

Joseph C. Logue
Walter J. Kleinfelder
Paul Lowy
J. Randal Moulic
Wei-Wha Wu

Techniques for Improving Engineering Productivity of VLSI Designs

The ability of the semiconductor industry to produce chips with higher and higher circuit densities has created a challenge for the product designer: to utilize this capacity and still develop chips rapidly at reasonable cost. A multi-faceted approach to VLSI design is described that significantly reduces product development time and resource from those required with existing methods. This approach is based on the use of PLA structures or macros. It consists of a hardware/software modeling technique, use of laser-personalizable PLAs for rapid modeling of PLA macros, and a method for repairing design errors (that may hide other errors) on the actual VLSI wafers with a laser tool. A two-pass VLSI design is therefore highly probable.

Introduction

Achieving a cost-effective VLSI (very large scale integration) chip design with a two-year design cycle requires that chip function be maximized while the time and number of iterations through the design cycle are minimized. However, as the degree of integration increases to thousands or tens of thousands of circuits per chip, it becomes more difficult to ensure that the desired function is correct and that the final chip performs that function after the first or even succeeding design passes. This is especially true when the design is optimized for circuit density and performance (this type of design is known as a *custom design*). Examples of such custom chips now available are the Motorola 68000 microprocessor [1] and the Zilog Z8000 microprocessor [2]. The claim has been made that today's microprocessors have from 25 000 to 70 000 transistors, which must certainly represent more than 10 000 circuits [1]. The question arises as to whether the VLSI design capability of the semiconductor industry can keep pace with its technical ability to fabricate an ever-increasing number of logic gates on a chip. It has been estimated that today's design cost for a VLSI chip is approximately \$100 per gate [3]. As our capability to put more gates on a chip increases, we must improve the pro-

ductivity of our chip designers; otherwise, the growth of VLSI and the semiconductor industry will be seriously limited.

One technique for reducing design time at the expense of circuit density, now becoming more prevalent, is the *masterslice*. Recent advances include a 5000-circuit wired chip of this type described by IBM [4] and an 1144-cell image chip supported by design automation marketed by Signetics [5]. It is possible not only to reduce the design time by automated means [6], but also to decrease the manufacturing time, since wafers consisting of pre-designed circuits are built and stockpiled for later personalization by metal interconnections. Regardless of these advantages, it still behooves designers to guarantee correctness in order to reduce the number of manufacturing passes. Each design cycle subsequent to the first requires a debug phase to determine problems, a redesign effort, and a build phase. Only after the design is functioning well can exhaustive testing under stress conditions be performed to guarantee that the chip meets specifications. The turnaround time of the build phase varies with several factors. These include the complexity of the pro-

Copyright 1981 by International Business Machines Corporation. Copying is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract may be used without further permission in computer-based and other information-service systems. Permission to *republish* other excerpts should be obtained from the Editor.

Table 1 Eight-bit microprocessor characteristics.

Technology	n-channel MOSFET
Chip size	7.92 × 7.92 mm
Type of design	Macro
Number of instructions	65
I/O	Programmed I/O, 8 interrupt levels, memory-mapped I/O, register-mapped I/O
Storage addressability	64K bytes, extendable to 512K bytes
Performance	900-ns cycle, 6.5 μ s per instruction (average)
Number of devices	12 100
Number of equivalent logic gates (2.5-way NOR circuits)	3500
Power	1.2 W
Power supplies	+5 V $\pm$ 10%, +8.5 V $\pm$ 10%
Module size	25 mm
Module pins	72

cess, the necessity to continue to process chips already in production on the line, and process problems. In some cases, the time to process prototype parts may be as long as four months.

Based on these facts and the need to achieve densities higher than those obtainable from a masterslice design, a multi-faceted approach that is geared toward custom design has evolved. It is one of obviously many approaches that can be taken. We do not claim this to be the best approach, even though it has worked very well for us. VLSI design is a rapidly evolving body of knowledge and the approach to be described is still in an evolutionary stage.

