7. An Example Project Chip and Starting Frame

The steps outlined in this guide have been applied to build an actual IC at Xerox PARC during the summer of 1978. Four summer students and four employees of Xerox PARC contributed designs, ranging from wafer processing evaluation testers to novel arithmetic and memory circuits. Ten such projects have been combined into a Multi-Project-Chip by Bob Hon. Figure 7.1 shows the overall layout and one mask layer of the chip. The chip includes a set of alignment marks and line-width testers laid out by Bob Hon and Dick Lyon and a general processing test chip laid out by Rick Davies. In this chapter the features of the starting frame are first presented, followed by a discussion of the general test chip. Finally in section 7.3, Robert Baldwin, the youngest student on the team, describes his experience as a novice who had to design his first IC without prior knowledge of the subject.

7.1 The Starting Frame

The multi-project chip is divided into two parts separated by one internal scribe line so that both parts are small enough to fit inside the cavity of a 40 pin DIP. Further, the layout is enclosed by exterior scribe lines placed around the periphery. The exterior lines differ from the interior lines in that the former are missing the outside "shoulder" (figure 7.1.1), the exterior lines of one pattern are completed by the overlap with the exterior line of the next chip. The scribe lines were designed simply to provide access to the Si substrate for the scribe tool to contact during wafer separation (see section 6.1).

The alignment marks designed were intended to unambiguously indicate which layers are to be aligned relative to each other. The marks consist of a number of "squares" and "fortresses".

