DTU Library

Low-power operation using self-timed circuits and adaptive scaling of the supply voltage

Nielsen, Lars Skovby; Niessen, C.; Sparsø, Jens; Berkel, C. H. van

[Link back to DTU Orbit](#)

Special Issue Papers———————————————————————

Low-Power Operation Using Self-Timed Circuits and Adaptive Scaling of the Supply Voltage

Lars S. Nielsen, Cees Niessen, Jens Sparsø, and Kees van Berkel

Abstract—Recent research has demonstrated that for certain types of applications like sampled audio systems, self-timed circuits can achieve very low power consumption, because unused circuit parts automatically turn into a stand-by mode. Additional savings may be obtained by combining the self-timed circuits with a mechanism that adaptively adjusts the supply voltage to the smallest possible, while maintaining the performance requirements. This paper describes such a mechanism, analyzes the possible power savings, and presents a demonstrator chip that has been fabricated and tested. The idea of voltage scaling has been used previously in synchronous circuits, and the contributions of the present paper are: 1) the combination of supply scaling and self-timed circuitry which has some unique advantages, and 2) the thorough analysis of the power savings that are possible using this technique.

I. INTRODUCTION

THE DOMINANT source of power dissipation in digital CMOS circuits is the dynamic power dissipation:

$$P_{\text{dynamic}} = a \cdot f_{\text{clk}} \cdot C_L \cdot V_{DD}^2 \qquad (1)$$

where f_{clk} is the switching frequency, C_L the total node capacitance in the circuit, a the average fraction of the total node capacitance being switched (also referred to as the activity factor), and finally V_{DD} the supply voltage.

For a given technology and application, the power consumption can be minimized by reducing V_{DD} and/or a.

Reducing V_{DD} leads to an increase in circuit delays. With good accuracy, the circuit delay can be estimated using the following equation [1], where μ is the mobility, C_{ox} the oxide capacitance, V_t the threshold voltage, and W/L the width to length ratio of transistors.

$$T_d = \frac{C_L \cdot V_{DD}}{\mu C_{\text{ox}}(W/L)(V_{DD} - V_t)^2}. \qquad (2)$$

The activity factor a can for example be reduced by avoiding glitches or by gating the clock. While synchronous circuits

require special design effort and clock gating circuitry, self-timed circuits inherently avoid redundant transitions. For this reason, self-timed circuits have attracted more attention in recent years, particularly in areas where the computational complexity is strongly data dependent. An example of work in this area is the error corrector for the Digital Compact Cassette (DCC) player developed at Philips Research Laboratories [2]. This circuit dissipates about 80% less power than its synchronous counterpart. For an introduction to self-timed circuit design we refer to [3]–[5].

Another advantage of self-timed circuits is the ability to exploit variations in fabrication process and operating conditions in the best possible way. The performance of the chip depends on actual circuit delays, rather than on worst-case delays.

This paper describes a technique that combines self-timed circuitry with a mechanism that adaptively adjusts the supply voltage to the minimum possible, taking into account: process variations, operating conditions, and data dependent computation times. Adaptive supply scaling or "just-in-time processing" has been studied both at Philips Research Laboratories [6] and at the Technical University of Denmark [7] and this paper combines the experiences of the two parties.

The idea of voltage scaling has been used previously in synchronous circuits, and the contributions of the present paper are: 1) the combination of supply scaling and self-timed circuitry which has several unique advantages, and 2) the thorough analysis of the power savings that are possible using this technique.

The paper is organized as follows. Section II presents the concept of adaptive supply scaling. Section III provides an analysis of the power savings that can be obtained, and Section IV extends the analysis to include the effects of velocity saturation and short-circuit currents. Section V presents a small demonstrator chip that has been fabricated and tested. Section VI concludes the paper and discusses some open questions.

II. ADAPTIVE SUPPLY SCALING

In this section a system architecture for adaptive supply scaling is proposed and suitable applications are discussed. In the end of the section the approach is related to voltage scaling in synchronous circuits.
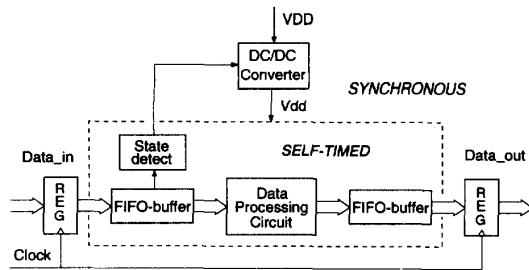
Fig. 1.   Self-timed circuit in synchronous environment using adaptive supply scaling.

