Radioport: A Radio Network for Monitoring and Diagnosing Computer Systems

Hans Eberle

Radioport: A Radio Network for Monitoring and Diagnosing Computer Systems

Hans Eberle

SMLI TR-2002-117 October 2002

Abstract:

A radio network is described for configuring, monitoring, and diagnosing the components of a computer system. Such a network offers several advantages: (a) It improves the robustness of the overall system by not having the monitoring functions rely on the interconnect of the monitored system; (b) by broadcasting information, it offers direct communication between the monitoring and monitored components thereby removing dependencies inherent to hierarchical and daisy-chained wired networks; (c) it does not rely on a physical interconnect thereby lowering implementation cost, offering non-intrusive monitoring, and improving reliability thanks to the lack of error- and failure-prone cables and connectors.

This report is an extended version of a paper presented at HOTI 2002, Stanford, California, August 2002. It received the “Most Interesting New Topic Award.”

Sun microsystems

M/S MTV29-01
2600 Casey Avenue
Mountain View, CA 94043

email address:
hans.eberle@sun.com

For information regarding the SML Technical Report Series, contact Jeanie Treichel, Editor-in-Chief <jeanie.treichel@eng.sun.com>. All technical reports are available online on our Website, http://research.sun.com/techrep/.

Radioport: A Radio Network for Monitoring and Diagnosing Computer Systems

Hans Eberle
Sun Microsystems Laboratories
2600 Casey Avenue
Mountain View, CA 94043
hans.eberle@sun.com

1. Introduction

We are currently experiencing a proliferation of wireless technologies driven mainly by mobile applications. Examples are 802.11 [4][11] and the emerging Bluetooth standard [1]. Thanks to decreasing implementation cost and ease of use, wireless technologies are also becoming increasingly attractive alternatives to existing wired solutions.

In this paper, I want to explore the applicability of a wireless network and, more specifically, of a radio network inside a computer system. Given the bandwidth limitations of radio communication, I am focusing on applications based on out-of-band signaling. Examples of such applications are configuring, monitoring, and diagnosing systems and system components.

Reliability, availability, and serviceability (RAS) are becoming increasingly important properties of computer systems, in particular, when running mission-critical applications. These properties rely on reliable monitoring and diagnosing functions. In today's systems, these functions often use interconnects which are shared with the system to be monitored. Typically, these interconnects use hierarchical or, even worse, daisy-chained wiring. Thus, a failure of the monitored system can easily impact the monitoring functions making it difficult to accurately determine the state of the system's components.

The outlined dependencies can be removed by an interconnect such as the radio network described here that is orthogonal to the physical system interconnect and that further offers direct communication between the monitoring and monitored components.

In addition to improving the robustness of a system, a radio network also offers ease-of-use thanks to the lack of a physical interconnect. Since connecting to the radio network is done in a non-intrusive way, applications can benefit that traditionally required a cable to be connected to the system.

The paper is structured as follows. Section 2 describes the Radioport technology and the advantages it offers. Section 3 outlines the applications of the Radioport technology. Section 4 describes a prototype system. Section 5 addresses future work. Finally, Section 6 contains the conclusions.

2. Radioport Technology

My vision is a system containing chips each equipped with a Radioport[1] that consists of a radio transceiver and a controller.

The Radioport technology uses a radio network with a range limited by the enclosure of the monitored computer system. The radio network connects each chip of a computer system via the Radioport. Ideally, the radio transceiver and the associated control logic are fully integrated into the chip to be monitored and diagnosed, and no external glue logic is required. The only external component needed is the antenna. It can be implemented as a trace on the printed circuit board or integrated into the package of the chip.

Of course, the Radioport technology is applicable at different system levels, not just at the chip-level as described. For example, it can be associated with a field-replaceable unit (FRU) such as a system board or IO device. In this case, a single-chip implementation housing the Radioport transceiver and controller would be the preferred implementation.

While in many applications the Radioport transceiver and controller are powered by the power supply of the system to be monitored and diagnosed, other applications might prefer an implementation that uses a local battery so that the Radioport can function independently without relying on the system's power supply.

It is the applications rather than the technology that makes Radioport a novelty. Several integrated solutions are available to implement low-cost radio links [15]. A step towards a single-chip Radioport implementation is the rfPIC family [10] of microcontrollers by Microchip Technology Inc. that offers a microcontroller with an on-chip radio transmitter. The Radioport technology is further related to the radio frequency identification (RFID) technology [2][12]. RFID chips are tags containing product-identifying information that can be read by radio communication. RFID applications are manifold including the replacement of conventional product bar codes.

