

Chapter 10

System-Level Design

Joel Björk, Peter Caldera, Paul Fugger, Karsten Einwich, Karl Hollaus, Wolfgang Pribyl, Jochen Reisinger, Willy Sansen, Andreas Wiesbauer, and Horst Zimmermann

Abstract This chapter gives an overview on the system design and its different methods to implement it in the whole development process. Various aspects that have to be taken into consideration to reach the described goals are discussed.

Keywords Design flow • Analogue design • Digital design • Design methodology • Simulation • Modelling

10.1 Introduction

The design of systems¹ in the More than Moore domain is tightly connected with the design of system-on-chip (SoC) and system-in-package (SiP). The close integration of digital and analogue functionality with software saves not only space but also cost and power. MEMS² and MOEMS devices, discrete components, and thin-film devices further extend the functionality available in a package.

This miniaturisation does, however, lead to increased interference between the parts, significantly reducing performance if it is not taken into account in the design stage. The heterogeneous nature of such systems greatly complicates the design work.

The design flow must be capable of incorporating a variety of hardware technologies, as well as software at a high abstraction level. It must also be able to exchange data between the different parts of the system at the very lowest levels to enable system-wide modelling and simulation.

¹The definition of a system, as used in this chapter, is ‘a combination of technologies integrated into a product’.

²Micro (opto) electromechanical systems.

J. Björk, P. Caldera (✉), P. Fugger, K. Einwich, K. Hollaus, W. Pribyl, J. Reisinger, W. Sansen, A. Wiesbauer, and H. Zimmermann
Technikon Forschungs und Planungsgesellschaft mbH, Villach, Austria
e-mail: peter.caldera@infineon.com; office@eniac-austria.net

10.2 Background

The electronic design automation (EDA) industry can be traced back to the beginning of the 1980s. Until that time, the development of the tools needed for hardware design was carried out in-house by the hardware manufacturers.

The market is dominated by a few, predominantly American, vendors offering suites of software composed of modules designed for specific purposes. Each vendor is committed to one or a few proprietary workflows but offers the ability to export or import data in (semi)standardised formats. This means that users have the opportunity, though limited, to mix tools from different vendors or, as is also common, produce their own tools. There are also a number of smaller vendors providing complementing modules, often for specific applications.

The business landscape for chip suppliers has changed considerably in recent years. Their position in the value chain has been extended from pure silicon suppliers to system suppliers, integrating not only digital circuits, but also software, see Fig. 10.1, and More than Moore technologies.

One of the motivations for such a move can be seen in Fig. 10.2. It is estimated that the 265 billion dollar semiconductor industry would generate 1.340 billion dollars of value in the next tier.

This has led to a change in how the industry perceives their products. The designs are no longer monolithic, where all of the value is created in a single piece of silicon. Value is instead created from the combination of possibly several types of hardware with software. This extension of the value chain, in part driven by rapidly evolving SoC and SiP technologies, has radically altered the requirements

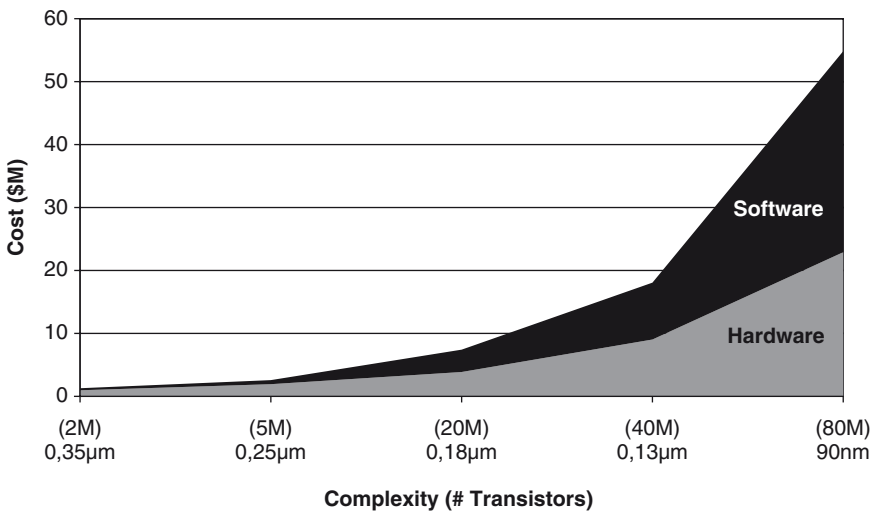


Fig. 10.1 Development of hardware and software costs over time for a typical ASIC (Source: Adapted from IBS, Inc)



Fig. 10.2 Semiconductor value chain – each market is drawn to relative scale

placed on system design tools. These tools must, just like the market they serve, expand to cover complete systems.

10.2.1 Diversity

More than Moore technologies are core enablers for ubiquitous computing, being part of everything from consumer products, to cars, to medical products. Consequently, the markets are more diverse than for the traditional chip industry and so are the lifetime profiles of the products.

- A mobile phone has a life expectancy of about 2 years and is mostly kept at room or pocket temperature. A malfunction is annoying but it is hardly life threatening.
- The anti-brake-lock controller of a car on the other hand is expected to outlast the car where it is fitted while enduring extreme environmental conditions. A failure here can lead to injuries or even loss of life.

The chip industry needs a way to parameterise these profiles and include them in the workflow.

For More than Moore technologies, system-level design will increasingly become the key design method. The costs for system-level design are relatively low, however, the impact of time-to-market and revenue are orders of magnitude higher than the impact of the subsequent design steps. Because of the low system-level design costs the market for tools is relatively small, thus such tools are poorly treated by the CAD vendors. State-of-the art system-level design tools are usually strongly focused on an abstraction level and/or application domain and are poorly integrated into the overall design flow. This becomes even worse for mixed-signal

and heterogeneous systems – the area where European companies have outstanding competencies.

10.2.2 System Design Goals

System design is mentioned here as the bridge between application design on the one hand and component and software design on the other.

Application design delivers a feature description in an application- and user-dependent language. Languages and formats are mainly defined by the market place. This means that the application-oriented community speaks many different languages when defining features and performance requirements. On the other side, the electronics design community also uses different languages, for example, Logic, Mixed Signal, RF and Sensors/Actuators.

