



Reproduced By GLOBAL
ENGINEERING DOCUMENTS
With The Permission of EIA
Under Royalty Agreement

JEDEC PUBLICATION

CMOS Semicustom Design Guidelines

JEP116

NOVEMBER 1991

ELECTRONIC INDUSTRIES ASSOCIATION
ENGINEERING DEPARTMENT



Obtained From
GLOBAL ENGINEERING DOCUMENTS
2805 McGraw Ave., Irvine, CA 92714



Documents
Document Sharing Hub

PLEASE!
DON'T VIOLATE
THE
LAW!

This document is copyrighted by the EIA and may not be reproduced without permission.

Organizations may obtain permission to reproduce a limited number of copies through entering into a license agreement with the EIA. For information, contact:

EIA Engineering Publications Office
2001 Pennsylvania Ave., N.W.
Washington, D.C. 20006
(202)457-4963

NOTICE

JEDEC Standards and Publications contain material that has been prepared, progressively reviewed, and approved through the JEDEC Council level and subsequently reviewed and approved by the EIA General Counsel.

JEDEC Standards and Publications are designed to serve the public interest through eliminating misunderstandings between manufacturers and purchasers, facilitating interchangeability and improvement of products, and assisting the purchaser in selecting and obtaining with minimum delay the proper product for his particular need. Existence of such standards shall not in any respect preclude any member or nonmember of JEDEC from manufacturing or selling products not conforming to such standards, nor shall the existence of such standards preclude their voluntary use by those other than EIA members, whether the standard is to be used either domestically or internationally.

JEDEC Standards and Publications are adopted without regard to whether their adoption may involve patents or articles, materials, or processes. By such action, JEDEC does not assume any liability to any patent owner, nor does it assume any obligation whatever to parties adopting the JEDEC Standards or Publications.

The information included in JEDEC Standards and Publications represents a sound approach to product specification and application, principally from the solid state device manufacturer viewpoint. Within the JEDEC organization there are procedures whereby a JEDEC Standard or Publication may be further processed and ultimately become an EIA Standard.

Inquiries, comments, and suggestions relative to the content of this JEDEC Publication should be addressed to the JEDEC Executive Secretary at EIA Headquarters, 2001 Pennsylvania Ave., N.W., Washington, D.C. 20006.

COPYRIGHT 1991

Published by

ELECTRONIC INDUSTRIES ASSOCIATION
Engineering Department
2001 Pennsylvania Ave., N.W.
Washington, D.C. 20006

PRICE: Please refer to the current
Catalog of EIA & JEDEC STANDARDS & ENGINEERING PUBLICATIONS
or call Global Engineering Documents, USA and Canada (1-800-854-7179)
International (714-261-1455)

Printed in U.S.A.
All rights reserved

CONTENTS

JEP116

1. Introduction	1
2. Planning a Semicustom Design	3
2.1. Product Specification	3
2.2. Design Methodology	3
2.3. Partitioning	4
2.4. Simulation Approach	5
2.5. Testing	5
2.6. Packaging	5
3. Internal Logic Design Considerations	7
3.1. CMOS Design Practices	8
3.2. Design Performance	12
3.3. Testability	19
3.4. Power Dissipation	21
4. External Design Interface Considerations	25
4.1. I/O Interface Levels	25
4.2. Output Drive	25
4.3. Performance Considerations	28
4.4. Noise	28
4.5. Power and Ground Pin Placement	33
4.6. Power Dissipation Considerations	36
5. Simulation	39
5.1. Purpose of Simulation	39
5.2. General Guidelines and Procedure	39
5.3. Functionality	41
5.4. Performance	41
5.5. Fault Simulation	42
6. Testing Considerations	45
6.1. Functional Testing	45
6.2. DC Parametric Testing	52
6.3. AC Parametric Testing	52
6.4. Board Level Testing	53
7. Packaging Considerations	55
7.1. Mechanical Considerations	55
7.2. Performance Considerations	55
7.3. Power Dissipation and Temperature	57
7.4. Board Assembly Considerations	59
8. Summary	63
9. Bibliography	65

Appendix A: Overview of Semicustom Packages	69
A.1 Types of Packages - and the Trade-Offs	69
A.2 Through Hole Packages: DIPs	69
A.3 Through Hole Packages: PGAs	69
A.4 Advances in Through Hole Packaging	71
A.5 Surface Mount Devices - General	72
A.6 Surface Mount Devices: Flatpacks	72
A.7 Surface Mount Devices: Leadless Chip Carriers (PQCC and CQCC)	74
A.8 Surface Mount Devices: Leaded Chip Carriers (PQCC and CQCC)	75
A.9 Advances in Surface Mount Packages	77
A.10 Packaging for High Power Dissipation	77
A.11 Socketing Surface Mount Devices	79
A.12 Noise Problems with Sockets	80
A.13 Surface Mount Packages for ASICs: What are the Tradeoffs	81
A.14 SMD Electrical and Thermal Characteristics	81
A.15 Packaging Problems and their Resolution	81
A.16 TAB Packaging	81
A.17 Package Selection Check List	83

(From JEDEC Council Ballot JCB-90-36 formulated under the cognizance of JC-44 on Semicustom ICs.)