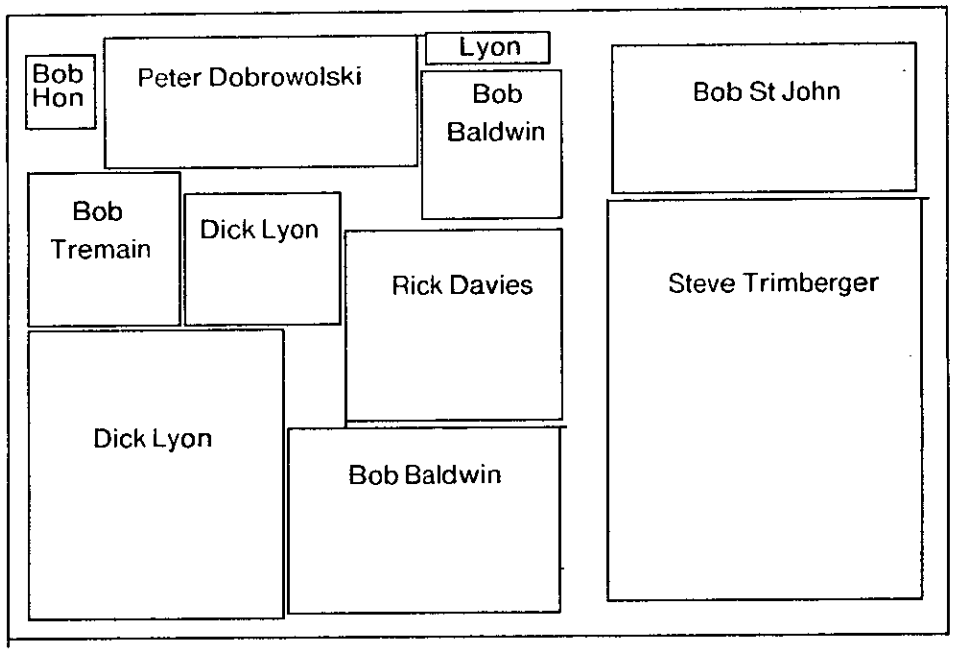

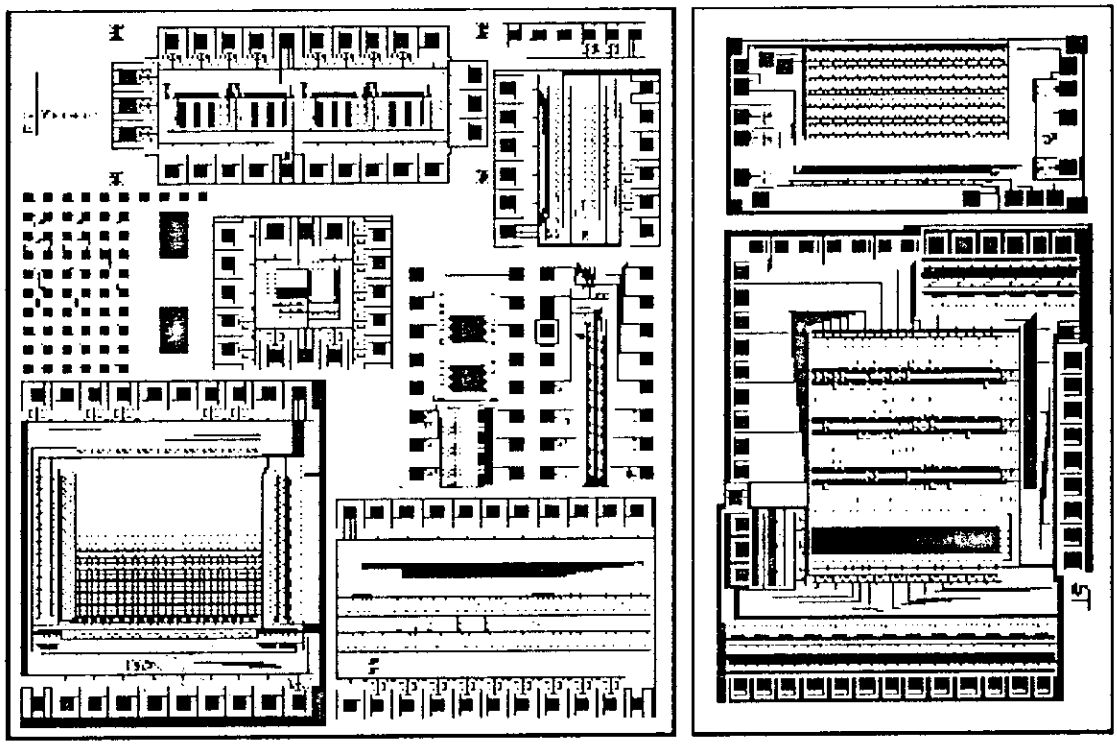

Figure 7.1. The Summer 1978 PARC Multi-Project Chip

Figure 7.1.1 Scribe Line Profile

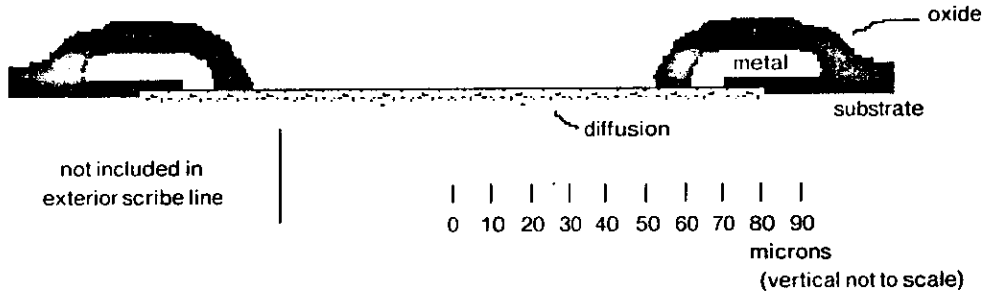

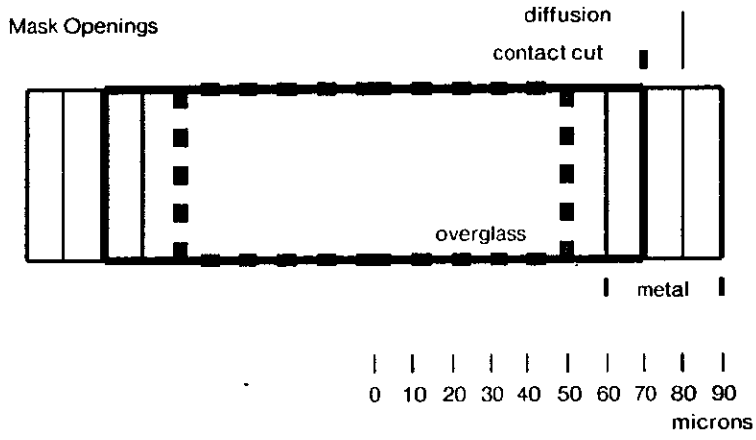

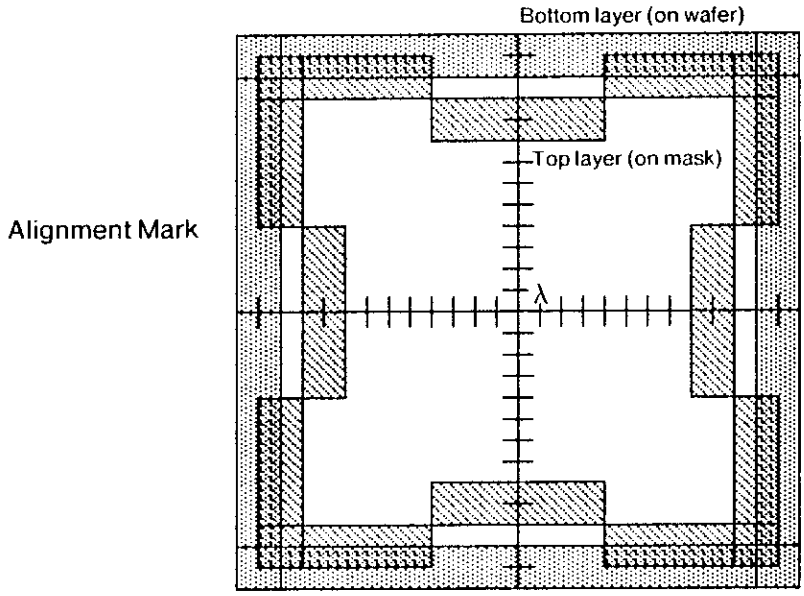

A square mark is placed on those layers which will serve as reference layers for masks in following fab steps. Each square has a corresponding fortress, located on a different mask, which will be aligned over it during the appropriate step (see figure 7.1.2). Each layer includes a large rectangle around the alignment marks to help the operator to locate them and to insure that the sequence is not shifted. The features are lines rather than areas, permitting the operator to align edges relative to one another. This makes the marks usable for clear as well as opaque working plate fields.

The fortress/square pairs are used in a left to right progression; a digit (omitted in the figures below for clarity) is placed in each fortress to indicate when it is to be used. A fortress is always aligned over a square and there is never more than one fortress per mask. The alignment sequence has the depletion mode implant, buried contacts (when used), and the polysilicon layer both aligned relative to the diffusion layer. The contact cuts are aligned relative to the polysilicon since there appeared to be more tolerance to misalignment between contact cuts and diffusion. The metal aligns to the contact cuts and the overglass to the metal layer. The following illustration represents the way an operator might align the mask for the polysilicon layer to the wafer. The pattern on the partially processed wafer is:

Figure 7.1.2 Alignment Marks for Mask Layers

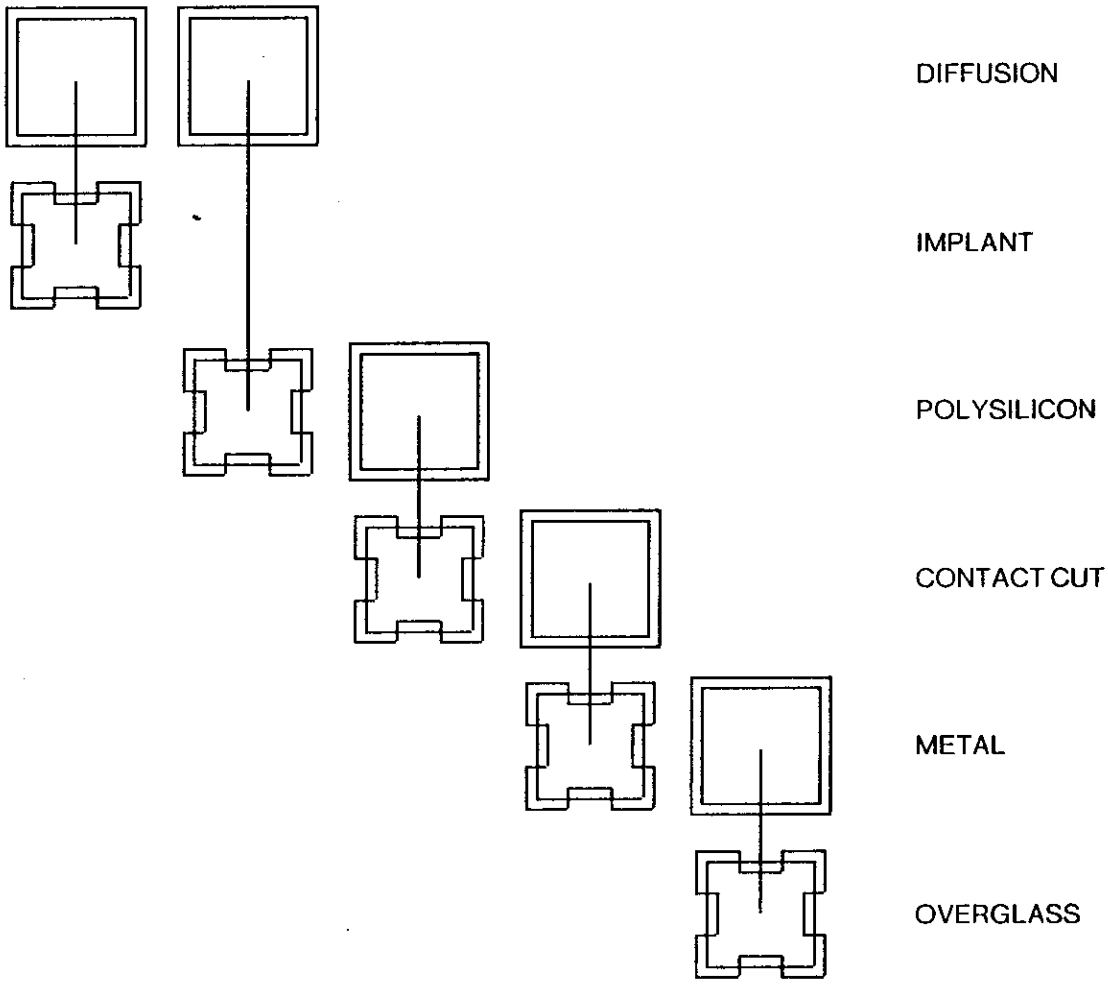

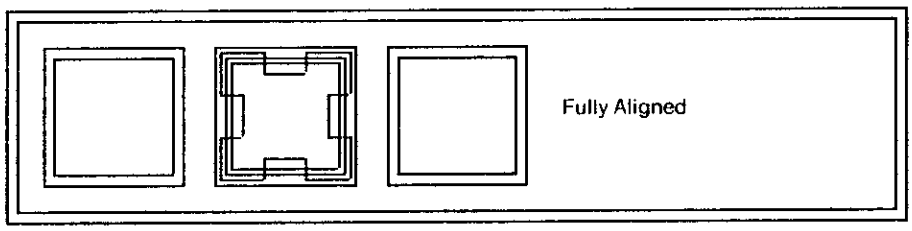

The third square will be used to align the contact cut layer; a fourth square is placed on the wafer during the contact cut step and will be used to align the metal layer.

The critical dimension marks are simple crosses made of lines. The line widths vary from layer to layer, and are typical of the feature dimensions found on the particular layer.

