




(86) Date de dépôt PCT/PCT Filing Date: 2002/02/25
 (87) Date publication PCT/PCT Publication Date: 2002/09/06
 (85) Entrée phase nationale/National Entry: 2003/07/25
 (86) N° demande PCT/PCT Application No.: US 2002/005568
 (87) N° publication PCT/PCT Publication No.: 2002/069552
 (30) Priorité/Priority: 2001/02/24 (60/271,124) US

(51) Cl.Int.⁷/Int.Cl.⁷ H04L 7/00, G06F 12/00, H04L 7/02,
H03D 3/24, H04B 10/08
 (71) Demandeur/Applicant:
INTERNATIONAL BUSINESS MACHINES
CORPORATION, US
 (72) Inventeurs/Inventors:
BARRETT, WAYNE M., US;
CHEN, DONG, US;
COTEUS, PAUL W., US;
GARA, ALAN G., US;
JACKSON, RORY D., US;
KOPCSAY, GERARD V., US;
...
 (74) Agent: ROSEN, ARNOLD

(54) Titre : TECHNIQUE DE CAPTURE DE DONNEES POUR SIGNALISATION HAUTE VITESSE
 (54) Title: DATA CAPTURE TECHNIQUE FOR HIGH SPEED SIGNALING



(57) Abrégé/Abstract:
 A data capture technique for high speed signaling to allow for optimal sampling of an asynchronous data stream. This technique

(72) **Inventeurs(suite)/Inventors(continued)**: NATHANSON, BEN J., US; TAKKEN, TODD E., US; VRANAS, PAVLOS M., US

(57) **Abrégé(suite)/Abstract(continued)**:

allows for extremely high data rates and does not require that a clock be sent with the data as is done in source synchronous systems. The present invention also provides a hardware mechanism for automatically adjusting transmission delays (Fig. 5) for optimal two-bit simultaneous bi-directional (SiBiDi) signaling.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 September 2002 (06.09.2002)

PCT

(10) International Publication Number
WO 02/069552 A1

(51) International Patent Classification⁷: H04L 7/00, 7/02, H03D 3/24, H04B 10/08, G06F 12/00

(21) International Application Number: PCT/US02/05568

(22) International Filing Date: 25 February 2002 (25.02.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/271,124 24 February 2001 (24.02.2001) US

(71) Applicants and

(72) Inventors: **BARRETT, Wayne, M.** [US/US]; 5522 Silas Dend Road, N.W., Rochester, MN 55901 (US). **CHEN, Dong** [CN/US]; 15 Scenic Drive, Apt. K, Croton-On-Hudson, NY 10520 (US). **COTEUS, Paul, W.** [US/US];


2742 Quinlan Street, Yorktown Heights, NY 10598 (US). **GARA, Alan, G.** [US/US]; 38 Marion Avenue, Mount Kisco, NY 10549 (US). **JACKSON, Rory, D.** [US/US]; 17 Lorraine Drive, Eastchester, NY 10709 (US). **KOPCSAY, Gerard, V.** [US/US]; 2950 Curry Street, Yorktown Heights, NY 10598 (US). **NATHANSON, Ben, J.** [US/US]; 701 Linden Avenue, Teaneck, NJ 07666 (US). **TAKKEN, Todd, E.** [US/US]; 118 Foxwood Circle, Mount Kisco, NY 10549 (US). **VRANAS, Pavlos, M.** [GR/US]; 30-A Bedford Avenue, Bedford Hills, NY 10507 (US).

(74) Agents: **GROLZ, Edward, W.** et al.; Scully, Scott, Murphy & Presser, 400 Garden City Plaza, Garden City, New York, NY 11530 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,

[Continued on next page]

(54) Title: DATA CAPTURE TECHNIQUE FOR HIGH SPEED SIGNALING



(57) Abstract: A data capture technique for high speed signaling to allow for optimal sampling of an asynchronous data stream. This technique allows for extremely high data rates and does not require that a clock be sent with the data as is done in source synchronous systems. The present invention also provides a hardware mechanism for automatically adjusting transmission delays (Fig. 5) for optimal two-bit simultaneous bi-directional (SiBiDi) signaling.

WO 02/069552 A1

WO 02/069552 A1

GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

DATA CAPTURE TECHNIQUE FOR HIGH SPEED SIGNALING

5

CROSS-REFERENCE

The present invention claims the benefit of commonly-owned, co-pending United States Provisional Patent Application Serial Number 60/271,124 filed February 24, 2001 entitled MASSIVELY PARALLEL SUPERCOMPUTER, the whole contents and disclosure of which is expressly incorporated by reference herein as if fully set forth herein. This patent application is additionally related to the following commonly-owned, co-pending United States Patent Applications filed on even date herewith, the entire contents and disclosure of each of which is expressly incorporated by reference herein as if fully set forth herein. U.S. patent application Serial No. (YOR920020027US1, YOR920020044US1 (15270)), for "Class Networking Routing"; U.S. patent application Serial No. (YOR920020028US1 (15271)), for "A Global Tree Network for Computing Structures"; U.S. patent application Serial No. (YOR920020029US1 (15272)), for "Global Interrupt and Barrier Networks"; U.S. patent application Serial No. (YOR920020030US1 (15273)), for "Optimized Scalable Network Switch"; U.S. patent application Serial No. (YOR920020031US1, YOR920020032US1 (15258)), for "Arithmetic Functions in Torus and Tree Networks"; U.S. patent application Serial No. (YOR920020033US1, YOR920020034US1 (15259)), for "Data Capture Technique for High Speed Signaling"; U.S. patent application Serial No. (YOR920020035US1 (15260)), for "Managing Coherence Via Put/Get Windows"; U.S. patent application Serial No. (YOR920020036US1, YOR920020037US1 (15261)), for "Low Latency Memory Access And Synchronization"; U.S. patent application Serial No. (YOR920020038US1 (15276)), for "Twin-Tailed Fail-Over for Fileservers Maintaining Full Performance in the Presence of Failure"; U.S. patent application Serial No. (YOR920020039US1 (15277)), for "Fault Isolation Through No-Overhead Link Level Checksums"; U.S. patent application Serial No. (YOR920020040US1 (15278)), for "Ethernet Addressing Via Physical Location for Massively Parallel Systems"; U.S. patent application Serial No. (YOR920020041US1 (15274)), for "Fault Tolerance in a Supercomputer Through Dynamic Repartitioning"; U.S. patent application Serial No. (YOR920020042US1

35

(15279)), for "Checkpointing Filesystem"; U.S. patent application Serial No. (YOR920020043US1 (15262)), for "Efficient Implementation of Multidimensional Fast Fourier Transform on a Distributed-Memory Parallel Multi-Node Computer"; U.S. patent application Serial No. (YOR9-20010211US2 (15275)), for "A Novel
5 Massively Parallel Supercomputer"; and U.S. patent application Serial No. (YOR920020045US1 (15263)), for "Smart Fan Modules and System".

BACKGROUND OF THE INVENTION

1. Field of the Invention

10

The present invention relates generally to a data capture technique for high speed signaling, and more particularly pertains to a technique to allow for optimal sampling of an asynchronous data stream. This technique allows for extremely high data rates and does not require that a clock be sent with the data as is done in source
15 synchronous systems.

The present invention also provides a hardware mechanism for automatically adjusting transmission delays for optimal two-bit Simultaneous Bi-Directional (SiBiDi) signaling.

20

2. Discussion of the Prior Art

A large class of important computations can be performed by massively parallel computer systems. Such systems consist of many identical compute nodes, each of
25 which typically consist of one or more CPUs, memory, and one or more network interfaces to connect it with other nodes.

The computer described in related U.S. provisional application Serial No. 60/271,124, filed February 24, 2001, for A Massively Parallel Supercomputer,
30 leverages system-on-a-chip (SOC) technology to create a scalable cost-efficient computing system with high throughput. SOC technology has made it feasible to build an entire multiprocessor node on a single chip using libraries of embedded components, including CPU cores with integrated, first-level caches. Such

packaging greatly reduces the components count of a node, allowing for the creation of a reliable, large-scale machine.

The present invention relates to the field of massively parallel computers used for various applications such as, for example, applications in the field of life sciences. More specifically, this invention relates to the field of high speed signaling, to either unidirectional signaling or Simultaneous BiDirectional (SiBiDi) signaling.

There are cases where large data transfers are required but the number of wires that can be used is limited. Simultaneous Bidirectional (SiBiDi) signaling allows the simultaneous transmission and reception of signals using the same wire. This reduces the number of wires by a factor of two. An example where large data transfers are needed but where the number of cables is severely constrained is a large parallel super computer with thousands of processors communicating through wires.

SiBiDi signaling operates by sending data on the same wire as it receives data. Therefore during reception one receives not only the desired data sent from the other end of the wire but also the data that one has just transmitted. Of course this corrupts the desired signal. However, since the data that was just transmitted is known one can "subtract it out". This is done by standard SiBiDi circuitry.

SUMMARY OF THE INVENTION

Accordingly, it is a primary object of the present invention to provide a data capture technique for high speed signaling, particularly to allow optimal sampling and capture of an asynchronous data stream without sending a clock signal with the data stream. The data is captured by sending serial bits of the data stream down a clocked delay line with a series of delay taps, and sampling all of the delay taps with a clock. Each delay tap output is compared with a neighbor delay tap output to determine if it is the same, and the comparisons are used to form a clocked string to generate a data history record which is examined to determine optimal data capture eyes by looking for data capture eyes where the data does not transition between adjacent delay taps, which are detected as optimal data capture eyes.

A further object of the subject invention is the provision of a hardware mechanism for automatically adjusting the transmission delays for optimal two-bit SiBiDi signaling to improve the signal quality of the two-bit SiBiDi signaling. A special hardware algorithm is implemented and each of the two bits is used in unidirectional channels in order to allow the hardware algorithms of the two nodes to safely exchange setting parameters during the set-up sequence. A unidirectional channel of the same frequency has half the bandwidth of the SiBiDi channel but it has considerably better signal quality.

10 **BRIEF DESCRIPTION OF THE DRAWINGS**

The foregoing objects and advantages of the present invention for a data capture technique for high speed signaling may be more readily understood by one skilled in the art with reference being had to the following detailed description of several embodiments thereof, taken in conjunction with the accompanying drawings wherein like elements are designated by identical reference numerals throughout the several views, and in which:

Figure 1 illustrates a data receive macro that can capture serial data at a 2Gbit rate and bring it into the local clock domain.

20

Figure 2 is a data send macro block which receives a data input 1 byte wide at 500 Mhz, and produces a data output of 2 data streams, each being 2 Gb/s serial data, which is transmitted over a SiBiDi (Simultaneous Bi Directional) differential data link, and is then the data input to the data capture macro of Figure 1.

25

Figure 3 illustrates an instage, 2-bit macro block.

Figure 4 illustrates an implementation of a clocked delay line wherein serial data passes through a combinatorial series of inverters, each of which adds an increment of delay.

30

Figure 5 illustrates a history logic block which is shown as an extension of the clock delay line, and shows one clock phase.

Figures 6A and 6B illustrate data and history sample MUXs, each of which corresponds to the MUX shown at the bottom of Figure 3, one for each of the leading and falling edge clock phases, while the MUX of Figure 6C receives the output signals at the bottom of Figure 5.

5

Figure 7 shows serial bit combining and byte align logic, to be utilized as a paired link capable of capturing 4 data bits per clock cycle, which combines two 2-bit macros (as shown in Figure 3) and finds the proper byte alignment between the 2 input data streams.

10

Figure 8 illustrates the eye detection process in an eye detection flow diagram,

Figure 9 illustrates a two-bit macro state diagram 1 of 2 which shows the state flow for the phase during which the eye position sample points are being determined.

15

Figure 10 illustrates a two bit macro state diagram 2 of 2 which shows the repetitive state flow during normal data capture operations.

20

Figure 11 illustrates a first embodiment wherein two differential data lines connect a pair of identical nodes 1, 2, and wherein each node has a unique ID, and each node operates with a 2-bit sender CPU and a 2-bit capture CPU.

25

Figure 12 illustrates a second embodiment of SiBiDi electrical communications between two nodes, Node 1, Node 2, wherein a single differential communication line connects the nodes.

The state machines of Figures 13A and 13B illustrate the steps taken by the node compute chip in the training of a synchronous Si-Bi-Di connection.

30

DETAILED DESCRIPTION OF THE INVENTION

Overview

The present invention is designed to be employed in implementing interconnections in a massively parallel supercomputer which solves two longstanding problems in

the computer industry; (1) the increasing distance, measured in clock cycles, between the processors and the memory and (2) the high power density of parallel computers built of mainstream uni-processors or symmetric multi-processors.

- 5 The present invention relates generally to a data capture technique for high speed signaling, and more particularly pertains to a technique to allow for optimal sampling of an asynchronous data stream. This technique allows for extremely high data rates and does not require that a clock be sent with the data as is done in source synchronous systems.

10

Serial Link Investigations

The target bandwidth for serial links connecting nodes of the massively parallel supercomputer is 1.4 Gb/s (each direction). This bandwidth must be bi-directional.

- 15 The bi-directional requirement can be handled in a number of ways. All cases share the constraint that they be low power and low cost. The implementation of choice will be integrated into an ASIC within a processing node. A particular challenge associated with this approach is the low power constraint. This coupled with the lack of relative phase information for the link transmission eliminates standard PLL clock and data recovery designs. In this case the phase must be extracted from the data
20 itself with high reliability without the use of a PLL.

Digital data capture

25 Overview

- This specification describes in detail a digital data capture technique. Figure 1 illustrates a data receive macro that can capture serial data at a 2Gbit rate and bring it into the local clock domain. The Goal is to do this reliably with low power
30 utilizing a small number of cells. Figure 2 illustrates a send macro block that is considerable simpler than the data receive macro. It will be described in the second section of this specification. This describes a DDR (double data rate) style data recovery that allows for an internal clock that is half the frequency of the bit time. This can be utilized in a SDR (single data rate) mode or extended to a quad data rate
35 scheme if desired.

Referring to Figure 1, the data input to the data receive macro are two data streams of 2 Gb/s input serial data, which represents a total data stream of 4 Gb/s, and the data output is one data stream, a byte wide (8 bits wide) at 500 Mhz. The other input signals are a 1 Ghz clock, a reset signal which resets the data capture macro to a known state, a train signal directing the macro to find optimal eyes (positions or stages along the multiple tap delay line (see Figures 3, 4 and 5) at which data is not undergoing transitions and therefore are likely to be the most accurate data capture positions least likely to have data errors) to recover data, an Idle bytes signal which is a predetermined idle data pattern which is received by the macro when data is not being received, a DDR mode which directs the macro to operate in a double data rate mode, and a Minimum distance which is a constraint parameter to find the optimal data or idle eye. The other output signals include a Valid Idle signal indicating a valid receipt of an idle pattern, an Eye_found signal which indicates that the optimal eye positions and parameters have been detected, a locked signal indicating that the optimal eye position is locked, and a Warning signal indicating that the optimal eye position is in danger of being lost, by being too close to one end of a multiple tap digital delay line.

The macro of Figure 1 is explained further with reference to Figures 3, 4, 5, 6 and 7.

The latency in the receive macro is between 7 and 12 bit times depending on the byte phase of the data. One can reduce the latency to 5 to 6 bit times by skipping the byte output. This is a reasonable approach for signal redriving where data content can be ignored.

Figure 2 is a data send macro block which receives a data input 1 byte wide at 500 Mhz, and produces a data output of 2 data streams, each being 2 Gb/s serial data, which is transmitted over a unidirectional or SiBiDi (Simultaneous Bi Directional) differential data link as described below, and is then the data input to the data capture macro of Figure 1. A further input is a Byte valid signal which indicates a valid data signal is being received and is to be sent rather than an idle signal, and further inputs are a 1 GHz clock signal and a reset signal which resets the data send macro to a known state.

Figures 3, 4 and 5 illustrate the data receive and capture. The data is captured by sending the data bits down a fast tapped delay line (see Figures 3, 4, 5) and sampling all the taps with the local clock. Each tap is compared with its neighbor (see Figures 3, 5) to see if it is the same. The aggregate of these comparisons form a clocked string that is combined with previous clocked strings to generate a history that can be used to determine the optimal sampling points. The optimal sampling points can be found from the history string (see Figure 5, Registers A, B, C, D) by looking for the regions where the data does not ever change between delay taps, which are referred to herein as "eyes". The history is periodically updated such as every local clock. The periodic update compensates for changing parameters, such as changes in the temperatures or voltages of different components. There will also be three additional "eye" pipelined registers (see Figure 5, Registers B, C, D) that are infrequently updated. This allows one to develop a capture scheme which has a programmable persistence period as well as being immune to isolated bit errors. The persistence time can be set arbitrarily long but must not be shorter than the maximum time necessary to reliably sample data edges. To accommodate bit sample times faster than the local clock period, both edges of the clock are used to capture DDR data (see Figures 3, 4). Each edge of the clock has its own associated capture registers and independent logic to find the optimal eye. This technique is therefore largely immune to asymmetries in the local and sending side clock duty cycles.

As the history registers will change, the optimal sampling point will also move. This updating should be done on a time scale shorter than the persistence time. This function is done in the histogram and sampling point determination unit.

This method of data capture involves a two stage initialization which proceeds after either a system reset or a separate "train" signal is asserted.

- Stage1: After reset or "train" signals (see Figures 1, 2), the history registers are flushed and a new history pattern is acquired in all 3 "eye" registers.

After acquiring a valid set of "eye" registers, the best sampling point is determined through a state machine sequence (see Figures 9, 10). This is done for each phase of the clock independently. These sampling points then are used and the two bits are forwarded to the next stage every system clock.

- Stage 2: The two bits are received and inserted into a shift register, and this shift register is used along with a barrel shifter to allow for appropriate nibble (1/2 byte or 4 bits) boundaries (see Figure 7). The boundaries are found through the use of unique idle nibble patterns during the initialization sequence.

The clocked delay line block:

Figures 4 and 5 illustrate an implementation of a clocked delay line wherein serial data enters the left hand inverter I, and passes through a combinatorial series of inverters, each of which adds an increment of delay. At each inverter output there are two register latches FF, Figures 3-5, one clocked by the positive edge of the clock, and the other by the negative edge. This allows the logic to capture data at twice the clock rate. One bank of latches captures the data eye for the positive clock phase while the other bank of latches captures the data eye for the negative clock phase. Both eyes are separately detected and sampled, such that each clock phase requires a separate circuit as shown in Figure 5. The independent positive and negative clocked logic circuits result in very little dependence on the duty cycle of the clock signal, particularly asymmetries in the local and sending side clock duty cycles.

This module has as its input the high-speed signal after the input receiver. The only other input to this module is the local clock that is fanned out equal time to all the flip-flops. The only outputs of this module are N+1 clocked delay taps. D[0:N]. Each tap is to be approximately 50ps with relatively good matching between rising and falling edges. The matching required between the falling delay versus the rising delay is approximately 20-30%. We require the clocks to be equal time to all neighboring latches to within ~10ps. This may be better achieved with a tapped clock line rather than a clock tree. Many of these data capture circuits may be implemented so power is critical.

This module is layout critical and therefore requires extra layout consideration.

For test chip purposes, the number of elements is fixed at 32. This gives a nominal total delay of approximately 1.6nsec, which is enough to capture DDR data at frequencies down to approximately 1Gb/s.

5 Referring to Figures 3 and 5, and particularly Figure 5, the output of each register FF (flip flop) is directed to an exclusive OR gate XOR which also receives an input from the next register FF in the delay line. Referring to the first and second stages of the delay line, since the data bit is inverted by the second inverter before it enters the second register FF, if the data bit does not undergo a transition between the two
10 consecutive stages, the first and second registers will hold opposite values, such that the first stage XOR will produce a 1, indicating there was no transition between the stages. Conversely, if the data bit undergoes a transition between the two consecutive stages, the first and second registers will hold the same value, such that the first stage XOR will produce a 0, indicating there was a transition between the
15 stages.

The system of Figure 5 is searching for a stable position or eye along the clocked delay line to detect the data whereat the data does not undergo transitions, which is indicated by a series of 1 outputs from a series of consecutive XORs, such that the
20 data detection eye should be aligned to the middle of a series of 1s. The output of each XOR is input to an AND gate, the output of which is input to a Register A, which is the first of a series of FF (flip flop) history registers A, B, C, and D. The first register A is sampled at the full 1 Ghz clock rate, and is periodically reset to high by a Set to high signal at a relatively slow clock rate e.g. > 1 millisecond (ms)
25 whereas the registers B, C, and D are sampled and updated at the same clock logic rate as Set to high.

The register A is set or reset to a high 1 output by the clock at a >1 ms clock rate, and after a reset if the output of the XOR is a 1, then the output of the AND gate is a
30 1, and the output of the register A is a 1 which is subsequently clocked (by an Update signal to the load (ld) input of the registers B, C and D) serially through the registers B, C and D. Conversely, if the output of the XOR is a 0, and the output of the Register A is set or reset to 1, then the output of the AND gate is a 0, and register

A outputs a 0, which is subsequently clocked serially through the registers B, C and D. The arrangement is such that once the output of register A is a 0, it remains a 0 until the register A is reset by the Set to high signal, such that the outputs of each of registers B, C and D are serially clocked to 0 and remain at 0 until the Register A is
5 reset to a 1 by Set to high signal.

