STRINGENT NEWS

Issue 2 **www.ida.liu.se/stringent** **September 26, 2005**

This is the second Newsletter from Stringent (Strategic Integrated Electronic Systems Research at Linköping University). The first Newsletter from November 2005 had Programmable RF Systems as theme and this time the theme is Low Power Design.

The newsletter is intended to keep in touch with you. We seek contact and cooperation with industry and institutes worldwide, in the field of electronics and electronic design. We are centred in the Nordic countries but consider ourselves European. The first newsletter received good response; thanks for emails and other contacts! I also noted that Dr. Hermann Eul, Infineon, referred to the Newsletter in his plenary talk at ISSCC in San Francisco in February.

Stringent is considered the largest electronics research centre in Sweden with about 50 people. It was founded 3 ½ years ago by combining the efforts from 4 professor chairs, lead by profs. Atila Alvandpour, Dake Liu, Zebo Peng and Lars Wanhammar.

I hope you enjoy our Newsletter!

— Christer Svensson, Professor and Stringent director

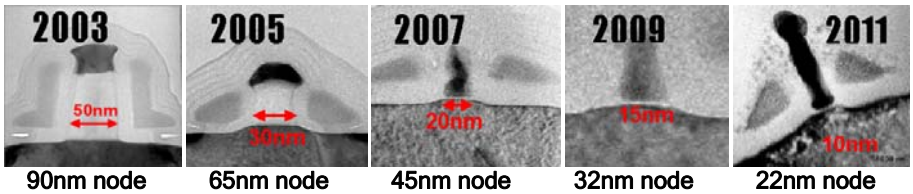

Figure 1 - Present and near-future MOS transistors (source INTEL)

Power consumption in a nano-CMOS perspective

Advanced transistors with channel lengths far shorter than 60nm (Fig. 1) have been experimentally fabricated and reported, suggesting that the technology fundamentals are still intact for decades to come. However, the semiconductor industry is experiencing a paradigm shift, where the historically excellent improvements in cost, performance, and density of integrated circuits are no longer assured by scaling of device dimensions only. The enormous capability of the present and future advanced CMOS technologies encourages increasingly large and complex chips. On the other hand, the existing design strategies, methodologies, system architectures, and circuit techniques are not sufficiently effective to support the provided device capability. This has

Great success for ESLAB

One of Stringent research groups, ESLAB (Embedded Systems Laboratory), under the leadership of Prof. Zebo Peng, has received a lot of attention recently.

The group received a new EU grant, "Verification and Validation of Embedded System Design workbench (VERTIGO)," together with, among others, ST Microelectronics (Italy), AerieLogic (France), and TransEDA Technology (UK).

They presented four papers at the DATE'06 conference in April, and were given the Best paper award for DATE'05 for their paper "Design Optimization of Time- and Cost-Constrained Fault-Tolerant Distributed Embedded Systems," by V. Izosimov, P. Pop, P. Eles and Z. Peng. Furthermore, Prof. Peng was appointed Vice Program Chair of DATE'07 in Nice and Program Chair of the12th IEEE European Test Symposium (ETS'07) in Freiburg.

Finally, they are publishing a new book, "Real-Time Applications with Stochastic Task Execution Times: Analysis and Optimisation" by S. Manolache, P. Eles and Z. Peng, with Springer.

We are proud of ESLAB!

IN THIS ISSUE

created several major design challenges, which if not addressed efficiently might set the limit on progress of microelectronics much earlier than the fundamentals.

The increasingly large on-chip power consumption is one the most serious issues. Today, low power operation is not only a critical requirement for mobile phones, laptops, and other battery-powered devices, but also a general and severe constraint for all high-performance VLSI chips! This is due to the limitations in capacity of conventional (low-cost) cooling techniques to remove the generated heat from the chips, as well as the increasing complexity and cost for supplying, distributing, and controlling large amount of currents into a small piece of silicon.

