

Detailed Publication List

George Xylomenos

Books and Dissertations

- [D.4] *Multimedia Technology and Multimedia Communications* (in Greek), George Xylomenos and George C. Polyzos, Klidarithmos, Athens, Greece, 2009.

Contents: Introduction, Multimedia applications, Multimedia systems, Audio, Image, Video, Information theory, Coding principles, Entropy coding, Audio coding, Image coding: JPEG, Video coding: H.26x, Video coding: MPEG, Multimedia synchronization, Multicasting, Multimedia transmission protocols, Streaming multimedia, Teleconferencing, Best effort services, Guaranteed Quality of Service, ATM networks, Multimedia in 3G networks.

- [D.3] *Distributed Systems with Java* (in Greek), 3rd Edition, John Cavouras, John Milis, George Xylomenos and Ekaterini Roukounaki, Klidarithmos, Athens, Greece, 2011.

Contents: Introduction, Event ordering and Global States, Global Predicates, Mutual Exclusion, Leader Election, Fault Tolerance, Atomic Transactions, Communication Mechanisms, Concurrency and Threads, Security Issues, Load Distribution, Distributed Naming Systems, Distributed Hash Tables, Distributed Storage Systems, Distributed File Exchange, Distributed Object Systems, Message Oriented Middleware, Web Services.

- [D.2] *Multi Service Link Layers: An Approach to Enhancing Internet Performance over Wireless Links*, George Xylomenos, Ph.D. Dissertation, University of California, San Diego, Dept. of Computer Science and Engineering, 1999.

Abstract: Wireless communications and Internet use have both experienced explosive growth rates during the 1990's. Unfortunately, the performance of Internet applications over wireless links is severely degraded by transmission errors. Previous approaches to those performance problems have not succeeded in bridging the gap between end-to-end application requirements and local wireless link quality so as to offer a universal solution.

Aiming to solve wireless problems at their source, we concentrated on the link layer of the network protocol stack. To understand what types of solutions can optimize Internet performance over wireless links, we performed an extensive simulation study of different applications over a variety of transport layer protocols, wireless links, network topologies and link layer mechanisms. Our analysis reveals that the best solution for each situation depends on underlying link properties but not on awareness of higher layer semantics. We found that the use of appropriate link layer error control mechanisms leads to tremendous improvements in application performance, for both loss intolerant and loss tolerant applications. However, different solutions are preferable to satisfy diverse application requirements.

In order to optimize the performance of heterogeneous mixes of Internet traffic over wireless links, we proposed a multi service link layer architecture that allows the simultaneous operation of multiple link layer mechanisms, each satisfying the needs of a particular class of applications. Our proposal can be extended to support future needs and customized for any wireless link with minimal effort. We evaluated this architecture by repeating our extensive simulations with all applications operating simultaneously over multi service link layers, using appropriate mechanisms for each type of traffic. Our results show that all applications improved their performance by virtually the same factor as when operating in isolation over their preferred link layer mechanisms.

Our architecture can be transparently integrated with the existing Internet to enhance application performance without any modifications to other protocols. It can also complement emerging Quality of

Service provisioning schemes so as to combine congestion and error management. Finally, it provides standardized performance metrics that can be used to dynamically create end-to-end services for adaptive protocols and applications.

- [D.1] *Operating System Kernel Simulator, C Implementation Manual* (in Greek), George Xylomenos, B.S. Dissertation, Athens University of Economics and Business, Dept. of Informatics, Athens, Greece, 1993.

Abstract: The computer system simulator was originally developed by J.C. Cavouras, as part of his Ph.D. dissertation which was published in 1978, using the PL/1 language on an IBM 370/168. A second version was developed by J.C. Cavouras in 1980, using the C language on a DEC VAX 11/780 under Version 7 UNIX, which however could not be executed in other systems. The author undertook the task of revising the C language version so as to allow it to be executed in any modern UNIX system, as part of his B.S. dissertation on operating systems in the Department of Applied Informatics of the Athens University of Economics and Business (AUEB), during the 1992-93 Academic Year. The basic deliverable of the dissertation is an ANSI C and a UNIX C version of the program, with minor differences among the two. Porting the program led to considerable modifications of its code in order to optimize its performance and improve its quality in terms of structure and readability. In parallel, all program messages were localized and the code was internally documented (in Greek). This manual consists of the external documentation of the code and some remarks on the structure and operation of the program which may be of assistance to its users during program execution and modification.

Refereed Journals

- [J.26] *IoT Group Membership Management Using Decentralized Identifiers and Verifiable Credentials*, Nikos Fotiou, Vasilios A. Siris, George Xylomenos and George C. Polyzos, *Future Internet*, vol. 6, no. 14, 173, 2022.

Abstract: Many IoT use cases can benefit from group communication, where a user requests an IoT resource and this request can be handled by multiple IoT devices, each of which may respond back to the user. IoT group communication involves one-to-many requests and many-to-one responses, and this creates security challenges. In this paper, we focus on the provenance that has been received by an authorized device. We provide an effective and flexible solution for securing IoT group communication using CoAP, where a CoAP client sends a request to a CoAP group and receives multiple responses by many IoT devices, acting as CoAP servers. We design a solution that allows CoAP servers to digitally sign their responses in a way that clients can verify that a response has been generated by an authorized member of the CoAP group. In order to achieve our goal, we leverage Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs). In particular, we consider that each group is identified by a DID, and each group member has received a VC that allows it to participate in that group. The only information a client needs to know is the DID of the group, which is learned using DNSSEC. Our solution allows group members to rotate their signing keys, it achieves group member revocation, and it has minimal communication and computational overhead.

- [J.25] *On the Impact of Audio Characteristics to the Quality of Musicians' Experience in Network Music Performance*, Konstantinos Tsioutas and George Xylomenos, *Journal of the Audio Engineering Society*, vol. 69, no. 12, 2021, pp. 914–923.

Abstract: Audio delay is a crucial factor in the Quality of Musicians' Experience (QoME) in Network Music Performance (NMP). Previous studies have explored the tolerance of musicians to delay and its dependence on the timbre of the instruments used and the tempo of the performance. Although their findings are intriguing, the small size of these studies makes the extraction of concrete conclusions quite hard. In order to shed more light on these issues, we undertook a larger scale NMP study with real musicians, assessing a wide range of subjective QoME variables against delay, and correlated these results with the audio characteristics of the instruments and the performance. Due to the large number of participants, our findings validate and extend previous studies with a wider array of QoME variables and audio characteristics.

- [J.24] *Multipath Congestion Control with Network Assistance*, Yannis Thomas, George Xylomenos and George C. Polyzos, *Computer Communications*, vol. 153, 2020, pp. 264-278.

Abstract: The use of multiple transport flows over distinct, if possible, paths, is a well-known technique for enhancing the performance and stability of data transfer. Multipath TCP (MPTCP), the most popular multipath transport protocol in-use, allows a single receiver to exploit multiple paths from a single sender. Nevertheless, MPTCP cannot fully exploit the potential gains of multipath connectivity, as it must fairly share resources with regular, single-path TCP, without knowing whether the available paths are distinct or share bottleneck links, due to IP's design choices. We introduce a hybrid congestion control algorithm for multipath transport that enables higher bandwidth utilization compared to MPTCP, while remaining friendly to TCP-like flows. Our solution employs (i) Normalized Multiflow Congestion Control (NMCC), a novel end-to-end congestion control algorithm and (ii) an in-network module that exposes topological information to the end-users in order to support the greedy friendliness technique. The end-to-end NMCC is architecture-independent and can be seamlessly integrated with MPTCP. The in-network module has been implemented for the PSI Information-Centric Networking architecture, but it can also be integrated with Multi-Protocol Label Switching (MPLS) and Software Defined Networking (SDN). Using an actual protocol implementation deployed on our testbed, as well as on a comprehensive packet-level simulator, we obtain experimental results which demonstrate clear gains for our design in terms of throughput and friendliness to other flows.

- [J.23] *Low Latency Friendliness for Multipath TCP*, Yannis Thomas, Merkourios Karaliopoulos, George Xylomenos and George C. Polyzos, *IEEE/ACM Transactions on Networking*, vol. 28, no. 1, 2020, pp. 248–261.c.69

Abstract: Efficient congestion control is critical to the operation of MPTCP, the Multipath extension of TCP. Congestion control in such an environment primarily aims at enhancing the cumulative TCP throughput over the available paths, while preserving TCP-friendliness by fairly sharing the available bandwidth with single-path TCP flows in each path. While most existing multipath congestion control algorithms fulfill the TCP-friendliness objective in their steady state, their throughput convergence latency is high, rendering them ineffective for short-lived flows. We have proposed Normalized Multipath Congestion Control (NMCC), an MPTCP congestion control algorithm that achieves TCP-friendliness faster, by normalizing the growth of individual sub-flow throughput rather than the throughput itself. As NMCC can become unfriendly when it experiences sparse congestion events, in this paper we introduce the extended NMCC (e-NMCC) protocol that caters for TCP-friendliness upon both throughput growth and throughput reduction epochs. We analytically characterize e-NMCC in terms of TCP-friendliness and responsiveness and compare it with alternative algorithms. Finally, we assess the performance of e-NMCC through experimentation with the htsim simulator and a real Linux implementation. Our results confirm that e-NMCC accelerates throughput convergence, thus ensuring TCP-friendliness regardless of connection duration and underlying network conditions.

- [J.22] *Improving Video QoE with IP-over-ICN*, Ioannis Doumanis, Alexander Phinikarides, George Xylomenos, Stuart C.M. Porter and Michael Georgiades, *International Journal of Network Management*, vol.30, no. 3, 2020.

Abstract: Information-centric networking (ICN) has long been advocating for radical changes to the Internet, but the upgrade challenges that this entails have hindered its adoption. To break this loop, the POINT project proposed a hybrid, IP-over-ICN, architecture: IP networks are preserved at the edge, connected to each other over an ICN core. This exploits the key benefits of ICN, enabling individual network operators to improve the performance of their IP-based services, without changing the rest of the Internet. This paper first provides an overview of POINT and outlines how it can improve upon IP in terms of performance and resilience. It then describes a trial of the POINT prototype in a production network, where real users operated actual IP-based applications. As part of the trial, we carried out experiments to evaluate the Quality of Experience (QoE) for video services offered via either HLS or IPTV, using either IP or POINT as a substrate. The results from the trial verify that the IP-over-ICN approach of POINT offers enhanced QoE to the users of these video services, compared to traditional IP, especially under exceptional network conditions.

- [J.21] *Adaptive Semi-Stateless Forwarding for Content-Centric Networks*, Christos Tsilopoulos and George

Xylomenos, *Computer Networks*, vol. 153, 2019, pp. 73–85.

Abstract: In Content-Centric Networks, users request content by issuing Interest messages, receiving Data messages in response. Interests leave state on routers on their path, to allow Data to follow the reverse path. Despite the advantages made possible by this state, its storage requirements raise scalability concerns. We propose an adaptive semi-stateless forwarding scheme where Interests are tracked only on a fraction of the routers. Between state-tracking routers, Interests gather reverse path information, which is used to deliver Data via Bloom filter-based forwarding. We show how the fraction of state-tracking routers can be dynamically adapted to reduce state while limiting communication overheads. Our scheme allows this state, along with all other required data structures, to comfortably fit into the fast memory of a practical router.

- [J.20] *Error and Congestion Control for Wireless Sensor Networks*, Charilaos Stais and George Xylomenos, *International Journal of Communication Systems*, vol. 32, no. 4, 2019.

Abstract: In the Wireless Sensors Network (WSN) field, a wide variety of sensors produce a heterogeneous traffic mix, targeting diverse applications with different reliability requirements. We focus on emergency response scenarios, where a mobile rescuer moves through a, possibly disconnected, network, trying to talk to diverse sensors. We assume two types of sensors, event sensors triggered by an event and periodic sensors activated at predefined time intervals, as well as two types of transmission, either using the highest bit rate available or using predefined bit-rates. Our Reliable Transport protocol for Sensor Networks with MOBILE Sinks (RT-SENMOS), takes into account all these parameters and tries to provide the best possible user experience under the current circumstances of the network, using a sink-driven approach where an application-specific sink is combined with generic sensors. RT-SENMOS was implemented and tested over a real network with emulated losses and compared against Rate-Controlled Reliable Transport (RCRT), a well-known sink-driven protocol. The results show that RT-SENMOS fully exploits the available bandwidth in all cases, while RCRT only manages to exploit 60% to 90% of it. Furthermore, RT-SENMOS adapts much faster to prevailing network conditions, while its protocol overhead, in terms of control messages exchanged, is much lower than that of RCRT.

- [J.19] *Access control as a service for the Cloud*, Nikos Fotiou, Apostolis Machas, George C Polyzos and George Xylomenos, *Journal of Internet Services and Applications*, vol. 6, no. 11, 2015.

Abstract: Cloud computing has become the focus of attention in the computing industry. However, security concerns still impede the widespread adoption of this technology. Most enterprises are particularly worried about the lack of control over their outsourced data, since the authentication and authorization systems of Cloud providers are generic and they cannot be easily adapted to the requirements of each individual enterprise. An adaptation process requires the creation of complex protocols, often leading to security problems and lock-in conditions. In this paper we present the design of a lightweight access control solution that overcomes these problems. With our solution access control is offered as a service by a third trusted party, the Access Control Provider. Access control as a service enhances end-user privacy, eliminates the need for developing complex adaptation protocols, and offers data owners flexibility to switch among Cloud providers, or to use multiple, different Cloud providers concurrently. As a proof of concept, we have implemented and incorporated our solution in the popular open-source Cloud stack OpenStack. Moreover, we have designed and implemented a Web application that enables the incorporation of our solution into Google Drive.

- [J.18] *H-Pastry: An Inter-domain Topology Aware Overlay for the Support of Name-Resolution Services in the Future Internet*, Nikos Fotiou, Konstantinos V. Katsaros, George Xylomenos and George C. Polyzos, *Computer Communications*, vol. 62, 2015, pp. 14–22.

Abstract: Overlay networks are widely used for locating and disseminating information by means of custom routing and forwarding on top of an underlying network. Distributed Hash Table (DHT) based overlays in particular, provide good scalability and load balancing properties. However, these come at the cost of inefficient routing, caused by the lack of adaptation to the underlying network, as DHTs often overlook physical network proximity, administrative boundaries and/or inter-domain routing policies. In this paper we show how to construct a DHT-based overlay network that takes all these aspects into account, so as to ease the global deployment of Future Internet architectures which require largescale name resolution, such as Information-Centric Networking (ICN) and the Internet of Things (IoT). Based

on the Pastry distributed object location and routing substrate and the Canon paradigm for multi-level DHTs, we developed H-Pastry, an overlay DHT scheme that harvests the scalability and load balancing features of DHTs, while also adapting to the underlying network topology, administrative structure and routing policies. We evaluate the performance characteristics of the proposed scheme through an extensive set of detailed simulations over realistic inter-network topologies. Our results show that H-Pastry substantially improves routing by reducing both overlay path stretch (by up to 55%) and routing policy violations (by up to 70%), compared to the Canonical (multi-level) Chord DHT. In addition, the design of H-Pastry keeps traffic within administrative boundaries as far as possible, reducing inter-domain hops by up to 27% compared to Pastry, while also creating excellent opportunities for the support of caching and multicast.

