Standards Development at Si2: A Synergistic, End-User Driven Approach

Sumit DasGupta
Senior Vice-President, Engineering
Silicon Integration Initiative (Si2)
Austin, TX 78759

I. Introduction:

The relentless evolution of semiconductor technology, in line with Moore's Law, continues to provide fresh challenges to the chip design process and which, technology pundits predict, will only worsen over time. They remind us of the ever-increasing need for:

1. Tightly integrated design flows with interoperating design and analysis tools that include accurate modeling of device and circuit behavior in the presence of large process variations,
2. "Plug-and-play" platforms, which are essential to create best-in-class flows, where tools from multiple sources may be integrated,
3. Design-for-manufacturing capabilities that allow semiconductor manufacturing process information to flow upstream into the design process and design intent to flow into manufacturing, and as part of late breaking news,
4. Enhanced capabilities in design tools and flows that can take into account more consistent power constraints to achieve much more efficient low-power or power-efficient chip designs than ever before.

It is generally agreed that these daunting challenges can best be resolved through an open, industry-supported process that defines such key infrastructure components of a design system as:

1. Open, syntactically and semantically well-defined API's and formats,

2. Clearly documented information model, along with a reference implementation, that allows the storage and sharing of design data, intent, and constraints, as well as,
3. Technology models to enable design and manufacturing despite the greater challenges of process variability.

This would allow EDA tools built in this environment, independent of its origins, to interoperate and enable the most advanced design methodologies for the chip designer. These same assumptions above serve as the basis for multiple parallel, focused, industry-supported projects, or coalitions, at Si2 to address the above challenges during their pre-competitive phase of development. While each coalition generally operates autonomously, there are often synergistic opportunities on topics of mutual interest.

II. OpenAccess Coalition

The longest-running project at Si2, the OpenAccess Coalition (OAC) is responsible for the continuing development and evolution of OpenAccess. In a short span of time, OpenAccess has become a well-received infrastructure standard (with a rich and supportive reference implementation) that enables "plug-and-play" interoperability of EDA tools, and whose components are depicted in Figure 1. It allows the opportunity of creating design flows that integrate best-in-class design tools from vendors or internal CAD groups without the severe penalties historically associated with creating and maintaining data translators between steps in the design flow.

The OpenAccess Roadmap, representing short and long-term directions, is managed by this coalition. The OAC consists of representatives from the end-user, semiconductor and EDA industries and currently has 35 members. They elect the

representatives of the OpenAccess Change Team (CT) from among the OAC

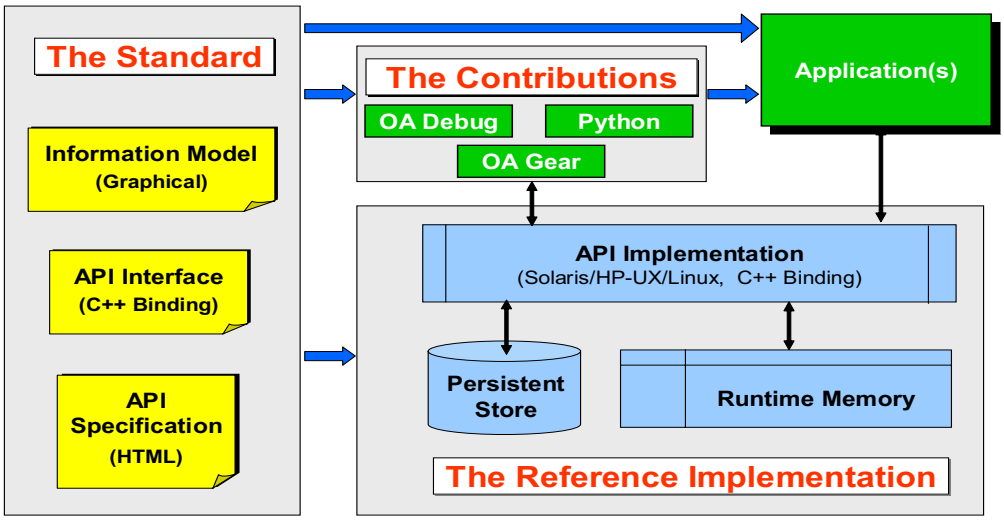

Figure 1: OpenAccess Components

members. The CT defines and approves both near-term and long-term goals that are embodied in the Roadmap document that is available to all through the Si2 web-sites. Some of its newest features are described in [1].

