TAM 5.5: A Zero-Overhead Self-Timed 160ns 54b CMOS Divider

Ted E. Williams, Mark A. Horowitz
Stanford University, Stanford, CA

This circuit demonstrates a self-timed iterating ring which attains the speed of a combinational array while using only a fraction of the silicon area. The stages in the ring compute mantissa quotient digits for a floating-point division operation. Unlike circuits which implement self-timing by using a matched on-chip clock generator to provide an internal clock for synchronous blocks, the circuit of this paper uses local control handshaking between fully asynchronous blocks and will operate correctly for any values of gate delays.[1] To avoid requiring matching path delays, completion information is embedded in the data throughout the design by using dual-monotonic wire pairs. The precharged function blocks use merged n-channel pull-down networks to choose which of the wires in each pair to set high.

A self-timed iterating ring is formed from an asynchronous pipeline by looping data from its output back to its input. The total latency and throughput tradeoff with the number of latches in the ring.[2] Minimal latency is achieved in this chip by directly concatenating precharged logic blocks into a looped domino chain without adding any explicit latches. The precharge (reset) signals for each block are controlled separately so each block can be used as an implicit latch without adding any additional transistors. The self-timed control is designed to precharge each block after data passes it, and to remove its precharge enabling its evaluation, before data loops around to its inputs again. A graph-based method is used to analyze the inter-block dependencies and aids in keeping the critical path solely within the combinational data elements.

The ring is organized as a series of adjoining stages each of which is internally composed of precharged blocks. A key to removing extra control dependencies which could degrade performance is to place enough stages in the loop to fully utilize and encompass the time taken by the control so its delay is completely hidden. This chip uses five stages to allow the control signals to enable each block 0.7 stage delays before its data arrives, which is measured at 2ns as shown in Figure 1. This margin ensures no control logic enters into the critical path even with some variances in the delays. Thus, the data flows continually at the same rate it would flow through an "unwrapped" combinational array implementing the same functions. While most previous asynchronous circuits have suffered delays due to handshaking and control, this method of self-timing adds zero control overhead to the latency of the raw function computation.

The block diagram of the self-timed iterative structure for implementing division is shown in Figure 2. The blocks of each stage in the iterative ring implement one step of a modified radix-2 SRT division algorithm.[3] Carry-save adders can be used for the computation of the next partial remainder because the SRT algorithm allows selecting signed quotient digits based on approximating the remainder by the top few bits. Only a 3b carry-propagate adder is necessary to resolve this approximation because the quotient-selection logic also examines the previous quotient digit to avoid certain boundary cases. While the traditional SRT algorithm is purely sequential, the implementation here, shown in Figure 3, overlaps the execution of adjacent stages by replicating the carry-propagate adders for each possible quotient digit so they can begin operation before the actual quotient digit arrives to choose the correct branch. This approach allows the delay through a stage to be the average rather than the sum of the propagation delays through the partial remainder and quotient digit-selection paths. The performance is improved by 30% over a standard sequential arrangement of the same blocks. The self-timed control in each stage uses C-elements (*last-of* gates, defined in Figure 4) to combine the completion-detector signals and produce the block reset signals.

After the input registers are loaded with the dividend and divisor, a go signal triggers the self-timed ring to begin iterating. As the bits in the quotient are determined, they are collected from each stage by five separate asynchronous shift registers composed of the cells in Figure 4. The ring loops a maximum of 11 times to fill the five shift registers with a total of 55b for a double-precision result. The iterations are terminated earlier if the remainder comparison on the right side of Figure 2 determines the partial remainder has remained unchanged during the last iteration. If the remainder repeats, then subsequent quotient digits would also repeat; hence, there is no need to compute them again, and the division-done signal can be generated early. Even when iterations terminate early, the full quotient is immediately available from the shift registers because the asynchronous design using C-elements ripples the quotient digits to their final positions as they arrive rather than waiting for a fixed number of clocks as would a synchronous shift register. The repeated quotient digits are also immediately available because they fill all of the positions behind them as they ripple through the shift registers.

The full-custom layout of the self-timed divider was performed using MAGIC and verified using the Stanford IRSIM simulator. The design is fabricated in 1.2μm CMOS technology through MOSIS. The die photo in Figure 5 shows an active area of 9.7mm^2 containing test registers surrounding a core iterating ring in the central 6.8mm^2. The five ring stages are columns which are mirrored appropriately to weave the datapath and achieve equal path lengths. The chips generate the correct data outputs over a wide range of operating conditions. Measured performance as shown in Figure 6 for operation at 5V and 35°C ambient is 160ns for worst-case data, and 45ns for best-case data requiring only two ring iterations. The outputs of this chip are still in redundant form, but a single 54b carry-propagate-add could both round and resolve the final results.

Since the design produces a done indication, the outputs can be used as soon as they are available without waiting for worst-case margins over the ranges of possible data values, temperature, voltage, and fabrication spread. The operating conditions determine performance, and Figure 7 shows measured speeds for varying temperature and voltages.