The outputs of each of the history registers B, C and D are input to a 2 of 3 logic element which produces a 1 or high (H) output if any 2 of its 3 inputs are 1s. The purpose of the 2 of 3 logic is to compensate for glitches in the data stream through
10 the digital delay line which might erroneously cause 1 of the 3 inputs to be a 0, such that an accurate output is produced in spite of data glitches. Moreover, the occurrences of a 2 of 3 logic detection can be counted and reported as an indication of the integrity of the data being received. The H outputs (0 to N-1) are inputs to the MUX in Figure 6C as indicated therein. In general a string of 1s in the H
15 outputs indicates a good candidate for a data sampling eye which should be centered in the middle of the string of 1s.

Figure 3 illustrates the instage, 2-bit macro block, most of which has been explained with respect to Figure 5. The Histogram and sampling point determination unit
20 includes the AND gates, history registers and 2 of 3 logics of Figure 5 and the MUX of Figure 6C and the State Diagrams of Figures 9 and 10. The Macro block also includes an MUX which receives as inputs all of the outputs of the FF registers of the digital delay line, and selectively passes those inputs as Data out under the control of the signals H[0], H[1], H[N-1] at the bottom of Figure 5.

25

The history block:

Figure 5 illustrates a history logic block which is shown as an extension of the clock delay line block, and shows one clock phase. An identical circuit is required for the
30 other clock phase. The inverter string shown Figure 5 is common to both clock phases.

Sampling Point Block:

The sampling point block is most easily described by a state diagram that determines the two optimal sampling points, one for each clock phase. As the sampling points will not be updated frequently (at least 50 clocks between updates), we can use a
5 multiple clock process to find the optimal sampling points.

Figures 6A and B illustrate data and history sample MUXs, with the sampling MUXs having inputs of respectively $D_even[N:0]$ and $D_odd[N:0]$, each of which correspond the MUX shown at the bottom of Figure 3, one for each of the leading
10 and falling edge clock phases, while the MUX of Figure 6C receives the output signals at the bottom of Figure 5 as explained previously.

Combining two 2-bit Macros

Figure 7 shows serial bit combining and byte align logic, to be utilized as a paired
15 link capable of capturing 4 data bits per clock cycle, which combines two 2-bit macros (as shown in Figure 3) and find the proper byte alignment between the 2 input data streams, each at 2 Ghz, input to the two 2-Bit Macros 70 (as illustrated in Figure 3). The two input signals can be considered to be the input signals of Figure 1 which are being combined with the proper byte alignment. The outputs of the two
20 2-bit macros of Figure 3 are two data streams, each at 1 Ghz, which are input to a Register 71 which delays and standardizes the 2 data streams, which are input to a 12 bit Shift Register 72 which is clocked at half speed $Clk/2$, which converts the 2 data streams to a 12 bit wide data stream at 500Mhz. These are input to a logic 73 and a Barrel Shifter 74 which has a 24 bit input of the two 12 bit wide data streams,
25 and essentially selects 8 bits of the 24 bits which are properly aligned, under the control of Logic 73 to determine the correct bit shift for the barrel register. The logic 73 uses a known training pattern to produce two 4-bit wide outputs which control the barrel shifter. The Logic 73 essentially keeps resending the same known data training pattern through the clocked delay line, under software control, until it
30 knows the correct bit shift for the barrel shifter. The Barrel Shifter selectively picks the best 8 properly aligned 8 bits of the 24 bit input, under control of the Logic to pass as the Byte output.

State Diagrams:**Eye Detection Flow**

Figure 8 illustrates the general flow of the eye detection process in an eye detection flow diagram, and starts with a Reset 80 which initializes the system to known values, followed by block 81 which waits for a sufficient number of data transitions to ensure a clean determination of an eye. Block 82 examines the 1s and 0s in the even eye history registers to determine an even eye, corresponding to the rising edge clocked data, and selects a first eye with the smallest delay through the clocked delay line. Block 83 does the same thing with respect to an odd eye, and examines the 1s and 0s in the odd eye history registers to determine an odd eye, corresponding to the falling edge clocked data, and selects an odd eye with the smallest delay through the clocked delay line.

There may be several different even phase and odd phase eyes corresponding to different positions along the delay line, and so after the smallest delay eye is detected, the flow diagram recycles from block 84 to block 82 to find the next pair of eyes with the next largest delay, and the logic control continues recycling to block 82 until the complete length of the delay line has been checked for corresponding even and odd phase eyes. Block 84 compares each next detected eye pair with the best previously detected eye pair, and retains the best eye pair, such that it selects the best eye pair of all of the candidate eye pairs, which function is performed by Logic 73 of Figure 7.

At this point, Logic 73 then waits in block 85 for an "align" packet, which is a known training pattern such as a known sequence of bits, to establish the byte boundary which is unknown at this point. After the byte boundary is established by the align training pattern by Logic 73 in Figure 7, then block 86, which also corresponds to Logic 73, sets the align inputs to the Barrel Shifter to align the Barrel Shifter 74 to select and pass the correctly aligned 8 bits as Byte out, and Logic 73 also produces the locked output signal. The Barrel Shifter is then ready during a normal data receive to pass 8 correctly aligned bits as the Byte out.

Block 87 indicates that the data sampling eyes are constantly being updated. A preferred realignment starts at the existing even and odd data sampling eyes, and then looks left and right of the existing eyes to determine the left and right eye edges, and then realigns the center of the even and odd phase eyes between their left and right edges, as explained with reference to Figure 9.

State Diagram for Training Eye Detection

Figure 9 illustrates a two-bit macro state diagram 1 of 2 which shows the state flow for the phase during which the eye position sample points are being determined, and corresponds generally to blocks 82 and 83 of Figure 8.

In Figure 9, the $H[N-1:0]$ inputs to the MUX of Figure 6C, which correspond to the outputs at the bottom of Figure 5, are sampled in sequence and are examined one at a time as the passed output H_{samp} of the MUX of Figure 6C. The first bit $H[0]$, which is either 0, indicating a data transition outside of an eye, or a 1, indicating no data transition possibly inside an eye, is examined in the sequence of steps of Figure 9, where a 0 is shown as $\text{Samp (sample)}=0$, and a 1 is shown as $\text{Samp}=1$. After the first bit $H[0]$ is examined through all of the steps S_0 – SUMMARY of Figure 9, then the second bit $H[1]$ is examined through the same sequence of steps, and etc. until the last bit $H[N-1]$ has been examined.

The states in Figure 9 in the Two Bit Macro State Diagram 1 of 2 are as follows:

- S_0 is the reset state. Control remains here while a reset is active.
- S_1 is an initialization state. After reset is released, the control waits here until an update counter expires, then progresses to S_2 if $\text{samp}=0$ (which indicates that the examined tap of Figure 5 produces a 0 and so is not in the eye) or S_3 if $\text{samp}=1$ (the examined tap output is a 1 and possibly part of an eye).

States S_2 - $S_3.5$ search for an even eye by incrementally searching through the even delay line history, which corresponds to the rising clock edge clocked data.

- S_2 searches for an even eye ($\text{samp}=1$) by incrementing through the even delay line history left to right. It finds the left end of an even eye. If found or if it hits the right end of the delay line it goes to state S_3 , else it stays in S_2 .

- S3 searches for the right end of the even eye, still searching right until it finds samp=0. Control remains in S3 while samp=1 since it is within the eye. When samp=0 is found; control goes to S3.5, a delay state necessary for control to work correctly in certain cases.
- 5
- S3.5 immediately transitions to S4.

States S4-S13 search for the odd eye, which corresponds to the falling clock edge clocked data. The search for the odd eye starts at the detected center of the even eye, and is more complex than the search for the even eye.

- S4 begins searching for the odd eye. For a normal mode, if the odd eye samp=0 it progresses to S7, else samp=1 and it progresses to S5. samp=0 means the odd eye is not aligned in the odd delay line history with the even eye and normal strategy is to search left and right picking the closest odd eye. There are two alternative modes. Search right where control goes to S13, or search left where control goes the S12.
- 10
- S5 means the odd eye is aligned with the even eye and the initial sample point is already in the odd eye. So this state searches left from the initial sample point by decrementing a sample pointer (used to select data points) until it finds the left end of the odd eye or the left end of the delay line. Then control goes to S6.
- 15
- S6 searches right seeking the right end of the odd eye. When it finds it or the right end of the delay line, control goes to the SUMMARY state, which is the state at the end of the search after step S13, where data sample pointers are set for normal processing of even and odd eyes.
- 20
- S7 means the odd eye is not aligned with the even eye and the initial sample point is outside any eyes in a noise area. This state searches left for a matching odd eye by decrementing the sample pointer. When samp=1 it has found the right edge of one odd eye which it remembers in MAX (a right edge register) and control goes to S8. Or it reaches the left end of the delay line finding no left odd eye in which case it goes to S11.
- 25
- S8 is intended to search for an unaligned right odd eye. It continues searching while samp=0 until samp=1 which indicates the left end of a right
- 30

odd eye is found and remembered in MIN (a left edge register) and control progresses to S9. If the right end of the delay line is reached before finding samp=1, there is no right eye, so control goes directly to left eye processing in S10.

- 5
 - S9 is where the MIN and MAX distance from the even eye is compared. If MIN is closer, control passes to S6. If MAX is closer, control passes to S10.
 - S10 searches for the left end of the odd eye. Control remains in S10 while samp=1. When samp=0 or the left end of the delay line is reached, control passes to SUMMARY.
- 10
 - S11 searches for an unaligned right odd eye when there is no left odd eye. Control remains in S11 while samp=0. When samp=1, control goes to S6. If the right end of the delay line is reached before finding samp=1, there are no odd eyes. This is an error condition which is detected and indicated by warning indicators.
- 15
 - S12 searches for an unaligned odd eye left of the even eye in the delay lines. Control remains in S12 while samp=0 and goes to S10 when samp=1. If the left end of the delay line is reached before samp=1 is found, no left eye exist and control opts for looking right in S13, unless control came into S12 from S13, in which case there are no odd eyes. This is an error condition detected
- 20
 - and indicated by warning indicators and control goes to SUMMARY.
 - S13 searches for an unaligned odd eye right of the even eye in the delay lines. Control remains in S13 while samp=0 and goes to S6 when samp=1. If the right end of the delay line is reached before samp=1 is found, no right eye exists and control opts for looking left in S12 unless control came into
- 25
 - S13 from S12, in which case there are no odd eyes. This is an error condition detected and indicated by warning indicators and control goes to SUMMARY.
 - SUMMARY is the state where the eye data sampling points are fixed for normal operation.

30

State Diagram for Normal Operation with Eye Sample Point Realignment

Figure 10 illustrates a two bit macro state diagram 2 of 2 which shows the repetitive state flow during normal data capture operations. In this phase, the logic is to capture serial data and convert it to a byte parallel format.

5

Figure 10 shows the normal data capture run states. Control normally resides in RUN0. Periodically as determined by a clock counter the done update signal enables control to progress into RUN1. States RUN1 and RUN2 increment and decrement to the extremes of the even eye, which may have changed. The new limits are remembered. Similarly, states RUN3 and RUN4 increment and decrement to the extremes of the odd eye and the limits are remembered. State RUN5 uses the findings of RUN1 through RUN4 to calculate new data sampling points, which get latched into use. Control then returns to RUN0 for another update period.

15

SiBiDi

The present invention also provides a hardware mechanism for automatically adjusting transmission delays for optimal two-bit simultaneous bidirectional SiBiDi signaling.

20 The SiBiDi (simultaneous BiDirectional) IOcell "subtraction" of the transmitted signal is more successful if the signal that needs to be subtracted changes at times where the desired received signal does not change. This can be achieved by delaying transmission by an appropriate amount (some fraction of the clock cycle). But delaying transmission at one end means that the data will arrive at the other end
25 delayed. Then the circuitry at the other end will have to readjust its transmission delay so that its own "subtraction" is optimal. Therefore one needs to find a pair of delay settings, one for each circuit at each of the two ends of the wire, so that the "subtraction" gives equally good quality results for both ends.

30 Furthermore, in order to achieve this, the two ends need to exchange information regarding the quality of the local subtraction for each choice of transmission delay on the other end. But this information cannot be exchanged using the same signaling transmission technique that is being optimized. If a delay setting is bad, it

may corrupt the data sent that describe how bad it is. The present invention describes a hardware mechanism for automatically adjusting the transmission delays for optimal two-bit SiBiDi signaling.

- 5 Figure 11 illustrates a first embodiment wherein two differential data lines connect a pair of identical nodes 1, 2, and wherein each node has a unique ID, and each node operates with a 2-bit sender CPU and a 2-bit capture CPU.

10 The method uses a "safe communication" set-up phase to communicate the results of each set of transmission delays. The 2 bit sender/capture units are used for safe communication by using a unidirectional setting for the IOcells (wherein transmission is in one direction only to minimize noise) and only one of the 1 bit parts of the units.

- 15 Figure 11 illustrates the electrical connections between a pair of nodes, and shows two differential data lines (each composed of 2 wires to enable differential signaling). The arrows indicate the direction of the unidirectional signals during the safe communication set-up phase. Otherwise, the electrical connections are bi-directional during normal communications.

20

The Sel A, B boxes are MUX's wherein Sel = 0 chooses the upper path from the 2-bit sender unit and to the 2-bit capture unit, and Sel = 1 the lower path.

- I) Set a READY register (not shown) to 0.
- 25 1) IOcell A = unidirectional transmitter mode.
IOcell B = unidirectional receiver mode.
Sel A = 1 (chooses upper path).
Sel B = 1 (chooses upper path).
- 30 2) Set sender delay mode to zero delay.
3) Begin training, which has been described hereinabove to find and detect good data capture eyes.

- 4) Save eye parameters for the safe unidirectional set-up phase communication in the middle of a good data capture eye.
- 5) Send first bit of ID. In this embodiment, the unique IDs of each of nodes 1 and 2 determine which node is a master and which node is a slave, with the higher ID automatically being the master node.
- 5 Wait until you receive the other node's first bit of ID.
Compare first ID bits.
If equal repeat.
If local is less than neighbor's then set PRIORITY = 0
- 10 If local is larger than neighbor's then set PRIORITY = 1
- 6) IOcell A = SiBiDi mode.
IOcell B = SiBiDi mode.
Sel A = 0
Sel B = 0
- 15 7) Sender delay mode = 0
Begin training.
Look for eyes.
Save local eye parameters.
- 8) Go to safe mode as in step 1).
- 20 9) Send your local eye parameters to the other node.
Receive the other node's eye parameters
- 10) Compare the parameters and save the minimum.
- 11) Compare current minimum eye parameters with previous minimum eye parameters and save the maximum of the two together with the local sender delay mode of the maximum. This mode is considered to be the optimum mode and is designated as OPT_SENDER_MODE.
- 25

In this first embodiment, an 8 tap delay line is assumed, so that each node has a possibility of 8 different delays ranging from zero delay to the maximum delay in 8 steps, so the number of possible combinations is $8 \times 8 = 64$. Stated differently, for each of 8 delays at one node, there are 8 possible delays at the second node. So 64

30

possible combinations must be tested to select the optimum combination. Step 12 simply cycles through all 64 combinations, one at a time.

Go back to step 6) and repeat for a total of 64 times using the following sender delay

5 mode sequence:

If PRIORITY = 0 then the neighbor changes modes first. The local sequence is:

0 for 8 steps

1 for 8 steps

10 2 for 8 steps

3 for 8 steps

4 for 8 steps

5 for 8 steps

6 for 8 steps

15 7 for 8 steps

If PRIORITY = 1 then the local sequence is:

0 for 1 step

1 for 1 step

2 for 1 step

20 3 for 1 step

4 for 1 step

5 for 1 step

6 for 1 step

7 for 1 step

25 Repeat 8 times

12) Go to SiBiDi operation as in 6) with sender delay mode = OPT_SENDER_MODE.

13) Set the READY register to 1 to indicate that the system is optimized and ready for normal SiBiDi communications.

30

If at any step there is a failure so that step 14 is not reached, then the node has failed. The failed node can be identified by the contents of the READY register.

Figure 12 illustrates a second embodiment of SiBiDi electrical communications between two nodes, Node 1, Node 2, with one distinction from the first embodiment being that a single differential communication line connects the nodes. A second distinction is that this embodiment operates with 1-bit data rather than 2-bit data as in the embodiment of Figure 11. Another distinction is that in this embodiment, the Master node and Slave node are designated, as by control software, rather than being based upon assigned IDs. This embodiment also assumes a 16 tap delay line rather than an 8 tap delay line as in Figure 11.

10 *Overview:*

An Initial Alignment Procedure (IAP) is a sequence of steps whereby each synchronous signal of each port determines the optimal transmit delay line setting (for its Outstage). The Massively Parallel Supercomputer described in U.S. provisional application Serial No. 60/271,124 describes a massively parallel computer having 32 x 32 x 64 nodes connected as a three dimensional torus wherein each node connects to 6 adjacent nodes. Each node has 6 ports with 20 synchronous signals per port, such that all 120 synchronous signals (6 ports x 20 sync signals/port) on a node computer chip at a node of the supercomputer are able to perform this individual training independently. All could occur in parallel, or just one at a time. (all under software control). Training is done on both directions of a SiBiDi link at the same time; which allows for the necessary ISI (Inter Symbol Interference) and near end noise (with environmental noise). Referring to Figure 12, the high level flow of the IAP Sequence is:

25

1. Software action: Identifies one side of a synchronous link as "master" and the other side as "slave" by writing to the IAP Control register of each node computer chip.

30

2. Hardware action: The master side (side A in Figure 12) communicates with the slave side (side B in Figure 12) to start the training. This is somewhat complicated since information must be communicated across a link before the link is fully trained. (See "Communication Across an Untrained Link" below)

3. Hardware action: Each side of an individual link has a state machine (as shown in Figures 13A and 13B) that runs through all possible delay line settings and compares the results to find the optimal delay line setting. Changing the delay setting on one side influences the eye on both sides, so the system needs to run through all 16x16 combinations. (Note: the Outstage (data send) delay line has 16 settings, as explained previously). For each loop of the delay line training, the Instage (data capture) macro receives a pseudo-random data stream from the other side and seeks to find the eye and presents the eye size information for analysis.
4. Software action: Read the IAP Status registers to determine the success/failure of the training. The exact delay line settings and eye-size margins that were achieved may be read via other link-specific status registers, which are software accessible.

Link Training Sequence:

The state machines of Figures 13A and 13B illustrate the steps taken by the node compute chip in the training of a synchronous Si-Bi-Di connection. Each side of the link utilizes the following registers:

DTR - *Delay Tap Register* - Controls the delay line in the Outstage. (Valid settings are 0-15). Two additional "working" copies are used during the IAP Sequence: Mst-DTR and Slv-DTR.

LBDTR - *Local Best Delay Tap Register* - Holds the DTR value that corresponds to the best yet seen eye size during the training. At the end of the training, the contents of the LBDTR are permanently loaded into the DTR.

MBESR - *Mutual Best Eye Size Register* - Holds the best yet seen eye size during the training (based on the minimum of the side A and side B eye sizes for a given step in the training).

Noise generator macros can be enabled during the link training sequence as a way of artificially adding more noise to simulate a very noisy environment and guaranteeing more vertical voltage margin, which relates to the size of the eye. Software begins

the IAP Sequence by writing the "Start" bit in the IAP Control register, and identifies the chip as Master (side A) or Slave (side B).

Communication Across an Untrained Link:

5

It is necessary to perform communication between the two sides of a link prior to the link having been fully trained. To ensure the most reliable data transfer possible, the following procedures are utilized:

- 1) Data is sent in only one direction at a time.
- 10 2) Data is sent at a slower data rate. A 1:8 ratio can be used. (i.e. holding a '1' or '0' for eight bit times).

Prior to training, the two sides of a link have no predictable phase relationship. Therefore, if one side transmits a "110011", and if the sample point lines up with the switching data, then the data may be received as "111011" or "100011", etc. The
 15 transmission rate has to be slow enough to detect stable data across consecutive samples, and not be confused by the mis-samplings that may occur during transitions of 0->1 or 1->0 within the bit stream.

"Commands" sent between the Master and Slave are preceded by a long string of 1's followed by eight 0's. A command will appear as: many ones, 8 zeros, 8 bit times
 20 of the first bit of the command, 8 bit times of the second bit of the command, 8 bit times of the last bit of the command. The receiving side detects the 1 --> 0 transition and estimates the middle of the 8 bit-time window. (In reality, this may be the 3rd, 4th, or 5th bit of the 8 bit-time window; all of which should be stable and valid). Thereafter, every eighth bit is sampled to decipher the command/information.