- [J.17] *A Reliable Multicast Transport Protocol for Information-Centric Networks*, Charilaos Stais, George Xylomenos and Alexios Voulimeneas, *Journal of Network and Computer Applications*, vol. 50, 2015, pp. 92–100.

Abstract: In the past few years, many researchers have argued that the Internet should transition from its traditional endpointcentric architecture to an information-centric paradigm. One of the advantages of the information-centric model is that the network can easily aggregate requests for the same content and serve them via multicast. Indeed, most information-centric architectures proposed to date offer native support for multicast, promising a vast improvement in the efficiency of content distribution. However, designing efficient reliable transport protocols for multicast is a largely open issue, due to the problem of feedback implosion towards the sender as group size grows. In this paper we propose RMTPSI, a retransmission-based reliable error control protocol for multicast communication designed specifically for information-centric networks. We compare RMTPSI with existing approaches proposed for IP multicast and evaluate its performance via simulation, showing that our approach leads to more efficient content distribution and error recovery than previous solutions.

- [J.16] *A Survey of Information-Centric Networking Research*, George Xylomenos, Christopher N. Ververidis, Vasilios A. Siris, Nikos Fotiou, Christos Tsilopoulos, Xenofon Vasilakos, Konstantinos V. Katsaros and George C. Polyzos, *IEEE Communications Surveys and Tutorials*, vol. 16, no. 2, 2014, pp. 1024–1049.

Abstract: The current Internet architecture was founded upon a host-centric communication model, which was appropriate for coping with the needs of the early Internet users. Internet usage has evolved however, with most users mainly interested in accessing (vast amounts of) information, irrespective of its physical location. This paradigm shift in the usage model of the Internet, along with the pressing needs for, among others, better security and mobility support, has led researchers into considering a radical change to the Internet architecture. In this direction, we have witnessed many research efforts investigating Information-Centric Networking (ICN) as a foundation upon which the Future Internet can be built. Our main aims in this survey are: (a) to identify the core functionalities of ICN architectures, (b) to describe the key ICN proposals in a tutorial manner, highlighting the similarities and differences among them with respect to those core functionalities, and (c) to identify the key weaknesses of ICN proposals and to outline the main unresolved research challenges in this area of networking research.

- [J.15] *Handling Mobility in Future Publish-Subscribe Information-Centric Networks*, Nikos Fotiou, Konstantinos Katsaros, George C. Polyzos, Mikko Särelä, Dirk Trossen and George Xylomenos, *Telecommunication Systems*, vol. 53, no. 3, 2013, pp. 299–314.

Abstract: Future information-oriented Internet architectures are expected to effectively support mobility. PSIRP, an EU FP7 research project, designed, prototyped, and investigated a clean-slate architecture for the future Internet based on the publish-subscribe paradigm. PURSUIT, another EU FP7 research project, is further developing this architecture, which we refer to as Ψ , the Publish Subscribe Internet (PSI) architecture, extending it in various directions, including a deeper investigation of higher (transport and application) and lower layers (e.g., various link technologies, such as wireless and optical). In this paper we present the basics of the Ψ architecture, including the built-in multicast and caching mechanisms, with particular focus on mobility support. We discuss how the native, clean-slate, Ψ instantiation of the information-centric model can support mobility and also present an overlay variant of \square we have developed in order to provide an evolutionary path to adoption. Based on analysis and

simulation we demonstrate the advantages of the proposed architecture compared to well established solutions such as Mobile IPv6.

- [J.14] *Caching and Mobility Support in a Publish-Subscribe Internet Architecture*, George Xylomenos, Xenofon Vasilakos, Christos Tsilopoulos, Vasilios Siris and George C. Polyzos, *IEEE Communications Magazine*, vol. 50, no. 7, 2012, pp. 52–58.
Abstract: The Internet is straining to meet demands that its design had never anticipated, such as supporting billions of mobile devices and transporting huge amounts of multimedia content. The Publish-Subscribe Internet (PSI) architecture, a clean slate Information Centric Networking approach for the Future Internet, was designed to satisfy the current and emerging user demands for pervasive content delivery that the Internet can no longer handle. This article provides an overview of the PSI architecture, explaining its operation from bootstrapping to information delivery, focusing on its support for network layer caching and seamless mobility, which make PSI an excellent platform for ubiquitous information delivery.
- [J.13] *MultiCache: An overlay architecture for information-centric networking*, Konstantinos Katsaros, George Xylomenos and George C. Polyzos, *Computer Networks*, vol. 55, no 4, 2011, pp. 936–947.
Abstract: It has become apparent for quite some time that the Internet has evolved from a network connecting pairs of end-hosts to a substrate for information dissemination. While this shift towards information centric networking has been clearly demonstrated by the proliferation of file sharing and content delivery applications, it has not been reflected in a corresponding shift in network architecture. To address this issue, we designed MultiCache, an information-centric architecture aiming at the efficient use of network resources. MultiCache is based on two primitives: multicast and caching. It exploits overlay multicast as a means for content delivery and takes advantage of multicast forwarding information to locate, in an anycast fashion, nearby caches that have been themselves fed via multicast. We evaluate MultiCache against a widespread file sharing application (BitTorrent) with respect to both network resource consumption and end-user experience.
- [J.12] *Support of multiple content variants in the multimedia broadcast / multicast service*, George Xylomenos, Konstantinos Katsaros and Vasilis Tsakanikas, *International Journal of Communication Systems*, vol. 24, no 6, 2011, pp. 691–708.
Abstract: The Multimedia Broadcast/Multicast Service (MBMS) was designed to enable the mass distribution of multimedia content in 3rd Generation and beyond cellular networks. If such services are to become commercially viable, they must be able to efficiently support widely heterogeneous user requirements, for example, due to terminal limitations and cost constraints. This paper presents an MBMS extension that allows multiple variants of the same content to be economically distributed to heterogeneous receivers, explicitly taking into account the possibility of using either dedicated or common radio channels. We describe our extended multiple content variant MBMS model by explaining the modifications that it imposes on the standard MBMS model, as well as the manner in which it can be combined with layered coding. We also present an analytical evaluation of our approach against alternatives based on the standard MBMS in terms of control and user plane overhead, and compare the analytical predictions with detailed simulation results. Both the analysis and the simulations indicate that our proposal can indeed satisfy heterogeneous user requirements, while consuming considerably lower resources than the standard-based alternatives.
- [J.11] *Link Layer Adaptation for Shared Wireless Links*, George Xylomenos and Michael Makidis, *Mobile Networks and Applications*, vol. 13, no. 3–4, 2008, pp. 259–273.
Abstract: While traditional link layer protocols assume that they fully control the underlying link, in contemporary wireless networks the link may be dynamically shared by sessions belonging to different users and/or applications. To assess the impact of link sharing, we measure the File Transfer and Web Browsing throughput achieved over a Selective Repeat (SR) protocol, with or without contention from Media Distribution. Our results indicate that the optimal protocol settings strongly depend on the level of contention for the link. We therefore present two link layer protocols that adapt to the available bandwidth, our Adaptive Selective Repeat (ASR) protocol which dynamically modifies its retransmission timeouts, and the Radio Link Control (RLC) protocol specified for use by UMTS networks which does

not employ retransmission timers. We first repeat our performance measurements to determine the optimal settings for each protocol, and then compare the fine tuned versions of all protocols with respect to their File Transfer and Web Browsing throughput, as well as to the delay induced to the contending Media Distribution packets. Our results indicate that while both RLC and ASR are more stable than SR, the complex RLC does not match the performance of our simpler ASR.

- [J.10] *The Multimedia Broadcast / Multicast Service*, George Xylomenos, Vasilis Vogkas and George Thanos, *Wireless Communications and Mobile Computing*, vol. 8, no. 2, 2008, pp. 255–265.

Abstract: This article provides an overview of the multimedia broadcast/multicast service (MBMS) for universal mobile telecommunications system (UMTS) networks. We first outline the features of UMTS networks as defined by the 3rd generation partnership project (3GPP) in order to provide a background for the discussion to follow. We then present the overall MBMS architecture, the services that it provides to the users and the differences between the broadcast and multicast options. The implementation details of MBMS are explained in terms of the modifications needed to the network, the new signaling procedures required and the impact of MBMS on the radio part of the network. We then describe how the security architecture of UMTS is used to support content protection and key distribution for MBMS groups. We evaluate the prospects of MBMS by comparing it first with IP multicast and then with DVB-H, considered by many to be the closest competitor of MBMS. Finally, we conclude with a discussion of the technical and business challenges still faced by MBMS.

- [J.9] *A multi-service link layer architecture for the wireless Internet*, George Xylomenos and George C. Polyzos, *International Journal of Communication Systems*, vol. 17, no. 6, 2004, pp. 553–574.

Abstract: The performance of Internet applications over wireless links is disappointing due to the adverse effects of wireless errors on higher layer protocols and applications. This paper focuses on link layer enhancement mechanisms which attempt to hide these wireless impairments. We simulate file transfer and WWW browsing over TCP and continuous media distribution over UDP, in conjunction with various link layer schemes. Our results reveal that WWW browsing behaves differently than bulk file transfer, that some TCP aware enhancements have limited applicability and that UDP applications are best served by schemes inappropriate for TCP. We then describe a multi-service link layer architecture that simultaneously enhances the performance of diverse applications by supporting multiple error recovery schemes in parallel. In order to evaluate our architecture, we repeat our previous simulations with all applications executing simultaneously. The results reveal that with our approach each application achieves similar improvements as when it operates alone over its preferred link layer.

- [J.8] *Quality of Service Support over Multi-Service Wireless Internet Links*, George Xylomenos and George C. Polyzos, *Computer Networks*, vol. 37, no. 5, 2001, pp. 601–615.

Abstract: Internet application performance over wireless links is disappointing, since wireless impairments adversely affect higher protocol layers. In order to address these problems without global protocol modifications, we examine link layer enhancement schemes. Simulations show that different schemes work best for different applications. We have thus developed a multi-service link layer architecture that simultaneously enhances the performance of diverse applications by supporting multiple link mechanisms concurrently. Simulations show that our approach dramatically improves performance. We present various ways of embedding this architecture into the Internet, thus allowing applications to select themselves the appropriate trade-off between throughput, loss and delay.

- [J.7] *TCP Performance Issues over Wireless Links*, George Xylomenos, George C. Polyzos, Petri Mahonen and Mika Saaranen, *IEEE Communications Magazine*, vol. 39, no. 4, 2001, pp. 52–58.

Abstract: This article discusses the problems arising when the TCP/IP protocol suite is used to provide Internet connectivity over existing and emerging wireless links. Due to the strong drive towards wireless Internet access through mobile terminals, these problems must be carefully studied in order to build improved systems. We review wireless link characteristics using Wireless LANs and Cellular Communications systems as examples. We then outline the performance problems of the TCP/IP protocol suite when employed over those links, such as degraded TCP performance due to mistaking wireless errors for congestion. We present various proposals for solving these problems and examine their benefits and limitations. Finally, we consider the future evolution of wireless systems and the challenges that emerging systems will impose on the Internet protocol suite.

- [J.6] *Link Layer Support for Quality of Service on Wireless Internet Links*, George Xylomenos and George C. Polyzos, *IEEE Personal Communications*, vol. 6, no. 5, 1999, pp. 52–60.

Abstract: We have developed a novel link layer architecture that provides multiple Quality of Service points simultaneously over wireless Internet links. Our approach enhances the performance of diverse applications over error prone links. We discuss the performance problems of Internet protocols over wireless links, presenting as a case study our measurements on a wireless LAN, and argue that it is preferable to handle wireless impairments at the link layer. We present a simulation study of various link layer enhancements and their impact on TCP and UDP performance. Our results show that different approaches are preferable for each type of application. We thus propose a Multi Service Link Layer approach that supports multiple link layer mechanisms over a single link. Our scheme is transport protocol independent and customizable for the underlying wireless link technology. While our approach can be directly deployed on the existing Internet, it also provides support for future Quality of Service-aware protocols and applications. Our simulation results show that our approach improves uniformly the performance of both TCP and UDP applications.

- [J.5] *Internet Protocol Performance over Networks with Wireless Links*, George Xylomenos and George C. Polyzos, *IEEE Network*, vol. 13, no. 4, 1999, pp. 55–63.

Abstract: This article discusses the problems that arise when standard Internet protocols such as TCP are used over wireless links. We review wireless link characteristics with case studies drawn from commercial Wireless LANs and Cellular Telephony systems. We discuss problems with Internet protocols when employed over these systems, such as degraded TCP performance when wireless errors are interpreted as congestion losses. We survey various proposed approaches to mitigating such problems and examine their applicability. Finally, we look at the future of wireless systems and the new challenges that they will create for Internet protocols, and state some goals for further protocol enhancement and evolution, pointing out the need for better protocol integration across layers.

- [J.4] *IP Multicast Group Management for Point-to-Point Local Distribution*, George Xylomenos and George C. Polyzos, *Computer Communications*, vol. 21, no. 18, 1998, pp. 1645–1654.

Abstract: We examine the applicability of existing IP multicast mechanisms for Point-to-Point links such as wired and wireless telephone lines. We identify problems such as overhead due to IGMP leave latency and unnecessary probing of hosts, both important issues for power constrained mobile hosts and low bandwidth wireless links. We propose alternative mechanisms that preserve the IP multicasting model but employ join/leave messages to track group membership. We describe the implementation requirements of our mechanisms and compare them to existing ones with respect to performance, mobile power efficiency, interoperability, robustness and implementation complexity, demonstrating that the join/leave approach is uniformly superior for this environment.

- [J.3] *The Multimedia Multicasting Problem*, Joseph C. Pasquale, George C. Polyzos and George Xylomenos, *ACM Multimedia Systems*, vol. 6, no. 1, 1998, pp. 43–59.

Abstract: This paper explores the problems associated with the multicasting of continuous media to support multimedia group applications. The interaction between multicasting and the delivery of multiple time-correlated continuous-media streams with real-time delay requirements poses various new and interesting problems in research on communication protocols and architectures. We describe these problems, and identify where the opportunities are for effective solutions, all in the context of providing an overview of the current state of research in multimedia multicasting. The issues we discuss include quality of service, resource reservations, routing, error and traffic control, heterogeneity, and the use of hierarchical coding and open-loop control techniques.

- [J.2] *IP Multicast for Mobile Hosts*, George Xylomenos and George C. Polyzos, *IEEE Communications Magazine*, vol. 35, no. 1, 1997, pp. 54–58.

Abstract: We present alternative designs for efficiently supporting multicast for mobile hosts on the Internet. Methods for separately supporting multicasting and mobility along with their possible interactions are briefly described, and then various solutions to the combined problem are explored. We examine three different multicast delivery mechanisms and compare them based on their efficiency and impact on host protocol software.

