

High-throughput asynchronous datapath with software-controlled voltage scaling*

Yee William Li, George Patounakis, K. L. Shepard
Columbia Integrated Systems Lab, Columbia University, New York, NY 10027

Abstract

It is widely recognized that adaptive control of the power supply is one of the most effective variables to achieve energy-efficient computation. In this work, we describe the development of a high-performance asynchronous micropipelined datapath that provides robust interfaces across voltage domains, performing appropriate voltage level conversions and operating between stages with fanout-of-four delays differing by almost two orders of magnitude. With software-specified throughput requirements, the power supply of the datapath is scaled from 2.5 V to 650 mV using an on-chip dc-dc conversion system that combines linear regulators and switched-capacitor power supplies. Because of the asynchronous design style, the processor operates continuously during the voltage scaling transitions.

Introduction

Power consumption has become one of the most important issues in processor design, not only in portable, battery-powered applications, but in high-performance desktop and server applications because of packaging and cooling requirements. Dynamic (or adaptive) voltage scaling (DVS) has been widely studied as one of the most effective means of achieving energy-efficient design[1]. Our design uses software to specify datapath pipeline throughput requirements; a control system automatically scales the voltage to just achieve these requirements. Because of the asynchronous design style, the datapath operates continuously during the voltage scaling transitions. We employ an on-chip dc-dc converter that combines linear regulators and switched-capacitor power supplies, using only on-chip components.

Multirate signal processing applications, such as software radio[2], provide the ideal vehicle for exploring performance-power tradeoffs with adaptive voltage scaling. Vector (or stream) dataflow architectures are the natural choice for such applications and benefit considerably from deep pipelining. Building on earlier work[3] in exploiting the inherent latching properties of dynamic logic to build fine-grained asynchronous pipelines, we develop pipeline circuits which exploit self-resetting techniques to achieve high performance. These circuits introduce robust interlocking to allow for slow environments, to function in the presence of aggressive adaptive voltage scaling, and to handle level conversion between different voltage domains.

Overall chip architecture

The design contains three custom 1K-by-16b self-resetting, low-power SRAMs with pulsed wordline decoding[4]. Each SRAM has associated with it an address-generation unit that consists of a datapath and an asynchronous burst-mode controller designed using generalized C-elements (gC)[5]. Address generation and array access are pipelined such that the array can supply the datapath with operands without limiting throughput.

The prototype datapath in this testchip is a 16-bit carry-lookahead (tree) adder, implemented with seven (micro)pipeline stages. The basic asynchronous pipeline structure supports a mixture of static and dynamic logic with a uniform design for the pipeline controls. The pipeline circuits are designed to operate across all process corners from 2.5 V down to 650

*This work was supported by NSF under Grant CCR-0086007, by the DARPA/MARCO C2S2 Center, and by a gift from the IBM Corporation.

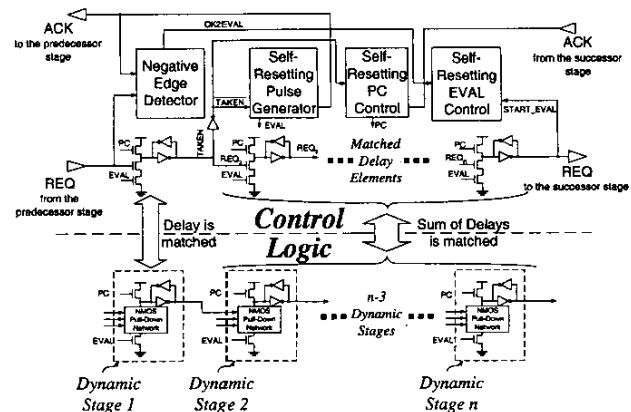


Figure 1: One pipeline stage.

mV and continue to correctly handshake with the SRAMs operating at 2.5 V. Additional pipeline stages are used to perform the voltage level conversions across voltage domains, in this case between the datapath and SRAMs. For maximum testability, scannable latches are included in each pipeline stage. In addition, we have placed pads on the critical signals of the pipeline to allow time-domain "picoprobing" of the waveforms on the testchip.

An instruction unit broadcasts the (VLIW-like) instruction word for the vector (or stream) operation to each of the units. The instruction word consists of starting and ending addresses for each of the SRAMs and the required throughput performance for the datapath unit. The power management system scales the supply for the datapath to meet the throughput requirement specified in the instruction word. An on-chip 40-MHz, 6-bit flash analog-to-digital converter (ADC) allows noninvasive transient monitoring of the power supply to test the power management system functionality.

