

SYSTEM MANAGEMENT OF A REDUNDANT
CLOCKING NETWORK

by

Charles E. Manush, III
S.B., Massachusetts Institute of Technology
(1974)

SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY


June, 1976

Signature of Author _____ **Signature redacted**
Department of Aeronautics and Astronautics
February 10, 1976

Certified by _____ **Signature redacted**
T~~h~~esis Supervisor

Accepted by _____ **Signature redacted**
Chairman, Departmental Graduate Committee

ARCHIVES



Thesis
Aero
1976
MS

SYSTEM MANAGEMENT OF A REDUNDANT
CLOCKING NETWORK

by

CHARLES E. MANUSH, III

Submitted to the Department of Aeronautics and Astronautics on February 10, 1976, in partial fulfillment of the requirements for the degree of Master of Science.

ABSTRACT

The timing references required to enable the units of a system to operate in synchronism are provided by a clocking network. The development of highly reliable computer configurations, specifically a fault-tolerant multiprocessor, has introduced the need for a similar improvement in the reliability of the clocking network. This thesis emphasizes the application of fault-tolerant concepts for improving the network's reliability.

In this thesis the attributes which enable a clocking network to achieve fault tolerance through the use of hybrid redundancy are defined and methods for the detection and correction of faults within the network are developed.

A clock receiver was developed which polled the transitions of clock signals instead of the logic levels. This was found to reduce the number of input clock signals required for the production of a fault-tolerant secondary clock signal. A method of testing was devised to detect and isolate faults using relatively uncomplicated circuits.

Thesis Supervisor: Albert L. Hopkins, Jr.
Title: Associate Professor of
Aeronautics and Astronautics

ACKNOWLEDGEMENTS

The author would like to express his gratitude to Professor Albert L. Hopkins for the guidance and encouragement he provided throughout the production of this thesis. He is indebted to Dr. T. Basil Smith for his invaluable suggestions concerning the clock receiver. Also, he is grateful to Mary Shamlan for her aid in typing the text of this thesis.

This report was prepared by the Charles Stark Draper Laboratory under Grant GJ-36255 and DCR74-24116 with the National Science Foundation.

The publication of this report does not constitute approval by the Charles Stark Draper Laboratory or the National Science Foundation of the findings or conclusions contained therein. It is published only for the exchange and stimulation of ideas.

TABLE OF CONTENTS

	<u>Page</u>
CHAPTER 1. CLOCKING	6
1.1 Fault Tolerance	6
1.1.1 Modular Redundancy	6
1.1.2 Voting	11
1.2 Clock Receiver	11
1.2.1 Logic Level Voting	11
1.2.2 Edge Voting	14
1.3 The Daly-McKenna Clock	21
1.4 Failure Modes	23
CHAPTER 2. CLOCK RECEIVER	29
2.1 Architecture	29
2.1.1 Edge Detection	29
2.1.2 Voting	29
2.2 Fault-Free Behavior	32
2.3 Fault Tolerance	35
2.3.1 Primary Clock Failure	37
2.3.2 Clock Receiver Failure	44
2.3.3 Multiple Failures	44
2.4 Masking	46
CHAPTER 3. FAULT DETECTION	47
3.1 The Delay Test	47
3.1.1 The Delay Selection Control	49
3.1.2 The Input Difference Detector	49

	<u>Page</u>
3.1.3 The Reference Clock Receiver	52
3.1.4 The Output Difference Detector	57
3.2 Failure Search	59
3.2.1 Input Clock Signal Failure	64
3.2.2 Receiver System Failure	75
3.3 Alternative Detection Test	80
3.3.1 Delay Test (Alternatives)	80
3.3.2 The Pattern Test	84
 CHAPTER 4. APPLICATION	 88
4.1 Multiprocessor Systems	88
4.2 The CARDS Multiprocessor	89
4.2.1 Clocking Requirements	97
4.2.2 The Error Latch	97
4.3 Testing the Clocking Network	100
4.3.1 The Independent Test	100
4.3.2 The Dual Test	100
4.4 Clock Frequency	104
4.5 Recommendations	105
4.6 Conclusions	108
 REFERENCES	 110

CHAPTER 1

CLOCKING

The use of a clock as a timing reference is essential for any synchronous system. The clocking network which drives the various units allows the system to maintain its members in synchronism. The reliability of the synchronization is dependent upon the quality of the clocking network. One method of improving this quality is by making the network fault-tolerant.

1.1 Fault Tolerance

Fault tolerance is achieved either by making the system insensitive to faults or by providing the system with the means to detect, locate and correct faults as they occur. The first approach results in the dynamic masking of errors and the prevention of their propagation through the system. An example of a fault-tolerant clocking network employing this approach is discussed later in this chapter. Fault tolerance is achieved with the second method through the system's ability to tolerate an initial fault (or faults) and to correct it before another occurs to propagate the error. A clocking network using this method for achieving fault tolerance is the principal theme of this thesis.

1.1.1 Modular Redundancy

Modular redundancy is a systematic structural approach for the achievement of fault tolerance through comparison and/or voting amongst replicated units, hereafter referred to as modules. Only replicated redundancy is considered in this thesis, other approaches having so far appeared to be impractical as far as a clocking system is concerned. A system which is not redundant is referred to as a simplex system. Such a system is not systematically capable of detecting or tolerating a faulty module. In a duplex system, the modules perform the same operation in pairs and compare results. This assures the detection of all faults that are exercised, or flexed in the course of operation. Tolerance is not achieved, since the faulty module can not systematically be distinguished from its twin. By arranging the modules in triads, a system is provided with the systematic capacity of tolerating the failure of a triad member, provided the voting mechanism has not failed. In

general, the failure rate of a module can be assumed to be much greater than the failure rate of the voting device. A system arranged in this manner is a Triple Modular Redundant (TMR) system.

A TMR system is capable of tolerating the failure of one triad member, that is, it possesses single-fault-tolerance. There are two approaches for increasing the number of faults that a system can survive. The first is by increasing the number of modules that are polled. For example, the NMR (N-tuple Modular Redundant) system shown in Figure 1.1 is capable of tolerating n faults through the application of majority voting, provided that the number of modules, N, is greater than or equal to 2n+1. The reliability of an NMR system for a given simplex reliability is shown in Figure 1.2 as a function of the degree of fault tolerance desired.⁵ Note that the reliability of a simplex module must exceed 0.5 for a given job to justify the use of an NMR arrangement.

The second approach is to replace the failed modules with stand-by spares. The most efficient utilization of a given number of modules, N, is to employ a hybrid redundancy configuration using a TMR arrangement with N-3 spares. The reliability of this configuration is equal to the sum of the probability that no module fails or that only one fails, times the reliability of the voter (Equation 1.1). The reliability of a hybrid system consisting of a TMR arrangement with S spares is expressed in Equation 1.2 with the assumption that the reliability of the voting and correcting mechanisms are essentially 1.0 for the job. The reliability of such a hybrid system with a given simplex reliability is shown in Figure 1.3 as a function of the number of spares available.⁵

$$R(\text{TMR}) = (R^3 + 3R^2(1-R))R_{\text{voter}} \quad (1.1)$$

$$R(\text{TMR} - S) = 1.0 - ((1-R)^{S+3} + (S+3)R(1-R)^{S+2}) \quad (1.2)$$

Both methods of increasing the level of fault tolerance have their drawbacks. From Figure 1.2, it can be seen that the reliability of an NMR system for a given simplex reliability increases less and less as the level of fault tolerance is increased. Consideration must also be given to the total number of units required to perform a given number of simultaneous tasks. An NMR system requires N times as many modules as a simplex system to perform the same number of jobs. A compromise must be made between the reliability of the system and efficient job utilization of the modules. The main disadvantage of this hybrid system is the necessity to detect, isolate and replace the failed unit before a second fails. The reliability of the correction procedure must be taken into account when considering the overall reliability of the arrangement.

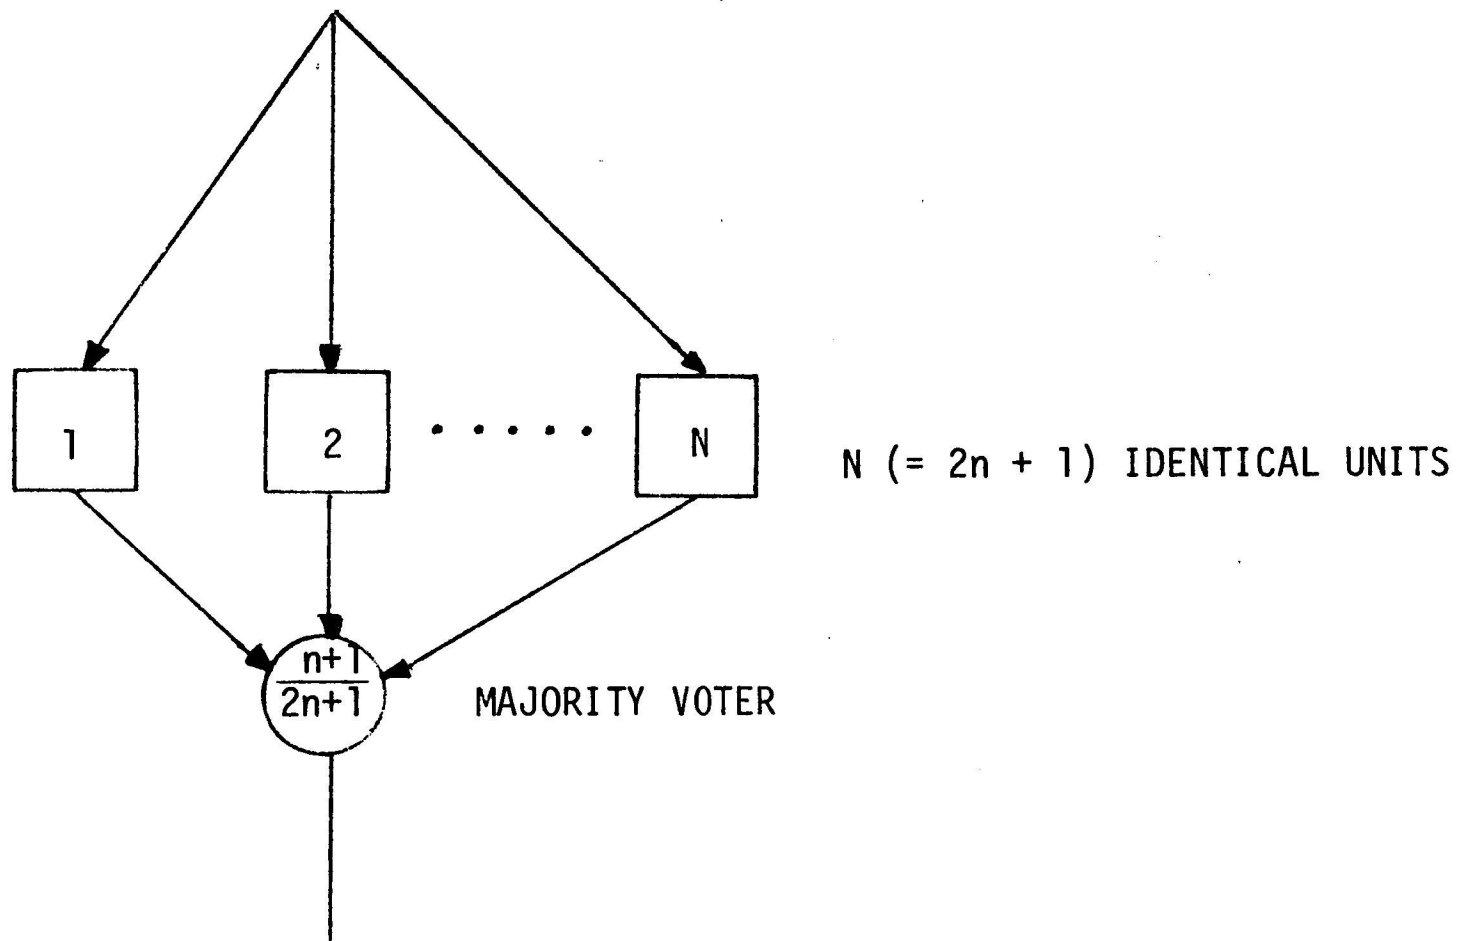


Figure 1.1 An NMR System.

NMR Reliability

Vs.

Simplex Reliability

Parameter: Number of Modules, $N = 2n + 1$




Figure 1.2

Hybrid Reliability

Vs.

Simplex Reliability

Parameter: Number of Spares, S

Number of Active Modules = $N = 3$




Figure 1.3

1.1.2 Voting

The voter used in NMR systems is a majority voter in which the output agrees with $n+1$ or more of the $2n+1$ inputs. Figure 1.4 shows a typical design for a two-out-of-three majority voter. The output of the voter agrees with the majority of the inputs. This voter is capable of tolerating one failed input provided that the two other inputs are identical. Since, in general, no two clocks are exactly alike, a simple majority voter should not be used to poll clock signals. Such a clock voter results in the possibility of abnormalities being produced in the output as the result of a single failed clock signal as illustrated in Figure 1.5. Abnormalities in the secondary clock signal may upset the timing sequence of the unit it serves.


For the purpose of "polling" several clock signals to obtain a fault-tolerant secondary clock signal, a more complicated voter than just a simple majority voter is needed. A voter designed to be used with clock signals is a clock receiver. A clock receiver may poll an input clock signal in one of two ways. It can respond to the state of the clock signal as with logic level voting, or to the transition of a clock signal from one state to the other, that is, by counting the edges of the signal.

1.2 Clock Receiver

In a clocking network devised by Daly and McKenna², primary clock signals are produced in one location and then distributed throughout the entire complex to various units in need of a system clock. At these units, the primary clock signals are used by a clock receiver to generate a fault-tolerant clock for the unit's use. The secondary clock signal (system clock) is produced by polling either the logic levels or the edges of the input clock signals. This is the basic structure that will be assumed throughout, while various modifications are discussed.

1.2.1 Logic Level Voting

In designing a clock receiver with logic level voting and single-fault tolerance, it is advantageous to use four input clock signals instead of three. There are two advantages to using four signals. First, the system clock changes state only after the majority of the unfailed signals have changed. Second, the clock receiver makes use of hysteresis to separate the input functions that cause the transitions of the system clock from one state to the other. The first advantage results from the fact that the system clock changes state only after at least three of the four input clock signals have changed. If one of the input clock signals fails, the system clock still changes after at least two of the three



INPUTS			OUTPUTS
A	B	C	X
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	1
1	0	0	0
1	0	1	1
1	1	0	1
1	1	1	1

Figure 1.4 A Majority Voter.

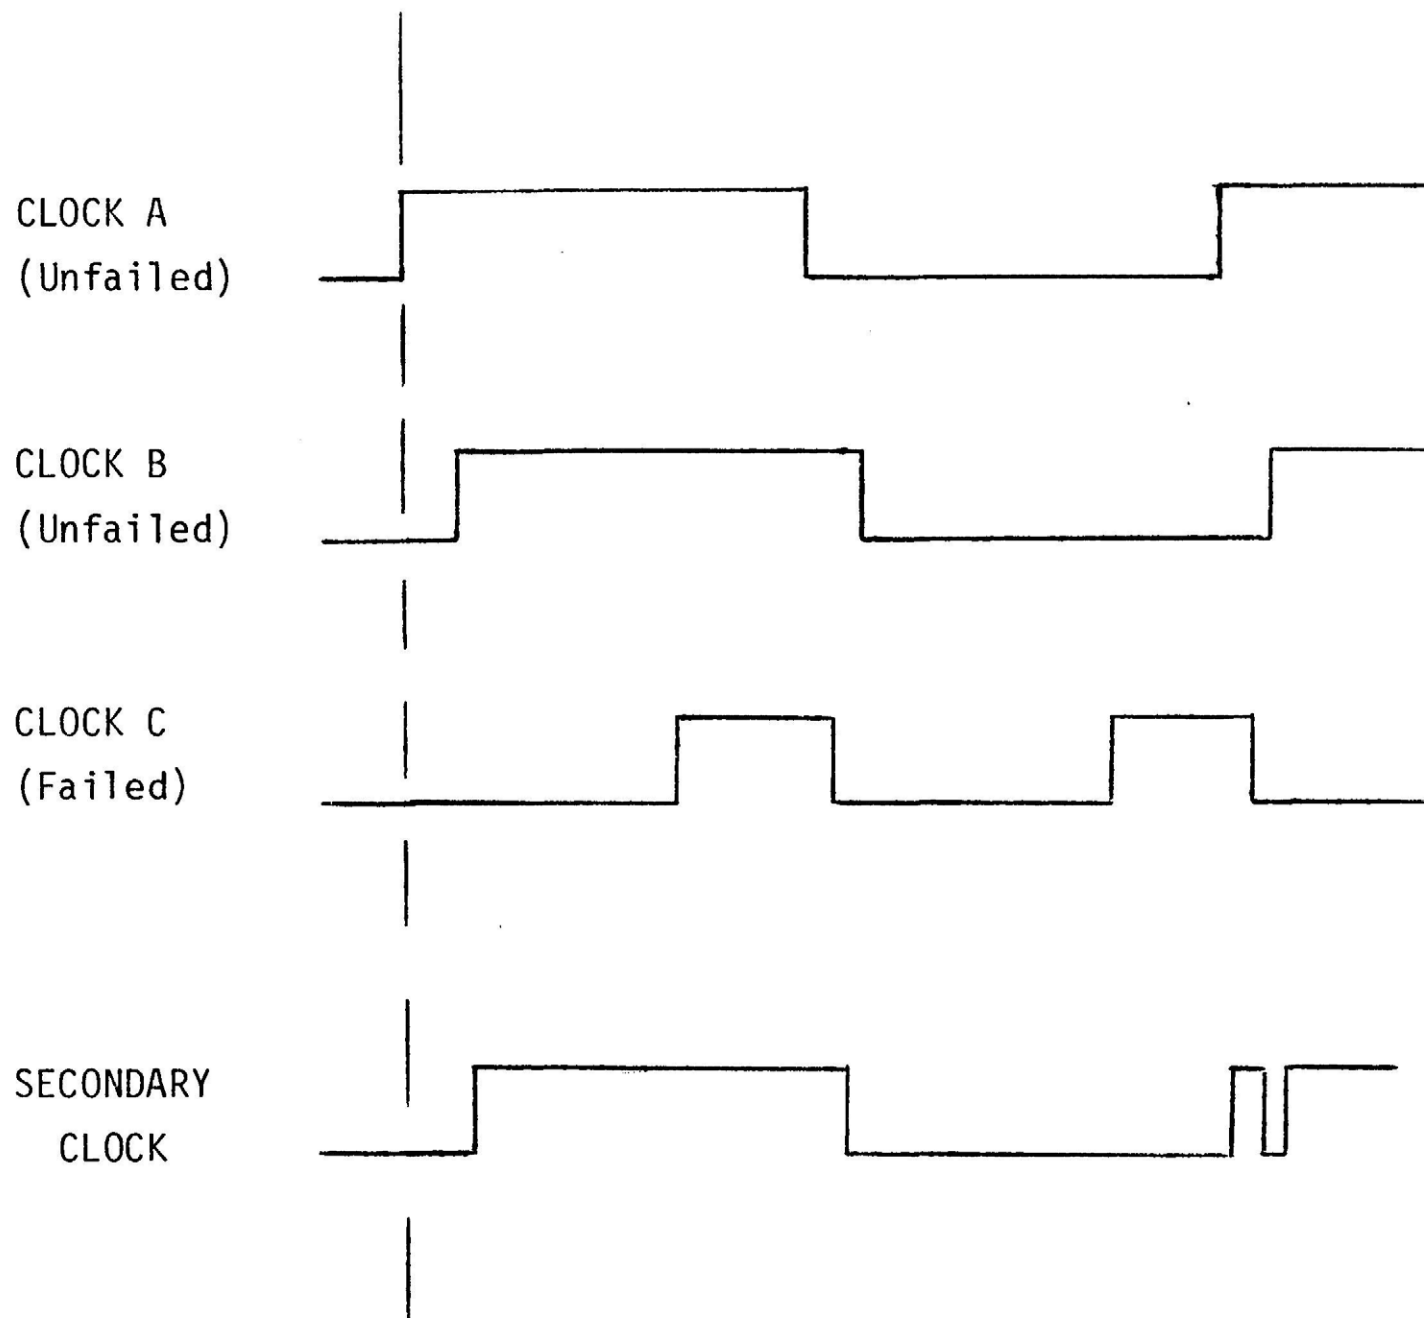


Figure 1.5 Abnormality Produced by a Failed Clock Signal.

unfailed clock signals have changed. The hysteretic property of the four-input receiver is illustrated in Figure 1.6. When there is no majority, the state of the system clock depends on its past history (the state remains unchanged). Figure 1.7 shows two examples of a single-fault-tolerant receiver employing hysteresis. The Boolean minterms are listed in Table 1.1 under the appropriate output expression for the receivers. Note that the expression for SYSCLK is more than unit distance from that for $\overline{\text{SYSCLK}}$. This provides protection against abnormalities within the system clock in the event of a clock failure. Since components do not respond instantaneously to a change to their inputs, additional safeguards may be needed. Figure 1.8 shows why the receiver of Figure 1.7a may need additional safeguards. The system clock becomes high after clock B goes high, delayed by an amount of time approximately equal to the propagation delay time of the voter. However, during this time, clock C goes low, resulting in the system clock becoming low. Since the system clock is used as one of the inputs to the majority voter, the system clock is in a racing situation until forced high by clock D becoming high. This is similar to the problem encountered with the three-input receiver.


The receiver shown in Figure 1.7b needs additional circuitry to prevent abnormalities in the voter's output from being inputted to the flip-flop. Such abnormalities may create a racing condition when inputted to a flip-flop. The racing condition resulting from a single clock signal failure can be avoided by adding a retriggerable one-shot or a similar device for producing pulses of a predetermined width. Figure 1.9 illustrates the manner in which such a receiver produces its system clock despite an input clock signal failure. The system clock ignores the second pulse in both the SET and RESET signals. The SYSCLK changes state after a majority of the unfailed input clock signals have changed. The receiver in Figure 1.7b with retriggerable one-shots must use four or more input clock signals to achieve fault tolerance. Figure 1.10 uses clocks B, C and D of Figure 1.9 to illustrate why the receiver does not use three clocks. A receiver that uses one-shots with logic level voters operates in a manner similar to one using edge voting in which a flip-flop is set when at least three (out of four) input clock signals have produced a leading edge and reset when at least three have produced a trailing edge.

1.2.2 Edge Voting

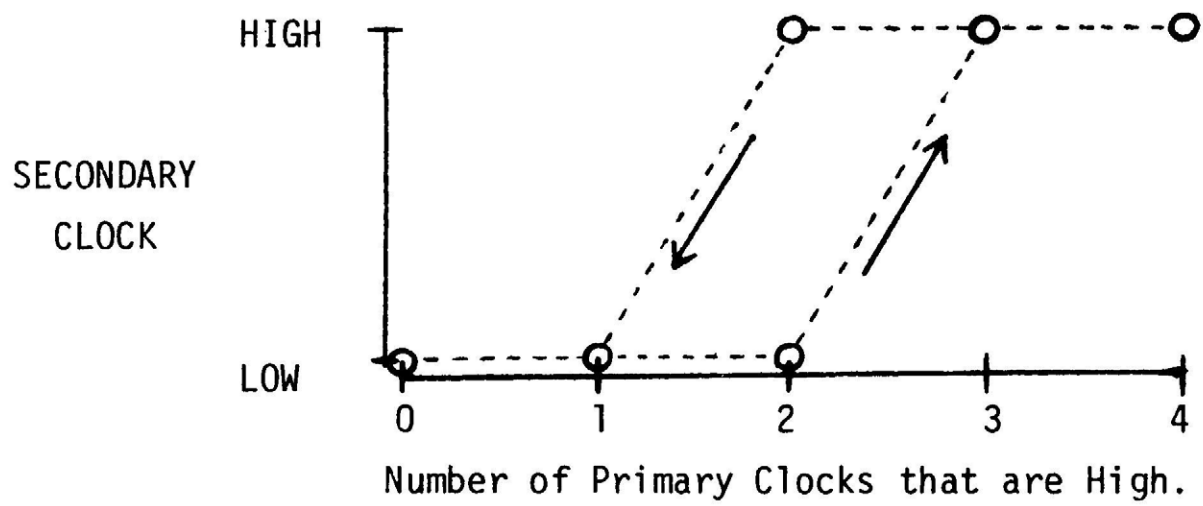
The edges of a clock signal are classified into two categories. The leading edge (L.E.) group are the state transitions of the clocks from their low level to their high level. The trailing edge (T.E.) group are the transitions from high to low logic level.

TABLE 1.1
 BOOLEAN MINTERMS OF FOUR-INPUT RECEIVER

SYSCLK	$\overline{\text{SYSCLK}}$	UNCHANGED
ABCD	\overline{ABCD}	\overline{ABCD}
$\overline{A}BCD$	$\overline{A\overline{B}CD}$	$\overline{A\overline{B}CD}$
$A\overline{B}CD$	$\overline{A\overline{B}CD}$	$\overline{A\overline{B}CD}$
$AB\overline{C}D$	$\overline{AB\overline{C}D}$	$\overline{AB\overline{C}D}$
$ABC\overline{D}$	$\overline{ABC\overline{D}}$	$\overline{ABC\overline{D}}$
		$ABCD$




a) Three-Input Clock Receiver.




b) Four-Input Clock Receiver.

Figure 1.6 The State of the Secondary Clock.

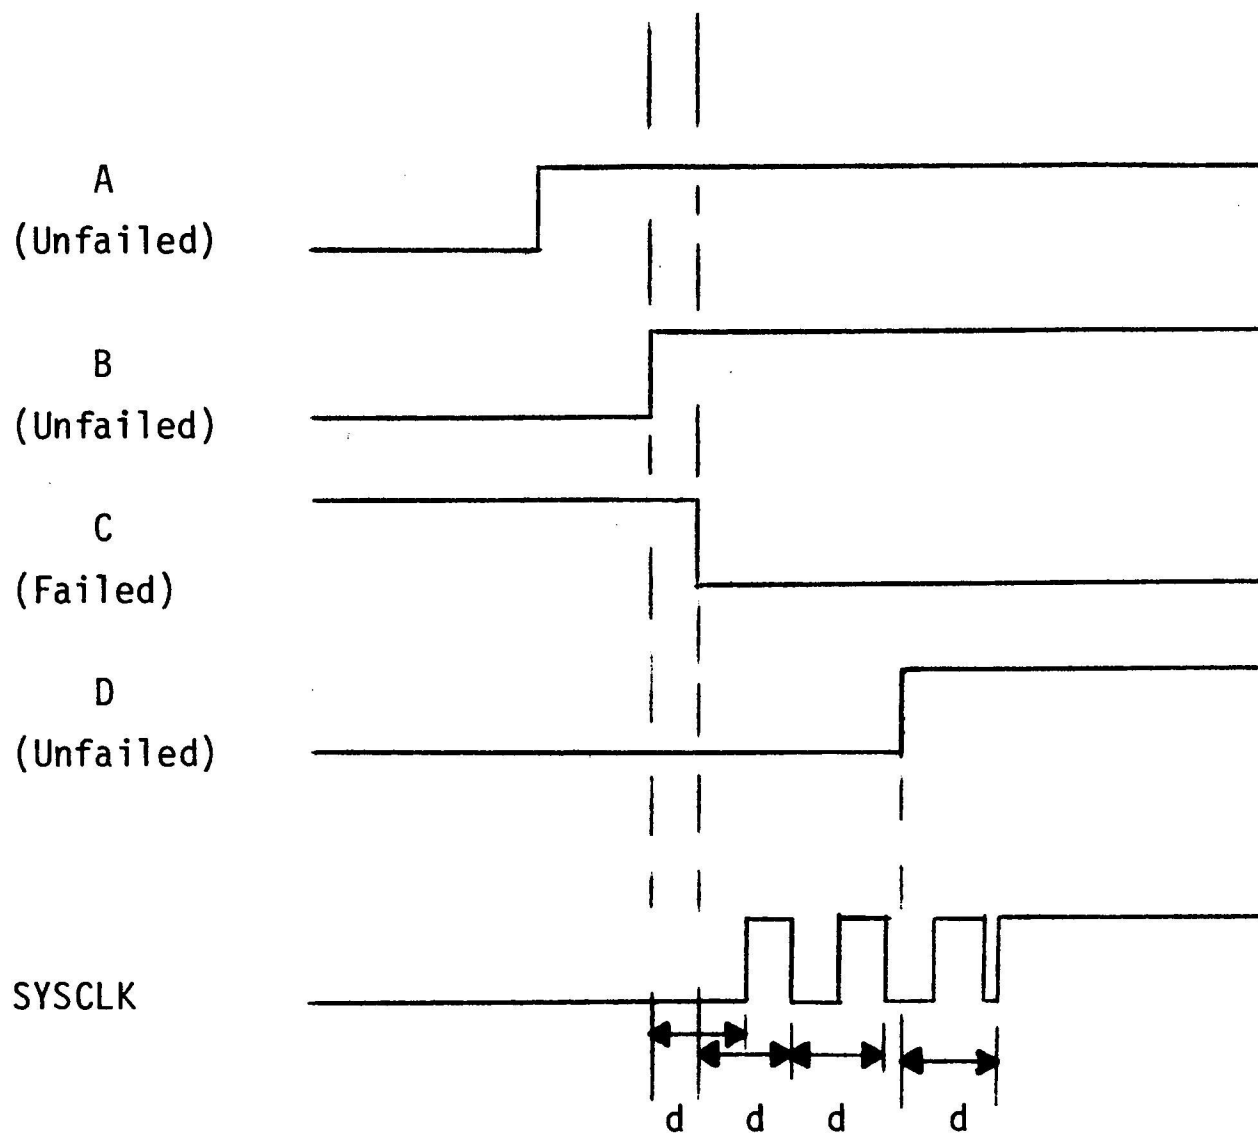


a) Receiver with Feedback.



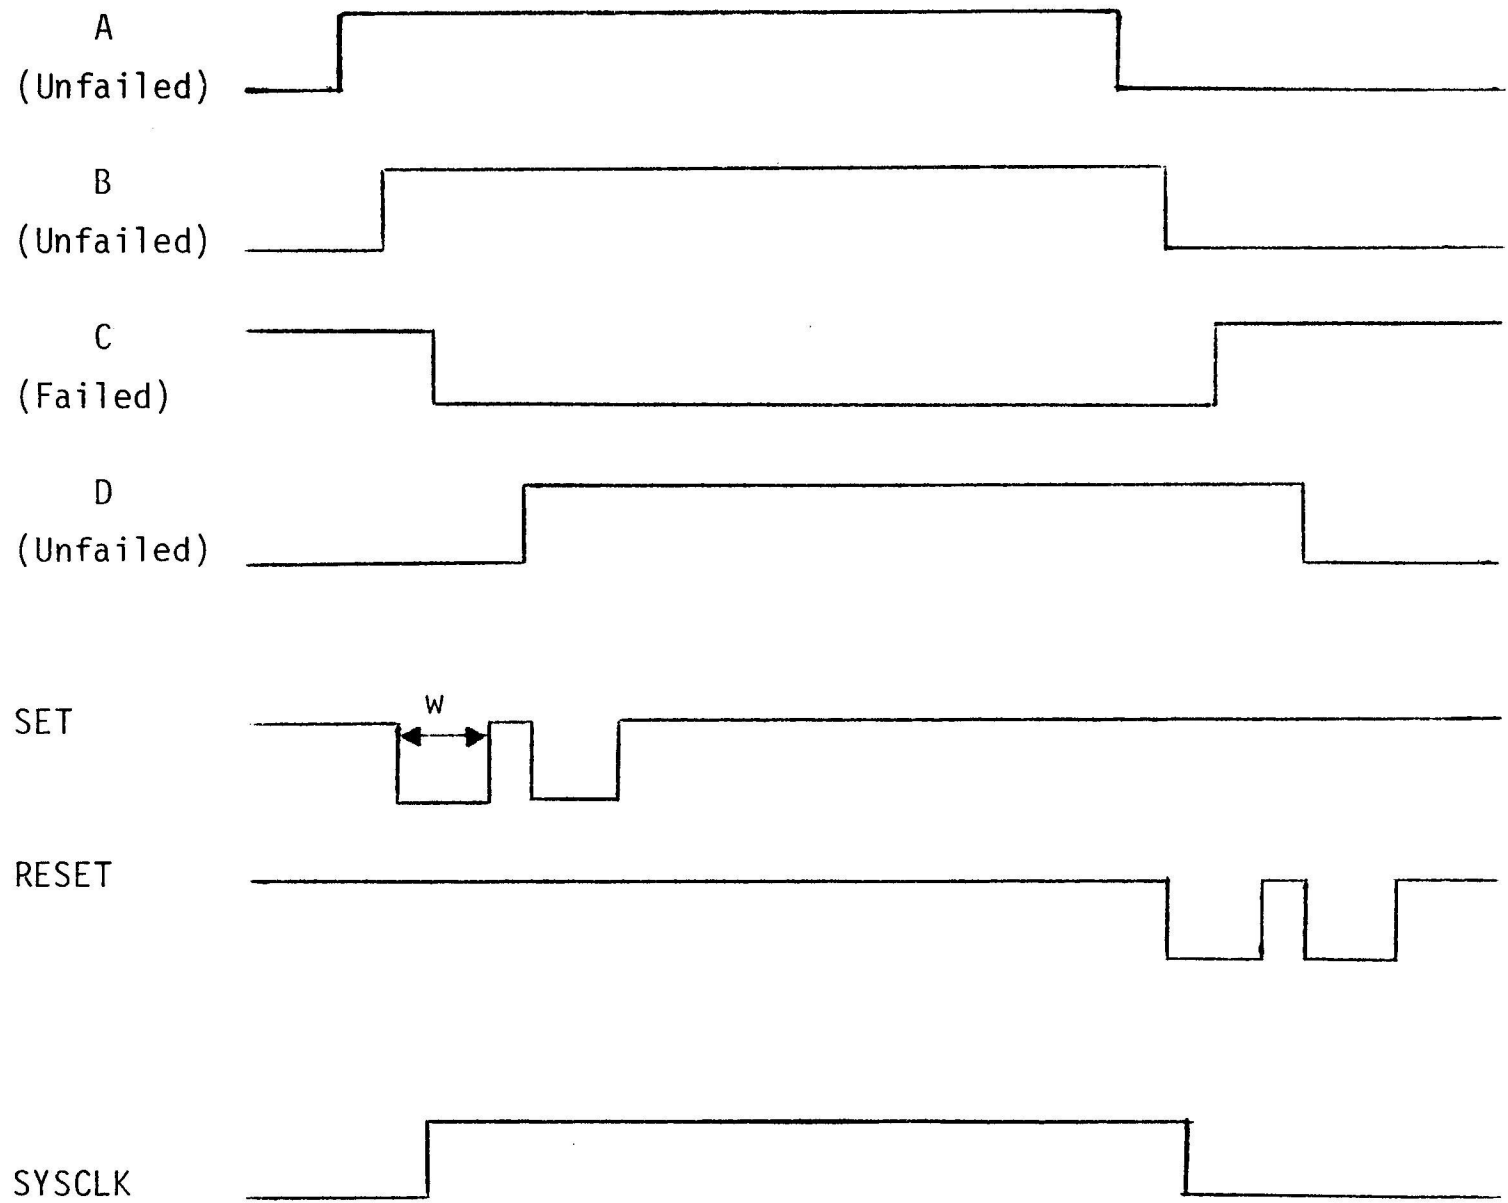
b) Receiver with R-S Flip-Flop.

Figure 1.7 Four-Input Clock Receivers.



d: Propagation Delay of the Major Voter

Figure 1.8 Racing Situation Due to a Failed Clock Signal.



w: Pulse Width of One-Shot.

Figure 1.9 Prevention of Abnormalities with One-Shots.

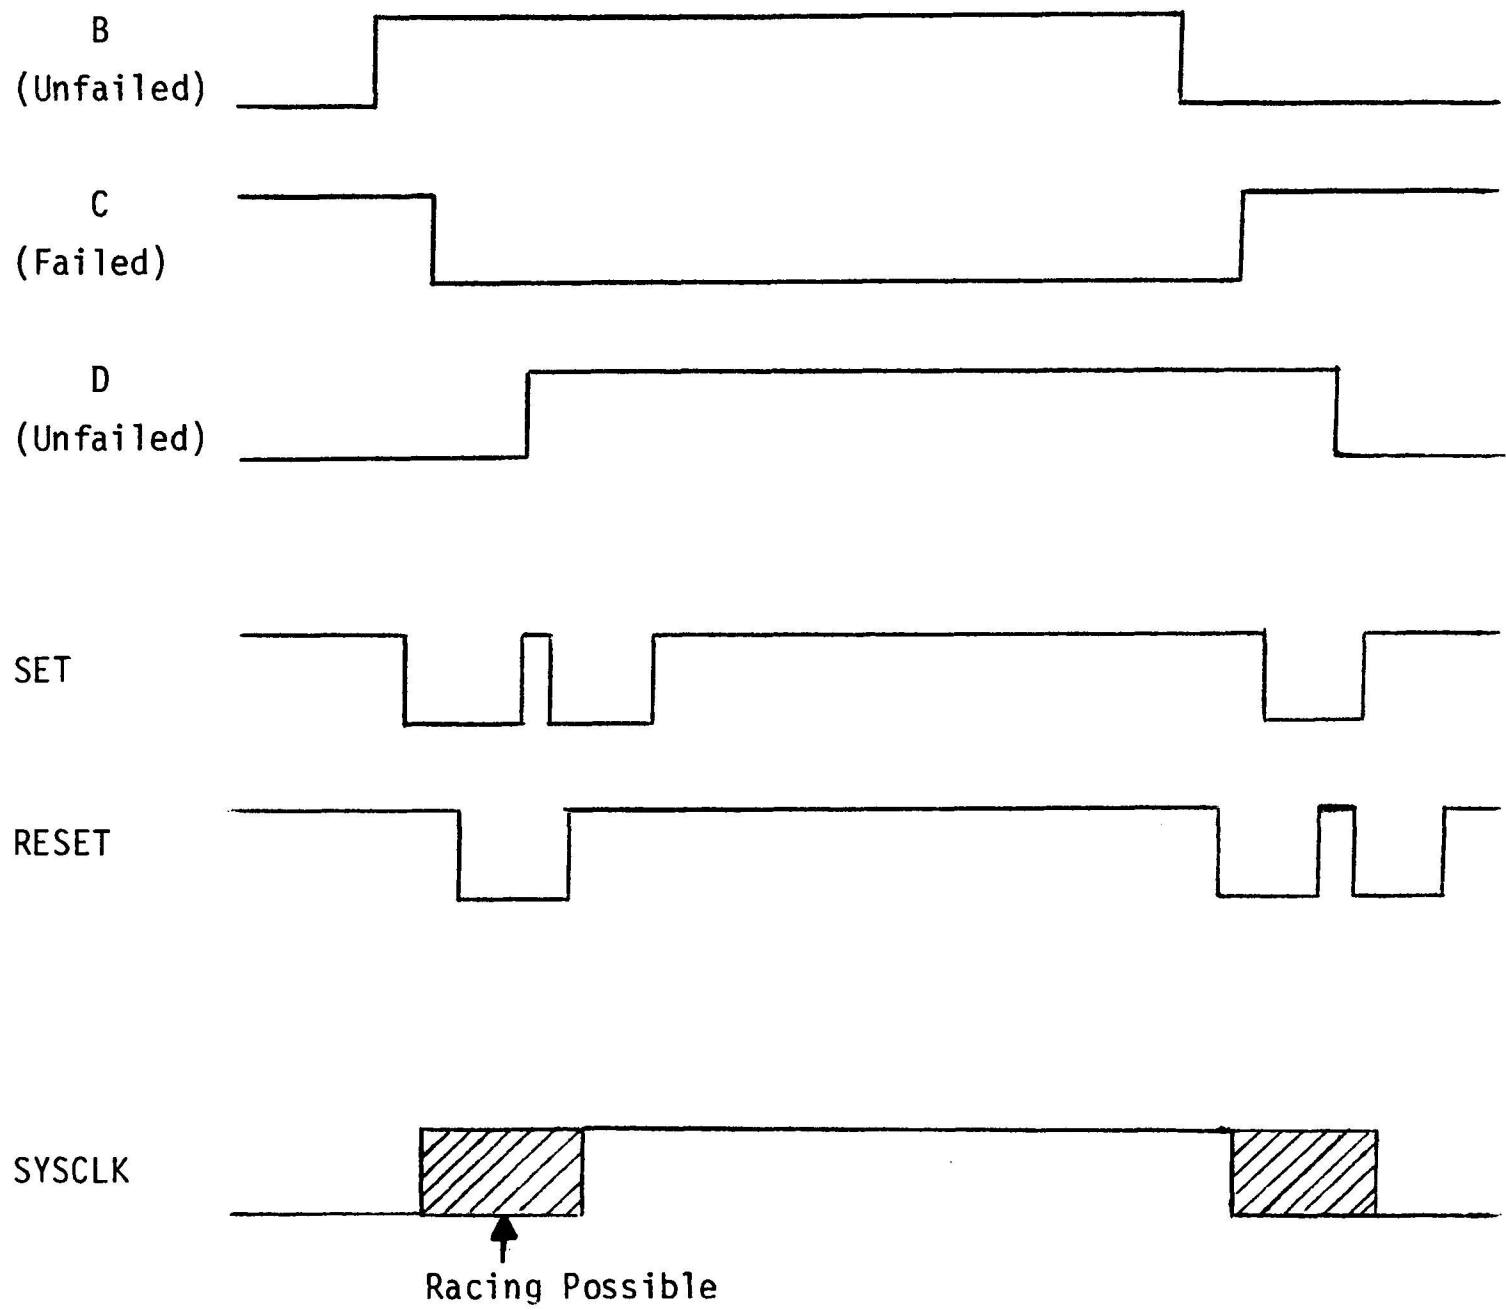


Figure 1.10 Possible Racing Situation.

