

Energy Delay Measures of Barrel Switch Architectures for Pre-Alignment of Floating Point Operands for Addition

R.V.K. Pillai, D. Al-Khalili[†] and A.J. Al-Khalili

Concordia University, Montreal, CANADA; [†]Royal Military College, Kingston, CANADA

Abstract

Significant pre-alignment is a pre requisite for floating point additions. This paper¹ addresses the architectural design and energy delay evaluation of a Low Power Barrel Switch for pre-alignment of floating point significands. Architectural energy delay analysis of Barrel Switch schemes suggests the suitability of transition activity scaled architectures for Low Power CMOS designs. Our energy delay estimates of operand pre-alignment Barrel Switches for the addition of IEEE single precision floating point numbers, taking into account the architectural as well as circuit implementation issues, suggests an energy delay reduction of better than 50% for transition activity scaled architectures for coefficients of parasitic loading exceeding 10. The corresponding reduction in power consumption is more than 55%.

I Introduction

Addition of floating point numbers essentially requires the alignment of significands in accordance with the difference between the exponents. In general, owing to the limited width of significand data fields, shifts beyond the width of significand are not necessary. For example, the probability that a valid shift condition exists is around 0.18 for the addition of IEEE single precision floating point operands, under such assumptions that the exponents of the numbers are independent and are uniformly distributed. If the assertion status of the nodes of the Barrel Switches are preserved during 'no valid shift' conditions, the resulting savings in dynamic power consumption can be substantial. The following paragraphs explain the architectural design and energy delay evaluation of a transition activity scaled Barrel Switch.

II Barrel Switch Architectures

The shifting of binary data through a number of bit positions can be implemented in different ways [1] - [5]. In general, a cascaded array of multiplexors can perform the

requisite amount of data shifts. Fig. 1 presents the block diagram of a Barrel Switch (BSI) which performs data alignment operations for the addition of IEEE single precision floating point data. In Fig. 1, alignment shifting is accomplished by routing the data through a cascaded array of 2X1 and 4X1 MUXs. Our investigations² suggest that this type of a shifter architecture is advantageous as far as energy




Fig. 1 - Barrel Switch - Scheme I (BSI)




Fig. 2 - Barrel Switch - Scheme II (BSII)

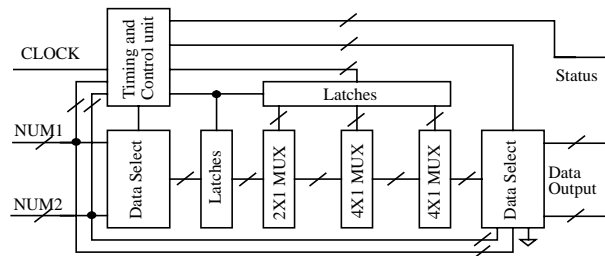


Fig. 3 - Control/Data flow Scheme of Barrel Switch (BSIII)

delay minimization of Barrel Switches is concerned. The data selection MUXs at the input of the shifter array present the significand of the smaller number for shifting. This block also performs the additional operation of routing the larger number as well as the sign of the smaller number to

1. This work had been supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada. The device models used for the energy delay analysis had been made available by the Canadian Microelectronics Corporation (CMC).

2. A single stage of nX1 MUXs can also perform the requisite shifts. Though such a scheme appears attractive owing to the reduction in the number of cascaded stages, the energy delay implications of such a scheme can be worse than that of multi stage shifters. Referring to the switch level diagram of a shifter MUX given in Fig. 6, it can be seen that the switching of the data select transistors constitute a charge sharing problem between the input nodes and the input of the level restoration inverter. With a large number of data select transistors, the delay accumulation due to this effect can be severe. The operational power demand of such schemes can also be worse than that of multi stage shifters owing to the effects of parasitic capacitances due to wiring complexity. The power/delay degradation of such schemes had been highlighted by other Researchers also [4].

the output. In Fig. 1, A and B represent the exponents while $NUM1$ and $NUM2$ represent the input floating point numbers. The shift control signals shown in Fig.1 are derived through the decoding of the bits of the exponent difference, $|A - B|$. For values of $|A - B|$ exceeding the shift range, the aligned significand must be set to zero. The ‘BSR’ signal shown in Fig.1 controls this operation. This Barrel Switch scheme is vulnerable to power losses due to spurious transitions [6] in the control/operand data paths due to the absence of delay balancing schemes. Fig. 2 presents a delay balanced architecture (BSII). The latches at the input of the shifter array provide the required delay balancing.

