

IEM926: An Energy Efficient SoC with Dynamic Voltage Scaling

Krisztián Flautner

David Flynn

David Roberts

Dipesh I. Patel

{krisztian.flautner, david.flynn, david.roberts, dipesh.patel}@arm.com
ARM Limited, 110 Fulbourn Road Cambridge, UK CB1 9NJ

Abstract

One of today's most successful embedded devices, the mobile phone, embodies a set of challenging design requirements: long battery life, small size, high performance and low cost. The single parameter that complicates the simultaneous fulfilment of all of these design goals is energy efficiency of the system, since batteries only hold a finite amount of charge. To operate within the allotted energy budget, systems must be optimized for energy consumption during design and also at run-time. Increasingly it is not sufficient to statically optimize for worst-case conditions but designers must enable systems to adapt to conditions at run-time. The Intelligent Energy Manager™ (IEM) technology provides an integrated solution for addressing energy management of SoC devices. In this paper we present data about the energy consumption characteristics of a multiple power-domain based SoC which includes PDA functionality built around an ARM926EJ-S core.

1. Introduction

Power consumption is arguably the most important feature of embedded processors with significant impact on the cost and physical size of the end device. Historically, low power consumption in embedded processors has been achieved through simplicity, limited use of speculation, and through the use of low-power sleep modes that reduce idle-mode power consumption. Embedded processors are now performing more sophisticated tasks, which require ever-higher performance levels. As a result, new processor designs are more dependent on sophisticated architectural techniques (such as prediction and speculation) to achieve high performance. Unfortunately, such techniques can also significantly increase the processor's power consumption.

One way to bridge the gap between high performance and low power is to allow the processor to run at different performance levels depending on the workload's requirements. An MP3 audio player, for example, requires about an order of magnitude less performance than an MPEG video player. The difference in performance requirements can be exploited to save energy with the use of dynamic voltage scaling (DVS). DVS exploits the fact that the peak frequency of a processor implemented in CMOS is proportional to the supply voltage, while the amount of dynamic energy required for a given workload is proportional to the square of the processor's supply voltage. Reducing the supply voltage while slowing the processor's clock frequency yields a quadratic reduction in energy consumption, at the cost of increased run time [5].

The IEM technology includes software components that can accurately predict the minimum necessary performance level of the processor for the running workload, thus a reduction of performance does not necessarily imply any degradation of quality [3]. In this paper we show the potential energy savings that can be achieved on a real SoC using dynamic voltage scaling.

2. The IEM926 test chip

The IEM926 test chip was explicitly designed to support DVS and fast clock-switching and includes on-chip peripherals that are similar to the ones found on PDA devices. The test chip includes the following components, graphically illustrated in Figure 1:

- ARM926EJ-S processor with caches (16K I and D).
- 16K I and D Tightly Coupled Memories (TCMs).
- 240, 180, 120, 60, 0 MHz processor performance levels.
- A DMA subsystem.
- The Intelligent Energy Controller prototype.
- SDRAM and Flash memory controllers and basic peripherals (including on-board audio) to support a minimal Linux environment.
- Interface to National Semiconductor's PowerWise™ controller to support open- and closed-loop DVS [6].

The SoC is partitioned into three power domains. The system bus and peripheral bus subsystems are in a single power domain supplied with a fixed 1.2V. The CPU domain is the only domain whose frequency and voltage can be varied dynamically and it includes a separate power domain for the TCMs which can be placed in a low-power state retention mode while the main processor is powered off. The design includes clamps between the TCMs and the core to

