Simplified Design Rules for VLSI Layouts

Richard F. Lyon, Xerox Palo Alto Research Center

The complexities of detailed layout design rules that change over the life of an evolving LSI/VLSI technology have forced a reconsideration of layout design methodology. Even during the design period for a single chip, a target process may scale, making the design inefficient or obsolete before the chip is produced. The single most important change in any generic process is the resolution, or feature size, which scales down steadily with time. Therefore, rules parameterized by a single variable (feature size, scale, or resolution) can provide longevity of designs, which will remain workable and efficient as long as the process does not change radically. These simplified, single-parameter design rules for generic processes are the subject of this article.

Simplified design rules enable restructuring of design, data translation and management, maskmaking, and wafer fabrication jobs to allow much more independence. Clever people who are not associated with any semiconductor company are already becoming independent IC designers, and are using simplified design rules similar to those described below.

The implementation system referred to in this article should be interpreted as a service that receives design files from a large number of designers and arranges for mask preparation, wafer processing, and packaging, much like the MPC79 system (Conway, Bell, and Newell, 1980). (See also the Silicon Foundry article in this issue.) The MPC79-type service provides a simple interface between the designer and the various required mask/fab services.

What Are Design Rules?

Design rules specify to the designer certain geometric constraints (dimensional inequalities) on his layout artwork (design file) such that the patterns on the processed wafer will preserve the topology and geometry of his designs (with high probability), so that functional digital systems can be built. Connected regions will stay connected, disjoint regions will stay disjoint, and electrical length/width ratios of field-effect transistors (for example) will stay reasonably close to the as-drawn values. These crucial properties cannot be guaranteed if the design rules are violated.

Design rules should always be specified in terms of the layout artwork (*i.e.*, as rules for the designer), rather than in terms of some intermediate artifact (*e.g.*, pattern generator tape, reticle, or mask). The implementation process (data handling, mask making, and wafer fabrication) should produce final patterns as close as possible to the drawn patterns, in terms of effective electrical widths and other functional properties (not apparent optical width). Some IC fab lines have not taken this approach in formulating their rules; therefore, their rules may be harder to interpret.

Stricter rules and additional statistical information will be needed for "analog" circuit design (which requires better control of parasitics and electrical parameter tolerances). Because access to such information often requires a captive and cooperative fab facility, the independent digital designer should try to avoid needing this data.

Motivating Strategy for Lambda Rules

A single parameter, λ (the Greek letter *lambda*), has been chosen by Mead and Conway (1980) to represent the "resolution" of a typical lithographic step in wafer processing. Design rules can be formulated in terms of λ, using simple rules for the interpretation of "resolution."

Basically, λ may be thought of as a bound on the width deviation of a feature on the final processed wafer from its ideal as-drawn size; it may also be thought of as a bound on the misalignment of any feature (feature-center) on one mask level from its ideal position relative to a feature on any other mask level. In the worst case, these effects combine to cause *feature-edges* on different mask levels to deviate as much as 2λ from their ideal relationship.

Typical feature-width and alignment control for a process should be much better than the worst case implied by the value of λ. Line widths, which may be critical to high-performance devices, should be much better controlled for certain mask levels than this discussion may imply.

Design rules are usually stated as minimum distances between specified feature-edges. Most rules for the required separation of feature-edges may be derived from a pair of simple "meta-rules," stated here:

"Fatality" meta-rule: *Relative movement of feature-edges by amounts less than 2λ should not be obviously fatal.*

"Degradation" meta-rule: *Relative movement of feature-edges by amounts less than 1λ should not cause significant degradation in performance.*

Thus, to preserve the topology of features on one layer, we have the immediate rules of 2λ minimum width (separation of opposite edges) for any part of any feature, and 2λ minimum space between separate features. However, these "2λ" rules may not be strict enough for all layers, for various reasons relating to the process. In processes with many layers, stricter rules (larger minimum sizes and clearances) typically apply to later steps (upper layers) and to relations between layers that

