Analog Integrated Circuits and Signal Processing, 36, 175–181, 2003


# A Dynamic Thermal Management Circuit for System-On-Chip Designs*

**HERMING CHIUEH,[1,†] JEFFREY DRAPER[2] AND JOHN CHOMA, JR.[2]**

[1]*Department of Communication Engineering, National Chiao Tung University, HsinChu 30050, Taiwan*
[2]*Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089, USA*
*E-mail: chiueh@ieee.org; draper@isi.edu; johnc@usc.edu*

*Received April 27, 2002; Revised December 12, 2002; Accepted January 25, 2003*

**Abstract.** A novel fully integrated dynamic thermal management circuit for system-on-chip design is proposed. Instead of worst-case thermal management used in conventional systems, this design yields continual monitoring of thermal activity and reacts to specified conditions. With the above system, we are able to incorporate on-chip power/speed modulation and integrated multi-stage fan controllers, which allows us to achieve nominal power dissipation and ensure operation within specification. Both architecture and circuitry are optimized for modern system-on-chip designs. This design yields intricate control and optimal mangement with little system overhead and minimum hardware requirements, as well as provides the flexibility to support different thermal mangement algorithms.

**Key Words:** thermal management, system-on-chip, VLSI system design

## 1. Introduction

Increases in circuit density and clock speed in modern VLSI designs have brought thermal issues into the spotlight of high-speed integrated circuit design. Local overheating [1] in one spot of a high-density circuit, such as CPUs and high-speed mixed-signal circuits, can cause a whole system to crash due to resulting clock synchronization problems, parameter mismatches or other coefficient changes due to the uneven heat-up on a single chip [2].

Passive heat dissipation mechanisms, such as heat sinks and fans, are widely used in system design. Recently, advanced computer systems and circuit designs have incorporated active mechanisms to detect and properly handle an over-heating event [3]. Such a capability guarantees the system will operate within a certain package temperature specification to avoid failure. The ACPI (Advanced Configuration and Power Interface) standard is an example specification for active power and thermal management in personal computer systems [4,5]. However, the ACPI standard is quite limited, as it simply supports extra control to turn on or off a cooling mechanism and shift the alert level that is fed back to the system.

As die size and power density increase in this system-on-chip (SoC) era, the management of package temperature is no longer sufficient to solve the problem. Uneven heat-up and temperature offset on chip [1,6,7] has become a major factor and limits the system performance. A good example is the recent Intel recall on Pentium III 1.13 GHz CPUs [8–10]. Recent research has focused on predicting on-chip temperature offset [1] and electro-thermal simulations [7,11,12] to provide thermal distribution information to circuit simulation to achieve more accurate circuit behavior predictions. Such research is valuable for circuit design, but post-fabrication approaches to addressing on-chip temperature offsets are also needed as die size and power density increase. Without such an approach, some circuit behavior becomes unacceptable, which makes management and control of on-chip temperature offset as important as the reduction of package temperature.

In this research, we propose a dynamic thermal management circuit to provide a watchdog for system-on-chip designs. This circuit is optimized in architecture

*This research was supported by DARPA contract F30602-98-2-0180, USA and by National Science Council grant NSC-92-2215-E-009-042-, Taiwan.
†Author to whom correspondence should be addressed. Tel.: +886-35-712121 ext. 54597, +996-918422677, Fax: +886-35-5710116.

and circuit implementation to fit system-on-chip designs. The following items describe the technical justification of the thermal management design for SoC that we take into consideration. First, since an on-chip monitoring mechanism is included, complicated electro-thermal or numerical thermal simulation [7,11,12] can be omitted. However, an analytical model [1,6,7] providing sufficient information like temperature range and quality guidelines for circuit designers is beneficial. Second, architecture and circuit implementation will be constrained to be compatible with the system's process (most likely to be a digital process), and minimum extra system resources should be used. Therefore, an interrupt-based system will be implemented, and reprogramming to provide flexibility and simplify the architecture will be supported. Finally, with respect to system integration, a complex cooling system that requires extra processing steps is not chosen, although this proposed system has the potential to cooperate with such novel micro-machining cooling methods [13]. Instead, a pure digital design for a fan controller is attractive if the circuit block is small enough. Target systems with power management can take advantage of such a cooling mechanism when combined with thermal management systems.

