CHAPTER 4
PRACTICAL CONSIDERATIONS IN IC PATTERN PREPARATION

Implementing custom integrated circuits is an expensive process, both in terms of dollars and effort. The monetary expense may be so high that it is impractical for a designer working on a research budget to have chips made for a single circuit or system. One answer is to amortize the cost of the masks over several projects by putting more than one design on a single chip. If as many as ten different designs can be placed on a single 1/4 inch square chip, the cost of masks may only increase from about $7000 to $8500. A minimum run of about 20 3" wafers (Si gate NMOS) costs roughly $3000 at any of a number of fab firms; such a run yields 1600 chips. The cost in manpower is similarly high – it might require one man-month to administrate such a multi-project chip, in addition to the design effort. Good design tools can reduce the human effort necessary, but of course building the requisite tools is a time consuming process. Luckily, in many research groups manpower is easier to come by than real dollars.

This chapter focusses on the process of creating multi-project chips (*MPC's*).

4.1 Merging Many Projects

Over the past two years much progress has been made in streamlining the process of merging many individual IC designs into the specification for a single mask set. Early efforts were carried out in and *ad hoc* fashion, with many costly errors made. With each multi-project chip, techniques and support systems have been improved. A recent chip set, called MPC79, was organized by researchers at PARC-SSL using a new form of "VLSI implementation system" that enabled the remote-entry of designs by a large community of designers scattered throughout the country. MPC79 contained 82 designs on 12 die types, and its total implementation time was 29 days. For more details on that chip set see [Conway, et al 1980].

At least in the near future, most first-time multi-project chips will be assembled by hand by a few coordinators. As in any endeavor involving many participants (the designers), there are a great number of purely administrative details to be attended to. It is imperative that a convenient means of *communication* be available. Messages and design files must be transported on a regular basis, and as final deadlines approach, the lack of a rapid means of communication can become an insurmountable barrier. In a small, intra-departmental effort, written messages and a central design facility are probably adequate. For an MPC involving geographically diverse locations and differing design systems, some form of electronic message and file transport capabilities are essential.

In the remainder of this section, we will outline the steps of the merging process, elaborating as necessary.

1. Creation and checking of design files

   This process may be left up to the individual designers, or may be monitored by the coordinators. Large efforts, like MPC79, favor the former approach, while smaller chip sets allow the coordinators to actively aid in the checking process. In the absence of design rule checking programs, designers should be assigned to cross-check someone else's design. Many trivial errors can be eliminated in this way. We have found that there must be a *central depository* for design files, where the official version of each design resides. Participants are only allowed to change files with the consent of the coordinators, who are responsible for any actions that an update requires. These activities include such things as updating logs, running new checkplots, or perhaps notifying human design rule checkers that a new version is in place.

2. Submission of final designs

   If the coordinators have been involved in checking designs, this step simply means freezing the state of the central depository. If not, when the appointed deadline is reached the designs should be collected and moved to a protected location so that no more changes can be made.

3. Conversion to standard format

   Once the designs are frozen, they must be converted to a standard format (e.g. CIF 2.0). Alternatively, the coordinators may require that designs be submitted in a standard format, perhaps with other restrictions applied (for example, they may require that each design have the lower left-hand corner located at 0,0). From each standard design, information about location and bounding rectangle can be obtained.

4. Merging

   The frozen designs are now packed into one or more chips (see Section 4.2 below). Typically 3 to 10 projects will fit onto one chip; each must include a number of standard features (see Section 4.3), which are merged in as any other project. Designs expressed in a graphics language are easily merged by translating and rotating the various designs, and then concatenating the design files (note that symbols may have to be given unique prefixes in order to avoid naming conflicts).

5. Conversion to mask format

   Each multi-project chip is now available as a single large design file. These files must be converted into PG or e-beam format and written on a mag tape. At the same time, features may be over- or under-sized to meet fab requirements (see Section 3.2.4). Instructions detailing die sizes, location, step and repeat distances, mask materials, and more (see Appendix A) are written. They, the mag tape, and a check are sufficient to have wafers made.

