.

Workshop on Random Graphs and Random Processes

King's College London

Unfortunately, due to the current uncertainties, the workshop will not be held on Friday 3rd April 2020.

We are hoping to hold the event at a later date, to be confirmed subsequently.

Workshop on Random Graphs and Random Processes

King's College London
Friday 3rd April 2020

The workshop, the 4th in its series, looks at recent work in the area of random structures and algorithms and random processes on networks. Typical examples are properties of discrete random structures, random walks and interacting particle systems, threshold behaviour, design of randomized algorithms, and applications.

Talks: Strand Campus (Strand, London, WC2R 2LS), King's Building, room K2.31
Coffee: room K2.40

Morning Session 10:00–13:00
(Coffee from 9:40)

10:00-10:40	Stefanie Gerke Royal Holloway	The k-th shortest s-t path in a complete graph with random exponential edge weights
10:40-11:20	Milan Vojnovic LSE	Convergence Rates of Gradient Descent and MM Algorithms for Statistical Ranking Models

Coffee

11:40-12:20	George Giakkoupis Rennes	Optimal Time and Space Leader Election in Population Protocols
12:20-13:00	Olivier Bénichou LPTMC, Paris	Geometry-induced super-diffusion in driven crowded systems

Lunch Break 13:00–14:00

Afternoon Session 14:00–17:30
(Coffee from 13:45)

14:00-14:40	Claire Matieu CNRS Paris	Two-sided matching markets with correlated random preferences have few stable pairs
14:40-15:20	Marcel Ortgiese U. Bath	Voter models on inhomogeneous random graphs
15:20-15:50	Alan Frieze CMU	Minimum-weight combinatorial structures under random cost-constraints

Coffee

16:10-16:50	Manuela Fischer ETH Zurich	Local Glauber Dynamics
16:50-17:30	Nikolaos Fountoulakis U. Birmingham	Dynamic inhomogeneous random structures

Coffee breaks are organised, but not lunch.

Organisers: Colin Cooper, Frederik Mallmann-Trenn, Tomasz Radzik (Informatics Department)
Reimer Kuehn (Mathematics Department)

Abstracts

Olivier Bénichou. *Geometry-induced super-diffusion in driven crowded systems.*

Determining the response of a medium whilst in the presence of a tracer particle (TP) driven by an external force is a ubiquitous problem. It is for example at the heart of active microrheology, which has become a powerful experimental tool for the analysis of different systems in physics, chemistry, and biology . At the theoretical level, it has been the subject of many studies, most of which focus on analysing the velocity of the TP. Behavior beyond this force-velocity relation has been addressed in recent Molecular Dynamics simulations of active microrheology of glass-forming liquids and revealed superdiffusive fluctuations associated with the position of the TP. Such anomalous response, whose mechanism remains elusive, has been shown to occur only in systems close to their glass transition, suggesting that this could be one of its hallmarks. I will show that the presence of superdiffusion is in actual fact much more general, provided that the system is crowded and geometrically confined. I will rely on an analytical solution of a minimal model consisting of a driven TP in a dense, crowded medium in which the motion of particles is mediated by the diffusion of packing defects, called vacancies. This model represents a combination of two paradigmatic models of non-equilibrium statistical mechanics; asymmetric (for the TP) and symmetric (for the host particles in the system) simple exclusion processes. Through examining such non glassy systems, our analysis predicts a long-lived superdiffusion which ultimately crosses over to giant diffusive behavior. We find that this trait is present in confined geometries, for example long capillaries and stripes, and emerges as a characteristic response of crowded environments to an external force.

Stefanie Gerke. *The k-th shortest s-t path in a complete graph with random exponential edge weights.*

Manuela Fischer. *Local Glauber Dynamics.*

Sampling constitutes an important tool in a variety of areas: from machine learning and combinatorial optimization to computational physics and biology. A central class of sampling algorithms is the *Markov Chain Monte Carlo* method, based on the construction of a Markov chain with the desired sampling distribution as its stationary distribution. Many of the traditional Markov chains, such as the *Glauber dynamics*, do not scale well with increasing dimension. To address this shortcoming, we propose a simple local update rule based on the Glauber dynamics that leads to efficient parallel and distributed algorithms for sampling from Gibbs distributions. Concretely, we present a Markov chain that mixes in $O(\log n)$ rounds when Dobrushin's condition for the Gibbs distribution is satisfied. This improves over the *LubyGlauber* algorithm by Feng, Sun, and Yin [PODC'17], which needs $O(\Delta \log n)$ rounds, and their *LocalMetropolis* algorithm, which converges in $O(\log n)$ rounds but requires a considerably stronger mixing condition. Here, n denotes the number of nodes in the graphical model inducing the Gibbs distribution, and Δ its maximum degree. In particular, our method can sample a uniform proper coloring with $\alpha\Delta$ colors in $O(\log n)$ rounds for any $\alpha > 2$, which almost matches the threshold of the sequential Glauber dynamics and improves on the $\alpha > 2 + \sqrt(2)$ threshold of Feng et al.