- [J.1] *Evaluation and Selection Criteria for Software Requirements Specification Standards*, E.A. Giakoumakis and George Xylomenos, *IEEE/BCS Software Engineering Journal*, vol. 11, no. 5, 1996, pp. 307–319.
Abstract: Various organisations have published proposals to prescribe the form and content of software requirements specification documents; the standards were designed to support the specific needs of these organisations and the intricacies of their development projects. To help third parties in taking advantage of this body of work, a set of criteria are proposed and discussed that can be used to evaluate such standards, according to the unique characteristics of specific combinations of organisations and software development projects, and then the question of how the criteria can be applied in an evaluation, selection and tailoring process, depending on the circumstances, is discussed. Finally, the criteria are demonstrated by applying them on some published standards, to help interested organisations to preselect those that seem most appropriate for their needs.

Refereed Conferences

- [C.77] *Selective Content Disclosure using Zero-Knowledge Proofs*, Nikos Fotiou, Vasilis Kalos, Yannis Thomas, George Xylomenos, Vasilios A. Siris and George C. Polyzos, *Proceedings of the Global Information Infrastructure and Networking Symposium (GIIS) 2022*.
Abstract: Information-Centric Networking (ICN) is a Next Generation Internet architecture that facilitates content sharing. ICN natively supports content multi-sourcing, allowing content items to be stored in multiple storage nodes. In order to ensure data integrity, data owners can sign their content items. This, however, prevents storage nodes from sharing partial content items. We present a data sharing scheme where data owners store structured data items (e.g., IoT measurements) in semi-trusted storage nodes. We allow data consumers to express interest for a portion of a data item and we enable storage nodes to “hide” the remaining item without invalidating its integrity. We achieve our goal by leveraging BBS+ digital signatures that support selective data disclosure through Zero-Knowledge Proofs. We define a protocol for data owners to issue authorizations in the form of Verifiable Credentials, which indicate which parts of the data a consumer is allowed to access, and a protocol for consumers to send these authorizations inside ICN Interests. This allows storage nodes to implement fine grain access control, without having access to the secrets of the data owners, while data consumers can still verify the authenticity and integrity of the partially revealed data. In addition to its security advantages, our solution requires significantly less storage and communication overhead compared to an approach that relies on commonly used digital signature algorithms.
- [C.76] *Audio delay in Web Conference tools*, Konstantinos Tsioutas and George Xylomenos, *Proceedings of the Workshop on Web Engineering and Collaborative Music Learning (WECML) 2022*.
Abstract: Music collaboration over the Internet, known as Network Music Performance (NMP), remains a challenge for researchers and engineers, since transmission, switching and audio processing delays hinder the synchronization of the participating musicians. Although widely available Web-based voice and video communication tools are not designed for real-time remote musical performances, during the pandemic many musicians worldwide had to use them, due to the lack of widespread NMP-oriented tools. In this paper we provide measurements for the end-to-end audio delay of a number of Web conference tools, using real network conditions, either in a LAN or in a WAN setting, and compare them to the corresponding delays exhibited by an NMP-specific tool.
- [C.75] *Multimodal Assessment of Network Music Performance*, Konstantinos Tsioutas, Konstantinos Ratzos, George Xylomenos and Ioannis Doumanis, *Proceedings of the Workshop on Multimodal Affect and Aesthetic Experience (MAAE) 2021*.
Abstract: The most common method of assessing the Quality of Musician’s Experience (QoME) in Network Music Performance (NMP) is to perform a subjective study, where the participants evaluate their experience via questionnaires. Translating experiences into metrics is not an exact science, though: in our recent study on the effects of audio delay and quality in the QoME of NMP, the responses had a high variance and were inconsistent. To strengthen our confidence in the results of the subjective study, we analyzed video recordings of the participants using machine learning. Specifically, we used Facial Expression Recognition (FER) to detect the emotions felt by the participants and then compare

them with their questionnaire responses. In addition to pointing out interesting phenomena that were not apparent from the questionnaires, this multimodal analysis showed analogies between the emotions felt (as captured by FER) and the emotions expressed (as captured by the responses).

- [C.74] *Assessing the QoME of NMP via Audio Analysis Tools*, Konstantinos Tsioutas and George Xylomenos, *Proceedings of the International Conference on Signal Processing and Multimedia Applications (SIGMAP) 2021*.

Abstract: Analyzing the Quality of Musicians' Experience (QoME) in Network Music Performance (NMP) typically involves having musicians perform NMP sessions and then assessing their experience via questionnaires. Such subjective studies produce results with wide variances, making the extraction of solid conclusions difficult. For this reason, we complemented a subjective study on the effects of delay in the QoME of NMP with an analysis of the audio captured during the study using automated tools. Specifically, we used signal processing techniques to analyze the captured audio, in order to detect tempo evolution during each performance and examine its correlation with delay. Our results indicate that musicians in real NMP settings are more tolerant to delay than previously thought, holding a steady tempo even with one way delays of 40 ms.

- [C.73] *Securing Named Data Networking routing using Decentralized Identifiers*, Nikos Fotiou, Yannis Thomas, Vasilios A. Siris, George Xylomenos and George C. Polyzos, *Proceedings of the IEEE Semantic Addressing and Routing for Future Networks Workshop (SARNET) 2021*.

Abstract: Named Data Networking (NDN) is a realization of the Information-Centric Networking (ICN) paradigm, where routing is based on content identifiers rather than on network location identifiers. The routing state in NDN can grow exponentially, not only due to the huge number of content identifiers (as opposed to network addresses) but also because it is difficult to detect "fake" routing advertisements. For example, in contrast to IP-based routing, a potentially valid routing entry in NDN can be advertised from multiple network locations, making NDN susceptible to Denial-of-Service attacks at the routing layer. In this paper, we leverage Decentralized Identifiers (DIDs) to build self-verifiable content advertisements. With our solution, any router can verify that a content advertisement originates from an authorized entity, without requiring any trusted third party. We implement our solution and we evaluate it in a scenario where filtering is implemented by the edge routers. We show that our solution reduces fake routing advertisements with minimal computational overhead.

- [C.72] *An empirical evaluation of QoME for NMP*, Konstantinos Tsioutas, George Xylomenos and Ioannis Doumanis, *Proceedings of the IFIP International Conference on New Technologies, Multimedia and Security (NTMS) 2021*.

Abstract: For Network Music Performance (NMP), low end-to-end delay is probably the most critical factor affecting the Quality of Musicians' Experience (QoME), as inflated delays inhibit the synchronization between musicians. Even though customized NMP tools can greatly reduce application-layer delays, generally available Internet connections impose considerable network-layer delays to NMP sessions in wide-area settings. In order to empirically assess the delay tolerance of NMP, multiple pairs of musicians were asked to perform in a controlled network environment, while we varied the end-to-end delay. The musicians were then asked to evaluate a number of metrics. The statistical analysis of the results indicates that the tolerance of actual musicians to delay is higher than previously thought, thus expanding the applicability of NMP to longer distances.

- [C.71] *In-network packet-level caching for error recovery in ICN*, Yannis Thomas, George Xylomenos and George C. Polyzos, *Proceedings of the International Workshop on Content Caching and Delivery in Wireless Networks (CCDWN) 2020*.

Abstract: In-network packet-level caching is one of the salient features of Information-Centric Networking (ICN) architectures, offering reduced communication latency, sender load, and network traffic. The literature on ICN caching is quite extensive, exploring in depth different aspects of caching, such as the positioning of storage resources in the network, the caching policies of individual caches, as well as the interaction of caching and routing. However, many researchers have questioned the value of in-network caching, arguing that the limited storage resources of ICN routers cannot eliminate a significant amount of redundant transmissions, hence the performance advantages of in-network caching may be similar to those of traditional edge-caches. Nevertheless, there is another potential use of in-network

caching that poses less challenging storage requirements: exploiting it to enhance the efficiency of error recovery by speeding-up retransmissions of lost packets. In this paper, we investigate the impact of in-network packet-level caching on transport-layer error recovery. We first sketch a retransmission-oriented packet-level caching scheme, highlighting the differences in the performance requirements compared to general-purpose caches. Then we model and assess its expected gains in terms of latency and network traffic reduction, and finally we provide a reality check of the approach based on current network requirements and hardware specifications.

- [C.70] *Quality of Musicians' Experience in Network Music Performance: A Subjective Evaluation*, Konstantinos Tsioutas, George Xylomenos, Ioannis Doumanis and Christos Angelou, *Proceedings of the Audio Engineering Society Convention 148*, 2020.

Abstract: In Network Music Performance (NMP), audio quality and audio delay are considered to be the most critical variables affecting the Quality of Musicians' Experience (QoME). In order to quantify the extent to which these parameters affect QoME, we executed a pilot study where eight musicians performed music in pairs in a controlled NMP setting and were asked to evaluate eight variables related to perception, while the end-to-end delay and the quality of the exchanged audio were varied. We present the design and execution of this experiment and discuss its results and their implications for the tolerance of musicians to increased delay and degraded audio quality.

- [C.69] *Keyword-based information retrieval for the WoT*, George Xylomenos, Evangelos Zafeiratos and Marios Prokopakis, *Proceedings of the ACM/IEEE Workshop on Hot Topics on Web of Things (HoTWoT) 2019*.

Abstract: The Internet of Things (IoT) is expected to contain huge numbers of "things" producing vast amounts of information. To turn these raw data to useful services for the Web of Things (WoT), we have previously proposed KIOT, a keyword-based scheme for gathering and processing IoT information using Information-Centric Networking (ICN) techniques. In KIOT data items, such as sensor readings, can be named with arbitrary sets of keywords, while users can retrieve all data items matching a desired set of keywords and (optionally) process them with arbitrary functions. In this paper we focus on a prototype implementation of the data retrieval part of KIOT. To maximize flexibility in diverse settings, our implementation automatically configures the network and its routing tables, allowing arbitrary sets of keywords to be used for both data items and queries. Our implementation can be used on any IoT device supporting Java, and is also available for large scale testing over emulated networks using Mininet.

- [C.68] *Named Functions at the Edge*, George Xylomenos, George Pavlou, Ioannis Psaras and Ioannis Karakonstantis, *Proceedings of the IEEE Symposium on Computers and Communications (ISCC) 2019*.

Abstract: As end-user and edge-network devices are becoming ever more powerful, they are producing ever increasing amounts of data. Pulling all this data into the cloud for processing is impossible, not only due to its enormous volume, but also due to the stringent latency requirements of many applications. Instead, we argue that end-user and edge-network devices should collectively form edge computing swarms and complement the cloud with their storage and processing resources. This shift from centralized to edge clouds has the potential to open new horizons for application development, supporting new low-latency services and, ultimately, creating new markets for storage and processing resources. To realize this vision, we propose Named Functions at the Edge (NFE), a platform where functions can i) be identified through a routable name, ii) be requested and moved (as data objects) to process data on demand at edge nodes, iii) pull raw or anonymized data from sensors and devices, iv) securely and privately return their results to the invoker and v) compensate each party for use of their data, storage, communication or computing resources via tracking and accountability mechanisms. We use an emergency evacuation application to motivate the need for NFE and demonstrate its potential.

- [C.67] *IP Over ICN Goes Live*, George Xylomenos, Yannis Thomas, Xenofon Vasilakos, Michael Georgiades, Alexander Phinikarides, Ioannis Doumanis, Stuart Porter, Dirk Trossen, Sebastian Robitzsch, Martin J. Reed, Mays Al-Naday, George Petropoulos, Konstantinos Katsaros, Maria-Evgenia Xezonaki and Janne Riihijärvi, *Proceedings of the European Conference on Networks and Communications (EUCNC) 2018*.

Abstract: Information-centric networking (ICN) has long been advocating for radical changes to the IP-based Internet. However, the upgrade challenges that this entails have hindered ICN adoption. To break this loop, the POINT project proposed a hybrid, IP-over-ICN, architecture: IP networks are preserved at the edge, connected to each other over an ICN core. This exploits the key benefits of ICN, enabling individual network operators to improve the performance of their IP-based services, without changing the rest of the Internet. We provide an overview of POINT and outline how it improves upon IP in terms of performance and resilience. Our focus is on the successful trial of the POINT prototype in a production network, where real users operated actual IP-based applications.

- [C.66] *IPTV Over ICN*, George Xylomenos, Alexander Phinikarides, Ioannis Doumanis, Xenofon Vasilakos, Yannis Thomas, Dirk Trossen, Michael Georgiades and Stuart Porter, *Proceedings of the Packet Video Workshop, 2018*.

Abstract: The efficient provision of IPTV services requires support for IP multicasting and IGMP snooping, limiting such services to single operator networks. Information-Centric Networking (ICN), with its native support for multicast seems ideal for such services, but it requires operators and users to overhaul their networks and applications. The POINT project has proposed a hybrid, IP-over-ICN, architecture, preserving IP devices and applications at the edge, but interconnecting them via an SDN-based ICN core. This allows individual operators to exploit the benefits of ICN, without expecting the rest of the Internet to change. In this paper, we first outline the POINT approach and show how it can handle multicast-based IPTV services in a more efficient and resilient manner than IP. We then describe a successful trial of the POINT prototype in a production network, where real users tested actual IPTV services over both IP and POINT under regular and exceptional conditions. Results from the trial show that the POINT prototype matched or improved upon the services offered via plain IP.

- [C.65] *A Keyword-based ICN-IoT Platform*, Onur Ascigil, Sergi Reñé, George Xylomenos, Ioannis Psaras and George Pavlou, *Proceedings of the ACM Information-Centric Networking (ICN) Conference 2017*.

Abstract: Information-Centric Networking (ICN) has been proposed as a promising solution for the Internet of Things (IoT), due to its focus on naming data, rather than endpoints, which can greatly simplify applications. The hierarchical naming of the Named-Data Networking (NDN) architecture can be used to name groups of data values, for example, all temperature sensors in a building. However, the use of a single naming hierarchy for all kinds of different applications is inflexible. Moreover, IoT data are typically retrieved from multiple sources at the same time, allowing applications to aggregate similar information items, something not natively supported by NDN. To this end, in this paper we propose (a) locating IoT data using (unordered) keywords combined with NDN names and (b) processing multiple such items at the edge of the network with arbitrary functions. We describe and evaluate three different strategies for retrieving data and placing the calculations in the edge IoT network, thus combining connectivity, storage and computing.

- [C.64] *Edge-ICN and its application to the Internet of Things*, Nikos Fotiou, Vasilios A. Siris, George Xylomenos, George C. Polyzos, Konstantinos V. Katsaros and George Petropoulos, *Proceedings of the IFIP Workshop on Information Centric Fog Computing (ICFC) 2017*.

