Optimal design of performance measurement experiments for complex, large-scale networks (DOENET)

John Schormans[1], Andrew Moore[2], Steven Gilmour[3] and Jonathan Pitts[1]

1 - Department of Electronic Engineering, Queen Mary, University of London, 2 - Computer Laboratory, University of Cambridge, 3 - Department of Mathematical Sciences , Queen Mary, University of London.

Part 1: Previous Track Record

This project proposal grows out of recent success in interdisciplinary research at Queen Mary, and the recognition that the problem domain - optimal design of network measurements - is necessarily a cross-discipline research problem. Our track record of tackling engineering problems by taking a multidiscipline approach includes the joint work of Pitts and Arrowsmith (see *Complexity in large-scale communications networks* below), and the EPSRC *Bridging the Gaps* award (announced last quarter 2007, EP/F033133/1 *Building a New Community: Modelling, Visualisation and Verification of Large Scale Systems*), to facilitate multidisciplinary research between the departments of Electronic Engineering, Engineering, Mathematical Sciences and Computer Science; in EP/F033133/1 Schormans is PI, and Pitts a CI.

Network measurements In a series of recent papers Schormans at al have explored the limits to measurement accuracy when performing packet level network measurements. Focusing on relatively simple scenarios it has been found that even for smooth, static traffic (the easiest case to measure) there are practical load limits beyond which measurement accuracy necessarily degrades so rapidly that measurements will have little (or no) practical value, [1]. Internationally this work was the first to begin to quantify this problem [2]. This experience of theoretical modelling is complemented by Moore, [3], whose work on practical experimental platforms for monitoring and measurement has received considerable EPSRC support, e.g. GRIDprobe (GR/S07759), UKLight-MASTs (GR/T10510/02) and 46PaQ (GR/S93714/02) (on which Schormans was CI at Queen Mary).

The Design Of Experiments (DOE) thread of the project is led by Prof Steven Gilmour. DOE has been widely applied to research in the Biological Sciences and many aspects of Industrial Engineering but it has not yet been significantly applied to networking. Gilmour heads The Statistics Group at Queen Mary, which has an international reputation for its work in DOE - a large proportion of the UK's experts are based at Queen Mary, including also Bogacka and Coad. The importance of this research, and the strength of the group, were recently recognised by the award of a £470K EPSRC grant to work on Unifying Approaches to DoE (PI: Gilmour, EP/C541715/1), and by the Defra sponsored project Design of Experiments Website (see the DoE Website http://biometrics.hri.ac.uk/experimentaldesigns/website/hri.htm).

DOE techniques have been very successfully applied in linear and static environments. Most work on nonlinear systems has so far assumed static processes, or at least deals only with static aspects of the processes, i.e. a set of inputs (treatments) is fixed by the experimenter for each and one or more outputs (responses) observed from that run. The first work on DOE for models which are solutions of nonlinear differential equations was published in the field of chemical kinetics, with one of the authors being in the Statistics Group at Queen Mary. Between 2002 and 2007 Gilmour has published 9 journal papers in the area of nonlinear design, e.g. [4]. This work includes finding optimal designs for estimating the stationary point on a response surface, the parameters in Michaelis-Menten and higher order enzyme kinetic models, random treatment effects and parameters in a fractional polynomial model. All of these are nonlinear models or nonlinear functions of the parameters in a linear model. Although the models involved are fairly small-scale compared with those arising in networks, having a single input variable and no feedback, this will provide an excellent starting point.

Complexity in large-scale communications networks Communications networks are an excellent example of large-scale complex systems (e.g. see *The Internet as a Large-Scale Complex System*, a volume in the Santa Fe Institute Studies in the Sciences of Complexity). It is well known that interactions between users, applications and the shared network all cause highly variable, complex behaviour to emerge. These aspects of network modelling have formed the basis of much of Pitts' earlier work, particularly with Arrowsmith and Mondragon (GR/L78659; GR/R83453; GR/R30136; EP/C520246).

To build parsimonious models of large-scale networks Pitts has worked on the use of non-linear dynamical (chaotic) traffic and networks models to map underlying behaviour of packet networks. This research is inter-disciplinary with Prof Arrowsmith - head of the Dynamical Systems and Statistical Physics Group, School of Mathematical Sciences, Queen Mary. Funding has been through a series of projects: 1) Chaotic Control for Fast Resource Management in ATM Networks (GR/L78659); 2) Sources of Long Range Dependence in Internet Traffic (GR/R83453); 3) Small-World Modelling of Internet Behaviour (GR/R30136) and 4) Topologies and Congestion Invariants (EP/C520246). (1) and (2) were inter-disciplinary between the Department of Electronic Engineering and the School of Mathematical Sciences at Queen Mary, and (2) was an interdisciplinary MathFIT project involving the Maths Dept. of University of York and BT Exact. Projects

(3) and (4) were directly funded by EPSRC. In this area of research BT has supported three CASE PhD student supplements (BT) and Lucent a PhD studentship.