Fig. 2.   Normalized throughput versus supply voltage for the DCC error corrector.

A. System Architecture

The proposed system using adaptive supply scaling in combination with self-timed circuitry is shown in Fig. 1. The system consists of the data processing circuit itself, two FIFO-buffers, a state detecting circuit, and a DC–DC converter for scaling down the supply voltage. The converter can be anything from a resistive device (a transistor on the chip) to a more sophisticated lossless device.

The actual design of the DC–DC converter is beyond the scope of this paper. Current research in low-power portable electronics includes design of low-voltage on-chip DC-DC converters, and efficiencies above 90% have been reported. A few pointers to this literature are: [8],[9].

The self-timed circuit is operating in a synchronous environment and the requirements are therefore, that the input buffer never runs full, and that the output buffer never becomes empty. With this requirement, synchronization problems will not occur at the synchronous/asynchronous interface.

The state detecting circuit monitors the state of one of the buffers, for example the input buffer as shown in Fig. 1. In this case, if the buffer is running empty, the circuit is operating too fast, and the supply voltage can be reduced. If, on the other hand, the buffer is running full, the supply voltage must be increased. The alternative is to monitor the output FIFO, and the state of the buffer must then be interpreted in a complementary way: a buffer running full should lead to a lower supply voltage, and vice versa. In this way the supply voltage will be adjusted to the actual workload, at all times maintaining the throughput requirements at lowest possible supply voltage.

The synchronous embedding shown in Fig. 1 was used for illustration purposes. Adaptive supply scaling may be used in a wider range of applications. Furthermore the architecture in Fig. 1 uses two FIFO's. In many cases one of the FIFO's may be omitted, because of particular characteristics of the algorithm/application to be implemented, or because buffering is provided by the environment [6]. This leaves the other FIFO to be part of the feed-back loop.

B. Suitable Applications

Adaptive supply scaling is particularly useful in systems with highly sequential algorithms that perform a large number of computation steps per data item, and where the computation time is d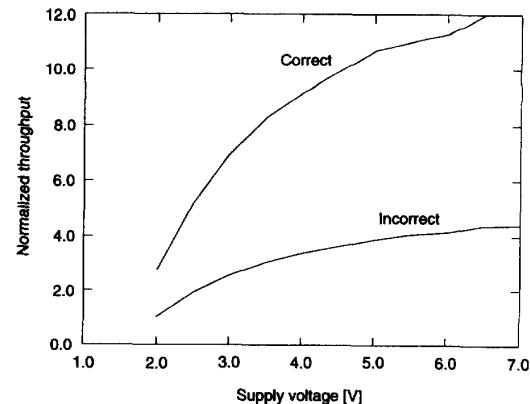ata dependent. In addition, many systems are designed for worst-case conditions in order to guarantee response time, and therefore they possess a great unused speed potential. A safety margin of 2.5 is common in synchronous circuits, to accommodate variations in process and operating conditions. The idea is to convert this speed potential into a corresponding power saving, by reducing the power supply until the delay of the computation just fits the available time slot. The FIFO-buffers allow for averaging, which enables the system to exploit data dependencies.

Two factors limit the usefulness of the approach, 1) the FIFO-buffers add to the latency as seen by the environment, and 2) V_{dd} should only vary at a slow rate relative to the internal operational speed of the circuit, otherwise it may interfere with the operation of the circuit (signal levels, noise margin, etc.). In many applications latency is not a critical issue, and even in real time audio systems a latency in the order of a few milliseconds is acceptable. The limitations on the dynamics of V_{dd} makes the technique most suitable for applications with moderate throughput requirements where the external and internal frequency of operation differ by one or more orders of magnitude.

Examples of algorithms/applications that are particularly suited for adaptive supply scaling are, sampled audio systems, floating-point units, and error correction. For instance, the DCC error corrector described in [2] processes code words of 32 bytes and the processing time of a code word depends critically on its correctness. The measured throughput for correct code words is three times that for incorrect code words (cf. Fig. 2). Given that over 95% of the code words are correct, the DCC error corrector can operate below 2 V most of the time. Only a sequence of incorrect codewords would scale up V_{dd}.

C. Relation to Existing Techniques

Having presented the key ideas, it is relevant to relate the approach to voltage scaling in synchronous circuits.

One approach is to derive the supply voltage from the clock frequency as described in [10]. Here, a self-regulating voltage reduction circuit adjusts the internal supply voltage to the

lowest value compatible with chip speed requirements, taking temperature and technology parameters into account. This is done using a phase locked loop where the clock signal is compared with the output of an on chip ring oscillator, whose delay-V_{dd} properties are assumed to be proportional to those of the actual circuit.

