Design considerations of a static LSSD polarity hold latch pair

by A. Correale

There are many considerations relating to the design of a static LSSD polarity hold latch pair. High performance, low power dissipation, small size, and stability are some of the major requirements for a good design. The engineering trade-offs needed to ensure that all goals are met are discussed.

Introduction

Static LSSD [1] polarity hold latches have been used in conjunction with random logic circuits for LSI and VLSI microprocessors [2] and other IBM designs. The latch configuration proposed in [3] provides the foundation for these circuits.

As with any design, there are many considerations and trade-offs which need to be taken into account to ensure a successful design. The intent of this paper is to document the technical details associated with the design of this latch type and to identify performance improvements to the basic latch structure while minimizing power dissipation and area.

In the following sections, the latch basic operation is first presented, including the function of each of the devices. Then, design goals, latch device size considerations, and physical layout are discussed in subsequent sections. This is followed by descriptions of the L1-L2 latch pair charge redistribution phenomenon and sections on latch output considerations, push-pull output buffer polarity choice, and latch scan-output considerations. The power dissipation and performance characteristics of the L1 and L2 latches are discussed in the section on power/performance characteristics. Finally, a summary section recaps the highlights of the paper.

Latch basic operation

The basic latch pair is made up of two similar latches interconnected via a low (zero-V_t) threshold transfer device (an enhancement device having a threshold voltage near ground potential). Each latch output utilizes an inverting push-pull circuit to eliminate latch degradation due to loading. The latch pair schematic design is shown in **Figure 1.**

With respect to Fig. 1, the application of a valid clock signal to the gate of the zero-V_t transfer device (device 8) allows data present on the data input of the L1 latch to be transferred to the latch internal node (node K).

The LSSD scan-input port (device 9) is held inactive during normal operation and is utilized only during testing. Since the zero-V_t device is used for data transfer rather than an enhancement device, the transfer of data occurs at a

lower gate voltage, thereby improving performance and allowing a lower gate Least Positive Up Level (LPUL) of 3.5 volts (assuming a V_{dd} minimum limit of 4.5 volts).

The data at the K node of the latch attains a voltage which is the lesser of the voltages of the data at the input or one threshold drop below the gate voltage while the gate is active. Since the zero-V_t device threshold is approximately zero volts, the voltage at the K node can attain the gate voltage of the zero-V_t transfer device when loading an up-level.

The data now at the K node then propagates through two inverters, reaching the B node. The depletion device connecting the K and B nodes acts as a high-impedance device to limit current flow from K to B nodes and as a low-impedance feedback path required for data retention.

This can be seen by evaluating the mode of operation of the feedback device. When K-node voltage is high ($\simeq$3.5 volts) and B-node is low ($\simeq$0 volts), the device is operating in the saturated region. Hence,

$$i_{SAT} = 0.5\gamma(W/L)(V_{gs} - V_t)^2 = 0.5\gamma(W/L)(V_t)^2.$$

When the B-node voltage is high ($\simeq V_{dd}$) and K-node is low ($\simeq$0 volts), the device is operating in the linear region. Hence,

$$i = \gamma(W/L)(V_{gs} - V_t - V_{ds}/2)V_{ds},$$

but $V_{ds} = V_{gs}$; therefore,

$$i = \gamma(W/L)(V_{ds}/2 - V_t)V_{ds}.$$

As the threshold voltage is $\simeq -2$ volts for this depletion device and V_{ds} is positive,

$$i_{linear} \simeq 0.5\gamma(W/L)(V_{ds} + 2|V_t|)V_{ds}.$$

Comparing both expressions,

$$i_{SAT} = 0.5\gamma(W/L)(V_t)^2$$

and

$$i_{linear} = 0.5\gamma(W/L)(V_{ds} + 2|V_t|)V_{ds}.$$

We can see that the device operating in the linear region has a higher current-supplying capability than when it is saturated. For a $V_{ds} = 3.5$ volts and a $V_t = -2$ volts, i_{linear} is 6.5 times greater than i_{SAT}.

Hence, the device appears as a high-impedance device when looking into the K node (saturated) and a lower-impedance device when looking into the B node. This characteristic is very important to overall operation of the latch. The consideration of device size is discussed in subsequent sections of this paper.

The data now present at the B node of the latch are fed back to K node via the feedback device and the clock pulse can now be removed, leaving the latch in a stable state.

