad P_x \text{ a.s.}$$

for each $x \in S$ (where $P_x(\cdot) \triangleq P(\cdot | X(0) = x)$).

In the presence of Harris recurrence, it is known that one-dependent regenerative structure can always be identified within the process X . In particular, one can find randomized stopping times $(T(n) : n \geq 0)$ for which $0 \leq T(0) < T(1) < \dots$ and such that:

- (1) $((X(T(j-1) + s) : 0 \leq s < T(j) - T(j-1)) : j \geq 0)$ is a sequence of 1-dependent random elements where $T(-1) \triangleq 0$;
- (2) $((X(T(j-1) + s) : 0 \leq s < T(j) - T(j-1)) : j \geq 0)$ is a sequence of identically distributed random elements.

See [1] for details. Such regenerative structure permits one to easily develop strong laws and central limit theorems for time averages of the form $t^{-1} \int_0^t f(X(s)) ds$. Furthermore, the associated centering and scaling constants can be expressed in terms of “cycle quantities”, thereby providing straightforward estimation of these

quantities within the context of Monte Carlo simulation of these processes; see [2]. Such regenerative estimation can then be used to develop asymptotically valid confidence interval procedures for steady-state simulation of such Harris recurrent processes.

One important mathematical tool that is often available in the context of Harris recurrence is that of renewal theory.

Definition 2 A stochastic process $X = (X(t) : t \geq 0)$ is said to be *wide-sense regenerative* if there exist random times $(\beta(n) : n \geq 0)$ for which $0 \leq \beta(0) < \beta(1) < \dots$ and such that:

- (1) $(X(\beta(n) + u) : u \geq 0)$ is independent of $\beta(n)$ for $n \geq 0$;
- (2) $(X(\beta(n) + u) : u \geq 0) \stackrel{\mathcal{D}}{=} (X(\beta(0) + u) : u \geq 0)$ for $n \geq 0$.

In the presence of wide-sense regeneration, renewal equations can be developed. In particular, if X is non-delayed (so that $\beta(0) = 0$) and f is nonnegative, one can express $Ef(X(t))$ in terms of the renewal equation

$$Ef(X(t)) = Ef(X(t))I(\beta(1) > t) + \int_{[0,t]} Ef(X(t-s))P(\beta(1) \in ds)$$

for $t \geq 0$. Renewal theory can then be applied to study issues such as rates of convergence to equilibrium.

In the discrete time setting, it is known that all Harris recurrent Markov chains are both one-dependent regenerative and wide-sense regenerative (and we may take $\beta(n) = T(n)$ for $n \geq 0$ if we wish). In continuous time, this issue has not been resolved, and remains open.

Open Problem: Is every Harris recurrent Markov process wide-sense regenerative?

Settling this issue one way or the other would represent a major step forward in our understanding of the probabilistic and analytical structure of Harris recurrent Markov processes.

REFERENCES

- [1] K. Sigman, *One-dependent regenerative processes and queues in continuous time*, Math. Oper. Res. **15** (1990), 175–189.
- [2] P. W. Glynn, *Some topics in regenerative steady-state simulation*, Acta Applicandae Mathematicae **34** (1994), 225–236.

Peer to Peer Communication in Networks – Issues, Models, and Analysis

BRUCE HAJEK

(joint work with Laurent Massoulié, Sujay Sanghavi, Kuang Xu, Ji Zhu)

Distributed protocols for peer to peer file sharing, streaming video, and video on demand have revolutionized the way the majority of information is conveyed over the Internet. The peers are millions of computers, acting as both clients and servers, downloading and uploading information. Information to be shared is broken into pieces, and the pieces are traded among peers in the network. There can be turnover in the set of pieces of information being collected and/or in the set of peers collecting the information. Coding, in which groups of pieces are combined to form new pieces, can enhance the collection process. The systems are distributed and scalable. The theory for understanding peer-to-peer systems has lagged far behind our ability to mathematically model, predict, and optimize system performance. This presentation discusses stochastic models, mathematical results, and challenges relating to the performance of peer to peer communication in large networks.

1. INTRODUCTION

Peer-to-peer (P2P) communication in the Internet is communication provided through the sharing of widely distributed resources. Our focus is on peer-to-peer networks of unstructured type, such as *BitTorrent* [1], meaning there is no specific network topology to be formed by the participating peers. Use of network coding for replication of a single file is described in [2]. In [11] the same constants are obtained as in [2] but without network coding and without exchange of piece lists as in BitTorrent [1]. The work [7] considers file replication with continuous arrival of peers. The model of our paper is close to the one in [7]. Each peer receives a set of pieces upon arrival. The paper [7] studies the fluid limit model derived from the theory of density dependent jump Markov processes [4]. Simpler two-state models were derived and studied in [10] and [13]. An instability phenomenon discussed here was discovered independently by Norros et al. [9]. The paper [9] discusses interesting mechanisms for selecting rarest pieces that may stabilize the system.

2. A SYSTEM WITH A FIXED POPULATION OF PEERS

The paper [11] investigates the dissemination of multiple pieces of information in large networks where users contact each other in a random uncoordinated manner, and users upload one piece per unit time. The underlying motivation is the design and analysis of piece selection protocols for peer-to-peer networks which disseminate files by dividing them into pieces. It first investigates one-sided protocols, where piece selection is based on the states of either the transmitter or the receiver. It shows that any such protocol relying only on pushes, or alternatively only on pulls, is inefficient in disseminating all pieces to all users. It proposes a hybrid one-sided piece selection protocol – INTERLEAVE – and shows that by

using both pushes and pulls it disseminates k pieces from a single source to n users in $10(k + \log n)$ time, while obeying the constraint that each user can upload at most one piece in one unit of time, with high probability for large n . An optimal, centralized protocol, such as the one given in [8], takes $k + \log_2 n$ time in this setting. Moreover, efficient dissemination is also possible if the source implements forward erasure coding, and users push the latest-released coded pieces (but do not pull).

The paper [11] also briefly investigated two-sided protocols where piece selection is based on the states of both the transmitter and the receiver. Users carry out pushes/pulls, but have knowledge of the target's current state. We consider an initial state where n distinct pieces are present in the system, one in each user. Completion from such an initial state has been previously studied, often under the alternate title of "all-to-all communication". For this state, consider the following two-sided pull protocol:

ADVOCATE: If the user does not already possess the target's initial piece, it downloads that piece. Else it pulls a random piece from among those present in the target but absent in the user.

In this protocol each user acts as an advocate for its initial piece. If each user is restricted to downloading at most one piece in every time slot, an optimal central protocol takes at least $n - 1$ time slots to complete. The following theorem shows that the ADVOCATE protocol completes in time very close to this optimal, with high probability.

Theorem ([11]) *Starting with each user having one unique piece, the ADVOCATE protocol operating under soft constraints finishes in $n + O(\log n)$ time with high probability.*

In the above theorem the preconstant 1 of n is the best possible. The above theorem means that for large n the fraction of wasted time slots is negligible. A conjecture, backed by preliminary simulations [5], is the following:

Conjecture *Starting with each user having one unique piece, the ADVOCATE protocol operating under soft constraints finishes in n time steps with probability converging to one as $n \rightarrow \infty$. That is, at most one extra time step is needed. (This may also hold without invoking the advocate mechanism.).*

3. A SYSTEM WITH CONTINUOUS-ARRIVALS OF PEERS

The following is a composite of models in [6, 7, 14].

- C = set of subsets of $\{1, \dots, K\}$, where K is the number of pieces
- A peer with set of pieces c is a *type c peer*
- A type c peer becomes a type $c \cup \{i\}$ peer if it downloads piece $i \notin c$
- The detailed Markov state is $\mathbf{x} = (x_c : c \in C)$, with x_c = number of type c peers, and $|\mathbf{x}|$ is the number of peers in the system
- Exogenous arrivals of peers of type c form a rate λ_c Poisson process, and such processes for different c are mutually independent

- Each peer contacts other peers, chosen uniformly at random from among all peers, for opportunities to download a piece (i.e. pull) from the other peers, according to a Poisson process of rate $\mu > 0$.
- Downloads are modeled as being instantaneous. This assumption is reasonable in the context of the previous assumption.
- Random useful piece selection – when a peer of type c has an opportunity to download a piece from a peer of type s , the opportunity results in no change of state if $s \subset c$. Otherwise, the type c peer downloads one piece selected at random from $s - c$, with all $|s - c|$ possibilities having equal probability.
- There is one fixed seed, which contacts peers for opportunities to upload a random useful piece to the peers according to a Poisson process of rate U_s .
- After obtaining a complete collection of pieces, a peer remains in the system an exponentially distributed amount of time with parameter γ as a *peer seed*, and then departs.

Given a state \mathbf{x} and $c \in C$, let $T_c(\mathbf{x})$ denote the new state resulting from the arrival of a new type c peer. Given $c \in C$, $1 \leq i \leq K$ such that $i \notin c$, and a state \mathbf{x} such that $x_c \geq 1$, let $T_{c,i}(\mathbf{x})$ denote the new state resulting from a type c peer downloading piece i . Finally, let $T_d(\mathbf{x})$ denote the new state resulting from the departure of a peer seed. The positive entries of the generator matrix $Q = (q(\mathbf{x}, \mathbf{x}') : \mathbf{x}, \mathbf{x}' \in S)$ are given by:

$$\begin{aligned} q(\mathbf{x}, T_c(\mathbf{x})) &= \lambda_c \\ q(\mathbf{x}, T_{c,i}(\mathbf{x})) &= \frac{x_c}{|\mathbf{x}|} \left(\frac{U_s}{K - |c|} + \mu \sum_{s:i \in s} \frac{x_s}{|s - c|} \right) \\ &\quad \text{if } x_c > 0 \text{ and } i \notin c. \\ q(\mathbf{x}, T_d(\mathbf{x})) &= \gamma x_K \end{aligned}$$

One of the models of Massoulié and Vojnović [6, 7] is the special case of the above with $\gamma = \infty$ and $U_s = 0$ (no peer seeds or fixed seeds). The model is a refinement of the two-dimensional models of [14, 10]. Massoulié and Vojnović [6, 7] applied the theory of density-dependent jump Markov processes (see [4]) to analyze their model. They found that there is a finite resting point of the fluid ordinary differential equation. The model is a bit like an $M/GI/\infty$ system because the total service rate tends to scale linearly with the number of peers.

Norros et al. [9] and, independently, Ji Zhu and I, found that the actual system behavior, in some regimes, is not accurately predicted by assuming all pieces are in equal supply. Rather, if the rate of the fixed seed U_s is small enough and the departure rate of peer seeds γ is large enough, the system can get into a state such that all the peers are missing the same piece, and most of the peers have all the pieces except that missing piece. That is the main intuition behind the following result. (A proof of this result for the special case $\gamma = \infty$ and $\lambda_s = 0$ for $s \neq \emptyset$ is given in [3].)

Theorem (with Ji Zhu, in preparation) Given $(\lambda_c : c \in C)$, U_s , μ , and γ , the system is positive recurrent if either

- $\gamma \leq \mu$ and for $1 \leq k \leq K$, it is possible for type k pieces to enter the system.
- $0 < \mu < \gamma \leq \infty$ and the following condition holds for each $k \in C$:

$$(1) \quad \lambda_{total} < \left[U_s + \sum_{k \in C} \lambda_C(K + 1 - |C|) \right] \left(\frac{1}{1 - \mu/\gamma} \right)$$

Conversely, the system is transient if $\gamma > \mu$ and, for some k ,

$$\lambda_{total} > \left[U_s + \sum_{k \in C} \lambda_C(K + 1 - |C|) \right] \left(\frac{1}{1 - \mu/\gamma} \right).$$

We remark that the stability condition (1) is equivalent to

$$\sum_{C:k \notin C} \lambda_C < \left(U_s + \sum_{C:k \in C} \lambda_C(K - |C|) \right) + \left[U_s + \sum_{k \in C} \lambda_C(K + 1 - |C|) \right] \left(\frac{\mu/\gamma}{1 - \mu/\gamma} \right)$$

4. LIVE-STREAMING WITH A FIXED POPULATION OF PEERS

The interesting work of Zhao, Liu, and Chiu on the optimal piece selection policy for live-streaming was reviewed in the presentation. The following comments on this model are taken from [12]. Zhao, Liu, and Chiu [15] simulated and formulated a mean field approximation. The state of a peer is the set of buffer positions filled. The mean field approximation taking into account the entire buffer occupation as a state (so 2^N states for a buffer of length N) matches simulation results better than assuming buffer positions are occupied independently of each other. Zhao, Liu, and Chiu found that a mixture of greedy and rarest first strategy performs much better than either strategy. Kuang Xu [12] showed uniqueness of the mean field fixed points for those strategies. He also showed that if a tit for tat constraint is added, then there are multiple fixed points, and the earliest deadline first strategy is not always the optimal greedy strategy for a peer interacting with an arbitrary distribution of buffer occupancy for contacted peers. Kuang numerically calculated the mean field game equilibria for small buffer sizes.

ACKNOWLEDGEMENT

Discussions with Da Ming Chiu are gratefully acknowledged.

REFERENCES

- [1] B. Cohen, *Incentives build robustness in BitTorrent*, P2PECON Workshop (2003).
- [2] S. Deb and M. Médard, *Algebraic gossip: A network coding approach to optimal multiple rumor mongering*, Allerton Conference on Communication, Control, and Computing (2004).

- [3] B. Hajek and J. Zhu, *The missing piece syndrome in peer to peer communication*, February 2010, available on arXiv:1002.3493.
- [4] T. G. Kurtz, *Approximation of population processes*, volume 36 of CBMS-NSF Regional Conference Series in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, Pa. (1981).
- [5] J. Martin, *Personnal communication*, March 2010.
- [6] L. Massoulié and M. Vojnović, *Coupon replication systems*, In SIGMETRICS '05: Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems, pages 2–13, New York, NY, USA, ACM (2005).
- [7] L. Massoulié and M. Vojnović, *Coupon replication systems*, IEEE/ACM Trans. Networking **16** (2008), 603–616.
- [8] J. Munding, R. Weber, and G. Weiss, *Optimal scheduling of peer-to-peer file dissemination*, J. Scheduling (2007).
- [9] I. Norros, H. Reittu, and T. Eirola, *On the stability of two-chunk file-sharing systems*, available on arXiv:0910.5577 (2009).
- [10] D. Qiu and R. Srikant, *Modeling and performance analysis of BitTorrent-like peer-to-peer networks*. In SIGCOMM '04, 367–378, New York, NY, USA, ACM (2004).
- [11] S. Sanghavi, B. Hajek, and L. Massoulié *Gossiping with multiple messages*, IEEE Transactions on Information Theory, **53** (2007), 4640–4654.
- [12] K. Xu, *Modeling and analysis of peer-to-peer (p2p) live video streaming*, Bachelor of Science Thesis, University of Illinois at Urbana-Champaign (2009).
- [13] X. Yang and G. de Veciana, *Service capacity of peer to peer networks*, In IEEE INFOCOM, 1–11 (2004).
- [14] X. Yang and G. de Veciana, *Performance of peer-to-peer networks: Service capacity and role of resource sharing policies*, Performance Evaluation **63** (2006), 175–194.
- [15] B.Q. Zhao, J. Lui, and D. M. Chiu, *Exploring the optimal chunk selection policy for data-driven P2P streaming systems*, In Ninth IEEE Int. Conf. on Peer-to-Peer Computing (P2P '09), 271–280, September 2009.

Shortest weight routing in complex networks

GERARD HOOGHIEMSTRA

(joint work with Shankar Bhamidi, Remco van der Hofstad)

We consider random graphs with i.i.d. weights on the edges and focus on distances between two randomly chosen nodes in these graphs. Here the distance is defined as the number of edges on the path with minimal total weight.

A surprising universal picture emerges: there is a central limit theorem (CLT) for many random graphs and edge weights (but not for all). This CLT has been proven for the configuration model (including the case where the degree distribution has a heavy tail) in [2], and for the Erdős-Rényi graph in [2]. In both cases the weights are i.i.d. exponential with mean 1. The mean and variance are asymptotically equal and of order $\log n$, where n denotes the number of vertices of the underlying random graph.

For the complete graph the CLT is true (with adapted asymptotic mean and variance) for i.i.d. weights of the form E^s , with E exponential and $s > 0$ ([4]).

However, on the complete graph with edge-weights E^{-s} , $s > 0$, the hopcount remains bounded and the CLT is no longer valid. In this case we were able show

that the hopcount converges in probability to the closest integer to $s + 1$, larger than or equal to 2 ([3]).

In many applications, edge weights represent cost structure of the use of an edge, such as actual economic costs or congestion costs across edges. Actual time delay experienced by vertices in the network is given by hopcount H_n , which is the number of edges on shortest-weight path. One of the most important questions is *How does weight structure influence hopcount and what is the weight of the minimal path?*

The problem with exponential edge weights has received tremendous attention on the complete graph; it is our intention to extend these results to general (random) graphs. We start by considering the configuration model.

Let n be number of vertices. Consider an i.i.d. sequence of degrees D_1, D_2, \dots, D_n with given distribution. We pay special attention to power-law degrees, i.e.,

$$P(D_1 = k) = c_\tau k^{-\tau},$$

where c_τ is constant and $\tau > 2$.

The configuration model is constructed from a given degree sequence for which we take i.i.d. sequence D_1, D_2, \dots, D_n as defined above. Assign to vertex j degree D_j and assume that the total degree $L_n = \sum_{i=1}^n D_i$ is even. Incident to vertex i we have D_i ‘stubs’ or half edges.

- (1) Connect stubs to create edges as follows: Number stubs from 1 to L_n in any order. First connect first stub uniformly at random with one of *other* $L_n - 1$ stubs.
- (2) Continue with second stub (when not connected to the first one) and so on, until all stubs are connected...

The following theorem has been proven in [1]

Theorem 3. *Let H_n be number of edges between two uniformly chosen vertices on CM with i.i.d. exponential edge weights. Assume that $D \geq 2$ a.s. and $\nu = \frac{E[D(D-1)]}{E[D]} > 1$. For $\tau > 3$ or $\tau \in (2, 3)$,*

$$\frac{H_n - \alpha \log n}{\sqrt{\alpha \log n}} \xrightarrow{d} Z,$$

where Z is standard normal, and

$$\begin{aligned} \alpha &= \frac{\nu}{\nu-1} > 1 && \text{for } \tau > 3, \\ \alpha &= \frac{2(\tau-2)}{\tau-1} \in (0, 1) && \text{for } \tau \in (2, 3). \end{aligned}$$

It is interesting to compare this theorem to two older results concerning the *graph distance* between two uniformly chosen vertices, which were proven in [5] and in [7], respectively. In these theorems \tilde{H}_n denotes the graph distance between two uniform vertices.