In many cases, different kinds of standards play a very important role, for example:

- Transmission standards for mobile communication (EDGE, GPRS, etc.)
- Media standards for multimedia (compact disc, etc.)
- Security standards for chip cards
- Interface standards for computing systems (USB, etc.)
- Standards for electromagnetic radiation limits

Goal number 1: The creation of a clear ‘translated’ specification of features and functions that can be read into the system design flow:

- No risk that specification items are lost
- No risk that the specification might be misinterpreted (translation error)
- Enables a comparison (verification) of the input specification with the implemented system capabilities

Goal number 2: In all cases, it is key that the system implementation is verified against relevant standards. The other side of the bridge – the components (& software) side – will be defined in more detail within this chapter.

10.2.2.1 Application Example

The mobile phone is without doubt the most ubiquitous electronic device today. Over one billion devices were produced in 2007. It is also the main driver for the development of SiP and More than Moore technologies. Figure 10.3 shows, to scale, how mobile platforms have evolved over a period of 6 years.

The system includes RF transceiver, BAW³ filters, processor, persistent and volatile memory, MEMS devices, and high-voltage parts. The lower end of the

³Bulk acoustic waves.




Fig. 10.3 Evolution of mobile phone platforms – the value below each picture indicates the number of components needed to realise the phone

market is increasingly being served by SiP solutions while higher end smart phones are being built using components from several manufacturers.

10.2.3 The Design Flow

The design flow is the structured process used to turn an idea into a product. This process has become more and more complicated as the functionality of the hardware has increased while the physics of deep-submicron processing has added new restraints in every generation.

The details of this process vary from organisation to organisation. A general description must, therefore, be created at a generic level. The specification of the product is translated into a high-level language such as SystemC. Subsequent steps then model and simulate parts of the system at more and more detailed levels. Any changes are updated across all the abstraction levels, referred to as back annotation. This process is then repeated until a design has been achieved that meets the specifications. An overview of the abstraction levels and the languages used is provided in Fig. 10.4.

The majority of the efforts in these flows is concentrated on the area shown in a darker colour in Fig. 10.4. These tools provide very accurate results using commercially reasonable computing resources, as the majority of the value of a monolithic system is created here.

However, the majority of the design effort in a More than Moore design flow is expected to take place at the system level and the very lowest abstraction levels.

10.2.3.1 The Hell of Nanophysics

With the advent of modern deep-submicron technologies, design is becoming more and more complex. The most dramatic effect can be described as ‘digital goes analogue’. Because of the behaviour of deep-submicron devices (non-ideal device characteristics such as subthreshold, gate leakage, etc.) even pure digital circuits

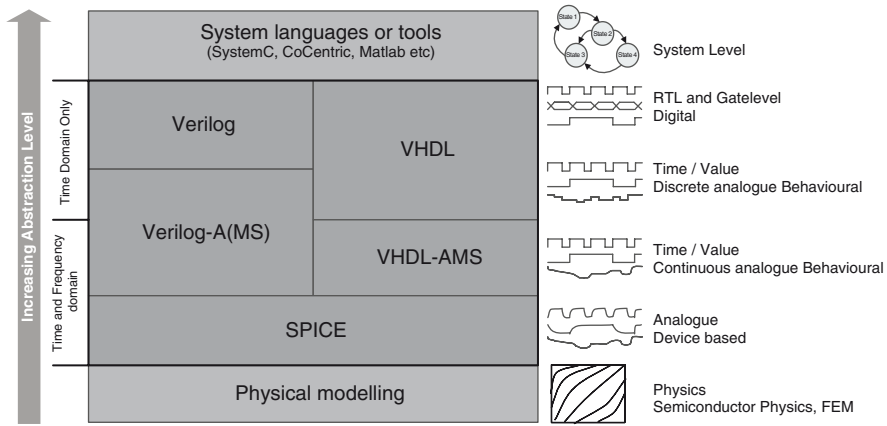


Fig. 10.4 Abstraction levels in design and simulation

have to be simulated with consideration given to different analogue effects, including back annotation from the ever-increasing influence of routing resistance, capacitance, and even inductance.

These device effects also deteriorate the performance of analogue blocks. To compensate for this, more digital correction and calibration techniques have to be applied in close co-operation with the analogue designers. The effects mentioned earlier call for a completely new simulation and verification strategy, which is best tackled in a hierarchical way, starting at the system level.

10.2.3.2 More than Moore

In design flows of the future it will, in general, not be enough to start the design work by creating a model from a paper specification. Furthermore, the specification must be developed iteratively by verifying numerous application scenarios and use cases. Therefore, an executable platform will become more and more essential. This platform must permit the execution of real-time use cases for the overall system including the environment.

The overview of such a design flow is presented in Fig. 10.5. The exact composition and components depend greatly on the type and complexity of the product.

10.2.4 The State-of-the-Art Technology

State-of-the-art design tools clearly serve one physical domain, operating with different abstraction layers. Two different design flows are required to simulate a simple SiP with a 60-nm digital and a 120-nm analogue part. This means that two

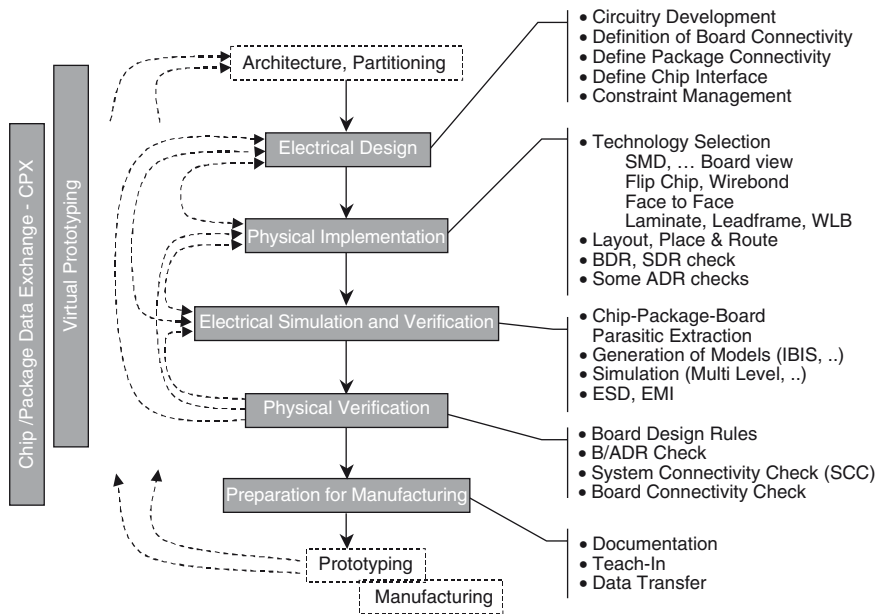


Fig. 10.5 Example of More than Moore design flow

different simulator environments are needed. To simulate the complete system, the two environments have to be coupled using TCP/IP or shared memory, which increases the complexity of the system and dramatically increases the need for computing resources.