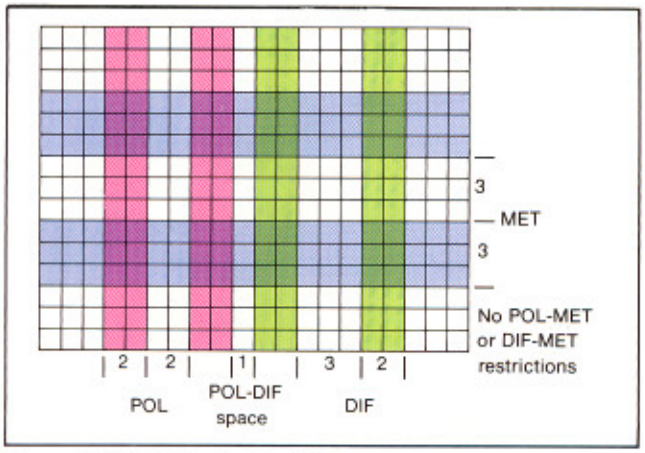

Figure 1. NMOS Conductor Width and Spacing Minimums

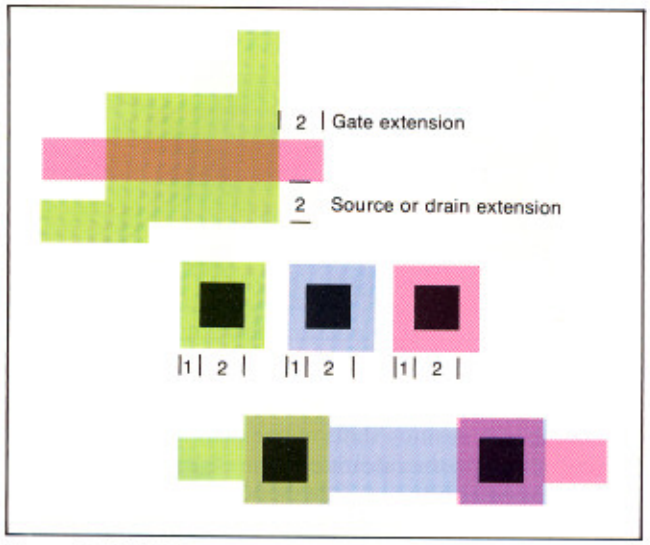

Figure 2. NMOS Transistor Rules
Figure 3. NMOS Contact Size Rules and Typical Contacts

are "far apart" vertically.

It is possible to draw features on any layer which are not topologically significant and which have edges closer than 2λ to each other (narrow stubs and notches); these should be avoided, whether or not they are considered to violate the rules, since they cannot be counted on to be resolved on the finished wafer. Such features, and other subtle geometric problems, present a stumbling block to strict formalization of the rules. For example, because sharp corners cannot be resolved reliably, should the rules include a minimum radius for rounded corners of features? What about internal angles (corners of complemented features)? Accepted practice is to allow corners no sharper than right angles, and not to rely on sharp corners for functional or electrical properties.

Example: The Basic Silicon-Gate nMOS Process

In the simplest silicon-gate nMOS process, there are three mask levels for conductors (red, green, and blue; or POL, DIF, and MET; where POL over DIF also makes enhancement-mode field-effect transistors with POL as the gate) and one mask level for opening contact windows through the insulating oxide between the layers (black, or CUT). See Mead and Conway (1980) for a description of the process (including implants for depletion-mode transistors, described below). Design rules for this basic process and some of its varations are discussed in detail below.

Width and Spacing Rules for Conducting Layers

The POL layer (polysilicon, usually the best-controlled layer) should obey the basic "2λ" rules for widths and spaces. The other two conducting layers were found by Mead and Conway to require more width and or space than the "2λ" rules. Figure 1 shows width and spacing rules for the conducting layers.

On the DIF layer (source and drain and interconnect diffusions and channel regions), extra spacing is required (in most processes) to prevent inadvertent conduction through depletion regions, which, at high voltage, can extend significantly beyond the normal edge of a diffused region. Therefore, 3λ is the DIF spacing rule.

