

VPN Gateways over Network Processors: Implementation and Evaluation

Yi-Neng Lin¹, Ying-Dar Lin¹, Yuan-Cheng Lai², Chuan-Hung Lin¹

¹Department of Computer Science, National Chiao-Tung University, Taiwan

²Department of Information Management, National Taiwan University of Science and Technology, Taiwan
ynlin@cs.nctu.edu.tw, ydlin@cs.nctu.edu.tw, laiyc@cs.ntust.edu.tw, nestlin@cs.nctu.edu.tw

Abstract

Networking applications, such as VPN and content filtering, demand extra computing power in order to meet the throughput requirement nowadays. In addition to pure ASIC solutions, network processor architecture is emerging as an alternative to scale up data-plane processing while retaining design flexibility. This article, rather than proposing new algorithms, illustrates the experience in developing IPsec-based VPN gateways over network processors, and investigates the performance issues. The external benchmarks reveal that the system can reach 45Mbps for IPsec using 3DES algorithm, which improves by 350% compared to single XScale core processor and parallels the throughput of a PIII 1GHz processor. Through the internal benchmarks, we analyze the turnaround times of the main functional blocks, and identify the core processor as the performance bottleneck for both packet forwarding and IPsec processing.

Keywords: VPN, Gateway, Network Processor, Bottleneck, Implementation.

1 Introduction

Today's networking applications, such as virtual private network (VPN) [1] and content filtering that offer extra security and application-aware processing, have demanded more powerful hardware devices to achieve high performance. The most straight-forward way to tackle this problem is to increase the clock rate of a general purpose processor, though some disadvantages, such as the cost and the technology limit, accompany. Moreover, the low efficiency is also expected since the processor, as its name suggests, is not specifically designed for the processing of networking packets.

Another solution to this problem is to employ the concept of *offloading*, that is, to shift the computing-intensive tasks from the core processor to a number of additional processors. The Application-Specific Integrated Circuit (ASIC) [2] has been a possible candidate to serve as an additional processor. Nonetheless, this workaround might not be preferred in two aspects. First, since the

functionalities are fixed once tapped out, it needs to be redesigned for any modifications. Second, the development period is so time-consuming that the time-to-market requirement may not be met.

Network processors [3] are now embraced as an alternative solution to remedy the above-mentioned problems because of its re-programmability, specifically designed instructions for networking purpose and the hardware threads with minimal, if not zero, context switch overhead. In this work, we explored the feasibility of implementing VPN, which is a computation intensive application, over the Intel IXP425 [4] network processor featuring an XScale core, *multiple hardware contexts* and *coprocessors*, and tried to figure out the performance and possible bottlenecks of the implementation. The VPN mechanism, which is usually based on IPsec [5], comprises several processing stages such as packet reception (Rx) and transmission (Tx), encryption and decryption, authentication and table lookups, each of which needs a certain amount of processing. We analyzed the detailed packet flow and described how to offload packet processing overhead to coprocessors. Some efforts have also been done to port the VPN application from ordinary PC to IXP425 in the meantime. We then externally and internally benchmarked the prototype. The former characterized performance figures of the implementation, while the latter carried out the in-depth analysis of the observations, such as system bottlenecks that were left unexplained in the external benchmarks. The Xscale is identified as the bottleneck for both Forwarding and IPsec processing.

Some related works researching the bottlenecks of network processors can also be found in the literature: Spalink et al. [6] presented the results of simple IP forwarding and Lin et al. [7] implemented DiffServ, both over Intel IXP1200. Nevertheless, our work differs from theirs in that (1) no coprocessor was involved in their implementations; (2) both the control-plane and part of the data-plane processing were handled in the core processor of IXP425 while the core of IXP1200 took care of the control-plane packets only, and (3) computation intensive VPN application was considered, as compared with simple forwarding and memory intensive classification of these two studies. Other works regarding implementations over

*Corresponding author: Yi-Neng Lin; E-mail: ynlin@cs.nctu.edu.tw

network processors include [8][9][10]. However, none of them investigates the coprocessor effects, and few of them investigate the bottlenecks of their design.

This article is organized as follows. We first describe the hardware and software architectures of IXP425. Next, we elaborate the details of the design and implementation of VPN over IXP425. Then we present the results and observations from the external and internal benchmarks. Some conclusive remarks of this article are made finally.