4.2 Physical Constraints

The designs are arranged to minimize the area of the chip bearing in mind a number of important factors. Optical equipment limitations at the mask house make it difficult to generate masks for chips larger than 10mm by 10mm. Defect-free reticles become harder to generate as the chip size increases, thus it is disproportionately expensive to make masks as the chip gets large. If e-beam masks are available, these size restrictions are not so critical. The 10mm x 10mm size limit is somewhat misleading because of an additional restriction imposed by current packaging technology. The *cavity size* of a standard 40 pin dual inline package (DIP) is about 7.5mm x 7.5mm, thus projects should be limited to this size unless there is access to special packages. The 10mm x 10mm chip must be subdivided to meet this constraint by placing *interior scribe lines* between projects; these scribe lines must extend all the way across the chip (and thus across the wafer), that is, interior "tees" are not allowed.

When packing the designs, it is wise to consider the wire-bonding apparatus that will be used to connect the circuit pads to the package pins. Depending on the type of equipment used (see Section 5.3), it may be difficult to wire bond configurations where wires must run long distances or at acute angles. A particularly troublesome layout places a small project in the middle of a large chip. In this case long wires must be run to the periphery; they are subject to sagging and subsequent short circuiting. Long wires also protrude above the top of the package, because of the trajectories that they must follow.

The overall size of the chips is impacted by the expected *yield*, which is the fraction of the IC's that function correctly. As the *active area* (the area containing active devices but excluding empty space, bonding pads, etc.) grows the yield decreases geometrically. Typical yields for a 6mm by 6mm circuit, assuming standard defect densities, are about 20-40%; in industry, yields much below this figure are not acceptable profit-wise. Designers in a research environment may well be able to tolerate low yields since a project is rarely dependent on all of the devices on the chip working. Huge projects may still be feasible, as a yield of a only few percent gives the designer enough chips to verify his design, measure the performance and demonstrate feasibility.

4.3 The Starting Frame

At some point the designer has several complete IC designs that are ready to be turned into chips. Before his design can be realized, several important details have to be taken care of which are not part of the actual circuit design process. Among these are the physical placement of several projects and test patterns on the multi-project chip plus the the addition of some extra features required by the fabrication line. Test patterns are useful for an evaluation of the quality of the wafer processing, for checking standard circuit parameters, as well as post mortem debugging should a chip fail to perform correctly. Most of these relatively fixed, universally required features (e.g.

CD's, fiducial and parity marks, alignment marks and scribe lines) can be collected at each research site and grouped together in a *starting frame* (see also [Mead & Conway 1980]). This starting frame provides a set of "symbols" (or whatever construct is appropriate in the local design system) that can be combined with the individual design projects to provide the masks for a complete multi-project chip.

The most important features that must be added to the net circuit is the set of alignment marks, which are needed to register subsequent layers on the IC with one another. Alignment marks take many forms:

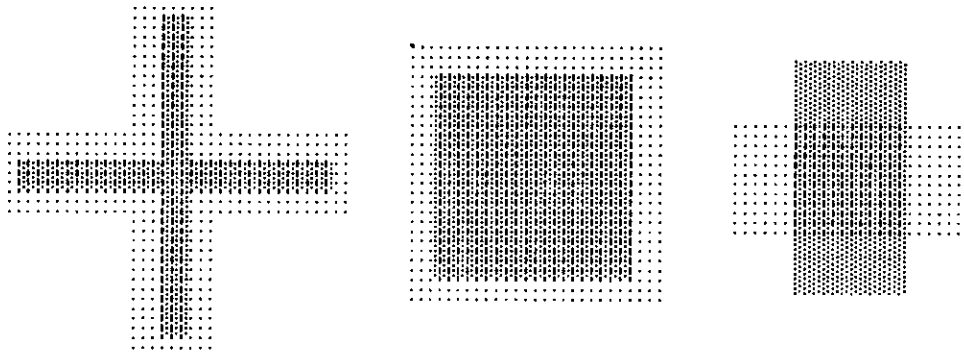

light area on wafer dark area on mask

– but their purpose is the same. There is little magic in designing alignment marks; in fact, almost any reasonable features will do. However, a carefully designed set can mean the difference between good devices and those that are only marginal.

When the designer is deciding on which alignment marks to use it may help to consider the following scenario. The fabrication line operator puts a partially processed wafer onto the movable (x, y, rotation) stage of the alignment machine. The next mask is held over the wafer and the operator looks through a microscope from above. First of all, is it possible to locate the alignment marks? The designer can help by providing "black and clears" on which he has indicated the location of the alignment marks. A line or box enclosing the marks may also draw the operator's attention amidst the confusion of the other features. After the marks are located, is it possible for the operator to successfully align with them? Consider the case where the alignment marks consist of a large square on the wafer and a small square on the mask.