The Change Team sponsors working groups (WG) specifically targeted to recommend solutions for specific problems. WG's currently underway include the Debug WG that is driving the development of a much-desired database debugger to support both EDA developers and chip designers, Contributions Process WG that is defining the infrastructural needs to catalyze greater amount of contributions from the community and the Joint Data Model WG jointly staffed by the OAC and the Open Modeling Coalition, details of which are deferred to the next section.

Perhaps the most exciting recent development within this coalition is the clarification in the definition of the API into what is defined as: "core", the design of which requires more thorough reviews, implementation and regression testing by the integrator; and, "non-core" which does not require as much rigor and hence can be

contributed by members of the Si2 community without the support of the integrator. Most of these enhancements would exploit, as a first step, "extension objects" which are one of the key innovations in OpenAccess. It is expected that this new strategy and structure would be a key enabler to even faster growth and adoption of OpenAccess, thus further confirming its status as an industry-accepted standard platform.

III. Open Modeling Coalition

The Open Modeling Coalition (OMC) [2] was begun in response to a request from several members of Si2 who design leading-edge processor or ASIC chips. The challenges of process variability mentioned before were affecting their ability to create accurate library models to support detailed timing, noise and other signal integrity analyses. OMC was formed to define an infrastructure with open, standard interfaces and a flow to create robust library models for sub-65 nm process nodes.

In response to the above requirement from leading chip makers, OMC has taken a comprehensive approach to the problem of library modeling. It has defined an architecture, that is posted on the Si2 web-site, to accommodate the need for both static, table-based modeling based on extensions to the Liberty format and a dynamic modeling approach based on the IEEE1481 standard. To support the former approach, there exists a WG that continues to evolve extensions to the Cadence-donated Effective Current Source Model (ECSM) [3] as Si2 standards. Recent extensions cover power and noise; the extension to support statistical timing is now under development jointly between the Statistical WG and the ECSM Change Management Group (CMG). Parallel to the work of the OMC, the Liberty Technical Advisory Board (LTAB) continues to advance the Composite Current Source (CCS) model. There is growing interest among user and IP companies in converging both static formats into a single standard.

The dynamic approach is based on the use of executable rules written in the Delay Calculation language (DCL) and is intended to cover modeling of library