Experimental platforms To bring together networks and DOE it is necessary to have an experimental platform. This platform must enable the modelling and testing of scenarios of sufficient complexity to be representative of real-world networks. In Pitts and Schormans, the investigatory team have considerable experience in techniques for building network simulation platforms, including the use of accelerated simulation techniques, in which Schormans has supervised 3 PhD's: Lu (2006), Ariffin (2006) and Ma (2003). Within the UK, Pitts has led research that is unique in its approach to the challenges of very large-scale networks. EPSRC project GR/T18615 has used high-performance computing to address benchmarking, representation and reproducibility in simulated networks of arbitrarily large scale; this was the first time UK supercomputers have been used for large-scale networks research addressing methodological issues.

The investigators

This project benefits from the considerable experience that the investigators have in working together.

Dr John Schormans (PI, Queen Mary) has worked in networks for around 20 years, focussing on probabilistic methods for simulation and analysis of packet based networks. Recently he has concentrated on the limits of measurability in packet networks, and this work has identified the critical relationships between a network's capacity, applied traffic load and burstiness, and the inaccuracy in the returned measurements, e.g. [1, 2]. The discovery of these relationships provided the impetus to investigate DOE for network measurement. He is CI on Moore's EPSRC project 46PaQ (focused on technological aspects of measurements), and is jointly supervising (with Gilmour) a Queen Mary sponsored PhD student researching the initial application of DOE to packet level measurements.

Dr Andrew Moore (PI, Cambridge) is at the Computer Laboratory, Cambridge University. Previously he was with the Dept of Computer Science at Queen Mary, working on new optically-switched networks, measurement-based management of network resources, and passive network monitoring and characterization. He is PI on the UKLIGHT/MASTS (Measurement at all scales in time and space) project to investigate instrumentation in the UKLIGHT network, PI on the 46PaQ project investigating monitor-based performance analysis for new network protocols and CI on the EPSRC GRIDprobe project, based at Cambridge University. Earlier he was an Intel Research Fellow at the University of Cambridge Computer Laboratory where he worked on network monitoring, performance and characterization topics for over 15 years. His PhD, at the Cambridge University Computer Laboratory, was on measurement-based network resource management.

Prof Steven Gilmour (CI, Queen Mary) has worked in the design and analysis of experiments for almost 20 years. His research interests were originally in industrial experiments, especially fractional factorial and response surface designs. While these interests have continued, he has expanded his interests to both other methodological topics and other areas of application, authoring 31 research papers in statistical methodology and 18 papers which apply statistical methodology to experiments. He is Editor of the *Journal of the Royal Statistical Society, Series C (Applied Statistics)* and a member of the Research Section Committee of the Royal Statistical Society. Prof Gilmour also heads the Statistics Group at Queen Mary, which has expanded in the last four years to become one of the UK's leading research groups in the subject: 7 academic staff, 2 post-doctoral RAs and 12 research students, who are in both the School of Mathematical Sciences and in the cross-faculty Centre for Statistics (also lead by Gilmour).

Prof Jonathan Pitts (CI, Queen Mary) has worked on the analysis, simulation and measurement of network performance for around 20 years. His contribution to this project is based on his experience in a) solving the problems associated with modelling network complexity, and in particular his interdisciplinary work with David Arrowsmith in Maths on nonlinear dynamical models for networks; b) the construction of simulation based experimental platforms, and most recently pioneering use of UK supercomputers for simulating large scale networks (GR/T18615/01). Pitts and Schormans have collaborated for two decades; Pitts is PI and Schormans CI on GR/T18615.

[1] Hasib M, Schormans JA and Timotijevic T: Accuracy of Packet Loss Monitoring over Networked CPE. IET Communications Vol 1, Issue 3, p. 507-513, June 2007.

[2] Timotijevic T. and Schormans J.A.: Bandwidth overhead of probe technology guaranteeing QoS in packet networks. IEE Electronics Letters, 15th May, 2003, Vol. 39, No. 10, pp 816-818.

