Technology-migratable ASIC library design

by T. R. Bednar
R. A. Piro
D. W. Stout
L. Wissel
P. S. Zuchowski

A library strategy has been developed to enable IBM Microelectronics ASIC development to keep pace with rapid technology enhancements and to offer leading-edge performance to ASIC customers. Library elements are designed using migratable design rules to allow designs to be reused in future advanced technologies; and library contents, design methodology, test methodology, and packaging offerings for the ASICs also are consistent between current and future technologies. The benefit to the ASIC customer is an ASIC with a rich library of logic functions, arrays, and I/Os for today's designs, and with a ready migration path into future designs.

Introduction

Before 1993, IBM ASICs were available only to internal customers. The ASIC libraries were usually jointly defined and developed by IBM in Burlington and by the various IBM customer organizations, resulting in libraries without a clearly defined strategy across technology generations. Library development effort was repeated from one technology generation to another, and, from the customer perspective, ASIC designs could not readily be migrated from one technology generation to another.

In 1993, the IBM Microelectronics Division began marketing ASICs to customers outside IBM. The requirements of this market have stimulated a new strategy for creating ASIC libraries that emphasizes consistency with industry standards, and efficient migration of customer designs to future generations of technology. This standardization allows the ASIC libraries themselves to be migrated from a 0.4-μm technology forward into 0.35-μm and 0.25-μm technologies and beyond; customers can migrate ASIC designs into advanced technologies with minimal redesign effort, thereby reducing design cycle times.

This paper focuses on the evolution of the ASIC library design strategy to meet customer expectations for ease of use and early access to new technologies. Customer access to a technology is limited by the availability of an ASIC library in a new technology. By using a migration strategy which emphasizes reuse of library elements, the ASIC developer can deliver new technologies to customers with minimized time and resources.

Customer expectations

To be competitive in today's market, an ASIC product must provide the customer a rich library of logic functions, input and output circuits (I/Os) supporting multiple interface standards, memory arrays, and embedded high-level special functions. All of these elements must be supported with models which allow logic synthesis and simulation, static timing analysis, test vector generation, and physical design to be performed on multiple tool platforms. IBM ASICs support multiple vendor design tool platforms, as well as IBM EDA tools used for in-house

0018-8646/96/$5.00 © 1996 IBM

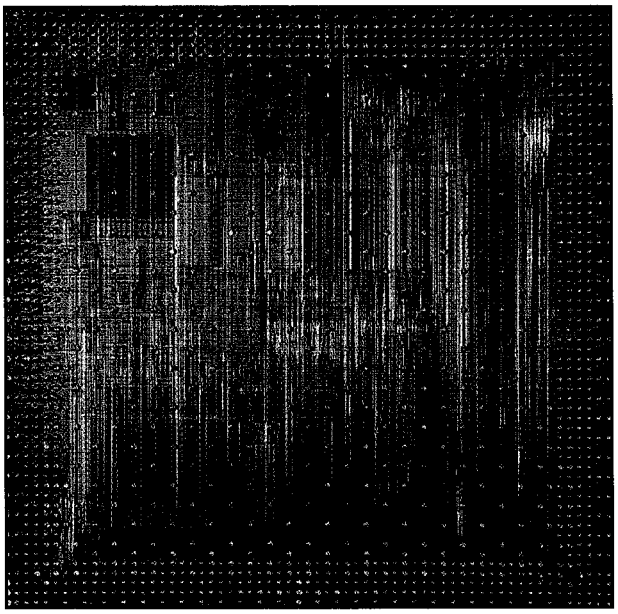

Figure 1

Photograph of 14.5-mm CMOS 5X ASIC.

designs. The models for the library are delivered to the customer in the form of an ASIC design kit.

Customers expect a robust test strategy that guarantees high-quality delivered parts. The IBM ASIC full-scan design-for-test methodology includes built-in self-test (BIST) for array test and reduced-pin-count testers with boundary scan [1–3]. Full-scan design, BIST, and boundary scan drive a large set of requirements in the ASIC design methodology; the methodology implications are described elsewhere in this issue [4]. They also drive requirements into the ASIC libraries for functions to support the test strategy.

The ASIC library and all other aspects of ASIC design (design methodology and tools; test strategy; ASIC packaging; product data book) must have a commonality between implementations in existing technology and in future technology. Long-term relationships between customers and ASIC suppliers are desirable to both customer and supplier not only for the efficiency of learning the interaction process, but also for the possibilities of creating even higher levels of performance in the ASIC design as customer and supplier learn to jointly push the limits of ASIC design.

Customers embrace the need for a certain "overhead" in their ASIC designs for test, and readily accept the burdens of design-for-test. However, they expect to be able to apply to successive designs the learning they have achieved in design-for-test during their first ASIC design with the supplier.

An ASIC in a future technology may require a lower power supply because of the finer features of the new technology. Customers will benefit from the lower power dissipation of the reduced supply, but they require the ability to continue to interface their designs with other electronic components. The ASICs must be able to comply with these interface levels.

