Experiences and Challenges in System Design

Jan M. Rabaey
EECS Department, University of California at Berkeley
Berkeley, CA. USA

Abstract

This paper examines the desirability and viability of state-of-the-art system design at the university. To put the problem in perspective, a number of actual designs exercises, executed at UC Berkeley over the last decade, will be examined. A number of potential models for success are proposed and analyzed. While this paper might provide some insight, it is surely hoped that it might serve as the basis for an ongoing discussion that might lead to an (inter)national infrastructure for system-design in the long term.

1. Introduction

University research in design methodology for integrated circuits lays at the core of the success of today's semiconductor industry. Jump-started by the VLSI System approach, preached by Mead and Conway, universities paved the way for the development of a structured and mostly automated approach to IC design. Concepts such as higher levels of abstraction in the specification process, silicon compilation, and logic and behavioral synthesis emerged from the university community and have a lasting impact on how ICs are being designed at present. At the core of these developments was a vibrant group of active and enthusiastic designers, enabled by cheap access to technology (MOSIS), simple standards (CIF), and a shared tool-set (MAGIC, SPICE, IRSIM, LAGER).

Yet, we observe that today this design activity has all but vanished (with the exception of some of the larger schools), and funding for design-related projects and the corresponding methodologies has dwindled. At the same time the design industry is ringing the alarm bell concerning an impending productivity crisis, and is sending a strong message that innovation at the system-design level is a necessity, lest it will become increasingly harder to exploit the tremendous opportunities offered by the continuous miniaturization of the semiconductor devices.

In this paper, we will evaluate the reasons for this crisis, make the argument for a continued presence of universities in the system-design forum, present a number of past and current design projects at Berkeley, and use this information to evaluate the potential of system-design at universities in the future.

2. A System-Design Crisis

The complexity of integrated circuits has followed a fairly predictable trend over the last decade, and has tracked

Die Area: 2.5x2.5 cm
Voltage: 0.6 V
Technology: 0.07 µm

	Density	Access Time
	(Gbits/cm2)	(ns)
DRAM	8.5	10
DRAM (Logic)	2.5	10
SRAM (Cache)	0.3	1.5

	Density	Max. Ave. Power	Clock Rate
	(Mgates/cm2)	(W/cm2)	(GHz)
Custom	25	54	3
Std. Cell	10	27	1.5
Gate Array	5	18	1
Single-Mask GA	2.5	12.5	0.7
FPGA	0.4	4.5	0.25

Figure 1 . Projected Integration Capabilities of CMOS technology by the year 2010.

the exponential growth curve, predicted by Moore's law. While improved design methodologies and higher levels of abstraction have succeeded in effectively managing this increased complexity so far, it seems that we are at the verge of (or for some industries have passed) a break-point that may have a profound impact on the way design is going to be conducted. Typically, the design of a system meant the assembly of a number of modules (off-the-shelf or application-specific) on a substrate. Each of these components fulfilled a give function in the system an could be designed relatively independently. The increasing integration density now makes it possible to integrate a multiplicity of these functions onto a single die (the *system-on-a-chip*), requiring a much closer interaction between the components, and introducing a whole new set of constraints.

To put the potential of tomorrow's system integration into perspective, let us briefly consider the projected capabilities of CMOS technology by the year 2010, as shown in Figure 1 Such a technology makes it possible to put widely diverse and complex functions on the same die, as is illustrated in the integrated sensor-system of Figure 2. This circuit combines an advanced imager, an FPGA for adaptive signal formatting, an SIMD for advanced two-dimensional signal-processing, and a high-performance CPU (with associated memory) for recognition and compression functions. Designing this part within a reasonable amount of time and with a limited set of designers while still ensuring that the final part operates correctly, requires the combination and coordination of a wide range of skills. Observe that a large range of the components on the die are programmable, making software an essential component of the design process.