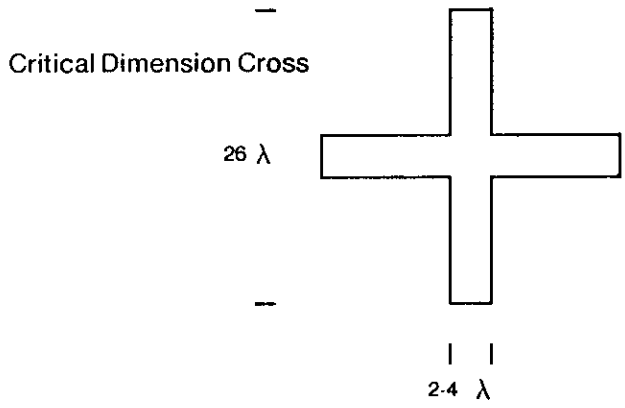

A set of features used to monitor the quality of the working plates and fabrication process was also included. This *etch test pattern* consists of a set of nested "L"'s with the same spacing between L's as the width of the feature; two different sizes were placed to check the quality of the mask and the quality of the photolithographic process.

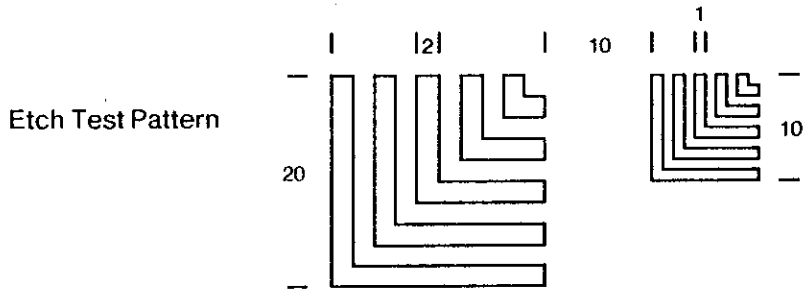

all dimensions in λ (= 3 microns in 1978)

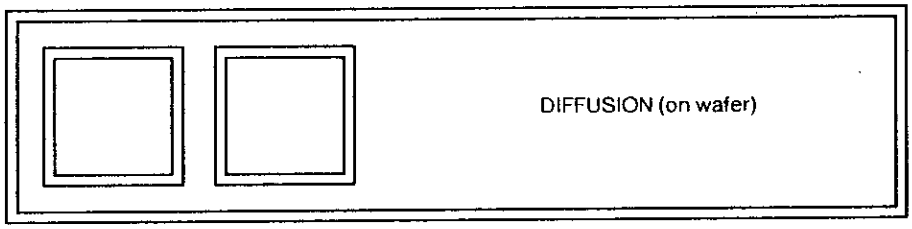

The operator must now align the polysilicon mask which contains one fortress and one square.

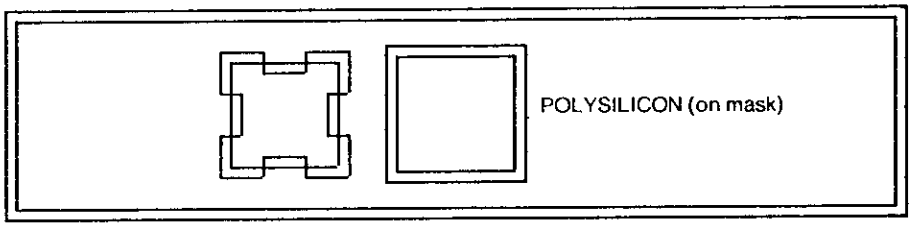

The diffusion and depletion implant steps are already complete, thus the first fortress/square pair has been used. The operator lines up the second pair using the gross alignment mark for guidance. The following diagram shows the mask in partial alignment with the features in place on the wafer.

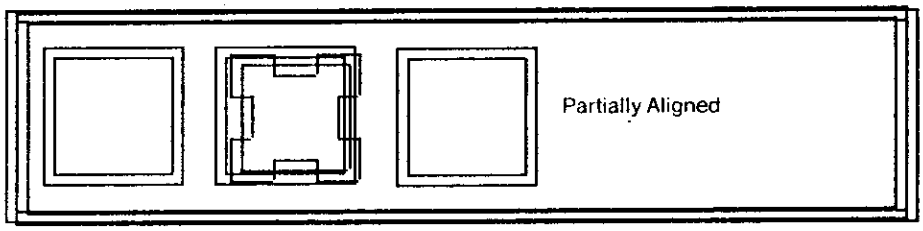

Final corrections are made using the fortress/square pair.

Measurements on each mask layer provide a check on the dimensional correctness of the working plates while measurements on the wafer are used to verify that the fab line performed as anticipated (e.g. that lines over- or under-etched as expected).

Appendix A contains a copy of the information sent to the mask house.

7.2 Test Patterns

[section contributed by Rick Davies, Xerox PARC]

The starting frame contains a number of simple test structures to answer the following two questions.

> Was the wafer properly processed? Specifically are all the layers properly patterned, are gate oxide and deposited oxide films of acceptable dielectric integrity, are contact holes properly opened, etc.?
>
> What are the first-order device and circuit performance characteristics such as transistor threshold voltages, extent of short- or narrow-channel effects [Dennard 1974, Wang 1978], polysilicon sheet resistivity, and inverter propagation delay obtainable with the process?

The test structures described here are general enough so that it is assumed that they will prove useful to most participants in a multi-project chip. This should not deter any designer from adding his own special test structures. It may be desirable in the future to add additional patterns to the common test structure, to test such parameters as the quality of the buried contacts connection of diffusion to polysilicon or the limits of wafer processing (e.g. At what spacing do metal lines begin to show bridging?).

The test pattern consists of several separate regions which are described below. It occupies 2 mm x 2 mm and has this general layout:

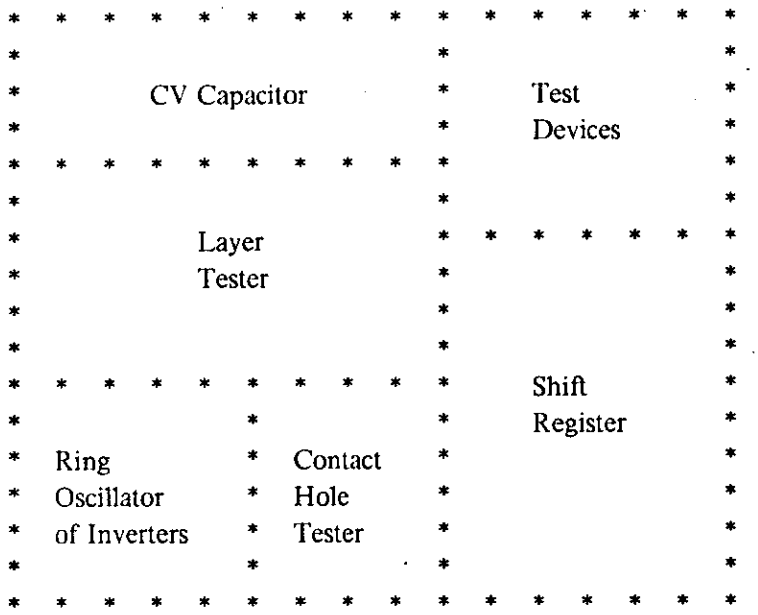

7.2.1 Layer Tester (Fig. 7.2.1)

This is a long serpentine metallization path that runs between two interdigitated metal combs and lies over a serpentine of polysilicon and active transistor area. It tests the following features.

a. *Metal bridging.* There is a 25,000 μm periphery of minimum-spaced metal lines (9 μm, 3λ spacing) between the serpentine and combs. Conductance between the serpentine and either comb indicates bridging caused by failure to properly image the pattern in photoresist and then etch it in the aluminum.