Given the above considerations, a circuit based on our previous research [14–16] with significant architecture enhancements is proposed. Those enhancements are described as follows. First, the number of temperature sensors has been increased to fit the need of more complex systems to monitor the temperature in several locations throughout the system. Second, the updated architecture provides simultaneous monitoring of multiple temperature sensors instead of the previous approach of single-sensor monitoring at a time. Third, circuits to monitor temperature offset between sensors and thresholds for interrupts that provide alerts other than package temperature have been added. Fourth, the threshold values have been expanded to have upper and lower limits for each sensor in order to achieve a more robust monitoring capability. Last, we have integrated a multi-channel, multi-stage fan controller, which we developed as an active cooling mechanism for maintaining a consistent package temperature.

These enhancements are aimed at solving thermal problems specific to SoC designs. The proposed thermal management subcomponents are encapsulated into a single IP block to foster use by the SoC market, in which the IP-based design approach has become very popular. The resulting discrete IP block facilitates verification of the architecture and thermal management algorithm through small low-cost test prototypes without compromising the applicability of the approach to SoC designs.

In Section 2, the function and architecture of the Thermal Management Circuit are described. In Section 3, the implementation plan of this system is addressed. A summary and conclusion follow in Section 4.

## 2. Function and Architecture

### *2.1. Architecture*

The architecture of the thermal management circuitry is divided into two portions: the thermal management circuit blocks and the system integration blocks. The former represent the designed thermal management system, and the latter represent the interface to the target system. The designed thermal management system could be applied to different SoC designs. However, the system integration portion is modified to fit different target systems as well as prototyping implementations.

The block diagram of the dynamic thermal management circuit is shown in Fig. 1. The thermal management circuit blocks are the white boxes with shadows; the gray boxes represent the system integration blocks. The function of every block is described in Section 2.1.1 and Section 2.1.2.

#### *2.1.1. Thermal Management Circuit Blocks*

- *Temperature Acquisition Unit*: This unit is simply an interface to acquire temperature from sensors. This circuit could be very different when applied to

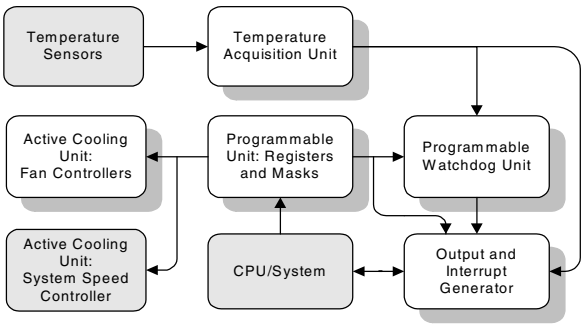


*Fig. 1.* Block diagram of the dynamic thermal management circuit.

different temperature sensors. In our prototype system, a commercial temperature sensor with one-bit serial output will be used. The major function of this circuit is to convert and latch the temperature input to parallel digital values. In our prototype design, four sensors with 16-bit precision are supported. Sensor placement will be optimized by application of the developed analytical model [1,6] to compensate for the temperature offset between the heat source and sensor. Thus, the same analytical model will also predict the maximum reading error with respect to the highest junction temperature. These offsets will be processed in the Temperature Acquisition Unit in order to provide a complete thermal analysis of the target to the thermal management system.

- *Programmable Unit*: This unit contains 8 threshold registers to program the high and low threshold values for each temperature sensor. Two threshold registers for upper and lower bounds are provided for offset temperatures between different sensors. With these threshold values, the watchdog unit can generate interrupts for desired situations. Three fan-speed registers provide the setup for the integrated fan controllers. Interrupt mask and offset mask registers indicate which interrupts should be enabled and which set of temperature sensors should be included for offset temperature monitoring. Finally, decoding circuitry and necessary configuration registers provide the communication signals between the processor and other circuit blocks.
- *Watchdog Unit*: This unit contains two monitoring circuits: the threshold monitor for each temperature sensor, and the offset temperature monitor. Both circuits are designed to minimize circuit area while providing sufficient speed to compare the sensors provided in the system.
- *Output and Interrupt Generator*: This unit provides data outputs that are read by the system CPU, like temperature value, offset temperature value, and interrupt types.
- *Active Cooling Unit-Fan Controller*: There are two active cooling units: the integrated fan controller and the system speed controller provided by the system. The integrated fan controller circuit is based on our previous pure-digital fully integrated design [15].