State-of-the-art simulators work independently, one after another. We start with chip and package development and simulation. When a chip is finished, board design and component engineering begin their own optimisation. When a board is finished, system integrity and EMI are simulated and tested.

The performing of a chain of simulations one after another is time consuming, and the risk of having to start again from the beginning, if model and parameter assumptions fail en route, is very high. Nowadays, the impact of such risks on manufacturing is reduced by adding extra components that act as spare parts for later improvements if a system should fail any end tests.

A few software packages are available for the simulation of PCBs. Full 3D field solvers are very slow and thus not suitable for the design of large systems. Highly non-linear or large complex problems are still not a routine task in the design process at present. Software packages that apply approximation techniques are particularly dedicated to the analysis of PCBs and are usually faster than full 3D field solvers. However, some physical effects are not taken into account and substantial errors are exhibited. Unfortunately, the implemented approximation methods in state-of-the-art products are not made public, being kept secret by tool suppliers, which completely impedes fair validation.

The closed nature of the tools makes the construction of tool chains a laborious and often bespoke work. Little reuse is possible between projects and a lot of information is lost in translation, reducing the accuracy of the results and impeding automation.

10.2.4.1 Current Trends in Microelectronics

More than Moore will still be the dominant way forward in the nanoelectronics industry in the coming years, driven by Flash and DRAM technology. The importance of More Moore is steadily increasing as a number of physical and financial effects can be foreseen, which will slow the competitive advantage found in the next generations of nanometer-scale CMOS:

- *Leakage power* starts to dominate unless larger gate delays than performance scaling predictions are introduced.
- *Voltage headroom* shrinks and makes analogue and radio frequency design very difficult. New digital architectures providing functions previously only available in the analogue domain are able to mitigate this problem somewhat. This is primarily achieved through faster analogue to digital converters requiring less filtering of the input signal and switched mode amplification techniques. Moving to SiP solutions is another possibility, but it brings problems at the system level with increased integration complexity and power requirements.
- *Increasing variability* reduces predictability at the device level. Materials can no longer be treated as homogenous bulk materials, as very small local variations during the manufacturing process affect the properties of transistors and interconnects. This greatly increases the need for more precise process control to retain yield levels.
- *Signal integrity* decreases as the interaction between devices and interconnects increases. The design and simulation of digital circuit starts to increasingly resemble the analogue design flow with manual layouts of critical subcircuits and the use of circuit simulators. The effect is decreased productivity in the design and simulation stage.
- *Layout restrictions* imposed by lithography effects will continue to affect the freedom of the designs.
- *Global interconnect delay and variance* endangers the global synchronous concept.

All these factors combined with very high investment costs for sub 45-nm technologies are driving the chip suppliers to look for competitive advantages outside of the more Moore track.

10.2.4.2 Digital Design

Digital design in the context of system design, including the design of mixed signal systems, must be separated into two main parts: design and implementation, and functional verification.

Design and Implementation

The design process includes the task of describing hardware at a high abstraction level using, for example, SystemC or a hardware description language (HDL), such as Verilog, VHDL. The implementation process uses software tools to translate the high-level description into lower levels of abstraction, such as gate-level netlists or layout data. This flow is referred to as RTL2GDS.

One of the main drivers for the development of any electronic product is time-to-market. On the digital design side, reuse helps to meet short design cycles. IP-based design (intellectual property) means the use of commercially available IP, such as microcontrollers and interfaces, as well as company internal IP such as company proprietary interfaces and hardware filter structures. Another way to meet tight schedules is concurrent development of hardware and software. This means that computational efforts are moved from hardware to software, i.e., algorithms are executed as a program in a microcontroller system and not in hardware. Although this means that development tasks are shifted from hardware to software, it offers easy extension or update of a system and the possibility of separate evaluation for hardware and software. With respect to system projects, software development has in general already become the dominating effort.

Another key parameter for a successful product is cost competitiveness. Since one of the main cost contributors for an IC is tester usage (this can be between 20 and 50% of the overall development costs) measures must be taken to keep this time as short as possible – automated test circuit creation is available for boundary scan, memory BIST (built-in self-test), and scan test. These methods achieve high test coverage in acceptable tester times. Another direct contributor to the costs is yield. To reach high yields for digital designs, methods such as redundancy insertion for memory blocks as well as for other parts are applicable. Strict fulfilment of timing constraints, so-called timing closure, helps a system become resilient to voltage and temperature effects.

For many products, such as mobile devices, low power consumption is a must. On the digital design side, many different techniques can be used to achieve a low-power design, e.g., voltage scaling and power management in power islands and selective shutdown of the clock, called clock gating.

Functional Verification

Functional verification is the task of verifying that the logic design conforms to the specification. In everyday terms, it attempts to answer the question: ‘Does this proposed design do what is intended?’. Functional verification is NP-hard, or even worse, and no solution has been found that works well in all cases. However, it can be approached using various verification strategies:

Directed tests using simulation-based verification, called dynamic verification, are widely used to ‘simulate’ the design, since this method scales up very easily. Stimuli are provided to exercise each line in the HDL code. A test bench is built

to functionally verify the design by providing meaningful scenarios to check that given a certain input, the design performs to specification. The best usage of directed tests is in a regression environment with self-checking tests (go/no-go results).

Test bench automation is used to perform pseudo-random testing. A stress test for digital components can be implemented by using specially developed hardware verification languages, such as ‘e’, ‘VERA’ or ‘TestBuilder C++’. Millions of test vectors are generated in a semi-automatic way.

Property checking is a ‘formal’ mathematical approach to verification. A system is described by means of state predicates describing inputs, outputs, data transport and temporal expressions that make the link to cycle-based timing behaviour. Property checking tools calculate all the possible transitions and check whether these transitions fulfil the state predicates and the described timing behaviour.

All these methods are part of a common verification process, which is iterative (see Fig. 10.6).

Conclusion

Whereas design tasks are being increasingly driven by the reuse of IP delivered in any HDL, verification has to manage the exponential increase of complexity with rising design sizes. Figure 10.7 shows the relationship between complexity, represented by lines of RTL code, and the verification effort, represented by simulation cycles for two generations of the same processor architecture. Nowadays a split of 20–30% design and up to 80% verification effort is realistic and widely accepted as a matter of fact.