25 Referring to Figures 13A and 13B, all commands are indicated in capital letters, and "Same" in the blocks of Figure 13B indicates the same or a corresponding block as in Figure 13A.

Referring to Figure 13A, which is the Master side, at stage (0) Wait for a Start software command, then reset the registers DTR's, LBDTR, MBESR. At stage (1)
 30 the Master waits to receive a BEGIN command from the Slave, and if not sends a BEGIN to the Slave, and waits to receive a BEGIN reply from the Slave. If not the

Master waits and resends BEGIN to the Slave. If the Master receives a BEGIN from another Master, it aborts.

At stage (4), if Yes, the Master sends a TRAIN command, indicating the Master is about to start synchronization and then pauses.

- 5 At stage (5), the Master transmits a random data bit stream to enable capture of the eyes.

At stage (6), the Master waits for capture of the eyes and evaluates information on each eye such as the eye size.

- 10 At stage (7), the Master waits to receive data on the eye size, and if not, waits (e.g. 1 usec) and sends data on the eye size and again waits to receive data on the eye size.

When received, the Master updates the MBESR and LBDTR registers and increments the Mst-DTR, and if a wrap (counter overflow) increments the SLV-DTR register, and stages (4)-(9) are repeated for all 256 combinations.

- 15 If Yes, at stage (10), the Master sends an END command to end Eye-Training.

At stage (11) the Master awaits receiving an END command from the slave.

If Yes, at stage (12) the Master loads the DTR with LBDTR registers, and resets the instage, which is a set-up node.

At stage (13), the optimal eye parameters are used to transmit random data.

- 20 Stages (12) and (13) use optimal eye parameters to transmit data, and then the optimal eye parameters are re-evaluated, and if successful are locked in place.

At stage (14), the Master asserts a Reset Glitch signal to reset and re-evaluate data capture, checks the eye size against minimum eye size data, and updates an IAP Status Register.

- 25 The operation of the Slave Side of Figure 13B should be apparent from the above description of the Master Side.

While several embodiments and variations of the present invention for a data capture technique for high speed signaling are described in detail herein, it should be apparent that the disclosure and teachings of the present invention will suggest many alternative designs to those skilled in the art.

5

CLAIMS

Having thus described our invention, what we claim as new and desire to secure by Letters Patent is:

- 1 1. A data capture method to allow optimal sampling and capture of an
2 asynchronous data stream without sending a clock signal with the data stream
3 comprising:
4 capturing the data by sending serial data bits of the data stream down a
5 clocked delay line with a series of delay taps;
6 sampling all of the delay taps with a clock;
7 comparing each delay tap output with a neighbor delay tap output to
8 determine if it is the same;
9 using the comparisons to form a clocked string to generate a data history
10 record;
11 examining the data history record to determine optimal data capture eyes by
12 looking for data capture eyes where the data does not transition between adjacent
13 delay taps, which are detected as optimal data capture eyes.
- 1 2. The method of claim 1, including periodically updating the data history
2 record to compensate for changing parameters.
- 1 3. The method of claim 1, wherein the clocked string is combined with previous
2 clocked strings to generate the data history record.
- 1 4. The method of claim 1, wherein the serial data enters the clocked delay line,
2 and is clocked through a combinatorial series of inverters, each of which adds an
3 increment of delay, and each inverter output is directed to history registers.
- 1 5. The method of claim 4, wherein each inverter output is directed to even
2 history registers and odd history registers, the even history registers are clocked by a
3 positive edge of the clock, and the odd history registers are clocked by a negative
4 edge of the clock, to allow logic to capture the serial data at twice the clock rate, the
5 even history registers are used to detect an even data capture eye for the positive

6 clock phase, and the odd history registers are used to detect an odd data capture eye
7 for the negative clock phase.

1 6. The method of claim 5, wherein an even eye multiplexer receives all of the
2 outputs of the even history registers and an odd eye multiplexer receives all of the
3 outputs of the odd history registers.

1 7. The method of claim 4, wherein the history registers include a first history
2 register clocked at a first clock rate and serially arranged second, third and fourth
3 data history eye registers serially receiving the output of the first history register and
4 clocked at a second clock rate.

1 8. The method of claim 4, wherein
2 the clocked delay line comprises a delay line register at the output of each
3 inverter, and
4 the output of each delay line register is directed to an exclusive OR gate
5 XOR which also receives an input from the next delay line register in the clocked
6 delay line, and since the data bit is inverted by the next delay line inverter before it
7 enters the next delay line register, and if the data bit does not undergo a transition
8 between the consecutive stages, each register and the next register will hold opposite
9 values, such that the XOR gate will produce a 1, indicating there was no data
10 transition between the consecutive stages, and conversely, if the data bit undergoes a
11 transition between the consecutive stages, each register and the next register will
12 hold the same value, such that each stage XOR gate will produce a 0, indicating
13 there was a data transition between the consecutive stages.

1 9. The method of claim 8, wherein the output of each XOR gate is input to an
2 AND gate, the output of which is input to a first history register, which is the first of
3 a series of four history registers, and the first history register is sampled and reset to
4 high at first clock rate, and the second, third and fourth history are sampled and
5 updated at a higher second clock rate.

1 10. The method of claim 5, including searching for an even data eye, by
2 incrementally searching through the even history registers, to search for leading and
3 ending edges of the even data eye, and the search for the odd eye starts at the
4 detected center of the even eye, and then searches for the leading and ending edges
5 of the second data eye.

1 11. The method of claim 5, wherein:
2 in a first stage, the history registers are periodically reset and flushed and a
3 new history record is acquired, after which the best data eye is determined for each
4 phase of the clock independently, the best data eyes are then used to send and
5 capture data bits which are forwarded to a next stage every system clock;
6 and in a second stage, the forwarded data bits are inserted into a shift register
7 which is used along with a barrel shifter to select and pass correctly aligned data bit.

1 12. The method of claim 5, wherein the data sampling eyes are constantly being
2 updated and realigned, which starts at existing even and odd data sampling eyes, and
3 then looks left and right of the existing eyes to determine the left and right eye
4 edges, and then realigns the center of the even and odd eyes between their left and
5 right edges.

1 13. The method of claim 8, wherein the outputs of the XOR gates are sampled in
2 sequence and examined one at a time to determine if each output is either a 0,
3 indicating a data transition outside of an eye, or a 1, indicating no data transition
4 possibly inside an eye, which sequential sampling and examination is repeated in
5 sequence for each delay tap output of the shift register.

1 14. The method of claim 1, wherein the delay tap outputs are sampled by a first
2 circuit clocked by a positive edge of the clock, and by a second circuit clocked by a
3 negative edge of the clock, an even data capture eye is detected for the positive
4 clock phase, and an odd data capture eye is detected for the negative clock phase
5 independently of the detection of the even eye.

1

1 15. A mechanism for automatically adjusting transmission delays for optimal
2 simultaneous bi-directional (SiBiDi) signaling between two nodes to improve the
3 signal quality of the simultaneous bi-directional signaling over a communication
4 line, wherein during a set-up sequence, parameter setting data is sent in a
5 unidirectional communication over the communication line to allow the two nodes
6 to more accurately exchange the parameter setting data during the set-up sequence,
7 whereby the unidirectional communication has better signal quality to more
8 accurately exchange the parameter setting data.

1 16. The mechanism of claim 15, wherein during the set-up sequence, data is sent
2 at a slower data rate than during SiBiDi signaling, and in which a 1:n ratio is used,
3 holding a '1' or '0' for n bit times.

1 17. The mechanism of claim 15, wherein each node has a possibility of n
2 different delays ranging from a minimum or zero delay to a maximum delay in n
3 steps, so the number of possible combinations of delays is n x n, so that n x n
4 combinations are tested to select an optimum delay combination, and the mechanism
5 cycles through all n x n combinations, one at a time.

1 18. The mechanism of claim 15, wherein one differential data line connects the
2 two nodes, and each node operates with a 1-bit sender CPU and 1-bit capture CPU.

1 19. The mechanism of claim 15, wherein two differential data lines connect the
2 two identical nodes, and each node operates with a 2-bit sender CPU and 2-bit
3 capture CPU.