In summary, customers expect an ASIC design kit with a rich variety of functions. They expect support for performing their design using industry-standard design tools they have previously used. ASIC users will usually accept unique methodologies from the ASIC supplier in areas such as design-for-test, but expect the requirements to be stable during the period of interaction with the supplier. And they expect the supplier to offer technologies and libraries with performance that allows their chip designs to be competitive in the marketplace, with these technologies and libraries allowing migration of the ASIC designs and the investment of design expertise in the future.

Technology features

The present generation of IBM ASIC products were designed in the IBM 0.4-μm CMOS 5S and 0.35-μm CMOS 5X technologies, with anticipated migration to 0.25-μm CMOS 6S technology and beyond [5]. The CMOS 5S and CMOS 5X technologies share a common interconnect process for up to five levels of metal wiring, but CMOS 5X introduces a more aggressive device for improved performance. Both technologies include silicided diffusions and polysilicon for reduced device resistance and a local interconnect level. The local interconnect allows direct connections between both polarities of device diffusions and the gate polysilicon without requiring a separate contact level. This level improves the density of circuit layouts, particularly complex logic functions or storage elements with cross-coupled inverter pairs (SRAMs, registers, and latches). The technologies also include a diffused resistor, which is used for impedance control in the output stage of I/O pad drivers.

The layout design rules for the CMOS 5S and CMOS 5X technologies are not identical. However, a common set of migratable design rules (MDRs) were defined which encompass the layout restrictions of both technologies. Circuits laid out to the MDRs are then transportable between the technologies without modification.

Logic cells from the ASIC library are constrained to have a width which is an even multiple of the fundamental "track" unit of 1.8 μm; customizing the cells to the unique nonmigratable design rules of CMOS 5X only rarely allows a cell to be a full two tracks smaller in CMOS 5X than its implementation in MDRs. The potential

performance benefit in customizing CMOS 5X layouts to nonmigratable rules is seldom more than one percent. In arrays, where a leaf cell is repeated many times, the use of nonmigratable rules is of modest benefit in area and performance; however, the time-to-market benefit to the ASICs library of using MDRs far outweighs the benefits of customization to nonmigratable rules.

The common features of the CMOS process technology, and the use of MDRs, allow layouts for CMOS 6 and beyond to be created from CMOS 5 layouts by scaling down the lithographic dimensions.

Die template and package features

The ASIC die template offerings include 17 die sizes with wirable circuit counts up to 1.6 million gates. Die sizes range from 3.5 mm to 18.2 mm [6]. Each die template features a peripheral I/O ring and uses the IBM patented flip-chip technology for attachment to one of the following package types: ceramic quad flat pack; ceramic ball grid array and ceramic column grid array; and tape ball grid array. Packaged signal pin counts of up to 748 are possible.

In addition to a standard single-power-supply template for on-die power distribution, a second dual-supply template is also offered with either the ceramic ball grid array or ceramic column grid array packages. The design of the dual-supply template allows the designer complete freedom in defining each signal pin location as being associated with either supply voltage. This approach lends itself readily for technology migration where changes in supply voltages are common. For example, the CMOS 5X process has a nominal supply voltage of 2.5 V. A design could utilize the dual-supply template to allow communications with other designs using a 2.5-V supply and designs using the existing 3.3-V standard. Other applications of the dual-supply template include the use of small-swing differential I/O buffers for high-frequency applications.

The die is attached to the package with solder bumps. These bumps are arranged on the die in a regular pattern which is dense near the I/O cells at the die perimeter, and sparse over the gate-array area (**Figure 1**). At the die edge, bumps are used both for signals and power; centralized over the gate-array cells, the bumps are used only for power. The top metal layer is used to route wide lines (pad transfer wiring) from the bumps to the I/O cells, and is also used for power buses over the gate array. The bump footprint and pad transfer wiring were developed in conjunction with the packaging designers to create a solution that provided the best signal escape for all packages offered in the menu.

The I/O cells are 256 tracks long by 32 tracks wide. The minimum pitch for solder bumps is 125 tracks; however, a pitch of 128 tracks is used to make the bump pitch an

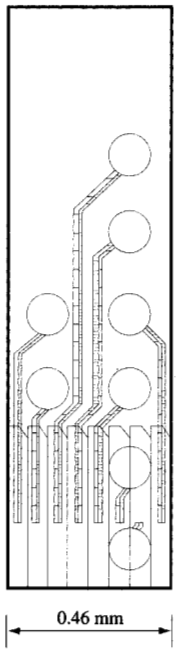

Figure 2

A pattern cell of eight I/O cells, their associated solder bumps, and top-layer interconnect.