Theorem 4. *When $\tau > 3$ and $\nu > 1$,*

$$\frac{\tilde{H}_n}{\log_\nu n} \xrightarrow{P} 1.$$

Theorem 5. *When $\tau \in (2, 3)$,*

$$\frac{\tilde{H}_n}{\log \log n} \xrightarrow{P} \frac{2}{|\log(\tau - 2)|},$$

Below we present a short discussion on the differences between the results for the hopcount and for the graph-distance.

- (1) Notice that the introduction of edge weights have a remarkable effect on the distances: with edge weights distances scale as $\log n$ for all $\tau > 2$, whereas without edge weights the distance scales in the same way for $\tau > 3$, whereas for $\tau \in (2, 3)$, distances are ultra small and scale as $\log \log n$.
- (2) It is known that the fluctuation of the graph distance around the centering is bounded in probability [5] and [6]. Including the edge weights the asymptotic variance of the hopcount is of the same order as the asymptotic mean. Moreover the asymptotic variance of the hopcount is equal to the asymptotic mean of the hopcount.
- (3) For the proof of the first theorem we compare the neighborhood of a uniform vertex to branching process, and use wealth of results on FPP on trees.
- (4) Finally, we observe a surprisingly universal behavior for FPP on configuration model. Universality is a leading paradigm in statistical physics. Only few examples exist where universality can be proved rigorously.

The above observations lead to the following key question: To what extent is universality true for processes on random graphs models?

To shed some more light on this key question consider the complete graph K_n with edge weights E_e^s , where E_e are i.i.d. exponentials. Then the following theorem holds ([4]):

Theorem 6. *Let H_n be number of edges of shortest path between two uniform vertices. Then, with Z standard normal,*

$$\frac{H_n - s \log n}{\sqrt{s^2 \log n}} \xrightarrow{d} Z.$$

On the other hand: consider again the complete graph K_n but now take the i.i.d. edge weights E_e^{-s} , $s > 0$. Then according to [3]

Theorem 7. *For almost all $s > 0$, the hopcount converges in probability to the nearest integer of $s + 1$.*

This raises the following open problem what are the universality classes for FPP on the complete graph?

We finish with a tightness-proof of the hopcount H_n on the complete graph with edge weights E^{-s} , $s > 0$.

Observe that the maximum of n independent exponentials scales like $\log n$, and that therefore, **whp**,

$$(1) \quad W_n \geq H_n \log^{-s}(n(n-1)/2).$$

On the other hand W_n is at most equal to the minimal weight of all two-edge paths between 1 and n , so that:

$$(2) \quad W_n \leq \min_{2 \leq j \leq n-1} ((E'_j)^{-s} + (E''_j)^{-s}),$$

It is not hard to verify, that (2) implies that **whp**,

$$(3) \quad W_n \leq \frac{C}{(\log n)^s}.$$

Together this yields tightness of H_n .

REFERENCES

- [1] S. Bhamidi, R. van der Hofstad, and G. Hooghiemstra, *First passage percolation on random graphs with finite mean degrees*, Annals of Applied Probab. **20** (2010), 1907–1965.
- [2] S. Bhamidi, R. van der Hofstad, and G. Hooghiemstra, *First passage percolation on the Erdős-Rényi random graph*, Preprint (2010), arXiv:1005.4104B.
- [3] S. Bhamidi, R. van der Hofstad, and G. Hooghiemstra, *Weak disorder in the stochastic mean-field model of distance II*, arXiv:1009.4025, 2010
- [4] S. Bhamidi and R. van der Hofstad, *Weak disorder asymptotics in the stochastic mean-field model of distance*, Preprint (2009), arXiv:1002.4362v1 [math.PR].
- [5] R. van der Hofstad, G. Hooghiemstra, and P. VanMieghem, *Distances in random graphs with finite variance degrees*, Random Structures & Algorithms, **27** (2005), 76–123.
- [6] R. van der Hofstad, G. Hooghiemstra, and D. Znamenski, *Distance in random graphs with finite mean, and infinite variance degrees*, Electronic Journal of Probability, **12** (2007), 703–766.
- [7] I. Norros, and H. Reittu, *On the power-law random graph model of massive data networks* Performance Evaluation, **55** (2004), 3–23.

Equilibrium analysis of large scale stochastic games

RAMESH JOHARI

A range of modern systems, such as markets, communication systems, and online social systems, are characterized by *dynamic* interactions among a *large population* of agents. While game theoretic methods appear to hold promise for providing insight into incentives of agents in such systems, they often struggle on two fronts. First, computing equilibria in large dynamic games rapidly becomes intractable, due to the curse of dimensionality. Second, equilibrium concepts for dynamic games make strong rationality assumptions about players' behavior; these may not hold even with moderate numbers of players.

These challenges have spurred renewed recent interest across economics, operations research, and control theory in an approximation methodology that we term *mean field equilibrium* (MFE). In MFE, each player reacts to only the long run average state of other players. This construction simplifies both computational and rationality requirements.

We present an overview of recent results concerning mean field equilibrium. First, we describe MFE in dynamic stochastic games with strategic complementarities (joint work with S. Adlakha); formally, these are games where the payoff of a player has increasing differences between her own state and the aggregate

empirical distribution of the states of other players. Such games naturally arise in models of interdependent security, or in collaborative filtering systems. We prove that such games possess MFE, and further, that a natural, myopic best response dynamic converges to MFE. We provide an analysis of sensitivity of equilibria to parameter changes.

We also demonstrate that under appropriate continuity conditions on model primitives, MFE is a good approximation to equilibria of general stochastic games with finitely many players (joint work with S. Adlakha and G.Y. Weintraub).

We devote the bulk of the talk to surveying an application of MFE to dynamic auctions with learning. In particular, we consider a setting with a sequence of repeated second price auctions, where bidders do not know their valuations. Each time a bidder wins, they receive an additional observation correlated with their true valuation. Thus bidders have a tendency to bid above their current expected valuation, since they have a value for learning. How much do they overbid? Characterizing this overbidding behavior structurally is intractable in auctions with finitely many players, but we demonstrate that a simple description of MFE can be constructed in a limit where the number of players grows large. (Joint work with K. Iyer and M. Sundararajan.)

REFERENCES

- [1] S. Adlakha and R. Johari, *Mean field equilibrium in dynamic games with complementarities* (2010). Available from: http://www.stanford.edu/~rjohari/uploads/MFE_comp_TR.pdf
- [2] S. Adlakha, R. Johari, and G.Y. Weintraub, *Equilibria of dynamic games with many players: existence, approximation, and market structure* (2010). Available from: http://www.stanford.edu/~rjohari/uploads/SE_market_TR.pdf.

Some recent observations regarding growth collapse processes and their generalizations

OFFER KELLA

(joint work with Onno Boxma, Marc Yor, Andreas Löpker and David Perry)

This talk summarizes four recent studies related or motivated by growth collapse processes. Growth collapse processes are processes that increase and from time to time collapse to a fraction of their pre-collapse times. When the increase is linear, the decrease is to a fixed fraction of the pre-collapse time and the epochs between collapse times are independent and identically distributed (i.i.d.) exponential random variables, these models are used to model TCP window size control. Other phenomena that exhibit this kind of behavior are stock indexes when one speeds up time and observes the increases followed by crashes or sand piles which build up until they collapse to a fraction of their former height.

The first study is [9] where the aim is to generalize earlier results for the steady state distribution of growth collapse processes. It begins with a stationary setup with some relatively general growth process and it is observed that under certain expected conditions point and time stationary versions of the processes exist as

well as a limiting distribution for these processes which is independent of initial conditions and necessarily has the marginal distribution of the stationary version. It is then specialized to the case where an i.i.d. structure holds and further specialize to the case where the growth process is a nondecreasing Lévy process and in particular linear and the times between collapses form an i.i.d. sequence. Known results can be seen as special cases for example when the inter-collapse times form a Poisson process or when the collapse ratio is deterministic. Finally a relation between these processes and shot-noise-type processes is commented on and it is observed that under certain conditions, the steady state distribution of one may be directly inferred from the other.

The second study which seems unrelated is [11] and was actually motivated by observing that a growth collapse process satisfies a stochastic linear equation where the free process (and also the integrator) has bounded variation on finite intervals. In this paper a new representation is given for the solution for a stochastic linear equation of the form $X_t = Y_t + \int_{(0,t]} X_{s-} dZ_s$ where Z is a càdlàg semimartingale and Y is a càdlàg adapted process with bounded variation on finite intervals. As an application the case where Y and $-Z$ are nondecreasing, jointly have stationary increments and the jumps of $-Z$ are bounded by 1 is studied. Special cases of this process are shot-noise processes, growth collapse (additive increase, multiplicative decrease) processes and clearing processes. When Y and Z are in addition independent Lévy processes the resulting X is called a generalized Ornstein-Uhlenbeck process.

The third study is [10] where a growth collapse model in a random environment is considered, where the input rates may depend on the state of an underlying irreducible Markov chain and at state change epochs there is a possible downward jump to a level which is a random fraction of the level just before the jump. The distributions of these jumps are allowed to depend on both the originating and target states. Under a very weak assumption an explicit formula is developed for the conditional moments (of all orders) of the time stationary distribution. Special cases are then considered and it is shown how to use this result to study a growth collapse process in which the times between collapses have a phase type distribution.

The fourth study is [3]. In this paper is about generalizing existing results for the steady state distribution of growth collapse processes with independent exponential inter-collapse times to the case where they have a general distribution on the positive real line having a finite mean. In order to compute the moments of the stationary distribution, no further assumptions are needed. However, in order to compute the stationary distribution, the price that one is required to pay is the restriction of the collapse ratio distribution from a general one concentrated on the unit interval to minus-log-phase-type distributions. A random variable has such a distribution if the negative of its natural logarithm has a phase type distribution. Thus, this family of distributions is dense in the family of all distributions concentrated on the unit interval. The approach is to first study a certain

Markov modulated shot-noise process from which the steady state distribution for the related growth collapse model can be inferred via level crossing arguments.

Some (certainly not all) earlier literature is given in [1, 2, 4, 5, 6, 7, 8, 12, 13, 14] and references therein.

REFERENCES

- [1] E. Altman, K. Avrachenkov and C. Barakat, *A stochastic model of TCP/IP with stationary random losses*, IEEE/ACM Transactions on Networking **13** (2005), 356-369.
- [2] E. Altman, K. Avrachenkov, A. A. Kherani and B. J. Prabhu, *Performance analysis and stochastic stability of congestion control protocols*, Proceedings of IEEE Infocom (2005).
- [3] O. Boxma, O. Kella, O. and D. Perry, *On some tractable growth collapse processes with renewal collapse epochs*, (2010), submitted, available at <http://pluto.msc.huji.ac.il/~mskella/pubs.html>.
- [4] O. Boxma, D. Perry, W. Stadje and S. Zacks, *A Markovian growth collapse model*, Adv. Appl. Probab. **38** (2006), 221-243.
- [5] V. Dumas, F. Guillemin, and P. Robert, *A Markovian analysis of additive-increase, multiplicative-decrease (AIMD) algorithms*, Adv. Appl. Probab. **34** (2002), 85-111.
- [6] I. Eliazar and J. Klafter, *A growth-collapse model: Lévy inflow, geometric crashes and generalized Ornstein-Uhlenbeck dynamics*, Physica A **334** (2004), 1-21.
- [7] I. Eliazar and J. Klafter, *Growth-collapse and decay-surge evolutions, and geometric Langevin equations*, Physica A **367** (2006), 106-128.
- [8] F. Guillemin, P. Robert and B. Zwart, *AIMD Algorithms and exponential functionals*, Ann. Appl. Probab. **14** (2004), 90-117.
- [9] O. Kella, *On growth collapse processes with stationary structure and their shot-noise counterparts*, J. Appl. Probab. **46** (2009), 363-371.
- [10] O. Kella and A. H. Löpker, *A Markov-modulated growth collapse model*, Prob. Eng. Inf. Sci. **24** (2010), 99-107.
- [11] O. Kella, and M. Yor, *A new formula for some linear stochastic equations with applications*, Ann. Appl. Probab. **20** (2009), 367-381.
- [12] J. S. H. Leeuwaarden and A. H. Löpker, *Transient moments of the window size in TCP*, J. Appl. Probab. **45** (2008), 163-175.
- [13] A. H. Löpker and J. S. H. van Leeuwaarden, *Hitting times for multiplicative growth-collapse processes*, EURANDOM report No. 2007-049 (2007), available at <http://www.eurandom.tue.nl/EURANDOMreports.htm>.
- [14] T. Ott, J. Kemperman and M. Mathis, *The stationary behavior of ideal TCP congestion avoidance* (1996), available at www.teunisott.com.

Diameter in weighted random graphs

MARC LELARGE

(joint work with Hamed Amini, Moez Draief)

Given a graph $G = (V, E)$, the distance $\text{dist}(a, b)$ between two nodes a and b in V is the number of edges in E in the shortest path connecting these two vertices. The diameter of G , denoted by $\text{diam}(G)$, is the maximum graph distance between any pair of vertices in V , i.e.

$$\text{diam}(G) = \max\{\text{dist}(a, b), a, b \in V, \text{dist}(a, b) < \infty\}.$$

To define the weighted diameter of a graph $G = (V, E)$, we assign to each edge $e \in E$ a weight w_e which will be an exponential random variable with parameter

one independent of everything else. For any $a, b \in V$, a path between a and b is a sequence $\pi = (e_1, e_2, \dots, e_k)$ where $e_i = (v_{i-1}, v_i) \in E$ and $v_i \in V$ for $i \in [1, k]$, with $v_0 = a$ and $v_k = b$. We write $e_i \in \pi$ to denote the fact that edge e_i belongs to the path π . For $a, b \in V$, we define

$$\text{dist}_w(a, b) = \min_{\pi \in \Pi(a, b)} \sum_{e \in \pi} w_e,$$

where the minimum is taken over all paths between a and b in the graph. The *weighted diameter* is given by

$$\text{diam}_w(G) = \max\{\text{dist}_w(a, b), a, b \in V, \text{dist}_w(a, b) < \infty\}.$$

For $n \in \mathbb{N}$, let $(d_i)_1^n$ be a sequence of non-negative integers such that $\sum_{i=1}^n d_i$ is even. By means of the configuration model [4], we define a random multigraph with given degree sequence $(d_i)_1^n$, denoted by $G^*(n, (d_i)_1^n)$ as follows. To each node i we associate d_i labeled half-edges. All half-edges need to be paired to construct the graph, this is done by randomly matching them. When a half-edge of i is paired with a half-edge of j , we interpret this as an edge between i and j . The graph $G^*(n, (d_i)_1^n)$ obtained following this procedure may not be simple, i.e., may contain self-loops due to the pairing of two half-edges of i , and multi-edges due to the existence of more than one pairing between two given nodes. Conditional on the multigraph $G^*(n, (d_i)_1^n)$ being a simple graph, we obtain a uniformly distributed random graph with the given degree sequence, which we denote by $G(n, (d_i)_1^n)$, [7].

For $r \in \mathbb{N}$, let $u_k^{(n)} = |\{i, d_i = k\}|$ be the number of vertices of degree k and $m^{(n)}$ be the total degree defined by

$$m^{(n)} = \sum_{i=1}^n d_i = \sum_{k \geq 1} k u_k^{(n)}.$$

We assume that the sequence $(d_i)_1^n$ satisfies the following regularity conditions analogous to the ones introduced in [8]:

- (i) There exists a distribution $p = \{p_k\}_{k=0}^\infty$ such that $u_k^{(n)}/n \rightarrow p_k$ for every $k \geq 0$ as $n \rightarrow \infty$;
- (ii) $\lambda := \sum_{k \geq 0} k p_k \in (0, \infty)$;
- (iii) $\sum_{i=1}^n d_i^2 = O(n)$,
- (iv) for some $\tau > 0$, $\Delta_n := \max_{i \in V} d_i = O(n^{1/2-\tau})$.

Let d_{\min} denote the minimum degree of the graph $G(n, (d_i)_1^n)$, i.e. $d_{\min} = \min\{k \mid u_k^{(n)} > 0\}$. We define $q = \{q_k\}_{k=0}^\infty$ the size-biased probability mass function corresponding to p , by $q_k = \frac{(k+1)p_{k+1}}{\lambda}$, and let ν denote its mean, i.e.

$$\nu = \sum_{k=0}^\infty k q_k \in (0, \infty).$$

The condition $\nu > 1$ is equivalent to the existence of a giant component in the configuration model and we will consider this case only in what follows.

The typical graph distance has been studied in this framework by Van der Hofstad, Hooghiemstra and Van Mieghem:

Theorem 8. [9] *For a and b chosen uniformly at random in the giant component of $G(n, (d_i)_1^n)$, we have*

$$\frac{\text{dist}(a, b)}{\log n} \xrightarrow{p} \frac{1}{\log \nu}.$$

Let $G_q(z)$ be the probability generating function of $\{q_k\}_{k=0}^\infty$, i.e. $G_q(z) = \sum_{k=0}^\infty q_k z^k$. The extinction probability of the branching process with offspring distribution q , is the smallest solution β in $[0, 1]$ of the fixed point equation $\beta = G_q(\beta)$. We also define $\beta_* = G'_q(\beta)$.

The diameter has been studied by Fernholz and Ramachandran:

Theorem 9. [6] *We have*

$$\begin{aligned} \frac{\text{diam}(G(n, (d_i)_1^n))}{\log n} &\xrightarrow{p} \frac{1}{\log \nu} - \frac{\mathbf{1}(d_{\min} = 2)}{\log q_1} - 2 \frac{\mathbf{1}(d_{\min} = 1)}{\log \beta_*} \\ &= \frac{1}{\log \nu} + \frac{2 - \mathbf{1}(d_{\min} \geq 2) - \mathbf{1}(d_{\min} \geq 3)}{|\log \beta_*|} \end{aligned}$$

Concerning weighted distances, the typical distance has been investigated by Bhamidi, Van der Hofstad and Hooghiemstra:

Theorem 10. [3] *For a and b chosen uniformly at random in $G(n, (d_i)_1^n)$ with $d_{\min} \geq 2$ and with i.i.d. exponential 1 weights on its edges, we have*

$$\frac{\text{dist}_w(a, b)}{\log n} \xrightarrow{p} \frac{1}{\nu - 1}.$$

A similar result holds for Erdős-Rényi random graph [2].

Our main result gives the asymptotic behavior of the diameter of $G(n, (d_i)_1^n)$ (extending the result of [5] dealing with regular graphs):

Theorem 11. *Consider a random graph $G(n, (d_i)_1^n)$ with i.i.d. exponential 1 weights on its edges. Then*

$$\frac{\text{diam}_w(G(n, (d_i)_1^n))}{\log n} \xrightarrow{p} \frac{1}{\nu - 1} + 2 \frac{\mathbf{1}(d_{\min} \geq 3)}{d_{\min}} + \frac{\mathbf{1}(d_{\min} = 2)}{(1 - q_1)} + 2 \frac{\mathbf{1}(d_{\min} = 1)}{1 - \beta_*}.$$

These results allow us to analyze an asynchronous randomized broadcast algorithm for random regular graphs. Our results show that the asynchronous version of the algorithm performs better than its synchronized version [1]: in the large size limit of the graph, it will reach the whole network faster even if the local dynamics are similar on average.