The original latch circuit, as presented in [3], utilized enhancement mode transistors for transfer gates and used the latch internal nodes as outputs. With the advent of zero-V_t devices, this basic latch structure can be modified to improve its overall power/performance characteristics. The use of zero-V_t devices for data transfer as described above allows the transfer of data to occur at lower gate voltages, thereby improving performance. In addition, the use of a push-pull output buffer which uses the first inverter of the basic latch structure as its predrive, and a zero-V_t device for the output pull-up, provides load isolation to the latch, thereby improving latch setup time and output response. Since the buffer isolates the capacitive load, the power of the latch internal stages can be reduced while the overall performance is improved.

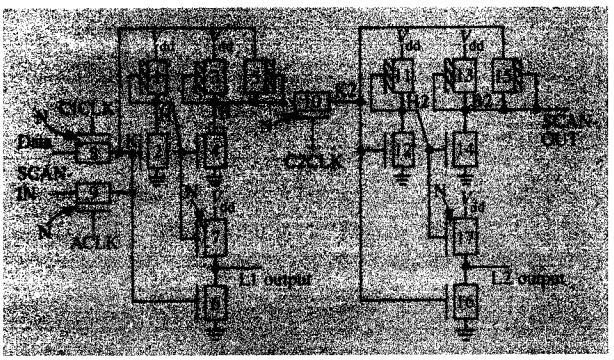

Figure 1

L1-L2 latch-pair schematic.

The push-pull buffer configuration used dissipates very low power and provides very good performance. As a result, the overall latch power/performance is greatly improved. More detailed discussions regarding the characteristics of the latch and the push-pull buffers are given in later sections.

Design goals

The device sizes selected for the latch implementation must provide good performance at the lowest power level and be insensitive to process variation while occupying the smallest area. These goals are common to all circuit designs and are sometimes mutually exclusive. For example, in an NMOS technology, good performance is usually associated with high power dissipation circuits requiring large devices and occupying large areas.

Since the latch is being used as a storage element whose output performance is essentially determined by its output buffer, the power of the latch can be minimized. The use of the zero-V_t device in the push-pull buffer dissipates very low power while providing good performance. Hence, for this latch configuration, it appears that all goals can be met.

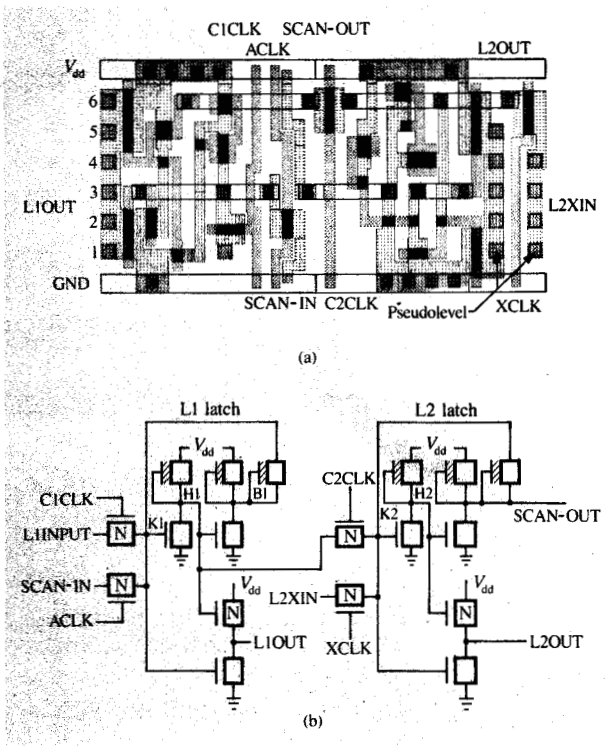

Figure 2

L1/L2 latch: (a) physical layout; (b) schematic.

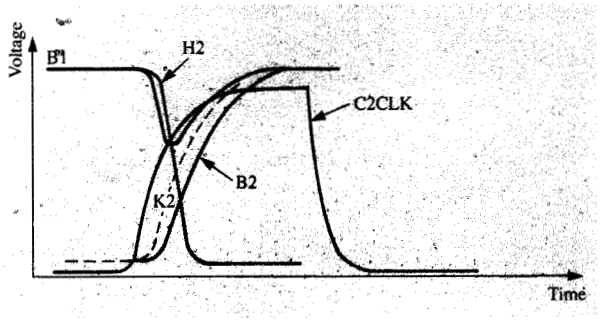

Figure 3

L1-L2 latch charge redistribution waveforms.

Table 1 Best-case to worst-case power spreads.