#### *2.1.2. System Integration Blocks*

- *Temperature Sensors*: Many kinds of temperature sensors can be used in this design. Our previous on-chip temperature sensor design is one option for system-on-chip design. For the prototype, we are using a commercial part with a system management bus interface [17].
- *CPU/System*: For pure system-on-chip design, the thermal management circuitry should be directly mapped into a CPU special-purpose register and interrupt space. In this prototype design, a memory-mapped approach will be implemented to emulate the proposed architecture. This approach also supports a flexible off-chip hardware and software platform for testing the circuit.
- *Active Cooling Unit-System Speed Controller*: For complete dynamic thermal management systems, the processor should be able to use the offset temperature data to tune the speed of different execution units to maintain the offset temperature within specification. Tradeoffs for slowing down some execution units are necessary in a critical temperature situation to prevent system failure. The mechanisms provided in the SoC implementation or processors should cooperate with this circuit to provide the function of managing the offset temperature.

### *2.2. Operating Modes*

The operation of the thermal management system can be divided into three modes from the point of view of the processor. They are the programming, data acquisition, and interrupt modes. Each mode requires different timing and data order definitions, which will be implemented in the programmable unit of the system. The basic functionality of each mode is described below.

- *Programming Mode*: This mode provides the function to program the threshold registers for temperature sensors and offset temperatures, mask registers for interrupt and offset temperature monitoring, and fan stage assignments for the integrated fan controllers. To conserve address space, the multiple temperature sensor registers will be mapped to the same address, with the configuration register contents specifying which set is actually being accessed.
- *Data Acquisition Mode*: This mode provides the capability for the processor to read data and status from the thermal management circuit in a polling fashion. Information like current temperatures, offset temperatures, setups, and interrupt status can be acquired by

the processor as often as it wishes to flexibly support different thermal management algorithms.
- *Interrupt Mode*: Interrupts are provided for designated alert conditions, and interrupt type information is also provided when the interrupt service routine reads the interrupt type register.

### *2.3. System Integration*

The thermal management circuit architecture for an SoC design is proposed in the previous sections. However to prove the validity of the complete architecture, some attention to detail must be given to the system integration blocks (gray boxes in Fig. 1). The technical decision and justification for these blocks are given here with a detailed implementation following in Section 3.

To reduce circuit complexity and die size in this area-constrained prototype chip, an off-the-shelf processor with conventional interface signals will be used for the prototype of the SoC design. In our previous design, a PowerPC interface was implemented [14], but in this prototype, a simpler memory-mapped interface will be implemented for more flexible hardware/software support. The basic function and architecture of the thermal management system are not affected since only system integration portions of the proposed designed have been modified due to the prototyping limitations as discussed in Section 2.

Instead of using on-chip temperature sensors from previous research [17], we will use external commercial parts for monitoring temperatures. Although on-chip sensors provide direct temperature readout without constraining the data transfer protocol due to pin limitations, as the number of sensors increases, the die size limit makes using on-chip sensors impractical for this prototype. Furthermore, the interface to the external sensors requires very few pins, and the sensors are not the focus of this prototype.

## 3. Prototype Implementation

A prototype implementation of the proposed design is presented in this section. Due to the limitations of an area-constrained prototype TinyChip [18] and cost of integrating the proposed IP to a complete SoC design, the prototype thermal management system is implemented separately from the processor (computer system). In Fig. 2, a detailed block diagram of a prototype design is presented. Block diagrams of the offset temperature monitor and threshold temperature monitor are shown in Figs. 3 and 4. Both designs achieve minimum area with sufficient speed to respond to system temperature changes.

As shown in Figs. 2–4, the proposed SoC IP is implemented in a discrete fashion while adhering to the IP-based design methodology. Using this approach, the proposed thermal management architecture is verified using the external temperature sensors and cooling mechanisms; such parts and processors are often treated as "hard" IPs in modern SoC design flows. Once the architecture and management algorithms are verified, this design can be easily integrated to IP-based platform design flows. The following remarks address the compatibility of the prototype implementation with the proposed architecture:

- The SoC computer system will be replaced with a hybrid design consisting of a commercial processor and a prototype TinyChip. Since the proposed design requires a special register and address mapping in the processor, a bus interface circuit between the Thermal Management Unit and processor is implemented to replace the special register and address mapping. The signal assignment between bus interface and Thermal Management Unit will accurately reflect the proposed design.
- Since the target temperature reading is on the processor, a matching hybrid temperature sensor part for the processor is used to replace the on-chip one. This situation introduces an extra System Mangement Bus Interface [17] between the Thermal Management Unit and temperature sensor used. This mechanism provides the ability to measure the target processor's temperature and does not impede the concept of the proposed design since the on-chip temperature sensor is implemented in our previous research [19], matching the qualification defined in Section 2.
- Even with the added blocks and replacement parts needed for an initial prototype implementation, the signal assignment and design specification is still valid for SoC design. The prototype implementation is simply used to verify the architecture as described in Section 2.
- In both designs of offset temperature monitor and threshold temperature monitor, the speed of the temperature sensors and speed of the programmable unit have been defined to use a single comparator circuit