Design and verification

To achieve increased productivity in design and still maintain reasonable circuit densities, a macro design approach is utilized. Macros are collections of circuits which, when physically implemented, fit closely together; e.g., registers, bus multiplexers, PLAs, etc. We make use of PLA macros wherever possible because of their design advantages, as described by R. L. Golden *et al.* [7], and their modeling advantages, which we now present in detail. To illustrate our methodology, the verification procedure for an eight-bit microprocessor chip (Table 1) will be described.

After initial design, a dual effort was undertaken whereby a prototype was designed and built using TTL (transistor-transistor logic) modules and field-programmable PLAs, and then tested exhaustively. In parallel with this effort, the logical design of the chip in its ultimate technology (n-channel MOSFET) was entered into a computer data base and simulated. This logical represen-





Figure 1 Microprocessor (a) logic data flow and (b) physical macro placement.

tation was also used to automatically check against the physical design and to create test patterns. However, this parallel approach created two concerns: Manpower and schedule constraints would be exceeded, and the LSI data base would not faithfully reflect the prototype design.

Our solution to the first concern was to initially define functionally equivalent macros for both designs, maintain

I/O correspondence along macro boundaries, and keep the designs identical to avoid duplication of design effort. The macros ranged from 50 to 500 equivalent logic gates (they are not limited by these numbers) and may be categorized as PLA or random macros. PLA macros consisted of input partitioning circuits, AND and OR arrays, and output buffers. Random macros were made up of registers, multiplexers, and other primitive logic functions. There were 19 PLA macros, which made up about one-half the circuits on the chip, including the complex "control" section (14 PLA macros). Figure 1(a) is a logic flow diagram of the microprocessor, showing macro partitioning. PLA macros are shown as shaded and random macros as unshaded. The corresponding physical outlines of these macros as they lay out on the chip are depicted in Fig. 1(b).

After the partitioning (and hence the preliminary design of the chip) had been completed, design of the prototype commenced. Most PLA macros were implemented in the prototype as field-programmable PLA modules (FPLAs) [8]; see Table 2. Of the 19 PLA macros, 17 mapped directly into FPLAs. The largest control PLA was implemented as seven FPLAs, and one PLA (the ALU), although functionally equivalent, was redesigned as two PLAs in series. It should be pointed out that the laser-personalizable PLAs (LPLAs) and the laser tool described in this paper were not available to us at this time. Had these been available, all PLAs in the LSI version would have mapped directly to the prototype design, saving modules and design time. During the debug and test of the prototype, corrections to this part of the design were made in minutes by burning in new PLAs. Other macros were implemented with available TTL register, multiplexer, and unit logic integrated circuit modules. These macros made up the data flow, which has similar connectivity, and therefore required fewer changes than the control section. To illustrate this point, the reader is referred back to the flow diagram in Fig. 1(a) to see how an instruction is executed. The eight-bit operation field of an instruction is loaded into the instruction register and decoded with the current state of the machine by the control PLA section to generate approximately 75 signals to control the rest of the chip. On the other hand, the eight-bit (byte) data field of the instruction is merely routed from the data buffer register to the A register and then to the ALU PLA, where it is processed (*e.g.*, during an ADD operation) to form a new byte-wide field. The registers, multiplexers (not shown), and parity-check circuits are connected similarly for each bit, and the controlling inputs of the byte multiplexers are from the control PLA section. Therefore, by design, the most complex portions of the chip are indeed the PLA macros; however, these can be changed most easily in the prototype model.

Table 2 FPLA and LPLA characteristics.

