Performance Measurement and Management for Two-Level Optimization of Networks and Peer-to-Peer Applications

J. Sventek[§], L. Mathy[†], D. Hutchison[†], H. De Meer[‡]

[§]University of Glasgow
Department of Computing Science
joe@dcs.gla.ac.uk

[†]Lancaster University
Computing Department
{laurent, dh}@comp.lancs.ac.uk

[‡]University College London
Department of Electronic & Electrical Engineering
h.demeer@ee.ucl.ac.uk

1 Investigators' Track Records

University of Glasgow

The Networked Systems Measurement and Control (NSMC) Group in the Department of Computing Science, University of Glasgow, consists of 5 academics, 2 research associates and 8 research students. The department was rated 5 in RAE 2001. The group conducts research into architectures, models, algorithms, measurement, and control of networked/distributed systems. The group is especially interested in large-scale systems, based on both wireless and wireline interconnection technologies, as well as high-performance system interconnects.

Prof. Joe Sventek obtained his B.A. in Mathematics from the University of Rochester and his PhD in Nuclear Chemistry from the University of California. He is currently the Professor of Communication Systems in the Department of Computing Science at the University of Glasgow. Prior to joining Glasgow, he had a distinguished career pursuing research into networked and distributed systems and managing research teams at Lawrence Berkeley Laboratory (1979-1986), Hewlett-Packard (1987-1999), and Agilent Technologies (1999-2002). His research interests include programmable networks, embedded systems, closed-loop network management, and distributed system architectures. He has several publications on these topics, as well as holds four patents (with three other patents pending) in these particular areas. He has been the general chair for TINA99 and Middleware 2001, programme chair for COOTS98, TINA99, and Middleware 2000, and a member of programme committees too numerous to mention. He is an advisor to the TeleManagement Forum Board, is an adviser to the Wiley Series in Communications Networking and Distributed Systems, and was on the editorial board of the IEE/BCS/IOP Distributed System Engineering Journal.

Professor Sventek led the Communications Solutions Department in Agilent Laboratories in a focused research programme of active measurements of IP-based network flows and using those measurements to drive proactive network management activities, such as "real-time" traffic re-engineering; this experience will be critical to the successful design and deployment of the network-level measurement system. Glasgow/NSMC is also a partner in a proposed EU 6th Framework STREP Project which is focused on embedding such active measurements of flows across broadband access networks onto the line cards of network elements for measuring flow QoS performance parameters.

Lancaster University

The Networking and Distributed Computing group is well known for its work on communications protocols and services, and in middleware developments, with a strong emphasis on Quality of Service (QoS) and on multimedia and multipeer applications. Lancaster University was a pioneer in the area of QoS architectures during the mid 1990s, and in more recent years has done leading edge work on multimedia caching, multimedia indexing and content management (MPEG-7), IPv6 and mobility, and application level multicast and overlay networks. The Department has considerable support from the EPSRC (the UK national research council), from the European Commission, and directly from industry.

Lancaster has a strong interest and a long history in the design and implementation of programmable network devices. Indeed, the group has been responsible for the design of three programmable router architectures, namely LARA [1] (Linux kernel-based programmability), LANode (combination of application-level and network-level programmability - http://www.activenet.lancs.ac.uk/lanode), and LARA++ (network programmability in Windows 2000 - http://www.landmarc.net/people/stefan/LARA++.ps). In the Alpine project (funded by BTExact), Lancaster was concerned with the design, using application-level active networking techniques (ALAN), of a scalable distributed algorithm for the control and self-organising maintenance of application-level overlay networks.