The MET layer (aluminum metal interconnect) is the last circuit layer to be patterned, and must reliably cover the non-planar terrain left by structures of the previous layers; therefore, to be consistent with a typical level of conservativeness in the industry, Mead and Conway have chosen 3λ for both minimum width and minimum spacing of MET regions.

POL and DIF lines which do not meet intentionally to form a contact or a transistor should be kept separated by 1λ . If they should shift between 1λ and 2λ , and begin to overlap one another, the only effect is a narrowing of the diffused width and a (possibly large) mutual capacitance through the gate oxide between them where they overlap. The MET layer does not interact with DIF and POL except at contacts (and via parasitic capacitance).

Rules for Transistors

Where POL crosses DIF, the source and drain diffusion is masked by the gate oxide remaining under the POL region; the source, drain, and channel are thereby self-aligned to the gate. The resulting MOSFET has minimum length and width determined by the minimum width of POL and DIF lines, respectively.

POL and DIF are constrained in two other important ways where they form a transistor (see Figure 2). First, POL must cross DIF completely, or the transistor will be shorted by a diffused path between source and drain; thus 2λ of POL overlap is required beyond the edges of the DIF region. Similarly, DIF must extend beyond the POL gate so that diffused regions definitely exist to carry charge into and out of the gate region; thus 2λ of DIF extension is needed to preserve the source and drain regions.

Rules for Contacts Between Layers

The only other interactions between layers involve contacts. Several rules assure adequacy of contact area between layers and safe separation of contact cuts from uninvolved conductors. Features on the CUT mask specify windows in the intermediate oxide to expose POL (if it exists), or DIF (if there is DIF and not POL), so that subsequently placed MET will contact the exposed silicon area. Contacts are made using a minimum of 2λ square of CUT, with a big enough conductor region (POL or DIF) under it to assure a reasonable area of coverage, as shown in Figure 3. Thus, 1λ of DIF or POL

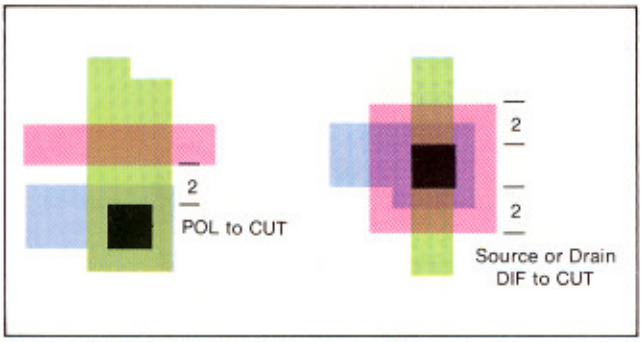

Figure 4. NMOS Contact Clearance Rules

surrounding a CUT is the rule to prevent the contact area from decreasing with position shifts of less than 1λ, and to prevent etching through to the substrate by the edge of the conductor. The same overlap rule applies to the MET region that overlies the CUT window, to assure adequate contact area and periphery. The aluminum step coverage at the contact window edge is a potential source of problems such as metal migration and breakage; the overlap and large periphery length help assure low resistance and high current-carrying capacity.

To prevent a CUT window from accidentally allowing contact to an uninvolved conductor, 2λ clearance is required, as shown in Figure 4. Mead and Conway (1980) state this rule for POL-to-CUT clearance, as it relates to placement of a transistor gate near a contact to the source or drain diffusion. DIF-to-CUT clearance should also be observed in the unusual case of a contact to POL over a channel region, which must be separated adequately from the source or drain diffusion. Some processes suggest no use of CUT over a channel region, to avoid a significant threshold shift due to the different work function of the resulting silicon-aluminum alloy gate material. When CUT over channel is not used, sufficient DIF-to-CUT clearance will always be assured by combinations of other rules.