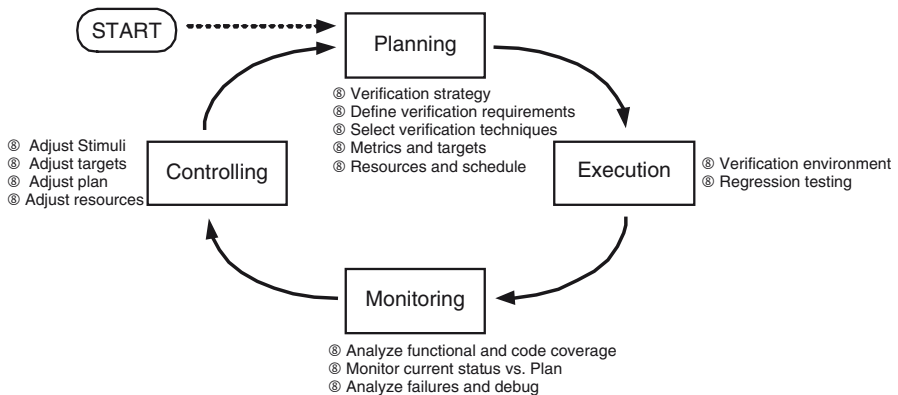


Fig. 10.6 Verification process

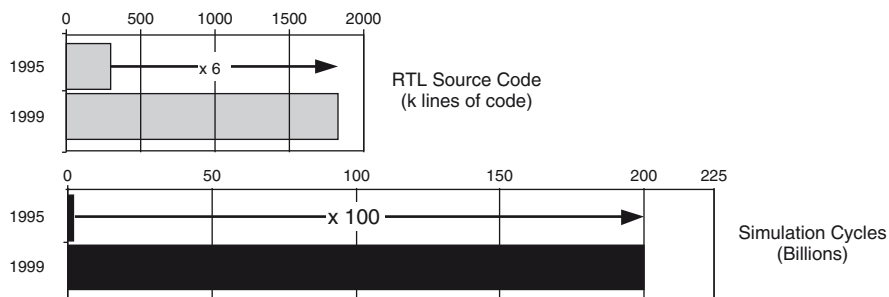


Fig. 10.7 Verification effort – circuit vs. simulation complexity growth over a 5-year period

10.2.4.3 Analogue Design

Numerical simulation became popular with the publication of the SPICE simulator in 1971. SPICE is an open source circuit simulator released under the BSD Copyright.⁴ SPICE and the different commercial and open source variations thereof are by far the most commonly used simulators for analogue circuit analysis. The program works by taking netlists and converting them into non-linear differential algebraic equations that are then solved. Supplying SPICE models for components has become the de facto standard in the industry.

In the eighties, the activities were focused on digital simulation techniques that resulted in two competing HDLs – VHDL and Verilog and the underlying simulation techniques. In the nineties, these HDLs were extended to the analogue domain (VHDL-AMS/Verilog-AMS). In addition, a couple of new concepts such as behavioural modelling at equation level (instead of macromodelling in SPICE or S-Parameter) and physical domains were introduced. In 2000, the activities went back to the digital domain and were focused increasingly on verification. One result of this activity is the standardised HDL SystemVerilog, which is based on Verilog and has its strength in the verification of digital implementations. Thus, concepts such as assertion-based verification and constraint randomisation were introduced.

From a More than Moore viewpoint, these activities only focus on the implementation level. At the system level, we currently have a much more heterogeneous and non-standardised situation. As a result there are a couple of proprietary and very specialised tools such as Matlab (Mathworks), CocentricSystemStudio (Synopsys), SPW (CoWare) or ADS (Agilent). A number of these tools were derived from the Berkley OpenSource framework Ptolemy I. With the introduction of Ptolemy II at the end of the nineties, Berkley shifted their system-level activities to an academic level. In 1999, an HDL based on the programming language C++, called SystemC, was introduced. The first version of SystemC focused on the implementation level (cycle accurate RTL), however, with the progression to

⁴<http://embedded.eecs.berkeley.edu/pubs/downloads/spice/index.htm>

SystemC 2.0 system-level concepts became more dominant. Thus, a couple of new concepts and abstract modelling techniques such as the separation of communication and behaviour, gradual refinement and transaction level modelling (TLM) were widespread. Thus, SystemC grew in importance for the system level design of large digital systems.

However, our world is not digital. At least the interfaces to the environment and to the user are analogue. As analogue components in general cannot be shrunk in the same way as digital components, and the analogue performance of the devices becomes worse, the ‘small’ analogue part can dominate the die area, yield, power consumption and design costs. To overcome this more and more analogue functionality will be shifted to digital functionality and/or imperfections will be digitally compensated. Thus, in general, we will have digitally assisted analogue components. This results in a much tighter interaction between analogue and digital, which will become a new challenge for system-level design.

This was the motivation for recent activities to extend the SystemC language to the analogue mixed signal domain (SystemC-AMS).

Initiatives such as SystemC-AMS are trying to close the gap to the digital domain at the higher abstraction levels. The accuracy of SPICE and S-parameter simulators is starting to deteriorate as devices are getting smaller. Multiphysics FEM simulators are, therefore, being explored to improve accuracy in what is commonly referred to as the nanophysics hell, described earlier, but they are still some years away from providing reliable results in a commercial setting.

A system-level approach for the More than Moore domain means that the design flow also needs to account for software, discrete components, and M(O)EMS devices. It must be able to describe not only the function of the component but also the restraints, interfaces and physical properties of the device at several different abstraction levels.

10.2.4.4 System Design

The most important feature of system design methodologies is the ability to cover more than one hierarchy level. This means that, for example, the top-level system model consists of a set of highly aggregated models. The complexity of these models is always a compromise of simulation complexity/simulation time and the level of details covered. For example, a simulation model for a CPU to check the implementation of software algorithms will never be at transistor level – even simple timing conditions will not be covered.

But there can be a pressing need to get the capability to ‘dive’ to a more detailed level. For example, the non-linearity effects of an ADC have a strong influence on the signal-to-noise ratio (SNR) of a converter system – in this case the model for the ADC can be a simple quasi-static model for the evaluation of first order system effects – but for the SNR check a very detailed model of second order effects has to be applied. The resulting design view is depicted in Fig. 10.8, which allows a designer to work on several layers with different resolutions simultaneously.




Fig. 10.8 Design views – (a) represents an ideal environment, (b) is the current state and (c) represents the multi level, mixed resolution approach

10.2.4.5 Productivity Gap

An extension of Moore’s law predicts that, taking economics into consideration, the number of transistors that is possible to put on a single chip is doubled every 18 months. The productivity of a hardware designer is, however, only doubled every 2 years.

The amount of software needed to provide the functionality for silicon systems doubles every 10 months while software productivity takes 5 years to double.

