Next Generation Embedded System Design: Issues, Challenges, and Solutions

RAVI SHANKAR

Florida Atlantic University, Boca Raton, FL

1. INTRODUCTION

The SoC (System on a Chip) design cost during the past decade has stayed low and flat. The cost, however, is projected to increase exponentially, by 25 fold over the next ten years, from about $ 30 M today to $ 800 M by 2014 [ITRS, 2001]. The flat curve of the past ten years itself was the beneficiary of breakthrough EDA (engineering design automation) innovations, without which the exponential increase would have started a decade ago. Much of the consumer electronic revolution would have been missed. ITRS (International Technology Roadmap for Semiconductors) expects more innovations to come, but the current challenges are cross-disciplinary, not multi-disciplinary, requiring team synergy. ITRS projects the need for ESL (electronic system level) tools to address this. Dataquest has seconded this by predicting ESL as the primary growth area for EDA companies in the next few years. However, experiences of the EDA industry, as chronicled in the weekly issues of the Electronic Engineering Times, do not support this contention. Though there have been a few visionary attempts by the EDA industry, they may have underestimated the complexities of the challenge. Failures have a way of rarefying the field. However, not everything is lost. SystemC and UML seem to have arrived, in their new versions, albeit with a little baggage, but much potential. European initiatives may provide the blueprint for initiatives by US Federal agencies to lead the system design revolution – let us hope, quite similar to the MOSIS initiative that launched a radical change in the VLSI thought process.

2. ISSUES

There are three major issues to contend with:

1. It is a complex product. The future SoC may be a NoC (network on chip) [Jantsch and Tenhunen, 2003]. NoC envisions systems with hundreds of processors with local and shared memory, communicating with each other via asynchronous communication links, similar to the switches and routers used in the Internet. We need mechanisms to subdivide the application space among these processors and perform what-if scenarios to satisfy both performance and quality-of-service measures [Jantsch, 2004]. Subsequent to that, we need to abstract and model the process interaction at high levels to study their concurrency, communication, and synchronization issues [Magee and Kramer, 2000]. And, finally, in the same environment, we need to design both software and hardware, using both top-down and bottom-up methods, with verification well integrated into the flow [Muller et al., 2003].

2. Product development cost must be reduced. If not, the consumer demand will not be sufficient enough to provide a healthy return on investment to the EDA and the Systems (or high-tech) companies to justify the investment in further innovation. Historically, system development tools of the past have been focused on software since the hardware platforms were standardized and stable. EDA tool companies may have shied away from this domain because of the low cost, high volume model, quite unlike their modus operandi of high cost and low volume. On the other hand, today's embedded system development (ESD) companies may not have the expertise to deal with traditional hardware concepts of concurrency, communication, synchronization, hierarchy, and mixed mode design. System design companies may not appreciate the need for such wizardry by their tool vendors, and thus, may be unwilling to pay the price.

3. The integration issues are cross-disciplinary, that is, many changes made in one domain or discipline may effect the performance in another field. Analog and RF designers are proficient in understanding and accounting for loading and interference effects, but this is a new concept for software and digital hardware developers. Stakes are also high since a mismatch generally leads to design iterations or a compromise on the product features. The cross-disciplinary issue is acute in the universities as well because of the academic silos.

3. STAKEHOLDER CHALLENGES

The main players are the EDA Companies, the Systems Companies, the Academia, and the Funding Agencies.

1. EDA Companies: These companies are still focused on back-end tools, because of the large revenue from such tools. System level tools, as marketed by the ESD companies today, are inexpensive and may not provide a path of significant revenue growth for the EDA companies. Further, certain visionary ESL tools developed by EDA companies failed because of a few shortcomings: steep learning curve; no component libraries; no back-end integration, and no good reference designs and application notes. The past few years also have seen raging debates on the relative merits of SystemC and SystemVerilog, the two languages put forward by the EDA industry for standardization. Development of ESL tools will require the joint effort of EDA and ESD companies, the current progress towards which is fairly slow.