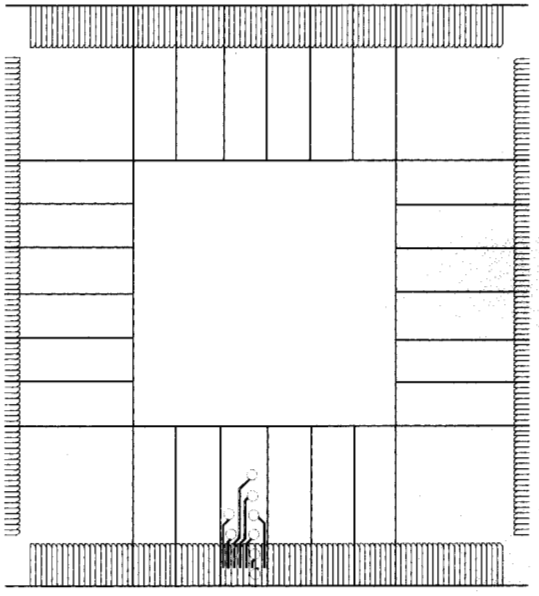

Figure 3

Construction of die image using repeated instances of Figure 2 pattern cells plus corner cells.

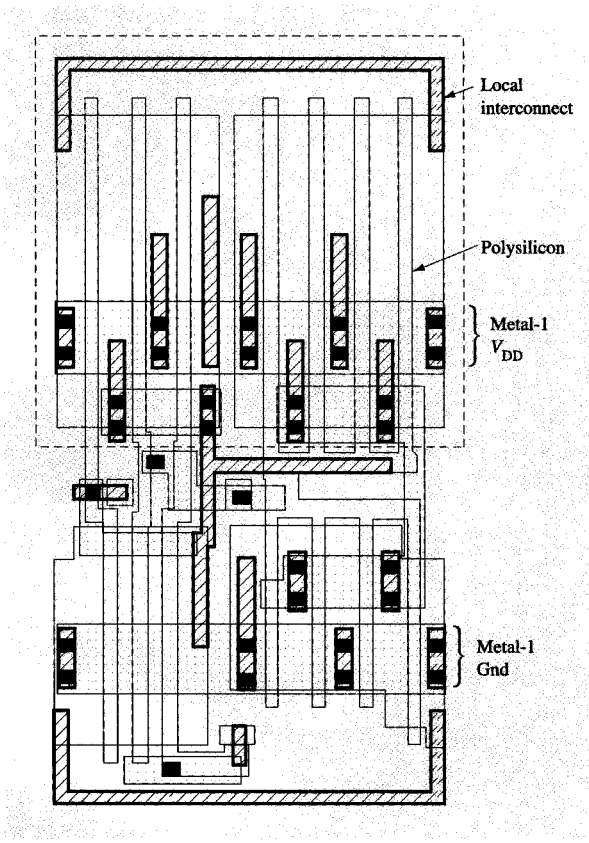

Figure 4

Standard-cell AND3.

exact multiple of the I/O cell pitch. This allows I/O cells to associate with solder bumps in a repeating pattern of eight I/O cells and eight solder bumps. **Figure 2** illustrates one of these pattern cells. The eight narrow rectangles at the bottom are the I/O cells containing the input and output buffers; the circles and the associated wire paths represent the bumps and the top-layer metal connecting the bumps to eight I/O cells.

The solder bumps and pad transfer routing of all four sides of the die are built up by repeating this pattern of eight I/O cells and solder bumps across all edges of the die, with unique pattern cells at the corners of the die. **Figure 3** illustrates features of the die template for a small die, showing how the 0.46-mm-wide pattern cell of Figure 2 is repeated across all edges of the die, with unique pattern cells in the corner.

The base footprint is found on the smallest die (3.5 mm). Each successive die image template is generated by adding two more pattern cells to each edge of the previous die size, generating a series of die templates in 0.92-mm increments. This modular approach allows for expedient design of both die templates and packages.

The design approach described above facilitates the reuse of packages across technologies while allowing customers to maintain plug compatibility.

Gate-array/standard-cell intermix

The CMOS 5 generation of ASICs makes use of the gate-array/standard-cell intermix methodology introduced by IBM in 1987 [7]. The layouts of standard-cell elements are customized on both the diffused layers and the metallization layers. **Figure 4** depicts the layout of the AND3 standard cell, with the unshaded diffusion and polysilicon layers forming the transistors, and the shaded metallization layers (local interconnect, contact, and metal-1) interconnecting the transistors.