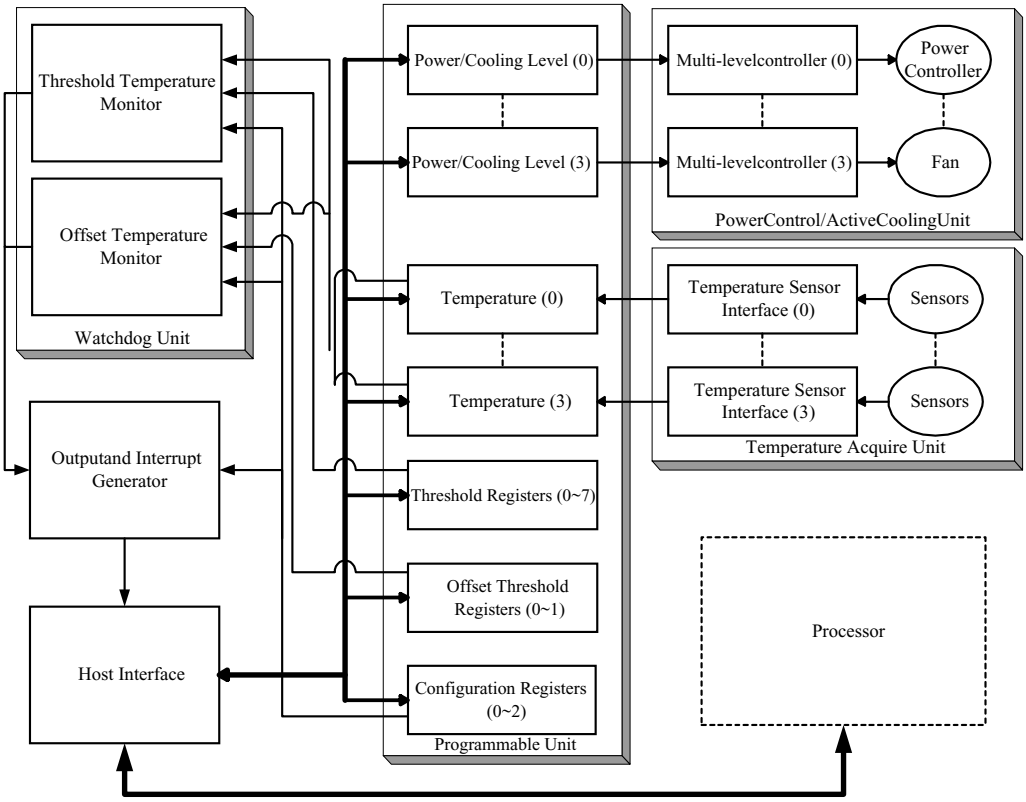


*Fig. 2.* Detailed block diagram of thermal management system.

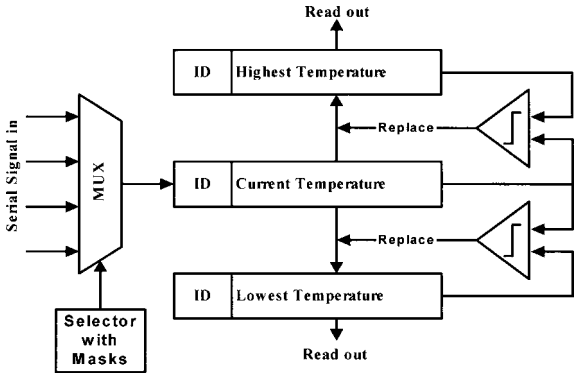


*Fig. 3.* Offset temperature monitor.

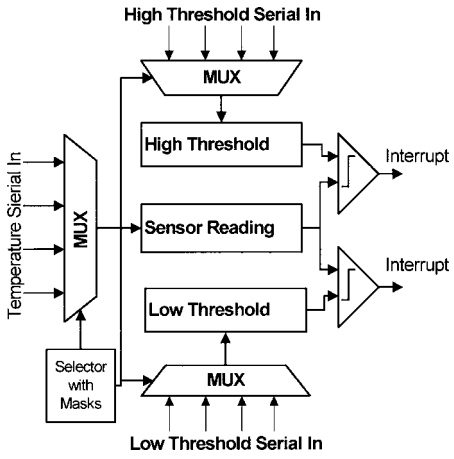


*Fig. 4.* Threshold temperature monitor.

to monitor multiple sensors using serial I/O, thus an implementation of more then 4 channels in this design can be done with very little extra circuitry.

