A low-power transmission-gate-based 16-bit multiplier for digital hearing aids

Conference Paper**Author(s):**

Carbognani, Flavio; Buergin, Felix; Felber, Norbert; Kaeslin, Hubert; Fichtner, Wolfgang

Publication date:

2008

Permanent link:

<https://doi.org/10.3929/ethz-b-000014831>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

Analog Integrated Circuits and Signal Processing 56(1-2), <https://doi.org/10.1007/s10470-007-9086-0>

A low-power transmission-gate-based 16-bit multiplier for digital hearing aids

Flavio Carbognani · Felix Buergin · Norbert Felber ·
Hubert Kaeslin · Wolfgang Fichtner

Received: 15 September 2006 / Revised: 11 January 2007 / Accepted: 29 June 2007 / Published online: 9 October 2007
© Springer Science+Business Media, LLC 2007

Abstract The most widespread 16-bit multiplier architectures are compared in terms of area occupation, dissipated energy, and EDP (Energy-Delay Product) in view of low-power low-voltage signal processing for digital hearing aids and similar applications. Transistor-level simulations including back-annotated wire parasitics confirm that the propagation of glitches along uneven and re-convergent paths results in large unproductive node activity. Because of their shorter full-adder chains, Wallace-tree multipliers indeed dissipate less energy than the carry-save (CSM) and other traditional array multipliers ($6.0 \mu\text{W}/\text{MHz}$ versus $10.9 \mu\text{W}/\text{MHz}$ and more for $0.25 \mu\text{m}$ CMOS technology at 0.75 V). By combining the Wallace-tree architecture with transmission gates (TGs), a new approach is proposed to improve the energy efficiency further ($3.1 \mu\text{W}/\text{MHz}$), beyond recently published low-power architectures. Besides the reduction of the overall capacitance, minimum-sized transmission gate full-adders act as RC-low-pass filters that attenuate undesired switching. Finally, minimum size TGs increase the V_{dd} to ground resistance, hence decreasing leakage dissipation (0.55 nW versus 0.84 nW in CSM and 0.94 nW in Wallace).

Keywords Hearing aids · Low power · Multiplier · Transmission gate · Switching activity

1 Introduction

The strict specifications for digital hearing aids impose tight constraints on both area occupation and energy consumption, while relaxing the timing requirements. Thanks to their compact size and relatively large energy storage [1] (around 100 mAh), zinc-air batteries dominate the market of the energy suppliers for these applications. In order to decrease the dissipation and to exploit the battery lifetime fully, digital signal processing units, DSPs or ASICs, are supplied down to 0.9 V or even lower [1, 2], depending on the technology.

Multipliers represent at the same time basic modules and relevant energy sinks in digital signal processing; for this reason, many efforts have been spent in trying to increase their efficiency. Many works in the past, [3, 4] among others, focused on decreasing the consumption of the single full-adder cell. While this improves the overall multiplier efficiency, it overlooks the most significant limitation to the design of low-power multipliers, namely the large extent of spurious activity, due to unevenly long and re-convergent paths [1]. All the solutions proposed in literature so far to suppress glitches can be classified in three main families:

- (1) the shortening of full-adder chains in the multiplier matrix that penalizes, for instance, array architectures over Wallace trees [5];
- (2) the equalization of the internal delays, as in the Leapfrog architecture [6];
- (3) the alignment of the internal signals by means of self-timed circuits [7].

The first approach is simple and attractive; its efficiency is confirmed in this work and the new proposed multiplier is also based on the Wallace-tree architecture. As opposed

F. Carbognani (✉) · F. Buergin · N. Felber ·
H. Kaeslin · W. Fichtner
D-ITET, Integrated Systems Laboratory, ETH Zurich,
Gloriastrasse 35, Zurich 8092, Switzerland
e-mail: carbo@iis.ee.ethz.ch

to that, the efficiency of the second and the third techniques is strongly dependent on the specific technology; they require a very careful transistor-level design and dedicated calibration. The third approach is analyzed in more detail in Sect. 6, where a direct comparison to the proposed multiplier is drawn.