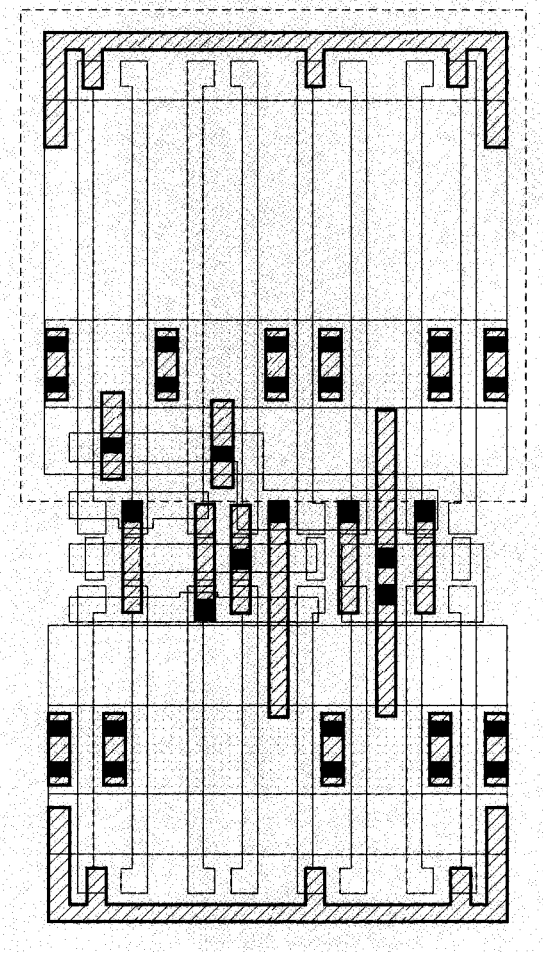

Figure 5

Gate-array AND3.

Gate-array elements are customized only on the metallization layers. **Figure 5** depicts the layout of the AND3 gate-array cell. The transistors formed by the diffusion and polysilicon layers are the standard template for all gate-array cells. Only the metallization layers interconnecting the transistors are unique for this function. The availability of local interconnect allows this cell to leave many metal-1 tracks open for global routing [8].

The identical boundary conditions and power bus location of the cells in Figures 4 and 5 allow them to be placed side by side. The all-level customization of the standard cells generally makes them denser and faster than their gate-array counterparts, and therefore desirable for use in critical paths of a chip design. However, the gate-array elements may be used in less critical paths, and also in paths which may be subject to a later design change.

In the intermix methodology, unused area is filled with transistors of the gate-array template. **Figure 6** shows the standard-cell AND3 of Figure 4 with transistors of the gate-array template placed in the adjacent unused space. Even with five metal layers, designs are routing-limited to less than 100% silicon usage. If a design change to the logic is required after the initial hardware has been built, additional gates can be implemented in these unused transistors with only metal mask changes, as shown in **Figure 7**.

Gate-array I/O functions are included in the library to allow quick cycle time changes to the I/Os as well. This feature allows customers whose external application conditions have changed to accommodate these changes with metal-only changes to the design. The impedance matching and slew rate of the drivers can be altered in this scheme to match changes in the application's off-chip net configurations.

Library deliverables

Access to the library is obtained by the delivery of a design kit from the ASIC supplier to the customer. This design kit includes documentation and models to support selected design methodologies. Software tools have been written for the automatic generation of both documentation and models to speed up development of the design kit for a new technology.

The documentation process has been simplified by the automation of databook generation. The product databook contains detailed information about the complete ASIC offering, including timing information and functional descriptions of each library element. Timing information is automatically extracted from the timing model and inserted into a databook template. Because library element names, pin names, and library element function are consistent across technologies, very little manual editing of the databook is required between technologies.

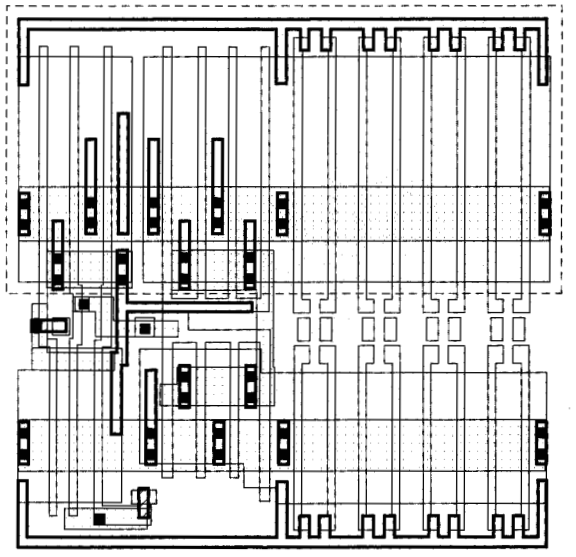

Figure 6

Standard-cell circuit with unused adjacent area filled with unused gate-array transistors.

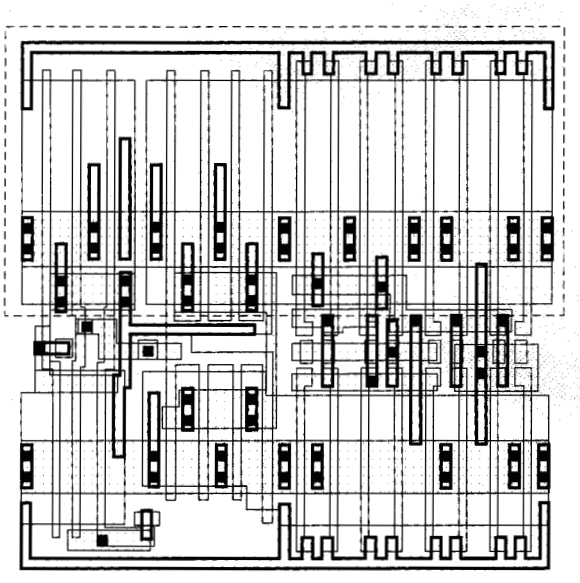

Figure 7

Unused gate-array transistors converted into useful circuits through a metal-only design change.

