Panel: The Heritage of Mead & Conway

What Has Remained the Same, What Was Missed, What Has Changed, What Lies Ahead

Organizer: Marco Casale-Rossi, Synopsys, USA
Panel Moderator: Prof. Alberto Sangiovanni-Vincentelli, UCB, USA
Panelists: Prof. Luca Carloni, Columbia University, USA; Dr. Bernard Courtois, CMP, France;
Prof. Hugo de Man, K.U. Leuven & IMEC, Belgium;
Dr. Antun Domic, Synopsys, USA; Prof. Jan M. Rabaey, UCB, USA

Abstract

Thirty-two years ago, Electronics Magazine honored Carver Mead and Lynn Conway with its Achievement Award for their contributions to VLSI chip design. The 'Mead & Conway methods' were being taught at 100+ universities all over the world, and "not only have helped spawn a common design culture so necessary in the VLSI era, but have greatly increased interaction between university and industry so as to stimulate research by both." Concepts such as simplified design methods, new, electronic representations of digital design data, scalable design rules, 'clean' formalized digital interfaces between design and manufacturing, and widely accessible silicon foundries suddenly enabled many thousands of chip designers to create many tens of thousands of chip designs. Today, as Moore's Law – a term coined by Carver Mead – has brought as from 10 microns to 10 nanometers, what is the heritage of Mead & Conway?

UCB Professor Alberto Sangiovanni-Vincentelli will moderate an industry and research panel, to discuss what has remained the same, what was missed, what has changed, and what lies ahead.

Panel Moderator's Introduction

Prof. Alberto Sangiovanni-Vincentelli, UCB, USA

It is a great honor for me to moderate this panel where long-time friends (and certainly Richard Newton is much missed in this group) debate about the legacy of a revolutionary set of ideas and one of the originators gives historical remarks. My own perspective on the MC methodology is that indeed the meta principles are the pillars of its success over the years, namely simplification, interdisciplinarity, collaboration and orthogonalization of concerns. In 1979 at Berkeley, Richard Newton and I were young Assistant Professors eager to change the way design was done with massive use of software tools based on optimization and verification algorithms. The MC methodology offered a framework where tools played an essential role in supporting designers. While Richard and Carlo Sequin were teaching the first MC class in Berkeley, Richard and I became convinced that while the MC approach was a great way to popularize the design of integrated circuits among neighboring disciplines such as computer architecture, it lacked attention towards optimization and re-use. The attention devoted to the tall thin designer missed the need of "local" optimization at each layer of abstraction. Providing an intellectual framework for a top-down successive refinement design flow was a major accomplishment of the MC way, but it did create a certain level of confidence in computer scientists and of engineers in neighboring disciplines in being able to design competitive integrated circuits that was not warranted. Simplification in the design process, while excellent for explaining methodologies and tools, may create problems when trying to bring novel integrated circuits to the market.

The orthogonalization push versus separating design and manufacturing was a great success of the methodology albeit contradicting in part the figure of the tall, thin designer. It did allow the creation of EDA companies and computer programs to verify whether a mask set satisfied a set of design rules that could guarantee the yield of the manufacturing process (ECAD that eventually merged with SDA to form Cadence was offering design rule checking software that had and still has great success). VTI and Silicon Compilers were the results of MC teaching and research.

The question is whether the lessons from the MC approach are still valid today and will be a good path for the future. I do believe that the intellectual insight on the design process and on how to look at it for simplifying it, will be with us forever.

The panelists will argue about the specific characteristics of the MC approach. Luca Carloni will stress three main characteristics of the methodology: i) interdisciplinarity; ii) vertically integrated design

978-3-9815370-0-0/DATE13/©2013 EDAA

methodology; and iii) the need for open standards. Bernard Courtois will dwell into: i) simplification of the design rules and ii) decoupling of design from manufacturing. Hugo de Man will provide a historical perspective and the impact of the MC approach in Europe. Antun Domic identified in separation, abstraction simplification and education the major pillars of the MC approach and argued that these pillars will still be valid for many years to come. Jan Rabaey will conclude pointing out how the separation of concerns between manufacturing and design played a fundamental role in the semiconductor industry but warns that for More than Moore technologies this is not a done deal. He also reminds us of the chapters in the VLSI book where energy minimization and synchronization were discussed, very hot topics today.