	FPLA ^a	New FPLA ^b	LPLA
Inputs	16	24	34 ^d
Outputs	8	16 ^d	44 ^d
Internal feedbacks	0	16 ^c	22
Product terms	48	72	80 or 160
Array bits	1920	9216	16 320
Partitioning	1-bit	2-bit	2-bit
Maximum cycle time (input to output delay)	50 ns	200 ns	100 ns
Power dissipation	0.6 W (typical)	1.1 W (maximum)	1.85 W (maximum)
Power supply	+5 V	+5 V	+5 V
Testing	Standard	Standard	LSSD ^e

^aSignetics 82S100

^bNEC μ PB450D

^cJK latches

^dLatched

^eLevel-sensitive scan design [6]

The prototype consisted of four wired cards of about 50 modules each. One engineer and one laboratory technician designed, built, and thoroughly tested the prototype in six months, expending a total of eight man-months of effort. At the completion of this effort, we were able to verify the prototype design by running in excess of six million instructions on the model in about five minutes. In addition, almost half of the 3500 circuits (the directly mapped PLA portion) had been designed for use "as is" in the LSI version.

Our second concern, ensuring that the LSI data base matched the now-functional prototype, was alleviated by the following approach. First, new PLA modules for the prototype were personalized from tables in the LSI data base and tested exhaustively in the prototype to ensure that the PLA data were correct. Next, extensive simulations were run on a model derived from the LSI data base. The simulator program of the IBM Engineering Design System (a collection of design-automation programs) provides a memory model which can be "loaded" with instructions that are then used in the LSI model simulation. Other features of the simulator program are that it keeps track of all selected PLA product terms and all switched nets. Instructions were added to the memory model and simulated until every product term was selected, and every net and every I/O were pulsed at least once. In so doing, a subset of the six million instructions used to test the prototype were run against the software model. The subset consisted of 35 programs of approximately 100 instructions each, requiring about five minutes apiece to run, for a total of about three hours of computer time (on an IBM 370/168 System). It was usually possible




Figure 2 Schematic of the laser tool.

to complete all these simulations within one day. Third, a log of all changes to the prototype was kept, and the update of the LSI data base was verified. Finally, any macro which differed significantly in implementation from the prototype (such as the ALU) was exhaustively simulated by itself. Since simulation was used to verify that the computer data base matched the prototype, the data base was suitable for automatic logical/physical checking.

Two engineers and one clerical specialist expended eight man-months to document and simulate the LSI design. This effort commenced two months after the start of the prototype design and lasted five months. The logic design effort (prototype plus conversion to target technology), as described in this paper for the first-pass design of the 3500-circuit microprocessor, was completed with one and one-third man-years of effort in seven calendar months.

The results of our dual approach provided us with first-pass chips from manufacturing that contained only three design defects. These design errors were attributed to omissions of logical/physical checks (an interchanged net, an incorrect circuit type, and a missing voltage connection). Fortunately, with the aid of the laser tool to be described later, we were able to break a connection on the chip and a suitable TTL patch was added to the card containing the LSI chip, so that complete chip testing could proceed. Another result reported previously [7] was that we achieved higher densities than those obtainable with a masterslice approach. A comparison between the microprocessor and a masterslice in the same technology showed that the macro approach afforded a 1.8-fold increase in density and a 1.7-fold reduction in power for the same or better performance.

The reader will note that the chip described here does not even approach the circuit densities of the most recent custom microprocessors such as the Motorola 68000 [because the 68000 is manufactured using the more advanced technology HMOS (high-density short-channel MOS)]. The question of extendability is, therefore, a valid one. One limitation of our verification approach is that the hardware model for a chip must not exceed 6 or 7 cards (about 350 modules) or it will become unwieldy. Therefore, as the target technology improves, so must the prototype technology. This has happened for the case of the PLA macros where the LPLA (to be described in the next section) has been developed. Its capability (see third column of Table 2) significantly exceeds that of the FPLA. Similarly, advances in off-the-shelf modules for the random macros, *e.g.*, greater function, more I/O, and programmability, would be desirable in the future to prevent the hardware model from becoming too large. Another limitation is the possibility that because of human error the LSI data base may not reflect precisely the same function as the prototype, despite the precautions outlined previously. We believe this risk is worth taking because of the schedule and resource savings, and therefore we recently applied our methodology to the design of a complex logic chip containing twice as many circuits as the eight-bit microprocessor and approaching the density of currently available microprocessors. In this case, the chip functioned properly when received from manufacturing with the exception of four minor logical errors which were not all discovered until the chip had been tested thoroughly in its ultimate application. Interestingly, three of the four errors were traced to a mismatch between the LSI design and the prototype. Nonetheless, the first-pass hardware is adequate for stress testing and at the time this paper is being written, a two-pass design appears highly probable.