Dr. Laurent Mathy is a Lecturer in the Computing Department at Lancaster University. He spent the 1995-1996 academic year at the Center for Integrated Computer Systems Research (CICSR), the University of British Columbia, Vancouver, Canada, as a visiting scholar. He also was a research engineer in the Research Unit in Networking (RUN) of the University of Liege, Belgium, from 1993 to 1995. He was awarded his Ph.D. in Computer Science from Lancaster University in January 2000. His research interests include multimedia networking, programmable networks, overlay structures and network control technologies. He was Lancaster's research leader on the BT URI Alpine project concerned with advanced programmable network technology and services. He has served on the Technical Program Committee of several workshops and conferences in the field of Networking (ACM Multimedia, NGC, IWQoS, EuroMicro, MIPS (formally IDMS-PROMS)) and has served on the steering committee of MIPS. He has published many papers in the field of networking for multimedia and is a guest editor of a forthcoming special issue of Computer Communications on "Network Support for GRID Computing". He is the recipient of the Young Researcher Award of CFIP'99.

Prof. David Hutchison is Professor of Computing at Lancaster University and has worked in the areas of computer communications and distributed systems for more than 18 years. He has completed many UK and European funded research contracts and published over 125 papers as well as writing and editing books in these areas. His research interests cover architecture, services and communication protocols for distributed multimedia systems, including Quality of Service (QoS) mechanisms, programmable networking and also multimedia content distribution networks. He participates in several UK and European collaborative projects, in which an integrating theme is QoS for multimedia communications. He is a programme committee member for many international conferences and workshops, including IEEE Infocom and ACM SIGCOMM, IEEE OpenArch, IWQoS, NGC and QofIS, and is a member of the editorial board of the ACM Multimedia Systems Journal, the Computer Networks Journal and Computer Communications.From mid 2001, he has been a member of the UK Grid Technical Advisory Group, and chair of the Grid Network Team responsible for advising on the communications infrastructure to be used by e-Science applications funded by the DTI and the UK Research Councils.

University College London

The Networking and Distributed Systems group in the E&E Eng. Department of UCL (EE-UCL) was founded in August 2001 and is interested in IP-Networking, Peer-to-Peer Networks, Home Networking, Quality of Service and Performance Evaluation and consists of 8 doctoral students, one research associate and several MSc and MRes students in addition to the head of the group. EE-UCL has consistently been one of the top-rated EE Departments in the UK and has a track record of excellence in a broad range of disciplines. The Department and the research group has been and continues to be involved in research funded by a number of sources: direct industrial funding (BTExact), funding from the UK Government (EPSRC), and funding from European research programs for projects like Fabric, Fain, Tequila, or Android.

Dr. Hermann de Meer has led several national and international projects in Modelling and Performance Evaluation, Networking and Distributed Systems. His main focus of research is on Active Networking and Programmable Networks, Quality of Service for multimedia communications in heterogeneous networks, quantitative methods for Performance Evaluation, Peer-to-Peer and Ad Hoc Networks, interoperability of Home Networks and self-organizing systems. He has been awarded research grants from EPSRC, Deutsche Forschungsgemeinschaft (DFG) and the European Union. Dr. De Meer is currently appointed as Reader for Communication Networks and Distributed Systems at the Electronic and Electrical Engineering Department of UCL. He has been Assistant Professor at Hamburg University, Germany, Visiting Professor at Columbia University in New York, U.S.A., and Research Fellow of DFG. He has published more than 50 fully reviewed papers at conferences and in journals and is co-authoring a book on Queueing Networks and Markov Chains by John Wiley. He has served on many programme committees such as IWQoS, IEEE OpenArch, QofIS, IWAN or ITC and peer reviewed for journals and magazines like IEEE Internet, ACM Multimedia Systems Journal, Journal of Network and System Management, IEEE Computer, Intern. Journal of Parallel and Distributed Systems and Networks, Performance Evaluation or Transactions on Networks.

Dr. de Meer has been largely involved in Application Level Active Networking and Peer-to-Peer Networking throughout the last three years. This research has been funded by BTExact and has resulted in several international publications and invitations to give tutorials at international conferences such as ITC 2003. In addition, he has been involved as a co-investigator in an EPSRC funded project on Language Engineering for Programmable Networks, which has also resulted in several publications and which has attracted industrial sponsorship as well (Eastman Kodak and BTExact). Further research on Programmable Networks dates back to 1998 and 1999 when Dr. De Meer was associated with the Comet group at Columbia University, which had spearheaded research on Programmable Network throughout the late Nineties, and was working on the Genesis project. Dr. De Meer's research group has also a long standing history in Performance Modelling and Optimization.