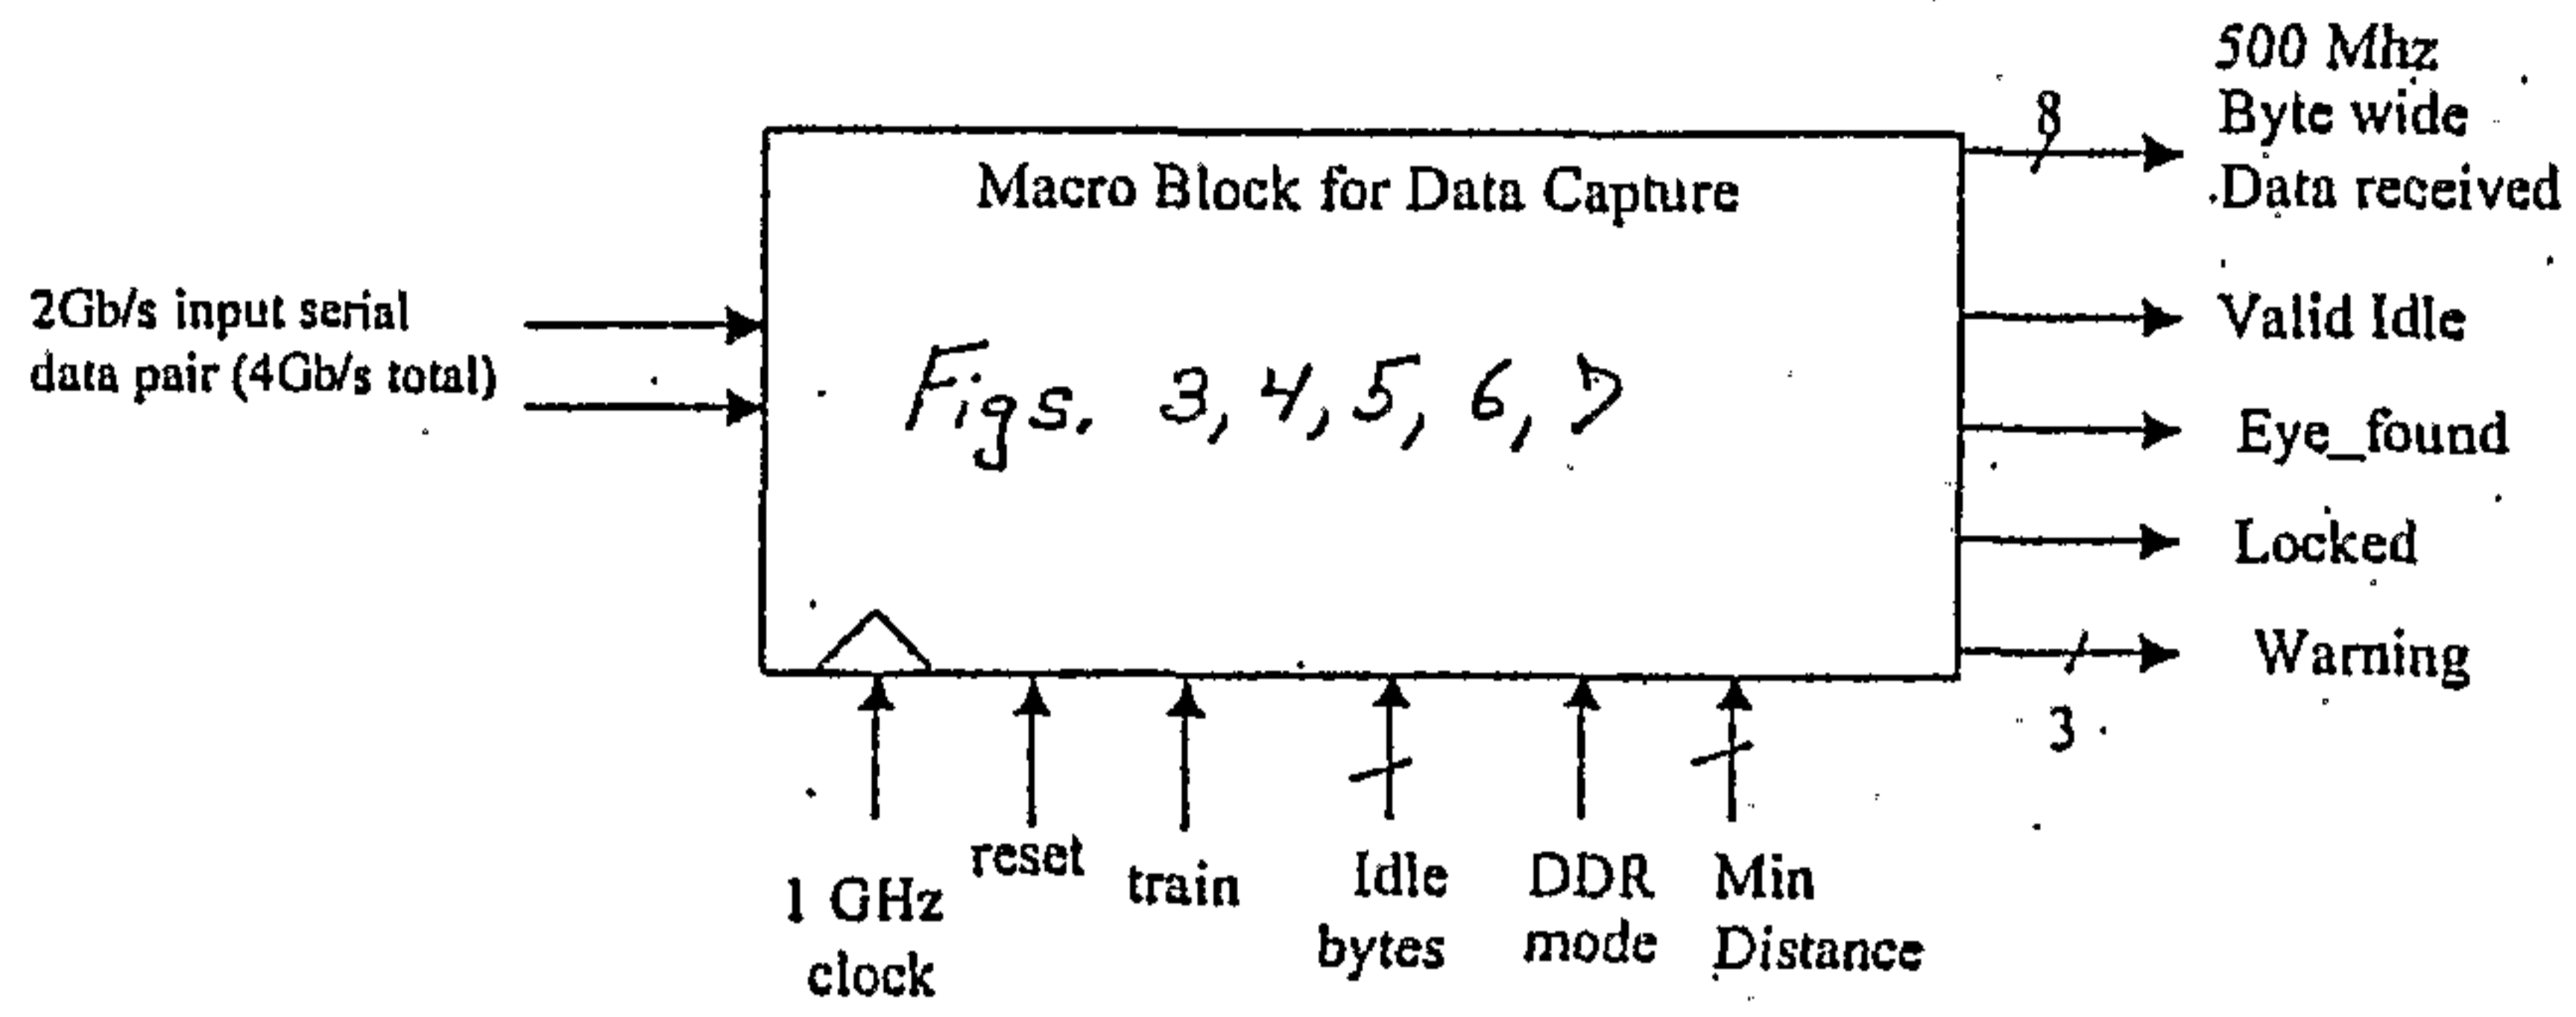


Figure 1 Data Receive Macro Block

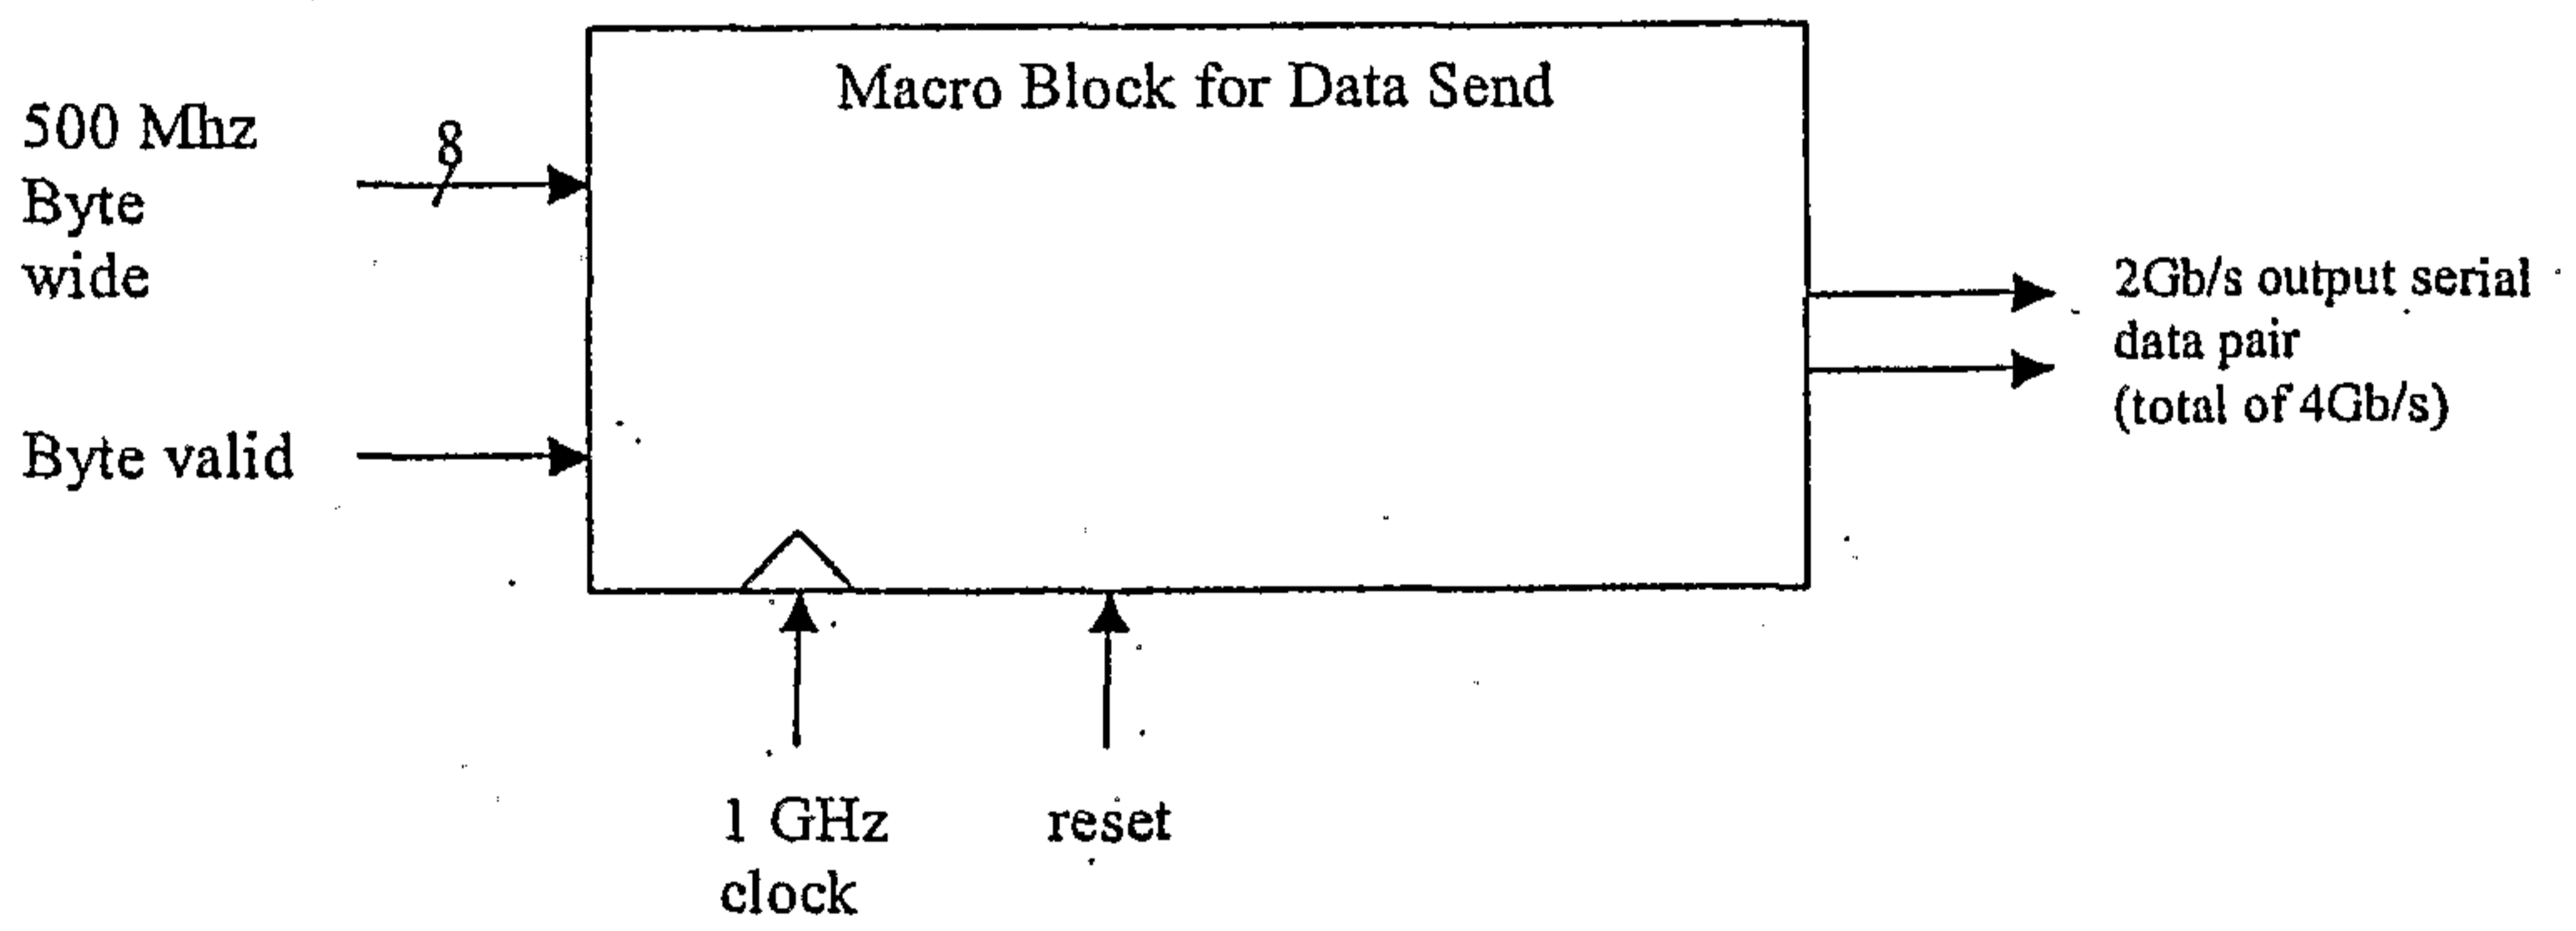


Figure 2 Data Send Macro Block