2 Hardware and Software Architecture of IXP425

2.1 Hardware Architecture of IXP425

The hardware block diagram of IXP425 is depicted in Figure 1. The core of IXP425 is a 533MHz XScale processor handling system initialization and software objects execution. Three buses interconnected by two bridges provide the connectivity among components on IXP425.

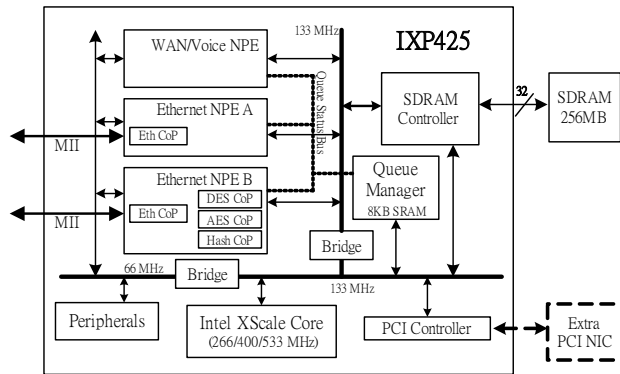


Figure 1 Hardware Architecture of IXP425

To assist the XScale core in processing networking packets, three 133MHz programmable network processor engines (NPEs) are used to execute in parallel the code image stored in internal memory for providing functions such as MAC, CRC checking/generation, AAL2, AES, DES, SHA-1 and MD5, in cooperate with a number of application-specific coprocessors. The support of hardware multithreading with single cycle context switch overhead further makes NPEs more tolerant to long memory accesses and thus reduces the number of processor stalls. The communication between the XScale core and NPEs is handled by a hardware queue manager using interrupt and message queue mechanisms. The queue manager also contains 8KB SRAM divided into 64 independent queues manipulated as circular buffers for allocating free memory space to incoming packets and for locating packets in the memory. The SDRAM can be expanded up to 256MB for storing tables, policies and OS applications in addition to

packets. A PCI interface is available for an additional PCI NIC. Some peripheral controllers, like USB and UART controllers, are also equipped into IXP425 for better extensibility.

2.2 Detailed Packet Flow in IXP425

The processing flow of an ordinary packet is elaborated below referring to Figure 1. Upon the arrival of a packet at the interface of an NPE, it is partitioned into several 32byte segments and stored at the Receive FIFO of an Ethernet coprocessor which in turn performs MAC-related operations. The NPE then moves those segments into corresponding addresses in SDRAM allocated by the queue manager, which then interrupts the XScale of the reception for further processing. During normal processing procedures such as IP and other higher layer protocol stacks at XScale, chances are that some authentic and cryptographic operations are needed. The XScale core may handle them either by itself or by *offloading* the computation overhead to appropriate coprocessors residing in NPE B. In the latter scenario, the coprocessors are directly invoked by NPE B, requested by the XScale, to process a certain data segment in SDRAM, where a message queue implemented in the queue manager is exploited to pass the request. The queue manager is informed by NPE B upon the completion of the operations and then interrupts the XScale.

2.3 Software Architecture of IXP425

The software architecture is divided into two portions, namely the platform independent (applications and some higher level components such as networking protocol stacks in OS) and dependent parts (mainly device drivers). This design is favorable especially when an OS migration from a certain H/W platform to IXP425 is demanded, that is, the developers need to focus only on the dependent part, namely the development of drivers. When implementing device drivers, a set of software libraries collectively referred to as *AccessLibrary* can be used to drive devices such as NPEs, coprocessors, peripherals, etc. The *AccessLibrary* also provides utilities to implement some OS-related functions such as mutual exclusion.

The software processing flow is described as follows with library functions adopted from the *AccessLibrary*. During the boot time a function named *IxNpeDI* is called to download the corresponding code image into the instruction cache of each NPE. Then two functions, *IxQmgr* and *IxNpeMh*, are called to initialize the queue manager as well as the message handler responsible for the communication between NPEs and XScale. The Ethernet-related functions, *IxEthAcc* and *IxEthDB*, are used to receive and transmit Ethernet frames, while the *IxCryptoAcc* function is incorporated for possible cryptographic operations during packet processing.

3 Design and Implementation

In this section, we first introduce basic operations in a VPN environment and then analyze its packet processing flow in order to identify possible bottlenecks as offloading candidates. Finally, we describe how to implement a VPN gateway over IXP425.