Abstract: While research on Information-Centric Networking (ICN) flourishes, its adoption seems to be an elusive goal. In this paper we propose Edge-ICN: a novel approach for deploying ICN in a single large network, such as the network of an Internet Service Provider. Although Edge-ICN requires nothing beyond an SDN-based network supporting the OpenFlow protocol, with ICN-aware nodes only at the edges of the network, it still offers the same benefits as a clean-slate ICN architecture but without the deployment hassles. Moreover, by proxying legacy traffic and transparently forwarding it through the Edge-ICN nodes, all existing applications can operate smoothly, while offering significant advantages to applications such as native support for scalable anycast, multicast, and multi-source forwarding. In this context, we show how the proposed functionality at the edge of the network can specifically benefit CoAP-based IoT applications. Our measurements show that Edge-ICN induces on average the same control plane overhead for name resolution as a centralized approach, while also enabling IoT applications to build on anycast, multicast, and multi-source forwarding primitives.

- [C.63] *Securing Information-Centric Networking without negating Middleboxes*, Nikos Fotiou, George Xylomenos and George C. Polyzos, *Proceedings of the IFIP International Conference on New Technolo-*

gies, *Multimedia and Security (NTMS) 2016*.

Abstract: Information-Centric Networking is a promising networking paradigm that overcomes many of the limitations of current networking architectures. Various research efforts investigate solutions for securing ICN. Nevertheless, most of these solutions relax security requirements in favor of network performance. In particular, they weaken end-user privacy and the architecture's tolerance to security breaches in order to support middleboxes that offer services such as caching and content replication. In this paper, we adapt TLS, a widely used security standard, to an ICN context. We design solutions that allow session reuse and migration among multiple stakeholders and we propose an extension that allows authorized middleboxes to lawfully and transparently intercept secured communications.

- [C.62] *RT-SENMOS: Sink-Driven Congestion and Error Control for Sensor Networks*, Charilaos Stais, George Xylomenos and Evangelos Zafeiratos, *Proceedings of the IFIP International Conference on New Technologies, Multimedia and Security (NTMS) 2016*.

Abstract: A very common problem in the sensor networks area, where large numbers of sensors transmit data to a single sink node, is reliable transport. As sensors become cheaper and more powerful, sensor networks face the challenge of controlling all these devices in a efficient manner, which naturally leads to a sink-driven transport scheme. In this paper we propose RT-SENMOS, a sink-driven reliable transport protocol for sensor networks which places the responsibility for transmission rate allocation, congestion avoidance and error control to the sink. RT-SENMOS is intended to be integrated within a specific application which will set its operating parameters depending on its needs, operating on top of UDP/IP to avoid the need for kernel/superuser access. As RT-SENMOS is fully sink-controlled, it enables the use of simple and inexpensive fixed sensors, which offload all protocol intelligence to a more expensive but reusable sink. We present the design of the protocol, comparing it with a similar approach, called RCRT, and evaluate its performance using a real implementation.

- [C.61] *Multi-Flow Congestion Control with Network Assistance*, Yannis Thomas, George Xylomenos, Christos Tsilopoulos and George C. Polyzos, *Proceedings of the IFIP Networking 2016 Conference*.

Abstract: A well-known technique for enhancing the performance and stability of content distribution is the use of multiple dissemination flows. Multipath TCP (MPTCP), the most popular multiflow protocol on the Internet, allows receivers to exploit multiple paths towards a single sender. Nevertheless, MPTCP cannot fully exploit the potential gains of multipath connectivity, as it must fairly share resources with (single-flow) TCP, without a clear understanding of whether the available paths do share any bottleneck links. In this paper, we introduce a hybrid congestion control algorithm for multisource and multipath transport that enables higher bandwidth utilization compared to MPTCP, while remaining friendly to TCP-like flows. Our solution employs (i) an in-network module that offers essential topological information and (ii) Normalized Multiflow Congestion Control (NMCC), a novel end-to-end congestion control algorithm. While NMCC is architecture-independent and the in-network module can be adapted for Multi-Protocol Label Switching (MPLS) or Software Defined Networks (SDNs), our prototype was implemented on the Publish-Subscribe Internetworking (PSI) architecture, which offers centralized path formation and source routing. Using an actual protocol implementation deployed on our test-bed, we provide experimental results which validate the effectiveness of our design in terms of performance, adaptation to shifting network conditions and friendliness to other flows.

- [C.60] *Protecting medical data stored in public Clouds*, Nikos Fotiou and George Xylomenos, *Proceedings of the International Conference on Information and Communication Technologies for Ageing Well and E-Health (ICT4AWE) 2016*.

Abstract: Public Clouds offer a convenient way for storing and sharing large amounts of medical data. Nevertheless, using a shared infrastructure raises significant security and privacy concerns. Even if the data are encrypted, the data owner should share some information with the Cloud provider, in order to enable the latter to perform access control; given the high sensitivity of medical data, even such limited information may jeopardize end-user privacy. In this paper we employ an access control delegation scheme to enable the users themselves to perform access control on their data, which are stored in a public Cloud. To selectively provide access to these data without sacrificing their confidentiality we rely on encryption: our system encrypts data before storing them in the Cloud and applies proxy re-encryption so as to encrypt data separately for each (authorized) user.

- [C.59] *Object-oriented Packet Caching for ICN*, Yannis Thomas, George Xylomenos, Christos Tsilopoulos and George C. Polyzos, *Proceedings of the ACM Information-Centric Networking (ICN) Conference 2015*.

Abstract: One of the most discussed features offered by Information-centric Networking (ICN) architectures is the ability to support packet-level caching at every node in the network. By individually naming each packet, ICN allows routers to turn their queueing buffers into packet caches, thus exploiting the network's existing storage resources. However, the performance of packet caching at commodity routers is restricted by the small capacity of their SRAM, which holds the index for the packets stored at the, slower, DRAM. We therefore propose Object-oriented Packet Caching (OPC), a novel caching scheme that overcomes the SRAM bottleneck, by combining object-level indexing in the SRAM with packet-level storage in the DRAM. We implemented OPC and experimentally evaluated it over various cache placement policies, showing that it can enhance the impact of ICN packet-level caching, reducing both network and server load.

- [C.58] *Evaluating the Impact of Network I/O on Ultra-Low Delay Packet Switching*, George Baltas and George Xylomenos, *Proceedings of the IEEE Symposium on Computers and Communications (ISCC) 2015*.

Abstract: Low latency is a crucial requirement for demanding conferencing applications, such as Networked Music Performance (NMP), the collaboration of musicians in real time. Modern conferencing systems employ a Selective Forwarding Unit (SFU) to transparently duplicate and forward media streams between participants. Since an SFU does not process the media streams, so as to reduce delay, its latency is mainly determined by the underlying network I/O mechanism that moves packets to/from the network hardware and user space. Such mechanisms are usually based on POSIX sockets, which were not designed for high performance networking. We designed and implemented pktswitch, a minimal, socket-based SFU and measured its performance. We then modified pktswitch to employ netmap, a framework for fast packet I/O, to overcome the performance bottlenecks imposed by the socket-based design. The modified implementation handles packets in user space, with minimal kernel interaction. We describe and contrast the two implementations and then compare their performance in terms of packet processing overhead and delay. Our results show that the netmap-based implementation reduces packet processing overhead by 76% and delay by 89% compared to the socket-based implementation, thus allowing an SFU to host much higher loads (e.g. more users with more media streams and higher bit rates) without introducing delays.

- [C.57] *RT-SENMOS: Reliable Transport for Sensor Networks with Mobile Sinks*, Charilaos Stais and George Xylomenos, *Proceedings of the IEEE Workshop on Distributed Mobile Systems & Services (DMSS) 2015*.

Abstract: Gathering information efficiently from a, possibly fragmented, sensor network presents a serious problem in disaster recovery applications. Unless a transmission control mechanism exists, a sink can be flooded with information sent by sensors, or sensor transmissions may be lost on their way to the sink. Consequently, there is a necessity for a reliable protocol that automatically and speedily adapts to losses, congestion and network changes due to sink mobility. This paper describes RT-SENMOS, a reliable transport protocol for controlling sensor transmissions based on sink-assigned rates. The sink decides how to share the available bandwidth among the sensors and also determines the reliability to be achieved in each case. Our protocol operates on top of UDP/IP, therefore it can be directly integrated into a disaster recovery application that will set its parameters depending on the situation. Moreover, as it is fully sink-controlled, it enables the use of simple and inexpensive fixed sensors, which offload all protocol intelligence to a more expensive but reusable mobile sink. We present the design of the protocol, comparing it with similar approaches, and evaluate its performance using a real implementation.

- [C.56] *Exploiting Path Diversity for Networked Music Performance in the Publish Subscribe Internet*, Yannis Thomas, George Xylomenos and George C. Polyzos, *Proceedings of the International Conference on Information, Intelligence, Systems and Applications (IISA) 2015*.

Abstract: The high bandwidth and low delay requirements of Networked Music Performance (NMP) often lead to doubts about its feasibility in wide area networks. Multipath routing with Quality of Service (QoS) constraints is one way to meet these requirements, as it offers bandwidth aggregation while keeping latency low. However, multipath routing with QoS constraints is not inherently supported in

IP networks, requiring complicated extensions to network operation. In contrast, the Publish-Subscribe Internet (PSI) architecture, an Information-Centric Networking (ICN) approach, natively supports multicast, centralized path selection and source routing, all the main elements needed for multipath QoS routing. In this paper, we discuss the implementation of multipath QoS routing in PSI and experimentally evaluate its benefits for NMP services.

- [C.55] *The MusiNet project: Addressing the challenges in Networked Music Performance systems*, D. Akoumi-anakis, C. Alexandraki, V. Alexiou, C. Anagnostopoulou, A. Eleftheriadis, Y. Mastorakis, A. Modas, A. Mouchtaris, D. Pavlidi, G. C. Polyzos, P. Tsakalides, G. Xylomenos and P. Zervas, *Proceedings of the International Conference on Information, Intelligence, Systems and Applications (IISA) 2015*.

Abstract: This paper presents the progress in the MusiNet research project, which aims to provide a comprehensive architecture and a prototype implementation of a Networked Music Performance (NMP) system. We describe the Musinet client and server components, and the different approaches followed in our research effort in order to culminate in the most appropriate scheme in terms of delay and quality for the audio and video streams involved. We also describe the MusiNet user interface, which allows an integrated communication between the participants and the proposed NMP system.

- [C.54] *IP Over ICN - The Better IP?*, Dirk Trossen, Martin J. Reed, Janne Riihijärvi, Michael Georgiades, Nikos Fotiou and George Xylomenos, *Proceedings of the European Conference on Networks and Communications (EUCNC) 2015*.

Abstract: This paper presents a proposition for information-centric networking (ICN) that lies outside the typical trajectory of aiming for a wholesale replacement of IP as the internetworking layer of the Internet. Instead, we propose that a careful exploitation of key ICN benefits, expanding previously funded ICN efforts, will enable individual operators to improve the performance of their IP-based services along many dimensions. Alongside the main motivation for our work, we present an early strawman architecture for such an IP-over-ICN proposition, which will ultimately be implemented and trialed in a recently started H2020 research effort.

- [C.53] *On the Inter-domain Scalability of Route-by-Name Information-Centric Network Architectures*, Konstantinos V. Katsaros, Xenofon Vasilakos, Timothy Okwii, George Xylomenos, George Pavlou and George C. Polyzos, *Proceedings of the IFIP Networking 2015 Conference*.

Abstract: Name resolution is at the heart of Information-Centric Networking (ICN), where names are used to both identify information and/or services, and to guide routing and forwarding inside the network. The ICN focus on information, rather than hosts, raises significant concerns regarding the scalability of the required Name Resolution System (NRS), especially when considering global scale, inter-domain deployments. In the route-by-name approach to NRS construction, name resolution and the corresponding state follow the routing infrastructure of the underlying inter-domain network. The scalability of the resulting NRS is therefore strongly related to the topological and routing characteristics of the network. However, past work has largely neglected this aspect. In this paper, we present a detailed investigation and comparison of the scalability properties of two route-by-name inter-domain NRS designs, namely, DONA and CURLING. Based on both real, full-scale inter-domain topology traces and synthetic, scaled-down topologies, our work quantifies a series of important scalability-related performance aspects, including the distribution of name-resolution state across the Internet topology and the associated processing and signaling overheads. We show that by avoiding DONA's exchange of state across peering links, CURLING results in deployment costs proportional to the total number of downstream customers of each Autonomous System. This translates to a 62-fold global state size reduction, at the expense of a 2.78-fold increase in lookup processing load, making CURLING a feasible approach to ICN name resolution.

- [C.52] *Towards improving the efficiency of ICN packet-caches*, Yannis Thomas and George Xylomenos, *Proceedings of the International Workshop on Quality, Reliability and Security in Information-Centric Networking (Q-ICN) 2014*.

Abstract: Abstract: In-network packet-level caching is one of the most promising features offered by Information-centric Networking (ICN) architectures. In ICN, routers can use their queueing buffers as temporal storage units, thus allowing on-path caching by exploiting the network's storage resources.

Packet-caches can be highly beneficial for content delivery, but they are also known to have three significant weaknesses: (i) packetgranularity produces huge cache indexes, (ii) Zipf-like content popularity penalizes the hit-ratio at core nodes and (iii) any discontinuity in the stored packets disrupts RTT-based congestion control. This paper presents OPC, a novel caching management strategy designed to support wire-speed in-network caching, while dealing with the above problems. OPC works at the objectlevel, thus reducing indexing requirements, is destined for access routers, thus avoiding the small hit-ratios of caches at core-nodes, and stores contiguous groups of packets, thus easing RTT-based congestion control.

- [C.51] *I-CAN: Information-Centric Future Mobile and Wireless Access Networks*, George C. Polyzos, Vasilios A. Siris, George Xylomenos, Giannis F. Marias and Stavros Toumpis, *Proceedings of the International Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness (Q-SHINE) 2014*.

Abstract: This short paper describes the objectives and initial results of project I-CAN: Information-Centric Future Mobile and Wireless Access Networks. I-CAN seeks to radically advance the integration of cellular and wireless access technologies by developing and evaluating architectures and procedures for future access networks based on Information-Centric Networking (ICN).

- [C.50] *Ultra Low Delay Switching for Networked Music Performance*, George Baltas and George Xylomenos, *Proceedings of the International Conference on Information, Intelligence, Systems and Applications (IISA) 2014*.

Abstract: Low latency is essential for videoconferencing applications such as presence and collaboration between remote participants. In modern videoconferencing systems, the Selective Forwarding Unit (SFU) has the role of transparently duplicating and forwarding media streams between participants, hence it must be able to process large volumes of incoming packets at very high rates. SFU performance is heavily affected by the network I/O mechanisms employed to move packets from the Network Interface Card (NIC) to user space, and then move the copies back to the NIC. Traditional mechanisms, such as POSIX sockets, are not designed for high performance networking and prove to be a major bottleneck in such scenarios, by increasing packet latency and undermining the SFU's scalability. In this paper, we present a novel SFU platform which was designed to handle the ultra-low latency requirements of Networked Music Performance (NMP) applications, that is, the collaboration of musicians in real time. We implement a prototype SFU based on POSIX sockets and outline its performance bottlenecks. To overcome them, we turn to the netmap framework for fast packet I/O, which provides direct but safe access to the NIC buffers. We argue that ultralow latency videoconferencing is a natural application for netmap and thus design and implement a netmap-based SFU.