Asynchronous micropipelines

Figure 1 shows a single stage of a linear pipeline; the top half of the figure contains the control circuits (or local "clock" generators) for the pipeline. In the layout, this resembles a "spine" that runs down the side of the datapath with the area and power overhead of the controller amortized over an entire datapath slice. Adjacent pipeline stages are interlocked by means of the request (REQ) and acknowledge (ACK) signals. PC and EVAL control signals are sent to the stages of the pipeline. Assertion of the REQ signal indicates that a new data token is ready to be evaluated by the successor stage while the pulsed ACK going to the predecessor stage notifies the latter that it can change its output.

Assuming the logic blocks are implemented as domino logic with their pFET clocked by PC and the evaluation foot device clocked by EVAL, such a decoupling defines three functional "phases" for the domino stage, precharge, evaluation, and hold[6]. Each stage cycles through these three phases; after evaluate completes, the stage "self-resets" into the hold

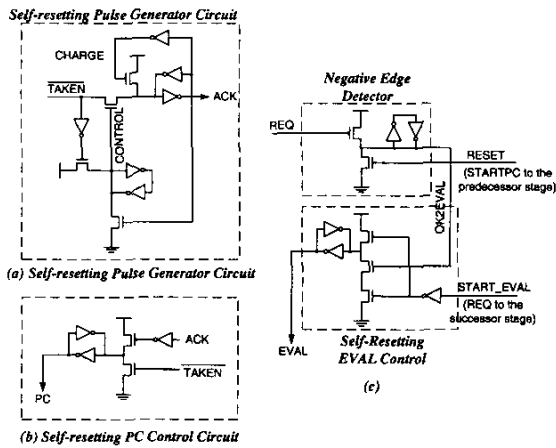


Figure 2: Negative edge detector and self-resetting EVAL control circuits of pipeline controller.

state. When the successor stage evaluates, the current stage is triggered to precharge and then subsequently “self-resets” into the evaluate state. The evaluation of a given stage triggers the predecessor stage to complete its entire next cycle: precharge, evaluation of a new data item, and hold.

One micropipeline stage consists of n ($n \geq 1$) dynamic pull-down networks. For optimal performance, n should be the same across stages (i.e., balanced). The controller has n domino buffers which are sized to match the evaluation delay of the corresponding logic stages. The outputs of the first and last of these dynamic buffers along with the request from the preceding stage and the acknowledgement from the successor stage are processed by four modules within the controller (as shown in Figure 1) independently described below.

Self-resetting pulse generator. This circuit[7] acts on the $\overline{\text{TAKEN}}$ signal, converting a $0 \rightarrow 1 \rightarrow 0$ event on $\overline{\text{TAKEN}}$ into a pulse which constitutes the ACK signal back to the predecessor stage as shown in Figure 2(a). The pulse triggers the precharge of the previous stage.

Self-resetting PC control. This circuit acts as a “pulse-catcher” for the ACK signal from the successor stage and is implemented as shown in Figure 2(b). Precharging starts once the ACK pulse is captured and it is deasserted by $\overline{\text{TAKEN}}$ going to logic one when precharge finishes.

Self-resetting EVAL control and negative edge detector. The EVAL signal is asserted when precharge completes and is deasserted at the end of evaluation, putting the current stage in hold. Logically, this function could be implemented using an inverter. However, slow precharge of the predecessor stage would lead to false evaluation of old data. To avoid this problem, the EVAL should be asserted only after the predecessor stage is precharged, which is accomplished with the more complex circuit structure of Figure 2(c).

Static logic can replace levels 2 to $n - 1$ (as shown in Figure 1) for $n \geq 3$. However, the last stage has to be modified to “clock” the evaluate foot device by the REQ_n signal, rather than EVAL, so as to maintain the monotonicity of the output signal. Beginning the pipeline stage with domino logic prevents the corruption of data when the predecessor is precharged. For $n = 2$ pipelines, the first domino stage can be eliminated since the predecessor will not be able to precharge before valid data is successfully captured in the latch.

The circuit in Figure 3 is used to provide low-latency voltage conversion, that is, to convert a digital signal with a logic one value of V_A to a signal with a logic one value of V_B . This circuit differs from the “traditional” level shifting circuit[8] by the addition of devices M_4 and M_5 and the associated feedback. These devices ensure the pFET pull-up networks




Figure 3: Enhanced voltage level-shifting circuit

are transiently disabled when the state of $\text{OUT}/\overline{\text{OUT}}$ is changing as a result of the switching of $\text{IN}/\overline{\text{IN}}$, reducing the latency of the level converter. Low latency voltage conversion is necessary to prevent the voltage interfaces from becoming a throughput bottleneck. Furthermore, to ensure the pipeline functions correctly across voltage domains, certain elements of the controller must always be run at the full V_{DD} supply.