This gap between theoretical possibilities and practical capabilities, illustrated in Fig. 10.4, has been named the productivity gap or the silicon system design gap (Fig. 10.9).

This gap will continue to grow as each technology or layer adds nodes and interfaces to what can be described as a design mesh. Each addition of a node to the design flow brings new requirements and dependencies. The linear workflows of digital hardware design are broken up by frequent iterations with inputs from several different sources. This will drastically increase the amount of data transferred between nodes as new nodes are added and the interaction between nodes at both a functional and physical level increases.

The solution to reducing the gap within each node is to increase the abstraction at the system level while reuse of, particularly analogue and mixed signal, pre-qualified IP reduces the work as the abstraction level decreases.

As in any mesh, the communication between the nodes has a major influence on the overall performance. The nodes in a system design flow are inherently heterogeneous, and so is the data exchanged between them.

Each node uses a format specific for its task, often differing between vendors. Therefore, connecting two tools often requires the implementation of a translation node between the tools. Such a node only targets the communication between two other specific nodes. Adding a new node, for example, an alternative technology for

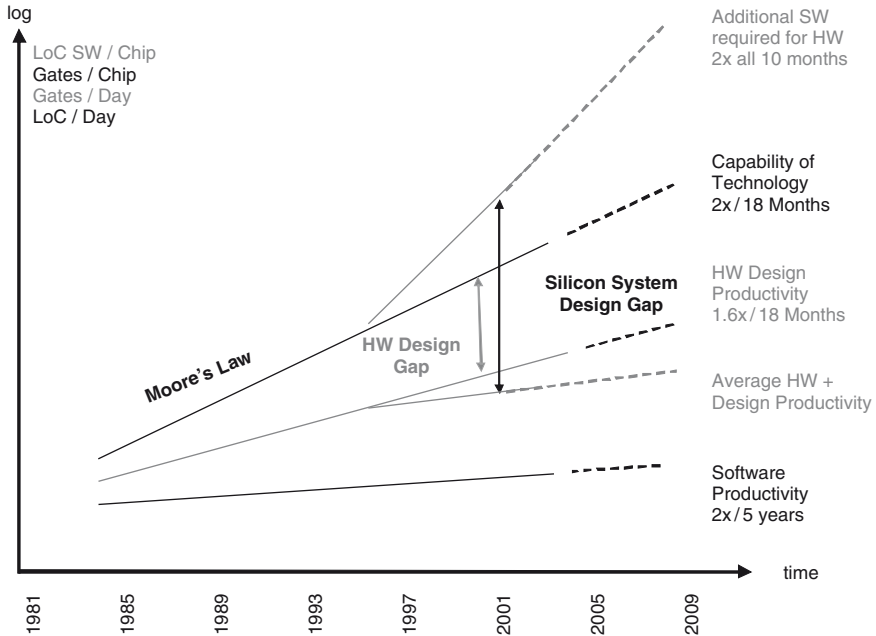


Fig. 10.9 The productivity gap

a core component can, therefore, be a cumbersome task with limited reusability. As the complexity of system-level design workflows increases there is an additional need for a standardised way of exchanging data within the flow.

Increased complexity means that simulations require additional computing resources. The move to smaller device sizes also means that simulations previously carried out in the digital domain must be carried out in the analogue domain. The performance penalty in such cases is severe; a factor of 15,000 is not uncommon.

10.2.5 Research Subjects/Future Trends

The new technologies permit the integration of complete systems into a single chip or package. Thus, the integrated systems will become more and more heterogeneous and the design teams will become larger and larger. This leads to new problems:

- A paper specification is not sufficient for clear communication.
- The used technologies are not well suited to all the system parts.
- The interaction between analogue, micromechanical, and digital hardware/software becomes tighter and tighter as the analogue/mechanical parts have to be simplified and digitally assisted to achieve the required performance using

devices with worse parameters, and to reduce the die area of the non-shrinkable analogue/mechanical parts.

- The system-level verification becomes more complicated as several independent standards are being supported and the percentage of reused IPs increases.

There are several ongoing activities aimed at overcoming these problems. Methods for the development of executable specifications will be evaluated, and new modeling and simulation methods must be developed to permit an overall system-level simulation of real-life use cases. New verification methods must be developed that are suitable for system-level verification and which take care of the heterogeneity of an overall system and the increasing IP usage.

In particular, the ongoing SystemC/SystemC-AMS activities are a promising solution for this problem. However, research work is still required to raise these new techniques to a widespread industrial standard.

- Multiphysics simulation
- Error propagation
- Multi-technology
- Multi-scale: device (nm) to board (dm)
- Costs
- Analogue and digital design for deep-submicron technologies

The accuracy of simulators has improved remarkably in recent years. Simulations of a production mature technology in combination with design-for-manufacturability aware workflows regularly produce first-time-right designs.

The full analogue simulation of a complex system is, however, prohibitive from a resource point of view.

10.2.5.1 Analogue Circuits in Deep-Submicron and Nanometer CMOS

Deep-submicron and nanometer CMOS leads to ever thinner gate-oxide thicknesses. On average, the gate oxide is about 1/50 of the channel length. As a result, gate leakage starts showing up. In the future, different gate dielectrics will be used to suppress this leakage current. Many other effects are also visible, which render the realisation of analogue designs much more cumbersome if not impossible.

They are as follows:

- The voltage gain is severely reduced.
- The thermal and $1/f$ noise levels are increased.
- The random mismatch is increased by factors such as line-edge roughness and polygate depletion.

More research is thus needed on these parameters.

The supply voltage seems to stabilise around 1 V. This means that the SNR cannot be made very high. Two circuit solutions have been proposed hitherto. They are the use of continuous-time sigma-delta converters and the use of pulse-width

modulation converters. The latter solution converts amplitude into time. The subsequent digital circuitry readily exploits the high-density integration capabilities of nanometer CMOS.

In addition to these two techniques, digital calibration has to be used to correct the analogue values. This applies to a wide variety of applications in both analogue and RF.

Nevertheless, the speed is increasing, inversely proportional to the channel length. Receivers thus realise analogue-digital conversion functions closer to the antenna. On the other hand transmitters require higher supply voltages. As a consequence, they always end up on different chips, using different technologies.

More attention also has to be paid to microwave- or millimetre-wave IC design, which combines analogue, RF, and transmission-line design techniques. Communication systems require such blocks in all disciplines, such as wireless communications, automotive electronics, consumer electronics, etc.

Not only SNR and ADCs are affected by the ever lower supply voltages. Distortions in voltage-mode analogue filters are also a serious problem. A solution for this may be current-mode filters.