This mechanism obviously has some resemblance to the method described in this paper. However, there are some important differences and advantages that originate from the very nature of self-timed circuits — the handshaking that signals when computations have finished:

- In the self-timed approach the feedback is based on the actual delays in the circuit. This makes it more robust. The ring oscillator may provide a good match for static CMOS circuitry, but for circuits including pass-logic, memories and other irregular parts the matching may be less accurate. This means that a safety margin has to be introduced, reducing the power savings.
- The self-timed approach takes advantage of data dependencies, and that can contribute significantly to the power savings (cf. Section III-D).
- The feed-back signal controlling the DC-DC converter is easily derived from the FIFO's, and the FIFO's further smoothens fluctuations in workload, which again tends to filter out fluctuations in V_{dd}.

III. ANALYSIS OF POWER SAVINGS

In this section, the power savings made possible by the use of adaptive supply scaling will be estimated based on first order approximations of circuit delays. In Section IV the analysis is extended to include the effects of short-circuit currents and velocity saturation. First, the fabrication process, operating conditions, and data dependencies are considered with a lossless DC-DC converter, and second, the power loss related to the converter is taken into account. In order for the results to be independent of the fabrication process, the supply voltage V_{DD} will be normalized with respect to the threshold voltage V_t.

It should be noted that all estimations in the analysis will be based on a self-timed circuit with a constant throughput requirement, and no comparison between self-timed circuits and synchronous circuits is made.

A. Power Versus Delay

A circuit designed for worst-case conditions, allows for supply scaling, when worst-case conditions are not present. Operating the circuit at a fixed supply voltage V_{DD}, leads to the power consumption $P(V_{DD})$ and scaling the supply voltage to V_{dd}, leads to $P(V_{dd})$. The power reduction γ can thus be expressed:

$$\gamma = \frac{P(V_{dd})}{P(V_{DD})} = \left(\frac{V_{dd}}{V_{DD}}\right)^2. \tag{3}$$

In the typical case (typical process and operating conditions), the supply voltage can be reduced until circuit delays $T_{d,\mathrm{typ}}$ match those determined by the worst-case conditions
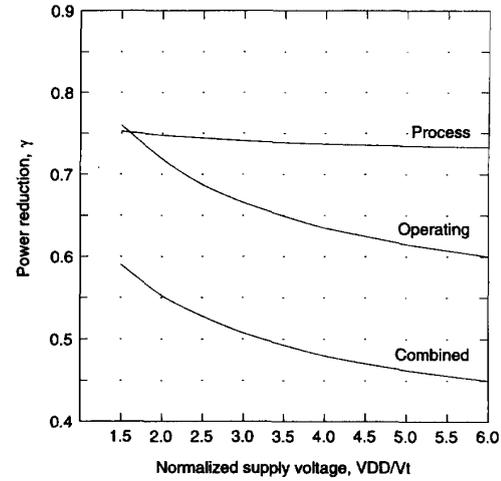

Fig. 3. Power savings due to variations in process, in operating conditions, and both.

$T_{d,\mathrm{worst}}$. Using (2) the reduced supply voltage is found by solving the following equation for V_{dd}:

$$T_{d,\mathrm{typ}}(V_{dd}) = T_{d,\mathrm{worst}}(V_{DD})$$
$$\Rightarrow \frac{V_{dd}}{\mu_{\mathrm{typ}}C_{\mathrm{ox}}(W/L)(V_{dd} - V_{t,\mathrm{typ}})^2}$$
$$= \frac{\alpha_{VDD}V_{DD}}{\alpha_T\alpha_p\mu_{\mathrm{typ}}C_{\mathrm{ox}}(W/L)(\alpha_{VDD}V_{DD} - \alpha_{\mathrm{th}}V_{t,\mathrm{typ}})^2}. \tag{4}$$

A number of coefficients are introduced in this equation to accommodate process variations and operating conditions:
Operating:

$$\alpha_{VDD} = V_{DD,\mathrm{worst}}/V_{DD}$$
$$\alpha_T = \mu_{\mathrm{typ}}(T)/\mu_{\mathrm{typ}}(T_0).$$

Process:

$$\alpha_{\mathrm{th}} = V_{t,\mathrm{worst}}/V_{t,\mathrm{typ}}$$
$$\alpha_p = \mu_{\mathrm{worst}}(T_0)/\mu_{\mathrm{typ}}(T_0).$$

where T is the temperature. As it can be seen, only variations in μ, V_{DD} and V_t are included in the analysis, and a distinction between the influence of operating conditions and process variations is made.

