

Report on

The COST (EU)-NSF (USA)

Workshop on

EXCHANGES & TRENDS IN NETWORKING

NeXtworking'03

Edited by

Ioannis Stavrakakis, University of Athens

Ibrahim Matta, Boston University

Michael Smirnov, Fraunhofer FOKUS

--- VERSION 1---

(Later versions will appear at the workshop site)

www.di.uoa.gr/~NeXtworking

Sponsored by:

The COST-IST program of the European Union

The National Science Foundation (NSF) of the United States of America

July 2004

Introductory Comments by the Organizers

This workshop was the first attempt to bring together the Networking research communities from the USA and Europe in a common forum, to exchange research ideas and trends in Networking. Unique features of this meeting were the joint and balanced participation of researchers from the European and USA research communities and the high concentration of top, mature and highly respected researchers from both communities.

Specific workshop objectives were to:

- Discuss key, new and upcoming networking technologies and open issues by focusing on the driving fundamental aspects and principles.
- Provide the opportunity to the networking communities in Europe and USA to get together and have a first-hand exposure to the advancements and strengths in the other side of the Atlantic.
- Provide a forum that could help the workshop-sponsoring funding agencies in Europe and USA identify the key and necessary future networking technologies and direct research funding accordingly.

The workshop was made possible with the generous financial support from the **COST-IST program of the European Union** and the **National Science Foundation (NSF) of the United States of America**. **Boston University (USA)** and **Fraunhofer Fokus (Germany)** hosted the NSF and COST-IST funding, respectively, while the **National and Kapodistrian University of Athens (Greece)** undertook the overall organization of the workshop. The participation of all 48 speakers was generously supported, while registration and some additional coverage were provided to almost every other attendee as well.

Ioannis Stavrakakis, *General Workshop Chair and TPC co-Chair*

Ibrahim Matta, *TPC co-Chair*

Michael Smirnov, *TPC co-chair*

1	Description of the NeXtworking'03 Report.....	4
1.1	Who could use this report.....	4
2	Description of the Workshop	5
2.1	Workshop Objectives	5
2.2	Workshop venue and support.....	5
2.3	Workshop Organizing and Technical Program Committee	6
2.4	List of Participants	7
2.5	Technical Program Focus and Organization	9
2.6	Overall Workshop Feedback and Conclusions	11
3	Summary of Workshop Findings and Recommendations.....	13
4	Summary of Technical Challenges	17
5	Research Challenges	20
5.1	Research Challenges in WIRELESS NETWORKING: General Aspects, Ad-Hoc and Sensor Networks.....	20
5.2	Research Challenges in TOMOGRAPHY, MEASUREMENTS, and MODELING 23	
5.3	Research Challenges in PEER-TO-PEER AND OVERLAY NETWORKS: Overlay networking and Content Distribution.....	26
5.4	Research Challenges in LARGE SCALE NETWORKS AND FUTURE NETWORK ARCHITECTURES	27
5.5	Research Challenges in PROGRAMMING THE INTERNET.....	28
6	Workshop Session Summaries	30
6.1	Session title: WIRELESS NETWORKING: General Aspects	30
6.2	Session title: WIRELESS NETWORKING: Mobile Ad-Hoc Networks.....	32
6.3	Session title: WIRELESS NETWORKING: Sensor Networks	32
6.4	Session title: NETWORK TOMOGRAPHY AND TRAFFIC MODELING, MEASUREMENTS AND CONTROL: Network Traffic Modeling and Control.....	38
6.5	Session title: PEER-TO-PEER AND OVERLAY NETWORKS: Peer-to-Peer Networking.....	40
6.6	Session title: PEER-TO-PEER AND OVERLAY NETWORKS: Overlay networking and Content Distribution.....	45
6.7	Session title: NETWORK TOMOGRAPHY AND TRAFFIC MODELING MEASUREMENTS AND CONTROL: Network Tomography	49
6.8	Session title: NETWORK TOMOGRAPHY AND TRAFFIC MEASUREMENTS AND CONTROL: Network Measurements	51
6.9	Session title: LARGE SCALE NETWORKS AND FUTURE NETWORK ARCHITECTURES	52
6.10	Session title: PROGRAMMING THE INTERNET.....	53
7	Post-Workshop Questionnaire Results.....	56

1 Description of the NeXtworking'03 Report

The *NeXtworking'03* report consists of the following two parts:

- A. **The workshop description and findings** containing:
 - a. The description of the workshop's scope, format and participation.
 - b. The summary of the findings and recommendations that take into consideration the workshop presentations and discussions.
 - c. The detailed description of the research challenges.
 - d. Technical session summaries prepared by the session chairs after the workshop.

- B. **The workshop supporting electronic material** consists of:
 - a. The workshop proceedings (abstracts and slides of the workshop presentations).

This document contains Part A of the report. Part B is available at the *NeXtworking* workshop site at www.di.uoa.gr/~NeXtworking.

1.1 Who could use this report

This report is addressed to anyone interested in networking research and particularly to the following constituencies:

- A. Networking research-funding agencies may find the information and recommendations contained in this report useful in setting up their funding priorities. In addition, this report may trigger more serious efforts to fund joint EU-USA research collaborations.
- B. Established Networking researchers may benefit from this report in setting up their research agenda.
- C. Junior Networking researchers and graduate students may be assisted in focusing their research efforts along the research directions identified in this report, as well as be educated in the research and presentation process by studying the presentation material and talks given by mature researchers in the field.

2 Description of the Workshop

2.1 Workshop Objectives

The purpose of this workshop was to provide a forum for exchanges of research ideas and trends in Networking.

Unique features of this meeting were:

- The joint and balanced participation of researchers from the European and USA research communities.
- The high concentration of top, mature and highly respected researchers from both communities.

Specific workshop objectives were to:

- Discuss key, new and upcoming networking technologies and open issues by focusing on the driving fundamental aspects and principles.
- Provide the opportunity to the networking communities in Europe and USA to get together and have a first-hand exposure to the advancements and strengths in the other side of the Atlantic.
- Provide a forum that could help the workshop-sponsoring funding agencies in Europe and USA identify the key and necessary future networking technologies and direct research funding accordingly.

Post-workshop objectives were to:

- Prepare a final electronic version of the workshop proceedings (abstracts and presentations).
- Prepare a comprehensive report on the workshop findings to be released to the community.
- Establish *NeXtworking* as the premium forum for joint European-USA exchanges and trends in Networking on a bi-annual basis. The X has been added to the workshop name to stress the fact that this Networking research forum focuses strongly on the future (NeXt) and the trends in the field; the (larger) X seen alone captures the other focus of this workshop on exchanges of research ideas between (geographically) distinct communities (Europe and USA).

2.2 Workshop venue and support

The workshop took place on June 24, 25 and 26 of 2003 at the Venetian “Great Arsenal”, situated in the picturesque Venetian harbor of Chania, Greece. All lunches and dinners were organized by the workshop and offered additional opportunities for research and cultural exchanges. Despite the escape challenges of the fascinating venue, the workshop’s participation level was close to 100%.

The workshop was made possible with the generous financial support from the **COST-IST program of the European Union** and the **National Science Foundation (NSF) of the United States of America**. **Boston University (USA)** and **Fraunhofer Fokus (Germany)** hosted the NSF and COST-IST funding, respectively, while the **National and Kapodistrian University of Athens (Greece)** undertook the overall organization of the workshop. The participation of all 48 speakers was generously supported, while registration and some additional coverage were provided to almost every other attendee as well.

2.3 Workshop Organizing and Technical Program Committee

General Workshop Chair:

Ioannis Stavrakakis, University of Athens

TPC Co-Chairs:

- *Ibrahim Matta, Boston University,*
- *Michael Smirnov, Fraunhofer FOKUS,*
- *Ioannis Stavrakakis, University of Athens,*

TPC members (EU):

- *Remy Bayou, European Commission,*
- *Chris Blondia, Antwerp University*
- *Ernst W. Biersack, Institut EURECOM*
- *Olga Casals, Polytechnic University of Barcelona*
- *Marco Conti, CNR-ITT*
- *Jon Crowcroft, University of Cambridge*
- *Christophe Diot, INTEL, Sprint Labs*
- *Serge Fdida, University Pierre et Marie Curie*
- *Luigi Fratta, Politecnico di Milano*
- *David Hutchison, Lancaster University*
- *Jean-Pierre Hubaux, EPFL*
- *Gunnar Karlsson, Royal Institute of Technology*
- *Jim Roberts, France Telecom R&D*
- *Ian Wakeman, University of Sussex*
- *Peter Wintlev-Jensen, European Commission*

TPC members (USA):

- *Ian F. Akyildiz, Georgia Institute of Technology*
- *Ken Calvert, University of Kentucky*
- *Constantinos Dovrolis, Georgia Institute of Technology*
- *Anthony Ephremides, University of Maryland*
- *Mario Gerla, UCLA*
- *Sugih Jamin, Univ. of Michigan/Univ. of Tokyo*
- *Edward Knightly, Rice University*
- *Jim Kurose, University of Massachusetts*
- *Jorg Liebeherr, University of Virginia*
- *Ness Shroff, Purdue University*
- *Don Towsley, University of Massachusetts*
- *Taieb Znati, University of Pittsburgh / NSF*

2.4 List of Participants

The total number of participants was 89 out of which 49 gave a presentation (28 from EU and 21 from USA).

EU: (speakers)

1. *Remy Bayou, European Commission,*
2. *Chris Blondia, Antwerp University*
3. *Ernst W. Biersack, Institut EURECOM*
4. *Erdal Cayirci, Istanbul Technical University*
5. *Israel Cidon, Technion*
6. *Marco Conti, CNR-ITT*
7. *Jon Crowcroft, University of Cambridge*
8. *Michel Diaz, LAAS-CNRS*
9. *Serge Fdida, University Pierre et Marie Curie*
10. *Anja Feldmann, TU-Muenchen*
11. *Luigi Fratta, Politecnico di Milano*
12. *Per Gunningberg, Uppsala University*
13. *Jean-Pierre Hubaux, EPFL*
14. *Perti Jauhiainen, European Commission*
15. *Gunnar Karlsson, Royal Institute of Technology*
16. *Peter Key, Microsoft Cambridge*
17. *Francesco Lo Presti, Universita' dell'Aquila*
18. *Jim Roberts, France Telecom R&D*
19. *Pablo Rodriguez, Microsoft Cambridge*
20. *Anthony Rowstron, Microsoft Cambridge*
21. *Kave Salamatian, LIP6/UPMC*
22. *Michael Smirnov, Fraunhofer FOKUS,*
23. *Ioannis Stavrakakis, University of Athens,*
24. *Patrick Thiran, EPFL*
25. *Christian Tschudin, Uppsala University*
26. *Gabor Vattay,*
27. *Ian Wakeman, University of Sussex*
28. *Peter Wintlev-Jensen, European Commission*

USA: (speakers)

29. *Ian F. Akyildiz, Georgia Institute of Technology*
30. *Victor Bahl, Microsoft Research*
31. *Ken Calvert, University of Kentucky*
32. *Hao Che, University of Texas at Arlington*
33. *Christophe Diot, INTEL, Sprint Labs*
34. *Constantinos Dovrolis, Georgia Institute of Technology*
35. *Anthony Ephremides, University of Maryland (6/25)*
36. *Lixin Gao , Umass*
37. *Mario Gerla, UCLA*
38. *Ramesh Govindan, USC*
39. *Edward Knightly, Rice University*
40. *Balachander Krishnamurthy, ATT Research*
41. *Jim Kurose, University of Massachusetts*

42. *Brian Levine, University of Massachusetts – (Umass)*
43. *Jorg Liebeherr, University of Virginia*
44. *Ibrahim Matta, Boston University,*
45. *Ravi Mazumdar, University of Purdue*
46. *Vishal Mishra, Columbia University*
47. *Don Towsley, University of Massachusetts*
48. *Magda Zarki, UC Irvine*
49. *Taieb Znati, University of Pittsburgh / NSF*

Additional attendees

50. *Kihong Park, Purdue*
51. *Vassilis Tsaoussidis, Northeastern University*
52. *David Yau , Purdue*
53. *Yevgeni Koucheryavy, Tampere Univ.*
54. *Guy Leduc, Universite de Liege*
55. *Lars Wolf, Technical University of Braunschweig*
56. *Moshe Sidi, Technion*
57. *George C. Polyzos, Athens University of Business & Ec*
58. *Christos Douligeris, University of Piraeus*
59. *Nikos Pronios, Intracom*
60. *Vasilis Siris, FORTH*
61. *George Stefanou, University of Athens*
62. *Michael Paterakis, Technical University of Crete*
63. *Sophia Tsakiridou, Technical University of Crete*
64. *Leonidas Georgiadis, University of Thessaloniki*
65. *Lazaros Polymenakos, AIT (Athens Information Technology)*
66. *Athanasios Tyropanis, AIT (Athens Information Technology)*
67. *Georgios Smaragdakis, Boston University*
68. *Leonidas Tzevelekas, University of Athens*
69. *Kostantinos Oikonomou, Intracom S.A.*
70. *Nikos Laoutaris, University of Athens*
71. *Athanasios Vaios, University of Athens*
72. *Antonios Panagakis, University of Athens*
73. *Pandelis Balaouras, University of Athens*
74. *Vassilios Stoidis, Technical University of Crete*
75. *Anna Satsiou, Technical University of Crete*
76. *Spiros Psyxis, Technical University of Crete*
77. *Petros Kaklamanis, Technical University of Crete*
78. *Panagiotis Mouziouras, Technical University of Crete*
79. *Ioannis Mproustis, Technical University of Crete*
80. *Stratis Idraios, Technical University of Crete*
81. *Pantelis Koukousoulas, Technical University of Crete*
82. *Matoula Magiridou, Technical University of Crete*
83. *Dhmosthenis Anthonelidis Technical University of Crete*
84. *Aggelos Vlavianos, Technical University of Crete*
85. *Panagiwiths Afraths, Technical University of Crete*
86. *Giannis Atzarakh, Technical University of Crete*
87. *Ioannis Mparmpounakis TSI - Technical University of Crete*
88. *Vaggelis Kokkinos, TEI of Crete*
89. *Georgios Liodakis, TEI of Crete*

2.5 Technical Program Focus and Organization

The workshop technical program was focused on **three major themes**:

- Wireless
- Architectures
- Modeling.