Depletion device size, artwork (device width/length) (μm)	*Power dissipation* (μW)		
	BC	*WC*	*WC-BC*
4.3/3.0	234	935	701
5.9/4.0	274	812	538
7.4/5.0	294	764	470
8.9/6.0	307	744	437

Latch device size considerations

The depletion device lengths are specified above the technological minimum limits to improve best-case to worst-case power/performance spreads [2]. **Table 1** illustrates the best-case to worst-case power spreads for an inverter with a nominal power dissipation of 0.480 mW as a function of the depletion device geometry. As can be seen, a depletion device size of 8.9/6.0 offers a best-case to worst-case power dissipation spread improvement of 37.7 percent when compared to a device size of 4.3/3.0. Since performance spreads are closely tied to power dissipation spreads in an NMOS technology, a tighter best-case to worst-case power spread results in tighter performance spreads.

Hence, for the same nominal power dissipation, the worst-case power and performance are greatly affected by the device size chosen. Since the "goodness" of the design is measured by its worst-case characteristics, device size selection is of prime importance.

The depletion device sizes (devices 1, 3, 11, 13) are selected to provide the best performance at the lowest possible worst-case power dissipation. The internal nodes of the latch have relatively low values of capacitance (less than 0.2 pF) and hence high power circuits are not required. Typical nominal power dissipation of 0.1 mW to 0.15 mW provides good performance. It should be noted that for such low-power latch inverter stages the maximum L2 B-node external capacitance should be limited to 0.1 pF. Since this internal node is used only for LSSD scan-out to an adjacent latch, this restriction is easily satisfied. Performance characteristics are described in a later section.

The feedback device size (devices 5, 15) is selected to offer the best performance while limiting the current when setting the latch. Typical device size is width = 5 μm, length = 12 μm.

Physical layout

The physical layout and the schematic representation of a typical static L1-L2 latch pair with push-pull output buffers are shown in **Figures 2(a)** and **2(b)**, respectively. Signal wiring runs on metal, parallel to circuit metal power buses. Each latch pair can share a common power supply rail with an adjacent circuit. This results in saved area, since separation between circuit power busing is eliminated.

As seen in Fig. 2(a), the latches were designed with variable input/output locations and with little or no metal. The variable I/O positions permit different circuit and latch

types to be connected without the need for additional circuit separation for submetal (polysilicon, diffusion) wiring, while the lack of metal within the circuit permits the use of metal as a circuit interconnect, resulting in efficient area utilization.

The outputs of the L1 latch are available on every wiring field, whereas the L2 latch output is available in every wiring field except field 6. The L1 data input is limited to wiring field 1 for this configuration, whereas the secondary L2 data input is available on wiring fields 1 through 4. Additional information regarding circuit physical layout is available in [2]. The L1-L2 latch pair occupies an area of 8328 μm square, which satisfies the design goal.

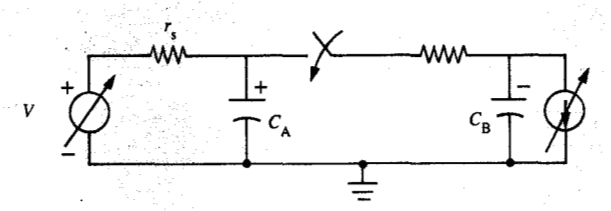

Figure 4

Equivalent circuit to illustrate charge redistribution.

L1-L2 latch pair charge redistribution

It is important to evaluate the loading effect of the L2 latch on the performance of the L1. **Figure 3** illustrates the timing diagram for an L1-L2 pair. As shown by this diagram, the use of a very-low-power L1 inverter stage results in a charge redistribution between L1-L2 pairs when the C2 clock is activated. That is, the voltage of the B node of the L1 latch (B-node transfer) "dips" from V_{dd} to approximately 2.5 volts as the C2 clock turns on the L2 data transfer device (this assumes that the L2 latch is currently in a "0" state).

For a B-node transfer (as illustrated in Fig. 1), the size of device 3 determines the severity of this charge redistribution (voltage dip, for a fixed L2 capacitive loading). This phenomenon can be simply described by the circuit in **Figure 4.**

Referring to Fig. 4, C_A is a very small internal L1 latch capacitance; C_B is meant to model the large L2 input capacitance; the series resistor and voltage source are meant to represent the L1 depletion device; the switch represents the zero-V_t transfer device; and the voltage-sensitive current source represents the leakage through the L2 high-impedance feedback device to ground via the B2-node pull-down device.