FIGURE 1. Components of the IEM926 SoC

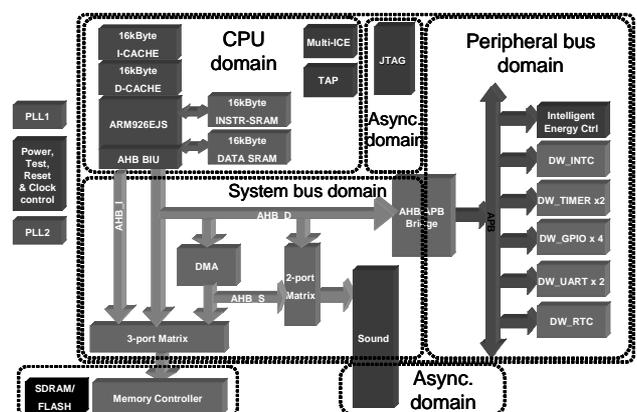
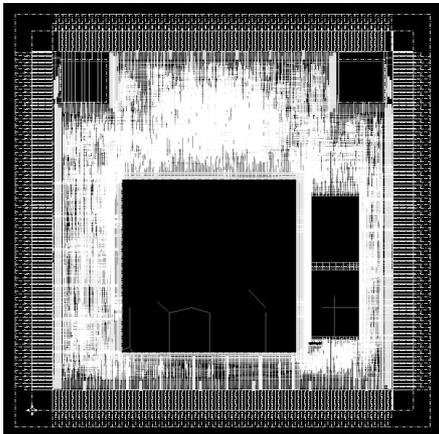


FIGURE 2. IEM926 die photo



support this mode of operation, however when running, both the TCM and core run at the same voltage and frequency.

The test chip was manufactured using the TSMC 0.13G process. A picture of the 5x5mm die is shown in Figure 2 without the processor (middle box), two PLLs (top left and right corners) and instruction and data TCMs (middle right box). The system includes two PLLs: one controls the frequency of the processor and another provides a fixed frequency for the peripherals.

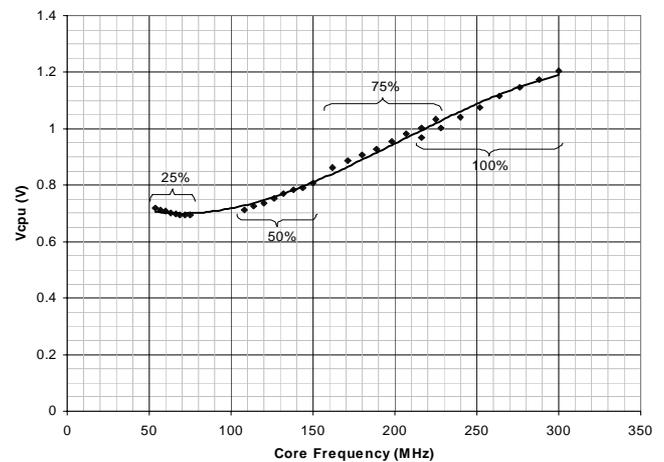
2.1 Clocking strategy

The two main challenges of the SoC design were to support fast switching between the available frequency levels and to support dynamic frequency changes on a core with only synchronous bus interfaces. The first issue was solved by the use of frequency division of one of the PLLs running at 480MHz to four frequency levels: 240, 180, 120, and 60 MHz. On the bench the chips successfully operate at 300, 225, 150, and 75 MHz by running the PLL at 600MHz. To simplify the system design, the system bus and peripherals run at a fixed 25% of the peak frequency of the processor. Generating a frequency at 75% of peak is a challenge with a single PLL, further complicated by the need for synchronous interfaces to the buses. The solution employed in this chip relies on a skewed clock that has an uneven duty cycle (3/8 comprised of 1:1.5, 1:1.5, 1:2 core to PLL clock ratios), ensuring that bus and core clocks are aligned on the rising edge of each bus clock transition. In the following figures, the datapoints corresponding to a wide variety of frequencies were generated by under- and over-driving the 480MHz PLL by -10% to +25% in 5% increments and then dividing by the four ratios (1, 3/4, 1/2, 1/4).

2.2 Voltage levels

Figure 3 shows the minimum voltage levels sufficient for sustaining a wide range of frequencies on the core at room temperature. The peak frequency of the core is set between 215 and 300 MHz and scaled to 75%, 50%, and 25%. Theoretical models suggest a linear relationship between voltage and frequency. Our measurements broadly confirm these expectations with two important differences:

FIGURE 3. Core voltage vs. frequency



voltages for frequencies corresponding to 75% of peak lie above the linear predictions and voltages for minimum (25%) frequencies do not substantially decrease below the levels at 50% and in fact show an increase for low frequencies.