Panelists' Position Statements

Prof. Luca Carloni, Columbia University, USA

The heritage of "Mead & Conway" starts with a book that has been studied in hundreds of universities and influenced thousands of engineers. The book, innovative in its format and organization as much as in its content, brings together the essential background concepts from different disciplines in Electrical Engineering and Computer Science before proposing a truly innovative design methodology. The methodology has guided students to combine theory and practice in learning VLSI design while inspiring them to work collaboratively on hands-on projects that require thinking in a very large scale.

Thirty years later what has remained the same includes: (1) the importance for interdisciplinary approaches to research and development, (2) the continuous quest for new vertically-integrated scalable design methodologies, and (3) the need for open standards and interchange procedures that foster innovation by enabling collaborative engineering across institutions and beyond geographic constraints.

Early in the preface of "Introduction to VLSI Systems" Carver Mead and Lynn Conway state their goal of filling a gap in the literature by introducing "all EE/CS students to integrated systems and architecture design." In doing so, they have formed a bridge between the two fields overcoming traditional boundaries that had until then confined device physics and integrated circuit design to EE departments and digital system architecture to CS departments. Thirty years later other traditional boundaries must be crossed to address new challenges across all fields of engineering. The continuous progress of semiconductor technologies has produced a chain reaction that brought us the personal computer, the mobile phone, the Internet, and the smartphone and leads to a future where Ubiquitous Computing will permeate all aspects of human endeavour. In this context interdisciplinary approaches are a must and a growing number of engineering students need to master skills that are traditionally taught in separate university departments. For example, the development of many state-of-the-art distributed embedded systems found in cars or airplanes requires knowledge of control and signal processing, taught in Electrical Engineering departments, as well as knowledge of formal methods, real-time systems, and hardware/software co-design, taught in Computer Science departments. The list of examples extends beyond engineering, given the impact that electronics and computing increasingly have on all disciplines, to the point of enabling the birth of new fields such as biomedical engineering.

In the specific disciplines of integrated circuits and computing systems, Moore's Law continues to increase the transistor density, but power-dissipation and wire-delay constraints have forced the emergence of heterogeneous multi-core systems-on-chip. An SoC is a distributed concurrent system hosting a growing variety of programmable cores, configurable cores, specialized hardware accelerators, and memory modules, all interconnected by an on-chip network. The combination of heterogeneity, concurrency, and distribution makes SoCs very difficult to design, program, and validate. In this specific context, the heritage of Mead & Conway lives on due to the continuous need to develop new vertically-integrated design methodologies, which requires reinventing the stack of levels of abstractions to tame design complexity while unleashing performance scalability. On the other hand, given the continuous progress of the semiconductor industry in the realm of nanometer processes, the emergence of new disruptive technologies like 3D-integration and chip-scale optical communication, and the crucial role played by software in all its forms, from device drivers and operating systems to virtual machines and user applications, it has become critical to accommodate the modeling of heterogeneity across all levels of the design flow.

Finally, as the Internet has accelerated the process of globalization, simplifying cultural and professional interactions across geographic distances, major changes are to be expected in the way engineers collaborate to conceive, design, prototype, and test complex systems. Fast access to large-scale amounts of information combined with the evolution of social networks, the progress in data mining, data visualization, and human-computer interaction, and the rise of additive manufacturing offer still unknown opportunities for collaborative engineering. While it is impossible to predict exactly what lies ahead, the role of open standards, intermediate formats, interchange procedures, and

automation tools will be instrumental in an increasingly interdisciplinary and collaborative environment.

Dr. Bernard Courtois, CMP, France

We had 30+ great years. Where should we go now?

