Clock Generation and Distribution for the First IA-64 Microprocessor

Simon Tam, *Member, IEEE*, Stefan Rusu, *Member, IEEE*, Utpal Nagarji Desai, Robert Kim, Ji Zhang, and Ian Young, *Fellow, IEEE*

***Abstract*—The clock design for the first implementation of the IA-64 microprocessor is presented. A clock distribution with an active distributed deskewing technique is used to achieve a low skew of 28 ps. This technique is capable of compensating skews caused by within-die process variations that are becoming a significant factor of the clock design. The global, regional and local clock distributions are described. A multilevel skew budget and local clock timing methodology are used to enable a high-performance design by providing support for intentional clock skew injection and time borrowing. By providing test access port interface to the deskew architecture and the incorporation of the on-die-clock-shrink, this design is equipped with two very powerful post-silicon timing debug tools that are critical to high-performance microprocessor design and enabled quick time-to-market.**

***Index Terms*—Active deskew, clock, clock distribution, clock generation, clock skew.**

I. Introduction

THE MICROPROCESSOR described in this paper is the first-generation implementation of the Intel IA-64 64-bit microprocessor. It is realized using the explicitly parallel instruction computing (EPIC) instruction set architecture that maximizes performance through hardware and software synergy. The microprocessor core contains 25.4 million transistors and is fabricated on a 0.18-μm, six layer metal CMOS process, packaged in an organic land grid array (OLGA) package using the controlled collapsed chip connection (C4) technology and operates at 800 MHz [1].

Demand for high performance and high level of integration requires a significant increase in silicon real estate that resulted in an increase of the clock loading. Aggressive technology scaling coupled with larger die size have made within-die process variation a prominent factor in determining the on-die clock skew as the clock skew due to within-die variation has become a significant component of the total skew budget. Aggressive clock skew control also has the obvious benefits to the microprocessor performance, as any clock skew reduction will be directly added to the available cycle time.

Many techniques have been used in the clock design for high-performance microprocessors [2]. A balanced-tree structure for the distribution of the primary clock is one example. Managing a perfectly balanced clock design in a complex microprocessor is difficult due to the many floor plan and the clock loading constraints. In addition, a balanced tree structure does not eliminate the effects of within-die process variation that affects the clocking element across the branches differentially. For these reasons, an active deskewing scheme is used in conjunction with a combined balanced clock tree and clock grid. The active deskewing technique also provides the means to correct critical paths on silicon.

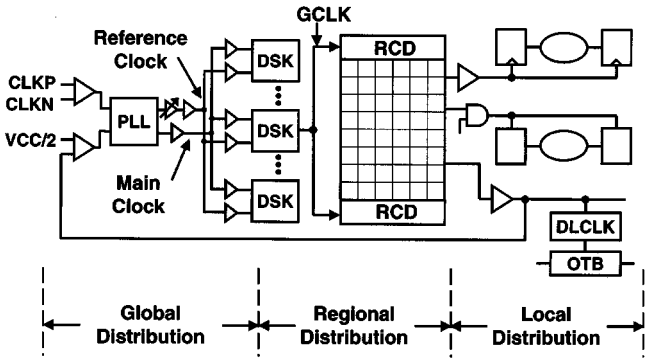

Fig. 1. Clock distribution topology.

A high-performance microprocessor design should provide enough flexibility for the local clock implementation to support intentional clock skew and time borrowing. A multilevel clock skew model and a local clock timing methodology provide the supports needed.

The later part of the paper describes the on-die-clock-shrink (ODCS) [3], which is a key timing debug feature. This feature has the capability to manipulate the clock edges at any desired clock cycle thus permitting fast and efficient identification of timing critical paths. The combination of the ODCS and the distributed active deskewing architecture provided valuable capabilities to the post-silicon timing debug.

The clock system architecture is shown in Fig. 1. The clock topology is partitioned into three segments:

1) Global distribution comprising the clock synthesis by the on-die phase-locked loop (PLL) and the distribution of the core clock and the reference clock from the PLL clock generator to the deskew buffers (DSK);
2) Regional distribution consists of the clock distribution from the DSKs to 30 clock regions (regional clock grids);
3) Local distribution comprising the local clock buffers (LCBs) taking the input from the regional clock grid and the local interconnect to support the clocked elements and the opportunity-time-borrowing (OTB) delay clocks generation.

Manuscript received March 25, 2000; revised June 27, 2000.

The authors are with Intel Corporation, Santa Clara, CA 95052 USA (e-mail: simon.tam@intel.com).

Publisher Item Identifier S 0018-9200(00)09440-3.

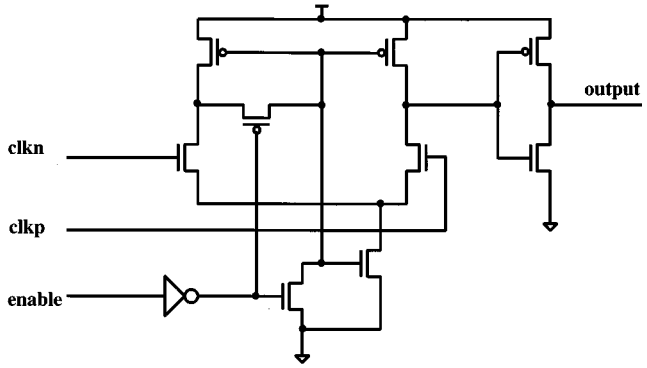

Fig. 2. Differential clock input buffer schematic drawing.