The former irregularity is explained by the clocking technique employed on the SoC: at the 75% peak frequency point the core is actually operating slightly (a little over 6%) faster than 75% due to the interface with synchronous buses. The higher actual frequency in turn necessitates a higher operating voltage, which explains the divergence. The irregularity at low frequencies is as yet unexplained but is likely to be caused by the level-shifters employed in the system. We have also observed that the voltage characteristics when caches are turned off are substantially the same as in the graph above, thus in this case, the sense-amplifiers are not the cause of the lower limit on voltage scaling.

3. Power and energy

Figure 4 shows that there is a linear relationship between the core's frequency and the amount of work done per unit time in a processor-bound workload. As expected, running at 25% of peak frequency causes this workload to

FIGURE 4. Performance vs. frequency

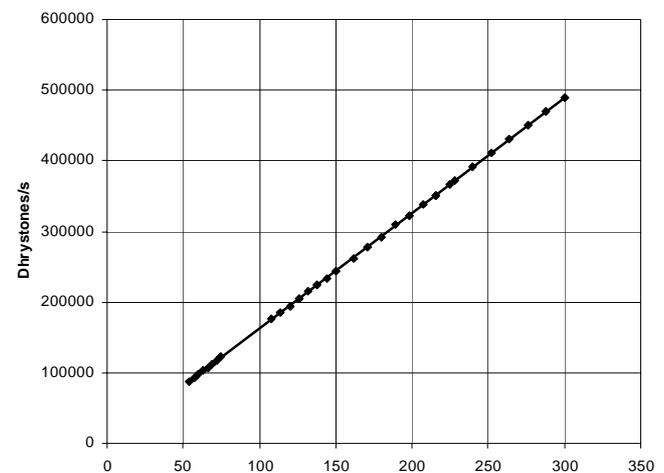
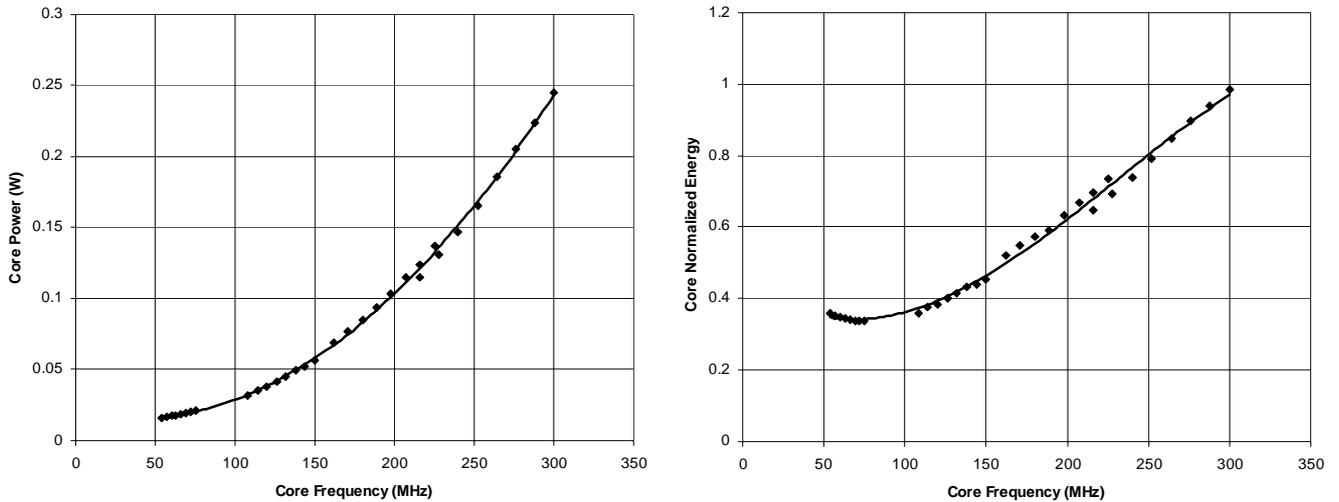


FIGURE 5. Power and energy consumption at different frequency and voltage levels



run four times longer. In general, bus-bound applications exhibit a flatter slope, meaning that due to the uneven scaling of bus frequencies, with reduced frequency the workloads’ run-time increase at a lower rate than that of processor-bound applications.