Throughput (T) is defined as the number of data items processed by the pipeline per unit time. When there are only a few data items in the pipeline, the throughput is said to be data-limited. When the number of data items in the pipeline becomes too high, the pipeline becomes congested and the throughput is limited by the rate at which empty stages (or holes) can move from right to left. These throughput expressions are given by:

$$T_{\text{data-limited}} = \frac{K}{Gt_f} ; T_{\text{hole-limited}} = \frac{G-K}{Gt_r} \quad (1)$$

where K is the average number of data tokens in the pipe and G is the number of pipeline stages. t_f , the forward latency, is defined as the time it takes one data token to move from one stage to its successor. t_r , the reverse latency, is the time it takes a hole to move from one stage to its predecessor.

When Equations 1 are equal, the throughput is optimal and we can calculate K_{optimal} and maximum throughput.

$$K_{\text{optimal}} = \frac{G}{1 + \frac{t_r}{t_f}} ; T = \frac{1}{t_r + t_f} = \frac{1}{t_{\text{cycle}}} \quad (2)$$

where t_{cycle} is the circuit cycle time.

Dynamic voltage scaling is used to adapt this raw throughput capability to the sample rate demands of the signal processing application. More power must be burned to accommodate high bandwidth (high sample rate) signals but the “intrinsic bandwidth” of the pipelines (as characterized by t_{cycle}) can be reduced (saving power) in the case that low bandwidth (low sample rate) signals are being processed. Note that the ability to perform this optimization continuously and without having to stop execution is a feature of the asynchronous nature of the chip and is not easily achieved with synchronous techniques.

Power management system

The power management system is responsible for efficiently scaling the supply voltage for the datapath to just meet the performance target specified in the instruction word. A synchronous state machine accomplishes this by a monotonic search starting from the voltage established for the previous instruction. The voltage-to-performance conversion is achieved via a scaled down replica slice of the unit being regulated and a counter to capture the number of replica “ticks” during a controller clock cycle.

An equally important aspect of the power management system is the design of efficient dc-dc converters to generate the required supply voltages. Dc-dc downconversion from a V_{dd} supply can be generally accomplished in one of three ways: “buck” converters, switched capacitor dividers, and linear regulators. In theory, the “buck” converters[9, 10, 11]

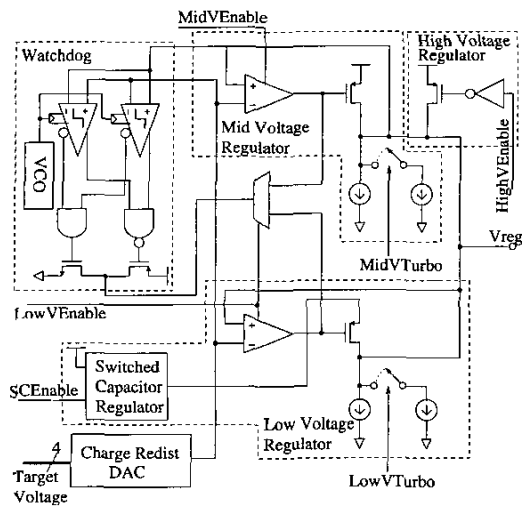


Figure 4: "Hybrid" voltage regulator.

can achieve 100% efficiency if all the components are ideal. Unfortunately, efficient "buck" converters require pins for off-chip inductors; therefore, they are not practical for fine-grain voltage domains.

Linear regulators are the most easily integrable dc-dc converters because they consist of only transistors, but they have poor efficiencies at low output voltages. Furthermore, a linear regulator's op amp requires quiescent current that must be considered when the load is drawing little current. The bias current of the linear regulator must be increased if a fast response time is required.

Switched capacitor voltage dividers (SCVDs) can trade efficiency for integrated chip area and can achieve higher efficiencies than linear regulators at low voltage. The efficiency of an *ideal* SCVD is inversely proportional to the output voltage ripple; therefore, it is proportional to the size of the switching capacitors and frequency of switching for a fixed load current. Real SCVDs incur a power dissipation overhead due to real CMOS switches and implementation details of the on-chip switching capacitors. Increasing the values of the switched capacitors increases the efficiency only at the cost of increased area.