II. Global Clock Distribution

The front-end global distribution consists of the clock-input buffer, an on-die PLL, and the global clock distribution. The microprocessor receives differential clock inputs running at the front-side-bus frequency. A high-performance PLL [4] generates a high-frequency clock running at twice ($2\times$) the core clock frequency. A divide-by-two circuit generates a high-frequency 50% duty- cycle core clock, and a reference clock. These two clocks are routed from the PLL clock generator via two identical and balanced H-trees to eight deskew clusters. Each deskew cluster consisting of four distinct DSKs (Fig. 1), each of which is used for the active deskewing.

One input to the PLL is the system clock (i.e., external clock, delivered to the microprocessor differentially). The other input to the PLL is the feedback clock and is generated on-die and operates at the same frequency as the system clock and performs the front-side-bus clocking. Core frequency multiplication is accomplished with the on-die digital dividers acting on the high-frequency core clock.

A clock-input buffer converts the differential system clock (HSTL) into a single-ended clock with CMOS levels. To ensure that the front-side-bus (FSB) timing accounts for the delay across the clock-input buffer, an identical clock input buffer is inserted in the PLL feedback clock path (Fig. 1). This ensures the FSB clock will be matched to the system clock for the bus transactions. Since the PLL feedback clock is single ended and has CMOS levels, the other input to the feedback clock differential input buffer is connected to $V_{CC}/2$. Measurement has shown this pseudodifferential operation introduces very small additional skew. A simplified schematic of the differential input buffer is shown in Fig. 2. This differential clock input buffer is developed based on a fully complementary self-bias design [5] and exhibits good common-mode range.

The core clock and the reference clock are distributed from the PLL clock generator to the DSK clusters via two identical and balanced H-trees. A schematic drawing of the global core clock tree is shown in Fig. 3. The global clock tree is implemented exclusively in the two highest-level metal layers. To reduce capacitive noise coupling and to ensure good inductive return path [2]; the global H-tree is fully shielded laterally with V_{CC}/V_{SS}. In addition, inductive reflections at the branch points are minimized by properly sizing the metal widths for impedance matching. To achieve this, a two-dimensional

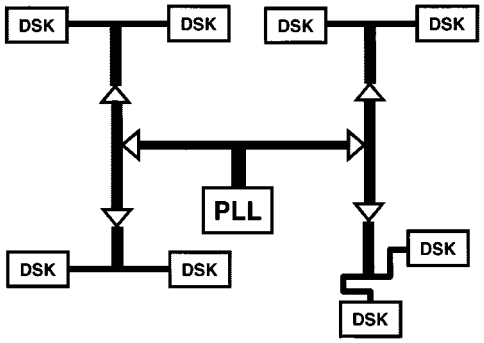

Fig. 3. Schematic drawing of the global H-tree.

RLC and Return Path Extraction	Line Delay	Gate Delay
Addition of return path resistance	+ 27 %	- 3%
Addition of Inductance	+ 8 %	-10 %
Difference from RC Model	+ 35 %	- 13%

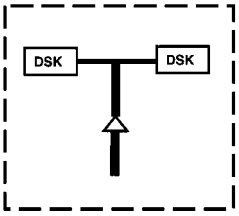

Fig. 4. Inductance modeling and delay optimization.

(2-D) interconnect simulator computes the global clock tree inductance, capacitance, and the return path resistance and they are used as inputs to a distributed *RLC* network model for optimization. Fig. 4 illustrates an example of the inductance optimization at one of the branch point. In this example, after the incorporation of the line inductance and the return path resistance, it was discovered that these factors affected the line delay and the clock repeater delay differently. Namely, the wire delay increased due to line inductance and the return path resistance while the clock repeater delay decreased due to faster transition rates caused by the inductive effects. This example illustrates that for high-performance clock design, it is critical to incorporate the inductance in the analysis.

Although the reference clock resembles the global clock in many aspects, an intentional delay has been added to the root of the reference clock. This additional delay is chosen such that the timing of the reference clock at the input of the DSK will coincide with the total delay of the global clock at the local distribution. Since this delay is incorporated at the root of the reference clock, it introduces no additional skew. The limited span, the reduced buffering employed in its distribution, fixed loading and symmetry makes the reference clock significantly less sensitive to the within-die process variations.