Figure 5 shows the ARM926EJ-S core’s power consumption and energy use (including on-chip cache and RAM structures) when running Dhrystone on a wide range of frequency and voltage levels. Energy consumption is normalized to the amount used at the statically characterized maximum operating point (240MHz at 1.2V). Our results show that a factor of 10 (90%) power and more than a factor of two (65%) energy saving is achievable by running the cores at their minimum levels (25% of peak frequency). However, there is very little pay back on running the core below the half-frequency point since voltage cannot be significantly reduced and consequently the energy consumption remains about the same. On the other hand, if heat management (thus average power consumption) is an issue, then more power savings can be achieved by further lowering the frequency—this behaviour is shown at the bottom of the power curve.

Our measurements match the theory: the power consumed during a workload is proportional to the frequency times the square of the voltage at which it is run. Since energy is the integral of power consumption, the longer execution time due to lowered operating frequency cancels out the frequency term and thus energy consumption is proportional to the square of the operating voltage.

The amount of energy saved for a given workload depends on the peak frequency and voltage levels of the core. Table 1 illustrates our results when running a work-

load at the minimum 25% (second column) and 50% (third column) for three maximum frequencies. Results in the second column show that there is more energy reduction if the peak frequency and voltage levels are higher. One reason for this is that our hardware does not function below 0.7V and this voltage level can already be achieved at 50% of maximum frequency for the 240MHz and 216MHz configurations. This implies that the 25% frequency levels of these configurations do not significantly reduce energy consumption any further.

However, even on cores with lower minimum voltage levels, the primary benefit of voltage scaling is towards the top end of the frequency range. This is a consequence of the scaling equations and the quadratic relationship between energy consumption and operating voltage [4]. The third column of Table 1 shows the energy savings for workloads running at half of maximum frequency. While the difference between the reported energy savings in each row is less than in the first case, the trend is clear: higher maximum frequency and voltage enables more energy savings when slower operating levels are used.

3.1 Operating margins

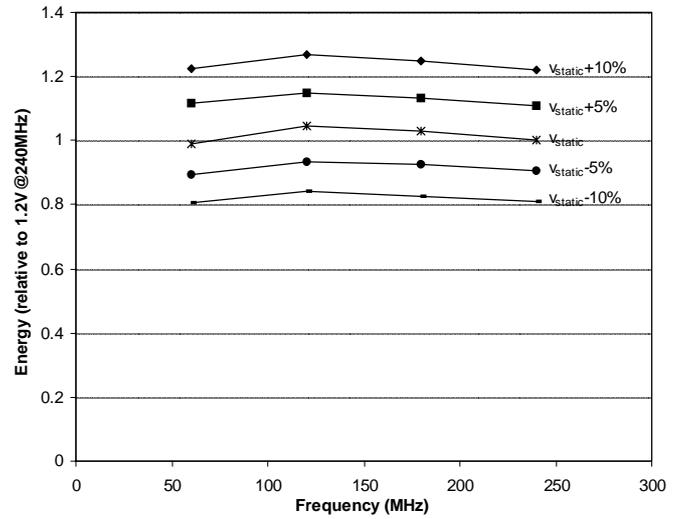
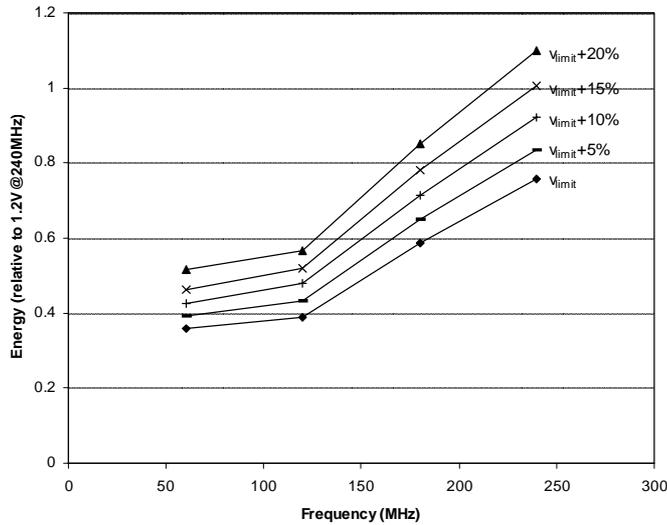
Data in the previous sections were collected for minimum voltages at room temperature. However, there is no guarantee that the same voltage levels would be sufficient to run the processor at the specified frequencies under different conditions (or that different chips would behave the same way). To deal with uncertainty and variations due to the ambient environment, silicon, IR-drop, etc. designers include operating margins in the voltages that are specified for each frequency level.