b. *Metal step coverage.* The 12 μm-wide serpentine passes 266 times over a 6 μm-wide region of polysilicon, over gate oxide, between 6 μm-wide source-drain diffusions. This should provide a good indication of step-coverage quality (assuming that the above bridging test passed), measured as a low end-to-end impedance (< 100Ω). Metal running over diffusion and polysilicon is the worst-case condition for metal step coverage; the limited solid angle provided by the evaporation source can make it difficult to transport aluminum to the vertical features in small-geometry device structures.

c. *Gate-oxide dielectric integrity.* The presence of approximately 100,000 μm^2 of polysilicon over gate oxide provides a test for pinholes and shorts between gate and transistor-channel area. This is equivalent to the gate area of about 1000 transistors of typical transistor geometry. To pass this test there must not be measurable conductance between the polysilicon and the diffusion.

d. *Deposited-Oxide dielectric integrity.* The presence of approximately 100,000 μm^2 of metal over active area (polysilicon gate or source-drain diffusion) provides a test for

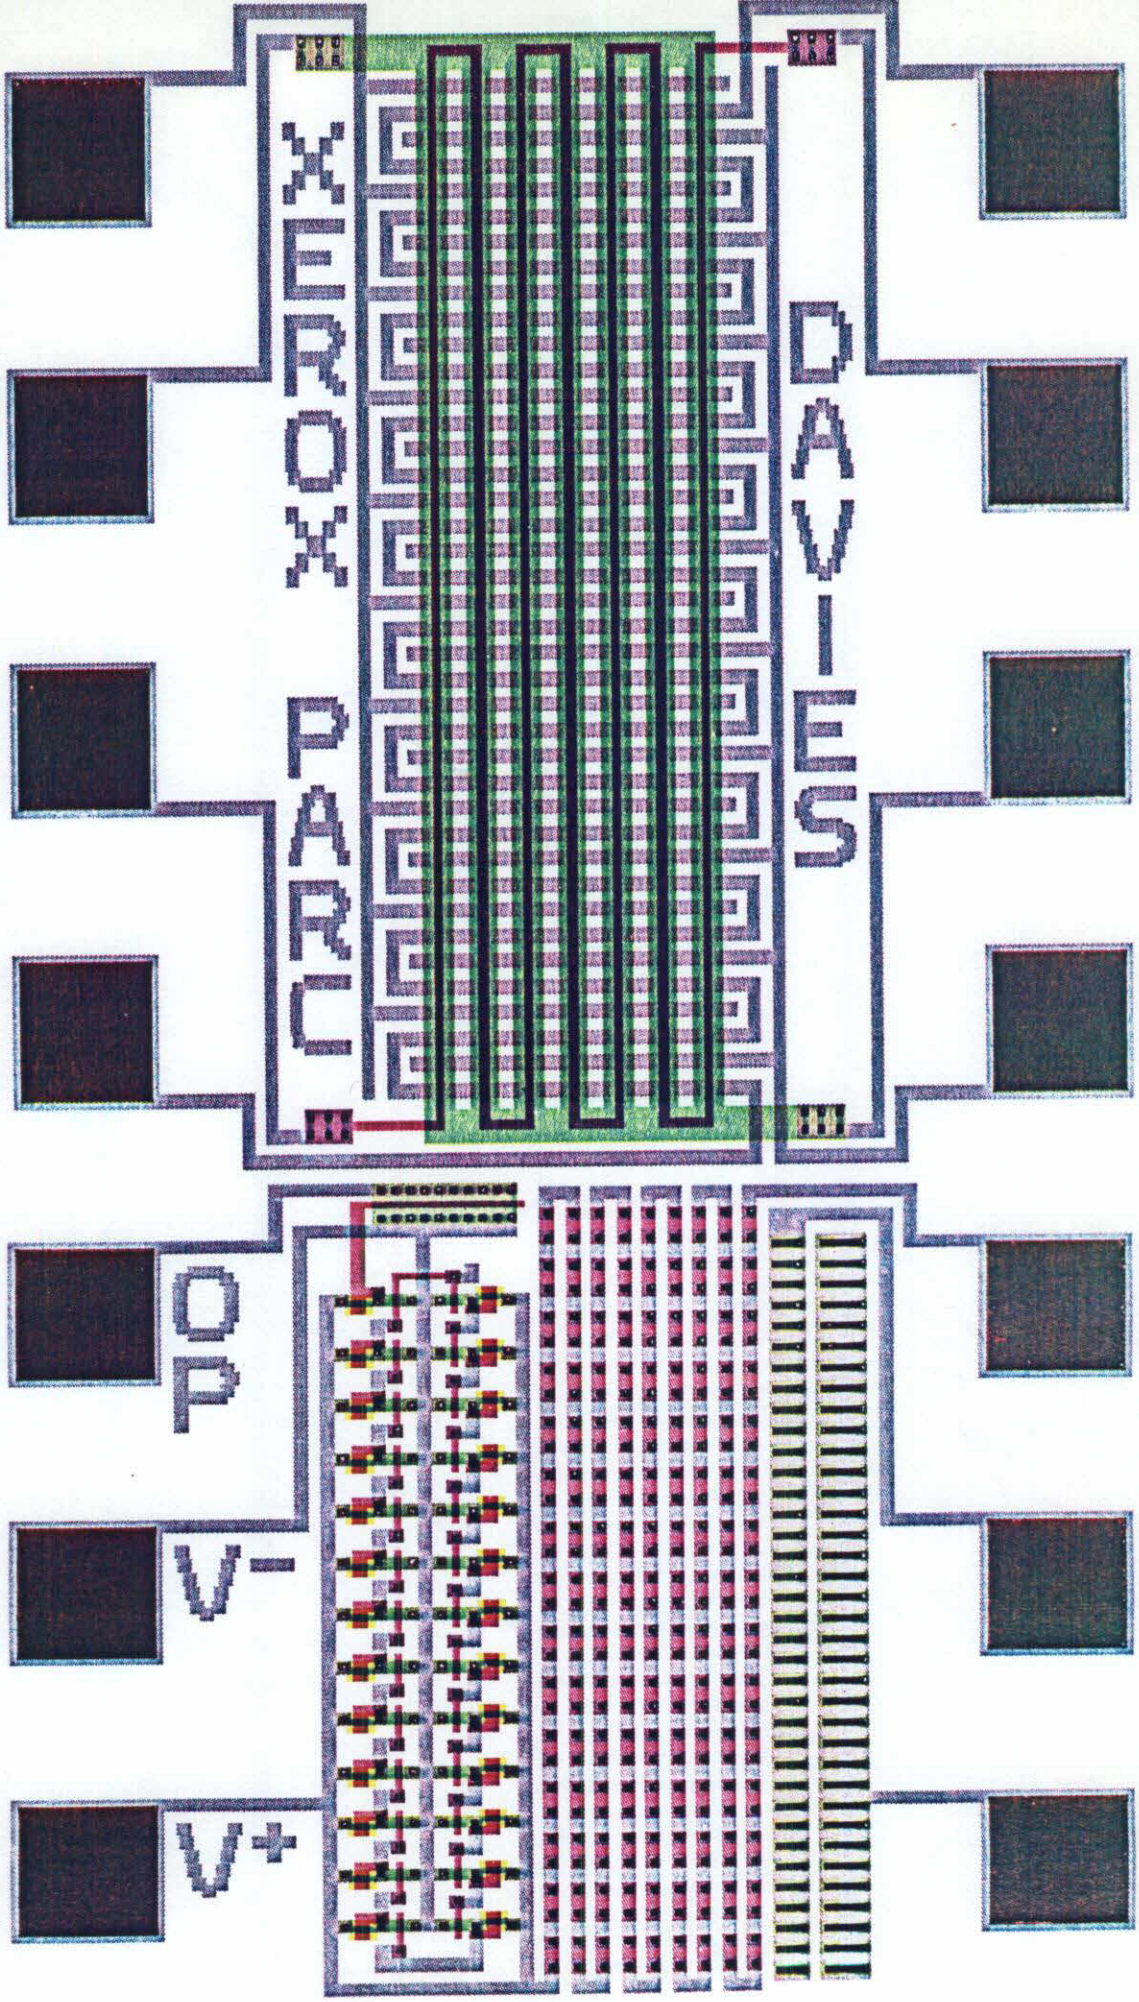

Figure 7.2.1 Test Structures

pinholes and shorts to the metallization; this could result from improper annealing of the metal causing spiking through the deposited oxide layer. No conductance should be observed between the metal and either diffusion or polysilicon.

e. *Polysilicon and aluminum sheet resistivity.* Although the most accurate resistivity measurement uses a Van der Pauw structure (separate forced-current and sensed-voltage terminal pairs), the present structure provides a quick estimate by inspection of the end-to-end resistance. About 700 squares are present in either level.