REFERENCES

- [1] H. Amini, M. Draief, and M. Lelarge, *Flooding and Diameter in Weighted Random Graphs*.
- [2] S. Bhamidi and G. H. R. van der Hofstad, *First passage percolation on the Erdős-Rényi random graph*, Arxiv (2010).

- [3] S. Bhamidi, R. van der Hofstad, and G. Hooghiemstra, *First passage percolation on random graphs with finite mean degrees*. To appear in *Annals of Applied probability* (2009).
- [4] B. Bollobás, *Random Graphs*, Cambridge University Press (2001).
- [5] J. Ding, J. H. Kim, E. Lubetzky, and Y. Peres, *Diameters in supercritical random graphs via first passage percolation* (2009). <http://arxiv.org/abs/0906.1840>
- [6] D. Fernholz and V. Ramachandran, *The diameter of sparse random graphs*, *Random Structures and Algorithms* **31** (2007), (4) 482–516.
- [7] S. Janson, *The probability that a random multigraph is simple*, *Combinatorics, Probability and Computing* **18** (2009), (1-2) 205–225.
- [8] M. Molloy and B. Reed, *The size of the giant component of a random graph with a given degree sequence*, *Combinatorics, Probability and Computing* **7** (1998), 295–305 (1998).
- [9] R. van der Hofstad, G. Hooghiemstra, and P. V. Mieghem, *Distances in random graphs with finite variance degrees*, *Random Structures and Algorithms* **27** (2005), 76–123.

Markov couplings, stochastic orders, and stochastic relations

LASSE LESKELÄ

Stochastic monotonicity is a key methodology for approximating random variables and processes whose distributions are hard to compute explicitly [7, 8, 10]. For example, the stochastic ordering of two stationary Markov jump processes can be established without the explicit knowledge of the stationary distributions, by verifying that the generators of the systems preserve the order (Massey [6], Whitt [11]).

In this talk I will discuss my recent paper [5] which introduces a new definition of stochastic relation. This definition is motivated by the observation that for the stochastic ordering of two stationary random dynamical systems, it suffices to show that the generators of the systems preserve some, not necessarily reflexive or transitive, subrelation of the order. A relation between two measurable spaces S_1 and S_2 is a measurable set $R \subset S_1 \times S_2$, and we denote

$$x \sim y$$

if $(x, y) \in R$. The stochastic relation generated by R is defined in terms of couplings. Recall that a coupling of two random elements X and Y is a bivariate random element (\hat{X}, \hat{Y}) such that \hat{X} is distributed according to X and \hat{Y} is distributed according to Y . For random variables X and Y we denote

$$X \sim_{\text{st}} Y$$

if there exists a coupling (\hat{X}, \hat{Y}) of X and Y such that $\hat{X} \sim \hat{Y}$ with probability one. The relation \sim_{st} is called the *stochastic relation* generated by \sim . The stochastic relation \leq_{st} generated by a partial order \leq on a space S corresponds to the usual strong stochastic order on S .

One of the main results of the paper is the following functional characterization [5, Theorem 2.5], which can be regarded as a measurable version of Strassen's coupling theorem [9]. The conjugate of a set $B \subset S_1$ is defined by $B^{\rightarrow} = \cup_{x_1 \in B} \{x_2 \in S_2 : x_1 \sim x_2\}$ (see Figure 1) and the conjugate of a function f on S_1 by $f^{\rightarrow}(x_2) = \sup_{x_1: x_1 \sim x_2} f(x_1)$.




FIGURE 1. Relational conjugate of a set $B \subset S_1$.

Theorem. For any stochastic relation generated by a closed relation R between two complete separable metric spaces S_1 and S_2 , the following are equivalent:

- (1) $X \sim_{\text{st}} Y$,
- (2) $\mathbb{P}(X \in B) \leq \mathbb{P}(Y \in B^{\rightarrow})$ for all compact $B \subset S_1$,
- (3) $\mathbb{E} f(X) \leq \mathbb{E} f^{\rightarrow}(Y)$ for all upper semicontinuous compactly supported $f : S_1 \rightarrow \mathbb{R}_+$.

The following results are obtained by applying the above theorem:

- (1) necessary and sufficient conditions for a pair of probability kernels to preserve a stochastic relation, and
- (2) an iterative algorithm for computing the maximal subrelation of a given relation that is preserved by a pair of probability kernels (see [4] for a numerical implementation).

Coupling results based on this theory have a broad range of applications. In this talk I will discuss applications to multilayer call centers [3] and load balancing in parallel queues [5, Section 5.3]. Besides queueing systems, such coupling techniques have also been seen to be valuable when analyzing interacting particle systems (e.g. [1]).

REFERENCES

- [1] R. Delgado, F. J. López, and G. Sanz, *Local conditions for the stochastic comparison of particle systems*, Adv. Appl. Probab. **36** (2004), 1252–1277.
- [2] A. Economou, *Necessary and sufficient conditions for the stochastic comparison of Jackson networks*, Probab. Eng. Inform. Sci. **17** (2003), 143–151.
- [3] M. Jonckheere and L. Leskelä, *Stochastic bounds for two-layer loss systems*, Stoch. Models **24** (2008), (4) 583–603.
- [4] L. Leskelä, *Computational methods for stochastic relations and Markovian couplings*. In Proc. 4th International Workshop on Tools for Solving Structured Markov Chains (SMC-Tools) (2009). <http://arxiv.org/abs/0906.0153>

- [5] L. Leskelä, *Stochastic relations of random variables and processes*, J. Theor. Probab. **23** (2010), (2) 523–546. <http://arxiv.org/abs/0806.3562>.
- [6] W. A. Massey, *Stochastic orderings for Markov processes on partially ordered spaces*, Math. Oper. Res. **12** (1987), (2) 350–367.
- [7] A. Müller and D. Stoyan, *Comparison Methods for Stochastic Models and Risks*, Wiley (2002).
- [8] M. Shaked and J. G. Shanthikumar, *Stochastic Orders*, Springer Series in Statistics. Springer, New York (2007).
- [9] V. Strassen, *The existence of probability measures with given marginals*, Ann. Math. Statist. **36** (1965), (2) 423–439.
- [10] R. Szekli, *Stochastic Ordering and Dependence in Applied Probability*, Springer, New York (1995).
- [11] W. Whitt, *Stochastic comparisons for non-Markov processes*, Math. Oper. Res. **11** (1986), (4) 608–618.

Nonlinear large deviations: Exact Asymptotics

DAVID McDONALD

(joint work with Ivo Adan, Bob Foley)

Rare events are important in fields such as finance, inventory, reliability, manufacturing and telecommunications, and often correspond to catastrophes. Simulation is difficult precisely because the rare events are rare. Here we study large deviations that are more likely to result from the accumulation of many slightly unusual steps (light-tailed) rather than a single giant step (heavy-tailed). The aim is to develop methodology for determining how rare a particular large deviation is; i.e. the probability of the rare event, and how the system evolves as the large deviation develops. In particular, we are primarily interested in the unscaled process—not the fluid limit and not the log probability of the rare event. Among other things, these two pieces of information can help a designer decide whether, and how, to redesign a system to make a particular large deviation less likely.

We are particularly interested in analyzing large deviations where a large deviation in one part of the system was triggered by an earlier large deviation in another part of the network. In such situations, the most probable path (fluid limit) changes direction.

If the earlier large deviation is not a problem or even desirable, the system manager can be sideswiped by the later catastrophic large deviations. To illustrate this large deviation we propose a simplistic model called Subprime. We wish the reader to know that the model is not serious, but the message is.

Suppose new mortgages in good standing arrive at rate λ . Mortgages in good standing are paid off at rate $\mu_2 r_{20}$ or go into arrears at rate $\mu_2 r_{21}$. Mortgages in arrears convert to mortgages in good standing at rate μ_1 although this conversion may be the result of a short sale. Let x represent the number of mortgages in arrears; i.e. the number of customers at node 1. Let y be the number of mortgages in good standing; i.e. the number of customers at node 2. Let X denote the Markov process with states (x, y) . The flow rates along the arcs when both queues




FIGURE 1. Transition rates for the Subprime network

are positive is shown in Fig. 1. The steady state of this Jackson network is

$$\pi(x, y) = \left(1 - \frac{\lambda_1}{\mu_1}\right) \left(\frac{\lambda_1}{\mu_1}\right)^x \left(1 - \frac{\lambda_2}{\mu_2}\right) \left(\frac{\lambda_2}{\mu_2}\right)^y \text{ where } \lambda_1 = \frac{\lambda r_{21}}{\mu_1 r_{20}}, \lambda_2 = \frac{\lambda}{\mu_2 r_{20}}$$

and where we assume $\lambda_1 < \mu_1$ and $\lambda_2 < \mu_2$ for stability.

Note that large deviations in the first coordinate correspond to a large number of mortgages in arrears, which would be undesirable. Large deviations in the second coordinate correspond to a large number of mortgages in good standing, which would seem to be desirable.

This simplistic model is simply a two node Jackson network with a single server at each node, and the stationary distribution π is known, which will make it easier to illustrate several points. However, our methodology is designed for situations where π is unknown, and we address that later.

Consider the following folk theorem: The large deviation path from the origin to a rare event can be obtained by observing the time reversal from the rare event back to the origin.

The folk theorem has been applied rigorously for a few networks [2, 8, 10]. Assume that we have an irreducible, discrete time Markov process X on a countable state space with a stationary distribution π . Let X^* denote the time reversal of X . Let x_0 be a state with $\pi(x_0) > 0$ and x_ℓ be a sequence of states with $\pi(x_\ell) \rightarrow 0$ as $\ell \rightarrow \infty$. Let E_ℓ^* denote the set of paths that start at x_ℓ and go to x_0 without returning to x_ℓ . Let F_ℓ^* be a set of paths such that

$$(1) \quad \lim_{\ell \rightarrow \infty} P_{x_\ell}(X^* \in F_\ell^*) = 1$$

Theorem 12 (Large Deviation Folk Theorem). *Let E_ℓ denote the time reversal of paths in E_ℓ^* . Let F_ℓ denote the time reversal of paths in $F_\ell^* \cap E_\ell^*$; i.e. determine the first time the path in $F_\ell^* \cap E_\ell^*$ hits x_0 and take this as time 0 for the time reversal. Then*

$$(2) \quad \lim_{\ell \rightarrow \infty} P_{x_0}(X \in F_\ell | X \in E_\ell) = 1$$

To apply this theorem to our Subprime example first define X to be the discrete time chain whose uniformization is the Subprime process (also denoted by X). Let $x_\ell = (\ell, 0)$ and $x_0 = (0, 0)$. Using Kurtz’s theorem [7] and Corollary 1 in [4], there

is a fluid limit $p^*(t); 0 \leq t \leq T$ that is piecewise linear connecting the points:

$$\begin{aligned}
 p^*(0) &= (1, 0) \\
 p^*\left(\frac{1}{(\mu_1 - \mu_2 r_{21})}\right) &= \left(0, \frac{\mu_1 + \lambda - \mu_2}{\mu_1 - \mu_2 r_{21}}\right) \text{ and} \\
 p^*(T) &= (0, 0)
 \end{aligned}$$

with $T = (\mu_2 r_{20} - \lambda)^{-1}$ and that

$$(3) \quad \lim_{\ell \rightarrow \infty} P_{(\ell, 0)} \left\{ \sup_{0 \leq t \leq T} \left| \frac{X^*[\ell t]}{\ell} - p^*(t) \right| \leq \epsilon \right\} = 1$$

for all $\epsilon > 0$. Consequently, define $F_\ell^* = \{X^* : \sup_{0 \leq t \leq T} |\frac{X^*[\ell t]}{\ell} - p^*(t)| \leq \epsilon\}$. Thus, defining $p(t) = p^*(T - t)$ we have

$$\lim_{\ell \rightarrow \infty} P_{(0, 0)} \left\{ \sup_{0 \leq t \leq T} \left| \frac{X[\ell t]}{\ell} - p(t) \right| \leq \epsilon \mid X \text{ hits } (\ell, 0) \text{ before } (0, 0) \right\} = 1$$

Proof of Folk Theorem.

$$\begin{aligned}
 P_{x_0}(X \in F_\ell \mid X \in E_\ell) &= \frac{\pi(x_0)P_{x_0}\{X \in F_\ell\}}{\pi(x_0)P_{x_0}(E_\ell)} \\
 &= \frac{\pi(x_\ell)P_{x_\ell}\{X^* \in F_\ell^* \cap E_\ell^*\}}{\pi(x_\ell)P_{x_\ell}(X^* \in E_\ell^*)} \\
 &\rightarrow 1 \text{ because of (1).}
 \end{aligned}$$

□

We can apply the folk theorem to the Subprime network even if we didn't know the steady state since the Subprime chain is time reversible. More generally if we know the steady state π we can calculate the time reversal. We discuss what to do if we don't later. Assume $\mu_1 + \lambda > \mu_2$ then by stability

$$\mu_1 > \mu_2 - \lambda > \mu_2 - \mu_2 r_{20} = \mu_2 r_{21}$$

i.e. the chain drifts northwest. Therefore the time reversal starting from $(\ell, 0)$ drifts northwest until it hits the y axis. Then it jitters down the y axis. Hence the large deviation path from $(0, 0)$ has two segments. The first is the time reversal of the subprime chain with respect to the measure

$$\left(1 - \frac{\mu_2 r_{21}}{\mu_1}\right) \left(\frac{\mu_2 r_{21}}{\mu_1}\right)^x$$

which is equivalent to running the film of the jitter down the y axis backwards. This results in a jitter up the y axis to a point at roughly $\ell(\mu_1 + \lambda - \mu_2)/(\mu_1 - \mu_2 r_{21})$. The network seen by the loan manager is as in Figure 2. The loan manager sees mortgages in good standing increasing at rate $\mu_2 r_{20} - \lambda$ while mortgages in arrears are arriving at rate $\mu_2 r_{21}$ which convert to mortgages in good standing at rate μ_1 .

The second segment starts at roughly $\ell(\mu_1 + \lambda - \mu_2)/(\mu_1 - \mu_2 r_{21})$. The large deviation heads south-east after reversing the arrows of the Subprime chain. The loan manager sees Figure 3.




FIGURE 2. Transition rates for the large deviation path as it jitters up the y axis




FIGURE 3. Transition rates for the large deviation path as loans turn bad too fast to convert them.

Lots of good mortgages turn bad at rate μ_1 and only convert to good mortgages at rate $\mu_2 r_{21}$. The mortgage manager has been sideswiped by an excessive number of new good mortgages turning bad.

Unfortunately we usually cannot directly calculate the time reversal since the steady state is unknown. Suppose the government increases lending when the number of good loans gets too small. Then the product formula fails. Can we still use the folk theorem? The answer is yes by using an approximate time reversal.

Let $\mathbb{Z}_+ = \{0, 1, 2, \dots\}$ and $\mathbb{Z} = \{\dots, -1, 0, 1, \dots\}$. Uniformize the CTMC to a DTMC with transition kernel K on the state space $\mathcal{S} = \mathbb{Z}_+ \times \mathbb{Z}_+$. Assume that there is a Markov additive kernel K^∞ on $\mathcal{S}^\infty = \mathbb{Z} \times \mathbb{Z}_+$; i.e.

$$K^\infty((x_0, y_0); (x_1, y_1)) = K^\infty((0, y_0); (x_1 - x_0, y_1))$$

and a region \blacktriangle ($\blacktriangle = \{(x, y) : x \leq 0\}$ in the subprime example) such that $K(i; j) = K^\infty(i; j)$ for points $i = (x_0, y_0), j = (x_1, y_1)$ where $i \in \mathcal{S}$ and $j \in \blacktriangle^c$. Let $\Delta = \mathcal{S} \cap \blacktriangle$.

Now find a left invariant measure ψ for K^∞ which has the same asymptotics as π on Δ as $y \rightarrow \infty$. For the subprime network $\psi(x, y) = (\lambda_1/\mu_1)^x \cdot (\lambda_2/\mu_2)^y$; i.e. ψ is proportional to π . The approximate time reversal

$$\overleftarrow{K}(i; j) = \frac{\psi(j)}{\psi(i)} K^\infty(j; i)$$

is defined for i in $\mathcal{S} \setminus \Delta$ and j in \mathcal{S} . Let \overleftarrow{X} represent the approximate time reversal.

Let $h(x, y) = \pi(x, y)/\psi(x, y)$ Then h is harmonic for \overleftarrow{K} in $\mathcal{S} \setminus \Delta$. If $i \in \mathcal{S} \setminus \Delta$ then

$$\begin{aligned} \sum_{j \in \mathcal{S}} \overleftarrow{K}(i; j)h(j) &= \frac{1}{\psi(i)} \sum_{j \in \mathcal{S}} \frac{\psi(i)}{\psi(j)} \overleftarrow{K}(i; j)\pi(j) \\ &= \frac{1}{\psi(i)} \sum_{j \in \mathcal{S}} \pi(j)K^\infty(j; i) = \frac{1}{\psi(i)} \sum_{j \in \mathcal{S}} \pi(j)K(j; i) = \frac{\pi(i)}{\psi(i)} = h(i) \end{aligned}$$

We now prove a neat lemma. Assume \overleftarrow{X} drifts toward Δ (west in the Subprime example). Let $\overleftarrow{\tau}$ be the first time the x -component of \overleftarrow{X} is in Δ . By assumption $\overleftarrow{\tau} < \infty$ a.s. $h(0, y) \rightarrow C$ as $y \rightarrow \infty$ by construction and is therefore bounded. It follows that $h(\overleftarrow{X}(n \wedge \overleftarrow{\tau}))$ is a uniformly integrable martingale and $h(x, y) = E_{(x,y)}h(\overleftarrow{X}(\overleftarrow{\tau}))$; i.e.

$$\frac{\pi(x, y)}{\psi(x, y)} = E_{(x,y)}h(\overleftarrow{X}(\overleftarrow{\tau})).$$

The uniform integrability follows as below plus the structure of \overleftarrow{X} .