The state of the secondary clock signal is raised high when a majority of the primary clock signals have produced a leading edge. After the secondary clock goes high, the counting period begins for clock signals that have produced a trailing edge. When the majority of the clocks produce a trailing edge, the state of the secondary clock is lowered. This transition (of the secondary clock) begins the counting period for the leading edges and the process begins again.

The clock receiver proposed in Chapter 2 needs only three input clock signals since no more than one edge of a signal is counted during the counting period for that type of edge. The counting period for leading edges begins when the secondary clock is lowered and terminates when that clock is raised high again. The trailing edge count starts after the secondary clock is raised and ends when it is lowered. Thus, the number of primary clock signals needed for a single-fault-tolerant clock receiver can be reduced from four to three by using edge voting in place of logic level voting.

1.3 The Daly-McKenna Clock


Primary clock signals in most clocking networks are produced independent of one another. However, Daly and McKenna² have outlined an approach in which the production of the clock signals is interrelated, as a result of which they maintain phase-lock despite n clock signals out of $3n+1$ failing. The clocks synchronize themselves internally by making each primary clock signal a function of all the primary clocks. Figure 1.11 illustrates the case for single-fault tolerance which requires four primary clocks. The QUORUM generates two functions Q_2^4 and Q_3^4 which are expressed in Equations 1.4 and 1.5. The logic flow for the DIGITAL DIFFERENTIATION is shown in Figure 1.12. The flip-flop is set after a delay of approximately Δt by Q_2^4 changing low. This can also be expressed as \bar{Q}_2^4 changing high (Eq. 1.6). The flip-flop is reset after a delay, Δt , by the function Q_3^4 becoming high.

$$Q_2^4 = A_1A_2 \vee A_1A_3 \vee A_1A_4 \vee A_2A_3 \vee A_2A_4 \vee A_3A_4 \quad (1.4)$$

$$Q_3^4 = A_1A_2A_3 \vee A_1A_2A_4 \vee A_1A_3A_4 \vee A_2A_3A_4 \quad (1.5)$$

$$\bar{Q}_2^4 = \bar{A}_1\bar{A}_2\bar{A}_3 \vee \bar{A}_1\bar{A}_2\bar{A}_4 \vee \bar{A}_1\bar{A}_3\bar{A}_4 \vee \bar{A}_2\bar{A}_3\bar{A}_4 \quad (1.6)$$

If all inputs are equal ($A = A_1 = A_2 = A_3 = A_4$), we essentially have an interconnected network of crude oscillators whose frequencies are centered around a frequency of $1/2\Delta t$. Equations 1.5 and 1.6 require that only three clocks must be working properly to keep those three clocks synchronized and to force the fourth



Q: QUORUM
 DIG. DIF.: DIGITAL DIFFERENTIATION
 FF: FLIP-FLOP

Figure 1.11 Daly-McKenna Clock.

clock back into synchronization if possible. If not, there are three unfailed primary clocks left but no fault tolerance. Daly and McKenna have determined that for a clock to tolerate n failures, the number of clock elements, N , must be greater than or equal to $3n+1$ (Eq. 1.7). Using the minimum value for N , the QUORUM expressions are defined as Q_y^n and Q_x^n where x and y are given by Eq. 1.8 and 1.9.

$$N \geq 3n+1 \quad (1.7)$$

$$x = n+1 \quad (1.8)$$

$$y = 2n+1 \quad (1.9)$$

The basic guidelines for a system employing the Daly-McKenna approach are as follows:

1. Each A_i ($i = 1, 2 \dots N$) is the output of an R-S flip-flop.
2. A_i is set by Q_x^n going from low to high, or after a delay t , by Q_x^n going from high to low.
3. A_i is reset by Q_y^n going from high to low, or after a delay t , by Q_y^n going from low to high.

Figure 1.13 shows the single-fault-tolerant clock used in CERBERUS, an experimental fault-tolerant multiprocessor built at the C.S. Draper Laboratory. This system was designed using the principles just discussed.⁴ This design, unlike that of Figure 1.12, requires no external starting logic. Phase-locking in both systems is achieved by the hysteretic nature of the design. Since all four outputs are used as the inputs to each clocking element, the clocks synchronize themselves through the quorum functions. A fault in a primary clock element therefore affects only the output of that clock element. The other clock elements are unaffected since their inputs are polled.

1.4 Failure Modes

The basic theme of this thesis is the design of a fault tolerant clocking network in which the clock receivers are capable of tolerating an initial clock signal failure, and then replacing the failed signal with an unfailed spare. In order to facilitate the discussion of the effects of a failed clock signal upon the network, the range of possible failures is limited to seven general examples. The behavior of all clock signal failures is assumed to be a variation or combination




Figure 1.12 Clock Element.

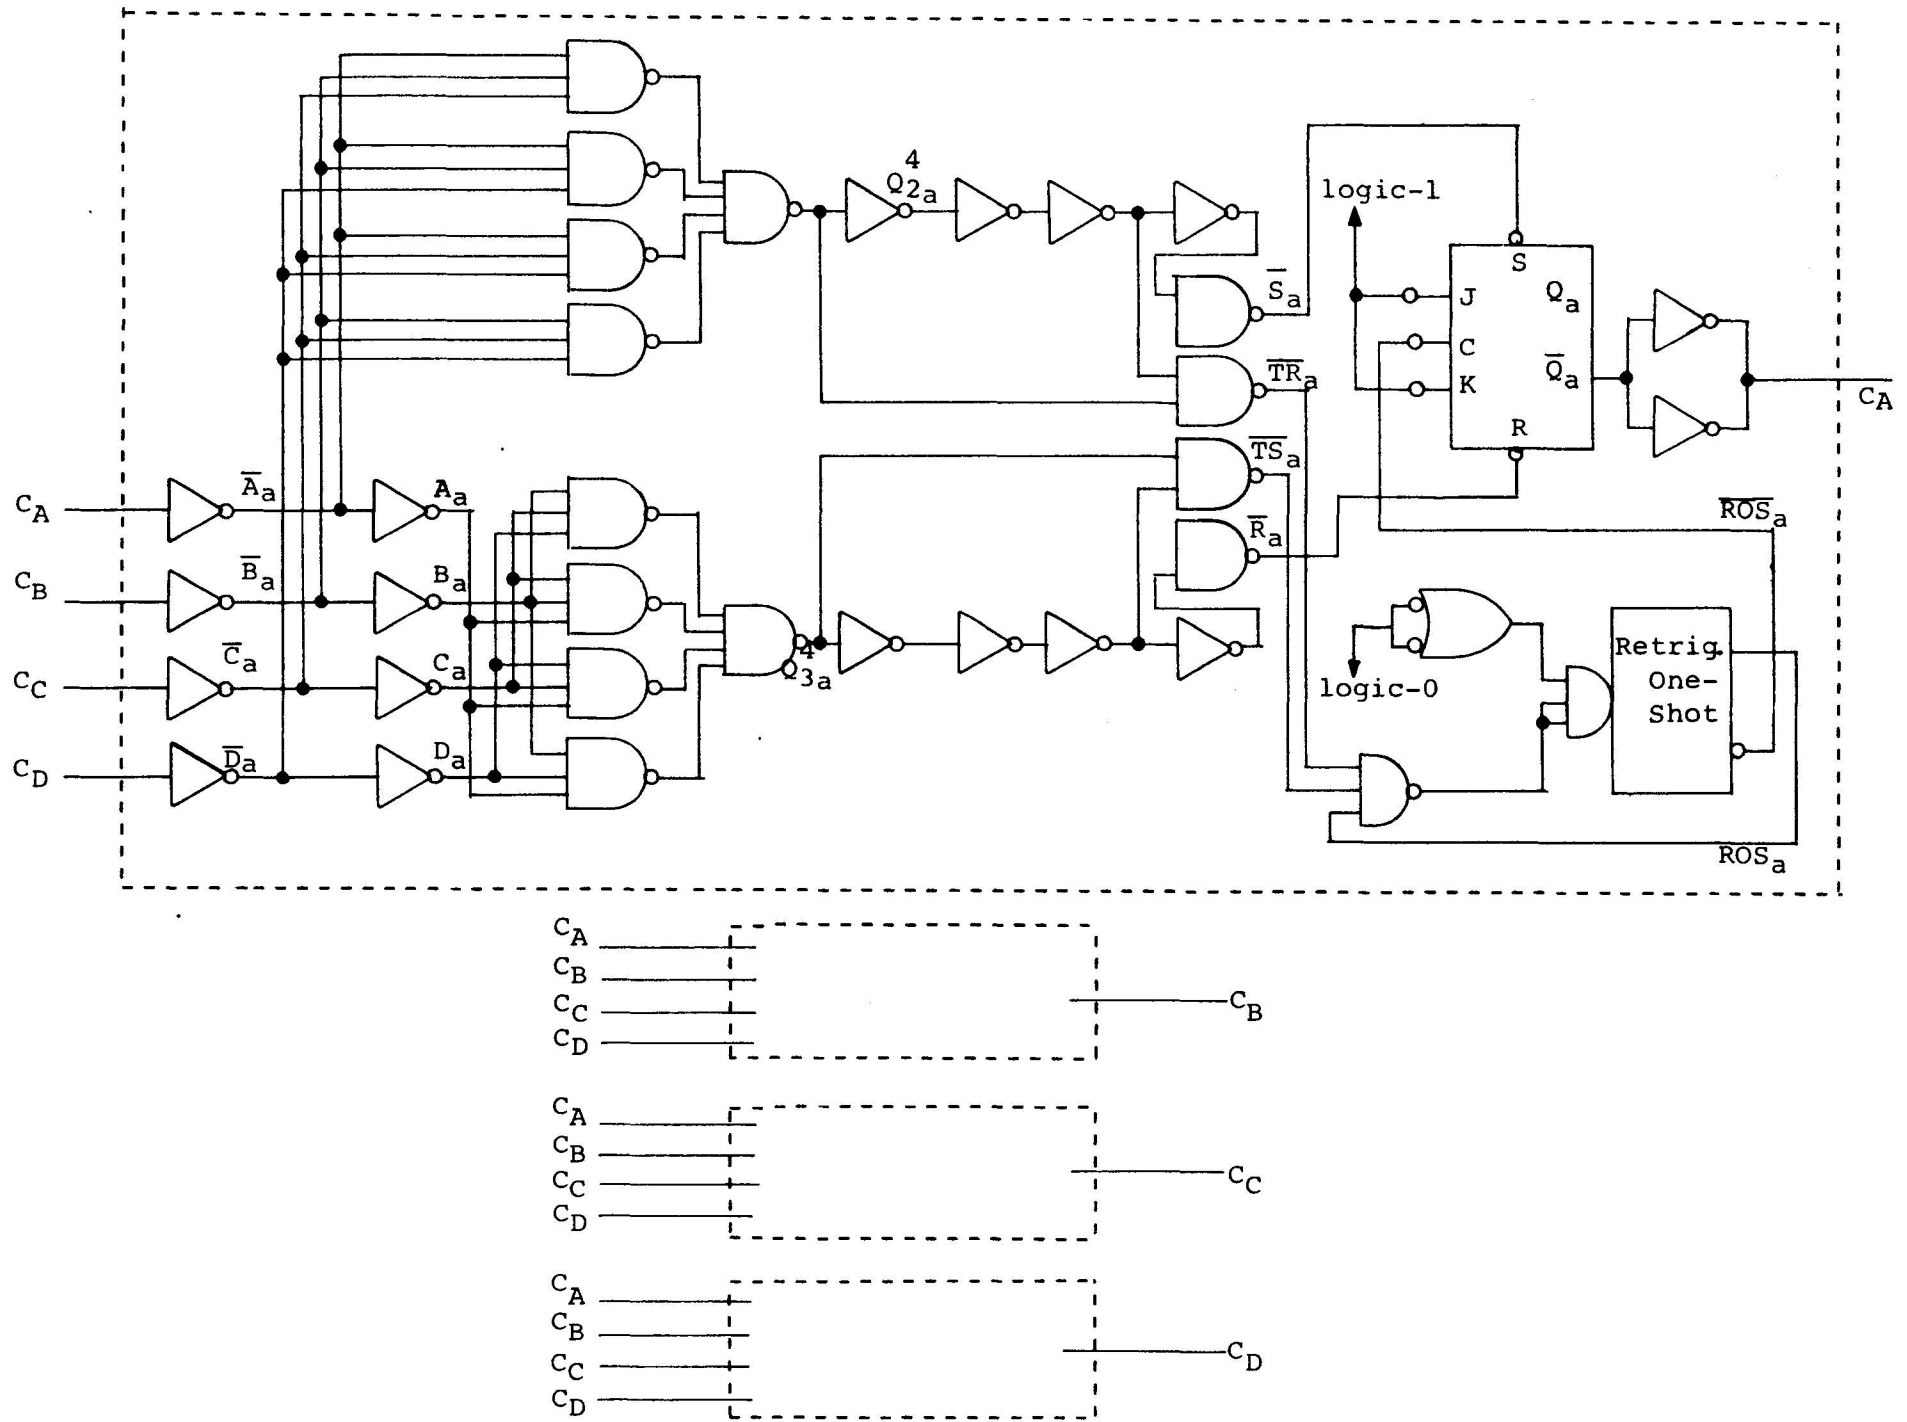


Figure 1.13 McKenna Clock

of the following examples:

1. Stuck-at-logic 0.
2. Stuck-at-logic 1.
3. Metastable.
4. High frequency oscillation (racing).
5. Noise pulses.
6. Phase shift (out-of-phase).
7. Frequency drift (varying frequency).

In the first two examples, the state of the clock signal remains at one of the logic levels. A metastable condition results when the signal remains between the two logic levels for an indefinite period, after which it may go to a logic level. A high frequency oscillation (MODE 2) results from a racing condition in which the frequency is a function of propagation delays and whose duration may be indefinite. The MODE 1 (metastable) and MODE 2 behavior of a failed clock signal usually occurs when the signal is the output of a flip-flop whose inputs are faulty. These two failure modes, MODE 1 and MODE 2, are shown in Figure 1.14 and Figure 1.15 respectively. The fifth example of a clock signal failure is one in which very short pulses (spikes) are produced at random. The study of this type of failure is important when a flip-flop is the recipient of the signal. Noise in the input to a flip-flop can do one of three things depending on the dimensions of the noise spike: change the state of the output, cause the output to become unstable (i.e. MODE 1 and MODE 2), or be ignored completely. A phase shift failure occurs when a clock signal becomes out of phase (by an amount exceeding some tolerance limit). A clock signal failing in this manner changes its state before (lead) or after (lag) the unfailed clock signals. The last example of a clock signal failure occurs when the frequency of a clock begins to vary significantly. (The frequency of all clock signals varies slightly about some nominal frequency, but phase lock is retained if the signal is unfailed).




Figure 1.14 Selected Trajectories of Mode 1 Unstable State.




Figure 1.15 Typical Response with Mode 2 Unstable State.

CHAPTER 2

CLOCK RECEIVER

The purpose of a clock receiver is to produce a fault-tolerant system clock signal from several phase-locked primary clock signals. The clock receiver proposed in this chapter is operationally dependent upon the edges of the primary clocks rather than their logic levels, for reasons discussed in the preceding chapter.

2.1 Architecture

Figure 2.1 shows the proposed clock receiver. The clock receiver may functionally be divided into two sections. The edge detection section consists of the six left-most flip-flops plus the inverters to their left. The voting section consists of the two majority voters and an R-S flip-flop.

2.1.1 Edge Detection

The type of flip-flop used in the edge detection section of the clock receiver is shown in Figure 2.2. The information at the data input terminal (D) is transferred to the output (Q) by the leading edge of the signal to the clock terminal (C), provided that the input to the reset terminal (R) is high. A low at the reset terminal results in a low output (Q) by "overriding" any high from the data terminal that has been transferred or is being transferred. In the clock receiver, the data input is kept high, so that a high is transferred to the output by a leading edge. To transfer data information on a trailing edge, the signal to the clock terminal (C) is inverted. Each of the three primary clocks is connected to two flip-flops as shown in Figure 2.1. This allows one triad of flip-flops to be triggered by the leading edges of the primary clock signals (FF_{LA} , FF_{LB} , FF_{LC}), and the other by the trailing edges (FF_{TA} , FF_{TB} , FF_{TC}). In order to use the flip-flops to detect the edges of the primary clock signals, the flip-flops must be cleared between triggering edges. The method of clearing these flip-flops is discussed in Section 2.2.

2.1.2 Voting

The voting section contains a majority voter for the leading edge-triggered triad, and a majority voter for the trailing edge-triggered triad. The leading

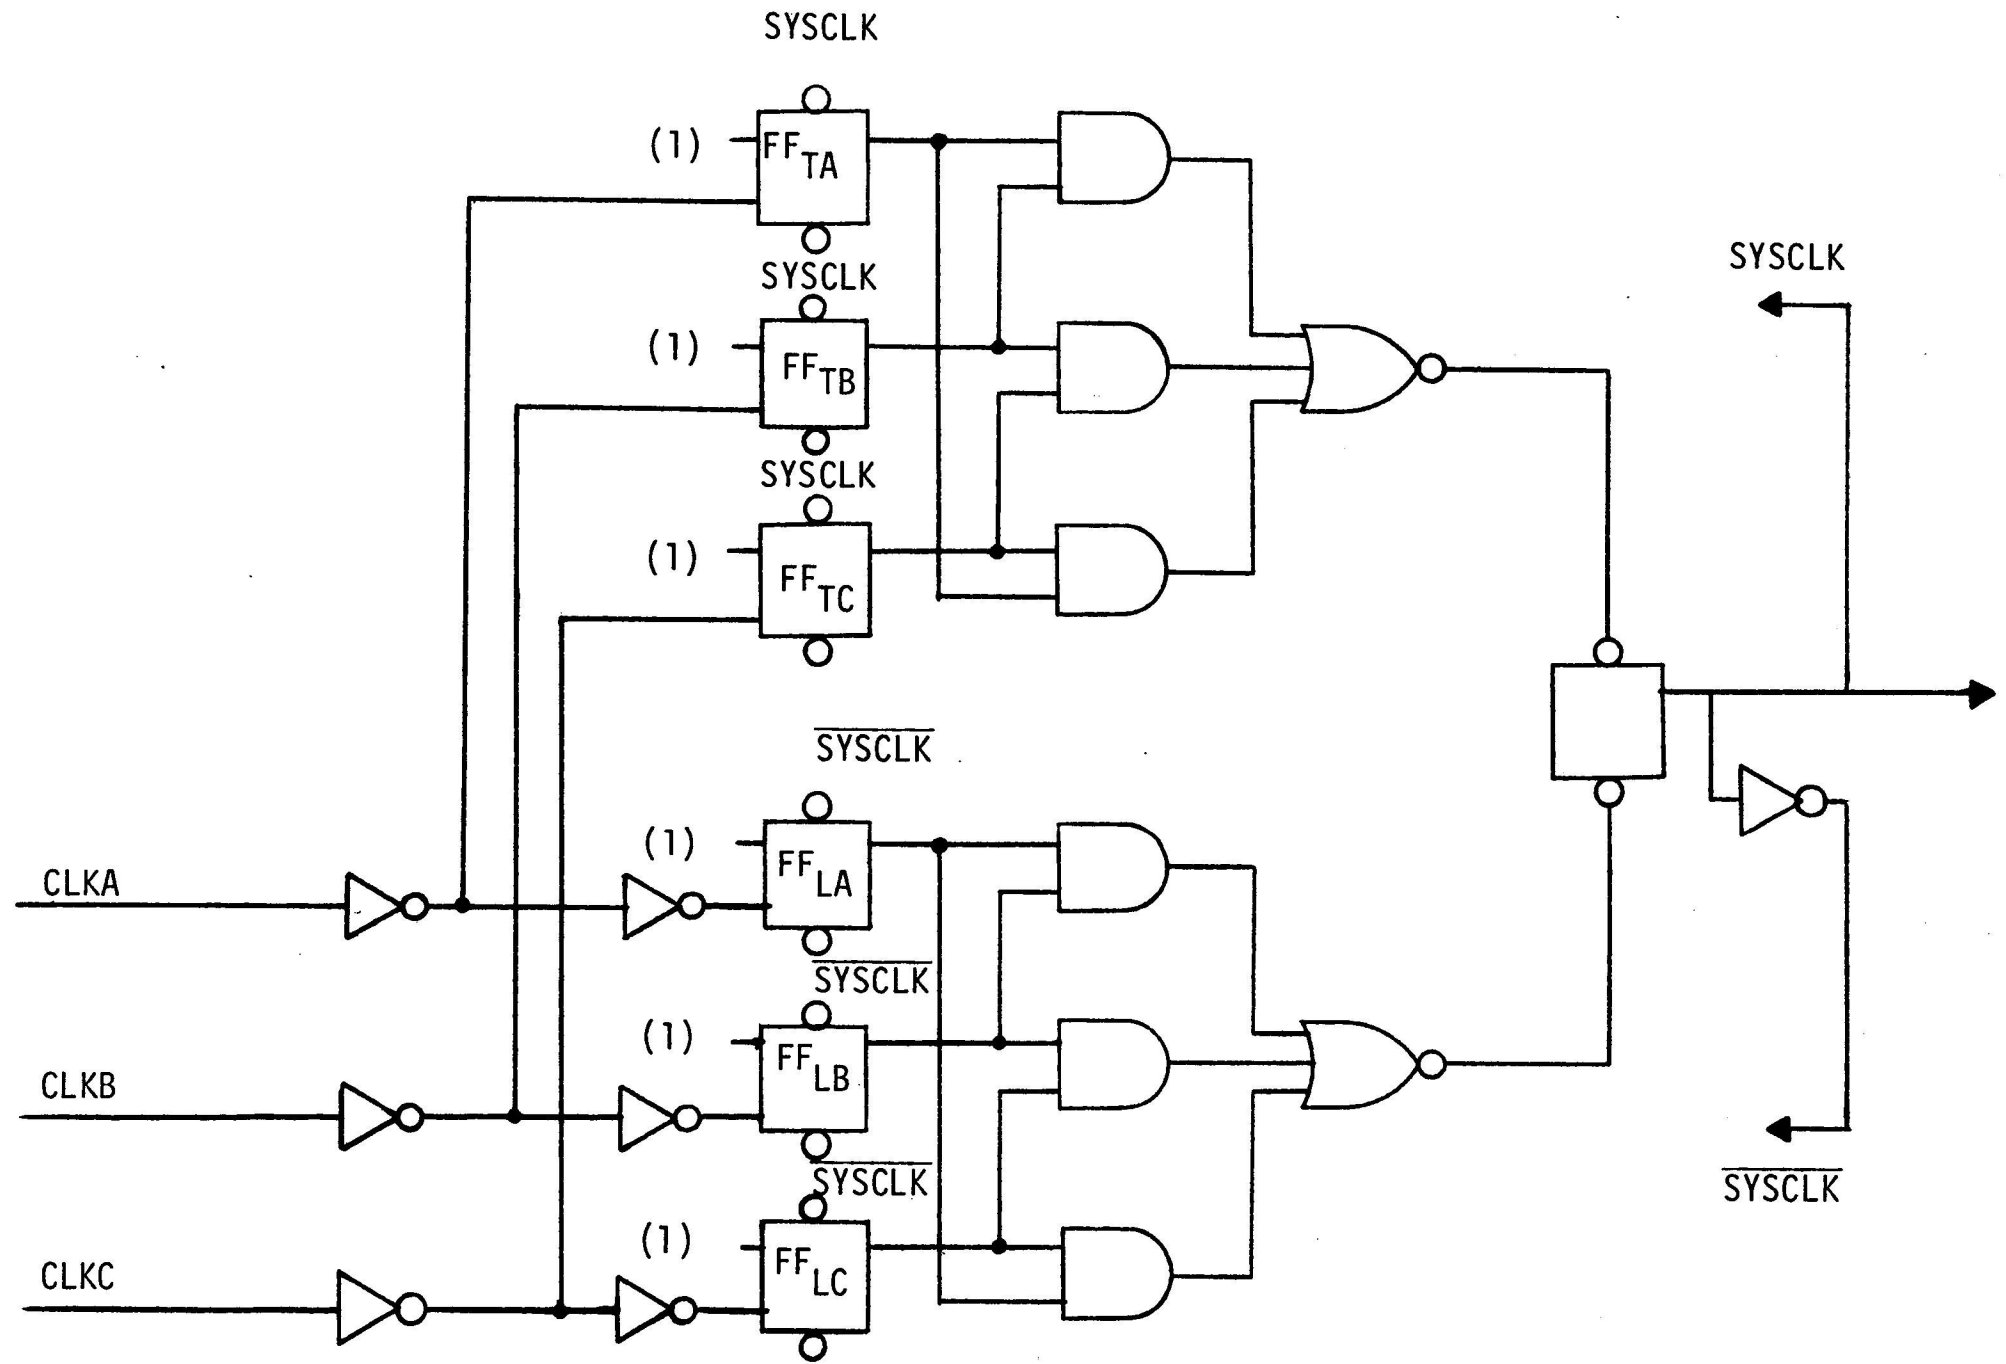


Figure 2.1 The Clock Receiver.




Figure 2.2 Edge-Triggered Flip-Flop.




Figure 2.3 A Voter Within the Clock Receiver.

edge (SET) voter receives the outputs of the three leading edge-triggered flip-flops producing a low output when at least two of the flip-flops' outputs are high. The trailing edge (RESET) voter produces a low when at least two of the trailing edge-triggered flip-flops are high. Figure 2.3 shows the logic design of this type of voter.

The output of the SET voter is connected to the set terminal (S) of the R-S flip-flop. The output of the RESET voter goes to the reset terminal (R). The state of the flip-flop's output is changed by a low input to the set (S) or to the reset (R) terminal. A low to the set terminal raises the output, the secondary clock signal, high whereas a low to the reset terminal clears the output. The state of the output remains unchanged when both terminals are high.

The output of the clock receiver (SYSCLK) is fed back to clear the flip-flops of the T.E. (trailing edge) triad. The inverted output ($\overline{\text{SYSCLK}}$) is used to clear the L.E. (leading edge) triad.

2.2 Fault-Free Behavior

To facilitate the discussion on the behavior of the system clock (SYSCLK), the following assumptions are made.

1. Fault-free primary clock signals are equal.
2. The propagation delays of like (fault-free) components are equal.
3. A component's propagation delay is the same for a high to low transition as for a low to high transition.

Figure 2.4 illustrates the procedure by which the secondary clock signal (SYSCLK) is produced. The following quantities are defined:

CLK _i :	An input (primary) clock signal (i=A,B,C).
Q _{Li} :	A L.E. flip-flop's output (i=A,B,C).
Q _{Ti} :	A T.E. flip-flop's output (i=A,B,C).
SET:	The L.E. voter's output.
RESET:	The T.E. voter's output.
t _i :	Time at which CLK _i goes high (i=A,B,C).
Δt _i :	Amount of time that CLK _i remains high (i=A,B,C).
T _i :	Cycle period of CLK _i (i=A,B,C).
Δ1:	Propagation delay of edge-triggered flip-flop.
Δ2:	Propagation delay of a voter.
Δ3:	Propagation delay of R-S flip-flop.
Δ4:	Propagation delay required to clear a flip-flop.

In the case of fault-free operation, we have:

$$t_A = t_B = t_C$$

$$\Delta t_A = \Delta t_B = \Delta t_C$$

$$T_A = T_B = T_C$$

The procedural steps illustrated in Figure 2.4 can be expressed as follows:

Initial Conditions

$$\text{CLK}_i = Q_{L_i} = Q_{T_i} = \text{SYSCLK} = 0$$

$$\text{SET} = \text{RESET} = 0$$

The timing analysis given these initial conditions is as follows: (Steps numbered for reference purposes.)

$$\text{CLK}_i \rightarrow 1 \quad \text{at } t_i \quad (2.1)$$

$$Q_{L_i} \rightarrow 1 \quad \text{at } t_i + \Delta 1 \quad (2.2)$$

$$\text{SET} \rightarrow 0 \quad \text{at } t_i + \Delta 1 + \Delta 2 \quad (2.3)$$

$$\text{SYSCLK} \rightarrow 1 \quad \text{at } t_i + \Delta 1 + \Delta 2 + \Delta 3 \quad (2.4)$$

$$Q_{L_i} \rightarrow 0 \quad \text{at } t_i + \Delta 1 + \Delta 2 + \Delta 3 + \Delta 4 \quad (2.5)$$

$$\text{SET} \rightarrow 0 \quad \text{at } t_i + \Delta 1 + \Delta 2 + \Delta 3 + \Delta 4 + \Delta 2 \quad (2.6)$$

$$\text{CLK}_i \rightarrow 0 \quad \text{at } t_i + \Delta t_i \quad (2.7)$$

$$Q_{T_i} \rightarrow 1 \quad \text{at } t_i + \Delta t_i + \Delta 1 \quad (2.8)$$

$$\text{RESET} \rightarrow 0 \quad \text{at } t_i + \Delta t_i + \Delta 1 + \Delta 2 \quad (2.9)$$

$$\text{SYSCLK} \rightarrow 0 \quad \text{at } t_i + \Delta t_i + \Delta 1 + \Delta 2 + \Delta 3 \quad (2.10)$$

$$Q_{T_i} \rightarrow 0 \quad \text{at } t_i + \Delta t_i + \Delta 1 + \Delta 2 + \Delta 3 + \Delta 4 \quad (2.11)$$

$$\text{RESET} \rightarrow 1 \quad \text{at } t_i + \Delta t_i + \Delta 1 + \Delta 2 + \Delta 3 + \Delta 4 + \Delta 2 \quad (2.12)$$

$$\text{CLK}_i \rightarrow 1 \quad \text{at } t_i + T_i \quad (2.13)$$

After step 2.12, the clock receiver has returned to its original state. Each step (2.1 through 2.12) is repeated at intervals of T_i , the cycle period.

As the input clock signals become high (2.1), their leading edges trigger the L.E. flip-flops high (2.2). This results in the output of the SET voter going low (2.3), which in turn, raises SYSCLK high (2.4). In this manner, the leading edges of the input clock signals produce a leading edge in the secondary clock after a delay equal to the sum of $\Delta 1 + \Delta 2 + \Delta 3$. As the output of the clock receiver (SYSCLK) becomes high, the inverted output ($\overline{\text{SYSCLK}}$) becomes low. Since the inverted secondary clock signal is sent to the RESET terminals of the leading edge-triggered flip-flops, the outputs of these flip-flops (Q_{L_i}) become low (2.5). Once these outputs are low, the output of the SET voter becomes high again (2.6). In a

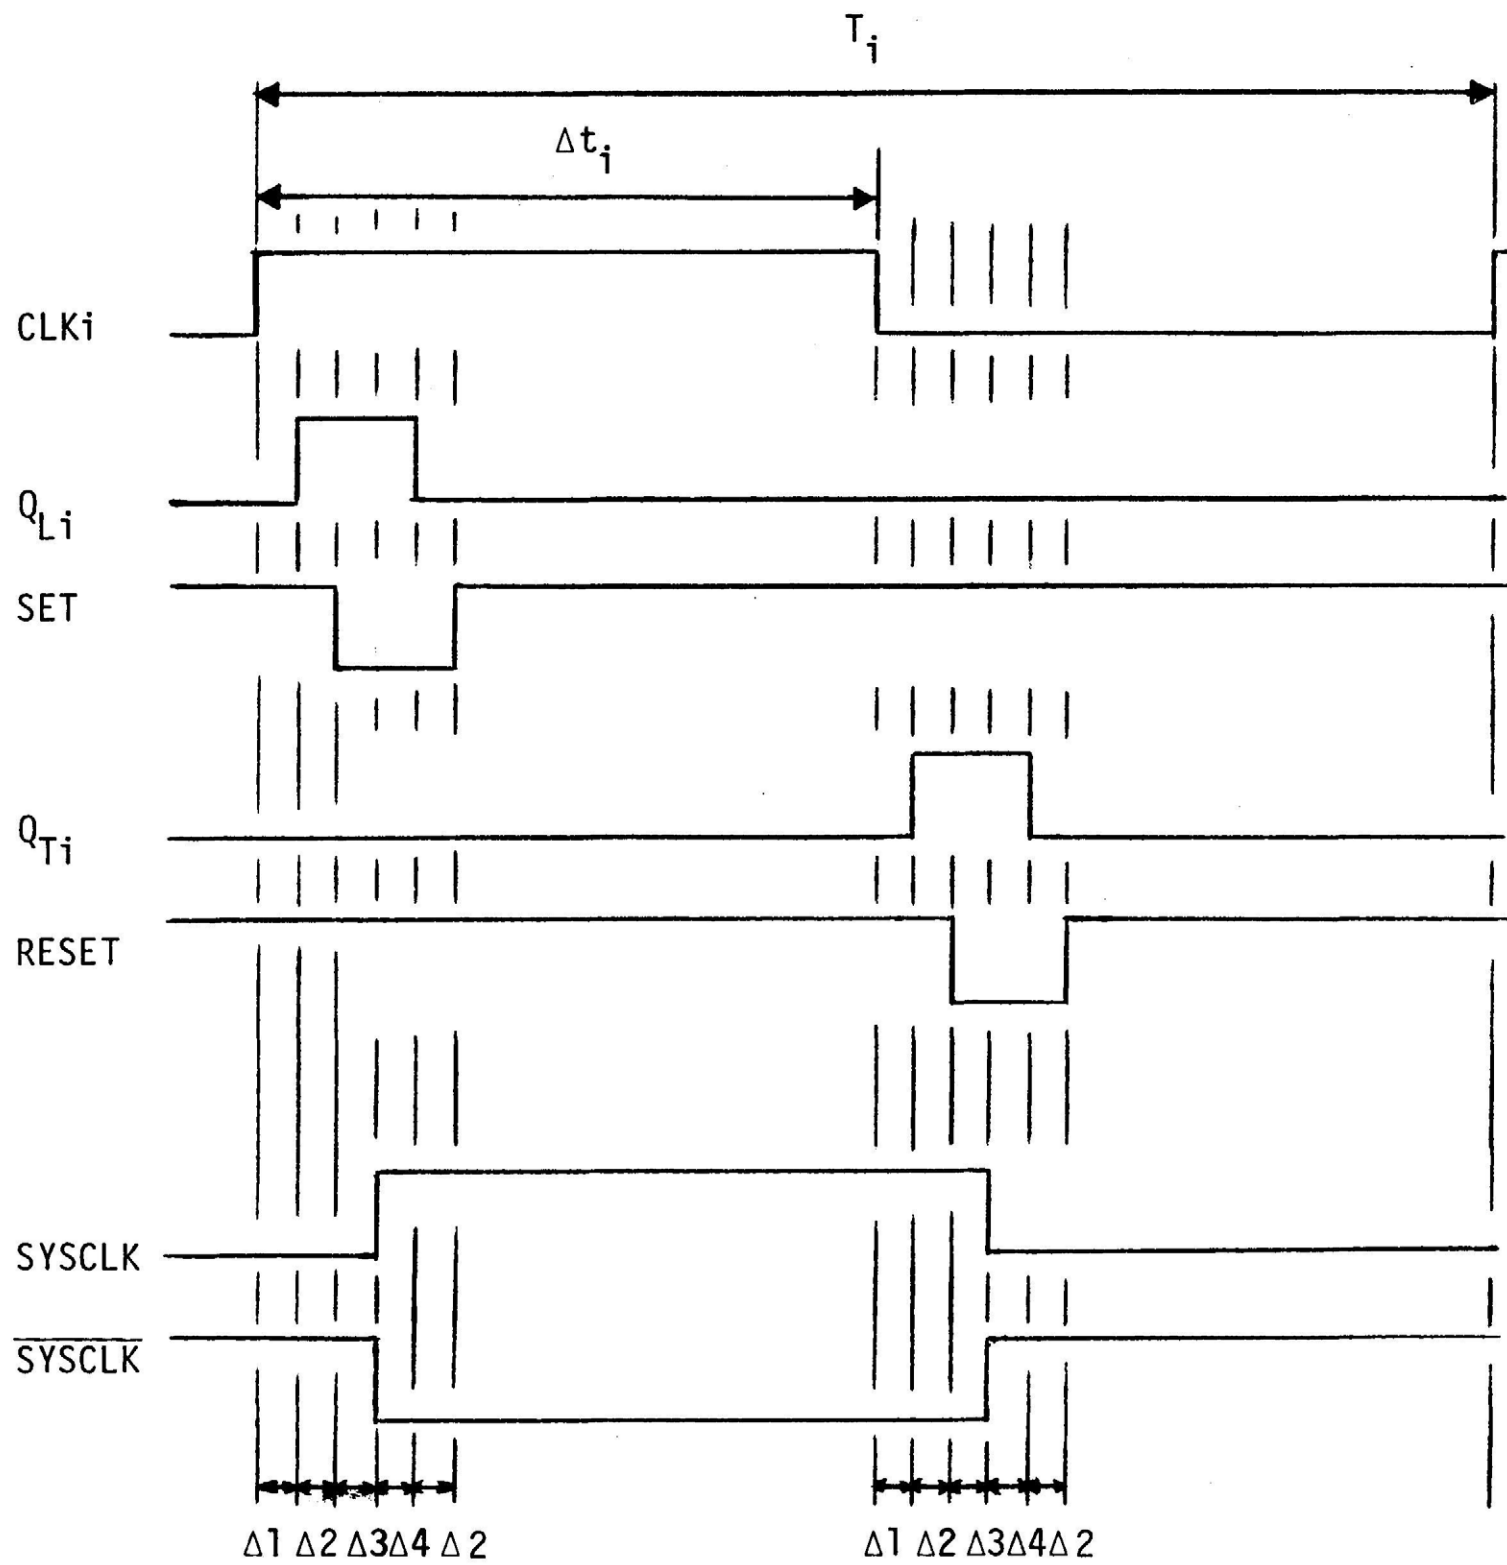


Figure 2.4 The Production of the System Clock.

similar fashion, the trailing edges of the input clock signals produce a trailing edge for the secondary clock. The whole procedure begins again with the next group of leading edges (2.13). It should be noted that the inputs to the R-S flip-flop are not low at the same time. These inputs remain low for the sum of the propagation delays, Δ_2 , Δ_3 and Δ_4 , and must alternate because of the clearing feedback.