A new idea is presented in this paper; it may be considered as a fourth approach. Transmission gates are used to act as RC-low-pass filters that suppress glitches while reducing the overall capacitance. It will also be shown that minimum-size transmission gates present the relevant feature of reducing the leakage consumption. Meanwhile, the EDP (Energy-Delay Product) is kept quite low, which preserves flexibility in the choice of the working frequency (up to few MHz).

The rest of this paper is organized as follows: Section 2 presents a short introduction to the operation of multiplication itself for a better understanding of the different architectures, which are reviewed and compared in Sect. 3. A method to estimate the spurious activity is presented in Sect. 4. It allows the investigation of the reasons of the uneven energy dissipation in the various architectures. In Sect. 5 the new structure is introduced. Its performance is the topic of Sect. 6. Section 7 investigates the benefits of the mixed topology of CMOS and transmission gate full-adders, while Sect. 8 draws the conclusions.

2 Signed two's complement multiplication

The following notation for a 16-bit unsigned multiplication [8] is used in the next section:

$$Z = \sum_{j=0}^{15} \left(\sum_{i=0}^{15} (X_i Y_j) 2^{(i+j)} \right) \quad (1)$$

where Z represents the product, X_i the i -th bit of the multiplicand and Y_j the j -th bit of the multiplier.

The modified Baugh-Wooley algorithm (see paragraph 11.3 of [9]) enables the extension of unsigned to signed multiplication.

When the Booth radix-4 recoding [10] is applied, Eq. 1 transforms into the following:

$$Z = \sum_{j=0}^7 \left(\sum_{i=0}^{15} (X_i Y_{b_j}) 2^{(i+2j)} \right); Y_{b_j} \in \{-2, \dots, 2\} \quad (2)$$

where Y_{b_j} represents the j -th operand of the multiplier after Booth recoding. As can be noticed from Eq. 1 and Eq. 2, Booth recoding allows the number of partial products, hence the number of additions, to be halved. Yet, the precalculation of Y_{b_j} and the multiplication of the multi-

plicand by -2 , -1 and 2 require extra logic, which is paid for in terms of power dissipation (and area occupation).

Depending on the way the terms $X_i Y_j$ in Eq. 1 and the terms $X_i Y_{b_j}$ in Eq. 2 are generated and summed up, various multiplier architectures are possible. In the next section, the RCM (Ripple-Carry Multiplier), the CSM (Carry-Save Multiplier), the CSM featuring Booth recoding and the Wallace-tree architecture [8] will be examined and compared to each other. The first one is the most straightforward extension of the concept of ripple-carry adder to the multiplier architecture. The last three represent the most widespread multiplier architectures in a variety of applications [11].

3 Review of multiplier architectures

RCM and CSM differ only in the carry propagation mechanism [8], as illustrated in Fig. 1. In the RCM, the carry bits ripple from right to left, from the LSB to the MSB; in the CSM, they further descend one row. CSM has also been implemented with radix four Booth recoding.

The Wallace-tree multiplier is substantially different: the partial terms $X_i Y_j$ of Eq. 1 are added all at once before entering the main full-adder network [8]. That results in a quite irregular architecture (Fig. 2), which, nevertheless, makes it possible to shorten the longest path up to the final addition. The latter can be carried out according to well known adder topologies; in this work, a final RCA (Ripple Carry Adder) has been chosen, trading speed for energy, as shown in [12].

The four multiplier architectures were simulated with Spectre at transistor-level in a 0.25 μm CMOS technology. In order to satisfy the requirements imposed by hearing aid applications, the supply voltage has been fixed at 0.75 V, less than one third the nominal voltage (2.5 V). All multipliers are fed with the same set of random stimuli vectors; the input signals are aligned in time and they are glitch-free. These assumptions can be considered valid with very good approximation for time-shared units, where input data emanate from registers. All the presented architectures

Fig. 1 Sections of two basic multiplier architectures: Ripple-Carry multiplier (left) and Carry-Save multiplier (right)

Fig. 2 Section of the Wallace-tree multiplier

have undergone the place and route phase. Standard cell density has been set at 90%, to minimize the effect of wire parasitics. These have been back-annotated prior to transistor-level simulations for more accurate results.