III. Regional Clock Distribution

The regional clock distribution encompasses the DSK, the regional clock driver (RCD) and the regional clock grid. There are 30 separate clock regions each consisting of the above three elements. Since 32 DSKs are available, two of the DSKs are unused in the current implementation. The DSK is a digitally controlled delay-locked-loop structure whose function is to eliminate the skew between its inputs. It will be described later. The deskewed core clock (GCLK in Fig. 1) continues to a clock region and is buffered by the RCDs before it reaches the regional clock grid. It should be noted that the reference clock, which begins at the

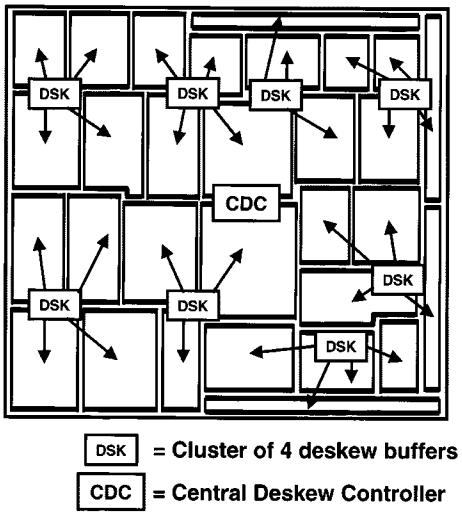

Fig. 5. Regional clocks.

core PLL clock generator terminates at the DSKs and is not required further downstream.

Fig. 5 illustrates the 30 regional clocks in this design. Every regional clock is mapped to a specific DSK. This distributed DSK architecture permits independent skew reduction due to within-die process variations. The RCD design is cell based and is modular in nature. For each clock region, a group of RCD library cells is replicated to cover the area. The RCD also supports features like scan and clock gating.

The regional clock grid is implemented using metal 4 (M4, x-direction) and metal 5 (M5, y-direction). It effectively floats over one or more functional units. As with the global clock network, the regional clock grid contains full lateral shielding to ensure low capacitance coupling and good inductive return paths. The regional clock grid utilizes up to 3.5% of the available M5 and up to 4.1% of the available M4 routing over a region. The RCDs are located at the top and bottom of the clock grid. The maximum dimension of a clock region is chosen based on the underlying clock loading and the span of the region.

A binary distribution network is used to connect the DSK to the RCD. These routes are realized using the top layer metals and with complete lateral shielding for performance. Contrary to conventional clock distribution, the use of deskewing in this design permits these routes to be slightly unmatched without introducing any skew to the system.

IV. Local Clock Generation

The local clock distribution constitutes the final segment of the clock distribution. It consists of the LCBs and the local clock routings that are embedded within a functional unit. The LCBs within a clock region receive the input directly from the regional clock grid. Local clocking needs such as the delay clocks that are needed for the OTB domino circuits [6] and the clocked sequential elements are supported by the LCBs.

While the global and the regional clock distribution is the responsibility of the global clock owner, the local clock is the responsibility of the functional unit designers. Local clock design involves the insertion of LCBs and local interconnects to drive the clocked elements. A family of LCB cells is provided as library cells. Each library clock cell is designed to drive a specific load range. In this manner, when applied according to their specifications, the local clocks will have the same delay across all the paths. When necessary, additional LCBs may be inserted for intentional time borrowing. This is a valuable technique for resolving critical timing paths. Since the regional clock grids effectively float over the functional units, consumption of the regional clocks by the LCBs is accomplished with short point-to-point routes.

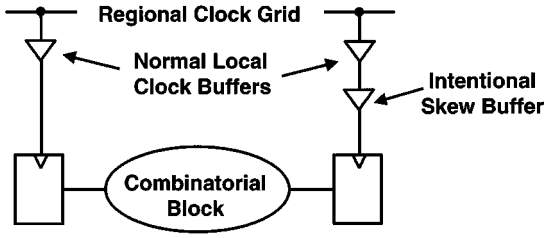

Fig. 6. Local clock distribution.

V. Deskew Architecture and Deskew Buffer Design

As described earlier, a reference clock is distributed along with the core clock from the PLL clock generator to the DSKs. The availability of the reference clock and the incorporation of distributed active deskewing greatly enhanced the overall skew management and enabled a high-performance microprocessor design.

In a typical design, techniques such that balanced trees, load matching with dummy devices, and clock path length tuning are used to minimize the clock skew across the die. These techniques, however, are unable to compensate for skews caused by on die process variations. In addition, it is unlikely that a balanced clock tree that encompasses the entire clock distribution network from the clock generator to the clocked elements can be achieved. A local path tuning approach can be used but it would require detailed analysis of all the clock routes after the implementation is complete. Optimization at this late phase of design will have significant impact on the overall design schedule, not to mention that nonsymmetric clock path tuning will be more sensitive to future process shifts. Therefore, a technique that can compensate the within-die process variations and load mismatches, while being flexible is of paramount value.

A distributed clock deskewing technique is used to satisfy these objectives. The premises of this approach are:

1) distribute a separate reference clock across a region significantly smaller than that which is required for the core clock distribution;
2) ensure the loading for the reference clock is identical at all its end points;
3) make available feedback clocks from the end points of the core clock distribution for skew comparison and adjustment.

Fig. 7 shows the DSK architecture consisting of the DSK and a clock region. As described earlier, each clock region is served by a dedicated DSK. The local feedback clock is sampled from the regional clock grid and sent back to the DSK. Other inputs to the DSK are the global clock and the reference clock.