2.3 Fault Tolerance

The clock receiver tolerates and masks any fault that occurs in the edge detection section or the input (primary) clocks. This results from the fact that the SET voter and the RESET voter are two-out-of-three majority voters. Each voter consists of three AND gates and a NOR gate as shown in Figure 2.3.

The following relationships concerning Figure 2.3 are listed and numbered for reference purposes.

$$P_{AB} = I_A \cdot I_B \quad (2.14)$$

$$P_{BC} = I_B \cdot I_C \quad (2.15)$$

$$P_{CA} = I_C \cdot I_A \quad (2.16)$$

$$O_V = \overline{P_{AB} \vee P_{BC} \vee P_{CA}} \quad (2.17)$$

$$= \overline{I_A \cdot I_B \vee I_B \cdot I_C \vee I_C \cdot I_A} \quad (2.18)$$

A fault in the edge detection section or in one of the input clocks can affect only one of the voter's inputs (I_A , I_B or I_C). For example, assume that a fault has occurred which can affect I_A . P_{BC} , the product of I_B and I_C , can not be influenced by the fault (2.15). The two other products, P_{AB} and P_{CA} , are functions of one good input and one faulty input (2.14 and 2.16). These products are influenced by the behavior of I_A only when the good input is high. The possible outputs of a voter if the state of I_A is unknown are listed in Table 2.1. The condition of I_A is irrelevant if I_B is equal to I_C . The effects of I_B and I_C not being equal are discussed later.

Most faults that occur in the voting section directly affect the system clock (SYSCLK) and thus can not be tolerated. This is due to the fact that a voter is not redundant.

The occurrence of the first fault usually results in the system clock either failing or losing its fault tolerance. The fact that a clock receiver appears to be functional is no assurance that its fault-tolerant capability is intact. The development of further faults is likely to alter the system clock from its expected (fault-tolerated) behavior.

TABLE 2.1
RECEIVER VOTER WITH AN UNKNOWN INPUT

I_A	I_B	I_C	P_{AB}	P_{BC}	P_{CA}	O_V
UNK	LOW	LOW	LOW	LOW	LOW	HIGH
UNK	LOW	HIGH	LOW	LOW	UNK	UNK
UNK	HIGH	LOW	UNK	LOW	LOW	UNK
UNK	HIGH	HIGH	UNK	HIGH	UNK	LOW

2.3.1 Primary Clock Failure

The clock receiver masks the failure of any one of the three input clock signals (CLKA, CLKB and CLKC). For example, if CLKA fails, the remaining two clocks are sufficient to insure the validity of the system clock (CLKA's failure is tolerated). Figure 2.5 shows the effect on the clock receiver of removing CLKA (CLKA=0). Comparing Figure 2.5 with Figure 2.1, it can be seen that two flip-flops and four AND gates have been removed. The outputs of these "removed" components are low and will remain low because CLKA was removed. The other primary clock signals, CLKB and CLKC, allow the clock receiver (Figure 2.5) to perform in the same manner as discussed in Section 2.2.

A primary clock failure may be considered as the improper production of triggering edges since the primary clocks are used with edge-triggered flip-flops. These flip-flops are affected by a failed clock signal in one of the following ways:

1. The clock signal contains no edges capable of triggering a flip-flop. This results in the output of the flip-flop being low.
2. The clock signal contains edges capable of triggering a flip-flop, but not at the proper time. The flip-flop's output is determined by the state of the clear input signal at the time the edge is produced.
3. The clock signal triggers a flip-flop in a manner which results in instability. The flip-flop's output is considered to be indeterminate.

In Section 1.4, seven illustrative failure modes were described to help show the range of primary clock failures. These examples are now used to help illustrate the behavior of a flip-flop triggered by a failed clock signal. The first two, stuck-at-logic 0 (s-a-0) and stuck-at-logic 1 (s-a-1), are failures which result in the clock signal containing no triggering edges. The outputs of both flip-flops triggered by this clock signal are low since they are still being cleared, but are no longer being triggered. Note that a clock signal which is s-a-1 has the same effect on a flip-flop as it would if it was s-a-0 (or if the clock signal were removed as in Figure 2.5). Thus, if two clock signals fail, one in the s-a-0 condition and the other in the s-a-1 state, the system clock fails, whereas, if logic level voting were used, these two clock failures would be tolerated, thus making the fault detection process more complex and less reliable.

Referring to Figure 1.14, a Mode 1 type clock signal failure (metastable) may result in no triggering edges (3,4,5) or in the flip-flop being triggered at the wrong time (1,2). There is also a finite probability that the fault may propagate

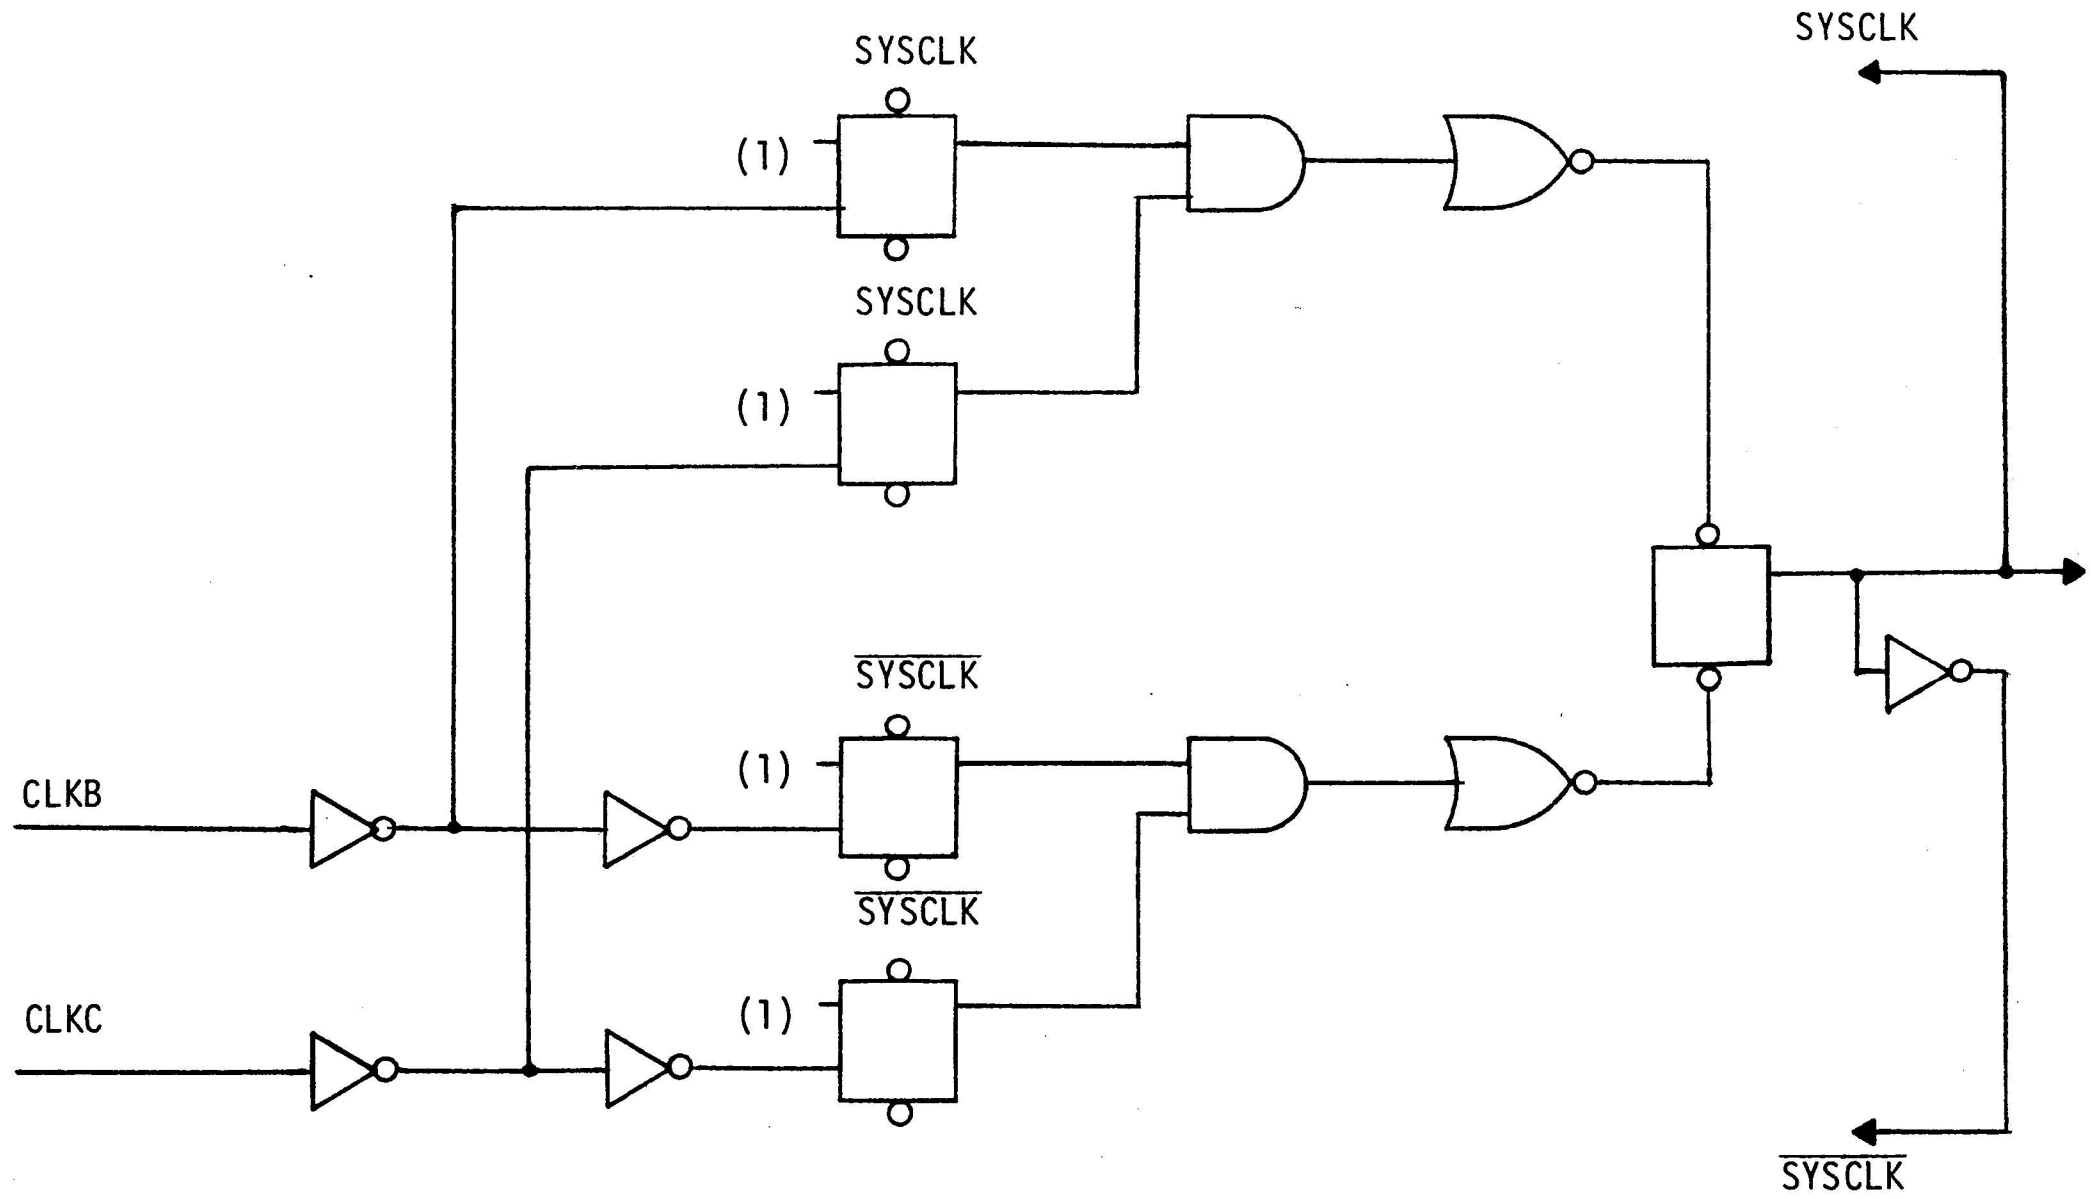



Figure 2.5 The Clock Receiver with One Signal Removed.

through the flip-flop. The effect of a clock signal producing no edges capable of triggering a flip-flop has already been discussed. To facilitate the discussion on triggering edges arriving at the wrong time (relevant to the two unfailed clocks), the clear input signal to an edge-triggered flip-flop is divided into four (time) regions (see Fig. 2.6). In Region #1, the clear input is high, so any edge which arrives during the period will trigger the flip-flop high. The flip-flop's output remains high as long as the clear input remains high, once it has been triggered. In Region #3, the clear input is low, so any clock edge which arrives during this time is ignored. An edge arriving in Region #2 produces a "spike" in the output of the flip-flop whose height may vary from the nominal high to the nominal low level, depending on its arrival time. In this region a low clear input overrides the transfer of a high from the data input to the output. In the last region, the flip-flop's output must be considered indeterminate since the clear input is neither high nor low. In the last case, and for any case in which the clock failures are propagated through the flip-flop, the output is considered indeterminate, except when the clear input is low.

A Mode 2-type signal failure may also affect a flip-flop in any one of the three ways. An edge-triggered flip-flop has a voltage threshold that must be exceeded for a certain amount of time to enable the detection of an edge. If this amount is exceeded, the edge triggers the flip-flop. However, if this amount is not exceeded, the flip-flop either may not be triggered, or its output may become unstable (indeterminate).

A failure which results in a clock signal having noise (short pulses) can affect a flip-flop in any of the three ways, depending on the width of the noise pulses. However, most noise is not of sufficient width to affect an edge triggered flip-flop except, perhaps, for Emitter Coupled Logic. A clock signal out of phase with the unfailed clocks can produce triggering edges at improper times. The output of the flip-flop is dependent upon the region in which the edge arrives. A clock whose frequency drifts from its nominal frequency will produce triggering edges, unless its frequency becomes so high that the edges of the clock can not trigger the flip-flop. Table 2.2 summarizes the effects of clock failures upon an edge-triggered flip-flop.

As previously mentioned (Section 2.3), unfailed voter inputs (the flip-flops' outputs) are not always equal. Consider I_A to be the output of a flip-flop whose driving clock signal has failed. I_A will be low if this failure results in the flip-flop receiving no triggering edges. If the flip-flop receives triggering edges in the third region (when the clear input is low), I_A will again be low. Table 2.3 lists the possible outcome of the voter given that the clock failure results in I_A



EXAMPLES:

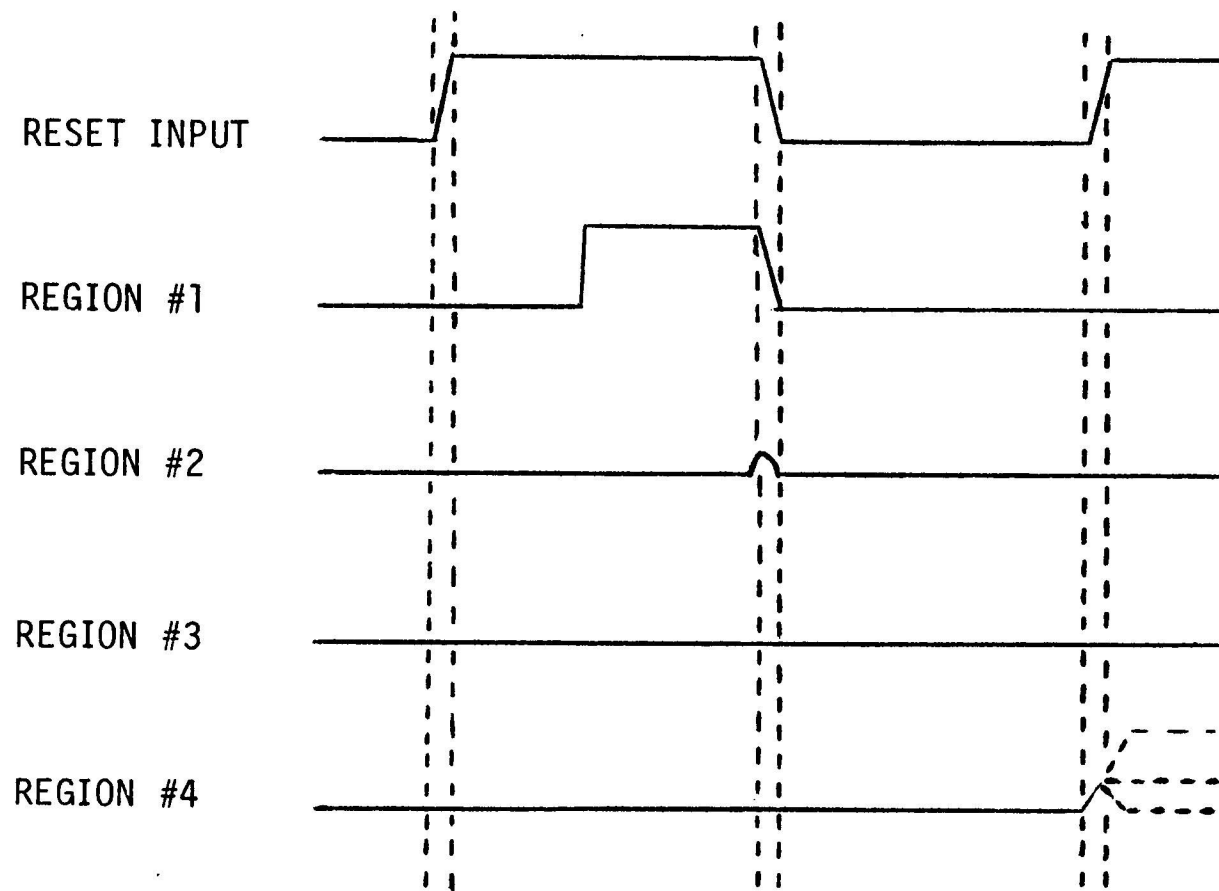


Figure 2.6 The Reset Feedback.

TABLE 2.2

FLIP-FLOP'S OUTPUT RESULTING FROM A CLOCK
SIGNAL FAILURE

FAILURE	NO TRIGGERING EDGES	IMPROPER TRIGGERING EDGES	FAULT PROPAGATION
s-a-0	possible		
s-a-1	possible		
Mode 1	possible	possible	possible
Mode 2	possible	possible	possible
Noise	possible	possible	possible
Phase Shift		possible	
Frequency Drift	possible	possible	possible

being low. From this table, it can be seen that the state of the system clock will change only when both of the unfailed clocks have changed.

A triggering edge arriving in region #1 triggers its flip-flop before the two unfailed clocks trigger theirs. Using CLKA again as the failed clock, I_A would go high before either I_B or I_C . Table 2.4 lists the possible outcome of the voter given that the clock failure results in I_A becoming high before the two unfailed inputs. Since only two inputs need to be high in order to lower the voter's output, the state of the system clock may be changed by either of the unfailed voter inputs. A triggering edge which arrives in Zone #2 (the falling edge of the clear input) triggers its flip-flop after the other two have been triggered. The two good voter inputs are already in the process of changing the state of the system clock when the bad input goes high. The effects of a triggering edge in Region #4 (the rising edge of the clear input) may trigger its flip-flop (Table 2.4 applies) or it may not trigger its flip-flop (Table 2.3 applies) or it may cause its flip-flop to behave in an unpredictable manner (fault propagation). Table 2.1 shows the voter's outcome for the case when the state of the bad voter input is unknown. The output of the voter is unknown from the time the first good voter input becomes high until the time the second good input becomes high. Thus, it is possible, given that unfailed clock signals are only approximately equal, that a failed clock signal may cause an abnormality in a voter's output during the time that the output is unknown.

The outputs of the two voters are used as the inputs to an R-S flip-flop. If Table 2.3 applies to the voter's output, then the system clock follows the second unfailed clock signal. If Table 2.4 applies, then the system clock follows the first unfailed clock signal. If Table 2.1 applies, then the system clock may follow any one of the three clock signals, or it may fail. It is desirable to keep the possibility of the last case as small as possible. An unfailed R-S flip-flop requires special combinations of the width and the voltage level of an input pulse in order to cause the system clock to fail. If abnormalities are present in a voter's output (an input to the R-S flip-flop) as the result of a single signal failure, the system clock will fail (contain abnormalities) only when these abnormalities have the right properties (duration and voltage level). Thus, a single failed clock signal can cause the system clock to fail only if the fault is propagated through the edge triggered flip-flop. If TTL logic is used, a flip-flop forced to fail by a faulty input will produce an output similar to the one shown in Figure 1.15. The frequency of this output is a function of the propagation delays within the flip-flop. If proper precautions are taken, a racing condition within the edge-triggered flip-flop resulting from a failed clock signal will not produce the necessary R-S flip-flop input to cause the system clock to fail. It is important to note that a single primary clock failure is not apt to produce this failure mode

TABLE 2.3
RECEIVER VOTER WITH A S-A-0 INPUT

I_A	I_B	I_C	P_{AB}	P_{BC}	P_{CA}	O_V
LOW	LOW	LOW	LOW	LOW	LOW	HIGH
LOW	LOW	HIGH	LOW	LOW	LOW	HIGH
LOW	HIGH	LOW	LOW	LOW	LOW	HIGH
LOW	HIGH	HIGH	LOW	HIGH	LOW	LOW

TABLE 2.4
RECEIVER VOTER WITH A S-A-1 INPUT

I_A	I_B	I_C	P_{AB}	P_{BC}	P_{CA}	O_V
HIGH	LOW	LOW	LOW	LOW	LOW	HIGH
HIGH	LOW	HIGH	LOW	LOW	HIGH	LOW
HIGH	HIGH	LOW	HIGH	LOW	LOW	LOW
HIGH	HIGH	HIGH	HIGH	HIGH	HIGH	LOW

in several receivers at the same time.

2.3.2 Clock Receiver Failure

Failures in the clock receiver are divided into three categories according to the effect that the failure has on the output of a component. They are as follows:

1. Stuck-at-logic 0 (zero).
2. Stuck-at-logic 1 (one).
3. Indeterminate (UNK).

The output of an edge-triggered flip-flop is a voter input (Section 2.3). The discussion on the failure of a single clock signal (Section 2.3.1) may be applied to the failure of an edge triggered flip-flop, with the understanding that only one voter (one input to the R-S flip-flop) is involved. Table 2.3, Table 2.4, and Table 2.1, respectively, may be applied to the three failure categories listed above for the case of a flip-flop triggered by CLKA failing.

Most failures in the voting section of the clock receiver will directly affect the system clock. Table 2.5 lists the effect on the system clock for a given voting section component in each of the three failure categories. The only failure tolerated in the voting section is for one of the AND gates to be stuck-at-logic 0.

2.3.3 Multiple Failures

The system clock may or may not be able to tolerate more than one failure in its clock receiver or the input (primary) clocks. In a system designed with single-fault tolerance, it is necessary to be able to detect, locate and correct the first failure before a second one occurs. The possibility of more than one failure occurring simultaneously is considered to be remote, unless they result from the same cause (i.e. bridging, physical damage).

From the discussion of single-failure tolerance (Sections 2.3.1 and 2.3.2), a failure may cause unfailed components to behave as if they had failed. For example, if a clock fails in the s-a-0 position, the outputs of the two flip-flops triggered by the failed clock are stuck-at-logic 0, even though the two flip-flops have not failed. It is not possible to distinguish between these "pseudo" failures and the actual failure that caused them. Figure 2.5 shows the minimum number of components that must remain good to insure the validity of the system clock. Thus, in certain cases, more than one failure can be tolerated. One approach to fault detection is to induce a test failure leaving the system clock vulnerable to an actual failure. The test failure and an actual failure would not be tolerated whereas the test failure alone would be. This method is explored in the next chapter.

TABLE 2.5
 FAILED COMPONENT WITHIN THE VOTING SECTION

FAILED COMPONENT	COMPONENT'S OUTPUT	EFFECT ON THE SYSTEM CLOCK
One of the AND gates in the set voter	Stuck-at-logic 0 Stuck-at-logic 1 Indeterminate	None (failure tolerated) Stuck-at-logic 1 Indeterminate
The NOR gate in the set voter	Stuck-at-logic 0 Stuck-at-logic 1 Indeterminate	Stuck-at-logic 1 Stuck-at-logic 0 Indeterminate
One of the AND gates reset voter	Stuck-at-logic 0 Stuck-at-logic 1 Indeterminate	None (failure tolerated) Stuck-at-logic 0* Indeterminate
The NOR gate in the reset voter	Stuck-at-logic 0 Stuck-at-logic 1 Indeterminate	Stuck-at-logic 0* Stuck-at-logic 1 Indeterminate
The R-S flip-flop (TTL 7474)	Stuck-at-logic 0 Stuck-at-logic 1 Indeterminate	Stuck-at-logic 0 Stuck-at-logic 1 Indeterminate

*With TTL (7474 as R-S flip-flop) a low set pulse overrides the reset input (s-a-0).

2.4 Masking

The clock receiver not only tolerates a primary clock failure (subject to conditions just discussed) but also covers up any evidence of its occurrence. The system clock continues to behave as if no failure had occurred. Such a failure is said to have been masked. The following chapter explores various means by which a failure, masked or unmasked, can be exposed.

CHAPTER 3

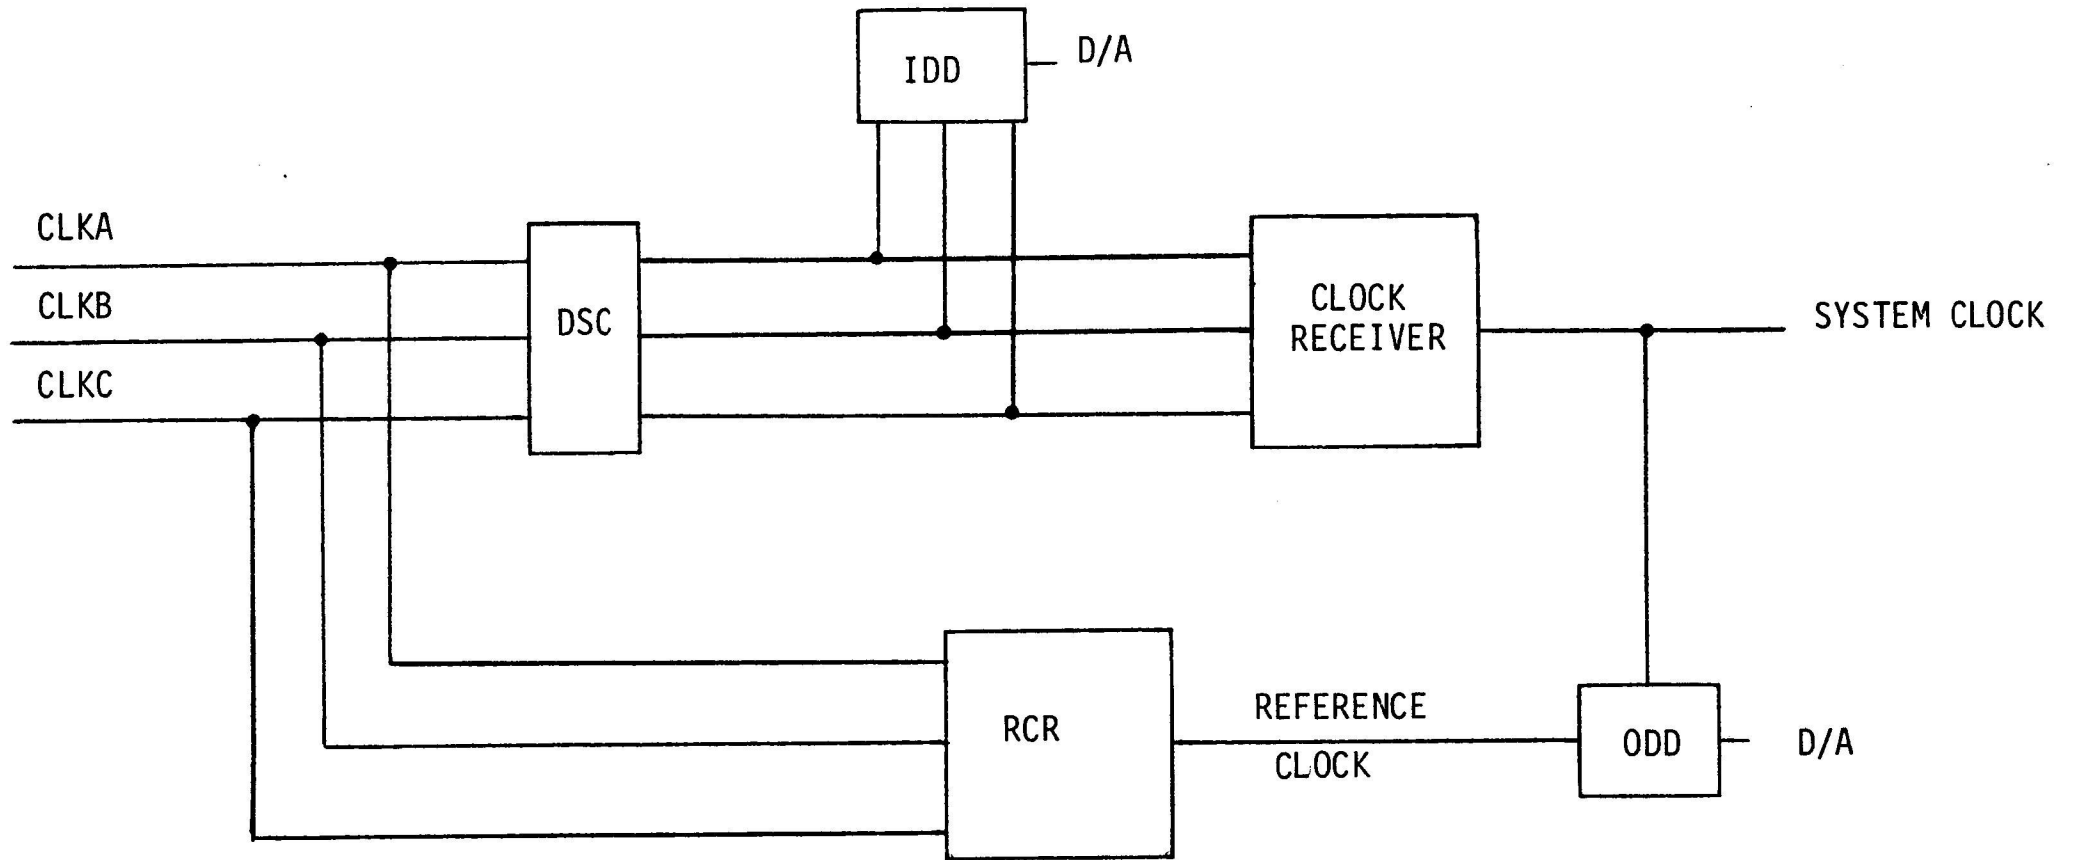
FAULT DETECTION

The clock receiver proposed in the preceding chapter tolerates any failure to itself not in the voting section or to a clock signal inputted to it. In the process of being tolerated, a failure is often masked. That is, its existence can not be determined by simply monitoring the behavior of the system clock. To expose such a failure, a special simulated failure can be generated, which, combined with an actual failure, results in the system clock deviating from its expected (fault-free) behavior. The simulated failure in the absence of an actual failure is tolerated and does not affect the performance of the system clock. Two methods of generating simulated failures are discussed in this chapter along with various schemes of exploiting their diagnostic capabilities. In the first method a simulated failure is achieved by delaying the signal from an input clock. Since the receiver uses three input clock signals to produce the system clock, the clock receiver can be subjected to eight simulated failure conditions. The second approach involves substituting a special "Fail Clock" for an input clock signal. Designs will be shown for various testing components, meant to serve as a guide to help illustrate the underlying concepts involved.

3.1 The Delay Test

The Delay Test uses simulated failures produced by delaying the input clock signals to expose failures. There are several ways of implementing this test. A schematic diagram of one version is illustrated in Figure 3.1. Other variations of the Delay Test are discussed later in this chapter. The arrangement shown in Figure 3.1 is discussed first to illustrate the underlying principles embodied in all versions of the Delay Test.

The Delay Selection Control (DSC) produces the simulated failure conditions by selecting which input clock signals are delayed. The Input Difference Detector (IDD) is used to indicate disagreement among the input signals to the Clock Receiver. The Reference Clock (REF) is used as the standard against which the performance of the system clock (SYSCLK) is evaluated under testing conditions. In this version, The Reference Clock is produced by a second clock receiver (Reference Clock Receiver) using only the undelayed input clock signals. The Output Difference Detector (ODD)



D/A - DISAGREEMENT OR AGREEMENT

Figure 3.1 The Delay Test (Version 1).

detects the deviation of the system clock from its expected behavior as represented by the Reference Clock.

The results expected from the IDD and the ODD in the absence of an actual failure under each of the eight simulated failure conditions are listed in Table 3.1. The IDD and the ODD both indicate agreement if none of the signals is delayed. The IDD shows disagreement when one of the input clock signals is delayed, but the ODD continues to show agreement since a delayed input signal is tolerated. The IDD and the ODD both show disagreement when two signals are delayed. When all three input clock signals are delayed, the ODD indicates disagreement while the IDD shows agreement among the input signals. The method by which these results were determined is discussed in Section 3.2. A failure reveals its presence by causing one or more of the test results produced by the IDD and the ODD to differ from the expected results listed in Table 3.1.

3.1.1 The Delay Selection Control

Figure 3.2 shows one possible design for the DSC. The inverters on the left are used to delay the input clock signals by an amount approximately equal to the sum of their propagation delays. The amount of delay used depends on the tolerance required for the input clock signals and various testing components such as the IDD and the ODD. (See Section 3.2.1.) The circuitry on the right side determines whether an input signal is sent to the receiver delayed or undelayed. In this design, a low selection input would result in that input signal arriving at the Clock Receiver undelayed (except by the two NAND gates). A high selection input would result in the signal being delayed.

3.1.2 The Input Difference Detector

The IDD is designed to detect the presence of a failed input clock signal through its disagreement with the other input signals. When no signal failures exist, it can be used to monitor the DSC. In this case, the IDD can confirm that all three input signals are either undelayed or delayed by indicating agreement. Similarly, the IDD can confirm that either one or two signals are delayed by an indication of disagreement.

A failure to an input clock signal or to the DSC or the IDD is exposed by the failure of the IDD to produce the expected results. For the present, it will be assumed that both the DSC and the IDD are fault-free. There are two ways in which an input clock failure may be revealed.

1. The IDD indicates disagreement when all three input clock signals are either delayed or undelayed.

TABLE 3.1

TEST RESULTS EXPECTED IN A FAULT-FREE ENVIRONMENT

DSC	SIMULATED FAILURE CONDITION	IDD	ODD
DS-0	NO SIGNAL DELAYED	AGREEMENT	AGREEMENT
DS-A	CLKA DELAYED	DISAGREEMENT	AGREEMENT
DS-B	CLKB DELAYED	DISAGREEMENT	AGREEMENT
DS-C	CLKC DELAYED	DISAGREEMENT	AGREEMENT
DS-AB	CLKA & CLKB DELAYED	DISAGREEMENT	DISAGREEMENT
DS-AC	CLKA & CLKC DELAYED	DISAGREEMENT	DISAGREEMENT
DS-BC	CLKB & CLKC DELAYED	DISAGREEMENT	DISAGREEMENT
DS-ABC	ALL SIGNALS DELAYED	AGREEMENT	DISAGREEMENT

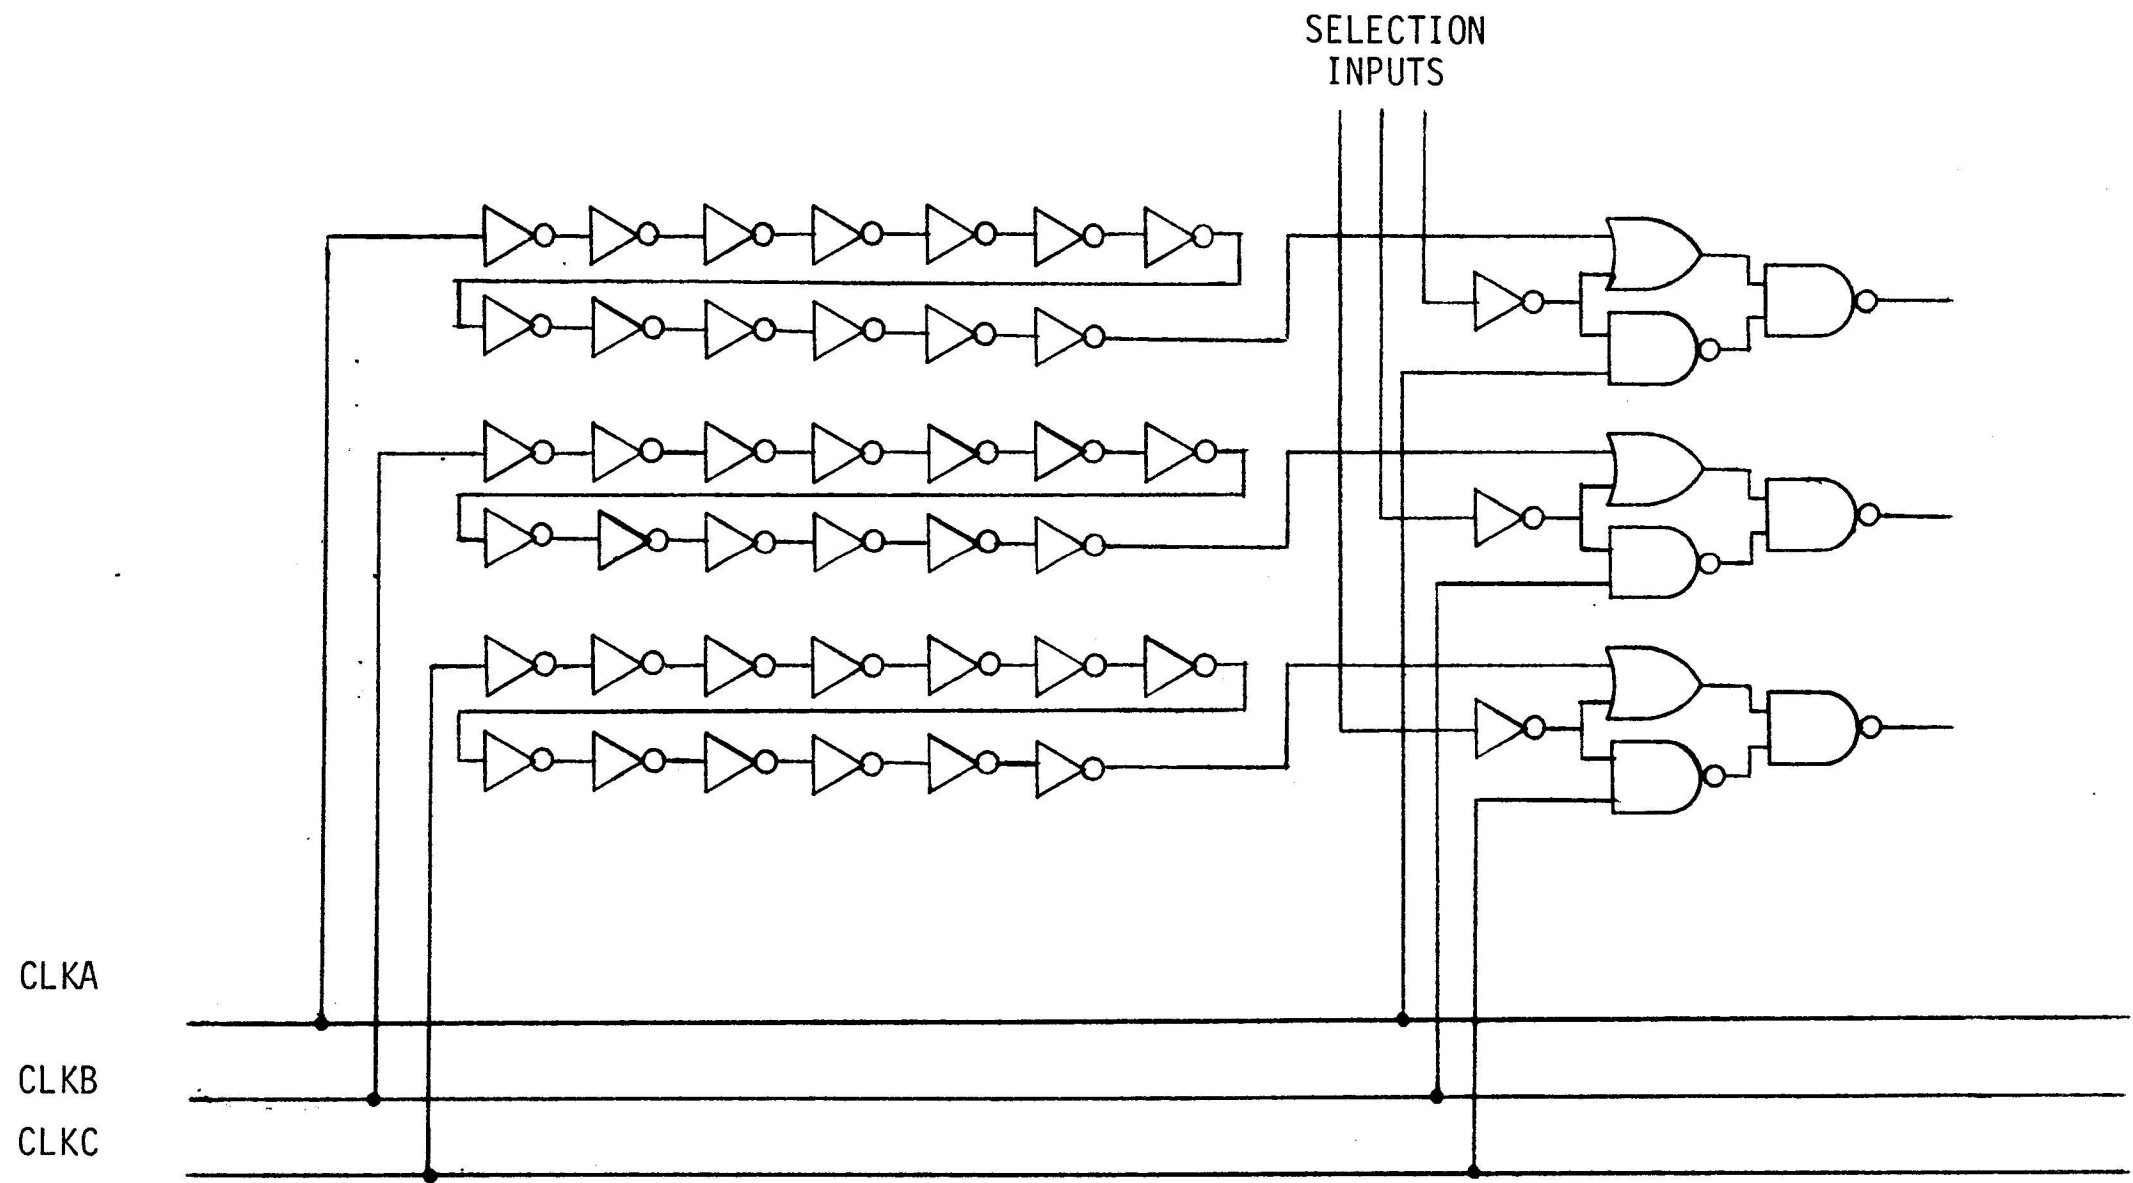


Figure 3.2 The Delay Selection Control.