Laser personalization of PLAs

As indicated in the introduction, our design approach to VLSI has evolved over a period of time and is still evolving. At the time the eight-bit microprocessor used as the example was in development, our laser tool and the LPLA were in a very early state of development. This is why field-programmable PLAs were used to implement the control circuits of the hardware simulator previously described. Today, the laser tool and the LPLA are being used to do hardware modeling with very rapid turnaround times, *i.e.*, personalization of an entire module in 20 minutes.

The laser tool used in IBM consists of a pulsed dye laser pumped by a pulsed ultraviolet nitrogen laser, beam optics, a precision *x/y/z* table combined with a $\theta/\phi_1/\phi_2$ tower assembly, a TV viewer, a laser energy monitor, and

a control unit. The tool is driven by an IBM 5100 trans-portable computer [9]. The dye laser produces several-nanosecond light pulses with a peak power of up to several kilowatts. These pulses are attenuated by an achromatic halfwave plate and delivered to the optical system by a Pockels cell used as a high-speed electronic shutter. Figure 2 shows the schematic of the laser tool.

A logical "1" can be personalized on a semiconductor chip by connecting a transistor to a crosspoint. This is accomplished by welding two layers of metal together. As shown in Fig. 3, the first laser pulse penetrates the top quartz and melts the top layer of metal, which erupts upward. As a result, a thin layer of metal coats the wall of the erupted hole. When the second laser pulse hits the bottom layer of metal, the molten metal erupts again to form a cone-shaped weld connecting the two layers of metal. The laser tool can make 2000 such connections in less than 10 minutes. The thicker the insulator on the metal, the higher the laser energy required to make a good weld. With reference to its reliability, a good weld can withstand more than 100 mA before it opens. Parts made from the LPLA contain some 1200 to 3200 good welds.

The LPLA used for building functional macros for modeling VLSI hardware, or for building early engineering models, consists of input latches, bit partitioning circuits, an AND array, an OR array, output latches, and off-chip drivers. The PLA can be programmed to operate as a large PLA or as two small PLAs. Table 2 shows the functional characteristics of the PLA. It performs both combinatorial and sequential logic and is equivalent to a 600-logic-gate masterslice. The laser personalization technique enables us to reduce power dissipation both by programming the load resistor values to half or three-quarters of their maximum value, depending on the circuit loading, and by disconnecting the unused circuits on the chip. This is accomplished, respectively, by welding selected resistor contact terminals and by welding current-source circuits to ground. In addition, the user has the following programmable options: gated or ungated input latches; one- or two-bit partitioning; up to a maximum of 160 product terms; JK, AND, GATED-CLOCK, or exclusive-OR output latches; three-state, push-pull, or open-collector off-chip drivers; normal or inverted output; and one or two clock designs. These features, plus the 16 320 programmable sites in the array, dictate a modest chip size of 6.85×6.85 mm. The chip is solder-ball mounted to a 100-pin, 28-mm metallized ceramic substrate with a hole in the center.

In our first attempt at personalization and testing, chips were pretested on the wafer, personalized, and tested functionally. The wafer was then diced, and the chips




Figure 3 Connection of two metal layers by means of a laser beam.

were packaged on the module and tested. However, the yield and turnaround times of this method were unsatisfactory. To improve the situation we developed a new technique we call *in-line lasering and testing* (ILLT).