- [C.49] *The MusiNet project: Towards unraveling the full potential of Networked Music Performance systems*, D. Akoumianakis, C. Alexandraki, V. Alexiou, C. Anagnostopoulou, A. Eleftheriadis, V. Lalioti, A. Mouchtaris, D. Pavlidi, G. C. Polyzos, P. Tsakalides, G. Xylomenos and P. Zervas, *Proceedings of the International Conference on Information, Intelligence, Systems and Applications (IISA) 2014*.

Abstract: The MusiNet research project aims to provide a comprehensive architecture and a prototype implementation of a complete Networked Music Performance (NMP) system. In this paper we describe the current status of the project, focusing on critical decisions regarding the system's architecture and specifications, the low delay audio and video coding techniques to be employed, the media relay design, and the synchronous and asynchronous collaboration algorithms to be adopted.

- [C.48] *Sink Controlled Reliable Transport for Disaster Recovery*, Charilaos Stais, George Xylomenos and Giannis F. Marias, *Proceedings of the International Conference on Pervasive Technologies Related to Assistive Environments (PETRA) 2014*.

Abstract: We present a reliable transport layer protocol for sensor networks, targeting disaster recovery applications where human or robotic rescuers try to gather information from a possibly fragmented sensor network by moving through the disaster area. The mobility of the information sink means that the protocol must quickly adapt to a constantly changing view of the network, where connections and disconnections are the norm. Our protocol is purely sink driven, that is, the sink controls congestion by rate limiting the sensors, choosing how to assign the available bandwidth to different sensor types and deciding on the level of reliability to be achieved. In addition, our protocol operates at the application

layer with minimal requirements from lower layers, allowing its integration with a disaster recovery application that will set its parameters depending on the disaster scenario. As a result, our protocol allows simple and inexpensive fixed sensors to be combined with expensive but reusable mobile equipment for disaster recovery purposes.

- [C.47] *Accelerating File Downloads in Publish Subscribe Internetworking with Multisource and Multipath Transfers*, Yannis Thomas, Christos Tsilopoulos, George Xylomenos and George C. Polyzos, *Proceedings of the World Telecommunications Congress (WTC) 2014*.

Abstract: We present mmFTP, a file transfer protocol for the Publish Subscribe Internetworking (PSI) architecture, which follows the Information-Centric Network (ICN) paradigm. mmFTP is designed to utilize diverse in-network resources: (i) it is receiver-driven, thus supporting on-path caching, (ii) it downloads files from multiple sources, thus utilizing off-path caching and (iii) it offers multipath transfers, thus exploiting path diversity and assisting network load-balancing. mmFTP combines these features into a single framework without complicating network operation. This is achieved by exploiting the functional organization of the PSI architecture which - among other aspects - separates routing from packet forwarding, delegates routing control to a logically centralized module and employs an explicit routing scheme for packet forwarding. In this paper we introduce the basic operation of mmFTP and present preliminary experimental performance results from a prototype implementation deployed in the PlanetLab testbed.

- [C.46] *Access control delegation for the Cloud*, Nikos Fotiou, Apostolis Machas, George C. Polyzos and George Xylomenos, *Proceedings of the IEEE CrossCloud Workshop 2014*.

Abstract: Cloud computing has become the focus of attention in the industry, from the point of view of both providers and customers, as well as researchers. However, security concerns still impede the widespread adoption of this technology. Most enterprises are particularly worried about the lack of control over their outsourced data since the authentication and authorization systems of Cloud providers are generic and they cannot be easily adapted to the requirements of each individual enterprise. An adaptation process requires the creation of complex protocols, often leading to security problems and lock-in conditions. In this paper we present the design of a lightweight solution that overcomes these problems. We have implemented and incorporated this solution in a popular open-source Cloud stack: OpenStack. Our solution eliminates the need for developing complex adaptation protocols, offers data owners the flexibility to switch among Cloud providers, or use multiple, different Cloud providers concurrently, and enhances end-user privacy.

- [C.45] *Reducing Forwarding State in Content-Centric Networks with Semi-Stateless Forwarding*, Christos Tsilopoulos, George Xylomenos and Yannis Thomas, *Proceedings of the IEEE INFOCOM, 2014*.

Abstract: Routers in the Content-Centric Networking (CCN) architecture maintain state for all pending content requests, so as to be able to later return the corresponding content. By employing stateful forwarding, CCN supports native multicast, enhances security and enables adaptive forwarding, at the cost of excessive forwarding state that raises scalability concerns. We propose a semi-stateless forwarding scheme in which, instead of tracking each request at every on-path router, requests are tracked at every d hops. At intermediate hops, requests gather reverse path information, which is later used to deliver responses between routers using Bloom filter-based stateless forwarding. Our approach effectively reduces forwarding state, while preserving the advantages of CCN forwarding. Evaluation results over realistic ISP topologies show that our approach reduces forwarding state by 54%-70% in unicast delivery, without any bandwidth penalties, while in multicast delivery it reduces forwarding state by 34%-55% at the expense of 6%-13% in bandwidth overhead.

- [C.44] *Are Information-Centric Networks Video-Ready?*, Christos Tsilopoulos, George Xylomenos and George C. Polyzos, *Proceedings of the Packet Video Workshop, 2013*.

Abstract: Video constitutes the majority of all Internet traffic and its share is expected to grow. Any future Internet architecture with a chance at success should provide some tangible benefits for video applications. Information-Centric Networking (ICN) architectures were designed with the specific goal of improving content distribution on the Internet; thus, this paper attempts to answer the obvious question: is ICN appropriate and ready for video traffic and, if not, what is missing or should be modified? To

this end, we consider two different ICN architectures, *Content-Centric Networking* (CCN) and *Publish-Subscribe Internetworking* (PSI), and examine their applicability to Video on Demand and Live Video Streaming applications. Our goal is to clarify what ICN already does well for video, what it still needs to do, and, most importantly, what it could or should do differently.

- [C.43] *Scaling Bloom filter-based multicast via filter switching*, Christos Tsilopoulos and George Xylomenos, *Proceedings of the IEEE Symposium on Computers and Communications (ISCC) 2013*.
Abstract: Stateless multicast forwarding with in-packet Bloom filters (iBF) has recently been proposed as a highly scalable way for supporting a large number of multicast groups. However, iBF multicast generates redundant traffic due to false positive forwarding decisions and it also scales poorly with multicast group size. In this paper we investigate scaling iBF multicast to arbitrary multicast group sizes, by partially sacrificing the networks fully stateless operation. We propose a switched-iBF multicast scheme that places multicast forwarding state at a few network nodes, so as to minimize redundant traffic regardless of the group size. We evaluate the scheme through simulations and find that switched-iBF multicast can scale to any group size while keeping redundant traffic below 1%-4% at the (minimal) cost of placing state at no more than 0.5%-2.5% of network nodes. We also compare the state requirements of switched-iBF multicast against other multicast schemes. Our evaluation shows that switched-iBF multicast achieves a tremendous reduction of multicast state in the range of 87%-99.6%. Hence, even though the system is no more fully stateless, it remains far more scalable than other approaches.
- [C.42] *Towards an Error Control Scheme for a Publish/Subscribe Network*, Charilaos Stais, Alexios Voulime-neas and George Xylomenos, *Proceedings of the IEEE International Conference on Communications (ICC) 2013*.
Abstract: Many proposals for the next generation of the Internet suggest moving from an end-point oriented to an informationcentric oriented architecture. Many of these proposals are based on the publish/subscribe paradigm, which lends itself naturally to native multicast support, a key factor for efficient content distribution. However, the design of efficient reliable transport protocols for multicast is a largely open problem, due to the problem of feedback implosion towards the sender as group size grows. In this paper we propose a hierarchical retransmissionbased error control scheme for a native publish/subscribe internetwork. We compare our protocol with similar approaches proposed for IP multicast and evaluate its performance against IP multicast with unicast-based error control.
- [C.41] *Networked Music Performance over Information-Centric Networks*, Charilaos Stais, Yannis Thomas, George Xylomenos and Christos Tsilopoulos, *Proceedings of the IEEE Workshop on Immersive & Interactive Multimedia Communications (IIMC) 2013*.
Abstract: Information-centric networking (ICN) constitutes an alternative to the conventional, IP-based, internetworking, with information itself being identified rather than the host where it resides. This approach introduces powerful tools and operations for content delivery, such as native support for multicast. Exploiting this native multicast capability is a very promising approach for multimedia applications such as Networked Music Performance (NMP), where a set of musicians located in different places wish to perform together in real time. While conferencing applications traditionally rely on a Multipoint Conferencing Unit (MCU) that receives media streams from each participant and then retransmits a mixed stream to each one, in NMP we would prefer direct communication between the participants, so as to reduce transmission delays and allow each participant to mix the incoming media streams in the desired manner. In this paper we introduce an ICN-based NMP application exploiting native multicast, and compare its performance with both MCU and non-MCU NMP variants, using both unicast and multicast.
- [C.40] *Efficient Real-time Information Delivery in Future Internet Publish-Subscribe Networks*, Christos Tsilopoulos, Ioannis Gasparis, George Xylomenos and George C. Polyzos, *Proceedings of the International Conference on Computing, Networking and Communications (ICNC) 2013*.
Abstract: In this paper we explain how efficient delivery of real-time information can be supported in the Publish Subscribe Internet (PSI), a network architecture proposal for the Future Internet. PSI departs from IP thinking with respect to the core abstractions made and the functional organization of the system. PSI places information at the heart of the network layer and decouples the forwarding, path formation and topology management functionalities. This design approach can be highly beneficial for

real-time communications, as it enables the network to apply sophisticated mechanisms for multicast tree construction, such as delivery of information over optimal (minimum cost) Steiner trees. Initial experiments with a proof-of-concept implementation of PSI indicate the feasibility of realizing such optimization policies. Our results show that significant bandwidth savings can be achieved at the cost of small, un-noticeable to the end-users, delays in flow establishment.

- [C.39] *Cloud computing for global name-resolution in information-centric networks*, Xenofon Vasilakos, Konstantinos Katsaros and George Xylomenos, *Proceedings of the IEEE Network Cloud Computing and Applications Symposium (NCCA) 2012*.

Abstract: Information-Centric Networking (ICN) is a novel paradigm for future Internet architectures. It exploits the current trend in Internet usage which mostly involves information dissemination. ICN architectures based on the publish/subscribe model use names for information in order to route requests and data, as well as to facilitate in-network caching, anycasting and multicasting for efficient content delivery. However, the number of named information objects is expected to be huge in the future Internet, raising serious concerns with respect to a global-scale deployment of ICN. Routing and forwarding will require vast amounts of state, which pushes storage, maintenance and processing demands to the limit. In this paper we discuss the feasibility of deploying the Data Oriented Networking Architecture (DONA) by leveraging cloud computing facilities. We identify the exact scalability concerns for DONA based on simulations over a realistic model of the current Internet topology and find that registrations for information objects lead to a state explosion. For this reason, we then discuss how cloud facilities can assist DONA deployment, focusing on various options for deploying DONA in the cloud and their suitability for different areas of the inter-network.

- [C.38] *Realistic Media Streaming over BitTorrent*, Charilaos Stais and George Xylomenos, *Proceedings of the Future Network and Mobile Summit 2012 Conference*.

Abstract: While BitTorrent was originally conceived as a Peer-to-Peer file exchange protocol, it has proved extremely successful for asynchronous content distribution, allowing content sources to support huge numbers of users with a modest amount of bandwidth. This has prompted many researchers to study the possibility of using BitTorrent to support real-time media streaming. In this paper we present a comparison of three proposed adaptations to BitTorrent for media streaming, using our detailed packet-level BitTorrent simulator. Unlike previous evaluations which assumed that the streaming media player would drop data that did not arrive on time, in this paper we use a more realistic model where the player stalls when data are not available, thus placing emphasis on delays rather than losses. Our experiments indicate that under this, more realistic, evaluation model, user level performance can be quite reasonable.

- [C.37] *On Inter-domain Name Resolution for Information-Centric Networks*, Konstantinos Katsaros, Nikos Fotiou, Xenofon Vasilakos, Christopher Ververidis, Christos Tsilopoulos, George Xylomenos and George C. Polyzos, *Proceedings of the IFIP Networking 2012 Conference*.

Abstract: Information-centric networking (ICN) is a paradigm that aims to better reflect current Internet usage patterns by focusing on information, rather than on hosts. One of the most critical ICN functionalities is the efficient resolution/location of information objects i.e., name resolution. The vast size of the information object namespace calls for a highly scalable and efficient name resolution approach. Currently proposed solutions either rely on a DHT structure, thus ensuring load balancing and scalability at the cost of inefficient routing, or on hierarchical structures, thus preserving routing efficiency at the cost of limited scalability. In this paper, we study in detail the tradeoff between state/signaling overhead versus routing efficiency for a generic name-resolution system based on a novel DHT scheme with enhanced routing properties, and compare it to DONA, an ICN architecture based on hierarchical resolution and routing.

- [C.36] *GlobeTraff: a traffic workload generator for the performance evaluation of future Internet architectures*, Konstantinos Katsaros, George Xylomenos and George C. Polyzos, *Proceedings of the IFIP International Conference on New Technologies, Multimedia and Security (NTMS) 2012*.

Abstract: Having identified important limitations of the current Internet architecture, several research initiatives have engaged in the design of new architectures for the Internet of the Future. New features

and protocols are designed, and in many cases, a *clean slate* approach is followed, advocating the replacement of almost the entire current protocol stack. In order to reliably evaluate the performance of such emerging protocols and architectures, we need a realistic evaluation framework reflecting current and forecasted traffic patterns. This is especially important for the emerging information-centric paradigm, where in-network caching plays an important role and network performance is heavily dependent on traffic characteristics. To this end, we have designed and implemented *GlobeTraff*, a traffic workload generator for the creation of synthetic traffic mixes. *GlobeTraff* supports several application traffic types, generated via models in the recent research literature, allowing the detailed parameterization of the respective models and the composition of the resulting traffic mix.

- [C.35] *Supporting Diverse Traffic Types in Information Centric Networks*, Christos Tsilopoulos and George Xylomenos, *Proceedings of the ACM SIGCOMM Information Centric Networking (ICN) Workshop 2011*.

Abstract: In this paper we focus on the issue of transferring diverse kinds of information through information-centric networks (ICNs). We argue that the one request per packet mode of operation suggested in the early development of ICN applications is not a good fit for some types of traffic, such as media streams and real-time notifications. To efficiently deliver all kinds of information, we argue that an ICN should not only identify information by its name, it should also be aware of the nature of its traffic. We classify information traffic types based on two characteristics: a) reliable vs. unreliable transfer and b) real-time vs. on-demand delivery. The combination of these two characteristics leads to three broad categories: a) channels, b) on-demand documents and c) real-time documents. To handle all traffic types, we propose two extensions to the CCN architecture: Persistent Interests and Reliable Notifications. We describe how these additions, together with a careful selection of information names, can efficiently support these three categories of information traffic types.