In short, the emergence of systems-on-a-chip forces a

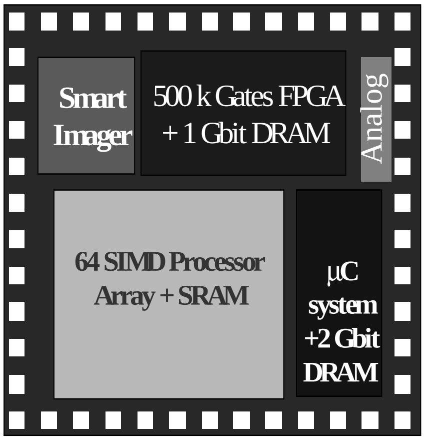

Figure 2 . Futuristic Integrated Sensor System

number of different worlds together (system specification and partitioning, chip design, board and software design), To make this effectively happen will require the introduction of higher levels of abstraction and a rethinking of the overall design flow. Techniques such as design reuse and component-based are already making some in-roads into the industry. Yet, to be successful, some fundamental problems regarding specification, abstraction, partitioning, modeling, synthesis, and verification have to be addressed. It is our belief that universities can and will have an important input in this process. A necessary condition for this to happen is that universities remain active in the system-design arena, or risk that their research in this area becomes altogerther irrelevant.

3. Universities and System Design

Hence, the worry about the waning system-design efforts at universities, at the core of which is clearly the complexity issue. Any real system brings together a wide variety of components, each with its own set of specifications, models, and tools. Mastering the complexity of such an environment while trying at the same time to get a degree in a limited number of years is non-trivial. Unfortunately, the design-automation industry has ensured that hurdle is placed as high as possible by making the learning process for any of their tools a challenge on its own. Just the task of integration, maintenance, and upgrading of a complete design environment and its module libraries is beyond the scope of virtually any school.

Another important aspect hindering system design is that the university culture is not sympathetic to the concept of large collaborative projects that are at the core of modern day system design. The award system for both students and faculty tends to focus on individual level.

4. System Design at Berkeley

To explore the challenges, pitfalls, and opportunities of complex system design at the university, a number of actual design efforts, executed at Berkeley over the last decade, will be presented. In the early to mid eighties, most of our efforts had gone towards the development of an integrated environment for the design of dedicated integrated circuit. The resulting LAGER environment made it possible for numerous students to create fairly complex application-specific integrated circuits in a very short time-span. Even more important, the overall training period for an unexperienced student to get familiar with the tool flow was short as well.

4.1 Connected Speech-Recognition System

In the late eighties, we took on the challenge to extend this methodology towards the system level. In line with our general philosophy, a driver example was selected to motivate, guide, and critique the methodology developments. In this case, a 60,000 word real-time connected speech recognition was chosen as the driver. A dedicated implementation of one specific algorithm, being a Hidden Markov based model, was selected.

The resulting design combined 3 triple-size VME boards (word processing, grammar processing (Figure 3) and memory board), multiple ASICS and a wide range of off-the-shelf components.

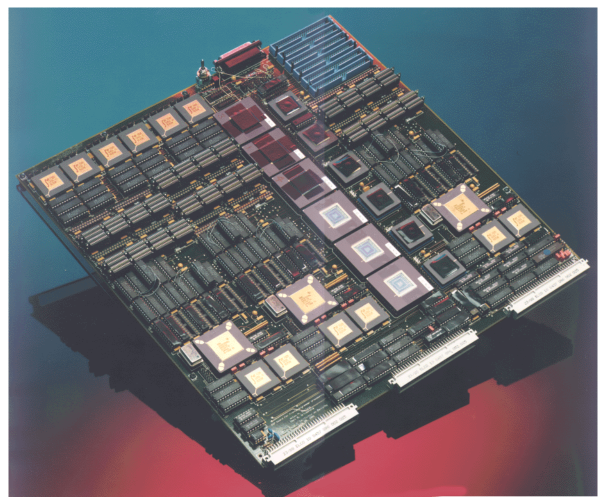

Figure 3 . Grammar-processing board of 60,000 word real-time speech recognition system.

The aggressive goals of the project exposed one important aspect of the system-design problem: *the evolving specification*. While the traditional CAD scenario's present a top-down design methodology, one quickly realizes that there is no such thing as a fixed specification. A large range of algorithmic parameters are negotiable and can be modified in response to architecture or hardware constraints. This negotiation process proceeds between different camps, however, and every iteration can take a substantial amount of time as it may require substantial re-verification. This pleads for the development of an algorithm-architecture-hardware co-design environment rather than a traditional waterfall design model.