Pretested modules having good peripheral circuits (the unpersonalized array cannot be tested) are stockpiled to be ready for laser personalization and testing at the module level. Personalization patterns are processed to obtain two separate sets of data: x - y coordinates for the laser personalization step and test patterns. A high-speed link is used to transmit these data from an IBM System/370 to the laser tool. The system is programmed such that if one PLA product term is in error, it either can be immediately repaired by repeating the laser pulses for that term, or deleted by adding welds such that the product term can never be selected, *e.g.*, personalizing all the crosspoints for that term. If a product term is deleted, the same personalization can be automatically applied to and retested at the next row of the array, so that even though the number of personalized crosspoints (and hence welds) is large, the yield will be good. This process is repeated again and again until all the product terms or rows are laser-personalized and test as "good." It takes only 20 minutes to complete a module, including module-alignment time. The module yield using the ILLT technique was improved by a factor of from two to seven over our initial method. The larger the number of welds on a chip, the greater the yield improvement.

The major advantage of LPLAs is the ability to produce high-function hardware rapidly in a heavy design change

environment. A disadvantage is their cost. FPLAs such as those used in the microprocessor model can also be used to achieve rapid turnaround time at less cost than the LPLA since production of laser-manufactured parts entails high overhead and low throughput. However, commercially available FPLAs do not contain the function of the LPLA, as can be seen by comparing the characteristics of a new higher-functionality bipolar FPLA [10] with the LPLA (see Table 2). This is because FPLA technologies such as fusible link or junction punch-through require additional circuitry for personalization. In deciding which type of PLA to choose for modeling, one must therefore consider not only the cost but also by how much the increased function will simplify hardware modeling.

For the future, the laser tool can be modified to achieve smaller beam size by using a tighter optical system and a different dye to change the wavelength of the laser output. The table within the laser tool also can be improved to obtain higher resolution and accuracy by implementing a more sensitive position detector and a better mechanical structure. These improvements and thinner quartz and metal layers on the semiconductor chip would make it possible to create smaller laser welds, *i.e.*, to decrease their size from about 5–7 μm to 1–2 μm , thus providing a future laser personalization technique compatible with 1- to 2- μm mask geometries.

Laser repair of VLSI chips

The laser tool for personalizing PLAs which we have described has also been used to rapidly repair and modify the wiring structures of fully fabricated chips. Design alterations have been made on microprocessor, memory, and masterslice logic chips without the need for totally recreating the parts with a new, corrected mask set.

The laser tool can form a weld between two overlying conductors separated by an insulating layer (Fig. 3), a situation which might exist if an interconnection between layers of metal was inadvertently left out of a via personalization mask [11, 12]. During the development cycle of the IBM 3081 CPU, a repair of this kind was made on a high-speed cache memory chip. A yield problem in early hardware was traced to a second-to-third-level-metal via which had been omitted from the mask set. In order to verify this as the problem, the conductors were "welded" in the appropriate location using the laser tool. Thus, the problem was solved in three days *vs.* a conventional turnaround time on the standard processing line of three months.

Selected conductor patterns can be deleted by utilizing the vaporization capability of the laser. This type of correction enabled us to proceed with the testing of the eight-

bit microprocessor with the first-pass chips. Specifically, one of the inputs to a "dot" OR circuit had the wrong polarity, making it active at the wrong time. The OR circuit was deactivated by breaking an aluminum segment of this input with the laser. Additional TTL modules were next added to the card containing the chip to perform the intended function. We were then able to run the six million instructions against the chip.

The deletion technique has also proved effective in diagnosing faulty IC designs. The designer can use the laser tool to "dissect" portions of the chip wiring and isolate complete sections of a chip. This was done on a masterslice chip designed as a logic level converter. In this case, the chip was mounted back-down in an empty module package and chip pads were wire-bonded to module pins. The device was then powered up for electrical testing on the laser tool. While monitoring circuit outputs, the laser was used to delete aluminum interconnections, one after another, until the faulty one was uncovered. The defect proved to be a short in the chip wiring which resulted from a mask error. Usually, metal conductors are removed with a single 10-ns high-power laser pulse (100–500 W peak power) directed at the feature, generally without damage to underlying structures.