1. INTRODUCTION

The design of ASIC circuits is becoming a significant part of system or product design, yet many problems continue to exist in current design practice. The guidelines in this document provide an explanation of common ASIC design problems and concerns and where possible offer solutions.

ASIC circuits, which include standard cells, gate arrays, and cell-based circuits, differ primarily in the technique used for placement of the cells and interconnect routing. The current mainstream of the industry is often assumed, with references to 5 volt power and logic constructed with CMOS FETs. The guidelines assume CMOS technology and do not discuss those issues unique to other technologies such as ECL and Gallium Arsenide. There is no attempt to consider mixed mode or analog arrays. The information in these guidelines may be of some help in designing with other technologies, but that is not the purpose. The design of the circuit follows the same process, differing primarily in the complexity of the cells available for use in the circuit design (Figure 1.1). Despite these variations, the design process for ASIC circuits follows an established pattern common to all vendors. The guidelines address issues of general concern to the ASIC circuit user as he proceeds with the design process.

These guidelines are of primary interest to ASIC designers, but will also benefit ASIC vendors, users, and others involved in the evaluation and procurement of semicustom circuits. The guidelines include the design approach, issues in deciding how to do the circuit, the design and simulation, the I/O interface, packaging, and testing. Although an attempt is made to be thorough in the coverage of these areas, some issues may have been overlooked and omitted. The text is a consensus of design information gathered by members of the ASIC Standards (JC-44) committee of the Electronics Industries Association.




Figure 1.1 Design Flow for a Semicustom Circuit

2. PLANNING A SEMICUSTOM DESIGN

The variety of design approaches and CAD tools available provide the designer of ASIC circuits many choices which determine the outcome of the design. Some designs require special tools or design methods. In many cases the approach can affect the schedule and the difficulty of the design. Careful review of the design approach is required by the designer to anticipate problems and select the tools needed for success in the design.

In order to provide documentation and a visual overview of the device being developed, it is recommended that the initial schematic and subsequent modifications be captured, ie., developed with Computer Automated Design/Engineering (CAD/CAE) tools which are available from many vendors. This interconnection/relational data can be submitted for further analysis in the same or different environment.

2.1. Product Specification

The most important document in the development of a product is the specification which defines the function of the device and the interface requirements to a system. This document should specify the following features:

1. operating voltages
2. operating and storage temperatures
3. maximum operating frequency
4. functional requirements
5. testing features
6. terminating impedance, input and output
7. output loading
8. reference specifications for environmental testing
9. reference specifications for reliability requirements
10. packaging specification

This specification should also define the reliability levels, and reference testing requirements (electrical and mechanical), and a qualified vendors list. A well conceived specification will save substantial time in the development of a product since it will allow many groups within a company to review the design and prevent last minute interface and cost problems from arising. The packaging requirements are especially important. For example, specifying a 100 pin plastic flat pack is not an adequate description of the package. 100 pin plastic flat packs are available in both square and rectangular form factors. They are available in both metric and English lead pitches.

Many companies will develop a device without a specification of this detail. However, they will be usually be susceptible to major delays in the development schedule without some formal description of the product and its impact on the manufacturing.

2.2 Design Methodology

The design methodology is a description of the principles and flow to be used to develop a product for the company. It addresses the coordination levels required between different stages of product development. This is particularly important when the end product will have a relatively long product life. Present and future technologies by which the product will be produced must be considered. The partitioning of a design should consider the possible replacement on one or more modules with new devices which utilize higher speed technologies as they become available. The procedure or methodology should be defined to answer the following questions:

1. Is design file portability required, and if required what will be the format to obtain portability?



2. Will logic synthesis be used?
3. What will be defined as the check points at which the design will be reviewed and by whom?
4. Is multichip simulation required? If so, are there simulation models for devices with which the chip under design must interface?
5. Is scan design required?
6. Is Built In Test (BIT) required?

2.3 Partitioning

Partitioning a design into ASIC circuits is the most important decision after determining the function of the design. The method of partitioning will have a great impact on the cost and longevity of a system. Some of the considerations are:

1. How can the system be partitioned to have the least effect on the critical delay paths?
2. How can the system be partitioned to effectively use the least expensive package?
3. How can the system be partitioned to segregate modules which can be migrated for future upgrades, or can be used in multiple systems or products?