2. The Systems Companies: The hardware-centric model has changed to a software-centric model in such systems companies. Software development today occurs with ESD tools, which attain their fast simulation speed by abstracting hardware and ignoring technological and architectural variations. Concurrency issues may get ignored by software developers. Consequently, design iterations, integration and verification may continue to be major productivity issues. Because of the pressure of quarterly deadlines, the engineers may be unwilling to invest time and effort to explore an immature technology. Further, systems companies may continue to view the technology issues as alien to their domain, to be taken care of by their IP (intellectual property) and ASIC

(application specific integrated circuit) providers. However, the DSM (deep submicron) design is more analog in its behavior and cannot be ignored at the system level.

3. Universities: The universities had it good for the past decade, with increasing enrolment in the computer science and engineering programs. Graduates got jobs soon after graduation. Thus, there may not have been any motivation to evaluate the future challenges and plan changes in the curriculum. Further, the US university model of individualism makes cross-disciplinary team effort a challenge among research peers of equal stature. The federal funding with its emphasis on research excellence in a single domain creates silos and discourages faculty from venturing out to address general systems issues. The teaching oriented faculty may not have the wherewithal to update the courses to include the latest topics. The practice oriented faculty, typically adjunct faculty and affiliated with a local company, may be out of the loop with regard to curriculum development. Further, inter-disciplinary communication is hampered by incompatible vocabulary, tools and methodologies. These may lead to communication breakdowns and misunderstandings.

4. Funding Agencies: SRC (Semiconductor Research Corporation) and NSF (National Science Foundation) have focused on mutually exclusive goals of technology and theory, and seem to have missed the all encompassing ‘practice’ world of embedded system design [Gupta, 2002]. However, there are many major cross-disciplinary intellectual challenges here and they should not be considered mere 'practice' issues. The EDA and Systems companies may have put their relevant plans on hold because of the economic downturn. Further, many systems companies may have sought lateral opportunities to extend the reach of their current product portfolio, postponing the reckoning to a later quarter.

4. SOLUTIONS:

First, here are some desirable incentive-based outcomes: The systems companies and the funding agencies must take the initiative to lead. The incentive is that otherwise the high end consumer products will cost disproportionately more; The universities must develop an integrated curriculum [Shankar and Suryaprasad, 2004] that require joint teaching teams representing both computer science and computer engineering perspectives. That should help the universities enhance their enrolment by being responsive to the industry needs; Students should be taught to take responsibility for their future – they need to demand changes in the curriculum. The new system level jobs are unlikely to be candidates for job export since they require cross-disciplinary team effort and cannot be automated significantly as yet; finally, there is a need for generalist faculty members to coordinate the efforts of specialist faculty members in various sub-fields. The orchestras have their elite conductors, but the academia does not have a way to reward such coordination, since no publications can ensue from such efforts. The GP (general practitioner) is making a come back in the medical field, but will never match the specialist in terms of prestige and financial return. The generalist faculty member may have to take on the job for the love of it. A department chairperson may be ideally suited for such a leadership role.

Second, here are some developments/ opportunities that hold promise:

Standard Components: The component technology (software and hardware IP) has promise [Wolf, 2001], despite the fact that the embedded whole is significantly more than the sum of its parts [Vahid and Givargis, 2002, and Magee and Kramer, 2000]. SoC and NoC have evolved from similar notions [Jantsch and Tenhunen, 2003]. Efforts at evolving standard components must continue.

Standard Languages: Both SystemC and UML working groups have released version 2 that respectively address the system level and real-time issues. Though both have a few good books [Bhasker, 2002, Grotker et al., 2002, and Douglass, 2004], good reference designs and libraries are still not available. SystemC seems to be suffering from a lack of interest and a core group of stakeholders. SystemC also has disadvantages of low simulation speed and duplication of effort. These issues can be researched and overcome, but the faculty members must first gain awareness and support of the funding agencies. UML's use for real-time applications and executable code generation [Douglass, 2004], overshadowed by UML's business applications, nevertheless is promising.

International Cooperation: Much has been written about the vertical model of integration in the Route 125 area companies of Boston and the horizontal model of integration in the Silicon Valley companies of California [Saxenian, 1996]. Saxenian compares them to states of competition and cooperation, respectively, and points out the superiority of the latter. In a broader context, Rifkin sees the European Dream, based on cooperation and not competition, as a better long term strategy than the American Dream based on individual effort [Rifkin, 2004]. This may be especially true in the embedded systems world – many excellent team-inspired books have evolved in Europe [See Muller et al., 2003, and Jantsch and Tenhunen, 2003, for example]. The European model of cooperation among select universities with focus on different subspecialties seems to be working. Asia has always been successful with attention to detail. Thus, Europe's ability for abstraction and collaboration, US's apparent bent for commercializing innovation, and Asia's penchant for commoditization, may be harnessed in a synergistic manner.