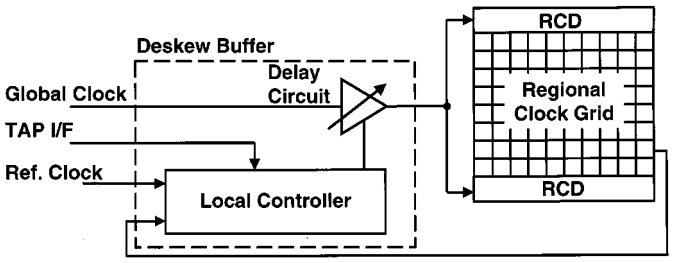

Fig. 7. Deskew buffer (DSK) architecture.

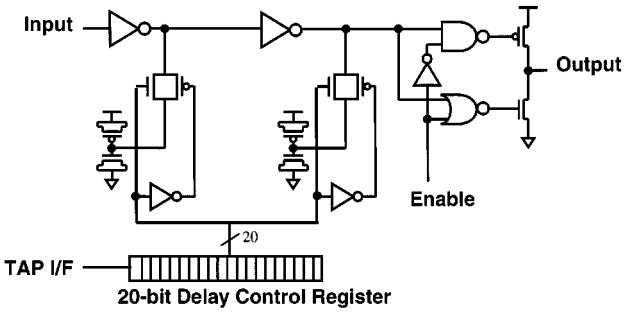

Fig. 8. DSK variable delay circuit.

A phase detector residing within the DSKs local controller will sample the phase difference between the reference clock and the local feedback clock. Adjustments to the core clock delay are accomplished through the variable delay circuit. This continues until a minimum phase error is achieved. Therefore, any load mismatches and within-die variations in the core clock distribution are automatically compensated. Since all the clock regions utilize the same reference clock, the residual skew of the reference clock, the uncertainty of the phase detector, and the mismatches of the feedback clocks will determine the overall skew across these regions. Simple analysis will show that the reduced span of the reference clock and the matched load will result in significantly lower skew. Phase detector uncertainty is made very small by a symmetric circuit design and layout and by allowing longer time for the phase comparison. Strategic DSK placements relative to the clock regions also ensure short feedback clock delays. Coupling this with selective network tuning will ensure that the feedback delay mismatches are negligible. In comparison, in a conventional approach, the entire clock network, the clock loading, and within-die process mismatches will determine the overall skew of the design.

Fig. 8 shows the DSKs variable delay buffer. It is a digitally controlled analog delay line consisted of a 20-bit delay control register, a two-stage variable delay circuit, and a push-pull style output buffer. The 20-bit delay control register forms a 20-steps linear delay coding. A 20-bit delay control was selected based on the delay step-size resolution and the total buffer delay range. Delay adjustment can be accomplished by shifting a "1" from one end of the register to decrease its delay or by shifting a "0" from the opposite end to increase its delay. In addition to the input derived from the DSK local controller, the delay control register also accepts input from the test access port (TAP) interface. This feature permits a manual adjustment of the DSK delay through the TAP interface which can be used for post-silicon *in-situ* timing optimizations. The variable delay circuit is

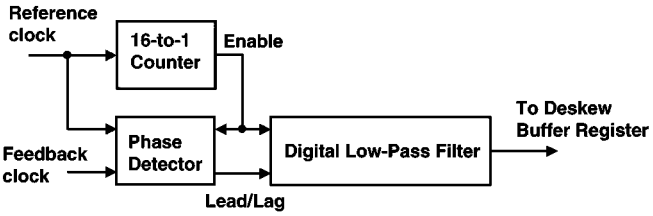

Fig. 9. Deskew buffer local controller.

constructed of CMOS inverters and two arrays of passive loads [7]. The delay across the inverters will vary in accordance to the setting stored in the delay control register. Advantages of this design over a starving inverter approach are linear delay steps and more symmetric layout. The push-pull style output stage consists of twelve parallel drivers of which any can be enabled individually via mask options to match the extracted loading of each region. This allows one standard design to accommodate a wide range of regional clock loads. The measured delay range of the DSK is 170 ps with a step size of 8.5 ps.

Fig. 9 shows the DSKs local controller architecture. The phase detector is an R–S latch design with high emphasis on circuit and layout symmetry for low offset. To ensure high resolution and to avoid metastability, the phase comparison is performed once every 16 cycles by an enable mechanism that samples the reference clock and the feedback clock once every 16 cycles. By not having to make a phase comparison every clock cycle, the phase detector is given plenty of time to resolve the phase difference. Output of the phase detector is a "Lead/Lag" signal which is further filtered by a six-tap digital low-pass filter before sending the result to the delay control register. The purpose of the filter is to ensure that any phase comparison noise will not propagate to the delay control register. The filter also detects a steady-state Lead/Lag pattern that signals the completion of the deskew operation. At this point, the phase difference between the reference clock and the feedback clock will be less than one DSK delay step size. The DSK local controller will stop the deskew operation and set the delay control register such that the feedback clock is always faster than the reference. This last procedure will ensure that the overall residual skew due to the DSK will be limited to one delay step size.