7.2.2 CV Capacitor.

A 200μm x 900μm MOS capacitor with a diffusion guard-ring is provided for analysis of the process parameters Q_{SS} (density of fixed charges at the oxide-silicon interface), N_{SS} (density of trapping states at the interface) and the gate oxide thickness [Grove 1967]. A minor amount of final wafer preparation may be required to form a suitable ohmic backside substrate contact for reliable measurements.

Measurements on this structure would allow separate determination of the implant dose and interface characteristics components of the threshold voltages of the enhancement and depletion NMOS transistors. This structure may also be used to test gate oxide dielectric integrity; 180,000 μm^2 are present.

7.2.3 Contact Hole Tester

The following two contact-hole tests are incorporated on this test pattern (Fig. 7.2.2).

a. A series connection of 270 metal-polysilicon contacts (6μm x 6μm) tests for failure to make electrical contact between metallization and the underlying layer. A measured resistance significantly above the expected impedance corresponding to the parasitic 135 squares of connecting polysilicon indicates poor quality ohmic contacts. This could be caused by improper imaging of the pattern in the photoresist (in particular a scum residue might have been left in the bottom of a hole), improper etching of the oxide, or by metal breakage around the rim of the etched contact hole.

b. A special tester consisting of 2 columns of 36 metal-diffusion contacts compares contact holes which are properly centered over a diffused area and others which overlap the field oxide region. The latter is used to test whether one could overlap contact holes onto the field oxide, in order to save the area otherwise consumed by alignment tolerance. Because the phosphorus-doped SiO_2 (sometimes called *P-glass*, see section 3.2 for a description of the fabrication process) deposited before contact hole definition etches much faster than does the thermally grown field oxide, the contact holes should open before the field region is etched through.