B. Process Variations

To estimate the amount of power that can be saved due to a typical process outcome, (4) is solved with a 15% variation on both μ and V_t, leading to $\alpha_p = 0.85$ and $\alpha_{\mathrm{th}} = 1.15$. These values stem from the technology used for the demonstrator circuit described in Section V and are representative for typical 1 μm CMOS processes. The result is shown in Fig. 3 labeled "Process".

The figure shows, that the power reduction is approximately constant over the supply voltage range, and that the dissipated power, for the typical case, is 3/4 of that dissipated in the worst-case. In case of the best fabrication process, the

dissipation will be approximately half of that in the worst case (not shown in Fig. 3).

C. Operating Conditions

Operating conditions influence circuit delays through α_T (temperature) and α_{VDD} (supply voltage) in (4). When temperature rises, mobility decreases [11]:

$$\mu(T) = \left(\frac{T}{T_0}\right)^{-\frac{3}{2}} \mu(T_0) = \alpha_T \cdot \mu(T_0). \qquad (5)$$

The exponent is an empirical value, and values ranging from 1.5 to 2 are reported in [11]. In order not to overestimate the possible power savings, the value 1.5 is used in this analysis. Using $T_0 = 300$ K as typical operating condition and $T = 350$ K as worst case gives $\alpha_T = 0.80$, and with a 10% tolerance on the supply voltage, $\alpha_{VDD} = 0.9$. With these numbers, an estimation of the power dissipation in the typical case, compared to that in the worst case, can be made. The result is shown in Fig. 3 labeled "Operating".

The combined effects of process variations and operating conditions, can also be found in Fig. 3 with the label "Combined". At $V_{DD} = 3V_t$, the power dissipation can be approximately halved.

D. Data Dependencies

A simple model is introduced to quantify the possible power savings due to data dependencies, i.e. variations in workload: For each input data, the system makes a sequence of computations, which is data dependent. Using this model, the workload can be expressed as a "cycle utilization" or "duty cycle" factor d, corresponding to the average number of computation steps divided by the worst-case number.

With duty cycle d, the circuit delay can be scaled by $1/d$, yielding a cycle utilization equal to 1. Including this dependency in (4), gives:

$$T_{d,\mathrm{typ}}(V_{dd}) = \frac{1}{d} \cdot T_{d,\mathrm{worst}}(V_{DD})$$

$$\Rightarrow \frac{V_{dd}}{\mu_{\mathrm{typ}}C_{\mathrm{ox}}(W/L)(V_{dd} - V_{t,\mathrm{typ}})^2}$$

$$= \frac{1}{d} \cdot \frac{\alpha_{VDD}V_{DD}}{\alpha_T\alpha_p\mu_{\mathrm{typ}}C_{\mathrm{ox}}(W/L)(\alpha_{VDD}V_{DD} - \alpha_{\mathrm{th}}V_{t,\mathrm{typ}})^2} \qquad (6)$$

from which V_{dd} can be derived. The reduction of V_{dd}, caused by the reduced workload, is not the only effect that will influence the power reduction γ. When $d < 1$, less work is being done, leading to a linear reduction of the power dissipation based on this effect alone. The power reduction can thus be expressed as a combination of the two effects:

$$\gamma = d \cdot \frac{P(V_{dd})}{P(V_{DD})} = d \cdot \left(\frac{V_{dd}}{V_{DD}}\right)^2 \qquad (7)$$

To estimate the influence of data dependencies on power reduction, the α-coefficients in (6) are all set to one, and V_{dd} is found and inserted into (7). In Fig. 4 the power reduction γ is plotted as a function of d for two examples: $V_{DD} = 3V_t$ and
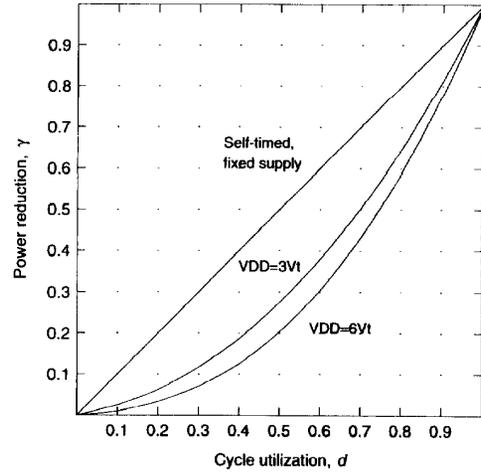

Fig. 4. Power savings due to data dependencies.

$V_{DD} = 6V_t$. For comparison the figure also shows the power reduction in a self-timed circuit with a fixed supply voltage.