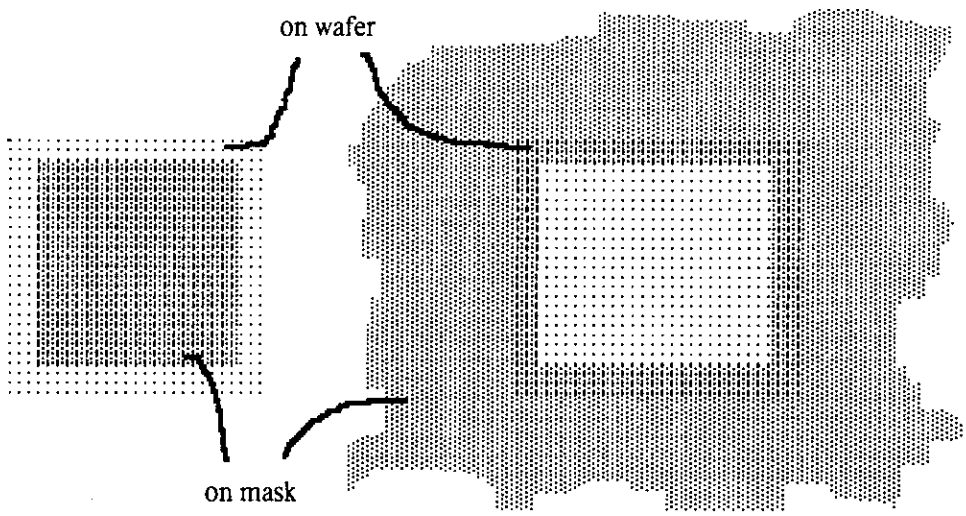

Clear Field Dark Field

The operator is supposed to center the small square over the large one. This system is fine if the small box is opaque on a clear field (see the explanation of working plate polarity in Chapter 5), but not when the small box is clear on an otherwise opaque mask. In the latter case the designer should have an alternate version of the alignment marks for opaque field masks, then one or the other set will work for the mask polarity used in the particular fabrication step. Another alternative is to design a set of marks that can be used regardless of mask polarity.

As the operator tries to line up the alignment marks, is it obvious which small square goes over which large one? A one square shift is certain to be disastrous. Some type of *gross* alignment mark should be provided to prevent shifting; again there are many alternatives – an enclosing box or a simple square which is superimposed over one already on the wafer, or even numbering the small/large square combinations. Furthermore, it is important to eliminate ambiguity regarding the layer to which the current mask is aligned. For example, in figure 4.3.1 there are two large squares in place on the wafer, one in diffusion, the other in polysilicon. The operator has two small squares on the contact cut mask to line up over the two large squares. Unfortunately the large ones are not exactly in line because of a small misregistration introduced in a previous step. The operator must decide whether to align to the diffusion or the poly feature, or perhaps to split the difference between the two. Whichever course the operator chooses may affect device operation. The *designer* should make this decision by providing only one pair of marks. (The actual set of alignment marks chosen for our starting frame are discussed in section 6.1.)

After the integrated circuits are fabricated one must determine whether or not they are functioning correctly. While the designer has the option of simply powering up his circuit and seeing if its input/ouput behavior is correct, a more satisfactory test method might use several *test structures* included on the chip (see Chapter 5). Simple structures like inverters can answer yes/no questions (Were the wafers processed to a minimum level of competence? Do individual transistors work?) and thus indicate whether more complete testing is warranted. More importantly, test

