

Introduction and Overview

Thucydides Xanthopoulos

Cavium Networks

Clock frequency is a major attribute of any microprocessor design. Early on, during product definition, it constitutes a major business or marketing decision and it is usually the result of a trade-off among customer needs, competitive landscape, and time-to-market. As soon as the frequency target is handed down the food chain to silicon implementation, it will affect all project design aspects from the day that the project is kicked off until it tapes out (and in most cases well beyond this point too). It is not surprising therefore that the job of generating, distributing, and analyzing the clocks in complex chips is considered to be an important and visible assignment. Clock design has traditionally been an area of innovation and has been in the spotlight in technical conferences and journals.

Why is clock frequency such an important microprocessor aspect? For a number of applications it is only loosely correlated with performance with other design aspects such as memory system, parallelism, and hardware acceleration being equally or even more effective. Nevertheless, it is a single number that is widely understood by both technical and nontechnical audiences and in certain situations has strong correlation with single-thread performance.

Clock frequency, although very important, is only one aspect of clock design. Other aspects include power dissipation, efficient clock signal distribution in large and complex chips, coping with variation and uncertainty, managing multiple clock domains in the context of highly integrated system-on-a-chip (SoC) designs, and multicore integration, providing good voltage/frequency scalability to support a wide product roadmap, tuning capabilities for yield enhancement and postsilicon optimization, and sophisticated active power management features.

The purpose of this book is to introduce a designer to important aspects of state-of-the-art clock design by exposing methodology steps and analytical modelling techniques, providing design examples and case studies and enumerating a long list of references for further study.

1.1 The Clock Design Problem

The plots of Figs. 1.1–1.3 provide good insight into the problem faced by clock designers today. Moore’s plot (Fig. 1.1) states that integration keeps on increasing at historical rates. Transistor density is spearheaded by large multicore server processor chips with large caches, integrated memory controllers and multiple high-speed *I/O* links. Although the datapoints in Fig. 1.1 only extend to 2008, in ISSCC 2009 a server processor [1] has been introduced that contains 2.3B transistors and exceeds all previously reported chips in transistor count. For the clock designer, increasing integration means more far reaching and complex clock distribution, more clock domains, and more testability and yield enhancement features.




Fig. 1.1. Microprocessor transistor number trend over time [Data compiled by the author from ISSCC proceedings 1973–2008 and publicly available vendor information. Trendline is “visually” fitted and signifies a rate of doubling every 2 years]

On the other hand, frequency and power seem to be levelling off (Figs. 1.2 and 1.3). Microprocessors in excess of 3–4GHz are rarely reported these days. Frequency stopped scaling due to excessive power dissipation penalties. Instead, the industry chose to keep on scaling performance through multicore integration. On the power side, the industry seems to converge to a fixed power envelope for both server (130W) and desktop (60–70W) processors. As far as clock design is concerned, this




Fig. 1.2. Microprocessor frequency trend over time [Data compiled by the author from ISSCC proceedings 1973–2008 and publicly available vendor information. Trendline is “visually” fitted and signifies a rate of doubling every 3 years]




Fig. 1.3. Microprocessor power trend over time [Data compiled by the author from ISSCC proceedings 1973–2008 and publicly available vendor information. Trendline is “visually” fitted and signifies a rate of doubling every 3.5 years]

translates to increased number of features and complexity to support fine-grain clock and power gating, variable voltages, and frequencies in addition to other sophisticated active power management related attributes.

In addition to the above evolving specifications, one must add increased variation in device and interconnect characteristics present in advanced process nodes and increased voltage and temperature variation due to higher integration and more complex interactions. We end up with a multidimensional design problem requiring substantial resources for each product generation.

1.2 Some Subjective Milestones in the History of Microprocessor Clocking

In the last 15–20 years there have been major innovations in microprocessor clocking that have resulted in large increases in frequency as well as substantial clock design methodology changes. This section lists some of those important milestones so that we can observe the industry trends and speculate on potential future directions.

1.2.1 Integrating the PLL

An integrated PLL for a microprocessor application was first reported in 1992 [2]. The motivation for an integrated PLL at the time was the desire to keep the processor and the external bus in phase so as to minimize timing constraints and reach the maximum possible system frequency given the synchronous nature of the system. Additional goals were to clock the processor at even higher frequencies ($2x$) than the bus in addition to running the VCO at twice the processor frequency with a factor of two postscaling for duty cycle fidelity.