It is notable that for large values of V_{DD} the execution frequency is proportional to the supply voltage (refer to (2)), and since the execution frequency scales with d, $V_{dd}(d)$ can be expressed:

$$V_{dd}(d) = d \cdot \frac{C_L}{\mu C_{\mathrm{ox}}(W/L)} \quad \text{for } V_{DD} \gg V_t \qquad (8)$$

which is linear in d. Combining this result with (7):

$$\gamma \approx d^3 \quad \text{for } V_{DD} \gg V_t. \qquad (9)$$

For $V_{DD} = 3V_t$ in Fig. 4, $(V_{dd}/V_{DD})^2 \approx d$ and therefore:

$$\gamma \approx d^2 \quad \text{for } V_{DD} = 3V_t.$$

In summary, the power reduction in a self-timed circuit with fixed V_{DD} is proportional to d, and the power reduction in a self-timed circuit with adaptive supply scaling can range from d^2 to d^3 when $V_{DD} > 3V_t$.

E. Circuitry for Supply Scaling

Adaptive supply scaling involves two power losses: one corresponding to the circuit overhead, and another to the efficiency of the DC-DC converter. The power loss in the circuit overhead (the FIFO-buffers and state detecting circuit) can be relatively small and is ignored in the analysis. The power loss in the DC-DC converter, on the other hand, can be quite significant, depending on the type of converter being used. In the analysis a resistive approach is used as the worst case and a lossless converter as the best case. Using a resistive approach the power saving γ is reduced to (cf. (3)):

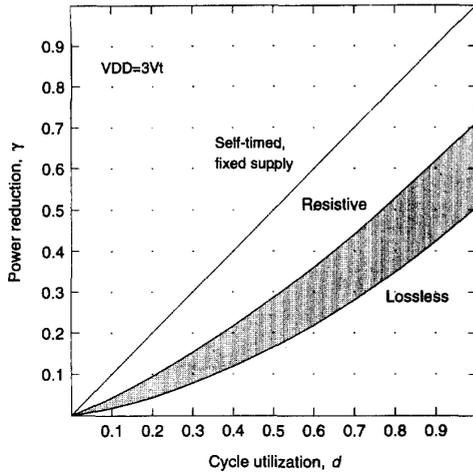$$\gamma = \frac{V_{dd}}{V_{DD}}. \qquad (10)$$

Fig. 5. The combined effects of process variations, operating conditions, and data dependencies for the typical case.

F. Summary — Combining the Effects

Fig. 5 shows the possible power savings, using both a lossless and a resistive converter, when all the effects described above are combined: 1) typical process, 2) typical operating conditions, and 3) data dependencies. A supply voltage of $V_{DD} = 3V_t$ is assumed (corresponding to $V_{DD} = 2.4$ V in a typical CMOS process). For larger values of V_{DD} the possible power saving increases, and for smaller values it decreases. As V_{DD} approaches $2V_t$ the circuit delays increase drastically, enabling only marginal power savings.

At $d = 1$ the supply voltage is reduced to $V_{dd} = 2.1V_t$ due to typical process and operating conditions. At $d = 0.35$ the supply voltage is reduced to $V_{dd} = 1.6V_t$. As a reference, Fig. 5 also shows the power savings in a self-timed circuit with fixed supply voltage.

Comparing a self-timed circuit using adaptive supply scaling with a self-timed circuit using a fixed supply, two interesting cases are:

- For a worst-case computation ($d = 1$) and $V_{DD} = 3V_t$, the power saving, using a resistive supply scaling, is a factor of 1.4. This is a lower bound on the saving.
- For a computation with data dependency ($d = 0.35$) and $V_{DD} = 3V_t$, the power saving using a lossless supply scaling is a factor of 3.6, and for $V_{DD} = 6V_t$ (= 5 V for $V_t = 0.83$ V) the power saving is a factor of 6.4.

This latter example corresponds to the cycle utilization factor of the DCC error corrector described in Section II-B. As this chip has a rather low cycle utilization, the factor of 6.4 may be considered as an upper bound on the possible power savings in general.

IV. REFINING THE ANALYSIS

In this section the analysis is extended to include the effects of short-circuit currents and velocity saturation. Both effects lead to additional power savings, but it should be noted that the impact on power reduction is strongly technology and design dependent.