Ten technical sessions were formed in the following areas:

- WIRELESS NETWORKING: General Aspects
- WIRELESS NETWORKING: Mobile Ad Hoc Networks
- WIRELESS NETWORKING: Sensor Networks
- NETWORK TOMOGRAPHY AND TRAFFIC MODELING, MEASUREMENTS AND CONTROL: Network Traffic Modeling and Control
- NETWORK TOMOGRAPHY AND TRAFFIC MODELING, MEASUREMENTS AND CONTROL: Network Tomography
- NETWORK TOMOGRAPHY AND TRAFFIC MODELING, MEASUREMENTS AND CONTROL: Network Measurements
- PEER-TO-PEER AND OVERLAY NETWORKS: Peer-to-Peer Networking
- PEER-TO-PEER AND OVERLAY NETWORKS: Overlay networking and Content Distribution
- LARGE SCALE NETWORKS AND FUTURE NETWORK ARCHITECTURES
- PROGRAMMING THE INTERNET

In addition, there was **one non-technical panel** on:

- EU (COST-IST) - USA (NSF) FUNDING PERSPECTIVES

And, **two summarizing panels**:

- EXECUTIVE SUMMARY ON NETWORKING CHALLENGES AND PRIORITIES – part I
- EXECUTIVE SUMMARY ON NETWORKING CHALLENGES AND PRIORITIES – part II

Participation in the workshop was by invitation only. Thus, it was possible to select the session chairs and speakers carefully to guarantee the coverage of the targeted area with the highest possible quality. Speakers were instructed to focus their presentations on:

- Major research challenges
- Innovative approaches addressing key problems
- Fundamental work addressing fundamental (class of) problems
- Participants' own recent work as it relates to important largely open problems and thoughts for follow up and applications in other domains.

Table I shows the workshop organization content-wise. A multi-step and multi-level of abstraction approach was followed. In the first step, the three major themes were identified as the workshop focus. These themes were then detailed and organized in ten technical sessions.

The presentations and discussions during these sessions were later summarized by the session chairs in two summarizing panels; a third panel summarized and discussed funding policies and priorities of the workshop supporting funding agencies from the EU and USA. Finally, the present report summarizes the entire event, focusing on the workshop findings and recommendations, also including summaries of the sessions and results from a post workshop feedback by the speakers regarding the identified research challenges.

Workshop Organization (content-wise)			
Wireless	General (Ephremides)	Panel on Wireless/Sensor Networking	Workshop Report (Findings & Recommendations)
	Ad-hoc (Gerla)		
	Sensor (Akyildiz)		
Architectures	P2P (Crowcroft)	Panel on Modeling & Architectures	
	Overlays (Biersack)		
	Large-scale (Kurose)		
Modeling	Programming (Calvert)	Panel on EU (COST-IST) & US (NSF) Funding perspectives	
	Measurements (Diot)		
	Tomography (Towsley)		
	Traffic (Roberts)		

Table I

The workshop organization work-wise is shown in Table II. There were 4 types of contributors to the workshop (TPC co-chairs, session chairs, speakers and other attendees) with contributions to the workshop before, during and after the event.

Before the workshop: The TPC co-chairs secured funding for the workshop, defined the format and topics and invited the TPC members, session chairs, speakers and other attendees; they also collected and published the presentation abstracts and slides. The session chairs invited some of the speakers, coordinated the presentations within their respective sessions, and submitted the abstract and slides of their talks and their session's summary. The speakers submitted the abstracts and slides of their talks and coordinated their talks with the session chair.

During the workshop: The TPC co-chairs ran the workshop, distributed some introductory questions to the participants and moderated the summarizing sessions. The session chairs gave their presentations, ran their sessions and prepared (during the breaks and the meals) and

presented a summary of their sessions in one of the summarizing panels. The speakers gave their presentations. Finally, all participants were involved in the lively discussions during the entire workshop.

After the workshop: The TPC co-chairs reimbursed the participants, collected feedback from the participants, prepared the present report and distributed and processed a questionnaire about the technical challenges identified in the workshop (results of this questionnaire will be available in the final version of this report). The session chairs submitted a summary with the highlights of their sessions, revised abstracts/slides of their presentations, provided feedback on the workshop and responded to the post-workshop questionnaire. Finally, the speakers submitted revised abstracts/slides of their presentations, provided feedback on the workshop and responded to the post-workshop questionnaire.

Workshop Organizatin (work-wise)			
	BEFORE	DURING (all participate)	AFTER
TPC co-chairs	Secure funding Define format, topics Invite session chairs and speakers Collect & publish abstracts / talks	Run the workshop Distribute questions Chair summary sessions	Reimburse participants Collect feedback Process Questionnaire Prepare the workshop report Disseminate
Session chairs	Invite some speakers Coordinate talks Submit abstract / talk Submit session summaries	Run their session Present their views Prepare and present summary of their session	Submit summary of session highlights Revise slides / abstract Provide feedback Respond to Questionnaire
Speakers	Submit abstract / talk Coordinate with session chairs	Present their views	Revise slides / abstract Provide feedback Respond to Questionnaire
Other Attendees	Be invited		Provide feedback

Table II

2.6 Overall Workshop Feedback and Conclusions

The participants felt that this workshop was a great event in all aspects. This sense was echoed both during and after the event. The quality of the speakers and their presentations, the stimulating discussions and the general organizational details were highly commented on. The revisiting of the key aspects in the summarizing sessions was commented as a good idea that helped further discuss the larger and more controversial issues that emerged during the

individual presentations. The only recommendation for improvement was the tighter control of the speakers' time and the further reduction of their allotted time in favor of more time for discussions. The participants felt that this event was of very good value and were favorable to the continuation of such events in the future. The unique feature of this event regarding the balanced mix of quality views from both sides of the Atlantic was appreciated, as it was the possibility that cross-Atlantic research funding opportunities would result from such efforts.

3 Summary of Workshop Findings and Recommendations

This section contains higher-level non-technical and technical findings and recommendations. A detailed description of the specific research challenges identified in the workshop is presented in Section 5, with a summarizing overview of the technical challenges in Section 4.

Over the past two decades, the Internet and wireless cellular networking technologies experienced an enormous growth and were widely embraced by the society; as a result, a good deal of our everyday functions at work, home and anywhere have been substantially enhanced. These technological successes are based on the introduction of a few new ideas of transporting information (e.g., packet switching techniques) and accessing common channels (random-access, multi-user schemes), supported by advances in enabling VLSI and fiber technologies and some killer applications (Web access for the internet and personal communications anytime and anywhere). The networking community is largely responsible for introducing the new networking paradigms, establishing their potential, as well as introducing the basic techniques and supporting protocols to turn them into functional and effective technologies.

Finding: The workshop presentations and discussions revealed a relative unawareness of the USA community regarding the advances of the EU community (and vice versa) in related and supplementary research fields. It was felt that knowledge of these advances and strengths could enhance their research.

Recommendation: There is a need for better communication between the two communities; research forums seeking to present a balanced view from both sides of the Atlantic, such as *NeXtworking*, should be encouraged.

Finding: It was strongly felt that there is expertise in the other side of the Atlantic that could greatly enhance a group's effort in a certain area. Very frequently the desire was expressed that groups pursue jointly certain problems to capitalize on the diverse experiences and strengths. Based on the information provided by the representatives from the funding agencies, as well as the experience of some researchers, there is presently no program to support a joint research effort of teams from both sides of the Atlantic. Existing programs support only travel and exchanges and not the research itself, while at the same time the proposal submission and evaluation process appears to be prohibitively complex for the level of support it provides, requiring an independent approval from the respective agencies that operate on different time scales and have different priorities and evaluation criteria.

Recommendation: Joint EU-NSF funding mechanisms that go beyond travel support should be established to foster serious international collaborations.

Finding: The workshop presentations and discussions revealed that important open problems would require multi-disciplinary research effort to be comprehensively addressed, as well as expertise obtained by studying similar problems in the past.

Recommendation: Researchers should create communities that actively communicate, formulate and cooperate to solve specific grand challenge problems.

EU vs. USA Finding & Recommendation: The Network of Excellence instrument in EU appears to be a good step in this direction. No similar instruments are currently available in USA.

Finding: Internet's success has led to the treatment of networking as a commodity and the perception that Networking research is done. While the networking community should also consider focusing on services and applications - especially in order to motivate and develop new networking paradigms – the effort to address the fundamental (new) networking problems cannot be diminished or abandoned. Governmental funding agencies, increasingly under pressure to demonstrate the social benefits of networking, appear to assume the position that the basic networking research is done and that further work would only be incremental.

Recommendation: Basic research should continue to receive the attention of the research communities, to address the great challenges appearing in an ever-changing networking environment. Funding agencies should continue to fund fundamental research.

EU vs. USA Finding & Recommendation: In the USA, basic research has been well served by NSF; to some extent, DoD funding agencies and the Industry (more so in the past) have also funded basic research. Recently, there is some trend by NSF to reduce the level of funding for basic research in networking; the USA Industry has reduced its effort substantially, while DoD's focus is heavily biased towards specific, narrowly focused aspects only. It was felt in the workshop that the USA funding agencies – and especially NSF – should continue to support basic research, as such funding has contributed to R&D efforts that gave in the past two decades a technological advantage to the USA in the networking field. In EU, funding for networking is primarily channeled through the industry-controlled R&D efforts funded by the IST program. Basic research is not central to this program and when present at all it constitutes a rather insignificant portion of the effort and the expected deliverable. Although the Future Emerging Technologies (FET) program addresses better basic research than the main IST program, the available funding for networking research is extremely limited. Mostly, this program seeks to fund innovative vision or core enabling technologies for the future, rather than basic research that could move forward emerging technology that could mature in a 5-year horizon. Thus, it appears that there is a critical space between the industry-led, product-oriented R&D effort (funded by IST) and the long-term, visionary and very high-risk research (funded by FET), which is not funded in EU. This precisely seems to be the basic research space traditionally funded by NSF that has supported the development of emerging technologies that mature in a 5-year horizon and has led to a competitive advantage.

Finding: Networking is viewed as one of the most important sectors of technology and economy today, offering a great deal of job opportunities. Highly knowledgeable and well-trained personnel are needed in order for the industry to be successful and competitive. Offering good quality courses in networking and engaging students in networking research can generate the best qualified personnel for this industry, capable of making good contributions to the advancement of the technology.

EU vs. USA Finding & Recommendation: In the USA, about 90% of the NSF research project funding is used to support and train graduate students. This is viewed as a tremendous opportunity for the best students to focus on basic research and the process of thinking and researching. Not only good students are given a great opportunity to make good contributions to basic research, but also they are superbly trained and position themselves to making significant R&D contributions in the future. NSF's commitment to education is highly

appreciated by the community. In EU, the typical funded project is not tailored to supporting graduate education, but rather to promote the competitiveness of the industry. In most cases, graduate students cannot handle these projects. Whenever graduate students are involved, the cost in time and effort paid is usually not compensated for by the experience with technology they gain, as there is little room for pursuing basic research and the process of researching. Education is supposed to be a concern of the national governments in EU, but funding for basic research training is rather limited and fairly diverse across EU. It is believed that research training and education should be addressed at the European level and be substantially increased.

Finding: Applications have been the driver of networking research and development in the last 30 years; for instance, Internet's boom is largely due to the web application. Applications are expected to continue to drive networking research.

Recommendation: Basic research should be linked to problems emerging when trying to effectively accommodate important applications that are likely to be adopted by the users, so that its impact on technology adopted by society is maximized.

Finding: Internet is a very large-scale, complex system run by distributed, independent entities. It is built on principles not well understood with very little fundamental research to help us understand its behavior, as it grew large and complex. It seems that the driving principle has been mostly that of "rough consensus and running code". While this has kept its operation simple and has facilitated its proliferation and evolution, it has made it very difficult to deliver Quality-of-Service (QoS) and exercise tighter control, as some emerging applications would require.

Recommendation: New theories should be developed to help us understand Internet's behavior. Such theories may need to utilize knowledge and experience from non-traditional and multi-disciplinary scientific fields.

Finding: While applications are considered to be the driver of new networking paradigms and theory provides a sound basis for understanding behaviors, experimentation and measurements appear to be an increasingly useful approach to gaining understanding about complex, large scale systems, such as the Internet, and to helping in building good models. Due to the distributed ownership and administration of the Internet, it is hard for individual researchers to access data recording various aspects of the system's behavior.

Recommendation: Applications, theory and experimentations should all be promoted and synergies among them should be developed. Governmental agencies should encourage strong cooperation between the research community and network operators to promote sharing of internal network information. Governmental agencies should promote the sharing of prototypes, software tools and data traces.

Finding: Theory and modeling appear to be on the defense. Recent quote on the DiffServ mailing list "we don't have the math, so let's not bother [trying to model]"

Recommendation: Model and understand first before building.

Finding: New exciting wireless networking applications are emerging, calling for new mostly application-specific networks and requiring flexible standards.

Recommendation: The research community should be involved in Standards more heavily to influence the adoption of more flexible standards and reduce the degree of technology ossification.

Finding: There is the belief that we tend to create faceted research communities; for example, new communities created for ad-hoc networks, sensor networks, etc. We often reinvent the wheel and fail to define and address grand challenges.

Recommendation: We need to maintain communication and work on identifying common fundamental problems. More effort is needed to understand commonalities in the problems. Experience from addressing similar problems in the past, or from related disciplines, should be reused, to help move forward and address the problems more effectively.

Finding: Important technologies that could allow for the efficient support of a wide range of applications by the Internet, such as QoS and multicast, are largely not delivered. Reasons for this include ossification, complexity, lack of business models, etc.