The first graph in Figure 6 shows the energy impact of the operating margins on the IEM926 processor running at the four different frequency levels that are achievable in a single configuration. For each frequency, the energy consumption of the Dhrystone workload is plotted using five different voltage levels: the limit voltage—below which the system fails to operate—and at 5%, 10%, 15%, and 20%

TABLE 1. Energy savings at different (f, v) points

Max. speed (MHz)	Workload speed (MHz)	Energy reduction	Workload speed (MHz)	Energy reduction
300	75	66%	150	54%
240	60	56%	120	48%
216	54	45%	104	45%

FIGURE 6. The energy impact of operating margins



above the minimum. The knee in the line at 120MHz shows the limited energy savings at the 60MHz level due to the hard limit on minimum voltage that is near the same level as at 120MHz. Typical tolerance levels on supply voltages are between 10%-15%, which translates into 20%-25% energy overhead when the processor is not running close to worst-case conditions. The energy consumption of the different configurations are normalized to the amount consumed at the statically characterized level (240MHz at 1.2V), which corresponds closely to the line with 15% voltage overhead.

The second graph in Figure 6 shows the energy consumption of the workload without voltage scaling. In these experiments, the operating voltage was kept at the statically characterized level and at levels 5% and 10% above and below for all frequency points. The results show that without voltage scaling the energy consumption for a workload is not reduced and in fact may increase at lower frequencies. We believe that this behaviour is due to on- and off-chip bus interactions and extra overhead incurred during some memory transactions. While the +5% and +10% voltage levels may be beyond the amounts that are incorporated into the processor's operating margins, such overshoots may be a function of the power regulator. Accurate power delivery is an important component of an energy efficient system as even a small increase over the necessary voltage level incurs significant energy overhead.

4. Conclusion and future work

Our results show that voltage scaling enables significant reduction of the energy consumption of the core implemented in a 130nm process. Our ongoing work quantifies the system-wide impact on energy consumption under real workloads, operating systems, and performance-setting pol-

icies. Our initial results indicate that when running at the peak level, the processor accounts for 75% of the energy used on the IEM926 SoC. Our data confirms that while designing with worst-case parameters may be necessary, actually running a chip with worst-case voltage levels wastes energy: in our case up to 25%. Our ongoing research explores on-chip structures [1] and microarchitectural techniques [2] for reducing operating margins.

5. Acknowledgements

The IEM926 design was done as a joint project between ARM, Synopsys, and TSMC. We thank Anwar Awad and Han Pin-Hung Chen of Synopsys for their implementation work.

References

- [1] S. Dhar, D. Maksimovic, and B. Kranzen. Closed-Loop Adaptive Voltage Scaling Controller For Standard-Cell ASICs. *Proceedings 2002 Int'l Symposium on Low Power Electronics and Design (ISLPED-2002)*, August 2002.
- [2] D. Ernst, N. S. Kim, S. Das, S. Pant, R. Rao, T. Pham, C. Ziesler, D. Blaauw, T. Austin, T. Mudge, and K. Flautner. Razor: A Low-Power Pipeline Based on Circuit-Level Timing Speculation. *Proceedings of the 36th Symposium on Microarchitecture (MICRO-36)*, San Diego, CA, 2003.
- [3] K. Flautner and T. Mudge. Vertigo: Automatic Performance-Setting for Linux. *Proceedings of the 5th Symposium on Operating Systems Design and Implementation (OSDI 2002)*, Boston, MA, 2002.
- [4] N. S. Kim, T. Austin, D. Blaauw, T. Mudge, K. Flautner, J. S. Hu, M. J. Irwin, M. Kandemir, and V. Narayanan. Leakage Current: Moore's Law Meets Static Power. *IEEE Computer*, December 2003.
- [5] T. Mudge. Power: A First Class Architectural Design Constraint. *IEEE Computer*, vol. 34, no. 4, April 2001.
- [6] <http://www.national.com/appinfo/power/powerwise.html>