A. Short-Circuit Dissipation

In (1) short-circuit dissipation was ignored as a contribution to dynamic power dissipation. This form of dissipation occurs during a gate-output transition, when both the n-path and the p-path conduct. The short-circuit current may be substantial (both transistor paths are in saturation), but this lasts only for the duration of the corresponding input transition. For carefully designed circuits the short-circuit dissipation is typically about 20% of the dynamic dissipation for a channel length of 1 μm and $V_{DD} = 5$ V. The short-circuit dissipation for $V_{dd} \geq 2V_t$ is given by [12]:

$$P_{\text{short}} = \frac{\beta}{12} \cdot (V_{dd} - 2V_t)^3 \cdot \frac{\tau}{T_p} \tag{11}$$

where β is the gain factor of a MOS transistor, τ the rise or fall time, and T_p the clock period. Both τ and T_p scale with the supply voltage and therefore:

$$P_{\text{short}} \sim V_{dd}^3 \quad \text{for } V_{dd} \gg 2V_t. \tag{12}$$

Hence, the power reduction due to down scaling of the supply voltage is even more attractive than implied by (3). For $V_{dd} < 2V_t$ short-circuit dissipation is negligible.

B. Velocity Saturation

Velocity saturation [13] is a phenomenon that is becoming more and more significant as technology is being scaled down. Due to velocity saturation the performance of CMOS circuits grows less than linearly with $V_{dd} - V_t$, as suggested by (2). For some technologies the velocities of electrons and holes in MOS channels tend to saturate beyond an electric field $\mathcal{E}$ of 2–6 V/μm. At $V_{DD} = 5$ V this effect may reduce saturation currents (and therefore the performance) by more than a factor of two! The good side of velocity saturation is that, when V_{dd} is scaled down the corresponding performance loss will be modest. This implies that the power savings can be substantially better than estimated in the previous section.

The significant impact velocity saturation can have on circuit performance, is well illustrated by the throughput of the DCC error corrector shown in Fig. 2. The technology used for this chip has a critical electric field $\mathcal{E}_c = 1.7$ V/μm which leads to substantial performance degradation at high supply voltages. In the figure this is seen by the rapid decline of the slope of the normalized throughput versus V_{DD} graph.

Modifying (2) for velocity saturation with L being the length of transistor channels we get [13, Eqn. 5.3.10]:

$$T_d' = T_d \cdot \left(\frac{V_{dd}}{L\mathcal{E}_c} + 1 \right)$$

$$= \frac{C_L \cdot V_{DD}}{\mu C_{ox}(W/L)(V_{DD} - V_t)^2} \cdot \left(\frac{V_{dd}}{L\mathcal{E}_c} + 1 \right). \tag{13}$$
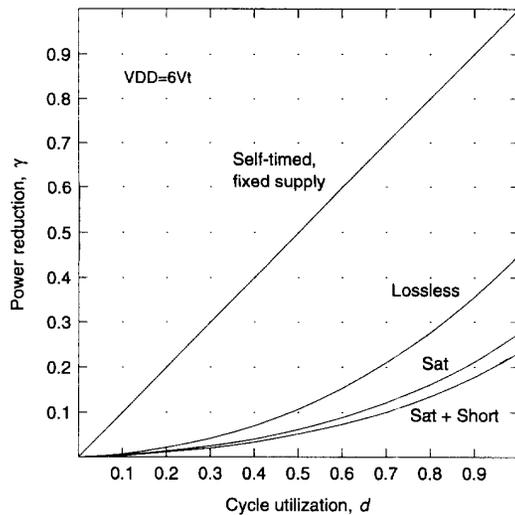
Fig. 6. Power reduction due to typical process and operating conditions when considering the effect of velocity saturation, labeled "Sat", and the combined effect of velocity saturation and short-circuit dissipation, labeled "Sat + Short". In all cases a lossless supply scaling is assumed.

C. An Example

As mentioned before the effects of short-circuit currents and velocity saturation depends very much on the technology and the particular design. To provide a quantitative analysis of the power savings we assume a technology with $V_{DD} = 6V_t$, $L\mathcal{E}_c = 2V_t$, and a design where short-circuit dissipation equals to 20% of the dynamic power dissipation at $V_{DD} = 6V_t$.

Using these figures and (13), the power savings due to the presence of velocity saturation may be estimated. The result is shown by the graph labeled "Sat" in Fig. 6. For comparison the figure also shows the power reduction in a self-timed circuit using a fixed supply voltage, and the estimated power reduction using adaptive supply scaling with a lossless DC-DC converter, when typical process, typical operating conditions, and data dependencies are considered (similar graphs are found in Fig. 5 for $V_{DD} = 3V_t$). The figure shows a substantial power reduction. At worst case computation ($d = 1$) the velocity saturation leads to additional power savings of 40%.