2. The IDD indicates agreement when one or two of the input clock signals are delayed.

In order for a failed signal to avoid detection, it is necessary that the IDD produce the expected results in spite of the failure. To help illustrate the general concept for the IDD, two relatively simple designs are shown in Figure 3.3. Both designs consist of two sections, the Phase Shift Function (PSF) and the Pulse Width Gauge (PWG). The Phase Shift Function is low only when the logic level of all three input signals are the same, otherwise, it is high. Table 3.2 lists the possible outcomes for the PSF given that CLKA has failed in the described manner. The PSF produces a pulse (high logic level) when there is disagreement among the input signals. Since unfailed input clock signals are only approximately equal, it is necessary to establish a tolerance range within which the signals may be considered to be equal. The Pulse Width Gauge is used to eliminate those disagreement pulses in the PSF whose duration fails to exceed the tolerance limit. Only those pulses which exceed the tolerance are capable of triggering the flip-flop and revealing the presence of disagreement among the input clock signals. The minimum width required is approximately equal to the sum of the propagation delays of the AND gates (or inverters) plus the pulse width required to trigger the flip-flop. The amount of tolerance required is primarily dependent on the quality of synchronization that the primary clocks can achieve. Figures 3.4 and 3.5 illustrate the method used by the PWGs shown in Figure 3.3 to eliminate the PSF pulses that fail to exceed the IDD's tolerance limit.

3.1.3 The Reference Clock Receiver

The Reference Clock Receiver (RCR) is identical in design to the Clock Receiver. The input signals to both receivers are the same if the Clock Receiver's input signals are undelayed. This situation allows the two receivers to check the integrity of each other to a considerable degree. The outputs of the two receivers are in agreement (within a tolerance zone determined by the ODD) if one of the following conditions applies.

1. The two receivers tolerate any and all failures.
2. The outputs of both Clock Receivers have failed but still remain in agreement.
3. The ODD fails to show disagreement due to failure within itself.

TABLE 3.2

THE PHASE SHIFT FUNCTION (ONE SIGNAL FAILURE)

INPUT CLOCK SIGNALS			PHASE SHIFT FUNCTION (PSF)
FCLKA	CLKB	CLKC	
LOW	LOW	LOW	LOW
LOW	LOW	HIGH	HIGH
LOW	HIGH	LOW	HIGH
LOW	HIGH	HIGH	HIGH
HIGH	LOW	LOW	HIGH
HIGH	LOW	HIGH	HIGH
HIGH	HIGH	LOW	HIGH
HIGH	HIGH	HIGH	LOW
UNKNOWN	LOW	LOW	UNKNOWN
UNKNOWN	LOW	HIGH	HIGH
UNKNOWN	HIGH	LOW	HIGH
UNKNOWN	HIGH	HIGH	UNKNOWN (Inverted)





Figure 3.3 The Input Difference Detector.

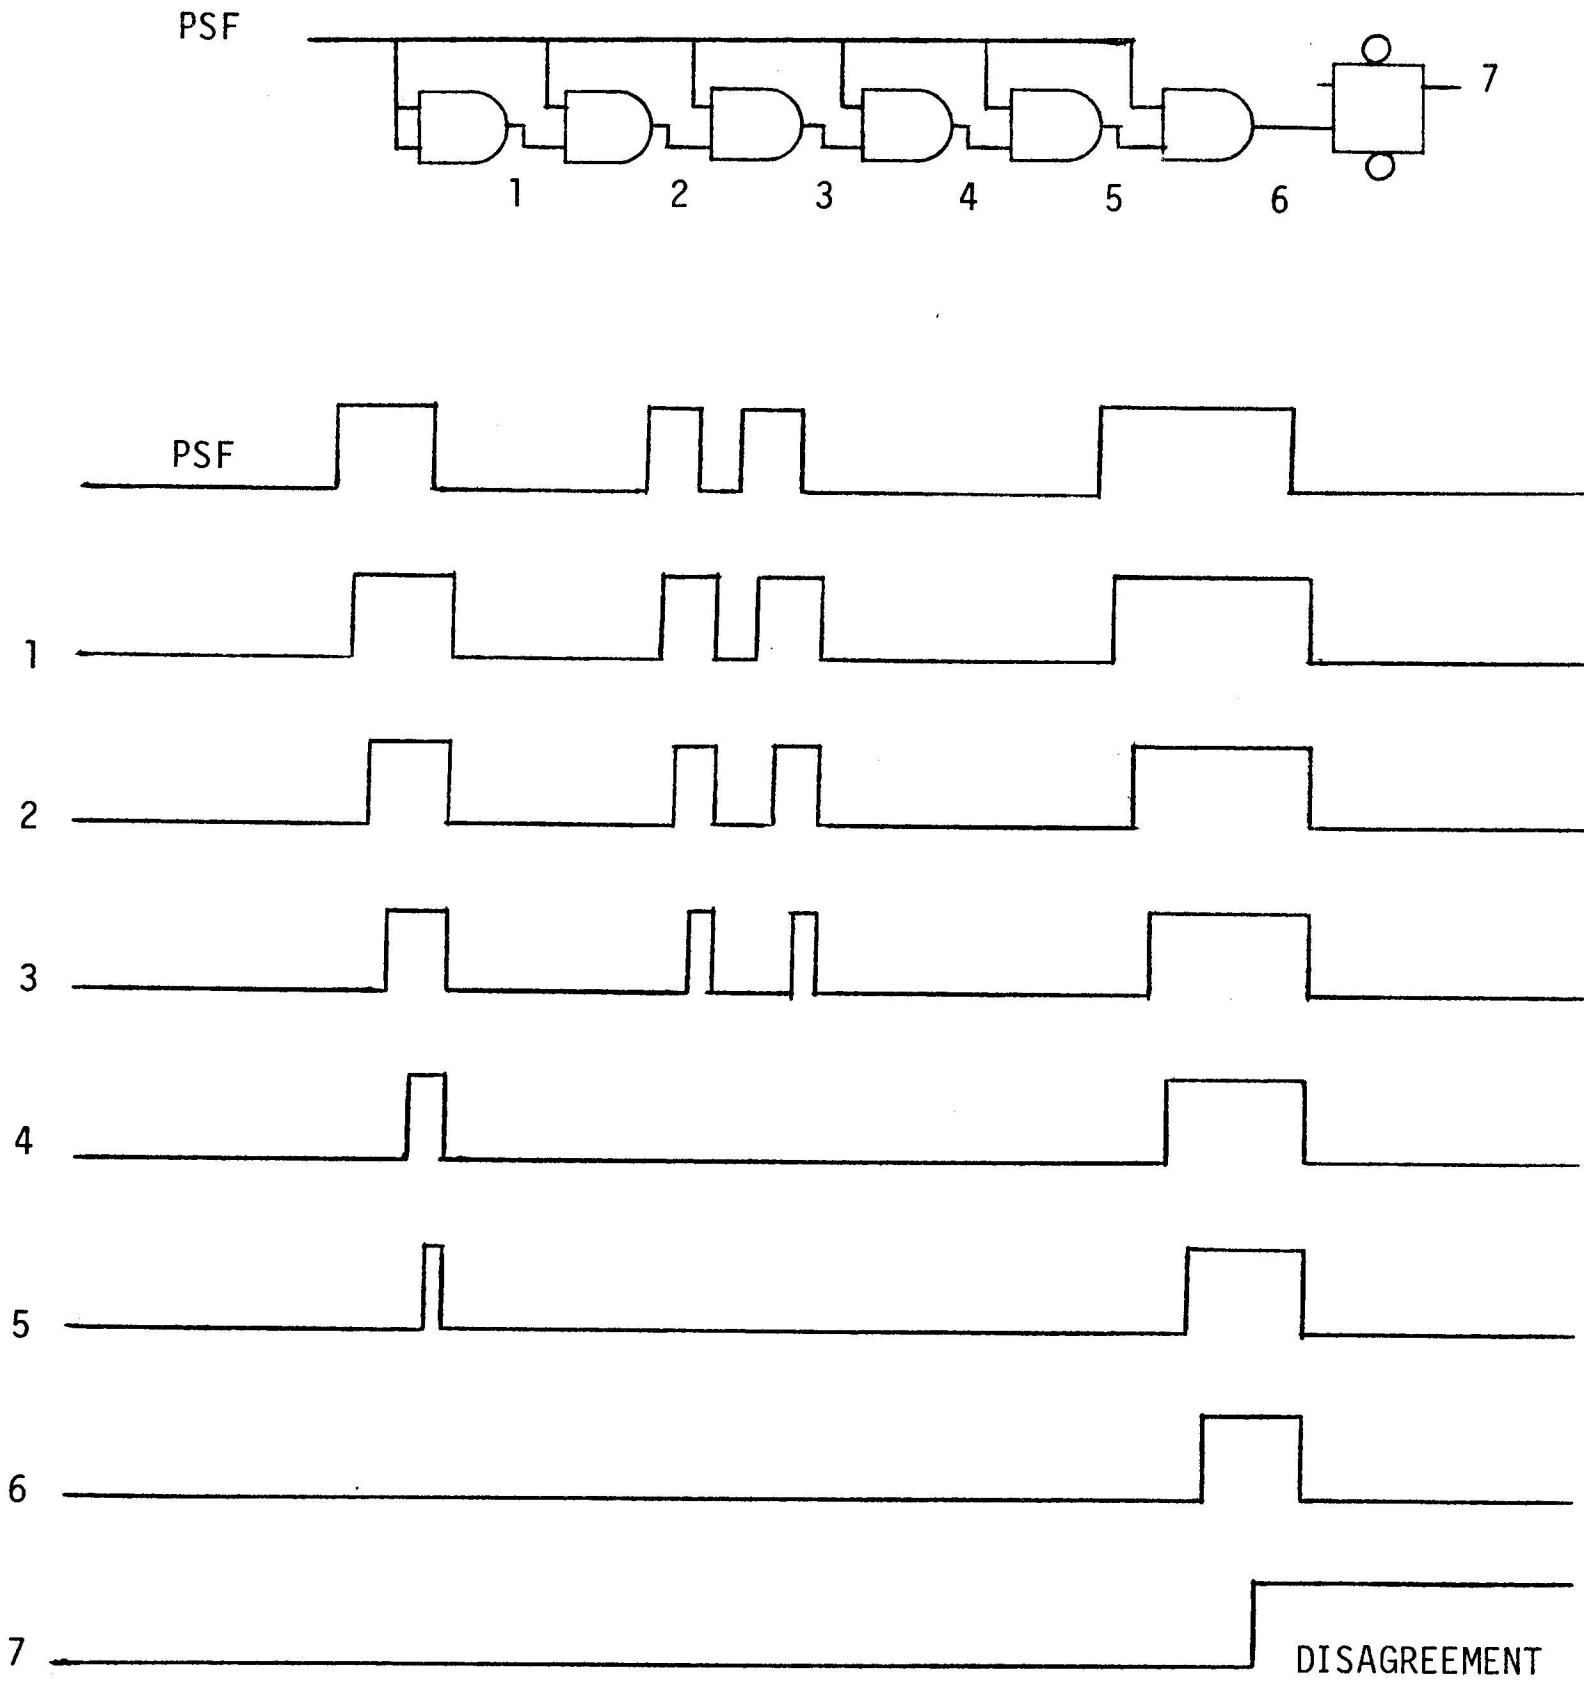


Figure 3.4 Pulse Width Gauge #1.

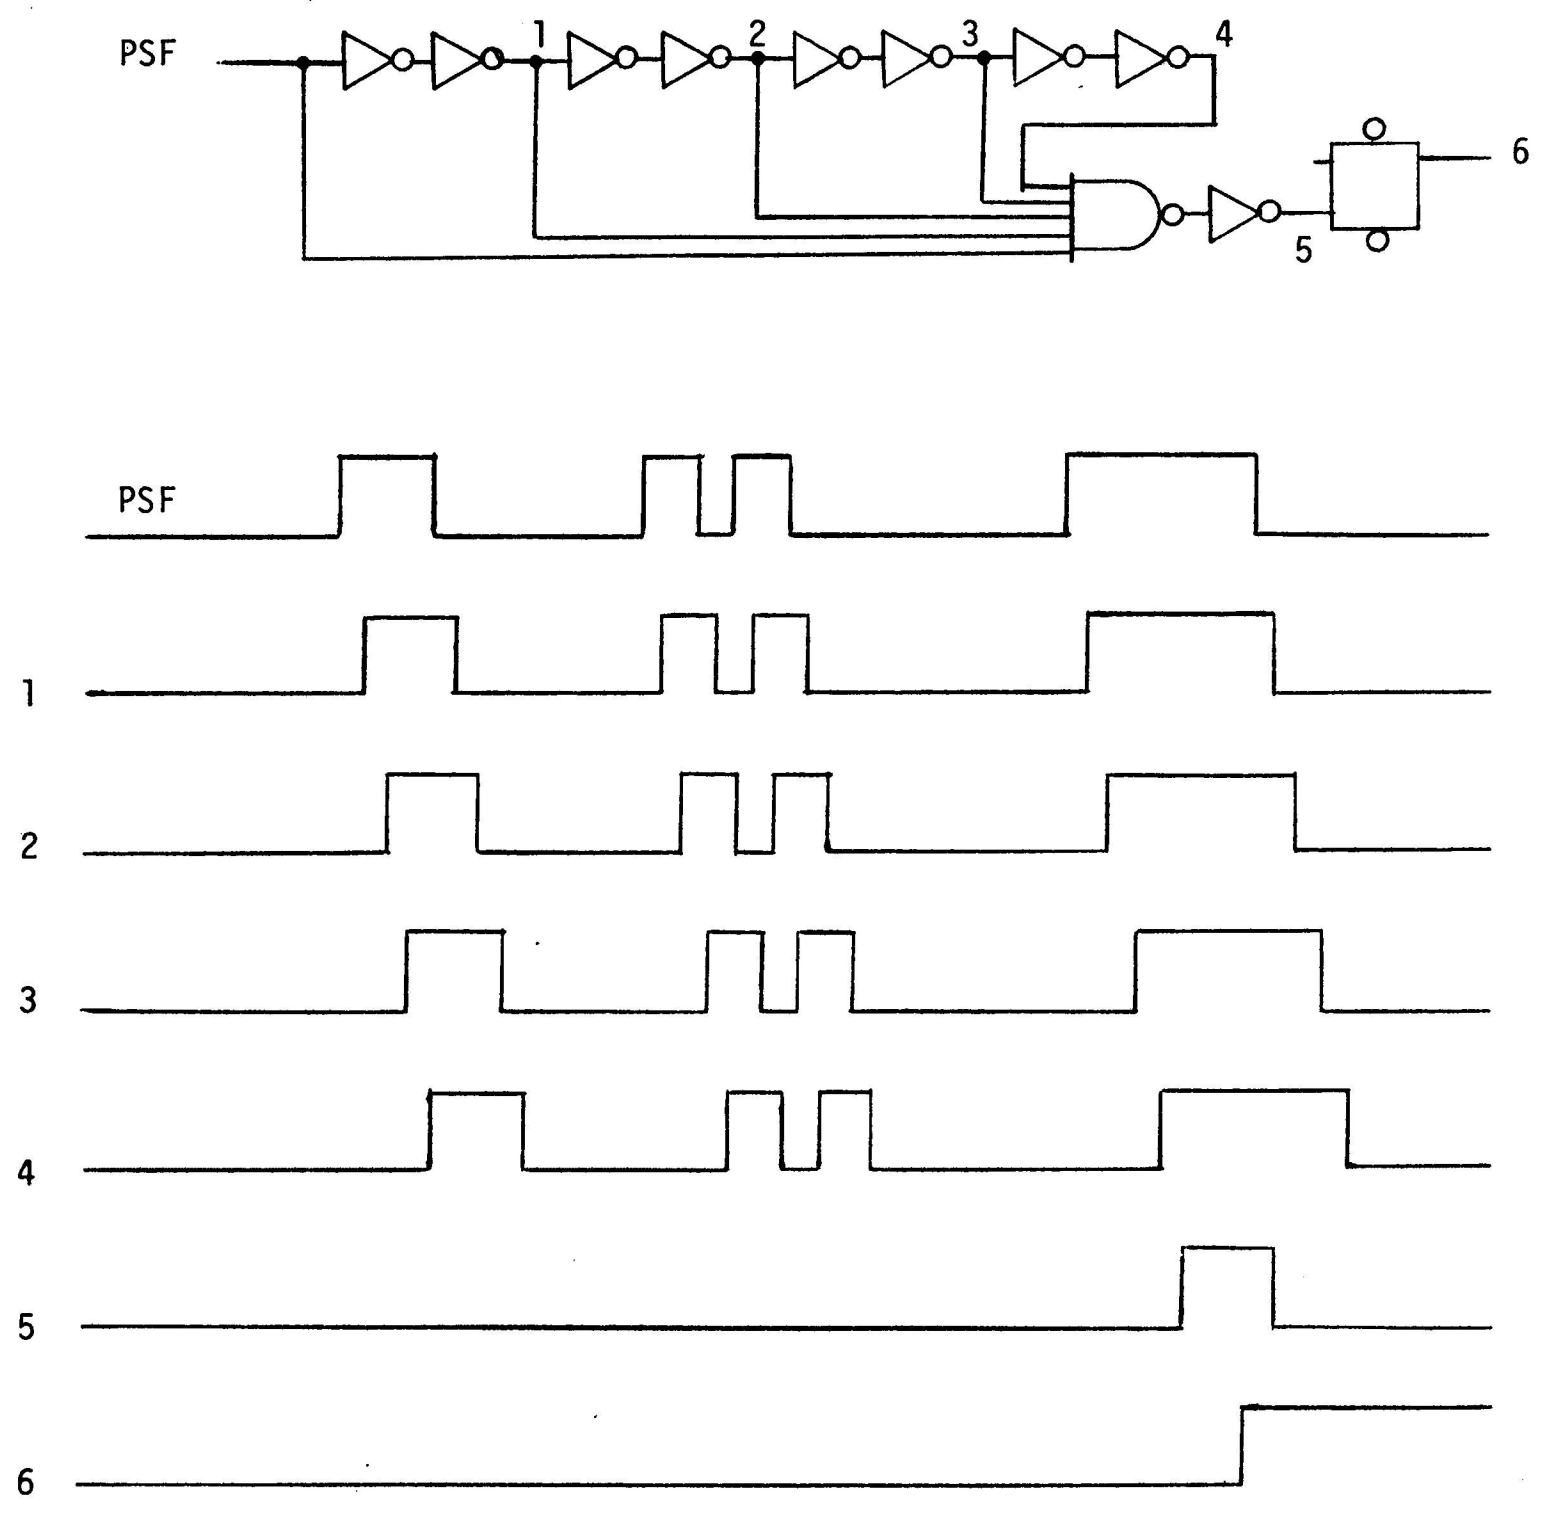


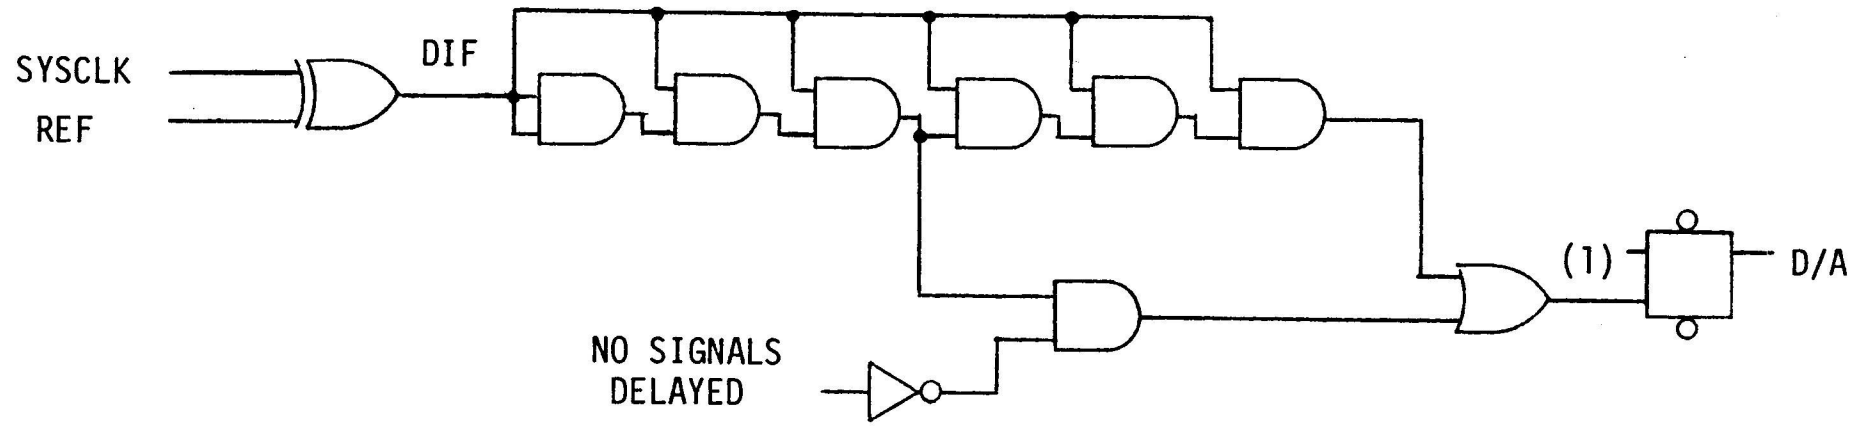
Figure 3.5 Pulse Width Gauge #2.

The third condition would be revealed if disagreement was the expected result. Thus, this possibility can be eliminated if the ODD indicates disagreement when two or three of the input signals are delayed. The second condition requires that both receiver outputs fail, either simultaneously, if no more than one signal is delayed, or during the time that the ODD normally shows disagreement. The time that clock receiver spends being tested is much less than the time it spends working. Thus, the result expected from the ODD most of the time is agreement (DS-0). The possibility of both outputs failing and the ODD not showing disagreement is remote, except in the case of the failures being correlated. If the two receiver outputs fail while disagreement is the expected result, both outputs must still be in agreement when the test is over. The failure of an input clock signal results in a condition which resembles a correlated failure with each receiver. The presence of the failed signal is detected by the IDD. In addition, the failure of a single clock signal is not sufficient to cause either of the receivers' outputs to fail. At least one additional failure is required to occur in order to alter an output. The likelihood of the second condition is extremely small given that no input clock signal has failed. Therefore, if the outputs of the two receivers are in agreement, the first condition (the receivers tolerate any and all failures) is assumed to apply.

The purpose of the Reference Clock (REF) is to represent the fault-free behavior of the system clock. The use of the RCR's output as the Reference Clock is conditioned on the indication of agreement from the IDD under the no signals delayed condition (DS-0). The Reference Clock under this circumstance is reliable due to the fault-tolerant capacity of the Reference Clock Receiver. The system clock (SYSCLK) is compared with the REF under the simulated failure conditions to detect any deviation.

3.1.4 The Output Difference Detector

The ODD is designed to detect disagreement between the system clock and the REF by a method similar to that used in the IDD. An example of a design for an ODD is illustrated in Figure 3.6. The tolerance limit of the ODD is reduced when the simulated failure condition is DS-0 since the input signals to both receivers are essentially the same. The variation in the parameters of the input clock signals (the quality of synchronization) need not be taken into account. Any difference between the system clock and the REF not resulting from a failure is due to the inherent difference between the two receivers. A larger tolerance limit is used when one or more signals is delayed since the variation among the input clock signals must be considered in addition to the inherent difference of the two receivers. The larger tolerance limit of the ODD is slightly greater than the



SYSCLK	REF.	DIF.
LOW	LOW	LOW
LOW	HIGH	HIGH
HIGH	LOW	HIGH
HIGH	HIGH	LOW

Figure 3.6 Output Difference Detector.

tolerance limit of the IDD since the latter need not take into account the difference between the two receivers. The amount that a signal is delayed must be slightly greater than twice the tolerance limit of the IDD for reasons that are discussed in Section 3.2.1.

3.2 Failure Search

A failed input clock signal may be replaced by a stand-by spare, but a failure within the Clock Receiver or the testing circuitry associated with the receiver makes it necessary to remove the entire receiver system (clock receiver plus associated testing components). A failure need only be identified as one of the following:

1. FCLKA - failed input clock signal (CLKA).
2. FCLKB - failed input clock signal (CLKB).
3. FCLKC - failed input clock signal (CLKC).
4. RSF - receiver system failure.

This classification can be achieved by monitoring the results from the IDD and the ODD under the simulated failure conditions. The results expected if no failures are present were listed in Table 3.1. Figures 3.7 through 3.10 help illustrate the manner in which these results were derived. The tolerance limit and the amount that the signals are delayed are exaggerated in order to emphasize the underlying principles that are involved. Three unfailed and undelayed input clock signals are shown in Figure 3.7. The PSF (in the IDD) produces a high level pulse when there is disagreement among these signals. If one of the PSF's disagreement pulses exceeds the IDD's tolerance limit, T-1, the IDD indicates disagreement. In this particular case (Figure 3.7), the PSF fails to produce such a pulse, thus agreement is the result given by the IDD. The system clock changes its state only after two input signals have changed theirs (Chapter 2). The Reference Clock also changes state after the second input signal has changed. The REF signal remains the same for all four diagrams since its input signals are not changed by the DSC. The difference between the system clock and the REF in the absence of failures is due to the differences between the two receivers. This difference is too small to result in the DIF producing any pulses that exceed the ODD's lower tolerance limit, T-2, provided that this limit was selected correctly. The lower tolerance limit is used only when no input signals are delayed. Thus, in the absence of any failure, both the IDD and the ODD indicate agreement under the DS-0 condition.

In Figure 3.8, CLKA has been delayed by the DSC (DS-A). The delaying of signal CLKA(DCLKA), causes pulses to appear in the PSF whose widths exceed T-1 resulting in an indication of disagreement from the IDD. The system clock changes its state after the second undelayed signal has changed. DCLKA is a simulated




Figure 3.7 Waveforms - No Input Signal Delayed.




Figure 3.8 Waveforms - One Input Signal Delayed.




Figure 3.9 Waveforms - Two Input Signals Delayed.




Figure 3.10 Waveforms - All Input Signals Delayed.

failure which is tolerated by the Clock Receiver. The DIF produces no pulses whose widths exceed the ODD's higher tolerance limit, T-3. The ODD continues to indicate agreement while the IDD shows disagreement when an input clock signal is delayed.


In Figure 3.9, two input clock signals are delayed (DS-AB). Pulses in the PSF exceed T-1 resulting in the IDD indicating disagreement. The system clock does not change its state until after one of the delayed signals has changed. The difference between the system clock and the REF is now sufficient for the ODD to detect the disagreement.

Finally, in Figure 3.10, all three input clock signals are delayed (DS-ABC). The pulses in the PSF no longer exceed T-1 so the IDD indicates agreement. The system clock changes its state after two of the delayed signals have changed. Thus, the SYSCLK lags behind the REF sufficient to result in an indication of disagreement by the ODD.


3.2.1 Input Clock Signal Failure

The IDD is capable of detecting a failed input clock signal, but not isolating the source. The failed signal is identified through the effect it has on the flip-flops that it drives. The location of these pseudo-failed flip-flops can be determined by analyzing the results obtained from the ODD under the simulated failure conditions. The input signal used to drive these flip-flops becomes the prime suspect.


The deviation of the results of the IDD from those expected (Table 3.1) is due to a failure or failures within the IDD, the DSC or the input clocks. The question of a failure to the IDD or the DSC is discussed later (Receiver System Failure). A failed input signal is detected usually through its disagreement with the other signals. In Figure 3.11, the PSF is shown for CLKA being unfailed, stuck-at-logic 0 and out-of-phase (by an amount that exceeds the IDD's tolerance limit). Both failure modes produce disagreement pulses in the PSF that exceed the tolerance limit. The IDD's capacity for detecting failed input clock signals can be improved by examining the IDD results under all the simulated failure conditions. An example of a failure mode that may not be detected by the IDD under the DS-0 condition is shown in Figure 3.12. To detect such a failed signal, it is necessary to obtain information while the failed signal is delayed (DS-A) and while the two unfailed signals are delayed (DS-BC). All three PSF's shown in Figure 3.12 contain no pulses with sufficient width for the IDD to indicate disagreement. However, the result expected from the IDD under these conditions is an indication of disagreement. Thus, the presence of the failed signal is revealed by agreement being indicated when disagreement was expected.



a) Unfailed.




b) Stuck-At-0.




c) Out-Of-Phase.


Figure 3.11 The PSF With a Failed Signal.



a) No Signals Delayed.



b) The Failed Signal Delayed.



c) The Two Unfailed Signals Delayed.

Figure 3.12 Failed Signal Detected by Agreement.

The occurrence of more than one input signal failure presents no problem to the IDD's detection ability with the exception of three correlated signal failures. This is unlikely to happen if the input signals are produced separately. The occurrence of two signal failures may be detected not only by their disagreement with the unfailed signal but also with each other. Three uncorrelated signal failures would probably be detected by their disagreement among themselves. However, since they are uncorrelated, it is not likely that all three will fail before the IDD detects the presence of a failure.

The probability of the IDD detecting a failed signal depends not only on the quality of the IDD's components but also upon the amount of time tolerance required for the input clock signals. The tolerance limit of the IDD is a function of the synchronization that can be attained among the input clocks. The closer the parameters of the input clocks are made the more the tolerance limit can be reduced. This reduction of the tolerance limit means that the definition of an unfailed signal is more precise. The IDD's detection capability under the DS-0 condition is improved with the small time tolerance.

The isolation of a failed input clock signal is accomplished through an examination of the ODD results under the simulated failure conditions. From Table 3.1., agreement is the expected result when none or just one of the input signals is delayed. Disagreement is the result expected when two or three of the signals are delayed. Delaying an input clock signal causes the flip-flops served by that signal to pseudo-fail. The triggering edges of the delayed signal arrive at these flip-flops after the clear input has become low, thus preventing the flip-flops' outputs from being raised. The outputs of the flip-flops are used as the inputs to the two (inverted) majority voters of the Clock Receiver. Labeling the inputs to the SET voter (the outputs of the leading edge flip-flop triad) as I_A , I_B and I_C with CLKA, CLKB and CLKC being the driving signals respectively, the (inverted) output of the SET voter is expressed in Eq. 3.1. The output expression when a single input signal is delayed is expressed in Equations 3.2, 3.3 and 3.4 with CLKA, CLKB and CLKC being delayed respectively. By delaying each input signal separately, each term in Eq. 3.1 is tested.

$$\overline{\text{SET}} = I_A I_B + I_B I_C + I_C I_A \quad (3.1)$$

$$\overline{\text{SET}} = 0 I_B + I_B I_C + I_C 0 = I_B I_C \quad (3.2)$$

$$\overline{\text{SET}} = I_A 0 + 0 I_C + I_C I_A = I_C I_A \quad (3.3)$$

$$\overline{\text{SET}} = I_A I_B + I_B 0 + 0 I_A = I_A I_B \quad (3.4)$$

A failed input clock signal can influence the output of the voter only when one of the unfailed signals has been delayed. A failed input signal produces either a TYPEONE voter input or a TYPEZERO voter input. The voter input is classified as TYPEONE if the flip-flop triggered by the failed clock signal raises its output before those driven by the two unfailed input clock signals. It is classified as TYPEZERO if the triggering edges of the failed clock signal arrive after those of the two unfailed input clock signals. Since two unfailed signals are sufficient to drive the Clock Receiver (Chapter 2), the clear input is low when the triggering edges of the failed signal arrive. If the failed signal has no triggering edges at all (s-a-0 or s-a-1), the voter input is TYPEZERO also. The voter input of an unfailed signal is classified as UNFAILED since its edges are produced properly. Several input clock signals along with the SET voter inputs they produce are shown in Figure 3.13. In all five cases, at least two clock signals are unfailed, say CLKB and CLKC. The Clear Input is used as a reference since the Clock Receiver can tolerate CLKA failing. CLKA is unfailed in Figure 3.13a, thus the voter input is UNFAILED. Figure 3.13b shows the voter input produced when CLKA is delayed. Figure 3.13c and 3.13e show two failed signals which produce a TYPEONE voter input while Figure 3.13d shows one which results in a TYPEZERO voter input. The other two input signals, CLKB and CLKC, are unfailed, so the voter inputs associated with them resemble that of Figure 3.13a. As far as the SET voter is concerned, there is no difference between an UNFAILED voter input and a TYPEONE input as long as the other two voter inputs are UNFAILED. This results from the fact that the SET voter raises the system clock after the second input becomes high.

Figure 3.14 compares these three types of voter inputs. The tolerance limit of the IDD is the maximum amount of disagreement that can occur among the three input signals without being detected by the IDD. The tolerance zone within which the arrival of a leading edge results in an UNFAILED voter input is shown in relation to the clear input and the two unfailed signals. The maximum width of the tolerance zone is twice the tolerance time limit. One edge can arrive either before or after the other two by an amount less than the tolerance limit without causing the IDD to show disagreement. Also shown are the regions in which the arrival of a leading edge would produce a TYPEONE or a TYPEZERO voter input for the SET voter. A similar discussion can be used for the inputs of the RESET voter.