Nikolaos Fountoulakis. *Dynamic inhomogeneous random structures.*

Alan Frieze. *Minimum-weight combinatorial structures under random cost-constraints.*

Recall that if the lengths of edges in the complete graph are drawn independently from $[0, 1]$ then the expected value of the minimum spanning tree is asymptotically equal to $\zeta(3)$. We consider analogous problems where edges have not only a random length but also a random cost, and we are interested in the length of the minimum-length structure whose total cost is less than some cost budget. For trees and arborescences we determine asymptotically accurate estimates of the minimum cost. These problems are NP-hard and we describe polynomial time algorithms. (Joint work with Tomasz Tkocz.)

For several other classes of structures, we determine the correct minimum length structure as a function of the cost-budget, up to constant factors – shortest paths, minimum matchings, Traveling Salesperson Problem. Moreover, we achieve this even in the more general setting where the distribution of weights and costs are arbitrary, so long as the density $f(x)$ as $x --> 0$ behaves like cx^γ for some $\gamma \geq 0$; previously, this case was not understood even in the absence of cost constraints. We also handle the case where each edge has several independent costs associated to it, and we must simultaneously satisfy budgets on each cost. In this case, we show that the minimum-length structure obtainable is essentially controlled by the product of the cost thresholds. (Joint work with Wesley Pegden, Greg Sorkin and Tomasz Tkocz.)

George Giakkoupis *Optimal Time and Space Leader Election in Population Protocols.*

In the population protocol model, n agents represented as identical finite state machines perform pairwise interactions, with the pair of agents interacting in each step chosen at random. A significant amount of work has been devoted to the study of the time and space complexity of the leader election problem, where all agents start from the same state, and the goal is to reach and maintain a global state where exactly one agent is in a dedicated leader state. It is known that $\Omega(\log\log n)$ states per agent are needed to complete this task with fewer than $\tilde{\Theta}(n)$ expected interactions per agent (Alistarh et al., 2017), and that $\Omega(\log n)$ expected interaction are required regardless of the number of states (Sudo and Masuzawa, 2019). On the algorithmic side, Gasieniec and Stachowiak (2018) gave the first protocol that uses an optimal $\Theta(\log\log n)$ number or states and elects a leader in $O(\log^2 n)$ expected interactions per agent. The number of interactions was subsequently improved to $O(\log n \log\log n)$ (Gasieniec et al., 2019). In this talk, we will present the first leader election population protocol that is both time and space optimal, electing a leader with $O(\log n)$ expected interactions per agent and using $\Theta(\log\log n)$ states. Joint work with Petra Berenbrink and Peter Kling.

Claire Matieu. *Two-sided matching markets with correlated random preferences have few stable pairs.*

Stable matching in a community consisting of N men and N women is a classical combinatorial problem that has been the subject of intense theoretical and empirical study since its introduction in 1962 in a seminal paper by Gale and Shapley.

In this paper, we study the number of stable pairs, that is, the man/woman pairs that appear in some stable matching. We prove that if the preference lists on one side are generated at random using the popularity model of Immorlica and Mahdian, the expected number of stable edges is bounded by $N \ln(N)+N$, matching the asymptotic value for uniform preference lists. If in addition that popularity model is a geometric distribution, then the number of stable edges is O(N) and the incentive to manipulate is limited. If in addition the preference lists on the other side are uniform, then the number of stable edges is asymptotically N up to lower order terms: most participants have a unique stable partner, hence non-manipulability. Joint work with Hugo Gimbert and Simon Mauras

Marcel Ortgiese. *Voter models on inhomogeneous random graphs.*

The voter model is a classical interacting particle system modelling how consensus is formed across a network. We analyse the time to consensus for the voter model when the underlying graph is a scale-free random graph. Moreover, we generalise the dynamics to include a 'temperature' parameter. The interplay between the temperature and the structure of the random graph leads to a very rich phase diagram, where in the different phases different parts of the underlying geometry determine the time to consensus.

Milan Vojnovic. *Convergence Rates of Gradient Descent and MM Algorithms for Statistical Ranking Models.*

We will present new results on the convergence rates of gradient descent and MM algorithms for estimating the parameters of the Bradley-Terry models of ranking outcomes, including models of paired comparisons, choice, and full-ranking outcomes. We will consider the maximum likelihood (ML) estimation and the maximum a-posteriori probability (MAP) estimation under a prior distribution commonly used in practice. An appealing feature of the MAP estimation is that an optimal solution always exists, unlike to the ML estimation. Our results provide tight characterisations of the convergence rates and shed light on how they depend on the properties of a graph defined by comparisons in the input training data. Perhaps surprisingly, we will show that there exist instances for which the MAP estimation can be arbitrarily slower than the ML optimisation. We will then propose an accelerated optimisation algorithm, which resolves the identified slow convergence issue, by leveraging the spectral properties of a graph Laplacian matrix. It turns out that our accelerated MM algorithm is a simple modification of the classical ML algorithm, which is commonly used for learning ranking models. (Joint work with Se-Young Yun and Kaifang Zhou.)