It is also possible to repair metal-to-metal via connections in which the resistance of the connection is too high due to incomplete formation of the via during processing. Vias like this can be locally sintered by using a medium-power laser pulse (20–50 W peak power). The advantage of this technique is that only the via is exposed to the high temperature, typically 300–400°C. The method was applied to a 64K ($K = 1024$) RAM chip during its development cycle when testing indicated high-resistance paths in one portion of the chip. Several chips were corrected with the laser to provide hardware for the developers. Meanwhile, the problem was isolated to the via processing steps and solved.

One disadvantage of the vaporization and melting of metal films by the laser is that dimensions cannot usually be controlled to better than 2 or 3 μm . Rough elevated edges are frequently formed. If the deletion or weld to be made is beneath a passivation layer, that layer is frequently blown away, leaving considerable microcracking in the surrounding area. The resulting chip topography is typically such that further planar processing steps are prohibited; hence, corrections must be made on completely finished devices. Another limitation of this type of correction is that a line to be deleted must be visible and not overlaid by another metallization level. Welds can only be made between existing overlapping patterns. Additionally, the nature of these corrections limits their us-

age to short-lived development parts because the aluminum metallurgy is frequently exposed to the environment and is thus subject to corrosion. Microcracking of the dielectric layers could also present long-term stability problems. To circumvent these limitations, a series of techniques called *laser microsurgery* have been developed [13] in which 1-2- μm geometries can be attained.

Laser microsurgery provides the VLSI designer with a method of implementing changes to a design at any stage of wafer fabrication in a manner analogous to the way wire-wrap or solder changes are made to a card, but with standard silicon planar processes. The schematic of the laser tool (Fig. 2) shows that the system has all the basic components of an artwork generation system: a computer-controlled positioning table, a focused source of ultraviolet energy (the laser), and an aperture or exposure size control. Therefore, it is possible to locally expose photosensitive materials such as photoresists. The laser tool can perform on-chip pattern generation, providing the circuit designer with a freedom not found in the previous weld/delete method, or for that matter, in current batch processing. The designer can define new and unique shapes on the surfaces of chips during any stage of processing *without* the use of a master mask by exposing a resist with the laser.

Since the wavelength of our dye laser is 420 nm (Stilbene-3 dye), it is well within the spectral sensitivity of most resists. The energy density needed to expose a film such as AZ-1350J [14] is $60 \text{ mJ}/\text{cm}^2$. The objective lens (Lietz 32X) of the laser tool focuses the laser beam to a 1- μm -diameter spot so that a 10-ns, 60-mW laser pulse suffices to expose the film. At low power levels (60-500 mW peak power) the diameter of the exposed spot is proportional to the incident laser power, providing an aperture capacity. The spot diameter varies from about 1 to 9 μm at a rate of 70 mW per micrometer of diameter. The developed resist image forms vertical sidewalls due to the monochromatic coherent nature of the laser (see Fig. 4). By successively pulsing the laser and controlling table motion, complex patterns can be exposed.

An example of the use of laser microsurgery to correct a minor defect is the unwanted short in the logic level converter chip previously cited. Once the defect had been determined (by the high-power laser deletion techniques) wafers of this design were stopped in the process cycle. Following exposure and development of the metal-pattern resist image in which the defect resided, the resist was exposed again in the region of the defect using the laser tool. This was done on each chip site of the wafer. Following development, the aluminum pattern was etched as usual, the short being removed at the same time as



Figure 4 Scanning electron microscope image of laser-exposed spot in an AZ-1350J resist film.

the rest of the excess metal. The wafer could then be further processed with passivation quartz and pad metallurgy as if no defect had ever existed. Such a change could be made by regenerating the mask or by creating a special contact mask to expose just the defect area. A designer's choice between these alternatives would be driven by turnaround time. Layout and mask generation can take days to weeks, whereas the laser exposure can be accomplished in a matter of hours. The real-time television monitoring of the laser tool, in conjunction with the instantaneous positional readout of the x - y table, allows the designer to make *in situ* measurements right on the TV screen, and then program the 5100 so that the laser will write the new pattern, step to the next chip, and repeat. Corrections can be made at any level of the device processing cycle, *i.e.*, during ion implantation, diffusion, oxide isolation, via etch, or metallization. Deletions and additions to the interconnection patterns of a chip can take place after processing is completed. The addition of conductors is accomplished by the use of a stencil or metal-lift-off process. If a metallic film is evaporated over the exposed and developed resist film, the metal in the bottom of the exposed region comes in contact with the substrate to form the conductor. Due to the vertical sidewalls of the resist image, the metal is noncontinuous at the pattern edge. The photoresist is then dissolved away, removing the residual metal and leaving the conductor behind.