[3] A.W. Moore and D. Zuev, Internet Traffic Classification Using Bayesian Analysis Techniques in the Proceedings of the ACM SIGMETRICS June 2005, Banff, Canada.

[4] Gilmour, S.G. and Trinca, L.A.: Fractional polynomial response surface models. Journal of Agricultural, Biological and Environmental Statistics, v. 10, 2005, 50-60.

Part 2: Proposed Research and its Context

A - Background

Optimal measurement of packet level performance (loss and delays) is a challenging open problem in engineering mathematics. All packet-level network measurements are numerical experiments, in which stochastic processes are sampled, and the sampling is constrained by the resources available (e.g. bandwidth). Current measurement methods are not actually designed to provide the maximum information from a minimal data set, and they can be prone to very large errors in estimated end-to-end delay (mean & jitter) and packet loss rates. Our project aim is to address this by applying the Statistical Design of Experiments (DOE) to network measurements.

Packet level measurement is critical to many aspects of broadband networking, e.g. guaranteeing Service Level Agreements (SLAs) between network and service providers and customers, measurement-based admission control algorithms and network tomography. Many current UK packet level measurement research projects are focused on the technological aspects of measurement. UKLight/MASTs (GR/T10510), 46PaQ (GR/S93714) and ESLEA (GR/T04465) are looking at monitoring in IP WAN packet networks and concentrate on the challenges arising from the scale of the technology itself: bandwidth in Gbps, large networks and storage space, and new network protocols and algorithms that all have to interwork. Moore has been involved in many key UK/EU funded initiatives (including the EU (IST) 6QM project [QM603], and UKLight/MASTs, 46PaQ, and ESLEA - Exploitation of Switched Lightpaths for e-Science Applications).

Recent work has addressed optimal inverting of sampled data, e.g. [Hohn06] in which the authors develop a flow-sampling method, and concentrate on stationary queueing models, looking at the distribution of the number of packets per flow, and the spectral density of the packet arrival process. Earlier related work can be found in [Duff05]. Other very recent research (October 2007) has addressed optimal packet probing patterns, [Bacc07], and discovered that traditional Poisson sampling is not always optimal, developing a more general class of distribution (Gamma renewal processes) that can minimise the mean-square error in the sampled data.

There is a growing awareness of the significance of the errors that arise when measuring network performance. The presentation by Filsfil, [Fils06], (a Cisco Systems engineer) at a Cisco hosted symposium on *Measuring Internet Quality* reports that magnitude of measurement error is very poorly understood at present, and that simple approaches to improving it are still focused on straightforward adaptations of existing ideas. The limitations inherent in simply increasing sample rate are very well discussed in the SIGMETRICS'05 paper of Roughan [Roug05]; he finds innate correlation to be the critical (and largely not-understood) limitation on simply taking more samples. Further work, [Roug06], gives a very thorough analysis and discussion of the possibilities available in varying the sampling pattern (rather than just the rate).

Despite a wide variety of measurement tools having been reported in the literature, there "…has been very little analysis of the accuracy of these tools or their impact on the network…" (quoted from [Somm05] at SIGCOMM'05). This has been partly addressed in a series of recent papers by Schormans et al, which has explored the inaccuracy inherent in packet level measurement, [Timo04], [Scho05], [Hasi07][1]. It has been discovered that even for static traffic in simple buffering scenarios there are practical load limits beyond which measurement accuracy degrades very rapidly. These loads limits are likely to be well within the normal operating specification of packet networks, e.g. 70% load on a VoIP access link. Precise numerical results are very situation specific. However results to date show that very significant levels of expected absolute error are inherent in sampled results for even a very simple example; e.g. mean delay in a VoIP access link (we have found absolute error in the measured mean delay can rise to 100's of msecs when a link at around 80% utilisation is probed over a full busy hour). Absolute error in measured mean values increases with the number of end-to-end links, load and traffic burstiness, and is generally (approximately) inversely proportional to bandwidth. Furthermore, it is much harder to measure tail probabilities (e.g. for packet loss probability or delay jitter) than it is to measure mean delays. One of the significant results of the work reported in [Hasi07] is that the *burstiness of the loss process* also has a critical effect, requiring far more samples than might otherwise have been allowed for.

It is possible that router technology will advance over the next few years to the point where a standard interface could be defined through which passive data (e.g. for packet queue lengths, losses and waiting times) can be made available externally. End-to-end measurements would then be available by combining data from all the nodes in the route across a path (there are already some hybrid methods that essentially try to do this now). However, from the point of view of this project any such situation is still a numerical experiment, in which data must be sampled and processed, and as pointed out in [Hohn06], generating traffic statistics does not scale well with bandwidth. Any such technique would still be optimised by best design of sampling experiments.