Some of the original problems that the designers faced were the digital noise which is highly prevalent on a microprocessor die, in addition to low quality passive devices and overall sensitivity to process, voltage, and temperature variations. The supply noise problem was addressed by the adoption of a differential CML-based VCO with high supply noise rejection capability.

Integrating the PLL was a major step in general purpose microprocessor clocking because it paved the way for the increased frequencies and complex clocking schemes that were to follow.

1.2.2 Clock Distribution Moves to the Forefront: The Dawn of the GHz Race

The original DEC Alpha microprocessor [3] moved clock design to the forefront. It was operating at 200MHz, approximately a factor of 2 faster than other processors at the time (Fig. 1.2). The DEC Alpha design introduced the low-skew and high-power grid clock distribution coupled with detailed RC-based skew simulation and construction of skew contour maps that are taken into account during timing closure: The Alpha pipeline was based on level-sensitive latches. Race-through in this context is a major functional concern. The radial profile of clock skew from the center of the chip to the periphery was taken into account while floorplanning the pipelines in order to guarantee that the skew would improve the functional race margins.

The design also featured important contributions from a process and physical design perspective. The process technology featured a thick low resistance metal 3 layer used exclusively for power, clock, and a handful of critical signals. On-chip decoupling capacitors built with thin oxide devices were also used in close proximity to the clock drivers in order to address Ldi/dt concerns arising from driving a highly capacitive final clock node with a very fast edge rate. A 10:1 ratio of decoupling

capacitance to switching capacitance was maintained throughout the design. Many of these contributions are still in use today in current processor designs.

Arguably, the DEC Alpha designs started the GHz race among microprocessor vendors that culminated in the deep pipelines and multiGHz designs of present high performance chips.

1.2.3 Delay Lock Techniques

Simple first-order mechanisms to achieve phase lock [4] have been used extensively in processor designs in the last 10–15 years in order to simplify the clock distribution problem: A large distribution throughout a big die is broken into pieces tailored for each chip partition, and each partial distribution is phase locked using delay locked loops. There are multiple such examples discussed in Chap.2. Such a partitioning helps with design time, power dissipation, testability, and manufacturing yield. It is definitely possible to achieve the same goal with higher order systems (i.e. distributed PLLs) and this has been demonstrated in the literature [5, 6]. Yet, nothing beats the simplicity of a digital delay line controlled by a basic finite state machine. More details on DLL design will be presented in Chap.6.

1.2.4 Exploiting Inductance for Oscillation and Distribution

The notion of being able to return energy back to the clock generator has been rather intriguing and holds a lot of promise. Resonant clock drivers have been originally introduced as an off-chip solution in the context of powering adiabatic circuits [7, 8]. In the past several years, there has been renewed interest in this technique due to the fact that digital frequencies have become consistent with resonant frequencies of fully integrated passive devices. The strong motivation is the potential of saving substantial clock power by using LC resonance for clock pulse generation. LC techniques have been recently augmented with transmission-line-based techniques (traveling waves [9], standing waves [10] and salphasic [11]) that address both clock generation and low (or highly predictable) skew distribution. Commercial applications of these techniques are already prevalent. Chapter 4 explores this subject in more detail.

1.2.5 Variable Frequency (and Voltage)

A server processor design [12] introduced the idea of a constant power envelope and variable voltage/frequency. This active power management scheme uses an integrated ammeter that monitors incoming current, a clocking scheme that can generate variable frequencies with fast adjustment time and an on-chip micro-controller that monitors power/ temperature and controls frequency and core supply voltage through an external regulator. This technique is discussed in more detail in Chap.2.

Variable frequency clocking methods are rapidly becoming mainstream [13, 14] as part of sophisticated power management methodologies designed to control thermal design power in large multicore chips. All digital methods ensure repeatability in a production environment.