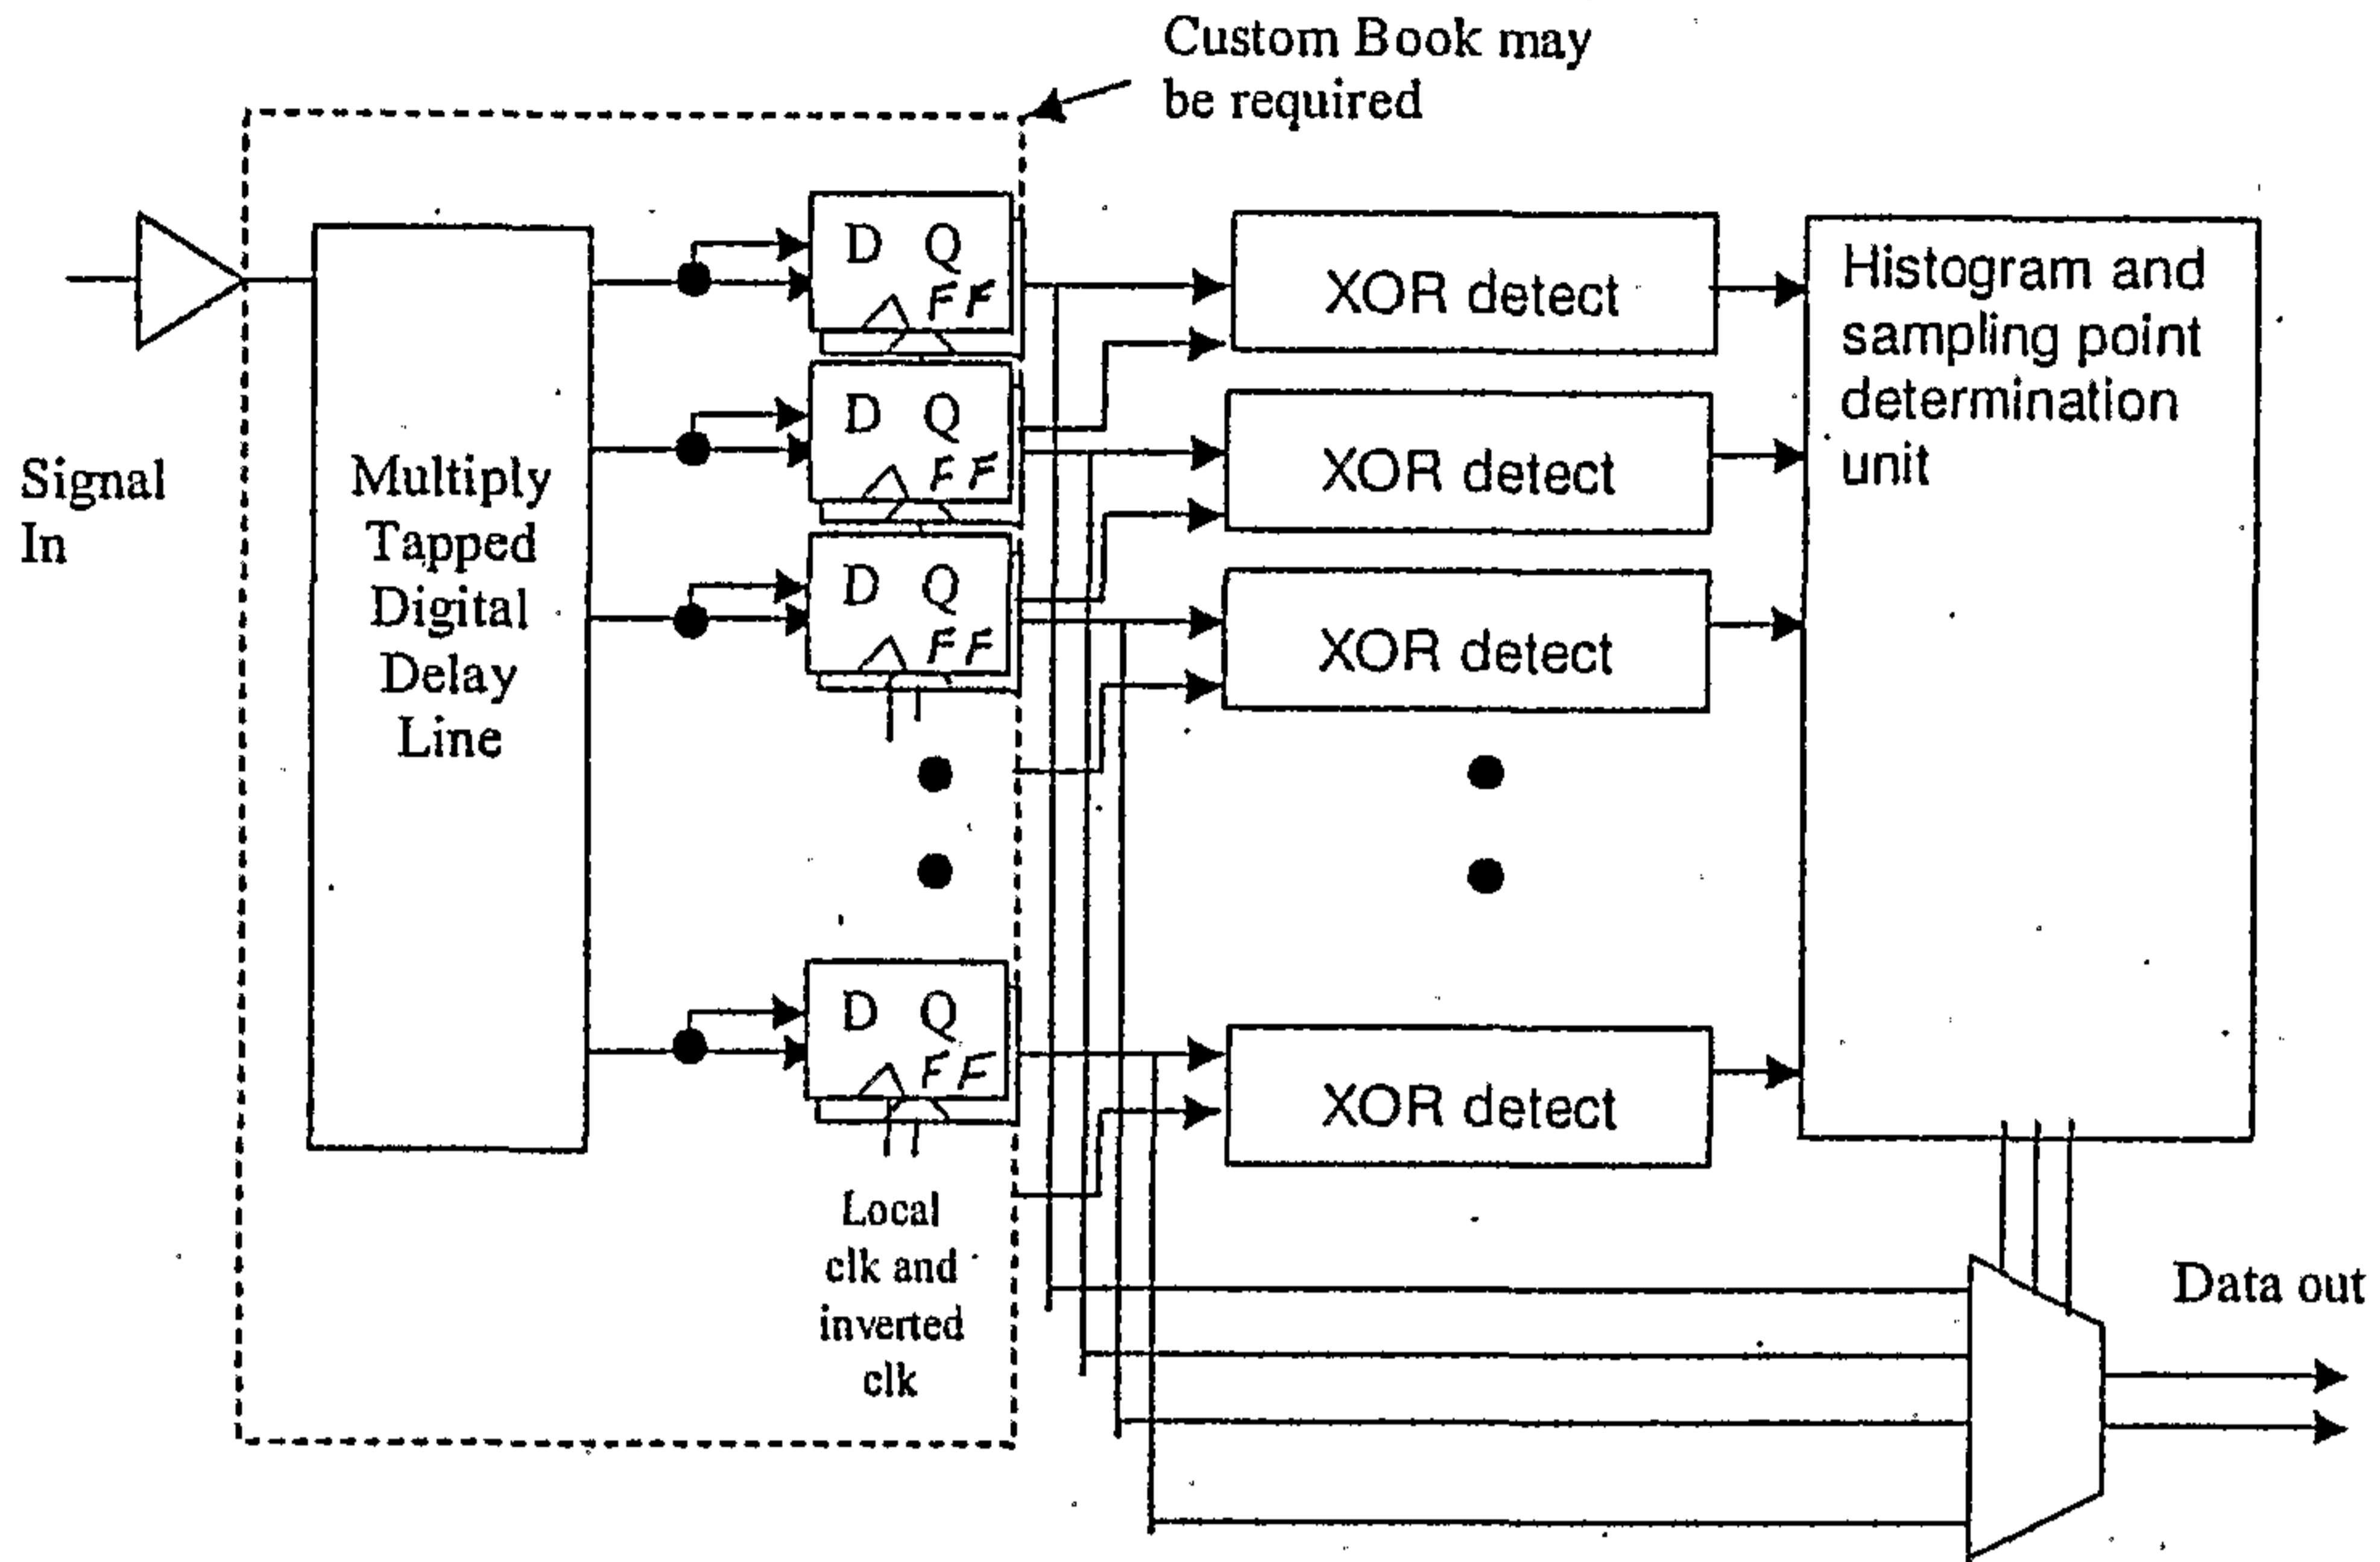


Figure 3 Instage, 2-bit Macro Block

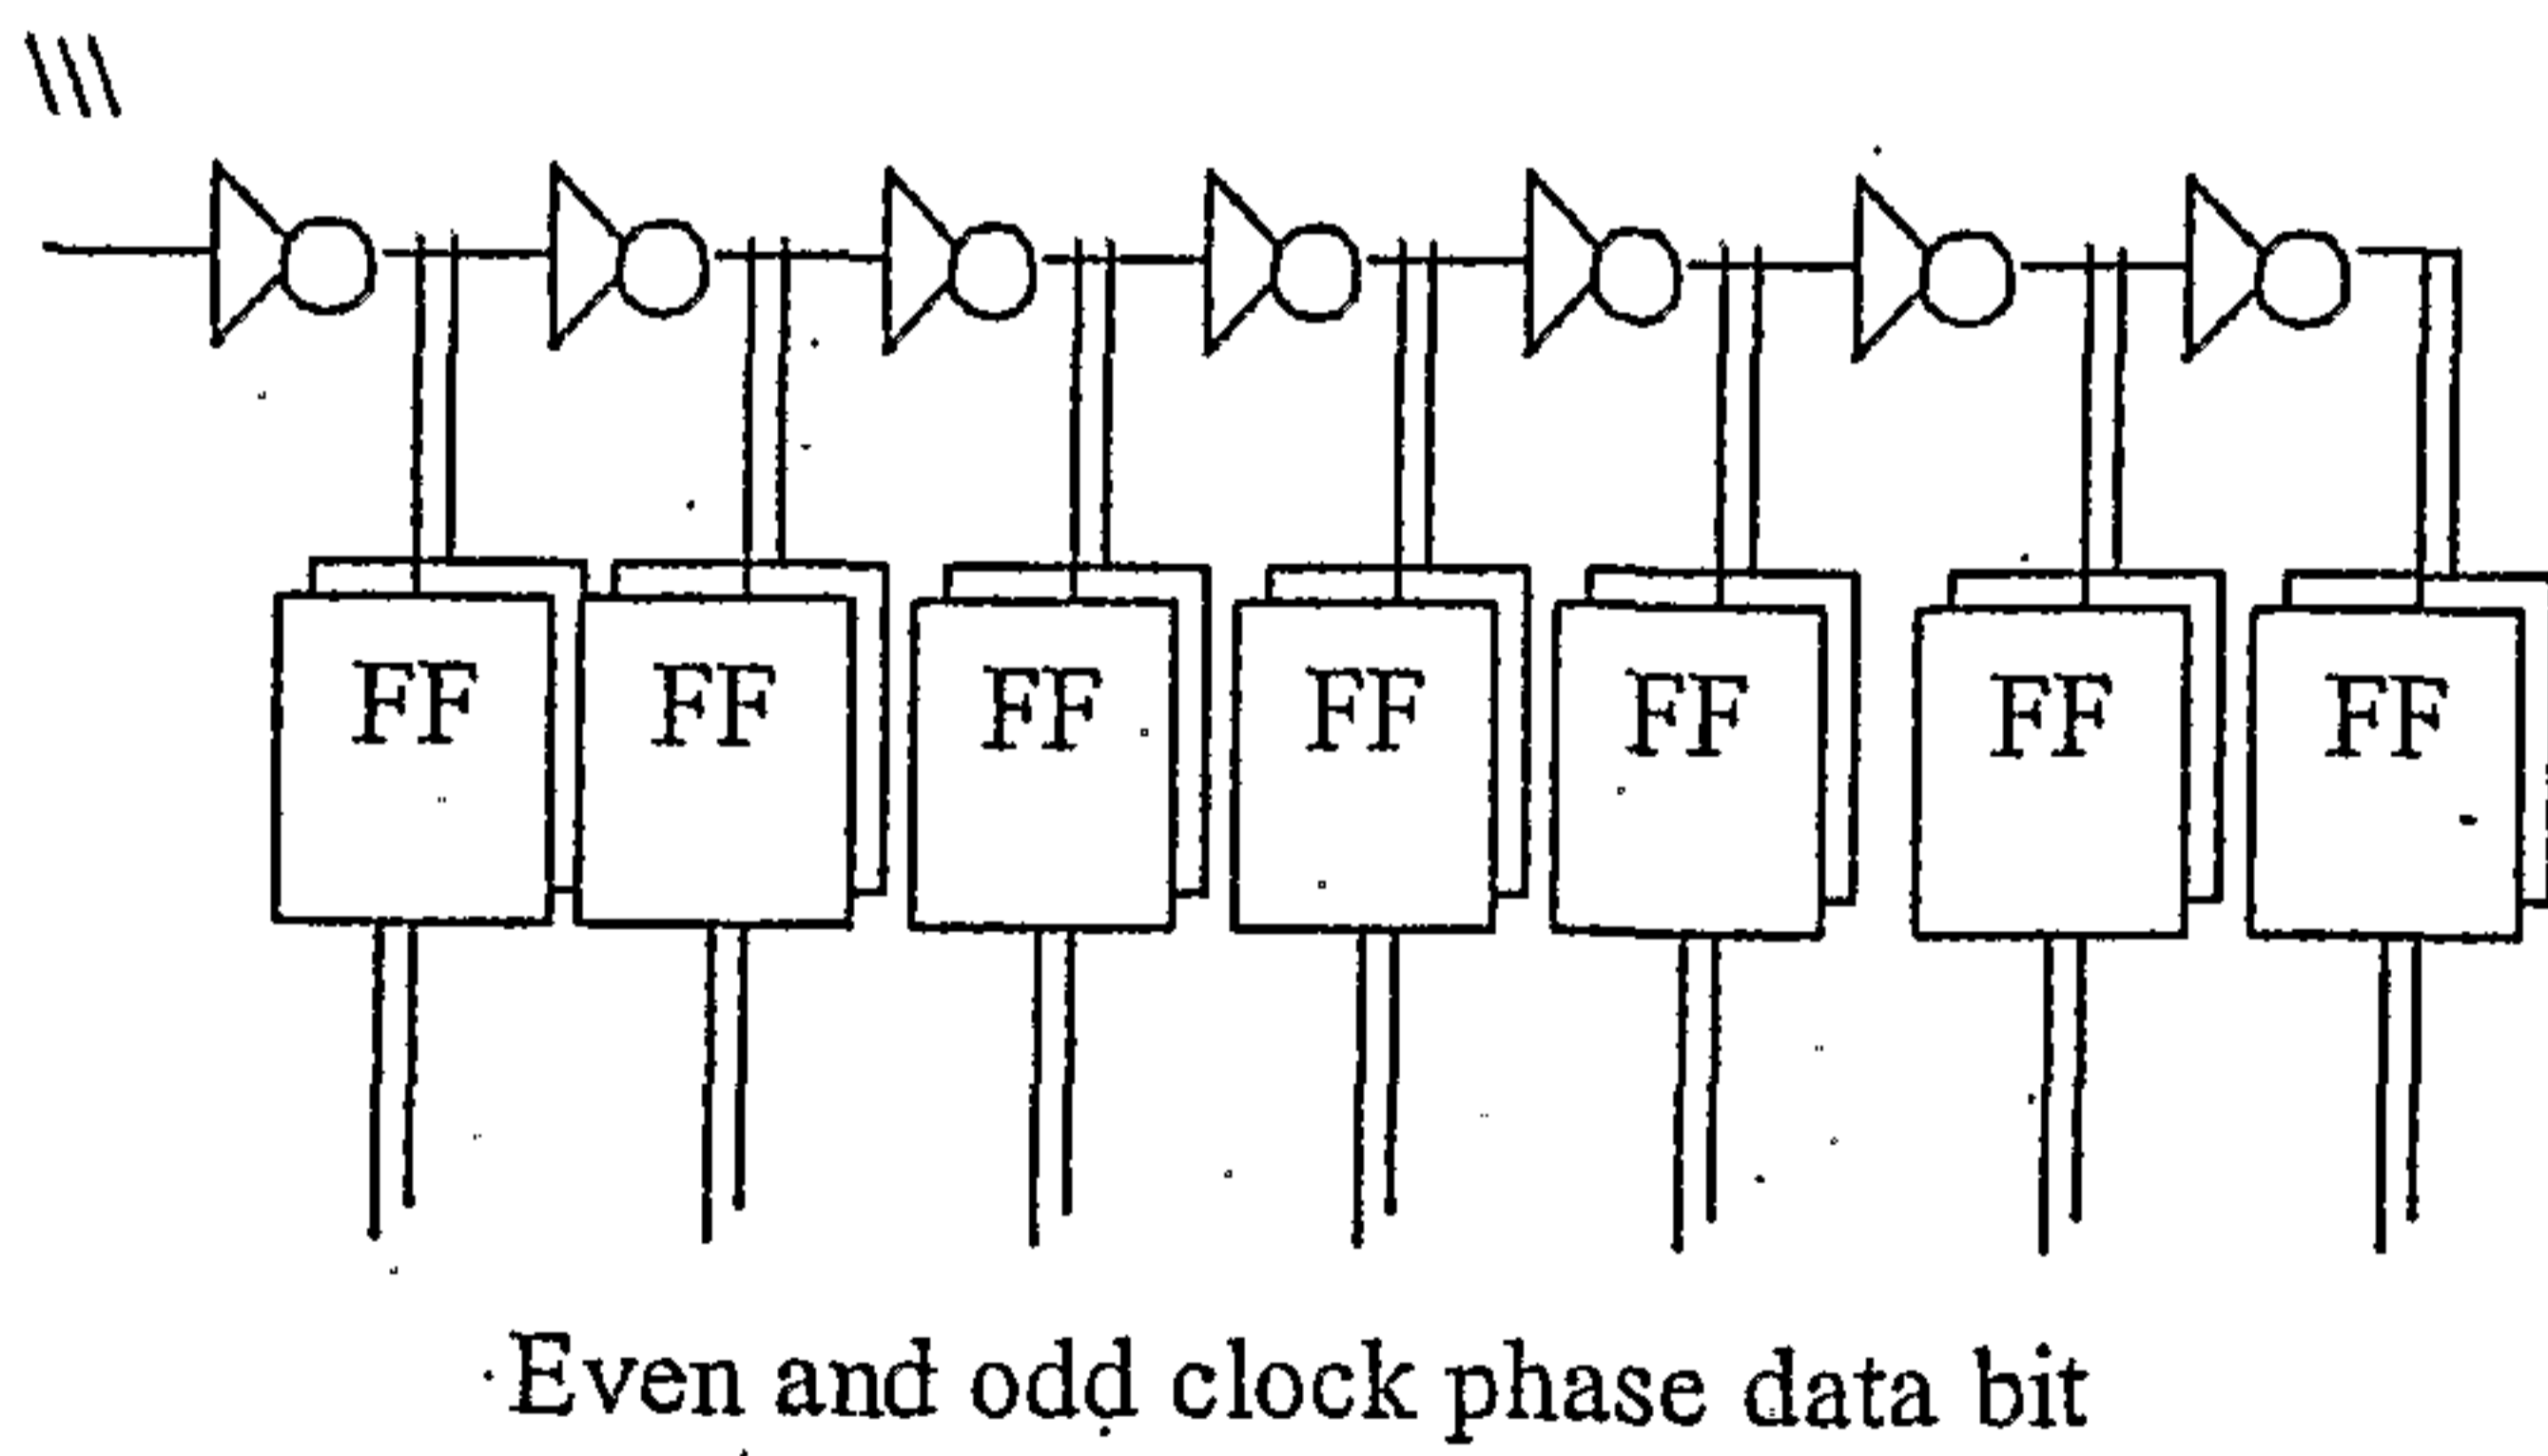


Figure 4 Clocked Delay Line




Figure 5 History Logic