Initially, C_A is charged to $+V$, signifying that L1 is set at a "1." C_B is initially at 0 volts, signifying that L2 is reset at a "0." Upon closing the switch (activating the C2 clock), charge from C_A and the finite current source begin charging C_B. The voltage across C_A, however, begins to fall as the capacitor C_B begins to charge. The degree to which this voltage drops is a function of the size of C_A and C_B, the ability of device 3 to supply charge to these capacitors, and the impedance of the L2 feedback device. This charge redistribution effect can have significant consequences, the worst being the resetting of the L1 latch, thereby losing data.

Care must be taken in selecting the proper device sizes for the L1 latch transfer stage, the high impedance of the L2 feedback device, and the overall L2 input capacitance. Since this voltage dip affects the overall L2 performance, all L2 performance simulations should be done using the L1 latch.

To lessen the degree of the voltage dip associated with an L1-L2 data transfer, the following four steps can be taken:

1. Increase the power dissipation of the L1 latch inverter being used for transfer. This allows more current (charge) to be supplied, thereby resulting in a smaller voltage dip.
2. Increase the size of the parasitic capacitance of the L1 latch (C_A). This enables more charge to be available, thereby reducing the voltage dip due to capacitor divider action.
3. Decrease the capacitance seen looking into the L2 latch (C_B). This helps reduce the voltage dip due to capacitor divider action.
4. Provide the highest-impedance feedback device for the L2 latch. This limits current flow.

Of these four steps, the first is the most easily accomplished. The second requires additional area and results in performance degradation of the L1. The third is limited by the physical size of the devices of the L2 latch. Although steps two and three are not the best alternatives, designing the L1 and L2 stages with these in mind makes for a better design. Step four requires a larger area but, more significantly, adversely affects overall L2 performance.

Consistent with step one, the L1 inverter stage being used for transferring data to the L2 dissipates ≈50 percent more power than the other L1 inverter stage. Hence, for a B-node transfer, device 3 would dissipate 50 percent more power than device 1. On the other hand, for an H-node transfer, device 1 would dissipate more power than device 3. This statement is true for very-low-power designs. For designs using high power, this may not be necessary; sufficient current will be available to compensate for the charge redistribution effect.

The zero-V_t transfer gates (devices 8, 9, 10) should have an aspect ratio (width/length) of at least 5 to ensure good performance and dc voltage levels.

The high-impedance feedback device should be selected such that the best performance is attained. An aspect ratio of approximately 1:4 works nicely. In some instances, the L2 high-impedance feedback device can have an aspect ratio of

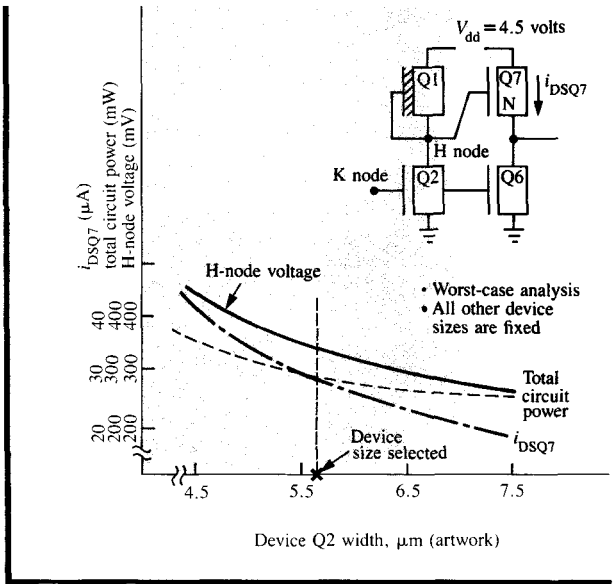

Figure 5

Latch push-pull dc characteristics versus Q2 device size.

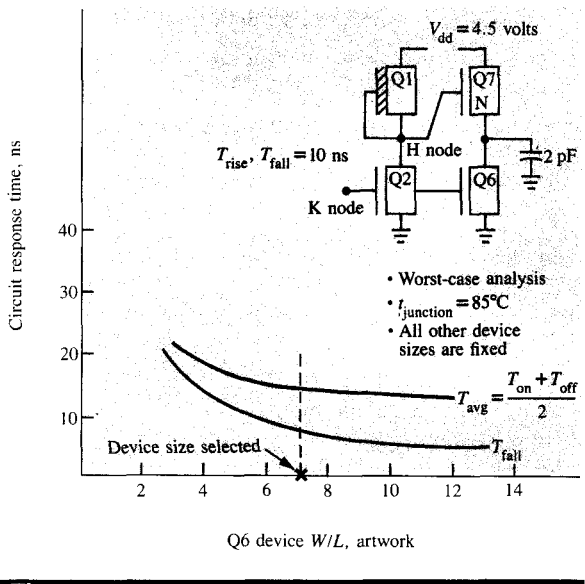

Figure 6

Latch push-pull buffer response versus Q6 device size.