To check that $h(\overleftarrow{X}(n \wedge \overleftarrow{\tau}))$ is a uniformly integrable martingale we use Proposition IV-3-16 in [9]. To do so we must check

$$(4) \quad \int_{\tau < \infty} |h(\overleftarrow{X}(\overleftarrow{\tau}))|dP < \infty$$

$$(5) \quad \lim_{n \rightarrow \infty} E_{((x,y))} \left(\chi_{\{\overleftarrow{\tau} > n\}} |h(\overleftarrow{X}(n))| \right) < \infty$$

(4) follows immediately from the fact that h tends to a constant for $y \rightarrow \infty$. (5) follows by the observation that by time reversal with respect to π :

$$\begin{aligned} \lim_{n \rightarrow \infty} E_{((x,y))} \left(\chi_{\{\overleftarrow{\tau} > n\}} |h(\overleftarrow{X}(n))| \right) &= \frac{1}{\psi(x, y)} \lim_{n \rightarrow \infty} E_\pi \left(\chi_{\{\tau > n; X(n) = (x, y)\}} \right) \\ &\leq \frac{1}{\psi(x, y)} \lim_{n \rightarrow \infty} P_\pi (\tau > n) = 0 \end{aligned}$$

where τ is the first time X is in Δ .

If \overleftarrow{X} drifts northwest then as (x, y) tends to ∞ the y component of $\overleftarrow{X}(\overleftarrow{\tau})$ tends to ∞ so

$$\lim_{x \rightarrow \infty} \frac{\pi(x, y)}{\psi(x, y)} = C.$$

In the subprime case this just gets us that

$$\lim_{(x,y) \rightarrow \infty} \frac{\pi(x, y)}{(\lambda_1/\mu_1)^x \cdot (\lambda_2/\mu_2)^y} = C;$$

but of course we already knew that. As we will see below, if the government gets involved, ψ is still asymptotically of product form

$$\psi(x, y) = (\lambda_1/\mu_1)^x \cdot H(y)(\lambda_2/\mu_2)^y$$

where $H(y) \rightarrow 1$ as $y \rightarrow \infty$.

We now sketch how to construct an invariant measure. Suppose K^∞ is close to random walk kernel \mathbb{R} on $\mathbb{Z} \times \mathbb{Z}$. Suppose $\pi(0, y)$ decays like $\exp(-\theta_2 y)$ with $\exp(-\theta_2) = \lambda r_{21} / \mu_2 r_{20}$ as it does for the Subprime example since the exact asymptotics can be obtained by an analysis of the jitter case up the y axis. Find a pair (θ_1, θ_2) such that $\eta(x, y) = \exp(-\theta_1 x - \theta_2 y)$ is left invariant for \mathbb{R} and such that $\overleftarrow{\mathbb{R}}$ drifts west where $\overleftarrow{\mathbb{R}}$ is the reversal with respect to η . Define the kernel

$$\mathbb{Q}_{y_1, y_2} = \sum_x e^{-(\theta_1 x + \theta_2 (y_2 - y_1))} K^\infty((x, y_2); (0, y_1)).$$

In general \mathbb{Q} will not be substochastic. Then by the proposition below there exists a harmonic function $H(y)$ for \mathbb{Q} such that $H(\ell) \rightarrow 1$ as $\ell \rightarrow \infty$.

Since H is harmonic for \mathbb{Q} :

$$\begin{aligned} H(y_1) &= \sum_{y_2 \geq 0} \mathbb{Q}_{y_1, y_2} H(y_2) \\ &= \sum_{y_2 \geq 0} \sum_x e^{-(\theta_1 x + \theta_2 (y_2 - y_1))} K^\infty((x, y_2); (0, y_1)) H(y_2) \end{aligned}$$

i.e.

$$H(y_1) e^{-\theta_2 y_1} = \sum_{y_2 \geq 0} \sum_x e^{-(\theta_1 x + \theta_2 y_2)} H(y_2) K^\infty((x, y_2); (0, y_1)).$$

Hence $\psi(x, y) = H(y) e^{-(\theta_1 x + \theta_2 y)}$ is left invariant for K^∞ and has the required asymptotics in y .

H will exist under the following conditions:

- \mathbb{Q} is *not* substochastic but still has radius of convergence $R > 1$.
- \mathbb{Q} is “close” to a matrix \mathbb{J} where

$$\mathbb{J}_{y_1, y_2} = \sum_x e^{-(\theta_1 x + \theta_2 (y_2 - y_1))} \mathbb{R}^\infty((x, y_2); (0, y_1)).$$

By *close*, we mean that $\sum_{i \geq 0} \sum_{j \geq 0} |\mathbb{Q}_{ij} - \mathbb{J}_{ij}| < \infty$.

- Moreover

- (6) $\mathbb{J}_{ij} = p_{j-i}$ for $i, j \in \mathbb{Z}_+$ where
- (7) $\{p_x, x \in \mathbb{Z}\}$ is a proper probability distribution
- (8) there exists $x < 0 < y$ with $p_x > 0$ and $p_y > 0$,
- (9) $d \doteq \sum_{x \in \mathbb{Z}} x p_x$ exists and is finite, and
- (10) $d \neq 0$

The following Proposition is proved in [5]: Let $\mathbb{A} \doteq \mathbb{Q} \wedge \mathbb{J}$, which is also substochastic. Suppose that \mathbb{Q} is an irreducible, nonnegative matrix with convergence parameter $R > 1$ and that \mathbb{Q} is close to \mathbb{J} where \mathbb{J} satisfies (6) through (10). If the

drift defined (9) satisfies $d > 0$, then

$$(11) \quad \mathbb{A}^n \vec{1} \downarrow \vec{a} = \mathbb{A} \vec{a} > 0$$

$$(12) \quad \vec{a}_y = \vec{1}$$

$$(13) \quad \mathbb{Q}^n \vec{a} \uparrow H = \mathbb{Q}H > 0, \text{ and}$$

$$(14) \quad \lim_{y \rightarrow \infty} H(y) = 1.$$

Thus, H is a nonnegative harmonic function with $H(\ell) \rightarrow 1$.

We conclude by remarking that exact asymptotics with non-linear path behaviour was first described in [3] although the conditions to check that this behaviour applies are not given explicitly. Moreover the work [6] is closely related to ours. We believe the idea of an approximate time reversal has potential for calculating exact asymptotics for a variety of networks and for simulation. After all it is a lot easier to start a simulation at the rare event $(\ell, 0)$ and let it wander back to $(0, 0)$.

REFERENCES

- [1] I.J.B.F Adan, R.D. Foley, D.R. McDonald, *Exact asymptotics for the stationary distribution of a Markov chain: A production model*. Queueing Systems: Theory and Applications, **62** (2009), 311–344.
- [2] V. Anantharam, P. Heidelberger, P. Tsoucas. *Analysis of Rare Events in Continuous Time Markov Chains via Time Reversal and Fluid Approximation*. Research Report, RC 16280, Computer Science Division, IBM T. J. Watson Center, Yorktown Heights, New York, Nov. (1990).
- [3] A.A. Borovkov and A.A. Mogul'skii. *Limit Theorems in the Boundary Hitting Problem for a Multidimensional Random Walk*. Siberian Mathematical Journal **4** (2001), 245-270.
- [4] R. Foley, D. McDonald, *Join the Shortest Queue: Stability and Exact Asymptotics*, Annals of Applied Probability, 11. **11** (2001). 569–607.
- [5] Foley, R.D., McDonald, D.R. *Constructing a harmonic function for an irreducible non-negative matrix with convergence parameter $R > 1$* . Accepted subject to revisions London Mathematical Society.
- [6] Irina Ignatiouk-Robert and Christophe Loree, *Martin boundary of a killed random walk on a quadrant* Ann. Probab. **38** (2010), 1106-1142.
- [7] Kurtz, T.G., *Strong approximation theorems for density dependent Markov chains*, Stoch. Proc. Appl. **6** (1978), 223-240.
- [8] K. Majewski, K. Ramanan, *How large queue lengths build up in a Jackson network*. preprint (2008).
- [9] J. Neveu, *Discrete-Parameter Martingales*, North Holland Publishing Co. (1975).
- [10] V. Nicola, T. Zaburdenko, *Efficient Importance Sampling Heuristics for the Simulation of Population Overflow in Jackson Networks*. ACM Transactions on Modeling and Computer Simulation, **17** (2007)

On the stability of population processes

ILKKA NORROS

(joint work with Hannu Reittu)

Consider the following idealized peer-to-peer file-sharing system. The file consists of k chunks that can be transferred separately. There is one permanent seed peer possessing all chunks. Every non-seed peer makes, with $Exp(1)$ distributed intervals, a contact to a randomly selected peer (seed included). If the contacted peer has a chunk that the contacting peer is missing, a download can happen instantaneously. New peers arrive according to the $Poisson(\lambda)$ process. When a non-seed peer has collected all chunks, it leaves the system immediately. Our aim is to find algorithms for the contacting peer that provide stable operation with any input rate λ despite the ‘non-altruism’ of the peers. For simplicity, we focus on the ‘toy model’ case $k = 2$.

The non-seed peer population is partitioned into classes according to the chunks a peer possesses, and often also according to its state in the algorithm it follows. This type of models for file-sharing systems were introduced by Massoulié and Vojnovic [2], making them instances of *population processes* as studied by Kurtz [1]. The models in question are Markovian queueing networks whose transition intensities scale linearly when both the input rates and node occupancies are multiplied by $N \rightarrow \infty$. In this scaling, they approach deterministic dynamic systems described by autonomous ordinary differential equations. Kurtz’ little monograph [1] gives limit theorems in both LLN and CLT regimes concerning the behaviour of population processes on fixed time intervals. However, he did not study the connection between the stability of a population process and the stability of the dynamical system obtained as the ‘large system limit’ of the former. Still today, little seems to be known about this question (unless something important has escaped our attention). Nevertheless, together with each system, identified by the chunk selection algorithm, we analyse its large system limit, since studying the stability of the dynamical system is easier than that of the stochastic system and can yield valuable insights concerning the latter.

We first show that a ‘Plain Random Contact’ system with the trivial greedy algorithm is unstable (when $\lambda > 1$): one of the chunks is eventually present as at most one non-seed copy, while the number of copies of the other grows linearly. This result has been recently proven by Hajek and Zhu also in the case of many chunks — see Hajek’s talk in this workshop.

As our first attempt to alleviate the ‘rare chunk phenomenon’, we consider the ‘Deterministic First Chunk’ system where an incoming peer first decides randomly which chunk it wants to download first, and downloads nothing in its encounters until it finds a peer possessing that chunk. Interestingly, the large system limit of this system is unstable, but the stochastic one is conjectured, by simulation evidence, to be stable despite of showing rather strong irregular fluctuations. If this conjecture is true, we thus have an example showing that the instability of a population process cannot be inferred from the instability of its large system limit.

We are still missing an example of a population process with ‘stability discrepancy’ in the opposite direction.

Arrivals to the Plain Random Contact system are similar to those of the classical Pólya urn, where upon picking a ball a new one with same color is added. To keep the system in balance, the arrivals should rather be those of Bernard Friedman’s urn, where the added ball is of the opposite colour. We first tried to improve the performance of the Deterministic First Chunk system by letting the arriving peer make one test contact and decide to download the opposite chunk first. The resulting system seemed to be stable, but showed strong, rather regular oscillations. The next algorithm finally provided the desired performance. In [4] it was generalized for many chunks as follows:

Enforced Friedman algorithm: Make contact to three peers simultaneously. If there are ‘minority chunks’ possessed by exactly one of the three, download a randomly chosen minority chunk, otherwise download nothing.

We proved that the corresponding dynamical system is globally asymptotically stable in the two-chunk case. However, proving the stability of the stochastic Enforced Friedman system has remained a challenge for us.

REFERENCES

- [1] T. Kurtz, *Approximation of population processes*, CBMS-NSF Regional Conference Series in Applied Mathematics **36** (1981).
- [2] L. Massoulie and M. Vojnovic, *Coupon replication systems*, ACM SIGMETRICS Performance Evaluation Review **33** (2005), 2–13.
- [3] I. Norros, H. Reittu and T. Eirola, *On the stability of two-chunk file-sharing systems*, submitted for publication. Preprint: arXiv:0910.5577.
- [4] H. Reittu, *A stable random-contact algorithm for peer-to-peer file sharing*, Proceedings of IWSOS 2009, Lecture Notes in Computer Science **5918**, Springer (2009). DOI: 10.1007/978-3-642-10865-5_16.

Short and Long-term behavior of Stochastic Systems through Mean Field Asymptotics

ALEXANDRE PROUTIERE

(joint work with Charles Bordenave, David McDonald)

In statistical physics, Mean Field Asymptotics (MFA) have been traditionally used to analyze the behavior of stochastic systems comprising a large number of interacting particles. By letting the number of particles tends to infinity, the system dynamics are captured in most cases by a set of deterministic ODEs or PDEs, see e.g. Boltzmann equation [2] describing the evolution of the statistical distribution of particles in rare gases. Beyond statistical physics, MFA have been successfully applied in many contexts, including population dynamics [6], genetics [7], econophysics [4], computer networks [3]. MFA are primarily used to approximate dynamics of large stochastic systems, but have also been instrumental in the analysis of the underlying limiting ODEs or PDEs, for instance in solving

Cauchy problems related to these limiting equations [10, 11]. Probabilistic methods have been developed to justify the convergence of stochastic systems towards their deterministic counterpart in the mean field regime. The most common argument to derive mean field convergence involves the construction of appropriate martingales associated with the Markov processes describing the evolution of the empirical distribution of the particle system, see [6, 5]. A different approach to justify convergence to the mean field regime relies on the concept of propagation of chaos originated with Kac's Markovian models of gas dynamics [9], see [12] for a comprehensive review.

Using the aforementioned standard approaches, convergence towards the mean field regime is limited to short time horizons only. More precisely, if $f^N(t)$ denotes the empirical distribution at time t of the stochastic system with N particles, we obtain

$\lim_{N \rightarrow \infty} \mathbb{P}[\sup_{t \leq \tau_N} \|f^N(t) - f(t)\|] = 0$, where $f(t)$ denotes the particle distribution in the mean field regime, i.e. a solution of the limiting ODE or PDE, at time t , and where τ_N is a time horizon that should remain $o(\log(N))$. In particular, MFA cannot say anything about the long-term behavior of stochastic systems. We may construct examples of stochastic systems with a unique steady-state distribution, but whose corresponding limiting mean field ODE or PDE does not converge to any equilibrium point. In some cases, see e.g. [8], it is possible to show convergence to the mean field regime of stochastic systems in steady state. To do so, one has to assume that the stochastic system with N particles possesses a unique steady state distribution π^N , that the sequence of probability measures π^N is tight, and that the limiting ODE has a unique attractive equilibrium point $f(\infty)$. Under such strong assumptions, we have $\lim_{N \rightarrow \infty} \pi^N = \lim_{N \rightarrow \infty} \lim_{t \rightarrow \infty} f^N(t) = \lim_{t \rightarrow \infty} \lim_{N \rightarrow \infty} f^N(t) = f(\infty)$.

In general, MFA asymptotics cannot help predicting ergodicity conditions of the initial stochastic systems. One may wonder under which conditions the following equivalence holds: the system in the mean field regime is globally stable if and only if the initial stochastic systems are ergodic when N is large enough. We investigate this equivalence in the case of systems of coupled queues. We consider a system of N infinite queues, and sharing a single resource (e.g. a radio channel) in a decentralized manner using slotted Aloha [1], a classical random multi-access protocol. Packets arrive in queue i according to some Markovian process of intensity λ_i , and in each slot, in queue i is non-empty, a packet is transmitted with fixed probability p_i . When two or more queues transmit simultaneously, none of the transmissions is successful. Deriving ergodicity conditions of the Markov chain describing the evolution of queue lengths has long eluded probabilists, and only a few simple systems (comprising very few queues, $N = 2$ or 3) can be analyzed, see e.g. [13]. When N grows large, the transient behavior of the system can be approximated using MFA. In the mean field regime, the distribution of queue lengths satisfies a set of ODEs. By studying the conditions under which these ODEs form a globally stable dynamical system, we hope to get a good approximation of the ergodicity

conditions of the original stochastic systems. Indeed we show that the ergodicity conditions converge when N tends to ∞ to the global stability conditions of the corresponding ODEs in the mean field regime. The proof of this equivalence essentially relies on some structural and monotonicity properties of the coupled queue systems.

REFERENCES

- [1] N. Abramson, *The ALOHA System - Another alternative for computer communications*, Proc. of AFIPS Conf., Fall Joint Computer Conference (1970).
- [2] L. Boltzmann, *Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen*, Sitzungsberichte der Akademie der Wissenschaften **66** (1872), 275–370.
- [3] C. Bordenave, D. McDonald and A. Proutiere, *Random multi-access algorithms, a mean field analysis*, Proc. 43th Allerton conference (2005).
- [4] S. Cordier, L. Pareschi and G. Toscani, *On a kinetic model for a simple market economy*, Journal of Statistical Physics **120** (2005).
- [5] R. Darling and J. Norris, *Differential equation approximations for Markov chains*, Probability Surveys **5** (2008) 37–79.
- [6] S.N. Eiler and T.G. Kurtz, *Markov Process: Characterization and Convergence*, John Wiley and Sons Inc., New York (1986).
- [7] W.H. Fleming and M. Viot, *Some measure-valued markov processes in population genetics theory*, Indiana Univ. Math. J. **28** (1979), 817–843.
- [8] C. Graham, *Chaoticity on path space for a queueing network with selection of the shortest queue among several*, Journal of Applied Probability **37** (2000), 198–211.
- [9] M. Kac, *Foundations of kinetic theory*, Proc. of the Third Berkeley Symposium on Mathematical Statistics and Probability **3** (1956), 171–197, University of California Press.
- [10] T.G. Kurtz, *Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes*, Journal of Applied Probability **7** (1970), 49–58.
- [11] S. Meleard, *Probabilistic interpretation and approximations of some Boltzmann equations*, Soc. Mat. Mexicana, Mexico (1998), 1–64.
- [12] A.S. Sznitman, *Propagation of chaos*, In: Ecole d’ete de probabilites de Saint Flour XIX, Springer Lectures Notes in Mathematics, **1464** (1991), 165–251 .
- [13] W. Szpankowski, *Stability conditions for some multiqueue distributed systems: Buffered random access systems*, Advances in Applied Probability **26** (1994), 498–515.

