ITT 1: Networks
26 - 30 January 2015

SAMBa Executive Team and others

March 20, 2015

1 Theme and Participants

The major theme of the ITT will be '*networks*'. This does not mean that all these problems are network problems. It does not mean that they all require ideas from 'network science' to solve. Part of the challenge of the ITTs is the mathematical maturity and combination of methods that we might need in order to distill open-ended situations posed by project partners into mathematical problems that might motivate mathematical work at the level of a PhD.

The problems posed below build upon background work that has taken place between academic staff in the department and staff from the companies Det Norske Veritas – Germanischer Lloyd (DNV GL) and British Telecom (BT). General information about these companies can be found on the web at `www.dnvgl.com` and `www.bt.com` respectively. The following core participants will attend.

ITT		Total
Academics		
SAMBa students:	Matthew Durey, Dorotya Fekete, Marcus Kaiser, Sam Moore, Matthew Parkison, Matthew Thomas, Kirsten Williams.	7
SAMBa management:	Susie Douglas, Andreas Kyprianou, Paul Milewski.	3
SAMBa academics (key):	Chris Budd, Jon Dawes, Evangelis Evangelou, Antal Jarai, Finn Lindgren, Tony Shardlow.	6
SAMBa academics (invited):	??	5
Non-SAMBa Ph.D. students:	??	5
Non-SAMBa Academics:	Sam Johnson (Warwick)	1
Industrial		
DNV GL:	David Worthington, ??	2
BT:	Keith Briggs, Botond Virginas, Stephen Cassidy	3
Retired Shell senior mathematician:	Hennie Poulisse	1
Week Observers		
USP Saõ Carlos:	Poti + ??	2
Day Observers		
Smith Institute:	Tim Boxer	1
Tesco:	Maren Eckhoff	1
Wessex Water:	??	?
Met Office:	??	?

2 DNV GL

Finn Lindgren, Tony Shardlow and Evangelos Evangelou
DNV GL are offering one main topic which is the use of Baysian networks and event trees to construct risk measures, specifically in connection with oil rig safety procedures.

During the ITT, DNV GL would like to address the issue of building a Baysian network from a current event trees that are encapsulated in existing software. A 'worked example' will be presented for which students will be asked to re-cast existing results as a Baysian network. In formulating such risk problems in the setting of Baysian networks, DNV GL are specifically interested in: sensitivity, reliability, robustness, how to compare measurements against reality, inference, inter-dependencies, how to compare one Baysian network against another.
Prior to the ITT, the following preparatory material has been suggested for integration with the Student Symposium.

- Research paper (something to do with the Railway Institue - Finn please fill) that can be studied and presented in advance.
- Internal DNV GL document: *Technology leadership: risk reliability and human factors* DNV Doc. No./Report No.: 17QNCNF-1/
- Book: *Baysian Networks* Olivier Pourret (Editor), Patrick Naïm (Editor), Bruce Marcot (Editor)
- Data: spreadsheet containing risk scenarios for oil rigs
- Software: a programme called Ginie, which is free to use, as well as some DNV GL specific software.

In preparation, DNV GL and Sam Johnson have expressed their willingness to engage with the student-led symposium by coming to talk to them about the nature of their interest in Baysian networks, the structure of the spreadsheet they have and how it currently interacts with existing software.

3 BT

BT have a history of engagement with Bath, and with several other UK Universities. The current discussions on which this ITT is largely based began in February 2014 with a bilateral discussion day at BT's Research and Innovation HQ which is outside Ipswich (North-East of London). Around ten Bath staff and a similar number of BT research staff attended that day.

The discussions there led to the selection of problems that are set out below. Each problem has been subsequently worked on by a lead Bath academic and a lead BT researcher. We have involved a student in each problem already: the wireless signal propagation project was taken up by a Bath MSc student (on the Modern Applications of Mathematics course), supervised by Chris Budd. Antal Jarai and Jonathan Dawes worked with summer internship students on the other two problems.

3.1 Graph colouring problems arising in wireless networks - Antal Jarai

One of the problems BT is interested in is radio channel allocation for wireless devices in a way that minimizes interference (wireless access points or LTE base stations). Devices are represented as nodes of a graph, and two nodes are connected with an edge, if they are close enough to cause interference. Possible radio frequency bands are represented by colouring the nodes with different colours. Two desirable features — that are, however, in conflict with each other — are to use as few colours as possible and to have as little interference as possible. Another desirable feature is that nodes behave autonomously, that is, when they update their colour, they only see the colours of their neighbours.

