



ITC 2011 - 23rd International Teletraffic Congress

September 6-9, 2011, San Francisco, USA

Technical Program

Tuesday, September 6

08:00 - 18:00 *Registration Table Open*
 08:00 - 09:00 *Light continental breakfast*
 08:45 - 09:00 O: *Opening Remarks*
 09:00 - 10:00 K-1: *Keynote Presentation: Software Defined Networks*
 10:00 - 10:30 AM Break
 10:30 - 12:30 TS-1: *Routing*
 12:30 - 13:30 *Lunch*
 13:30 - 15:30 TS-2: *Wireless Networks I*
 15:30 - 16:00 PM Break
 16:00 - 17:30 TS-3: *Traffic Classification and QoE*
 18:00 - 19:30 R: *Reception*

Wednesday, September 7

08:00 - 17:00 *Registration Table Open*
 08:00 - 09:00 *Light continental breakfast*
 09:00 - 10:00 K-2: *Keynote Presentation: Scheduling in Networks*
 10:00 - 10:30 AM Break
 10:30 - 12:30 TS-4: *Content Delivery*
 12:30 - 13:30 *Lunch*
 13:30 - 15:00 TS-5: *Optical Networking*
 15:00 - 16:00 S: *Student Poster Session/PM Break*
 16:00 - 17:30 TS-6: *Resource Management*
 18:00 - 22:00 D: *Social Event and Dinner*

Thursday, September 8

08:00 - 17:00 *Registration Table Open*
 08:00 - 09:00 *Light continental breakfast*
 09:00 - 10:00 K-3: *Keynote Presentation: Research Challenges for Modern Data-Center Networks*
 10:00 - 10:30 AM Break
 10:30 - 12:30 TS-7: *Wireless Networks II*
 12:30 - 13:30 *Lunch*
 13:30 - 15:40 TS-8: *Overlay, Addressing, and Control*
 15:40 - 16:00 PM Break
 16:00 - 17:00 P: *Panel: Economics and Pricing of Mobile Data Services*
 17:00 - 17:30 C: *Closing Session (Best Paper Award)*

Friday, September 9

08:00 - 11:00 *Registration Table Open*
 08:00 - 09:00 *Light continental breakfast*
 09:00 - 12:30 TM-1: *Tutorial: Network Survivability Modeling and Quantification,*
 TM-2: *Tutorial: Building an End-End Nationwide IPTV Service,*
 W1-M: *Workshop: Cnet*
 W2-M: *Workshop: DC-CaVeS*

10:30 - 11:00	AM Break
12:30 - 13:30	<i>Boxed Lunch</i>
13:30 - 17:00	TA-3: Tutorial: Cooperative and Green Heterogeneous Wireless Networks , TA-4: Tutorial: Traffic Models for Quality of Experience Assessment , W1-A: Workshop: Cnet W2-A: Workshop: DC-CaVeS
15:00 - 15:30	PM Break

Tuesday, September 6

Room: [Grand Ballroom](#) (for all sessions)

8:45 AM - 9:00 AM

O: Opening Remarks 

9:00 AM - 10:00 AM

K-1: Keynote Presentation: Software Defined Networks 

Nick McKeown (Stanford University, USA) (Bio)

In the past two decades, enormous innovation has taken place on top of the Internet architecture. Email, e-commerce, search, social networks, cloud computing, and the web as we know it are all good examples. While networking technologies have also evolved in this time, we believe that more rapid innovation is about to happen. We are about to witness a revolution in the networking towards so-called "Software Defined Networking" (SDN). SDN enables innovation in all kinds of networks - including data centers, wide area telecommunication networks, wireless networks, enterprises and in homes - through relatively simple software changes. SDN thus gives owners and operators of networks better control over their networks, allowing them to optimize network behavior to best serve their and their customers' needs. For instance, in data centers SDN can be used to reduce energy usage by allowing some routers to be powered down during off-peak periods.

The SDN approach arose out of a six-year research collaboration between our teams at Stanford University and the University of California at Berkeley. Essential to SDN are two basic abstractions: a general abstraction of packet forwarding, and a global view of the network upon which control and management tools can be built. Almost 50 of the largest network owners and equipment vendors have come together to create the Open Networking Foundation to standardize interfaces to implement the SDN abstractions, starting with OpenFlow (a software interface to packet forwarding). ONF leads the ongoing development of the OpenFlow standard (www.openflow.org).

I will describe the background leading to SDN, and the current status of SDN technology, deployments and standardization around the world.

10:30 AM - 12:30 PM

TS-1: Routing 

Chair: Kohei Shiimoto (NTT, Japan)

Modeling and Performance Evaluation of an OpenFlow Architecture

Michael Jarschel (University of Wuerzburg, Germany); Simon Oechsner (University of Wuerzburg, Germany); Daniel Schlosser (University of Wuerzburg, Germany); Rastin Pries (University of Wuerzburg, Germany); Sebastian Goll (University of Wuerzburg, Germany); Phuoc Tran-Gia (University of Wuerzburg, Germany)

The OpenFlow concept of flow-based forwarding and separation of the control plane from the data plane provides a new flexibility in network innovation. While initially used solely in the research domain, OpenFlow is now finding its way into commercial applications. However, this creates new challenges, as questions of OpenFlow scalability and performance have not yet been answered. This paper is a first step towards that goal. Based on measurements of switching times of current OpenFlow hardware, we derive a basic model for the forwarding speed and blocking probability of an OpenFlow switch combined with an OpenFlow controller and validate it using a simulation. This model can be used to estimate the packet sojourn time and the probability of lost packets in such a system and can give quick hints to developers and researchers on questions how an OpenFlow architecture will perform given certain parameters.

OpenFlow MPLS and the Open Source Label Switched Router

James Kempf (Ericsson Research, USA); Scott Whyte (Google, USA); Jonathan Ellithorpe (Stanford University, USA); Peyman Kazemian (Stanford University, USA); Mart Haitjema

(Washington University in St. Louis, USA); Neda Beheshti (Ericsson Research, USA); Stephen Stuart (Google, USA); Howard Green (Ericsson Research, USA)

Multiprotocol Label Switching (MPLS) is a protocol widely used in commercial operator networks to forward packets by matching link-specific labels in the packet header to outgoing links rather than through standard IP longest prefix matching. However, in existing networks, MPLS is implemented by full IP routers, since the MPLS control plane protocols such as LDP utilize IP routing to set up the label switched paths, even though the MPLS data plane does not require IP routing. OpenFlow 1.0 is an interface for controlling a routing or switching box by inserting flow specifications into the box's flow table. While OpenFlow 1.0 does not support MPLS, MPLS label-based forwarding seems conceptually a good match with OpenFlow's flow-based routing paradigm. In this paper we describe the design and implementation of an experimental extension of OpenFlow 1.0 to support MPLS. The extension allows an OpenFlow switch without IP routing capability to forward MPLS on the data plane. We also describe the implementation of a prototype open source MPLS label switched router, based on the NetFPGA hardware platform, utilizing OpenFlow MPLS. The prototype is capable of forwarding data plane packets at line speed without IP forwarding, though IP forwarding is still used on the control plane. We provide some performance measurements comparing the prototype to software routers. The measurements indicate that the prototype is an appropriate tool for achieving line speed forwarding in testbeds and other experimental networks where flexibility is a key attribute, as a substitute for software routers.

