The X initiative model: Collaborative resolution of 'shared red bricks' in ITRS

Aki Fujimura and Andrew B. Kahng describe the X Architecture on-chip interconnect and how it is being used to address design for manufacture.

The International Technology Roadmap for Semiconductors (ITRS) clearly indicates that nanometer-scale technology challenges demand unprecedented cooperation and information sharing across the semiconductor industry supply chain. The roadmap specifies a rapidly growing number of technology requirements – known as 'red bricks' due to their coloring in roadmap charts—for which there is no known solution. Only through shared efforts can the semiconductor and semiconductor supplier industries solve these tough technical challenges and enable continued advancement and innovation along the Moore's Law curve.

One industry consortium – the X Initiative – has brought together companies from every segment of the semiconductor supply chain to jointly address design-for-manufacturing (DFM) issues for a novel on-chip interconnect architecture known as the X Architecture. In doing so, the X Initiative has pioneered a new model for the kind of open collaboration that is increasingly necessary to resolve ITRS red bricks.

In addition to predicting the key industry trends that will emerge over the next 15 years, the ITRS identifies the technical capabilities that need to be developed for the industry to continue on the path of Moore's Law (i.e., the number of transistors on an IC doubles every 18-24 months).

Figure 1, from the 2003 update of the ITRS, illustrates the aggressive acceleration in gate length anticipated over the roadmap's 15-year span. Both the printed and physical gate-length trends are forecast to continue scaling forward by about 0.7x per three-year cycle, from this year's 90-nm physical gate length for high-performance logic devices.

Out of this exhaustive roadmapping effort have arisen two key findings: a) more aggressive scaling will be needed if Moore's Law is to continue; and b) the fundamental limits of the materials used in the traditional planar CMOS process are at the 15-year horizon, giving rise to the need to address post-CMOS device requirements as well. It should come as no surprise, given the scope of the roadmap and the rapid rate of change within the industry, that the ITRS now identifies an increasing number of areas where no solutions currently exist (red bricks). In fact, the number of these red bricks has grown to such an extent that the industry appears to be on a collision course with a 'red brick wall' in the not-too-distant future.

Solving any red brick is expensive and time-consuming, and requires significant R&D investment. Yet, although the ITRS is a shared roadmap, red bricks are typically vetted by individual companies or supplier segments. It is very likely that many red bricks actually stem from companies trying to pursue old processes or follow old trends, rather than seeking synergy with other links in the semiconductor supply chain.

The effect is analogous to automotive suppliers who are specifying requirements and creating new technologies in a vacuum, as if super tires, seats or engines alone could drive development of the next highly advanced car.

Fig 1: ITRS-projected gate-length trends

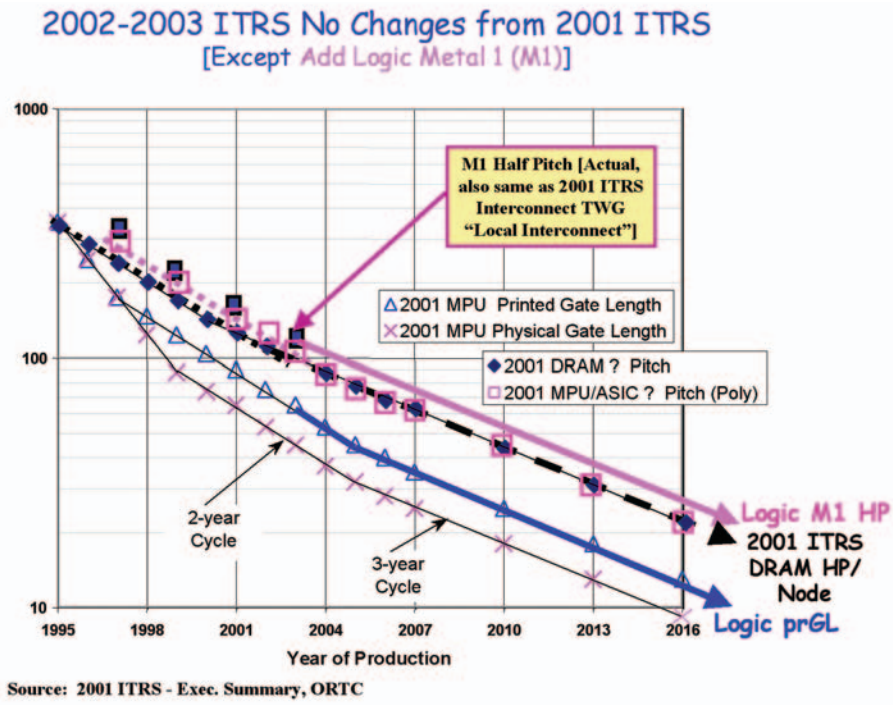

The sheer number and complexity of red bricks in the ITRS mandate new approaches predicated on new, open collaborative strategies that are focused on solving what are inherently global optimization and risk management issues. Not only is it economically wasteful and a misappropriation of resources, but also it is technologically impossible for each supplier segment to attempt to continue Moore's Law all by itself.