An example of line additions to a 256K CCD memory chip is shown in Fig. 5. Figure 5(a) shows the exposed and developed resist pattern where lines are to be added. The results of the lift-off step are shown in Fig. 5(b). Corrections to this chip began after the personalization of the first-level metallurgy. A contact (or via) from the metal to a diffused region was omitted on the masks. By using a




Figure 5 (a) Laser-patterned resist film for the addition of new lines. Aluminum lines have been deleted as shown. (b) New aluminum lines added to the chip by metal evaporation and lift-off.




Figure 6 LSI productivity for the eight-bit microprocessor. The numbers refer to man-years and each division on the horizontal scale is one month. The net savings were 7.25 man-years and eight months.

two-step microsurgery process, an opening in the oxide was made with one resist pattern and an etch step. A metal contact and interconnection pattern was then defined with a second resist layer and the metal lift-off step. Other alterations to the first-level personality included the removal of portions of lines and their subsequent rerouting by use of metal lift-off, simple deletions of unwanted lines, and the addition of an interconnection omitted on the mask. The wafer then received further processing, interlayer passivation, and second-level metal. Testing at the second level uncovered further defects on both the first and second levels of metallization. Line additions were made to the second level, and in three places, deletions for diagnostic purposes were made on buried first-level-metal lines. In all, microsurgery was repeated on the same wafer thirteen separate times, with electrical testing between each step to point out new faults. After the thirteen corrections were made, all significant portions of the chip had been made operational and functionally tested. The corrections were then incorporated in a second pass of the design, which yielded totally functional parts. The corrections and the associated problem determination and testing were accomplished over a period of three months. Assuming that the corrections could also be accomplished with only three passes through the usual process steps, between 10 and 15 months would be required to obtain functional chips.

The laser resist exposure capability has also been combined with reactive ion etching (RIE) techniques to extend the range of "after-fabrication" corrections into the depth of the silicon substrate. In a process developed by T. J. Schopen [15] of IBM, the laser was used to expose a small resist pattern on the surface of a multi-level metal-lurgy chip. RIE was then used to etch a narrow vertical channel through the personalization layers (two metal and

three quartz layers) and into the silicon substrate to a depth of $\approx 2 \mu\text{m}$. This process was used to provide an "air" isolation region between two previously butted transistor collectors which had their oxide isolation pattern omitted from the mask artwork. In this sense, the laser tool can provide three-dimensional correction capability to a chip design within the wiring structures and in the silicon substrate itself.

Results and conclusions

A comparison of the PLA/laser personalization approach discussed in this paper with an estimate of a conventional approach for the microprocessor we have designed is shown in Fig. 6. The result was a savings of 7.25 man-years and eight months of processing time, and a two-pass design (excluding yield-enhancement passes). The savings reflect the use of PLA macros, the dual hardware-software modeling approach, and the ability to correct an error on the chip with the laser tool. An LPLA chip for hardware modeling can be personalized rapidly with good yields and meets present and future capacity requirements. Laser microsurgery is being used to eliminate time-consuming passes through the process line, and as a diagnostic tool to aid in the rapid isolation and definition of errors.

We conclude that the use of PLA structures, particularly in the control portion of a VLSI microprocessor, improves design productivity, shortens design time, simplifies the physical design of the chip, and facilitates the implementation of design changes. The combination of a hardware simulator in conjunction with computer simulation greatly increases the probability that the VLSI chip will perform its intended function within specifications. The use of the laser tool is invaluable in correcting errors in VLSI chips. All of the above factors combine to provide operable first-pass hardware so that millions of instructions can be run to detect very subtle code sequence problems prior to final release of the masks to manufacturing. A two-pass design is therefore highly probable.