As stated earlier, the operational power demand of Barrel Switches can be reduced by inhibiting power consuming transitions within the shifter array during ‘no shift’ conditions. The scheme shown in Fig. 2 can be modified to this effect. Fig. 3 presents the control/data flow scheme of the proposed Barrel Switch (BSIII). The timing and control unit performs the following operations. (1) Computes the exponent difference, which eventually determines the magnitude of shifting required. (2) Decodes the exponent difference for controlling the shifter MUXs. (3) Evaluates ‘no shift’ conditions, viz. equality of exponents, exponent difference greater than shifting range and zero operands. (4) Evaluates the relative magnitudes of exponents. (5) Supplies delay balanced and inhibit controlled clocks for latching the inputs of the shifter array. (6) Generates control signals for effecting output data selection. The input data select block effects significand selection for shifting while the output data select block facilitates presentation of the aligned floating point numbers to the output. Apart from performing alignment shifts, the architecture also supports the evaluation of certain status signals that can be of interest in succeeding stages. A condition identical to ‘no shift’ exists for floating point adders as well, for a certain class of input operands such that the process of addition is not required. The pre computation of ‘no add’ conditions can be concurrently performed during the shifting operation, by making use of certain signals that are mandatory for the operation of the shifter. This signal can be useful for effecting power management/control of the significand adder. Apart from this signal, presentation of status signals like zero/infinity operands, ‘NaN’ (not a number) etc. can also be envisaged.

As stated earlier, the data latches at the input of the shifter array are enabled by properly delayed and ‘inhibit controlled’ clocks. These clocks are muted during those situations when the ‘no shift’ condition exists. This type of a control ensures the preservation of the assertion status of the nodes of the shifter array during ‘no shift’ conditions as well as delay balancing by virtue of which the operational power demand of the proposed scheme (BSIII) can be significantly less than that of conventional schemes.

III Circuit Realization

In floating point additions, the magnitude of data alignment shift is always decided by the difference between the expo-

nents whereas the selection of an appropriate significand for performing the requisite amount of shift is decided by the relative magnitudes of the exponents. Figs. 4 and 5 give the gate level representations of circuits that perform numerical




Fig. 4 - Evaluation of ‘end around carry’ for $|A - B|$ computation




Fig. 5 - Generation of sum bits for $|A - B|$ computation

comparison and subtraction of exponents. The triangular blocks in these figures represent 2×1 MUXs. The circuits given in Figs. 4 and 5 together, essentially constitute a 1’s complement subtracter which evaluates the absolute value of the difference between the exponents. In Fig. 5, the sum bits are evaluated only for the 5 LSB bit positions, which is fairly sufficient for performing the required shifts. In 1’s complement addition, the end around carry is used for effecting the required correction/complementation operation. The end around carry also reveals the relative magnitudes of the operands. For an addition of the type $A - B$ involving two exponents A and B , an end around carry of 1 indicates that $A > B$. In other words, the end around carry can be used for the selection of significands for shifting. This type of an approach is advantageous as far as the minimization of power and area are concerned.




Fig. 6 - 4×1 multiplexer

Comparisons of the type $|A - B| > \text{shift range}$ can also be performed by using comparators of the type shown in Fig. 4. Evaluation of ‘no shift’ condition can be realized through a logical OR operation of the conditions $A = B$, $A = 0$, $B = 0$

and $|A - B| >$ shift range. The ‘no shift’ condition can be used for muting the clocks which enable the latches at the input of the shifter array. Pass gate logic structures are ideal for the implementation of shifting as well as data selection MUXs. Fig. 6 shows the circuit diagram of a 4X1 MUX using NMOS pass transistors for data selection.

IV Energy Delay Analysis

The time averaged power consumption at the output of a CMOS logic structure is given by, $P = P_g (1 - P_g) f C_L V_{DD}^2$, where f is the operating frequency, P_g is the probability of finding a logic high at the node under consideration, C_L is the capacitive loading at the node and V_{DD} is the power supply voltage. The load capacitance at the output of a gate is proportional to the fanout of the gate. The total energy consumption (during one cycle of operation) due to signal dynamics at circuit nodes, in any logic structure can be expressed by the following relation.