During the ITT, BT would like to explore the performance of various colouring heuristics. One direction that is being explored at the moment is to use Markov chains on graph colourings and knowledge of their mixing properties to guarantee a certain level of performance. Some previous works that will be made available to the students:

- Exact optimization (of both channel and power allocation) for small networks: (i) M.R. Tijmes and K.M. Briggs: Optimizing the design of 802.11b wireless cities, BT technical report (7 pages) (2008); (ii) M.R. Tijmes: Simultaneous optimization of channel and power allocation for wireless cities, BT, University of Twente technical report (65 pages) (2008).

- M. Wiltshire, K.M. Briggs and A.A. Járai: Self-organizing networks: optimizing channel allocation with local knowledge only. Report on UG summer project, BT, University of Bath (10 pages) (2014).

3.2 BT – Domestic WiFi signal propagation - Chris Budd

BT will offer a problem on the reception of WiFi signals in a domestic environment. A set of data will be provided by BT of the strength of WiFi signals as measured in various places around a house. Students will be invited to construct a ray tracing model for the WiFi system which will give results in agreement with this data. BT are particularly interested in the effect on the WiFi signals of placing furniture and other obstacles in the room, and the students will be asked to extend their ray tracing models to see what happens in various scenarios. The students will also be asked to learn about more sophisticated models for WiFi propagation, but not to implement them during the ITT.
Prior to the ITT, the following preparatory material has been suggested for integration with the Student Symposium.

- PhD Thesis *Indoor Radio Channel Propagation Modelling by Ray Tracing Techniques* by David I. Laurenson

- Articles on FemtoCell systems and Rayleigh fading provided by Chris Budd

- Data: spreadsheet containing measurements made by BT

3.3 Self-organised network dynamics - Jon Dawes

A key part of the development of future communication networks (in many application areas) is their ability to self-organise. This capability is also referred to as having an 'autonomous system' or 'distributed control' heuristic / mechanism.

Within telecoms, this self-organising property is seen as a central feature of 4G networks and even further ahead, so called 'Long-Term Evolution' (LTE).

Over the summer we looked at two problems proposed by BT in this area. We will refer to these problems as the 'shortest path finding problem' and the 'bucket problem'.

In the form that they are stated below, they are both gross simplifications and there are many avenues for developing and extending these basic situations. We expect also that there is also useful literature on these problems available, but we haven't had the time to look in detail for it. A proper literature review would therefore be a very good first step in tackling these problems in a systematic way.

3.3.1 Shortest path problem

Fix an undirected unweighted graph without self-loops or multiple edges. Each node would like to learn how far away every other node is. In order to do this, nodes ask neighbouring nodes to send them information. The simplest information that a node could pass on would be its list of shortest distances to all other nodes. The receiving node can then use this to update its list of shortest distances. This process continues until sending more messages does not change the list of shortest distances that any node holds.

Natural questions to ask about about this process include:

- How many messages does it typically take before the process is complete?
- How can we construct a coarse-grained characterisation of this process on a large, but finite, graph?
- What happens if nodes are not picked to update their local list equally at random, but in some other sense?
- How do these answers depend on the number of edges in the graph, and
- on the type of graph (random, geometric random, lattice, 'small-world', clustered)?

Work in the summer looked at one possible coarse-grained description which appears to be useful, but is probably far from optimal. In addition, it also considered only the 'initial value problem' where initially a node knows their list of neighbours (which, naturally, have a shortest path of length 1 to the node) but nothing more: every node knows nothing about the paths to nodes that are not immediate neighbours.

Another useful and interesting scenario would be to consider starting from the equilibrium state where distances between every pair of nodes are known to both nodes in the pair, but on a graph in which nodes are dropping out and being added at random. In this setting the interest is in how the knowledge of paths to the newly added nodes are learnt by message-passing across the network.

In the simplest versions of these problems, the only information passed is the distances to other nodes. One could extend this to include more detailed information, for example the detailed route of the paths themselves, or the last time that a node acquired new information, and which of its neighbours provided that new information.

3.3.2 Bucket problem

Discussions with BT also concentrated on a very general formulation of a problem that was dubbed the 'bucket problem' for the following abstract description. Place a finite number of buckets (i.e. circular discs) centred on points selected at random in the unit square. Suppose that raindrops fall 'at random' in the square. How large should the buckets be in order to catch as much rain as possible? This is a general kind of resource allocation problem: one interpretation could be that the centres of buckets are mobile phone masts, or so-called femtocells that mobile phones are using to connect to to make calls. Another interpretation would be that the buckets are the call-out ranges of telecoms engineers who are maintaining the fixed network of wiring, and the rainfall is the unpredictable, random maintenance requirements of the equipment.