Acknowledgments:

The authors thank L. Ho, and especially Y. Shimazu for contributions. The work was partially funded by DARPA and an SBIR contract through Silicon Engines Inc.

References:

[1]Santoro M. and M. Horowitz, "A Pipelined 64x64b Iterative Array Multiplier," ISSCC DIGEST OF TECHNICAL PAPERS, pp. 36-37, Feb. 1988.

[2]Williams T.,"Latency and Throughput Tradeoffs in Self-Timed Asynchronous Pipelines and Rings," Stanford Tech Report CSL-TR-90-431, May 1990.

[3]Williams T. and M. Horowitz, "SRT Division Diagrams and Their Usage in Designing Custom Integrated Circuits for Division," Stanford Tech Report CSL-TR-87-326, Nov. 1986.

91CH2960-3/91/0000-0098$01.00 © 1991 IEEE

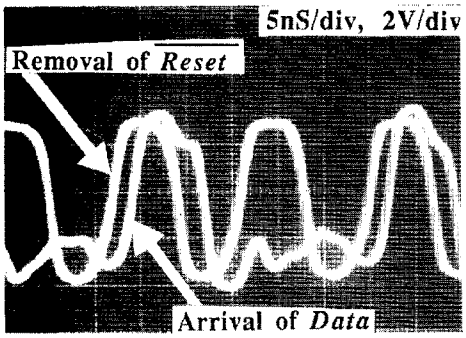

Figure 1: The self-timed reset of each stage is removed 2ns before data is applied. Verifying control introduces zero overhead.

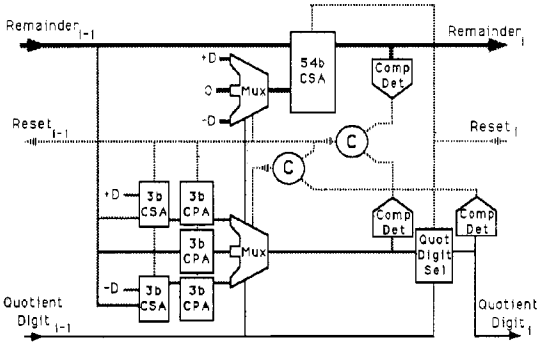

Figure 3: Internal structure of each stage in the ring implements an SRT division step with overlapped execution.

FIGURE 5 – *See page 296*

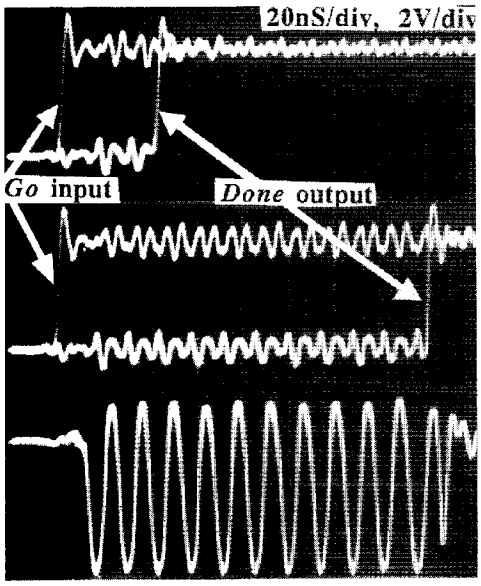

Figure 6: Above: latency varies from 45 to 160ns depending on data. Below: self-timed precharge signal for one internal stage

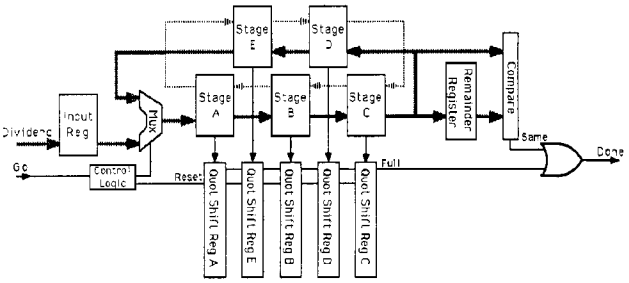

Figure 2: Block diagram for the division circuit using a ring of domino stages which iterates using self-timing. Dotted lines show control signals; shaded lines are dual-monotonic datapaths.

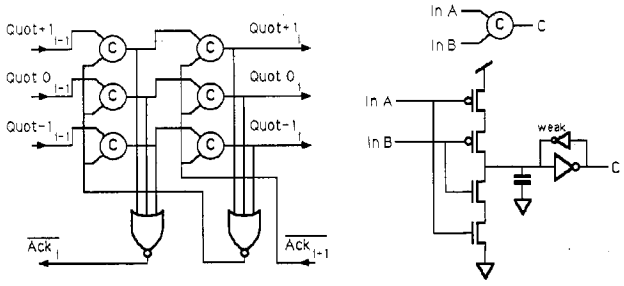

Figure 4: A cell of the asynchronous shift registers for capturing quotient digits on a triple-monotonic wire set. Each quotient digit arrives by setting one of the three wires high, and is followed by a "spacer" where all three are again low. A static C-element is defined at right.