The possible types of voter input that can result from delaying an input clock signal are listed in Table 3.3. As previously stated, the amount that the signals are delayed (delay time) is slightly greater than twice the IDD's tolerance limit or the maximum tolerance zone. This insures that when an unfailed signal is delayed, the voter input (SET and RESET) is TYPEZERO. A triggering edge which arrives in the tolerance zone will arrive in the TYPEZERO region when the signal is delayed. It is

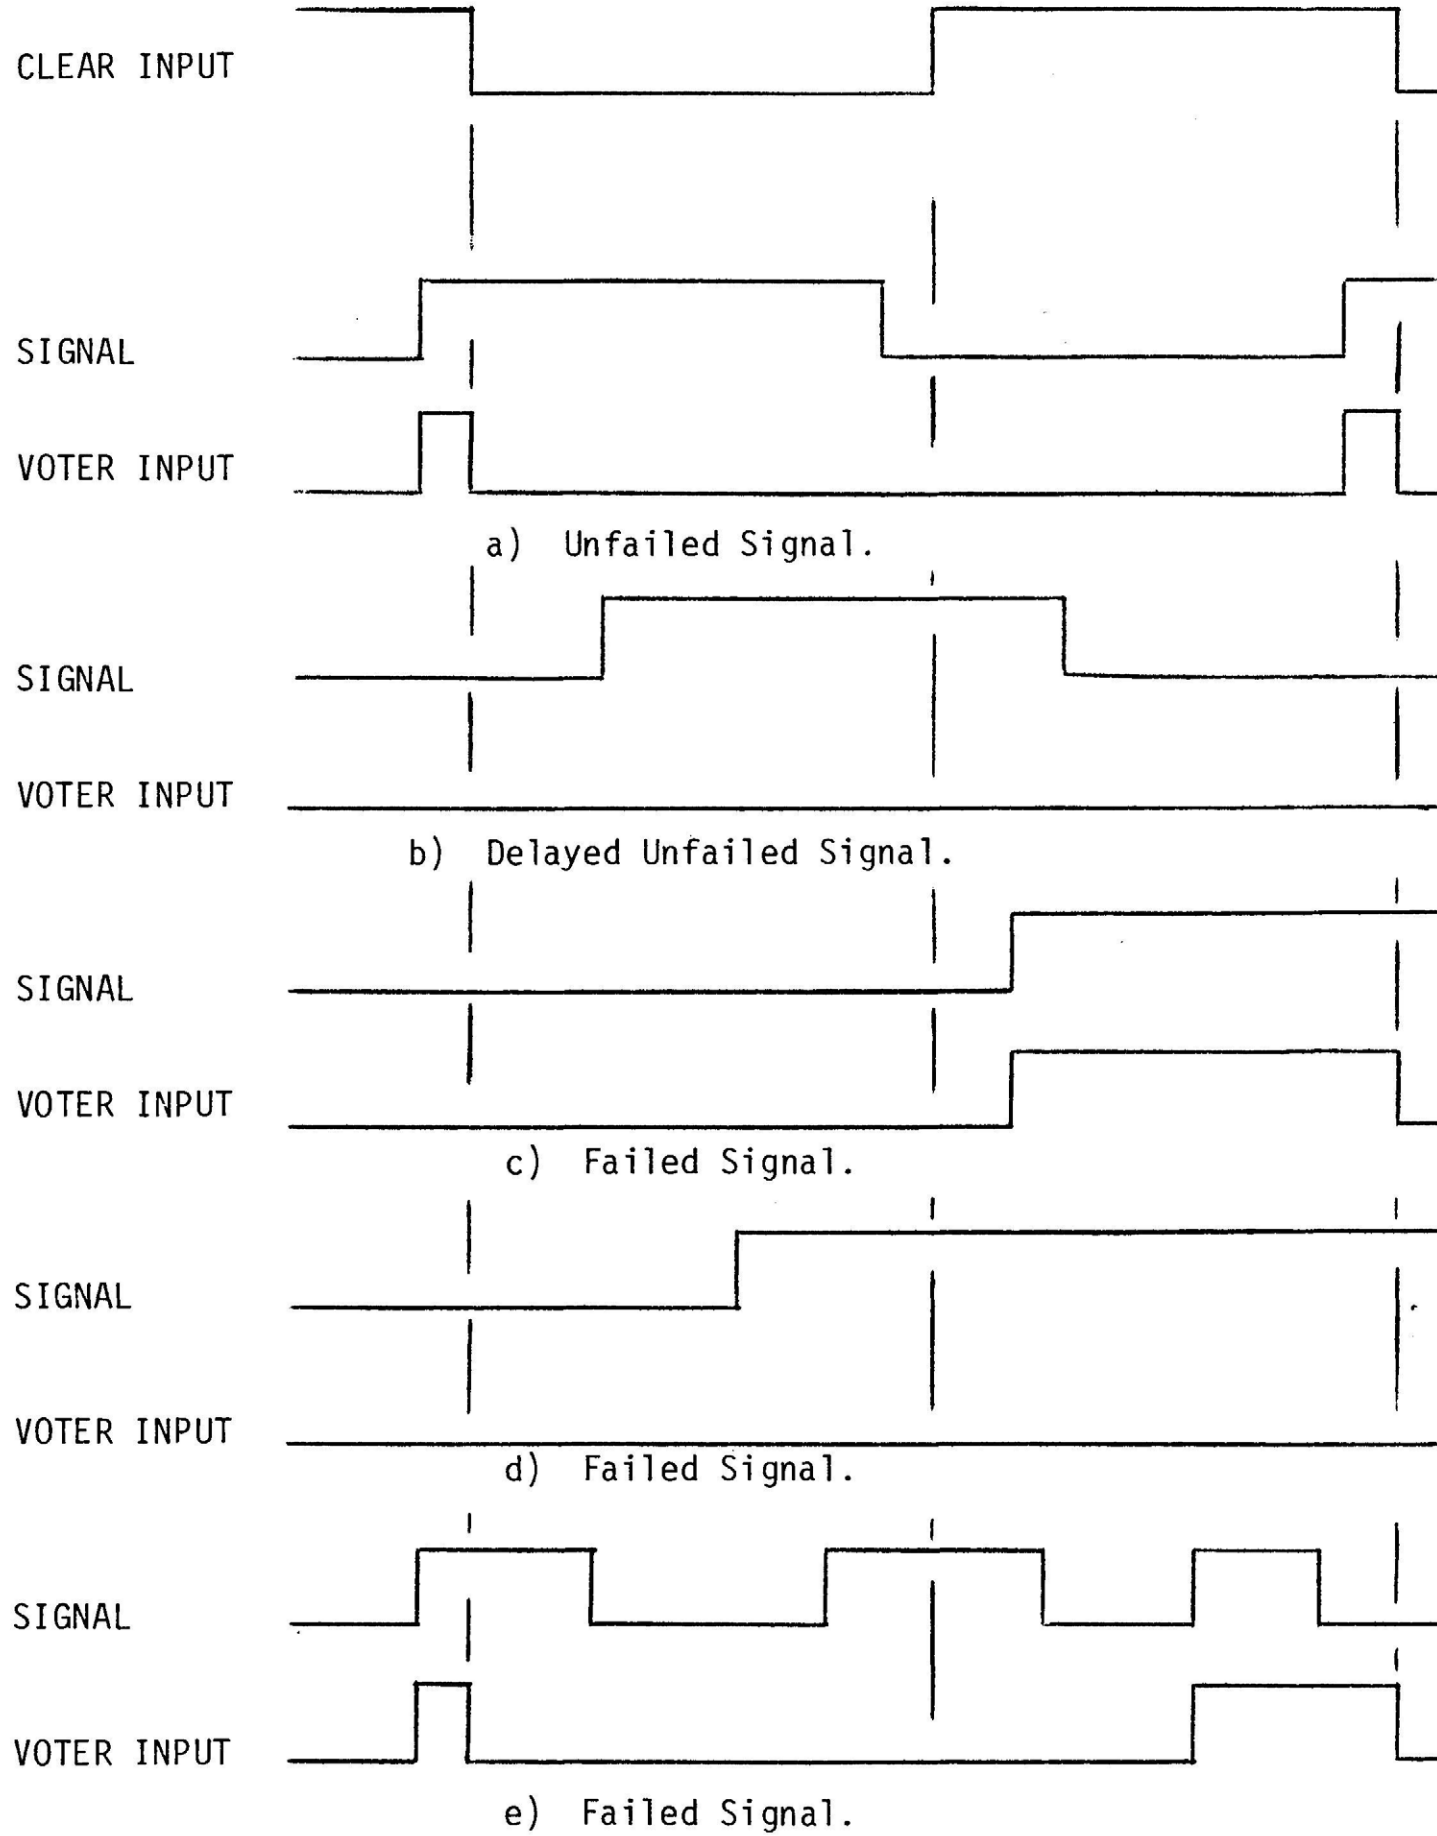


Figure 3.13 Examples of Set Voter Inputs (Only one Input Shown).

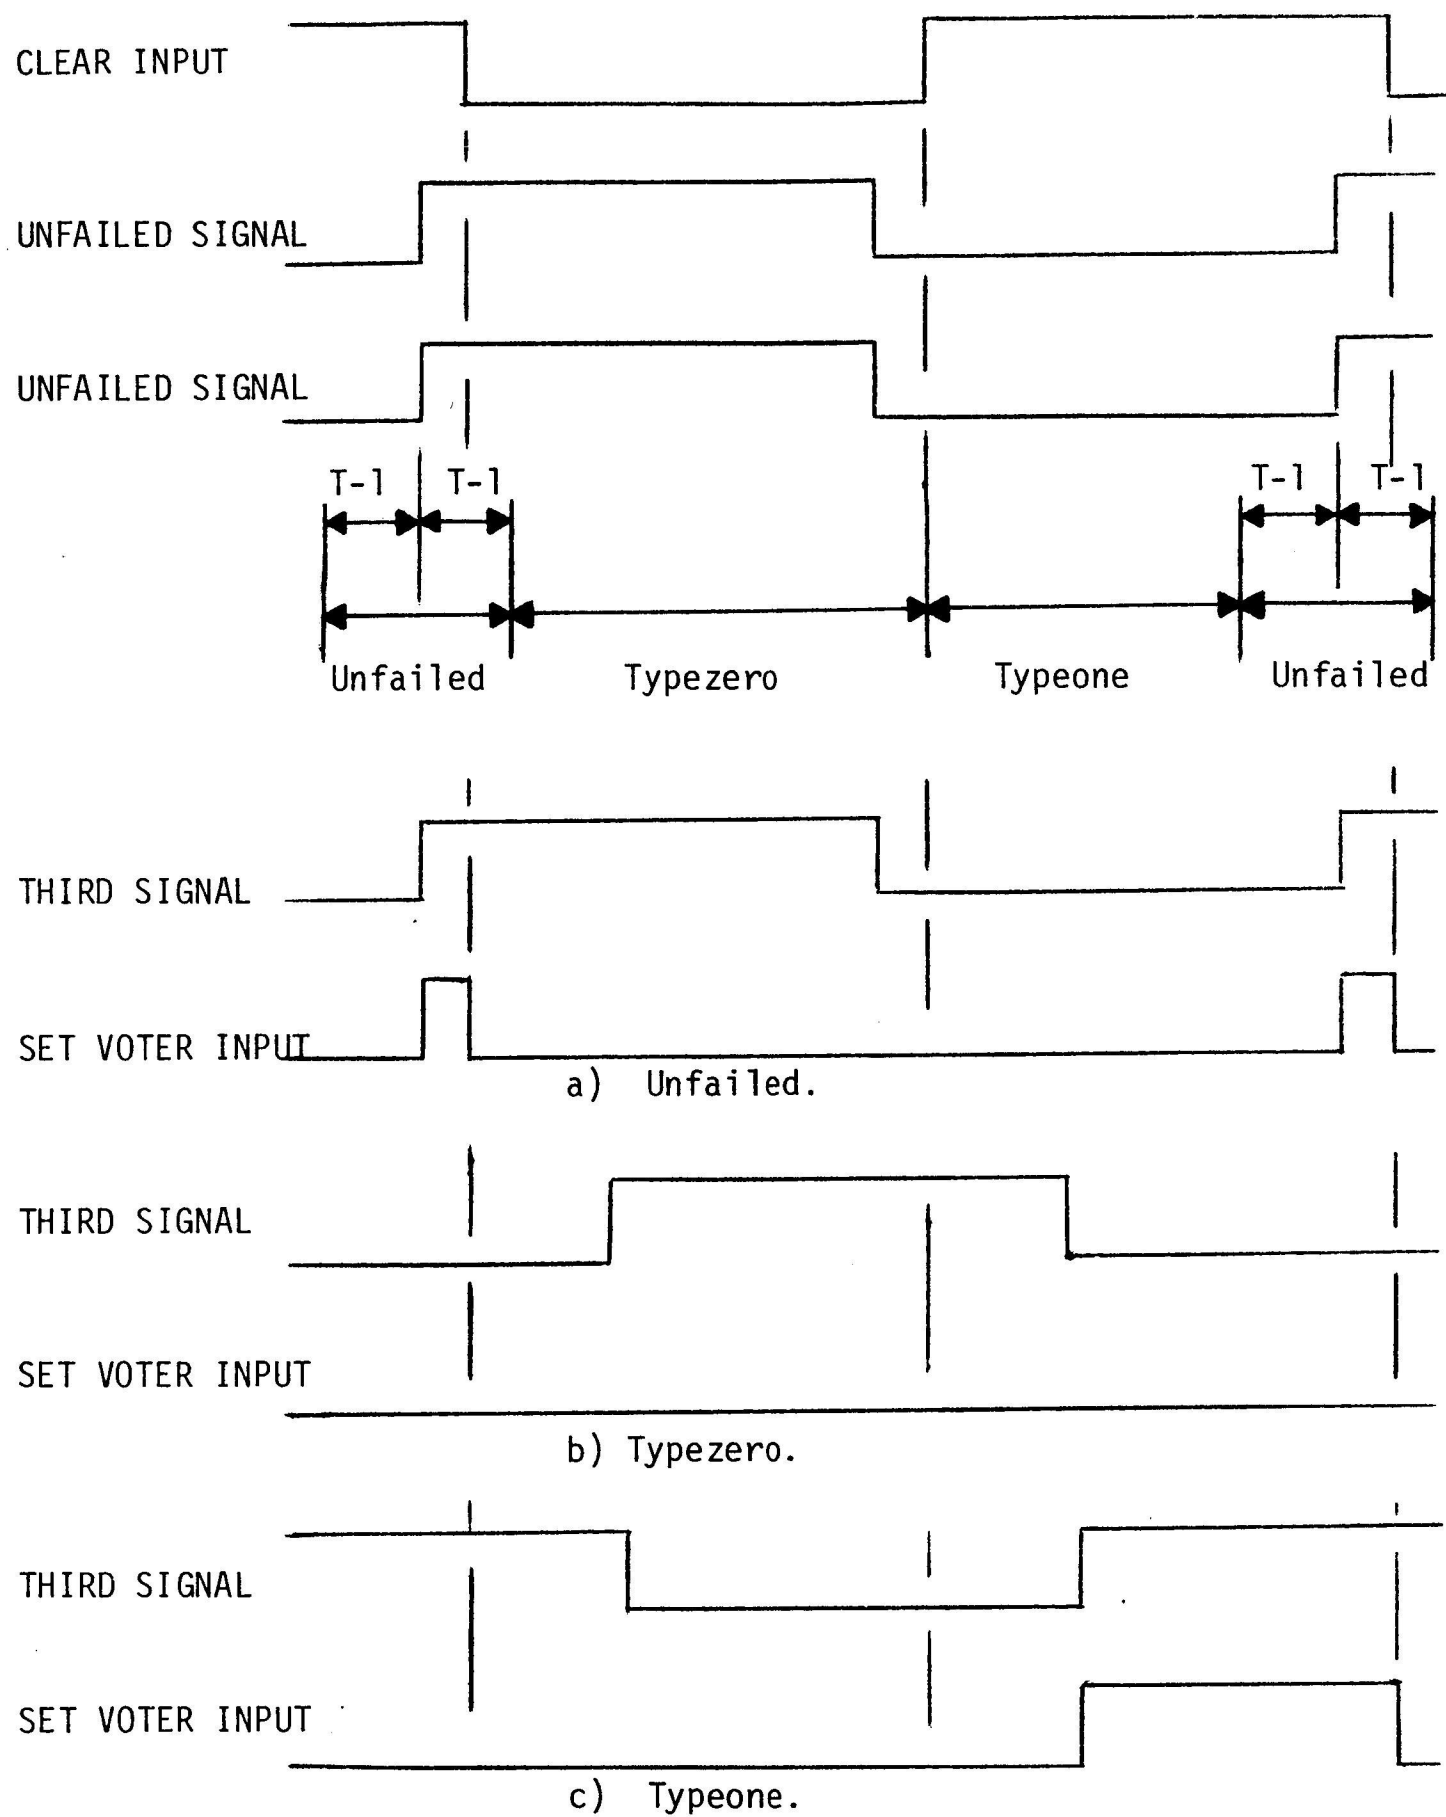


Figure 3.14 Set Voter Inputs.

TABLE 3.3

VOTER INPUT PRODUCED BY DELAYING A SIGNAL

CASE	VOTER INPUT UNDELAYED SIGNAL	VOTER INPUT DELAYED SIGNAL
A	UNFAILED	TYPEZERO
B	TYPEONE	TYPEONE
C	TYPEONE	UNFAILED
D	TYPEONE	TYPEZERO
E	TYPEZERO	TYPEZERO
F	TYPEZERO	TYPEONE

desirable that a triggering edge which arrives in the TYPEONE region will arrive in either the UNFAILED or the TYPEONE region. However, since the delay time must be greater than twice the IDD's time tolerance limit to insure that an unfailed edge arrives in the TYPEZERO region when delayed, the possibility exists that a TYPEONE voter input can become TYPEZERO when the signal is delayed (Case D). A signal which produces a TYPEZERO voter input will produce either a TYPEZERO or a TYPEONE when delayed.

The output of a flip-flop can be defined in terms of these three types of voter input. Table 3.4 lists the ODD results for various combinations of voter inputs with redundant combinations omitted. Table 3.4 includes all the possible combinations given that no more than one input signal failure has occurred. The ODD indicates disagreement when two of the inputs to a voter are both either TYPEONE or TYPEZERO. Since the simulated failures result in a TYPEZERO voter input, no two voter inputs are TYPEONE. The voters in the Clock Receiver do not distinguish between a TYPEONE and an UNFAILED voter input as long as there is not more than one failed input clock signal.

Table 3.5 lists the ODD results obtained under the simulated failure conditions for each case of Table 3.3. The IDD indicates the presence of a failed input signal in all the cases except Case A. Case A represents the voter inputs produced with three unfailed signals. Case A gives the expected ODD results as listed in Table 3.1. In Table 3.5, CLKA is the failed input signal while CLKB and CLKC are both unfailed. Case D gives the same results as did Case A, but the IDD indicates the presence of a signal failure. The triggering edges of this signal (undelayed) arrive before those of the unfailed signals by an amount which exceeds the IDD's time tolerance but is less than the delay time minus the tolerance limit. If the region is made small enough, a failed signal is not likely to confine its edges to this area for several test cycles. Thus, the delay time used for the input signals should be made as small as possible. The lower limit of the delay time is set by the requirement that an UNFAILED voter input becomes TYPEZERO when delayed.

The other failed input signals (Cases B,C,E and F) cause the results from the ODD to differ from those of Table 3.1. The simulated failure conditions under which these deviations occur are dependent upon which signal has failed and the manner in which it fails. For FCLKA, these deviations can occur either when an unfailed signal is delayed (DS-B and DS-C) or when an unfailed signal and the failed signal are delayed (DS-AB and DS-AC). From Table 3.5, the ODD result deviates from the expected result under the DS-B and DS-C condition if FCLKA produces a TYPEZERO voter input (Case E and Case F). Two voter inputs are TYPEZERO and the ODD indicates disagreement. If when FCLKA is delayed, it produces either a TYPEONE or UNFAILED

TABLE 3.4

RESULTS FROM THE ODD FOR A GIVEN SET OF VOTER INPUTS

VOTER INPUT COMBINATIONS			ODD RESULT
I_A	I_B	I_C	
UNFAILED	UNFAILED	UNFAILED	AGREEMENT
TYPEZERO	UNFAILED	UNFAILED	AGREEMENT
TYPEONE	UNFAILED	UNFAILED	AGREEMENT
TYPEZERO	UNFAILED	TYPEZERO	DISAGREEMENT
TYPEONE	UNFAILED	TYPEZERO	AGREEMENT
UNFAILED	TYPEZERO	TYPEZERO	DISAGREEMENT
TYPEZERO	TYPEZERO	TYPEZERO	DISAGREEMENT
TYPEONE	TYPEZERO	TYPEZERO	DISAGREEMENT

TABLE 3.5

COMPARISON OF THE RESULTS FROM THE ODD
FOR EACH CASE OF TABLE 3.3

TEST CONDITION	ODD RESULTS					
	A	B	C	D	E	F
DS-0	a	a	a	a	a	a
DS-A	a	a	a	a	a	a
DS-B	a	a	a	a	d	d
DS-C	a	a	a	a	d	d
DS-AB	d	a	a	d	d	a
DS-AC	d	a	a	d	d	a
DS-BC	d	d	d	d	d	d
DS-ABC	d	d	d	d	d	d

a - agreement
d - disagreement

voter (Case B,C and F), the IDD will indicate agreement under the DS-AB and DS-AC condition.


To summarize, the presence of a failed input clock signal is detected through the IDD. The voter inputs (one to each voter) produced by the failed signal do not behave in their normal manner. The triggering edges of the failed clock are being produced incorrectly. If these edges trigger their flip-flops before the edges of the unfailed signal trigger theirs, the voter input (the output of the pseudo-failed flip-flop) is TYPEONE. If the triggering edges of the failed signal are produced after those of the unfailed signals or not at all, the voter input is TYPEZERO. By delaying one of the unfailed signals or the failed signal with an unfailed signal, the failed signal is revealed.

3.2.2 Receiver System Failure

The Receiver System consists of the Clock Receiver and the test components associated with it. These test components are responsible for checking the validity of the Clock Receiver and its input clock signals. A failure within one of these components may impair the performance of this task. Thus, the validity of the test must also be confirmed.

In Chapter 2, the Clock Receiver was discussed in terms of two operations, edge-detecting and voting. A failure within the edge-detecting section affects the output of only one flip-flop. That is, only one voter input to one majority voter (SET or RESET) is incorrect. By contrast, a failed input signal could result in a voter input to each majority voter being incorrect. The most common failure modes for a flip-flop (TTL) are stuck-at-logic 0 (TYPEZERO voter input), stuck-at-logic 1 (TYPEONE voter input) and a racing condition (high frequency oscillation). A voter input, say I_A for the SET voter, can be detected if it is s-a-0, but not if it is s-a-1. Its ability to be detected in a racing situation depends on the width of the pulses it produces.

The SET voter raises the system clock high when two of its inputs become high while the RESET voter lowers it when two of its inputs are high. If I_A (SET) is TYPEZERO (or if I_A becomes high after the UNFAILED voter inputs), the failed voter input is exposed by delaying either CLKB or CLKC (DS-B or DS-C). This results in the leading edges of the system clock lagging behind those of the reference clock by an amount sufficient for the ODD to show disagreement. However, if I_A is TYPEONE, the failure can not be exposed since the ODD under the simulated failure conditions produces the expected results. This difficulty may be resolved either by redesigning the ODD as shown in Figure 3.15 or as outlined in Figure 3.16. By using the ODD shown in Figure 3.15, a TYPEONE voter input is exposed when CLKA and either CLKB or



PWG - PULSE WIDTH GAUGE

SYSCLK	REF.	PSF	1	2
LOW	LOW	LOW	LOW	LOW
LOW	HIGH	HIGH	HIGH	LOW
HIGH	LOW	HIGH	LOW	HIGH
HIGH	HIGH	LOW	LOW	LOW

Figure 3.15 An Expanded ODD.

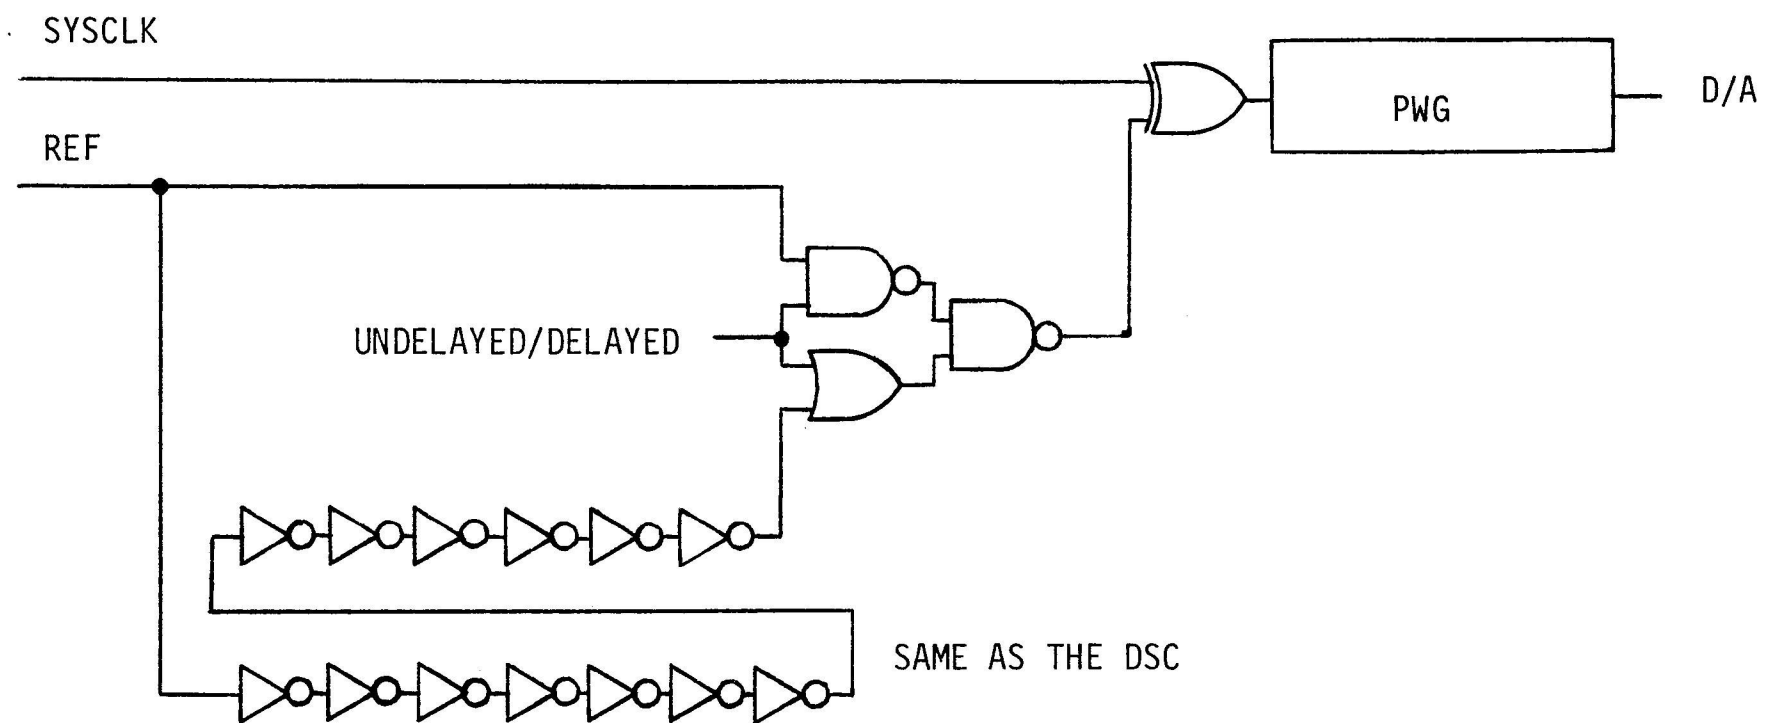


Figure 3.16 The Dual Mode ODD.

CLKC are delayed (DS-AB or DS-AC). The leading edges of the SYSCLK and the REF are produced at approximately the same time. The upper part of the ODD will show agreement which is not the expected result under these failure conditions. The trailing edges of the system clock lag behind those of the reference clock sufficient for the lower section to indicate disagreement, the expected result. The original ODD, unable to show separate results, gives only an indication of disagreement. The separate results produced by the ODD in Figure 3.15 allows the detection of TYPEONE voter inputs resulting from a receiver failure. The ODD shown in Figure 3.16 allows the REF signal to be used either delayed or undelayed. By delaying the REF signal, we are essentially advancing the SYSCLK relative to the ODD. This in turn is the same as advancing the input clock signals. Thus, the simulated failures not only produce a TYPEZERO voter input but also a TYPEONE. The presence of two TYPEONE voter inputs to a voter results in the ODD showing disagreement. If only one of these TYPEONE inputs results from a simulated failure, the expected result was agreement and so the failure is revealed. This ODD also makes it possible to isolate the failed signal in Case D. Table 3.6 lists the results expected under the expanded simulated failure conditions.

A failure in the voting section of the Clock Receiver is exposed in one of two ways. First, it may cause the SYSCLK to be altered enough for the ODD to detect the disagreement. With the exception of an AND gate stuck-at-zero, all failures directly affect the system clock. For this reason, the ODD uses a smaller tolerance limit when no input clock signals are delayed (DS-0). This allows the ODD to monitor the behavior of the system clock closer than would be possible with a single tolerance limit. The second method is by analyzing the results from the ODD under simulated failure conditions. A s-a-0 AND gate is detected when the input clock signal not associated with the gate is delayed. The voter containing the failed gate lags behind its counterpart in the RCR resulting in the ODD showing disagreement.

A failure within the RCR that is tolerated remains masked since its input signals are not altered by the simulated failure conditions. An intolerated failure within the RCR is exposed provided that the Reference Clock is sufficiently altered for the IDD to show disagreement under the DS-0 condition. A failure within the ODD can be exposed by the failure of the ODD to produce the expected results. Any deviation of the ODD's results in the absence of an input clock signal failure is due to a Receiver System Failure (RSF).

In order for the IDD to produce the expected results, it is not only necessary that no failed input signals are present, but also that no unmasked failures are present in the IDD or the DSC. When the IDD results differ from the results expected, it may not be possible to distinguish between a failed input clock signal

TABLE 3.6

SIMULATED FAILURE CONDITIONS

SIMULATED CONDITION	REF STATE	DSC	IDD RESULT	ODD RESULT
DS-0	UNDELAYED	DS-0	A	A
DS-A	UNDELAYED	DS-A	D	A
DS-B	UNDELAYED	DS-B	D	A
DS-C	UNDELAYED	DS-C	D	A
DS-AB	UNDELAYED	DS-AB	D	D
DS-AC	UNDELAYED	DS-AC	D	D
DS-BC	UNDELAYED	DS-BC	D	D
DS-ABC	UNDELAYED	DS-ABC	A	D
AS-0	DELAYED	DS-ABC	A	A
AS-A	DELAYED	DS-BC	D	A
AS-B	DELAYED	DS-AC	D	A
AS-C	DELAYED	DS-AB	D	A
AS-AB	DELAYED	DS-C	D	D
AS-AC	DELAYED	DS-B	D	D
AS-BC	DELAYED	DS-A	D	D
AS-ABC	DELAYED	DS-0	A	D

A - AGREEMENT

D - DISAGREEMENT

and a failure within the IDD or the DSC. If it is caused by a failed signal, the signal can be isolated by the ODD and replaced. This solution can be confirmed if the IDD now indicates the absence of input signal failures. However, if the deviation was due to a Receiver System Failure within the DSC or the IDD, the IDD results would be unchanged when the "suspected" input clock signal is replaced. The presence of more than one signal failure poses a problem in that one failed signal may be replaced by another failed signal resulting in the removal of an unfailed Receiver System. Therefore, in a clocking network, the condition of the spare input clock signals must be routinely monitored.

3.3 Alternative Detection Test

In the first version of the Delay Test (Figure 3.1), the IDD detects the presence of a signal failure, but requires the results from the ODD to isolate the source. Another arrangement is to modify the IDD to enable it to detect and isolate a failed input clock signal. This arrangement plus other suggestions are briefly discussed in this section. The purpose of these alternatives is to allow flexibility in meeting the requirements of a particular clocking network to permit trade-offs between cost and reliability, hardware and software requirements, etc.

The Delay Test produces simulated failures by delaying the input clock signals. The other method of producing simulated failures, the Pattern Test, substitutes a special "Fail Clock" for an input clock signal. The effect of this Fail Clock upon the flip-flops it drives is to make the output alternate between a TYPEONE and TYPEZERO voter input.

3.3.1 Delay Test (Alternatives)

Two different arrangements of the Delay Test are outlined in Figures 3.17 and 3.19. In Figure 3.17, the Reference Clock Receiver in the original version (Figure 3.1) is replaced by a simple majority voter for producing the Reference Clock. An unfailed majority voter will produce a reliable output as long as there are no input signal failures. The results from a modified IDD (MIDD) are used to confirm this condition. The MIDD has the ability to detect and to identify the failed clock signal. The ODD is no longer responsible for isolating the failed signal but can be used to check the verdict of the MIDD. A simple design for a MIDD is outlined in Figure 3.18 to illustrate its general workings. The results expected from the MIDD are listed in Figure 3.18.

In Figure 3.19, the reference clock is selected from among the input clock signals by a Reference Clock Selector (RCS). The IDD results are used to determine if a signal has failed. If one has occurred, the RCS uses the ODD to compare each

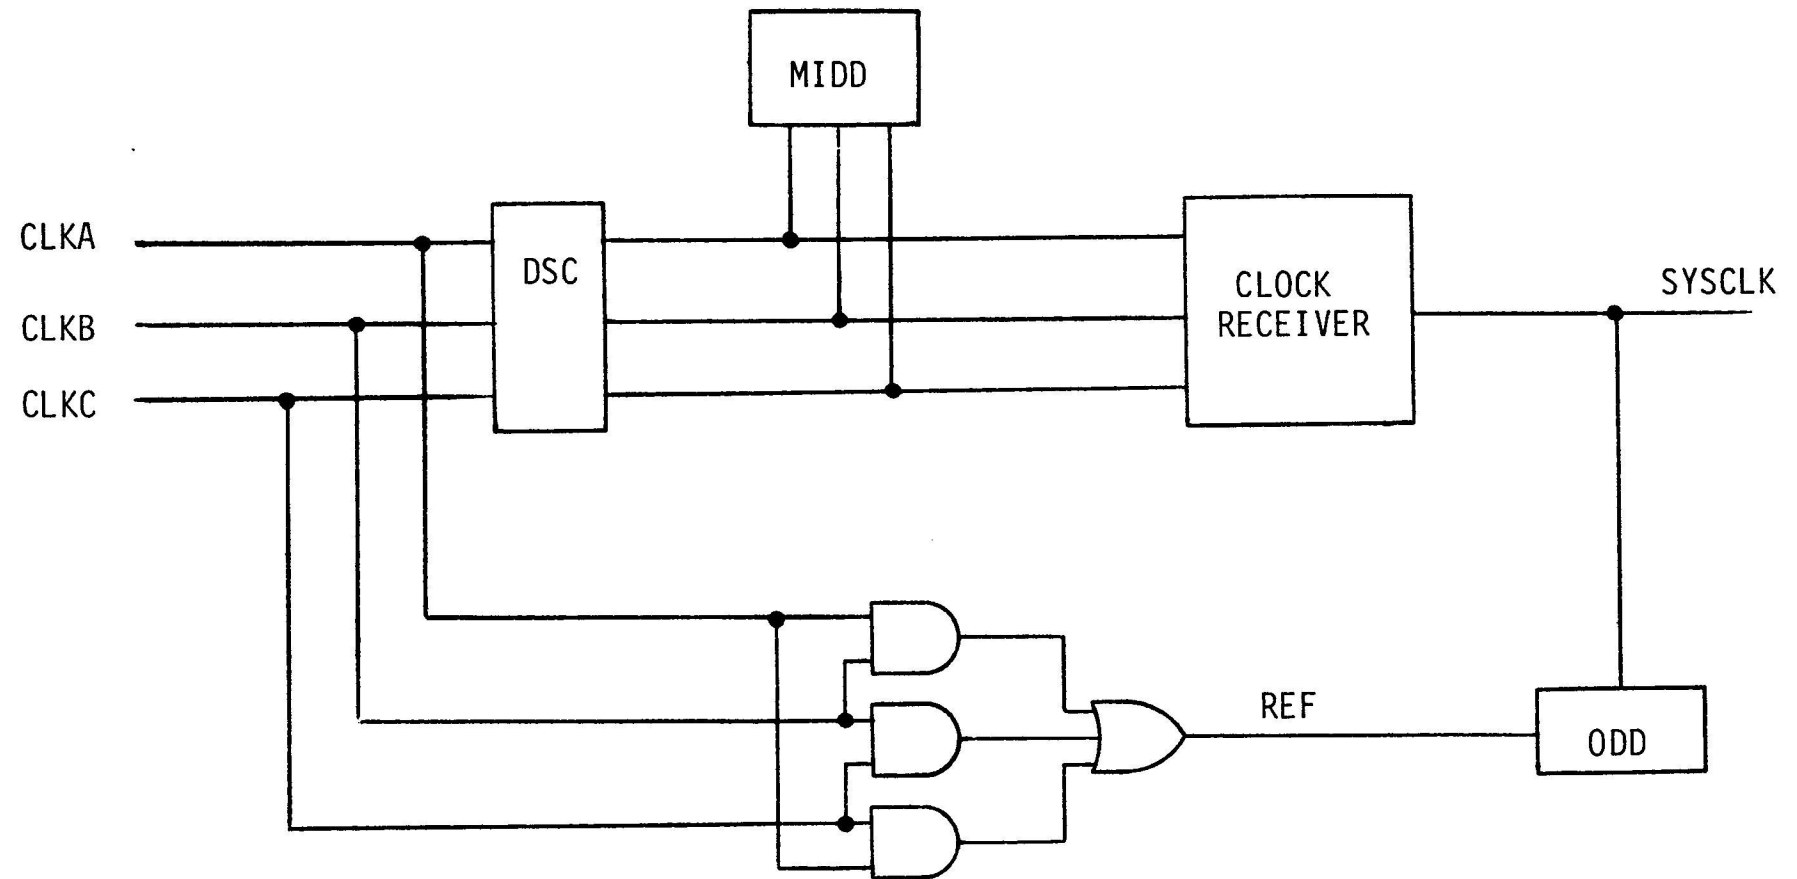



Figure 3.17 The Delay Test (Version #2).



EXPECTED RESULTS

FAILURE CONDITION	1	2	3
DS-0	A	A	A
DS-A	A	A	D
DS-B	D	A	A
DS-C	A	D	A
DS-AB	A	D	A
DS-AC	D	A	A
DS-BC	A	A	D
DS-ABC	A	A	A

A - AGREEMENT D - DISAGREEMENT

Figure 3.18 Modified IDD.

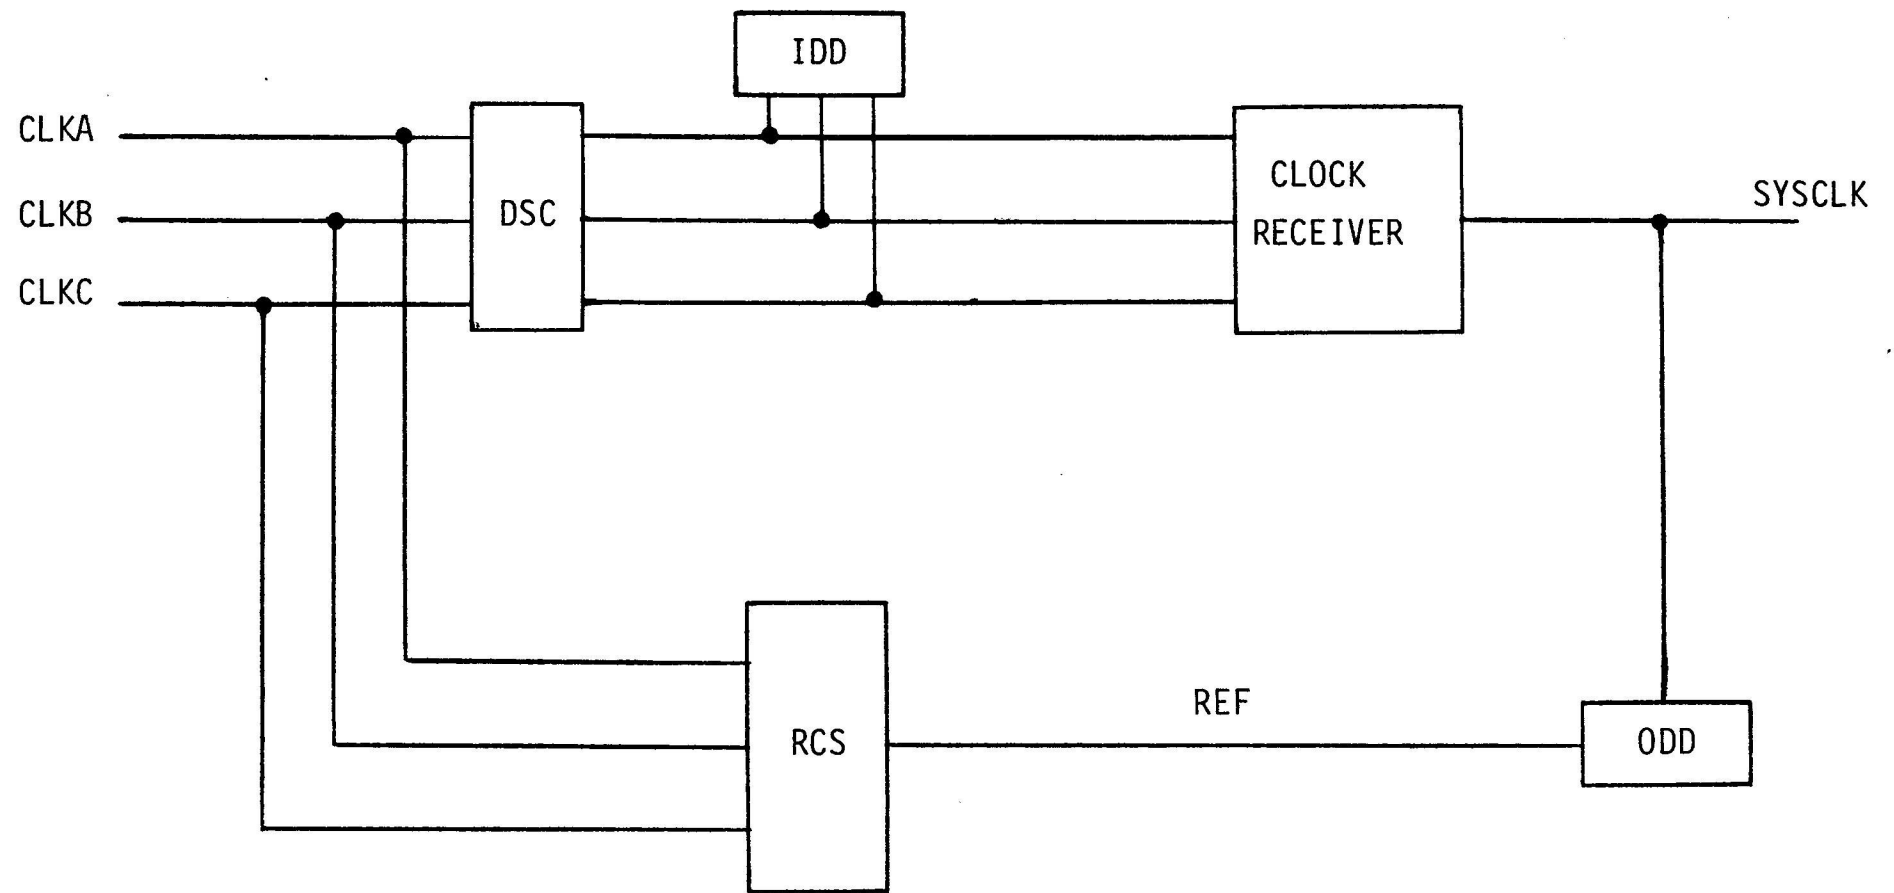


Figure 3.19 The Delay Test (Version #3).

input clock signal with the system clock under the DS-0 condition. The system clock is fault-tolerant when none of its input signals is delayed. By comparing each input signal with the system clock, the failed signal can be isolated. The smaller tolerance limit of the ODD can not be used since any one of the input signals can be used as the reference clock under the DS-0 condition. This requires that the variations among the input signals be taken into account when determining the tolerance limit. For some clocking networks, it may be desirable to use a MIDD instead of an IDD to improve the reliability of detecting and isolating a failed signal. The signal that is suspected as a failure by the MIDD can be checked by using it as the reference clock and comparing it to the system clock.