3.1 VPN Briefing

Virtual Private Network (VPN) provides secure transmission over un-trusted networks. Normally the IPSec protocol is adopted as the underlying technique due to the popularity of the Internet Protocol. It supports data authentication, integrity and confidentiality, in which two VPN gateways are employed as endpoints that construct tunnels for secure data transmission. Since VPN bandwidth between networks depends mainly on the processing capability of the endpoints, improving the performance of the gateways will be decisive to the VPN throughput.

3.2 Identifying Offloading Candidates

To resolve the performance issue, we analyze the VPN packet processing flow in order to identify possible candidates to be offloaded to coprocessors. A detailed incoming IPSec packet flow was displayed in Figure 2. It consists of three main blocks, namely the packet reception, IPSec processing and packet transmission. Their operations are elaborated below.

Once an Ethernet frame is received by the physical interface, checking for frame check sequence screens out broken frames and then remaining frames are filtered in accordance with possible destination MAC address configurations. Reception is accomplished after the frame is moved into memory, followed by a classification recognizing it as an IPSec packet. At this time, some table lookups for processing rules and cryptographic parameters are performed and payload of this IPSec packet is decrypted or checked for authentication. Finally, a new packet

decrypted from the original IP payload may be further processed by higher-level protocols, or be transmitted according to the routing table.

Tasks suitable to be offloaded to coprocessors can be identified by two characteristics: whether those tasks are repeated routines or computation intensive ones. As mentioned earlier this section, we know that IPSec processing, especially the cryptographic operation, is computation intensive. Hence, we decide to pick the cryptographic processing as an offloading candidate. Another candidate comprises the packet transfer, CRC checking/generation, MAC filtering, and packet movement between NPE and memory, since the procedures are precisely the same for every packet. From the hardware block diagram in Section 2, clearly the IXP425 meets well the requirements of the identified candidates.

3.3 Implementation

We adopt the NetBSD [11], a secure, highly portable and open-source OS derived from 4.4BSD, as our operating system. Clean design between platform dependent and independent parts makes it a good implementation target for new hardware platform. Following relates three major components in prototyping a security gateway over IXP425.

3.3.1 Operating System Porting

The most efficient way to porting an OS to a new platform is to refer to the port of another similar platform and then implement drivers for devices of the target platform [12]. To port NetBSD over IXP425, therefore, we adopt the “EVBARM” port in NetBSD. It supports various evaluation boards that are based on XScale or other ARM-based core processors, so that only system-level modifications have to be done to enable normal operations on IXP425. Example modifications include the CPU identification, setup of board-specific memory map, and system initialization procedures.

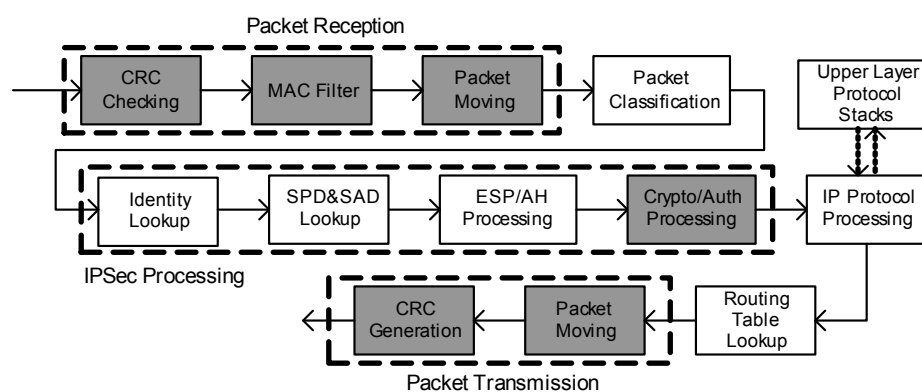


Figure 2 Processing Flow of an Inbound IPSec Packet. Shaded Blocks Are Candidates to Be Offloaded

3.3.2 Driver Development

A number of drivers for devices such as UART, NPEs and coprocessors need to be implemented for communication between the operating system and those devices. This effort can be alleviated with the help of the AccessLibrary introduced in Section 2. Besides drivers, we have to modify two OS dependent modules, namely OSSSL and IxOSServices, in AccessLibrary to ensure proper operations of the OS-related services.