The many advantages offered by the Radioport technology are discussed in more detail in the following paragraphs.

Direct Communication

Radioport technology provides a direct connection between a system controller in charge of monitoring and the components to be monitored. There are no indirections that are, for example, found in architectures that use daisy-chained or hierarchical wiring. In such systems, a failure of an intermediate communication node makes any dependent node inaccessible thereby preventing the system controller from determining the state of all components.

An example of a scan path using traditional wiring techniques is shown in Figure 1a and an alternative design using wireless communication is shown in Figure 1b. In this example, the wired scan paths connecting the chips contained on a single daughter card are linked serially in the form of a daisy chain. A bridge on the mother board connects the daisy chains of the daughter cards. A large system might employ additional levels of hierarchy (not shown) in that the daisy chains are connected in a tree-like structure where the daisy chains are the leaves and the

[1] Radioport is the code name of a Sun Microsystems Laboratories internal project.

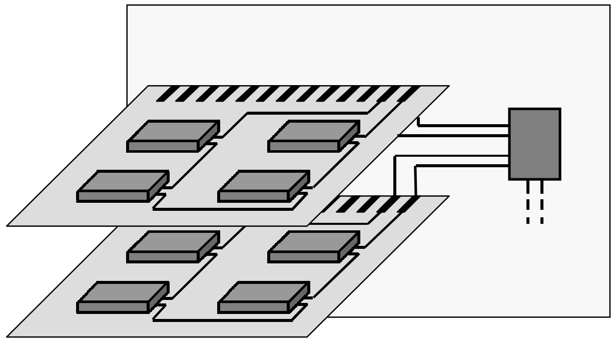

(a)

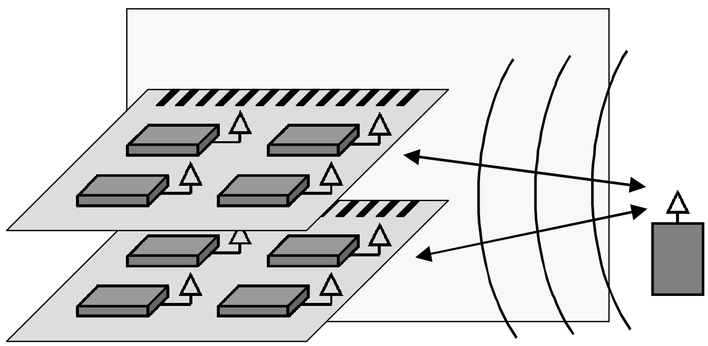

(b)

Figure 1: Wired scan path (a) vs. wireless scan path (b).

bridges are the nodes. Such a topology contains many unwanted dependencies. Figure 2 illustrates examples of possible failure scenarios that illustrate these dependencies:

- If a failure of a chip on the daughter card breaks the daisy chain, all chips that come later in the daisy chain become inaccessible even though they might be fully functional (Figure 2a).
- Assuming a hierarchy of bridges, any failure of a bridge causes the corresponding sub-tree to become inaccessible even though the nodes and leaves of the sub-tree might be fully functional (Figure 2b).
- System availability can be improved by employing redundancy. Still, if components and wires are simply replicated, as is standard practice today, similarities in failure behavior can be expected.

Non-Intrusiveness

Radioport technology is non-intrusive since it does not rely on any physical cabling. Thus, it is a particularly attractive alternative for any task that requires a physical communication link to be temporarily connected to the system. Connecting cables typically requires manual intervention which can be error-prone and time-consuming. Radioport technology removes these shortcoming. Moreover, it offers an efficient way to automate these tasks.