Counting and sampling matchings on infinite graphs

JUSTIN SALEZ

(joint work with Charles Bordenave, Marc Lelarge)

A matching on a graph $G = (V, E)$ is a subset of pairwise disjoint edges $M \subseteq E$. In the so-called monomer-dimer model [12, 16], a random matching \mathcal{M}_G^z is drawn on G according to the Gibbs-Boltzmann distribution with parameter $z > 0$:

$$(1) \quad \mathbb{P}(\mathcal{M}_G^z = M) = \frac{z^{|M|}}{P_G(z)}, \quad \text{where} \quad P_G(z) = \sum_M z^{|M|}.$$

In particular, \mathcal{M}_G^1 is just a uniformly chosen matching on G , and \mathcal{M}_G^z converges in distribution to a uniformly chosen maximum matching as $z \rightarrow \infty$. The normalizing factor $P_G(z)$ is known as the matching polynomial of G ; it has been extensively

studied in algebraic graph theory [11, 15]. The marginal distribution viewed by a given vertex $v \in V$ can be described by the $[0, 1]$ -valued number

$$(2) \quad \pi_z(G, v) = \mathbb{P}(v \text{ is unmatched in } \mathcal{M}_G^z).$$

This defines a $[0, 1]$ -valued functional π_z on the class of finite rooted graphs, and it happens to satisfy the following local recursion [12, 10, 3] :

$$(3) \quad \pi_z(G, v) = \left(1 + z \sum_{w \sim v} \pi_z(G - v, w) \right)^{-1}.$$

Our first result states that this formula can be used to define π_z on infinite graphs, in a way that continuously extends the finite case definition. Continuity is here understood with respect to the topology of local convergence introduced by Benjamini and Schramm [4], and studied further by Aldous and Steele [2], and Aldous and Lyons [1]. It provides a natural and convenient framework for studying asymptotical properties of large diluted graphs [5, 9, 14, 8]. Although our results hold in slightly larger generality, for the sake of clarity we here restrain ourselves to the class \mathcal{G} of all locally finite rooted graphs with at most exponential growth, which contains in particular all graphs with bounded degree. We let \mathcal{H} be the space of analytic functions on $\mathbb{C} \setminus (-\infty, 0]$, with uniform convergence on compact sets.

Lemma 1 (The fundamental local recursion).

- (1) For any $z > 0$, the recursion (3) admits a unique solution $\pi_z: \mathcal{G} \rightarrow [0, 1]$.
- (2) For any fixed $(G, v) \in \mathcal{G}$, $z \mapsto \pi_z(G, v)$ extends analytically to $\mathbb{C} \setminus (-\infty, 0]$.
- (3) The resulting mapping $(G, v) \in \mathcal{G} \mapsto \pi_{(\cdot)}(G, v) \in \mathcal{H}$ is continuous.

This has strong implications for random matchings, which we now list. The first one is the existence and uniqueness of the monomer-dimer model on any graph $G \in \mathcal{G}$. On the regular lattice $G = \mathbb{Z}^d$, this is well known [12, 6].

Theorem 13 (Random matchings on infinite graphs). *Let $G \in \mathcal{G}$ and $z > 0$. As M ranges over all finite matchings of G , the cylinder-event marginals*

$$\mathbb{P}(M \subseteq \mathcal{M}_G^z) := z^{|M|} \prod_{k=1}^{2|M|} \pi_z(G \setminus \{v_1, \dots, v_{k-1}\}, v_k),$$

are consistent and independent of the ordering $v_1, \dots, v_{2|M|}$ of the vertices matched by M . They thus determine the law of a random matching \mathcal{M}_G^z on G , which continuously extends the finite case definition (in the sense of local weak convergence).

Theorem 14 (Scaling limit for the matching polynomial). *Let $(G_n)_{n \geq 1}$ be a sequence of finite graphs converging locally weakly under uniform rooting to a random rooted graph (G, v) in \mathcal{G} . Then,*

$$\frac{1}{|V_n|} \log P_{G_n}(z) \xrightarrow{n \rightarrow \infty} \int_0^z \frac{1 - \mathbb{E}[\pi_s(G, v)]}{s} ds$$

analytically on $\mathbb{C} \setminus (-\infty, 0]$, under the only requirement that the degree of a uniformly chosen vertex in G_n is uniformly integrable as $n \rightarrow \infty$.

Motivated by the asymptotics of the matching number

$$\nu(G) = \deg P_G = \max\{|M|; M \text{ matching on } G\},$$

we now let $z \rightarrow \infty$. The limit $\pi_\infty(G, v) = \lim_{z \rightarrow \infty} \pi_z(G, v)$ turns out to exist for any $(G, v) \in \mathcal{G}$. Moreover, $\pi_\infty : \mathcal{G} \rightarrow [0, 1]$ is the largest solution to the recursion

$$(4) \quad \pi_\infty(G, u) = \left(1 + \sum_{v \sim u} \left(\sum_{w \sim v} \pi_\infty[G - u - v, w] \right)^{-1} \right)^{-1},$$

with the conventions $0^{-1} = \infty, \infty^{-1} = 0$. Unfortunately, several distinct solutions may exist, and this results in the non-continuity of π_∞ with respect to local convergence. However, continuity turns out to hold “on average”, i.e. when looking at a uniformly chosen vertex v . Consequently, we obtain :

Theorem 15 (Scaling limit for the matching number). *Let $(G_n)_{n \geq 1}$ be a sequence of finite graphs satisfying $|E_n| = O(|V_n|)$ and converging locally weakly under uniform rooting to a random rooted graph (G, v) in \mathcal{G} . Then,*

$$\frac{\nu(G_n)}{|V_n|} \xrightarrow{n \rightarrow \infty} \frac{1 - \mathbb{E}[\pi_\infty(G, v)]}{2}.$$

In particular, in the case where the limit (G, v) is a unimodular Galton-Watson tree with degree distribution π , (4) can be solved explicitly [7] and we obtain:

$$\frac{\nu(G_n)}{|V_n|} \xrightarrow{n \rightarrow \infty} \min_{x \in [0,1]} \left\{ 1 - \frac{1}{2}x\phi'(1-x) - \frac{1}{2}\phi(1-x) - \frac{1}{2}\phi\left(1 - \frac{\phi'(1-x)}{\phi'(1)}\right) \right\},$$

where $\phi(x) = \sum \pi_n x^n$ is the generating function of π .

For example, the random weak limit of a sequence of Erdős-Rényi graphs with parameter c/n on n vertices is a unimodular Galton-Watson tree with degree distribution $\pi = \text{Poisson}(c)$. In this case, $\phi(x) = \exp(cx - c)$, and we recover precisely the well known formula by Karp and Sipser [13]. More generally, a unimodular GWT with degree distribution π (where π has finite mean) is a random rooted tree obtained by a Galton-Watson branching process where the root has offspring distribution π and all other genitors have offspring distribution $\hat{\pi}$, where

$$(5) \quad \forall n \in \mathbb{N}, \hat{\pi}_n = (n + 1)\pi_{n+1} / \sum_{k \geq 1} k\pi_k.$$

Although this tilted offspring distribution might seem unnatural at first sight, unimodular GWTs are in fact the most reasonable random weak limit for a sequence of locally tree-like random graphs (see [1]).

REFERENCES

[1] D. Aldous and R. Lyons, *Processes on unimodular random networks*, Electronic Journal of Probability **12** (2007), 1454–1508.
 [2] D. Aldous and J. M. Steele, *The objective method: probabilistic combinatorial optimization and local weak convergence*. In Probability on discrete structures, volume 110 of Encyclopaedia Math. Sci., pages 1–72. Springer, Berlin (2004).

- [3] M. Bayati, D. Gamarnik, D. Katz, C. Nair, P. Tetali *Simple deterministic approximation algorithms for counting matchings*, Proc. of the 39th Annual Symposium on Theory of Computing (2007).
- [4] I. Benjamini and O. Schramm, *Recurrence of distributional limits of finite planar graphs*, Electron. J. Probab. **6** (2001), no. 23, 13 pp. (electronic).
- [5] I. Benjamini, O. Schramm, and A. Shapira, *Every minor-closed property of sparse graphs is testable*. In R. E. Ladner and C. Dwork, editors, STOC, pages 393–402, ACM (2008).
- [6] J. Berg, *On the absence of phase transition in the monomer-dimer model*, Technical report, Amsterdam, The Netherlands (1998).
- [7] C. Bordenave, M. Lelarge, and J. Salez, *The rank of diluted random graphs*, to appear in Annals of probability.
- [8] C. Borgs, J. Chayes, J. Kahn, L. Lovász, and W. Notes, *Left and right convergence of graphs with bounded degree*, Arxiv preprint arXiv:1002.0115 (2010).
- [9] G. Elek, *On limits of finite graphs*, Combinatorica **27** (2007), (4) 503–507.
- [10] C. D. Godsil, *Matchings and walks in graphs*, J. Graph Theory **5** (1981), (3) 285–297.
- [11] C. D. Godsil, *Algebraic combinatorics*, Chapman and Hall Mathematics Series. Chapman & Hall, New York (1993).
- [12] O. J. Heilmann and E. H. Lieb, *Theory of monomer-dimer systems*, Comm. Math. Phys. **25** (1972), 190–232.
- [13] R. Karp and M. Sipser, *Maximum matchings in sparse random graphs*, Proc. of the Twenty-second Annual Symposium on Foundations of Computer Science, IEEE 364–375 (1981).
- [14] R. Lyons, *Asymptotic enumeration of spanning trees*, Combin. Probab. Comput. **14** (2005), 491–522.
- [15] L. Lovász and M. D. Plummer, *Matching theory*, AMS Chelsea Publishing, Providence, RI (2009). Corrected reprint of the 1986 original [MR0859549].
- [16] L. Zdeborová and M. Mézard, *The number of matchings in random graph*, Journal of Statistical Mechanics (5)(2006), P05003.

Distribution of cost functionals in spatial network models: Scaling limits and Monte-Carlo methods

VOLKER SCHMIDT

(joint work with Frank Fleischer, Catherine Gloaguen and Florian Voss)

We consider stationary point processes on random geometric graphs and Palm calculus in order to analyse the distribution of shortest path lengths and other cost functionals in spatial stochastic networks [3,6,8]. More precisely, we model location and type of network stations by marked point processes on the edge sets of random tessellations. Our model can be applied e.g. to telecommunication networks involving road systems [4]. Based on this modelling approach it is then possible to define cost functionals like typical connection lengths which can be analysed by Monte Carlo simulation and, simultaneously, by asymptotic methods.

In particular, we show how the distribution of such cost functionals can be estimated using efficient simulation algorithms for the typical cell of random tessellations [1,2,5,7]. Furthermore, we derive scaling limits for the distribution of these cost functionals for infinitely dense and sparse networks, respectively. In both cases the limit distribution is a parametric distribution whose parameters are known or depend on the underlying tessellation [5,9].

Finally, both approaches can be combined in order to obtain parametric distributions for connection lengths in spatial networks which can be used e.g. to efficiently analyse existing and future telecommunication networks [4].

REFERENCES

- [1] F. Fleischer, C. Gloaguen, H. Schmidt, V. Schmidt and F. Schweiggert, *Simulation algorithm of typical modulated Poisson-Voronoi cells and application to telecommunication network modelling.*, Japan Journal of Industrial and Applied Mathematics 25 (2008), 305-330.
- [2] F. Fleischer, C. Gloaguen, V. Schmidt and F. Voss, *Simulation of the typical Poisson-Voronoi-Cox-Voronoi cell.*, Journal of Statistical Computation and Simulation 79 (2009), 939-957.
- [3] C. Gloaguen, F. Fleischer, H. Schmidt and V. Schmidt, *Analysis of shortest paths and subscriber line lengths in telecommunication access networks.*, Networks and Spatial Economics 10 (2010), 15-47.
- [4] C. Gloaguen, F. Voss and V. Schmidt, *Parametric distributions of connection lengths for the efficient analysis of fixed access network.*, Annals of Telecommunications 66 (2011) (to appear)
- [5] F. Voss, C. Gloaguen, F. Fleischer, V. Schmidt, *Distributional properties of Euclidean distances in wireless networks involving road systems.*, IEEE Journal on Selected Areas in Communication 27 (2009), 1047-1055.
- [6] F. Voss, C. Gloaguen, F. Fleischer and V. Schmidt, *Densities of shortest path lengths in spatial stochastic networks.*, Stochastic Models 27 (2011) (to appear).
- [7] F. Voss, C. Gloaguen and V. Schmidt, *Palm calculus for stationary Cox processes on iterated random tessellations.*, Proceedings of the 5th Workshop on Spatial Stochastic Models for Wireless Networks (SpaSWiN), Seoul 2009
- [8] F. Voss, C. Gloaguen and V. Schmidt, *Capacity distributions in spatial stochastic models for telecommunication networks.*, Image Analysis and Stereology 28 (2009), 155-163.
- [9] F. Voss, C. Gloaguen and V. Schmidt, *Scaling limits for shortest path lengths along the edges of stationary tessellations.* Advances in Applied Probability 42 (2010) (to appear)

Three metrics for stochastic networks: capacity, queue-size and complexity

DEVAVRAT SHAH

(joint work with David Tse, John Tsitsiklis, Damon Wischik, Yuan Zhong)

1. MODEL, A CLASS OF POLICIES

We consider a collection of N queues operating in discrete time, indexed by $\tau \in \{0, 1, \dots\}$: $Q_n(\tau)$ be work in queue n , $1 \leq n \leq N$ at time τ , $\mathbf{Q}(\tau) = [Q_n(\tau)]$ be the vector of queue-sizes and initially it is $\mathbf{Q}(0)$. Let $A_n(\tau)$ be the total amount of work arriving to queue n and $B_n(\tau)$ be the cumulative potential service provided to queue n , up to time τ respectively, with $\mathbf{A}(0) = \mathbf{B}(0) = \mathbf{0} = [0]$. We consider single-hop network (for simplicity of exposition). Let $d\mathbf{A}(\tau) = \mathbf{A}(\tau + 1) - \mathbf{A}(\tau)$ and $d\mathbf{B}(\tau) = \mathbf{B}(\tau + 1) - \mathbf{B}(\tau)$. Then basic Lindley recursion is

$$(1) \quad \mathbf{Q}(\tau + 1) = [\mathbf{Q}(\tau) - d\mathbf{B}(\tau)]^+ + d\mathbf{A}(\tau)$$

where the $[\cdot]^+$ is taken componentwise. The fundamental ‘switched network’ constraint is that there is some finite set $\mathcal{S} \subset \mathbb{R}_+^N$ such that

$$(2) \quad d\mathbf{B}(\tau) \in \mathcal{S} \quad \text{for all } \tau.$$

For simplicity, we shall consider $\mathcal{S} \subset \{0, 1\}^N$. We will refer to $\boldsymbol{\pi} \in \mathcal{S}$ as a schedule, and \mathcal{S} as the set of allowed schedules. The departure from queue n up to time τ is

$$(3) \quad D_n(\tau) = \sum_{s=0}^{\tau} dB_n(s) \mathbf{1}_{\{Q_n(s) > 0\}},$$

where $\mathbf{1}_{\{x\}} = 1$ if $x = \text{true}$ and 0 otherwise.

A policy needs to choose schedule $d\mathbf{B}(\tau) \in \mathcal{S}$ in each time slot. The specific class of policies of interest are the so called maximum weight (MW) introduced (in basic version) by Tassiulas and Ephremides [10]. In the basic version, the schedule is chosen as follows (ties broken randomly or as per a fixed rule):

$$(4) \quad d\mathbf{B}(\tau) \in \operatorname{argmax}_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \mathbf{Q}(\tau),$$

where we use notation: $\mathbf{u} \cdot \mathbf{v} = \sum_n u_n v_n$ for $\mathbf{u}, \mathbf{v} \in \mathbb{R}^N$. More generally, given an increasing function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ with $f(0) = 0$, $f(x) \rightarrow \infty$ as $x \rightarrow \infty$, the MW- f policy chooses schedule

$$(5) \quad d\mathbf{B}(\tau) \in \operatorname{argmax}_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot f(\mathbf{Q}(\tau)).$$

The specific class of policies of interest are those induced by $f(x) = x^\alpha$ for $\alpha \in (0, 1)$; we shall denote this policy as MW- α policy. We shall assume the following.

Assumption 1. *We assume that $\mathcal{S} \subset \{0, 1\}^N$ is monotone: if $\boldsymbol{\pi} \in \mathcal{S}$, then for any $\boldsymbol{\rho} \in \{0, 1\}^N$ with $\rho_n \in \{0, \pi_n\}$, $\boldsymbol{\rho} \in \mathcal{S}$.*

To evaluate the performance of the system, we model the uncertainty in the system (in form of arrivals) by means of appropriate stochastic model. Specifically, we assume the following.

Assumption 2. *Let $\mathbf{A}(\cdot)$ be a random process with stationary increments. Let there be $\boldsymbol{\lambda} \in \mathbb{R}_+^N$ so that for any $r \in \mathbb{Z}_+$*

$$(6) \quad \mathbb{P}\left(\sup_{\tau \leq r} \frac{1}{r} \left| \mathbf{A}(\tau) - \boldsymbol{\lambda} \tau \right| \geq \varepsilon_r\right) \leq \delta_r,$$

where ε_r, δ_r go to 0 as $r \rightarrow \infty$.

The specific instances of $\mathbf{A}(\cdot)$ that satisfy Assumption 2 and we shall consider are: (a) $A_n(\cdot)$ is Bernoulli process with parameter $\lambda_n \in [0, 1]$ independent across n ; (b) $A_n(\cdot)$ is Poisson process with parameter $\lambda_n \in \mathbb{R}_+$ independent across n .

2. RESULTS

We describe the known results about interplay between three performance metrics of scheduling policies: (a) capacity, (b) average queue-size as well as exponential tail bounds on queue-sizes, and (c) complexity of implementation of the policy.

2.1. Capacity. This is about how well the network resource is utilized by the scheduling policy. While there are various ways to understand this, we shall define the effective resource utilization by studying the net departure rate. Specifically, define

$$(7) \quad d_n = \liminf_{\tau} \frac{1}{\tau} D_n(\tau).$$

The net departure rate is $\bar{d} = \sum_n d_n$. An algorithm is called optimal in terms of capacity if for any given $\lambda \in \mathbb{R}_+^N$, the induced net departure rate $\bar{d} = \bar{d}(\lambda)$ is maximal possible with probability 1. In [9], the following is established (by means of fluid model):

1. For any policy with probability 1, $\bar{d}(\lambda) \leq (\sum_n \lambda_n) - \bar{q}(\lambda)$ where $\bar{q}(\lambda)$ is the value of the following optimization problem:

$$(8) \quad \begin{aligned} & \text{minimize } \sum_n r_n \quad \text{over } \mathbf{r} \in \mathbb{R}_+^N \\ & \text{subject to } \zeta \cdot \mathbf{r} \geq \zeta \cdot \lambda - 1, \quad \forall \zeta \in \mathcal{D}, \end{aligned}$$

where $\mathcal{D} = \{\zeta \in \mathbb{R}_+^N : \zeta \cdot \pi \leq 1 \quad \forall \pi \in \mathcal{S}\}$.

2. The net departure rate induced by MW- α is, with probability 1, $\bar{d}_\alpha(\lambda) = (\sum_n \lambda_n) - \bar{q}_\alpha(\lambda)$ where $\bar{q}_\alpha(\lambda)$ is the value of the following optimization problem:

$$(9) \quad \begin{aligned} & \text{minimize } \sum_n r_n^{1+\alpha} \quad \text{over } \mathbf{r} \in \mathbb{R}_+^N \\ & \text{subject to } \zeta \cdot \mathbf{r} \geq \zeta \cdot \lambda - 1, \quad \forall \zeta \in \mathcal{D}, \end{aligned}$$

where $\mathcal{D} = \{\zeta \in \mathbb{R}_+^N : \zeta \cdot \pi \leq 1 \quad \forall \pi \in \mathcal{S}\}$.