When partitioning a system into ASICs, there are a number of criteria which guide the scalpel as to where the "incisions" will be made, namely:

1. Avoid partitions which "cut" critical paths. It is desirable to prevent critical paths from spanning chip boundaries where possible unless access to these nodes is necessary to drive external components.
2. Maximize gate count per partition. Note that this may not be the maximum gate count available in a given technology, but may be constrained by package. For example, to minimize cost, it may be desirable to ensure that all arrays can fit into low cost plastic flatpacks as opposed to PGAs.
3. Minimize I/O count per partition.
4. Partition by function. Should analog components be integrated, or kept off chip? What about memory, etc?

Often, the constraining factor when partitioning a system is not available gate count of the ASICs, but rather available I/O. However, this I/O to gate ratio is of critical importance. Early work by Rent of IBM on System Partitioning as related to boards is still relevant to partitioning into ASICs, and is currently, with modernization and substantial supporting design cases, widely used today.

Rent's rule states that the number of the I/Os, P , is proportional to the number of blocks G used to construct the ASIC (or board, etc.) raised to a power, K_2 , times a constant, K_1 . Mathematically:

$$P = K_1 G^{K_2}$$

The constant K1 is the number of pins or terminals of the block, while K2 is a constant which has been determined to be in the range 0.57 to 0.75 at the board level. Analysis done on various circuit designs indicates that the number of I/Os relative to the number of gates is a bimodal distribution. For small complexity circuits (<5K gates) K2 is around 0.33 to 0.50, while for larger complexity circuits K2 is about 0.28 to 0.5.

2.4 Simulation Approach

Simulation tools can be found on a PC, workstation, dedicated hardware accelerators and mainframe levels. The simulation flow is a function of the design philosophy of the user. A detailed discussion of this can be found in Section 5.0 of this document. The flow usually consists of four levels:

1. Behavioral level: Usually performed with C or Fortran languages. This level examines the general algorithm and total system function. The system partitioning is performed.
2. Functional level: Usually performed on smaller computing equipment. This level checks the implementation with fixed delays for each logic element. At this level, simulation is fast and easily iterated.
3. Silicon Simulation level: Usually performed on a workstation or mainframe and incorporates the nominal values for cell and routing delays and varies with each manufacturers library. Logic optimization programs are also run at this point. Simulation is performed at speed and test vectors are developed for parametric tests and fault coverage.
4. Post layout level: The actual path and cell delays are used to verify the critical paths and maximum operating frequencies. Particular attention should be given to the vendor's clock distribution layout scheme.

2.5 Testing

Testing method will greatly impact the cost and success of the final product and can assure that future field problems be held to a minimum cost. Lengthy test patterns can assure a high quality product but will increase the cost of the product. The trade off between cost and quality assurance is a difficult decision.

Testing issues relative to device development can be found in Section 3.3 and testing issues related to interfacing with device testers is discussed in Section 4.7 of this document.

2.6 Packaging

The first criteria for package selection should be the intended interface technology between the silicon device and the board. The package provides the environment for the silicon device and determines how efficiently the chip can dissipate heat. The criteria for package selection are usually:

1. By the number of leads required for signal and power pins. For high current applications, additional pins must be allocated for power and ground
2. By power dissipation and ambient operating temperatures
3. By cost and manufacturing capability
4. By speed of operation. For fast transitional edges on the I/O, internal ground planing and packaging methods which reduce inductance may be necessary

A complete discussion of these package considerations can be found in Section 7.0



Documents

Document Sharing Hub

Avoid asynchronous timing interfaces inside ASIC circuits (See Sections 3.2.6 and 3.2.7).

Remember that in CMOS circuits, rise and fall times propagate at different rates (See Section 3.2.9).

Avoid spikes and marginal pulse widths (See Section 3.2.9).

Add test circuits to make testing faster and easier (See Section 3.3).

Consider power dissipation as a potential problem in CMOS ASIC circuits, especially in very fast and very large circuits (See Section 3.4).

Use output drivers matching the output requirements to achieve speed requirements and minimize noise or power (See Section 3.4.3).

Some vendors offer operation at a low battery voltage and under stand-by (no clock change) with reduced performance (See Sections 3.4.4 and 3.4.5).

Provide sufficient power and ground pins to minimize noise caused from simultaneously switching output buffers (See Section 4.5).

Separate output buffer power and ground pins from power and ground for internal logic where possible (See Section 4.5).

3.1. CMOS Design Practices

3.1.1. Cell Inputs

When a library cell has unused inputs, care must be taken that these pins are properly handled to avoid noise or failure. The input of a cell cannot be left floating between a "0" and "1" state since this could cause problems by drawing dc current through the cell and generating noise, or worse, adversely affecting the functional operation of the circuit.