[1] *Also [Roug05] studies sources of inaccuracies due to sample correlation and initial-conditions.*

B - Programme and Methodology

Project objectives In accord with our overall aim, the objectives are:

OBJECTIVE 1 **To apply DOE to evolving experimental units**

Our domain of interest, packet based communications networks, is neither linear nor static, and for that reason a key objective is the application of DOE to evolving experimental units. This will require reparameterisation, so that specific changes in the experimental unit over time are represented as factors in the experiment and we will be interested in the effects and (possibly complex) interactions of several of these factors simultaneously. Objective 1 is expressed in Tasks 2, 3, and 4.

OBJECTIVE 2 **To develop an algorithm that evaluates the performance of a single target packet flow from the measured performance of the aggregate traffic**

Objective 2 is expressed in Task 5.

OBJECTIVE 3 **To build an experimental platform**

The success of objectives 1 and 2 requires a very precise control over the factors, levels and interactions. To enable the success of the 1st two objectives, we will build a simulation based experimental platform. Objective 3 is expressed in Task 1.

Programme of work

The following tasks constitute the programme of work, and are split between the RS, 2 PDRAs and the academic team. In summary the project is planned to proceed as follows: we employ 2 PDRAs; PDRA1 is for the first year only, to build the experimental platform (Pitts/Schormans); the second PDRA (Gilmour/Schormans) follows the first, and uses the platform PDRA1 developed as the tool for DOE application to networking scenarios. In parallel with this Moore, in Cambridge, has a single PhD student (Moore/Schormans) working on the problem of obtaining the performance of a single flow based on the measured performance for the whole of the traffic in aggregate. This is important: users experience quality of service through what happens to their data flow. Project management is described after the tasks.

TASK 1 Experimental platform for large-scale networks PDRA1, Pitts, Schormans, Month 1 – 12

This task is to build the platform for the simulation experiments. It is the only task for PDRA1, and uses supervisory input from Pitts and Schormans.

TASK 1 a) Construction of the experimental platform PDRA1, Pitts, Schormans, Month 1 – 9

We will take advantage of the existing experimental platform built for the project *Topology and Congestion Invariants in Global Internet-scale Networks* (EP/C520246/1), which currently consists of a controller (head) and 3 compute engines. In this way we are able to proceed with this project without the need to buy another controller, although we will be buying additional computer nodes and disk storage for the substantial volumes of measurements (see Justification Of Resources). We therefore concentrate on addressing the most critical problem associated with WAN simulation – ensuring that the simulation tool has enough RAM. Any form of swapping (virtual memory) severely degrades the speed of operation, and this project will avoid it by appropriate design and dimensioning of the platform.

A significant challenge both for the construction of the testbed, and in the application of design of experiment methods to network measurements, is inherent network complexity. Our solution is to take advantage of the results coming out of our current project *Topology and Congestion Invariants in Global Internet-scale Networks* (EP/C520246/1 – Mondragon PI, Pitts CI). Mondragon and Pitts have successfully developed a number of ideas that now permit us to determine which source-destination (s-d) pairs are significant, and focus attention accordingly (an important discovery from EP/C520246/1 is that most s-d pairs are not important from a traffic performance measurement perspective). The techniques provide parsimonious methods for network representation, including further developments of the idea of *Betweenness Centrality*. Mondragon recently presented a subset of these ideas and innovations at the Royal Society[2]. Additionally Schormans has considerable recent experience in developing aggregate traffic models, beginning with [Scho01] and including 3 recently completed PhDs (latest [SA2006]). Overall the importance of all this work to this project is that we can use it to cut down the network "state-space", allowing us to build (and experiment with) controlled-scale network models *that are a valid representation of the whole.*

TASK 1 b) Validation of the platform PDRA1, Pitts, Schormans, Month 9 –12

Recent projects (e.g. GR/T18615/01 and EP/C520246/1) are going to be used to provide a range of known scenarios against which the experimental platform can be validated.

---

[2] *"Networks: modelling and control" Organised by Dr Keith Briggs, Professor Frank Kelly FRS, and Professor Mike Smith,* see www.royalsoc.ac.uk/downloaddoc.asp?id=4514).

TASK 1 c) Deliverable 1: instructions for use of the experimental platform PDRA1, Pitts, Month 11 – 12

The operating instructions will be made available to the whole team and particularly PDRA2 at end of year 1.