In the original version (Figure 3.1), a failed signal is detected by the IDD and isolated by the ODD. For the version shown in Figure 3.17, a failed signal is detected and isolated by the MIDD. In the last version (Figure 3.19), a failed signal is detected by the IDD and isolated by the ODD through the comparison of each input signal with the system clock (DS-0). The design used is dependent upon the requirement of the clock network. In some systems, more than one design of the Delay Test may be employed in a clocking network. The implementation of a Receiver System into a clocking network is explored in Chapter 4.

3.3.2 The Pattern (High-Low) Test

A TYPEZERO voter input to each voter occurs when the driving input signal is delayed. A TYPEONE voter input results when the signal is advanced (the REF and the other signals are delayed). The Pattern Test achieves these results by substituting a special "Fail Clock" for one of the input signals. A voter input to each voter can be held low (TYPEZERO), if the driving signal is replaced by a stuck signal (s-a-0 or s-a-1). If the driving signal is replaced by a signal oscillating at a higher frequency, the flip-flops driven by this signal are triggered before those driven by the other signal (TYPEONE). The Fail Clock may be produced by alternating a high frequency burst with a low period, both lasting for several cycles of an input clock. Figure 3.20 shows one design for producing the Fail Clock by using an odd number of inverters as a crude oscillator. The high frequency period of the Fail Clock results in the output of a flip-flop driven by the Fail Clock being high unless the clear input is low. The low period results in the output being low.

Figure 3.21 shows one possible arrangement for the Pattern Test. The Test Input Selector (TIS) selects which input signal is replaced by the Fail Clock. The IDD confirms that the substitution has been made. The IDD is capable of detecting signal failures only when no substitutions have been made. The ODD is tested by

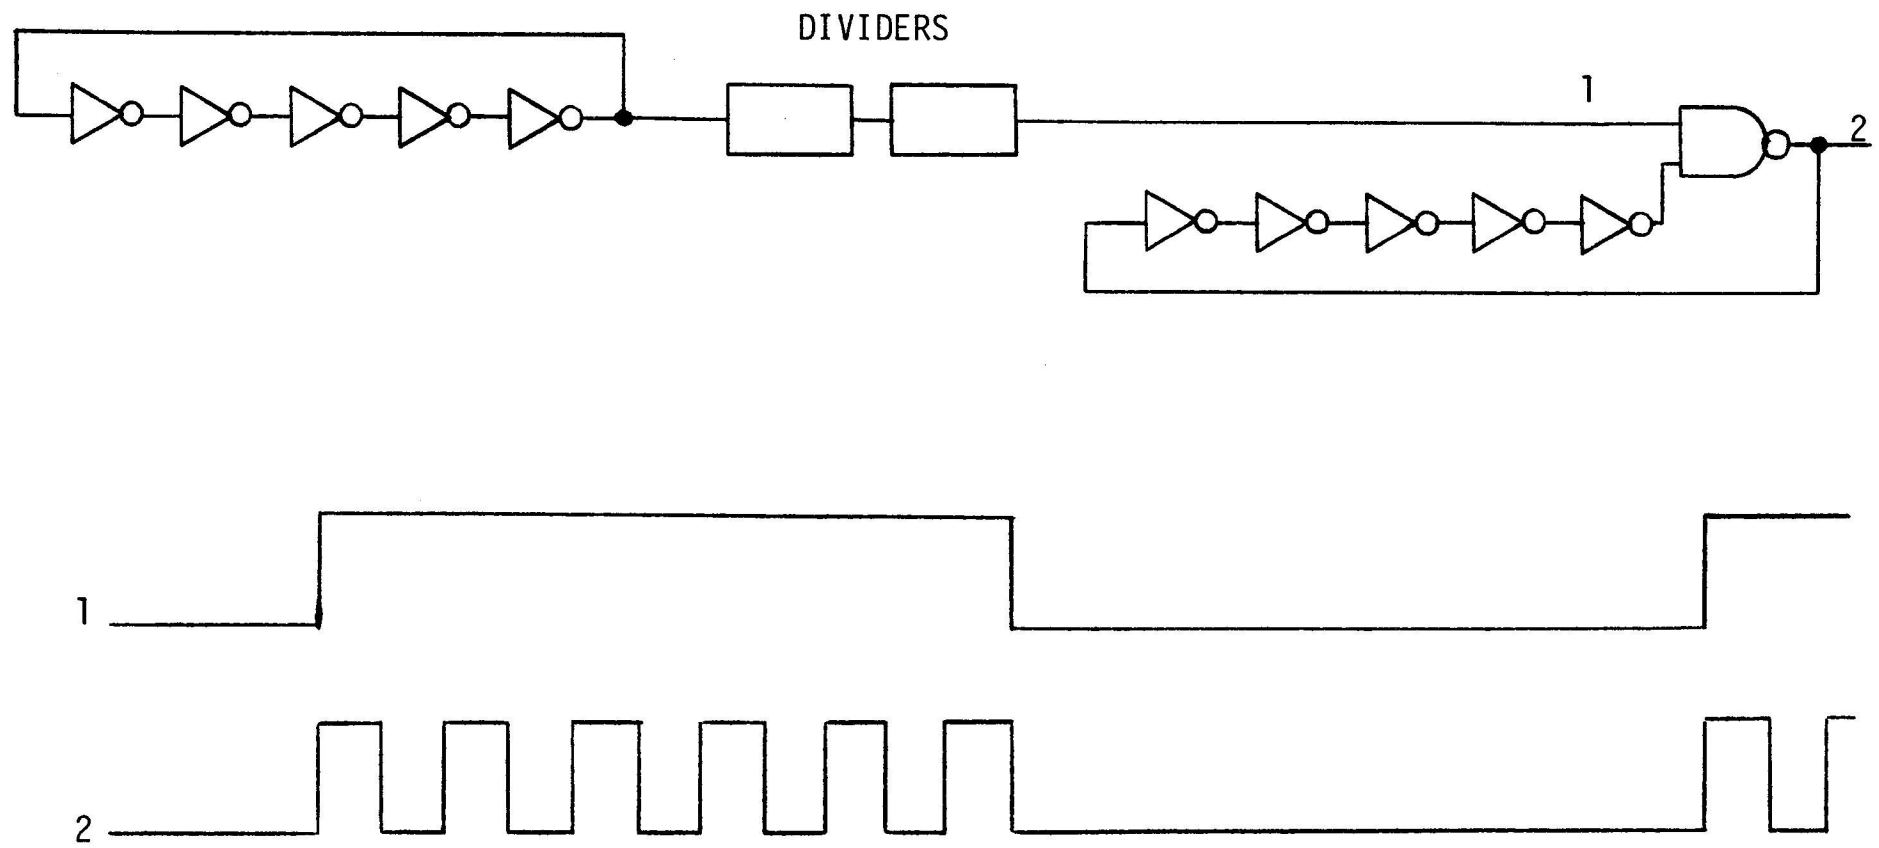


Figure 3.20 The Fail Clock.

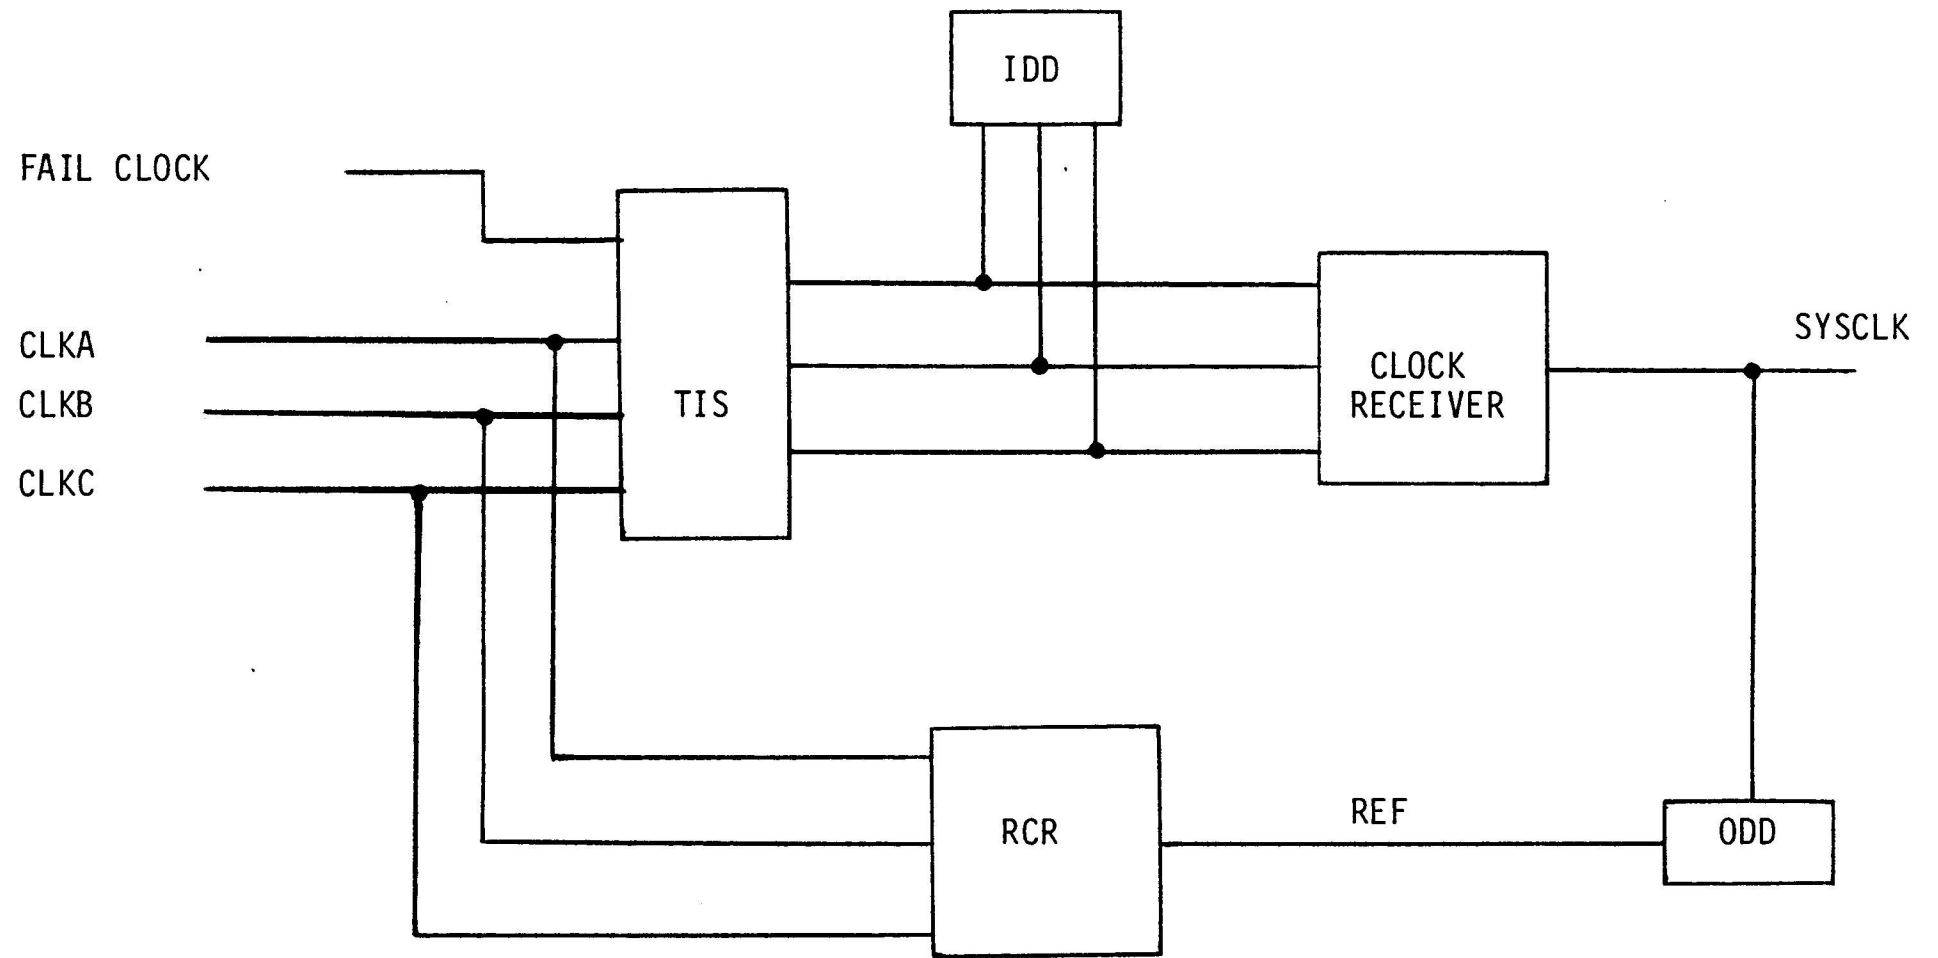


Figure 3.21 The Pattern Test.

substituting the Fail Clock for two or three input clock signals at the same time. Table 3.7 relates the four test conditions with those of the Delay Test (Table 3.6).

TABLE 3.7
COMPARISON OF THE PATTERN TEST WITH THE DELAY TEST

FAIL CLOCK SUBSTITUTED FOR	LOW LEVEL PERIOD (TYPEZERO)	HIGH FREQUENCY PERIOD (TYPEONE)
NO SUBSTITUTIONS	DS-0, AS-0	DS-0, AS-0
CLKA	DS-A, AS-BC	AS-A, DS-BC
CLKB	DS-B, AS-AC	AS-B, DS-AC
CLKC	DS-C, AS-AB	AS-C, DS-AB
ALL THREE	DS-ABC	AS-ABC

CHAPTER 4

APPLICATION

A clocking network is used to provide a timing reference for a synchronous computer system. For a fault-tolerant system, it is necessary to provide the clocking network with a fault-tolerant capacity to maintain the overall fault tolerance of the system. The Clock Receiver and the testing concepts discussed in this thesis can be incorporated into a clocking network to achieve fault tolerance through the detection, isolation and correction of faults as they occur. The CARDS system, a multiprocessor system currently under development at the C.S. Draper Laboratory, was chosen to illustrate how a fault-tolerant clocking network might be incorporated into a fault-tolerant system.

4.1 Multiprocessor Systems

The processors in a multiprocessor system are capable of executing their programs either collectively (usually in groups of two or three) or individually. As discussed in Chapter 1, fault tolerance can be achieved through comparison and/or voting with replicated units. In this case, replicated units are processors performing the same operation. When working collectively, the processors are in "tight synchronism" if they compare the results of each microstep at the end of each microinstruction. Tight synchronism requires the use of a time reference to insure that the processors are executing the same microinstruction. A single system clock is not used as the time reference since its failure would jeopardize the entire multiprocessor system. Instead, each processor has its own system clock produced by a Clock Receiver from the distributed primary clock signals. The influence of a failed system clock is limited to the processor unit it serves. The failure of a primary clock signal is tolerated by the Clock Receivers and then replaced by the clocking network using the fault detection techniques discussed in Chapter 3.

1. The processors execute individually.
2. The processors execute in pairs in tight synchronism.
3. The processors execute in groups of three (triads) in tight synchronism.

Figure 4.1 shows the basic elements of a multiprocessor in which the processors execute individually. All processors use the memory and the Input/Output Controller, but not at the same time. This configuration makes the most efficient use of the processors' computation capability (job utilization) at the cost of reliability. This arrangement is not able to tolerate all possible types of processor faults.

In the second mode, the processors execute in pairs. Figure 4.2 shows an example of a multiprocessor employing this configuration. Under the conditions discussed in the first chapter for a duplex system, this arrangement is more reliable than that of Figure 4.1. This is accomplished by the processors within a pair comparing their results to detect errors. However, this arrangement requires twice the number of processors to perform the same amount of work as that shown in Figure 4.1.

The third configuration is that of a Triple Modular Redundant (TMR) system. Figure 4.3 illustrates the basic concepts involved in this type of configuration. The units are arranged in groups of three (triads), but there are various ways of connecting and reconnecting these units. The third configuration, unlike the first two, continues to produce valid results after the single failure of any member of any triad. The number of failures tolerated and the overall reliability of the system can be increased by providing the multiprocessor with the ability to detect, isolate and replace failed units. This arrangement is essentially the same as the TMR System with Replacements discussed in Chapter 1. The CARDS multiprocessor employs such an arrangement to achieve high reliability.

4.2 The CARDS Multiprocessor

The CARDS system has the ability to group and regroup its processor and memory modules into triads composed of unfailed units. Figure 4.4 shows a schematic view of CARDS. An expanded view of a processor module and a memory module are shown in Figures 4.5 and 4.6 respectively. The processor module contains an error latch, two Guardsmen and bus control logic in addition to the microprocessor and scratchpad memory. The memory module also contains bus control logic along with two Guardsmen. A Guardsman or Bus Guardian Unit (BGU), proposed by Smith⁶, is used to select when, and on which bus line, the module transmits. It also serves as an addressable isolation unit which can remove the module it serves from the multiprocessor if the module fails. Figure 4.7 shows the essential elements of the Guardsman. The information concerning the bus line selection is clocked into the Output Control Registers by the Guardsman's clock along with the Guardsman's address. The data within the Output Control Registers controls the Bus Isolation Gates (BIGs) which, in turn, enable the module to transmit on a particular line. The Bus Isolation Gates are controlled by two Guardsmen in order to protect the system from a single failed

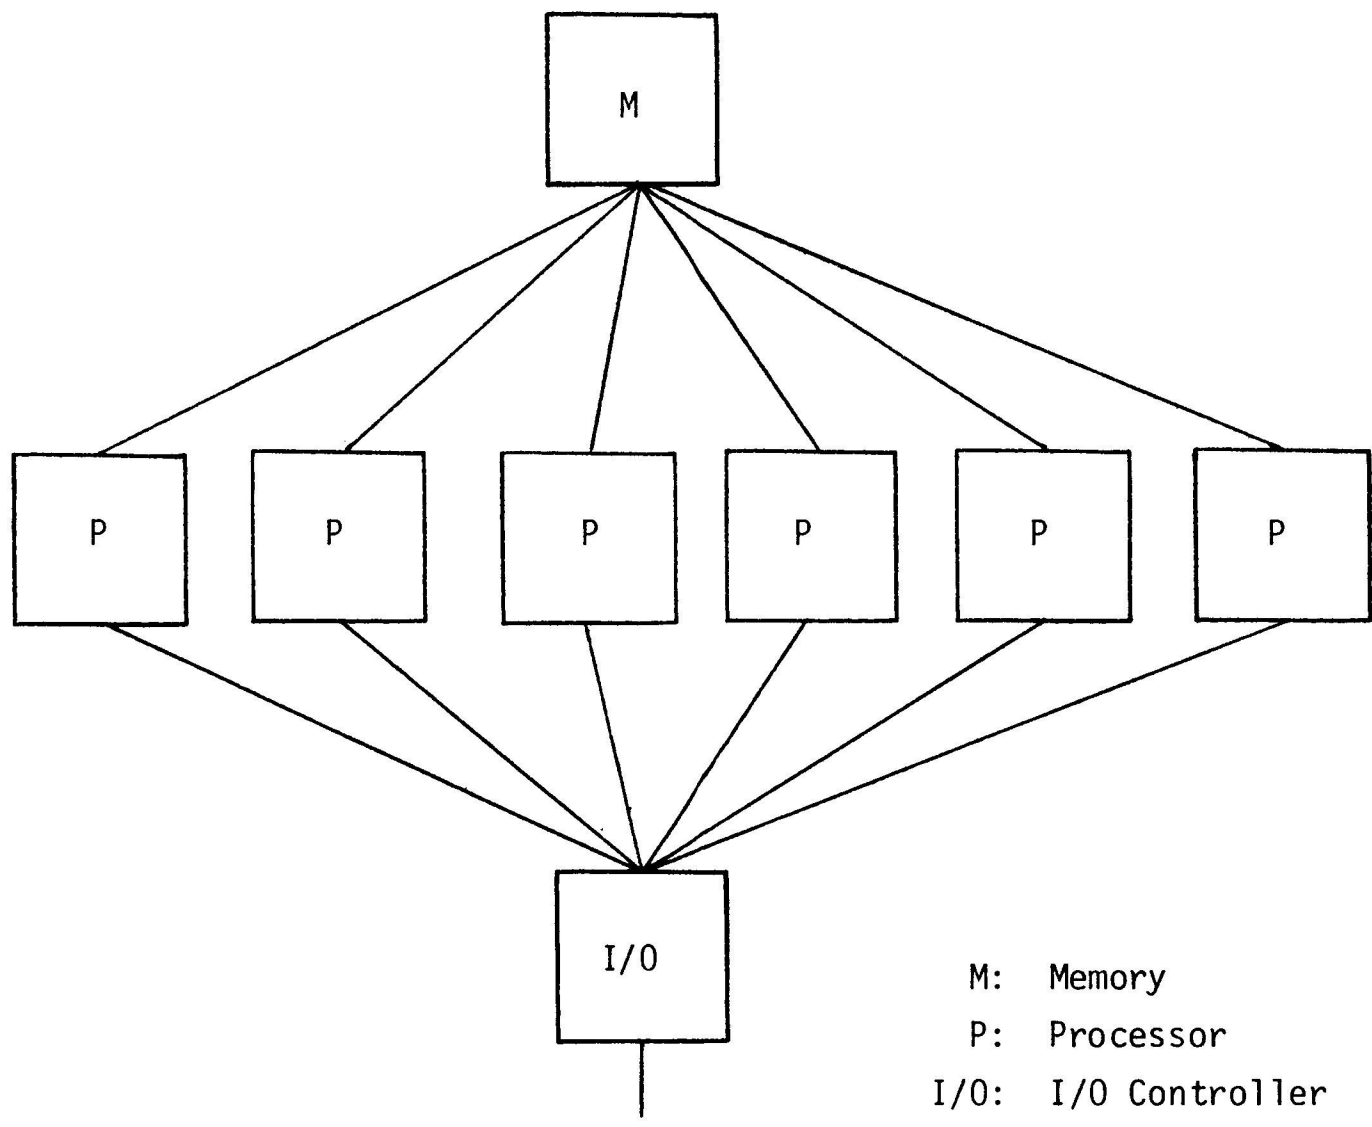


Figure 4.1 Multiprocessor (Simplex).

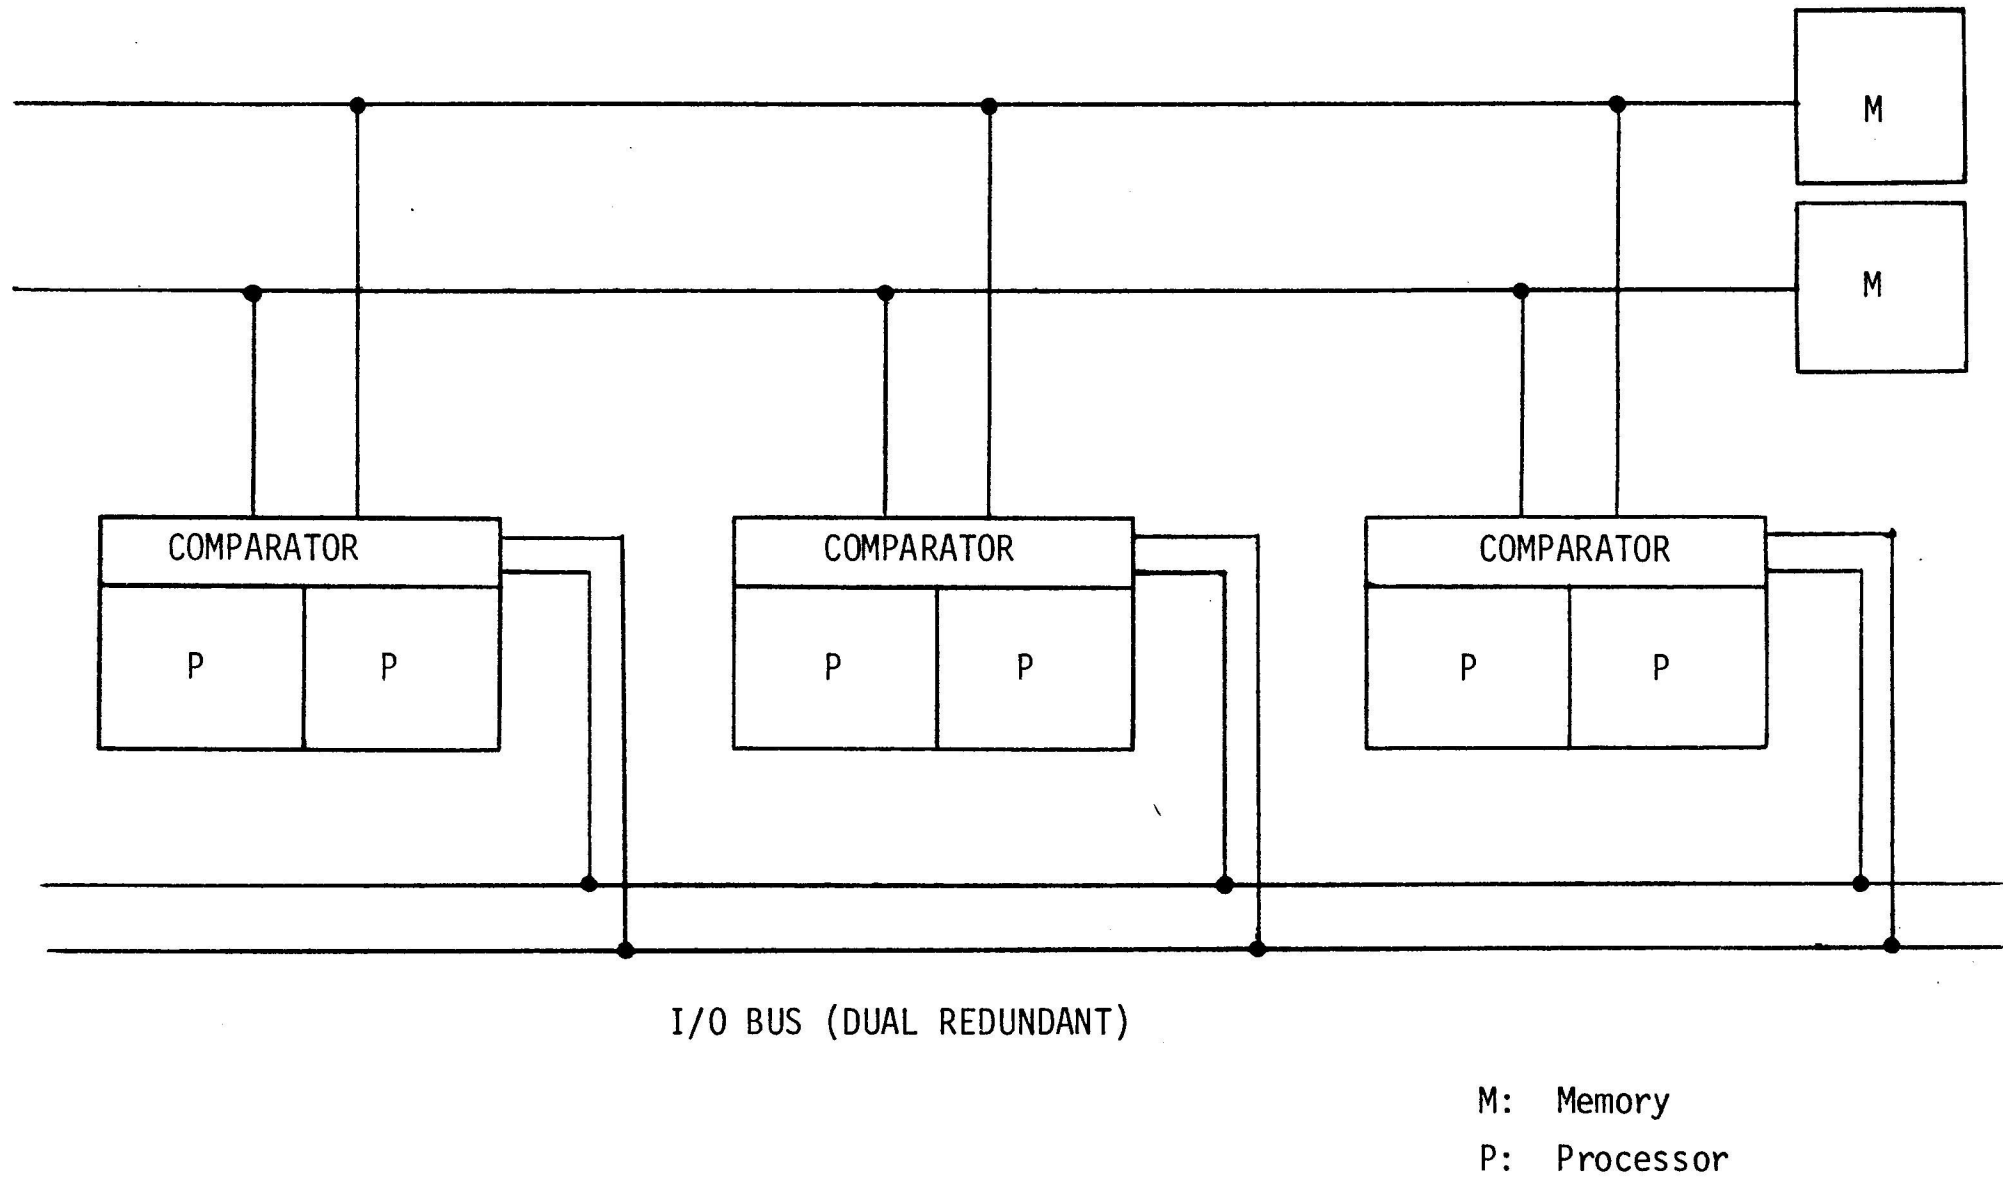


Figure 4.2 Multiprocessor (Dual Mode).

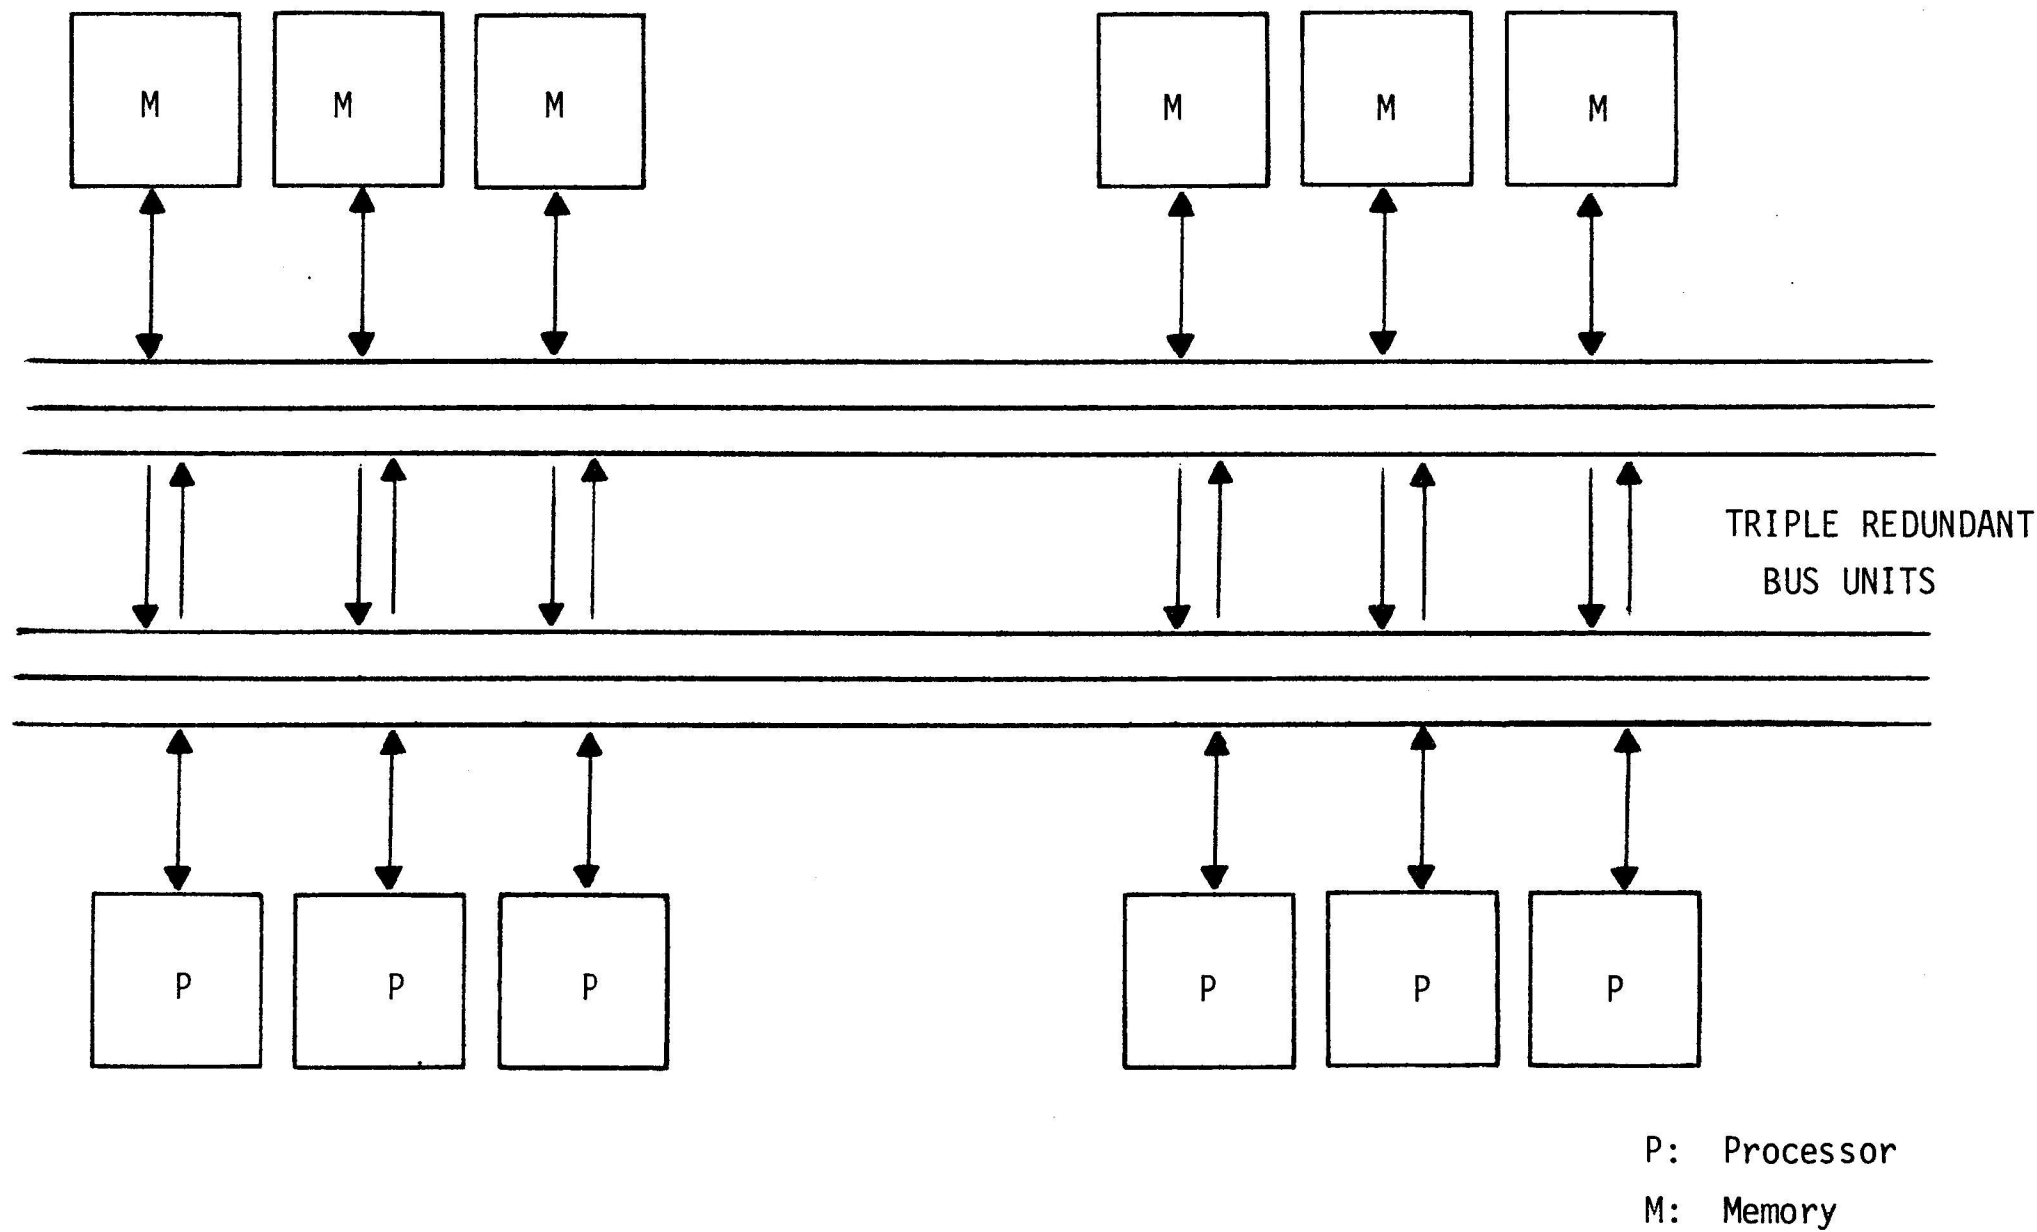


Figure 4.3 Multiprocessor (TMR Arrangement).