1.2.6 Frequency Increase (or Supply Lowering) Through Resiliency

Commercial designs have substantial frequency margins to address issues such as lack of total coverage in production tests, unanticipated corner cases and noise patterns, noncompliant system specifications and device aging. This margin can be substantially reduced or even eliminated if the underlying hardware has error detection and correction capabilities. The margin can then be used as performance benefit by increasing frequency, power benefit by dropping voltage below the specified V_{MIN} , or even yield enhancement by populating existing frequency bins with parts that under different scenarios would not make it.

The Razor technique [15] addresses the frequency/voltage margin issue by instituting the capability of timing error detection and correction in the processor pipeline. Performance can be maximized (or power minimized) by increasing the frequency (or lowering the supply) up to the point where the overhead of error correction will start exceeding the performance or power benefit. The Razor technique is addressed in more detail in Chaps. 3 and 7.

It is not easy to predict the future but given the current industry trends one can conclude that the clocking system will be part of an increasingly sophisticated active power management scheme: Highly sophisticated firmware threads implementing complex control algorithms will be running in parallel with the application. They will be receiving input such as on-chip and system temperature, current measurements, error rates, moving averages of architectural events, and cues from the application and they will control supply voltage, clock frequencies, and higher level architectural events such as clock and power gating, instruction issue rate, and pipeline stalls. To some extent, this is already happening.

1.3 Overview of this Book

Chapter 2 introduces the fundamental setup and hold constraints, and defines basic clock attributes such as skew, jitter, latency, and duty cycle distortion. It introduces basic clock distribution methods such as balanced tree, central spines, and grids and examines them from a performance and power perspective. Numerous case studies from commercial microprocessors are presented and a number of advanced topics such as global and local skew compensation, on-die attribute measurements, various techniques for locating critical paths and synchronization methods are discussed.

Chapter 3 constitutes a detailed discussion of clocked elements (level-sensitive latches and flip-flops) from the viewpoint of latency, hold time requirements, power dissipation, and testability. The focus is primarily on state-of-the-art designs with advanced topics such as process variation and reliability addressed in detail.

As mentioned in Sect. 1.2.4, exploiting inductance for clock generation and distribution makes perfect sense: Inductance can produce oscillations with lower energy since an LC-based system inherently recycles energy between the capacitive and inductive elements. Moreover, oscillator phase noise is less because all active-device-related noise sources do not exist. Chapter 4 presents detailed background

information on integrated inductors and transmission lines. Furthermore, it discusses examples of LC-based oscillators and transmission-line-based clock generation and distribution schemes.

Jitter analysis is very important in clock system design. If not properly managed, jitter can be the limiting factor in both core and *I/O* clocking. Chapter 5.1 defines all jitter types relevant to clock design and establishes their relationship to phase noise using Parceval's theorem. Furthermore, it enumerates all noise sources inside a clock generator and clearly shows with numerical examples how a PLL transfers jitter from input to output. Based on this analysis, it establishes the importance of reference clock phase noise and jitter regarding the quality of the multiplied output clock. The domain seamlessly moves from frequency to time using mathematical "filter" functions to transform phase jitter to period jitter which is more relevant for core clock generation. Since jitter has a random component, which is theoretically unbounded, the chapter establishes an MTBF-based statistical analysis for estimating the effect of jitter on critical paths. A serial link discussion is also presented, which shows how reference clock jitter can be removed from the total link budget under certain conditions.

Chapter 6 is an attempt at textbook-like coverage of digital delay locked loops that are used extensively in clocking systems. The chapter is design-oriented and contains detailed discussions and analyses on all DLL components and presents a number of relevant applications. It also contains a detailed analysis of metastability in the context of phase detection and a simple analytical model for supply-induced jitter on long delay lines and/or clock buffers.

Advanced process nodes exhibit large variation and uncertainty in device and interconnect parameters. Chapter 7 presents methods of addressing this issue on the design front, manufacturing, and postsilicon tuning front and also by using resilient methods involving hardware timing error detection and correction.

Finally, Chap. 8 addresses process, voltage, and temperature variation issues from a physical design perspective. Clock skew components in the context of setup and hold constraints are redefined with a statistical approach and all variation sources are taken into account. Sources of transistor and interconnect variation are enumerated, quantified, and explained. Methods of accounting for voltage and temperature variations are discussed. In the end, important physical design guidelines are presented to minimize uncertainty and variation in clock-related circuits.