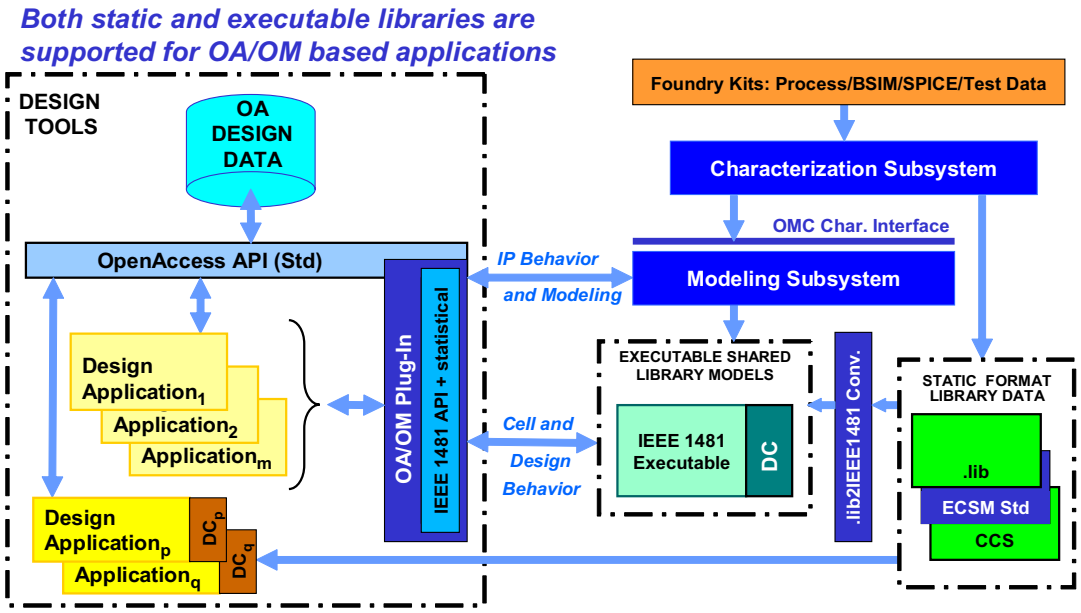

Figure 2: Open Modeling Architecture

elements like simple NAND gates all the way to IP blocks like microprocessor cores. The key feature of the dynamic approach is the ability to conveniently utilize a single delay calculator for all applications. A common characterization interface allows either approach above to utilize a common solution for circuit characterization. Figure 2 provides a high-level view of the OMC-proposed architecture for delay calculation.

One important activity is occurring in the Joint Data Model (JDM) WG which is staffed by members of the OAC and OMC. The purpose of this WG is to define a plug-in to connect the OpenAccess API with the IEEE1481 API so that design information such as, fan-in/fan-out, parasitic data, etc, can be transferred to the delay calculator and results of the delay calculation can be fed back to the requesting client application. This is an example of the synergy among the coalitions at Si2.

IV. Design-to-Manufacturing Coalition

Problems associated with manufacturing chips in nanometer technologies (90nm and below) have been well-documented in the past few years. The International Technology Roadmap for Semiconductors (ITRS) predicted this pattern and continues to warn of impending problems. These problems arise out of many factors, such as:

1. The greater complexity that comes naturally from the larger designs,
2. The lack of ability to inject greater knowledge of design intent to distinguish between different parts of design data based on its criticality,
3. The significantly greater variability in the underlying manufacturing process and the increasing effect of minor variations on the parametric properties of the miniaturized features,
4. The sequential nature by which design data is processed in manufacturing and, last but not the least,
5. The incomplete means by which process information is fed into the design process.

Point solutions have been suggested based on several recently developed ideas, such as, Reticle Enhancement Technology (RET), Critical Area Analysis and Statistical Analysis for random and parametric yield issues affecting timing, power, etc, to name just a couple.

The vision of the DTMC [4] is to define an open-standard IT infrastructure that will provide the means for enhanced communication of information across the IC supply chain and against which design and manufacturing applications can be integrated. This will provide a basis for more complete communication of information to improve design and manufacturing cycle times, more effective integration of the entire flow, larger choice of applications with less difficulty to insert them into the flow, and faster introduction of new design technology and transfer into production.

Given the nascent nature of the industry and the wide swath it cuts, including design companies, IDM's, fabless semiconductor companies, IP companies, mask makers,