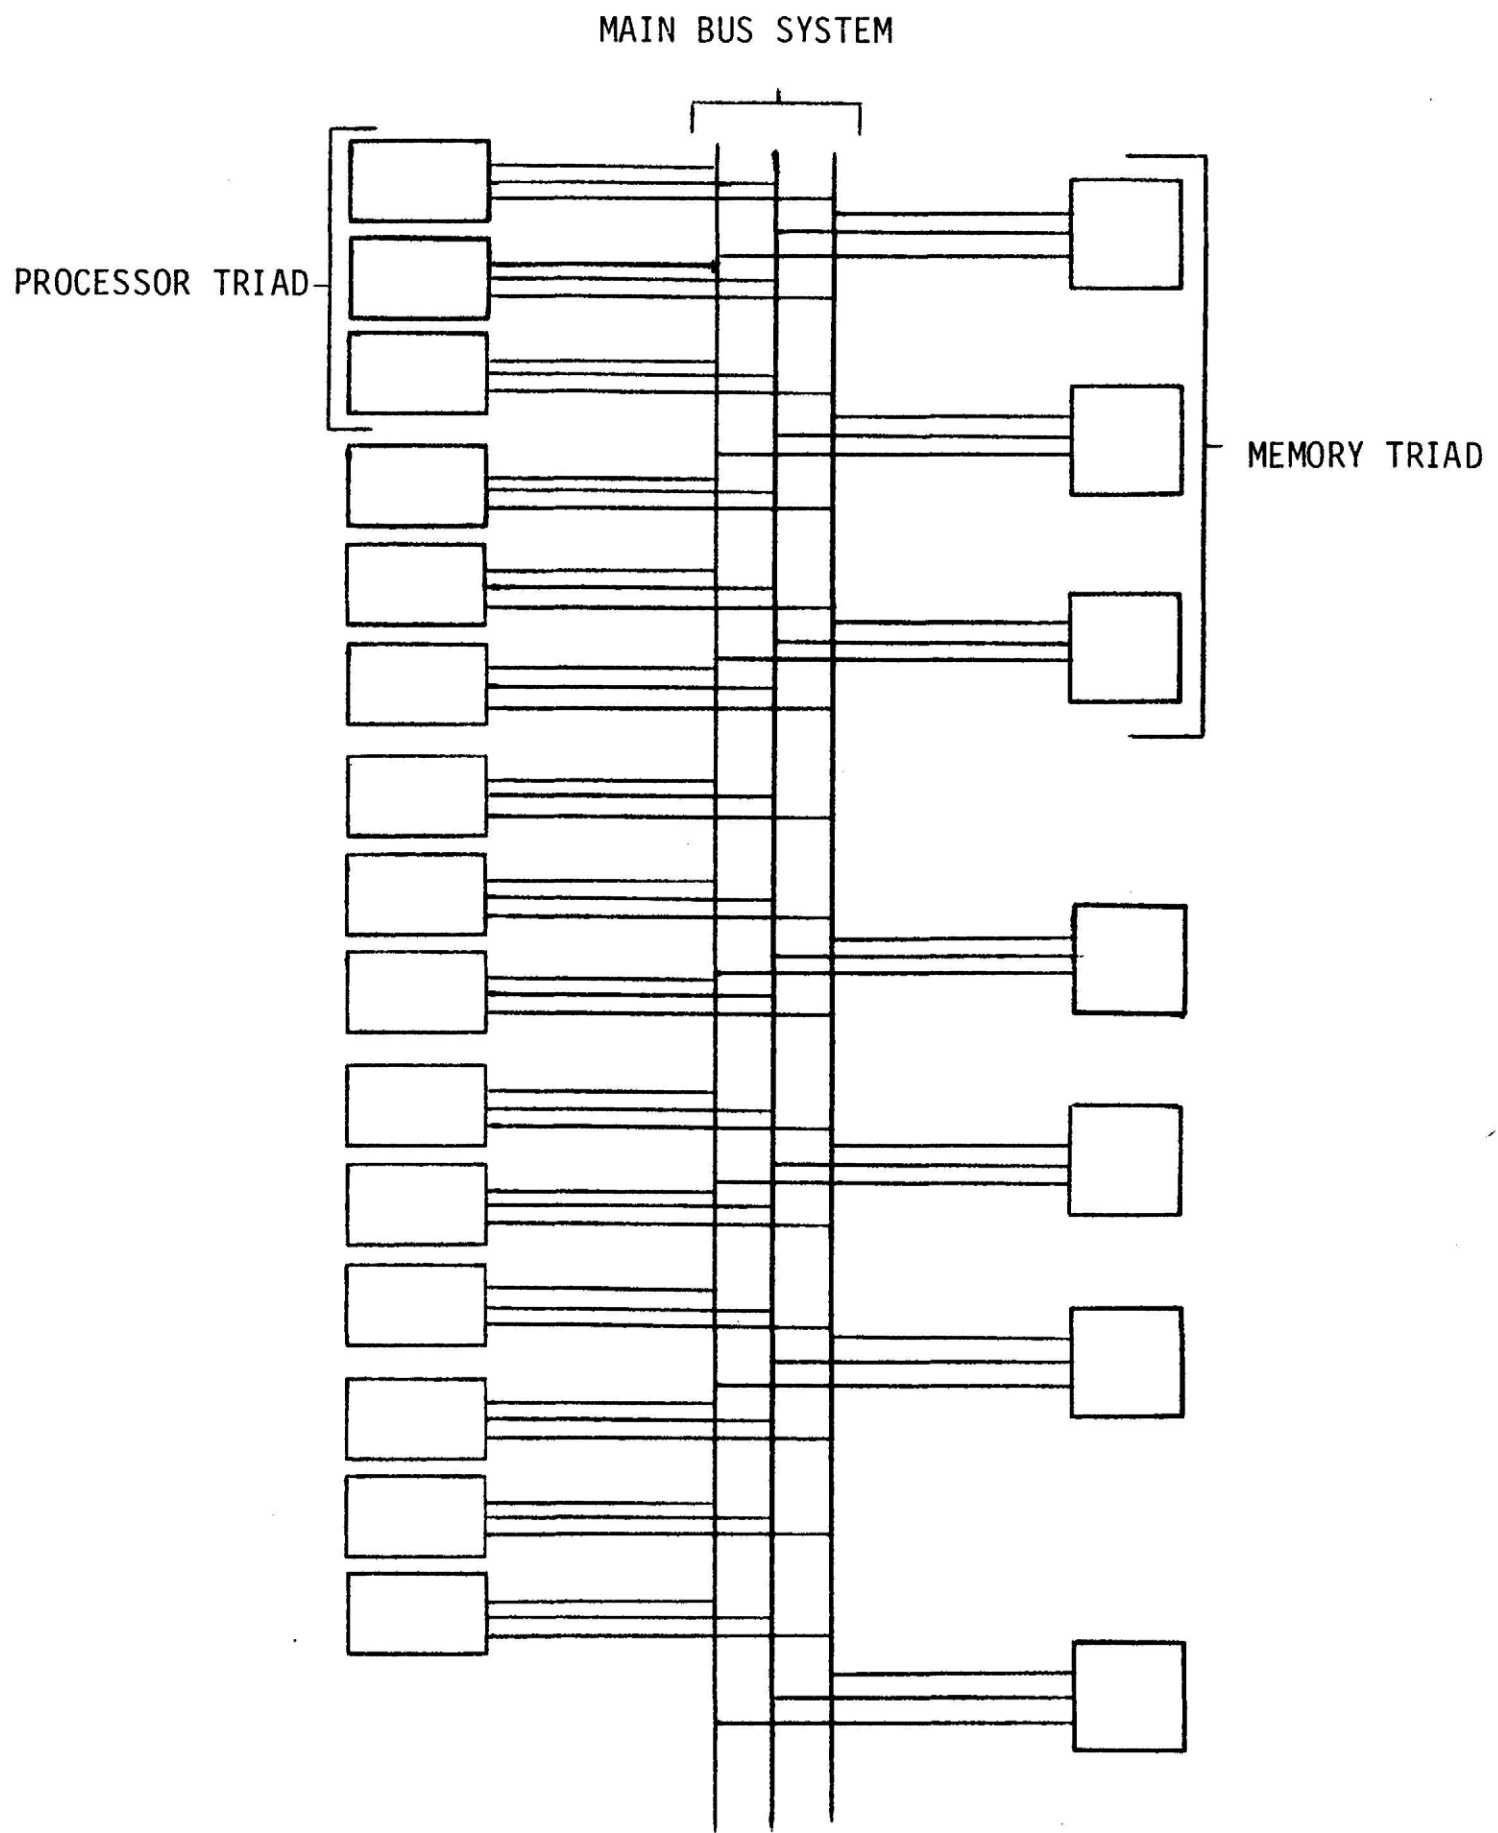


Figure 4.4 The CARDS System (Overview).

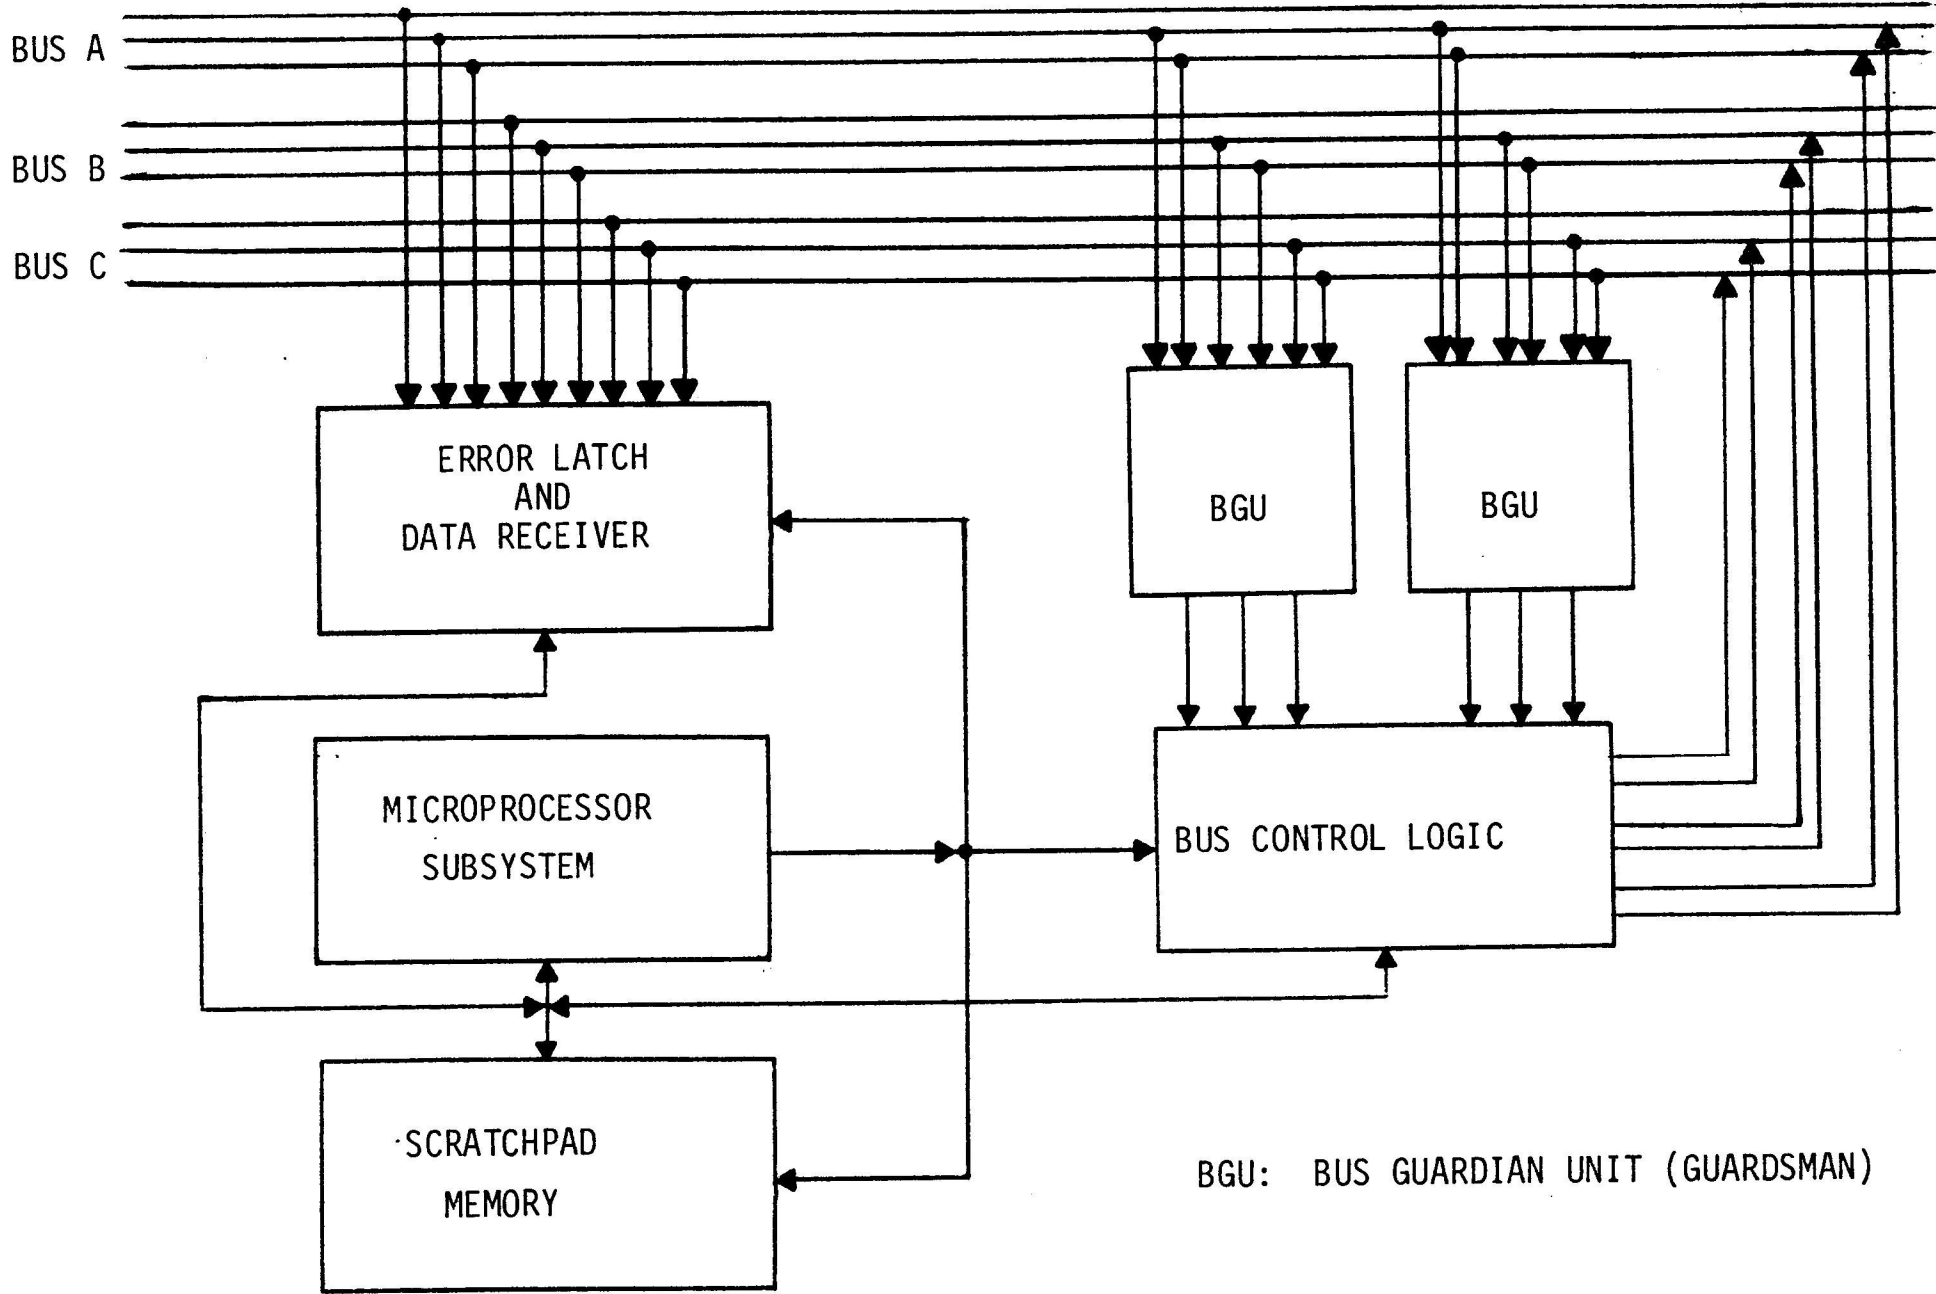


Figure 4.5 Processor Module.

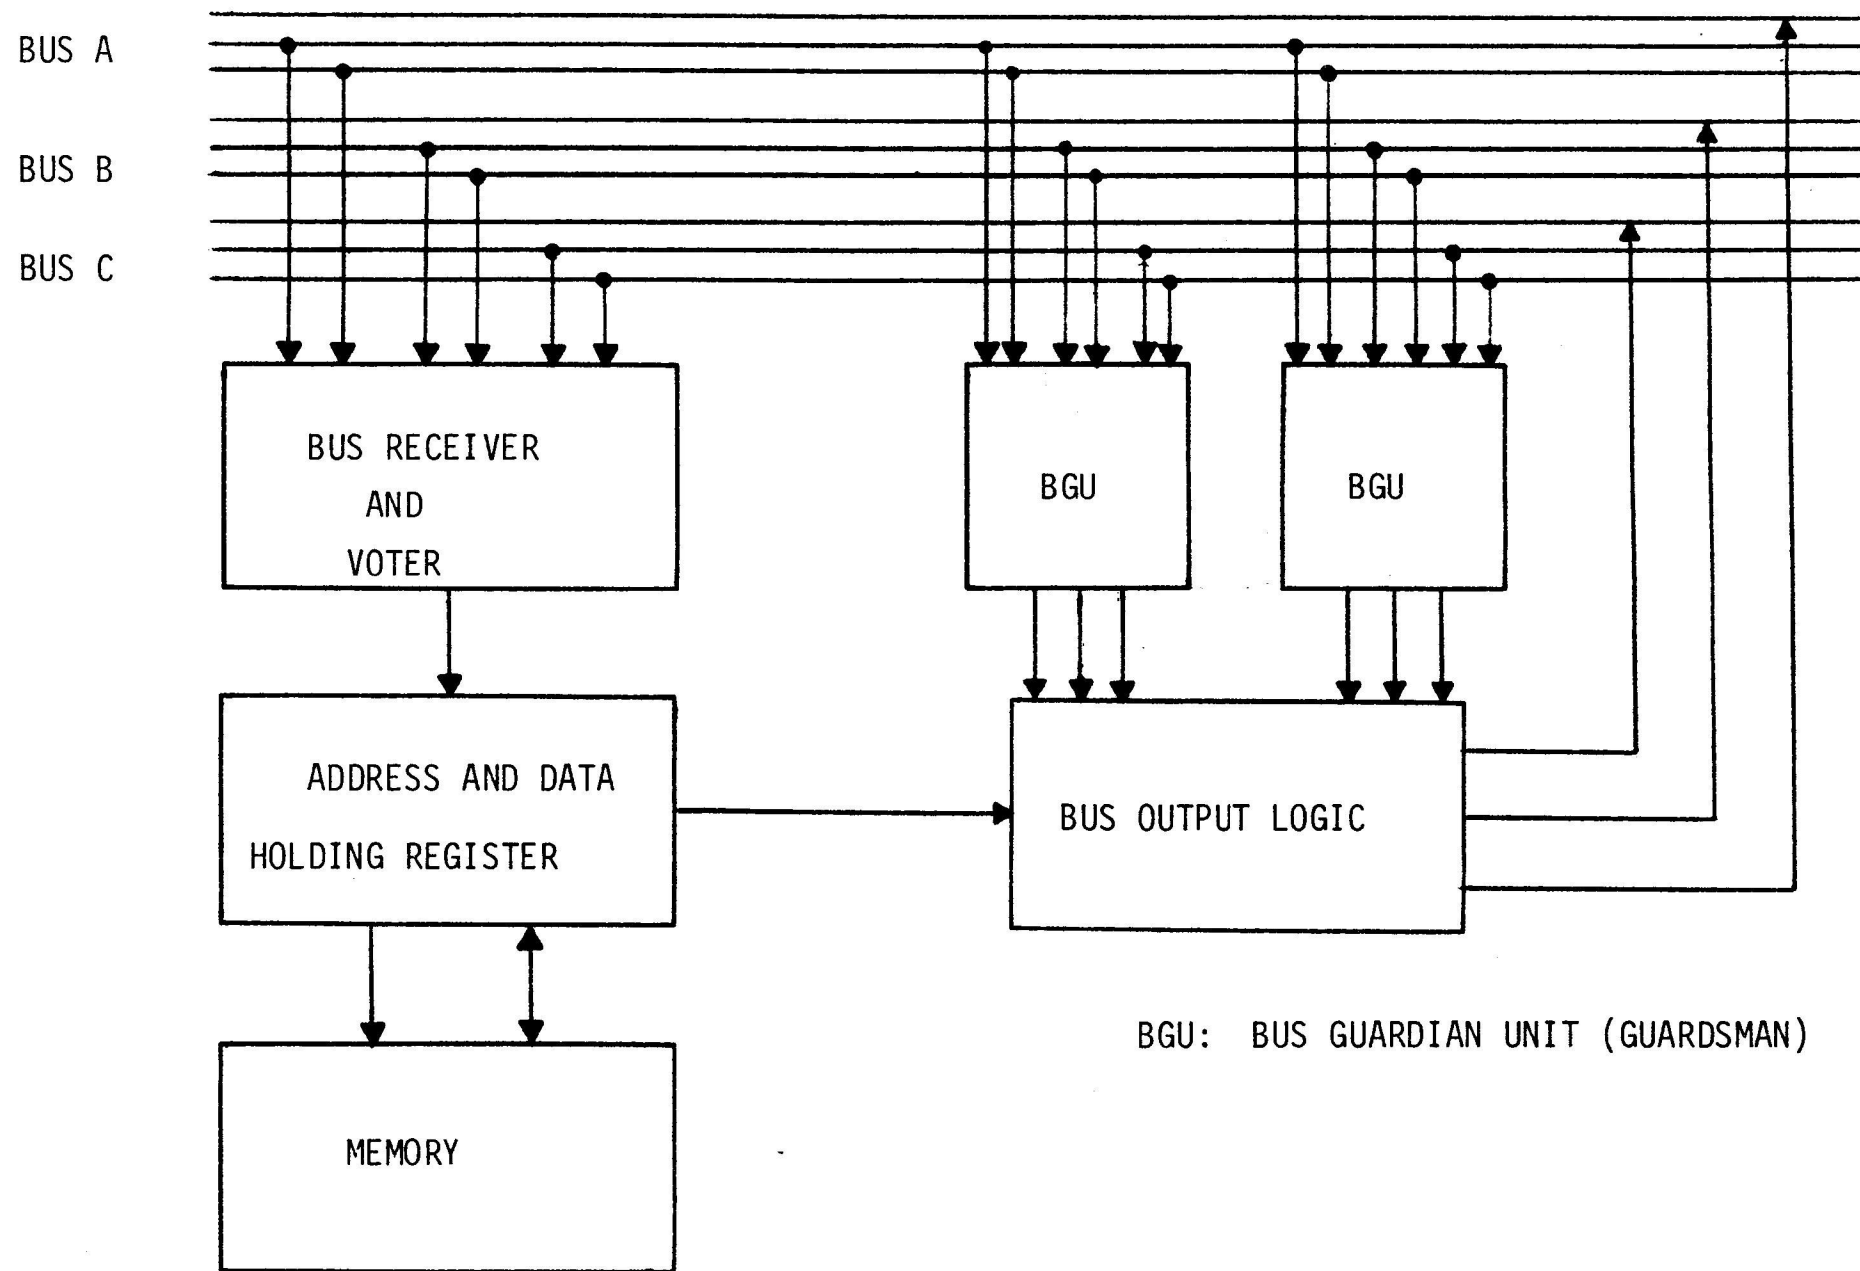


Figure 4.6 Memory Module.

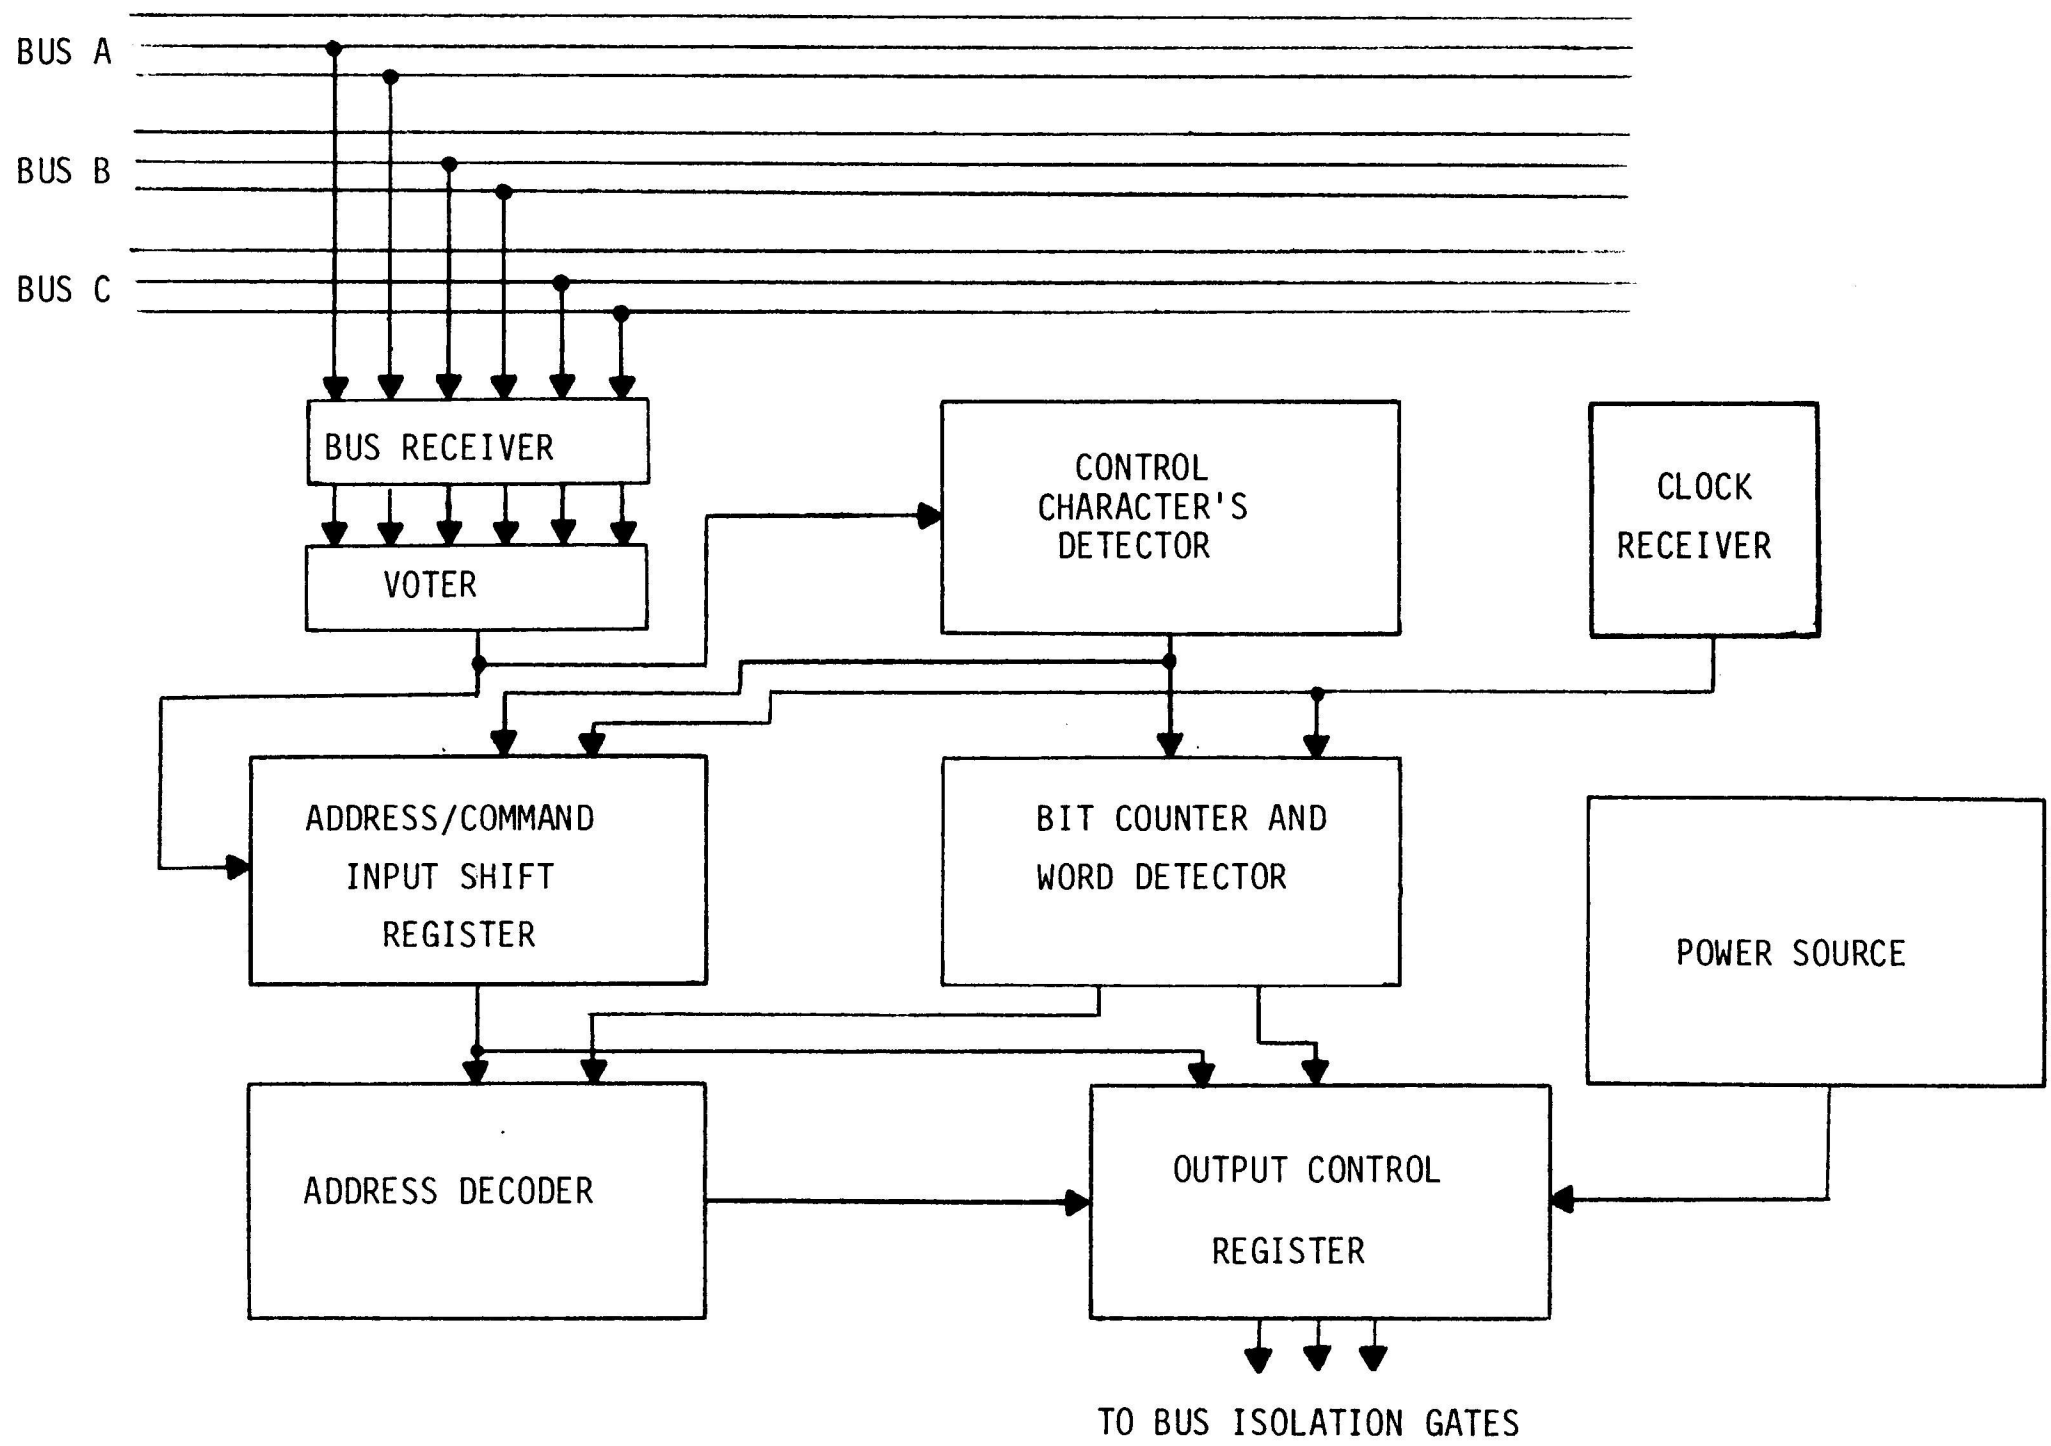


Figure 4.7 Guardsman.

Guardsmen that might allow its module to transmit at the wrong time and/or on the wrong line.

The error latch in each processor monitors all bus transactions and records any error which occurs on the bus line, if associated with the activities of its triad. The processors of this triad can determine by reading the contents of their error latches which bus line and/or transmitting module contains the fault.

4.2.1 Clocking Requirements

As presently constructed, the CARDS system uses fifteen processor modules, seven memory modules and three bus lines. Since two Guardsmen are needed to control a module's access to the bus lines, the CARDS system has a total of forty-four Guardsmen. A system clock is distributed by the clocking network to each microprocessor, memory element and Guardsman. The system clocks of each microprocessor and memory element are produced by separate clock receivers so that the failure of a clock receiver affects only the module it serves. To prevent a module from inadvertently transmitting over a bus line as the result of a single clock receiver failure, the system clock for each BGU are produced independently. Thus, two clock receivers are required for each module, one for each Guardsman. The system clock for a microprocessor or a memory element can be obtained from the same clock receiver used by one of its Guardsmen. The principal elements of a processor triad and a main memory triad are shown in Figure 4.8. Six Guardsmen, and therefore six clock receivers, are needed for each triad.

The testing procedures are designed to exercise specific components by subjecting them to a stimulus and noting the response. This testing is controlled and evaluated by a processor triad. The evaluation of the testing is performed by the error latch in each processor of the testing triad.

4.2.2 The Error Latch

The error latches of a processor triad monitor the information on the bus line when data is being transmitted or received by that triad. A faulty module within a triad which fails to agree with its partners will result in the data carried on one bus line differing from that on the other two lines. Figure 4.9 shows a schematic outline of an error latch. The TMB Bus Receiver, Voter and Comparator monitors the processor to memory communications while the FMB Bus Receiver, Voter and Comparator covers the memory to processor link. The flags in the error latch can be read by the microprocessor. Since each processor has its own error latch, a failure becomes known to all unfailed processors in the triad to which the failure pertains. The error indication provided by the error latch helps determine the identity of the

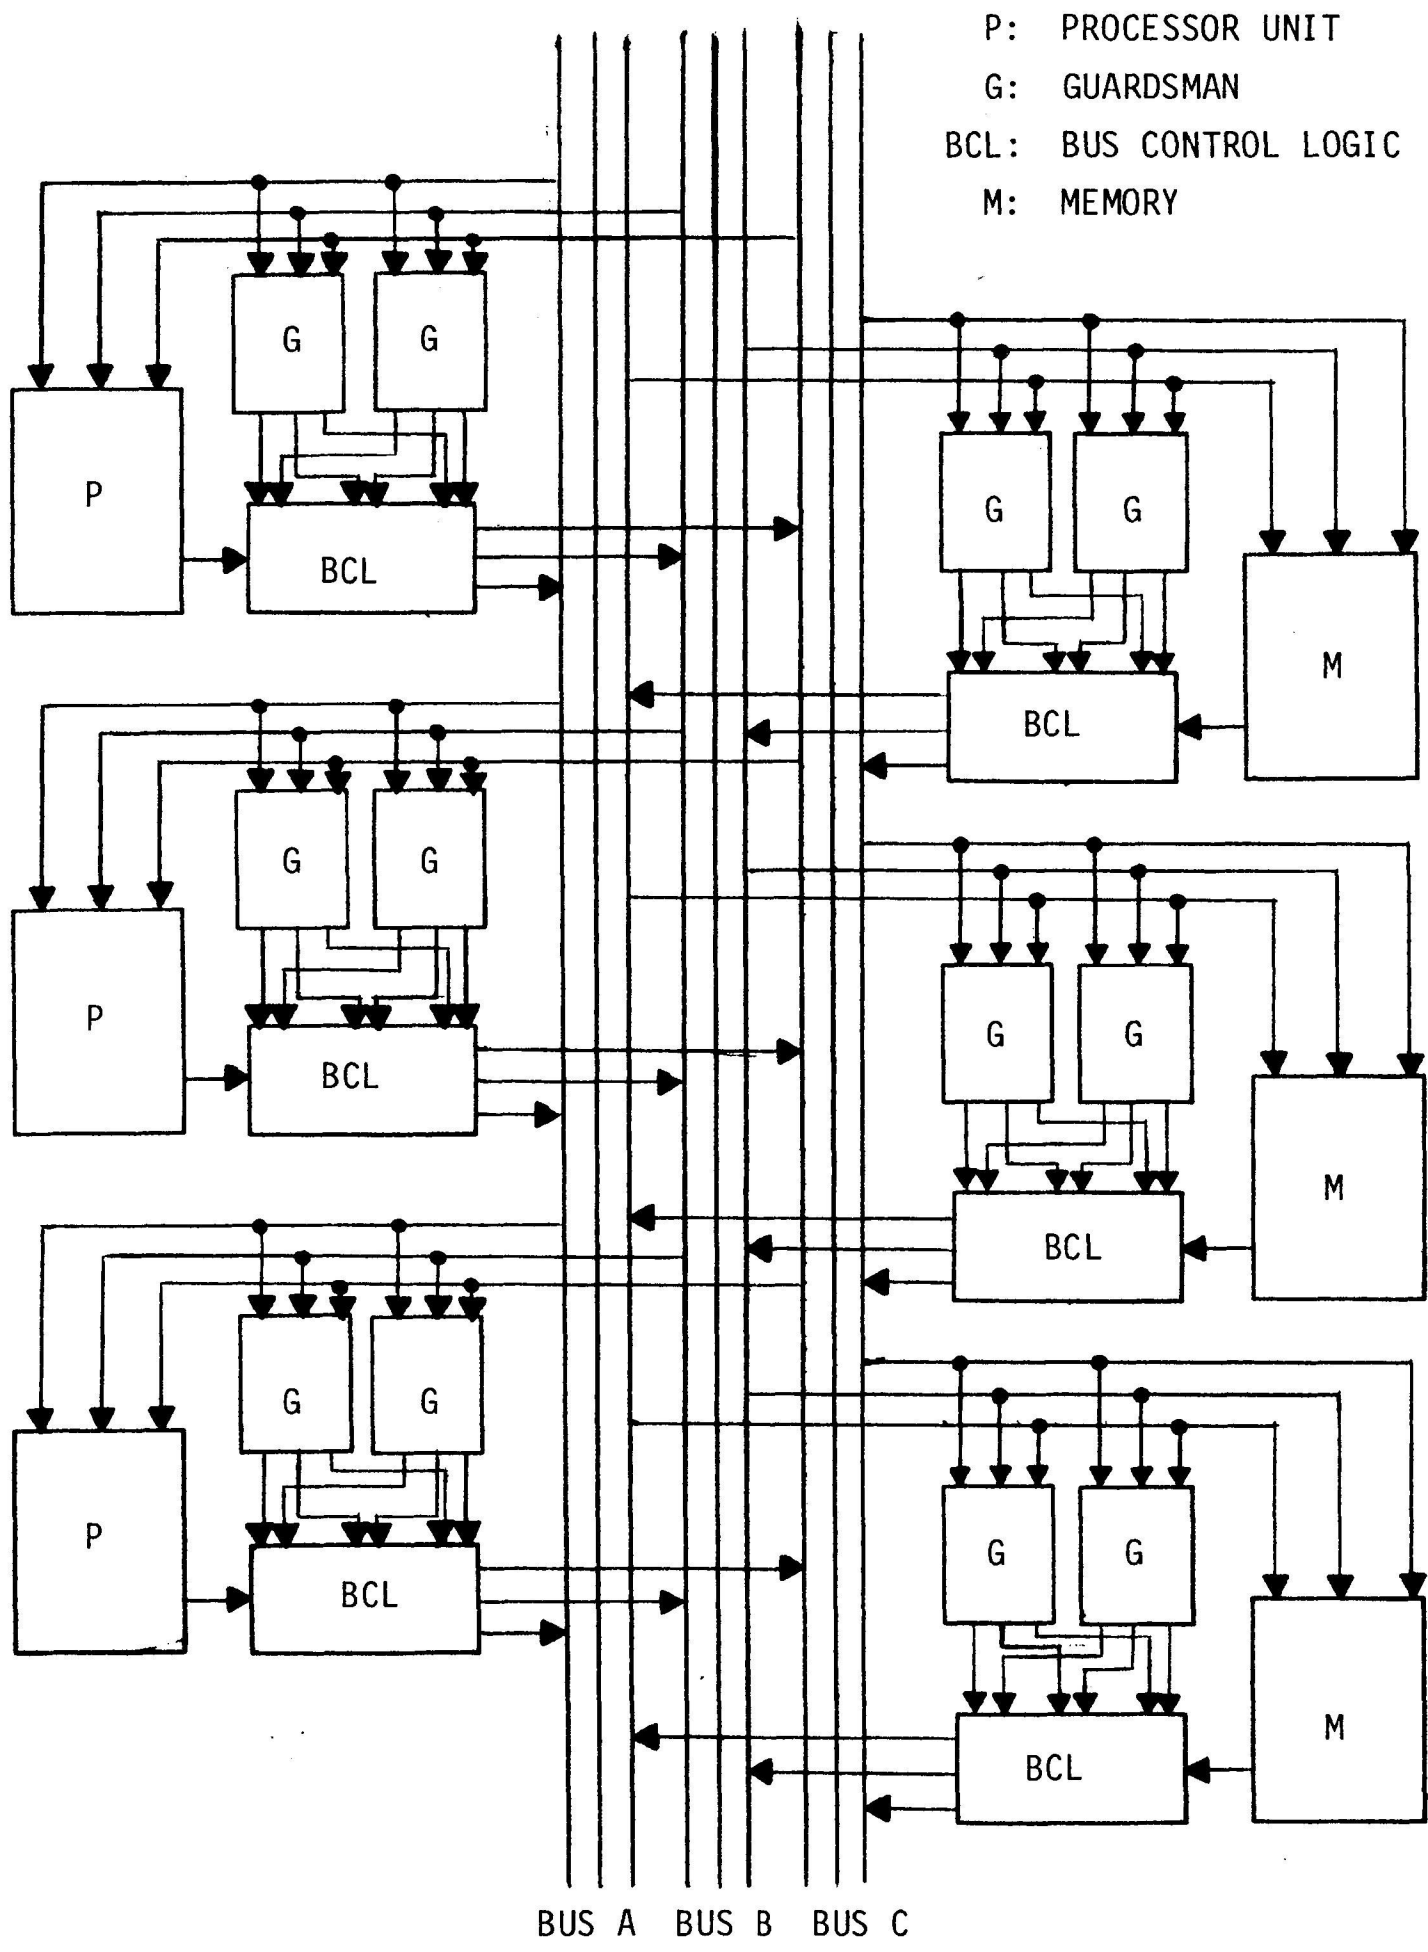


Figure 4.8 Processor and Memory Triads.