2 Motivation

Peer-to-Peer (p2p) networking is a new, evolving communications paradigm aimed at building large distributed applications focused on the ideas of community or common interest; these applications do not have any notions of centralized management and are removed from the traditional client/server paradigm by accepting that every device in the peer group is symmetric [2]. Any device that can generate a digital heart-beat, including desktop computers and servers, PDAs, cell phones and other connected devices, could form part of a peer group and allow its resources to be shared amongst its members to deliver a common service or deploy a common application of interest.

Any community of like-minded users can establish a peer group, exploiting as many (or as few) of their systems to deliver the desired service to the group. This grass-roots ability to easily exploit the Internet in new ways without operator involvement has led to a steady increase in the number of p2p application systems appearing, and to increases in the volume of network traffic generated by such applications.

Similar to *ad hoc* mobile systems, peer groups can be structured into dynamically evolving communications topologies, known as overlay networks, over an underlying communications infrastructure. Peer group members come and go at any time, potentially appearing at different points of attachment. This will have an impact on the services provided and hence applications built upon these services need to be tolerant of such dynamic membership.

These inherent dynamics in p2p applications have consequences for Internet Service Providers (ISPs), whose clients use p2p applications, as the ISPs can then experience rapidly varying traffic patterns, which can adversely impact the network performance perceived by all applications (p2p and others). Traditional network management configuration tools support quite static forms of network engineering, usually based upon offline analysis of accumulated traffic matrices; the rapidly varying traffic patterns expected from p2p applications are not addressed by such tools. Consequently, network operators are in need of more dynamic tools that provide fine control over the flow of (p2p) traffic within their network [3,4].

One promising approach that has been proposed for providing this control at the application layer [5] is based upon Application-Level Active Network (ALAN) concepts [6]. The fundamental components of an ALAN system are “proxylets” and execution environments for proxylets (“EEPs”). Proxylets represent dynamic code modules which specify protocol handling. Single copies of these are stored on a central server and are identified using a URL reference. Communication between proxylets is achieved via remote method invocation (RMI). The EEPs acts as proxies and are therefore located at strategic points in networks (mostly access and edge points) and represent separate Java virtual machines allowing proxylets to be loaded, run, modified and stopped. Proxylets appear to represent micro services like application-level routing, discovery, available capacity etc., which can be combined together to build a dynamic overlay of interconnected application proxies. In short, the proylets provide the basic hooks for establishing and maintaining (through the use of application-specific proxylets) an ISP controlled overlay mesh connecting the members of a p2p community, while not only understanding the semantics of the data transfers but controlling the associated traffic flows.

In the more traditional approach [7], programmable network elements in the physical network can be used by a network operator to dynamically re-engineer the traffic to meet performance objectives for both the p2p community and other network users. Such use is integrally linked with QoS measurement techniques driving analysis components in the network management system to assert control over the programmable network elements. Such programmable network elements will likely be deployed at network edges and boundaries, as well as possibly at a few strategic points inside networks.

While programmable network elements operate at the network layer by interacting directly with network level protocols (e.g. routing protocols) and making decisions based on traffic aggregates, the ALAN approach operates at the application level by interacting directly (and possibly transparently) with applications and making decisions based upon application traffic semantics while taking a "black box" view of the underlying physical network. Both application-level and network-level programmability have specific advantages and we therefore believe that the combined used of both approaches has the potential to provide powerful control methods over highly dynamic network traffic. Such a two-layer programmability approach for performance measurement, management and optimization in networks supporting p2p applications is the focus of our proposed research.