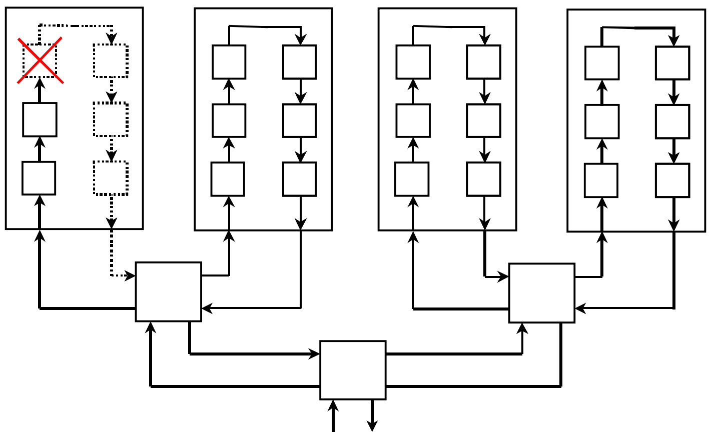

(a)

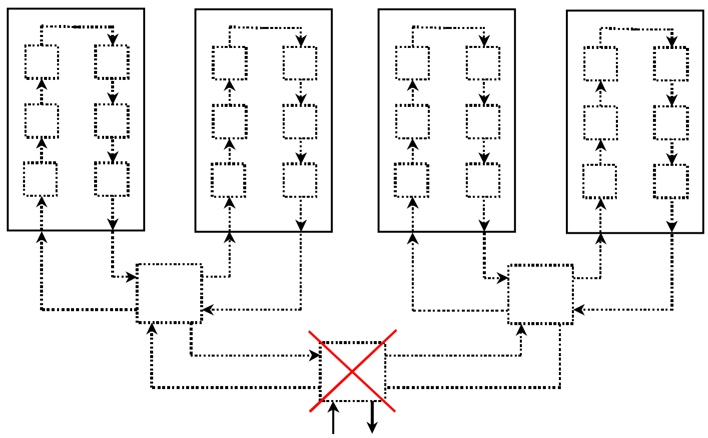

(b)

Figure 2: Failure scenarios for daisy-chained and hierarchical wiring.

Cost

Radio communication has the potential to reduce system cost since it replaces a physical interconnect based on wires and connectors.

Reliability

Physical interconnect technologies based on cables and connectors are inherently unreliable. Causes are manifold: wires are improperly connected in the first place or become loose through wear and tear, while connectors and their contacts become corroded and worn. Thus, a radio interconnect has the potential to be more reliable than a physical interconnect.

3. Applications

In this section, I give a few examples of applications that demonstrate the advantages offered by Radioport technology.

3.1. Wireless Sensors

Operating electronic components at temperatures or voltages exceeding their specified ranges can dramatically reduce their reliability and possibly cause damage. To prevent this from happening, sensors are deployed to monitor the temperature and voltage levels of chips and subsystems such as disk drives or power supplies. If the sensors report a reading outside the specified operating temperature or voltage, an alert is generated that can, for example, trigger the shutdown of a component before damage occurs, and initiate the fail-over to a redundant component before unreliable operation is observed.

Connecting the sensors by means of a wireless network rather than a wired network results in a more robust and cost-effective implementation for the reasons given above.

Moreover, wireless sensors could play a useful role during the debugging phase of the product development cycle. For example, wireless self-powered temperature sensors could be temporarily mounted in various locations and added to arbitrary components to conduct a thermal analysis.

3.2. Wireless JTAG

The JTAG standard [6] is a serial interconnect that adds testing capabilities to system components and, in particular, to chips. It is primarily intended for performing boundary scan tests that test the logic internal to the component as well as the connectivity external to the component with respect to other components. Furthermore, JTAG interfaces are widely used to configure components such as EEPROMs and FPGAs.

A wireless version of the JTAG standard offers several advantages. A wireless JTAG offers a broadcast medium that offers direct communication between the JTAG master and the JTAG slaves without the need for daisy-chained wiring or hierarchical wiring as was illustrated in Figures 1 and 2.

Considering the use of JTAG for configuring components, a wireless version offers non-intrusive operation in that no cable needs to be attached thereby making it easier to automate the task and avoiding cabling problems.

3.3. System Configuration

Increasingly, systems are being built-to-order. On the manufacturing floor, configuration management needs to be in place to track and check the configurations of the systems being assembled. Today, this task is often accomplished by using bar codes to identify system components. Often, manual intervention is required to read the bar codes.