Two possible solutions exist:

1. Connect "1" or "0" states to inputs depending on which will allow required operation of the cell ("1" for NAND and "0" for NOR). The previous stage is not slowed by the additional load of the redundant input.
2. Change the logic to use cells where all the inputs are used.

Connecting package pins to internal cells should not be done. Vendors provide input buffers to provide the interface between input pins and the internal core logic. There are several reasons that these buffers are necessary:

1. The input buffers have sufficient electrical protection to prevent damage from electrostatic discharge.
2. The input buffers are designed to prevent latch up from being induced by transient conditions.
3. The input threshold of the input buffers is designed with the noise margin required to meet the input specifications.

In cases where the delay of the input buffer is too large and would not allow the circuit to operate at speed, the deletion of the input buffer should be avoided because of the risk involved. If a faster technology is not available, the designer should try the suggested solutions for problems with critical paths in section 3.2.1.

3.1.2. Floating Nodes

Floating nodes occur when the driving circuit is in a high impedance state. The signal can drift to a level between the high and low states, causing both the "p" and "n" transistors driven by the signal to be on at the same time. Current then would flow from the power bus to ground. Noise on the floating node could then be amplified by the "p" and "n" transistors driven by the node. The amplified noise is seen on the power and ground lines. Should the power to ground current be large enough or concentrated in a path, the current density will increase the temperature locally, and may reduce the reliability of the part.

For example, a three-state bus inside the ASIC may float for periods of time when the bus is not used to carry data. The floating state results in a signal level between the acceptable high and low states, causing current to flow in the cells driven by the bus. The condition can be corrected in several ways:

1. Reconfigure the enable signals driving a three-state bus so that one is always active.
2. Convert the circuitry driving the bus to selector logic.
3. Pull the bus to a high or low state by adding a clamp circuit which would hold or drive the bus during the float periods. This can be a marginal design practice for high performance circuits, but is often used in CMOS designs.




Figure 3.1 Avoiding Floating Buses

In cases where the delay of the input buffer is too large and would not allow the circuit to operate at speed, the deletion of the input buffer should be avoided because of the risk involved. If a faster technology is not available, the designer should try the suggested solutions for problems with critical paths in section 3.2.1.

3.1.2. Floating Nodes

Floating nodes occur when the driving circuit is in a high impedance state. The signal can drift to a level between the high and low states, causing both the "p" and "n" transistors driven by the signal to be on at the same time. Current then would flow from the power bus to ground. Noise on the floating node could then be amplified by the "p" and "n" transistors driven by the node. The amplified noise is seen on the power and ground lines. Should the power to ground current be large enough or concentrated in a path, the current density will increase the temperature locally, and may reduce the reliability of the part.

For example, a three-state bus inside the ASIC may float for periods of time when the bus is not used to carry data. The floating state results in a signal level between the acceptable high and low states, causing current to flow in the cells driven by the bus. The condition can be corrected in several ways:

1. Reconfigure the enable signals driving a three-state bus so that one is always active.
2. Convert the circuitry driving the bus to selector logic.
3. Pull the bus to a high or low state by adding a clamp circuit which would hold or drive the bus during the float periods. This can be a marginal design practice for high performance circuits, but is often used in CMOS designs.




Figure 3.1 Avoiding Floating Buses

Another floating node problem can occur due to insufficient noise margin on an input to a chip, where input noise causes the logic to switch. There are three possible solutions to the problem:

1. Increase the noise margin by using a buffer with a threshold further away from the signal level. Vendors often offer both CMOS and TTL compatible input buffers which do have significantly different thresholds.
2. Reduce the noise externally by filtering or biasing until the noise magnitude is lower than the noise margin.
3. If an input signal is slowly changing, a Schmitt trigger circuit can be used to shift the transition threshold and thus reduce the transition time between states.

3.1.3. Unused Cell Outputs

Some cells, such as a flip flop or counter, have more than one output, and at times not all outputs are needed. Leave these outputs open because:

1. The power dissipation will be minimized.
2. If the output signal feeds back into the cell, the speed will be reduced by loading the output.

3.1.4. Tying Cell Outputs Together

In TTL designs, outputs of some TTL components can be tied together to produce AND or OR functions. CMOS logic has low impedance output, both for high and low output states, and should not be used such that a node is pulled simultaneously high and low. Excessive current flow from this condition can result in circuit failure. When a design is being converted from TTL logic to a CMOS ASIC, and has gate outputs tied together in wired AND and OR configurations, these logic functions must be added to the schematic and implemented as logic gates.

3.1.5. Fan-out Limitations

Excessive loading of the output of a cell causes long rise and fall times, which increase current flow from Vdd to ground during the transitions of the signal. Excessive current could lead to a reliability problem where the transition time is very long. Also, the cells driven by the slow output are subject to noise propagation and spurious switching when the input signal level is between the high and low states.