Table 1 Compilable arrays.

Array type	*Maximum total bits*	*Maximum words*	*Maximum bits per word*
ROM	256K	8192	64
One-port SRAM	256K	8192	128
Two-port SRAM	128K	4096	128
Multi-port register	5120	256	80

The same approach was taken for model generation. Synthesis, simulation, test, and physical design models can all be quickly generated using databases that change little from technology to technology. The information that does change between technologies, such as timing information, is extracted from the timing model and inserted into the other model types. This approach has also led to higher model quality. With an automated system and a common database, there are few steps requiring error-prone manual data entry.

To allow as many customers as possible to maintain their existing design methodologies, the migratable ASIC library supports three unique synthesis/simulation methodologies and two unique physical design methodologies. The synthesis/simulation methodologies include an industry-tool Verilog® path, an industry-tool VHDL path, and an IBM-tool VHDL path. Physical design methodologies include an industry standard tool and an IBM tool methodology [9].

Library function content

The internal logic library can be broken down into several categories: primitive logic, complex logic, unique functions, and latches. The primitive library consists of inverters and two- to four-input NAND, AND, OR, NOR, XOR, and XNOR functions. These functions are easily targeted by synthesis tools, and are available with as many as four drive strengths to optimize performance, power, and input capacitance in various net configurations.

The complex logic elements in the library consist of a series of compound AND–OR, AND–NOR, OR–AND, OR–NAND functions. These elements are provided in various input pin combinations, and offer significant advantages in density and performance over the construction of these functions with primitive elements.

The unique logic segment of the library includes a variety of functions fundamental to ASIC design. This category includes multiplexors, decoders, full adders, and comparators. A full set of functions are also provided to support clock tree design and optimization, including clock drivers, clock splitters, clock choppers, and terminators of varying load used for clock tree balancing.

The library includes a variety of sequential elements which offer functionality familiar to customers with logic design experience based on non-scan edge-triggered flip-flops; these elements also support IBM's full-scan design-for-test methodology.

A full-scan version of the edge-triggered D flip-flop is available to support customers who choose the simplest design-for-test implementation of an ASIC. A full-scan master–slave register is available to support the many customers who choose a more sophisticated design-for-test implementation which shares the library clock splitter with a group of registers. These registers also support the customers who have a long-established experience base in full-scan design-for-test. A register with the data output from the master stage is available to support full-scan implementation of transparent latches.

Compilable memory arrays

Most ASICs require embedded memory capability for such applications as local control or data caches, where very fast access by the surrounding logic is essential. Typical ASICs include multiple instances of arrays of varying configurations. Supporting this variety of array requirements requires a compiler which can create unique arrays with little design resource.

IBM ASICs development provides compilers for a variety of array types, designed for optimum density, with performance capability consistent with the overall logic library. The four compilable array types provided in this design system are summarized in **Table 1**.

All of the compilable memories are designed with compatible naming conventions and operating characteristics to make working with multiple array types as simple as possible for the designer. Clocking conventions are consistent among the array types, both during functional operation and also during test. Similarly, pin names and pin active states are consistent where possible.

The one-port and two-port SRAMs and the multi-port register array have individual bit write masks, blocking overwrite of any segment of the data field during a write operation. The self-timed-restore architectures of the RAM and ROM allow both array access and restore to be initiated from a single edge of the array clock, reducing the restrictions on the array clock pulse widths or duty cycle. Data and address inputs are latched internally, minimizing the setup and hold times required on these pins. Data outputs are latched, thus retaining their data until overwritten by a subsequent read cycle.

Multiple column-decode options are supported for all four array types, which allows the designer to vary the aspect ratio of the array to meet the floorplan requirements of the chip. The amount of global wiring blocked by the array is also variable for the SRAMs. In their densest configurations, the SRAMs completely block the metal-2 global wiring level, forcing all vertical wiring

over the array to be done on the metal-4 level. Additional wiring porosity in the array is available for ASICs with heavy vertical wiring utilization or with limited access to the metal-4 level. However, the array width must be increased to produce porosity on the metal-2 layer. The 25% and 50% porosity options increase the width by 33% and 100%, respectively. The designer can make a trade-off between the relative density of the array and the global wirability of the overall ASIC.

The compilable ROM has recently been added to the IBM ASICs menu [10]. It uses the local interconnect level, which is among the metallization levels used to create gate-array elements from the raw transistor template, to personalize the cell. Thus, the ROM contents can be updated with a metal-only change to the chip design, minimizing the turnaround time.

The SRAMs and ROM also include built-in self-test (BIST), which generates a complete pattern set for testing the memory array while requiring only a clock sequence from the chip tester. For the SRAMs, the BIST applies multiple test patterns to the array, and can be run at functional speed, thereby providing an ac test of the SRAM access and cycle times. The BIST for the ROM has clocking requirements similar to those of the SRAMs, but its internal operation is different. Since the ROM contents are fixed, each address of the array must be read and compared to its expected contents. This is accomplished by using compression circuitry to compress each data word read sequentially from the array into a test signature. After the entire contents of the array have been read, the resulting compression signature is scanned from the chip and compared in the chip tester to the signature calculated for the unique personalization of the ROM.