- [C.34] *A comparison of streaming extensions to BitTorrent*, Charilaos Stais, George Xylomenos and Angelos Archodovassilis, *Proceedings of the IEEE Symposium on Computers and Communications (ISCC) 2011*.

Abstract: Peer-to-peer (P2P) protocols have been widely accepted by users and operators alike as efficient mechanisms for non real-time content distribution. It is therefore reasonable to extend these protocols to also handle more demanding applications, such as multimedia streaming. Many researchers have proposed modifications to the well known BitTorrent protocol in order to adapt it to the needs of such applications. In this paper we present findings from our experiments with three proposals for the integration of multimedia streaming into BitTorrent: the fixed-size window approach, the high-priority set approach and the stretching window approach. We evaluate these proposals under identical circumstances using our detailed packet-level BitTorrent simulator, showing that while all approaches are reasonably capable of supporting multimedia streaming, their different design choices have a pronounced effect on their streaming performance.

- [C.33] *VoPSI: Voice over a Publish-Subscribe Internetwork*, Charilaos Stais, Dimitris Diamantis, Christina Aretha and George Xylomenos, *Proceedings of the Future Network and Mobile Summit 2011 Conference*.

Abstract: Information-centric networking constitutes an alternative to the conventional, IP-based inter-networking, with information being identified rather than the host where it resides (which is the case for IP networking). This approach appears to be very promising for the next generation Internet. However, many challenges and critical issues remain to be addressed, associated with the range of applications that can be supported by the new architectures. Specifically, it is unclear whether informationcentric networking abstractions can support conversational applications which are very important in the Internet, and even more important in other telecommunication networks. In light of those reflections, we present the design, prototype implementation and performance evaluation of a simple voice application for the Publish-Subscribe Internet architecture developed in PSIRP and now being further refined in PURSUIT, two EU FP7 research projects on clean slate Future Internet design.

- [C.32] *Bootstrapping a Publish/Subscribe Information Centric Network*, Christos Tsilopoulos, Dimitris Makris and George Xylomenos, *Proceedings of the Future Network and Mobile Summit 2011 Conference*.

Abstract: In this paper, we focus on the bootstrap operation of a publish/subscribe information centric network. We consider a set of interconnected network nodes and describe how they organize themselves into a fully functional network by publishing and subscribing to control plane information. Network bootstrap includes establishing point-to-point communication between network elements, exchanging topological information and setting up the rendezvous system. We showcase the network's publish/subscribe service with a complete example.

- [C.31] *DHTbd: A Reliable Block-based Storage System for High Performance Clusters*, George Parisis, George Xylomenos and Theodore Apostolopoulos, *Proceedings of the IEEE/ACM CCGrid 2011*.

Abstract: Large, reliable and efficient storage systems are becoming increasingly important in enterprise environments. Our research in storage system design is oriented towards the exploitation of commodity hardware for building a high performance, resilient and scalable storage system. We present the design and implementation of DHTbd, a general purpose decentralized storage system where storage nodes support a distributed hash table based interface and clients are implemented as in-kernel device drivers. DHTbd, unlike most storage systems proposed to date, is implemented at the block device level of the I/O stack, a simple yet efficient design. The experimental evaluation of the proposed system demonstrates its very good I/O performance, its ability to scale to large clusters, as well as its robustness, even when massive failures occur.

- [C.30] *Design and Evaluation of a Socket Emulator for Publish/Subscribe Networks*, George Xylomenos and Blerim Cici, *Proceedings of the Future Internet Symposium 2010*.

Abstract: In order for a Future Internet architecture to be globally deployed, it must ensure that existing applications will continue to operate efficiently on top of it. As part of the Publish Subscribe Internet Routing Paradigm (PSIRP) project, we have explored various options for making endpoint centric applications based on the Sockets Application Programming Interface (API) compatible with the information centric PSIRP prototype implementation. We developed an emulator that mediates between the client/server socket calls and the publish/subscribe PSIRP calls, transforming the exchange of packets to distribution of publications. To assess the overhead of our emulator, we measure the execution time of a simple file transfer application in native socket mode, in emulated socket mode and in native publish/subscribe mode.

- [C.29] *On the incremental deployment of overlay information centric networks*, Konstantinos Katsaros, Charilaos Stais, George Xylomenos and George C. Polyzos, *Proceedings of the Future Network and Mobile Summit 2010 Conference*.

Abstract: It has long been realized that the proliferation of information-centric applications and services must be reflected in a corresponding shift of the underlying Internet architecture. Even though users increasingly focus on the desired information, the underlying network still focuses on the end-points providing/consuming this information and in many cases this mismatch has resulted in an inefficient utilization of network resources, as demonstrated by peer-to-peer (P2P) and file sharing applications. In view of this situation, many research projects have focused on the investigation of alternative networking models centered around information. However, less attention has been paid to the transition process from the current end host centric model to an information centric one. In this paper, we propose an overlay multicast-enabled, publish-subscribe architecture and focus on its gradual deployment both inside administrative domain boundaries as well as across the Internet. Our simulation results demonstrate the benefits for individual network operators as they gradually adopt our new networking model, and shed further light on the extent of deployment required within an administrative domain in order for our approach to perform optimally.

- [C.28] *A hybrid overlay multicast and caching scheme for information-centric networking*, Konstantinos Katsaros, George Xylomenos and George C. Polyzos, *Proceedings of the Global Internet 2010 Symposium*.

Abstract: It has long been realized that the use of the Internet has moved away from its original end-host centric model. The vast majority of services and applications is nowadays focused on information itself rather on the end-points providing/consuming it. However, the underlying network architecture still focuses on enabling the communication between pairs of end-hosts, leading to a series of problems, such as the inefficient utilization of network resources, demonstrated by the proliferation of peer-to-peer (P2P) and file sharing applications. In essence, the prevailing end-to-end nature of the current Internet

architecture prohibits network operators from controlling the traffic carried by their networks, leaving this control entirely to end users and their applications. In this paper, we investigate the potential benefits of MultiCache, an overlay network architecture aiming at handing control back to network operators. In MultiCache proxy overlay routers enable the delivery of data either via direct multicast, or via multicast fed caches residing at the leaves of multicast delivery trees. We study crucial aspects of our architecture, paying special attention to the properties of our distributed caching scheme, and investigate the feasibility of a progressive deployment of the proposed functionality over the existing Internet.

- [C.27] *MultiCache: an incrementally deployable overlay architecture for information-centric networking*, Konstantinos Katsaros, George Xylomenos and George C. Polyzos, *Proceedings of the IEEE INFO-COM 2010 Work In Progress Symposium*.

Abstract: It has been long realized that the Internet is evolving from a network connecting pairs of end hosts to a substrate for information dissemination. While this shift towards information centric networking has been clearly demonstrated by the proliferation of file sharing (e.g., BitTorrent) and content delivery (e.g., YouTube) applications, it has not been followed by a corresponding shift in network architecture. As a result, even though such applications are attractive to both content providers, due to their lower bandwidth requirements, and to end users, due to their reduced download times, they plague the underlying network with redundant packet transmissions, a significant part of which takes place over costly inter-domain links. In essence, the end-to-end nature of the current Internet architecture prevents network operators from controlling the traffic carried by their networks, delegating such control to end users and their applications. In this paper, we propose MultiCache, an information centric architecture aiming at the efficient use of network resources that is based on two primitives: multicast and caching. To this end, we revisit overlay multicast as a means for content delivery, and take advantage of multicast forwarding information to locate, in an anycast fashion, nearby caches that have been themselves fed by multicast sessions. Our architecture is evaluated against a widespread file sharing application (BitTorrent) with respect to both network resource savings and end user experience.

- [C.26] *Monitoring and Modeling Simple Everyday Activities of the Elderly at Home*, George Papamathiakis, George C. Polyzos and George Xylomenos, *Proceedings of the IEEE Consumer Communications and Networking Conference (CCNC) 2010*.

Abstract: We present our work on a sensor-based smart system that is automatically trained to recognize the activities of individuals in their home. In this paper we present and analyze a method for recognizing indoor the everyday activities of a monitored individual. This method is based on the data mining technique of association rules and Allen's temporal relations. Our experimental results show that for many (but not all) activities, this method produces a recognition accuracy of nearly 100%, in contrast to other methods based on data mining classifiers. The proposed method is accurate, very flexible and adaptable to a dynamic environment such as the "Smart Home" and we believe that it deserves further attention.

- [C.25] *BLAST: Off-the-shelf hardware for building an efficient hash-based cluster storage system*, George Parissis, George Xylomenos and Dimitris Gritzalis, *Proceedings of the IFIP International Conference on Network and Parallel Computing (NPC) 2009*, pp. 148–155.

Abstract: During the past few years, large, reliable and efficient storage systems have become increasingly important in enterprise environments. Additional requirements for these environments include low installation, maintenance and administration costs. In this paper we propose a hash-based storage approach, combined with block-level operating system semantics. The experimental evaluation confirms that the proposed approach is viable and can offer a cost-effective storage solution.

- [C.24] *A BitTorrent Module for the OMNeT++ Simulator*, Konstantinos Katsaros, Vasileios P. Kemerlis, Charilaos Stais and George Xylomenos, *Proceedings of the IEEE MASCOTS 2009*, pp. 361–370.

Abstract: In the past few years numerous peer to peer file sharing, or more generally content distribution, systems have been designed, implemented, and evaluated via simulations, real world measurements, and mathematical analysis. Yet, only a few of them have stood the test of time and gained wide user acceptance. BitTorrent is not just one such system; it holds the lion's share among them. The reasons behind its success have been studied to a great extent with interesting results. Nevertheless, even

though peer to peer content distribution remains one of the most active research areas, little progress has been made towards the study of the BitTorrent protocol, and its possible variations, in a fully controllable but realistic simulation environment. In this paper we describe and analyze a full featured and extensible implementation of BitTorrent for the OMNeT++ simulation environment. Moreover, since we aim to establish a realistic simulation platform, we show our enhancements to a conversion tool for a popular Internet topology generator and a churn generator based on the analysis of real BitTorrent traces. Finally we present the results from the evaluation of our prototype implementation regarding resource demands under different simulation scenarios.

- [C.23] *Router assisted overlay multicast*, Konstantinos Katsaros, Nikolaos Bartsotas and George Xylomenos, *Proceedings of the Euro-NGI Conference on Next Generation Internet Networks 2009*.

Abstract: While multicasting is considered valuable for content distribution, it is not widely supported on the Internet. Content providers have instead turned to peer assisted content distribution in order to efficiently serve large numbers of clients via unicast, thus removing the bandwidth bottleneck from their side. The redundant unicast transmissions of the same packet are not avoided however, they are just distributed between the peers. Since peer assisted content distribution represents a major fraction of total Internet traffic, a more efficient distribution scheme would be of great interest to users and network operators alike. For this reason, we reconsider overlay multicast as a potential solution for mass content distribution. We present an overlay multicast scheme inspired by Scribe that exploits co-operative access routers so as to improve the multicast content distribution trees produced. We investigate the properties of our scheme compared to both regular Scribe and IP multicast over Internet-like network topologies, via a full fledged simulation platform that can be used as a basis for the realistic evaluation of multicast based content distribution applications.

- [C.22] *Overlay Multicast Assisted Mobility for Future Publish/Subscribe Networks*, Konstantinos Katsaros, Nikolaos Fotiou, George C. Polyzos and George Xylomenos, *Proceedings of the ICT Mobile Summit 2009*.

Abstract: This paper shows how the multicast nature of proposed future publish/subscribe network architectures can assist mobility. Publish/subscribe is an information-centric networking paradigm. Unlike the current send/receive based Internet architecture which favours the sender of information, the publish/subscribe paradigm leads to a more balanced relationship between entities. Publish/subscribe is considered a promising architecture for the future Internet as it can potentially address various current Internet problems, such as spam and (Distributed) Denial of Service attacks. In a pub/sub architecture all data are transmitted via multicast and end host identification, which has a less critical role to begin with, is decoupled from location identification. In such an environment, fast mobility can be supported effectively, particularly for continuous media distribution (such as mobile TV) and other types of real-time multimedia applications. In this paper, a prototype overlay multicast scheme based on Scribe is extended to support mobility. We compare this scheme against an alternative solution based on Mobile IPv6 and demonstrate the effectiveness of the proposed approach and the basic trade-offs.

- [C.21] *Supporting Mobile Streaming Services in Future Publish/Subscribe Networks*, Konstantinos Katsaros, Nikolaos Fotiou, George C. Polyzos and George Xylomenos, *Proceedings of the Wireless Telecommunications Symposium (WTS) 2009*.

Abstract: The architecture of the current Internet was not originally designed to support either mobility or multicast. In particular, its coupling of host identification and location identification has hindered the provision of effective mobile services. At the same time, its lack of support for multicast distribution causes a multitude of redundant unicast transmissions, leading to an inefficient utilization of network resources. Both these limitations are especially apparent in the case of real-time continuous media distribution. The publish/subscribe paradigm has been proposed as a promising alternative to the current send/receive paradigm for a future Internet architecture. In future publish/subscribe networks, multicast will be the norm, and this change of the end-to-end communication semantics will lead to a networking environment more suitable for mobility. In the context of this paradigm, this paper considers a prototype architecture based on the Scribe overlay multicast scheme. Preliminary simulation results show that our publish/subscribe network implementation achieves better performance during mobility compared to Mobile IPv6 in all relevant metrics, such as hand-off delay (or resume time) and

loss of real-time traffic during disconnections, at the cost of a slight increase of the end-to-end delay due to the routing stretch imposed by the overlay.

- [C.20] *RLC protocol performance over TCP SACK*, Michael Makidis and George Xylomenos, *Proceedings of the Symposium on Modeling and Optimization in mobile, ad hoc and wireless networks (WiOpt) 2008*, pp. 135–140.

Abstract: The Radio Link Control (RLC) protocol, used in Universal Mobile Telecommunication System (UMTS) networks, is one of the most advanced and complex link layer protocols. Among its notable features are the absence of retransmission timers, which makes it tolerant to contention for the link, and the ability to abandon persistently lost frames, which makes it suitable for reliable transport layers. In order to assess whether it makes sense to use RLC with non-UMTS wireless links, especially in the face of the enhanced error recovery offered by TCP with the Selective Acknowledgment (SACK) option, we implemented it in the ns-2 simulator and measured the throughput achieved by File Transfer and Web Browsing over RLC, either with or without contention from a Media Distribution application. While we found that RLC adapts well to the available bandwidth, providing considerable performance gains with random losses, with bursty losses RLC hardly improved upon TCP SACK.

- [C.19] *Adaptive Link Layer Protocols for Shared Wireless Links*, George Xylomenos and Michael Makidis, *Proceedings of the ACM International Mobile Multimedia Communications Conference (MobiMedia) 2007*.