1:5. The sizes of devices 2 and 4 are selected to guarantee good dc design margin and ac performance.

Device 2 must be large enough to ensure an adequate down level to the zero-V_t pull-up device of the output buffer to limit power dissipation of that stage. Typical inverter stage gains are 7.6 and 3.1 for stages 1 and 2 of a latch with a B-node transfer, respectively. Typical inverter stage gains are 5.2 and 4.0 for stages 1 and 2 of an L1 latch with an H-node transfer, respectively.

The push-pull device sizes are selected to provide the best power/performance trade-offs. The use of the zero-V_t device in this configuration provides very good performance at essentially zero power (device leakage), while providing better up levels than those attainable by using an enhancement device. In order to keep the leakage current low, a nonminimum device length is selected. In addition, the device sizes of inverter stage 1 are selected to provide a low enough down level to keep the zero-V_t device essentially off. The pull-down device size of this push-pull buffer is selected to provide the best performance possible. Typical device sizes for this buffer are width/length ratios of 5 for both devices.

Figure 5 illustrates the latch push-pull dc characteristics versus predrive pull-down device size. A device width of 5.6 was selected due to physical layout restrictions. Layout permitting, this device size should have been specified at 6.5 to achieve optimum power dissipation. The ac characteristics of this buffer versus output device sizes are illustrated in **Figures 6** and **7**. Figure 6 shows the worst-case average delay and output fall time of the push-pull buffer as a function of the output pull-down device (Q6) *W/L* ratio for a fixed capacitive load of 2.0 pF. The device artwork *W/L* ratio selected was 7.1. This selection provides performance which is optimized. Figure 7 shows the worst-case average delay and output rise time as a function of the output pull-up device (Q7) *W/L* ratio for the same capacitive load. The device artwork *W/L* ratio selected was 5.7 due to layout restrictions. However, this device size results in an output rise time response which is nearly optimum.

Latch output considerations

The latch output utilizes an inverting push-pull buffer which shares the first inverter of the latch as its predrive circuit. This technique reduces the device count and saves silicon area. The L1-L2 data transfer is directly through a zero-V_t transfer gate; that is, no buffer is used between the latches comprising a latch pair. The L1 output to the L2 latch can be taken from either the B or H nodes, both of which are internal L1 latch nodes.

The B-node transfer results in noninverted data to the L2, whereas the H-node transfer results in out-of-phase data to the L2. The L2 latch also utilizes an inverting push-pull output buffer. To obtain an L2 output which is noninverting, the H node of the L1 is used for transferring data to the L2. For an inverting L2 output, the L1 latch B node is used for transfer to the L2 latch.

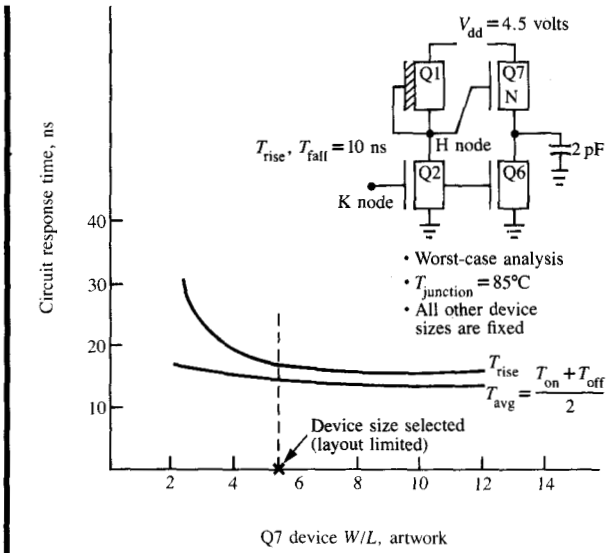

Figure 7

Latch push-pull buffer response versus Q7 device size.