More specifically, the simultaneous use of both approaches opens many integration issues. The most obvious of these is that if each programmable layer was to operate independently of the other, there is every likelihood that the two approaches may work at cross purposes, potentially introducing network instabilities.

Even if the two programmable layers are aware of each other and coordinate their actions, little is known regarding the correlation of p2p application-level traffic and the network-level traffic that it generates. Indeed, for example, timescales of traffic variability will be different in each layer; impact of changes in traffic patterns in one layer on the traffic load in the other layer is not well understood and therefore difficult to predict, etc. While the GridProbe project (http://www.cl.cam.ac.uk/~jac22/out/gridprobe.pdf) may eventually shed some light on this traffic correlation for traditional grid applications, it will not be able to measure, for instance, the impact of ALAN-triggered rerouting.

Finally, most studies of p2p or any other network traffic in general require the capture and analysis of complete traces of network traffic. However, capture, analysis and real-time synthesis of high-speed, high-volume traffic is difficult to achieve in an unobtrusive way and thus requires specialist techniques and devices. Despite recent advances in traffic measurement technology, much still needs to be done in order to build a satisfactory large-scale measurement infrastructure. In order for a network to react to the onset of p2p application behaviour, not only must the generated traffic patterns be known, but edge network elements must be able to detect the pattern in real-time. Innovative active measurement techniques, implemented in programmable edge devices, will be used to detect the onset of p2p behaviour.

It is important to note that while the correlation between application-level and network-level traffic will be determined based upon closed-world experimentation (i.e. all applications and all network edge nodes are measured), the application of these results is intended for an open-world scenario. The ability, at a single monitoring point, to detect the onset of p2p behaviour at the network level obviously depends only upon the measurements at that point and the knowledge of patterns that indicate onset of p2p activity.

3 Approach

We will measure application-level traffic patterns of selected p2p applications and compare these against measured network traffic patterns captured while the applications are active.

This correlated data will permit a number of areas to be explored:

- We will investigate integration issues in a two-layer programmable architecture; the goal is to understand how to synergistically optimize both programmable application and network levels.
- We will experiment with overlay optimization techniques within an ALAN infrastructure; the measurement system will permit us to quantify the optimization results.
- We will be able to detect network traffic patterns that indicate the onset of p2p activity.
- We will experiment with active measurement techniques in programmable edge devices to detect these patterns in real-time, and use such a detection capability to drive dynamic traffic re-engineering techniques.
- We will explore using this deeper knowledge regarding p2p traffic effects on network traffic to re-engineer the network measurement system as an optimized, p2p application

It is important to note that this project is focused on exploitation of network-level and application-level programmability, not the invention of another programmable platform. As such, we will be leveraging existing platforms (ALAN [8], LARA++ [1], and IPv6 extension header-based measurement [9]).

While much of the research will depend upon instrumentation and measurement of real applications on real networks, we also intend to perform large-scale simulation and large-scale emulation to further amplify the results of our work. The real experimentation will initially take place between the partner sites over the JANET network. Through our industrial collaboration with UKERNA (the JANET operator), we hope to add additional instrumentation to the JANET network after we have proven the efficacy and safety of our measurement system through experimentation on ourselves.

4 Research Goals and Issues

The overall research goals for the project are:

i. Correlate application-level and network-level traffic to detect network-level signatures that indicate onset of p2p activity
ii. Use programmable network elements to implement always-on, active measurements of the network-level traffic to detect these patterns in real-time and trigger appropriate traffic engineering to ensure acceptable performance levels in the network.
iii. Analyze the effect of the overlay network structure on the measured network-level traffic to drive optimization of the overlay network structure
iv. Understand how to synergistically optimize programmable network and application level structures simultaneously.

Research issues

The measurement system provides input and feedback to both the application-level and network-level programmable decision points. As already pointed out, the correlation of interactions of traffic at these two levels is an important research issue. This is particularly true from a stability point-of-view, as timescales of variability are likely to be different at each of the two levels.