By taking the deskew feedback clock from the regional clock grid and noting that the local clock is one buffer away, active skew compensation is achieved nearly throughout the entire clock distribution. Since the local clock design is incorporated with the full-chip timing model, skew effects at the local clocks are explicitly accounted for. Additional delay is added to the reference clock at its root so the reference clock timing will match the average delay of the feedback clocks.

The deskew operation commences during the initial microprocessor reset. The deskew for the 30 clock regions proceed in parallel. A global deskew controller monitors the progress and signals the deskew completion. Once this occurs, the DSK delay register settings will be fixed until the next power up sequence. This is in contrary to an alternative mode of operation where the deskew operation is allowed to continue during normal microprocessor operation. The tradeoffs between the two modes

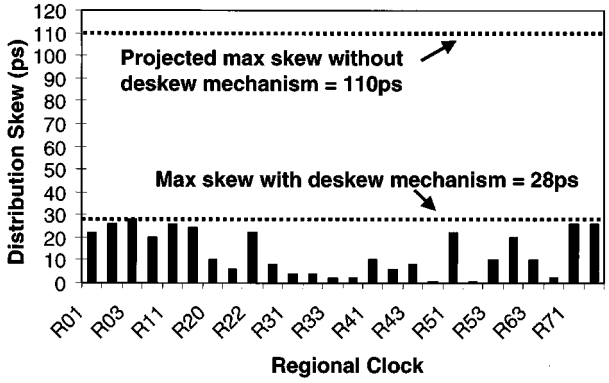

Fig. 10. Experimental skew measurements.

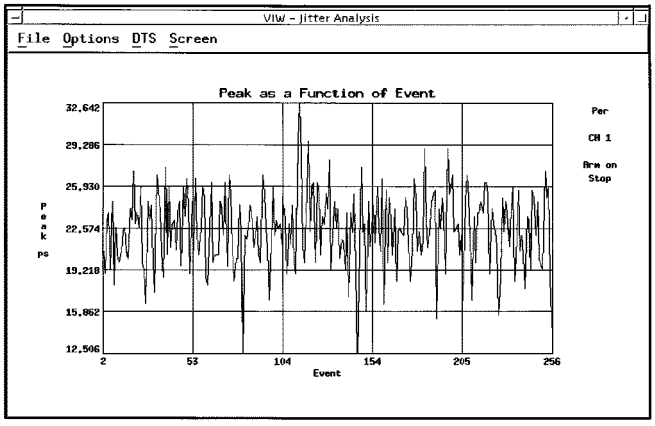

Fig. 11. Clock period jitter.

are the level of loop stability and the ability to further compensate for slow dynamic effects such as temperature variations and V_{CC} voltage drifts.

Direct access to the DSK delay register via the TAP interface permits manual adjustment of the DSK delay. This feature is extremely valuable during post-silicon timing debug. It can be used to support intentional clock skewing in order to discover or to correct critical timing paths. This feature was successfully employed during the first silicon.

Fig. 10 shows the measured skew. The total skew is 28 ps with deskewing and is four times larger with the deskew mechanism disabled. The data is collected using the laser voltage probing (LVP) technique that enables high-resolution signal probing from the backside of the silicon die [8]. Clock period jitter was measured by buffering the clock to the output (Fig. 11).

VI. Clock Skew Analysis Methodology

Full-chip timing analysis included all the local clock distributions whereas the global and regional distributions are modeled with a clock skew budget. A paramount concern in any high-performance design is the analysis of the hold-time margin. Fig. 12 illustrates the data transfer between two sequential elements. To ensure functionality, the minimum data delay time T_d must be greater than the clock skew T_{skew} between the driver and the receiver clocks and the hold time requirement of the receiving

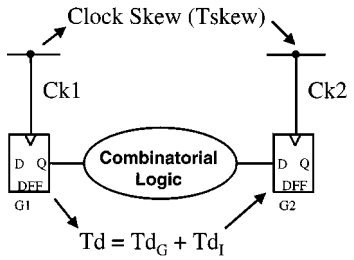

Fig. 12. Hold time analysis.

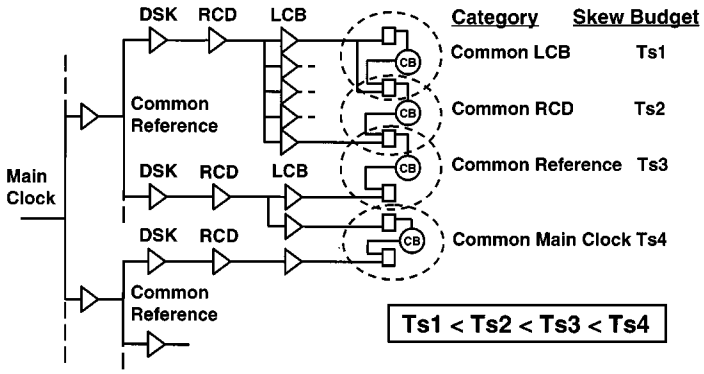

Fig. 13. Multilevel skew timing budget.

sequential element. By including the hold time with the data delay time, (1) shows the inequality which must be satisfied.

$$\min(T_d) \geq \max(T_{\text{skew}}). \tag{1}$$

A simple methodology is to apply one worst-case skew budget across all the paths. This however, will be overly conservative when the clock delay from the common point of clock convergence is small. The clocking topology used in this design conveniently leads to four common points of clock convergence (Fig. 13):

1) Common LCB applies where the driver and receive clocks are derived from a single LCB. This constitutes the best skew scenario.
2) Common RCD applies when both the driver and the receiver elements belong to same clock region. This is the second best scenario.
3) Common Reference applies when two clock regions can be mapped to the same DSK cluster which is the third-best scenario.
4) Common Main Clock applies when the common point of clock convergence is at the PLL clock-generator output. This corresponds to the worst-case scenario.