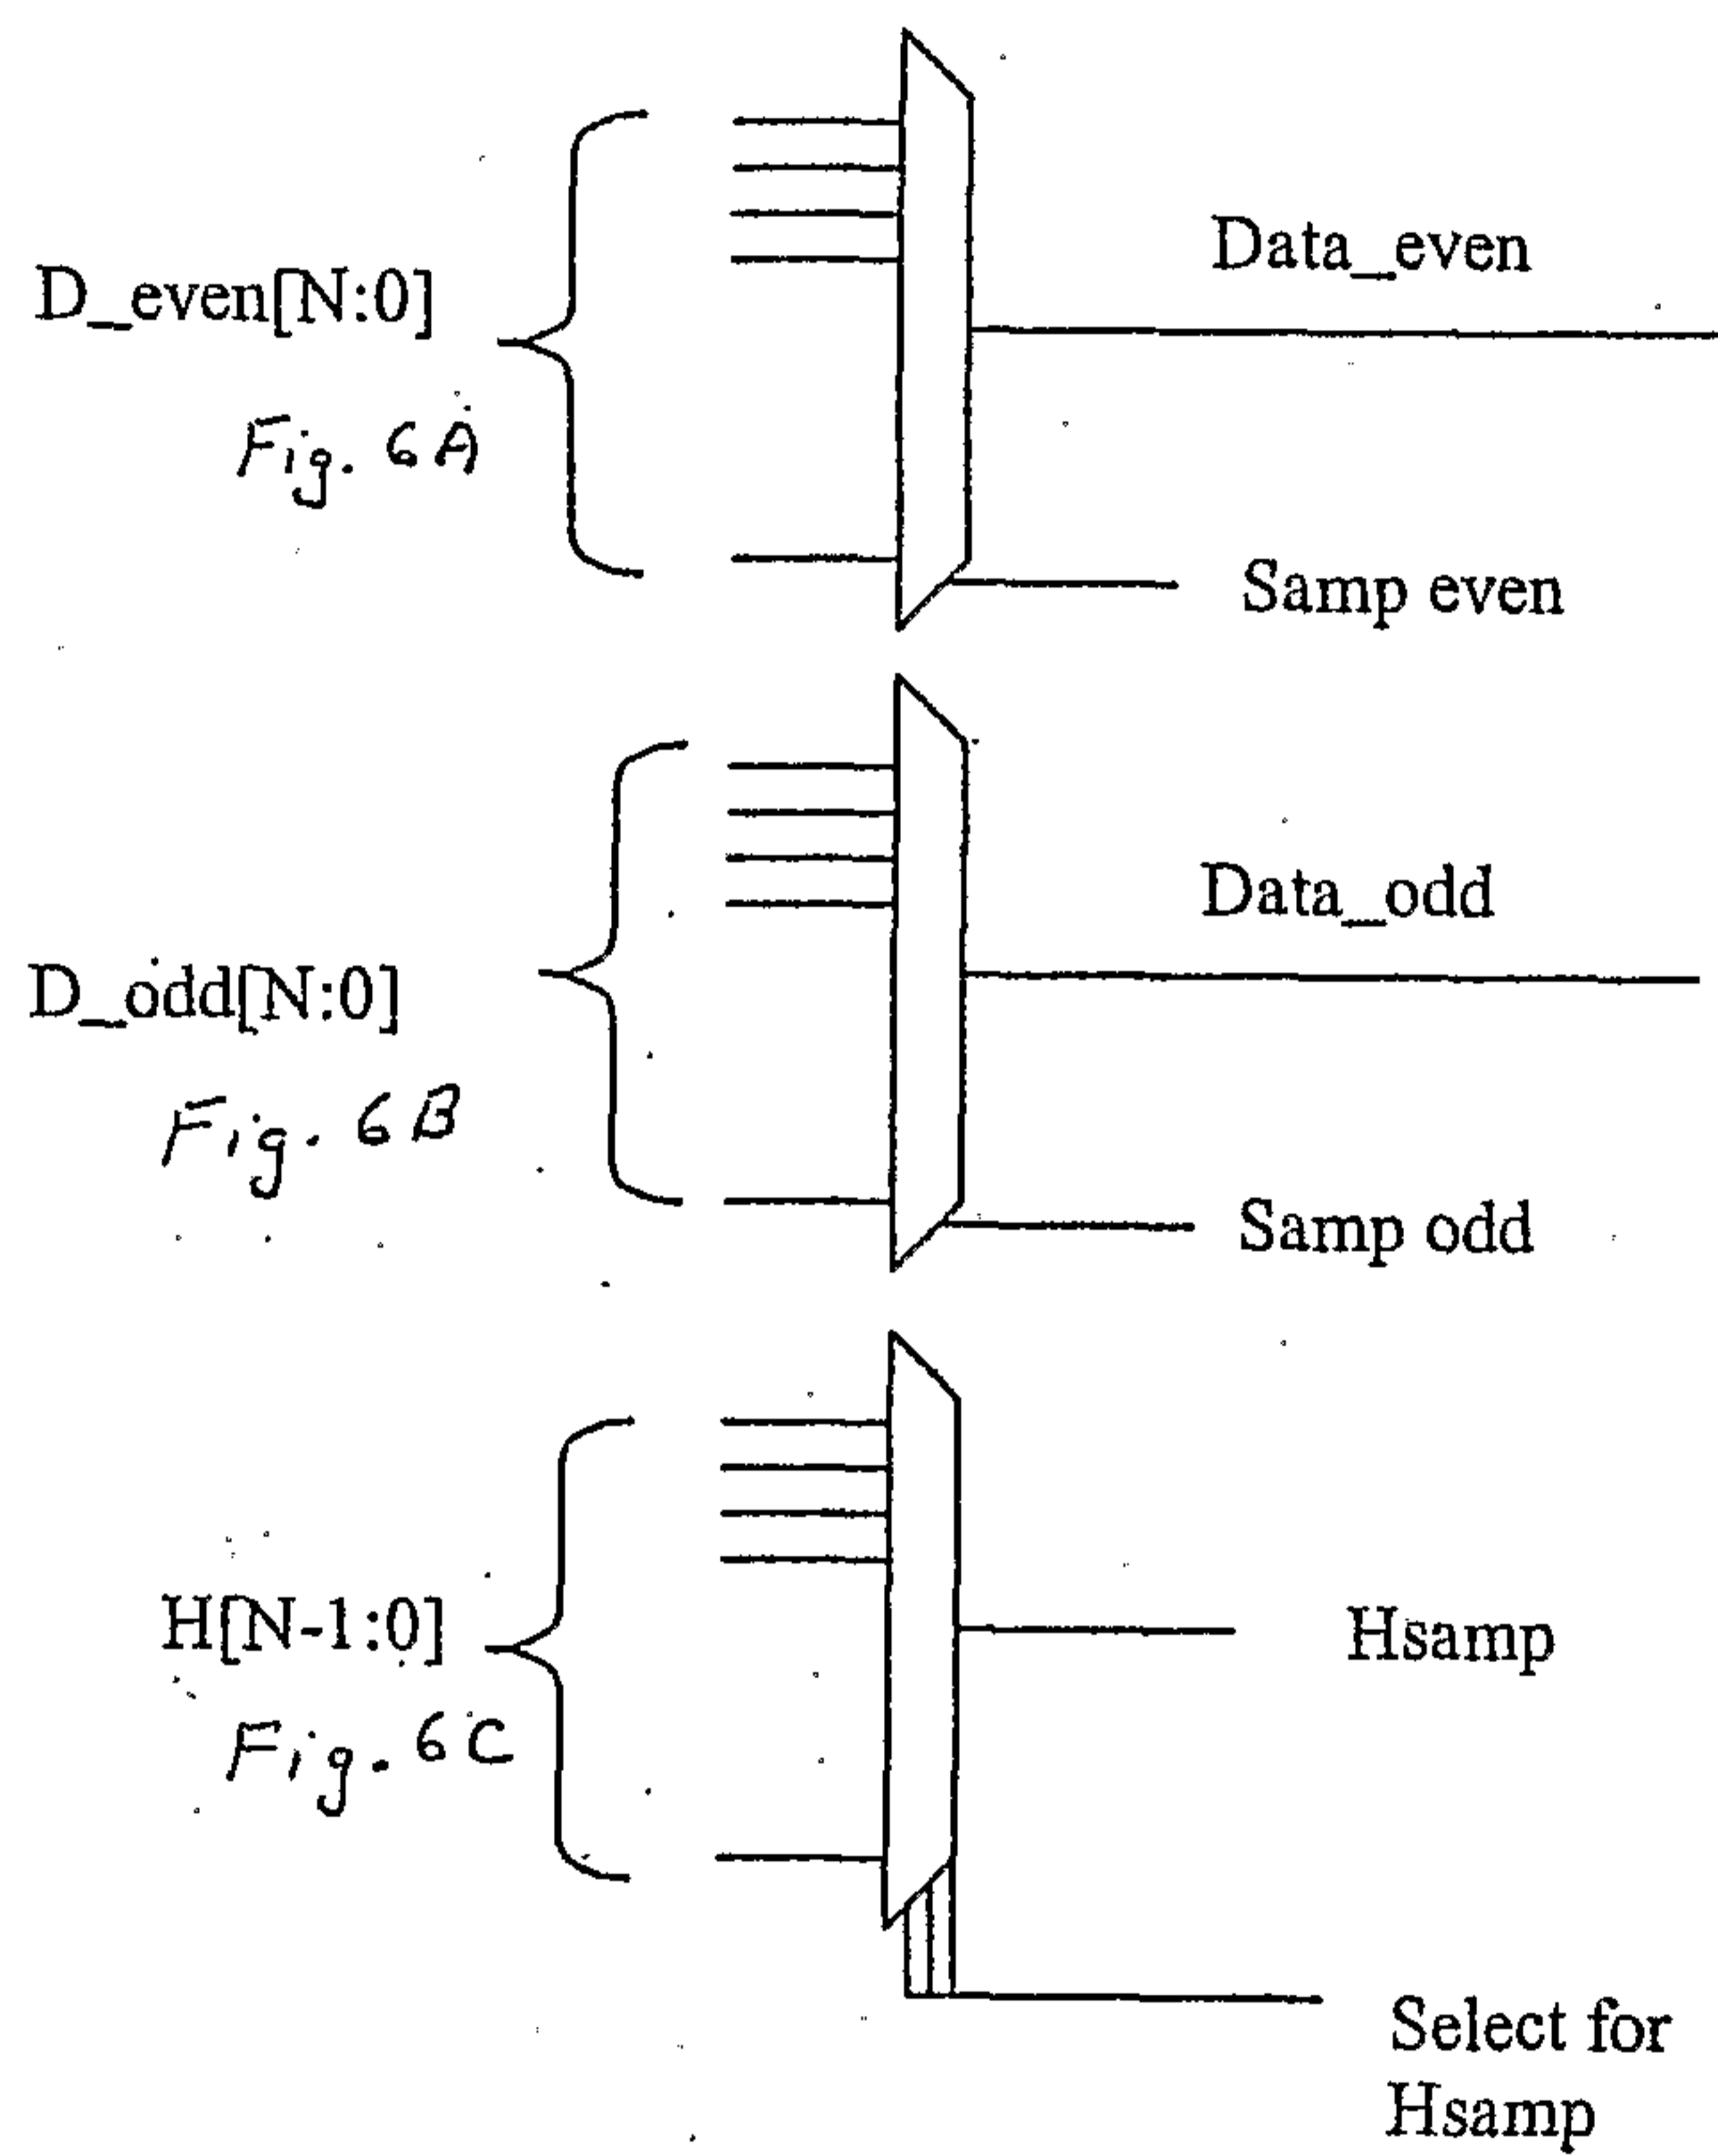


Figure 6 Data and History Sample Muxes


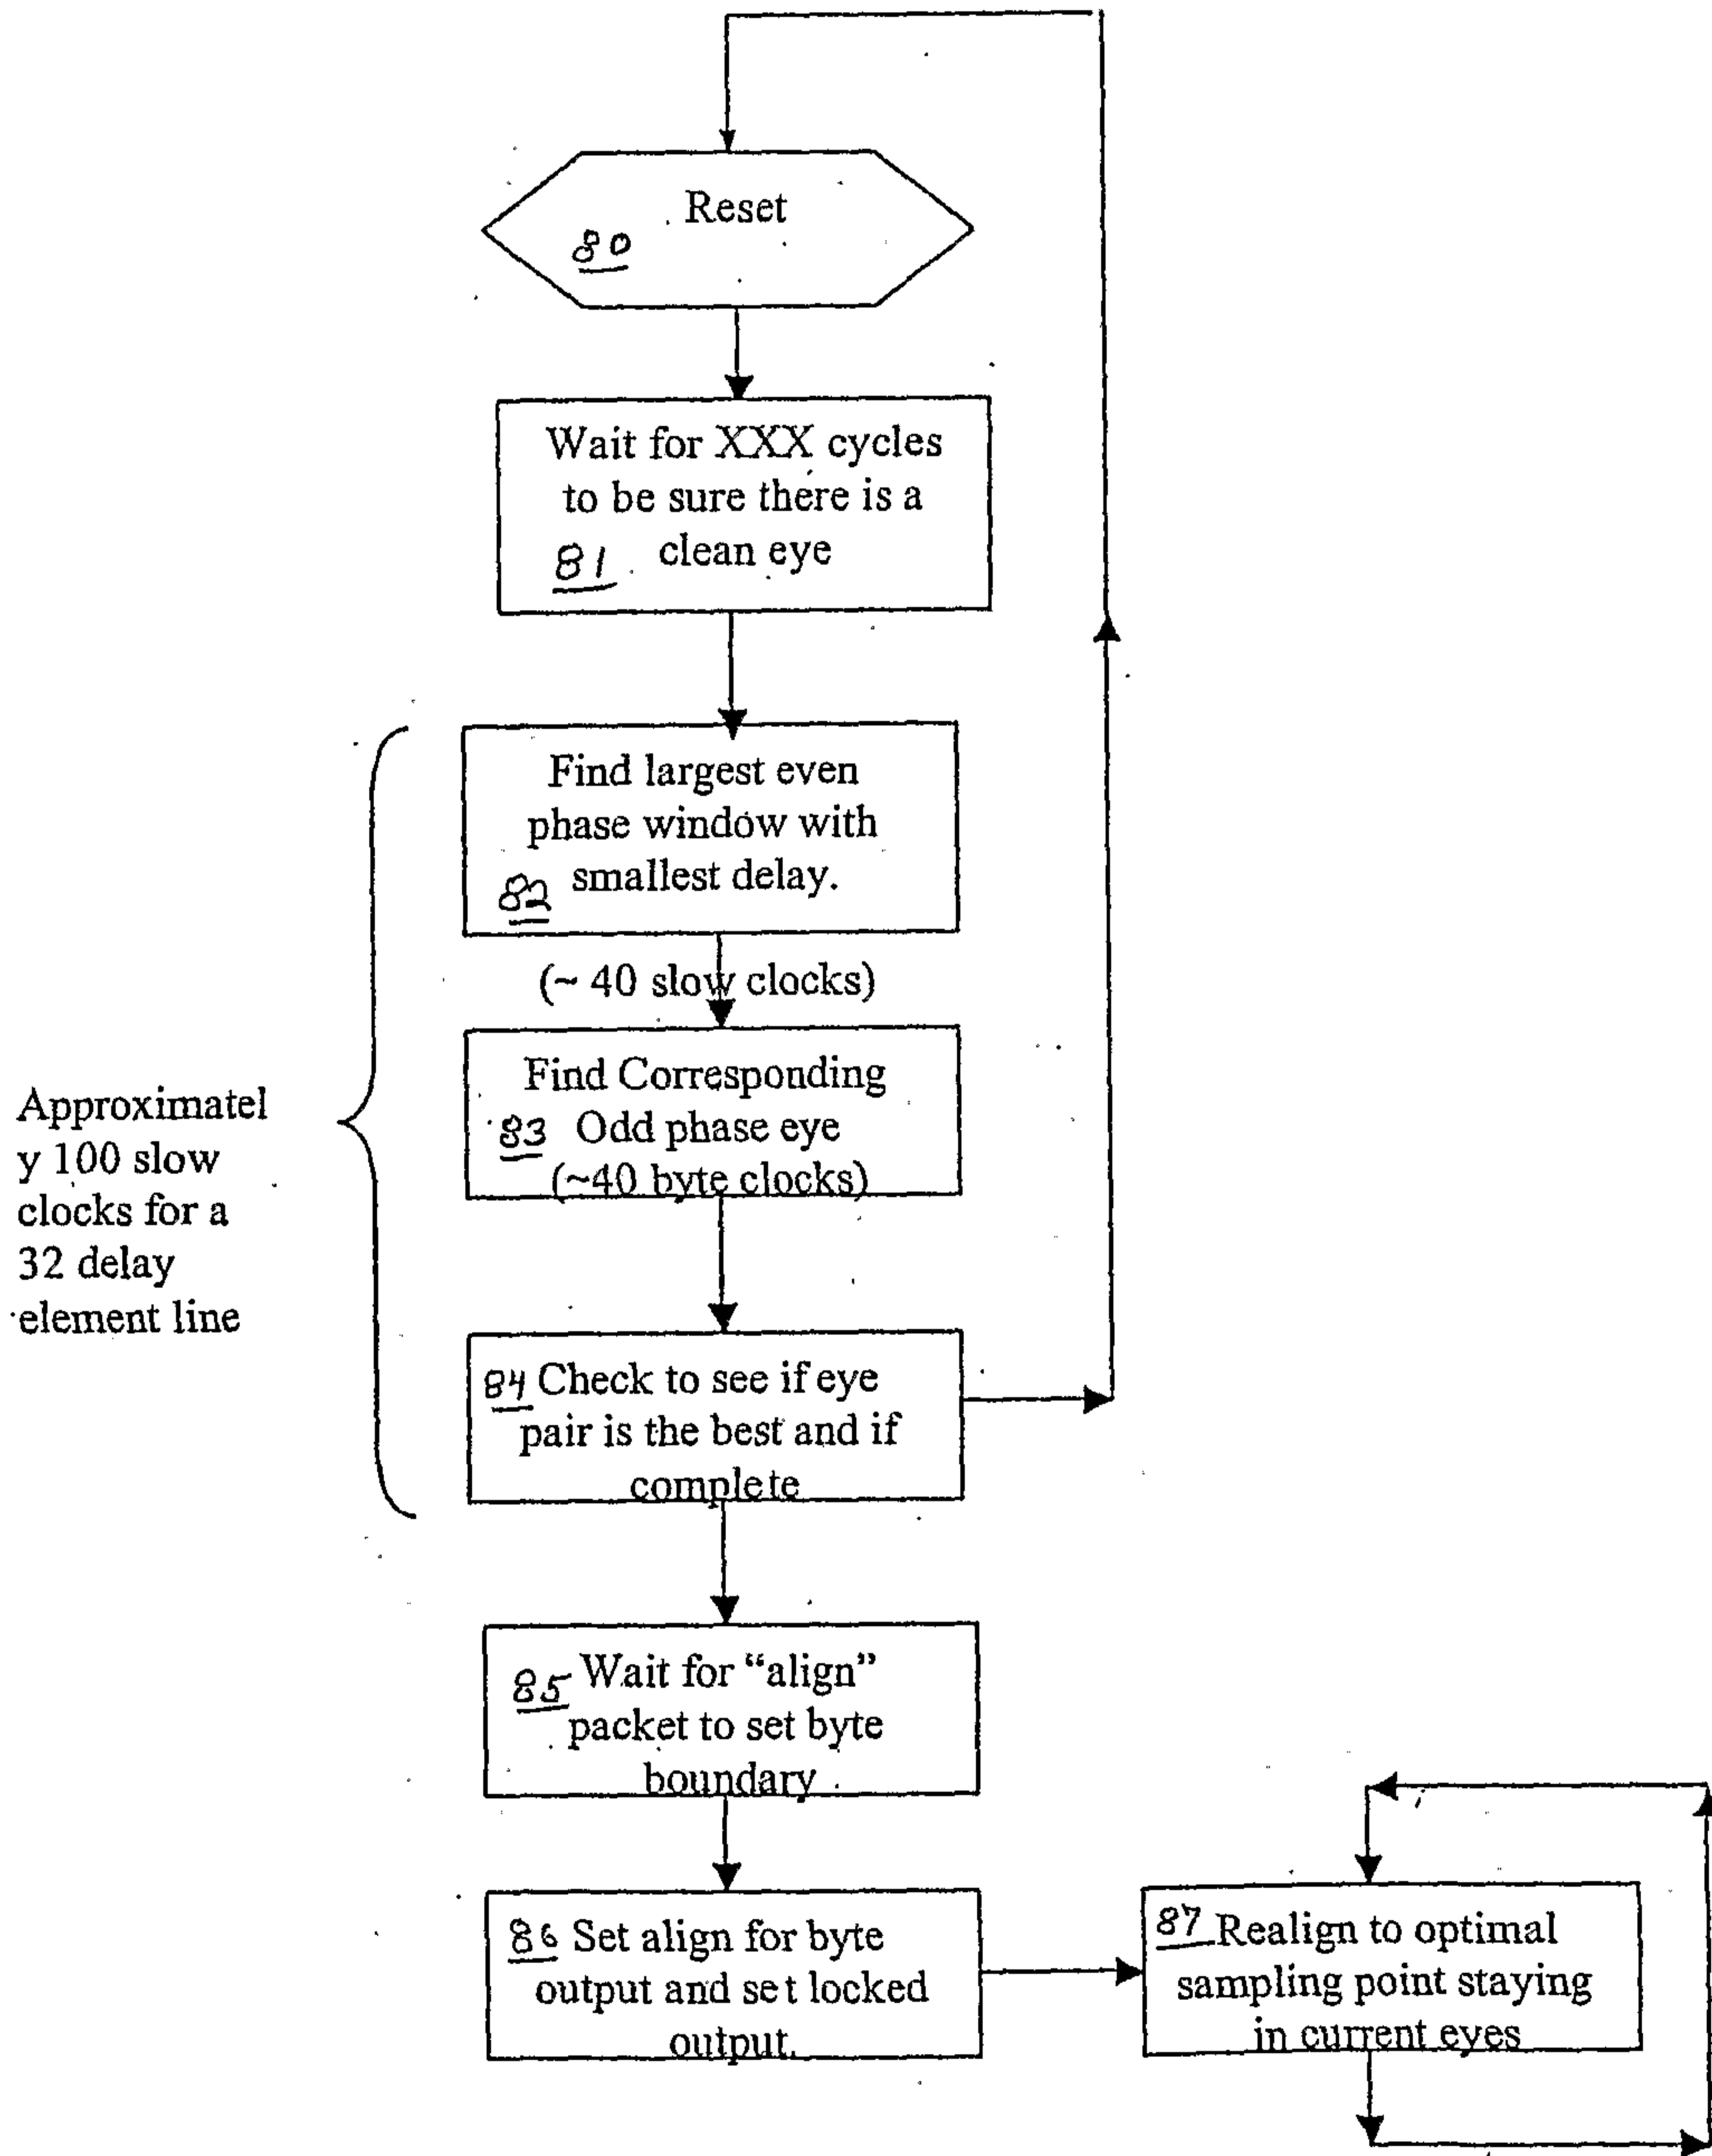
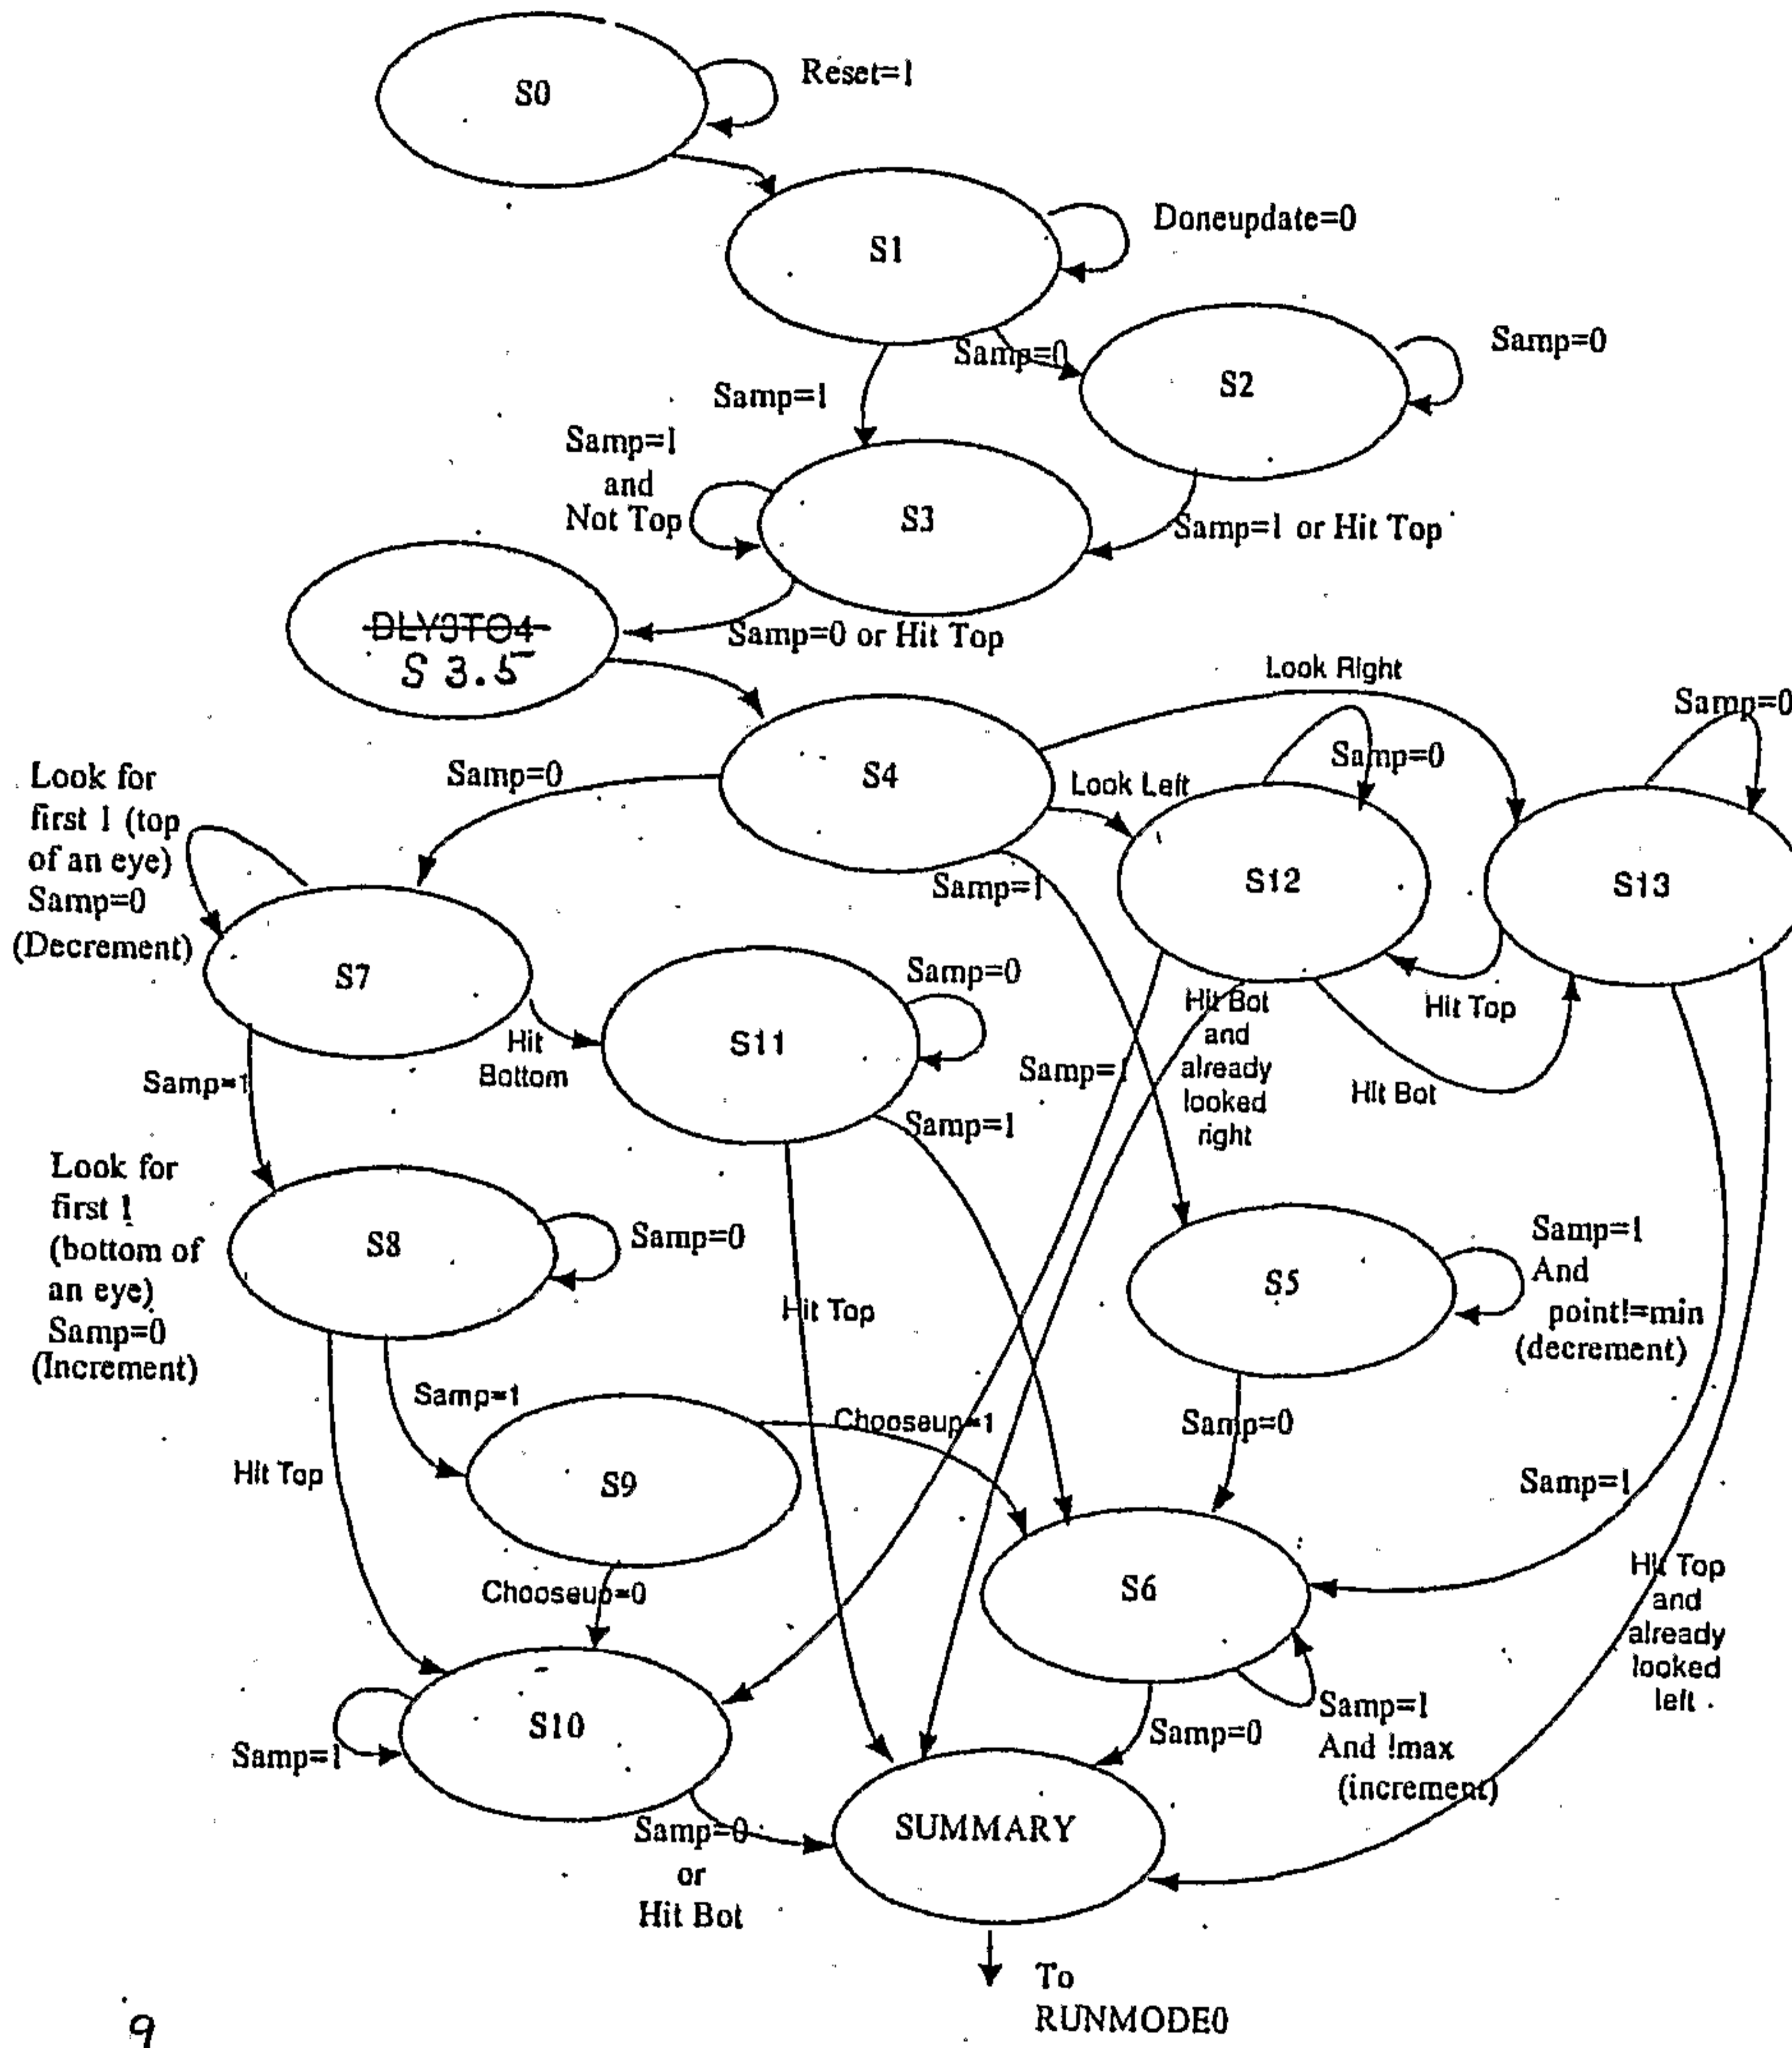