Recommendation: These are lessons for the future. We should try to quantify complexity and the other so-called “ities,” take into account incentives, programmability, etc. Academic researchers should get involve in standards early on to ensure enough flexibility for innovative solutions.

4 Summary of Technical Challenges

This section contains a summarizing overview of the specific technical challenges identified in the workshop. A detailed description of the research challenges is presented in Section 5.

The “*Science of Network Service Composition*” has clearly emerged during the workshop as *the* grand theme driving many of our research questions today. This driving force stems from the rise of sophisticated applications and new networking paradigms. Our society is becoming increasingly dependent on applications with sophisticated spatial-temporal semantics running in highly dynamic network settings. Examples include complex aggregate queries of tiny sensors embedded in the environment to non-intrusively monitor sensitive wildlife and habitats; surveillance tasks carried out by cooperating video sensors to monitor the security of public places; diagnostic queries of network-load monitors to ensure the healthy operation of the global Internet on which our economies are increasingly dependent; and many other examples.

Throughout the workshop, the discussions have often led to consensus on the technical challenges ahead of us to develop such a *Science of Network Service Composition*, and the great social benefits exemplified by the aforementioned applications. By “network service” we mean the interface through which a user, application or protocol, requests certain spatial-temporal properties for its data or query sent over the network. By “composition” we mean that the spatial-temporal properties local to the various constituent components of the network can be readily composed into global (end-to-end) properties without re-analyzing any of the constituent component in isolation or as part of the whole composite network system. The set of laws that would govern such composition is what will constitute that new science of composition.

Although the general idea of composing systems is not new, the combined heterogeneity and large-scale dynamic nature of network systems makes composition quite challenging in many respects. First, research is needed to identify the boundaries of those “components” that make up the whole system and their interfaces so they could exchange appropriate information and accordingly adapt their individual local control rules. These local control rules, when composed, must finally lead to global properties that satisfy end-to-end services. These constituent feedback (closed-loop) controllers would operate at different levels of the architecture, from application level to physical level. Second, these component boundaries and interfaces should be flexible and extensible to allow for programmability to fit the needs of a wider class of applications. Third, given the open nature of large-scale network systems, issues of security and accounting must be integral to the design.

Lastly, for all these aspects to come together in building a trustworthy network system, research is needed to advance our current modeling and analysis techniques. Although our networking community has made significant contributions in mathematical modeling and analysis, the theory has either often come too late following the design or been limited in its level of detail and “composition” power necessary to make it effective in dealing with large-scale dynamics and in being accessible to the users, service providers or network managers.

Thus, toward a new science of “network service composition”, the workshop categorizes the technical challenges into the following four major themes:

- **Cross-layering**
 - The term “cross-layering” refers to *cooperating* network and traffic controllers located within the same node and across all nodes of the network system. It was felt that more research is needed to capture the strong interaction among these various controllers located across all levels of the architecture, ranging from the application layer to physical layer, which in turn calls for more interdisciplinary research.
- **Programmability**
 - The term “programmability” refers to the ability of users to flexibly shape and extend the network architecture using an easy-to-use and expressive interface. The greatest benefit of increased programmability is to reduce ossification and increase extensibility of network systems to accommodate changing requirements.
- **Enabling technologies**
 - More research was felt needed in often-ignored areas of security, trust, incentives, pricing, and billing as enablers of innovative network systems.
- **Theory to model & understand**
 - All the above themes give rise to large-scale interactions. Fundamental theories are needed to model and understand network systems and their composition in the presence of uncertainties, and to characterize their performance limits and their complexity, reliability, maintainability, etc. (aspects referred to during the workshop as the often-ignored “*ities*.”)

A specific set of research issues related to each one of these four themes are summarized next.

Cross-layer Research: The scope of this research spans a range of possible solutions, from considering just new additional triggers between traditional layers of the protocol architecture to fully interleaving the layers. Throughout the discussions and especially in the context of wireless communication, it was felt that the networking community needed to more adequately consider issues related to the allocation of the physical spectrum and the availability of multiple radios and their impact on higher layers. New theories need to be developed to design protocols for the effective management of multiple resources under both space and timing constraints. Wireless environments are often perceived as limited in their resources. However, many characteristics, such as mobility and broadcast communication, can be exploited in self-organizing designs to achieve higher efficiencies. New theories will enable us to obtain bounds on such performance gains subject to cost constraints such as energy consumption.

Although there has been research on some limited forms of cross-layer interactions, it is clear that more work needs to be done in developing systematic formal methods to study the problem of interacting traffic/resource controllers in the face of uncertain time-varying views of the state of the system. Measures of stability and transient performance need to be defined and evaluated. We need to better understand how multiple controllers, where the (changing) output of one is input to another, can be composed and reasoned about. The global properties of such composition need to be readily computed from the local properties of constituent controllers. Such properties should include “safety” measures including predictability of performance, trust and progress. Since traffic/resource controllers must adapt to time-varying conditions, the architecture should support monitoring and measurement as an integral piece of the network system, and not as an afterthought. Research is needed in developing active as well as passive and inference (non-intrusive) techniques.

Programmability: Using measured views of the state of the network system, the user or protocol should be able to re-organize and tune the parameters of the underlying substrate to control the quality of the delivered end-to-end service. This kind of dynamic configuration and self-organization requires a new theory of “network service composition” which provides the rules needed to infer the properties of the composite system. Examples of compositions range from composing controllers within the same network node to composing controllers across network nodes and even across different networks. Research is needed to define a “specification language” that is expressive enough to describe different components of the network system, and that will include “typing” hierarchies to enable composition similar to type systems in general programming languages.

Programmability enables the architecture to adapt to specific applications, such as military or environmental sensor systems and dynamic content or overlay systems. Imagine we could program such network systems as we program applications today using a set of standard libraries, our own and third-party routines, which are linked (composed) into a single composite program. The challenge in programming large-scale network systems is to be able to identify the set of local invariants (properties) that individual controllers/components can provide, and then to develop the rules that could be used to report back global invariants of the composite (or lack thereof) to the user.

Enabling technologies: In developing traffic/resource controllers, the right incentives should be built in to encourage rational users/protocols to behave in such a way as to maximize the network welfare. The modeling of such incentives should be mathematically tractable to enable the study of the transient and steady-state behavior of the adaptive controller. Security against attacks should also be an integral piece of the developed models so as to study the tolerance of the network system to not only random disturbances but also to adversarial attacks. These issues of incentives and security and their effect on self-organization must be more adequately studied as they permeate through all levels of the architecture of open dynamic network systems.

Science of Network Service Composition: At the conclusion of the workshop, it became evident that there is a need for a new science where the basic assumptions in established related theories, such as graph theory, optimization theory, network calculus and control theory, need to be reconsidered and extended to adequately reflect the openness and complex dynamics of large-scale network systems. For example, although basic control theory has been used successfully to study certain network models, more advanced control theoretic techniques such as stochastic non-linear control techniques may be needed to study more complex/realistic network models. Some of these extensions may actually end up introducing new theories to effectively deal with multiple time scales of measurement and control common in large-scale dynamic network systems.

This new science of composition may borrow from several analysis techniques (e.g. control theory, game theory, network calculus, percolation theory, economics, queuing theory). In essence, different techniques may provide different “languages” by which certain properties of system components can be expressed and composed into larger systems. Modeling and analysis have often followed after the design of traffic controls, and time is ripe for them to guide the design of the expected cyber-infrastructure. The science of modeling and safe programming of network systems through hierarchies of composable properties should have more impact if we are to succeed in designing a resilient network of the future.

5 Research Challenges

The following five sections describe specific technical research challenges identified in the five tracks covered during the workshop, namely wireless networking, measurements and modeling, overlay networking, large-scale architectures, and network programming. Each of these research challenges is described along with its scope and impact.

5.1 Research Challenges in WIRELESS NETWORKING: General Aspects, Ad-Hoc and Sensor Networks

Cross-Layer Research

Description: Understand the impact of outputs from one layer on functions of another, identify parameters that are shaped by functions in other layers and develop a methodology for a cross-layer design approach that exploits the inter-layer relationships. Quantify the benefits and advantages from relaxing the rigid layered structure, as well as study the associated complexity and stability issues. Design methodologies that blur the lines between protocol layers and attempt to optimize across layer functionality. From considering just new/additional triggers between the layers, to fully interleaving the layers, to melting the layers. It is especially noted that the networking community has not considered adequately issues related to the physical layer and their impact on higher layers.

Scope / Impact: Cross-layer design is strongly advocated for wireless networks where the characteristics of the wireless channel permeate the functions of all protocol layers in the traditional protocol stack. Other characteristics of a wireless networking environment (such as mobility, resource constraints, etc.) also introduce strong inter-layer relationships and effects that need to be understood and considered in the protocol design.

Cross-layer considerations and designs:

- Will improve the efficiency at the system level by improving the efficiency of protocol functions that utilize critical information from other layers or generate less burdening input to affected functions in other layers. For instance, MAC and routing layers are strongly coupled and sub-optimally designed if cross-layer issues are not considered.
- From the applications' point of view, it will enhance the networks' capability to support a wide range of applications (sensor and ad-hoc networks, etc.)

Multi-Radio Wireless Systems

Description: Design wireless communication systems equipped with more than one radio systems. These radios would have diverse capabilities / characteristics. The most appropriate radio will be invoked to perform a certain function. Radios with different properties cooperate with each other to accomplish a certain task (set of functions).

Scope / Impact: Such systems will help reduce inefficiencies inherent in single-radio systems that are not designed to optimally perform the entirety of needed functions. As a result, a variety of critical issues in wireless networks can be better addressed, such as energy consumption, mobility management, last-hop quality of service, data link robustness and capacity improvements. Thus, a wider range of applications can be supported in a wider range of (varying) environments. Multi-radio systems is an example of a system-level optimization across various system components (software and hardware), and as such it is a cross-layer design.

Revisit or Develop New Fundamental Theories

Description: Reconsider the basic assumptions in established related theories and extend them to adequately reflect the new environments. For instance consider graph theory extended to deal with unreliable/"soft" edges. Some of these extensions may actually end up introducing new theories to effectively deal with multi-resource, multi-constrained environments.

Scope / Impact: Determine the capacity of multi-hop, ad-hoc, varying networks. Consider the joint optimization of processing, communication and information quality, etc. The evaluation of proposed solutions lacks the necessary yardstick for their assessment since the performance and fundamental limitations for wireless networks are not sufficiently understood. Such studies will provide this yardstick as well as help us design more effective networks.

"Opportunistic" Networking

Description: Identify potential advantages in traditionally limiting or problematic situations and explore opportunities for efficiency. Develop self-organizing notions taking into consideration such advantages.

Scope / Impact: Some of the recent and encouraging findings include the realization that the properties of the wireless channel, which are generally perceived as negative, can in fact be exploited positively to provide unique advantages to wireless networking. These include the broadcast/multicast advantage of the wireless medium that creates new opportunities for information sharing, the selective frequency fading of wireless channels that permit the use of OFDM (Orthogonal Frequency Division Multiplexing), and the time-variability of wireless channel quality that can be exploited in access control through the so-called multi-user diversity. Similarly the directivity of antennas, the mobility of users, etc. can be exploited. These opportunistic approaches could enhance the capabilities of networks, making them work in otherwise impossible situations.

Trustworthy Networks

Description: Address the numerous security and privacy aspects that hinder the use of mobile ad-hoc networks

Scope / Impact: Additional research issues in wireless networks are created in the area of security due to the ubiquitous nature of the wireless medium that gives rise to new threats or makes well-known threats more acute and difficult to combat. For example, a not so well looked at aspect is “sending info and hiding in a wireless network”. Besides content privacy, movement privacy is an issue. Taking care of such issues will result in more secure and trustworthy networks.

Application Specificity

Description: Consider carefully the main characteristics of the applications when designing new networks.

Scope / Impact: Application specificity permeates the design and study of a wireless network. We cannot have an architecture – like the Internet – which is ubiquitous across all wireless environments. Different applications (e.g. military sensor nets, an environmental net) may require a totally different new architecture. This calls for programmable networks that can serve a wider class of applications, which gives rise to tension between programmability and efficiency. Networks, which are more effective in accommodating specific applications, could be instantiated.

Mobility Considerations

Description: Characterize, predict and exploit (cf. opportunistic networking) mobility to improve efficiency; understand its impact on security and traditional protocol functions at various layers.

Scope / Impact: Mobility can help address scalability problems, design efficient forwarding schemes and trust protocols, etc. The ability to efficiently maintain “elastic” high-performance connections in the presence of mobility will enable wireless networks to perform as well as infrastructure-based ones.

Higher-level Multi-hop Issues for Symbiotic Networking: Cooperation, Charging, Trust and Reputation

Description: Given that each node has its own authority in cooperative networks, questions remain as to how can cooperation be encouraged or enforced; how can nodes with “good” reputation get rewarded; and how can such mechanisms change in the face of changing application requirements. We need to identify the unique wireless aspects compared to wired P2P environments. How can we leverage mobility? How can key revocation be done? Can we flexibly control the degrees of self-organization and cooperation?

Scope / Impact: Developing mechanisms for symbiotic networking is key for the widespread use of ad-hoc networks.

Lower-level Multi-hop Issues for Symbiotic Networking: New challenges ...

Description: Low-level issues in cooperative networks include those new challenges found in sensor networks, multi-hop cellular networks, ad-hoc, and multi-radio systems. Given the desire for global connectivity, we are faced with the challenge of integrating such networks with traditional wired networks; how to realize end-to-end functionalities (e.g. error control). Most studies have relied on heuristics. There is a need to develop a “science” so as to develop the appropriate protocols for each specific application.

Scope / Impact: Symbiotic networks, e.g. sensor networks, will be part of our lives. Examples include smart spaces and disaster relief networks. Developing the science to instantiate them for different applications and integrate them into existing infrastructure is key to their cost-effectiveness and success.

Spatial and Temporal Correlations and Issues

Description: Sensor networks give rise to unique needs for localization protocols to support location-dependent processing, and for data fusion algorithms to aggregate data inside the network.