Radioport technology offers an attractive alternative to the management of custom configurations. It automates the process of determining configurations. Furthermore, it allows for supplying much more detailed information on components than is possible with bar codes. This information can include component parameters, vendor information, and manufacturing information. In addition to reading information out of the component, Radioport technology makes it

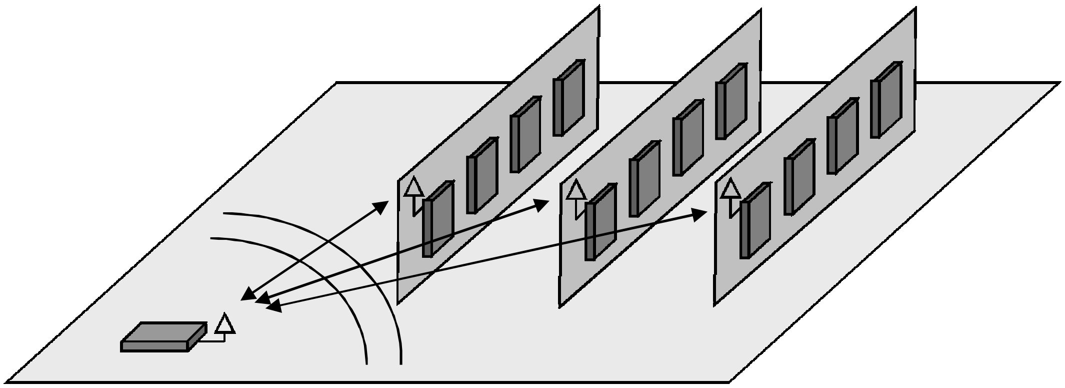

Figure 3: Memory subsystem using Radioport technology.

possible to also write information back into the component, for example, to maintain a log of the assembly process.

Configuration management has to be addressed not only on the manufacturing floor but also in the field. Before a system can be serviced, its configuration needs to be determined. Envision a service technician obtaining the configuration of the system simply by holding a PDA next to the system and reading out the configuration information via a radio link without connecting any cable or probe.

Traditionally, the configuration of a system can only be determined if the system is up and running, or, at least, if the system interconnect is operational. This need not be the case with Radioport technology.

4. Prototype System

We have built a prototype system that adds a radio link to dual inline memory modules (DIMMs). This application was chosen since it offered a cost-effective and practical way to demonstrate Radioport technology without requiring changes to the monitored system. It was not our intention to emphasize the applicability of Radioport technology to this specific application. Figure 3 shows the corresponding memory subsystem that uses Radioport technology for out-of-band signaling between a system controller and the memory modules.

A block diagram of the prototype card is shown in Figure 4. The system is implemented as an extender card that plugs into a DIMM socket of a motherboard on one side and connects to a regular DIMM card on the other side. The extender card contains a field-programmable gate array (FPGA), a microcontroller, a radio transceiver, and a three-port crosspoint switch connecting the motherboard, DIMM, and the FPGA. A picture of the prototype board is shown in Figure 5. In addition to the listed components, the picture shows the dipole antenna [9] [14] implemented with two 2.5 cm traces on the printed circuit board.

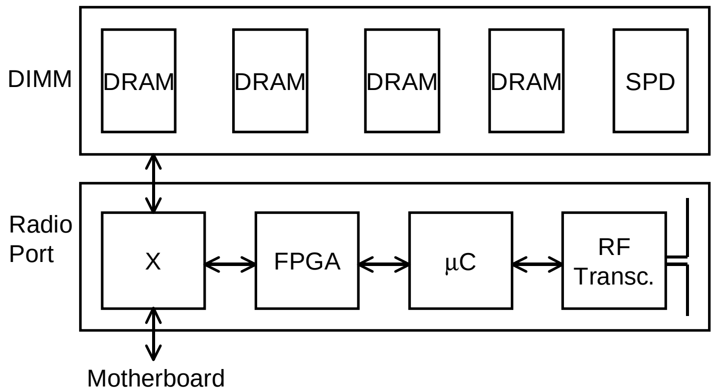

Figure 4: Prototype extender card adding a Radioport to a memory module.

The radio link is implemented with the TR1000 one-chip transceiver from RF Monolithics, Inc. [15]. The link operates at a frequency of 916.5 MHz and uses on-off keying (OOK) modulation. The transmission rate is 19200 bit/s — higher data rates are, of course, feasible. Transmitted characters are encoded according to a 8B/12B DC-balanced scheme.

The extender card communicates by radio with a system controller implemented with another extender card. Communication is based on a master/slave protocol whereby the system controller implements the master and the extender cards together with the DIMMs implement the slaves. Message exchanges use a simple request/acknowledge protocol: The master sends a command in a request packet and the slave returns a reply in an acknowledgment packet. That is, a single master initiates all packet transfers and slaves can only send packets after they were instructed to do so by the master.