Acknowledgments

The authors gratefully acknowledge the efforts of W. D. Benedict and R. W. Bartoldus for developing the hardware modeling capability, and S. E. Bello and P. W. Oman for their contributions to the design of the microprocessor. We thank W. A. Lyons for his early laser experiments, R. C. Flaker and J. P. Maslack for the design of the laser-personalizable bipolar PLA, and the many individuals in the IBM Burlington and Kingston facilities who contributed to and supported this work. The contributions of A. J. Griest to laser microsurgery are much appreciated.

References and notes

1. Edward Stritter and Tom Gunter, "A Microprocessor Architecture for a Changing World: the Motorola 68000," *IEEE Computer* **12**, No. 2, 43-52 (1979); Robert Sugerman, "Computers: Our 'Microuniverse' Expands," *IEEE Spectrum* **16**, 32-37 (1979).
2. Masatoshi Shima, "Demystifying Microprocessor Design," *IEEE Spectrum* **16**, 22-30 (1979).
3. Arthur L. Robinson, "Are VLSI Microcircuits Too Hard to Design?" *Science* **209**, 258-262 (1980).
4. C. M. Davis, G. A. Maley, R. G. Simmons, H. I. Stoller, R. H. Warren, and T. E. Wohr, "IBM System/370 Bipolar Gate Array Micro-Processor Chip," *Proceedings of the IEEE International Conference on Circuits and Computers ICC80*, Port Chester, NY, October 1-3, 1980, pp. 669-673.
5. Signetics 8A1200 ISL gate array, Signetics Corporation, 811 East Arques Ave., P.O. Box 409, Sunnyvale, CA 94086.
6. M. Feuer, K. H. Khokhani, and D. A. Mehta, "The Layout and Wiring of a VLSI Microprocessor," Ref. 4, *loc. cit.*, pp. 678-679.
7. R. L. Golden, P. A. Latus, and P. Lowy, "Design Automation and the PLA Macro," *IBM J. Res. Develop.* **24**, 23-31 (1980).
8. Signetics 82S100 FPLA, address as in Ref. 5.
9. M. P. Feder, J. F. Smith, and H. E. Liberman, "Nitrogen-Pumped Dye Laser Tool for Fabricating LSI Connections," *Digest of Technical Papers*, Conference on Laser and Electro-Optical Systems II, San Diego, CA, February 7-9, 1978; Optical Soc. of America, 2000 K St. NW, Washington, DC 20036.
10. NEC μPB450D , Nippon Electric Company Ltd., Tokyo, Japan.
11. P. W. Cook, S. E. Schuster, and R. J. von Gutfeld, "Connections and Disconnections on Integrated Circuits Using Nanosecond Laser Pulses," *Appl. Phys. Lett.* **26**, 124-126 (1975).
12. Lawrence Kuhn, Stanley E. Schuster, Peter S. Zory, Jr., George W. Lynch, and James T. Parrish, "Experimental Study of Laser Formed Connections for LSI Wafer Personalization," *IEEE J. Solid-State Circuits* **SC-10**, 219-228 (1975).
13. J. Randal Moulic and Walter J. Kleinfelder, "Direct IC Pattern Generation by Laser Writing," *1980 IEEE International Solid-State Circuits Conference ISSCC80 Digest of Technical Papers* **23**, 210-211 (1980).
14. [®]AZ is a registered trademark of American Hoechst Corp., Route 202-206 No., Somerville, NJ 08876. Their licensed distributor is the Shipley Corp., Newton, MA.
15. IBM's Data Systems Division Poughkeepsie Technology Applications Laboratory, unpublished results.

Received August 29, 1980; revised December 24, 1980

Walter J. Kleinfelder and J. Randal Moulic are located at the IBM Data Systems Division Poughkeepsie laboratories, P.O. Box 390, New York 12602. The other authors are with the IBM System Communications Division, Neighborhood Road, Kingston, New York 12401.