3G Meets the Internet: Understanding the Performance of Hierarchical Routing in 3G Networks

Wei Dong (University of Texas at Austin, USA); Zihui Ge (AT&T Labs - Research, USA); Seungjoon Lee (AT&T Labs Research, USA)

The volume of Internet traffic over 3G wireless networks is sharply rising. In contrast to many Internet services utilizing replicated resources, such as Content Distribution Networks (CDN), the current 3G standard architecture employs hierarchical routing, where all user data traffic goes through a small number of aggregation points using logical tunnels. In this paper, we study the performance implications of the interplay when 3G users access Internet services. We first identify a number of scenarios in which 3G users' service performance can be affected under hierarchical routing in comparison to an idealized flat routing. We then quantify this service impact by analyzing trace data obtained from a large-scale 3G network and a CDN provider. We find that the performance difference between hierarchical routing and flat routing increases when a 3G user accesses highly replicated service, and can further aggravates when the DNS caching is not properly managed under vertical handoff. For example, in our data analysis, the detour under hierarchical routing can cause a packet to travel extra distance by up to 1627km on the average case, which can lead to around 45.4% increase in round-trip latency. We also perform a measurement study to demonstrate that user mobility and web applications can lead to unexpected performance-impacting interactions, which can degrade the download throughput by up to an order of magnitude (0.9Mbps vs. 10.8Mbps).

Traffic Engineering for Multiple Spanning Tree Protocol in Large Data Centers

Trong Viet Ho (Université Catholique de Louvain, Belgium); Yves Deville (Université Catholique de Louvain, Belgium); Olivier Bonaventure (Université catholique de Louvain, Belgium); Pierre Francois (Université catholique de Louvain, Belgium)

The size of the capacity of data centers have been growing significantly during the last years. Most data centers rely on switched Ethernet networks. A drawback of the Ethernet technology is that it relies on the spanning tree protocol (or variants of it) to select the links that are used to forward packets inside the data center. In this paper we propose a Constrained-Based Local Search optimization scheme that is able to efficiently compute the optimum spanning tree in large data center networks. Our technique exploits the division of the data center network in VLANs. We evaluate its performance based on traffic matrices collected in data center networks and show good improvements compared to the standard spanning tree protocol with up to 16 VLANs.

A Study on Layer Correlation Effects through a Multilayer Network Optimization Problem

Iyad Katib (University of Missouri-Kansas City, USA); Deep Medhi (University of Missouri-Kansas City, USA)

Multilayer network design has received significant attention over the years. Despite this, the explicit modeling of three-layer networks such as IP/MPLS-over-OTN-over-DWDM in which the OTN layer is specifically considered has not been addressed before. In this paper, we present an optimization model for network planning of such multilayer networks that consider the OTN layer as a distinct layer with its unique technological sublayer constraints. More importantly, we present a comprehensive study to quantify the interrelationship between layers through change in unit cost of elements and capacity modularity, coupled with network demand. Focusing on the interrelation between the IP/MPLS and OTN layers, we provide a detailed numeric study that considers various cost parameter values of each layer in the network and analyze their impacts on individual layers and overall network cost.

1:30 PM - 3:30 PM

TS-2: Wireless Networks I

Chair: Luigi Fratta (Politecnico di Milano, Italy)

Backlog-Based Random Access in Wireless Networks: Fluid Limits and Delay Issues

Niek Bouman (Eindhoven University of Technology, The Netherlands); Sem Borst (Eindhoven University of Technology, The Netherlands); Johan van Leeuwen (Eindhoven University of Technology, The Netherlands); Alexandre Proutiere (KTH, Sweden)

We explore the spatio-temporal congestion dynamics of wireless networks with backlog-based random-access mechanisms. While relatively simple and inherently distributed in nature, suitably designed backlog-based access schemes provide the striking capability to match the optimal throughput performance of centralized scheduling algorithms in a wide range of scenarios. In the present paper, we show that the specific activity rate functions for which maximum stability has been established, may however yield excessive queue lengths and delays. The results reveal that more aggressive and persistent access schemes can improve the delay performance, while retaining the maximum stability guarantees in a rich set of scenarios. In order to gain qualitative insights and examine stability properties we will investigate fluid limits where the system dynamics are scaled in space and time. As it turns out, several distinct types of fluid limits can arise, exhibiting various degrees of randomness, depending on the structure of the network, in conjunction with the functional form of the activity rates. We further demonstrate that, counter to intuition, additional interference may improve the delay performance in certain cases. Extensive simulation experiments are conducted to illustrate and validate the analytical findings.

Bounds on QoS-Constrained Energy Savings in Cellular Access Networks with Sleep Modes

Balaji Rengarajan (Institute IMDEA Networks, Spain); Gianluca A. Rizzo (Institute IMDEA Networks, Spain); Marco G Ajmone Marsan (Politecnico di Torino, Italy)

Sleep modes are emerging as a promising technique for energy-efficient networking: by adequately putting to sleep and waking up network resources according to traffic demands, a proportionality between energy consumption and network utilization can be approached, with important reductions in energy consumption. Previous studies have investigated and evaluated sleep modes for wireless access networks, computing variable percentages of energy savings. In this paper we characterize the maximum energy saving that can be achieved in a cellular wireless access network under a given performance constraint. In particular, our approach allows the derivation of realistic estimates of the energy-optimal density of base stations corresponding to a given user density, under a fixed performance constraint. Our results allow different proposals to be measured against the maximum theoretically achievable improvement. We show, through numerical evaluation and simulation, the possible energy savings in today's networks, and we further demonstrate that even with the development of highly energy-efficient hardware, a holistic approach incorporating system level techniques is essential to achieving maximum energy efficiency.

Mitigating Signaling Overhead from Multi-Mode Mobile Terminals

Indra Widjaja (Bell Labs, Alcatel-Lucent, USA); Carl Nuzman (Bell Labs, Alcatel-Lucent, USA)

Modern cellular networks may be deployed using multiple radio-access-network technologies with multi-mode mobile terminals capable of selecting among different technologies. Such a network typically forms an overlay-underlay architecture where the underlay uses an older technology (e.g., 3G) while the overlay uses a newer technology (e.g., 4G). It has been observed that excessive signaling message updates can arise due to registration ping-pongs. Idle-mode Signaling Reduction (ISR) has been introduced as a mechanism to reduce update load. In this paper, we show that while ISR reduces update load, it also has the effect of increasing paging load. We investigate tradeoff between update and paging loads. Our analysis quantifies a threshold that is used to activate or deactivate ISR for each mobile terminal and results in significant signaling load reduction. We also provide a practical threshold-based algorithm that does not rely on knowledge of the structure of an overlay or terminal mobility.

Scheduling and Capacity Estimation in LTE

Olav Norvald R sterb  (Telenor Corporate Development, Norway)

Due to the variation of radio condition in LTE the obtainable bitrate for active users will vary. The two most important factors for the radio conditions are fading and pathloss. By including both fast fading and shadowing and attenuation due to distance we have developed analytical models to investigate obtainable bitrates for the basic resource unit in LTE. In addition we estimate the total cell throughput/capacity by taking scheduling into account. Particularly, the cell throughput is investigated for three types of scheduling algorithms; Round Robin, Proportional Fair and Max SINR where also fairness among users is part of the analysis. In addition models for cell throughput/capacity for a mixture of Guaranteed Bit Rate (GBR) and Non-GBR greedy users are derived. Numerical examples show that multi-user gain is large for the Max-SINR algorithm, but also the Proportional Fair algorithm gives relative large gain relative to plain Round Robin. The Max-SINR has the weakness that it is highly unfair when it comes to capacity distribution among users. Further, the model visualize that use of GBR with high rates may cause problems in LTE due to the high demand for radio resources for users with low SINR, i.e. at cell edge. Hence, for GBR sources the allowed guaranteed rate should be limited.