The use of buffers or cells with a larger drive capability will help maintain good rise and fall times. Most ASIC vendors specify the maximum fan-out for each cell type to limit rise and fall times. Clocks used to drive edge triggered circuits are especially sensitive to slow rise and fall times. Special attention should be given to clock drive circuits. By designing with sufficient drive for good rise and fall times, problems with noise and marginal performance can be avoided.

3.1.6. Clocking, Setting, and Resetting Flip-flops and Latches

The design of latches and flip-flops for CMOS ASIC circuits is done in a number of ways. One way is to duplicate the edge triggered TTL version of the cell with CMOS. Another popular CMOS technique is to use transmission gates driven by the clock to select either an input or feedback path. This latch circuit produces output similar to the edge triggered D latch in that a change in state occurs with the edge of the clock. The advantage of the transmission gate latches is that fewer transistors are needed. A third but less popular method for making latches is to make static latches with cross coupled NAND or NOR gates.




Figure 3.2 Timing Considerations With Flip-Flops

With these different designs, the clear and set inputs produce differing results. One signal may take precedence over the other. When the set and clear inputs are gated into the feedback loop of the latch, simultaneous set and clear inputs may force the Q and QN outputs of the latch to the same state. Under most designs, this may not be important, but it is important in testing because it is a valid test condition. Careful reading of vendor specifications can help avoiding problems in such cases.

3.1.7. Memory Cells (RAM and ROM)

Some gate arrays and most cell based libraries have ROM and RAM cells. Often these cells are configured by a compiler which offers a wide choice of sizes. Although the silicon area of memory cell is larger in a gate array or cell-based ASIC compared to the area for a standard memory part, there are some significant advantages to using ROM and RAM cells in an ASIC circuit:

1. The effective speed can be increased due to elimination of the delays going from chip-to-chip.
2. Access to proprietary ROM code can be limited.
3. Small memory cell sizes can be implemented which would not be available in a commercial part.

4. By combining memory and logic on the same ASIC chip, the number of package pins and circuit board area are reduced.
5. In cases where the addition of a memory cell does not significantly increase the die size of the part, there may be a cost reduction in including the memory with the ASIC circuit.

3.2. Design Performance

Although the technology used in the design and fabrication of an ASIC circuit determines the performance characteristics of the cells in the library, there are techniques for the designer to optimize the speed of an ASIC circuit. Libraries used in ASIC design offer a choice of cells which are designed for either cost or performance. By selecting faster cells and using parallel logic configurations or redundant logic, speed can often be increased.

3.2.1. Critical Paths

To operate at a specified clock rate, internal signals must propagate to the inputs of latches, flip-flops, counters and other storage elements before the next clock pulse. The signal must be ready under worst case conditions to meet the required set up and hold times specified by the vendor. The same requirement is true for input signals from outside the circuit. If the circuit cannot run at the desired speed because the propagation delay of a critical path is too slow, there are several techniques (See Figure 3.3) which can be used to increase the speed of a critical path:

1. Reduce the loads along a critical path by adding buffers to drive non-critical signals and reduce the fanout.
2. Reduce serial delays or number of logic levels by reconfiguration of the logic.
3. Balance the drive with the loads by using cells with higher output drive.
4. Arrange or change the floorplan of the design to reduce metal length in the critical path if possible.

3.2.2. Generation of Time Delays

Where an accurate delay is needed internally, using cell delays cannot generate a delay with precise accuracy. CMOS logic has delays that vary with temperature, voltage, and manufacturing process. This variation can be about 4:1 between maximum and minimum delay. In addition, the propagation delay of a cell can be significantly different for rising and falling signals.



Recommended Solutions include:

1. Use a clock with a higher frequency and counter to generate phase clocks or delayed clocks where time intervals are an integral number of clock pulses. Use a counter to generate intervals which are an integral number of clock cycle.
2. Change the system logic configuration to avoid the need for accurate time delays inside the ASIC circuit.


3.2.3. Maximum Toggle Rate

The maximum toggle rate is the highest frequency at which the circuit will operate reliably. Vendors may specify a maximum toggle rate for a flip-flop, for an input buffers, and for an output buffers. If the circuit must operate beyond the maximum frequency for any or these limitations, a faster technology must be used to build the ASIC circuit.

A. REDUCE NUMBER OF STAGES TO SPEED PATH A TO E



B. REDUCE FANOUT TO SPEED PATH AE



C. REDUCE CELL INPUTS TO SPEED PATH AE




Figure 3.3 Critical Timing Improvement

Where a small part of the ASIC circuit must operate at the high speed, this could be done external to the ASIC part using a faster technology. For example, in the case where an oscillator frequency is divided into a lower frequency clock, an external circuit can divide the clock to a lower frequency.