Abstract: The error prone nature of wireless links often necessitates the use of a link layer protocol to ensure acceptable application performance. While traditional link layers assume that they fully control the link, in most emerging wireless networks many sessions may dynamically share the link due to the presence of multiple contending users and/or applications. Such networks require link layers that can automatically adapt to bandwidth variations, offering good performance regardless of contention. To this end, we discuss two adaptive protocols, an Adaptive Selective Repeat (ASR) protocol that dynamically modifies its retransmission timeouts, and the Radio Link Control (RLC) protocol used by UMTS, an advanced protocol without retransmission timers. To assess the applicability of each approach, we measure the throughput achieved by File Transfer and Web Browsing over both protocols, with or without contention from a Media Distribution application, as well as the delay induced by these protocols to the contending application. Our results indicate that the complexity of RLC is not justified by its performance, as ASR nearly always outperforms it.

- [C.18] *Reducing the Transmission Power Requirements of the Multimedia Broadcast/Multicast Service*, George Xylomenos, Vasilis Tsakanikas and George C. Polyzos, *Proceedings of the IST Mobile & Wireless Communications Summit 2007*.

Abstract: The Multimedia Broadcast/Multicast Service (MBMS) was designed to support the economical distribution of multimedia content to large numbers of receivers in 3rd generation cellular networks. In this paper we present and evaluate an MBMS extension that reduces the transmission power requirements while increasing the number of potential users of such services, by supporting the distribution of multiple variants of the same content to heterogeneous receivers. We first describe the standard MBMS model, along with its state management and signaling procedures, as well as our extended MBMS model, in terms of the modifications that it imposes on the standard. We then present an analytical and simulation evaluation of the transmission power requirements of our approach against alternatives based on standard MBMS, showing that our approach maximizes the number of potential users, without excessive transmission power requirements.

- [C.17] *Multiple Layer Error Control over Wireless Links*, George Xylomenos and Giannis Vasalakis Kokkinakis, *Proceedings of the European Wireless (EW) 2007*.

Abstract: In this paper we compare the wireless performance of TCP with or without selective acknowledgments and in the presence or absence of a reliable link layer protocol, in order to determine whether link layer error control remains beneficial for TCP variants with improved error recovery capabilities, such as TCP SACK. We also examine whether there are adverse interactions between the two protocol layers, that is, whether link layer error control degrades the congestion control performance of the transport layer, and whether transport layer error recovery degrades the error control performance of the link layer. Our results show that, even with TCP SACK, link layer error control remains critical

for TCP performance over wireless links. Furthermore, our results show that with a TCP unaware link layer protocol we have a clear separation of concerns, whereby the link layer handles wireless losses and the transport layer handles congestion.

- [C.16] *Limitations of Fixed Timers for Wireless Links*, George Xylomenos, *Proceedings of the International Symposium on Parallel and Distributed Processing and Applications (ISPA) 2006*, pp. 159–170.
Abstract: Traditional reliable link layer protocols set their fixed retransmission timers under the assumption that they operate in isolation over the link. Emerging wireless networks however allow multiple link layer sessions to dynamically share the link. To assess the impact of this development, we examine the performance of Web Browsing over a Selective Repeat protocol with fixed retransmission timers, showing that the optimal retransmission timer values depend on the level of contention. We therefore propose an adaptive Selective Repeat protocol that modifies its retransmission timers based on prevailing conditions. Our measurements show that this adaptive scheme provides excellent Web Browsing performance regardless of the level of contention, under two very different wireless error models.
- [C.15] *A Multiple Content Variant Extension of the Multimedia Broadcast/Multicast Service*, George Xylomenos and Konstantinos Katsaros, *Proceedings of the IST Mobile & Wireless Communications Summit 2006*.
Abstract: This paper describes an extension of the Multimedia Broadcast/Multicast Service supporting the distribution of multiple variants of the same content to heterogeneous receivers. We first outline the standard MBMS model, along with its state management and signaling procedures, and then proceed to describe our extended MBMS model, detailing the modifications that it imposes on MBMS state management and signaling procedures. Finally, we compare via analysis and simulation our extension against an approach for achieving the same goal by using standard MBMS mechanisms in terms of the signaling overhead incurred, showing that our proposal is clearly superior.
- [C.14] *Supporting Multiple Content Variants in the Multimedia Broadcast/Multicast Service*, George Xylomenos and Konstantinos Katsaros, *Proceedings of the World Telecommunications Congress (WTC) 2006*.
Abstract: The Multimedia Broadcast/Multicast Service (MBMS) was recently standardized for use by 3rd Generation cellular networks, aiming to support the economical distribution of multimedia content to large numbers of receivers. This paper proposes an MBMS extension supporting the distribution of multiple variants of the same content to heterogeneous receivers. We first outline the standard MBMS model, along with its state management and signaling procedures, and then describe our extended MBMS model, explaining the modifications it imposes on standard MBMS. We then explain how our approach can be combined either with layered coding or transcoding for the generation of the multiple content variants to be distributed. Finally, we compare our proposal via analysis and simulation with some alternatives based on standard MBMS. Both the analytical and the simulation results indicate that our proposal increases the number of satisfied users without spending excessive resources, thus striking a good balance between the standards based alternatives considered.
- [C.13] *A Differentiated Services QoS Scheme Preventing Malicious Flow Behavior in Mobile Ad hoc Networks*, Margaritis Margaritidis, Christopher N. Ververidis, George Xylomenos and George C. Polyzos, *Proceedings of the European Wireless (EW) 2006*.
Abstract: The design of an efficient Quality of Service (QoS) scheme for a Mobile Ad hoc Network requires addressing all the challenging characteristics of such a network. Specifically, a QoS scheme for such an environment must have a light-weight implementation, in terms of both storage and processing requirements, must be scalable and keep the signaling overhead to a minimum, must focus on differentiated services instead of “hard” guarantees, and must provide the necessary incentives for the flows to be cooperative. This paper describes such a QoS scheme which manages to meet all the aforementioned requirements while still having a simple, yet effective, implementation.
- [C.12] *Analysis of a Multiple Content Variant Extension of the Multimedia Broadcast/Multicast Service*, George Xylomenos, Konstantinos Katsaros and George C. Polyzos, *Proceedings of the European Wireless (EW) 2006*.

Abstract: This paper describes and analyzes an extension of the Multimedia Broadcast / Multicast Service (MBMS) that supports the distribution of multiple variants of the same content to heterogeneous receivers. We first outline the standard MBMS model and then describe our extended MBMS model, detailing the modifications that it imposes on MBMS state management and signaling procedures. We then provide an analytical framework for the comparison of our extended model with both a single and a multiple service approach based on standard MBMS mechanisms. We apply this framework to a practical setting and show that our proposal is far more scalable in terms of signaling overhead than its standards based competitor.

- [C.11] *Adaptive Timeout Policies for Wireless Links*, George Xylomenos and Christos Tsilopoulos, *Proceedings of the IEEE International Conference on Advanced Information Networking and Applications (AINA) 2006*, vol.1, pp. 497–502.

Abstract: A considerable body of evidence indicates that the use of reliable link layer protocols over error prone wireless links dramatically improves the performance of Internet protocols and applications. While traditional link layer protocols set their timeout values assuming that they fully control the underlying link, some wireless networks allow multiple link layer sessions to co-exist over the same link. Since the optimal timeout values for a reliable link layer protocol depend on the available bandwidth, with dynamic link sharing such a protocol should ideally adapt its timeout values accordingly. We have thus designed an Adaptive Selective Repeat protocol that modifies its timeout values based on the policy used by TCP. We compare the performance of Web Browsing over Selective Repeat when using our adaptive timeout scheme with a range of parameters, against a manually tuned fixed timeout version. Our measurements show that these adaptive timeout policies outperform the fixed one, regardless of the level of contention, and that the best adaptive timeout policy in this setting is not the one used by TCP.

- [C.10] *Throughput Unfairness in TCP over WiFi*, Vasileios P. Kemerlis, Eleftherios C. Stefanis, George Xylomenos and George C. Polyzos, *Proceedings of the IEEE/IFIP Conference on Wireless on Demand Network Systems and Services (WONS) 2006*.

Abstract: This paper presents a measurement study of TCP performance at an operational WiFi deployment. After presenting the network topology and the tools used to generate and analyze traffic, we examine the throughput performance of competing TCP connections. We investigate how throughput is divided among the participating wireless hosts with respect to signal strength, traffic direction and use of the RTS/CTS mechanism. Our study shows that while competing clients with comparable signal strength are treated fairly, achieving similar throughput values, clients with lower signal strength are treated unfairly, relinquishing a larger share of the available bandwidth to clients with higher signal strength.

- [C.9] *Group Management for the Multimedia Broadcast/Multicast Service*, George Xylomenos, *Proceedings of the IST Mobile & Wireless Communications Summit 2005*.

Abstract: This paper compares the group management mechanisms used in the IP and the MBMS multicasting models. After outlining the design of each model, we describe the group management protocols that they employ. We then examine how the IP group management protocols can be adapted for MBMS and finally evaluate the group management approach adopted by MBMS. Our main findings are that IGMP v.2 is preferable for use with MBMS, that the join/leave group management approach of MBMS outperforms the query/report approach of IP and that the reliability of the MBMS approach can be enhanced by upcalls.

- [C.8] *Multi-Service Link Layer Enhancements for the Wireless Internet*, George Xylomenos and George C. Polyzos, *Proceedings of the IEEE Symposium on Computers and Communications (ISCC) 2003*, pp. 1147–1152.

Abstract: Although TCP and UDP application performance over wireless links may be substantially improved via link layer error recovery, different schemes are appropriate for each application class. We present a multi-service link layer architecture that simultaneously enhances the performance of diverse applications by supporting multiple error recovery mechanisms in parallel. We simulated concurrent file transfers and WWW browsing over TCP and continuous media distribution over UDP using our

architecture. The results show that each application achieves similar improvements as when it operates alone over its preferred scheme, despite the parallel execution of diverse applications.

- [C.7] *Wireless Link Layer Enhancements for TCP and UDP Applications*, George Xylomenos and George C. Polyzos, *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS), Workshop on Wireless, Mobile and Ad Hoc Networks (WMAN), 2003*, pp. 225–232.

Abstract: Internet application performance over wireless links is disappointing due to wireless impairments that adversely affect higher layers. This paper focuses on link layer enhancement mechanisms that hide wireless errors from the rest of the Internet. We simulated file transfer and WWW browsing over TCP and continuous media distribution over UDP, in conjunction with various link layer schemes. Our results reveal that WWW browsing has substantially different behavior than file transfer, that existing TCP enhancement schemes have limited applicability and that UDP applications are best served by schemes inappropriate for TCP. Therefore, multiple link layer solutions are needed to optimize the performance of diverse applications.

- [C.6] *Quality of Service Issues in Multi-Service Wireless Internet Links*, George Xylomenos and George C. Polyzos, *Proceedings of the International Workshop on QoS in Multiservice IP Networks (QoS-IP) 2001*, pp. 347–365.

Abstract: Internet application performance over wireless links is disappointing, due to wireless impairments and their adverse interactions with higher protocol layers. In order to effectively address these problems without the need for global protocol upgrades, we focus on link layer enhancement schemes. Simulations reveal that different schemes work best for different applications. We have thus developed a multi-service link layer architecture that can simultaneously enhance the performance of diverse applications by supporting multiple link mechanisms concurrently. Simulations confirm that this architecture provides dramatic performance improvements. The architecture can be embedded in various ways into the Internet. A critical issue for Quality of Service support over wireless links is the unpredictability of available resources. Our approach is based on fair sharing of the link before any measures are taken to improve the performance of individual traffic classes. This approach turns over the error control trade-off to the applications themselves.

- [C.5] *TCP & UDP Performance over a Wireless LAN*, George Xylomenos and George C. Polyzos, *Proceedings of the IEEE INFOCOM '99*, pp. 439–446.

Abstract: We present a comprehensive set of measurements of a 2.4 GHz DSSS wireless LAN and analyze its behavior. We examine issues such as host and interface heterogeneity, bidirectional (TCP) traffic and error modeling, that have not been previously analyzed. We uncover multiple problems with TCP and UDP performance in this system. We investigate the causes of these problems (radio hardware, device drivers, network protocols) and discuss the effectiveness of proposed improvements.

- [C.4] *Enhancing Wireless Internet Links for Multimedia Services*, George C. Polyzos and George Xylomenos, *Proceedings of the International Workshop on Mobile Multimedia Communications (MoMuC) '98*, pp. 379–384.

Abstract: We describe a novel link layer protocol architecture that aims to enhance the performance of Internet protocols over wireless links, in particular in order to effectively support interactive multimedia applications with varying Quality-of-Service requirements. The degraded performance of these links on the Internet and the inadequacy of existing approaches in overcoming these problems in a protocol independent manner motivate our solution. Our design provides multiple services and performance feedback to higher layers, thus supporting adaptive protocols and applications. In addition, it serves as the basis for Internet evolution towards Quality-of-Service provision, since it can be used to support relevant mechanisms at higher layers, regardless of the underlying hardware.

- [C.3] *Enhancing Wireless Internet Links*, George Xylomenos and George C. Polyzos, *Proceedings of the International Conference on Telecommunications (ICT) '98*, pp. 394–398.

Abstract: We describe a novel link layer architecture to enhance the performance of Internet protocols over wireless links. The degraded performance of these links on the Internet and the inadequacy of existing approaches of overcoming these problems in a protocol independent manner, motivate our solution. Our link layer provides multiple services and performance feedback to higher layers, thus

supporting adaptive protocols and applications. In addition, our approach can serve as the basis for future Internet evolution towards Quality of Service provision. We also describe our ongoing research and implementation directions.

- [C.2] *IP Multicasting for Point-to-Point Local Distribution*, George Xylomenos and George C. Polyzos, *Proceedings of the IEEE INFOCOM '97*, pp. 1382–1389.

Abstract: While support for IP multicasting continues to spread enabling new applications, an increasing number of hosts connects to the worldwide Internet via low bandwidth Point-to-Point links, such as wireline or wireless telephone lines. In this paper we discuss existing proposals for local and wide area IP multicasting and their implications for Point-to-Point links, identify problems with their integration in this environment, and propose alternative special purpose mechanisms to solve these problems. The main problems are overhead due to IGMP leave latency and unnecessary continuous probing of potentially power constrained hosts. Our solution is an alternative to IGMP mechanisms based on join/leave messages for tracking group membership over PtP networks. After presenting the implementation requirements of our proposed and the existing mechanisms, we compare them with respect to performance, interoperability, robustness and implementation complexity, demonstrating that our join/leave protocol is uniformly superior.

- [C.1] *IP Multicasting for Wireless Mobile Hosts*, George Xylomenos and George C. Polyzos, *Proceedings of the IEEE MILCOM '96*, pp. 933–937.

Abstract: We consider the problem of efficient multicast support for mobile wireless hosts in TCP/IP networks. We summarize techniques supporting multicasting and mobility, along with their constraints and interactions, and explore architectural alternatives for solutions to the combined problem that remain compatible with the existing architecture. We propose a new mechanism for group management, optimized for point-to-point links. We also discuss three proposed multicast delivery mechanisms and compare them with respect to efficiency on wireless networks and impact on host software.