A possible approach to efficient dc-dc downconversion is to use a hybrid voltage regulator scheme (as shown in Figure 4) to trade off area for power efficiency. The maximum core voltage of 2.5V is supplied by a large PFET power transistor with the gate tied to 0 V. Voltages down to 1.0 V are supplied by linearly regulating down from 2.5 V. The lowest voltage range (990 mV down to 650 mV) is supplied by linearly regulating from a 1.25V supply generated by a switched capacitor regulator. Control signals to select the appropriate regulator and allow for smooth transitions between regulator boundaries are generated by the same synchronous state machine that determines the target voltage. Furthermore, the state machine controls the frequency and magnitude of the pulses from the "watchdog" unit. At lower voltages, the digital logic is running at a much slower cycle time and, therefore, does not need to be monitored as frequently.

The switched capacitor regulator (SCR) is shown in Figure 5. Scaling techniques were employed in the SCR to minimize the overhead of switching the capacitors. Further energy savings were attained through low parasitic on-chip metal-insulator-metal (MIM) capacitors instead of the more dense MOS capacitors. Switching frequency scaling was also employed in the SCR by monitoring the minimum output voltage on the 1.25V supply. The simulated efficiency of producing approximately $\frac{V_{out}}{2}$ using this method is greater than 60% under most loads and thus more efficient than the ideal efficiency of a linear regulator (50%).

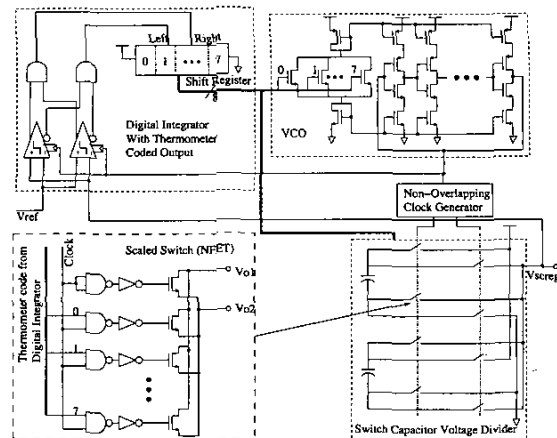


Figure 5: Switched capacitor voltage regulator.




Figure 6: Die photo of the asynchronous DVS prototype chip in a TSMC 0.25 μ m process

Approximately 25 pF of explicit thin-oxide on-chip decoupling capacitance on the supply node of the datapath is adequate to filter out most of the current fluctuations under normal pipeline operation. The asynchronous nature of the pipeline helps in this regard by "spreading" out the current demands of the digital logic and thus reducing the quiescent current in the op amps. Higher frequency changes in the load current are handled by a digital "watchdog" that continuously monitors the output voltage via clocked comparators and pulls the power transistor's gate to ensure the regulated voltage is within 100mV of the target. This approach provides for a more power-efficient design than increasing the large signal bandwidth (and thus the quiescent current) of the op amp in the regulator.

Results

Figure 6 is the die photo of the 9mm² chip as fabrication in the TSMC 0.25 μ m mixed-signal process. We present measured results on the full-supply performance of the datapath, performance-supply scaling, and regulator efficiency.

In Figure 7, we show the control signals PC and EVAL and the handshaking signal ACK of three consecutive pipeline stages. These signals are directly measured on-chip using GGB picoprobe Model 34A. Ringing in the signal is actually due to the relatively long ground wire of the probe. This is verified by a cleaner signal when the measurement is done using signal-ground probe. The signals were captured when the internal supply voltage was at 2.48V, showing a cycle time of 1.3ns.

Figure 8 shows the supply voltage measured from the ADC output and one of the PC signals measured on-chip. The system is running four instructions, each specifying a different performance. The system continues to function during supply-voltage transitions and the PC signal amplitude and period scale accordingly. Between instructions, the datapath is reset and the pipeline stops "ticking."




Figure 7: Full supply performance of the datapath adder.

The inset of Figure 8 shows the power and performance of the processor with voltage scaling of the datapath. The supply voltage is that measured by the on-chip ADC. At the full supply of 2.480V, the datapath adder runs at 1.3ns (770MHz) and burns 195mW. At the supply of 660mV, the circuit cycle time is about 21.06ns (47.5MHz) and power consumption is 850 μ W.