The gain of amplifiers is not only limited by the ever lower Early voltage (or ever lower intrinsic gain gm/gds) of ever shorter transistors, but also by the lower supply voltage. The low supply voltages tend to exclude common amplifier topologies for gain enhancement, such as cascodes and gain boosting with regulated cascodes. The use of cascaded amplifiers instead is not really a satisfactory alternative, because the bandwidth reduces with the increased number of amplifier stages.

Furthermore, switches with MOSFETs become a problem due to the low supply voltage because the ratio of off- and on-resistance becomes lower. Sample and hold circuits suffer from this and by gate tunnel leakage currents.

Band gap circuits for supply voltages lower than 1 V will be difficult to realise in a robust way. Hysteresis (memory effect) in a single transistor will lead to totally different circuit topologies for amplifiers. For example, analogue compensation or digital linearization will be employed at IC block level. Self-heating will become an issue once the transistor structures are down to 32 nm, requiring design methodologies and models that can be used to predict the impact on the performance of a circuit. Circuit layouts will become far more complicated since there will be a stronger influence (interdependence) from the circuit neighbourhood.

10.2.5.2 Layout and Modelling

Layout extraction and the modelling of parasitic devices and effects need to be improved.

In conclusion, there are tremendous problems that require a lot of research effort in order for analogue circuit design to keep up with the improvements in performance of digital circuits; otherwise, the performance of future systems-on-chip will be severely limited by the analogue parts. New basic circuit building blocks and new circuit architectures have to be developed and investigated before models for system-level design can be created.

10.2.5.3 Digital Design in Deep-Submicron CMOS

The simulation of digital deep submicron requires analogue functionality for a digital simulation. This means that synthesis of VHDL code and formal verification may not be sufficient to ensure functionality of a system. The reason for this may either be an ‘unpredictable’ CMOS as mentioned earlier or in the system integration itself. This means adding more and more functionality into a very small smart system, for example, into a SiP. This may also lead to non-predictable results such as high crosstalk, EMC and thermal problems or just reduced performance.

The only way to analyse this behaviour is with detailed multi-physics FEM simulation. FEM simulation is a nice solution and would help to solve system design problems, but again the problems simply tend to increase.

FEM simulation is a time-consuming process. It needs a tremendous amount of computing power, memory and also knowledge of the numerical process per se. Ways to increase computing power, such as domain decomposition, are described later. However, the detailed FEM analysis of a system would not help when we talk about system performance, which is simulated with a system simulator. System simulators handle noise sources, transfer functions but cannot handle S-matrixes from an FEM simulation of a subsystem, such as a package.

Therefore, to simulate future systems we need to add high-level simulators, such as SystemC, to deep physics analysis, such as FEM, on the one hand and define ways to handle these different abstraction layers of models and results on the other.

To summarise, we are focused on increased heterogenisation:

- Yesterday – Bulk of the value in one technology, handled by a simple tool chain of simulators. System integration is a step-by-step approach.
- Today – Value spread over several technologies including software; we talk of multi-technology simulation such as hardware software co-design.
- Tomorrow – Optimising costs through system space exploration, more flexibility, different layers of model abstraction and multi-physics simulation.

10.2.5.4 Executable Specifications

In design flows of the future it will, in general, not be enough to start the design work by creating a model from a paper specification. Furthermore, the specification must be developed iteratively by verifying numerous application scenarios and use cases. Therefore, an executable platform will become more and more essential. This platform must permit the execution of real-time use cases for the overall system including the environment.

For a number of application domains, it will no longer be possible to develop system specifications within one company. In particular for integrated automotives, a collaborative specification development between TIER2 (the circuit design house), TIER1 and the OEM (the car manufacturer) will become more and more essential. Thus, executable specifications must be exchangeable.

For other More than Moore applications, the size of the design team will keep increasing. Block-level specifications provided on paper will become ambiguous. It will become more difficult to specify the whole functionality, and it will be nearly impossible to verify the implementation against this written specification.

Also in this case an exchangeable executable specification, which can be used as a ‘golden’ reference and as a stimuli generator, will be essential.

10.2.5.5 System Simulation

The overall system simulation of real-time use cases is orders of magnitude too slow when using tools and languages that have their focus on the implementation and block level, such as SPICE, VHDL-/Verilog-AMS. For More than Moore systems, applying behavioural modelling techniques will also not solve this problem. New modelling and simulation techniques must be introduced to overcome this gap. The concept of Model of Computation (MoC) was introduced within Ptolemy and later became widespread with SystemC. Generally speaking, a model of computation *defines a set of rules that govern the interactions between model elements, and thereby specifies the semantics of a mode*. Therefore, SystemC implements a generic MoC, which permits the use and interaction of an arbitrary number of application-specific MoCs for digital system-level design. Thus, the user has the ability to apply abstract modelling techniques that permit an extremely fast simulation. The most known technique for SystemC is TLM, which permits simulation speedups in the area of 1,000 compared to classical VHDL/Verilog simulations. Recent SystemC-AMS activities will extend this concept to the analogue mixed-signal domain. Thus, SystemC-AMS provides a framework that permits the interaction of different analogue MoCs with each other and to the digital SystemC simulation kernel. This leads to advantages such as an order of magnitude higher simulation performance (Table 10.1), a significant reducing of the solvability problem of the analogue part, very good scalability and encapsulation of different model parts. Thus, a SystemC-AMS prototype supports a synchronous dataflow (SDF) modelling style to model non-conservative signal flow behaviour, especially for signal processing-dominated applications and linear conservative (network) descriptions including the interactions.

Table 10.1 Simulation times for a communication system (voice codec)

| Simulator | Real time | Simulation time |
|--|---------------|-----------------|
| SystemC-AMS simulation | 1-s real time | 1.5 h |
| VHDL RTL | 1 s | 300 h |
| Nano Sim | 1 ms | 15 h |
| SystemC-AMS simulation with reduced model complexity | 1 s | 90 s |

10.2.5.6 Finite Element Method

The finite element method (FEM) is a powerful method to solve boundary value problems, for instance, based on Maxwell's equations in the context of systems such as printed circuit boards (PCB), SoC or SiP.

For future systems to be able to achieve high accuracy and reliability, they will require an end-to-end simulation flow, i.e., from the electromagnetic field over the thermal field to the mechanical stress distribution. The electromagnetic field solution can be used to evaluate the electromagnetic compatibility, i.e., crosstalk between interconnects, and deliver the distribution of Ohm losses, which constitute a source for the thermal boundary value problem to be solved to highlight thermal hot spots [5], [6]. The thermal field represents the source for the mechanical stress analysis, providing information about the mechanical state of the problem [8].