LTE Virtualization: from Theoretical Gain to Practical Solution

Liang Zhao (University of Bremen, Germany); Ming Li (Hamburg University of Technology, Germany); Yasir Zaki (University of Bremen, Germany); Andreas Timm-Giel (Hamburg University of Technology, Germany); Carmelita Goerg (University of Bremen, Germany)

Virtualization of wireless networks has received more and more research attention recently. As a means to reduce the investment of mobile network operators and to improve the network performance, LTE virtualization is one of the most important study cases in the scope of wireless virtualization. To investigate the advantages of virtualization on the LTE air interface, this work starts with an analytical model of FTP transmissions in virtualized LTE systems, and an obvious multiplexing gain can be observed. First evaluations with the considerations of realistic simulation models and mixed services of FTP and VoIP traffic also validate the analytical analysis. For practical applications, an extended multi-party spectrum sharing model is proposed. As one of the key issues, spectrum budget estimation is further analyzed based on the characteristics of real time services, i.e.~VoIP, and it is validated by further simulations that the proposed mechanism can give a very close estimation on the instant spectrum requirements in one cell. Some of the other investigation possibilities in practice are also discussed at the end.

4:00 PM - 5:30 PM

TS-3: Traffic Classification and QoE

Chair: Ravi Mazumdar (University of Waterloo, Canada)

A Multi-task Adaptive Monitoring System Combining Different Sampling Primitives

Imed Lassoued (INRIA Sophia Antipolis, France); Chadi Barakat (INRIA Sophia Antipolis, France)

Traffic measurement and analysis are crucial management activities for network operators. With the increase in traffic volume, operators resort to sampling primitives to reduce the measurement load. Unfortunately, existing systems use sampling primitives separately and configure them statically to realize some performance objective. It becomes important to design a new system that combines different existing sampling primitives together to support a large spectrum of monitoring tasks while providing the best possible accuracy by spatially correlating measurements and adapting the configuration to traffic variability. In this paper, and to prove the interest of the joint approach, we introduce an adaptive system that combines two sampling primitives, packet sampling and flow sampling, and that is able to satisfy multiple monitoring tasks. Our system consists of two main functions: (i) a global estimator that investigates measurements of the different sampling primitives in order to deal with multiple monitoring tasks and to construct a more reliable global estimator while providing visibility over the entire network; (ii) an optimization method based on overhead prediction that allows to reconfigure monitors according to accuracy requirements and monitoring constraints. We present an exhaustive experimental methodology with different monitoring tasks in order to assess the performance of our system. Our experimentations are done on our MonLab platform that we developed for the purpose of this research.

MINETRAC: Mining Flows for Unsupervised Analysis & Semi-Supervised Classification

Pedro Casas (LAAS-CNRS, France); Johan Mazel (LAAS-CNRS, France); Philippe Owezarski (LAAS, France)

Driven by the well-known limitations of port-based and payload-based analysis techniques, the use of Machine Learning for Internet traffic analysis and classification has become a fertile research area during the past half-decade. In this paper we introduce MINETRAC, a combination of unsupervised and semi-supervised machine learning techniques capable of identifying and classifying different classes of IP flows sharing similar characteristics. The unsupervised analysis is accomplished by means of robust clustering techniques, using Sub-Space Clustering, Evidence Accumulation, and Hierarchical Clustering algorithms to explore inter-flows structure. MINETRAC permits to identify natural groupings of traffic flows, combining the evidence of data structure provided by different partitions of the same set of traffic flows. Automatic classification is performed by means of semi-supervised learning, using only a small fraction of ground-truth flows to map the identified clusters into their associated most-probable originating network service or application. We evaluate the performance of MINETRAC using real traffic traces, additionally comparing its performance against previously proposed clustering-based flow analysis methods and supervised/semi-supervised classification approaches.

Traffic Causality Graphs: Profiling Network Applications through Temporal and Spatial Causality of Flows

Hirochika Asai (The University of Tokyo, Japan); Kensuke Fukuda (National Institute of Informatics, Japan); Hiroshi Esaki (The University of Tokyo, Japan)

Traffic causality graphs (TCGs) are proposed for visualizing and analyzing the temporal and spatial causality of flows to profile network applications without inspecting packet payload. A key idea of TCGs is to focus on the causality of individual flows composed of different application protocols rather than a set of host flows. This idea enables us to analyze temporal interactions between flows, such as the temporal

manner of flow generation by identical application programs and interactions between incoming and outgoing flows. We demonstrate the effectiveness of TCGs for profiling network applications in case studies with ground truth datasets. The results show that the simple features of TCGs are discriminative for profiling network applications and that TCGs are also advantageous for profiling application programs, such as user agents of Web browsers and proxies that cannot be classified by existing approaches. One practical use of application program profiling is to identify a specific application program that uses the same protocol as other programs but has security problems. In addition to the TCG features, the visualization of TCGs reveals the causality of each flow, which consequently helps network operators to identify the root causes of other flows, such as malicious ones.

The Memory Effect and Its Implications on Web QoE Modeling

Tobias Hoßfeld (University of Wuerzburg, Germany); Sebastian Biedermann (Technische Universität Darmstadt & Computer Science Departement, Germany); Raimund Schatz (Telecommunications Research Center Vienna (FTW), Austria); Alexander Platzer (FTW, Austria); Sebastian Egger (Telecommunications Research Center Vienna (FTW) & Signal Processing and Speech Communication Laboratory, Graz University of Technology, Austria); Markus Fiedler (Blekinge Institute of Technology, Sweden)

Quality of Experience (QoE) has gained enormous attention during the recent years. So far, most of the existing QoE research has focused on audio and video streaming applications, although HTTP traffic carries the majority of traffic in the residential broadband Internet. However, existing QoE models for this domain do not consider temporal dynamics or historical experiences of the user's satisfaction while consuming a certain service. This psychological influence factor of past experience is referred to as the memory effect. The first contribution of this paper is the identification of the memory effect as a key influence factor for Web QoE modeling based on subjective user studies. As second contribution, three different QoE models are proposed which consider the implications of the memory effect and imply the required extensions of the basic models. The proposed Web QoE models are described on behalf of a) support vector machines, b) iterative exponential regressions, and c) two-dimensional hidden Markov models.

6:00 PM - 7:30 PM

R: Reception 

Wednesday, September 7

Room: Grand Ballroom (for all sessions)

9:00 AM - 10:00 AM

K-2: Keynote Presentation: Scheduling in Networks 

Jean Walrand (University of California-Berkeley, USA) (Bio)

Resource allocation in a large data center, cloud, wireless network, or switch is complex. An optimal allocation is NP-hard and requires full information. This talk reviews recently discovered practical schemes that can be proved to be epsilon-optimal in utility. The schemes are illustrated with applications to switches, ad hoc networks, general resource allocation problems, and processing networks.

The talk is based on the recent monograph: L. Jiang and J. Walrand: Scheduling and Congestion Control for Wireless and Processing networks. Morgan & Claypool, 2010.

10:30 AM - 12:30 PM

TS-4: Content Delivery 

Chair: Debasis Mitra (Bell Labs, Lucent Technologies, USA)

Modeling data transfer in content-centric networking

Giovanna Carofiglio (Bell Labs, Alcatel-Lucent, France); Massimo Gallo (Orange Labs, France); Luca Muscariello (Orange Labs - France Telecom, France); Diego Perino (Bell Labs, Alcatel-Lucent, France)

Content-centric networking proposals, as Parc's CCN, have recently emerged to define new network architectures where content, and not its location, becomes the core of the communication model. These new paradigms push data storage and delivery at network layer and are designed to better deal with current Internet usage, mainly centered around content dissemination and retrieval. In this paper, we

develop an analytical model of CCN in-network storage and receiver-driven transport, that more generally applies to a class of content oriented networks identified by chunk-based communication. We derive a closed-form expression for the mean stationary throughput as a function of hit/miss probabilities at the caches along the path, of content popularity and of content/cache size. Our analytical results, supported by chunk level simulations, can be used to analyze fundamental trade-offs in current CCN architecture, and provide an essential building block for the design and evaluation of enhanced CCN protocols.