3. Therefore, for any $\lambda \in \mathbb{R}_+^N$,

$$\lim_{\alpha \downarrow 0} \bar{q}_\alpha(\lambda) = \left(\sum_n \lambda_n \right) - \bar{q}(\lambda).$$

The above results implies that the MW- α policy is asymptotically optimal in terms of maximizing the departure rate as $\alpha \rightarrow 0^+$.

2.2. Queue-size on average. Here interest is in understanding the behavior of average queue-size induced by policy when the notion of stationary distribution as well as average queue-size (with respect to it) are well defined. To state results (in a clean form) about average queue-sizes, we shall assume arrival process to be independent Poisson. Given rate vector $\lambda \in \mathbb{R}_+^N$, define the load $\rho(\lambda)$ as

$$(10) \quad \text{minimize } \sum_{\pi \in \mathcal{S}} \alpha_\pi \quad \text{over } \alpha_\pi \geq 0, \quad \forall \pi \in \mathcal{S} \text{ subject to } \lambda \leq \sum_{\pi} \alpha_\pi \pi.$$

Clearly, $\mathbf{Q}(\cdot)$ is a Markov chain under MW- α policy when arrival process is Poisson. When $\rho(\lambda) < 1$ then it is positive recurrent with well defined, unique stationary distribution and $\rho(\lambda) < 1$ is necessary for this. For any MW- α policy with $\alpha > 0$, $\mathbb{E}[\sum_n Q_n]$ with respect to this stationary distribution is well defined (see [6]). Using the Foster-Lyapunov moment bound, it follows that for any system considered

here, under the MW-1 policy, the stationary average queue-size is bounded above as

$$(11) \quad \mathbb{E} \left[\sum_n Q_n \right] \leq C \frac{n^2}{1 - \rho(\boldsymbol{\lambda})},$$

where $C > 0$ is a universal constant. In [8] the following \mathcal{S} dependent lower bound on the net average queue-size under any policy is established:

$$(12) \quad \begin{aligned} & \text{minimize } \sum_n r_n \quad \text{over } \mathbf{r} \in \mathbb{R}_+^N \\ & \text{subject to } \zeta \cdot \mathbf{r} \geq \frac{\boldsymbol{\lambda} \cdot \zeta^2}{2(1 - \boldsymbol{\lambda} \cdot \zeta)}, \quad \forall \zeta \in \mathcal{D}, \end{aligned}$$

where $\mathcal{D} = \{\zeta \in \mathbb{R}_+^N : \zeta \cdot \boldsymbol{\pi} \leq 1 \quad \forall \boldsymbol{\pi} \in \mathcal{S}\}$.

The overall challenge is to identify $\chi_1(\mathcal{S})$ that depends on \mathcal{S} so that under the best policy with well defined stationary distribution for any $\rho \in (0, 1)$,

$$(13) \quad \sup_{\boldsymbol{\lambda}: \rho(\boldsymbol{\lambda}) = \rho} \mathbb{E} \left[\sum_n Q_n \right] = \Theta \left(\frac{\chi_1(\mathcal{S})}{1 - \rho} \right).$$

Clearly, the above statement implicitly conjectures existences of such $\chi_1(\mathcal{S})$ that determines the optimal (up to universal constants) dependence of average queue-size on scheduling constraints \mathcal{S} for any such system. The (11) and (12) provide upper and lower bound on such quantity. In [8] it is shown that the $\chi_1(\mathcal{S})$ (up to constants) is characterized by the lower bound (12) for the (general enough) instance of switched network model induced by wireless network model (see Section 2.4 for description of this model) with *regular enough* constraint graph structure describing \mathcal{S} . Indeed, we believe that bound of (12) is reasonably accurate.

2.3. Exponential tail probability. Here interest is in understanding the further detailed behavior of the stationary distribution of queue-size when it exists. Specifically, we shall assume that arrival process is Bernoulli with rate vector $\boldsymbol{\lambda} \in [0, 1]^N$ so that $\rho(\boldsymbol{\lambda}) < 1$. Again, in this setup the $\mathbf{Q}(\cdot)$ forms a Markov chain under MW- α policy for any $\alpha > 0$. It is positive recurrent with unique stationary distribution as long as $\rho(\boldsymbol{\lambda}) < 1$. In [6], it is shown that with respect to this stationary distribution, under MW- α policy

$$(14) \quad \limsup_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P} \left(\|\mathbf{Q}\|_{1+\alpha} \geq x \right) \leq -(1 - \rho(\boldsymbol{\lambda})) N^{-\frac{2+\alpha}{1+\alpha}}.$$

Further, for any policy it can be shown that for each $\rho \in (0, 1)$ there exists $\boldsymbol{\lambda}$ such that $\rho(\boldsymbol{\lambda}) = \rho$ and

$$(15) \quad \liminf_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P} \left(\|\mathbf{Q}\|_{1+\alpha} \geq x \right) \geq -(1 - \rho) N^{\frac{\alpha}{1+\alpha}}.$$

Similar to the average queue-size, the overall challenge is to identify $\chi_2(\mathcal{S}, \alpha)$ that depends on \mathcal{S} (and $\alpha > 0$) so that under the best policy with well defined

stationary distribution for any $\rho \in (0, 1)$,

$$(16) \quad \begin{aligned} & \sup_{\lambda: \rho(\lambda)=\rho} \liminf_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P} \left(\|\mathbf{Q}\|_{1+\alpha} \geq x \right) \geq -C_1(1-\rho)\chi_2(\mathcal{S}, \alpha) \\ & \sup_{\lambda: \rho(\lambda)=\rho} \limsup_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P} \left(\|\mathbf{Q}\|_{1+\alpha} \geq x \right) \leq -C_2(1-\rho)\chi_2(\mathcal{S}, \alpha), \end{aligned}$$

where $0 < C_1 \leq C_2$ are some universal constants. Clearly, the above statement implicitly conjectures existences of such $\chi_2(\mathcal{S}, \alpha)$ that determines the optimal (up to universal constants) dependence of exponential tail bound of $1 + \alpha$ norm of queue-sizes with respect to the stationary distribution. Indeed, the (14) and (15) provide upper and lower bound on such quantity.

2.4. Complexity. The summary of results thus far suggest that the MW- α class of policies are reasonably good: MW-0⁺ policy is (near) optimal in terms of capacity, it has good scaling in terms of average queue-size and very good behavior in terms of exponential tail probability. And most importantly, it is a *myopic* policy, i.e. it uses only current network state (queue-sizes) to make the scheduling decision. Therefore, it is only natural to wonder: *do we have a reasonable answer for the problem of designing scheduling policies for switched networks?*

To answer this question, it is important to understand the context where such policies will be used. Specifically, in the application domain like communication networks, the implementation of policies is highly constrained: (i) it should be very simple in terms of computation and data structure requirement so that it can operate at very high aggregate bandwidth (say making decision once every few nano seconds !) with existing limitations on the memory bandwidth with limited hardware requirement at low power; (ii) it should be preferably iterative and distributed so as to allow for architectures that are scalable. Such stringent constraints immediately lead to the following questions: is it possible to implement the MW policy with above requirements for arbitrary scheduling set \mathcal{S} ? if not, how does the performance suffer from implementation limitation ?

Before we start answering these questions rigorously, it is important to lay down constraints explicitly. The main problem is that the precise constraints are problem dependent and the above stated constraints, while provide flavor of the problem, do not define them explicitly. For that reason, collectively over the past decade or so, researchers have focused on few concrete applications where the constraints are quite clear. One such important problem has been the design of medium access in wireless networks. Formally, there are N wireless transmitters or nodes or queues, denoted by $V = \{1, \dots, N\}$. Let $E \subset V \times V$ represent the scheduling or transmission constraints: if $(i, j) \in E$ then no two nodes/queues can transmit/be served simultaneously. Thus, the space of all possible schedules is

$$\mathcal{S} = \{\mathbf{x} \in \{0, 1\}^N : x_i + x_j \leq 1 \forall (i, j) \in E\}.$$

The implementation of any policy for making transmission decision is highly constrained: (C1) each node must make decision (using few, preferably constant number of computations) to transmit or not in each time slot on its own based on local observations/historical information possibly summarized through limited data structure; (C2) nodes have (so called delayed carrier sense) information about transmission of other nodes in prior time slots (and hence if they transmitted, they know whether their transmissions were successful or not). Indeed, the popular implemented protocols that go by the names *Aloha*, *Random Backoff*, etc. satisfy C1 and C2.

Now in this setup, the maximum weight policy requires solving the so called maximum weight independent set problem in graph G (with weights dependent on queue-size). In general, this is hard problem and hence it is not clear if one can implement the maximum weight policy as is with stringent constraints like C1 and C2. This led to a long line of research finally resulting in its resolution very recently by Shah and Shin [3, 4] and Jiang and Walrand [2, 1]: they provide implementation so that the resulting policy keeps network Markov process positive recurrent as long as the system is underloaded (i.e. $\rho(\boldsymbol{\lambda}) < 1$). However, the performance of such implementation in terms of queue-size scaling is not clear.

As mentioned above, the MW-1 policy for example, induced average queue-size with respect to stationary distribution that scales as $O(N^2/(1 - \rho(\boldsymbol{\lambda})))$ when $\rho(\boldsymbol{\lambda}) < 1$. And the best bound one can obtain on average queue-size under the above mentioned implementations are scaling super polynomially in N , at the best.

In summary, MW is capacity achieving, provides small (polynomially scaling in N) average queue-size but may not yield to implementation that has (time-)computational complexity scaling polynomial in N (C1, C2 is too much to ask for !). The implementations of [3, 4, 2, 1] have polynomial complexity (actually, they satisfy C1, C2), is capacity achieving (positive recurrence when $\rho(\boldsymbol{\lambda}) < 1$) but may not yield small (polynomially scaling in N) average queue-sizes. This leads to the following fundamental question: *is it possible to have an implementation that is (i) capacity achieving, (ii) has small (polynomial in N) queue-sizes, and (iii) has polynomial in N computational complexity?*

Recently, Shah, Tse and Tsitsiklis [5] established that the answer to this question is *NO* for arbitrary G (assuming the standard computational hypothesis). The precise statement of the result, details as well as extensions can be found in [5].

3. CONCLUSION

In conclusion, we would like to take note of the key results and open problems surveyed here. Specifically, we surveyed performance of scheduling policies for switched networks in the context of three metrics: capacity, queue-size and complexity. The MW- α policy as $\alpha \downarrow 0$ achieves the maximal capacity, defined in terms of the effective net departure rate for any switched network. This is a strong evidence of the qualitative conjecture about goodness of MW-0⁺ policy

in the MW class of algorithms originally introduced by Shah and Wischik [7]. The MW family of policies induce stationary average queue-sizes when system is underloaded: they are reasonably good and close to fundamental, policy independent, lower bounds on them; similar qualitative result holds for the exponential tail probability. Though MW policies are myopic and have very good performance in terms of capacity and queue-size, in general their implementation can require solving computationally hard problems. Indeed, for general switched network, it is not possible to have any policy that is computationally simple and has good performance in terms of queue-size in complexity.

The questions of determining exact scaling of average queue-size (i.e. $\chi_1(\mathcal{S})$) and exponential tail probability (i.e. $\chi_2(\mathcal{S})$) remain important quest going forward – resolution of these will definitely advance the frontiers of methods for large complex queuing networks. And the grand challenge is to develop framework to understand the *pareto* boundary reflecting interplay between these three performance metrics.

REFERENCES

- [1] L. Jiang, D. Shah, J. Shin and J. Walrand, *Distributed random access algorithm: Scheduling and congestion control*.
- [2] L. Jiang and J. Walrand, *A distributed csma algorithm for throughput and utility maximization in wireless networks*. In Proceedings of 46th Allerton Conference on Communication, Control, and Computing, Urbana-Champaign, IL (2008).
- [3] S. Rajagopalan, D. Shah and J. Shin, *A network adiabatic theorem: an efficient randomized protocol for contention resolution*. In ACM Sigmetrics/Performance (2009).
- [4] D. Shah and J. Shin, *Randomized Scheduling Algorithm for Queueing Networks*, Arxiv preprint arXiv:0908.3670 (2009).
- [5] D. Shah, D.N.C. Tse and J.N. Tsitsiklis, *Hardness of low delay scheduling*. Submitted to IEEE Transaction on Information Theory (2009).
- [6] D. Shah, J.N. Tsitsiklis and Y. Zhong, *Qualitative properties of α -weighted policies*. In ACM Sigmetrics/Performance (2010).
- [7] D. Shah and D.J. Wischik, *Switched networks with maximum weight policies: Fluid approximation and multiplicative state space collapse*. Submitted to Annals of Applied Probability (2006, 2007).
- [8] D. Shah and D.J. Wischik, *Lower bound and optimality in switched networks*. In Allerton conference on communication, control and computing (2008).
- [9] D. Shah and D.J. Wischik, *Fluid model of switched networks in overload*. Submitted to QUESTA (2010).
- [10] L. Tassiulas and A. Ephremides, *Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks*, IEEE Transactions on Automatic Control **37** (1992), 1936–1948.

Scaling properties of queues under back-pressure algorithm

ALEXANDER STOLYAR

The back-pressure (MaxWeight) algorithm [3] has well-known optimality properties, in particular throughput maximization. It also has a well-known drawback: queues may scale "badly" as network size increases, both in terms of queue sizes and the number of queues. The issue of back-pressure (BP) algorithm bad scaling has been addressed in [1, 2], where a generic "remedy" is proposed: the algorithm

is "run" on virtual queues, as opposed to real ones, while real traffic uses conventional, per-next-hop FIFO queues. Paper [2] proves that this scheme guarantees stability, and demonstrates by simulation that it indeed dramatically reduces queue sizes and traffic delays. Work [2] also proves, in a very general setting, that under BP the queue length accumulation along each flow route grows at most quadratically with the max route lengths in the network. However, besides that upper bound, the question of how the queues scale under BP in specific network settings remains largely open.

It is of interest to understand fundamental scaling properties of BP algorithms. In this talk, we address this problem for a simple system – a single, rate λ , Poisson flow served by N queues in tandem, each is a $M/M/1$ system with unit service rate. We show that, perhaps somewhat surprisingly, the scaling behavior of BP is *not always* bad. Namely, if load is below some critical level, $\lambda < 1/4$ for our model, all queues remain uniformly stochastically bounded for all N . (In fact, we show that the stochastic bound has exponentially decaying tail.) When $\lambda > 1/4$, the queues increase to infinity with N .

We show that the asymptotic behavior of the sequence of finite systems, as $N \rightarrow \infty$, is essentially reduced to the behavior of the system with *infinite* number of queues in tandem. Such infinite-tandem model is within the framework of interacting particle systems – methods and results of the corresponding theory serve as our main analysis tools. In particular, the criticality of load $1/4$ is closely related to the fact that $1/4$ is the maximum possible flux (average flow rate) of a stationary *totally asymmetric simple exclusion process* (TASEP). In the subcritical case, $\lambda < 1/4$, a stationary regime exists such that only a finite (random) number of "left-most" queues can be greater than 1, the rest of the queues have at most one customer and the process "there" behaves like TASEP. In the supercritical case, $\lambda > 1/4$, each queue grows without bound with time; we prove this by contradiction: if this fact were not true, we could construct a stationary ("one-sided") TASEP with flux exceeding $1/4$ (which is impossible).

REFERENCES

- [1] L. Bui, R. Srikant, A.L. Stolyar, *Optimal Resource Allocation for Multicast Flows in Multihop Wireless Networks*, Philosophical Transactions of the Royal Society A **366** (2008), 2059–2074.
- [2] L. Bui, R. Srikant, A. L. Stolyar, *Novel Architectures and Algorithms for Delay Reduction in Back-pressure Scheduling and Routing*. In: Proceeding of INFOCOM'2009 Mini-conference (2009).
- [3] L. Tassiulas and A. Ephremides, *Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks*, IEEE Transactions on Automatic Control **37** (1992), 1936–1948.

Remarks on convergence to stationarity for unreliable queueing networks

RYSZARD SZEKLI

(joint work with Hans Daduna, Pawel Lorek)

We study unreliable Jackson networks, and give for them necessary and sufficient conditions for geometric speed of convergence to stationarity via spectral gap. Next we give bounds on spectral gap for networks and propose a strong duality for Möbius monotone chains.

Let $\tilde{\mathbf{X}}$ be unreliable Jackson network process with the infinitesimal generator $\tilde{\mathbf{Q}}$. Suppose that $\tilde{\mathbf{Q}}$ is bounded and the minimal service intensity $\underline{\mu} > 0$. If the routing matrix R is reversible, and regular then the following cond. are equivalent:

- $Gap(\tilde{\mathbf{Q}}) > 0$
- $(\pi_i)_{i \geq 0}$ is strongly light-tailed, $i = 1, \dots, J$
- $\inf_k \frac{1}{1 + \sum_{j=k+1}^{\infty} \frac{\lambda^{j-k}}{\mu_i^{(k+1)} \dots \mu_i^{(j)}}} > 0, i = 1, \dots, J.$

Consequently we have as a corollary the following result. Let $\tilde{\mathbf{X}}$ be unreliable Jackson network with generator $\tilde{\mathbf{Q}}$, and the corresponding transition semigroup (\tilde{P}_t) . Suppose the routing matrix R is reversible, and $R^k > 0$ for some $k \geq 1$.

If $\mu_i(n)/\lambda$ represents equilibrium rate of strongly light-tailed distribution for each $i = 1, \dots, J$, (which is then π_i) then equivalently

(i) for all $f \in L^2(\tilde{\mathbb{E}}, \pi)$

$$\|\tilde{P}_t f - \pi(f)\|_2 \leq e^{-Gap(\tilde{\mathbf{Q}})t} \|f - \pi(f)\|_2, \quad t > 0,$$

(ii) for each $\mathbf{e} \in \tilde{\mathbb{E}}$ there exists $C(\mathbf{e}) > 0$ such that

$$\|\delta_{\mathbf{e}} \tilde{P}_t - \pi\|_{tv} \leq C(\mathbf{e}) e^{-Gap(\tilde{\mathbf{Q}})t}, \quad t > 0.$$

Next we give one step correlation formulas for networks (see for details [1]), for example for arbitrary real functions f, g , and ξ which is the probability solution of the extended traffic equation, and for

$$W^{g,f}(\mathbf{n}) = [g(\mathbf{n} + e_i) f(\mathbf{n} + e_j)]_{i,j=0,1,\dots,J}.$$

we have

$$\begin{aligned} &\langle f, Q^{\mathbf{X}} g \rangle_{\pi} - \langle f, Q^{\hat{\mathbf{X}}} g \rangle_{\pi} = \\ &(\lambda/\xi_0) E_{\pi} \left(tr(W^{g,f}(X_t) \cdot diag(\xi) \cdot (R - \hat{R})) \right). \end{aligned}$$

This formula allows for comparison of spectral gaps. Consider ergodic Jackson networks with the same arrival and service intensities, and with state processes \mathbf{X} and $\hat{\mathbf{X}}$. Assume that for the *extended* routing matrices R and \hat{R} the stochastic solutions ξ of the traffic equation coincide. If $R \prec_{pd} \hat{R}$, then

$$Gap(Q^{\hat{\mathbf{X}}}) \leq Gap(Q^{\mathbf{X}}).$$

As a consequence we have the following result. Consider an ergodic Jackson network process \mathbf{X} with state dependent service intensities and with positive external arrival rates $\lambda_i > 0$ at all nodes $i = 1, \dots, J$.