The packet format is shown in Table 1. The packet size is limited to a total of 24 bytes. A short packet size was chosen to keep the packet retransmission rate low. Messages can span multiple packets. Multi-packet messages are, for example, needed to transmit long scan vectors.

Table 1: Packet format.

Field Name	Size [Byte]	Description
Preamble	2	sequence of alternating 1's and 0's; provided to train receiver PLL
Start Symbol	1	marks beginning of packet
Destination ID	0.5	0H: broadcast FH: master device ID 1H..EH: slave devices ID
Source ID	0.5	1H..FH: device ID
Packet ID	1	00H: one-packet message 01H..FFH: packet ID of a multi-packet message
Message Size	1	message length n ($n \leq 16$)
Message	n	message body
Checksum	2	ISO 3309 HDLC check sum [7]

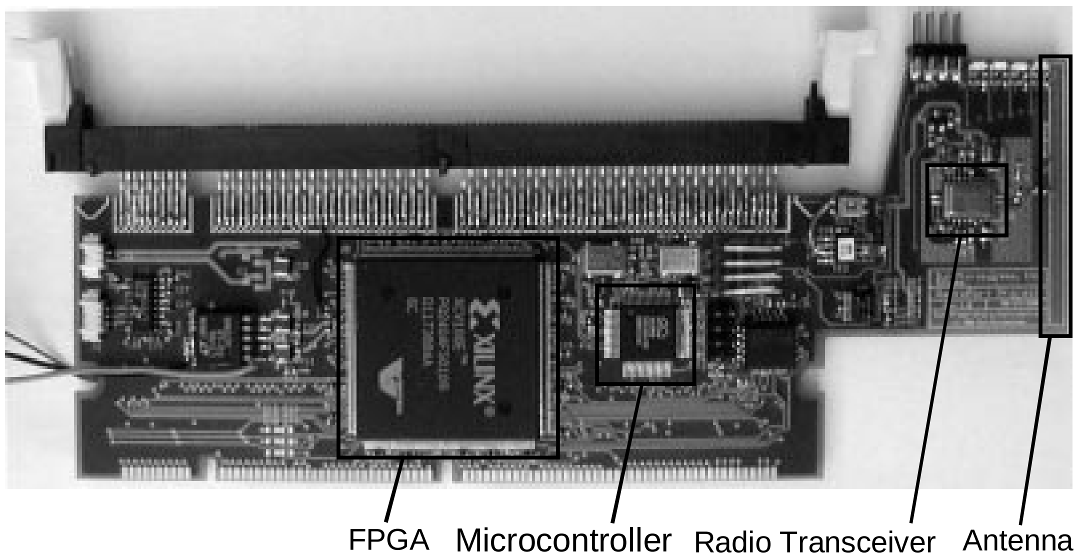

Figure 5: A prototype Radioport for DIMMs.

The extender card provides the following functions:

(1) Measure the voltages and temperatures provided by on-board sensors.

(2) Read the serial presence detect (SPD) [8] on the DIMM. The SPD is a serial ROM that contains the parameters specifying the DRAM chips and the organization of the DIMM.

(3) Perform a self-test of the DIMM. Tests consist of writing data patterns into memory locations and subsequently reading data back from these locations and checking the obtained data values [5].

(4) Perform a boundary-scan test [6]. This test involves two extender cards, one writing a test vector to the memory bus and another one reading the signal values back from the memory bus. By comparing the vectors that were written and read, the wiring of the DIMMs and the motherboard can be checked for open and shorted connections.

To execute one of the listed functions, the master broadcasts a packet containing the slave's ID and the function to be executed. The addressed slave executes the function and returns the results. Whereas functions (1) and (2) can be executed while the system is running, functions (3) and (4) can only be executed while the system is in reset state.

Other, more sophisticated access protocols might be needed, for example, if slaves should execute functions in parallel rather than sequentially or if there is no single master that can reach all slaves. Examples of such protocols are Space Division Multiple Access (SDMA), Time Division Multiple Access (TDMA), Frequency Division Multiple Access (FDMA), and Code Division Multiple Access (CDMA) [13].

Our prototype implementation easily covers a transmission range of up to a couple of meters. Initial tests do not indicate significant interference problems although more elaborate tests need to be conducted.

5. Future Work

Future work needs to address several issues. Firstly, higher integration needs to be achieved to reduce cost. Ideally, the Radioport is integrated into the chip that is being monitored and diagnosed and the antenna is integrated into the package of the chip.