3.3.3 Offloading the Cryptographic Operations

The last modification to kernel concerns the offloading of in-kernel IPsec cryptographic computations from XScale to NPE. In addition to the ordinary method requiring the kernel, and therefore the core processor, to perform encryption/decryption operations, NetBSD provides another option named *FAST_IPSec* that makes use of the Open Crypto Framework (OCF), for offloading. In OCF, the cryptographic operations can be handled by a *registered* function. The *FAST_IPSec* differs significantly from the original IPsec in that, in the processing flow of the later option, the XScale would not suspend during cryptographic operations. We exploit this technique by pre-registering the crypto driver, which drives the crypto coprocessor using functions in AccessLibrary, to the OCF so that all cryptographic overhead is shifted from the core to the coprocessor in NPE B.

4 System Benchmark and Bottleneck Analysis

In this section, we investigate the benefits from offloading by externally benchmarking the implementation using various offloading schemes. A number of internal tests are also conducted in order to observe what cannot be obtained in the external benchmarks.

4.1 System Benchmark Setup

4.1.1 Offloading Schemes Design and Benchmark Environment

To have a better understanding of the improvement from the network processor architecture as well as the offloading mechanism on packet Rx/Tx and IPsec processing, we design and benchmark systems of different offloading schemes, and compare their performance results. Four offloading schemes are adopted: (1) offload both crypto operations and packet Rx/Tx to the corresponding coprocessors; (2) offload crypto operations only; (3) offload Rx/Tx only, and (4) no offloading.

As for the external benchmark environments for packet forwarding and IPsec, we use SmartBits, which is a networking traffic generator and a performance analyzer, to generate the input traffic and to collect and analyze the results. For internal tests, some system utilities such as *vmstat*, *top* and *GProf*, are employed to obtain the system state as well as other internal behaviors such as CPU and memory utilizations.

4.2 Scalability Test

Scalability tests aim to derive the maximum throughput of the prototypes with different offloading schemes. Another gateway implementation using Pentium III 1GHz processor and 256MB SDRAM is also included for comparison between IXP425 and x86-based systems.

4.2.1 Packet Forwarding

Figure 3 shows the performance results of 1-to-1 packet forwarding under the condition of zero packet loss. The Rx/Tx is designed to be offloaded to up to 2 NPEs. From the figure we can see that throughput of the IXP425 offloaded by two NPEs parallels the one of Pentium III 1GHz. Both of them can support wired speed for packet lengths larger than 512 bytes. Besides, a performance improvement of up to 60% contributed by NPEs can be gained. We also observed that the maximum throughput occurs when the packet length is 1024 bytes, rather than other larger lengths. This is because the longer processing time, *mainly caused by more SDRAM accesses*, of larger packets counteracts the benefit from their reduced header processing overhead.

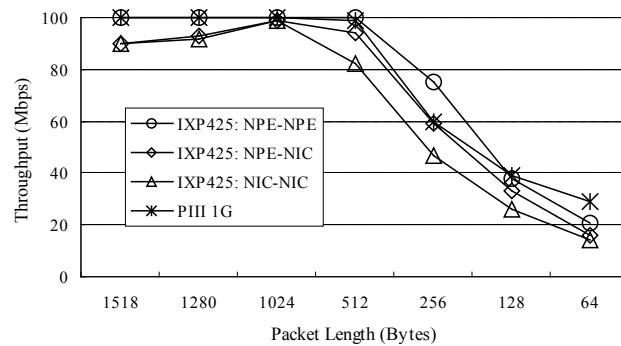


Figure 3 Throughput of Packet forwarding When Different Number of NPEs Is Used for Offloading

4.2.2 IPsec Processing

Figure 4 depicts the throughput of DES for different packet lengths. Some observations can be made. First, offloading IPsec processing to coprocessors in NPE B improves the performance of the prototype by 350%; in some cases IXP425 even outperforms the Pentium III 1GHz. Second, the maximum of throughputs occurs when the packet length is 1450 bytes, instead of 1518 bytes. This is because 1450 bytes is the largest length for a packet not to be fragmented when being encapsulated into an IPsec one. Third, the throughput of the IXP425 performing DES is similar to the one of 3DES, not shown in the figure, though the computation requirement of the former is almost triple of the later. The reason is that it is the XScale, not the coprocessor, which becomes the bottleneck.