Interference between neighbouring buckets reduces their effectiveness, so for a given set of bucket centres, how do we maximimise the total area covered by buckets subject to keeping interference to a minimum? A dynamical way of restating the problem is to propose a rule, say in discrete time increments, for adjusting the sizes (i.e. radii) of the discs based on current estimations of coverage and overlap. So this implicitly requires that discs learn the sizes of their neighbouring discs, in this geometric sense, and update their estimates of disc sizes in the locality from one time step to the next.

Similar issues to the shortest path problem arise for later work: what is the equilibrium state and how quickly does the system settle to equilibrium again after discs are added into the system or removed from it?

A further complication is that there may be good reasons to add to the model terms that drive the system towards selecting only a few discs rather than some kind of 'equal radius' solution with all the penalised overlaps distributed across the region. By using only a few discs the system may save on overheads, e.g. powering each transmitter, and so if use of an area ('rainfall') is low, it might be advantageous to reduce some disc radii to zero in order to save on costs. This is another 'policy' that one might wish to model in some motivating situations. How this interacts with other 'policies' (e.g. quality of service, robustness) is another general angle to explore.

Literature:

- Reports by Junghoon Yoon / Jonathan Dawes from this summer, including graphs summarising the dynamics, and some attempts at modelling that might, or might not, lead somewhere.
- Various papers on wireless routing, but by no means a complete list.
- Article by Bollobas and Riordan from the *Handbook of Graphs* introducing the random graph family $G(n, p)$.

4 Suggested Symposium structure

SAMBa students are obliged to build the student-led symposium themselves, but obviously this will not happen without some gentle persuasion and indication of how to do this. The table below should be a rough outline of some of the activities that can and should occurring during the 10 weeks of the symposium. They already have access to a number of key individuals who we would hope can contribute to some pre-ITT sessions, as well as some additional individuals who can give general background to the theme of the ITT as well as industrial mathematics at large.

It is obligatory that students are assessed on one presentation in the symposium per semester. The default setting for assessing students is that Paul and Andreas will be present to grade their talks against a template (yet to be built), but it is the expectation that other academics can be called in to help with this assessment process.

This timetable can be shuffled around (preferably by the students themselves) in due course as we/they confirm speakers and dates.

Week	Monday	Tuesday	Wednesday	Thursday	Friday
1	29/09 10:15 8W 2.28 (Preparation)		01/10 10:15 6E 2.2 (Preparation)		03/10 12:15 4W 1.7 (Wolfson) Seminar Numerical Analysis
2	06/10 8W 2.28		08/10 10:15 6E 2.2 Talk Antal Jarai	09/10 9:15 8W 2.12 Preliminary reading disussion	
3	13/10 10:15 8W 2.28 Talk Chris Budd		15/10 10:15 6E 2.2 Talk Finn Lindgren		17/10 13:15 8W 2.28 Talk (DNV GL) David Worthington
4	20/10 10:15 8W 2.28 Talk Chris Budd	21/10 13:15 4W 1.7 CNM Seminar Wessex Water	22/10 10:15 6E 2.2 Talk Jon Dawes	23/10 9:15 8W 2.12 Talk Finn Lindgren	
5	27/10 10:15 8W 2.28 Talk Tony Shardlow		29/10 10:15 6E 2.2 Talk Jon Dawes	30/10 9:15 8W 2.12 Talk Tony Shardlow	
6	03/11 10:15 8W 2.28		05/11 10:15 6E 2.2 Talk Tim Rogers	06/11 9:15 8W 2.12	
7	10/11 10:15 8W 2.28 Talk Alicia Kim		12/11 10:15 6E 2.2 Talk Mathew Penrose	13/11 9:15 8W 2.12 Talk Kirill Cherednichenko	14/11 13:15 8W 2.24 Marcus
8	17/11 8W 2.28 Talk Robert Scheichl	← Also Monday 17/11 14:15-16:05 8W 2.11 Peter Mörters	19/11 10:15 6E 2.2 Talk Mathew Penrose	20/11 9:15 8W 2.12 Talk Karim Anaya-Izquierdo	
9	24/11 14:15 8W 2.11 Talk Antal Jarai		26/11 10:15 6E 2.2 Talk Alexandre Stauffer	27/11 9:15 8W 2.12 Sam Matt P.	
10	1/12 10:15 8W 2.28 Matt D. Dorka		3/12 10:15 6E 2.2 Talk (BT) Botond Virginas	4/12 9:15 8W 2.12 Kirsten Matt T.	
11	8/12 10:15 8W 2.28 DNV GL Training Day		10/12 10:15 6E 2.2 Talk Henk Poulisse	11/12 9:15 8W 2.12 Talk Emmanuel Jacob	