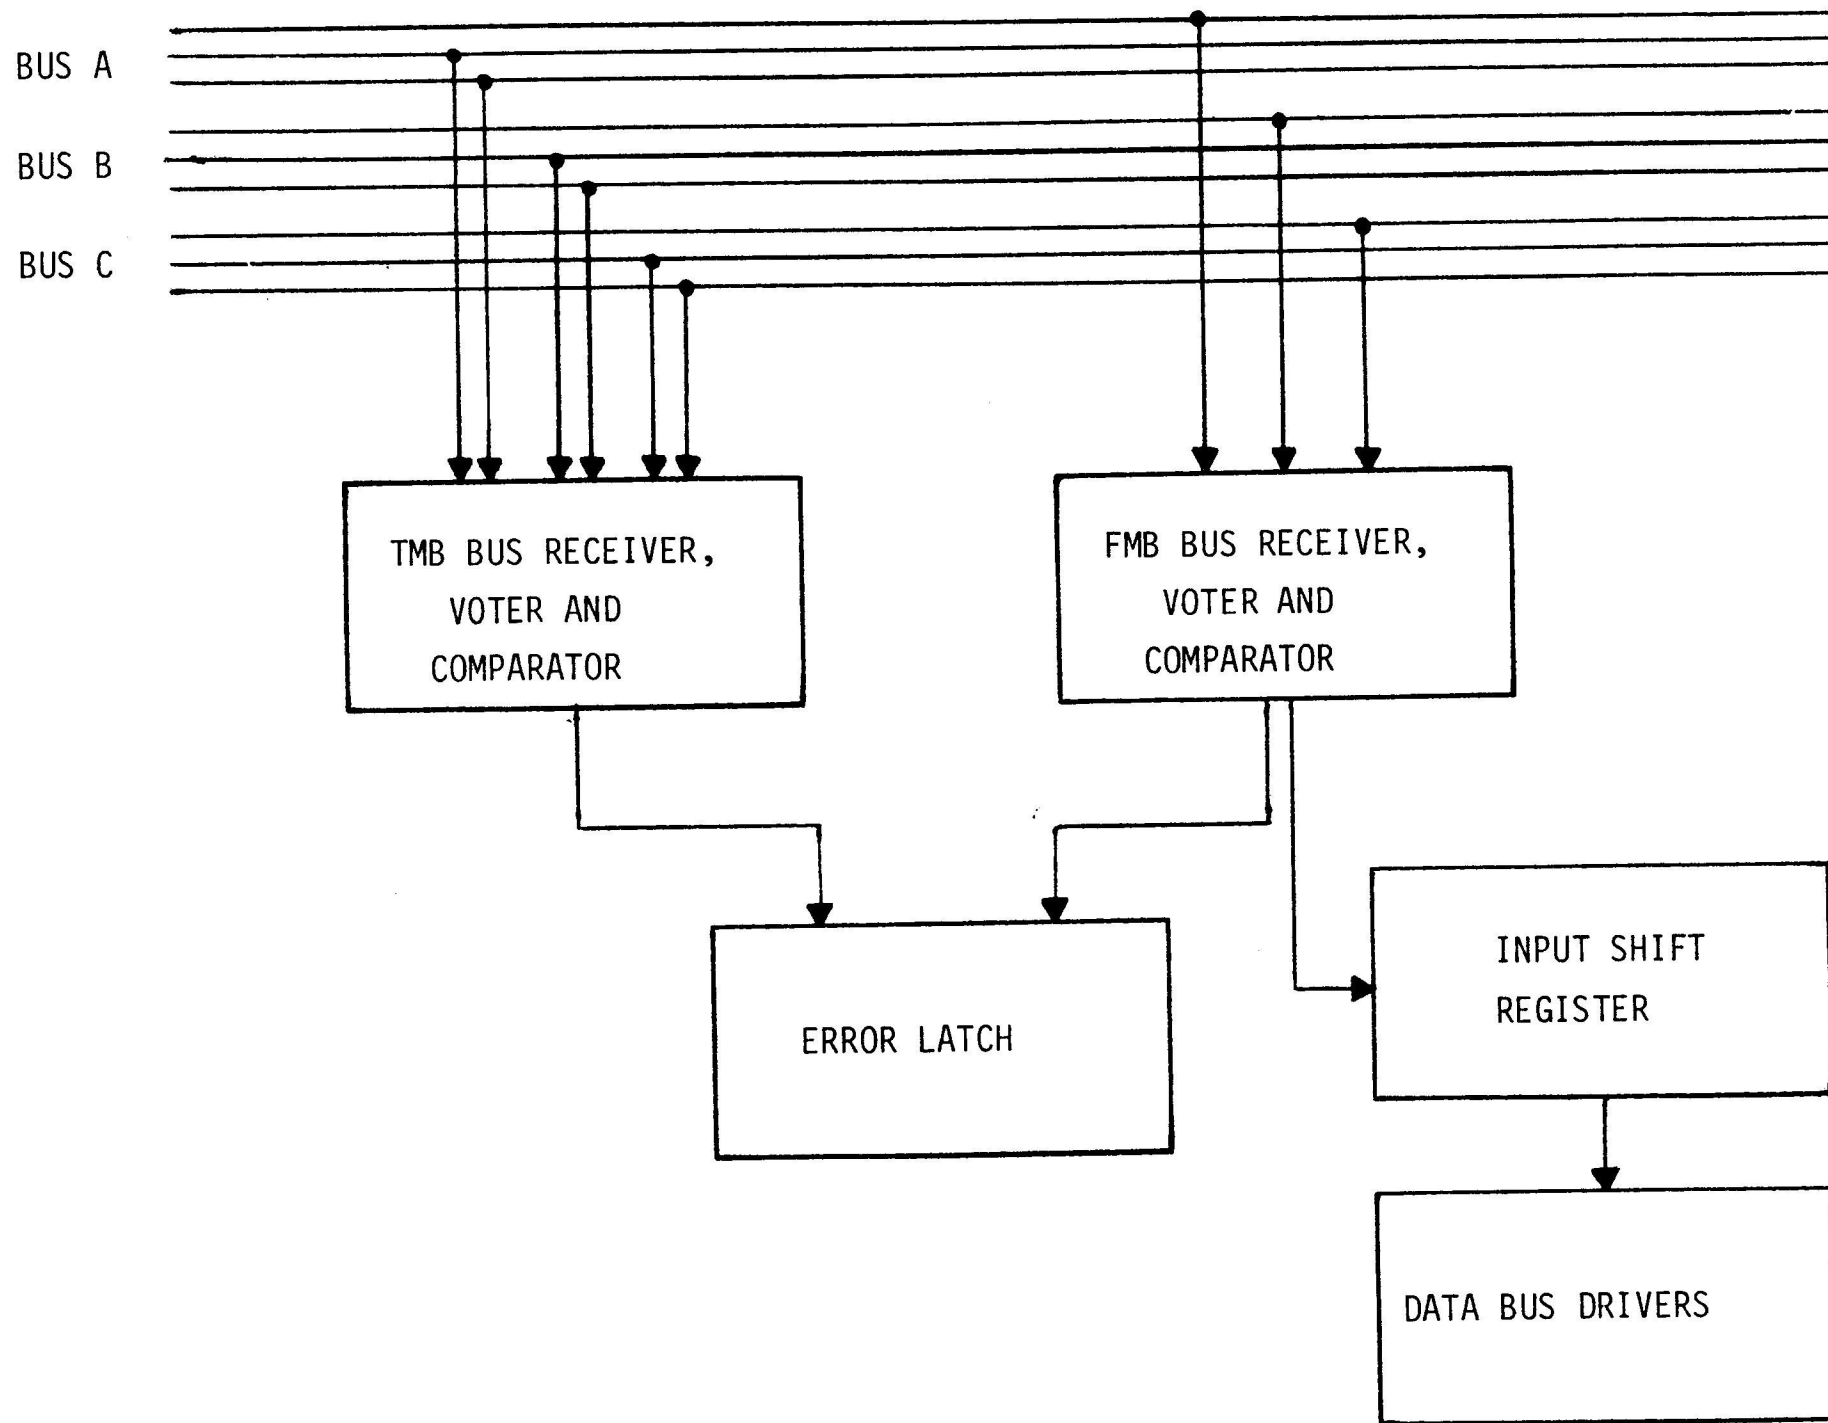


Figure 4.9 Error Latch.

failed unit.

Though the possibility of a bus line failure is smaller than that of either a processor or memory module failing, it must be considered because of its widespread effect. A failed bus line means that information transmitted over that line by an unfailed module in each triad will disagree with its unfailed partners. This situation can be resolved by transmitting the information from the suspected module over a different bus line. In this way, a faulty module can be distinguished from a bus line failure. In addition to detecting and isolating a failed module or bus line, the error latch can be used for a failure within the clocking network. This is explored in the next section (along with other means of providing information to the processors on the performance of the clocking network).

4.3 Testing the Clocking Network

There are two different philosophies involved with implementing a test for the clock receivers within a module. In the first, each clock receiver in the module is tested and evaluated independently from the other receiver. With the second approach, one receiver is subjected to the simulated failure conditions while the other module receiver serves as a reference clock receiver to help evaluate the results.

4.3.1 The Independent Test

For the Independent Test, the test components described for the Delay Test are employed. Figure 4.10 shows the clock receivers and their associated test components that are contained within a module if the Independent Test is used. The performance of a clock receiver and its input clock signals is obtained by analyzing the information provided by the IDD and the ODD under the simulated failure conditions. The results from the tests are written into the latches contained in each Guardsman. The information contained within the latches are then read by the microprocessor or memory element and evaluated. The status report on the receiver and its input clock signals are then sent to the other members of the processor triad.

4.3.2 The Dual Test

The schematic outline of the Dual Test is shown in Fig. 4.11. The output of a clock receiver is used as the reference clock when the other module receiver is tested. The conditions of the DSC's required to subject each of the clock receivers to the simulated failure conditions (Table 3.6) are listed in Table 4.1. By delaying all three input signals to the receiver providing the reference clock, the advanced signal (AS) simulated failure conditions can be produced. The obvious advantage of the Dual Test over the Independent Test is that fewer test components are required.

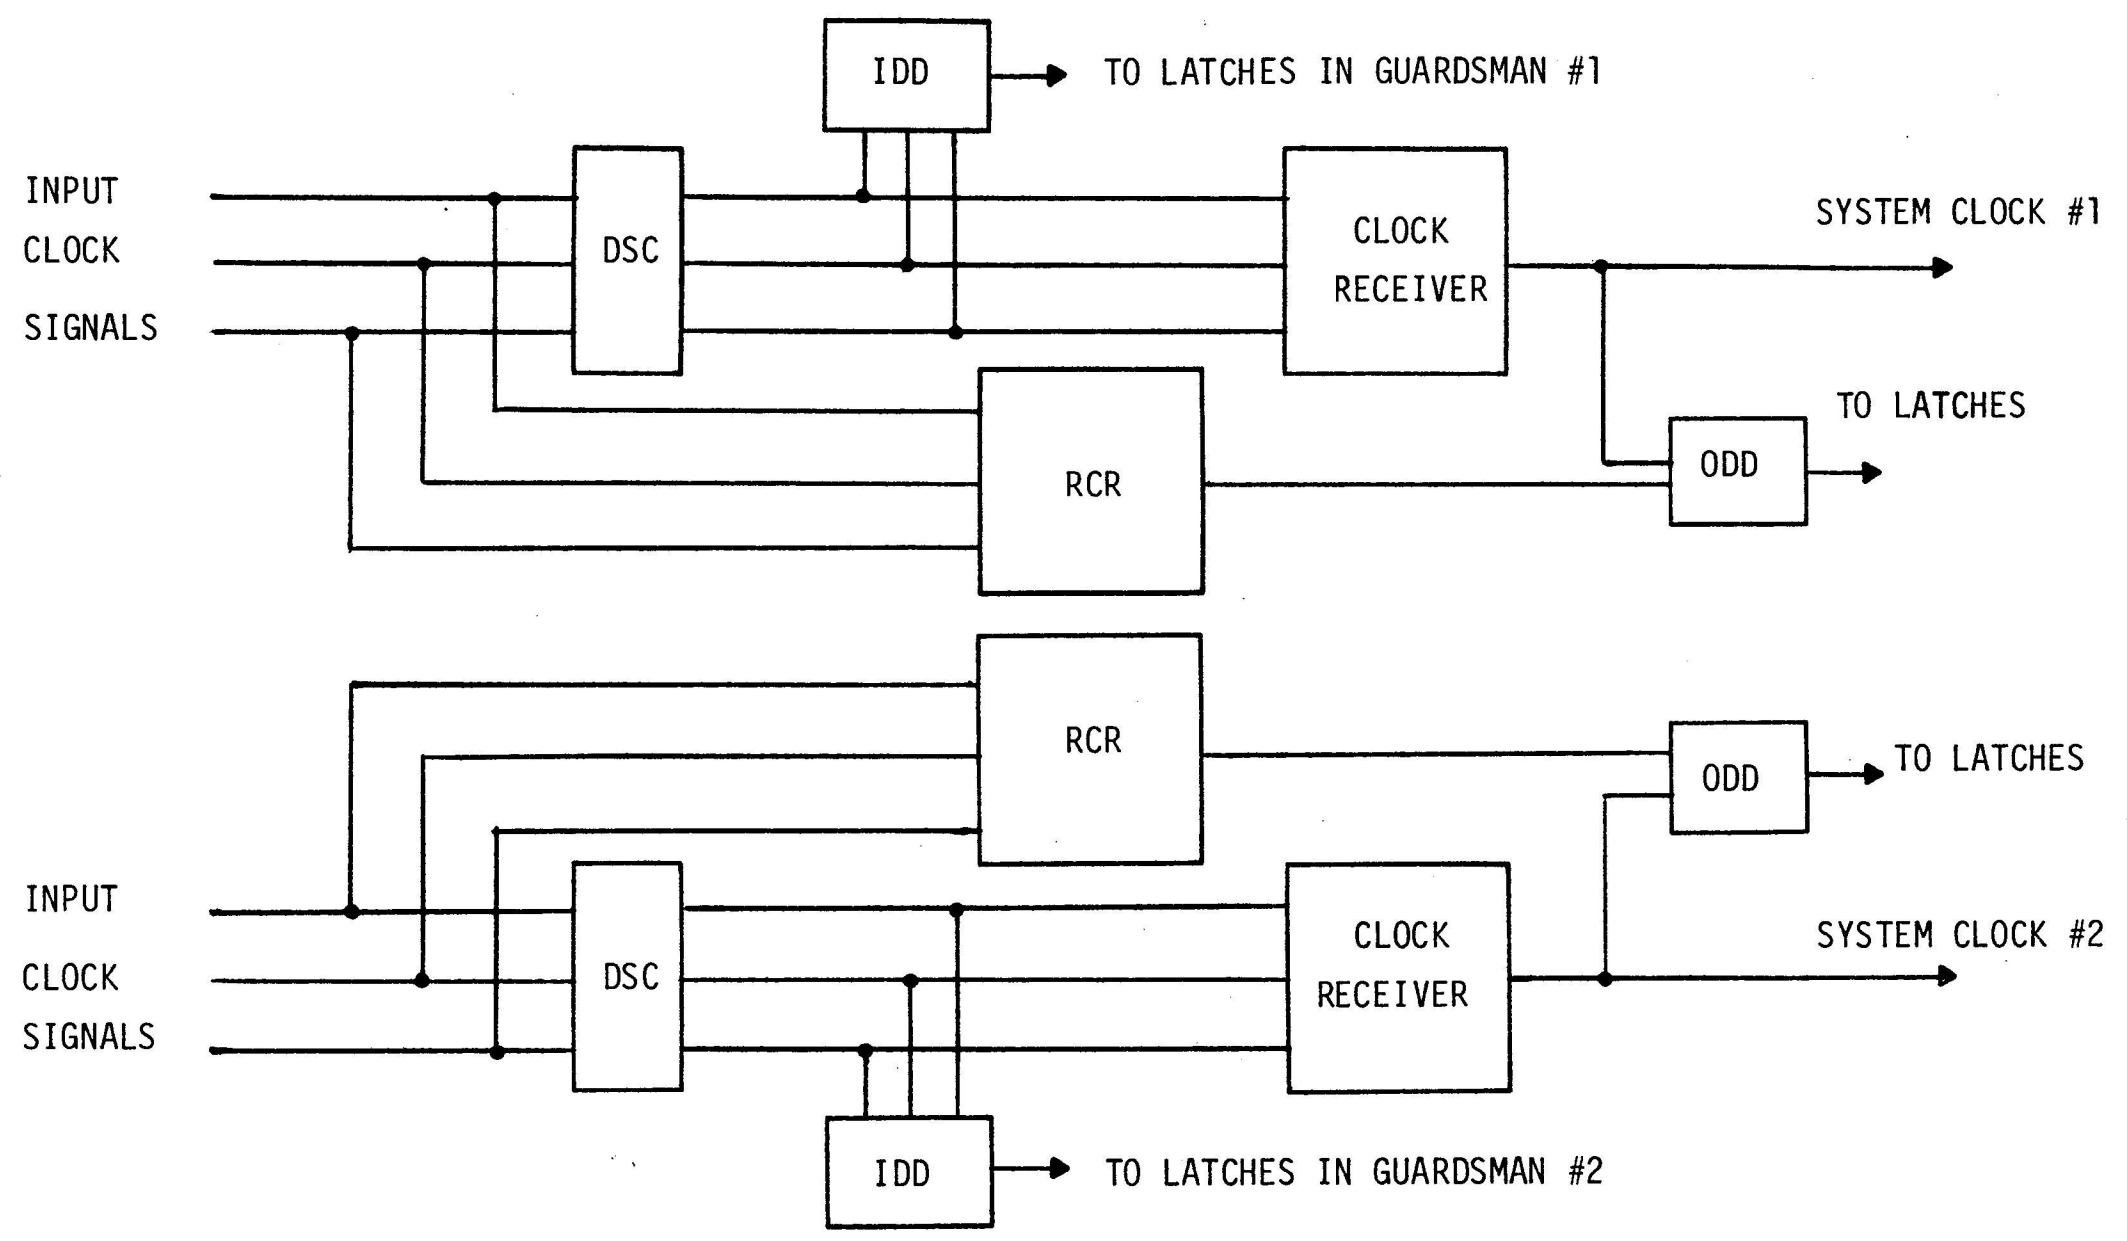


Figure 4.10 Independent Test.

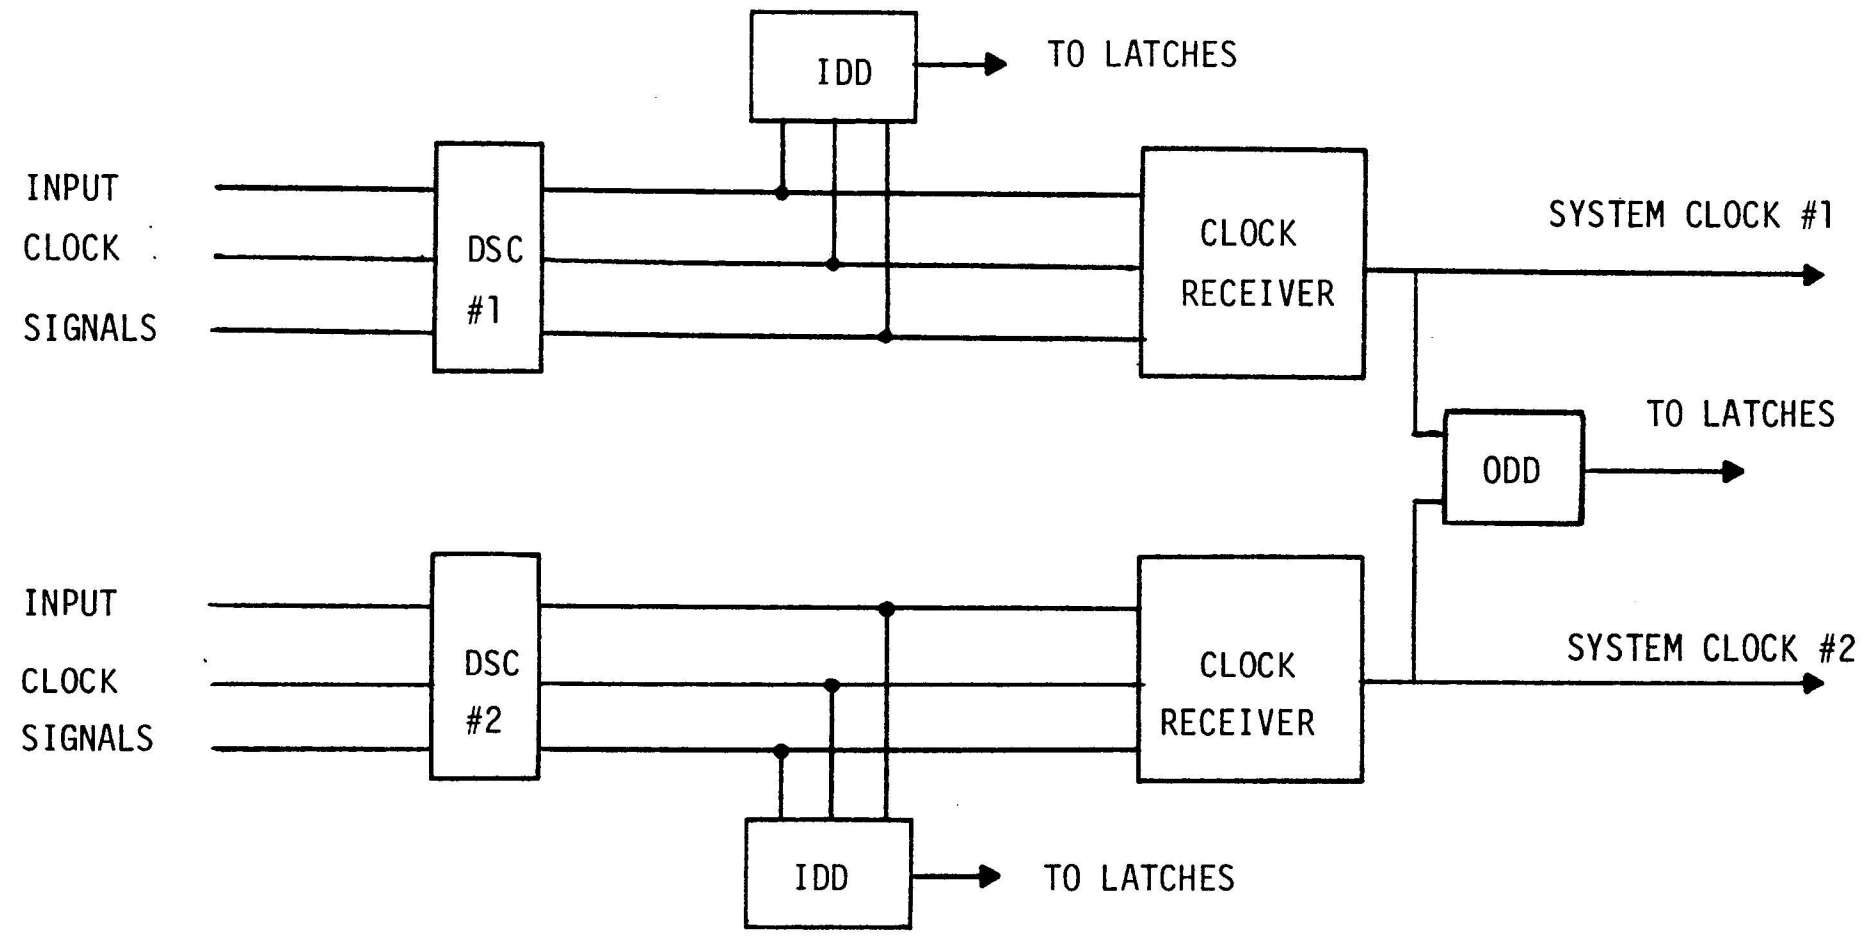


Figure 4.11 The Dual Test.

TABLE 4.1
THE DUAL TEST

SIMULATED FAILURE CONDITION	RECEIVER #1		RECEIVER #2	
	DSC #1	DSC #2	DSC #1	DSC #2
DS-0	DS-0	DS-0	DS-0	DS-0
DS-A	DS-A	DS-0	DS-0	DS-A
DS-B	DS-B	DS-0	DS-0	DS-B
DS-C	DS-C	DS-0	DS-0	DS-B
DS-AB	DS-AB	DS-0	DS-0	DS-AB
DS-AC	DS-AC	DS-0	DS-0	DS-AC
DS-BC	DS-BC	DS-0	DS-0	DS-BC
DS-ABC	DS-ABC	DS-0	DS-0	DS-ABC
AS-0	DS-ABC	DS-ABC	DS-ABC	DS-ABC
AS-A	DS-BC	DS-ABC	DS-ABC	DS-BC
AS-B	DS-AC	DS-ABC	DS-ABC	DS-AC
AS-C	DS-AB	DS-ABC	DS-ABC	DS-AB
AS-AB	DS-C	DS-ABC	DS-ABC	DS-C
AS-AC	DS-B	DS-ABC	DS-ABC	DS-B
AS-BC	DS-A	DS-ABC	DS-ABC	DS-A
AS-ABC	DS-0	DS-ABC	DS-ABC	DS-0

Another advantage is that the system clocks used by the two Guardsmen are compared directly with each other. This provides some insurance that the processor (or memory) and its two Guardsmen are synchronized, at least to within the tolerance of the ODD

The two receivers within a module can use either the same input clock signals or different combinations of signals, depending on the priorities of the network. For instance, if emphasis were placed on the toleration of primary clock failures (low quality clocks) to compensate for a short MTBF (Mean Time Between Failures) relative to the network's detection and correction abilities, the two receivers would use different combinations of primary clock signals with no more than one common signal. Two receivers using this arrangement can not be forced to fail by two (uncorrected) primary clock failures. This minimizes the probability that both system clocks fail in a similar manner which could allow the module containing the failed system clocks to gain access to the bus lines. However, this additional protection is at the cost of a reduction in the quality of the ODD's performance. The lower tolerance limit of the ODD is used only when the variations of the parameters of the input clock signals can be ignored. Thus, the lower limit can not be used with the Dual Test if the two receivers are driven by different clock signals. In the Independent Test, the clock receiver and the reference clock receiver are driven by the same input clock signals, allowing the lower ODD limit to be used under the working condition (DS-0). This is at the cost of an increase in the number of test components over that required by the Dual Test.

If the quality of the primary clocks and the correction process of the network are sufficient to insure an acceptable probability that no two uncorrected primary clock failures will occur, the two receivers can be driven by the same signals. In this case, the Dual Test has an advantage over the Individual Test since it requires fewer test components for the same testing effectiveness.

4.4 Clock Frequency

The use of medium-speed TTL technology for the construction of a clocking network for a multiprocessor is highly inefficient in terms of the potential operating speed of current processors. The frequency of the system clock produced by the receiver described in this thesis is limited by the propagation delays encountered within the clock receiver and the delay time required to produce a simulated failure. The frequency of a system clock within a network constructed using TTL component is on the order of 1 MHz, whereas the potential processor rate is at least an order of magnitude greater⁴.

Several methods can be employed to increase the operational speed of the clocking network. These include the use of high-speed logic components such as Emitter

Coupled Logic, Burst Generator and Frequency Multiplier. The propagation delays within the receivers and the delay time required for testing can be reduced through the use of high-speed components to construct the network (primary clocks, receivers, test and replace components). For instance, MECL III components have a propagation time at least an order of magnitude lower than the corresponding TTL components.

The CERBERUS system (Figures 1.13 and 1.7b) made use of a BURST GENERATOR to increase the system's operating speed. The purpose of the BURST GENERATOR was to produce a pulse train containing a given number of pulses with approximately known widths at the appropriate time. In CERBERUS, the pulse train was initiated by the arrival of a leading edge within the system clock. The pulse train ended before the arrival of the next leading edge. This allowed CERBERUS to operate at a higher frequency and still maintain some synchronization within the system.

There are several ways of increasing the speed of the network through frequency multiplication. One of the most promising methods is through the use of Phase-Locked Loop (PLL) circuitry⁴. Figure 4.12 shows a design for a PLL multiplier. The PLL multiplier reaches steady state when the frequency of the system clock is equal to f . This results in the phase difference between the two remaining at a constant value (V_p) which, in turn, keeps the voltage input to the Voltage Controlled Output (VCO) constant. The frequency of the output depends on the frequency divider. For example, a divide-by 4 causes the output of the PLL Multiplier to maintain a frequency (steady-state) four times greater than its input, the system clock. The parameters of the filter determines the response of the Multiplier to a change in the input signal. Thus, precautions must be taken to assure that the multiplier returns to steady state after the clock receiver has been subjected to a simulated failure condition.

4.5 Recommendations

The following recommendations are made for the construction of a clocking network for a multiprocessor system such as CARDS.

The choice of the number of primary clocks required by the network and the testing arrangement (Independent or Dual) should depend on the quality of the primary clocks and the procedure by which a failed input signal is replaced. If primary clock quality is sufficiently high then the Dual Test can be used with a smaller number of primary clocks. If this is not the case, then the Independent Test should be used with a larger number of primary clocks (same operation period). The numbers of possible combinations of input clock signals given four, five, six, seven and eight primary clocks, respectively, are listed in Table 4.2. Also included is the number of combinations with which a given combination contains no more than one common clock signal. The number of primary clocks employed by the network must

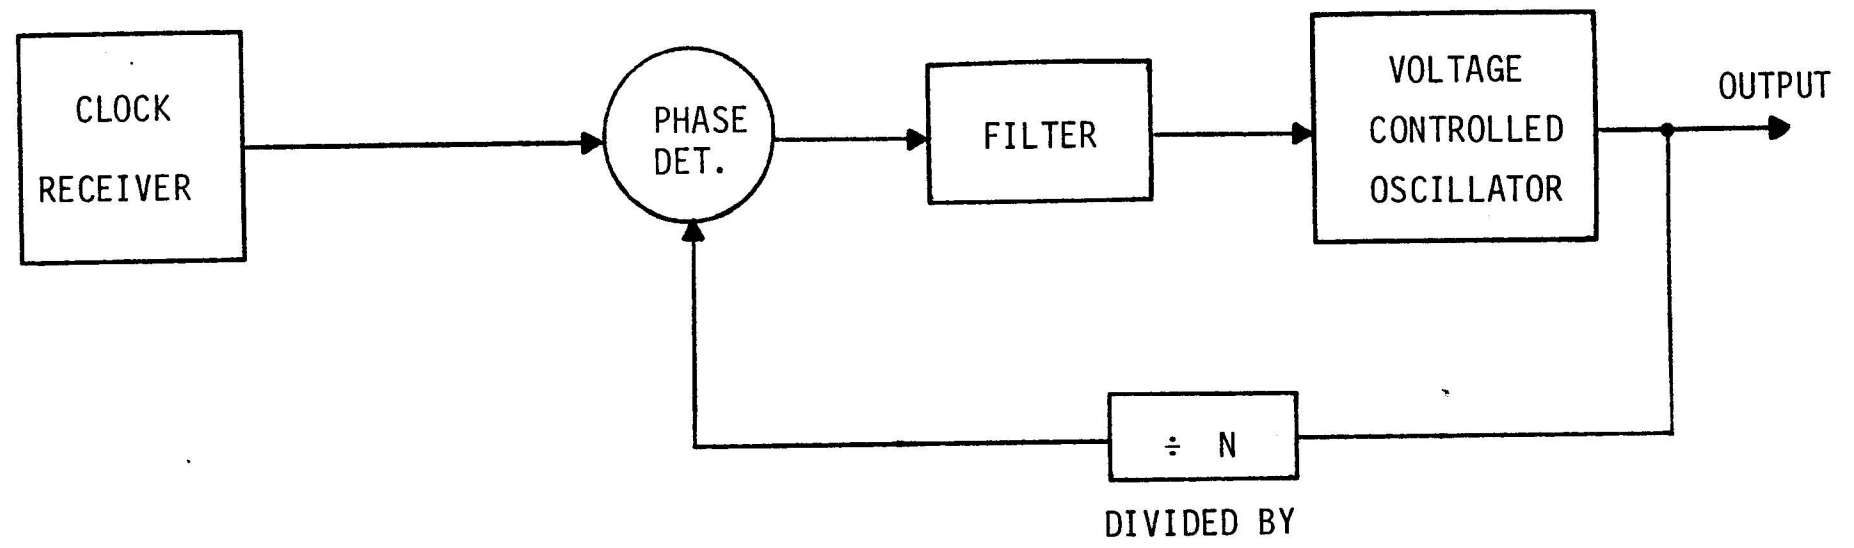


Figure 4.12 Phase-Locked Loop Multiplier.

TABLE 4.2

PROTECTION OF A MODULE'S RECEIVERS FROM
TWO SIGNAL FAILURES

NUMBER OF UNFAILED PRIMARY CLOCKS	INPUT CLOCK COMBINATIONS	COMBINATIONS WITH NO MORE THAN ONE COMMON CLOCK
FOUR	FOUR	NONE
FIVE	TEN	THREE
SIX	TWENTY	TEN
SEVEN	THIRTY-FIVE	TWENTY-TWO
EIGHT	FIFTY-SIX	FORTY

exceed four to permit the two receivers within a module to operate with no more than one primary clock in common.

The last four simulated failure conditions listed in Table 4.1 (AS-AB to ABC) can be eliminated to reduce the amount of time spent on testing the clocking network. No useful information concerning the performance of the IDD and the ODD are provided by these conditions. In addition, their role in detecting and isolating failures are minor since they essentially duplicate the DS-C, DS-B, DS-A and DS-0 conditions respectively.

The use of MIDD's (modified IDD's) through the system to monitor the behavior of the primary clocks can be used to supplement and to verify the results obtained from the testing of the receivers. The status of the stand-by clock signals should be made known to all the processor triads. This can be done by transmitting the results obtained by a processor triad to all the processors within the system.

4.6 Conclusions

This thesis has attempted to outline the basic concepts involved in the design of a highly-reliable clocking network. The use of hybrid redundancy to achieve fault tolerance was investigated with respect to its applicability to such a network. It was determined that a Triple Modular Redundant arrangement along with the capacity to detect, isolate and replace a fault was the most practical approach provided that the replicated units (the input clocks) and the method of testing were sufficiently reliable. It was found that a single-fault-tolerant system clock could be produced by the clock receiver described in Chapter 2 using only three input clock signals. This was accomplished by polling the edges of the clock signals instead of their logic levels, and without high frequency (noise) filters. The clock receiver was designed to use feedback to limit the time during which a signal can vote and to prevent a signal from voting more than once during a polling interval. The receiver was not only capable of tolerating the failure of an input signal but also of surviving several types of faults within itself.

A method of testing was devised that made use of a series of simulated failure conditions to expose faults within the clocking network. Simulated failures produced by delaying the input clock signals was explored in Chapter 3 to determine its practicality. The components needed to generate the failure conditions and to note the response of the network were designed using relatively uncomplicated circuits. It was shown that during the testing of the clocking network, the performance of these testing components could be evaluated. This resulted from the fact that these components were also stimulated during the testing sequence.

A synchronous system is served by a clocking network in a manner analogous to the function performed by the heart and cardiovascular system for the human body. The modules of a system deprived of a system clock cease to function in a fashion similar to the effect upon cells in a human body caused by blood deprivation. This thesis has strived to ensure that a valid system clock can be delivered to these modules in a reliable fashion. A clock network employing concepts developed in this thesis has been investigated in terms of its applicability to an existing fault-tolerant computational system (CARDS). It is felt that this thesis can serve as a guide to the implementation of a fault-tolerant clocking network in a fault-tolerant digital system or any system in which a highly-reliable clock is in demand.

REFERENCES

1. Daly, W.M., A.L. Hopkins, and J.F. McKenna, "A Fault-Tolerant Digital Clocking System", 1973 International Symposium on Fault-Tolerant Computers, June, 1973.
2. Daly, W.M., and McKenna, J.F., M.I.T. C.S. Draper Laboratory, Digital Development Memo #627, August, 1971.
3. Hopkins, A.L., "Computer Control for Manned and Automated Space Vehicles", C.S. Draper Laboratory Report E-2756, Cambridge, Massachusetts, March, 1973.
4. Krauss, H.R., "Clocking and Synchronization within a Fault-Tolerant Multiprocessor", C.S. Draper Laboratory Report T-564, Cambridge, Massachusetts, June, 1972.
5. Lala, J.H., "A Cost and Reliability Model of Partitioned Digital Systems", C.S. Draper Laboratory Report T-573, Cambridge, Massachusetts, February, 1973, pp. 19-32.
6. Smith, T.B., "A Highly Modular Fault-Tolerant Computer System", C.S. Draper Laboratory Report T-595, Cambridge, Massachusetts, November, 1973.