Selfish Content Replication on Graphs

Valentino Pacifici (Royal Institute of Technology, Sweden); György Dán (KTH, Royal Institute of Technology, Sweden)

Replication games are a model of the problem of content placement in computer and communication systems, when the participating nodes make their decisions such as to maximize their individual utilities. In this paper we consider replication games played over arbitrary social graphs; the social graph models limited interaction between the players due to, e.g., the network topology. We show that in replication games there is an equilibrium object placement for arbitrary social graphs. Nevertheless, if all nodes follow a myopic strategy to update their object placements then they might cycle arbitrarily long before reaching an equilibrium if the social graph is non-complete. We give sufficient conditions under which such cycles do not exist, and propose an efficient distributed algorithm to reach an equilibrium over a non-complete social graph.

Centrality-driven Scalable Service Migration

Panagiotis Pantazopoulos (National and Kapodistrian University of Athens, Greece); Merkourios Karaliopoulos (National and Kapodistrian University of Athens, Greece); Ioannis Stavrakakis (National and Kapodistrian University of Athens, Greece)

As social networking sites provide increasingly richer context, user-centric service development is expected to explode following the example of User-Generated Content. A major challenge for this emerging paradigm is how to make these exploding in numbers, yet individually of vanishing demand, services available in a cost-effective manner; central to this task is the determination of the optimal service host location. We formulate this problem as a facility location problem and devise a distributed and highly scalable heuristic to solve it. Key to our approach is the introduction of a novel centrality metric. Wherever the service is generated, this metric helps to a) identify a small subgraph of candidate service host nodes with high service demand concentration capacity; b) project on them a reduced yet accurate view of the global demand distribution; and, ultimately, c) pave the service migration path towards the location that minimizes its aggregate access cost over the whole network. The proposed iterative service migration algorithm, called cDSMA, is extensively evaluated over both synthetic and real-world network topologies. In all cases, it achieves remarkable accuracy and robustness, clearly outperforming typical local-search heuristics for service migration. Finally, we outline a realistic cDSMA protocol implementation with complexity up to two orders of magnitude lower than that of centralized solutions.

Gaussian Approximation of CDN Call Level Traffic

Andrzej Bak (Warsaw University of Technology, Poland); Piotr Gajowniczek (Warsaw University of Technology, Poland); Marcin Pilarski (Polish Telecom, Poland)

The Content Driven Networks (CDN) plays very important role in providing Internet services to the end users. There is an increasing competition between CDN providers that place their Web servers in various locations, closer to end users. The CDN solutions have to assure however that users are not only able to download different content via caching systems, but also that the CDN providers can ensure the quality adequate to the observed demand. To solve the QoS problem, one needs traffic model based on the observations and measurements taken from the network of the typical CDN provider that adequately captures the nature of CDN traffic. In the paper we analyzed data measured in Polish Telecom CDN network. We investigated the applicability of Gaussian models to model the call level traffic in the CDN network.

Modeling of Crowdsourcing Platforms and Granularity of Work Organization in Future Internet

Tobias Hoßfeld (University of Wuerzburg, Germany); Matthias Hirth (University of Wuerzburg, Germany); Phuoc Tran-Gia (University of Wuerzburg, Germany)

Beside of social media networks, crowdsourcing is one of the emerging new applications and business models in the Future Internet, which can dramatically change the future of work and work organization in the on-line world. The crowdsourcing technology can be viewed as 'Human Cloud' technique, in contrast to 'Machine Cloud Computing'. Using a crowd with a large number of internationally widespread workers and the flexibility of micro-payment services, crowdsourcing platforms like Amazon's MTurk and Microworkers can outsource traditional forms of work organization on a microscopic level of granularity to a large, anonymous crowd of workers, the human cloud. In such platforms work or tasks are organized at a finer granularity and jobs are split into micro-tasks that need to be performed by a human cloud. It is a need of analysis to understand the anatomy of such a platform and of models to describe the time-dependent growth of the human cloud, in order to predict the traffic impact of such novel applications and to forecast the growth dynamics. The purpose of this paper is a measurement-based statistical analysis of a crowdsourcing platform, using the Microworkers.com platform as example. The obtained results are then used to model the growth of such fast-changing environments in the Internet using well-known models from

biology. Based on the findings from the population growth, we develop a deterministic fluid model which is an extension of the SIR model of epidemics, in order to investigate the platform dynamics.

1:30 PM - 3:00 PM

TS-5: Optical Networking 

Chair: Poul E. Heegaard (Norwegian University of Science and Technology & NTNU, Norway)

Adaptive Optical Burst Switching

Thomas Bonald (Telecom ParisTech, France); Raluca Indre (Orange Labs, France); Sara Oueslati (Orange Labs & France Telecom Orange, France)

We propose a modified version of Optical Burst Switching (OBS) that adapts the size of switched data units to the network load. Specifically, we propose a two-way reservation OBS scheme in which every active source-destination pair attempts to reserve a lightpath and for every successful reservation, transmits an optical burst whose size is proportional to the number of active data flows. We refer to this technique as Adaptive Optical Burst Switching. We prove that the proposed scheme is optimal in the sense that the network is stable at all theoretically sustainable loads. We also evaluate the throughput and delay performance of Adaptive OBS through both analysis and simulation in order to assess the practical load ranges at which the network may operate.

Survivable Impairment-Aware Traffic Grooming in WDM Rings

Anteneh A Beshir (Delft University of Technology, The Netherlands); Fernando A. Kuipers (Delft University of Technology, The Netherlands); Ariel Orda (Technion, Israel); Piet Van Mieghem (Delft University of Technology, The Netherlands)

Wavelength Division Multiplexing (WDM) optical networks offer bandwidth using multiple, but independent wavelength channels (or lightpaths), each operating at several Gb/s. Since the traffic between users is usually only a fraction of the capacity offered by a wavelength, several independent traffic streams can be groomed together. In addition, in order to reverse the effect of noise and signal degradations (physical impairments), optical signals need to be regenerated after a certain impairment threshold is reached. We consider survivable impairment-aware traffic grooming in WDM rings, which are among the most widely deployed optical network topologies. We first show that the survivable impairment-aware traffic grooming problem, where the objective is to minimize the total cost of grooming and regeneration, is NP-hard. We then provide approximation algorithms (for uniform traffic), and efficient heuristic algorithms whose performance is shown to be close to the lower-bounds (for non-uniform traffic) both when (1) the impairment threshold can be ignored, and (2) the impairment threshold should be considered.

A Flow-aware MAC Protocol for a Passive Optical Metropolitan Area Network

Philippe Robert (INRIA, France); James Roberts (INRIA, France)

The paper introduces an original MAC protocol for a passive optical metropolitan area network using time-domain wavelength interleaved networking (TWIN) as proposed recently by Bell Labs. Optical channels are shared under the distributed control of destinations using a packet-based polling algorithm. This MAC is inspired more by EPON dynamic bandwidth allocation than the slotted, GPON-like access control generally envisaged for TWIN. Management of source-destination traffic streams is flow-aware with the size of allocated time slices being proportional to the number of active flows. This emulates a network wide, distributed fair queuing scheduler, bringing the well-known implicit service differentiation and robustness advantages of this mechanism to the metro area network. The paper presents a comprehensive performance evaluation based on analytical modelling supported by simulations. The proposed MAC is shown to have excellent performance in terms of both traffic capacity and packet latency.

Pipelining Multicast Scheduling in All-Optical Packet Switches with Delay Guarantee

Zhiyang Guo (Stony Brook University, USA); Yuanyuan Yang (Stony Brook University, USA)

In this paper, we study multicast scheduling in all optical packet switches. First, we propose a novel optical buffer called multicast-enabled Fiber-Delay-Lines (M-FDLs), which can provide flexible delay for copies of multicast packets using only a small number of FDL segments. We then present a Delay-Guaranteed Multicast Scheduling (DGMS) algorithm that considers the schedule of each arriving packet for multiple time slots. We also discuss some desirable features of DGMS in detail, such as guaranteed delay upper bound and adaptivity to transmission requirements. To relax the time constraint of DGMS, we further propose a pipelining technique that distributes the scheduling tasks among a sequence of sub-schedulers. The combinatorial logic circuit design of each sub-scheduler, which further reduces time complexity, is also provided. The performance of DGMS is tested extensively against statistical traffic models and real Internet traffic, and the results show that the proposed DGMS algorithm can achieve ultra-low average packet delay with minimum packet drop ratio.