3.1 Area

Table 1 shows a considerable homogeneity both in the number of gate-equivalents and area occupation for all the investigated architectures. The area in the table represents the final occupation on silicon after place and route, with 90% standard cell density.

The largest structure is CSM featuring Booth recoding, due to the overhead of the encoder, which is only partially compensated by the reduction of the full-adder matrix.

3.2 Energy

In Table 2 the dynamic and the static dissipations are shown. A quick look at Table 1 and Table 2 reveals that the substantial homogeneity in the area occupation is not reflected in the dynamic energy consumption: the energy varies considerably. RCM is the most energy hungry implementation, while Wallace represents by far the most efficient architecture. The reasons for this relevant inequality will be investigated in the next section.

As multipliers are computation-intensive arithmetic units, the proportion of the static consumption over the total dissipation is, as expected, negligible (around 0.1‰ at 1 MHz). Yet, possible long stand-by phases and the fast

Table 2 Comparison in terms of dynamic energy and static power consumption

Architecture	Dynamic energy [$\mu\text{W}/\text{MHz}$]	Static power [nW]
RCM	14.6	0.88
CSM	10.9	0.84
CSM_Booth4	11.0	1.08
Wallace	6.0	0.94

Table 3 Comparison in terms of worst-case delay and EDP

Architecture	Delay [ns]	EDP [$\text{aJ} \cdot \text{s}$]
RCM	234	3.42
CSM	142	1.55
CSM_Booth4	125	1.38
Wallace	91	0.55

progress in the fabrication technology demand special attention to be drawn towards leakage. The static dissipation is almost equally distributed among the various architectures. CSM featuring Booth recoding exhibits more leakage as a consequence of the increased amount of cells.

3.3 Delay and EDP

Table 3 collects the information about the worst-case delays and the EDPs of the studied multiplier architectures. Without surprise, the Wallace tree presents the shortest delay and the lowest EDP, while the RCM represents the slowest multiplier realization.

4 Spurious activity

The total switching activity (α_{TOT}) is proportional to the dynamic dissipation and indicates the average number of transitions per computation period that a signal undergoes. Note that in the single-edge-triggered clocking strategy, the clock has a switching activity of two.

The total activity α_{TOT} can be seen as the sum of two contributions:

$$\alpha_{TOT} = \alpha_F + \alpha_S \quad (3)$$

where α_F denotes the functional activity and α_S the spurious activity, caused by glitches.

In what follows, the value of α_F has been determined by monitoring the total switching activity of a gate-level simulation, after having forced the propagation delay of all standard cells to zero. As a matter of fact, the spurious activity is the consequence of different delays in reconvergent paths; hence, by forcing the cell delays to zero,

Table 1 Comparison in terms of number of cells, number of GEs and area occupation on silicon (90% cell density)

Architecture	No. of cells	No. of GEs	Area [μm^2]
RCM	512	2053	43000
CSM	497	2000	42000
CSM_Booth4	663	2037	45000
Wallace	531	2045	43000

Eq. 3 reduces to the mere functional activity ($\alpha_{TOT} = \alpha_F$), which can be easily determined. The value of α_{TOT} is instead deduced from transistor-level simulations, by counting the number of crossings through $V_{DD}/2$. The average spurious activity α_S results then by subtraction. This way of proceeding is just a first order approximation. A more accurate approach, based on statistics, will be sketched in Sect. 7.

Assuming that all the input signals arrive at the same time, the spurious activity in the multipliers originates from:

- (1) the different propagation delay of sum and carry bit inside the full-adders;
- (2) the uneven collection of the $X_i Y_j$ terms in the full-adder matrix;
- (3) the irregularity of the multiplier architecture.