Multi-port register arrays are also available for implementing very fast arrays up to 5120 bits. In the register arrays (RAs), each memory bit is a full-scan latch. Thus, the array core is larger than for an SRAM of corresponding size, but does not have the area overhead penalty of the BIST circuitry. For arrays with fewer than 1000 bits, an RA is more area-efficient than an equivalent SRAM.

RAs offer multi-port functionality beyond the capability of the SRAMs. Configurations of 1W–1R (one write port and one read port), 1W–2R, 2W–2R, 2W–3R, and 2W–4R are provided.

Since each bit of the RA is scan-testable, the use of many large RAs in an ASIC design results in a very high latch count on the chip. To minimize the impact of these latches on the test time for the logic around the RAs, a scan-bypass methodology has been implemented. The RA scan chain is scanned to its full depth only once to verify the integrity of the array. During subsequent scan operations to test surrounding logic, only latches at the input and output boundary of the RA are scanned. Latches in the array, which represent the greatest quantity of RA latches, are bypassed. This approach greatly reduces the number of patterns required during scan-based testing.

All of the compilable array types have been designed to allow rapid migration to future technology generations. Since circuit layouts were done in the migratable design rules, redesign is required only in subcircuits such as SRAM sense amplifiers that are sensitive to technology parameter variations.

Array compiler features

All of the memory array types described above are supported by a compatible set of compilers, written within a single framework. Logic simulation, synthesis, and test-generation models are compiled to unique configurations using code written in C. Physical and checking models are compiled using SKILL routines in a Cadence environment, then are streamed out to industry-standard GDS2 and EDIF formats.

A compiler is accessible by customers on the Internet through the World Wide Web. This compiler allows customers to compile and retrieve simulation and synthesis models for any configuration array as their needs require without the assistance of an IBM representative. Customers can assess ASIC architecture variations with minimal resource requirements for both the customer and the IBM ASICs organization. The proprietary mask and circuit schematic data for the arrays are not released to customers, but are retained in-house for the final configurations to be implemented on the customer's ASIC.

I/O circuits

The I/O library for the 3.3-V CMOS 5S technology has been designed to support 5.0-V and 3.3-V interface standards. The driver output impedance options range from 20 to 65 Ω, and each has three slew rates to allow the customer to trade off driver performance versus switching noise. To support the reduced-pin-count test methodology, each driver function is implemented as bidirectional I/O for use in I/O cells not contacted during test. Stand-alone drivers and input buffers are available for contacted pads. Optional pull-up or pull-down resistors are offered, as well as input buffers with and without hysteresis.

The I/O library for the 2.5-V CMOS 5X technology offers functions similar to those of the CMOS 5S library. A complete family of 2.5-V I/Os and a set of 3.3-V I/Os are available. The 2.5-V I/Os may be used on either the single-supply 2.5-V chips or the dual-supply chips. The 3.3-V I/Os may be used only on dual-supply 2.5-V/3.3-V chips.

The I/O library for CMOS 6S, also a 2.5-V technology, will include the same variety of 2.5-V and 3.3-V drivers as CMOS 5X, with the same package options.

Library design point

The key objective in defining the internal library design point is to quickly create designs which are both easily migratable and easily used, without sacrificing performance or density. This is accomplished through the choice of electrical and physical design characteristics.

Internal library designs must work well over a range of technologies with different supply voltages and device characteristics, particularly ratios of p-MOS to n-MOS device current. As a result, designs are analyzed not only for the technology in which they are designed, but also for those to which they are expected to migrate. Typical device ratios are established to be the best compromise between technologies, and designs that are highly ratio-dependent are generally avoided. Clock circuits are designed to minimize differences between rising and falling delay and transition rates. Critical cells are redesigned for specific technologies where required.

Physical design choices have been made to avoid unnecessary complexity without affecting density or wirability. Simple, symmetric cell boundary conditions reduce the incidence of chip-level cell-to-cell conflicts and allow cells to be mirrored to reduce wire lengths. All cells are self-contained, so separate placement of n-well or substrate contacts is not required. All designs adhere to strict standards for metal usage. With the local interconnect level available for intracell routing, only a minimum set of metal-1 tracks is reserved for use by the cell design. Additional tracks are used in preferential order when required, typically leaving six metal-1 tracks free for global routing. All pins are metal-2-accessible, but in many cases nets between adjacently placed cells can be wired on metal-1. The metal usage standards allow the creation of simple place-and-route models. Collectively, all such choices simplify library usage and reduce place-and-route turnaround time.