In the case where output rise and fall times are too slow due to a high capacitive load, there are several possible solutions:

1. Use an output buffer with a higher current drive, if one is available.
2. Use several outputs (buffers and package pins) connected in parallel. Observe vendor rules for parallel buffers to maintain reliability of the part.
3. Use several buffers, dividing the load between the buffers.
4. Use an external driver circuit outside the part to drive the load.

3.2.4. Maximum Clock Speed

The maximum clock speed is limited by the slowest circuit used to propagate the clock or use the clock signal. The clock rate may also be limited due to skew between signals which must change simultaneously. Factors which may limit the maximum clock speed include:

1. The input buffer toggle rate
2. The maximum oscillator frequency where an oscillator is used on the circuit
3. The output buffer toggle rate
4. The propagation delay between clocked latches or flip-flops
5. The maximum toggle rate for a latch or flip-flop
6. Set up and hold time requirements plus input propagation delays for inputs being latched by the clock
7. Set up and hold time requirements plus output propagation delays for outputs being latched by the clock externally

Most of these factors are limited by the manufacturing process technology, and are beyond the control of the designer. Only a faster technology will solve the problem in this case.

3.2.5. Avoiding Marginal Timing

Marginal timing occurs when timing paths are too slow or too fast to allow for skew and other variations in the timing not included in the simulation. The simulation may indicate that the timing is good, but simulation is done with all elements running at the fastest or slowest speed. In testing and in actual operation there can be input skew and variation in delay paths from local conditions affecting the timing. As a result system failures occur when signals are strobed or latched unless there is a margin for skew.

3.2.6. Set Up Time

The set up time must be adequate for the input of a latch or other storage element to be guaranteed stable at the clock time under all conditions. If, for example, the input of a D latch is changing too late for the signal to be reliably latched, this condition can cause a number of problems:

1. The output state is uncertain, depending on variables such as process parameters, voltage, temperature, or noise. The uncertainty will make production testing impossible.
2. Metastability problems may result because a changing input or spike on the clock may leave the latch partially switched with an output which does not converge to a high or low state before it is latched at the next clock.
3. Excessive power consumption will result if there is excessive switching time.
4. Noise on the output can result from excessive switching time and input noise, which is amplified by the gain of the cell.

There are two ways to solve the problem:

1. Make the input signal stable at an earlier time. This may be done by reducing the delay of the input data signal.
2. Shift the clock to a later time by adding delay. However, this will reduce the available time for propagation of the output of the storage element and the delay varies using logic cells.

3.2.7. Hold Time

The data input to a latch or other storage element must be held stable after the clock for the specified hold time to guarantee the expected output state under all operating conditions. If the hold time is not adequate the same problems as listed in section 3.2.6. may occur. There are two ways to solve the problem:

1. Hold the input signal stable for a longer period of time. This can be accomplished by delaying the input by adding delay or substituting a slower cell in the input data path. This is a marginal design practice and can only be done where there is ample margin for the set-up time.
2. Shift the clock to an earlier time if there is sufficient set-up margin.

3.2.8. Race Conditions

A race condition occurs in the circuit when a signal propagates fast enough to be clocked into the next stage in the same clock cycle. This is often caused by using latches of differing timing or having the clock on the second latch later than the first. There is usually a combination of clock skew and fast propagation through a flip-flop.

There are several ways to solve the problem:

1. Use a common clock driver for all related latches so they all change at the same time.
2. Avoid gating logic signals with the clock for some of the latches. Such practice can also be applied to the data path which will avoid the race condition.
3. Where a common clock cannot be used, clock the first latches in the sequence later than the those later in the sequence. In some cases the latches can be triggered on rising and falling edges of the clock.




Figure 3.4 Racing in a Latch

3.2.9. Spikes and Marginal Pulse Width

Because CMOS library cells are constructed with differing numbers of series devices in the pull-up and pull-down paths, the rising and falling signals have different propagation delays. Narrow pulses can get narrower propagating through logic, and reach a point where they are unreliable spikes or disappear completely. Circuits which generate narrow pulses must either be avoided, or carefully designed to ensure reliable operation over all operating conditions and manufacturing tolerances.

Where logic is used to generate a clock, a spike may occur where two signals change with a small overlap as in the examples in Figures 3.4 to 3.6. The overlap can produce a spike which does not have sufficient width (energy) to switch the cells driven by the clock in a predictable manner. This condition can cause a number of problems:

1. The output state is uncertain, depending on variables such as process parameters, voltage, temperature, or noise. The uncertainty makes production testing of the part impossible.
2. Metastability problems may result because a changing input may leave the latch partially switched with an output which does not converge to a high or low state before it is used.
3. Excessive power consumption will result if there is excessive switching time in recovering to a high or low state.
4. Noise on the output can result from excessive switching time and input noise which is amplified by the gain of the cell.