TASK 2 Response Surface Design PDRA2, Gilmour, Schormans, Month 12 – 24

Some work applying statistics and DOE to network measurements has already begun. With Queen Mary funding Gilmour and Schormans have been jointly supervising a PhD student (Ben Parker), and interesting results have emerged that are significant for this Task. Concentrating on measuring packet loss probability caused by buffer overflow, a particular initial problem was the identification of maximum likelihood estimators. A barrier to the effective introduction of DOE techniques arose because the analysis of the time-dependent behaviour in packet queues is extremely challenging, even using the combinatorial methods originally developed by Takacs. The solution we have adopted has been to use the first-order approximations that Schormans has worked on for many years (46 papers and 4 PhD theses to end 2007). [Hasi07] showed how to parameterise a finite buffering system as a Markov transition process. Seeking the minimum variance estimate of Θ (the Matrix representation of p, r – the probabilities of entering/leaving the buffer overflow state) – a network measurement system should observe the buffer every k^{th} time instant. So optimal probing is achieved (for an allowed fixed number, T+1, of transitions) by optimising 'k'. Given the underlying process (the buffer) has been active long enough to have reached steady-state, the Likelihood is:

$$L(p, r|\mathbf{y}) = P(Y_0 = y_0, Y_1 = y_1, \ldots, Y_T = y_T) = P(Y_0 = y_0).P(Y_1 = y_1 \,|Y_0 = y_0) \ldots P(Y_T = y_T \,|Y_{T-1} = y_{T-1} \ldots, Y_0 = y_0)$$

And the Fisher Information matrix is:

$$I(p, r|\mathbf{y}) = \begin{pmatrix} E[\,(-\,d_2 log(L(p, r|\mathbf{y}) / dp^2\,] & E[\,(-\,d_2 log(L(p, r|\mathbf{y}) / dpdr\,] \\ E[\,(-\,d_2 log(L(p, r|\mathbf{y}) / dpdr\,] & E[\,(-\,d_2 log(L(p, r|\mathbf{y}) / dr^2\,] \end{pmatrix}$$

which in this case is:

$$I(p, r|\mathbf{y}) = \begin{pmatrix} r / p^2(p+r) - 1 /(p+r)^2 + (Tr/(p+r)).(1/p + 1(1-p)) & -\,1/(p+r)^2 \\ -\,1/(p+r)^2 & p/ r^2(p+r) - 1 /(p+r)^2 + (Tp/(p+r)).(1/r + 1(1-r)) \end{pmatrix}$$

From this point there are a number of ways of proceeding to determine the optimum loss probing strategy for the buffer. This analysis can be extended to account for the bandwidth absorbed by the measurement process in active probing. With this in mind, this first task of PDRA2 will be to investigate response surface designs for both linear and non-linear scenarios, and their application to network measurement.

TASK 3 Reparameterisation PDRA2, Schormans, Gilmour, Month 18 – 30

An initial problem is the identification of the experimental unit. In using the experimental platform, the natural unit is a single complete run, whose inputs will be under our control, and can therefore be kept constant when using the experimental platform, so that they represent the treatments whose effects can be modelled. However in general they will not be fully under control in "live" networking situations. An unusual aspect of this task, which distinguishes it from standard response surface methodology, is quantifying the effect of evolving experimental units on the responses. An obvious idea is to break down the runs into small sub-runs over which the inputs can reasonably be expected to remain constant. However, this leads to highly complex correlations among experimental units, which invalidates, or greatly reduces the robustness of, most of the usual methods of DOE. Some initial work on the effect of these correlations has already been carried out by Parker, Gilmour and Schormans.

Our key insight, which we believe to be original, is that the complete runs should still be considered to be the experimental units, but the evolving inputs should be reparameterised in terms of aspects of a dynamic model which describe their evolution over time. Then each parameter of the dynamic model is a factor and specific combinations of these parameters are the treatments which can be applied to the experimental platform to mimic the behaviour seen in live networks. An important aspect will be to identify suitable levels for these factors in particular models.

TASK 4 Optimal Design for Evolving Experimental Units PDRA2, Gilmour, Schormans, Month 12 – 33

Key factors will certainly include the number and type of any applications using the network path whose performance is being measured, and the resulting overall load on the network's nodes and links. The appropriate levels to use will be harder to pin down, as the level of granularity will be small and will become ever finer as end-to-end bandwidth is increased. On the other hand, it will be advantageous to us that measurement is most critical over bottlenecks, and here the granularity (e.g. number of active applications active at time 't') will necessarily be coarser, reducing the number of possible, experimentally valid and interesting, levels.