- Assume that the extended routing matrix $R = [r_{ij}]_{i,j=0,1,\dots,J}$ has strict positive departure probabilities $r_{i0} > 0$ from every node $i = 1, \dots, J$.

- Assume further that the routing of \mathbf{X} fulfills overall balance for all network nodes with respect to the solution $\eta_i, i = 1, \dots, J$, of the traffic equation

$$\eta_j \sum_{i=1}^J r_{ji} = \sum_{i=1}^J \eta_i r_{ij}, \quad \forall j = 1, \dots, J.$$

then for the ergodic Jackson network process $\hat{\mathbf{X}}$ of independent birth-death processes, the nodes of which have the same service intensities and external arrival rates $\hat{\lambda}_i = \lambda_i$

$$Gap(Q^{\mathbf{X}}) \geq Gap(Q^{\hat{\mathbf{X}}})$$

Finally we study the failure coordinate of unreliable networks. We introduce for transition matrices on a finite partially ordered set the following monotonicity. \mathbf{P} is ***-Möbius monotone** if

$$C^{-1}\mathbf{P}C \geq 0,$$

for C representing the Möbius function of the ordering.

We present a partial order analog of a result by Diaconis and Fill [2], originally related to linearly ordered spaces.

Assume that there exists $\mathbf{e}_M \in E$ such that for all $\mathbf{e}, \mathbf{e} \preceq \mathbf{e}_M$ and that

- $\frac{\mu(\mathbf{e})}{\pi(\mathbf{e})}$ is Möbius monotone
- $\overleftarrow{\mathbf{P}}$ is *-Möbius monotone

Then there exists strong dual (μ^*, \mathbf{P}^*) on $E^* = E$ with link $\Lambda(\mathbf{e}^*, \mathbf{e}) = I(\mathbf{e} \preceq \mathbf{e}^*) \frac{\pi(\mathbf{e})}{H(\mathbf{e}^*)}$, where $H(\mathbf{e}^*) = \sum_{\mathbf{e} \preceq \mathbf{e}^*} \pi(\mathbf{e})$. Moreover, the dual is uniquely determined as

$$\begin{aligned} \mu^*(\mathbf{e}) &= H(\mathbf{e}) \sum_{\mathbf{e}' \succeq \mathbf{e}} \eta(\mathbf{e}, \mathbf{e}') \frac{\mu(\mathbf{e}')}{\pi(\mathbf{e}')} \\ \mathbf{P}^*(\mathbf{e}^*, \mathbf{e}) &= \frac{H(\mathbf{e})}{H(\mathbf{e}^*)} \sum_{\mathbf{e}' \succeq \mathbf{e}} \eta(\mathbf{e}, \mathbf{e}') \mathbf{P}(\mathbf{e}', \{\overleftarrow{Y}_1 \preceq \mathbf{e}^*\}), \end{aligned}$$

where $\overleftarrow{\mathbf{P}}$ corresponds to the time reversed process, $\{Y_n\}$ is a chain with transition matrix \mathbf{P} .

REFERENCES

[1] H. Daduna and R. Szekli, *Impact of routing on correlation strength in stationary queueing network processes*, Journal of Applied Probability **45** (2008), 846–878.
 [2] P. Diaconis and J.A. Fill, *Strong stationary times via a new form of duality*, The Annals of Probability **18** (1990), 1483–1522.

Gossip Algorithms Using (Weighted) Path Averaging

PATRICK THIRAN

(joint work with Florence Bénézit, Vincent Blondel, Alexandros G. Dimakis,
John Tsitsiklis, Martin Vetterli)

Gossip algorithms are distributed message-passing schemes designed to disseminate and process information over networks, in particular wireless sensor networks, in which case the network topology can be modeled by a random geometric graph or even a lattice. They have received significant interest because the problem of computing a global function of data distributively over a network, using only localized message-passing, is fundamental for numerous applications.

Every node or agent $i \in \{1, n\}$ in a network holds an initial state $x_i(0)$, and wants to learn the average $x_{ave} = (1/n) \sum_{i=1}^n x_i(0)$ of all the values in the network in a *distributed* manner. Ideally a distributed averaging algorithm should be efficient in terms of energy and delay, and not require too much knowledge about the network topology at each node, nor sophisticated coordination between nodes. In order to achieve this goal, each agent i holds at any iteration t an estimate $x_i(t)$ of the average x_{ave} . All the estimates are gathered in a vector of size n denoted by $\mathbf{x}(t)$. The goal of an average consensus algorithm can thus be restated as

$$(1) \quad \lim_{t \rightarrow \infty} \mathbf{x}(t) = \frac{\mathbf{1}\mathbf{1}^T}{n} \mathbf{x}(0) = x_{ave} \mathbf{1},$$

where $\mathbf{1}$ denotes the vector of all ones and T denotes transposition.

At each iteration t one or several nodes receive the estimates of one or several other nodes. Based on the received estimates, nodes update their own estimate according to certain update rules. The common feature between all the algorithms described is the linearity of these estimate update rules. For any algorithm and any of its iteration t , there is a matrix $\mathbf{W}(t)$ such that:

$$(2) \quad \mathbf{x}^T(t) = \mathbf{x}^T(t-1) \mathbf{W}(t).$$

The choice of the matrix $\mathbf{W}(t)$ depends on the particular averaging algorithm, see [1] for some examples. In a nutshell, there are three main constraints on the matrices $\mathbf{W}(t)$. First, consensus must be stable. If, at time t , all nodes agree on one estimate, then they should keep agreeing on this estimate for times greater than t . Because of Eq. (2), this mathematically translates to: $\mathbf{1}^T \mathbf{W}(t) = \mathbf{1}^T$. Second, in order to guarantee that the only possible consensus is *average* consensus, the average of the estimates at any time should be equal to x_{ave} , which yields that $\mathbf{W}(t) \mathbf{1} = \mathbf{1}$. The matrices $\mathbf{W}(t)$ must therefore be doubly stochastic. Finally, the estimates should contract to the true average, and the condition for convergence, which can in general be formulated as conditions on of the spectrum of a function of $\mathbf{W}(t)$, depends on the type of convergence; see Chapter 3.1 in [1].

The simplest gossip algorithm is *Pairwise Gossip*, where random pairs of connected nodes iteratively and locally average their values until convergence to the global average [7, 4]. Pairwise Gossip algorithms on random geometric graphs require a total expected *cost* of $\Theta(n^2)$ messages to converge to the global average

x_{ave} within a given degree of accuracy [4, 6]. In contrast, the obvious solution of averaging numbers on a spanning tree and flooding back the average to all the nodes requires only $O(n)$ messages (Note that any algorithm that averages n numbers will require $\Omega(n)$ messages).

Dimakis et al. [6] proposed *Geographic Gossip*, an alternative gossip scheme that reduces the cost to $O(n^{1.5}\sqrt{\log n})$ messages, with slightly more complexity at the nodes. Assuming that the nodes have knowledge of their geographic location and under some assumptions in the network topology, greedy geographic routing can be used to build an *overlay network* where *any pair* of nodes can communicate. The overlay network is a complete graph on which standard Pairwise Gossip converges with $O(n)$ iterations. At each iteration we perform greedy routing, which costs $\Theta(\sqrt{n/\log n})$ messages on a random geometric graph $G(n, r)$. In total, Geographic Gossip thus requires $O(n^{1.5}\sqrt{\log n})$ messages.

Path Averaging [2] is the same algorithm as Geographic Gossip, with the additional modification of *averaging all the nodes on the routed paths*. Averaging the whole route comes almost for free in multihop communication, because a packet can accumulate the sum and the number of nodes visited, compute the average when it reaches its final destination and follow the same route backwards to disseminate the average to all the nodes along this route.

In Path Averaging, the selection of the routed path (and hence the routing algorithm) will affect the performance of the algorithm. The mathematical analysis of Path Averaging with greedy routing is complex because the number of possible routes increases exponentially in the number of nodes.

Path Averaging on random geometric graphs combines gossip with random greedy geographic routing. A key assumption is that each node knows its location and is able to learn the geographic locations of its one-hop neighbors (for example using a single transmission per node). Also the nodes need to know the size of the space occupied by the network.

The algorithm operates as follows: at each time-slot one random node activates and selects a random position (target) on the unit square region where the nodes are spread out. Note that no node needs to be located on the target, since this would require global knowledge of locations. The node then creates a packet that contains its current estimate of the average, its position, the number of visited nodes so far (one), the target location, and passes the packet to a neighbor that is *randomly chosen among its neighbors closer to the target*. As nodes receive the packet, randomly and greedily forwarding it towards the target, they add their value to the sum and increase the hop counter. When the packet reaches its destination node (the first node whose nearest neighbors have larger distance to the target compared to it), the destination node computes the average of all the nodes on the path, and reroutes that information backwards on the same route.

Using the same overlay network construction as for Geographic Gossip, and the Poincaré inequality [5], we prove in [2] that on a random geometric graph $G(n, r)$ deployed on torus, with a sufficiently high connectivity radius $r > 15\sqrt{\log n/n}$, the cost drops to $O(n \log n)$ messages only with Path Averaging.

All the previous algorithms require some amount of coordination to route the packets back and forth between the two end-nodes of the chosen route at each iteration, which can be quite difficult in wireless networks, where the connectivity graph is often time-varying. Replacing the Path Averaging gossip algorithm described above by a one-way version of this algorithm eliminates this problem, but at the same time forces the matrix $\mathbf{W}(t)$ to be triangular, which in turns prevents it to be doubly stochastic.

As a solution, we propose in [3] an asynchronous generalization of the synchronous Uniform Gossip algorithm outlined in [7], which we call *Weighted Gossip*, where the states are written as ratios between *sums* $\mathbf{s}(t)$ and *weights* $\mathbf{w}(t)$. In other words, $\mathbf{x}(t) = \mathbf{s}(t)/\mathbf{w}(t)$ with the division performed element-wise. The update rule (2) is then replaced by an update rule on both $\mathbf{s}(t)$ and $\mathbf{w}(t)$, which is

$$\begin{aligned}\mathbf{s}^T(t) &= \mathbf{s}^T(t-1)\mathbf{D}(t) \\ \mathbf{w}^T(t) &= \mathbf{w}^T(t-1)\mathbf{D}(t),\end{aligned}$$

where the *diffusion matrices* $\{\mathbf{D}(t), t \geq 0\}$ form a stationary and ergodic sequence of stochastic matrices with positive diagonal entries, with the property that their expectation $\mathbb{E}[\mathbf{D}]$ is irreducible. The initial values are $\mathbf{s}(0) = \mathbf{x}(0)$ and $\mathbf{w}(0) = \mathbf{1}$. Observe that $\mathbf{D}(t)\mathbf{1} = \mathbf{1}$, but that $\mathbf{1}^T\mathbf{D}(t) \neq \mathbf{1}^T$: sums and weights are conserved through the iterations because of the former equality, but do not reach consensus because of the latter inequality. Nevertheless, we prove in [3] that the ratio $\mathbf{x}(t) = \mathbf{s}(t)/\mathbf{w}(t)$ verifies (1) with probability 1. Consequently, Weighted Gossip using $\{\mathbf{D}(t), t \geq 0\}$ converges to a consensus with probability 1.

REFERENCES

- [1] F. Bénézit, *Distributed Average Consensus for Wireless Sensor Networks*, PhD Thesis **4509** (2009), EPFL.
- [2] F. Bénézit, A.G. Dimakis, P. Thiran and M. Vetterli, *Order-optimal consensus through randomized path averaging*, IEEE Trans. Information Theory **56** (2010), 5150–5167.
- [3] F. Bénézit, V. Blondel, P. Thiran, J. Tsitsiklis and M. Vetterli, *Weighted Gossip: Distributed averaging using non doubly stochastic matrices*, Proceedings ISIT (2010), 1753–1757.
- [4] S. Boyd, A. Ghosh, B. Prabhakar and D. Shah, *Randomized gossip algorithms*, IEEE Transactions on Information Theory **52** (2006), 2508–2530.
- [5] P. Diaconis and D. Stroock, *Geometric bounds for eigenvalues of Markov chains*, Annals of Applied Probability, **1** (1991).
- [6] A.G. Dimakis, A.D. Sarwate and M.J. Wainwright, *Geographic Gossip: Efficient Aggregation for Sensor Networks*, Proceedings IPSN (2006).
- [7] D. Kempe, A. Dobra and J. Gehrke, *Gossip-based computation of aggregate information*, Proc. FOCS **44** (2003), 482–491.

Ordered random walks

VITALI WACHTEL

(joint work with Denis Denisov)

Let W be the k -dimensional Weyl chamber of type A, that is,

$$W := \{x \in \mathbb{R}^k : x_1 < x_2 < \dots < x_k\}.$$

Let $S = (S_1, S_2, \dots, S_k)$ denote a k -dimensional random walk with independent and identically distributed coordinates. We are interested in the behaviour of the exit time

$$\tau_x := \min\{n \geq 1 : x + S(n) \notin W\}, \quad x \in W,$$

and in the behaviour of paths of S conditioned on $\{x + S(k) \in W \text{ for all } k \in I\}$, where $I = \{1, 2, \dots, n\}$ or $I = \mathbb{N}$.

The case of $I = \mathbb{N}$ seems to be more interesting, since $\mathbf{P}(\tau_x = \infty) = 0$. Therefore, it is not at all obvious how to understand a process which never leaves W . Usually one performs a Doob transform with some positive harmonic function h . Recall that h is called harmonic, if $\mathbf{E}[h(x + S(1)), \tau_x > 1] = h(x)$. The question is how to find such a function. It is well known, that the Vandermonde determinant $\Delta(x)$ is harmonic for the Brownian motion. The corresponding process is called the Dyson Brownian motion. It is also known, that $\Delta(x)$ is harmonic for some special cases of nearest neighbour random walks, see [4]. The reason, why one could use Δ , is that in all these cases $\Delta(x + S(\tau_x)) = 0$ or, in other words, the random walk can not jump over the boundary of W . But if S is a random walk that can jump over the boundary, then the question on the harmonic function was open for quite a long time. Only recently, Eichelsbacher and König [3] have shown that $V(x) := \Delta(x) - \mathbf{E}\Delta(x + S(\tau_x))$ is harmonic for random walks with $\mathbf{E}|X|^{ck^3} < \infty$. In [1] we have shown the harmonicity of V under much milder condition: It suffices to assume that $\mathbf{E}|X|^{r_k} < \infty$ with $r_k = (k-1)1\{k \geq 4\} + (2+\varepsilon)1\{k=3\}$. Under this condition the relation

$$(1) \quad \mathbf{P}(\tau_x > n) \sim \kappa V(x) n^{-k(k-1)/4}$$

holds too. This result has been obtained earlier by Eichelsbacher and König [3] under their moment restriction. Since V is harmonic, we may perform the corresponding Doob transform:

$$\mathbf{P}^{(V)}(x + S_n \in dy) = \mathbf{P}(x + S_n \in dy, \tau_x > n) \frac{V(y)}{V(x)}, \quad y \in W.$$

It is not surprising that the rescaled conditioned random walk converges to the Dyson Brownian motion. More precisely,

- (a) the process $\frac{x\sqrt{n} + S(\lfloor nt \rfloor)}{\sqrt{n}}$ under $\mathbf{P}^{(V)}$ converges to $B(t)$ under $\mathbf{P}_x^{(\Delta)}$;
- (b) the process $\frac{x + S(\lfloor nt \rfloor)}{\sqrt{n}}$ under $\mathbf{P}^{(V)}$ converges to $B(t)$ under $\mathbf{P}_0^{(\Delta)}$.

The proof of well-posedness of V is based on martingale theory. To show all the asymptotic relations we use the standard Komlos-Major-Tushnady-coupling and

known formulae for the Brownian motion. (The use of coupling explains, why we need at least $2 + \varepsilon$ moments in the three dimensional case.)

Relying on the analogy to the central limit theorem, one can expect that any ordered random walk with finite second moments should converge to the Dyson Brownian motion. The main problem is how to construct ordered random walks, or, equivalently, how to find a harmonic function. (The function $V(x)$ is not defined for random walks with $\mathbf{E}|X|^{k-1} = \infty$.) Assuming that $\mathbf{P}(|X| > x) \sim x^{-\alpha}$ with some $\alpha \in (k - 2, k - 1)$ we studied the tail behaviour of τ_x : There exist a positive function U and an absolute constant θ such that

$$(2) \quad \mathbf{P}(\tau_x > n) \sim \theta U(x)n^{-\alpha/2-(k-1)(k-2)/4}.$$

Behind this relation lies a quite simple strategy. In order to achieve $\tau_x > n$, either the random walk on the top or the random walk on the bottom should jump away, i.e., $X_k(l) \approx \sqrt{n}$ or $X_1(l) \approx -\sqrt{n}$ for some $l \geq 1$. After such a jump we have a system of $k - 1$ random walks with bounded distances between each other and one random walk on the characteristic distance \sqrt{n} . This implies that the probability that all k random walks stay in W up to time n is of the same order as the probability that $k - 1$ random walks stay in W up to time $n - l$. But $\mathbf{E}[|X|^{k-2}] < \infty$. So we can apply our first result, which says that the latter probability is of order $n^{-(k-1)(k-2)/4}$. Since $\mathbf{P}(|X| > \sqrt{n}) \sim n^{-\alpha/2}$, we see that $\mathbf{P}(\tau_x > n)$ is of order $n^{-\alpha/2-(k-1)(k-2)/4}$.