Scope / Impact: Data fusion is very application specific, but the general problem is how to exploit temporal as well as spatial correlation to collect related events and efficiently communicate results to the user. There is a need for data structures that support spatial indexing and computing aggregates. Then a general-purpose programming system, based on these data structures, can be developed. Localization protocols are also important enabling technology to realize location-dependent functionalities without the need for a GPS system. For example, robust energy-aware geographic routing systems can be achieved using only local positioning systems. Exploiting temporal and spatial correlations enable the design of protocols that conserve energy; a significant constraint in battery-operated network nodes.

5.2 Research Challenges in TOMOGRAPHY, MEASUREMENTS, and MODELING

Measurement-based Modeling Theory and Control

Description: Theory that allows one to combine elements of descriptive modeling (e.g. HMM) and constructive modeling (e.g. queuing models) based on a mix of end-to-end and internal network measurements. This measurement-based characterization can be used for on-line traffic management.

Scope / Impact: A theory is needed to more accurately model the characteristics of network paths and elements (wired, optical, RF, acoustic, etc.) with respect to delay, loss, bandwidth, etc. The impact of both scale and heterogeneity should be investigated. There is a need to understand the tradeoff between intrusiveness and accuracy of measurements and how this affects various traffic management functions. Advances in this area would lead to better performing applications that are aware of the network state.

Scalable Measurement and Inference Techniques

Description: Inference for large sized networks and merging of topologies. Coordinate measurements at different layers and from multiple vantage points. Use of effective sampling/probing techniques.

Scope / Impact: A theory based on a richer network model and probing structures would make measurements and inference simpler and faster, thus yielding a distributed measurement infrastructure that is effective and less intrusive. Furthermore, a “measurement science” provides the tools needed for a “science of network design.”

Development of Benchmarks and Validation Tools

Description: Data is needed to validate hypothesis of prior models. Benchmarks need to be available for head-to-head comparison of various solutions.

Scope / Impact: For research projects to be community driven, members of the community should use a common set of benchmarks and data validation tools. A common set of benchmarks and validation tools will promote a sense of community behind specific research areas.

Composition of Large-scale Systems

Description: Modeling and analysis should guide, not follow. We need to learn from several domains (queuing theory, control theory, game theory, economics, percolation theory, large deviations, diffusion models, chaotic theory, information theory, even type systems of programming languages) to be able to compose and analyze large-scale networks.

Scope / Impact: A “science of design” is needed that allows us to model large-scale systems at different levels of abstraction (macro- and micro-levels) that can be combined into a unified model of the entire system. Advances in “science of design” will lead to a new network theory and hybrid analytical-numerical-simulation tools that can effectively handle multiple time scales of measurement and control, and non-linear interactions at different levels. This will undoubtedly lead to a robust/resilient cyber-infrastructure.

Distributed Algorithms

Description: Research is needed to compute theoretical bounds on the performance of large-scale systems and develop distributed implementations to achieve these bounds.

Scope / Impact: Research on distributed algorithms is needed to manage and coordinate resources distributed over large scale for purposes of self-stabilization, decentralized trust, etc. Advances in the theory of distributed algorithms will enable the fundamental understanding of how various algorithms apply to different kinds of networks (e.g., wired vs. ad-hoc) and how to accordingly tune their performance.

Economic Incentives in Resource Management

Description: Incentives are useful mechanisms to encourage rational users/protocols to behave in such a way as to maximize the network welfare. The concept of utility functions has been used successfully to understand incentives implicit in existing Internet protocols (e.g. TCP). A challenge is to apply this concept for other networks (e.g. wireless and QoS networks) where the utility function becomes non-convex.

Scope / Impact: Incentive-based resource management may be applied in various application domains, from new P2P to traditional Web.

New Network Theory

Description: Research on new network theories is needed to deal with multiple scales of control. Such theories will allow us to understand the dynamics of large-scale networks and how events (e.g. congestion) can propagate creating time-varying waves through the network.

Scope / Impact: It is important to create a theory that allows a network designer to apply the right analysis technique at the right time scale. For example, it is often the case that standard analysis (e.g. queuing theory and control theory) techniques can be applied under certain conditions (e.g. at the connection level as opposed to the level of packet dynamics, or when insensitivity of averages applies). A theory that encompasses new techniques (e.g. from game theory and economics) can enable the study of more complex phenomena resulting from interacting non-linear network elements or protocol layers.

5.3 Research Challenges in PEER-TO-PEER AND OVERLAY NETWORKS: Overlay networking and Content Distribution

Network Aware Overlays

Description: Robust large-scale P2P systems need to efficiently use underlying network services; this is a form of a cross-layer optimization.

Scope / Impact: Matching of overlays to network-layer infrastructure is necessary for improved latency, resiliency and resource utilization. This would involve network-aware distributed hash tables and overlay sockets, as well as network-level control and redundancy to enhance the overlay robustness, in particular to vulnerability attacks. Dynamic on-demand synergy between network infrastructure and an overlay is highly promising, though tremendously complex to achieve. Successful R&D in this area should have multidisciplinary nature with deep involvement of all stakeholders. The networking market may be mostly impacted by what will be traded at the network and overlay levels. Design trade-off between economy of scale and scalability has to be explored first.

Trustworthy Community-based Infrastructure

Description: Self-organization and rapid deployment of novel services within overlay networks where infrastructure, service components and content are donated mainly by motivated users themselves.

Scope / Impact: Develop adequate trust models that would work in dynamic infrastructure-less environments and be driven by user incentives and by self-organization rules. Develop trustworthy systems that are cost-effective in the sense of tradeoff between quantifiable error rate and security level.

Dynamic Content Systems

Description: Scalable, dynamic, personalized and interactively controlled distribution of digital content to potentially large and heterogeneous groups of receivers.

Scope / Impact: Organization and management of dynamic content has ever-increasing role in traditional CDNs, in push technologies for mobile and wireless, in information gathering systems, last but not least in large-scale multi-user interactive games. DCSs are found in both user- and control- overlays with similar requirements: global accessibility and local management (exactly as in DNS!). Successful research in DCS combined with neighboring areas may yield truly novel types of systems that will permanently and instantly gather, index, process, describe, distribute, retrieve and consume digital items of unlimited nomenclature coming from and consumed by unrestricted set of sources and sinks, while preserving SLA conditions including digital rights management when applicable.

Scalable Group Communication

Description: Application layer multicast as infrastructure-less, address-less, fully reliable, programmable and controlled group communication.

Scope / Impact: SGC is important for the Internet and other types of networks. After 15 years of development, native IPv4 multicast failed to take-off mainly due to the lack of incentives from network providers. When powered by motivated user community, SGC may become a wide spread network service. Nearly the same SGC technology may be used by overlays providing GRID services, and by an infrastructure itself, e.g. for distributed cache updates and synchronization; finally for network control plane tasks. Impact can be tremendous, especially when combined with self-organization, dynamic trust establishment and programmability.

5.4 Research Challenges in LARGE SCALE NETWORKS AND FUTURE NETWORK ARCHITECTURES

Control Plane Complexity

Description: Development of an understanding of complexity in the control plane, where management and security are essential ingredients, to support yet unknown Internet service architectures.

Scope / Impact: Control plane is a cross-layer issue, including downward hardware e.g. for signaling and QoS support, and upward overlays e.g. for overlay and network layer routing interaction. Understanding how to keep control plane complexity manageable will create ever increasing opportunities to solve:

- *Traditional problems:* robustness, adaptability, safety, predictability, evolvability, and security, and
- *New problems:* such as
 - revising a notion of network function that will allow for viewing a network as a Complex Functional System (*The CFS is characterized by a presence of a constant (invariant) task, performed by variable mechanisms, bringing the process to a constant (invariant) result*) that is characterized not only by complexity of its structure but also by the mobility of its component parts;
 - on-demand coupling of different functional systems that will allow true cooperative networking (seamless inter-working of networks with different technologies and business models).

Separation of Concerns

Description: De-coupling of node's address and its identifier/name and de-coupling of control and data planes.

Scope / Impact: Novel architectures, in which name and address are two independently designed concerns, could be a key for self-organizing networks (e.g. for routing); similarly, control and data plane functions should be separate concerns to isolate vulnerabilities. In general, separation of concerns improves all of the 'ities'. With this, networks would benefit from fine-grained [distributed] state management, consistent with the belief that "the larger the system and the higher the understanding the finer is the state management granularity".

Common Transport Platform

Description: New end-to-end transport layer protocols need to handle varying reliability, ordering and timing semantics of different classes of multimedia applications.

Scope / Impact: This issue is a transport specific CFS (Complex Functional System); this entails moving the different semantics from application to transport. The next step will probably be to organize them as a transport on-demand service relying on a common transport platform. The hope is to solve the long-lasting issues of the Internet – multimedia group communication and QoS guarantees for sensitive flows across heterogeneous networks (wired/wireless and cooperative).

5.5 Research Challenges in PROGRAMMING THE INTERNET

Programming the Internet

Description: Network programmability could speed up the evolution of network services, and improve the scalability of group applications through network assistance.

Scope / Impact: Programmability can help improve the Internet legacy architectures (QoS, routing, overlay and underlay architectures) as well as radical innovations (e.g. role-based architectures, non-IP communication other than client-server type). Programmability leads to flexibility needed to overcome ossification while keeping a stable fast path and improve scalability (e.g. through programmable fusion and hierarchical scaling).

Ambient Policy

Description: Creation, management and enforcement of policies, from a wide adoption of traditional security and QoS configuration policies to more active and proactive policies defining a choice in the behavior of the system.

Scope / Impact: From operational experience seems that policy-based approach is a promising theoretical framework for distributed management. However, radical re-thinking of all existing networking functionalities is needed so that they become policy-ready; ontology research should probably investigate other paradigms, e.g., self-organizing ontology's. A policy system, viewed as the nervous system of a network, shall play the primary role in defining principles of CFS (Complex Functional System).

Programmable Router

Description: High-end routers are becoming more and more decentralized. There is a need to map forwarding and routing applications onto more complex multi-processor architectures.

Scope / Impact: High-end routers continue to grow in terms of total throughput and interfaces supported. We need better system-level understanding of how to map highly-pipelined forwarding and routing computations onto what are essentially high-performance distributed systems comprising large numbers of channel interface modules, each with a hierarchy of processing elements with different levels of sophistication and throughput characteristics. Automated mapping of processing and routing requirements on today's complex router architectures avoids software failures, which are responsible for a large percentage of the observed failure rate.

Programming of Rule-Based Systems (RBS)

Description: Business objectives and preferences are expressed in policies whose enforcement requires that policies be expressed in a mechanism-friendly way (as rules). Being embedded into the operational structure, rules are hard to change. Novel approaches are needed to allow RBS programming.

Scope / Impact: RBS programming would range from programming a basic security system (e.g., access control), to on-demand creation of complex distributed behaviors involving hundreds of network devices exhibiting required pieces of behavior and collectively performing a needed system control or invariant function in CFS. RBS is a good candidate for non-conventional experiments with novel architectures; scalability of the approach is inherently good due to the use of fine-grained, self-organization mechanisms in controlling large-scale systems.

6 Workshop Session Summaries

6.1 Session title: **WIRELESS NETWORKING: General Aspects**

Chair: Anthony Ephremides

Speakers: Anthony Ephremides, Victor Bahl, Luigi Fratta, Chris Blondia

Wireless Networks include the traditional paradigm of cellular architectures that permit one-hop connectivity to the infrastructure and the relatively more recent paradigm of ad hoc networks. The latter include wireless LANs that have the property of full mesh connectivity among their members and the multi-hop ad hoc networks that require relaying through intermediate nodes. What they both share is independence from the infrastructure, to which they may connect via a gateway. In addition, there are networks that include satellites. These have special properties but share also several characteristics with the aforementioned network types.

In this section we outline what are the major research issues that are common to all wireless networks. It should be pointed out that a special case of ad hoc multi-hop networks includes a large number of sensor networks that utilize randomly deployed, battery-operated sensors. The special classes of wireless networks create additional requirements and raise new research issues that are detailed in other sections.

The key ingredients of wireless networking that generate new research challenges are:

- The use of the wireless medium at the physical layer
- The necessary sharing of that medium
- User mobility
- Reliance on portable energy sources.

In turn, these ingredients translate into the following networking requirements:

- Variable topology
- Interference management or access control
- Efficient use of available energy

It follows from these requirements that wireless networks must be designed so as to enable seamless communication under these stressful constraints. Hence, new and unexpected research issues arise, the resolution of which will enable wireless networks to perform as well as the infrastructure-based ones.

Specifically, the characteristics of the wireless channel permeate the functions of all protocol layers in the traditional OSI stack. Hence, cross-layer research is necessary to exploit the inter-layer relationships.

The wireless channel cannot guarantee connectivity. Whether a successful link can be established between two points depends on transmission power, bit rate, channel condition, and other-user interference. Thus, the use of graph theory that has been successful in studying the topology of wireline networks must be revised. A new graph theory is needed that is

node-centric and that is characterized by “soft” edges, the presence of which depends on network parameter choices, including the presence of adjacent links.

User mobility requires new methods for maintaining connectivity. Mobile IP has been successful in the cellular paradigm. In ad hoc, multi-hop networks, new techniques are needed that can create elastic connections. The nature of the wireless medium can be exploited here since power control and rate control along with judicious MAC protocol design can provide the needed elasticity.

Routing in multi-hop networks is inextricably coupled to MAC (a key example of cross-layer coupling). A given routing algorithm defines the flows between node-pairs. These flows are inputs to the MAC protocol to properly allocate bandwidth. The allocation of bandwidth modifies the link states of the affected links, which in turn modify the link metrics needed for routing. Hence, the selection of routes must reflect the effects of the rules of the MAC protocol.