With this sample prototype design, the proposed thermal management system for SoC design can be easily verified. Also, different approaches for a thermal management system can be easily implemented with the proposed architecture, since this system provides flexible ways for systems to read the temperature, set

the threshold value for interrupt generation, and measure temperature values from different sensors. This design can be used to implement but is not limited to the ACPI protocol. For instance, the temperature threshold can be set to any number of values to represent any number of critical situations. Fuzzy logic control and other algorithms requiring more levels of alerts can be applied. Also with the capability of actively acquiring temperature measures at any time, the CPU can verify a desired temperature response when it exectues a cooling action. With this feedback, actions like increase/reduce FAN speed and clock rates can be applied for more compex management algorithms.

## 4. Conclusion

A novel fully integrated dynamic thermal management circuit for system-on-chip design has been described. The architecture and design detail with its justification, as well as the final system integration for a complete thermal management system for SoC design was presented. The innovative temperature offset monitoring provides a mechanism for system-on-chip designs to monitor the temperature offset across the system and enhance stability. With proper handling of this information, the system not only prevents failure but also enhances performance by controlling each subcomponent's operation speed with feedback from thermal information. With minimum overhead in chip area and system resources, this design provides intricate control and optimal thermal management on chip, upon which a complete dynamic thermal management system for modern computer designs can be implemented.

## References


1. H. Chiueh, J. Draper, L. Luh, and J. Choma Jr., "A thermal evaluation of integrated circuits: On chip offset temperature measurement and modeling," in *Proc. 2nd Internationl Workshop on Design of Mixed-Mode Integrated Circuits and Applications*, 1998, pp. 109–113.
2. V. Szekely, M. Rencz, and B. Courtois, "Thermal testing methods to increase system reliability," in *Proc. 13th IEEE SEMI-THERM Symposium*, 1997, pp. 210–217.
3. J. Draper, J. Block, J. Koller, and C. Steele, "Thermal management in embedded systems using MEMS," in *Proc. Lecture Notes in Computer Science 1388 (IPPS/SPDP'98 Workshops Proceedings)*, 1998, pp. 900–901.
4. Compaq, Intel, Microsoft, Phoenix, and Toshiba, "Advanced configuration and power interface specification," July 27, 2000.
5. J. Steele, "ACPI thermal sensing and control in the PC," in *Proc. Wescon 98*, 1998, pp. 169–182.
6. H. Chiueh, J. Draper, L. Luh, and J. Choma Jr., "A novel model for on-chip heat dissipation," in *Proc. The 1998 IEEE Asia-Pacific Conference on Circuits and Systems*, 1998, pp. 779–782.
7. C.-H. Tsai and S.-M. Kang, "Substrate thermal model reduction for efficient transient electrothermal simulation," in *Proc. 2000 Southwest Sympoium on Mixed-Signal Designs*, 2000, pp. 185–190.
8. I. Fried, "Glitch prompts Intel to recall 1.13-GHz Pentiums," http://news.cnet.com.
9. I. Fried, "Hardware sites help Intel isolate chip problem," http://news.cnet.com.
10. S. Musil, "The week in review: Intel hits a speed bump," http://news.cnet.com.
11. J. W. Sofia, "Analysis of thermal transient data with synthesized dynamic models for semiconductor devices," in *Proc. 10th IEEE SEMI-THERM*, 1994, pp. 78–85.
12. V. Szekely, A. Poppe, A. Pahi, A. Csendes, G. Hajas, and M. Rencz, "Electro-thermal and logic-thermal simulation of VLSI designs." *IEEE Transactions on VLSI Systems* 5, pp. 258–269, 1997.
13. Goodson, Santiago, T. W. Kenny, Carruthers, and Towe, "Electrokinetic micro coolers," presented at *International Interconnect Technology Conference*, San Francisco, CA, 2000.
14. H. Chiueh, J. Draper, and J. Choma, Jr., "A programmable thermal management interface circuit for powerPC systems," in *Proc. 6th International Workshop on Thermal Investigation of ICs and Systems*, 2000.
15. H. Chiueh, L. Luh, J. Draper, and J. Choma Jr., "A novel fully intergrated fan controller for advanced computer systems," in *Proc. Southwest Symposium on Mixed-Signal Design*, 2000, pp. 191–194.
16. H. Chiueh, J. Draper, and J. Choma Jr., "Implementation of a temperature monitor interface circuit for powerPC systems," in *Proc. The 43rd Midwest Symposium on Circuits and Systems*, 2000.
17. SBS Implementers Forum, "System management bus (SMBus) specification," August 3, 2000.
18. MOSIS. http://www.mosis.com.
19. L. Luh, J. Choma Jr., J. Draper, and H. Chiueh, "A high-speed CMOS on-chip temperature sensor," in *Proc. European Solid-State Circuits Conference (ESSCIRC99)*, 1999, pp. 290–293.