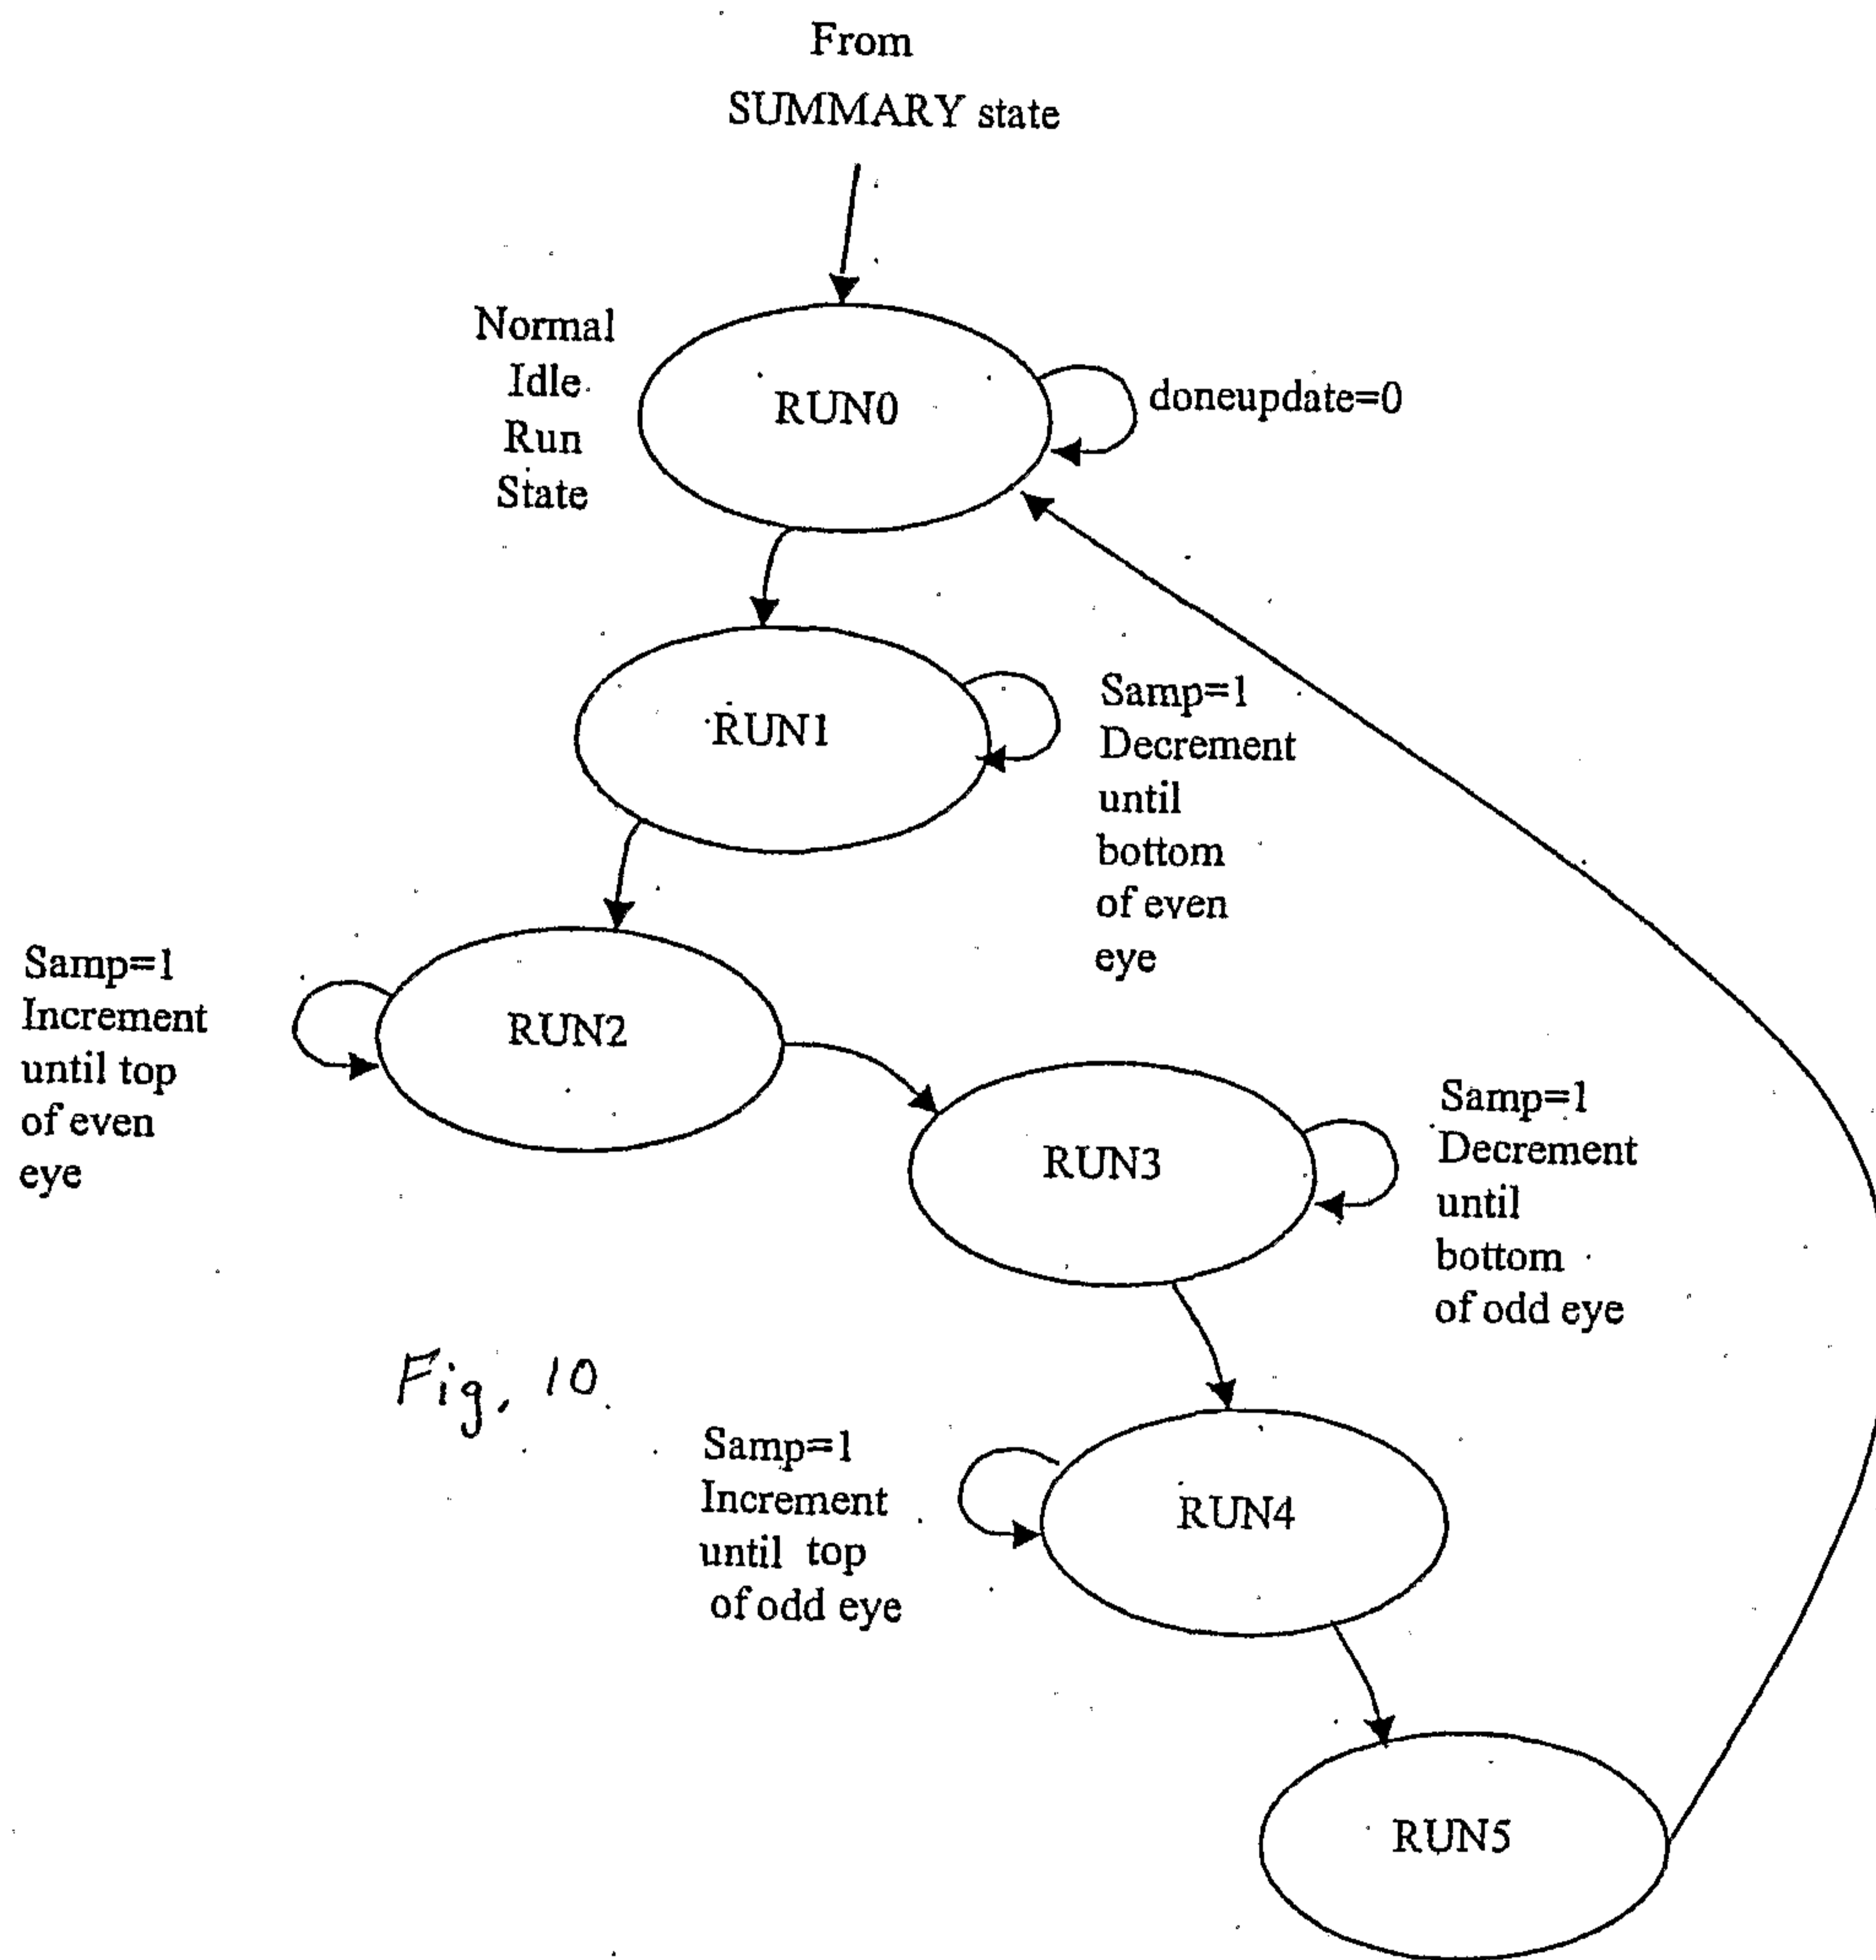
Figure 7 Serial Bit Combining and Byte Align Logic