The final complication of the rules for layer interactions involves a structure known as a "butting contact," which allows a compact path between DIF and POL through a small link of MET, as shown in Figure 5. A "butting contact" is made by removing one edge of POL overlap from a valid POL-CUT-MET contact, and then overlapping that contact with a valid DIF-CUT-MET contact, abutting the CUT squares to make a 2λ-by-4λ CUT region. Keep in mind that the gate oxide under the POL layer prevents contact with DIF below it (where there is not conductor anyway, just channel). With this contact, either the DIF-MET contact region or the POL-MET contact region might be diminished even for misalignments less than 1λ . The contact resistance can be about twice as high, in the λ-shift case, as in a series connection of the simpler DIF-MET and POL-MET contacts. Thus, this contact does not satisfy the "degradation" meta-rule in uses where its increased resistance causes performance degradation. Where contact resistance is critical (such as in a DC power distribution network or in series with a low-impedance output driver), more conservative contact structures should be used.

Rules for Depletion-Mode Implants

For better speed-power performance, and to allow inclusion of effective on-chip drivers without a separate supply voltage,

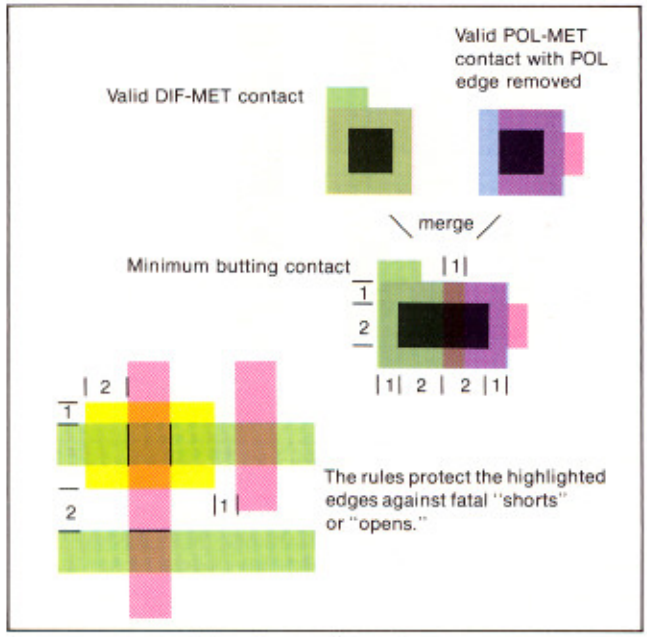

Figure 5. Butting Contact Rules
Figure 6. NMOS Depletion-Mode Implant Rules

most nMOS processes have added another mask level (yellow, or IMP, a donor implant) and associated process steps to allow depletion-mode transistors (*i.e.*, transistors with negative thresholds). Features on the IMP mask define regions in which the donor impurity will be implanted.

Figure 6 shows the rules for IMP overlap and clearance. It is important to cover the whole area of the POL-DIF intersection that defines a transistor. Moving the edge of the IMP feature relative to the transistor edges can cause parasitic enhancement devices in series or in parallel with the desired depletion device. A series device would be fatal (the combined threshold could be positive); therefore, 2λ of overlap is required beyond the source and drain edges. A parallel device would simply increase the effective resistance a bit by narrowing the negative-threshold region; therefore, only 1λ of overlap is required in the direction of the gate width.

It is just as important to observe clearances between IMP features and desired enhancement devices, especially to avoid a parallel parasitic depletion device which would be an effective short. Therefore, the clearance rules are also asymmetric: 1λ from an enhancement-device source or drain edge, and 2λ from the side of the enhancement-device channel.

To avoid having to remember and understand these asymmetric rules, Mead and Conway formulated a simpler set of compromise rules for 1.5λ of overlap and clearance on all device edges. Unfortunately, these rules are not as uniformly conservative as are the rest of Mead and Conway's rules; and because IMP rules are not critical to density in commonly-used gate logic, fab lines have not worked much on improving the resolution of the associated steps. Therefore, conservative designers will use "2λ-everywhere" IMP rules for all new designs that use modern nMOS processes.