$$E \propto \sum_{\forall g} P_g (1 - P_g) F_g \quad (1)$$

where F_g represents the fanout of the g th gate. The right hand side of the above equation represents the energy consumption measure of logic circuits. The energy delay product of the logic structure is given by

$$ED \propto \sum_{\forall g} P_g (1 - P_g) F_g \tau_{max} \quad (2)$$

where τ_{max} represents the delay of the critical path of the circuit. The above relation is useful for the evaluation of energy delay measures of gate level logic representations on the basis of signal probabilities, fanouts and circuit delays. The following equations give the activity driven (which is analogous to dynamic power consumption) energy measures of various Barrel Switch schemes. These equations had been derived through signal probability analysis as well as fanout considerations of the relevant schemes.

$$EM_I = 21.662 + 3.918n + 0.375N + \left[\frac{1}{\eta_1} (1.379n + 0.1788) + \frac{0.15N}{\eta_2} \right] (1 + k) \quad (3)$$

$$EM_{II} = 35.662 + 1.156n + 0.25N + \left[\frac{1}{\eta_1} (1.439n + 0.089) + \frac{0.15N}{\eta_2} \right] (1 + k) \quad (4)$$

$$EM_{III} = 63.231 + 0.333n + 0.25N + \left[\frac{1}{\eta_1} (0.3548n + 0.089) + \frac{0.15N}{\eta_2} \right] (1 + k) \quad (5)$$

where EM_I , EM_{II} and EM_{III} represent the energy measures of BSI, BSII and BSIII (Inhibit controlled Barrel Switch) respectively. In the above equations, k represents coefficient of parasitic loading, η_1 , η_2 represent the efficiencies of drivers whose fanouts are of the order n and N respectively, N

represents the width of the floating point number (including leading 1 and guard bits) and n represents the width of the significand. The validity of these models are restricted to the specific case of IEEE single precision floating point operands. The energy measures given by the above equations are specific for implementations using MUXs of the type shown in Fig. 6. For these MUXs (with device widths boosted to double the minimum sizes), the capacitances seen by the select lines are approximately 0.3 times the capacitances seen by the signal inputs, considering device models of 0.5 micron processes. For other types of MUXs, the above equations can be modified by appropriately scaling the terms multiplied by $(1 + k)$.

The energy measures given by the above equations take into account the architectural as well as circuit implementation issues. In general, signal probabilities, activity factors and fanouts capture the effects of the architecture as well as logic design. The parameter k , on the other hand, is solely dependant on the actual implementation. The higher the wiring complexity of the implementation, the higher the value of k . The coefficient of parasitic loading k and fanout F_g decide the value of the stage ratio (S) of drivers, as given by $S = \exp [(\ln ((1+k)F_g))/3]$ for a three stage buffer. The efficiencies of buffers (η) are functions of stage ratios as well as number of stages. The delay of BSI is 9τ plus driver delay while that of BSII and BSIII are respectively 10τ plus driver delay and 12τ plus driver delay, where τ represents the delay of a 2 input gate (a worst case estimate of which is 1.5 times the delay of a minimum sized inverter). A buffer delay of approximately 9τ can be anticipated from a typical three stage driver, having a stage ratio of around 4.4, optimally designed to drive the gate loads of around 26 MUXs of the type shown in Fig. 6 - considering a coefficient of parasitic loading of 10. The delay of latches is assumed to be equal to that of a 2 input gate.

V Results

Fig. 7 gives a plot of the percentage reduction in power consumption of the proposed Barrel Switch scheme, for various instances of parasitic loading. The dashed curve in Fig. 7 represents the percentage reduction in power consumption of BSIII with respect to that of BSII, while the solid line curve represents such a measure in comparison with BSI. The reduction in power consumption of the proposed scheme is better than 55% for a coefficient of parasitic loading of around 10^1 . Though the BSII scheme incorporates activity reduction through delay balancing, the power consumption of the delay balancing scheme offsets the power saving through activity reduction. For lower values of k , the

1. The following example highlights the significance of the value of k . In our energy delay analysis, parasitic loading is restricted to affect only the high fanout nodes, which are essentially excited by drivers. The select lines of shifter MUXs are good examples. For the type of MUXs shown in Fig. 6, with all device widths boosted to double the minimum required, the parasitic capacitance is around 1 pF for a k of around 13, for the CMC 0.5 micron process.

power consumption of BSII is less than that of BSI.

Fig. 8 gives the comparative energy delay reduction of BSIII against that of BSI and BSII. Here again, the dashed curve gives the relative reduction in comparison to BSII while the other curve gives that with respect to BSI. The reduction in energy delay is better than 50% for those values of k that are greater than 10.