Going forward, a more globally optimized allocation of R&D investments – in other words, a shared approach to solving red bricks – will be critical. The success of such efforts will hinge on the formation of deeper partnerships within and between each of the semiconductor supplier groups and technology industries represented in the ITRS.

Case in point – interconnect performance. The interconnect industry has traditionally tried to solve this issue with lower-resistivity interconnect materials, lower-permittivity dielectrics, thinner barrier layers, reduced dishing in CMP, etc. The X Initiative brought enabling layout architecture to the table, and then backed it up with an open collaboration within a broad consortium of design and manufacturing technology providers.

The X Initiative is an excellent example of a collaboration strategy in action, working to prove the benefits and support the adoption of a new approach to interconnect design – the X Architecture (see box below).

A collaboration model

The X Initiative anticipated the need for a collaborative effort that would span the entire supply chain – in this case, to bridge the disparate worlds of electronic design, mask fabrication, lithography and chip production – to accelerate the commercial availability of X Architecture chips. X Initiative members are working toward a set of shared solutions, co-developed and co-verified by the entire consortium.

The X Initiative was formed in 2001 in recognition of the fact that widespread adoption of the new architecture within the semiconductor industry mandated a collaborative effort to substantiate its manufacturability. It's becoming increasingly clear that when a major semiconductor process or design innovation is introduced, cooperation throughout the supply chain is required to ensure manufacturability and economic viability. This is a lesson learned from the efforts to commercialize some process and material advancements (e.g., low-k dielectrics, SOI), which might have been adopted more quickly through broader collaboration.

The greatest issues with any new chip technology include: the availability of required design and verification tools; photomask production and costs; the management of data requirements and data volume through the supply chain; and its impact on yield.

To address these concerns for the X Architecture, the X Initiative's charter is to educate the supply chain about the architecture, accelerate fabrication of X Architecture chips, and track and promote the technology's commercial proliferation. The X Initiative is thus focused on advancing the usage of the X Architecture by ensuring its support throughout the design and manufacturing cycle.

Looking at the ITRS in light of this charter, one can identify several shared red bricks that the X Initiative's efforts may help to resolve, including:

- Interconnect

Total interconnect length for active wiring
Local wiring pitch
Local A/R for copper

- Design

Dynamic power reduction (beyond scaling)
Standby power reduction (beyond scaling)

The results of X Initiative members' work to date show that a number of shared chal-

About the X Architecture

In the last five years, the dominant factor affecting chip timing and reliability has shifted from the transistor to the interconnect – i.e., the wiring that connects the devices. As design complexity increases in accordance with Moore's Law, the amount of interconnect rises concurrently.

The problem is that as more wires are packed onto the die, the smaller the conductor dimensions must become. Thus, the wire becomes more resistive, and approaches to combat this trend by increasing the wire height-to-width aspect ratio lead to greater coupling capacitance and delay variability. Increased signal delay leads to slower chips that consume more power. Moreover, because manufacturers can't fit as many chips onto a wafer, time to market and device costs soar.

The X Architecture is an innovative chip-design approach that proposes to help resolve these problems by introducing pervasive diagonal routing into advanced devices.

Prior to the development of the X Architecture, recent interconnect roadmapping has focused on process technologies and materials, with the manufacturing community aggressively pursuing copper, low-k dielectrics, phase shifting (PSM), optical proximity correction (OPC) and silicon-on-insulator (SOI), along with dummy fill and CMP models to control planarization and resistance variation.

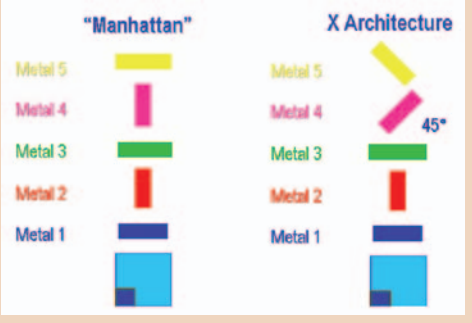

While these innovations have indeed enabled significant gains in device speed, performance, cost and area, the X Architecture's creators (innovators with Simplex Solutions, now Cadence) believed that if these materials and processes were going to deliver their full potential, physical design technology aimed at the interconnect performance gap needed to change as well.