Refereed Conferences (Short Papers, Posters and Demos)

- [S.11] *Aretousa: A Competitive Audio Streaming Software for Network Music Performance*, Konstantinos Tsioutas, George Xylomenos and Ioannis Doumanis, *Proceedings of the Audio Engineering Society Convention 146, 2019*.

Abstract: Many existing open source systems provide support for Network Music Performance (NMP), with each one catering to a specific system and usage scenario. As our research in evaluating the Quality of Experience (QoE) of NMP systems as perceived by musicians involves widely different scenarios and requires extensive instrumentation of the platform, we built a new NMP system, Aretousa. Our system offers a large number of configuration and monitoring options, without sacrificing latency, the most critical factor for NMP. To show that Aretousa provides flexibility while being competitive with the state of the art in terms of latency, we present measurements comparing it against JackTrip in multiple setups over a high speed research network.

- [S.10] *A Framework for Understanding and Defining Quality of Musicians' Experience in Network Music Performance Environments*, Konstantinos Tsioutas, Ioannis Doumanis and George Xylomenos, *Proceedings of the Audio Engineering Society Convention 146, 2019*.

Abstract: Considerable research has been conducted on network and system level metrics related to Network Music Performance (NMP). However, little empirical evidence is currently available for assessing the actual Quality of Musician's Experience (QoME) over NMP. We propose a research framework that integrates both subjective and objective aspects of musicians' experience, by explicitly considering the psychological state and profile of each musician, the environment acoustic variables, the music performance variables and the Quality of Service of the network as the key dimensions that impact QoME. We will use the proposed framework to drive empirical studies designed to explore the QoME of musicians performing over the Internet; this paper is a first step in this direction.

- [S.9] *ICN enabling CoAP Extensions for IP based IoT devices*, Nikos Fotiou, George Xylomenos, George C. Polyzos, Hasan Islam, Dmitriy Lagutin, Teemu Hakala and Eero Hakala, *Proceedings of the ACM Information-Centric Networking (ICN) Conference 2017 (Demo)*.

Abstract: The Constrained Application Protocol (CoAP) and its extensions, such as observe and group communication, offer the potential for developing novel IoT applications. However, a full-fledged CoAP-based application requires delay-tolerant communication and support for multicast: since these properties cannot be easily provided by existing IP networks, developers cannot take full advantage of CoAP, preferring to use HTTP instead. In this demo we show how proxying CoAP traffic over an ICN network can unleash the full potential of CoAP, simultaneously shifting overhead and complexity from the (constrained) endpoints to the network.

- [S.8] *Reduced Switching Delay for Networked Music Performance*, George Xylomenos, Christos Tsilopoulos, Yannis Thomas, George C. Polyzos, *Proceedings of the Packet Video Workshop, 2013 (Poster)*.
Abstract: Networked Music Performance (NMP), where musicians located in different places perform together in real time via a network, requires extremely low end-to-end delay. As part of an effort to reduce all aspects of NMP delay, we are focusing on the packet processing and replication delay at the centralized node which enables communication between the multiple NMP participants. We propose three approaches to dramatically reduce this delay, based on both hardware and software optimizations.
- [S.7] *Multisource and Multipath File Transfers through Publish-Subscribe Internetworking*, Yannis Thomas, Christos Tsilopoulos, George Xylomenos and George C. Polyzos, *Proceedings of the ACM SIGCOMM Information Centric Networking (ICN) Workshop 2013 (Poster)*.
Abstract: We present mmFTP, an information-centric and receiver-driven file transfer protocol for the Publish Subscribe Internetworking (PSI) architecture. mmFTP supports both multisource and multipath transfers, while requiring minimal complexity in terms of network operation. We describe the basic design and operation of mmFTP and present preliminary experimental performance results from a prototype implementation deployed in the PlanetLab testbed.
- [S.6] *Demonstrating Usage Diversity Over an Information-Centric Network*, George Parisis, Ben Tagger, Dirk Trossen, Christos Tsilopoulos, Yannis Thomas, Charilaos Stais and George Xylomenos, *Proceedings of the IEEE INFOCOM 2012 (Demo)*.
Abstract: Information-Centric Networking (ICN) is becoming increasingly popular: various architectures have been proposed to transform today's host-centric networks to information-centric ones. An important question that begs to be answered is whether the ICN paradigm can adequately support various types of existing network usage, as well as set the foundations for new application types. In this demonstration, we exhibit four different communication models via corresponding applications supported by our architecture. All applications are realized through the publish/subscribe service model exported by our ICN nodes.
- [S.5] *Demonstrating an Information-Centric Network in an International Testbed*, George Parisis, Ben Tagger, Dirk Trossen, Dimitris Syrivelis, Paris Flegkas, Leandros Tassioulas, Charilaos Stais, Christos Tsilopoulos and George Xylomenos, *Proceedings of the International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TRIDENTCOM) 2012 (Demo)*.
Abstract: Information-Centric Networking (ICN) has increasingly been attracting attention by the research community. In ICN the center of attention becomes the information itself and not the endpoints as in today's IP networks. In this demonstration we present applications that we developed as proof of concepts for our ICN approach. A video streaming as well as a voice and a HTTP over publish/subscribe application that run on top of our ICN prototype will be demonstrated running in an international testbed.
- [S.4] *Supporting Mobility in a Publish Subscribe Internetwork Architecture*, Varvara Giannaki, Xenofon Vasilakos, Charilaos Stais, George C. Polyzos, George Xylomenos, *Proceedings of the IEEE Symposium on Computers and Communications (ISCC) 2011 (Short Paper)*.
Abstract: Information-Centric Networking (ICN) is constantly gaining momentum within the Future Internet research community. In the PURSUIT research project we are developing a clean-slate Pub/Sub Internetworking (PSI or Ψ) approach with integrated seamless mobility support. The novel ICN mechanisms supported in Ψ , along with smartly placed in-network caches, enable the architecture to handle both mobile and fixed devices in a uniform way. This paper presents a blueprint for optimizing mobility support in Ψ without modifications to the architecture or add-on solutions. We demonstrate a

micro-mobility scenario that describes the functionality of Ψ 's core components in supporting mobility and then sketch our plans for future work and a proper assessment of these designs.

- [S.3] *Socket Emulation over a Publish/Subscribe Network Architecture*, George Xylomenos and Blërim Cici, *Proceedings of the Future Network and Mobile Summit 2010 Conference (Poster)*.

Abstract: The current drive towards new networking paradigms that make information the center of the network architecture, cannot ignore the fact that any new architecture will have to co-exist with the existing Internet for an extended period of time. In order for a new architecture to be globally deployed, it must ensure that existing applications will continue to operate, preferably without the need to even recompile them. As part of the Publish Subscribe Internet Routing Paradigm (PSIRP) project, we have explored the options for making existing applications based on the ubiquitous Sockets Application Programming Interface (API) compatible with the PSIRP prototype implementation, which is based on publish/subscribe principles throughout the protocol stack. We describe an emulator which mediates between the client/server socket calls and the publish/subscribe calls implemented by the PSIRP prototype. Our socket emulator allows existing applications, either TCP or UDP based, to run unmodified over an information centric network architecture that is radically different from the endpoint centric Internet architecture for which they were designed.

- [S.2] *Peer Assisted Content Distribution over Router Assisted Overlay Multicast*, George Xylomenos, Konstantinos Katsaros and Vasileios P. Kemerlis, *Proceedings of the Euro-NF Future Internet Architecture (FIA) Workshop 2008 (Short Paper)*.

Abstract: While multicasting is considered valuable for content distribution, it is not widely supported on the Internet, despite the emergence of scalable overlay schemes. Content providers have instead turned to peer assisted content distribution in order to efficiently serve large numbers of clients, thus removing the bandwidth bottleneck from their side but placing a heavy burden on the clients. Even if we assume that multicast will become prevalent in the future Internet, peer assisted content distribution will still be useful for asynchronously distributing very large amounts of data. We have thus designed a multicast variant of BitTorrent, paying special attention to the incentives required to ensure that peers will not only consume, but also contribute content. To provide a multicast substrate for our application in the current Internet, we also present an overlay multicast scheme inspired by Scribe that exploits co-operative access routers so as to improve the distribution trees.

- [S.1] *Monitoring and Modeling the Every-Day Behavior of People with Special Needs*, George J. Papamatthiakis, Elias C. Efstathiou, George C. Polyzos and George Xylomenos, *Proceedings of the International Conference on Distributed Computing in Sensor Systems (DCOSS) 2008 (Poster)*.

Abstract: We present preliminary work on a sensor-based system that we are currently implementing with the aim of monitoring the health of the elderly and people with special needs. The system functions both inside and outside the home of the monitored individuals, using sensors and GPS-enabled cellular phones. Its objective is to first learn the daily behavior of the monitored individuals, and then detect changes in their routines and health status, providing alerts and a preliminary diagnosis as quickly as possible when something out of the ordinary occurs.

Book Chapters

- [B.6] *An Improved Scheme for Protecting Medical Data in Public Clouds*, Nikos Fotiou and George Xylomenos, in *Communications in Computer and Information Science*, vol. 736, Carsten Røcker, John O'Donoghue, Martina Ziefle, Markus Helfert and William Molloy (Eds.), Springer, 2017. **Abstract:** Public Clouds offer a convenient way for storing and sharing the large amounts of medical data that are generated by, for example, wearable health monitoring devices. Nevertheless, using a public infrastructure raises significant security and privacy concerns. Even if the data are stored in an encrypted form, the data owner should share some information with the Cloud provider in order to enable the latter to perform access control; given the high sensitivity of medical data, even such limited information may jeopardize end-user privacy. In this paper we employ an access control delegation scheme to enable the users themselves to perform access control on their data, even though these are stored in a public Cloud.

In our scheme access control policies are evaluated by a user-controlled gateway and Cloud providers are only entrusted with respecting the gateway's decision. Furthermore, since medical data must often be shared with health providers of the user's choice, we rely on a proxy re-encryption technique to allow such sharing to take place. Our scheme encrypts data before storing them in the Cloud and applies proxy re-encryption using Cloud resources to encrypt data separately for each (authorized) user. Our proxy re-encryption scheme ensures that misbehaving Cloud providers cannot use re-encryption keys to share content with unauthorized clients, while delegating the costly re-encryption operations to the Cloud.

- [B.5] *Multicast Protocols & Algorithms*, George C. Polyzos and George Xylomenos, in *Encyclopedia of Computer Science and Engineering*, Benjamin Wah (Ed.), Wiley, 2009.

Abstract: Multicast refers to the transmission of data from one sender to an arbitrary set of receivers, a service useful for applications such as group conferencing, media distribution and update of replicated databases. From an application viewpoint, the importance of multicast is that it allows addressing a set of receivers via a common identifier. From a network viewpoint, the importance of multicast is that when the transmission paths between the sender and the receivers share some links, multicast can conserve resources by only sending the data once over each such link. The most important issues raised by multicast are routing data to multiple destinations with minimal delay and/or duplication, handling feedback from possibly numerous receivers in a scalable manner and providing appropriate Quality of Service for data delivery to each receiver.

- [B.4] *Wireless Multimedia in 3G Networks*, George Xylomenos and Vasilis Vogkas, in *Emerging Wireless Multimedia Services and Technologies*, Apostolos Salkintzis and Nikos Passas (Eds.), Wiley, 2005.

Abstract: The two most important aspects of 3G networks with respect to IP based multimedia services are the IP Multimedia Subsystem (IMS) and the Multimedia Broadcast / Multicast Service (MBMS). The IMS enables complex IP based multimedia sessions to be created with guaranteed QoS for each media component. Example applications include voice telephony and video conferencing. The IMS interoperates with both traditional telephony services and external IP based multimedia services. The MBMS provides native IP broadcast and multicast support in 3G networks, allowing high bandwidth services to be economically offered to multiple users. Example applications include video streaming via multicast and location based services via broadcast. The MBMS interoperates directly with IP multicasting. While both the IMS and the MBMS are IP based, their standardization is proceeding independently. It is however clear that their combination would allow numerous new services to be provided.

- [B.3] *Multi-Service Wireless Internet Link Enhancements*, George Xylomenos and George C. Polyzos, in *Wireless Communications Systems and Networks*, Mohsen Guizani (Ed.), Kluwer Academic, 2004.

Abstract: The deployment of several real-time multimedia applications over the Internet has motivated a considerable research effort on the provision of multiple services on the Internet. In order to extend this work over wireless links however, we must also take into account the performance limitations of wireless media. We survey various related approaches and conclude that link layer schemes provide a universal and localized solution. Based on simulations of application performance over many link layer schemes we show that different approaches work best for different applications. We present a multi-service link layer architecture which enhances the performance of diverse applications by concurrently supporting multiple link layer schemes. Simulations of multiple applications executing simultaneously show that this approach dramatically improves performance for all of them. We finally consider embedding this approach into a Quality of Service oriented Internet, discussing the traditional best-effort architecture, the Differentiated Services architecture and an advanced dynamic service discovery architecture.

- [B.2] *TCP/IP Performance over Wireless Networks*, George Xylomenos, George C. Polyzos, Petri Mahonen, and Mika Saarinen, in *High Performance TCP/IP Networking*, Mahbub Hassan and Raj Jain (Eds.), Prentice Hall, 2004.

Abstract: The current strong drive towards Internet access via mobile terminals, makes the inclusion of wireless systems such as Cellular Communications (CC) and Wireless Local Area Networks (WLAN)

into the mainstream Internet very desirable. CC and WLAN systems however raise a multitude of performance issues, since environmental conditions and terrestrial obstructions and reflections lead to high and unpredictable error rates. CC and WLAN systems mostly share the characteristics of traditional wireless systems (satellite and terrestrial microwave), such as high error rates. They also share some of the characteristics of wired systems, such as low physical layer propagation delays. As a result, in order to improve their performance, a synthesis of techniques for enhancing the performance of both wired and wireless links is required, that will also take into account the requirements of the TCP/IP protocol suite. In this chapter we present the characteristics and performance limitations of various existing and emerging wireless systems and survey a wide range of approaches for enhancing TCP/IP performance over such links. Although mobility is inherently associated with CC systems, the additional problems that it induces, such as communication pauses whenever mobile devices move between cells, are covered in a separate chapter.

- [B.1] *Internet Protocols over Wireless Networks*, George C. Polyzos and George Xylomenos, in *Multimedia Communications: Directions & Innovations*, Jerry D. Gibson (Ed.), Academic Press, 2000.

Abstract: We discuss the problems that arise when standard Internet protocols are used over wireless links, such as degraded Transmission Control Protocol (TCP) performance when wireless errors are interpreted as congestion losses. We use case studies drawn from commercial Wireless Local Area Networks and cellular telephony systems to illustrate these problems. Then, we survey some proposed approaches to mitigating such problems and examine their applicability. Finally, we look at the future of wireless systems and the new challenges that they will create for Internet protocols and state some goals for further protocol enhancement and evolution.