Interaction - A significant complication in the design of any experiment is the magnitude and significance of interaction between factors. Most prior use of DOE in computer systems design has been able to reduce these effects by limiting the size and scope of the experimental unit. This will not be easy with live measurements on complex broadband networks: the end-to-end paths will vary in length from the very small to the genuinely global, much of the traffic will be interactive and elastic and route selection is responsive to network load. The overall state of complex networks is affected by considerable interactivity: (i) for many applications, the rate of network usage will be dependent on the network congestion; (ii) user intensity will depend on availability and congestion; (iii) there will be long term "memory" in some of the underlying processes driving network behaviour, i.e. networks optimised to regimes of operation that are correlated over many timescales. Because of these characteristics we will make extensive use (see also Section A - *Background*) of the new ideas of Mondragon and Pitts that allow considerable topological simplification.

Given all these considerations, finding optimal designs for running the experimental platform requires new statistical methodology, even though the reparameterisation brings it closer to standard DOE.

TASK 5 Flow-based measurement RS, Moore, Schormans, Month 1 – 36

Cisco RFP-2007-009 (http://www.cisco.com/web/about/ac50/ac207/crc/rfp.html#009) states that "*There are few tools that relate end-user Quality of Experience (QoE) to measurable network behaviour. Historically, keeping per-flow information that might help connect individual resource consumption with QoE has been prohibitively expensive*", and single flow performance is noted as an unsolved problem.

Some recent work [Ying05] studies static buffer scenarios, specifically for loss probability; however this work depends on simple queue models as a means of traffic and performance modelling. Our approach is to build on Moore's experience in using machine learning techniques for packet / flow classification, and re-orient it towards learning the traffic characteristics that are critical in their influence on relative delay/loss performance: distributions of both packet inter-arrival times and packet length (per application type). This will allow us to work with the real distributions rather than with queue models that have inbuilt (and limiting) distributions chosen to provide analytic tractability.

The classification of network traffic has previously used features derived from streams of packets; such feature collections are often huge (200+), and can range in complexity from Fourier-Transformations and quartile statistics to mean and variance of packet inter-arrival time and the number of TCP SACK packets. Classification accuracy has been shown to be very good (e.g. >95%); however the disadvantage has been the complexity and costs associated with the collection, aggregation, and generation of the desired packet-derived features. Moore's prior experience is with lightweight application classification schemes, and in particular, supervised machine learning techniques, [Moor05, Jian07].

The machine learning is phase 1 of Task 5 – the 2nd phase of the work will be to parameterise application-based aggregate models. This builds on prior work of Schormans, who with PhD students (e.g. [Ma2003], [Leun03] and [SA2006]) developed a Markovian approach for parsimonious traffic representation. This earlier work designed a scheme for evaluating the parameters for strictly homogenous traffic sources, so the vital new step here is to use the flow statistics derived in the first phase of the Task, and re-apply these per flow/application.

The 3rd (final) phase of Task 5 is to evaluate the *relative* performance of the application flows via queue modelling. Here the simplicity of the Markovian aggregation technique is paramount: a generic tool like MATLAB is quite powerful enough to support these models and examine the resulting state space.

Task 5 is the sole task for the project PhD student. As both Moore and the student are to be based at Cambridge, the investigators need a strong culture of inter-working. This is in place; it was developed while all 4 were at Queen Mary. Regular meetings are planned, and described in the Project Management section of this proposal (see also Diagrammatic Project Plan).

TASK 6 Designing measurement experiments Schormans, Moore, PDRA2, Gilmour, Pitts, Month 33 – 36

This Task is a project summary activity, and draws together the findings of Tasks 2, 3, 4 and 5 for the benefit of the networking industry. The Task outputs will be:

- recommendations on how to design network measurement experiments
- the theoretical limits on the accuracy of the measurements, and
- recommended contents of the SLAs that guarantee measured performance bounds to customers.

This Task concludes with the completion of the 2nd project deliverable, *Designing Measurement Experiments*.

Project Management

This proposal is for a large project, requiring, in addition to the PI, 2 PDRAs, a PhD student and a significant input from 3 CIs. It is important that the investigators have a strategy to ensure that the project retains coherence, and this is summarised as follows:

YEAR 1: We plan weekly meetings between Schormans, Pitts and PDRA1; monthly meetings between PhD student, Moore and Schormans (normally at Cambridge).