There is no single solution for 'low power', but it must be considered throughout the entire design-flow from early system definitions and choice of system architectures to energy efficient digital, analog, and RF circuit technologies. The four division-members of STRINGENT have together a broad spectrum of knowledge and experience required to support such a multidisciplinary research on design for low power. In the following, some of the recent projects and results are described.

— *Atila Alvandpour, Professor*
Division of Electronic Devices

1.56 GHz On-chip Resonant Clocking in 130nm CMOS

The increase in power dissipation in high-performance microprocessors and other advanced VLSI chips presents a serious challenge for future advanced digital circuit design. It has been well recognized that a significant fraction of the total active power consumption in multi-GHz VLSI designs is due to clocking. Modern high-performance microprocessors can dissipate more than 25 % of the total core power in the

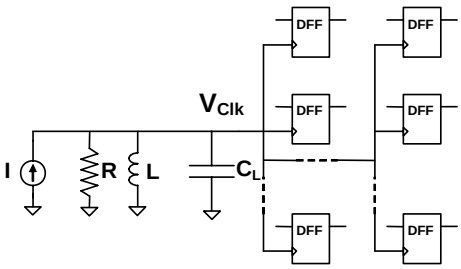

Figure 2 - Resonant clocking based on LC-tank oscillator.

clock distribution. Although conventional high performance clocking techniques are mature and robust, they are still based on a relatively rigid and traditional philosophy. This enforces power-hungry clock distribution networks based on conventional full-swing buffering, which leaves almost no room for low power.

As an alternative low power clocking solution energy-recovering clocking techniques have been investigated. The general idea behind the energy-recovering clocking is ideally a power dissipation-free (adiabatic) switching of the clock distribution. An approach that has gained increasing research interest is to utilize a high-Q LC-tank oscillator circuit to drive the entire global clock distribution without intermediate buffer stages as seen in Fig. 2. The inherent capacitive load from clocked elements and clock distribution networks act as the capacitive part in an LC-tank oscillator. The energy is thereby resonated between the inductor in the LC-tank and the clock load capacitance, and ideally no power is dissipated at resonance. Low resistive losses in both the distributed capacitance and the inductors are required in order to achieve large power savings. This therefore requires both careful design of the global clock distribution networks and the usage of high-Q inductors.

Utilizing an LC-tank based resonant clocking strategy we present the first successful chip implementation operating at GHz frequencies, using a fully integrated 1.56-GHz on-chip resonant clock oscillator (Fig. 3) in 130-nm CMOS (M. Hansson et al. CICC 2006). Detailed power measurements on the implemented test chip shows that the resonant clocking results in 57 % lower clock power and up to 29 % lower total chip power (Fig. 4), compared to the conventional clocking implemented in the same test-chip. Furthermore, clock jitter measurements show a worst-case peak-to-peak jitter of 28.4 ps across 0-to-80 % data activity in flip-flops and the data-path logic. In addition, the test-chip includes oscillation injection-locking capabilities, which enables a 50 % reduction of the worst-case peak-to-peak jitter to 14.5 ps (B. Mesgarzadeh et al. ESSCIRC 2006). Obviously several issues are left to be solved, such as improved inductors, low-loss interconnects, customized flip-flops, clock gating, and larger frequency tuning range. However, none of these issues seem to belong to fundamental barriers of CMOS technology.

— *Martin Hansson and Behzad Mesgarzadeh, Ph.D. Students*

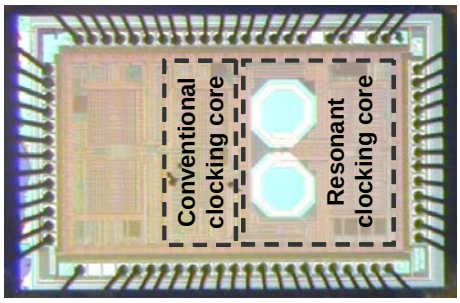

Figure 3 - Implemented test chip

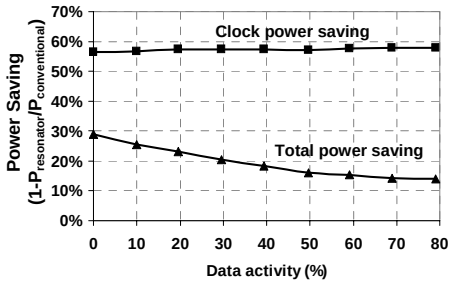

Figure 4 - Measured power saving vs. data activity (clock freq. 1.56 GHz).