**Herming Chiueh** received the B.S. degree from the Department of Electrophysics, National Chiao Tung University, Hsin-Chu, Taiwan in 1992, and the

M.S. and Ph.D. degrees from Department of Electrical Engineering, University of Southern California, Los Angeles, U.S. in 1994 and 2002. From 1996–2002, he was with Information Sciences Institute, University of Southern California, Marina del Rey, California, U.S. Currently, he is an Assistant Professor, Department of Communication Engineering, School of Electrical Engineering and Computer Science, National Chiao Tung University, Hsin-Chu, Taiwan. His research interests include system-on-chip design methodology, thermal management for VLSI, and power-aware integrated circuits.

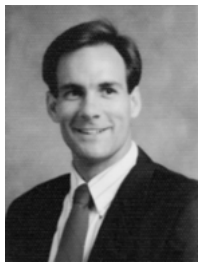

**Jeffrey Draper** holds a joint appointment as a Research Assistant Professor in the Department of Electrical Engineering at University of Southern California and a Project Leader in the Computational Sciences Division at USC Information Sciences Institute. Dr. Draper has led the VLSI effort on several large projects in the past 5 years and most recently directed the development of a 55-million transistor processing-in-memory (PIM) chip. Dr. Draper received his Ph.D. in Computer Engineering from the University of Texas at Austin. His research interests are PIM architectures, VLSI, thermal management parallel computer architectures, and interconnection networks.

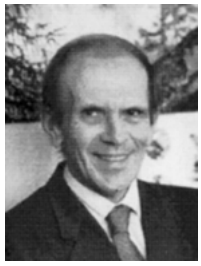

**John Choma** earned his B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Pittsburgh in 1963, 1965, and 1969, respectively. He is Professor of Electrical Engineering at the University of Southern California, where he teaches undergraduate and graduate courses in electrical circuit theory and analog integrated electronics. Prof. Choma consults in the areas of broadband analog and high-speed digital integrated circuit analysis, design, and modeling.

Prior to joining the USC faculty in 1980, Prof. Choma was a senior staff design engineer in the TRW Microelectronics Center in Redondo Beach, California. His earlier positions include technical staff at Hewlett-Packard Company in Santa Clara, California, Senior Lecturer in the Graduate Division of the Department of Electrical Engineering of the California Institute of Technology, lectureships at the University of Santa Clara and the University of California at Los Angeles, and a faculty appointment at the University of Pennsylvania.

Prof. Choma, the author or co-author of some 135 journal and conference papers and the presenter of more than sixty invited short courses, seminars, and tutorials, is the 1994 recipient of the Prize Paper Award from the IEEE Microwave Theory and Techniques Society. He is the author of a Wiley Interscience text on electrical network theory and a forthcoming text on integrated circuit design for communication system applications. Prof. Choma has contributed several chapters to five edited electronic circuit texts, and he was an area of editor of the IEEE/CRC Press Handbook of Circuits and Filters.

Prof. Choma has served the IEEE Circuits And Systems Society as a member of its Board of Governors, its Vice President for Administration, and its President. He has been an Associate Editor and Editor 'In' Chief of the IEEE Transactions On Circuits And Systems, Part II. He is an Associate Editor of the Journal of Analog Integrated Circuits and Signal Processing and a former Regional Editor of the Journal of Circuits, Systems, and Computers.

A Fellow of the IEEE, Prof. Choma has been awarded the IEEE Millennium medal and has received three awards from the IEEE Circuits and Systems Society; namely, the Golden Jubilee Award, the 1999 Education Award, and the 2000 Meritorious Service Award. He is also the recipient of several local and national teaching awards. Prof. Choma is a Distinguished Lecturer in the IEEE Circuits And Systems Society.