Rules for Buried Contacts

For increased density, some processes have added another mask level and more processing steps to create DIF-POL contacts without using the MET layer. These "buried con-

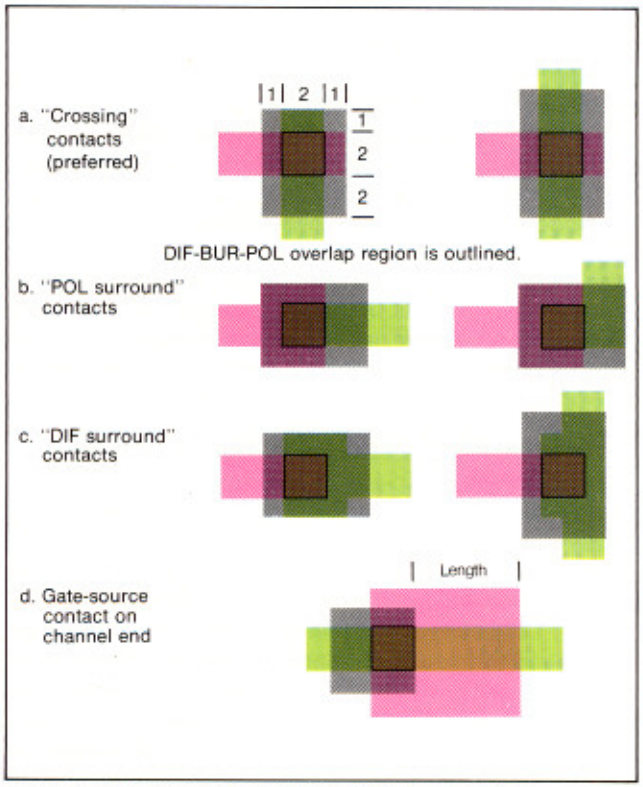

Figure 7. Rules and Variations for Buried Contacts

tacts" (gray, or BUR) have rather extensive roles, due to the effects of etching away thin oxide. Features on the BUR mask specify regions where thin (gate) oxide will be etched away before polysilicon is deposited on the wafer. The deposited silicon grows directly on underlying crystalline silicon, rather than on oxide, in regions where BUR and DIF coincide. When the polysilicon is patterned by etching, the etch may also form "pits" in the substrate in regions where BUR has removed oxide over DIF and there is to be no POL. In a later step, when POL and exposed DIF regions are doped by diffusion or by implantation, the contact region will be made conductive (but not as conductive as contacts to MET), and the pit region becomes part of the diffused conductor into the contact.

The rules for the BUR layer and its relationship to DIF and POL are based on insuring a reasonable area of contact, the absence of a possible transistor in series with the contact, the protection of nearby transistors from having their gate oxide perforated, and the prevention of other unwanted DIF-POL contacts.

As with the metallized contacts, a 2λ square of DIF-BUR-POL is the minimum area for reasonable contact conductance. This square can be formed several ways, as shown in Figure 7. In the "crossing" contacts of Figure 7a, the overlap area is independent of misalignments less than 1λ, whereas the overlap of the "POL surround" and "DIF surround" contacts of figures 7b and 7c can be reduced by a shift of DIF and POL edges, to the same extent as in the butting contact discussed above (*i.e.*, they may violate the "degradation" meta-rule, depending on their usage). The "DIF surround" contacts of Figure 7c are the least preferred form (due to the large area of "pit" that forms where there is BUR and DIF and not POL). The "crossing" contacts of Figure 7a are the most preferred (due to the alignment-insensitive overlap

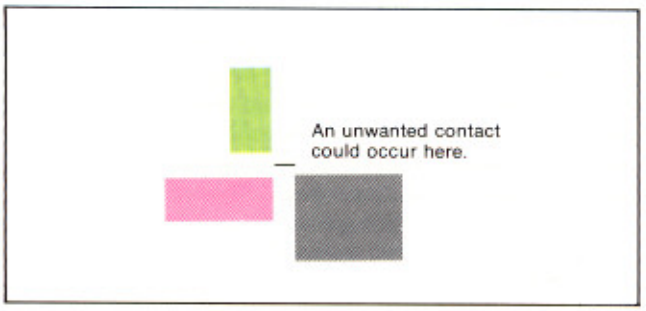

Figure 8. An Illegal Configuration of POL, DIF, and BUR

area). In any of the three forms, at least 1λ of POL and or DIF extension (as appropriate) is required beyond the overlap region.