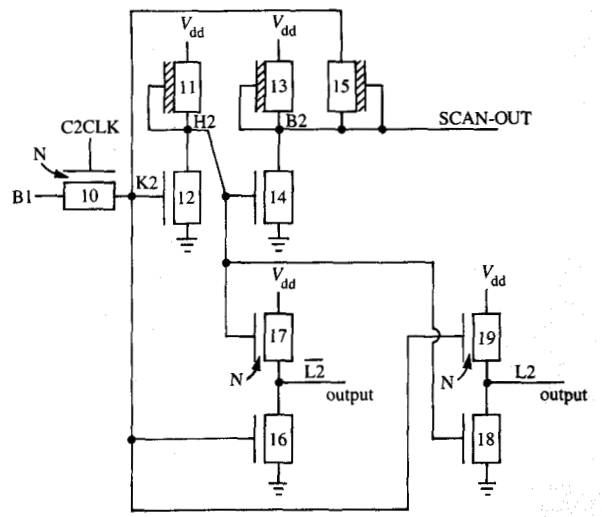

Figure 8

Latch with inverting and noninverting push-pull output buffers.

Although a noninverting L2 latch push-pull driver can be implemented, the resulting performance is severely degraded due to the charge redistribution effect which occurs in the L1 latch, as was previously discussed. The push-pull buffer polarity choice and rationale for such are discussed in the following section.

Push-pull output buffer polarity choice

An L2 latch inverting push-pull output has been used exclusively, because of performance considerations. **Figure 8** illustrates the L2 latch configurations for both inverting and noninverting push-pull output buffers. Note that the gate of the zero-V_t output device of the noninverting buffer is connected to the K2 node of the L2 latch. In this configuration the output rising response is a function of the rising transition of the zero-V_t gate voltage, which is connected to the K2 node.

Now, recall the resulting waveform of the K2 node of the L2 latch which is due to the charge redistribution associated with transferring a "1" from L1 to L2 (see Fig. 3). The K2-node voltage rises quickly to approximately 2.5 volts, which is the lowest point of the L1 output voltage dip. Further charging of the K2 node has the same response as that of the B node of the L1, which is recovering from the charge redistribution.

When the K2 node attains approximately 3.5 volts, the transfer device turns off and the remaining charging of the K2 node now must be supplied by device 13 via the high-impedance feedback device. This resulting waveform, as seen in **Figure 9**, is poor. The end result is that the output response of the buffer follows that of the K2 node, which is slow.

In the case of the inverting buffer configuration, the K2 node of the latch drives the output buffer pull-down device. This pull-down device turns on when the K2 node attains approximately 1.5–2.0 volts and therefore is insensitive to the slow response of the K2-node voltage above this value. The output falling is now associated with the input rising and is thus not following the input waveform. The output rising transition of this buffer configuration is associated with the input falling, which is not degraded by the charge redistribution effect. The H2 node of the latch is connected to the zero-V_t output buffer device in the inverting buffer configuration.

As the K2 voltage is falling, the H2 voltage rises very quickly because there is no charge redistribution effect on this node. The output of the buffer follows this fast-rising H2-node voltage, resulting in a fast output rising response.

Careful analysis of both buffer configurations revealed that the inverting buffer configuration resulted in performance which was approximately 2.0 times better than that obtained using a noninverting buffer configuration on the L2 latch. Performance data are presented in the power/performance characteristics section of this paper.

If both polarities are required concurrently, it is recommended that a separate inverting push-pull buffer be driven by the inverting push-pull buffer of the L2 latch.

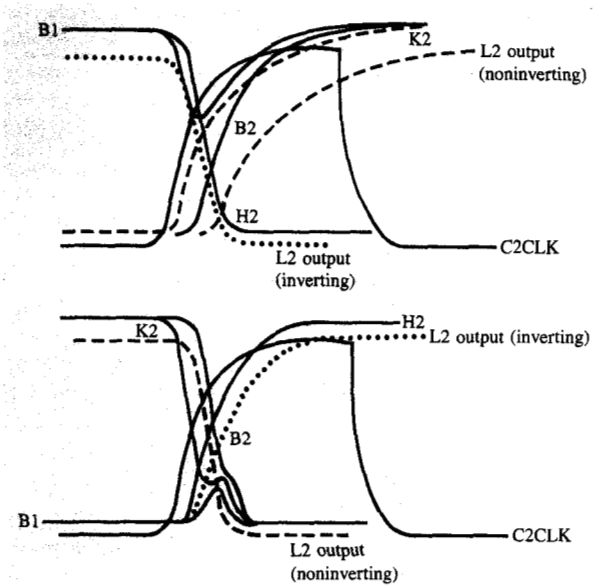

Figure 9

L1-L2 latch voltage waveforms for inverting and noninverting push-pull output buffers.