8
Figure 8 Eye Detection Flow Diagram



9
Figure 9 Two Bit Macro State Diagram 1 of 2



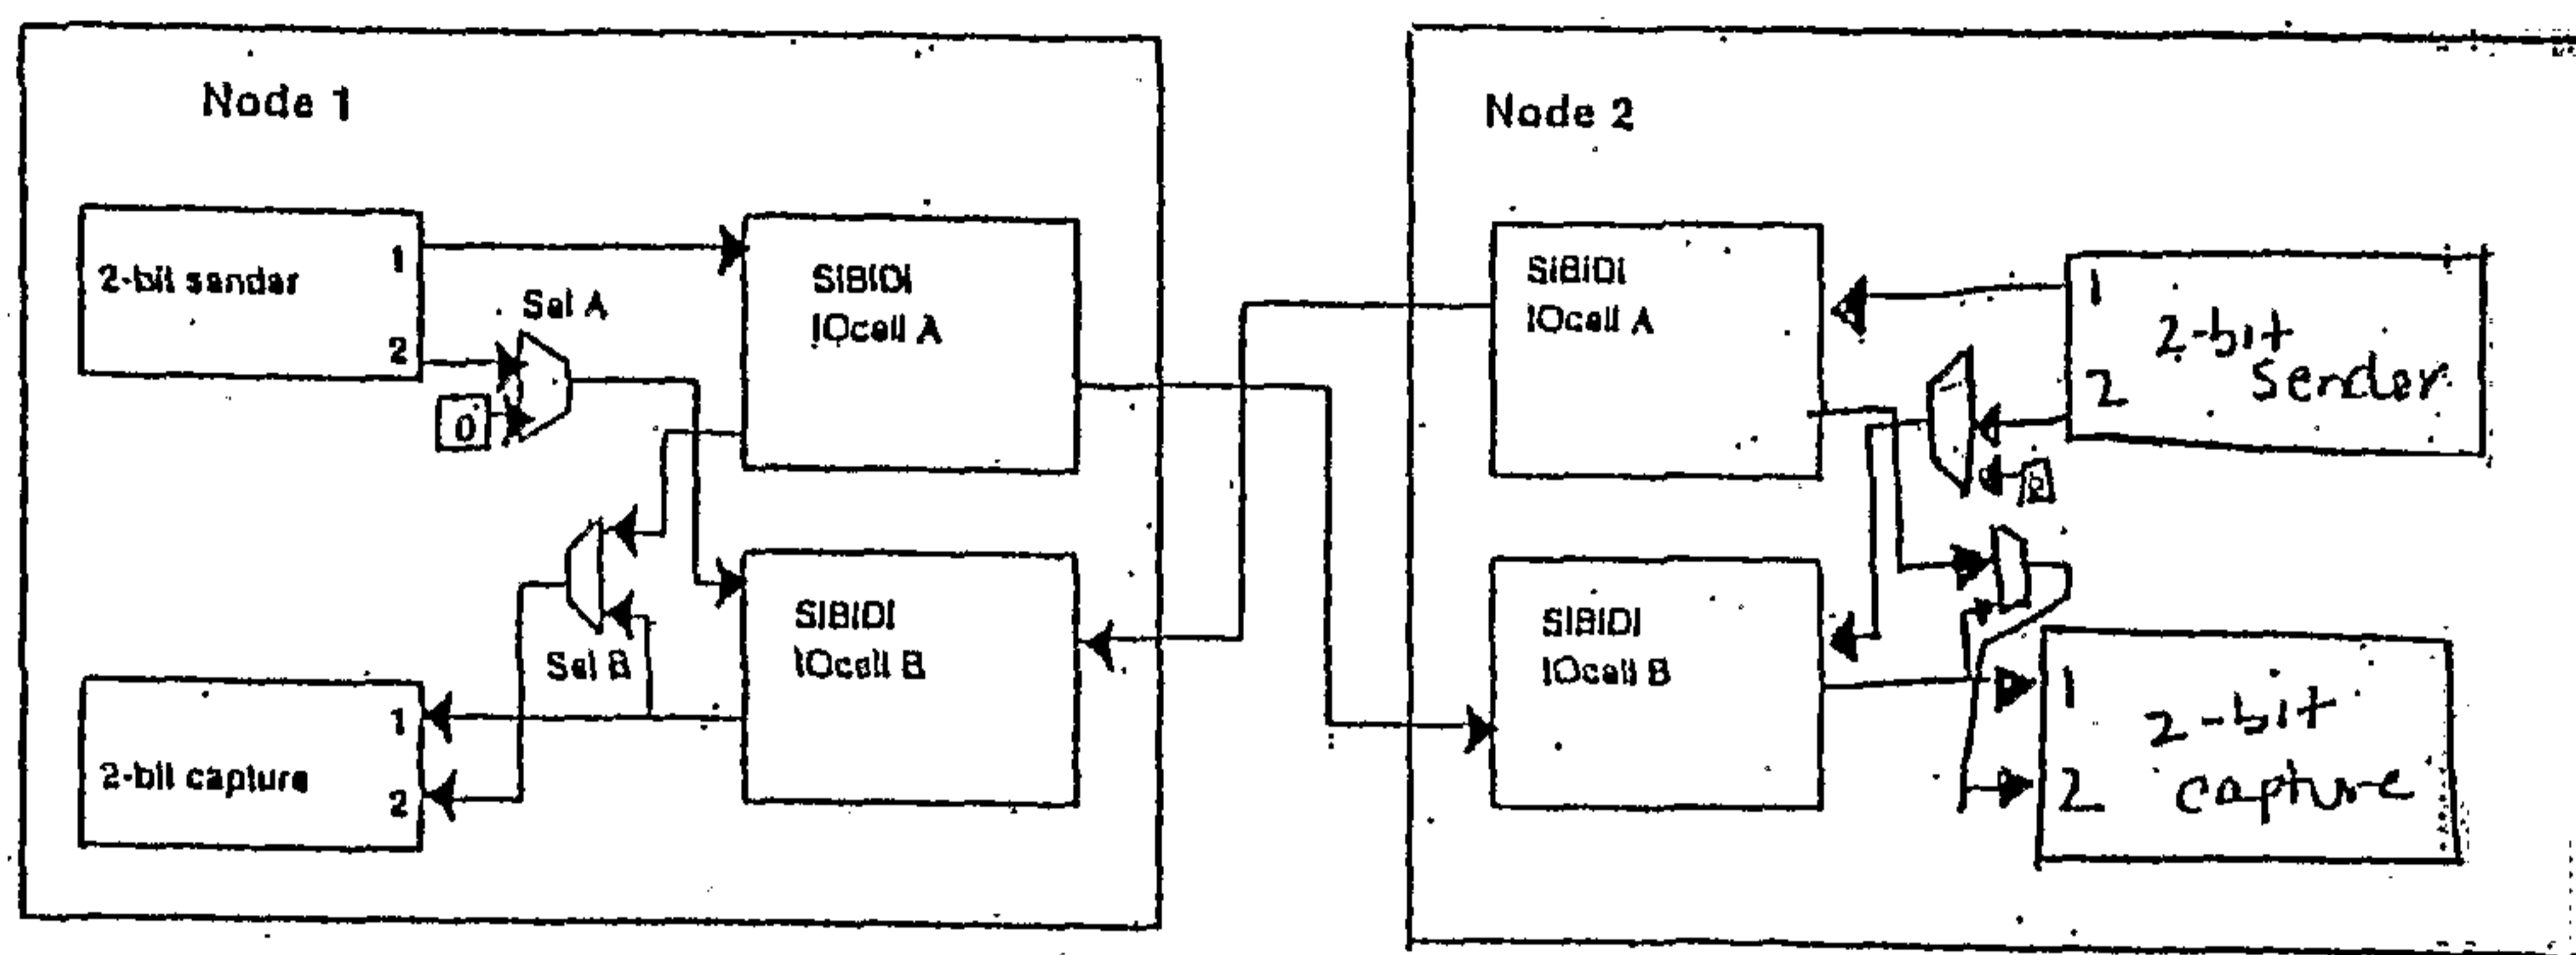


Figure 11

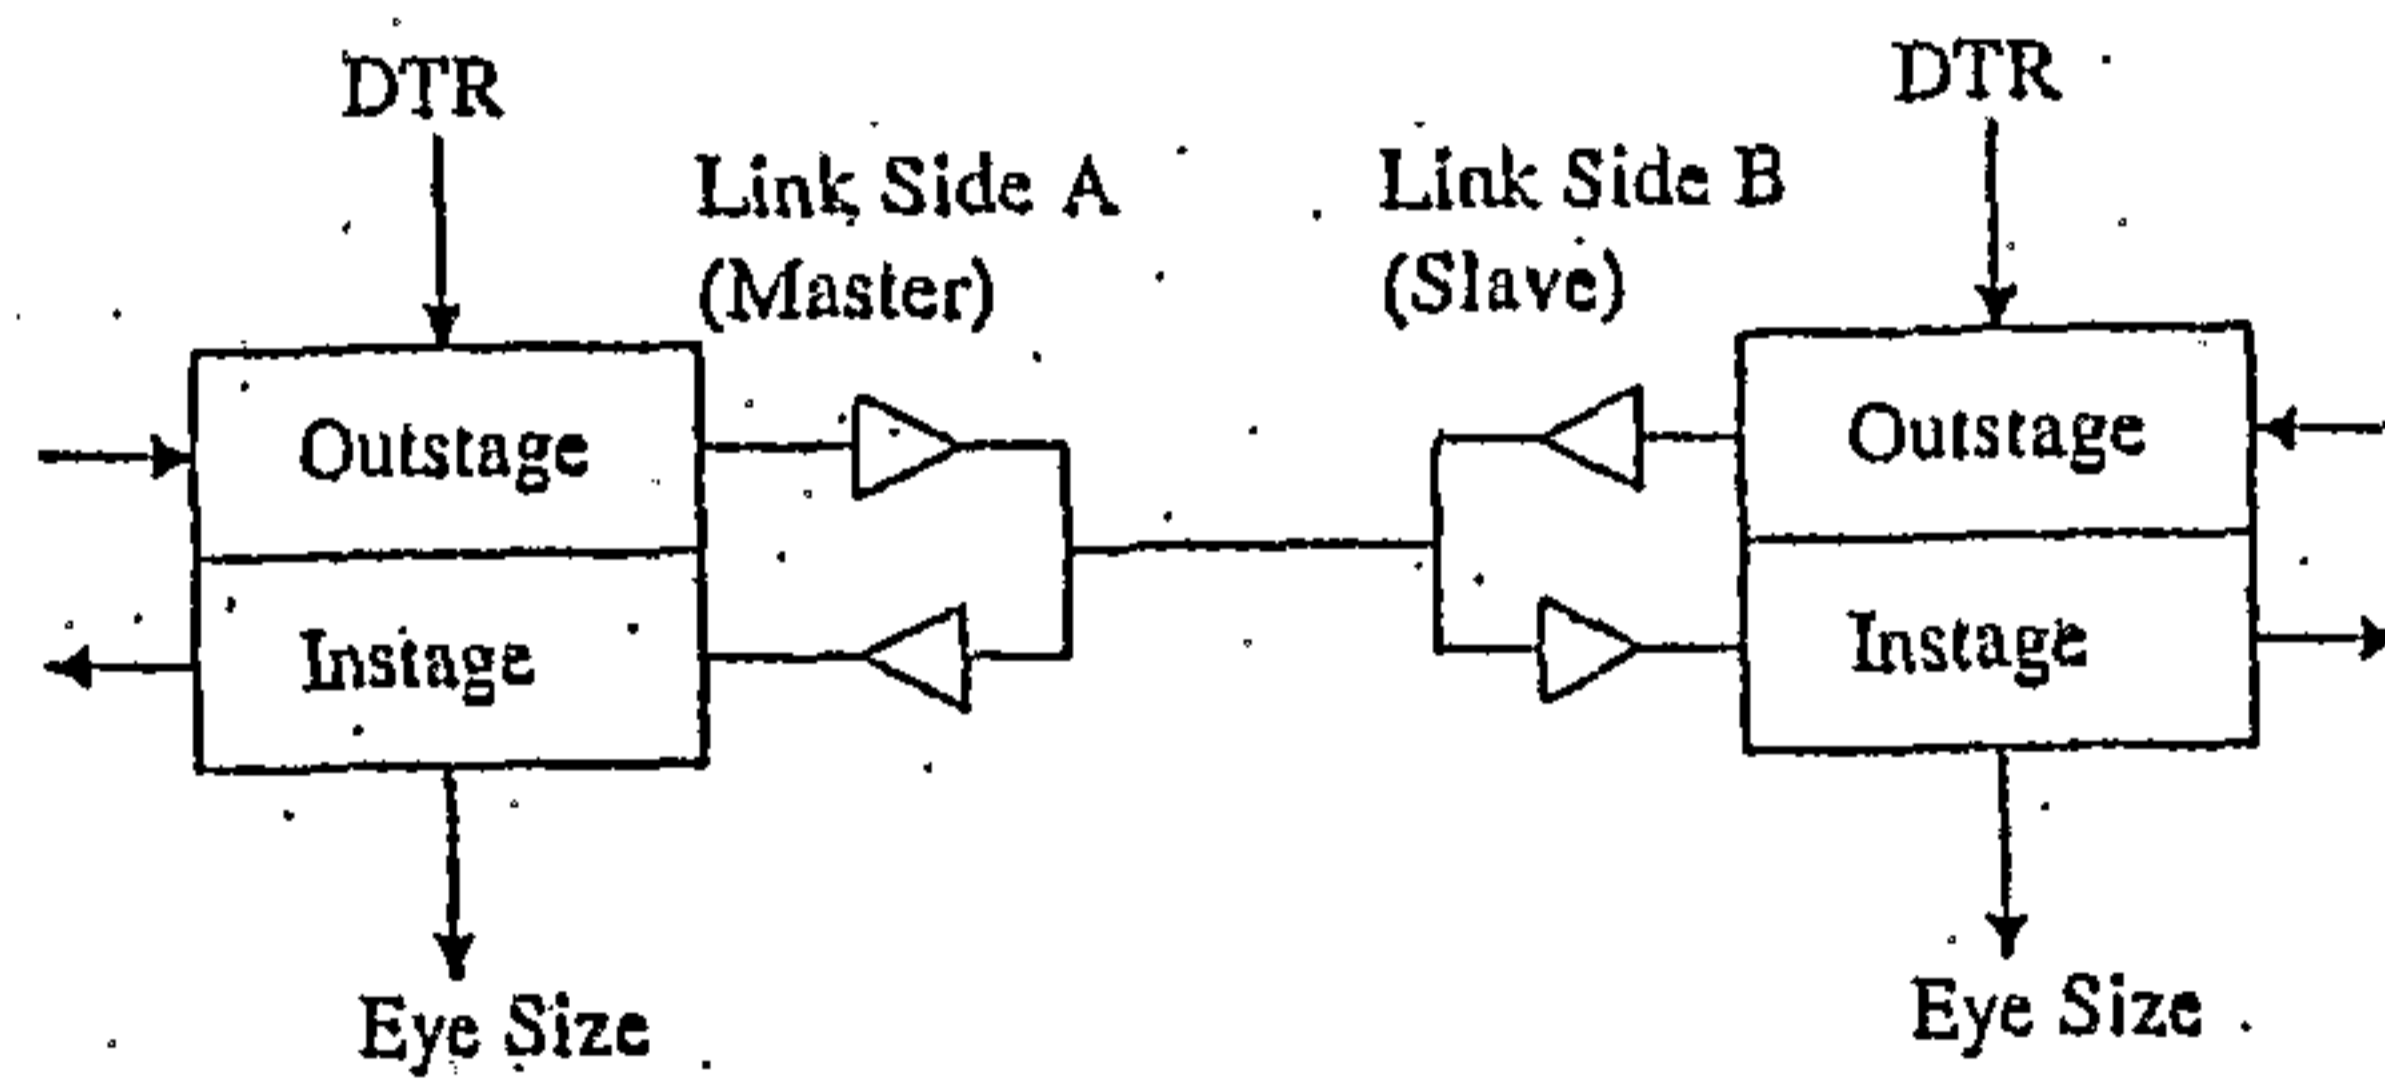


Figure 2:

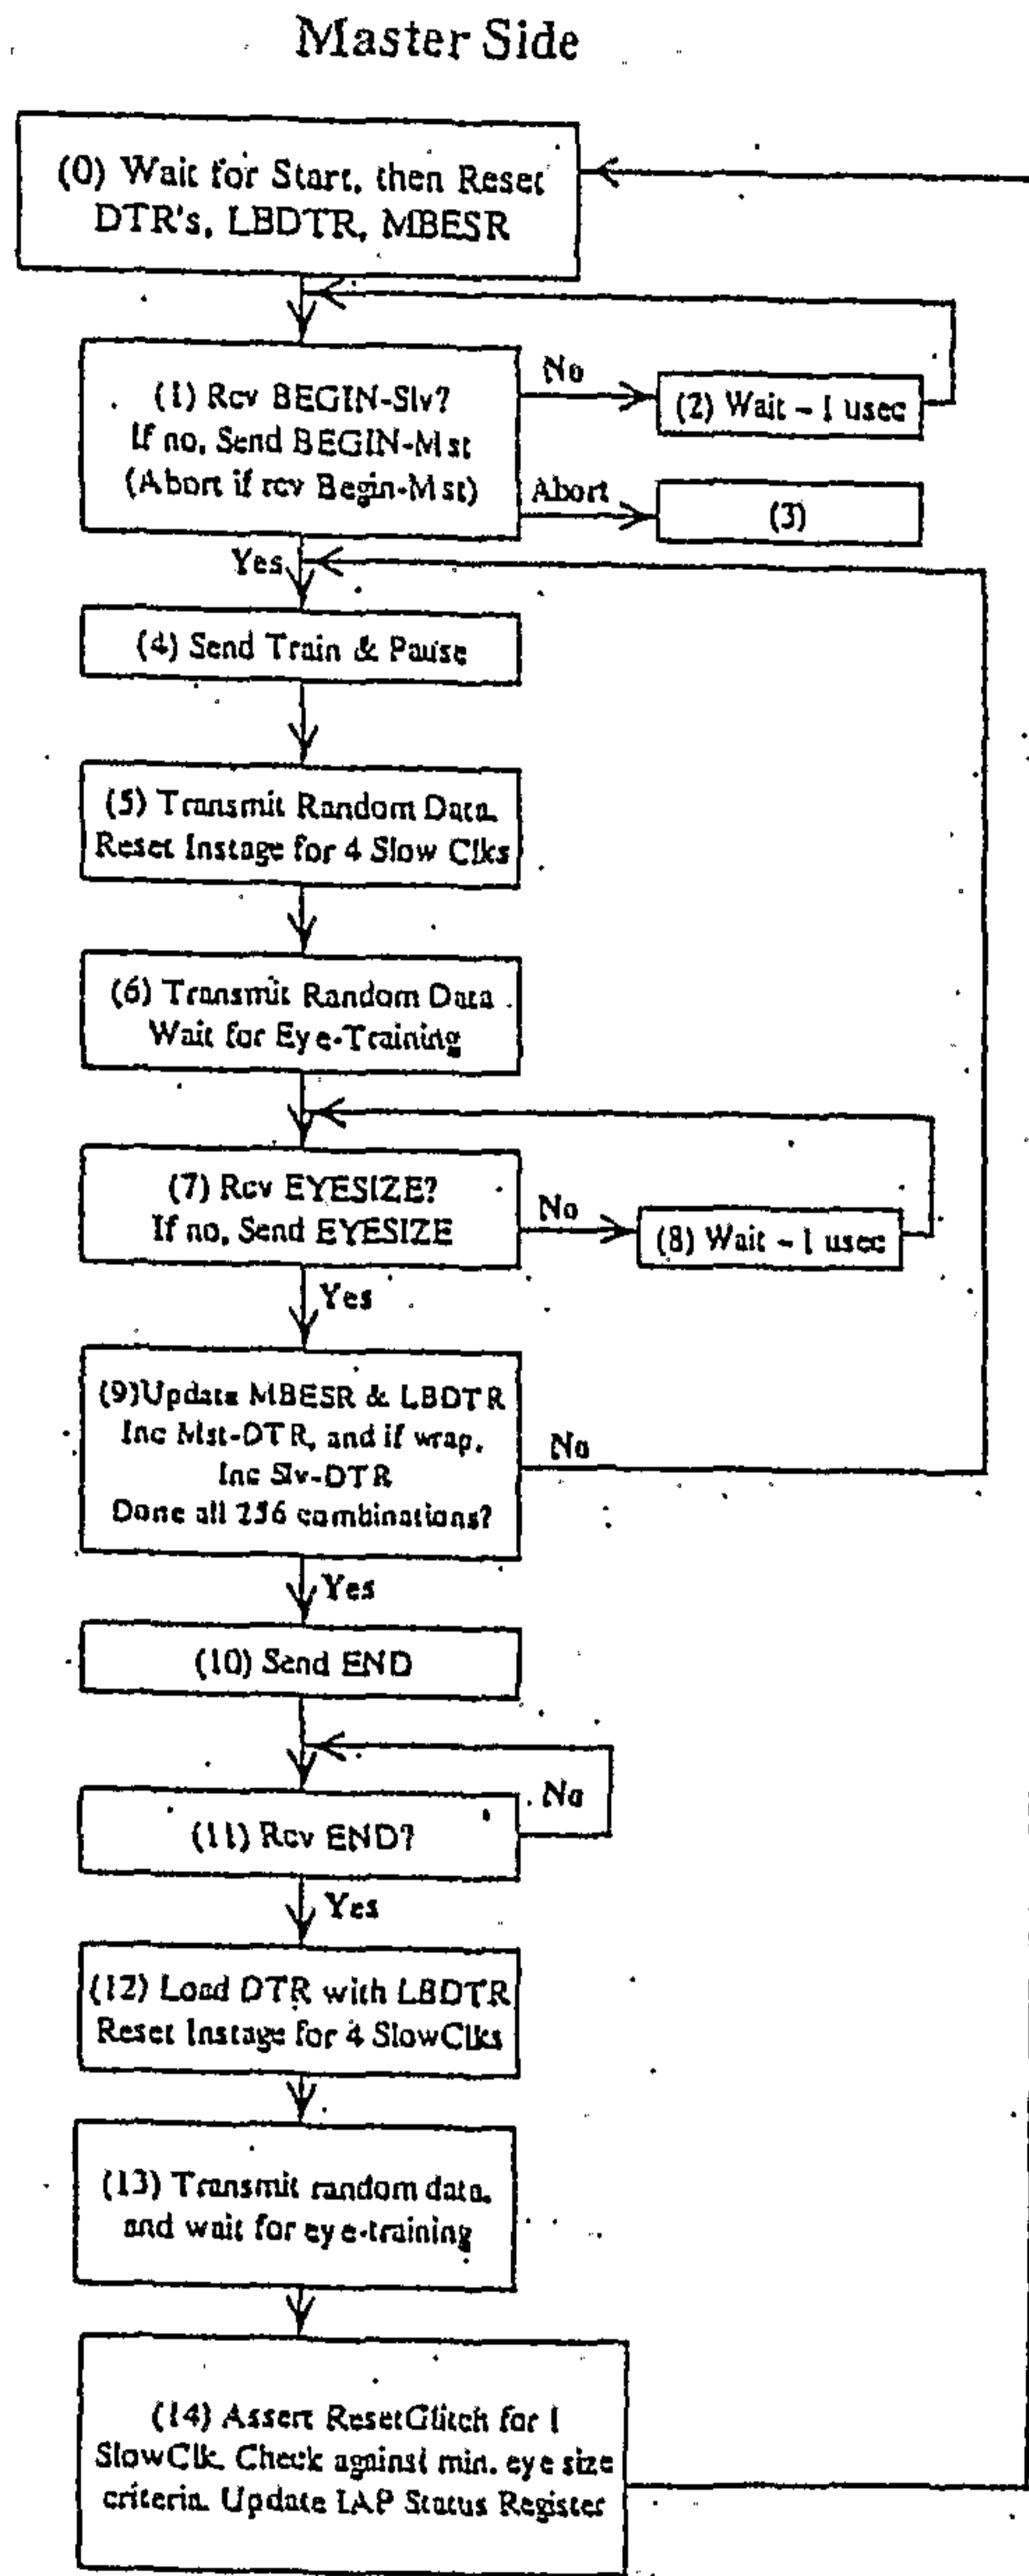


Figure 3A

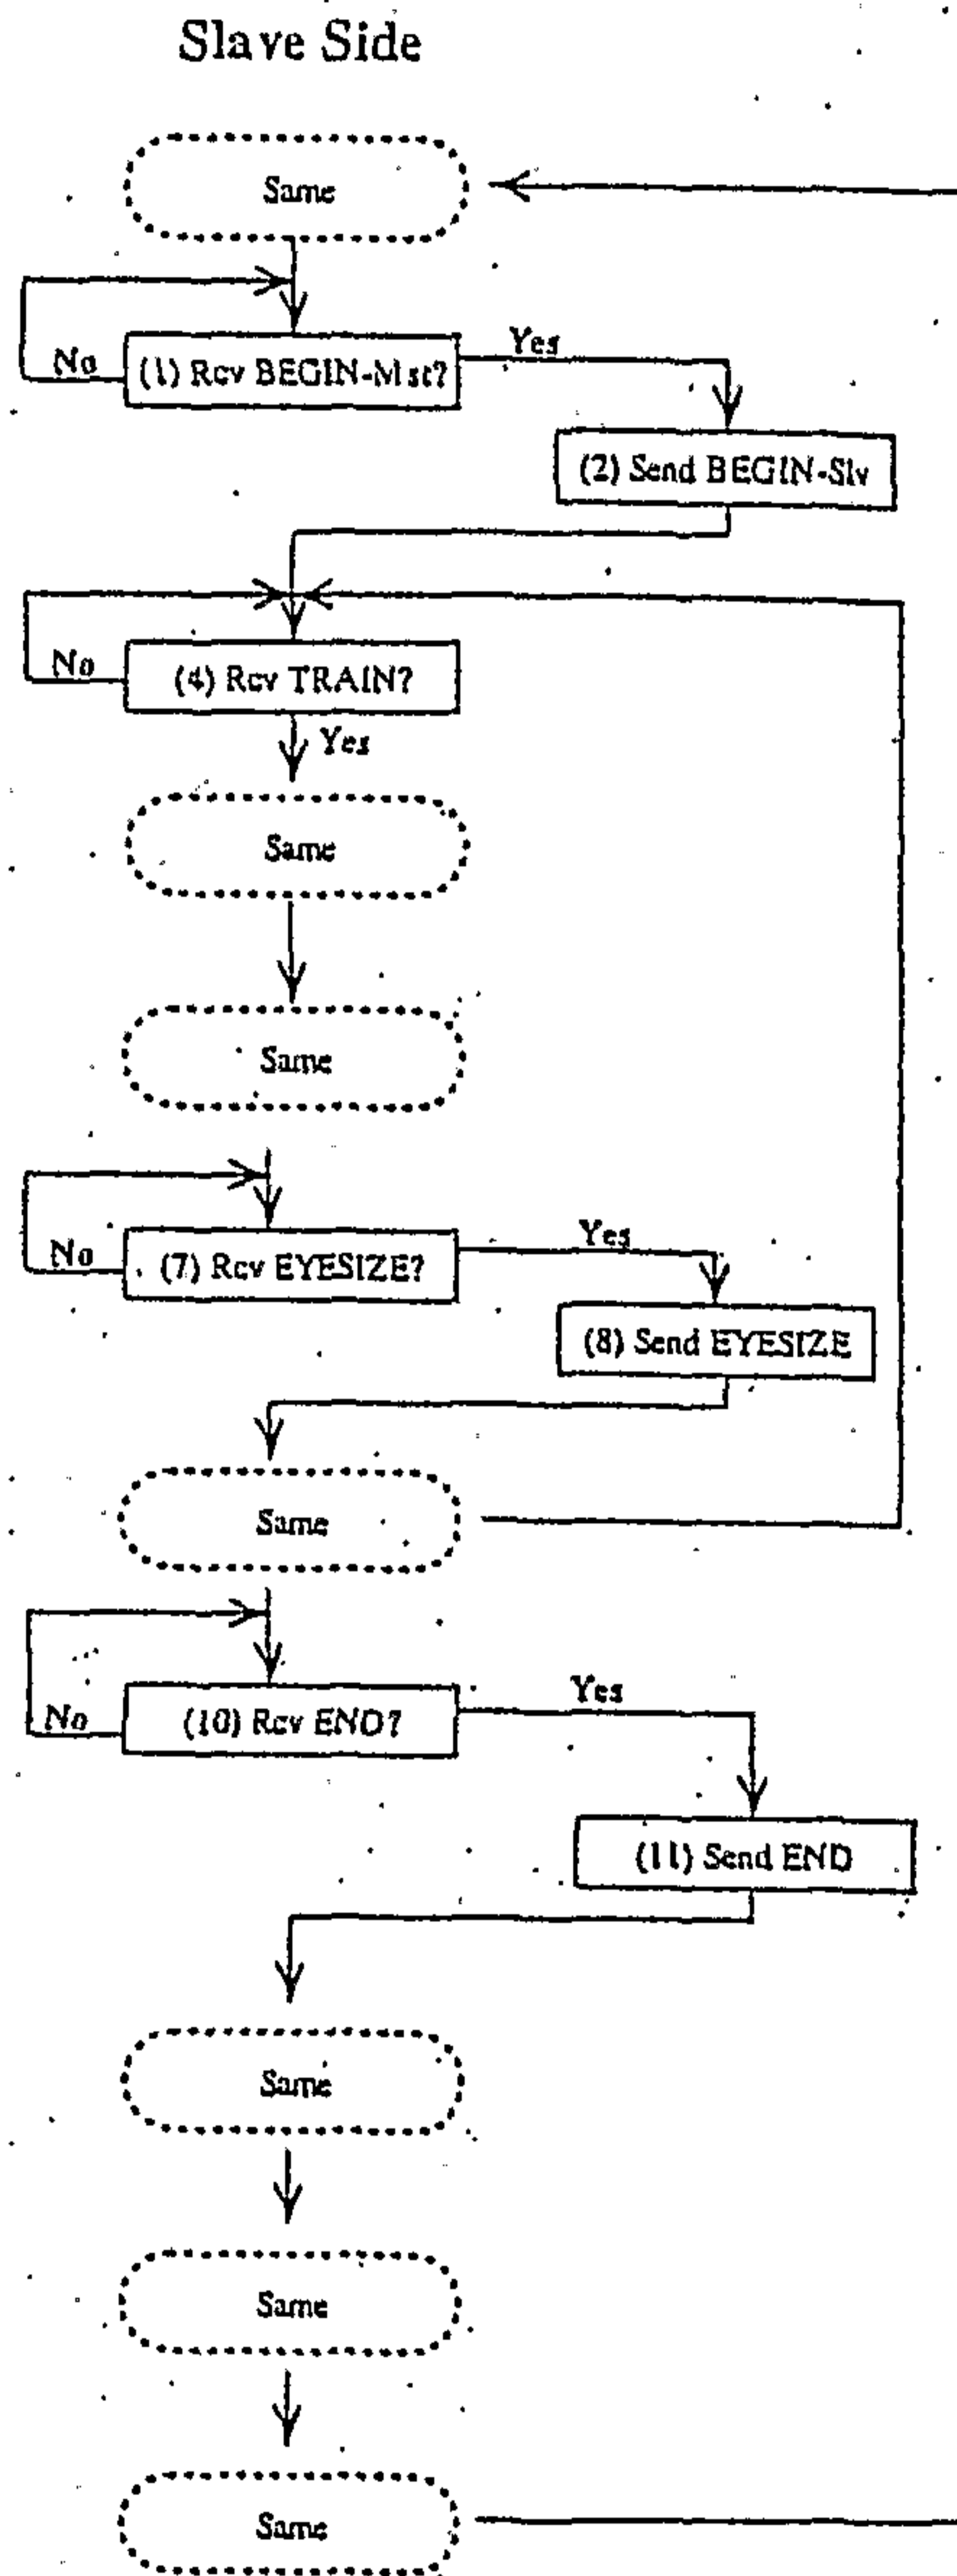



Fig. 13B



History Logic