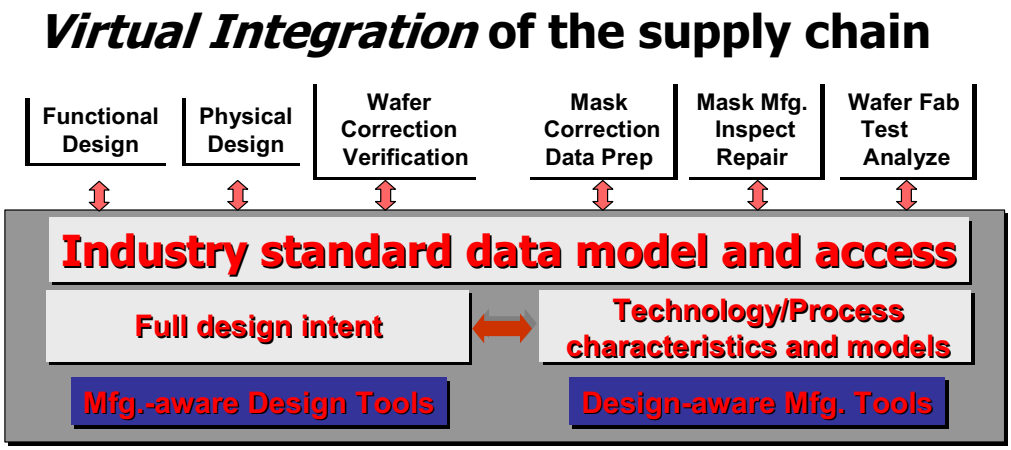

Figure 3: Design and Manufacturing Infrastructure

EDA companies and foundries, achieving consensus on a prioritized set of requirements is a foundational prerequisite. Thus, the DTMC has played the role of an industry voice, organizing open workshops and meetings as a starting point, to gather these necessary players together to collect a consistent set of requirements that span across the entire problem space. Figure 3 above provides a conceptual view of the desired environment where joint work with the OpenAccess Coalition to extend OpenAccess to satisfy the above architecture is a very enticing opportunity.

The first priority towards the DTMC vision now is to provide a comprehensive architecture of standards in support of design-for-manufacturing (DFM) models. Model-based (statistical) design has become paramount for designs at leading edge technology nodes due to the increased significance of manufacturing variations caused by the shrink in design feature size. The effect of design on manufacturability and of manufacturing process effects on design is now best analyzed through models of key manufacturing process steps. This need spans the entire design flow from early layout through tape-out. This requires that both EDA suppliers and foundries support this model-based approach. In addition, designers require portability and

choice of EDA tools and libraries, and this begs for the use of an open and common infrastructure across foundries and EDA suppliers. Foundries also wish to see common solutions for model-based design across all EDA to improve their business efficiency. However, the overall solution must take into account the requirement that each foundry, each EDA company, and each user must have the ability to protect its IP and also differentiate itself from its competition even as it endeavors to make the overall process more open, efficient, and streamlined.

Given the above conflicting business motivators, the DTMC has proposed an architecture for secure, black-box DFM models accessible only through a well-defined interface by any EDA application complying to that interface. The actual behavior of the models (the formulas, decision logic, etc. that represent the actual process behavior in software) may be completely hidden as required by their owners. Thus, models need be written only once by their supplier to be usable in an array of EDA tools, and an EDA tool need be written to only one common and open API to use models from an array of suppliers. Lastly, models can be supplied by foundries, library providers, EDA vendors, third parties, university researchers or the designer.

To complete this solution, the DTMC is also addressing a number of additional focus areas. These include: a taxonomy of DFM model classes, a dictionary of terms defining manufacturing characteristic parameters for each class of models, a file format for the exchange of manufacturing characteristic parameters for all classes of models, a standard for lithography measurement and verification rules, and an open library of model calibration tests usable to test and qualify these DFM models.

IV. Low Power Coalition

The newest of the Si2 coalitions, the Low Power Coalition (LPC) [5] is involved in the development of an infrastructure and flow to support the design of low-power or

power-constrained chips. This includes an open API and format to express power constraints, parsers, training materials, sample code and other enablers that are necessary for rapid adoption of the developed solution. From initial thoughts to a full-fledged coalition took less than six months. This underscores the immediacy of the need. The urgency of the challenge can best be expressed using the example in Figure 4 below obtained from the most recent edition of the ITRS which suggests that over the next decade while performance of power-efficient SoC's is expected to rise by ~500x, power consumption in these chips may only rise by ~60%. This fact ends the progress of Moore's law unless major advances in low power design are achieved. The LPC was formed by companies that understand that fact and intend to act on it. The coalition now has 18 very active members who are driving the