Another very important research issue is the operation of a dynamic two-layer optimization approach. For such an approach to be successful, each layer must export a concise representation of its view of the "network" (i.e. either the physical network or the overlay network) to the other layer, so that effects of optimization can be predicted and synchronized to provide a stable communication system. Both the issues of concise network representation and integration of dynamic optimization techniques in a two-layer programmable network architecture are still largely open. In particular, we will focus on predictive dynamic optimization, which requires techniques that ensure minimum disruption to the operation of the network to achieve maximum operational stability.

We therefore see that our approach will not only explore three distinct open research issues (measurements, and performance optimisation in both physical networks and overlays), but that it will also focus on their integration in a programmable network environment. This, we believe, fully justifies the presence of the three academic partners, each bringing expertise in one of the research issues.

5 Work Packages

The primary activities of the project will be:

a. A measurement system for network-level traffic resulting from p2p activity and anonymized traces from the operation of that measurement system; instrumented ALAN middleware yielding application-level traffic from p2p applications.
b. Correlation of application-level and network-level measurements from selected p2p applications to yield network traffic patterns indicating onset of p2p activity.
c. Exploitation of programmable network techniques to detect onset of p2p activity in real-time and to use such indications to drive network traffic engineering algorithms
d. In-depth investigation of application-level overlay optimization techniques such as predictive overlay topology modelling and self-organization, with the network measurement system providing true feedback on the efficacy of a particular overlay network structure.
e. Investigation of integration issues between network-level and application-level optimizations.

These will be realised through the following workpackage structure.

WP1 Measurement Systems, Measurements, and Correlation

- Design, deploy, and operate a monitoring system to collect network-level traces. The measurements must be able to drive the network-level traffic engineering techniques in WP2.
- Instrument ALAN middleware to provide application-level performance measurement data for correlation. The measurements must be able to drive the overlay network optimization research in WP3.
- Integrate ALAN middleware into selected p2p applications in order to drive the research.
- Correlate application-level measurements with network-level measurements to detect network-level traffic patterns indicative of onset of p2p activity.
- Working with our UKERNA partner, potentially deploy the monitoring system at appropriate points in the JANET network.

Deliverables:

Month 6: Report on the design of the network-level and application-level monitoring system.

Month 12: Operational network-level and application-level monitoring systems.

Month 30: Final report on measurement systems, measurements, and correlations.

WP2 Programmable Network Techniques for Detection and Traffic Engineering

- Investigate, design and implement active measurement techniques for monitoring the onset of p2p behaviour in real-time.
- Investigate and design network traffic engineering techniques in a programmable network environment.

Deliverables:

Month 18: Report on active measurement techniques applicable to detection of p2p onset and applicable network traffic engineering techniques.

Month 30: Design, implementation, and deployment of active measurement techniques and network traffic engineering techniques on the LARA++ platform.

Month 36: Final report on programmable network techniques for p2p onset detection and traffic engineering techniques.

WP3 Application-level overlay optimization techniques

- Investigate application-level overlay optimization techniques such as predictive overlay topology modeling and self-organization.

Deliverables:

Month 18: Report on applicable application-level overlay optimization techniques.

Month 36: Final report on performance of application-level overlay optimization techniques.

WP4 Integration Issues Between Network-level and Application-level Optimizations

- Investigate and design concise topology representation for both physical network and overlay network topologies, to be used as input to predictive algorithms.
- Investigate integration issues in a two-layer programmable architecture from the ALAN perspective.

Deliverables:

Month 36: Report on issues and mechanisms for two-level optimization of networks.

WP5 P2P Architecture for Network Monitoring System

- Design and implement a peer-to-peer version of the network monitoring system (PhD student thesis topic).

Deliverables:

Month 18: Monitoring system design based upon p2p architectural concepts.

Month 24: Initial implementation and deployment of monitoring system design.

Month 36: Final evaluation report.