Wireless networks will often require multiple radios in order to provide the needed flexibility to the nodes. The use of UWB (ultra wide band) techniques, multiple antennas, space-time coding, and innovative spectrum sharing and exploitation methods are ways in which the performance of wireless networks can be improved. Some of the recent and encouraging findings include the realization that the properties of the wireless channel, which are generally perceived as negative, can in fact be exploited positively to provide unique advantages to wireless networking. These include the broadcast/multicast advantage of the wireless medium that creates new opportunities for information sharing, the selective frequency fading of wireless channels that permit the use of OFDM (Orthogonal Frequency Division Multiplexing), and the time-variability of wireless channel quality that can be exploited in access control through the so-called multi-user diversity.

Additional research issues in wireless networks are created in the area of security due to the ubiquitous nature of the wireless medium that gives rise to new threats or makes well-known threats more acute and difficult to combat.

All-in-all, wireless networks have opened up new opportunities for worldwide connectivity and have revitalized the field of networking by posing many novel challenges. Last, but not least, is the realization that energy efficiency can be achieved at all layers of protocol design and that the emergence of standards, although useful and welcome, should not stifle unimpeded free thinking for future designs.

In closing, we must emphasize, that just as in wireline networking, performance and fundamental limitations for wireless networks remain unrevealed and thus the evaluation of proposed solutions lacks the necessary yardstick for their assessment. Network Information theory may in fact be easier to develop in wireless networks, since the theoretical investigation of the wireless medium is quite advanced and has scored already notable successes.

6.2 Session title: WIRELESS NETWORKING: Mobile Ad-Hoc Networks

Chair: Mario Gerla

Speakers: Mario Gerla, Edward Knightly, Jean-Pierre Hubaux, Marco Conti,
Per Gunningberg

6.3 Session title: WIRELESS NETWORKING: Sensor Networks

Chair: Ian Akyildiz

Speakers: Ian Akyildiz, Ramesh Govindan, Patrick Thiran, Erdal Cayirci,
Magda El Zarki

1. INTRODUCTION

The technological advances in the micro-electro-mechanical systems (MEMS) and the wireless communications have enabled the deployment of the small intelligent sensor nodes at homes, in workplaces, supermarkets, plantations, oceans, streets, and highways to monitor the environment. The realization of sensor networks to improve the efficiency of nearly every aspect of our daily lives by enhancing the human-to-physical world interaction is one of the most exciting potential sensor network applications utilizing these intelligent sensor nodes. However, this objective necessitates the efficient and application specific communication protocols to assure the reliable communication of the sensed event features and hence enable the required actions to be taken by the actors in the smart environment. In this session, the challenges and the existing solutions for the design and development of sensor network communication protocols are presented. More specifically, application layer, transport layer, network layer, data link layer, in particular, error control and MAC protocols, and physical layer issues as well as the localization protocols and the time synchronization algorithms are discussed in detail. Open research issues for sensor networks are also presented which will be outlined below.

2. SENSOR NETWORKS vs. AD HOC NETWORKS

While there are similarities between the sensor networks and traditional ad hoc networks the communication protocols and algorithms proposed for traditional wireless ad hoc networks may not be adequate in addressing the unique challenges posed by sensor networks. The major differences between sensor networks and ad-hoc networks are summarized below:

- The number of sensor nodes in a sensor network can be several orders of magnitude higher than the nodes in an ad-hoc network.
- Sensor nodes are densely deployed.
- Sensor nodes are prone to failures.
- Sensor nodes mainly use broadcast communication paradigm whereas most ad-hoc networks are based on point-to-point communications.
- Sensor nodes are limited in power, computational capacities, and memory.

- Sensor nodes may not have global identification (ID) because of the large amount of overhead and large number of sensor nodes. Sensor networks are deployed with a specific sensing application in mind whereas ad-hoc networks are mostly constructed for communication purposes.

3. APPLICATION LAYER ISSUES

Although many application areas for sensor networks are defined and proposed, potential application layer protocols for sensor networks remain largely unexplored. Three possible application layer protocols are known: Sensor Management Protocol, Task Assignment and Data Advertisement Protocol, and Sensor Query and Data Dissemination Protocol.

System administrators interact with sensor networks by using sensor management protocol (SMP), which performs the following administrative tasks:

- Introducing the rules related to data aggregation, attribute-based naming and clustering to the sensor nodes
- Exchanging data related to the location finding algorithms,
- Time synchronization of the sensor nodes,
- Moving sensor nodes,
- Turning sensor nodes on and off,
- Querying the sensor network configuration and the status of nodes, and re-configuring the sensor network, and
- Authentication, key distribution and security in data communications.

Another important operation in the sensor networks is interest dissemination through task assignment and data advertisement protocol. This interest may be about a certain attribute of the phenomenon or a triggering event. Another approach is the advertisement of available data in which the sensor nodes advertise the available data to the users, and the users query the data that they are interested in.

The Sensor Query and Data Dissemination Protocol (SQDDP) provides user applications with interfaces to issue queries, respond to queries and collect incoming replies. Note that these queries are generally not issued to particular nodes. Instead, attribute-based or location-based naming is preferred.

Although SCTL is proposed, different types of SQDDP can be developed for various applications. The use of SQDDPs may be unique to each application.

Moreover sophisticated Graphical User Interface is needed for sensor networks.

4. TRANSPORT LAYER RESEARCH ISSUES

The main objectives of the transport layer and its desired essential features to address the unique challenges posed by the characteristics of sensor networks paradigm are:

Reliable Transport: Based on the application requirements, the extracted event features should be reliably transferred to the sink. Similarly, the programming/retasking data for sensor

operation, command and queries for operation should be reliably delivered to the target sensor nodes to assure the proper functioning of the wireless sensor network.

Congestion Control: Packet loss due to congestion can impair event detection at the sink even when enough information is sent out by the sources. Furthermore, congestion control not only increases the network efficiency but also helps conserve scarce wireless sensor resources.

Self-Configurability: The transport layer protocols must be adaptive to dynamic topologies caused by node mobility/failure/temporary power-down, spatial variation of events and random node deployment in sensor networks.

Energy Awareness: The transport layer functionalities should be energy-aware, i.e., the error and congestion control objectives must be achieved with minimum possible energy expenditure.

Biased Implementation: The algorithms must be designed such that they mainly run on the sink with minimum functionalities at sensor nodes in order to conserve limited sensor resources and shift the burden to the high-powered sink.

Constrained Routing/Addressing: Unlike protocols such as TCP, the transport layer protocols for sensor networks should not assume the existence of an end-to-end global addressing. It is more likely to have attribute-based naming and data-centric routing which call for different transport layer approaches.

In summary, the transport layer mechanisms that can address the unique challenges posed by sensor networks are essential to achieve the potential gains of the collective effort of sensor nodes. Nevertheless, the existing and currently being developed solutions need to be exhaustively evaluated under the real sensor network deployment scenarios to reveal their shortcomings and hence, necessary modifications/improvements to obtain a complete transport layer solution for sensor networks.

5. NETWORK LAYER ISSUES

The design of the network layer used by the sensor nodes follows the following principles:

Power efficiency: The power consumption rate affects the life-time of the network. As a result, the routing protocol has to conserve as much power as possible.

Data-Centric: This concept enables information to be represented in a different way using attribute-based naming schemes. In addition, it may become central in designing a routing protocol.

Data Aggregation: This technique may enhance the performance of the routing protocol. It is useful only when it does not hinder the collaborative effort of the sensor nodes, so care must be taken.

Network Integration: The routing protocol should be easily integrated with other networks, e.g., Internet or networks that are within the smart environment.

The above principles serve as a guideline in designing a routing protocol for wireless sensor networks.

6. MEDIUM ACCESS CONTROL ISSUES

The distributed nature of the sensor networks and the application-oriented traffic properties make traditional MAC protocols impractical. The low cost requirements and the distributed nature of the sensor nodes constrain the energy consumption of all the layers. Hence, energy efficiency is of primary importance for the MAC layer protocol design. The MAC layer protocol should ensure that nodes transmit their information with minimum energy consumption which can be achieved by minimizing idle listening times and collisions among sensor nodes.

In traditional networks, per-node fairness is an important aspect of the MAC layer protocol due to the competitive nature of the nodes. Hence, MAC layer protocols should take a collaborative approach so that the overall information surpasses the individual capabilities of each node. In addition, the application specific information should be incorporated into the MAC approach to enhance the performance further.

The topological information of the network is another factor that should be incorporated into MAC layer protocol design. Since large number of sensor nodes can be deployed in sensor networks, the high density of the network should be exploited. The increasing density increases the number of nodes in reach of a sensor node which is both a disadvantage and an advantage to the MAC layer protocol. As the network density of the network increases, the number of nodes contending with each other increases resulting in higher collision probability. On the other hand, the connectivity of the network can be provided without compromising from the total energy consumption due to the high number of neighbor nodes. In addition, due the high density of the sensor nodes, the information gathered by each node is highly correlated. Exploiting the correlation between sensor nodes also in the MAC layer protocol can be a promising approach to further improve overall network performance. As far as the topology changes are concerned, the MAC layer protocol should also provide the procedures for nodes to join and leave the network while providing connectivity throughout the network.

7. ERROR CONTROL ISSUES

The use of FEC is the most efficient solution given the constraints of the sensor nodes. Although the FEC can achieve significant reduction in the bit error rate (BER) for any given value of the transmit power, the additional processing power that is consumed during encoding and decoding must be considered when designing an FEC scheme. The FEC is a valuable asset to the sensor networks if the additional processing power is less than the transmission power savings. Thus, the tradeoff between this additional processing power and the associated coding gain need to be optimized in order to have powerful, energy-efficient and low-complexity FEC schemes for the error control in the sensor networks. As researchers continue to investigate new FEC schemes for sensor networks, the designers must bear in mind that the new schemes may be application specific.

The link layer overall still remains a challenging area to work in since sensor nodes are inherently low-end. Combining the low-end characteristics of the sensor nodes with harsh

deployed terrains, collaborative approach and exploiting the correlation between sensor nodes, it calls for new medium access as well as error control schemes.

8. PHYSICAL LAYER ISSUES

For sensor networks, a small sized, low cost, ultra-low power transceiver is required. Certain hardware constraints and the tradeoff between antenna efficiency and power consumption limit the choice of a carrier frequency for such transceivers to the ultra high frequency range. The use of the 433 MHz ISM band in Europe and the 917 MHz ISM band in North America are also proposed. The main advantages of using the ISM bands are the free radio, huge spectrum allocation and global availability. They are not bound to a particular standard, thereby giving more freedom for the implementation of power saving strategies in sensor networks.

Another possible mode of inter-node communication in sensor networks is by infrared. Infrared communication is license-free and robust to interference from electrical devices. Infrared based transceivers are cheaper and easier to build. The main drawback is the requirement of a line-of-sight between the sender and receiver. This makes infrared a reluctant choice for transmission medium in the sensor network scenario.

Furthermore new channel models for indoor, outdoor underwater, deep space need to be developed. Moreover, the new antenna techniques (e.g., smart antennas) as well as the applicability of software radios must be explored. New modulation schemes, synchronization schemes, data encryption as well as FEC schemes on the bit level must be researched.

Physical layer remains a vastly unexplored domain of sensor networks.

9. LOCALIZATION PROTOCOLS

The purpose of a localization protocol is to enable sensor nodes to determine their locations, so objects can be located by the users. In addition, services can be provided when the users are moving. It is foreseeable in the near future that many of the objects will be tagged. These tags may enhance the quality of services to the users. One type of such services is real-time tracking and locating a person within a building with active badges.

Regardless if the objects are tagged or not, the sensor nodes may be randomly or strategically deployed. These sensor nodes must be aware of their locations in order to provide meaningful data to the sinks. In addition, location information may be required by the network and data link layer protocols.

The challenges in designing a localization protocol for sensor networks are as follows:

- Robust to node failures.
- Less sensitive to measurement noise.
- Low error in location estimation.
- Flexible in any smart environment.

Currently, there are two types of localization techniques that address these challenges: (i) beacon based and (ii) relative-location based.

The beacon based technique requires the sink to broadcast the relative location to the beacons; in turn, the beacons broadcast their locations to their neighbor nodes that are within the radius R . As for the relative-location based technique, the location of the sensor nodes are determined hop-by-hop relative to the sink. Both types of localization techniques may use range and angle estimations to estimate the location via received signal strength (RSS), time-of-arrival, time-difference-of-arrival, (TDOA) and angle of arrival (AOA).

Whether the beacon or relative-location based localization protocols are used, it is the location information that is required by the protocols in the transport, network, and data link layers. Each type of localization protocols offers different capabilities. Depending on the applications, both types of localization techniques may be used.

10. SYNCHRONIZATION ISSUES

The challenges such as robust, energy aware server-less, tunable service, provide a guideline for developing various types of time synchronization protocols that are applicable to the sensor networks. A time synchronization protocol may have a mixture of these design features. In addition, some applications in the sensor networks may not require the time synchronization protocol to meet all these requirements. For example, a data gathering application may require the tunable service and light-weight features more than the server-less capability. The tunable service and light-weight features allow the application to gather precise data when the users require it. In addition, the nodes that are not part of this data gathering process may not have to be synchronized. Also, the precision of the time does not need to be high, because the users may only need milliseconds precision to satisfy their needs.

As these design challenges are important for guiding the development of a time synchronization protocol, the influencing factors that affect the quality of the synchronized clocks have to be investigated. Although the influencing factors, temperature, phase noise, frequency noise, asymmetric delay, are similar to existing distributed computer systems, they are at different extreme levels.

Since sensor nodes are randomly deployed and their broadcast ranges are small, the influencing factors may shape the design of the time synchronization protocol. In addition, the links between the sensor nodes may not be reliable. As a result, the influencing factors may have to be addressed differently.

The requirements of sensor networks are different from traditional distributed computer systems. As a result, new types of timing techniques are required to address the specific needs of the applications. Since the range of applications in the sensor networks is wide, new timing techniques are encouraged for different types of applications. This is to provide optimized schemes tailored for unique environments and purposes.

11. CONCLUSIONS

The design-principles, open research issues of developing application, transport, network, data-link, and physical schemes as well as localization and timing techniques are described. They are to guide and encourage new developments in the sensor network area. As the technologies for sensor networks are advancing, the pervasive daily usage of these networks is foreseeable in the near future.