Latch scan-output considerations

Each L1 latch has two data input ports: one for inputting functional data and another for inputting LSSD data. The LSSD data port behaves exactly the same as the functional data port, which has already been described.

The basic concept of LSSD is to connect all memory elements (latches) via a separate test port as a shift register. This allows data to be shifted into all the latches from an external pin without regard for intermediate logic, thereby establishing the chip memory elements in a known state. Testing can then proceed to ensure the proper operation of all logic elements.

The L1 latches are connected to the L2 latches via the internal transfer gate (device 10, Fig. 1) of the latch pair. The scan-output (LSSD test output) of the latch pair or the functional push-pull output must be connected to the scan-input of another latch pair at the chip level until all L1-L2 latch pairs form a shift register chain. Multiple chains can be defined if desired.

With respect to the individual latch pairs, the scan-output is the B node of the L2 latch. It should be noted that the polarity of this output is a function of the L1 node of transfer. Hence, if an H node of transfer is used, the L2 latch scan-out is out of phase with the L1 scan-in data. It should also be noted that this scan-output node has a very limited drive capability, as it is an internal latch node. It is intended to drive the scan-input of an adjacent L1-L2 latch pair only.

Table 2 L1, L2 stage power dissipation at V_{dd} = 5.5 V.

	Output state = "0" (mW)	*Output state = "1"* (mW)
Nominal at 85°C	0.135	0.087
Maximum at 85°C	0.229	0.149

Excessive loading on this node results in nonstable operation. Therefore, for heavy capacitive loading (C greater than 0.1 pF), the L2 push-pull output should be utilized for the LSSD scan-output. Once again, polarity of the output must be taken into account.

Power/performance characteristics

Table 2 shows the nominal and maximum power dissipation at a junction temperature of 85°C and V_{dd} = 5.5 volts for both the L1 and L2 latch stages individually. This power dissipation also includes the power of the output push-pull buffer. Hence, the maximum power for an L1-L2 latch pair, including output buffers, is 0.378 mW, assuming a 50 percent duty cycle.

Tables 3 and **4** show the nominal and worst-case inverting push-pull output performance of the L1 latch stage and L2 latch stage, respectively, versus capacitive load. The delays assume data valid prior to the arrival of the clock signal.

T_{on} is defined as the delay for the output to fall to 1.5 volts with respect to the clock rising to the 1.5-volt level. T_{off} is defined as the delay for the output to rise to 1.5 volts with respect to the clock rising to the 1.5-volt level. T_{rise} and T_{fall} are the delays of the output rising from 1 volt to 3 volts and falling from 3 volts to 1 volt, respectively.

As shown by these tables, the L2 offers slightly better performance. This is due to device size differences of the L1 and L2 latch buffers, which were a result of physical layout constraints. **Table 5** shows the worst-case noninverting push-pull output performance of the L2 stage versus capacitive load when the output buffer is configured as illustrated in Fig. 8.

Comparing Tables 4 and 5, it appears that the delays of the noninverting push-pull buffer are comparable to, if not slightly better than, that of the inverting push-pull buffer. However, the output rise times of the noninverting push-pull buffer are significantly larger (approximately 3.8 times for C_{load} = 4 pF) than for the inverting push-pull buffer. **Table 6** shows a comparison of the delays for both the inverting and noninverting push-pull output buffers. To account for the rise time differences between these circuits, T_{offr} is defined as the delay for the output to rise to 3.0 volts with respect to the clock attaining the 1.5-volt level. Also, T_{avg} is defined as $(T_{on} + T_{offr})/2$.

Table 3 L1 stage inverting push-pull output performance (nominal/worst case).

Capacitance (pF)	T_{on} (ns)	T_{off} (ns)	T_{rise} (ns)	T_{fall} (ns)
1	15.0/20.7	16.9/27.3	11.2/22.1	6.2/9.3
2	18.6/25.8	19.0/30.5	15.6/26.2	9.3/13.1
4	25.9/36.0	23.3/36.8	24.5/34.3	15.6/19.7
6	33.2/46.2	27.6/43.1	33.4/42.5	21.8/26.3

Table 4 L2 stage inverting push-pull output performance (nominal/worst case).