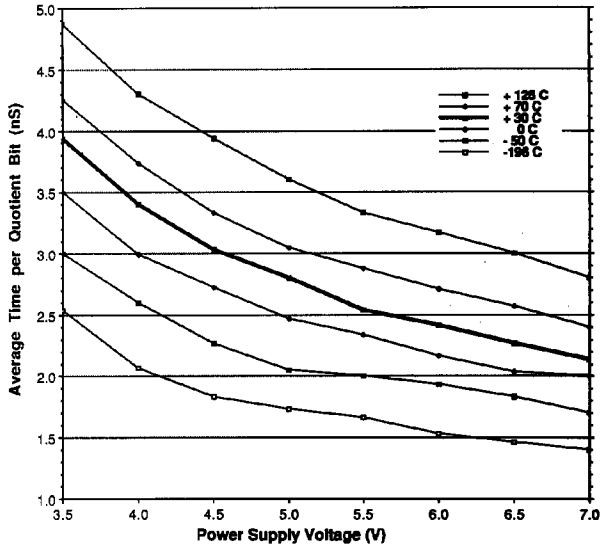

Figure 7: Measured performance per quotient digit under varying temperature and voltage.

TAM 5.5: A Zero-Overhead Self-Timed 160ns 54b CMOS Divider

(Continued from page 99)

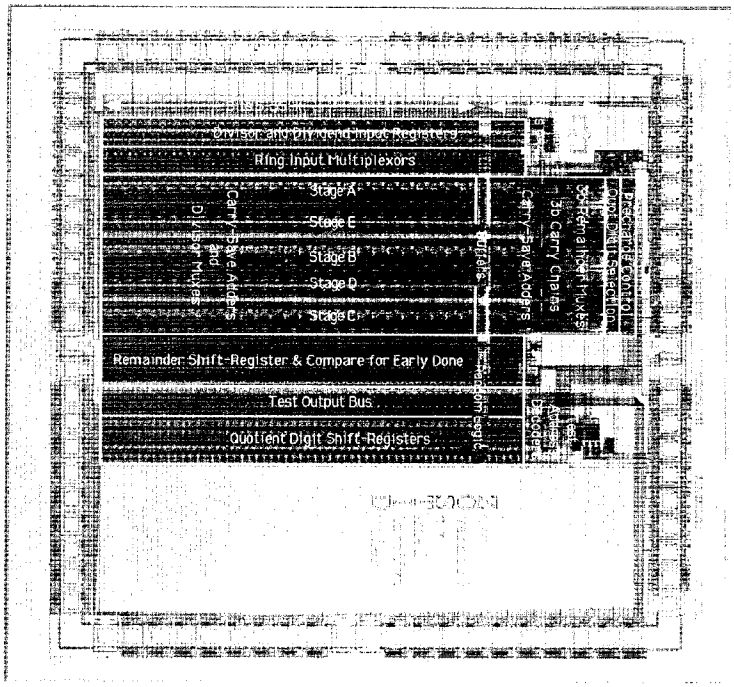

Figure 5: Micrograph of the zero-overhead self-timed 54b divider

TAM 5.6: A 100MIPS, 64b Superscalar Microprocessor with DSP Enhancements

(Continued from page 101)

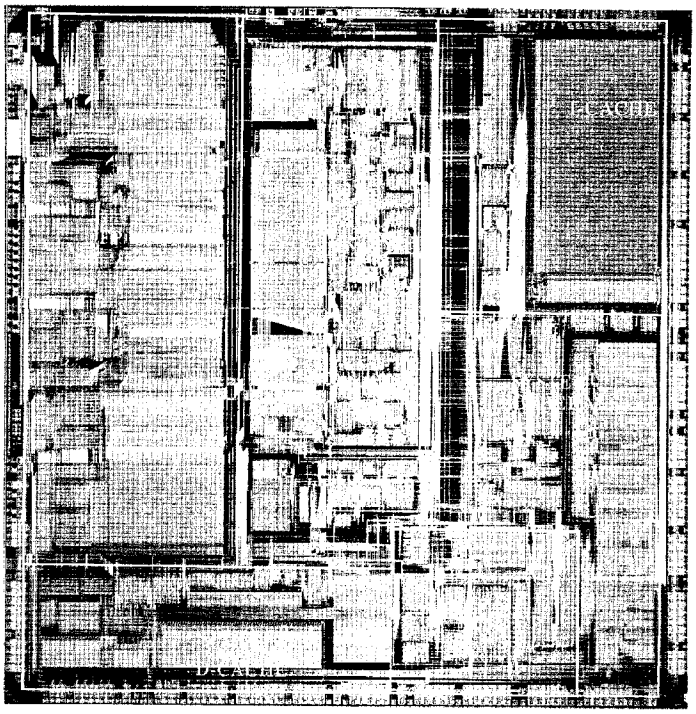

Figure 7: Chip micrograph