There are two ways to solve the problem:

1. Find a way to gate the signals with the data input instead of the clock input of the latch.
2. If signals must be gated with the clock, be sure they cannot change while the clock is active. Although clock gating is used, it often induces clock spikes and should be avoided.

A spike generated from simultaneous changes of inputs to some static logic is usually not a functional problem if it does not become part of a clock signal or an output of the chip except for the error messages generated by the timing verifier and simulation software. However, it is best to avoid these spikes, espe-

cially where power dissipation and noise are concerns of the designer. By adding a redundant term to the static logic, the spike is prevented. Figure 3.7 shows the circuit with the additional term in the equation and the Karnaugh map of the logic.

The width of a pulse may be too narrow for reliable operation of the circuit. This problem often occurs when a pulse is generated using a flip-flop with the Q output connected to the reset pin. Because rising and falling transitions have different delays through logic cells, the pulse could become narrower, or disappear, as it propagates through logic cells. Some designers will add delay between the Q and reset terminals to generate a wider pulse which is sufficient to reset the flip-flop and be used in elsewhere in the circuit. This practice is not recommended. The calculation and analysis of the delays with this solution can be avoided and a more reliable pulse generated if the width is derived from the edges of clock signals. This provides an accurate control of the pulse width.




Figure 3.5 A Circuit that Can Produce Spikes




Figure 3.6 How Spikes are Generated




Figure 3.7 A Circuit that can Avoid Spikes

3.2.10. Clock Skew and Clock Distribution

Clock skew is the variation in delay of the clock to the various places in the circuit where it is used. It must be kept to a minimum value to meet timing specifications of the circuit. Where different paths are involved, clock skew should be treated as an additional delay added to the delays of critical paths between clocked storage cells. In the design of a circuit with a very large number of cells using the same clock, the fanout for a single buffer may exceed the maximum value specified. There is also a concern that if all loads were driven from one point, the current density would be too high for the metal and contact size, causing a reduced lifetime of the part.

Many of the work station simulators used for ASIC design do not have the ability to simulate clock skew. The more sophisticated simulators can handle transmission line effect and distributed load on the interconnect which are much more accurate than the admittance model. Since the routing has a significant effect on skew, simulation of interconnect loading is best done after layout.

Several solutions have been used for clock distribution:

1. The clock can be distributed in a "tree" structure with the loads balanced to equalize the clock time at all loads. Because there will be some variation in the timing between branches of the tree, the loads should be arranged such that groups of loads requiring low skew within the group are driven by the same branch. The branches should be balanced with equal loads to minimize differences in delay between branches.
2. A few vendors have developed wide metal lines for the distribution of clocks. All clock inputs can then be connected together as a single node with a very large driver. This does not totally eliminate the skew, because the metal lines on the die have some small intrinsic delay due to the distributed impedance. Special software is required to do the clock routing with wide metal.
3. Some vendors offer a buffer with multiple outputs which do not require the wide metal and are internally connected with metal. The balancing of loads on the metal lines prevents problems with current density and minimizes clock skew.

Clock skew can also be a benefit for the designer. There are problems which occur with designs having large numbers of cells switching at the same time. As the cells switch through the transition state between the high and low levels, a current pulse flows from the Vdd to the ground terminals because both the p and n transistors are on. This current is in addition to the current used to charge the capacitive loads of all the cells which switch. The combination of these currents can cause severe noise spikes on the Vdd and ground lines. Since the noise is generated internal to the ASIC, there is little opportunity to use external filtering to reduce the noise. The clock skew can spread the switching over a longer period and reduce the peak noise. Where there is time to spread the clock switching an improvement in noise level may be achieved.

3.3 Testability

The complexity of ASIC designs often makes them difficult, or impossible, to test on production testers. By designing the circuit to be easily tested, the development of a test program will be much easier. Several companies have developed methods to reduce the testing problems by improving the ability to control and observe the internal nodes of the circuit. The most prominent techniques are Level Sensitive Scan Design (LSSD) developed by IBM, and the recently proposed boundary scan method described in IEEE standard 1149 which is developed for testing components assembled on a board. The boundary scan technique was developed by the Joint Test Action Group (JTAG), and JTAG is working toward world-wide acceptance of the boundary scan method of testing. Both techniques string together strategic nodes in the design so that they may be loaded or examined at various points during functional testing. ASIC testers are synchronous systems, expecting outputs at specific times. It is best to avoid using long asynchronous delays such as ripple counters.