Analog power consumption

Power consumption of digital systems is considered very well understood today, after 15 years of intensive research. However a similar basic understanding of analog power consumption is still lacking. An early attempt related analog power consumption to the required signal to noise ratio. This idea was later used for estimating the minimum power needed for sampling a signal, $Ps=12kTf_s2^{2n}$, where k is Boltzmanns constant, T is the absolute temperature, f_s is the sampling frequency and n is the resolution in the number of bits. We have used these same arguments as a basis for a theory for power consumption in AD-converters. We were able to show that the minimum power consumption of AD-converters is closely related to the minimum sampling power, mentioned above. What we did was to estimate the minimum power consumption needed by the key components of an ADC, samplers, comparators and amplifiers, based on speed and noise constraints. In this estimation we have not taken offset and matching problems into account, assuming that these can be mitigated through digital error correction. The impact of the need for matching was estimated to 5x in power consumption.

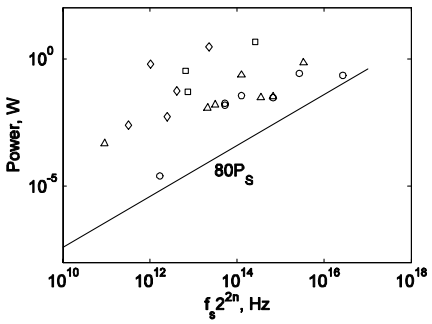

Figure 5 - Power consumption of AD-converters vs. sampling rate multiplied by 2^{2n}, where n is the number of bits. Data from ISSCC 2002 (squares and triangles) and ISSCC2006 (circles and diamonds) is compared to our minimum power estimation, 80Ps.

In figure 5 we show the power consumption of a large number of published ADCs versus a performance measure, f_s2^{2n}. In the same diagram we show our estimated minimum power for pipelined converters in a contemporary CMOS process (80Ps). We notice that most ADCs consume much more power than suggested by our estimation, but the best ones do approach our estimations. Our conclusion is that our attempt is promising and that there is a lot to gain in power savings of ADCs. We also note some improvement from 2002 to 2006, except for flash converters (diamonds and squares in the figure).

— *Christer Svensson, Professor, Electronic Devices.*

Low power FIR filters on DSPs

Normally power consumption is considered a hardware problem. However, it is quite possible to save power also through choice of algorithms or even choice of coefficients. Finite-length impulse response (FIR) filters is a common algorithm to realized on DSP processors. Many of the typical features of DSP processors such as cyclic addressing, multiply-and-accumulate (MAC) units, and hardware looping are motivated by FIR filters. Given a filter specification and a filter length it is often possible to change the coefficient values slightly without breaking the specification. As data switching is the main source for power dissipation in CMOS technologies it would therefore be advantageous to choose the coefficients such that the total number of switches between adjacent filter coefficients are small.

Using mixed integer linear programming (MILP) an approach has been proposed that minimizes the number of switched bits between adjacent filter coefficients in either two's complement or signed-magnitude representation. Example designs show that savings on the total number of switches can reach above 60% compared with simple rounding of the coefficients. This would correspond to significant savings for the multiplier and the coefficient bus without any overhead except for the off-line optimization.

Another alternative is to reorder the coefficients to minimize the total number of switches. This approach may also give large savings, especially if combined with simultaneous optimization of the coefficients. However, this may introduce drawbacks such as removal of input data correlation, more instructions and instruction fetching required and irregular data access. Hence, this should be evaluated further taking these drawbacks into consideration.

— *Oscar Gustafsson, Assistent professor, Electronic systems*

Energy Minimization in Time-Constrained Multiprocessor Systems

As a result of fast technological development and of an increasing demand for reliable embedded systems with highly complex functionality, new, sophisticated architectures, consisting of several interconnected processors are now widely used. Such systems have not only to satisfy strict time constraints but are very often working under a restricted energy budget.