6 Related Work

The nature of the research proposed in this case for support is such that the results could be deployed on any high-performance programmable platform, especially at the network programmability level. Consequently, our research and results would provide an ideal scenario to test the use, in real networking conditions, of the results of the research proposed in "High-level Languages for Programmable Networks", a proposal submitted to the Programmable Networks Call 3, on languages for network processor-based (NP-based) systems. Likewise, the technology developed in the NETKIT project (Programmable Network Call 2), proposing a component based programmable network platform for NP-based systems, could be exploited advantageously. It should also be noted that our proposed research will start by being carried out on the LARA++ programmable network platform which is the result of direct continuation of the LARA (Lancaster Active Router Architecture) ROPA project.

The proposed research in this case for support is also complementary to the GRIDProbe project, jointly funded under the e-science and Programmable Networks initiatives. This is because GRIDProbe is focused on technologies for measurements in very high-speed networks, while our research is aimed at making use of measurement output to detect and correlate traffic patterns and optimize network utilisation based on such patterns.

In this project, it will be crucial to address the different types of p2p technology that exist. These can be categorized in two main groups: unstructured (based on a logical mesh built amongst the participants) [10,11], and content-addressable (based on the control of content/content reference placement within the p2p network, driven by techniques such as distributed hash tables) [12,13]. Techniques for proxy "insertion" and overlay traffic control will have to be devised for each type of p2p technology.

Overlay networks is a general term used to represent the notion of logical interconnection of entities on a network. Of course, p2p networks fall under such a definition, but there are other types of overlay networks. Of particular importance to our work is the Resilient Overlay Networks effort (RON) [14], GROUPNET [15], and to a lesser extent the work on application-level multicasting [16].

Traffic measurements are currently receiving significant research attention [17]. Passive measurement techniques rely solely on network traffic sampling to deduce the characteristics of traffic patterns in a network; on the other hand, active measurement techniques use explicit probing to infer these patterns [18]. Recently, a hybrid measurement technique has been proposed for measuring QoS of IP network flows [9]. It should be noted that most of the research in this area has focused on network traffic patterns, with little effort, if any, devoted to application traffic measurements and its correlation with network traffic.

Finally, traffic engineering techniques are also relevant to this work. Traffic engineering techniques have been proposed to control routing within a domain [19], as well as between domains [20]. Other techniques rely on the use of explicit routing in an MPLS mesh [21]. It should be noted that all such techniques operate at very long timescales, so that new techniques are required for a more dynamic, responsive environment.

The near-term research projects on PlanetLab (http://www.planet-lab.org) are distinct from, yet complementary to, the research proposed herein. Each of the partners is a current member of PlanetLab, or will soon join; we will look to other PlanetLab members for applications that we can use to drive the work here. One particular project, PlanetProbe, is using the GRIDProbe hardware to passively monitor the PlanetLab network; comparison of our network measurement results with those obtained by PlanetProbe will be informative, and may assist us in establishing more complete coverage over equivalence classes of p2p applications than we can do on our own.

7 Resources

The project duration is three years. Three post doctoral research assistants are required, one for each university partner. Additionally, to re-engineer the monitoring system using peer-to-peer techniques, Glasgow requires a 3-year studentship.

All three RAs will be responsible for work package 1, with the Glasgow RA leading the effort; the PhD student will be active in this work package, as well. The Lancaster RA will lead the effort in WP2, with assistance from the Glasgow RA. The UCL RA will lead the effort in WP3, with assistance from the Glasgow RA. The Lancaster and UCL RAs will jointly lead the effort in WP4. The PhD student is responsible for WP5.

Glasgow is requesting 10% each of a system programmer, a technician, and a secretary; Lancaster is requesting 10% of an administrator and 5% of a technician; and UCL is requesting 10% of a programmer. The programmers and technicians will help with implementation aspects, as there will need to be considerable integration of software from other sources, as well as amongst the project partners. The administrative staff are needed to provide clerical support and coordinate our interactions as well as coordinating activities with our industrial collaborator, UKERNA. As we intend to make our software available in the public domain, there will be some documentation requiring clerical support.