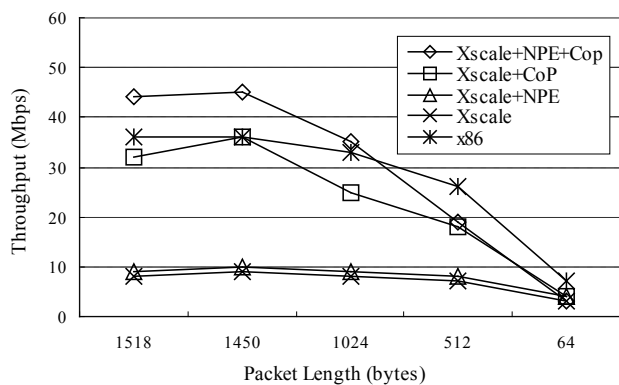


Figure 4 IPsec Throughput: the DES Case

4.3 Bottleneck Analysis

4.3.1 Bottleneck of Packet Rx/Tx

To proceed the bottleneck analysis, we considered four main functional units that could potentially affect system performance: bus, memory system, NPE and XScale. It is obvious that neither the bus nor the memory is a bottleneck because wired speed can be achieved for some larger packet lengths. The NPE is not a bottleneck either, since, as observed by the *netstat* utility, all packets are received and stored at the memory. The bottleneck can therefore be identified as the XScale since the packet processing is carried out mostly by it. Results show that the utilization of the XScale linearly advances as the traffic load increases.

4.3.2 Bottleneck of IPsec Processing

The bottleneck in the IPsec processing is known to be the XScale before offloading is applied, since the cryptographic calculation demands much computing power. However, the XScale is again found to be the bottleneck even after offloaded by the crypto coprocessors. Results show that when traffic load is 50Mbps exceeding the maximum system throughput of 46Mbps, the utilization of XScale approaches 100% and the success ratio of IPsec packets significantly drops to 22%. This can be explained by that the processor is so busy that incoming packets are dropped due to limited buffer space.

4.4 Turnaround Time Analysis of Functional Blocks

Figure 5 depicts the turnaround time analysis of the functional blocks in processing DES and 3DES packets. Functional blocks considered consist of the IP processing, IPsec preprocessing including identity and SAD/SPD lookups, and IPsec encryption. Three kinds of testbed configurations are conducted for testing DES and 3DES: IXP425 with the cryptographic operations offloaded to the coprocessor; IXP425 without offloading, namely XScale only; and PIII processor. From the figure we can see that cryptographic calculation accounts for a major portion, from 80% to 90%, in the packet processing time

before offloading. After offloading to the coprocessor, the time for cryptographic calculation is reduced from 700 μ sec to 100 μ sec. Notably both the IXP425 and single XScale configurations have the same IP processing and IP preprocessing periods because those tasks are executed only by XScale.

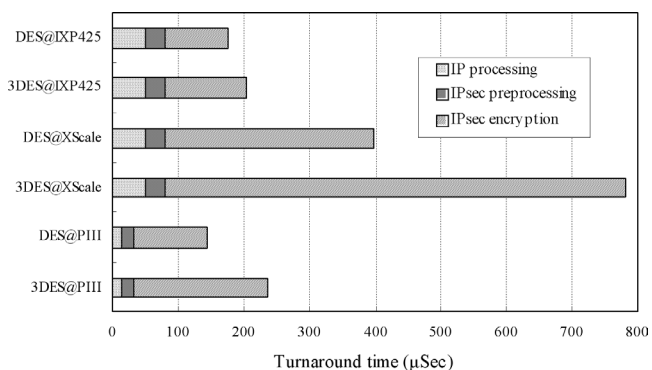


Figure 5 Turnaround Time of Functional Blocks

5 Conclusions and Future Work

In this work, we elaborate the implementation of a VPN gateway over the IXP425 network processor, where a number of coprocessors are provided for offloading computation intensive tasks from the Xscale core. We introduce the hardware and software architectures of the platform, analyze the VPN, i.e., IPsec, processing flow, and then identify the packet Rx/Tx as well as encryption/decryption as the ones to be offloaded to coprocessors. We realize the offloading design by implementing a number of drivers in NetBSD, and finally externally and internally benchmark the system in order to find possible performance bottlenecks.

The benchmark results show that the throughputs of IPsec processing and packet Rx/Tx have improvements of 350% and 60%, respectively, after offloading. However, the Xscale is again found to be the bottleneck for both packet Rx/Tx and IPsec processing.

Two issues are to be investigated in the future. First, more tasks may be offloaded to NPEs or to coprocessors. An example of this is the IPsec database lookup, which determines the policy to be applied to a certain IPsec packet. Second, the performance may be further improved if we call the related functions in the AccessLibrary directly for cryptographic operations, instead of going through the Open Crypto Framework.

References

- [1] T. Braun, M. Günter, M. Kasumi and I. Khalil, *Virtual Private Network Architecture, Technical Report IAM-99-001*, CATI, April, 1999.