3:00 PM - 4:00 PM

S: Student Poster Session/Break **Traffic-Centric Modeling of Future Wireless Internet Access Technologies***Paolo Pietro Pileggi (Universit  di Roma - Tor Vergata, Italy); Giuseppe Bianchi (University of Rome "Tor Vergata", Italy)*

Modeling the performance of wireless networks is essential for the success of the Future Wireless Internet. Results obtained from a previous WLAN study suggest that Traffic-Centric Modeling (TCM) is a viable option for developing useful wireless performance models. In this extended abstract, a novel TCM methodology is briefly introduced. We intend to use TCM to develop models of complex wireless systems, such as IEEE 802.16 and 3GPP LTE networks.

Pytomo: A Tool for Analyzing Playback Quality of YouTube Videos*Parikshit Juluri (University of Missouri-Kansas City, USA); Louis Plissonneau (Orange Labs, France); Deep Medhi (University of Missouri-Kansas City, USA)*

Online video services account for a major part of broadband traffic with streaming videos being one of the most popular video services. We focus on the user perceived quality of YouTube videos as it can serve as a general index for customer satisfaction. Our tool, Pytomo~\cite{pytomo}, is a tomography tool that is designed to measure the playback quality of videos as if it is being viewed by the user. We model the YouTube video player to estimate the playback interruptions as experienced by a user watching a YouTube video. We also examine some topology and download statistics (such as delay towards the server, download rates, buffering duration). We aim to analyse different DNS resolvers to get the IP address of the video server. This allows us to show how the DNS resolution impacts the performance of the video download, thus the video playback quality. As the tool is intended to run on multiple ISPs, we discover some interesting results in YouTube distribution policies. These results can be applied to any CDN architecture and should help Internet users to better understand what are the key performance factors of video streaming.

Multi-Layer Network Topology Design for Large-Scale Network*Taiju Mikoshi (Nihon University, Japan); Toyofumi Takenaka (Nihon University, College of Engineering, Japan); Ryuta Sugiyama (NTT, Japan); Akeo Masuda (NTT, Japan); Kohei Shiimoto (NTT, Japan)*

In this paper, we investigate the effectiveness of the conventional heuristic routing method that improves the computation time of MILP for the large-scale IP-optical network design and we propose a more effective heuristic routing method. The results of calculation experiments showed that proposed heuristic routing method is very excellent not only in the reduction of the computation time but also the accuracy of computation results.

On Modeling of Fluctuations in Quasi-Static Approach Describing the Temporal Evolution of Retry Traffic*Kohei Watabe (Osaka University, Japan); Masaki Aida (Tokyo Metropolitan University, Japan)*

We previously introduced a traffic model that describes the behavior of the retry traffic created by users who are impatient when waiting for a service to be provided. The behavior can be described in a simple form if it is assumed that the system offers infinitely fast processing (i.e. ideal). Moreover, we proposed the quasi-static approach that replicates the temporal evolution of traffic in finite speed (i.e. actual) systems. In the quasi-static approach, the difference between the behavior of the ideal system and that of the actual system is treated as stochastic fluctuation. However, work presented to date not verified that the quasi-static approach can express the traffic model. This paper calculates the temporal evolution of traffic in the M/M/1 based model with retry traffic by traditional Monte-Carlo simulations and the quasi-static approach. The results show that quasi-static approach is as good as the traditional approach in modeling the traffic.

Queuing Theoretic Approach to Server Allocation Problem in Time-delay Cloud Computing Systems*Taichi Kusaka (Aichi Prefectural University, Japan)*

Cloud computing is a popular computing model to support processing large volumetric data using clusters of commodity computers. It aims to power the next generation data centers and enables application service providers to lease data center capabilities for deploying applications depending on user QoS (Quality of Service) requirements. Because cloud applications have different composition, configuration, and deployment requirements, quantifying the performance of resource allocation policies and application scheduling algorithms, is important in cloud computing environments for different application and service models under varying load, network time- delay and system size. To obtain quantifying, the authors apply VCHS(Various Customers, Heterogeneous Servers)queuing systems.

On the Benefits of P2P Cache Capacity Allocation*Frank Lehrieder (University of Wuerzburg, Germany); Valentino Pacifici (Royal Institute of*

Technology, Sweden); György Dán (KTH, Royal Institute of Technology, Sweden)

Peer-to-peer (P2P) systems are responsible for a large fraction of inter-ISP transit traffic in the Internet. Hence, many Internet service providers (ISPs) have deployed P2P caches to decrease their P2P related transit traffic. In our work we consider the problem of allocating the limited upload capacity of a P2P cache between a set of overlay swarms. The goal of cache capacity allocation is to increase the amount of transit traffic that can be saved using the cache. Our preliminary results are based on analytical models and simulations and show that cache capacity allocation is a promising means of improving the efficiency of P2P caches. We are currently validating our results via experiments on Planet-Lab.

Rare Events in Network Simulation Using MIP

Emanuel Heidinger (EADS Innovation Works, Germany)

Switched Ethernet has become a serious candidate for real-time communication in industry and consumer electronics. Due to the unpredictable nature of switched networks, tools and techniques became necessary to determine worst case behaviour, such as network calculus, worst case scheduling analysis and network simulation. This paper proposes a network simulation technique to estimate the results expected from a deterministic network calculus analysis. Traffic will be generated from network calculus arrival curves. Packet generation will be triggered at points of time that provoke a worst case situation. These points will be identified by solving a mixed integer programming problem.

Performance Impacts Due to Number Portability Under Various Routing Schemes

Xuan Liu (University of Missouri-Kansas City, USA); Deep Medhi (University of Missouri-Kansas City, USA)

Number portability provides a user the capability to keep her telephone number as she moves to another provider. Number portability technology involves complicated architecture between different providers. From routing perspective, and in order to simplify the model, we consider that the traversal of a call to a ported number may involve three different networks: the Originating Network, the Donor Network, and the Recipient Network. In this work, we consider four routing schemes, Onward Routing (OR), Query on Release (QoR), Call Dropback/Return to Pivot (CD/RtP), and All Call Query (ACQ). for number portability and present preliminary results on their performance under various scenarios from a connection setup delay point of view.

Wideband Spectrum Sensing Experiments in Indoor Wireless Channels

Nicholas Misiunas (University of Massachusetts at Lowell, USA); Pratik Gandhi (University of Massachusetts at Lowell, USA); Raji Remany (University of Massachusetts Lowell, USA); Kavitha Chandra (University of Massachusetts Lowell, USA); Charles Thompson (Umass Lowell, USA)

In this work, spectrum sensing experiments are conducted in indoor wireless channels using state-of-the-art commercial software radio transceivers. The objective is to evaluate the performance of cooperative spectrum sensing of wideband channels and determine achievable gains against the limiting effects of multipath interference, frequency selective and correlated fading. A finite group of narrowband channels occupied by primary users will be synthesized into a wideband signal by the transmitting node, using quadrature mirror filter banks. The receivers located at different spatial positions implement sensing and detection on the narrowband channels deconstructed from the analysis filter bank. The statistics of the spatial power spectral density measurements are examined as a function of transmit power and spatial diversity afforded by dual sensing nodes. Preliminary results presented in this extended abstract are aimed at calibrating the sensing node by driving the indoor channel with an amplitude modulated chirp signal.