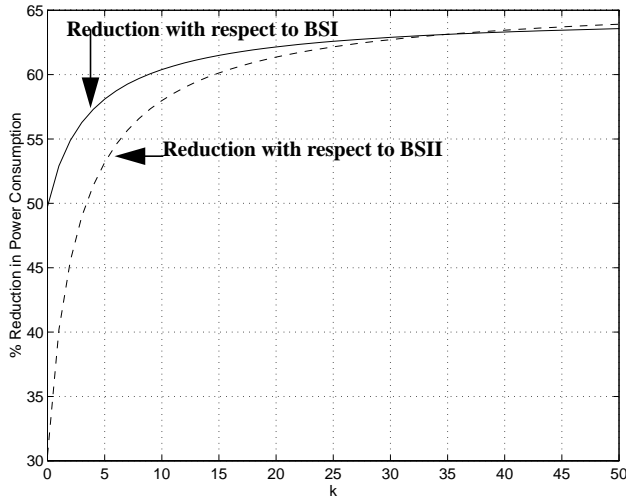


Fig. 7 - Percentage Reduction in Operating Power




Fig. 8 - Percentage Reduction in Energy Delay

VI Discussion

Table 1 highlights the significant components of the energy measures of various schemes. The values of the energy measures for shifter control as well as operand data path switching reveal the effects of transition activity reduction and delay balancing. The dominant component of power consumption is attributed to the switching activities of high fanout nodes and operand data path nodes, for all three schemes. Transition activity scaling of these nodes through ‘inhibit control’ is rewarding as far as design for Low Power operation is concerned. The delays of drivers which interface signals to the high fanout nodes, forms a major compo-

nent of the overall delay. Because of these reasons, the additional power/delay overheads attributed to the evaluation of extra control signals for effecting ‘inhibit control’ is insignificant in contrast to the savings in power wastage attainable through such a control. In general, the probability that a ‘no shift’ condition exist is given by $P(I) \approx \sum f(Z_i), \forall |i| > n$, where n represents the number of significant bits while $f(Z)$ represents the probability density function of $Z = A - B$. With independent, uniformly distributed exponents $P(I)$ is approximately 0.95 for IEEE double precision floating point format. The higher the value of $P(I)$, the better the chances for power reduction through ‘inhibit control’.

Table 1: Comparison of Architectural Energy Measures[†]

Operation	BSI	BSII	BSIII
Control Signal Evaluation	21.662	35.662	63.231
Shifter Control (Excitation of MUX Select Lines)	$1.05n(1 + k)/\eta_1$	$0.6n(1 + k)/\eta_1$	$0.106n(1 + k)/\eta_1$
Operand data path switching	$(0.375N + 3.918n)$	$(0.25N + 1.156n)$	$(0.25N + 0.333n)$

[†]The energy measures given in the above table are positive numbers. Scaling of these numbers with $fC_U V_{DD}^2$ (C_U represents a gate load of unity for any technology) will give dynamic power consumption.

VII Conclusion

The design of a Low Power, activity scaled (Inhibit Controlled) Barrel Switch for pre alignment of floating point significands for addition is presented. The energy delay advantage of this Barrel Switch renders it attractive for full custom or design synthesis implementations.

References

- [1] G. M. Tharakan and S. M. Kang, “A New Design of a Fast Barrel Switch Network”, *IEEE Journal of Solid-State Circuits*, Vol. 27, No. 2, pp. 217 - 221, February 1992.
- [2] Erdem Hokenek, Robert K. Montoye and Peter W. Cook, “Second Generation RISC Floating Point with Multiply - Add fused”, *IEEE Journal of Solid State Circuits*, Vol. 15, pp. 1207 - 1213, October 1990.
- [3] R. Pereira, J. A. Michell and J. M. Solana, “Fully Pipelined TSPC Barrel Shifter for High-Speed Applications”, *IEEE Journal of Solid State Circuits*, Vol. 30, No. 2, pp. 217 - 221, February 1992.
- [4] K. P. Acken, M. J. Irwin and R. M. Owens, “Power Comparisons for Barrel Shifters”, Digest of Technical Papers - 1996 International Symposium on Low Power Electronics and Design, pp. 209 - 212.
- [5] Raymond S. Lim, “A Barrel Switch Design”, *Computer Design*, pp. 76 - 79, August 1972.
- [6] A. P. Chandrakasan, S. Sheng and R. W. Brodersen, “Low-Power CMOS Digital Design,” *IEEE Journal of Solid State Circuits*, Vol. 27, pp. 473 - 483, April 1992.