To ensure complete etching of thin oxide from the contact region, the obvious rule is that BUR should extend 1λ beyond all sides of that region. Unfortunately, it is not quite that simple. To prevent the possibility of a parasitic enhancement device in series with the DIF line into the contact, the BUR region must extend 2λ in the direction of the DIF conductor, as shown in Figure 7. Extension of only 1λ beyond the other edges of the contact region is acceptable.

Another form of buried contact, useful as the gate to source connection on a depletion pull-up, is shown in Figure 7d. Here the DIF-BUR-POL overlap region is defined in one dimension by POL and BUR edges, so BUR extension does not apply at the edge defined by BUR. The transistor channel length should be measured from the BUR edge. (Butting contacts can also be used this way, in which case the channel extends all the way to the POL edge.)

Clearance rules for BUR can be extensive. We state here a simplified but overly restrictive "2λ-everywhere" rule. To prevent inadvertent gate-oxide shorts and other unwanted DIF-to- POL contact, BUR must be separated from unrelated DIF or POL regions by at least 2λ. Unconnected DIF and POL within a single BUR region (being shared by several contacts) must also be separated by at least 2λ.

More "complete" rules would allow DIF-to-BUR clearance and POL-to-BUR clearance of only 1λ, but would require many rules to keep the three layers from coming together to form an unwanted contact. For example, the rules would have to prevent the configuration shown in Figure 8; this is hard to do with constraints on feature-edges, even using *derived* features (*e.g.*, the logical AND of several drawn layers).

There are usually no further restrictions on BUR and MET, so buried contacts are often placed under metal lines. Some processes advise against MET over BUR, but these processes are not likely to survive, because they do not provide enough density advantage to justify the extra mask and processing steps.

Effect of Process Variations on Design Rules

Process options such as buried contacts, extra threshold values, second-layer polysilicon (either gate or interconnect), or second-layer metal (either interconnect or light-shield for imagers) may provide dramatic improvements in density or performance for a given process resolution, in applications that need or can use those features. Once the processing is understood, design rules can be formulated for any of these

features, following the techniques shown above.

Some industry experts have suggested that as processes move into the VLSI domain, simplicity and planarity will contribute increasingly to yield, and that therefore, buried contacts and additional layers of conductors are not likely to survive. (However, other experts think that integration will move into multi-layer, three-dimensional structures.)

The designer's choice of optional features depends on many factors, including process availability, design and maskmaking cost, desired design longevity, and the density advantage for the particular design.

Other, more subtle process variations are being tried by various fabrication and research groups. The goal is usually to increase the layout density of a particular class of circuits (especially memory and PLA structures) without scaling down the basic processing resolution, or lambda. Thus, the detailed design rules for any fab line will reflect their current process type. Some variations will cause *increases* in some of the minimum-distance rules. Thus, designs based on the simple rules stated here may actually become *larger*, because a larger value of λ will have to be used.

Alternative Design-Rule Models

We have used the "feature-edge separation" model of design-rules, and have used some informal statements of the applicability of these rules. Other possible models may be easier to formalize. For example, the "area constraint" model is a system of rules for areas of regions derived by logical operations on, and expansions of, other regions, starting with the mask features. One possible buried-contact spacing constraint might be that the area $(DIF+1)\cap(BUR+1)\cap(POL+1)\cap\neg(((DIF\cap BUR)+1)\cap((BUR\cap POL)+1)\cap((DIF\cap POL)+1))$ must be empty, where "*region* +1" means the 1λ expansion of *region* (the infinite union of disks of radius 1λ with centers in *region*), ∩ means "intersection" (AND), and ¬ means "complement" (NOT).

Multi-parameter design-rule sets are also interesting, and will be more usable when designs are represented simply as topology, and when layouts are created automatically in accordance with the rules. Separate parameters for alignment tolerance, alignment sequence, line-width control for the various layers, etc., might then be used to get an optimum match to a particular process, without requiring the designer to change his design at the topological level.