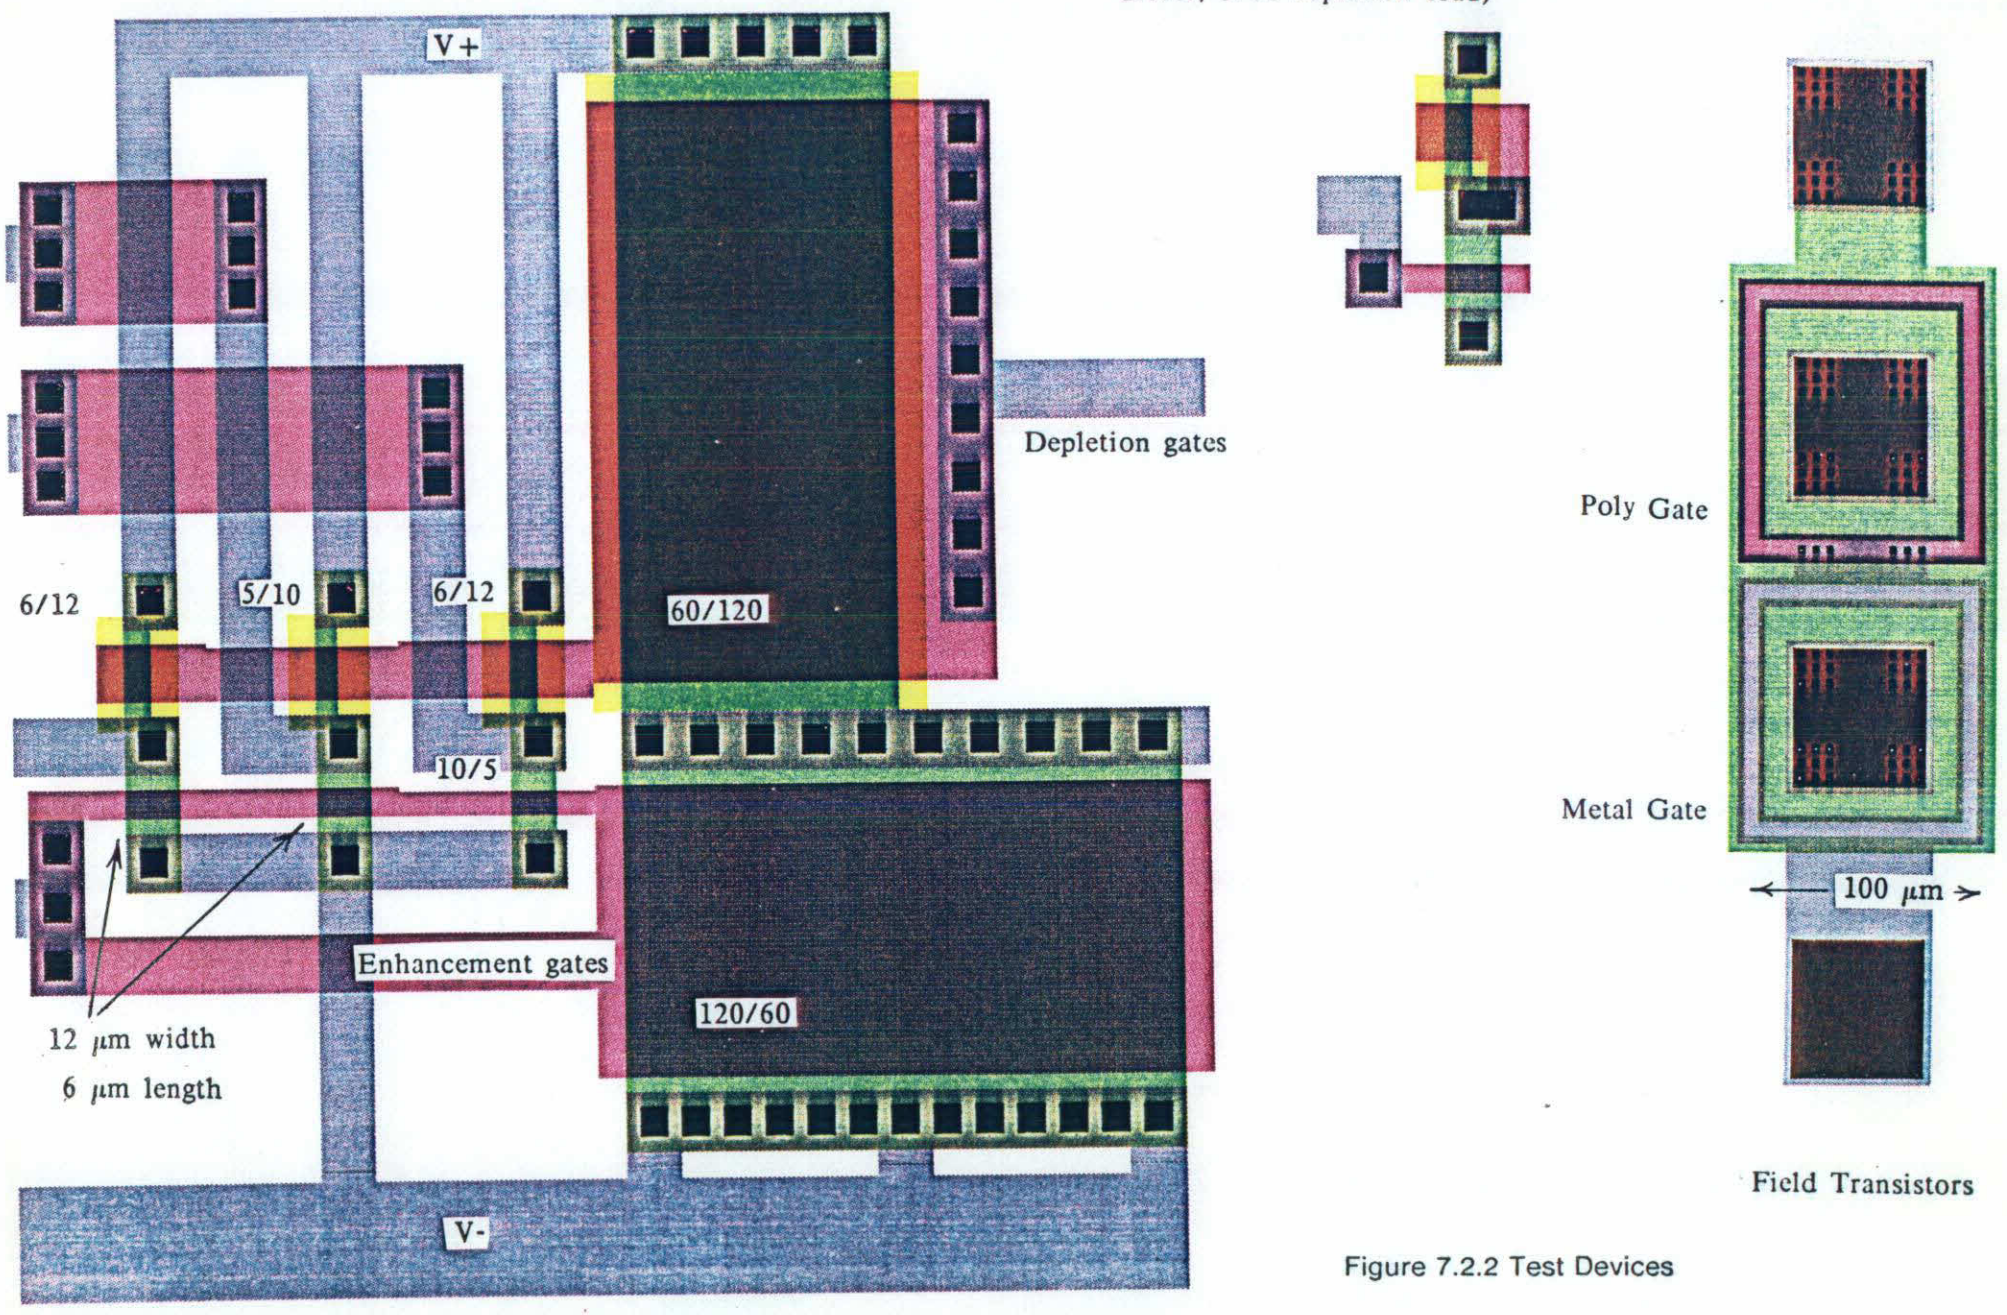

Figure 7.2.2 Test Devices

The two columns have identical bottom-wall diffusion area, diffusion periphery, and contact hole area over diffusion; the right one differs from the left only by the incorporation of a strip of field oxide in the middle of the contact hole. If the two columns reveal the same leakage characteristics to the substrate, then this overlap technique is probably acceptable.

7.2.4 Discrete devices.

Four enhancement and four depletion mode transistors are provided for dc testing (Figure 7.2.2). They are organized as four inverters for convenient transfer curve analysis, with uncommitted gates for the depletion-mode transistors to allow full testing of those devices. To minimize the number of bonding pads, the enhancement-gates are shared, as are the depletion-gates; all enhancement-sources and all depletion-drains are also shared. The four inverters are:

a. 12μm-width/6μm-length enhancement NMOS with 6/12 depletion NMOS, forming a standard 4:1 inverter with 2λ layout rules.

b. 12/6 enhancement and slightly narrowed 5/10 depletion load device. One expects a higher threshold in the narrowed channel because the channel potential is raised by the increased influence of edge effects. The use of scaling [Wang 1978] (which involves altering the fabrication process) would permit this smaller layout without disturbing the dc characteristics.

c. 10/5 enhancement and 6/12 depletion load devices. Short-channel effects should cause a lowering of the threshold voltage and produce increased output conduction in the enhancement device [Dennard 1974].

d. 120/60 enhancement and 60/120 depletion load devices. These devices should permit one to check device characteristics with little interference from peripheral effects.

Two transistors with closed layouts -- one with a metal and the other with a polysilicon gate -- on thick field oxide, are provided to test for isolation-region channeling or other parasitic leakage. A threshold voltage above about 25v should exist on each device. In both cases the transistor gate electrode overlaps the source and drain regions. The poly-gate structure makes gate-oxide devices at source and drain that are in series with the field-region under test.

7.2.5 Ring Oscillator of Inverters.

This structure (Fig. 7.2.1) tests ac device performance in a probe environment without hindrance from the inherent parasitic capacitances. Because it is an actual circuit, it should provide more reliable and directly usable information than making separate detailed ac device measurements and then using circuit equations. The ring oscillator is 25 stages long to provide a low frequency output signal. The average propagation delay is one half of one twenty-fifth of the inverse of the loop natural oscillation frequency. Inverters with 12/6μm enhancement drivers and 6/12μm depletion loads are used; a buffer/inverter taps the loop, giving one of the 25 loop inverters a fanout of two. The buffer in turn feeds an output transistor with W/L of 20/1; this may be used as a common-source output driver or could be used as a source-follower if desired.

7.2.6 Shift Register.

This 33-stage circuit is similar to the above ring oscillator, with the addition of a passgate in front of each inverter (they are 24/6μm enhancement drivers and 6/12μm depletion loads; the passgates are 6/6μm devices). The passgates are bussed in two phases that alternate between inverters (16 passgates per phase). The 33rd passgate is brought out to a separate bonding pad so that one can open the shift register loop. The complement of this signal is applied to another passgate which connects the shift register to a separate, inverting input buffer.

This configuration provides the following operations:

a. With both passgate phases and the control passgate line high, the circuit implements a 33-stage ring oscillator using passgate signal transmission. The measured average propagation delay may be compared to that of the simple ring oscillator (described in section 7.2.5).

b. With the control passgate line low and alternate clocking of the passgate phases, the circuit acts as a shift register which may be loaded with an arbitrary bit pattern. Raising the control passgate line while continuing two-phase clocking (at a rate below that of the loop self-oscillation frequency) forms a recirculating shift register.

7.3 An Example Project

[section contributed by Robert Baldwin, MIT]

The project illustrated in this section, named *AccumulatorTest*, was my first LSI design. It may be representative of what a beginner, who knows nothing about device physics or LSI, can expect to accomplish in three or four weeks. The organization of this section follows the design process chronologically. The four parts are: Conception, Refinement, Decomposition, and Implementation.

7.3.1 Conception and Refinement

This project started out with the desire to test a serial Adder/Subtractor cell that Dick Lyon had previously designed, but not implemented. The A/S cell takes two numbers in serial form, LSB first, and produces a serial output, which is either their sum or their difference, depending on the mode control input. Internally, it saves the carry (which represents a borrow if subtracting), and adds or subtracts it from the next pair of input bits. The A/S needs to know when to reset the carry, so it has an *LSBtime* signal input that marks the time slot of the LSB of each input. In order to be able to cascade the A/S's it also outputs an *LSBtime* signal, which marks the LSB of the result.