Estimating Optimal Cost of Allocating Virtualized Resource with Dynamic Demand

Haiyang Qian (University of Missouri-Kansas City, USA); Deep Medhi (University of Missouri-Kansas City, USA)

The dynamics of the demand impose unique challenges to the virtualized resource allocation problem. We construct mathematical programming models to optimizing the virtualized resource allocation with dynamic demand such that the server operational cost is minimized. The probabilistic aspect of the demand is considered here. We propose a simple index to estimate the optimum by synthesizing the characteristics of the demand and our optimization framework.

Collection of BCNET BGP Traffic

Tanjila Farah (Simon Fraser University, Canada); Sukhchandan Lally (Simon Fraser University, Canada); Rajvir Gill (Simon Fraser University, Canada); Nabil Al-Rousan (Simon Fraser University, Canada); Ravinder Paul (Simon Fraser University, Canada); Don Xu (Simon Fraser University, Canada); Ljiljana Trajkovic (Simon Fraser University, Canada)

This poster paper describes testbed for collection of BCNET Border Gateway Protocol (BGP) traffic. The BGP traffic was collected using special purpose hardware and software. Preliminary data collection was illustrated using the Wireshark and Walrus graph visualization tools.

4:00 PM - 5:30 PM

TS-6: Resource Management 

Chair: Hans van den Berg (TNO, The Netherlands)

Congestion in Large Balanced Multirate Links

Thomas Bonald (Telecom ParisTech, France); Jean-Paul Haddad (University of Waterloo, Canada); Ravi Mazumdar (University of Waterloo, Canada)

We obtain approximations for various performance measures in a multirate link sharing bandwidth under balanced fairness. For a large system, we obtain closed form expressions for the calculation of long run fraction of time that the system is congested, the probability that an arriving flow will not obtain its maximum bit rate and the average fraction of time that an arriving flow is not allocated its maximum bit rate while in the system. The techniques are based on local limit theorems for convolution measures.

Dispatching Problem with Fixed Size Jobs and Processor Sharing Discipline

Esa Hyytiä (Aalto University, Finland); Aleksi Penttinen (Aalto University, Finland); Samuli Aalto (Aalto University, Finland); Jorma Virtamo (Aalto University, Finland)

We consider a distributed server system with m servers operating under the processor sharing (PS) discipline. A stream of fixed size tasks arrives to a dispatcher, which assigns each task to one of the servers. We are interested in minimizing the mean sojourn time, i.e., the mean response time. To this end, we first analyze an M/D/1-PS queue in the MDP framework. In particular, we derive a closed form expression for the so-called size-aware relative value of state, which sums up the deviation from the average rate at which sojourn times are accumulated in the infinite time horizon. This result can be applied in numerous situations. Here we give an example in the context of dispatching problems by deriving efficient and robust state-dependent dispatching policies for homogeneous and heterogeneous server systems. The obtained policies are further demonstrated by numerical examples.

Approximate Fairness Through Limited Flow List

Addisu Eshete (Norwegian University of Science and Technology, Norway); Yuming Jiang (Norwegian University of Science and Technology (NTNU), Norway)

Most of router mechanisms proposed for fair bandwidth sharing lack either (1) simplicity due to complexity of intricate per flow management of all connections (e.g., WFQ, SFQ), (2) heterogeneity due to a design targeting a specific traffic type, e.g., RED-PD and Fair RED (FRED) or (3) robustness due to requirement for proper router configurations (e.g., CSFQ). All of these severely impact the scalability of the schemes. This paper proposes a novel router fairness mechanism, namely Approximate Fairness through Partial Finish Time (AFpFT). Key to the design of AFpFT is a tag field the value of which defines the position of the packet in an aggregate queue shared by all flows. The specific of tag computation depends on the router's role, edge or inner, to the flow. While gateways closest to traffic source manage all flows, successive or inner routers only manage a limited subset at flow level. The managed flows are usually of higher rates than fair share. Following the heavy-tailed Internet flow distribution, these flows are indeed the minority in the Internet. Using extensive simulations, we show that the scheme is highly fair and potentially scalable unlike other proposed schemes.

Measurement-Based Admission Control for Flow-Aware Implicit Service Differentiation

Jordan Augé (Université Pierre et Marie Curie, France); Sara Oueslati (Orange Labs & France Telecom Orange, France); James Roberts (INRIA, France)

It has previously been shown that the combined use of fair queuing and admission control would allow the Internet to provide satisfactory quality of service for both streaming and elastic flows without explicitly identifying traffic classes. In this paper we discuss the design of the required measurement based admission control (MBAC) scheme. The context is different to that of previous work on MBAC in that there is no prior knowledge of flow characteristics and there is a twofold objective: to maintain adequate throughput for elastic flows and to ensure low packet latency for any flow whose peak rate is less than a given threshold. In the paper we consider the second objective assuming realistically that most elastic and streaming flows are rate limited. We propose an MBAC algorithm and evaluate its performance by simulation under different stationary traffic mixes and in a flash crowd scenario. The algorithm is shown to offer a satisfactory compromise between flow performance and link utilization.

6:00 PM - 10:00 PM

D: Social Event and Dinner 

Thursday, September 8

Room: **Grand Ballroom** (for all sessions)

9:00 AM - 10:00 AM

K-3: Keynote Presentation: Research Challenges for Modern Data-Center Networks 

Jeff Mogul (HP Labs, USA) (Bio)

For many years, data-center networks were boring, and everyone was OK with that. But with the rise of innovations such as cloud-computing, massively-parallel applications such as MapReduce, convergence of storage traffic onto data networks, and the commoditization of network hardware, we now have lots of design choices. The large design space creates all sorts of interesting research challenges. In this talk, I will summarize recent research on data-center network designs, and describe some of the challenges for future research. These include physical topology design; multi-path routing and traffic engineering; support for zillions of virtual machines; security and privacy issues; and energy management.

10:30 AM - 12:30 PM

TS-7: Wireless Networks II 

Chair: Phuoc Tran-Gia (University of Wuerzburg, Germany)

A New Reliable Transport Scheme in Delay Tolerant Networks Based on Acknowledgments and Random Linear Coding

Arshad Ali (Telecom SudParis - UMR CNRS SAMOVAR & UPMC Paris VI, France); Eitan Altman (INRIA, France); Tijani Chahed (Telecom SudParis, France); Manoj Panda (INRIA Sophia Antipolis, France); Lucile Sassatelli (I3S - CNRS Universite de Nice UMR, France)

We propose a new reliable transport scheme in Delay Tolerant Networks (DTNs) based on the use of acknowledgments (ACKs) as well as coding. We, specifically, develop a fluid-limit model to derive expressions for the delay performance of the proposed reliable transport scheme and derive the optimal setting of the parameters which minimize the file transfer time. Our results yield optimal values for the number of outstanding random linear combinations to be sent before time-out as well as the optimal value of the time-out itself, which, in turn, minimize the file transfer time.

Topology Control and Channel Assignment in Lossy Wireless Sensor Networks

Dawei Gong (Stony Brook University, USA); Miao Zhao (State University of New York at Stony Brook, USA); Yuanyuan Yang (Stony Brook University, USA)

In wireless sensor networks (WSNs), a significant amount of packets are lost when transmitted over wireless links, leading to unnecessary energy expenditure. This lossy property of a link can be described by the packet reception ratio (PRR) over it. In the literature, it was shown that the PRR of a link is a non-decreasing function of its signal to interference-plus-noise ratio (SINR), which indicates that the PRR can be improved by either enhancing the received power or reducing the interference-plus-noise level. On the other hand, a number of topology control algorithms and channel assignment algorithms have been presented for WSNs to reduce interference. However, most of them simply use the number of interfering nodes to describe the level of interference, which is inaccurate thus cannot guarantee high PRR. In this paper, we propose a joint design of topology control and channel assignment for lossy WSNs, aiming at improving the PRR of each link in the network. We first construct a maximum PRR spanning tree, then adjust the transmitting power and channel of sensor nodes to further improve the PRR of links on the tree. This way, packet retransmission due to lossy links is minimized, which leads to performance improvement in terms of network throughput, energy efficiency and end-to-end packet delay. We formulate the joint design into an optimization problem and prove its NP-hardness. We then present heuristic algorithms to give practical solutions for the problem. We have carried out extensive simulations and the results show that network performance can be significantly improved by using the topology generated by our algorithms compared to the topologies generated by other schemes under the same traffic demand.