If the first point appears in all the above-mentioned architectures, the second one refers only to the CSM structures, whereas the Wallace tree suffers more from the third one. Table 4 demonstrates however that not all the discussed points play the same role in the generation of spurious activity. The table collects the activities in the traditional architectures, averaged over all the sum and carry bits, and gives valuable indications on the glitch propagation along the full-adder chains. First of all, it shows that the functional switching activity is practically the same in all structures. In RCM, CSM and Wallace almost the same number and type of standard cells are instantiated. The differences in the dynamic consumption are proportional, with good approximation, to the differences in the total activities of Table 4. It follows that the improved efficiency of Wallace is nothing else than the consequence of a full-adder matrix organization that is less prone to glitch generation. Namely, the assets are the collection of all AND terms before the main matrix and the shortening of the full-adder chains without generating additional glitches (as the encoder at the top of CSM featuring Booth recoding).

In the end, the Wallace tree is the most efficient multiplier among the analyzed traditional architectures, but it still suffers from a $\alpha_S/\alpha_{TOT} = 39\%$ spurious activity contribution over the total according to simulations.

Table 4 Average functional, spurious and total activity in traditional multiplier architectures

Architecture	α_F	α_S	α_{TOT}
RCM	0.41	1.40	1.81
CSM	0.42	0.88	1.30
CSM_Booth4	0.42	1.07	1.49
Wallace	0.40	0.25	0.65

5 New proposed architecture

Tables 1, 2, 3 depict a very clear scenario: in the given technology, among the traditional parallel multiplier architectures, the Wallace tree guarantees the most efficient choice in terms of both energy consumption and EDP, without inflating the area. Therefore, the new low-power topology is based on the Wallace tree architecture. The investigation has been focused on two main aspects:

- (1) decreasing the spurious switching activity further;
- (2) reducing the internal node capacitances.

The low-power full-adder cell of the given technology library incorporates MOS transistors with a width three to five times the minimum size. While oversized devices increase the maximum load that can be driven, and decrease the propagation delays, they do not appear mandatory in very-low-frequency applications. This is especially true inside the multiplier full-adder matrix, where the fanout of each sum/carry bit driver is always one. The only necessary precaution is then to keep the parasitic capacitances low, which is easily manageable during place and route.

A solution to decrease both the spurious activity and the overall switched capacitance is the usage of transmission gate cells with minimum size transistors in the multiplier architecture.

5.1 Transmission gate full-adder

Figures 3 and 4 show the schematic and the layout of the employed transmission gate full-adder [3]. Despite the fewer components (18 against 28) and the reduced transistor size (all minimum width), the transmission gate cell is only slightly smaller than the original CMOS mirror-adder (length of $13.5\ \mu\text{m}$ against $14.4\ \mu\text{m}$). A minor drawback of the implementation of transmission gate cells

Fig. 3 Schematic of the implemented transmission gate full-adder cell

Fig. 4 Layout of the implemented transmission gate full-adder cell

is actually the difficult integration in a dense layout. This is a consequence of the limitations imposed by DRC rules on the proximity of adjacent contacts on active regions, which are very frequent in such cells.

5.2 CMOS-transmission gate mixed topology multiplier

The proposed multiplier is based on the Wallace-tree architecture. The AND gates are implemented in level-restoring CMOS logic to decouple the multiplier from the input drivers. The full-adder cells in the final RCA are also level-restoring CMOS cells to preserve the driving capability. Therefore, from outside the same electrical behavior is maintained. All other stages make use of transmission gate cells. The proposed structure will be called TG-mult in the next sections.

Fig. 5 Schematic of the implemented latch-adder cell proposed in [1]. Dashed components represent the additions to the mirror adder

6 Results

6.1 A recent benchmark

The benchmark is a recently proposed multiplier architecture [1] that aligns all sum- and carry-bits at each partial product stage, by means of a self-timed circuit (a simple delay line). Special cells that implement the function of full-adders and latches at the same time are used, the so called latch-adders shown in Fig. 5. They enhance the mirror adder architecture by means of two pairs of enable transistors that can bring the cell into high impedance whenever data have to be retained in the semi-static latches located at the outputs.