The timing verification tool analyzes each path by tracing the clock starting the path and the clock ending the path and establishes the closest common point that determines the appropriate skew budget. Using multiple skew budgets permits a better optimization of the design by minimizing the skew penalty and without scarifying any design margin. Since the delay of the LCBs and the local clock routes are fully accounted for in the timing model, the impact of the local skew has been completely eliminated.

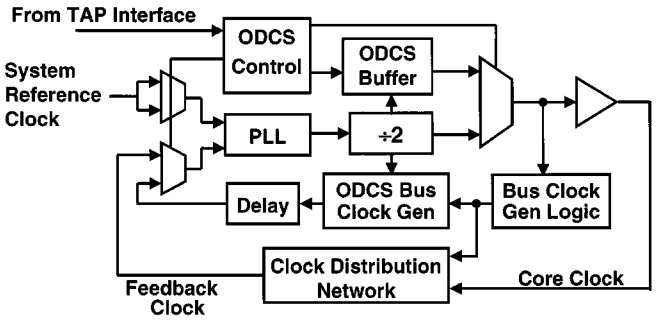

Fig. 14. On-die-clock-shrink (ODCS) architecture.

TABLE I
ODCS EXAMPLE—SHRINK CLOCK HIGH PHASE AT THE $N+1$ CYCLE

State	Rise Setting (arbitrary time unit)	Fall Setting (arbitrary time unit)
Current	10U	10U
N+1	10U	8U
N+2	8U	8U
N+3	8U	8U

VII. ON-DIE CLOCK SHRINK ARCHITECTURE

A successful microprocessor design is strongly influenced by the ability to perform efficient post-silicon timing debug. Techniques such as the PLL by-pass mode have been used in the past. The PLL by-pass mode intentionally bypasses the microprocessor's on-die PLL clock generator and relies on the external tester to supply the core clock. To apply this technique to a microprocessor that operates with frequencies exceeding 500 MHz will require the automatic test equipment (ATE) tester to supply at least two out-of-phase clocks and an on-die clock mixer to generate the high-frequency core clock. Accurate core-clock edge manipulation cannot be achieved with this technique due to the limited edge placement accuracy (EPA) of the ATE clock signal. For example, the EPA of a state-of-the-art ATE tester is approximately ± 100 ps per clock channel. When coupled to a $2\times$ on-die clock mixer and with two tester clocks, the core-clock edge accuracy will be ± 200 ps. This accuracy is of no value when applied to the post-silicon timing debug. The ODCS is the technique that retains the edge placement accuracy of the on-die PLL generated core clock while providing the capability to manipulate the clock edge at certain predetermined times so as to alter it instantaneous frequency.

Fig. 14 shows the architecture of the clock generation unit with the ODCS feature. The ODCS block consists of the ODCS controller, an ODCS delay buffer, and the ODCS bus block generation circuitry. The ODCS controller contains three registers that will specify the rise and the fall delays of the ODCS buffer at any three consecutive cycles (e.g., the Nth, $N+1$th , and $N+2$th). Nominally, the ODCS buffer stays with the Nth delay settings. The ODCS controller will determine the specific cycle when transitions from the Nth settings will occur. An example in which the high phase at the $N+1$ cycle is shrunk is shown in Table I with the corresponding waveform shown in Fig. 15.

Since the ODCS mode manipulates the core clock edges and in turn will affect the PLL feedback clock, there is the concern that the PLL may be affected. To ensure that the PLL clock generator is not affected by the ODCS operation, a dedicated PLL feedback path is utilized. The ODCS bus clock generation circuitry will generate PLL feedback clock that is unperturbed by the ODCS buffer. By inserting an additional delay to the ODCS feedback that corresponds to the normal PLL feedback clock delay, the FSB timing is unaffected. Any slight delay mismatch between the ODCS PLL feedback clock and the normal PLL feedback clock will not affect the PLL operation because the usage does not permit switching between these two modes dynamically.

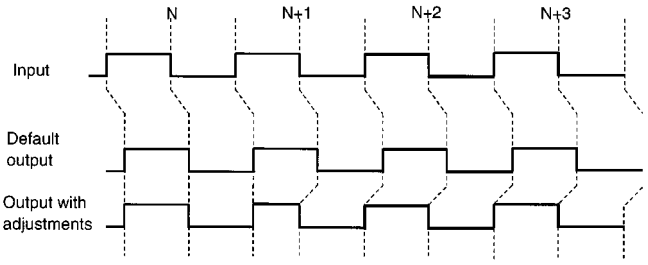

Fig. 15. ODCS manipulated clock waveform using the setting at Table I.