The function U is strictly superharmonic, i.e., $\mathbf{E}[U(x + S(1)), \tau_x > 1] < U(x)$. This implies that the measure

$$\mathbf{P}^{(U)}(x + S_n \in dy) = \mathbf{P}(x + S_n \in dy, \tau_x > n) \frac{U(y)}{U(x)}, \quad y \in W$$

is not a probability, $\mathbf{P}^{(U)}(W) < 1$. We loose a part of the mass because of one big jump. To reintroduce the lost mass we enlarge the state-space:

$$\hat{W} := W \cup W_1 \cup W_2,$$

where

$$W_1 = \{(x_1, x_2, \dots, x_{k-1}, \infty), x_1 < x_2 < \dots < x_{k-1}\}$$

$$W_2 = \{(-\infty, x_2, x_3, \dots, x_k), x_2 < x_3 < \dots < x_k\}.$$

On this enlarged set we can define a transition kernel, which does not loose mass. The asymptotic behaviour of the corresponding Markov process, which we also call 'ordered random walks', is quite clear. The process jumps at the random (but finite) time-point from W either to W_1 or to W_2 , and stays there for all times. The rescaled process converges to the limit, which can be described as follows: Take the $k - 1$ dimensional Dyson Brownian motion starting from zero. With some probability, say $p(x)$, we add one coordinate with constant value ∞ and with probability $1 - p(x)$ we add the coordinate with value $-\infty$. We see that, at least for this version of ordered random walks, the analogy to the central limit theorem is not complete. We leave the question open, whether one can find a harmonic function in the described situation.

REFERENCES

- [1] D. Denisov and V. Wachtel, *Conditional limit theorems for ordered random walks*, Electron. J. Probab. **15** (2010), 292–322.
- [2] D. Denisov and V. Wachtel, *Ordered random walks with heavy tails*. Work in progress.
- [3] P. Eichelsbacher and W. König, *Ordered random walks*, Electron. J. Probab. **13** (2008), 1307–1336.
- [4] W. König, N. O’Connell and S. Roch, *Non-colliding random walks, tandem queues, and discrete orthogonal polynomial ensembles*, Electron. J. Probab. **7** (2002), 1–24.

Auction Networks a Stochastic Network Environment

NEIL WALTON

(joint work with Peter Key)

In this abstract, we discuss the application of stochastic network theory to the theory of auctions. This summarises a presentation in the open problem session of the Mathematical Challenges in Stochastic Networks programme. Indeed, the problems discussed here do not represent a single unsolved problem but are more accurately a relevant space of models, problems and concepts which, in the opinion of this author, are amenable to the Stochastic Networks research community and could greatly benefit from their attention.

The problems discussed here are based on the observation of two data sets. Chiefly, from the auction data from Forza Motorsport 2 a car racing game for Microsoft’s Xbox 360 games console but, also, from auction data for adverts on Microsoft’s Bing search engine. Both data sets represent millions of auctions with hundreds of thousands of bidders. The participation and execution of auctions in both systems are fundamentally stochastic and the network formed through trade and competition is large and complex. Even so, simple microscopic effects can be observed and macroscopic phenomena are evident. We develop the components of a model which reflects different interests within the stochastic networks community, and then, we discuss qualitative phenomena which one would wish to observe in a formal mathematical analysis.

A Stochastic Auction Model: We consider an ascending price auction. Bidders, wishing to buy an item, arrive to the auction as a Poisson process. With each bidder, we associate an independent identically distributed bid. After r units of time the auction begins. From a price $\$r$ the auctioneer raises the price linearly in time until there are no remaining bidders. The item is then sold at this price to the last remaining bidder. The auctions of this sort should be analytically tractable by analysing point processes. The time dependent auction given here represents an auction where bidding extends the remaining time of an auction, such techniques are employed by companies such as eBay and Swoopo [2]

A Queueing Model: Consider a taxi stand. Taxis and customers both arrive at the stand as two independent Poisson processes and both are willing to wait for service but will renege after an exponentially distributed time. Normally, at a taxi stand service is first-come-first-served, but we can change this. Suppose customers

are willing to bid for quicker service. Such queueing analysis has been applied to the study of government corruption [3]. We suppose the queued customers bid for the taxi at the head of the queue. An auction of the type described above thus ensues. We note that the auction time to process the winning customer is given by the maximum bid within the queue. Thus, we could interpret the queueing policy as a longest processing time first policy. The long run return to the taxi company is the length of time the head of the line taxi spends during bidding minus the costs of taxi abandonments. Creating an optimal return to the taxi firm is thus a relevant problem. In addition, if we consider the prices paid if customers arrival and are served within batches the resulting payments form a “generalised first price auction”. These auctions were used by early internet search engines and currently a somewhat similar “generalised second price auction” is deployed by Goggle, Yahoo! and Microsoft [6].

A Network Model: Suppose now that there is a common pool of users. Users can set up auctions or bid on existing auctions. When an auction is set up by a user it is equivalent to sending a taxi into the above queueing model. Similarly bidding is equivalent to becoming a customer. The trading between users will form a network. When a user sells an item, he creates a half edge and when he bids on an auction he offers a half edge. When an item is sold between two users an edge is created between them. This creates a network. Rather than edges occurring at uniformly at random, like in the configuration model or Erdős-Renyi random graphs [4], the auctioneer choose which edge he prefers to connect. This is also different to preferential attachment models currently considered as the auctioneers choice of edge is locally optimal, not by a probabilistic weight [5, 1].

The above three components loosely describe how a closed stochastic network economy could be constructed, though certainly other models of network trade could be constructed. What sort of phenomenon and behaviours would we like to understand and what analysis do we believe would be amenable?

Variability in prices: We find in data that prices and user wealths are best described by heavy tailed distributions. The distributions are developed over time and there are distinct phases where prices are low due lack of competition and where prices are high but to the generation of wealth through the auction network. This phenomenon seems very analogous to percolation effects, where there is a critical point at which every node is influenced by the state of every other node. The heavy tail distribution in prices in turn effect performance of the auction queue described above. Auctions take longer and lead to high levels of abandonment. Thus the above queueing model is best studied when jobs have heavy tailed service requirements.

Network Structure: Data suggests that the degree distribution of the auction network is power law distributed and the network is highly connected (with a single giant component emerging even over short time periods). The game theoretic process by which edges form is a non-standard random graph enviroment and a successful mathematical analysis could represent a generalisation of traditional random graph models.

User behaviour: As pointed out at the open problem session, a large number of complex behaviours could be exhibited by users. Even so, data suggests behaviour is much more simple. Users increase and decrease their bids based on the success or failure. Users do not bid for items far beyond the first few auctions displayed. Such price probing is somewhat similar to rate control in a communications system. Perhaps such an analysis yield results?

As we have mentioned, this is certainly not a conjecture. These are observable phenomena. But hopefully this provokes thoughts to other areas of application, different theoretical models and different limiting regimes which may be relevant to this research community.

REFERENCES

- [1] A.-L. Barabási and R. Albert, *Emergence of scaling in random networks*, Science **286** (1999), 509–512.
- [2] J. Byers, M. Mitzenmacher and G. Zervas, *Information Asymmetries in Pay-Per-Bid Auctions: How Swoopo Makes Bank*, (to appear) Electronic Commerce (2010).
- [3] R. Hassin and M. Haviv, *To Queue or not to Queue: Equilibrium Behavior in Queueing Systems*, Kluwer Academic Publishers (2003).
- [4] M. Molloy and B. Reed, *A critical point for random graphs with a given degree sequence*, Random Structures & Algorithms **6** (1995), 161–180.
- [5] D. de Solla Price, *A general theory of bibliometric and other cumulative advantage processes*. J. Amer. Soc. Inform. Sci. **27** (1976), 292–306.
- [6] H.R. Varian, *Position auctions*, International Journal of Industrial Organization, Elsevier (2007).

Nonparametric inference for networks of queues

CORNELIA WICHELHAUS

(joint work with Roland Langrock and Susan Pitts)

We consider stochastic networks as interacting systems of nodes between which indistinguishable items (customers) move in order to receive service. Each node represents a single or multiserver and a waiting room in which customers queue for service. The external interarrival times, the sojourn times at the nodes and the routing decisions of the customers follow underlying stochastic rules. Important application fields of these systems with strongly dependent components are computer and telecommunication networks, the internet, manufacturing networks, supply chain networks and population dynamical systems.

In all these fields there is an increasing interest in statistical inference for system characteristics based on incomplete information of the working systems. However, whereas the probabilistic structure of stochastic networks (e.g. stability conditions, steady state behavior) has been widely explored over the last decades, statistical inference for these systems (e.g. parameter estimation, identifiability results) has not been developed properly yet. Difficulties arise due to the inner dependence structure of the networks which yields dependent multidimensional data. The aim is to establish a thorough statistical theory, including data selection, necessary and

sufficient criteria for identifiability as well as the construction and comparison of different classes of estimators.

A challenging problem of particular importance is the inference for the service time distributions at the nodes. The knowledge of the distribution of the operating times at the nodes is essential for specifying the performance behavior and for guaranteeing reliability and quality standards of the systems. In the literature, nonparametric approaches exist for single queues, see [1], [2], [3], [4], [5]. However, up to now there are no results for connected systems with two or more nodes. Thus, most of the currently applied techniques for service demand estimation in queueing networks focus only on the first and second order moments of the service times obtained e.g. by linear regression of measured utilization values. Appropriate techniques from nonparametric statistics must be derived for this new field of research. The mathematics behind is challenging, it is in the interplay of multivariate stochastic processes and inverse statistical problems.

In the talk we present a nonparametric approach for the service time estimation problem for networks with an arbitrary fixed number of nodes. We consider networks with general point processes as external arrival processes, with an infinite number of servers and a general service time distribution at each node. The routing topology is unknown. We assume that we are able to observe the external arrivals and external departures over a stretch of time whereas we cannot assign the departures to the arrivals and directly measure the individual sojourn times of the customers in the system. We give various identifiability criteria. The crucial point of our approach is the observation that the influence of the external arrival processes on the external departure processes can be described in a linear and time-invariant way. This makes it possible to apply cross-spectral techniques for multivariate point processes. We derive estimators for the routing probabilities and for the service time distributions at the nodes. The estimators are shown to be consistent and asymptotically normal. Simulation results for small networks show satisfactory results and illustrate the good quality of the estimators. In [7] we present the general theory and the results for feedforward systems, the generalization to networks with general routing topologies is given in [6].

REFERENCES

- [1] N.H. Bingham and S.M. Pitts, *Non-parametric estimation for the $M|G|\infty$ queue*, Ann. Inst. Statist. Math. **51** (1999), 71–97.
- [2] D.R. Brillinger, *Cross-spectral analysis of processes with stationary increments including the stationary $G/G/\infty$ queue*, Annals of Probability **2** (1974), 815–827.
- [3] P. Hall and J. Park, *Nonparametric inference about service time distribution from indirect measurements*, J. R. Statist. Soc. B **66** (2004), 861–875.
- [4] M.B. Hansen and S.M. Pitts, *Nonparametric inference from the $M/G/1$ workload*, Bernoulli **12** (2006), 737–759.
- [5] S. Pitts, *Nonparametric estimation of the stationary waiting time distribution function for the $GI|G|1$ queue*, The Annals of Statistics **3** (1994), 1428–1446.
- [6] S. Pitts and C. Wichelhaus, *Nonparametric inference for stochastic networks of general network topology using cross-spectral analysis*, University of Heidelberg, preprint (2010).

- [7] C. Wichelhaus and R. Langrock, *Nonparametric inference for stochastic tandem networks based on cross-spectral analysis of point processes*, University of Heidelberg, preprint (2009).

Participants

Prof. Dr. Venkat Anantharam
Department of Electrical
Engineering and Computer Sciences
University of California
Berkeley , CA 94720
USA

Prof. Dr. Soren Asmussen
Department of Mathematical Sciences
University of Aarhus
Building 530
Ny Munkegade
DK-8000 Aarhus C

Dr. Erik J. Baurdoux
Department of Statistics
London School of Economics
Houghton Street
GB-London WC2A 2AE

Dr. Bartek Blaszczyszyn
INRIA/ENS
CS 81321
23, avenue d'Italie
F-75214 Paris Cedex 13

Dr. Onno J. Boxma
Dept. of Mathematics & Computer Science
Eindhoven University of Technology
Postbus 513
NL-5600 MB Eindhoven

Prof. Dr. Andreas Brandt
Institut für Operations Research
Wirtschaftswissenschaftliche Fak.
Humboldt-Universität Berlin
Spandauer Str. 1
10178 Berlin

Prof. Dr. Hans Daduna
Fachbereich Mathematik
Universität Hamburg
Bundesstr. 55
20146 Hamburg

Dr. Krzysztof Debicki
Institute of Mathematics
Wroclaw University
pl. Grunwaldzki 2/4
50-384 Wroclaw
POLAND

Dr. Ton Dieker
School of Industrial and Systems
Engineering
Georgia Institute of Technology
765 Ferst Drive
Atlanta GA 30332-0205
USA

Prof. Dr. Serguei Foss
Department of Actuarial Mathematics
and Statistics
Heriot-Watt University
Riccarton
GB-Edinburgh EH14 4AS

Prof. Dr. David Gamarnik
Massachusetts Institute of Technology
Sloan School of Management
E 53 - BS 7
77 Massachusetts Avenue
Cambridge MA 02139-4307
USA

Dr. Ayalvadi Ganesh
Department of Mathematics
University of Bristol
University Walk
GB-Bristol BS8 1TW

Daniel Gentner

Institut für Stochastik
Karlsruher Institut f. Technologie (KIT)
Kaiserstr. 89-93
76133 Karlsruhe

Prof. Dr. Peter Glynn

Management Science & Engineering
Huang Engineering Center
475 Via Ortega
Stanford University
Stanford CA 94305-4121
USA

Prof. Dr. Bruce Hajek

Coordinated Science Laboratory
University of Illinois at
Urbana-Champaign
1101 W. Springfield Avenue
Urbana , IL 61801
USA

Dr. Gerard Hooghiemstra

Faculty of Electrical Engineering,
Applied Mathematics & Computer Sc.
Delft University of Technology
Mekelweg 4
NL-2628 CD Delft

Jevgenijs Ivanovs

EURANDOM
Technical University Eindhoven
P.O. Box 513
NL-5600 MB Eindhoven

Prof. Dr. Ramesh Johari

Management Science & Engineering
Huang Engineering Center
475 Via Ortega
Stanford University
Stanford CA 94305-4121
USA

Prof. Dr. Offer Kella

Department of Statistics
The Hebrew University of Jerusalem
Mount Scopus
Jerusalem 91905
ISRAEL

Prof. Dr. Takis Konstantopoulos

Department of Mathematics
University of Uppsala
S-75238 Uppsala

Prof. Dr. Günter Last

Institut für Stochastik
Karlsruher Institut f. Technologie (KIT)
Kaiserstr. 89-93
76133 Karlsruhe

Dr. Marc Lelarge

INRIA/ENS
CS 81321
23, avenue d'Italie
F-75214 Paris Cedex 13

Dr. Lasse Leskelä

Dept. of Mathematics & Systems Analy-
sis
Aalto University
PO Box 11100
FIN-00076 Aalto

Dr. Kurt Majewski

SIEMENS AG
CT T DE TC3
81739 München

Prof. Dr. Michel Mandjes

Korteweg de Vries (KdV)
Institute for Mathematics
University of Amsterdam
Plantage Muidersgracht 24
NL-1018 TV Amsterdam

Dr. James B. Martin

Department of Statistics
University of Oxford
1 South Parks Road
GB-Oxford OX1 3TG

Dr. Laurent Massoulié

Thomson Technology Paris Laboratory
46, quai A. Le Gallo
F-92648 Boulogne Cedex

Prof. Dr. David R. McDonald

Department of Mathematics & Statistics
University of Ottawa
585 King Edward Avenue
Ottawa , Ont. K1N 6N5
CANADA

Jennifer Mylosz

Fachbereich Mathematik
Universität Hamburg
Bundesstr. 55
20146 Hamburg

Prof. Dr. Ilkka Norros

VTT Technical Research Center
of Finland
P.O.Box 1000
FIN-02044 VTT - Espoo

Prof. Dr. Balaji Prabhakar

Dept. of Electrical Engineering and
Computer Science
Packard Building, Room 269
Stanford University
Stanford CA 94305-9510
USA

Dr. Alexandre Proutiere

Microsoft Research Systems and
Networking Group
Roger Needham Bldg.
7 J J Thomson Ave.
GB-Cambridge CB3 OFB

Prof. Dr. Kavita Ramanan

Division of Applied Mathematics
Brown University
Box F
182 George Str.
Providence , RI 02912
USA

Prof. Dr. Rhonda Righter

Department of Industrial
Engineering and Operations Research
University of California
Etcheverry Hall
Berkeley , CA 94720-1777
USA

Prof. Dr. Philippe Robert

INRIA
B.P. 105
F-78153 Le Chesnay Cedex

Justin Salez

Ecole Normale Supérieure
DMA
45 rue d'Ulm
F-75230 Paris Cedex 05

Philipp Schmidt

Department of Statistics
University of Oxford
1 South Parks Road
GB-Oxford OX1 3TG

Prof. Dr. Volker Schmidt

Institut für Stochastik
Universität Ulm
Helmholtzstr. 18
89081 Ulm

Prof. Dr. Devavrat Shah

Lab. for Information & Decision Systems
MIT
77 Massachusetts Avenue
Cambridge , MA 02139
USA

Dr. Alexander Stolyar
Bell Labs
600 Mountain Avenue
Murray Hill , NJ 07974-0636
USA

Dr. Ryszard Szekli
Institute of Mathematics
Wroclaw University
pl. Grunwaldzki 2/4
50-384 Wroclaw
POLAND

Prof. Dr. Peter G. Taylor
Dept. of Mathematics and Statistics
University of Melbourne
Melbourne VIC 3010
AUSTRALIA

Prof. Dr. Patrick Thiran
EPFL
School of Computer & Communication
Sc.
Station 14
CH-1015 Lausanne

Dr. Vitali Wachtel
Mathematisches Institut
Ludwig-Maximilians-Universität
München
Theresienstr. 39
80333 München

Neil Walton
Statistical Laboratory
Cambridge University
16 Mill Lane
GB-Cambridge CB2 1SB

Dr. Cornelia Wichelhaus
Institut für Angewandte Mathematik
Universität Heidelberg
Im Neuenheimer Feld 294
69120 Heidelberg

Andrea Winkler
Institut f. Mathematische Statistik
Universität Münster
Einsteinstr. 62
48149 Münster

Dr. Damon Wischik
Department of Computer Science
University College London
Gower Street
GB-London WC1E 6BT

Dr. Stan Zachary
Department of Actuarial Mathematics
and Statistics
Heriot-Watt University
Riccarton
GB-Edinburgh EH14 4AS

Dr. Ilze Ziedins
Department of Statistics
The University of Auckland
Private Bag 92019
Auckland 1142
NEW ZEALAND

Prof. Dr. Sergei Zuyev
Department of Mathematical Sciences
Chalmers University of Technology
S-41296 Göteborg

Prof. Dr. Bert Zwart
Centrum Wiskunde & Informatica (CWI)
Postbus 94079
NL-1090 GB Amsterdam