The I/O library design point is chosen to optimally support the multiple interface types required by the customer. However, I/O circuits do not migrate well because of limitations imposed by technology reliability and supply voltage issues. The objective is therefore to define families of circuit types which can be migrated across technologies with the least amount of redesign. For example, 3.3-V I/Os in a 3.3-V technology become the basis for 2.5-V designs in a 2.5-V technology. This is aided by the use of implanted resistors to control driver output impedance and slew rates, rather than environment sensing or feedback-based schemes which can require extensive circuit optimization. Logical behavior is preserved across technologies to simplify the migration of customer designs.

Library characterization and qualification

Qualification of an ASIC design system is much different from the qualification of a custom chip. For ASIC chips, each unique customer design is not individually characterized. Instead, delivery of a quality ASIC design system requires verification and qualification of all of the individual elements that may be used in that design system, ensuring that every element may be combined with the other elements to produce a design which will meet customer specifications. To achieve this goal, the following characterization and verification activities are required for each element of the library:

- Functionality verification.
- Timing verification.
- Reliability verification.
- Model verification.

The logical functionality of all of the logic, I/O, and compilable array elements in the library is verified on a test chip. Subsequent additions to the ASIC library are usually logically verified on a customer design, where a "green light" at the tester indicates that the circuit has performed as expected by the test model.

Performance characterization of the logic library is done on representative elements from the library placed in ring oscillators on the test chip. Elements are chosen which represent the various circuit topologies encountered in the library, including the various stacking increments of p-MOS and n-MOS devices, as well as passgate configurations. These ring oscillators are measured on a group of test chips chosen to have process parameters which span the allowed process skew. The resultant hardware delay measurements are compared to simulations of these circuits performed using the extracted process parameters. Any miscorrelation indicates an inability of the simulation device models to accurately predict the device operation at a specific point in the process, and the simulation models are corrected accordingly.

The timing models for the library elements are verified by comparing the circuit delays obtained from a range of hardware processed within the manufacturing tolerances to the range of slow process to fast process delays contained in the delay models. Any hardware delays which fall outside the range in the model indicate an error in the ability of the model to bound the product delay, and the model must be corrected.

A different characterization strategy is required for compilable arrays and embedded functions. These macros are bounded by characterization latches, and delays are

measured by varying the launch-capture timing of the boundary latches. This approach further eliminates the variable net delays associated with I/Os. The process of model correlation and delay model verification for these macros is similar to that described above for logic elements.

Reliability verification for an ASIC design system is performed by a combination of stressing representative hardware and analyzing the individual library elements. An initial stress of a quantity of samples of the test chip is performed at elevated voltage and temperature in order to identify sensitivities in the circuitry aggravated by these acceleration factors. A later stress is performed on representative customer hardware to verify the designs under use conditions before qualifying the design system for manufacture.

Given the myriad of possible configurations in which each library element could be used on a customer design, it is not practical to build and stress sufficient hardware to verify all possibilities. Therefore, reliability analyses of the individual elements must be performed to ensure robustness against hot-electron and electromigration effects. These analyses are performed through circuit simulations to determine the maximum currents in the devices for comparison to the reliability limits of the transistors and the intracell routing.

Summary

The marketplace demands from the ASIC supplier early access to emerging silicon technology. The ASIC customer's goal is to realize a return on a long-term investment in a design methodology and tool set which is supported by their supplier. This methodology is measured by how quickly and efficiently the ASIC customer can bring a new product to market.

This paper presents a strategy for the ASIC supplier to meet these demands. Fundamental to this strategy is a progression of silicon technology which enables circuit designs to be quickly and easily migrated forward into more advanced technologies. On this base, a library of circuit functions is developed in such a way as to remain near optimum in each technology, and to maintain a consistent look and feel to the chip designer. This reuse of circuit designs has the added benefit to the ASIC developer of reduced design risk in each library, and economies in characterization, quality, and reliability evaluations. This approach also features the reuse of package designs in multiple technologies to ease the introduction of new silicon technology into system designs. Finally, the migration strategy facilitates more effective communication of design data to designers through automated generation of library documentation, through reuse of simulation and test models, and through customer access to memory compilers via the Internet.