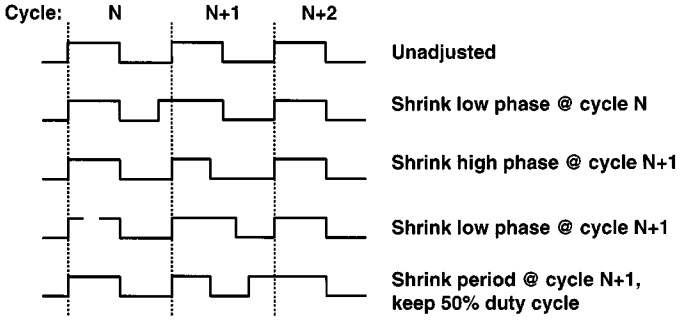

Fig. 16. ODCS capabilities.

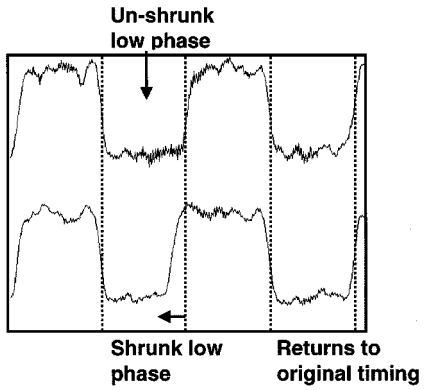

Fig. 17. Shrinking a low phase using ODCS.

Fig. 16 illustrates some of the waveforms that can be created using the ODCS feature. The total range for any edge manipulation in the ODCS mode is 200 ps with fourteen discrete linear steps. Figs. 17 and 18 show the measured LVP clock waveforms demonstrating a shrunk low phase and a shrunk high phase respectively. In both examples, the ODCS settings are constructed such that the clock waveform will return to normal after the manipulation.

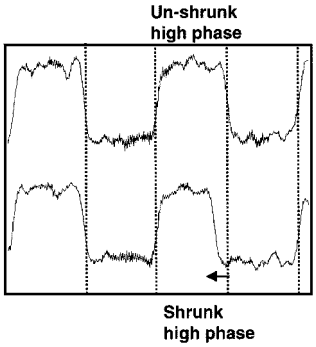

Fig. 18. Shrinking a high phase using ODCS.

VIII. Summary

In summary, this design utilizes an active distributed deskewing technique to achieve low clock skew. This technique compensates for process variations and design mismatches that cannot be achieved with conventional passive network tuning schemes. The measured clock skew across the microprocessor is 28 ps. This is about four times smaller than the case with the deskewing disabled. In additional to achieving low skew, the deskewing architecture also enables intentional clock skew injection to facilitate post-silicon timing debug and performance tuning through the TAP interface. A local clocking architecture and a multilevel skew budget and timing methodology enable a high-frequency design by explicitly supporting intentional local clock-skew injection and time borrowing. The ODCS feature is an indispensable tool for post-silicon timing debug for any high-frequency microprocessors.

Acknowledgment

The authors would like to acknowledge the following team members for their contributions: S.-S. Cheng, B. Kleveland, Y. Lu, J. Prak, and K. Wong.

References

[1] G. Singer and S. Rusu, "The first IA-64 microprocessor: A design for highly parallel execution," in *ISSCC Dig. Tech. Papers*, Feb. 2000, pp. 422–423.
[2] P. J. Restle and A. Deutsch, "Designing the best clock distribution network," in *1998 Symp. VLSI Circuits Dig. Tech. Papers*, 1998, pp. 2–5.
[3] S. Fischer *et al.*, "A 600 MHz IA-32 microprocessor with enhanced data streaming for graphics and video," in *ISSCC Dig. Tech. Papers*, Feb. 1999, pp. 98–99.
[4] I. Young *et al.*, "A 0.35 um CMOS 3-880MHz PLL N/2 clock multiplier and distribution network with low jitter for microprocessors," in *ISSCC Dig. Tech. Papers*, 1997, pp. 330–331.
[5] M. Bazes, "Two novel fully complementary self-biased CMOS differential amplifiers," *IEEE J. Solid-State Circuits*, pp. 165–168, Feb. 1991.
[6] D. Harris *et al.*, "Skew-tolerant domino circuits," *IEEE J. Solid-State Circuits*, pp. 1702–1711, Nov. 1997.
[7] G. Geannopoulos and X. Dai, "An adaptive digital deskewing circuit for clock distribution networks," in *ISSCC Dig. Tech. Papers*, 1998, pp. 400–401.
[8] T. Eiles *et al.*, "Optical probing of flip-chip packaged microprocessors," in *ISSCC Dig. Tech. Papers*, Feb. 2000, pp. 220–221.

Simon Tam (M'80) received the B.S., M.S., and Ph.D. degrees all in electrical engineering and computer sciences from the University of California, Berkeley.