The architecture proposed by Chong et al. in [1] is based on the CSM with Booth radix-4 recoding and constitutes a representative example of the third approach, mentioned in Sect. 1. The transistor sizes have been scaled down linearly to the employed 0.25 μm process.

6.2 Results: dynamic power

The final results (see Table 5) confirm the energy efficiency of [1] against the CSM, as stated by the authors. The key point is a remarkable suppression of spurious activity (as shown in Table 6), which actually survives only in the final RCA.

Yet, it presents significant difficulties in calibrating the delay line properly, such as to suppress the glitches while avoiding an excessive EDP. More than that, a large load overhead must be switched twice per period, namely all the transistor gate capacitances directly driven by the self-timed circuit. The Wallace tree is indeed more energy efficient than [1].

Table 5 Final results

Architecture	Area [μm^2]	Dynamic energy [$\mu\text{W}/\text{MHz}$]	Delay [ns]	EDP [$\text{aJ} \cdot \text{s}$]
CSM	42000	10.9	142	1.55
Wallace	43000	6.0	91	0.55
Chong05 [1]	49000	6.3	278	1.75
TG-mult	41000	3.1	223	0.69

Table 6 Average functional and spurious activity in various multiplier architectures

Architecture	α_F	α_S	α_{TOT}
CSM	0.42	0.88	1.30
Wallace	0.40	0.25	0.65
Chong05 [1]	0.42	0.03	0.45
TG-mult	0.40	0.19	0.59

The proposed multiplier is less aggressive towards glitches (about 10% α_{TOT} reduction compared to Wallace), but no energy overhead is necessary. The major drawback is the large increase in the delay, compared to the Wallace-tree architecture (223 ns against 91 ns), because of the inferior driving capabilities of TGs, especially at such low supply voltages. Yet, the maximum working frequency, about 4.5 MHz, is still compatible with the addressed applications, i.e., hearing aids. The related EDP is the second best in Table 5. In terms of area occupation it is the most compact implementation.

6.3 Results: static power

In terms of leakage dissipation, TG-mult is the only architecture that can guarantee a significant improvement. This is due to the enhanced overall resistance between V_{dd} and ground. Figure 6 demonstrates the final 35% leakage reduction compared to CSM, the most leakage efficient traditional architecture.

Fig. 6 Static dissipation in various multiplier architectures

7 Discussion

Two are the winning points of the proposed multiplier:

- (1) the reduction of the overall capacitance;
- (2) the filtering of glitches.

The first feature is the direct consequence of the downscale of the CMOS to minimum size. In low-voltage applications this generally leads to an increased efficiency, as shown in [13].

The second point is actually more interesting. A conducting transmission gate behaves, in first approximation, as a resistance (see Chapter 4 in [8]), whose value is not strongly dependent on the working region of the two transistors. Figure 7 shows that in the applied technology it reaches few hundreds of kilo-ohms. By connecting it to a typical node load of 10 fF, an equivalent RC filter with a time constant of few nanoseconds results. When several transmission gates are cascaded, the time constant easily reaches several nanoseconds, enough to filter out many glitches. In Fig. 8 the same internal carry bit signal is evaluated over a short time window in both the Wallace-tree architecture and the proposed TG-mult. The graphs show that the number of full-excursion glitches is two in the standard Wallace-tree implementation and goes to zero in TG-mult, thanks to the much slower signal transitions. The large channel resistances also contribute to the reduction of the leakage current.