Contrary to an analysis where the systems can be represented by circuits, scattering parameters, etc., a multi-physical analysis needs the electromagnetic field solution of the entire problem. The goal from the multi-physics point of view is the capability to analyse very accurately fairly large-scale electromagnetic, thermal and mechanical coupled problems (PCB, SiP, SoC, etc.). Knowledge of the electromagnetic, thermal and mechanical integrity is inevitable for leading edge system design exploiting the highest possible integration density.

Accurate simulations of systems comprising components made of non-linear materials need feasible material models from the computational effort point of view. A particularly challenging task represents the simulation of components made of materials exhibiting hysteresis [7], for instance, transformers with ferrite cores. The corresponding material models have to be a trade-off between accuracy and manageability for available computer resources. Because of the geometrical complexity, size and highly non-linear materials (Fig. 10.10) of these transformers the related boundary value problem at high frequency cannot be handled routinely by the finite element. Note that the permeability is frequency and temperature dependent (see also Sect. 10.2.5.12).

10.2.5.7 Thermal Simulators

The model representation for the thermal analysis is based on the principle that the conduction current density distribution represents the sources for the thermal field problem. A boundary value problem can be posed by considering all the thermal material parameters together with the thermal sources and suitable boundary conditions. Its solution yields the required temperature distribution.

Because of the fact that the bulk of the losses occurs in the chips (thousands of transistors that are assumed to be distributed arbitrarily), an additional type of thermal source for the thermal boundary value problem has to be taken into account. Thus, a thermal analysis considers mixed sources in general. The overall losses of individual chips in a SiP are known for different operating conditions.




Fig. 10.10 Complex permeability is frequency dependent (*left*); initial permeability is temperature dependent (*right*)

10.2.5.8 Mechanical Simulators

The mechanical boundary value problem has to be solved by considering materials with different elasticity modulus and thermal coefficients of expansion with the thermal field as source. Some of the materials involved are even non-linear. Thus, a FEM solver capable of analysing thermal stresses occurring in composed materials exhibiting different elasticity coefficients is required. All the materials involved have to sustain all the mechanical stresses that occur.

10.2.5.9 Domain Decomposition

Fast field solvers especially required for high-frequency problems are to be developed. The domain decomposition method is a promising way of coping with large problems that cannot be simulated as a whole [9], [10], [11]. Simply speaking, large problems are subdivided into smaller parts. The memory requirement of one part is essentially smaller than that required for the entire problem; hence, parts can be solved easily. To ensure the continuity of the electromagnetic field at the interfaces, appropriate interface conditions are interchanged iteratively during the solution process. Thus, domain decomposition methods allow efficient exploitation of parallel computing technology.

10.2.5.10 Time Domain Simulations

Investigations with FEM made in the frequency domain have shown a high memory requirement and long computation times in general.

The finite difference time domain (FDTD) method mainly prevails over FEM with respect to the memory requirement and computation time in the time-domain analysis. However, important drawbacks of FDTD are the staircase effect caused by the insufficient geometrical approximation of curvatures and the ill-defined assignment of material properties to the degrees of freedom.

Therefore, simulations of the electromagnetic field of structures in a system considering the full set of Maxwell's equations will be accurately carried out by FEM in the time domain.

It is expected that the discontinuous Galerkin technique, which is currently the focus of international research activities, will provide essentially faster solvers for FEM than the present ones.

10.2.5.11 Error Propagation, Sensitivity Analysis

The introduction of component models and variations in electrical parameters as a result of the manufacturing process introduces the challenge of passing errors on to the different stages in a simulation [12].

Sensitivity analysis is a powerful tool to study error propagation in the case of small perturbations. To accelerate the sensitivity analysis, the sensitivity matrix S for small perturbations, i.e., the Jacobian matrix, which maps the changes in the uncertain parameters onto changes of a signal, is needed. S can be calculated efficiently exploiting the reciprocity theorem valid for wave problems.

Once S has been determined, the signal change ΔU due to an arbitrary spatial material change can be approximated with sufficient accuracy by multiplying S with the vector whose entries are the material changes assigned to the finite elements within the limits, where a linearization can be assumed.

Linearization is no longer valid for large perturbations. Therefore, a parameter study of the impact of manufacturing uncertainties or of sources of interference on the operating performance will be carried out. In this case the basic wave problem has to be solved for each magnitude of perturbation by FEM.

10.2.5.12 Temperature Analysis of Strongly Coupled Problems

Some of the material parameters (electric conductivity, magnetic permeability and electric permittivity) of the electromagnetic boundary problem are temperature dependent. To express the mutual dependency of the basic boundary value problems, they can be regarded as strongly coupled problems. Thus, the electromagnetic and thermal problems have to be solved alternately, updating the material parameters successively until the process converges.

No commercial solver is capable of satisfactorily solving strongly coupled problems in the high-frequency range.

10.2.5.13 Modelling

Transistor modelling works due to old and continually improved SPICE model.

There are still no precise time domain models available for simple ceramic capacitors such as COG or X7R. Therefore, it is not possible to simulate the harmonic distortion of these components. Another example is a DSL transformer (see FEM simulation) where precise non-linear models do not exist.

There is also a lack of understanding about how physical properties of components are transferred to high-frequency SPICE subcircuits. This lack of analytic modelling methods forces the design engineer to rely on trial and error, with costly redesign when models fail to transfer all the relevant effects of a component into the design phase.

10.2.5.14 Methodology

Unified design database with different abstraction layers: To perform domain-specific system analysis (electrical, thermal, mechanical) interfacing to selected

(state of the art) analysis tools must be enabled. Therefore, export functionality from the SiP data model into data formats that can be imported into analysis tools must be created. Moreover, export functionality is required for miscellaneous purposes, such as linking to manufacturing. In addition, depending on the analysis and use model, analysis results must be fed back into the system. The required analysis types must be specified, the preferred simulation tools must be selected and correct interfacing must be defined and implemented. Interfacing will require mapping/translation from the unified data model into a format that can be interpreted by the selected tool.

IP reuse and IP protection: IP reuse of pre-qualified blocks will be essential. High-level models must be available for system integration, e.g., virtual car. Precompiled SystemC models provide IP protection and can act as executable specifications within a company or also offer the whole value chain.

Advanced simulator accuracy without increasing the simulation time: The accuracy of state-of-the-art simulation tools is too poor when applied within industrial environments for design and manufacturing. These environments only allow limited time for simulation runs, which is currently increasing dramatically as systems become more complex and the system operating frequencies go up. The core simulation components and numerical methods will have to be enhanced significantly to meet the challenging targets.