The workload imposed on such an embedded system is non-uniform over time. This introduces slack times during which the system can reduce its performance to save energy. Two system-level approaches that allow an energy/performance trade-off during runtime of the application are dynamic voltage selection (DVS) and adaptive body biasing (ABB). While DVS aims to reduce the dynamic power

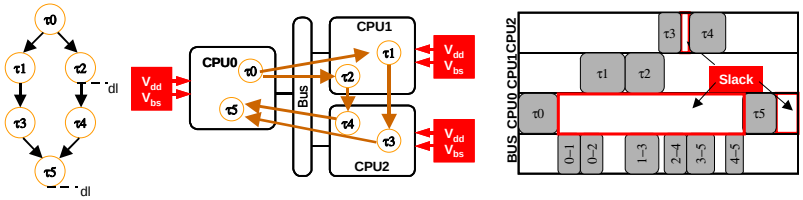

Task Graph **Mapping** **Schedule**

Figure 6 - Analysis and optimization flow

consumption by scaling down operational frequency and circuit supply voltage V_{dd}, ABB is effective in reducing the leakage power by scaling down the frequency and increasing the threshold voltage V_{th} through body-biasing.

The optimization flow is illustrated in Figure 6. The given applications are partitioned into task graphs that capture control and data dependencies. Tasks are mapped on the target hardware platform. Each task is analyzed and, depending on the processor where it is mapped, its worst-case execution time is computed. The task schedule is computed such that the precedence constraints and timing requirements of the application are met. An example schedule is illustrated in Figure 1. As we can see, when the processors are running at the highest frequency, the schedule contains a certain amount of static slack. We developed optimal algorithms for energy minimization that exploit the static slack using DVS and ABB. As a result, both dynamic and leakage energy is minimized.

This approach can be also used for the optimization of communication intensive systems. We have shown that, depending on the particular implementation style, communication energy can be saved by performing DVS and ABB on repeater based buses. In case of buses that are implemented with fat wires, applying DVS and dynamic swing scaling results in important savings.

The above-mentioned approach targets the energy minimization by exploiting the static slack. For some applications (for example an MPEG decoder), the actual execution time of the tasks, is very often shorter than their estimated worst-case, with variations of up to 10 times. Thus, an offline optimization based solely on the worst-case is too pessimistic and hampers the achievable energy savings. In order to take advantage of the dynamic slack that arises from variations in the execution times known only at runtime, it is useful to dynamically recalculate the voltage settings online, at the end of each task. Dynamic approaches, however, suffer from the significant overhead in terms of execution time and power consumption caused by the online voltage calculation.

We have proposed a quasi-static voltage selection technique that is able to exploit the dynamic slack and, at the same time, keeps the online overhead (required to recompute the voltage settings at run-time) extremely low. This approach is based on an offline calculation and storing in look-up tables of several voltage settings, depending on possible task finishing times. The online phase is responsible to adapt the task voltage settings at run-time, based on the information that is calculated and stored in the look-up tables.

— *Alexandru Andrei*
PhD Student, Embedded Systems

Donations to Stringent.

Stringent has received two donations recently, which we are very happy for.

We received an instrument donation from Ericsson. The donation includes instruments for RF, logic analysis and oscilloscopes, and constitutes an appreciable strengthening of our radio frequency laboratory.

We also received a donation from Texas Instruments, two Signal chain prototyping systems with DSPs. These will be very important for our signal processing lab. and our ability to test algorithms in real time.

Thanks very much Ericsson and

Contact Stringent

- **URL:** www.ida.liu.se/stringent
- **E-mail**
  Christer Svensson
  christer@isy.liu.se
- **Address**
  Dept. of Electrical Engineering,
  Linköping University,
  S-581 83 Linköping, Sweden
- **Telephone**
  +46 (0)13-28 12 23
- **Telefax**
  +46 (0)13-13 92 82

Editors: Christer Svensson
and Thomas Johansson