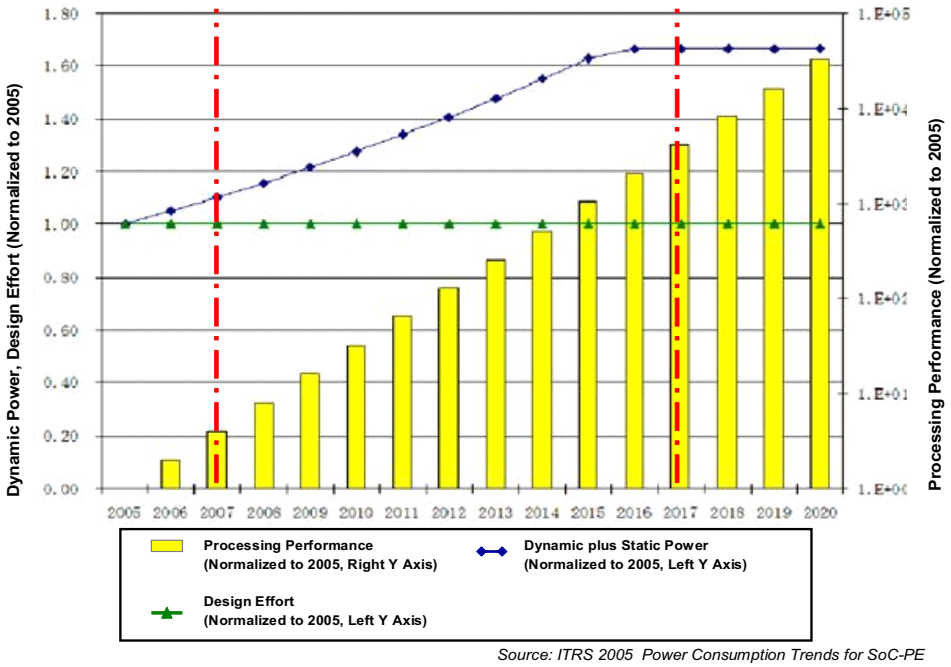

Source: ITRS 2005 Power Consumption Trends for SoC-PE

Figure 4: Power Consumption Projections

development of an ecosystem to satisfy the needs of the coalition members and the industry at large to solve the needs for low power design. Like all other coalitions,

there is a management group at the coalition level that provides overall business management while a Technical Steering Group (TSG), consisting of technical experts, is responsible for defining the roadmap for low power design improvements. The TSG will sponsor domain-specific WG's to focus on solving detailed technical issues.

The solution being pursued by the LPC is based on an initial technology contribution from Cadence called Common Power Format (CPF) version 1.0 which defines a file that specifies constraints and associated semantics to drive power requirements into a chip design flow from RTL to GDS. Activities to-date have included detailed training and review of the CPF 1.0 specification as a first step to further evolution and growth. CPF version 1.0 has been released by the LPC as an Si2 standard and in parallel, a WG has been underway to do a detailed comparison between CPF and UPF, the power specification developed under the sponsorship of Accellera. Results of this comparison are expected to be published shortly, hopefully as a catalyst for sharing information between these two formats. Looking ahead, additional WG's are soon to be convened to address several important topics. These include:

1. Refining the use model of the existing CPF version
2. Separating out the sections of the standard that refer to library constraints and semantics so that they can properly be applied to the Liberty format and both of its extensions,
3. Partnering with the OAC to determine potential extensions to the OpenAccess API and information model to support low power design, and last but not the least,
4. Extending CPF into areas that are not currently covered, such as, system-level low power design, and, analog low power design

V. Conclusions

This paper has described four major efforts underway at Si2. Each is at a different stage in its life cycle. Each coalition has its own structure, membership and goals which allow each sufficient autonomy within the overall structure of Si2 and yet there are sufficient areas of mutual interest which lead to joint WG's that connect one coalition with one or more of the others. However, they all share one common overarching principle, that standardization is based on well-defined and documented interfaces, formats, and flows that are supported by reference implementations, training materials, and other related documents. All of these products are designed to solve customer problems and the ultimate success of these solutions is measured by adoption by their customers and by their proliferation throughout the foundry, IP, EDA, and user companies.

References

[1] "Qualifying Applications Against Data Model 3", M. Guiney, Cadence Design Systems, 9th OpenAccess+ Conference, November, 2006[2] "Open Modeling (OMC) Overview", Timothy Ehrler, AMD, 9th OpenAccess+ Conference, November, 2006

[3] "ECSM: Complete Current Source Modeling for Sub-nanometer Geometries", V. Kariat, Cadence Design Systems, 9th OpenAccess+ Conference, November, 2006

[4] "FSA and Si2 Open DFM Workshop Summary", D. R. Cottrell, Si2, FSA and Si2 Open DFM Workshop #2, November 2006

[5] "Low Power Coalition @ Si2", S. DasGupta, Si2, 9th OpenAccess+ Conference, November, 2006