While the attempts of carrying the chip-level structured design methodology to the board-level resulted in some significant successes (such as a structured approach towards physical board designs), some of the approaches that worked for chips were not as successful for boards. An example of such is the development of module generators for sub-systems. The design-effort also exposed the problem of verification at the system-level, due to the diversity and the lack of models for the individual components, a problem that is surely hampering the systems-on-a-chip integration at present.

4.2 Reconfigurable Image Processing System

While the speech system presented an interesting set of challenges, it clearly accentuated one aspect of system design that has become more and more apparent since: to defer the cost of complex systems, it is important to keep them as flexible and adaptive as possible. The speech recognizer was almost obsolete at the time of completion due to evolutions in algorithmic research as well as technological progress. The role of software in systems is undeniable and becoming more and more important every day. The Spartan system is an early example of a system-design approach that stresses reconfiguration and adaptivity. This methodology has gained substantial following today, as is exemplified in the DARPA Adaptive Computing Systems (ACS) program.

Spartan (Figure 4) combines 8 reconfigurable PADDI-2 processor chips, and connects them to a SUN host and a video input-output sub-system. A Xilinx FPGA is used for the routing of data between the host, processor array, and video I/O system. A major part of the design effort for this system was related to the development, test, and verification of the software environment for the reconfigurable components and host interface.

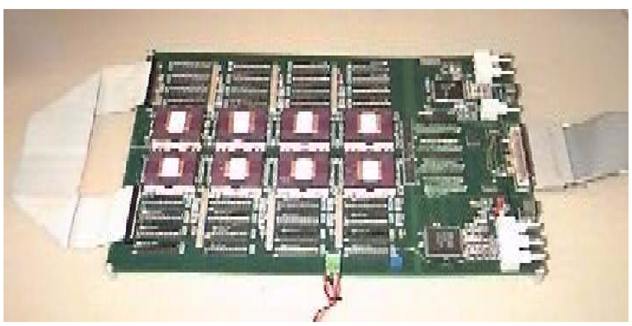

Figure 4 . Spartan reconfigurable processor board for image and video applications.

4.3 The Infopad System

The Infopad project superseded each of the above examples in scope, diversity, and size. While the original scope of the project was to evaluate low-power solutions for portable terminals and the potential of wireless access to high-bandwidth wired backbones, it quickly grew into a project that involved RF-design, wired and wireless networking, multimedia data management and transfer, user interface, and packaging (Figure 5). At its height, it involved almost 50 graduate students, 8 faculty and a number of staff, making it easily the largest integrated project ever conducted in our department.

While infopad was ultimately successful and judged as being a great model of system level-education by academia, industry, and students, the price of running such as major project at a university is considerable and out-of-the-scope

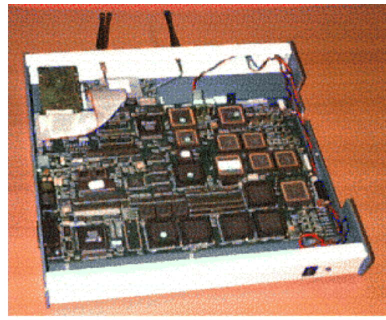

Figure 5 . Infopad basestation and terminal.

of even the largest universities. Yet the lessons learned and the technologies developed were so valuable that there is no doubt in the mind of the author that such projects indispensable if one wants to get to the core of the system-design problem. Finding a means to alleviate the cost and management overhead of large projects is the most important challenge to system-design at the universities (as well as companies...).

5. Plausible Scenarios

A number of scenario's are currently being explored at a variety of places to address exactly that problem.

A first example of such a scenario is the Berkeley Wireless Research Center, that is currently being set up in Berkeley. The goal of the center is the realization of the next-generation fully-integrated CMOS radio's. It brings together students and industrial visitors, and provides the critical mass and support needed for the undertaking of projects of this type.

Another model is the approach proposed by the MARCO program (under the SIA focus center header). This program forces a number of universities to work pro-actively together on the system-design problem. Bringing these schools together could ultimately lead to the emergence of a common system design platform (including tools, interfaces, and libraries), much in the style of the early VLSI days. Both these models will be discussed in more detail at the symposium.

6. Conclusions

Universities have been at the forefront of the development of novel design methodologies for integrated circuits. To address the impending system-design crisis will require a similar effort. A necessary condition to gain critical mass is that a substantial number of our universities get exposed to active system design. This is only possible through tightened industry-university and inter-university cooperation.