Two major outputs of the Mead/Conway methodology were to simplify the design rules, and to decouple the design from the process. This is today less and less true, considering the increase of the number of design rules along downsizing (400 DRCs on .35µ, 5000+ on 28 FDSOI), and the necessity of design for manufacturability rules.

One major consequence of these 2 outputs has been the establishment of Service Centers for the manufacturing of ICs like CMC, CMP, MOSIS and other National Services in various countries. Two questions today that should be addressed for the future are:

- how this manufacturing paradigm (in short: sharing the cost by sharing wafers) that has been the foundation of these Service Centers is going to evolve (will there be still a need of these Service Centers)?
- shouldn't we address more and more non-electronics communities?

These two questions are briefly addressed in the following.

The evolution of the manufacturing paradigm can be assessed according to a timescale: moving towards nanoelectronics, and next to beyond CMOS computing paradigms.

Nanoelectronics: downscaling leads to escalating costs of manufacturing. This means actually an increased need of these Service Centers, and further the need of a cooperation between these Service Centers. Such a cooperation was identified already a long time ago by CMC, CMP and MOSIS (announced at DAC in 2002).

Moving forward to beyond CMOS computing paradigms leads to ask whether these new computing paradigms will be still suitable for the "share of costs by sharing wafers" principle.

Here the answer is not obvious. The following is based on a quick review of these paradigms based on the synthesis published by the NANO-TEC EC Project (Ecosystems Technology and Design for Nanoelectronics).

Between nanoelectronics and beyond CMOS potential paradigms, we can distinguish:

- charge based state variable technologies
  - graphene: ✓ (answer is yes)
  - nanowires: ✓
- non charge based state variable
  - MEMS/NEMS: ✓
  - spintronics: ✓

Among beyond CMOS potential paradigms, we can distinguish:

- molecular electronics: ? (question mark)
- memristor-based computing: ?
- QCA: ?
- quantum computing: ?

Actually the question mark is mostly based on the possibilities to manufacture in large quantities.... If there are no cost-sharing options, or no large quantity possibilities, then we might observe a move from (centralized) Service Centers to a more decentralized scheme, similar to the NNIN in the USA. The consequence will be a question mark on the added value of these Service Centers. Their added value might shift to help commercial developments from the technology research labs.

Turning now to the second question, i.e. to address more and more non electronics communities, it may be answered by a requirement to address more and more societal needs, reinforced by the observation that electronics is becoming more and more a commodities industry (☹), like the computer science became more and more a consumer, commodities, industry.

A few examples of such communities are:

- health care, biomedical
- energy
- environment
- security

These communities might be addressed by pointing out how electronics (ICs, MEMS, etc.) can be designed to address their problems. Examples are electronics for biomedical applications [1] and electronics for the management of energy [2].

In conclusion, it can be said that we will soon be at a crossroad. The end of the CMOS roadmap will be soon in front of us.

There are (almost) present technologies like 3D, spintronics, to overcome the technology/financial cliff, but we should also address quite soon other avenues like non-CMOS based computing, and a stronger dissemination towards non electronics communities.

References:

[1] B. Courtois, Infrastructures for Education, Research and Industry: CMOS and MEMS for BioMed, *Invited paper* at the 12th World Multi-Conference on Systemics, Cybernetics and Informatics (WMSCI 2008), 29 June – 2 July 2008, Orlando, USA

[2] B. Courtois, Electronics for Energy Management *Keynote* at ASQED 2011, July 19-20 2011, Kuala Lumpur, Malaysia

Prof. Hugo de Man, K.U. Leuven & IMEC, Belgium

Before 1970 only a handful of universities had semiconductor labs where device physicist's and circuit people designed and fabricated bipolar analog circuits and wrote the first university ICCAD tools. In 1971 Intel produced the first 4 bit microprocessor and thus triggered the digitalization of the world. This drew the attention of the computer science community which was not very familiar with IC technology and design.