Precise electrical modelling of physical parameters: There is a lack of understanding about how physical properties of components are transferred to high-frequency SPICE subcircuits. This lack of analytic modelling methods forces the design engineer to rely on trial and error, with costly redesign when models fail to transfer all the relevant effects of a component into the design phase.

System-wide tolerance and error propagation: In the future, variances and error propagation will have to be included, thus reducing the need for costly Monte-Carlo simulation. There is currently no way of establishing how variations in components and errors from individual simulations propagate and affect the system as a whole. Costly margins and corrective circuits have to be built in to ensure the functionality of the system. A standardised framework to tie together the different parts of the simulation environment is needed.

New approaches to model and simulation parameter access across the development tool chain: Any change of the design and simulation process has to focus on model and parameter integration across the tool chain. New models will bring an even higher degree of parameter complexity in order to cover tolerances and system-wide error propagation. Parameter models that provide efficient and intelligent transfer of relevant parameters throughout the tool chain need to be investigated.

10.2.5.15 Interaction Across Boundaries

There are two main boundaries, already mentioned in the SRA as ‘side by side with ARTEMIS and EPOS’.

Artemis mainly deals with embedded systems; for system simulation, we talk about hardware software co-design. So the design environment or the design flow must be able to link with the needs and demands of ARTEMIS. One example of a future area of improvement is software development in the automotive industry. Today we expect 3,000 errors per one million lines of code, mostly timing related (Fabio Romeo, 2001). To achieve zero defects, close connection to both European Technology Platforms must be established, and models and evolution schemes must be developed.

EPOS, or system development in companies: As integrated systems become more intelligent, the understanding of the overall context becomes extremely important in the specification phase. Especially in the automotive or aerospace industry the system knowledge is distributed between several companies. Design houses are usually at a TIER 2 position. It is essential for them to specify the circuits together with TIER 1, and they together with the OEM. Because of the increasing complexity, an executable specification that permits verification in the next level environment will become mandatory. The main topic will be to handle the different aspects in size, such as nm to m, or time, ps to year, for these systems.

Besides extremely high simulation performance, protection of the internal system architecture is essential to permit such an exchange of executable models.

Vertically: We need to establish a tool chain in parallel to the value chain from the designer to the foundry and finally to the OEM. Examples for solving vertical integration are as follows:

- Front to end simulation
- User-dependent reliability profiles

Horizontally: A smooth problem-free model and IP exchange is urgently needed and has enormous potential for European industry. Examples for solving horizontal integration are as follows:

- Technology-independent – design migration
- Unified design environment
- IP reuse

10.2.5.16 Design Targets

Traditional design targets such as performance and cost are being joined by more specific targets such as design for manufacturability (DfM) and testability. There is also a desire to parameterise these targets to be able to optimise a certain target at a high abstraction level. The most important targets are as follows:

Design for manufacturability: In the future, DfM will not only have to deal with simple yield but also with redundancy aspects. System redundancy will be a completely new challenge for academia and industry. Yield optimisation will be performed in the system design stage using frontend to backend simulation able to analyse and optimise the complete production process.

Design for testability: The existing DfT methodology must be extended to system aspects. Therefore, in the future BIST will also have a thermal, mechanical and electrical aspect.

Design for reliability: DfR can be handled using multi-physics simulation; however, root causes for errors such as cracks or other mechanical and electrical impacts must be analysed in more detail. Also models for the different failures must be generated, tested and implemented. Also different ‘reliability profiles’ must be generated.

10.2.5.17 Interaction with Equipment

Equipment and component engineering will be essential in the future for precise models and new components. The detailed knowledge of the impact of equipment and components will allow their influences on the final system to be eliminated. So ‘overcome dirty analogue’ will evolve to overcome ‘dirty environment and equipment’.

10.2.6 Conclusions

The difference between system design and component design is mainly defined by the system architecture and the related partitioning. If the system complexity is low, the design tools and methodologies are mainly the same as for ‘non-system designs’. Functional abstraction level is not critical, as the complexity is not really relevant. There is no need to introduce a high-level description language and high-level simulation tools.

For the typical More than Moore objects, mainly SiPs, which include different silicon technologies and 3D stack-ups, there is no way to design the systems at quasi transistor level. One or even more levels of abstraction have to be introduced using modelling.

EDA vendors have started to provide multi-level tooling but we are still far away from adequate system design environments. A significant problem is the proprietary nature of the tools. Algorithms are well-guarded secrets and only limited access is granted to interfaces. Standard formats for data exchange are often not fully implemented or require vendor-specific extensions making it difficult to mix tools from different vendors. It also makes it difficult to conduct research in the area without support from an EDA vendor.

The drastic rise of complexity together with an increasing number of technologies and applications makes a single-vendor approach questionable. The semiconductor industry has reached a point where only large industry consortia can bear the costs of technology development. In the More than Moore field, the focus is shifting towards the combination of different hardware technologies and software in systems, and the exploration of new architectures to realise functions more effectively.

The competitive advantage of processing knowledge is, therefore, being overtaken by the ability to effectively combine and use the technologies.

Taking control of the design tools must, therefore, be seen as an important strategic goal of the European semiconductor industry. Several projects, supported by all major European semiconductor manufacturers and a large number of universities and research institutes, are investigating and creating the foundation for an open platform on which to build future design flows. The overall goals of these projects are to create a platform that supports reuse of tools as well as IP over a large range of technologies, while providing a platform for researchers and SMEs for developing and showcasing their research and products.

Underlying Literature

1. ISSCC 2008/Session 13/Mobile Processing/31.1
2. ISSCC 2007/Session 29/Analog and Power Management Techniques/29.2
3. ISSCC 2008/Forum 3: Architectures and Circuit Techniques for Nanoscale RF CMOS
4. <http://www.design-reuse.com/>
5. K. Hollaus, O. Bíró, P. Caldera, G. Matzenauer, G. Paoli, K. Preis and B. Weiss, "Electromagnetic Field Computation of Simple Structures on Printed Circuit Boards by the Finite Element Method," *IEEE Transactions on Magnetics*, vol. 42, pp. 815–818, April 2006.
6. K. Hollaus, O. Bíró, P. Caldera, G. Matzenauer, G. Paoli and G. Plieschenegger, "Simulation of Crosstalk on Printed Circuit Boards by FDTD, FEM and a Circuit Model," *IEEE Transactions on Magnetics*, vol. 44, pp. 1486–1489, April 2008.
7. Semiconductor Industry Association. The International Technology Roadmap for Semiconductors, 1999 Edition. SEMATECH: Austin, TX, 1999.