MG-Local: A Multivariable Control Framework for Optimal Wireless Resource Management

Chen Liu (University of Alberta, Canada); Mike MacGregor (University of Alberta, Canada); Janelle Harms (University of Alberta, Canada)

Competition for finite resources causes severe congestion and collisions in wireless networks. Without effective management, the network can become unstable, and users may experience very long delay, significant packet loss and poor throughput. In this paper, we propose a multivariable globalized-local (MG-Local) framework of resource management to find a balance between fair allocation and efficient utilization. This framework uses adaptive multivariable control to improve control effectiveness. Our design combines the advantages of both global and local optimization methods, and drives the system toward a global optimum by intelligently exploiting local information, without message passing. We compare the performance of our proposed method with four other approaches: an estimation-based multivariable control, single-variable control, distributed global-optimization, and CSMA/CA. Our experimental results show that our method significantly outperforms all four alternatives in terms of throughput, packet loss rate, end-to-end delay, and fairness.

Joint Mobile Energy Replenishment and Data Gathering in Wireless Rechargeable Sensor Networks

Miao Zhao (State University of New York at Stony Brook, USA); Ji Li (Stony Brook University, USA); Yuanyuan Yang (Stony Brook University, USA)

Recent studies have shown that energy harvesting wireless sensor networks have the potential to provide perpetual network operations by capturing renewable energy from external environment. However, the spatial-temporal profiles of such ambient energy sources typically exhibit great variations, and can only provide intermittent recharging opportunities to support low-rate data services. In order to provide steady and high recharging rates, and achieve energy-efficient data gathering from sensors, in this paper, we propose to utilize mobility for the joint design of energy replenishment and data gathering. In particular, a multifunctional mobile entity, called SenCar in this paper, is employed, which serves not only as a data collector that roams over the field to gather data via short-range communication but also as an energy transporter that charges static sensors on its migration tour via wireless energy transmissions. Taking advantages of the SenCar's controlled mobility, we give a two-step approach for the joint design. In the first step, the locations of a subset of sensors are periodically selected as anchor points, where the SenCar will sequentially visit to charge the sensors at these locations and gather data from nearby sensors in a multi-hop fashion. In order to achieve a desirable balance between the energy replenishment amount and data gathering latency, we provide a selection algorithm to search for a maximum number of anchor points where sensors hold the least battery energy, and meanwhile by visiting them the tour length of the SenCar is no more than a threshold. In the second step, we consider data gathering performance when the SenCar migrates among these anchor points. We formulate the problem into a network utility maximization problem and propose a distributed algorithm to adjust data rates, link scheduling and flow routing so as to adapt to the up-to-date energy replenishing status of sensors. The effectiveness of our approach is validated by extensive numerical results. When compared with solar harvesting networks, our solution can improve the network utility by 48% on the average.

Metering Re-ECN: Performance Evaluation and its Applicability in Cellular Networks

Ying Zhang (Ericsson Research, USA); Ingemar Johansson (Ericsson Research, USA); Howard Green (Ericsson Research, USA); Mallik Tatipamula (Ericsson Research, USA)

The Re-ECN protocol is a recently proposed congestion notification scheme for IP networks. Building upon Explicit Congestion Notification (ECN), which marks packets instead of drops them during congestion, Re-ECN requires the end users to re-insert the marking information back to the network as a feedback to allow the network react more effectively on resource management. In this work, we first propose an infrastructure based deployment strategy by incorporating Re-ECN mechanisms into the GTP-U protocols in cellular network without changing any end hosts. We conduct comprehensive analysis to assess the gain applications could achieve from using Re-ECN. We study the effect of Re-ECN on TCP in networks with a large number of short-lived and long-lived flows, different hops and variety of RTTs and loss rates. We compare it to other congestion avoidance mechanisms.

1:30 PM - 3:40 PM

TS-8: Overlay, Addressing, and Control 

Chair: Richard J Harris (Massey University, New Zealand)

Planarity of Data Networks

Rhys Bowden (University of Adelaide, Australia); Hung Xuan Nguyen (University of Adelaide, Australia); Nickolas J Falkner (University of Adelaide & The University of Adelaide, Australia); Simon Knight (University of Adelaide, Australia); Matthew Roughan (University of Adelaide, Australia)

The study of network topology has attracted a great deal of attention in the last decade. However, this study has been hampered by a lack of accurate data. Existing methods for measuring topology have many flaws, and arguments about the importance of these flaws have overshadowed the more interesting questions about network structure. The Internet Topology Zoo is a store of network data created from the information that network operators make public. As such it is the most accurate large-scale collection of network topologies available, and includes meta-data that we could never have derived from automated network measurements. With this data we can answer questions about network structure with more certainty than ever before, and in this paper we study a particular property of these networks -- planarity. A surprising number of graphs in the Zoo are planar, many more than can be explained through existing models (either optimization or randomization based) for graph formation. We speculate that planarity results from the requirement to build transparent networks that network operators can understand easily.

The Green-Game: Striking a Balance between QoS and Energy Saving

Aruna Prem Bianzino (Telecom ParisTech, France); Claude Chaudet (Telecom Paristech, France); Dario Rossi (Telecom ParisTech, France); Jean-Louis Rougier (TELECOM ParisTech)

& (formerly named ENST), France); Stefano Moretti (Université Paris Dauphine CNRS UMR, France)

The energy consumed by communication networks can be reduced in several ways. A promising technique consists in concentrating the workload of an infrastructure on a reduced set of devices, while switching off the others, and is referred to as resource consolidation. This technique is particularly suited to routing of data traffic when a network is lightly loaded, but the selection of the routers that can be safely switched off requires an accurate evaluation of the criticality of the devices. In this article, we define a measure of the criticality of the different interconnection devices that does not only take into account a network topology but also its traffic matrix. We model the scenario as a coalitional game, and show how the use of the Shapley value associated to the nodes can be used as criticality indexes. This index is compared with other classical indexes on realistic topologies. We find out that our index provides a robust and relevant criticality measure achieving a good tradeoff between energy efficiency and network robustness.

Dissemination of Address Bindings in Multi-substrate Overlay Networks

Jorg Liebeherr (University of Toronto, Canada); Majid Valipour (University of Toronto, Canada)

Self-organizing overlay networks have emerged as a new paradigm for providing network services. While most overlay networks are built over a single substrate network (mostly, the Internet), recently the construction of overlay networks over multiple heterogeneous substrate networks has received increased attention. Such networks seek to interconnect mobile or fixed devices using a diverse set of networking modalities. Here, a key challenge arises from the more complex address bindings, where a single logical identifier is bound to multiple substrate addresses. In this paper, we evaluate the design and inherent tradeoffs of mechanisms for exchanging information on address bindings in a multi-substrate overlay network. The evaluation is done using measurement experiments of an overlay network software system. The measurement data provides insights into the scalability of dissemination methods. An important finding is that gossip based address dissemination is less effective than an on-demand dissemination of address bindings.