In Figure 9, we show the simulated and measured efficiency of the power management system. Below 1.0 V, the switched-capacitor power supply is engaged to provide an efficiency "boost" at the lowest supplies. The heavy-loading curves are simulated with a large diode connected NMOS transistor. The medium-loading curves are simulated with the same type of load of about half the strength. The measured results reflect the actual load of the datapath; the efficiency "boost" due to the switched-cap supply below 1.0 V is evident.

Conclusions

In this paper, we have described the design of a high-performance asynchronous micropipelined datapath that provides robust interfaces across voltage domains, performing appropriate voltage level conversions and operating between domains with delays differing by almost two orders of magnitude. With software-specified performance, the power supply of the datapath is scaled from 2.5V to 650mV using an on-chip dc-dc conversion that combines linear regulators and switch-capacitor power supplies.

Acknowledgements

The authors gratefully acknowledge the contributions of A. Jose, K. Zhang, and B. Liu to the physical design of the chip. The authors further acknowledge S. M. Nowick and M. Singh for many helpful discussions.

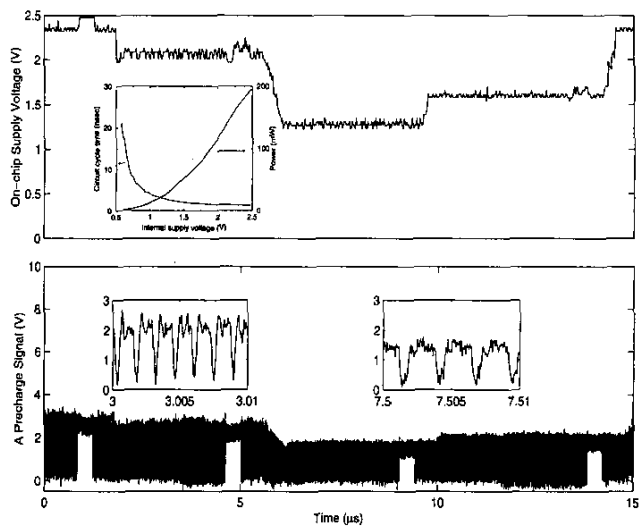


Figure 8: Datapath functionality with scaling supply.

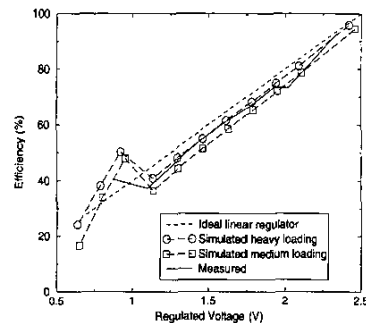


Figure 9: Measured and simulated efficiency of the power management system.

References

- [1] T. Burd and R. Brodersen, "Design issues for dynamic voltage scaling," ISLPED 2000.
- [2] J. E. Gunn, *et al*, "A low-power DSP core-based software radio architecture," IEEE J. Sel. Area. Comm. 17, 1999, pp. 574-590.
- [3] T. Williams and M. Horowitz, "A zero-overhead self-timed 160ns 54b CMOS divider," IEEE JSSC, Nov. 1991, pp. 1651-1661.
- [4] B. Amrutur and M. Horowitz, "A replica technique for wordline and sense control in low-power SRAMs," IEEE JSSC, 1998, pp. 1208-1219.
- [5] K. Y. Yun, "Automatic synthesis of extended burst-mode circuits using generalized C-elements," Proc. EuroDAC, 1996, pp. 290-295.
- [6] M. Singh, *et al*, "An adaptively-pipelined mixed synchronous-asynchronous digital FIR filter chip operating at 1.3 GHz," Proc. ASYNC 2002.
- [7] S. Schuster, *et al*, "Asynchronous interlocked pipelined CMOS circuits operating at 3.3-4.5 GHz," ISSCC 2000, pp. 292-293.
- [8] K. Usami and M. Horowitz, "Clustered voltage scaling technique for low-power design," Proc. Workshop on Low Power Design, 1995.
- [9] A. P. Dancy, R. Amirtharajah, and A. P. Chandrakasan, "High-efficiency multiple-output dc-dc conversion for low-voltage systems," IEEE Trans. VLSI, 2000, pp. 252-263.
- [10] A. J. Stratakos, *et al*, "A low-voltage CMOS dc-dc converter for a portable battery-operated power-supply regulator," Power Elec. Spec. Conf, 1994.
- [11] G.-Y. Wei and M. Horowitz, "A fully digital, energy-efficient adaptive power-supply regulator," IEEE JSSC, April, 2000, pp. 520-528.