He is a Principal Engineer at Intel Corporation's Intel Architecture Group—Enterprise Processors Division, where he has been engaged with the design of high-performance microprocessors with special emphasis on high-frequency clocking architecture and high speed circuits. Prior to joining the Intel Architecture Group, he was with the Intel Neural Network Group. He designed electrically programmable neural network chips using analog VLSI techniques and EEPROM technology. He was also with the Intel California Technology Development division engaged with the development of flash memory and EEPROM. He was the co-inventor of the Intel's ETOX flash memory cell. He joined Intel Corporation in 1983. He has been awarded 18 U.S. patents and has authored or co-authored 27 technical papers in the areas of microprocessor clocking, neural networks, flash/EEPROM memories, and VLSI device physics.

Stefan Rusu (M'85) received the M.S. degree in electrical engineering from the Polytechnic Institute, Bucharest, Romania.

He first joined Intel Corporation in 1984, working on data communications integrated circuits with emphasis on circuit and physical design. In 1988, he joined Sun Microsystems, working on microprocessor designs with focus on clock and power distribution, packaging, standard cell libraries, CAD, and circuit design methodologies. He re-joined Intel Corporation in 1996 working on the clock and power distribution, cell library, I/O buffers, and packaging of the first IA-64 microprocessor. He is presently a Principal Engineer in Intel's Enterprise Products Group, leading the technology and special circuits design group for all the IA-64 microprocessors. His technical interests are high-speed clocking, power distribution, I/O buffers, low-power design and high-speed circuit design techniques. He has published numerous technical papers and currently holds ten U.S. patents with several more pending. He has been a member of the ESSCIRC Technical Program Committee since 1998.

Utpal Nagarji Desai received the B.S. degree in electronics engineering from the Maharaja Sayajirao University, India, and the M.S. degree in electrical engineering from Auburn University, Auburn, AL.

He is a Senior Design Engineer at Intel Corporation's Intel Architecture Group—Enterprise Processors Division, working on the design of high-performance IA-32 and IA-64 microprocessors, with special emphasis on high-frequency clock synthesis and distribution. Prior to joining the Intel Architecture Group, he was with the Hand Held Products Group, working on Chipsets design for low-power IA-32 microprocessors. Besides clock synthesis and distribution, his areas of interest are high-speed circuits and neural networks. He has authored or co-authored seven technical papers in the areas of clocking and neural networks.

Robert W. Kim was born in Inchon, Korea, in 1973. He received the B.S. degree in electrical engineering from the University of California, Los Angeles, in 1996.

He joined Intel Corporation, Santa Clara, CA, in 1996. In 1997, he joined the Microprocessor Product Group and worked on the development of IA-64 microprocessor products.

Ji Zhang was born in Shanghai, R.O.C., in 1962. He received the B.S. degree in semiconductor physics from Fudan University, Shanghai, R.O.C., in 1984, and the M.S.E.E. degree from Purdue University, West Lafayette, IN, in 1994.

His interests include semiconductor device and electronic integrated circuit. He is currently working on clock generation and distribution in Intel's Enterprise Platforms Group.

Ian Young (F'98) was born in Melbourne, Australia. He received the B.E. and M.Eng.Sci. degrees in electrical engineering from the University of Melbourne in 1972 and 1975, respectively. He received the Ph.D. degree in electrical engineering from the University of California, Berkeley, in 1978.

He was a Consultant to American Microsystems, Inc., from 1976 to 1978, developing switched-capacitor circuits for telecommunications. He taught the Analog Integrated Circuit Design course at the Department of Electrical Engineering and Computer Science, University of California, Berkeley, in 1977. From 1978 to 1981, he was with Mostek Corporation, designing MOS integrated circuits for telecommunications applications. From 1981 to 1983, he worked as an Industrial Consultant designing analog/digital MOS integrated circuits. In 1983, he joined the Portland Technology Development Group at Intel Corporation, Hillsboro, OR, where he is currently an Intel Fellow and Director of Advanced Circuit and Technology Integration. He is responsible for defining and developing future circuit directions and optimizing the manufacturing process technology for high-performance microprocessor products. At Intel, he has been involved with the development of DRAMs and high-speed SRAMs, CMOS and BiCMOS digital circuit design techniques for microprocessors, phase-locked loop clock generation and distribution for microprocessors, interconnect capacitance and device modeling, and analysis of circuit design considerations in the development of process technology for six generations (1.0 through 0.18 μm) of advanced microprocessors. He holds 21 patents. He has been a member of the Program Committee for the Symposium on VLSI Circuits since 1991. During 1995 and 1996, he was Co-Chairman and Chairman of this Symposium's Program Committee, and in 1997, he was the Symposium's Co-Chairman. Since 1992, he has been a member of the International Solid-State Circuits Conference Program Committee, serving as the Digital Subcommittee Chairman since 1996.

Dr. Young has received three Intel Achievement Awards: in 1990 for phase-locked loop clock generator design, in 1992 for defining and implementing BiCMOS process technology and BiCMOS logic circuit techniques, and in 1996 for development of the high-performance transistors for the 0.25-μm technology. In 1994, he served on the Semiconductor Industry Association 1995–2010 National Roadmap Defining Technical Working Group. In 1994, he also served on the first Sematech TCAD project planning committee. In each of these activities, he represented circuit design and interconnect issues. He has been a Guest Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS for the December 1994, April 1996, and April 1997 special issues.