REFERENCES

Akyildiz, I. F., Su, W., Sankarasubramaniam, Y., and Cayirci, E., "Wireless Sensor Networks: A Survey," *Computer Networks (Elsevier) Journal*, vol. 38, no. 4, pp.393-422, March 2002.

6.4 Session title: NETWORK TOMOGRAPHY AND TRAFFIC MODELING, MEASUREMENTS AND CONTROL: Network Traffic Modeling and Control

Chair: Jim Roberts

Speakers: Jim Roberts, Peter Key, Jorg Liebeherr, Gabor Vattay, Ravi R. Mazumdar

This session included the following presentations:

J. Roberts: Important research problems in network traffic modelling and control

P. Key: Modelling and controlling ... the future?

J. Liebeherr: Statistical network calculus

G. Vattay: Interacting network elements

R. Mazumdar: Non-convex optimization and resource allocation in communication networks

In this section, we aim to identify the important issues raised and highlight the key research directions. In a final paragraph, we respond to a number of questions posed by the workshop chairs in the introductory session. Mathematical modeling of traffic and design of traffic controls are essential research activities in the development of a future QoS enabled network. Traffic modeling is necessary both to determine how much capacity the network should have to handle forecast demand and to specify the mechanisms by which that capacity is shared by different applications and users. It is often to be regretted that design of traffic controls and QoS mechanisms generally seems to precede traffic modeling work, often leading to sub-optimal network architectures and service models. Current QoS architectures rely on the notion of a traffic contract where users previously declare their traffic characteristics and the network performs admission control based on this declaration and the required performance targets. An extensive "network calculus" has been developed over recent years enabling the realization of deterministic performance guarantees. The work presented on statistical network calculus relaxes the worst case assumptions necessary for deterministic guarantees leading to considerably less conservative resource requirements (Liebeherr). Current research aims to extend single link modeling to the network context and to derive simple and efficient provisioning algorithms. An alternative point of view is to consider an a priori traffic description to be impractical and to rely essentially on measurement-based admission control. Admission control indeed appears as the key to quality of service and a number of promising approaches are being researched. End-point or distributed admission control relies on users

(or their "agents") testing availability by means of probe packets (Key). In another approach, admission control is deemed necessary for both elastic data flows and streaming flows (Roberts). This leads, in an integrated system, to simpler admission control since streaming traffic performance is protected by priority queuing while elastic flows are quite tolerant to admission control "errors". Use of implicit admission control and implicit service differentiation is advocated in an enhanced best effort architecture using a so-called "Cross-protect" router.

A large amount of interesting work has been performed lately on how one might use economic incentives to govern resource sharing, relying on user reactions to very simple network mechanisms like ECN marking. Under certain traffic assumptions, resource sharing under the control of TCP is shown to maximize a certain utility function. Alternative notions of utility result in other bandwidth sharing mechanisms. An interesting mathematical challenge arises when the utility function is not convex, as occurs in some wireless networks and wireline networks with QoS (Mazumdar). Consequences on the design of distributed algorithms are discussed.

A significant feature of network traffic modeling is the need to account for phenomena occurring over a wide range of time scales. It is frequently useful to practice a time scale separation when studying different kinds of congestion events. For example, one can study the detailed behavior of packet scale phenomena by assuming connections are permanent. This approach is used to study the behavior of the congestion control algorithms of TCP as multiple connections compete for bandwidth on a network path (Vattay). Applying techniques from the theory of chaos and non-linear systems it is shown for example how congestion can propagate through a network. An alternative modeling approach is to ignore the fine detail of packet scale interactions and to just assume these can be controlled to realize some kind of bandwidth sharing objective. A particularly interesting observation is that fair sharing (a common objective) is less effective than a policy like SRPT (shortest remaining processing time first scheduling) which favors short flows over long flows (Key).

The workshop chairs raised a number of issues on traffic modeling and control at the start of the workshop to which the talks and ensuing discussion bring some responses. Is standard queuing theory enough? - Yes, frequently; it is important to identify the appropriate time scale where the classical models apply (e.g., session arrivals are Poisson while packet arrivals are very complex); it is important to devise control schemes where standard insensitivity results can allow us to conclude that performance does not depend on detailed traffic characteristics. Is first course control theory enough? - Again, this is often true; it is noted in particular that the stability properties of TCP and related protocols have been studied using classical results from control theory; however, there remain open problems in this context (Key) and advanced techniques for studying interacting non-linear physical systems have their application (Vattay); additional mathematical tools that are increasingly used are game theory and economics (Key). How can we understand time scale interactions and design for them? - successful modeling relies on time scale separation; it is important however to take account of all time scales before drawing general conclusions; for example, to optimize the utility of bandwidth sharing for permanent connections may be sub-optimal when accounting for the fact that connections are of finite duration and need to be represented as a stochastic process at flow and application time scales. Are time varying averages enough or do we need to capture higher moments? - averages are sufficient when classical insensitivity results can be applied and this is clearly a useful result; time varying averages may be important if it is not reasonable to assume processes are stationary when for example some file transfers can last well beyond the busy period. What happens to effective bandwidth and good old Erlang? - It is indeed

important to remember the lessons of teletraffic theory as well as the modeling work performed for ATM; effective bandwidths are useful for sizing; they also have their role in the evaluation of currently proposed traffic controls (Liebeherr). Why are we insisting on modeling TCP and RED? - TCP is the key to congestion control in the current net and it is important to understand its behavior; RED and other active queue management mechanisms can make bandwidth sharing more efficient; use of ECN, where marking is necessarily determined by AQM, has significant advantages over packet loss; it is necessary to modify TCP for very high transfer rates in order to avoid loss of throughput occurring with the current AIMD congestion control algorithm. Why did it take us some time to be more careful? Unfortunately, traffic modeling still has relatively little impact on network design; this is typified by the recent quote from the end2end list on the non-usefulness of M/M/1 models (Key) and by a famous quote from the Diffserv mailing list "we don't have the math so lets not bother [trying to model...]"!

6.5 Session title: PEER-TO-PEER AND OVERLAY NETWORKS: Peer-to-Peer Networking

Chair: Jon Crowcroft

Speakers: Jon Crowcroft, Antony Rowstron, Don Towsley, Brian Levine

Other contributors: Sugih Jamin

Introduction/Overview

Ad Hoc wireless networks and peer-to-peer application networks and overlays share a common problem with the early Internet - the infrastructure is contributed by the users themselves, rather than by a central entity or set of entities such as government or commercial providers. This leads to a potential problem of free-riding, which has been well identified in the p2p world, but not yet observed so much in the Ad Hoc wireless world (unless "war-chalking" can be considered such!). Other problems with availability (of network and content), authenticity of content, privacy (theft of content and identity), and denial of service have been observed. Yet these types of systems have a "viral" nature (as observed by David Reed and others at the MIT Media Lab.) and as the metaphor shows, an ability to grow rapidly to provide useful services without a large scale investment (and associated risk) in deploying an infrastructure first (albeit they rest on some sort of infrastructure, where the phone lines (in the Internet case), or the Internet and home PCs (in the P2P case) or the handsets and laptops and home machines (in the community ad hoc wireless case), as well as the liberal supply of power:-).

So the problem ahead of us is: How do we create a set of systems that allow an exit strategy from the loose knit community network, to a commercial service, or else, allow some sort of aligning of incentives between peers as they switch role between provider and subscriber?

The sessions covered a number of critical topics where a good deal of research is still needed:

Firstly, we have the question of the "best-of-breed" in the actual distributed systems architecture of a p2p system: structured versus unstructured systems choices are discussed by Ant Rowstron.

Then we have the problem of search and query routing on P2P DHT versus efficient flood; event and group communications over p2p; locality and efficiency.

P2P blends into CDN and here we raise the question of "rights management" and anonymity: rights to use resources in and to/from a peer, rights over content, rights over identity/anonymity, and rights to charge for content and service, and the effect all these have on any and all given p2p technologies scaling and "peer-ness"...

Systems without a clear Provider/Subscriber (a.k.a client/server, consumer/business, or even master/slave) relationship have to rely on more distributed social means of organizing rights, obligations, responsibilities and cares: Incentives for sharing content and other resources (storage, transmission, reception, battery); mechanisms for exchanging tokens of trust, payment and reputation; mechanisms for protecting those tokens against theft, counterfeit; algorithms for policing p2p systems.

A general topic for overlays (and blurs into p2p and CDN) is "evolutionary paths for overlay technology" - here we could talk about competing overlays, overlay as function-not-layer, overlay APIs, overlay form v. function v. performance, etc...

Finally, p2p represents a new way to structure applications, and therefore generates new traffic patterns - these needs to be measured - there are many measurement projects underway on legacy internet traffic (HTTP/TCP etc) but few on many-to-many applications, and it seems like a rich seam to mine.

We should recall the myth of the Labyrinth of Minos, built by Daedalus, father of Icarus, and the overlay routing system designed by Ariadne so that Theseus could find his way out after killing the Minotaur (this being all in Crete). We could also discuss the use of wireless communication (Daedalus and Icarus flew to the mainland in part of the legend, well at least father did - son got too close to the sun and the glue holding the wires to the wings of his plane came unstuck, and he fell into the sea) - I always wondered what the point of James Joyce' Stephen Daedalus reference was).

Next we expand on these 6 problem areas.

1. Peer-to-peer: Problems and key issues - Structure!

Antony Rowstron, Microsoft Research Cambridge

In the last couple of years we have seen a large amount of research in to peer-to-peer systems, where I define a peer-to-peer system as one where all members/users/participants of the system have a symmetrical role in the system.

Fundamentally, all peer-to-peer systems can be described as either being structured or unstructured. Loosely a structured overlay places particular constraints on the other nodes in the system that it knows about (for example in the case of distributed hash tables (DHTs) this is based on their node identifiers). In general, an unstructured overlay (such as Gnutella) in general places no such restrictions.

One initial observation is that the understanding of the trade-offs between structured and unstructured overlays is not well studied, and I think many of the statements made in this space are little more than urban myth. People make vague comments about structured

overlays being less well suited to high churn rates when compared to unstructured overlays. The cost of maintaining structured overlays is higher than that of maintaining unstructured overlays. But where is the evidence?

My next observation is that peer-to-peer community largely overlooks the difficult and taxing question of security of the peer-to-peer systems. Given the environments that these systems are supposed to run in, clearly it is expected that some participants will run malicious nodes. It is interesting to observe that Gnutella appears relatively robust to malicious nodes (although I know of no real studies of this), and this is probably due to the high degree of redundancy achieved when flooding queries. However, many of the optimizations that are being proposed for gnutella appear to make it far more susceptible to malicious nodes – yet are rarely evaluated on this.

I think another interesting area which we are beginning to see some work on is the relationship between the overlay and underlay. What services can an underlay provide to make building overlays cheaper? What ideas used in overlays can be incorporated in underlays? It is important to tune the overlay to the underlay? At the moment also, overlay builds often ignore the services provided by underlays, for example how many application-level multicast systems are designed to exploit the islands of IP multicast?

I think also there are many challenges in the very trendy area of incentives. However, I have not seen many proposals that I find compelling. There seems to be several fairly difficult problems, such as auditing nodes to ensure that they are doing what they say they are doing. In many proposals it seems that the costs of doing often are too high to be realistic (and sometimes it is difficult to check - how do you know whether a node is routing your traffic? Or how much bandwidth it is contributing?). Systems that involve "reputations" seem open to many different attacks.

Then there are the issues related to applications - and understanding how to exploit p2p systems. I seem to read many papers recently extolling the virtue of using a p2p system (usually a DHT) and then using it in a way which makes little sense to me.

There are then finally the pragmatic issues: the requirement of better simulators with more realistic network topologies, able to scale to larger numbers of nodes and better understanding of workloads. There is the issue of how these systems run in the wild (and therefore validating the simulations) - and this needs large enough testbeds to be meaningful. Also, we lack tools to support the development of these applications; we could do with better debuggers and profiling tools.

So, before concluding, is there anything that we should not be doing? Well, personally, I feel that if you're going to design yet another DHT, it needs to have a really convincing motivation. Also, I'm continually being asked the question "why are all deployed p2p systems unstructured" (or, why are you doing research on structured ones?). The complexity of structured overlays means that they take longer to develop and understand, and we are now beginning to see several systems out there in the wild which use structured overlays. I think over the coming years unstructured overlays will come to dominate.

2. Recent Problems in Peer-to-Peer Content Search and Retrieval

Brian Neil Levine, University of Massachusetts Amherst

Peer-to-peer (P2P) protocols have been proposed for a wide variety of network services, but in practice are most commonly used for file sharing among Internet users. Such applications rank among the most popular sources of traffic on the Internet. Popular commercial P2P applications are used for sharing of audio and video content; they include Gnutella and Kazaa, and previously.

Approaches that have been taken for locating content are growing. Centralized approach is usually credited to the Napster protocol: one server maintains a centralized index for the resolution of search queries. In the distributed hash table (DHT) approach - used, for example by Chord, Pastry, and CAN among others - the index of available content and the task of resolving queries is distributed among peers. DHT-based protocols are similar to the Napster protocol in that queries are unicast (perhaps via several peers) to hosts that store indexes of content. Resolving a query means returning a pointer to the peer that is storing the actual content, and then the querying peer retrieves the desired file. In contrast, Gnutella nodes only answer queries when they locally store relevant content.

A different approach is possible as well, yet less studied. Peer-to-peer topologies can be organized into small world topologies based on autocorrelated content among neighbors. By autocorrelated, we mean that, neighbors in the graph are more likely to have similar content than non-neighbors. In this way, queries navigate directly and quickly toward content (not indexes) without flooding.

Like many problems in computer science, each content location approach has some advantage that is linked to the application scenario. For example, DHT protocols cannot be used to search for key word search of text documents stored by peers as there are too many keys to store per document. (Our own recent work has shown that the number of keys that must be stored for each peer is in fact fairly large even for mp3 music libraries.)