To address that issue, Carver Mead and Lynn Conway (MC) developed a visionary shared design culture for the CS academic community. It enabled students not only to design their own simple digital chip architectures but also to get them fabricated cheaply. It was based on three basic principles: the use of simple scalable design rules for nMOS technology, exploiting the inherent geometric regularity of datapaths and PLA controllers and, last but not least, the development of the Multi-Project Wafer (MPW) service MOSIS based on a common layout format (CIF2.0) and the use of ARPANET to submit designs. The methodology, described in the famous MC book in 1979, was soon adopted by numerous universities and transformed hundreds of EECS students into digital IC designers who got functional silicon back just as fabless companies do today. In Europe it triggered at first separate similar initiatives in the UK, Germany, Scandinavia, France and Belgium.

The impact has been and still is enormous but, in order to discuss the evolution, one must first answer what has changed during the past 30 years of More Moore. First of all, 3 micron nMOS technology evolved into 28nm CMOS technology today for which standard cell layout is the most natural and most productive layout style. Going from two- to multilevel interconnect eliminated area loss in routing channels and around 1980 Brayton et al. came up with multilevel logic synthesis tools making the need for structured data path and PLA design to disappear. By 1990 HDL based RT-level synthesis; technology mapping and high level synthesis took over, while back-end P&R tools and test insertion tools did the rest. Below 90nm, technology has become much more complex as static leakage and variability pop up and sub-wavelength lithography makes layout rules and mask cost to explode. So besides the top-down design strategy not much is left from the simple structured MC design methodology. So today we can state that the main heritage is the great idea of the MPW service.

This is why in 1989 the EC launched EUROCHIP: a Pan-European VLSI Prototyping MPW service for universities. However, in contrast to the original MC structured design based on university CAD, it was decided from the beginning to provide cheap access to professional EDA tools, to cell libraries, to layout rules and to device models of participating silicon foundries and to organize training courses to use them. In this way about 400 universities acquired IC design expertise of direct use to the predominantly mixed signal oriented European semiconductor industry. This strategy was also extremely beneficial to the many EE universities in Europe with top level analog design expertise as it allowed them to stay on track with Moore's law. In this way ASIC design skills also proliferated into SME's and that is why in 1992 the EC launched CHIPSHOP to provide MPW prototyping to SME's.

In 1995 both EUROCHIP and CHIPSHOP were merged into EUROPRACTICE which in addition offers low volume MPW prototyping to industry. Today EUROPRACTICE organizes more than 200 MPW runs/year and services 530 EU universities and 100 R&D labs as well as 500 SME's worldwide whereby EC financing is provided for technical support of EU academia while prototyping is paid for by the end users. In 2011 a total of 563 designs were processed in a wide variety of technologies ranging from 500nm down to 40nm CMOS.

However, what we see is a great diversification. Besides classical logic CMOS, EUROPRACTICE offers 13 RF CMOS, 5 Hi Voltage CMOS, 6 MEMS and 2 Photonics MPW services addressing a strong tendency towards training in More than Moore technologies for medical, automotive and wireless sensor applications. This trend is typical for the EU fablite and fabless semiconductor houses where the focus is less on commodity processors or memory but on providing added ambient intelligence value to classical products and services.

So, the main heritage of MC has historically been in creating a shared IC design culture in academia but its actual impact is the massive and cheap access for academia and SME's to the most advanced More Moore and More than Moore design and fabrication technologies.

Dr. Antun Domic, Synopsys, USA

I am sure every one of you has at least heard of the book Introduction to VLSI Systems, by Carver Mead and Lynn Conway, first published by Addison-Wesley in December 1979. If you are interested, a new copy of the first edition can be found on Amazon for only $61 plus taxes and shipping.

The impact of "Mead & Conway" on our industry has been immense. I do not have the time to describe in detail the many contributions of Carver Mead and Lynn Conway to EDA, but I would like to highlight two of them, the multi-project chip and the lambda rules

concepts, that we owe in particular to Lynn Conway, and that I have personally experimented.