Answering the question about the proportion of energy savings actually due to RC low-pass filtering is not difficult, at least in very first approximation. Table 5 shows an almost 50% dissipation gain towards Wallace. Table 6 indicates about 10% α_{TOT} decrease compared to Wallace. Combining these two pieces of information gives the answer: about

Fig. 7 Equivalent resistance associated with the implemented transmission gate

Fig. 8 Effects of transmission gate low pass filtering in the TG-mult (bottom), compared to the Wallace-tree architecture (top)

20% of the total energy savings are actually due to RC low-pass filtering, the remainder being mainly a consequence of the overall capacitance reduction. The model proposed in Sect. 4 is just a rough approximation. A more accurate modeling of the spurious activity should take into consideration the distribution of the glitch voltage excursion. Figure 9 shows the distribution averaged over a few sample signals in Wallace and TG-mult. The RC filtering presents a typical effect of lowering the mean value and enlarging the standard deviation of the glitch voltage swing.

As a matter of fact, the RC low-pass filtering increases the signal transition times; yet, no consequences are expected in the TG-mult final CMOS-based RCA. At such low voltages the contribution of the cross-over currents to the overall energy is at all effects negligible [14].

Fig. 9 Distributions of the glitch voltage excursion in Wallace (top) and TG-mult (bottom)

8 Conclusion

Various multiplier architectures have been reviewed and compared in terms of area occupation, energy consumption and EDP.

On the basis of these results, a new CMOS-transmission-gate mixed-topology multiplier has been presented. Its superior dynamic and static energy efficiency has been demonstrated in comparison to established architectures and a recently published low-power multiplier.

In particular, the capabilities of transmission gates to act as low-pass filters and to suppress glitches have been investigated. The proper combination of CMOS and transmission gate inside macro-cells enables large energy savings.

A final remark in the context of low-power low-voltage is the following. Gates that switch rapidly are more glitch prone before settling to their final state and, hence, they are likely to waste energy. It seems more desirable to allow ramp times to grow larger, e.g., through the instantiation of transmission gates, so that consecutive signal edges running in opposite directions do not result in full-swing or close-to-full-swing voltage excursions (see Fig. 9).

References

- Chong, K.-S., Gwee, B.-H., & Chang, J. S. (2005). A micropower low-voltage multiplier with reduced spurious switching. *IEEE Transactions on VLSI Systems*, 13, 255–265.
- Mosch, P., van Oerle, G., Menzl, S., Rougnon-Glasson, N., van Nieuwenhove, K., & Wezelenburg, M. (2000). A 660-μW 50-Mops 1-V DSP for a hearing aid chip set. *IEEE Journal of Solid-State Circuits*, 35(11), 705–712.
- Alioto, M., & Palumbo, G. (2002). Analysis and comparison on full adder block in submicron technology. *IEEE Transactions on VLSI Systems*, 10, 806–823.
- Chang, J.-H., Gu, J., & Zhang, M. (2005). A review of 0.18-μm full adder performances for tree structured arithmetic circuits. *IEEE Transactions on VLSI Systems*, 13, 686–695.
- Bisdounis, L. (2002). *Designing CMOS Circuits for Low Power* (ch 5). Netherlands: Kluwer Academic Publishers.
- Mahant-Shetti, S. S., Balsara, P. T., & Lemonds, C. (1999). High performance low power array multiplier using temporal tiling. *IEEE Transactions on VLSI Systems*, 7, 121–124.
- Carbognani, F., Buerger, F., Felber, N., Kaeslin, H., & Fichtner, W. (2006). A self-timed 16-bit multiplier for low-power low-frequency applications. In *Proceedings of the Midwest Symposium on Circuits and Systems (MWSCAS'06)*, (pp. 433–437), San Juan, Puerto Rico.
- Rabaey, J. M. (1996). *Digital Integrated Circuits*. Upper Saddle River, NJ: Prentice Hall.
- Parhami, B. (2000). *Computer Arithmetic*. New York, NY: Oxford University Press.
- Booth, A. D. (1951). A signed binary multiplication technique. *Quarterly Journal of Mechanics and Applied Mathematics*, 4, 236–240.
- Yeo, K.-S., & Roy, K. (2005). *Low-Voltage, Low-Power VLSI Subsystems*. New York, NY: McGraw-Hill.
- Meier, P. C. H., Rutenbar, R. A., & Carley, L. R. (1996). Exploring multiplier architecture and layout for low power. In *Proceedings of IEEE Custom Integrated Circuits Conference (CICC'96)*, (pp. 513–516), San Diego, CA, USA.