To enhance the usefulness of the test chip and reduce the external testing hardware required, I made one of the inputs the result of the last operation. All of the numbers are supposed to be 24 bits long (to conform to the format of another project), so a 23-bit shift register is needed between the A/S's output and one of its inputs (the A/S contributes one bit of delay itself). This leaves the chip with a single serial input, which can be added to or subtracted from the accumulator. In addition, there should be some way to initialize the accumulator to zero, thus a *clear* signal is provided.

Two more refinements were made to make testing easier. First, a multiplexer was added so that the A/S could be tested even if the shift register failed. Second, a circuit was added to synchronize the control signals (*mode* and *clear*); see section 7.3.3 (Implementation).

The whole chip has the following block diagram:

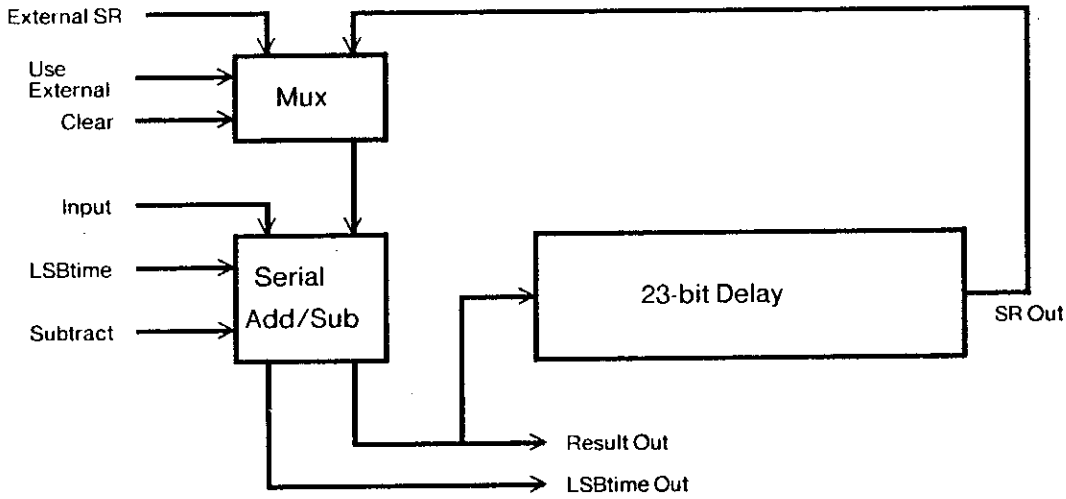

At the next higher level of abstraction, the chip has the following connections to the outside world:

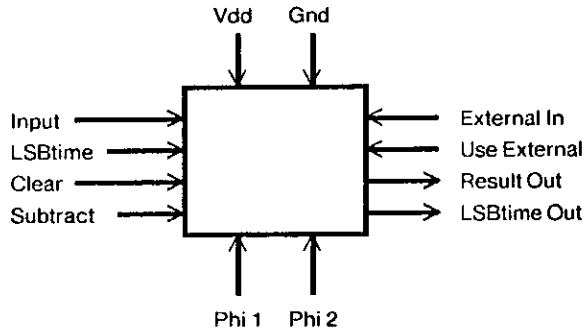

7.3.2 Decomposition

The I/O was done through two columns of pads placed on either side of the project. Since the Cell Library (see Appendix E) contained pads and drivers, I didn't have to design them myself.

Each block in the above functional diagram needs to be decomposed until it is clear how to implement it in silicon. First the decomposition of the A/S will be given, then the SR, and finally the Mux.

The first thing that was decided about the adder was the timing. A non-overlapping two-phase clock is employed. All the computation is done during phase Phi1, and latched during Phi2.

During Phi1 the adder computes the sum/difference bit which is used as the result, and the carry/borrow bit which is used during the next bit time. This suggested that the adder should be implemented with two logic blocks, and a set of double inverters to produce the true and complement of the inputs (a function block performs a three-input XOR function, and a 2-of-3 majority gate generates the carry). The outputs of these blocks are then latched by a second set of double inverters. Double inverters are used wherever possible in order to take advantage of existing library cells, and to reduce the number of cells used in the design. The A/S also computes *LSBtime Out* by delaying *LSBtime In* by one clock cycle (i.e., doubly inverted during phase1, and latched during phase2). The following diagram shows the general layout of the adder.

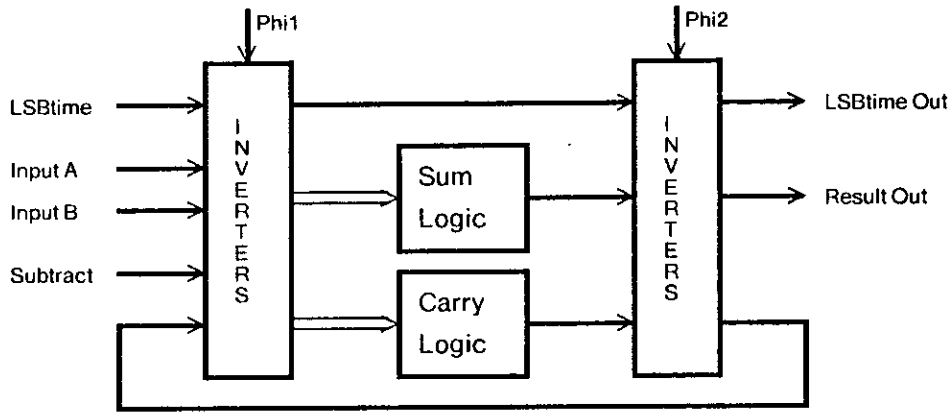

Thus the entire Adder/Subtractor can be layed out in a simple and regular way. For this project, layout optimization would have been pointless and wasteful of time. Often in such experiments it is appropriate to trade density and speed for shorter design time, increased reliability, and greater ease of modification. As in software optimization, only a few frequently used critical cells (subroutines) in the system need be optimized in order to realize most of the potential benefit.

To give the project an overall squarish shape, the shift register was folded so that the data weaves back and forth between the clock lines. This layout is composed of eleven and one-half double cells stacked on top of each other; each contains one complete weave (i.e. two bits of delay). I would like to acknowledge the help of Dick Lyon in laying out the final version of DoubleDelayCell, a descendant of the library's InverterPair and BackwardInverterPair.

The multiplexer selects either *SR Out* or *External In* to be used as one of the inputs to the adder. Its decomposition (which includes the logic for clearing *SR Out*) is best described by a circuit diagram:

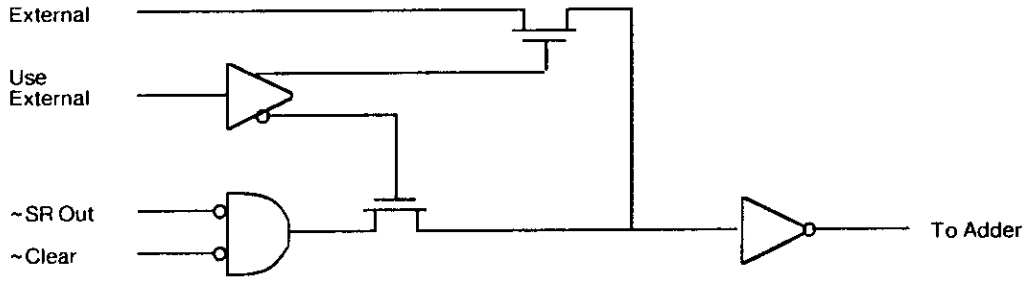

The inverters are implemented using the library's InverterPair cell, and the NOR gate is implemented as an InverterPair with an alternate path to ground for the input node.

7.3.3 Implementation

Because library cells were used wherever possible, very little needs to be added about the implementation. The detailed layout of each cell can be seen in figure 7.3.1. The basic blocks were stacked on top of each other in the order: Mux, SR, and A/S.