A MOSIS model: We need a MOSIS-Like institution to democratize and provide easy access to entrepreneurial faculty, whether they are at small or large, well established universities. We hope DARPA and NSF will take the lead, as they have done in the past with regard to VLSI (very large scale integration) and MEMS (microelectromechanical systems). This proposed EDIS (embedded development information system) would provide access to a common set of resources: libraries, designs, text books, simulation engines and accelerators, methodologies, and EDA tools that are community-based and developed. Faculty and students alike will be able to use the same for an easy transition to the field. Further, we also recommend that NSF adopt the model of the federal department of education in funding not one or two top-ranked proposals, but the top eight to ten proposals. This will ensure friendly competition among several cooperating teams, which is perhaps the best of both the worlds. These proposals should involve several universities (to ensure that the influence on curriculum is diffuse, but perhaps as a recognition of the existence of entrepreneurial faculty members everywhere) and industry mentors who commit to intellectual participation (not necessarily matching funds). The selection criteria should substitute the publication record with real-life successes (such as product development, innovation, and course development) to identify faculty members who are practice/applied research oriented. Funding in engineering should respect and honor such engineering 'practice' as a legitimate alternative to intellectual pursuits, so we can all contribute to a successful system design revolution.

5. SUMMARY

These are exciting and taxing times in the field of real-time embedded systems. With strategic federal funding and easy access to a common information system, cooperation and competition among universities can be enhanced. This can be exploited to lay the foundation for another wave of innovation, commercialization, and job growth in the computer industry.

ACKNOWLEDGEMENTS

I am grateful for the interaction I have had with senior managers and engineers at several EDA and high tech companies. I am also grateful to Dr. Michael VanHilst for discussing cultural differences between Europe and the US, as well as for help in evolving some of these solutions.

REFERENCES

BHASKER, J., 2002, *A SystemC Primer*, Star Galaxy Publishing, Allentown, PA, 2002

DOUGLASS, 2004, W., *Real-Time UML, 3rd Edition,* Addison-Wesley, 2004

GROTKER, T., LIAO, S., MARTIN, G., AND SWAN, S., 2002, *System Design with SystemC*, Kluwer Academic Publishers, Boston, MA, 2002

GUPTA, R., 2002, The neglected community, Editorial, *IEEE Design & Test of Computers*, Special Issue on System-level design methods, November- December 2002, pp. 2.

ITRS, 2001, *International Technology Roadmap for Semiconductors 2001 Edition,* Semiconductor Industry Association, ITRS, Austin, TX, 2001.

JANTSCH, A., AND TENHUNEN, H., 2003, Eds., *Networks on Chip,* Kluwer Academic Publishers, Boston, MA, 2003

JANTSCH, A., 2004, *Modeling Embedded Systems AND SOC'S: Concurrency and Time in Models of Computation,* Morgan Kaufmann, San Francisco, CA, 2004

MAGEE, J., and KRAMER, J., Concurrency: State Models and Java Programs, John Wiley and sons, Chichester, England, 2000

MULLER, W., ROSENSTIEL, W., AND RUF, J., 2003, Eds., *SystemC: Methodologies and Applications*, Kluwer Academic Publishers, Boston, MA, 2003

RIFKIN, J., 2004, *The European Dream: How Europe's Vision of The Future is Quietly Eclipsing The American Dream*, Jeremy P. Tarcher, Penguin Group, New York, NY, 2004.

SAXENIAN, A., 1996, *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Harvard University Press, Cambridge, MA, 1996

SHANKAR, R., and JAYADEVAPPA, S., 2004, A New SystemC-based Foundation for the CE Curriculum, *European Workshop on Microelectronics Education (EWME)*, Lausanne, Switzerland, April 2004, pp. 33-37.

VAHID, F., AND GIVARGIS, T., 2002, *Embedded System Design: A Unified Hardware/Software Introduction*, John Wiley & Sons, Inc., Danvers, MA, 2002

WOLF, W., 2001, *Computers as Components: Principles of Embedded Computing System Design*, Morgan Kaufmann, San Francisco, CA, 2001