The effect of short-circuit dissipation can be estimated using (11). Combining this with the effect of velocity saturation the graph labeled "Sat + Short" in Fig. 6 is obtained. As V_{dd} is scaled down the short-circuit current is reduced, but even for worst case computations ($d = 1$) the scaling of the supply voltage is significant enough to make short-circuit dissipation negligible ($P_{\text{short}}(V_{dd})/P_{\text{short}}(V_{DD}) = 0.02$). It is now possible to reduce the power dissipation by a factor of 13 at $d = 0.35$.

V. THE DEMONSTRATOR CHIP

A test chip has been fabricated via Eurochip in a 1.0 μm CMOS process. The chip contains a system for adaptive supply scaling, identical to the self-timed circuit in Fig. 1. It is noted that the supply scaling is performed off-chip, thereby allowing for experimentation with different circuit configurations.
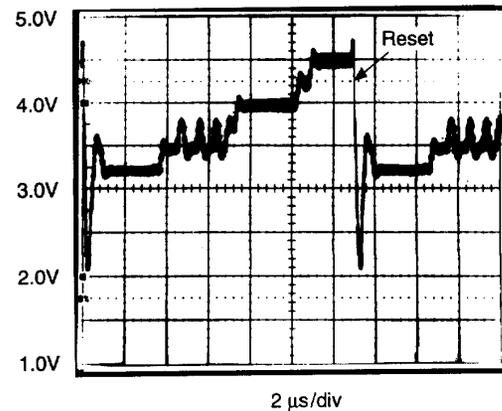

2 μs/div

Fig. 7. Oscilloscope snap shot of adaptive supply scaling.

The input FIFO is a 10 word deep buffer implemented using the latch type in [4]. From this, 9 state bits are generated in the state detecting circuit and fed to the external supply scaling circuit. The off-chip scaling circuit, used in the test setup, is a simple D/A converter, which scales the supply voltage linearly depending on the number of data words in the input FIFO.

The main circuit is a delay-insensitive circuit that implements a 16-bit dual-rail ripple-carry adder. In the adder circuit two Cascode Voltage Switch Logic (CVSL) function blocks [14] are used, one for the sum function and one for the carry. The scheme of indication in the ripple adder is identical to the one used by Martin in [15], which utilizes the carry-kill and carry-generate properties of the full-adder. Using this approach, the delay of the addition is data dependent, ranging from 1 to 16 times the delay of one full-adder. The circuit is therefore suitable for validation of the power estimations, based on data dependencies, found in Section III.

Fig. 7 shows an oscilloscope picture obtained by cyclically applying input data composed of sequences of operands that cause the carry to ripple 4, 8, 12, and 16 positions, respectively. Each cycle is initiated by a reset and the operands are input at a 16 MHz rate.

The figure shows that the first, third, and fourth sequence lead to stable supply voltages, whereas the second sequence (where the carry ripples 8 positions) requires a supply voltage different from the discrete voltages available by the supply scaling in this configuration. Therefore the supply voltage fluctuates between two adjacent supply voltage levels.

VI. CONCLUSION

In this paper we have described a technique that can increase the power savings for self-timed circuits as much as an order of magnitude or more. The technique is called "adaptive supply scaling" or "just-in-time processing" and is particularly useful in systems that implement sequential algorithms with a data dependent computation time.

The fabricated test chip nicely demonstrates the feasibility of adaptive supply scaling, but is clearly not a practical application (the relative cost of the overhead circuitry, is prohibitive in this case). For "adaptive supply scaling" or "just-in-time

processing" to pay off, larger subsystems must be considered. An example is the DCC error corrector discussed above. On the other hand the subsystem can also be too big. For example, it is sometimes the case that substantial but *local* variations in workload only lead to minor variations in workload at the subsystem level. The granularity is therefore a topic that requires further research. Other open questions are: 1) Are circuits that operate on multiple adaptively scaled supply voltages stable? If not in general, under which conditions? 2) What are minimal and optimal buffer sizes, given the time constants involved? 3) Is latch-up a problem after a sudden drop in V_{dd}?

ACKNOWLEDGMENT

REFERENCES

[1] A. P. Chandrakasan, S. Cheng, and R. W. Brodersen, "Low-power CMOS digital design," *IEEE J. Solid-State Circ.*, vol. 27, no. 4, pp. 473–483, April 1992.

[2] C. H. van Berkel, R. Burgess, J. Kessels, A. Peeters, M. Roncken, and F. Schalij, "Asynchronous circuits for low power: a DCC error corrector," *IEEE Design & Test*, vol. 11, no. 2, pp. 22–32, 1994.