The network-level measurement activities require a full-packet-capture system, with GPS time synchronization hardware, at each site; we are requesting the funds necessary to purchase the GigE monitoring products supplied commercially by Endace Systems (http://www.endace.com/gigemon.htm). The active measurement techniques require programmable network elements at each site; we are requesting the funds necessary to purchase high-end PC systems upon which we will use the LARA++ programmable network element system developed at Lancaster (http://www.landmarc.net/people/stefan/LARA++.ps). An additional set of commodity systems are requested to support intense experimentation with various overlay network structures.

In addition, funds are requested to support machines that will be used for development. Travel funds are requested to present papers at conferences and workshops on Network Measurement, Peer-to-peer Computing, and Network Management. There will be regular meetings at Glasgow, Lancaster, UCL and UKERNA. Funds are also requested to provide computer consumables, toner, cartridges, laser printing, backup and archive services, and paper at the three sites in support of the research programme.

8 Relevance to Beneficiaries

Service Providers: As the number of production applications based upon p2p technologies and architectures grows, the traffic patterns on service providers' networks will change accordingly. The results of this work will enable the service providers to better resource and manage their networks for all their customers (p2p users or not). This, in turn, will allow network operators to enforce service differentiation in real-time, which will be a vital tool in the very competitive networking business of the future.

Operational Support System Vendors: It is expected that the network-level measurement and control techniques explored by the project will have direct application to OSS vendors' need to provide more automation in their products.

P2P Application Developers and Users: The results in overlay network structure optimization will benefit the developers and the users of p2p applications, even if the underlying networks are not programmable.

9 Exploitation and Dissemination

The investigators have a good track record of publishing papers at major conferences and in journals and will continue to do so. They also have strong links with various *de facto* standards groups, and will continue to influence those bodies.

Dissemination will be done at UK, European and wider international levels. In the UK we will contribute (as in the past) to the annual Cosener's House networking event, or its successor. European activities include the annual Eurescom workshop where many of the telcos participate, and we would want to expose the results of our work to these organisations. International events include Opensig (which we have participated in since its inception -- and organised the 2001 event in London), the IEEE Openarch conference which runs alongside IEEE Infocom, and also IWAN.

Exploitation will also be the subject of considerable activity by the partners. UKERNA is an active collaborator in this project, and will intercept the output of the project as it matures and provides needed functionality for appropriate portions of the JANET network. In addition to targeting the telcos through the Eurescom events, we will approach specific companies including BT and Telekom Austria with whom we already have strong links. We realise that this involves 'selling' the ideas and the outcomes to organisations that may tend to be sceptical about the likely benefits. Other organisations that we will attempt to influence include ISPs, and equipment providers such as Intel (through Intel Research in the UK and in the US).

Prof. Sventek, as an advisor to the TeleManagement Forum board, will proactively push the results of this work into the appropriate working groups within the TMF. Through his strong commercial links with Agilent and Hewlett-Packard, he will influence their product directions, where appropriate. As a member of the Strategic Advisory Board to the National eScience Centre, he will also proactively push the results of this work into the eScience community, in general, and the eDIKT activities, in particular.

Another body to which we will expose our work is the IETF FORCES group: this will probably be done through strong contacts that we have at IBM Research in Zurich, and with Intel Research.

The software developed will be made available in the public domain, as will anonymized versions of the network traces obtained.

10 Project Management

The partners will hold quarterly, face-to-face meetings to review progress on the project. The partners, RAs, and PhD student will meet monthly by phone conference to discuss more short-term management issues. All will also attend semi-annual, face-to-face meetings to summarize results to date and to review the next sets of milestones.

The workplan for the project is shown below. The work allocation specifies where the main focus of the respective RAs and the PhD student will be.

References

[1] R. Cardoe, J. Finney, A. C. Scott, W. D. Shepherd, "LARA: a prototype system for supporting high performance active networking", Proceedings of the First International Working Conference on Active Networks -- IWAN '99 (June 1999).