- [2] M. John and S. Smith, *Application-Specific Integrated Circuits*, Addison-Wesley Publishing Company, ISBN 0-201-50022-1, June, 1997.
- [3] P. C. Lekkas, *Network Processors: Architectures, Protocols and Platforms (Telecom Engineering)*, McGraw-Hill Professional, ISBN 0071409866, July, 2003.
- [4] Intel IXP425 Network Processor, <http://www.intel.com/design/network/products/npfamily/ixp425.htm>
- [5] R. Atkinson, *Security Architecture for the Internet Protocol, RFC1825, IETF Network Working Group*, August, 1995.
- [6] T. Spalink, S. Karlin, L. Peterson and Y. Gottlieb, *Building a Robust Software-Based Router Using Network Processors, Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, 2001.
- [7] Y.-D. Lin, Y.-N. Lin, S.-C. Yang and Y.-S. Lin, *DiffServ Edge Routers over Network Processors: Implementation and Evaluation, IEEE Network, Special Issue on Network Processors*, July, 2003.
- [8] H. Bos and K. Huang, *On the Feasibility of Using Network Processors for DNA Queries, Proceedings of the 3rd Workshop on Network Processors and Applications*, February, 2004.
- [9] F. De Bernardinis, L. Fanucci, T. Ramacciotti and P. Terreni, *A QoS Internet Protocol Scheduler on the IXP1200 Network Platform, Proceedings of the 3rd IEEE International Workshop on System-on-Chip for Real-Time Applications*, June, 2003.
- [10] M. J. Rashti, H. R. Rabiee, A. Foroutan and M. Lavasani, *A Multi-Dimensional Packet Classifier for NP-Based Firewalls, Proceedings of the 2004 International Symposium on Applications and the Internet*, January, 2004.
- [11] The NetBSD Project, <http://www.netbsd.org>
- [12] L. Kesteloot, *Porting BSD UNIX to a New Platform*, January, 1995.

Biographies



Yi-Neng Lin received his BS, MS and PhD degrees all in Computer Science from National Chiao Tung University, Hsinchu, Taiwan, in 1999, 2001 and 2007, respectively. His research interests include architectural analysis of network processors and embedded systems, MAC protocol design in wireless networks, and backhaul technologies in telecom networks. Dr. Lin is now with Fiberlogic Communications, Inc.



Ying-Dar Lin received the BS degree in computer science and information engineering from National Taiwan University, Taipei, Taiwan, in 1988, and the MS and PhD degrees in computer science from the University of California, Los Angeles (UCLA), in 1990 and 1993, respectively. He joined the faculty of the Department of Computer Science, National Chiao Tung University (NCTU), Hsinchu, Taiwan, in August 1993 and has been a Professor since 1999. He was the director of the Institute of Network Engineering during 2005–2007. He is also the founder and director of Network Benchmarking Lab (NBL), cohosted by Industrial Technology Research Institute (ITRI) and NCTU since 2002, which reviews the functionality, performance, conformance, and interoperability of networking products ranging from switch, router, WLAN, to network and content security, and VoIP. In 2002, he cofounded L7 Networks Inc., which addresses the content networking markets with the technologies of deep packet inspection. At NCTU, he currently directs, or codirects, Computer and Network Center (2007), NBL (2002), Realtek-NCTU Joint Lab (2006), and D-Link NCTU Joint Lab (2007). His research interests include design, analysis, implementation, and benchmarking of network protocols and algorithms, wire-speed switching and routing, quality of services, deep packet inspection, network processors and SoCs, and embedded hardware software codesign. From 2008, he is on the editorial board of IEEE Communications Magazine and IEEE Communications Surveys and Tutorials. He was on the program committee of ICCCN'07 and a program cochair of International Computer Symposium'07.



Yaun-Cheng Lai received his PhD degree in the Department of Computer Science at National Chiao Tung University in 1997. He joined the faculty of the Department of Computer Science and Information Science at National Cheng Kung University in August 1998. He then joined the faculty of the Department of Information Management at National Taiwan University of Science and Technology in August 2001 and has been a professor since February 2008. His research interests include performance analysis, protocol design, wireless networks, and Web-based applications.



Chuan-Hung Lin is a software engineer at L7 Networks Inc. in Taiwan. He is responsible for developing the mechanism for networking detection and content management. He received his BS and MS degrees in 2002 and 2004 both in Computer Science from National Chiao-Tung University.