In many p2p systems, data is exactly replicated among the peers participating in the system. Replicated data, while providing scalability and fault-tolerance, introduces the problem of source selection. After determining the locations of a desired file, a client must decide where to download from in order to receive the file quickly. This problem has been studied mainly in the context of mirrored Web data, where it is called the server selection problem. Various solutions have been proposed and validated with experiments on the Internet. However, many of the existing techniques rely on assumptions that render them inapplicable in the dynamic setting of peers. For instance, selection strategies based on experience with specific hosts do not apply when hosts are not likely to be encountered more than once. In addition, selection strategies that rely on network-layer assistance are not feasible. Some new approaches that have been tried include those based on decision trees and MDPs. Techniques from distributed information retrieval can be employed when topics are replicated at peers, but not content (e.g., articles on a particular news topic). Our own research has shown that such techniques cannot choose peers with good network performance.

3. Rights Management v. Anonymity Jon Crowcroft, Cambridge

P2P is largely used for sharing files, typically music and increasingly now film and games. This has led to major social upheaval in the music industry. The technology for efficient content distribution based in an infrastructure already exists, although it is not used for

commercial streaming or distribution of copyright material much yet (e.g. Akamai/ Inktomi etc). P2P clearly also functions as a content distribution system, but is largely rooted in assumptions of strong anonymity and the ability this confers on users to carry out potential copyright theft. There is no special reason why we couldn't use the scaling properties of P2P but include the ability to pay (or otherwise recompense content creator or owner). However, there is very little work in this area yet.

4. Incentives for Cooperation in Anonymity Systems

D. Figueiredo, J. Shapiro, D. Towsley, Umass

Like many peer-to-peer applications, anonymous communication systems are vulnerable to freeriders, peers who use the system while providing little or no service to others. To complicate matters, the identity of the freeriders is obscured by the very service anonymity systems are designed to provide, which limits the efficacy of conventional approaches for promoting cooperation (e.g., reputation mechanisms). We argue that the design constraints imposed by anonymity systems lead naturally to the notion of a currency that can be exchanged for service in order to provide incentives for cooperation. Based on this idea, we describe an approach based on the use of anonymous digital cash payments between those who use the service and those that provide the service. We illustrate its application to a well-known peer-peer anonymous protocol (onion routing). We argue that it can be applied with only a slight additional overhead and at the same time preserving its architectural simplicity. We believe that a similar approach can be applied to other peer-peer anonymous protocols and, more generally, to any peer-peer application.

5. Deployment Issues of Unstructured P2P Networks

Brian Levine, Umass

We used an end-host multicast program we wrote to multicast live the recently concluded NetGames 2003 workshop. The largest problem we had was with the prevalence of firewalls and NATs. I'll briefly describe our end-host multicast topology construction algorithm and discuss how our algorithm was defeated by firewalls and NATs. We have some ideas on ways to address this problem, though none that are very satisfactory. The firewall and NAT problems will be faced by all p2p networks that tries to optimize their topology.

Longer term, I am interested in the authentication and authorization of p2p participants: How to prevent an impostor from participation or injection of bad data into the p2p network? I will explore the use of forward secure signature to ensure non-repudiation.

One potential use of a p2p network is to support multiplayer gaming. Here the paramount problem is latency, both latency on the physical network and join/leave latency on the p2p overlay. To reduce the number of interacting players, area of interest management has often been proposed as a solution. Area of interest management is usually implemented using multicast sessions, one per area of interest. As players move between areas of interest, the join and leave latency and scalability become an issue.

I think the above are the immediate three issues we need to solve for p2p networks.

6. Measurement of P2P Apps

Report on The COST (EU)-NSF (USA) workshop on Exchanges and Trends in Networking

44

Nextworking'03, June 23-25,2003, Chania, Crete, Greece

Brian Levine, UMass

Only a few measurement studies have been performed of real P2P file sharing applications, and even fewer empirical studies exist of how well DHT or other protocols support P2P file sharing in practice.

In our previous work, we took sample measurements of the Napster and Gnutella file systems, including node availability and shared file lists. From this data, we were able to infer the skewed popularity of shared and transferred files. Leibowitz, et al. have reported similar results by sniffing raw traffic seen on an Israeli ISP.

Markatos took measurements from three Gnutella clients at separate geographic locations for one hour and analyzed the effects of caching search queries. Due to the high temporal locality of queries observed, a simple query caching scheme was shown to reduce query traffic by as much as a factor of two. In our recent results, we apply query caching to Chord and show that it has less of an effect on balancing load.

Sripanidkulchai also claimed that simple caching schemes can help reduce the amount of query messages flooding the network. The study mentioned that the performance of a cache will ultimately be affected by the consistency of cached query results. This drawback also applies to cached queries in Chord.

We traced all users of a centralized p2p sharing network. In this recent work, we detailed many workload characteristics of this trace, including: the popularity trends of file transfers over time; the correlation between user library sizes and downloads; the skewed popularity and inter-arrival time distributions of queries; the skewed popularity of file keys; and the distribution of users downloading and serving files. Although several measurement studies of P2P networks have been performed in the past, this is one of the only studies to obtain a complete view of the usage characteristics of the system.

These measurement studies are important for several reasons. First, p2p applications are new and provide us with an opportunity to track its changing characteristics over time as connections to the home and office, and to mobile devices, change over time. Second, applying observed measurements to the evaluation of proposed protocols is an important method of evaluations. For example, in our recent work, we evaluated the effectiveness of the Chord protocol as a protocol for file sharing using measurements from a real P2P application.. We found that Chord does not succeed in distributing the index evenly among nodes. The power-law skew of keys from real queries and shared files results in a work load that is also skewed. We also consider the notion of caching in Chord to balance loads.

6.6 Session title: PEER-TO-PEER AND OVERLAY NETWORKS: Overlay networking and Content Distribution

Chair: Ernst Biersack

speakers: Ernst Biersack, Vishal Misra, Jorg Liebeherr, Pablo Rodriguez

E. W. Biersack: Issues in Peer to Peer Systems and Content Distribution

We have seen in the last few years a large amount of research in P2P systems, with most of the efforts concentrated on the design of distributed hash tables (DHTs). Other issues have received very little attention such as Hierarchical DHTs, Topology Aware DHTs, and Content Distribution via P2P Systems.

Hierarchical DHTs :

Inspired by hierarchical routing in the Internet, we propose hierarchical DHTs where peers are organized in disjoint groups [1]. Each group maintains its own overlay network and intra-group lookup service. A top-level overlay is defined among the groups. Within each group, a subset of peers are labeled as "superpeers". There are various issues related to Hierarchical DHTs that need to be investigated such as deployment and maintenance in face of peers joining and leaving the system.

Topology Aware DHTs :

Topological considerations are of paramount importance in the design of a P2P lookup service. TOPLUS [2] is a lookup service for structured peer-to-peer networks that is based on the hierarchical grouping of peers according to network IP prefixes. However, topology aware overlays pose a number of new problems such as (i) non-uniform population of ID space or (ii) correlated node failures.

Scalable Content Distribution using P2P systems:

There exist a large number of open and closed-loop schemes for distributing a multimedia stream simultaneously to a large number of users that all aggregate the client requests and transmitted single stream via multicast. The use of multicast transmission has also been advocated in the context of P2P systems and the native multicast distribution via IP multicast is replaced by an overlay multicast distribution using the nodes of the P2P system.

We question the use of multicast overlays for content distribution in P2P systems and argue that pull-based approaches involving parallel download [3] are more resilient to changes in node availability and the available network bandwidth resources. The use of parallel download for video streaming poses new challenges due to the real-time nature of the medium, that imposes time constraints and also constraints on the order in which data are downloaded.

Vishal Misra: Security and overlay networks

P2P systems pose a number of security vulnerabilities such as (i) routing attacks that misroute requests, (ii) storage and retrieval attacks that prevent the delivery of requested data, and (iii) other attacks such as DoS that overload the P2P system e.g. by flooding it with requests or large number of joins/leaves.

However, the overlay can also be used to prevent DoS attacks by shielding the target from the attacker. One such example is SOS [4] that utilizes an overlay network architecture to provide resilience to Denial of Service attacks. While SOS is a first step, more research is needed to explore issues such as spread spectrum overlay routing or techniques to detect malicious overlay nodes.

Jorg Liebeherr: Two challenges for building large self-organizing overlay networks:

Application-layer overlay networks, which organize sets of applications in virtual peer networks, have emerged as a new direction in networking research for deploying new network services. A fruitful area of research has been the search for solutions where a large number of peers can quickly (in the matter of seconds) self-organize in an overlay network without central control or coordination.

There are two issues related to overlay networks that deserve to be explored more deeply: (i) Economy-of-scale versus increased scalability and (ii) Programming Overlay Networks.

For illustration, examples are drawn from a currently ongoing project, called HyperCast, which builds an overlay network as a Delaunay triangulation graph [5] and which uses a construct, called overlay socket, which intends to simplify the task of overlay network programming.

Economy-of-scale versus increased scalability.

Overlay networks as logical graphs with a regular topology that ignore the underlying network infrastructure, self-organize quickly, even if the number of peers is very large. However, such overlay networks can be poorly matched to the network-layer infrastructure, yielding high latencies and poor utilization of resources. Each overlay network solution trades-off economy-of-scale for increased scalability. Research on overlay topologies has only started to explore the design space of this trade-off.

Programming Overlay Networks.

Research on overlay networks has focused on the design of protocols to maintain and forward data in an overlay network. However, less attention has been given to the software development process of building application programs in such an environment. Clearly, the complexity of overlay network protocols calls for suitable application programming interfaces (APIs) and abstractions that do not require detailed knowledge of the overlay protocol, and, thereby, simplify the task of the application programmer. A particular challenge is the development of programming paradigms that can evolve together with overlay network topologies.

Pablo Rodriguez: Open Issues in Content Distribution

Content distribution research was started more than five years ago. It first commenced with Web caching, cache cooperation and caching infrastructures. Then Akamai turned caching into a service for content providers, and suddenly, Content Distribution Networks became one of the most important advances in Internet technologies over the last years. However, with the dot com crash and the economic meltdown, many companies that had started ambitious Content Distribution projects put them aside or abandoned them completely. Similarly, research in Content Distribution went into a somehow stalled phase. The general idea that most interesting problems in content distribution were solved coupled with the decreasing industry interest in content distribution technologies, pushed many researchers into other areas.

CDNs form an important part of today's Internet and are used to offload the most popular Web sites on the Web, to stream live events to millions of simultaneous receivers, and to take care of flash crowd events such as the Olympics, or the World Cup. Not only are CDNs used to offload the public Internet, they are also used inside enterprises, financial corporations or large retailer stores to enable remote learning and training of their employees through out the world. Today, Content Distribution Networks should take a step forward to accommodate new technologies, augment their efficiency, and lower their deployment and maintenance costs so they can be used in a broader range of applications and environments.

CDNs emerged from the need for scalability and better end-user experience in the Web, the most popular application at the time. As such, most of the research focused in the delivery of static Web content. There are a number of challenges for CDNs.

Support of new applications Delivery of dynamic content delivery is still a major challenge not yet fully resolved. Being able to replicate and distribute the intelligence behind most e-commerce sites into the network edge while providing scalability, consistency, and reconciliation is a problem that is yet to find a good solution. Given that dynamic and personalized content accounts form a considerable percentage of the total content, if CDNs were to provide an efficient solution for dynamic content delivery, their impact would be much bigger than nowadays.

The fact that many homes are rapidly being upgraded with high-speed data connections is fostering the appearance of new interactive applications. Up to now, CDNs have mostly focused on the delivery on content from the server to the users. However, little attention has been paid to systems where information also travels in the opposite direction, i.e. from the end-user into the CDN. Such systems include distributed gaming platforms, or remote storage systems where users upload content into the network from their multimedia devices (e.g. cameras, recorders). Providing a scalable architecture that enables efficient content upload and a real-time architecture will encourage more interactive applications to use CDNs.

Support of mobile users:

The mobile Internet is becoming more and more of a reality. The deployment of GPRS, CDMA, UMTS, and WLANs is making ubiquitous mobile data access available all over the world. These new wireless networks create many opportunities for intelligent wireless content distribution architectures:

There is a real need to transform this centralized architecture into an intelligent distributed architecture where content is pushed closer to the mobile users, e.g. all the way to the base stations.

Wireless content distribution techniques need to be revisited to improve wireless user's experience, taking into account the link properties of next generation wireless networks (e.g. high burstiness, long delays, frequent disconnections, large channel acquisition times, etc) [6].

Wireless links, unlike terrestrial links, can not be easily upgraded or over-provisioned. Providing differential QoS over the wireless bearer, both for users and applications, is one of the biggest challenges faced by wireless mobile networks. Content distribution, on the other hand, can be used to provide better QoS decisions. The fact that proxies can see all wireless traffic and can identify the end-users as well as the applications, allows CDNs to implement much better QoS strategies, especially when embedded in the wireless architecture.

In addition to QoS, another technique that can improve the efficiency of Wireless links is Wireless multicasting. As opposed to terrestrial multicasting, Wireless multicasting needs to deal with a complete set of constraints in terms of losses, mobility, delay, connectivity, and throughput. Therefore, new content distribution techniques should be devised to relieve the scarce resources of wireless networks.

Cost of CDN deployment and operation

One important reason why CDNs have not been massively deployed by ISPs and enterprises is the high cost of deploying and maintaining a CDN. CDNs require large human resources to deploy, control, manage, and maintain the CDN architecture. Administration and control of these networks usually happens in a very centralized location and requires 24 hours a day, 7 days a week monitoring and manual intervention in many cases. One of the big challenges to facilitate the deployment of CDNs is to distribute the control and administration of the CDN so the maintenance and roll-out costs drop drastically. As a result, advances in self-healing, self-configuring, autonomic computing, and similar techniques are essential for the success of CDNs.