YEAR 2: The platform development will have concluded, so Pitts not formally required again until year 3, but remains available ad-hoc; weekly meetings between Schormans, Gilmour and PDRA2; monthly between PhD student, Moore and Schormans.

YEAR 3: In year 3 we plan 2 meetings per month between Schormans, Gilmour and PDRA2; monthly meetings between Schormans, PhD student and Moore. Pitts used again for last ¼ of the 3rd year (see Diagrammatic Project Plan).

ALL 3 YEARS PI to call status meetings of all the team every other month (see diagrammatic project plan). In addition to these formal project management meetings, the investigators will invite members of the wider community to regularly meet for off-site brainstorming.

Timeliness, Novelty and Adventure

Novelty in this proposal arises as the result of getting beyond the implications and challenges associated with the measurement technology itself. In regarding all network measurement (whatever the tool or algorithm) as examples of numerical experiments, we have opened the way to their optimal design. If the project works out as planned, we will have solved the critical problems of the design of these experiments.

This is an adventurous idea; while DOE has been applied extensively in static, and rather simple problems in computing (indeed Raj Jain wrote an excellent book on the subject[3]), it has never before been applied to a domain as complex and non-linear as that of wide-area networking. There is no guarantee that this will work.

There are a number of ways in which this project is timely. 1) The "traditional" best-effort model of packet data networking is now widely becoming seen as not adequate to support extensive networked e-commerce; SLAs require accurate measurement to support and guarantee them. 2) In the last 5 years analytic work on measurement for queue models of buffers (Schormans, Roughan etc) has shown that such measurements are likely to be inaccurate for a number of reasons. 3) In the last 2 years the results of investigations into the topological structure of networks (e.g. at Queen Mary under Mondragon, Pitts) has shown that we can represent very large and complex networks in a parsimonious fashion; this allows us to develop platforms for network modelling that are manageable, and yet still representative of the whole.

C - Relevance to beneficiaries

The academic networking research community will benefit from this project as it will open a whole new approach to designing network measurement experiments, e.g. see our Letter of support form Prof Matthew Roughan. The DOE community will benefit from the opportunity to extend statistical DOE into a novel application area, a complex example of DOE for non-linear systems.

There is also great potential for commercial significance. It has long been understood that accurate measurement is essential to ensure networked IT applications are properly supported, [Cisc05], [Cisc07], [Mort04], [Verm99]. The outputs of this project are of potentially huge importance to the commercial future of networking; see our Letter of Support from Dr Ramakrishnan of AT&T.

D - Dissemination and exploitation

To help disseminate the results of this project we will hold a workshop, which can be funded under the current Queen Mary EPSRC *Bridging the Gaps* project (EP/F033133/1 *Building a New Community: Modelling, Visualisation and Verification of Large Scale Systems*). The purpose of this workshop is twofold: to raise awareness of the fundamental measurement accuracy problem amongst industrial colleagues (some of whom will already be serving on the Industrial Liaison Committee of EP/F033133/1); additionally we already have some significant expressions of interest in attending such a workshop, including Willinger from AT&T Labs, and from network engineers based at Motorola (Jose Gil), Vodafone (Martin Russell) and Cisco (Keith Jones).

Another colleague who is a co-investigator in *Bridging the Gaps* (EP/F033133/1) is Prof Peter McOwan (Dept of Computer Science, Queen Mary); Peter has made very significant contributions to the public understanding of computer science (see http://www.dcs.qmul.ac.uk/~pmco/), in particular through his work on Computer Science for Fun (CS4fn). This project proposal - Optimal Design of Performance Measurement Experiments – would lead to a very interesting opportunity to add to CS4fn that we intend to take advantage of. We foresee this as follows: many people are now taking the opportunity, created by computer networking,

[3] *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modelling*

to play internet-based games. Such games require fast download speeds else we are confronted by the cry of the frustrated teenager that World Of Warcraft just isn't functioning properly. How does any user know if they are receiving the download speed they believe they are paying for – via packet level measurement, which is the subject of this project.

Other than this, dissemination of research findings will be via the usual route of journal papers and conferences. The conferences that are particularly important here are SIGMETRICS and SIGCOMM. These are not normally held in the UK; within the DOE area, we plan to present at the Annual Spring Research Conference (SRC) on Statistics in Industry and Technology. The SRC is an annual conference jointly sponsored by the Institute of Mathematical Statistics (IMS) and the Section on Physical and Engineering Sciences of the American Statistical Association (ASA/SPES), and held at Iowa State University in Ames, Iowa, USA. We have therefore allowed for some international travel in the Justification for Resources.