At MIT Lincoln Laboratories, I have had a chance to use MOSIS (Metal Oxide Semiconductor Implementation Service) to fabricate and test the first CMOS chip (A. Domic, et al., Report on the Lincoln Boolean Synthesizer, ICCAD 1983). MOSIS stem from MPC79, an idea of Lynn Conway to test the "new Mead & Conway design methods" on the field. During the fall of 1979, Lynn Conway and her colleagues at Xerox PARC organized a "live demonstration" of the full flow: 124 designers from universities around the world contributed 82 chip designs in CIF2.0 format; Hewlett-Packard agreed to donate wafer fabrication, and the packaged chips were shipped back to the designers in January 1980. MPC79, and then MOSIS were based on internet – then ARPANET – and contributions were sent electronically. As of today, MOSIS has fabricated more than 50,000 chip designs for hundreds of universities and research centers. The idea has been extremely successful, and has spread in multiple instances around the world: many of you have probably used Bernard Courtois' CMP (Circuits Multi-Projets) to verify on silicon your ideas at a reasonable cost.

At Digital Equipment, we used lambda rules to design the layout of four generations of microprocessors, at 1.5, 1.2, 1, and 0.75 microns, including the first generation of the Alpha; thanks to lambda rules, layout engineers could move smoothly from one technology node to the next, and layout generators fueled productivity from 50 transistors a day – in the good days – to 500 and more. By hiding the physics of silicon, the separation between design and fabrication by means of abstraction has allowed the birth of modern silicon foundries, and has made the many opportunities offered by semiconductor technology understandable, and therefore accessible to a wider set of engineers – including application, system, and software designers. I am convinced that this is an extremely valuable achievement that we should strive to preserve, especially today that the sheer complexity and the intricacies of nanometer lithography threaten this separation, as well as another important achievement: cell-based design methodology.

New challenges lie ahead, both technical and economical: the pace of innovation has accelerated, there is more and more integration, and we need to cope with the impact of new semiconductor technology. However, the heritage of Mead & Conway – separation, abstraction, simplification, and last but not least, education – are important values even for the years to come.

Prof. Jan M. Rabaey, UCB, USA

It was in the late seventies that Hugo De Man (my then adviser) pushed a plastic-bound document in my hands – "Jan, you should read this". It was a pre-print of what would later become the famous Mead-Conway VLSI book. It definitely provided for some spellbinding reading.

The most important legacy of the book was indeed that it helped create one of the first "orthogonalization of concerns" (to quote Alberto Sangiovanni-Vincentelli) abstractions in the semiconductor industry, separating manufacturing from design. As such it opened the door for thousands of enthusiastic students and engineers alike to get their go at the design of an integrated circuit. The rest is history. As amply stated in the other contributions, the Mead-Conway concepts formed the base of the development of the IC fab concept and gave rise to the creation of the EDA and fabless semiconductor industries. It put forward the idea that design process could be automated, and that the artifact could be "compiled" from a higher-level description language. While it took the academic and industrial research community quite some time and some false starts to truly make it happen, it ultimately did. High-level design synthesis is a reality today.

Yet, it is ironic that just as we approach the 35th year anniversary of Mead-Conway, its fundamental assumptions – at least as far as they concern silicon processing, maybe under serious threat in the next decade. With IC manufacturing becoming ever more complex and running into some serious challenges, the manufacturing abstraction may become untenable. This in term has the potential to topple the whole stack of cards that is built on top of it.

How to maintain the separation between manufacturing and design is one of the most serious challenges facing the information-processing community. Many avenues are open including innovative manufacturing strategies (such as self-assembly), novel materials and devices, different models of computation (such as neuro-inspired – another one of Carver's bets), or introspection on the function and nature of information processing in a future world. Plenty of things to be excited about.

Finally, one of the most important and often understated parts of the Mead-Conway legacy is that after more than 30 years and with the design process of integrated circuits going through many transformations, a large fraction of it still remains fresh and relevant to today's design of information-processing devices. The text did not shy away from trying to address some fundamental questions, such as the role of synchronization (in the Chuck Seitz Chapter), massively parallel architectures and bounds on energy dissipation. With the manufacturing abstraction challenged, these topics have become more relevant than ever.