Based on new algorithms that leverage the power of today's sophisticated computing resources, the X Architecture targets chips with five or more metal layers. With this new design scheme, the primary direction of the interconnect in the fourth and fifth metal layers is rotated by 45 degrees in relation to the conventional orthogonal, or "Manhattan," architecture.

Unlike Manhattan-based routing, X Architecture designs include routing in any of eight directions—allowing more direct connection between any two transistors on a chip. This illustration shows the 45-degree angles employed in designs created using the X Architecture, which enables both large- and small-scale diagonals.

The promise of the X Architecture is that by employing diagonal routing in a pervasive manner, designers can reduce wire length by an average of 20 percent and vias by 30 percent—simultaneously improving chip speed, power and cost, enhancing signal integrity and reliability, and increasing not only the probability of successful first silicon, but the manufacturing yield as well.

Fig 2: Design to manufacture flow

lenges are on the road to resolution due to this collaborative approach.

For the past two years, X Initiative members, who now number more than 40, have worked together to make major strides toward bringing the X Architecture into the realm of mainstream manufacturing. The consortium's key focus has been to demonstrate the manufacturability of X Architecture-based designs by helping bridge the gap that has traditionally existed between design and manufacturing. These efforts have enabled the semiconductor industry to see the value of addressing manufacturing challenges early in the process of new technology introduction.

This commitment was important right from the beginning because the introduction of the architecture and the initiative met not only with a great deal of interest, but also with skepticism from some quarters, as is often true of new technologies. X Initiative members have repeatedly encountered certain key myths and misperceptions associated with the implementation of this new design architecture. As table 1 illustrates, collaboration has been the key to clearing up these misperceptions through jointly delivered results.

Since its inception in June 2001, the consortium has made steady and significant progress. The following summarizes the key milestones achieved thus far, each of which has helped advance the X Architecture and illustrate the value of the X Initiative's collaborative approach.

Table 1: X Architecture manufacturability – open collaboration separates myth from fact

X Architecture mask costs will be prohibitively high.	**Based on results to date, neither DuPont Photomasks nor Dai Nippon Printing (DNP) sees any reason for X Architecture masks to cost more than the Manhattan equivalent. Mask writers such as Etec's raster-based laser system and NuFlare's vector e-beam with triangular aperture eliminate sensitivity to diagonals. Thus, there is no impact to write time or cost.**
Diagonal interconnect will cause an unacceptably high level of false defects in mask/reticle inspections.	**A 130-nm mask was inspected in September 2002 using KLA-Tencor's TeraStarTM reticle inspection system, which employs a completely new algorithm for die-to-database inspection. TeraStar handled the diagonal structures without identifying any false defects.**
Long, pervasive diagonals will create unacceptable lithography process windows and critical dimension (CD) variability.	**In wafer lithography experiments to date, conducted jointly by multiple X Initiative members, process windows and CD uniformity have been more than acceptable for both dense and isolated diagonals.**
Design rules for diagonals must be much looser than those for orthogonal (Manhattan) interconnect.	**The design-rule test chips created by X Initiative members at 130 nm and 90 nm have established that design rules for diagonals are comparable to those for Manhattan architectures.**
New materials such as copper and low-k dielectrics will elevate defects and CD uniformity problems when diagonal routing is employed.	**Applied Materials produced the first 90-nm and 65-nm test chips, validating the manufacturability of copper/low-k chips using the X Architecture.**

Key X Initiative Milestones

October 2001 – DuPont Photomasks created the first X Architecture photomask, in collaboration with other X Initiative members and using existing mask-manufacturing software and equipment, including: 180-nm design data from Cadence; the ALTA 3700 raster-based advanced laser mask-pattern-generation system from Applied Materials' Etec Systems division; Zygo Corporation's AutoKMS CD metrology system; and KLA-Tencor's 363 UV mask inspection system.

The purpose of the experiment was to determine a) that existing photomask equipment could be employed to create X Architecture masks, and b) the impact of the new architecture on data sizes, figure counts, write times, CD metrology and reticle inspection. While the diagonal features of the X Architecture design data did create more complicated figures, write times using the Etec equipment were not materially affected because laser systems are not sensitive to geometrical shapes, and CD measurements were easily attained using the system. The X Architecture masks were inspected using the KLA-Tencor 363 UV system with a 0.186-micron pixel. Both die-to-die and die-to-database modes were tested, using the highest sensitivity setting. The die-to-die inspection ran very well, even at this highest sensitivity. The die-to-database inspection, not unexpectedly, flagged some false defects—a common occurrence with new technologies, and one that was resolved in later experiments using KLA-Tencor's TeraStar reticle inspection system (see Sept. 2002).