Further, different RF transmission techniques should be studied with respect to signal integrity and signal interference. Ultra-wideband techniques [3] seem particularly suitable for Radioport technology as they require mostly digital logic and a minimum amount of RF electronics thereby offering low complexity and low cost. These techniques are also attractive as they are relatively immune to multipath cancellations and result in low energy densities.

Finally, more sophisticated protocols need to be explored. For example, when monitoring systems, support for multiple masters is desirable so that monitored components can initiate transfers of messages such as error messages without being polled by a single master. A further feature to be considered is support for secure communication. For instance, when configuring components, the integrity and possibly the confidentiality of the configuration data needs to be guaranteed, and the sender needs to be identified and authenticated.

6. Conclusions

An interconnect technology based on radio communication offers an attractive alternative for implementing system functions requiring out-of-band signaling. Such functions include system configuration, monitoring, and diagnostics. When implementing these functions, the Radioport technology described in this paper offers higher reliability, improved ease-of-use, and lower implementation cost than existing wired technologies.

Acknowledgments

Neil Wilhelm suggested the use of radio communication inside computer systems. José Cruz, Wyant Chan, Keith Hargrove, Mark Koch, and Scott Powers implemented the DIMM extender card. Nils Gura provided valuable feedback on the draft version of this paper.

References

[1] J. Bray, Ch. Sturman, “Bluetooth 1.1: Connect without Cables,” 2nd Edition, Prentice Hall PTR, December 2001.

[2] L. Di Jasio, “Using KEELOQ to Validate Subsystem Compatibility,” Application Note 825, Microchip Technology, www.microchip.com, 2002.

[3] J. Foerster, E. Green, S. Somayazulu, D. Leeper, “Ultra-Wideband Technology for Short- or Medium-Range Wireless Communication,” Intel Technology Journal, 2nd Quarter 2001.

[4] M. Gast, M. Loukides, “802.11 Wireless Networks”. O’Reilly and Associates, May 2002.

[5] A. van de Goor, “Testing Semiconductor Memories, Theory and Practice,” ComTex Publishing, Gouda, The Netherlands, 1998.

[6] IEEE Standards Committee, “IEEE Standard Test Access Port and Boundary-Scan Architecture,” IEEE Std 1149.1-1990, IEEE, 345 East 47th Street, New York, NY 10017-2394, July 1990.

[7] International Organization for Standardization, ISO Standard 3309-1979, “Data Communication – High-level data link control procedures – Frame structure,” 1979.

[8] JEDEC Solid State Technology Association, “Serial Presence Detect Standard,” JEDEC Standard No. 21-C, www.jedec.org.

[9] T. Lee, “The Design of CMOS Radio-Frequency Integrated Circuits,” Cambridge University Press, New York, 1998.

[10] Microchip Technology Inc., “rfPIC12C509AG/509AF – 8-Bit CMOS Microcontroller with UHF ASK/FSK Transmitter,” Data Sheet DS70031A, www.microchip.com, 2001.

[11] B. O’Hara, A. Petrick, “The IEEE 802.11 Handbook: A Designer’s Companion,” Standards Information Network IEEE Press, January 1999.

[12] S. Sarma, D. Brock, D. Engels, “Radio-Frequency Identification and the Electronic Product Code,” IEEE Micro, Nov.-Dec. 2001, pp. 50-54.

[13] T. Rappaport, “Wireless Communications, Principles and Practice,” Prentice-Hall, 1996.

[14] D. Pozar, “Microwave and RF Design of Wireless Systems,” Wiley, New York, 2001.

[15] RF Monolithics, Inc., “TR1000 Data Sheet,” www.rfm.com.

About the Author

Hans Eberle is a Senior Staff Engineer at Sun Microsystems Laboratories in Mountain View, California, where he works in the areas of networks and security. From 1993 until 1998 he was an Assistant Professor of Computer Science at ETH Zurich, Switzerland, teaching classes for EE and CS students and conducting research on switched system area networks. From 1988 until 1993 he was a Principal Engineer at the Systems Research Center of the Digital Equipment Corporation in Palo Alto, California, working on multiprocessors, ATM networks, and high-speed DES encryption. He received a diploma in Electrical Engineering and Ph.D. in Technical Sciences from ETH Zurich in 1984 and 1987, respectively.