13. Buergin, F., Carbognani, F., Felber, N., Kaeslin, H., & Fichtner, W. (2006). 29% power saving through semi-custom standard cell re-design in a front-end for hearing aids. In *Proceedings of the Midwest Symposium on Circuits and Systems (MWSCAS'06)*, (pp. 610–614), San Juan, Puerto Rico.
14. Carbognani, F., Buergin, F., Felber, N., Kaeslin, H., & Fichtner, W. (2005). Two-phase clocking and a new latch design for low-power portable applications. In *Proceedings of the Internal Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS'05)* Leuven, Belgium, (pp. 446–455).

Flavio Carbognani was born in Parma (Italy) in 1977. He studied Telecommunication Engineering at the University of Parma. In 2003 he received his Master Degree, discussing a thesis about lifetime prediction models of power semiconductor devices, prepared at the Institute of Technology Zurich (ETH), Switzerland. He joined the Integrated Systems Laboratory (IIS) as a research assistant in 2002, where he is currently pursuing the Ph.D. degree. His primary research interest is low power design for low-frequency applications.

Felix Buergin was born in Basel (Switzerland) in 1976. He studied Electrical Engineering at the Swiss Federal Institute of Technology Zurich (ETHZ). In 2002 he received his Master Degree, writing a thesis about a Sigma-Delta A/D Converter for high quality audio applications at the Integrated Systems Laboratory of ETHZ. At present, he is pursuing the Ph.D. degree at the same laboratory; his primary research interest is the low power design, focused on hearing aid applications.

Norbert Felber was born in Trimbach, Switzerland, in 1951. He received the Dipl. Phys. (M.Sc.) in 1976 from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland. Subsequently, he was a Research Assistant at the Labora-

tory of Applied Physics, ETH, where he received the Dr. sc. nat. (Ph.D.) degree in 1986. In 1987, he joined the Integrated Systems Laboratory (IIS) of ETH, where he is currently a Research Associate and Lecturer in the field of VLSI design and test. His research interests are in telecommunications, digital signal processing (digital filters, audio, video, pattern recognition, and image processing), optoelectronics, measurement techniques, and device characterization.

Hubert Kaeslin (M'86) received the Dipl. Ing. and Ph.D. degrees in Electrical Engineering from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 1978 and 1985, respectively. He currently is a lecturer at ETH, where he also heads the Microelectronics Design Center. His research interests include dedicated VLSI architectures, low-power and self-timed (GALS) design, hardware description languages, high-level design methodology, electronic design automation, and graph theory.

Wolfgang Fichtner (M'79-SM'84-F'90) received the Dipl. Ing. degree in Physics and the Ph.D. degree in Electrical Engineering from the Technical University of Vienna, Austria, in 1974 and 1978, respectively. From 1975 to 1978, he was an Assistant Professor with the Department of Electrical Engineering, Technical University of Vienna. From 1979 through 1985, he was with AT&T Bell Laboratories, Murray Hill, NJ. Since 1985 he has been Professor and Head of the Integrated

Systems Laboratory, ETH Zurich, Switzerland. In October 1999, he became Chairman of the Information Technology and Electrical Engineering Department, ETH Zurich. In 1993, he founded ISE Integrated Systems Engineering AG, a company in the field of Technology CAD, which was acquired by Synopsys, Inc. in 2004. His research activities cover Physics-based simulation of semiconductor devices and technologies in microelectronics and optoelectronics, physical characterization and electronic measurement in deep-sub-micrometer and nanotechnologies, as well as design and test of digital integrated circuits. Wolfgang Fichtner is member of the Austrian and Swiss National Academies of Engineering. In 2000, he was the recipient of the IEEE Andrew S. Grove Award for his contributions to Technology CAD.