Auxiliary Rules for Macroscopic Features

Processing lines often pose auxiliary geometric constraints beyond those needed to ensure topological correctness. These rules do not always scale with the other rules; macroscopic rules relating to bonding pads are a prime example. Large MET regions are used as bonding pads for wires connecting the chip to the package pins. Typical rules for these regions are 100-micron (0.1-mm or 4-mil) minimum length and width, and 100-micron minimum spacing, with 25-micron clearance to active circuitry (conductors not connected to the bonding pad).

An additional mask layer (brown, or PAD) is usually used to pattern windows in a passivating layer of oxide (overglass) or nitride, to allow the bonding wire to contact the bonding pad. The features that specify such windows should stay at least 10 microns inside the MET region, to avoid etching other oxides. If the masking and etching steps were done carefully enough, this clearance would need to be only 2λ.

Other "rules" often describe how to construct scribe lines, alignment marks, identifying numbers, etc. These "non-rules" (or rules for non-design items in the "starting frame") are really descriptions of additional artifacts to be included in the final mask layout, along with the actual system design. These "non-rules" are outside the scope of this article, and need not be known to the independent designer who has access to a good implementation service that will take care of the additional artifacts. For designers operating without an implementation system, or for someone who wants to operate such a system, the information in *Guide to LSI Implementation* (Hon and Sequin, 1980) is invaluable.

Stretching and Shrinking to Create Masks

Depending on types of resist, exposures, etching techniques, etc., the final features on the wafer may systematically deviate in width by as much as several microns from the width on the mask. The designer does not want to keep track of the inherent process stretches and shrinks of various fab facilities; he just wants to get what he drew. Therefore, stretch-and-shrink compensation for systematic process effects is the responsibility of the implementation system that transforms a design file into masks for a particular fab line. The implementer must cause appropriate stretches and shrinks when translating a design file to patterns on glass plates (masks), to compensate as closely as possible for what the fab process will do. This compensation may be done by software modification of the design data, or by optical methods at the maskmaking facility, according to the implementer's specification.

The fab line does not need to know what design rules the designer used, and the designer does not need to know what stretches and shrinks a fab line invokes in printing mask patterns on silicon. But the implementer needs to know both, to select a suitable fab line (or to tell the designer what minimum value of λ is available), and to specify masks correctly.

Geometry specified by a layout interchange format (for example, CIF2.0) is usually defined as "final desired geometry." The conversion to mask geometry is typically done during translation into a machine-dependent output format such as a pattern-generator tape, or a raster-scan format (for machines such as electron-beam mask-exposure systems).

Various subtle problems may come up in maskmaking, such as blooming due to multiple flashing on pattern generators, or reduced diagonal clearance between rectangles that have been stretched by simply increasing length and width. These are not concerns for the designer, but are problems that the implementer should keep in mind when trying to answer the designer's question: "Can you handle maskmaking and fabrication services to get nMOS chips built for me with λ equal to X microns, with depletion loads, and without buried contacts?" That is all the designer needs to know.

Design Rules for Other Processes

Single-parameter design-rule sets, based on the same concept of λ and the same meta-rules, are being developed for other processes. Rules for CMOS/SOS are nearing final form at

JPL, and have benefited from discussions with many SOS fabricators; rules for several bulk CMOS processes are being developed at Hewlett-Packard and at UC Berkeley. Rules for I^2L layouts are being developed by Dick Oettel of Boeing. Rules and meta-rules for a class of junction-isolated bipolar digital/linear processes (with more complicated structures than those in our "planar" view presented above) are being developed by the author and others at Xerox PARC, guided by the detailed design-rule development work of Glaser and Subak-Sharpe (1977).

Tricks for Production Parts

It is typical practice in the semiconductor industry to reduce the die size of a mature part by optically shrinking the masks or reticles as the process resolution improves. However, the stretches and shrinks due to processing should be reconsidered carefully, because improved etching techniques, etc., can cause severe changes. It will usually be wiser to scale the design file, and make a fresh pass through the implementation phase.