3.3.1. Initialization Requirements


The circuit must be completely initialized so that all nodes reach a known state, otherwise the circuit is not initialized, and it cannot be tested by the commercial test equipment used for ASIC circuit production. The initialization must be completed in a reasonable number of clock cycles so the functional test sequence can proceed. If there are any free running counters or clock division logic with no reset provision, reset logic must be added to the circuit. The testers used to functionally test gate array and cell-based circuits require all nodes to reach a known state during an initialization procedure. Any latch, flip-flop, or counter bit not set to a known state will propagate the unknown state causing outputs of the circuit to be in an unpredictable state.

When simulation of the circuit results in unknown states, a means of clearing the unknown condition must be found, either by a sequence of input simulation vectors or a change to the circuit. Where there is no signal to completely reset all nodes of the circuit, a reset can be generated from a combination of input states that cannot occur in normal circuit operation. Where there is an available package pin, a special test input pin can be added to initialize the circuit.

The initialization sequence can be repeated to divide the simulation into independent sections which can be run independently of each other. These sections can be run separately to verify the functionality of various parts of functions of the circuit. This reduces simulation time during development. Later, the sections can be combined to produce a complete functional test.


Another approach gaining in popularity is to use one of the scan techniques (see section 6.4) to serially load all the cells which initializes the circuit. This method has the advantage that software programs can be used to automatically generate test vectors for functional test. It also has the disadvantages of requiring several additional pins, adding to the logic, and slowing the performance.

A free running counter in a circuit cannot be tested unless it is reset. For example, often an oscillator frequency is divided down by a free running toggle to produce a symmetrical clock signal. Such a divider is a form of a counter and must be reset to initialize the circuit for test.



**SINCE THE LOGIC VALUES IN FLIP-FLOPS
CANNOT BE FIXED FOR SIMULATION BY
APPLICATION OF CLOCK PULSES, TESTING
OF ALL CIRCUITS CONNECTED TO THIS
CIRCUIT BECOMES IMPOSSIBLE.**

Figure 3.8 A Circuit that cannot be Initialized



**MAKE THE CIRCUIT ALLOW INITIAL
SETTING (OR RESETTING)**

Figure 3.9 A Circuit that can be Initialized

State machines can be difficult to initialize if there is no provision for reset. Some state machines take a long time to change states because there is a counter or many intermediate states creating delay. Initialization is difficult if it is not done by forcing the logic to an initialized state from any possible starting condition and perhaps to a disallowed state.

2.2 Scan Testing

Scan testing has a number of variations, but the principle is that all internal storage cells such as latches, counter bits, etc. are designed such that they can be loaded and read via test pins. A test clock is used to load and unload the data for testing. Under normal operation the scan test circuitry is turned off and has no effect on the function of the integrated circuit.

The major advantages of scan testing are that the test condition can be easily set with the scan input and states can be quickly read by the tester at any point in the test. Most vendors offering scan testing also offer automatic test pattern generation (ATPG) as part of their CAD tools. The testing covers all of the internal nodes with as close to 100 % fault coverage as possible. The test vectors will not aid in the simulation, which must also be done to verify that the circuit does perform the required functions.

Scan testing has the disadvantages of increasing the logic and pin count which may result in a larger array die size. The power is also increased which may be a consideration. The additional pins required may force the part into a larger and more expensive package. There is also a performance penalty because the scan cells are slower than the equivalent non-scan cells.

2.3 Power Dissipation

Although CMOS is noted for low power dissipation, there are applications such as battery operation where power utilization must be kept as low as possible to maximize battery life. As gate counts get larger and speeds get faster, it becomes easier to exceed the maximum current or power limitations of the part. So, with the large numbers of simultaneously switching outputs driving large capacitive loads, the peak instantaneous current can exceed the limitations of the power and ground lines on the die. Excessive power or ground current can cause metal migration (a movement of metal due to high current which reduces the current carrying capacity of an interconnection), and lead to significant reduction of the reliability of the parts.

2.3.1 Static Power

Even a minimum static (no dynamic operation) power must be achieved by the ASIC design because of a requirement for long life battery operation. Battery life is determined to a large extent by static power consumption of the ASIC circuit where dynamic activity is infrequent. Static power dissipation results from the following current flow in the circuit:

1. Junction leakage - current flowing through diffusion junction areas throughout the circuit. This current varies with temperature and voltage as well as with the diffusion area on the circuit. The diffusion area is fixed for a particular design. If speed is not a consideration, it may be possible to reduce the supply voltage to minimize the junction leakage current.
2. Source-drain leakage - current flowing through transistors when they are in the "off" condition. This current is normally much smaller than the junction leakage and is not a significant factor in the power dissipation. The only way for a designer to reduce this current is to reduce the supply voltage.
3. Pull-up and pull-down resistor dissipation - current flowing through resistors used to pull buffer outputs high or low. Normally such circuits are avoided in low-power CMOS design because the CMOS logic can be designed such that the outputs are always pulled either high or low.