Two-level Cache Architecture to Reduce Memory Accesses for IP Lookups

Sunil Ravinder (University of Alberta, Canada); Mario A Nascimento (University of Alberta, Canada); Mike MacGregor (University of Alberta, Canada)

Longest-prefix matching (LPM) is a key processing function of Internet routers. This is an important step in determining which outbound port to use for a given destination address. The time required to look-up the outbound port must be less than the minimum inter-arrival time between packets on a given port. Lookup times can be reduced by caching address prefixes from previous lookups. However all misses in the prefix cache (PC) will initiate a traversal of the routing table to find the longest matching prefix for the destination address. This table is stored in memory so a traversal requires multiple (perhaps many) memory references. These memory references become an increasingly serious bottleneck as line rates increase. In this paper we present a novel second level of caching that can be used to expedite lookups that miss in the PC. We call this second level a "dynamic substride cache" (DSC). Extensive experiments using real traffic traces and real routing tables show that the DSC is extremely effective in reducing the number of memory references required by a stream of lookups. We also present analytical models to find the optimal partition of a fixed amount of memory between the PC and DSC.

The Implication of Overlay Routing on ISPs' Connecting Strategies

Xun Shao (Osaka University, Japan); Go Hasegawa (Osaka University, Japan); Yoshiaki Taniguchi (Osaka University, Japan); Hiroataka Nakano (Osaka University, Japan)

The Internet is composed of many distinct networks, operated by independent Internet Service Providers (ISPs). There are primarily two kinds of relationships between ISPs: transit and peering. ISPs' traffic and economic relationships are mainly decided by ISPs' routing policy. However, in today's Internet, overlay routing, which changes traffic routing at the application layer to better satisfy the applications' demands, is rapidly increasing, and brings challenge to the ISPs' settlement interconnection researches. The goal of this paper is to study the economic implications of overlay routing on ISPs' peering incentive, costs and strategy choice. For this purpose, we introduce an ISP interconnection business model based on a simple ISPs' network. We then study the overlay traffic patterns in the network in various conditions. Combining the business model with traffic patterns, we study the ISPs' economic issues such as incentive to upgrade peering link and cost reduction conditions with various overlay traffic patterns and settlement methods. We also analyze the bilateral Nash equilibrium (BNE) strategy of ISPs in the network. At last, we give an numerical example to explain our results.

A Saturated Tree Network of Polling Stations with Flow Control

Paul Beekhuizen (EURANDOM, The Netherlands); Jacques Resing (Eindhoven University of Technology, The Netherlands)

We consider a saturated tree network with flow control. The network consists of two layers of polling stations, all of which use the random polling service discipline. We obtain the equilibrium distribution of the network using a Markov chain approach. This equilibrium distribution can be used to efficiently compute the division of throughput over packets from different sources. Our study shows that this throughput division is determined by an interaction between the flow control limits, buffer sizes, and the service

discipline parameters. A numerical study provides more insight in this interaction. The study is motivated by networks on chips where multiple masters share a single slave, operating under flow control.

4:00 PM - 5:00 PM

P: Panel: Economics and Pricing of Mobile Data Services

Moderator: [Iraj Saniee](#), Bell Labs, Alcatel-Lucent, USA

Panelist:

[Mung Chiang](#), Princeton University

[Ramesh Johari](#), Stanford University

[Marcus Weldon](#), ALU Bell Labs

[Roger Neal](#), NYU-Poly Media Lab

There is growing interest and anticipation regarding the future of wireless data services, since wireless is emerging as a primary mode of access to the Internet. While the early model of flat pricing and "all you can eat" for data services was appropriate to jump-start a nascent industry, the immense popularity of wireless data services signal maturity. The so-called wireless "bandwidth crunch" is likely to put an end to the all-you-can-eat bucket pricing models. There are many directions to proceed, and service differentiation in terms of quality and price provide a natural evolutionary path.

Concurrently, technological advances in commercial wireless communication systems known as 3rd or 4th Generation Networks, such as UMTS, WCDMA or LTE, have not only provided more resources (spectrum) for higher data rates to the end users, they have also made the management of the spectrum more efficient and granular: Small slices of the LTE spectrum can be allocated and managed in small time windows. Since the wireless spectrum is naturally both a scarce and a perishable resource, it is not hard to see that differentiated allocation of (the spectrum) resources among users, above and beyond basic services (such as real-time voice and messaging) provides significant scope for optimizing revenues of future wireless data services.

The participants in this panel will present individual perspectives from academia and industry, both as laymen and specialists, on this topic and respond to questions from the audience.

5:00 PM - 5:30 PM

C: Closing Session (Best Paper Award)

Friday, September 9

9:00 AM - 12:30 PM

TM-1: Tutorial: Network Survivability Modeling and Quantification

[Poul E. Heegaard](#) (Norwegian University of Science and Technology, Norway) and [Kishor S. Trivedi](#) (Duke University, USA) ([Bios](#))

Room: Jackson

This tutorial will consider the network's ability to survive major and minor failures in network infrastructure and service platforms that are caused by undesired events that might be external or internal. Survive means that the services provided comply with the requirement also in presence of failures. The network survivability is quantified as defined by the ANSI T1A1.2 committee -- that is, the transient performance from the instant an undesirable event occurs until steady state with an acceptable performance level is attained. The goal of this tutorial is to provide an introduction to the concept and definition of survivability and to demonstrate approaches to model and quantify the survivability in networks. Examples are taken from the survivability of virtual connection over an IP network.

TM-2: Tutorial: Building an End-End Nationwide IPTV Service

[K.K. Ramakrishnan](#) (AT&T Labs-Research, USA) ([Bio](#))

Room: Sansome

This tutorial will examine the architecture and protocols needed for building an end-end nationwide consumer IPTV service based on the Internet Protocol suite. The tutorials will describe models for live channel tuning and interactive viewing of content based on long-term analysis of user behavior.

W1-M: Workshop: Cnet 2011

International Workshop on Modeling, Analysis, and Control of Complex Networks

Room: Montgomery

Technical program is at <http://www.i-teletraffic.org/itc23/workshops/cnet2011-program/>

W2-M: Workshop: DC-CaVeS

3rd Workshop on Data Center - Converged and Virtual Ethernet Switching

Room: [Washington](#)

Technical program is at <http://www.i-teletraffic.org/itc23/workshops/dc-caves-program/>

1:30 PM - 5:00 PM

TA-3: Tutorial: Cooperative and Green Heterogeneous Wireless Networks

Rose Qingyang Hu (Utah State University, USA), Shilpa Talwar (Intel Corporation, USA), Pingping Zong (Verizon, USA) ([Bios](#))

Room: [Sansome](#)

The tutorial discusses the need for alternative strategies design and plan Heterogeneous Network (HetNet), where low power nodes are overlaid within a macro network. HetNets have drawn high research interests from both academia and industry. They have also attracted the attention of the standardization bodies recently, such as 3GPP LTE and IEEE 802.16m. The tutorial will explore a broad scope of technical areas that are under investigation in the context of HetNets. These areas include node/client cooperation, interference management, mobility, green radio, QoS, yield management, security, applications and services.

TA-4: Tutorial: Traffic Models for Quality of Experience Assessment

Markus Fiedler (Blekinge Institute of Technology, Sweden) ([Bio](#))

Room: [Jackson](#)

The tutorial aims at building bridges between traffic modelling and analysis -- which is at the core of ITC -- to QoE assessment. As a means of classification and prioritisation, generic relationships between QoE and QoS are given and discussed. The tutorial will focus on a set of models amongst others from the areas of mobile video streaming, network monitoring, seamless communications, service chains, and virtualisation, which allow to understand how technical choices affect QoE. The models are motivated by user studies on the one hand and by measurements taken in operational networks on the other hand.

W1-A: Workshop: Cnet 2011

International Workshop on Modeling, Analysis, and Control of Complex Networks

Room: [Montgomery](#)

Technical program is at <http://www.i-teletraffic.org/itc23/workshops/cnet2011-program/>

W2-A: Workshop: DC-CaVeS

3rd Workshop on Data Center - Converged and Virtual Ethernet Switching

Room: [Washington](#)

Technical program is at <http://www.i-teletraffic.org/itc23/workshops/dc-caves-program/>