Figure 4.3.1 Alignment of Contact Cut Mask

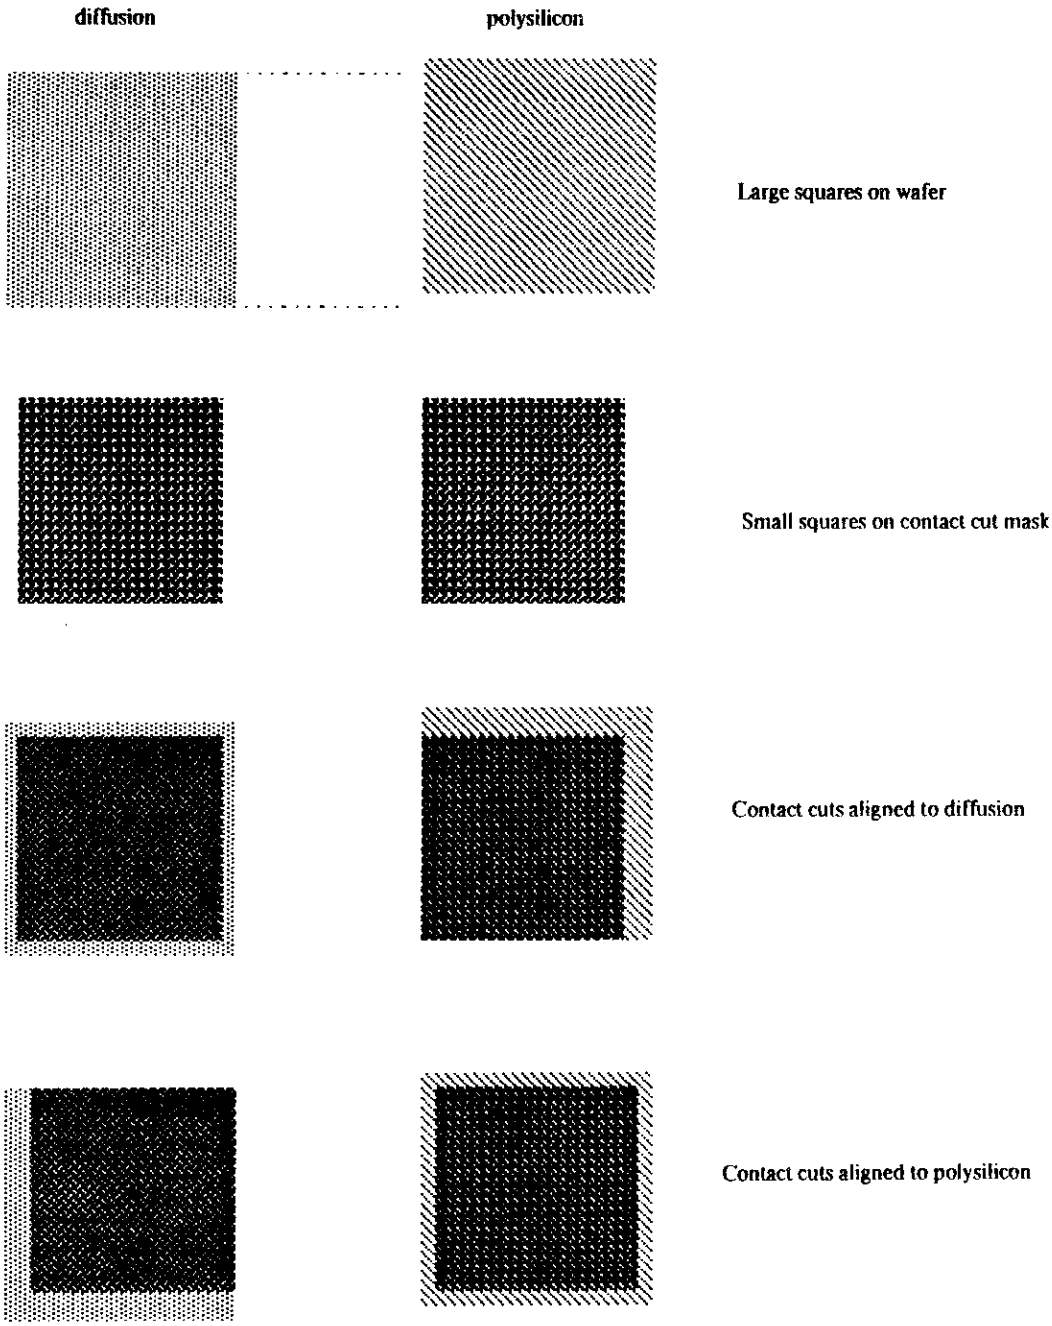

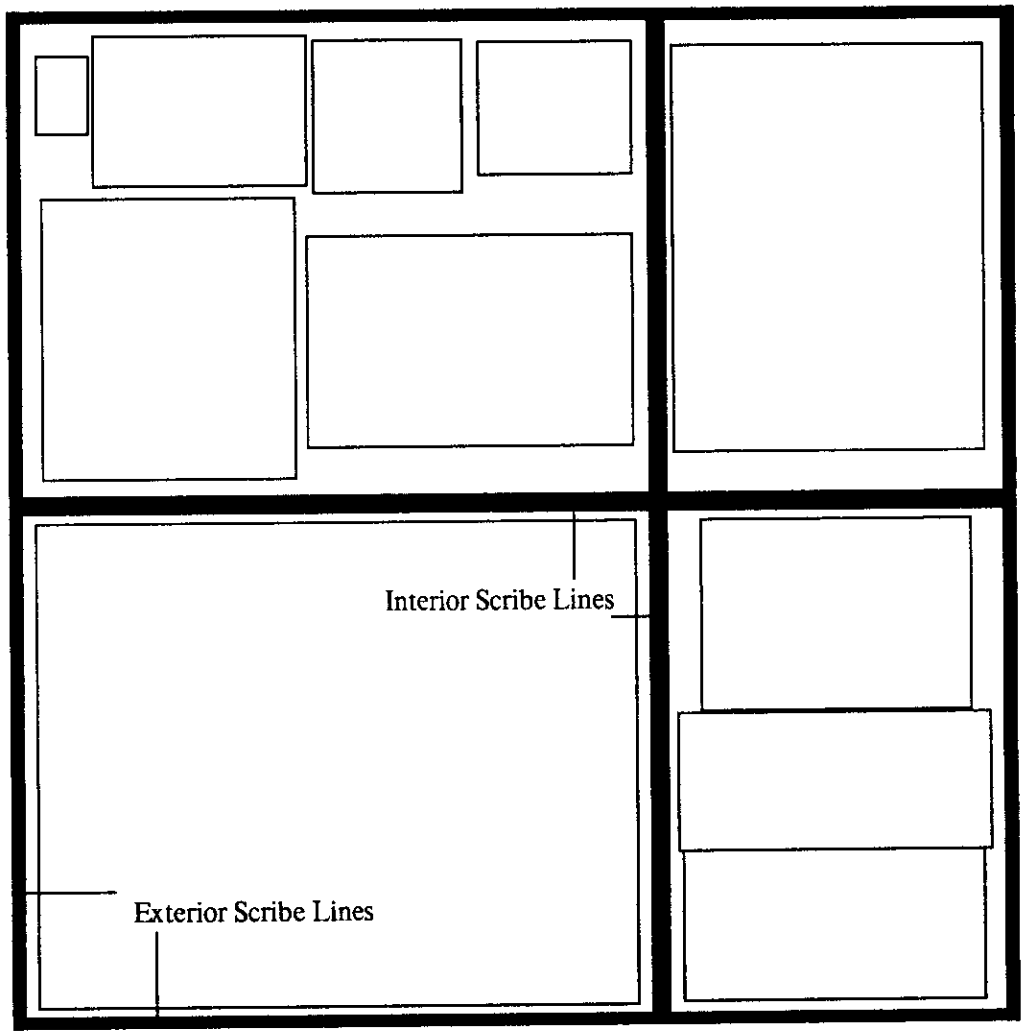

Figure 4.3.2 Overall View of a Multi-Project Chip

patterns can provide information which is useful in determining why a batch of chips does not work, or performs poorly. Properly designed test patterns can show wafer processing problems, or eliminate this cause, narrowing the search to the area of design errors.

Ultimately the various designs and the starting frame have to be combined into a single IC description. This involves merging of a number of files, usually in a geometric design language, into a single file containing all of the integrated circuit designs. Fiducials and parity marks may need to be added outside of the area occupied by the designs and the starting frame. The chip pattern is repeated on the surface of the silicon wafer many times; 2", 3", and 4" wafers are commonly available – a 3" wafer holds about 45 10mm by 10mm chips. *Exterior scribe lines* (see figure 4.3.2) are placed around the periphery of the area occupied by the project set. The purpose of the scribe lines is to provide a "lane" down to the silicon substrate in which the diamond-tipped scribe tool (see Section 5.2) will ride. The wafer will be broken into chips (also called *dies*), as defined by the scribe lines, following fabrication.

References

[Conway, et al 1980]
L. Conway, A. Bell, M. Newell, R. Lyon, R. Pasco, *Implementation Documentation for the MPC79 Multi-University Multiproject Chip-Set*, Xerox-PARC Report, January 1980.

[Mead & Conway 1980]
C. A. Mead and L. A. Conway, *Introduction to VLSI Systems*, Addison-Wesley, Reading, MA., 1980.