Further stretches and shrinks, that intentionally cause a working IC to deviate from the original design, can be used in some cases to improve performance. This violation of design intent will typically be done only in a production environment, with a captive and stable fab line; the aim might be to move the "simplified" design closer to the ultimate performance possible from the process, or to shift its position slightly along a speed/power trade-off curve (for example, by shrinking polysilicon lines to shorten channel lengths).

Conclusions

This discussion of design rules has pointed out several ways in which the mere *formulation* of the rules can affect the methodology of design and construction of integrated systems. The "simple rules" paradigm discussed here is in conflict, in several ways, with the traditional "vertically integrated" paradigm of the semiconductor industry. The basis for the "simple rules" paradigm is summarized below:

1. Rules are simple geometric constraints that ensure preservation of topological intent on any of a family of "generic" processes; *i.e.*, the rules are *fab-line independent*.
2. Single-parameter rules *scale* trivially, and remain relatively efficient over the lifetime of a generic process. Chip and subsystem designs remain useful for long enough to make the "design library" concept valid.
3. To take advantage of the longevity of designs possible with these rules, designers should *avoid complicated process options* that may be temporary aberrations from the generic process.
4. Products that achieve market success through *process innovation* are the province of vertically integrated IC manufacturers, while products that achieve success through *design innovation* are becoming the province of independent designers.
5. The designer can deal with a *clean interface* if an implementation system is available to deal with other services. The implementation system will only offer truly generic process capabilities (in terms of multi-source availability), which are compatible with one set of simple rules.
6. The designer is *in control* of all levels of his design, from architecture to geometry, and need not be restrained by arbitrary extra rules, for example on what circuits he can try; yet he need not understand maskmaking and wafer fab to specify desired final patterns correctly.
7. Beyond actual circuitry, the designer must know other constraints such as minimum bonding-pad sizes (which the implementation system will specify if the system is to do the wire bonding); otherwise, the designer can do as he wishes, *independent of the usual practice at the fab line* that will be used.

Simple interfaces and explicit statement of the design constraints, separated into layout design rules and other categories, are already enabling a dramatic increase in the number of practicing integrated-system designers.

The geometrical design rules discussed in this article are just one example of the types of rules a designer uses to help manage the complexity of the design task. Other sets of "design rules" include electrical (*e.g.*, metal conductor current-handling capability), logic (*e.g.*, avoid use of NAND gate structure), and others. Together, these rules form a *methodology* which helps the designer cope with the complexity of system design. One exemplary methodology that encompasses circuits, logic techniques, storage devices, system timing, physical communication, etc., is presented in Mead and Conway (1980).

References

Bell, A.G. November, 1980. "The Role of VLSI Implementation Systems in Interfacing the Designer and Fabricator of VLSI Circuits," *Proceedings of the International Telecommunications Conference*.

Conway, L.A., A. Bell, M.E. Newell. 2nd Qtr., 1980. "MPC79: A Large-Scale Demonstration of a New Way to Create Systems in Silicon," *LAMBDA*, Vol. I, No. 2, pp. 10-19.

Hon, R.W. and C.H. Sequin. 1980. *A Guide to LSI Implementation*, Xerox PARC Technical Report SSL79-7, Palo Alto, CA.

Glaser, A.B. and G.E. Subak-Sharpe. 1977. *Integrated Circuit Engineering*, Reading: Addison-Wesley.

Mead, C.A. and L.A. Conway. 1980. *Introduction to VLSI Systems*, Reading: Addison-Wesley.

About the Author

Richard F. (Dick) Lyon holds a B.S. in engineering and applied science from Caltech ('74) and an M.S. in electrical engineering from Stanford University ('75). He has a wide range of experience in communications-related fields including studies on the theory and implementation of enumerative codes for channel spectrum shaping, design and testing aspects of arithmetic circuits for digital signal processing, and the characteristics of coded, noncoherent channels. At Stanford Telecommunications, Inc., from 1975 to 1977, he designed and built a variety of analog and digital real-time communication and navigation signal-processing equipment.

Since April 1977, Dick has been with the VLSI System Design Area of Xerox Palo Alto Research Center. There he has been applying his digital system building experience to the problem of special-purpose architectures for signal processing with custom LSI and VLSI, and to speech recognition and synthesis.