February 2002 – Toshiba presented details of the first design implemented in the X Architecture – a 200MHz RISC processor core featuring approximately 750,000 random logic gates, as well as several SRAM and custom blocks, and targeted to 0.18-micron CMOS processes. Initial results indicated an overall 20-percent wire-length reduction and 10-percent area savings compared with conventional Manhattan orthogonal routing. Static timing analysis confirmed 20-percent performance improvements in all blocks of the design.

March 2002 – DNP, in conjunction with other X Initiative members, demonstrated the viability of vector-based mask-pattern generation equipment for producing X Architecture photomasks. DNP employed equipment from Toshiba Machines (now NuFlare Technologies) to successfully write an X Architecture test mask without significant runtime or data volume increases.

This work served to confirm a new supply chain for X Architecture photomasks.

April 2002 – ASML confirmed the manufacturability of the X Architecture, simulating lithography performance of 0.18-micron X Architecture design data and successfully completing proof-of-concept wafer exposures of diagonally oriented 0.25-micron interconnect structures based on 0.18-micron design rules.

These first-ever wafer results indicated that interconnect layers could be successfully made using the X Architecture approach, and marked a milestone in demonstrating the ability to manufacture X Architecture interconnect designs.

September 2002 – X Initiative members collaborated to deliver the first 130-nm X Architecture masks. KLA-Tencor completed characterization of its TeraStar reticle inspection system using a DuPont Photomasks-created mask and a Cadence-supplied X Architecture 130-nm (0.13-micron) design. Both die-to-die and die-to-database modes were tested, using the highest sensitivity. Both inspection modes ran very well, and the inspection results with the X Architecture design were comparable to Manhattan designs at the 130-nm process node.

December 2002 – X Initiative members collaborated with Nikon Corporation to validate a new design-to-wafer supply chain at 130 nm for the X Architecture. Nikon employed its NSR-S205C step-and-repeat exposure system, using a krypton fluoride (KrF) 248-nm excimer laser, to process X Architecture wafers, and validated that the results met the company's manufacturing criteria.

February 2003 – STMicroelectronics delivered first silicon results by producing the first 130-nm design-rule test chip based on the X Architecture. This work established design rules for the X Architecture equivalent to those for Manhattan designs.

June 2003 – Applied Materials produced the first 90-nm design-rule test chip for X Architecture interconnect designs. The fabricated test chip validated design rules and the manufacturability of X Architecture interconnect designs for copper/low-k chips using existing maskmaking and wafer-processing technologies, paving the way for use of X Architecture designs in manufacturing advanced devices on 300-mm wafers.

October 2003 – Toshiba Corporation completed the X Initiative's pre-production roadmap by producing the first functional silicon for the X Architecture at the 90-nm process node.

December 2003 – World-leading semiconductor foundry UMC announced that it is ready to accept X Architecture designs for fabrication at the 180-nm, 150-nm and 130-nm process nodes.

February 2004 – Applied Materials, Cadence Design Systems and Canon collaborated to produce the first 65-nm design-rule test chip for X Architecture, demonstrating the scalability of this innovative architecture into future process nodes. In addition, Infineon Technologies joined the X Initiative, and validated its X Architecture manufacturing readiness with the successful fabrication of a 130-nm test chip.

May 2004 – Taiwan's TSMC joined the X Initiative and produced a 130-nm test chip. The foundry announced that it is working with select customers to leverage the performance, cost and power advantages of the X Architecture.

June 2004 – In collaborating with Cadence Design Systems and members of the X Initiative, Toshiba Corp. launched the industry's first commercial SoC built on the X Architecture. Toshiba's new TC90400XBG chip is designed for digital-media and home-entertainment applications. Toshiba's first customer for this chip will integrate the device into digital TVs for the European market.

The cumulative result of these and other ongoing collaborations is that the manufacturability of this new design technology has been established early in its development cycle – preparing the entire supply chain to design and manufacture these new chips. The findings show a clear and logical progression of the X Initiative members' work, and as the X Architecture technology is commercialized in 2004, this collaboration will not only act as a catalyst to drive ongoing advancements, but will also serve as a model of what the industry can accomplish when companies work together to overcome mutual problems and challenges.

Approaches like this exemplify innovation in efforts to speed resolution on ITRS shared red brick challenges, without creating new technical or manufacturing challenges that could slow time to market, hinder productivity and yield, or diminish fab ROI and profits. The X Initiative is a collaborative entity with collective product-development and marketing muscle far beyond that of any of its individual members – resulting in new roadmap - advancement opportunities for the members, their customers and the industry as a whole.

Aki Fujimura, X Initiative Steering Group Member is CTO, New Business Incubation, Cadence Design Systems & Andrew B. Kahng, is from the University of California, San Diego, Dept. of Computer Science and Engineering and Advisor to the X Initiative.

Fig 3: X Initiative timeline