[2] A. Oram (Ed.): "Peer-to-Peer: Harnessing the Benefits of a Disruptive Technology", O'Reilly, 2001.

[3] I. Stoica, D. Adkins, S. Zhang, S. Shenker, S. Surana: "Internet Indirection Infrastructure", ACM SIGCOMM`02, August 2002.

[4] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, H. Balakrishnan: "Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications", ACM SIGCOMM`01, San Diego, September 2001.

[5] H. De Meer, K. Tutschku, P. TranGia: ``Dynamic Operation in Peer-to-Peer Overlay Networks'', Praxis der Informationsverarbeitung und Kommunikation, (PIK Journal), Special Issue on Peer-to-Peer Systems, June 2003.

[6] M. Fry, A. Ghosh, "Application level active networking", Computer Networks, Vol. 31, No. 7, 1999.

[7] A. Campbell, H. De Meer, M. Kounavis, K. Miki, J. Vicente, and D. Villela. "A Survey of Programmable Networks". ACM Computer Communications Review, pp. 7-23, (April 1999).

[8] A. Campbell, M. Kounavis, D. Villela, J. Vicente, H. De Meer, K. Miki, and K. Kalaichelvan. Spawning Networks. IEEE Network, (July/August 1999).

[9] D. Pezaros, D. Hutchison, F. Garcia, R. Gardner, J. Sventek, "An IPv6-based Measurement Infrastructure for Assessing Internet Traffic Behaviour", submitted to the Special Issue of the IEE Proceedings-Communications on Internet Protocols, Technology and Applications (VoIP).

[10] M. Ripeanu, A. Iamnitchi, I. Foster: "Mapping the Gnutella Network", IEEE Internet Computing, January/February 2002.

[11] E. Cohen, S. Shenker: "Replication Structures in Unstructured Peer-to-Peer Networks", ACM SIGCOMM, Pittsburgh 2002.

[12] H. Balakrishnan, M.F. Kaashoek, D. Karger, R. Morris, I. Stoica: "Looking up Data in P2P Systems", Communications of the ACM, Vol 43, No. 2, February 2003.

[13] S. Ratnasami: "A Scalable Content-Addressable Network", PhD Thesis, UC Berkeley, October 2002.

[14] D. Andersen, H. Balakrishnan, F. Kaashoek, R. Morris, "Resilient Overlay Networks", ACM SOSP, Banff, Canada, October 2001.

[15] P. Smith, S. Simpson, D. Hutchison, "Peer-to-Peer Networking for Discovering Programmable Resources", NGC 2002, Boston, October 2002.

[16] L. Mathy, R. Canonico, S, Simpson, D. Hutchison, "Scalable Adaptive Hierarchical Clustering", IEEE Communications Letters 6(3), March 2002.

[17] Proceedings of the ACM SIGCOMM Internet Measurement Workshop, Marseille, France, November 2002.

[18] V. Paxson, "Towards a Framework for Defining Internet Performance Metrics", Proceedings of INET '96, Montreal, June 1996.

[19] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, "NetScope: Traffic Engineering for IP Networks", IEEE Network 14(2), March 2000.

[20] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen, O. Bonaventure, "Traffic Engineering with BGP", IEEE Communications, May 2003.

[21] X. Xiao, A. Hannan, B. Bailey, L. Ni, "Traffic Engineering with MPLS in the Internet", IEEE Network 14(2), March 2000.

Project Workplan

	WP	LU RA	UCL RA	UG RA	UG PhD	Months					
						6	12	18	24	30	36
1	Measurements & Correlation	×	×	×	×	■	■			■	
2	Real Network Optimization	×		×				■		■	■
3	Overlay Network Optimization		×	×				■			■
4	Simultaneous Optimization Issues	×	×								■
5	P2P-based Monitoring System				×			■	■		■

■ **Deliverable**