[3] I. E. Sutherland, "Micropipelines," *Communications ACM*, vol. 32, no. 6, pp. 720–738, June 1989.

[4] J. Sparsø and J. Staunstrup, "Delay-insensitive multi-ring structures," *Integration, the VLSI J.*, vol. 15, no. 3, pp. 313–340, Oct. 1993.

[5] G. Birtwistle and A. Davis, Ed., *Proc. Banff VIII Workshop: Asynchronous Digital Circ. Design, Banff, Alberta, Canada, Aug. 28–Sept. 3, 1993.*, Springer Verlag, Workshops in Computing Science, 1995. Contributions from: S. Furber, "Computing without clocks: Micropipelining the ARM processor." A. Davis, "Practical asynchronous circuit design: Methods and tools." K. van Berkel, "VLSI programming of asynchronous circuits for low power." J. Ebergen, "Parallel program and asynchronous circuit design."

[6] C. H. van Berkel and C. Niessen, "An apparatus featuring a feedback signal for controlling a powering voltage for asynchronous electronic circuitry therein," European Patent Application 92203949.0, Published June 1993.

[7] L. S. Nielsen and J. Sparsø, "Low-power operation using self-timed circuits and adaptive scaling of the supply voltage," in *1994 Int. Workshop Low Power Design*, April 22–27, 1994. Unpublished proceedings.

[8] J. G. Kassakian and M. F. Schlecht, "High-frequency high-density converters for distributed power supply systems," *Proc. IEEE*, vol. 76, no. 4, pp. 362–376, 1988.

[9] F. Goodenough, "Off-Line and one-cell IC converters up efficiency," *Electronic Design*, pp. 55–64, June 27, 1994.

[10] P, Macken, M. Degrauwe, M. van Paemel, and H. Oguey, "A voltage reduction technique for digital systems," in *ISSCC 1990 Dig. Tech. Papers*, pp. 238–239, 1990.

[11] L. A. Glasser and D. W. Dobberpuhl, *The Design and Analysis of VLSI Circuits*. Reading, MA: Addison-Wesley, 1985.

[12] H. J. M. Veendrick, "Short-circuit dissipation of static CMOS Circuitry and Its Impact on the Design of Buffer Circuits," *IEEE J. Solid-State Circ.*, vol. SC-19, no. 4, pp. 468–473, Aug. 1984.

[13] Y. P. Tsividis, *Operation and Modeling of the MOS Transistor.* New York: McGraw-Hill Int., 1988.

[14] L. G. Heller and W. R. Griffin, "Cascode voltage switch logic: A differential CMOS logic family," *IEEE ISSCC Dig. Tech. Papers*, Feb. 1984.

[15] A. J. Martin, "Asynchronous datapaths and the design of an asynchronous adder," *Formal Methods in System Design*, vol. 1, no. 1, pp. 119–137, July 1992.

Lars S. Nielsen was born in Herning, Denmark in 1966. He received the M.Sc. degree in electrical engineering in 1992 from the Department of Computer Science, Technical University of Denmark, Lyngby, Denmark.

He is currently working towards the Ph.D. degree in electrical engineering at the Technical University of Denmark, and his research interests include self-timed circuits and low-power CMOS VLSI design.

Cees Niessen received the M.Sc. degree in electrical engineering from the Delft University of Technology.

He is a chief scientist at the IC Design Center of Philips Research Laboratories in Eindhoven, The Netherlands. His research interests include high level synthesis for digital signal processing and design for low power. In the latter area he is now coordinating an activity aiming at low-power design methods and circuitry for portable applications.

Jens Sparsø was born in Silkeborg, Denmark, in 1955. He received the M.Sc. degree in electrical engineering from the Technical University of Denmark in 1981.

Since 1982 he has been with the Department of Computer Science, Technical University of Denmark, where he became Associate Professor in 1986. He is teaching courses on VLSI and digital systems design, and his research interests are architecture and design of VLSI systems, i.e., design methods, circuit techniques and the interplay between technology and system architecture. Current activities involve the design of self-timed circuits.

Kees van Berkel received the M.Sc. (honors) degree in electrical engineering from Delft University of Technology and the Ph.D. degree from Eindhoven University of Technology.

He is a senior scientist at the IC Design Center of Philips Research Laboratories Eindhoven, The Netherlands. Currently he coordinates the work on VLSI programming at Philips Research and manages project EXACT ESPRIT 6143. His research interests include VLSI programming, VLSI architectures, compilers, asynchronous circuits, CMOS circuits, and low-power.

Dr. van Berkel is author of the book *Handshake Circuits - An Asynchronous Architecture for VLSI Programming.*