Capacitance (pF)	T_{on} (ns)	T_{off} (ns)	T_{rise} (ns)	T_{fall} (ns)
1	14.1/18.9	15.4/20.3	9.2/16.0	5.3/8.3
2	18.3/27.2	17.4/24.1	12.9/19.3	8.7/14.3
4	26.8/40.7	21.3/29.6	20.2/26.6	15.4/23.4
6	35.3/52.2	25.1/33.1	27.5/34.9	22.0/31.3

Table 5 L2 stage noninverting push-pull output performance (worst case).

Capacitance (pF)	T_{on} (ns)	T_{off} (ns)	T_{rise} (ns)	T_{fall} (ns)
1	23.1	18.1	74.8	6.8
2	29.9	22.0	87.3	9.4
4	38.7	29.7	100.3	14.0
6	45.4	36.7	109.9	18.9

Table 6 L2 stage worst-case push-pull output performance (inverting versus noninverting).

Capacitance (pF)	T_{on} (ns)		T_{offr} (ns)		T_{avg} (ns)	
	Inv.	*Noninv.*	*Inv.*	*Noninv.*	*Inv.*	*Noninv.*
1	18.9	23.1	32.2	58.7	25.5	40.9
2	27.2	29.9	38.2	103.7	32.7	66.8
4	40.7	38.7	51.2	120.5	46.0	79.6
6	52.2	45.4	62.2	134.2	57.2	89.8

As shown by Table 6, the average delay of the L2 latch using a noninverting push-pull buffer is between 1.6 and 2 times that of an inverting push-pull buffer. The detrimental effects on the performance of the circuit being driven by such slow rise times, which result from the use of the noninverting push-pull buffer, have not been taken into account. Hence, the use of the inverting push-pull output buffer offers superior performance over that of the noninverting buffer.

All analysis was done using the IBM Advanced Statistical Analysis Program (ASTAP) [4]. All process and design parameters were specified at their worst-case limits (including power supply and temperature) for the worst-case analysis. All process and design parameters were specified at their nominal values except power supplies and temperature, which were set to worst case for the nominal analysis.

Summary and conclusions

This paper has described the many design considerations of a static LSSD polarity hold latch pair. Discussions ranging from basic operation to concerns regarding latch stability were presented. The original design goals of good performance, low power dissipation, and small size were met by using high-performance output buffers consisting of zero-V_t pull-up devices driven from low-power inverters comprising the latch. Emphasis was placed on the charge redistribution effect associated with using very-low-power (0.1 mW) internal circuits and the rationale for using only inverting push-pull output buffers. Power and performance characteristics, along with device size selection of the latch, were also presented. A section describing the latch physical layout demonstrated the flexibility with respect to I/O placement and showed that the latch is an entity which is easily integrable into VLSI designs.

In summary, although there are many considerations in the design and use of this static polarity hold latch design, its use results in low power dissipation while delivering good performance, LSSD testability, and occupation of a small silicon area.

Acknowledgment

The author would like to acknowledge K. F. Eng for his consultation during the preparation of this paper and Miss D. M. Schussler for typing the manuscript.

References

1. E. B. Eichelberger and T. W. Williams, "A Logic Design Structure for LSI Testability," *Proceedings of the 14th Design Automation Conference*, New Orleans, 1977, pp. 462–468.
2. Anthony Correale, "Physical Design of a Custom 16-Bit Microprocessor," *IBM J. Res. Develop.* **26,** 446–453 (1982).
3. G. Williams II, "Single Input Bistable Circuit," *IBM Tech. Disclosure Bull.* **18,** No. 10 (March 1976).
4. *Advanced Statistical Analysis Program (ASTAP), Program Reference* Manual, Order No. SH20-1118-0, available through IBM branch offices.

Received August 8, 1983; revised February 20, 1984

Anthony Correale *IBM Communication Products Division, P.O. Box 100, Kingston, New York 12401.* Mr. Correale joined IBM's System Products Division at East Fishkill, New York, in 1974, after graduating from City College of New York with a B.S. in electrical engineering (*cum laude*). He received his M.S. in electrical engineering from Syracuse University in 1982. His work assignments in East Fishkill included activity in the systems integrations area and circuit design of an 8-bit custom FET microprocessor, until his transfer to the System Communications Division at Kingston in 1978. Since that time, he has been the lead engineer on a VLSI custom 16-bit FET microprocessor. He is now a development engineer in the Custom Microprocessor Chip Design Department in the Communication Products Division, responsible for the design and development of state-of-the-art VLSI designs. Mr. Correale has received a Kingston Quality Award and an IBM Outstanding Innovation Award for his circuit design work done on VLSI microprocessors.