References

1. D. L. Wheater, P. Nigh, J. T. Mechler, and L. Lacroix, "ASIC Test Cost/Strategy Trade Offs," *Proceedings of the IEEE International Test Conference*, 1994, pp. 93–102.
2. R. W. Bassett, M. E. Turner, J. H. Panner, P. S. Gillis, S. F. Oakland, and D. W. Stout, "Boundary-Scan Design Principles for Efficient LSSD ASIC Testing," *IBM J. Res. Develop.* **34,** 339–354 (1990).
3. R. W. Bassett, B. J. Butkus, S. L. Dingle, M. R. Faucher, P. S. Gillis, J. H. Panner, J. G. Petrovick, and D. L. Wheater, "Low Cost Testing of High Density Logic Components," *IEEE Design & Test of Computers* **7**, 15–28 (April 1990).
4. J. J. Engel, T. S. Guzowski, A. Hunt, D. E. Lackey, L. D. Pickup, R. A. Proctor, K. Reynolds, A. M. Rincon, and D. R. Stauffer, "Design Methodology for IBM ASIC Products," *IBM J. Res. Develop.* **40,** 387–406 (1996, this issue).
5. C. W. Koburger III, W. F. Clark, J. W. Adkisson, E. Adler, P. E. Bakeman, A. S. Bergendahl, A. B. Botula, W. Chang, B. Davari, J. H. Givens, H. H. Hansen, S. J. Holmes, D. V. Horak, C. H. Lam, J. B. Lasky, S. E. Luce, R. W. Mann, G. L. Miles, J. S. Nakos, E. J. Nowak, G. Shahidi, Y. Taur, F. R. White, and M. R. Wordeman, "A Half-Micron CMOS Logic Generation," *IBM J. Res. Develop.* **39,** 215–227 (1995).
6. R. Gregor, C. Ng, J. Libous, E. Carter, R. Beaudoin, A. Chu, D. Grindel, J. Kinney, M. Lee, L. Mentes, J. Oppold, M. Russel, A. Secor, and G. Yenik, "A One Million Circuit CMOS ASIC Logic Family," *Proceedings of the IEEE Custom Integrated Circuits Conference*, 1993, pp. 23.1.1–23.1.4.
7. Larry Wissel, Douglas Stout, and Nathan Buck, "Gate Array Library Design Using Local Interconnect," *Proceedings of the IEEE Custom Integrated Circuits Conference*, 1996, pp. 23.5.1–23.5.4.
8. R. Hornung, M. Bonneau, B. Waymel, J. Fiore, E. Gould, R. Piro, J. Martin, L. McAllister, and S. Tom, "A Versatile VLSI Design System for Combining Gate Array and Standard Cell Circuits on the Same Chip," presented at the IEEE Custom Integrated Circuits Conference, May 4–7, 1987.
9. J. H. Panner, R. P. Abato, R. W. Bassett, K. M. Carrig, P. S. Gillis, D. J. Hathaway, and T. W. Sehr, "A 300K-Circuit ASIC Logic Family CAD System," *Technical Bulletin TR-19.90507*, IBM Microelectronics Division, Essex Junction, Vermont, February 1990.
10. Robert L. Barry, John D. Chickanosky, Francesco M. Masci, Ronald A. Piro, Steven F. Oakland, Michael R. Ouellette, Douglas W. Kemerer, Maria R. Noack, and William C. Leipold, "A High-Performance ROM Compiler for 0.50 μm and 0.36 μm CMOS Technologies," presented at the IEEE International ASIC Conference and Exhibit, September 18–22, 1995.

Received October 23, 1995; accepted for publication February 14, 1996

Thomas R. Bednar *IBM Microelectronics Division, Burlington facility, Essex Junction, Vermont 05452 (tbednar@vnet.ibm.com).* Mr. Bednar received the B.S. degree in electrical engineering from the University of Notre Dame in 1982, and the M.S. degree in electrical engineering from the University of Vermont in 1988. He joined IBM in Essex Junction in 1982. He has worked on the definition, design, and qualification of several generations of CMOS and biCMOS circuit libraries and products. Mr. Bednar is currently the technical team leader for CMOS 5 and 6 ASIC product development.

Ronald A. Piro *IBM Microelectronics Division, Burlington facility, Essex Junction, Vermont 05452 (rpiro@vnet.ibm.com).* Mr. Piro attended Vermont Technical College and the Massachusetts Institute of Technology. In 1982 he joined IBM in Essex Junction in the logic products area. Mr. Piro has designed I/O circuits and SRAMs in several generations of CMOS technology; he is now the technical team leader for compilable SRAMs and ROMs for ASIC products.

Douglas W. Stout *IBM Microelectronics Division, Burlington facility, Essex Junction, Vermont 05452 (dstout@vnet.ibm.com).* Mr. Stout received the B.S.E. and M.S.E. in electrical engineering from the University of Michigan in 1981 and 1982, respectively. He joined the logic products group at IBM in Essex Junction in 1983. Mr. Stout has designed internal and I/O circuits and die images in several generations of technology. He is now the technical team leader for libraries and I/Os for ASIC products.

Lawrence Wissel *IBM Microelectronics Division, Burlington facility, Essex Junction, Vermont 05452 (wissel@vnet.ibm.com).* Mr. Wissel received the B.S. degree in engineering physics from Cornell University in 1975. He served as an officer in the U.S. Navy until 1979, when he joined IBM in Essex Junction. Mr. Wissel has designed RAMs, ROMs, I/Os, and die image templates in several generations of CMOS and biCMOS technology.

Paul S. Zuchowski *IBM Microelectronics Division, Burlington facility, Essex Junction, Vermont 05452 (paulz@vnet.ibm.com).* Mr. Zuchowski received the B.S. degree in electrical and computer engineering from Clarkson University in 1989. He then joined IBM in Essex Junction. In 1995, he received the M.S. degree in microelectronic manufacturing from Rensselaer Polytechnic Institute. Mr. Zuchowski has worked on full-chip physical verification, physical model generation methodology, and physical design methodology for library elements, cores, die templates, and ASIC chips. He is now the technical team leader for the methodology integration team for ASIC products.