All the control signals are synchronized with *LSBtime* by passing them through pass gates controlled by *LSBtime.*

The layout of power and clock lines was not random. I chose to run both power and clocks in vertical metal, because I knew there would be horizontal feedback signals in the A/S and the SR, which were most conveniently run in poly. Vertical power lines made it easy to design horizontal pullups, which in turn led to low wide inverters, suitable for vertical stacking. The spacing between the metal lines allowed inverters after each clock line, and a function block between clock phases.

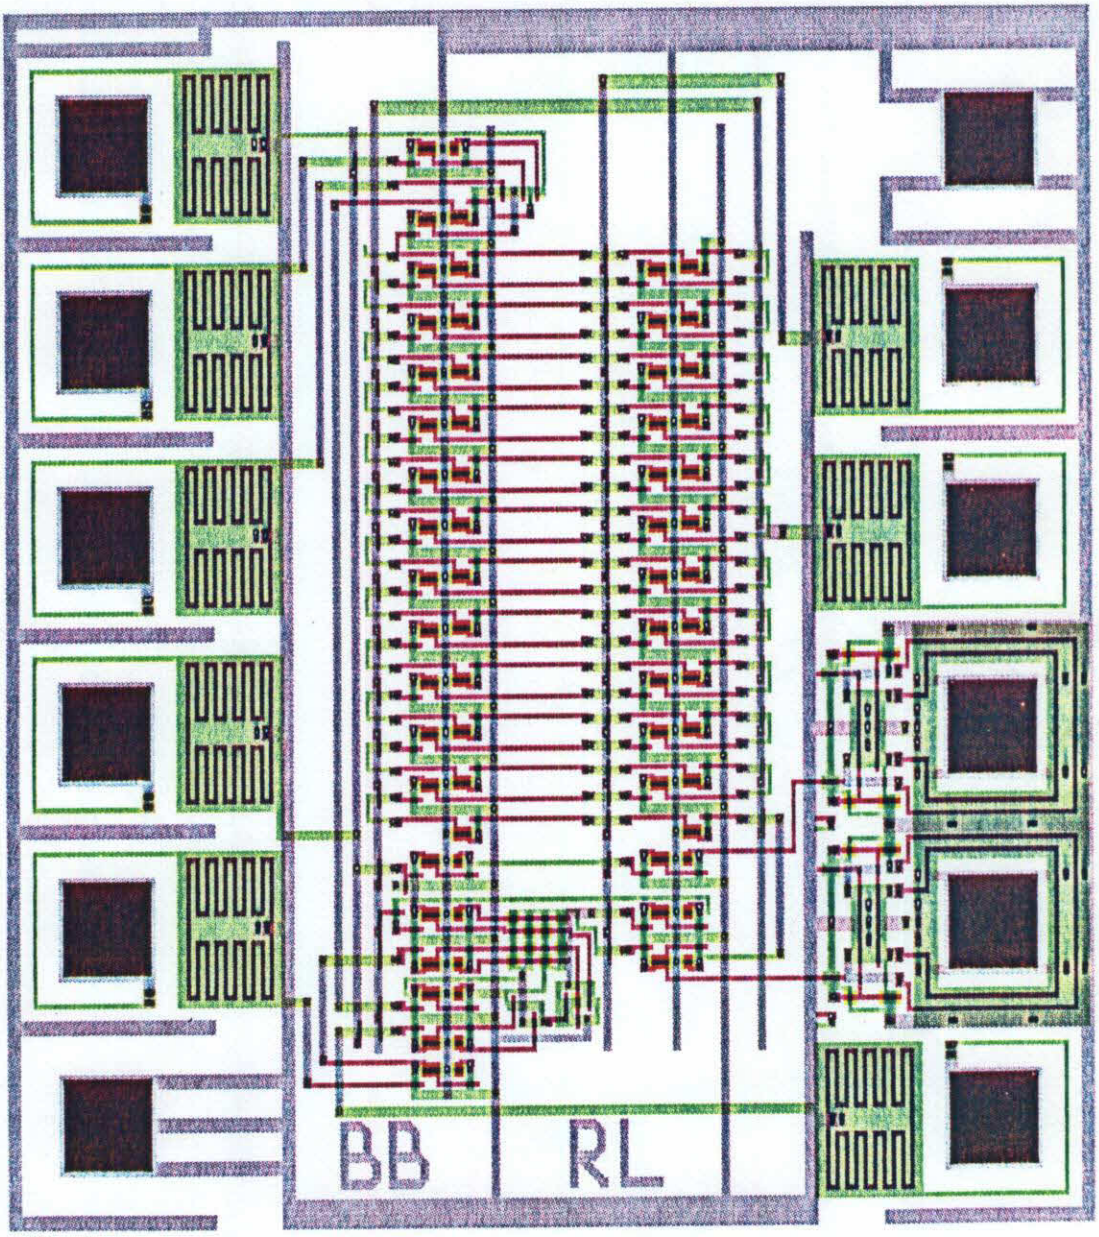

Figure 7.3.1 An Example Project: AccumulatorTest

The following example checkplot of the three-input XOR gate illustrates these ideas:

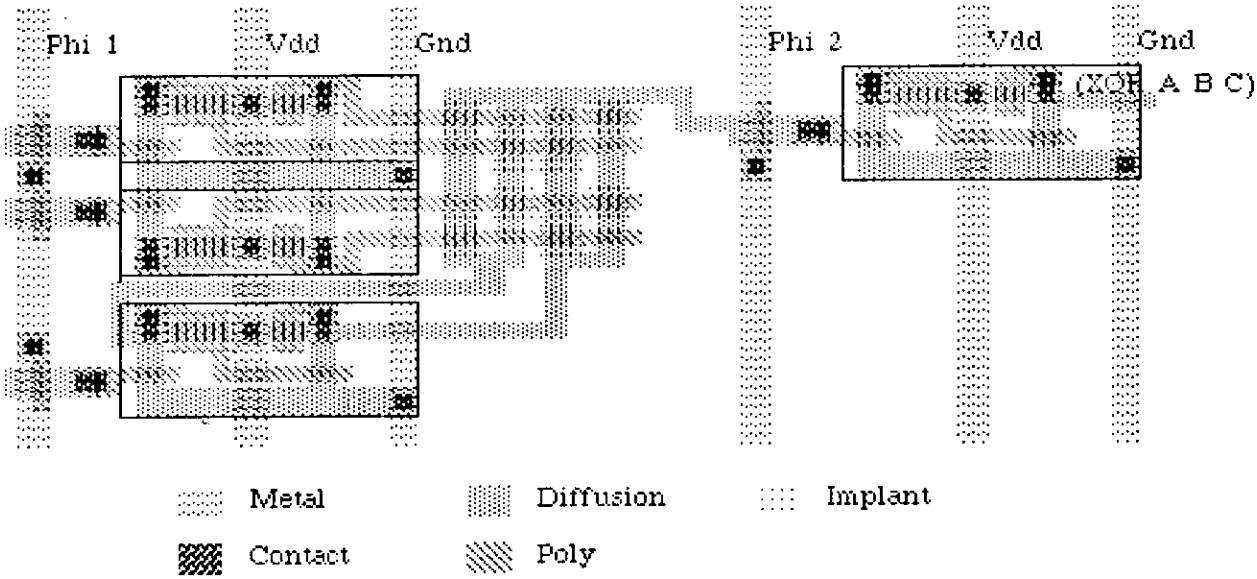

When checking the completed design, one major source of problems was found to be the mis-use of library cells. One such error was a misunderstanding of the interface to the library cell PadIn, which I had connected with poly instead of diffusion. Other bugs were due to treating library cells as black boxes, and wiring them up without checking for resulting design rule violations. Several times I forgot that butting contacts have metal over them, and ran other metal lines too close; for example, there is a butting contact in the InverterPair, which is only 2 lambda away from the connection point. I located these errors by plotting the project without the metal layer to check for poly-diffusion-implant violations, and without the poly and diffusion layers to check for metal-metal and metal-contact violations.