References

Bacc07 Baccelli, F, Machiraja, S, Veitch, D and Bolot, J: *On Optimal Probing for Delay and Loss Measurement*. Proc. IMC'07, San Diego California, October 2007. Copyright ACM.

Cisc05 *Service Level Management: Best Practices White Paper*. Cisco White Paper, 2005, Available: http://www.cisco.com/warp/public/126/sla.pdf

Cisc07 *Measuring delay, jitter, and packet loss with Cisco IOS SAA and RTTMON*. Available: http://www.cisco.com/warp/public/126/saa.pdf

Duff05 Duffield, N, Lund, C and Thorup M: *Estimating Flow Statistics from Sampled Flow Distributions*. IEEE/ACM Transactions on Networking, Vol.13, No.5, Oct 2005.

Fils06 Filsfils C and Ward D: *Measuring Internet Quality*, Available: http://www.cisco.com/web/about/ac50/ac207/proceedings/MIQ.ppt

Hasi07 Hasib M and Schormans JA and Timotijevic T: *Accuracy of Packet Loss Monitoring over Networked CPE*. IET Communications Vol 1, Issue 3, p. 507-513, June 2007.

Hohn06 Hohn, N and Veitch D: *Inverting Sampled Traffic*. IEEE/ACM Trans. on Networking, Vol.14, No.1, Feb 6.

Jian07 Jiang H, Moore A, Ge Z, Jin S and Wang J: *Lightweight Application Classification for Network Management* in the Proceedings of the SIGCOMM Workshop on Internet Network Management 2007: The Five-Nines Workshop, August, 2007, Kyoto, Japan.

Leun03 Leung C.M.: *Non-intrusive measurement in packet networks and its application*. PhD thesis, University of London, 2003, supervised by J A Schormans.

Ma2003 Ma, A.H.I.: *Accelerated Simulation of Power-Law Traffic in Packet networks*. PhD thesis, University of London, 2003, supervised by J A Schormans.

Moor05 Moore A and Zuev D, *Internet Traffic Classification Using Bayesian Analysis Techniques*. Proceedings of the ACM SIGMETRICS June 2005, Banff, Canada.

Mort04 Morton, A: *IP/MPLS Network Performance and QoS*. IEE Telecommunications Quality of Service (QoS2004), Savoy Place, London, March 2004.

QM603 http://www.6qm.org/overview.php

Roug05 Roughan M.: *Fundamental Bounds on the Accuracy of Network Performance Measurements*. SIGMETRICS'05, June 6–10, 2005, Banff, Alberta, Canada. Copyright 2005 ACM.

Roug06 Roughan M.: *A Comparison of Poisson and Uniform Sampling for Active Measurements*. IEEE Selected Areas in Communications, VOL.24, NO.12, Dec. 2006.

SA2006 Ariffin, Sharifah: *Accelerated Simulation for non-FIFO Scheduling in Packet networks*. PhD thesis, University of London, 2006, supervised by J A Schormans.

Scho01 Schormans, J.A., Liu, E., Cuthbert L.G. and Pitts, J.M.: *A hybrid technique for the accelerated simulation of ATM networks and network elements*. ACM Transactions on Modelling and Computer Simulation, Vol.11, No.2, April 2001, pp182-205.

Scho05 Schormans, J.A. and Leung, C.M.: *Measurement for guaranteeing QoS in broadband multiservice networks*. Invited chapter (8) in "Mobile and wireless systems beyond 3G: managing new business opportunities", ISBN 159140544-0.

Somm05 Sommers, J, Barford, P, Duffield,N and Ron, A: *Improving accuracy in end to end packet loss measurements*. SIGCOMM'05, August 21-26, Philadelphia, USA.

Timo04 Timotijevic T., Leung C.M. and Schormans J.A.: *Accuracy of measurement techniques supporting QoS in packet based Intranet and Extranet VPNs*. IEE Proceedings Communications Special Edition on VPNs and Broadband Technology Vol. 151, No.1, Feb. 2004. Special Edition on VPNs and Broadband Technology.

Verm99 Verma, D.: *Supporting Service Level Agreements on IP Networks*. MacMillan Technical Publishing, ISBN 1-57870-146-5. 1999.

Ying05 Ying, X and Guerin R: *Individual QoS Versus Aggregate QoS: A Loss Performance Study*. IEEE/ACM Transactions on Networking, Vol.13, No.2, April 2005.