REFERENCES

- [1] L. Garces-Erice, E. W. Biersack, K. W. Ross, P. A. Felber, and G. Urvoy-Keller, "Hierarchical P2P Systems", Proceedings of ACM/IFIP International Conference on Parallel and Distributed Computing (Euro-Par), Klagenfurt, Austria, 2003.
- [2] L. Garces-Erice, K. W. Ross, E. W. Biersack, P. A. Felber, and G. Urvoy-Keller, "Topology-Centric Look-Up Service", Proceedings of COST264/ACM Fifth International Workshop on Networked Group Communications (NGC), Munich, Germany, 2003.
- [3] P. Rodriguez and E. W. Biersack, "Dynamic Parallel-Access to Replicated Content in the Internet", IEEE/ACM Transactions on Networking, 10(4):455--464, August 2002.
- [4] A. Keromytis, V. Misra, and D. Rubenstein, "SOS: Secure Overlay Services", Proc. ACM SIGCOMM'02, (Pittsburgh, PA), August 2002.
- [5] J. Liebeherr, M. Nahas, and S. Weisheng, "Application-layer multicasting with Delaunay triangulation overlays", IEEE Journal on Selected Areas in Communications, 20(8):1472--1488, October 2003.
- [6] P. Rodriguez and V. Fridman, "Performance of PEPs in Cellular Wireless Networks", Proc 8th Int. Workshop on Web Content Caching and Distribution, September 2003.

6.7 Session title: NETWORK TOMOGRAPHY AND TRAFFIC MODELING MEASUREMENTS AND CONTROL: Network Tomography

Chair: Don Towsley

Speakers: Don Towsley, Constantinos Dovrolis, Francesco Lo Presti, Kave Salamatian

Network Tomography:

Network tomography refers to a class of measurement-based inference problems that include the following:

- Determining internal network characteristics such as: available bandwidths, link loss rates, link delay statistics, and network topology, based solely on the use of end-to-end path measurements.
- Determining the source destination traffic matrix based solely on traffic counts taken at some or all of the links within the network.

The first class of problems further breaks down according to whether the goal is to determine the path characteristics, such as available bandwidth or detailed link-level characteristics such as link loss rates.

The session on network tomography focused primarily on the first class of problems. Overviews of progress made in inferring path characteristics and link-level characteristics. Techniques exist now for characterizing the available path bandwidth, link-level loss and delay characteristics, and distribution tree topology. However, a number of challenges were raised that require resolution before these techniques can be placed into the mainstream. We list some of these below

Path characterization: Much progress has been made in developing estimators for the available bandwidth on an end to end path. However, there are some issues/questions still to be resolved. Foremost, is the question: exactly what does “available bandwidth” mean? What is the time scale associated with? Is it stable or stationary? Once obtained, is it possible to predict how it will change in the future?

Link-level characterization: Considerable progress has been made in developing a theoretical framework within which to develop multicast-based estimators of link-level metrics such as loss rates and delay statistics. Further progress has been made in extending these techniques to the case where only unicast measurements are possible. This has produced techniques based on packet pairs and stripes that either assume perfect correlation over common segments or infer this correlation as part of the measurement process. In spite of this considerable progress, many questions/issues need to be resolved. First and foremost is scalability. The proposed techniques have only been evaluated for very small networks consisting of 10 – 100 nodes. Revolutionary developments are needed to cover networks of thousands to hundreds of thousands of nodes. Second, although it is possible to identify the topology of a path-based tree spanning a collection of nodes, there still remains the difficult question of merging several such tree topologies together into one single consistent topology.

In addition, the two problem areas share several common challenges:

- integrating the developed techniques into applications,
- extending the algorithms to allow the use of internal measurements along with end-to-end measurements, and
- extension to new technologies such as wireless and optical.

Another interesting question regards how to extend the tomographic methods to the following class of problems: given detailed packet-level observations at a point within a network, what can be said about the structure/performance of segments of the network not containing the observation point through which the observed packets have or will traverse?

Last, most of the work in this area is based either purely on descriptive models (e.g., hidden Markov models) or constructive models (e.g., G/G/1/B queuing model). This presents a challenge of developing a framework within which to combine elements from descriptive and constructive models as needed. An even more fundamental challenge is to develop a measurement-based modeling theory that can be used to guide and develop parsimonious models for the purpose of inferring network or application performance metrics based on measurements.

6.8 Session title: NETWORK TOMOGRAPHY AND TRAFFIC MEASUREMENTS AND CONTROL: Network Measurements

Chair: Christophe Diot

Speakers: Christophe Diot, Balachander Krishnamurthy, Anja Feldmann

Measurement is a somewhat recent and emerging area in which our community needs to build expertise and experience. It is however commonly accepted that monitoring will become soon a requirement on most networks (for applications ranging from network management to QoS control).

The number of problems to investigate through measurement is huge and well identified by the community. Problems range from infrastructure equipment design to measurement techniques through observation of data and design of traffic models.

On the infrastructure side, monitoring links at 10Gbps is not simple. Managing a distributed infrastructure of hundreds of monitoring stations that monitor different metrics on various ASes is also not straight forward. The design of a robust monitoring infrastructure, ultimately usable for management and operational purpose, is a real challenge.

Monitoring also has a cost preventing systematic deployment of monitoring equipment. Therefore, techniques such as traffic sampling and property inference will be required. Very little has been done yet in these areas.

Last, once data are collected it is important to understand the significance of the data in order not to draw any inexact or unduly generalized conclusion from those data. Understanding the data we collect will allow the community to make significant progress in the fundamental understanding of network properties. It will ultimately make it possible to model the traffic, routing, and complex network themselves.

Several factors slow down the progress made thanks to research driven by measurement:

Data are the property of network operators and service providers. For obvious privacy issues, these data are not made available to the academic community. Progress would be much faster

if data would be available. We believe that the research community would benefit from building strong relationship with Network operators. In addition, the source of network traffic data is multiple and some of them are easy to obtain. The research community should not focus though on commercial ISP traces. Any traffic is useful to study for what it represents.

Our community does not have a strong history or sharing traces and software reuse. However, this would be extremely useful in understanding the significance of data. Reproduce an experiment on a different set of data should receive better considerations from Program Committee. Also, publishing a negative result should not be banned, as well as proving a previous work was wrong or inexact. Initiatives such as Planetlab are making a step in this direction but a lot more needs to be achieved

Last, governmental institutions could help making traces or tools available to the research community. We believe in the value of COST-IST and NSF as imprimateurs. Governmental institutions could for example encourage network operators to share their data, researchers to share their tools and their data, etc. These institutions could also be an example by providing their own data to the research community.

To conclude, let us notice that measurement was introduced in the Internet after most damages were made. We consequently believe that monitoring and measurement should be included in all networking research domains such as sensor nets, meshed networks, etc.

6.9 Session title: *LARGE SCALE NETWORKS AND FUTURE NETWORK ARCHITECTURES*

Chair: Jim Kurose

Speakers: Jim Kurose, Michel Diaz, Serge Fdida, Gunnar Karlsson, Israel Cidon,
 Lixin Gao

Several themes emerged from the presentations and discussion in the session on “Large Scale Networks and Future Network Architectures”:

A Focus on Control. Several of the presentations addressed issues in the control plane, where attention to the issues of robustness, complexity of control, adaptability, modes of failures, safety, predictability, manageability, evolvability, and security are important. Indeed, several speakers argued that the most important problems were in the control plane, rather than the data plane. One thread woven through several of the talks was that of complexity – how to best handle the management and control complexities that come from rapid growth in system size, technological heterogeneity, service offerings, and administrative demands. Gunnar Karlsson argued that systems should be designed with simplicity in mind, and described an Internet service architecture that allows scalable implementation, a clear basis for charging, and easy management. The notion of “keep it simple” was echoed in Israel Cidon’s presentation of Keep-It-Simple Signaling (KISS), a highly optimized, hardware-assisted signaling approach for unicast connections. Cidon argued that the introduction of hardware-assisted signaling and QoS routing functions may be the missing link to achieve full QoS support and the ultimate converged IP network. Jim Kurose’s presentation also focused on control, arguing that many of the basic characteristics of network control are still poorly understood, and that the development of such an understanding should be an important future

research goal for the community. He cited, as examples, the fact that the differences between hard-state and soft-state signaling, and the interaction between control operating at different layers (e.g., overlay and network-layer routing) are two fundamentally important questions that are still poorly understood.

Control: developing fundamental insights at the network layer. The presentations by Serge Fdida and Lixin Gao also emphasized control and the need to develop a fundamental understanding of networking principles, but focused their discussion to the network layer. Fdida argued that indirection should play a central role in routing for self-organizing networks. With such indirection, a node's identifier and its address can be decoupled, allowing services (such as mobility, P2P, and content distribution) to be flexibly implemented in an address-independent manner. Gao's talk argued that by better understanding variability of control plane traffic (including the interaction between control/data plane traffic), we can better detect faulty management (e.g., BGP) configurations, enhance performance (e.g., by separating control and data plane interactions), as well as detect attacks against the infrastructure.

QoS and the need for end-end protocols. Michel Diaz's talk emphasized the need to support new multimedia-based distributed applications. He argued for the need to develop new end-to-end transport-layer protocols (whose various services are currently being implemented primarily in the application-layer in a per-application fashion). These protocols will need to handle the varying reliability, ordering, and timing semantics of different multimedia applications. SCTP and DCCP were cited as two ongoing efforts that are seeking to build such new end-end transport-layer services.

In addition to noting the topics that were discussed in detail in this session, it is interesting to note the aspects of network architecture that were not discussed in detail. For example, there was no discussion of optical networking and circuit switching, and surprisingly little explicit discussion of security and network management capabilities. Several people noted that network management was (at least implicitly) in almost all of the talks. The relatively small amount of discussion about network security was probably mostly a function of participants' research interests, since there was general agreement that security is an extremely important topic, and must be a key aspect of network architecture.

6.10 Session title: PROGRAMMING THE INTERNET

Chair: Ken Calvert

Speakers: Ken Calvert, Ian Wakeman, Christian Tschudin, Michael Smirnov, Hao Che

This session focused on issues related to extending the shared infrastructure of the network to provide user control through some form of programmability. Speakers included: Ian Wakeman, University of Sussex ("Risky Business: Developing an Economic Infrastructure for Third Party Computation"); Christian Tschudin, University of Basel ("Analog Programs: Mobile Code for Fibers and Ethers"); Michael Smirnov, Fraunhofer Fokus ("Rule-based Systems Programmability"); Hao Che, University of Texas at Arlington ("Network Processor Programming"); and Ken Calvert, University of Kentucky ("Questions About Programming the Internet").

Take-away points:

Generally there are two paradigms for research on network programmability: one focused on understanding/improving the existing Internet artifact, the other looking beyond, changing assumptions and considering different environments. (The coexistence and contrast between these two paradigms showed up in other areas of the Workshop also.) The former paradigm focuses on extending the functionality of the "waist of the [current] hourglass" in a backward-compatible way. The latter focuses on replacing it with a different—possibly radically so—architecture.

Motivations for network programmability include speeding evolution of the network service, and improving the scalability of group applications through network assistance.

Another issue this area has in common with other areas of research is the problem of managing and scaling the trust relationships that are necessary, not only for network programmability but also for ad hoc networks and end-system-based overlay services. Creation, management, and enforcement of policies was a related theme that emerged repeatedly over the course of the workshop. Reputation-based systems look promising for providing strategy-proof incentive structures, but there is more work to be done. In the absence of such systems, approaches to programmability intended for deployment in the present network must take into account the "ugly" issues of billing, charging structures, etc.

In the context of the current Internet, issues of trust and incentive suggest a two-level structure for user-controllable services: Ultra-lightweight services (i.e. services that can be implemented at wire speed and have resource requirements similar to IP forwarding) are available on a per-packet basis, while more heavyweight services are set up to apply to particular users' packets after bilateral signaling/negotiation.

In addition, top-end routers are becoming more and more decentralized and continue to grow in terms of total throughput and interfaces supported. We need better system-level understanding of how to map highly-pipelined forwarding and routing computations onto what are essentially high-performance distributed systems comprising large numbers of channel interface modules, each with a hierarchy of processing elements with different levels of sophistication and throughput characteristics.

Another programming approach of interest is rule-based systems, in which behavior is defined by openly-exposed, possibly-dynamic sets of rules. The rule-action paradigm is already in use in firewalls and other policy-intensive domains. Thus, approaches for creating, composing, reasoning about and auditing rule-based services could have broad applicability across a number of areas including security and routing.

A more radical approach is to view computation as a side-effect of communication -- for example, as a result of interference among multiple transmissions in free space or across common channels (cf. recent work on "network coding"). Such an approach mandates fundamental research into analog models of computation suitable for realization in optical or wireless channels, as well as new ways of thinking about control and data separation.

Summary of Open Issues:

- Trust and incentive structures: how can users and providers' interests be aligned so that network (or overlay, or service gateway, or ...) programmability adds value for both?

- Models of programmability: a broad spectrum of areas, including lightweight fast-path per-packet programmability, rule-based systems, genetic approaches to protocol program development, system-level programming and management techniques for network nodes implemented as collections of nodes with hierarchies of computation capabilities.
- Further applications of programmability.

7 Post-Workshop Questionnaire Results

A questionnaire was sent out to the workshop speakers one year after the workshop (June 2004) referring to the various research challenges in networking as identified based on the *NeXtworking'03* workshop presentations, abstracts, discussions and session chairs summaries. The material presented in Section 5 of this report, was sent out as part of the questionnaire and the speakers were asked to grade the various research challenges as listed there, as follows.

Each issue is to be characterized by the metrics: *importance* and *problematic*. **Importance** refers to the level of significance of the issue and **Problematic** refers to the level of difficulty of reaching a reasonable solution to the problem.

The grades to be used are:

- 1 → lowest
- 2 → lower than average
- 3 → average
- 4 → above average
- 5 → highest

At the time of the submission of this report, these questionnaires are being collected. They will then be processed and the results will be reported at the workshop site, together with the final edition of the present report (www.di.uoa.gr/~NeXtworking).