References

- [1] S. Rusu, S. Tam, H. Muljono, J. Stinson, D. Ayers, J. Chang, R. Varada, M. Ratta, and S. Kottapalli, "A 45nm 8-core enterprise Xeon[®] processor," in *Digest of Technical Papers IEEE International Solid-State Circuits Conference (ISSCC 2009)*, 2009, pp. 56–57.
- [2] I. A. Young, J. K. Greason, and K. L. Wong, "A PLL clock generator with 5 to 10 MHz of lock range for microprocessors," *IEEE Journal of Solid-State Circuits*, vol. 27, no. 11, pp. 1599–1607, Nov. 1992.

- [3] D. Dobberpuhl, R. Witek, R. Allmon, R. Anglin, D. Bertucci, S. Britton, L. Chao, R. Conrad, D. Dever, B. Gieseke, S. Hassoun, G. Hoepfner, K. Kuchler, M. Ladd, B. Leary, L. Madden, E. McLellan, D. Meyer, J. Montanaro, D. Priore, V. Rajagopalan, S. Samudrala, and S. Santhanam, "A 200-MHz 64-b dual-issue CMOS microprocessor," *IEEE Journal of Solid-State Circuits*, vol. 27, no. 11, pp. 1555–1567, Nov. 1992.
- [4] M. Johnson and E. Hudson, "A variable delay line PLL for CPU-coprocessor synchronization," *IEEE Journal of Solid-State Circuits*, vol. 23, no. 5, pp. 1218–1223, 1988.
- [5] G. A. Pratt and J. Nguyen, "Distributed synchronous clocking," in *Proceedings of Sixteenth Conference on Advanced Research in VLSI*, 27–29 March 1995, pp. 316–330.
- [6] V. Gutnik and A. Chandrakasan, "Active GHz clock network using distributed PLLs," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 11, pp. 1553–1560, Nov. 2000.
- [7] C. L. Seitz, A. H. Frey, S. Mattison, S. D. Rabin, D. A. Speck, and J. L. A. van de Snepscheut, "Hot-clock NMOS," in *Proceedings of Chapel Hill Conference VLSI*, 1985, pp. 1–17.
- [8] R. Feynman, T. Hey, and R. Allen, *Feynman Lectures on Computation*. Westview Press, Boulder, CO, 2000.
- [9] J. Wood, S. Lipa, P. Franzon, and M. Steer, "Multi-gigahertz low-power low-skew rotary clock scheme," in *Digest of Technical Papers IEEE International Solid-State Circuits Conference (ISSCC 2001)*, 2001, pp. 400–401, 470.
- [10] F. O'Mahony, C. Yue, M. Horowitz, and S. Wong, "A 10-GHz global clock distribution using coupled standing-wave oscillators," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 11, pp. 1813–1820, Nov. 2003.
- [11] V. L. Chi, "Salphasic distribution of clock signals for synchronous systems," *IEEE Transactions on Computers*, vol. 43, no. 5, pp. 597–602, May 1994.
- [12] S. Naffziger, B. Stackhouse, and T. Grutkowski, "The implementation of a 2-core multi-threaded Itanium family processor," in *Digest of Technical Papers IEEE International Solid-State Circuits Conference (ISSCC 2005)*, 2005, pp. 182–183, 592.
- [13] R. Kumar and G. Hinton, "A family of 45nm IA processors," in *Digest of Technical Papers IEEE International Solid-State Circuits Conference (ISSCC 2009)*, 2009, pp. 58–59.
- [14] A. Allen, J. Desai, F. Verdico, F. Anderson, D. Mulvihill, and D. Krueger, "Dynamic frequency-switching clock system on a quad-core Itanium[®] processor," in *Digest of Technical Papers IEEE International Solid-State Circuits Conference (ISSCC 2009)*, 2009, pp. 62–63.
- [15] D. Ernst, N. S. Kim, S. Das, S. Pant, R. Rao, T. Pham, C. Ziesler, D. Blaauw, T. Austin, K. Flautner, and T. Mudge, "Razor: a low-power pipeline based on circuit-level timing speculation," in *Proceedings of 36th Annual IEEE/ACM International Symposium on MICRO-36 Microarchitecture*, 2003, pp. 7–18.