

The IEEE Rounding for Multiplier with Redundant Operands

M. I. Ferguson
fergie@cs.ucla.edu

Miloš D. Ercegovac
milos@cs.ucla.edu

Computer Science Department
University of California, Los Angeles

Abstract

We present a design for a multiplier with redundant operands which conforms to the IEEE-754 floating-point standard. The design consists of a multiplier core and a rounding unit which conforms to the rounding modes specified by the IEEE standard and introduces a novel technique for calculating the sticky bit directly from the operands. We simulate a single-extended precision design and evaluate the delay versus a conventional design to be roughly equivalent, with extra savings possible with minor modification. The area is also estimated to be roughly 15% larger than a conventional design of the same precision.

1 Introduction

The use of redundant representations in the design of arithmetic units such as multipliers, dividers, and accumulators is a well-established technique for reducing the critical path in implementation. However, its use has been internal and the inputs/outputs are assumed to be in conventional form. Conversion of the redundant result to a conventional one is a time-consuming step since it involves propagation of carries over the working precision. To avoid the conversion step arithmetic schemes which accept redundant operands and produce redundant results have been proposed [1], [2] [3]. The authors have developed a radix-4 fully redundant multiplier core [4] and extend its design in

this paper. To perform a relevant comparison, we comply with the IEEE Floating-Point Standard [5].

The main thrust of this research is to understand the effects of using redundant form operands in floating-point multiplication. In order to evaluate this, a baseline design must be developed in parallel, so that contrast can be achieved. The baseline design is a floating-point multiplier using two 32-bit conventional form operands of IEEE single-extended precision. The scope of this paper is restricted to the multiplier core and rounding unit, as the exponent adjustment is not representation specific. The design of the multiplier and specification of rounding modes is first developed using operand-independent blocks in section 2, then representation dependent parameters are discussed in section 3 and those contrasts are evaluated in section 4.

2 Baseline Design

The multiplication operation is denoted $z = x \times y$ where x and y are the operands whose significands are represented by X and Y , respectively and are in the range $[1, 2)$. The m -bit operands are of the form $M = M_0.M_1M_2\dots M_{m-1}$. The final result is obtained by performing normalization and rounding on an intermediate result P , $Z = \text{RN}(P) = \text{RN}(X \times Y)$ where $\text{RN}()$ indicates a fused normalize and round operation. The intermediate result P , is in Carry/Save form

$(P_{0..2m-3}^c/P_{-1..2m-2}^s)^1$: Since $1 \leq X, Y < 2$ the result is in the range $[1, 4)$. Three special digit positions are recognized, $L = P_{m-1}$, $G = P_m$, and $T = \text{OR}(P_{m+1} \cdots P_{2m-2})$ where the indices indicate bit location **after** normalization. L specifies the digit of least precision, G is a guard digit and T is the sticky bit which indicates the tie condition where the result is exactly between two m -digit representable values. An overview of the carry/save form of the multiplier can be found in figure 1, the conventional form differing in that there is no "radix convert" block as the multiplicand X is used directly. Another minor difference is that the result index 0 is implied with a normalized conventional number, but not in carry/save form because there could be a 2 in position 1.. The outputs in the figure indicated by the wide arrows are busses too complicated to show schematically. The **A** and **B** busses from the radix convert block are the carry/sum bits as described in section 3. The bus coming from the recode block are the digits of the recoded multiplier in unary code $\{+1, -1, +2, -2\}$ which encode the range of recoded multiplier digits $\{-2, \dots, 2\}$, zero indicated as the null set of those four bit lines.

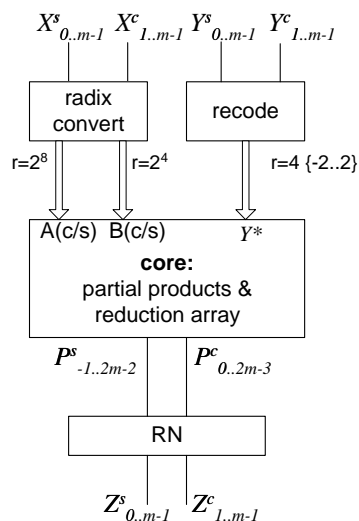


Figure 1. Overview of the carry/save form of the multiplier.

¹The index can range from index 0 to $2m-2$ if no recoding is done on the multiplier.

Core: The operands for the core are multiplicand X and multiplier Y . The multiplier is recoded to the set $\{-2, \dots, 2\}$. Partial products are reduced using the TDM method developed in [6]. The multiplicand is partially recombined as in [4] to reduce the complexity of the reduction array, this is discussed further below.

Rounding Unit: There are three rounding modes specified by the IEEE 754 include **Round to zero**, **Round to \pm infinity** and **Round to nearest (tie even)**. We concentrate here on Round to nearest as it is the default rounding mode. In this mode we add rnd to L where $rnd = G \text{ AND } (T \text{ OR } L)$. Equivalently a 1 is added to G and then L is forced to zero if $T = 0$. However, since the index of these digits isn't known until after normalization, a 2 might have to added at G instead. The choice between these is determined by whether or not P overflowed to the range $[2, 4)$. The two choices can be pre-computed, the selection between them being performed later based on the carry bit into position m , denoted C_m . The method used to combine the normalization, and adjustment of L is based on original work by Santoro [7] and expanded upon in [8]. The method for determining the sticky bit T is to subtract 1 from the carry/save output in the lower half of P by adding a string of 1s. If the result remains -1 , then $T = 0$. This leads to a simple tree of OR gates to determine T with the inclusion of P_m in the case of P_{ovf} . The structure for determining C_m is a simple carry-lookahead design with all the logic for computing intermediate sums and carries removed. An overview of the carry/save output version of the rounding unit can be seen in figure 2, the conventional form varies only in that the outputs of the Pre-Compute block, the Mux and the R1 Shift are not in carry/save form.

3 Representation Specific Issues

The use of redundant form operands has several effects on both the multiplier core as well as the RN block.

Partial Recombination: Naively, since the multiplicand has two rows of bits (X_s, X_c) , there must be twice as many partial products generated, which then feed into the reduction array. However,

Module	Conventional	Carry/Save
Recode	1	7
Radix convert	-	6
Generator	3	10
Reduction Array	30	38
Sticky	6	6
Carry	5	5
SEL	1	1
pre-compute	24	5
MUX	1	1
R1Shift	1	1
LADJ	1	1
Critical Path	60	63

Table 1. Results of T-Spice calculations of blocks. Units are ns.

X_0	X_1	X^-	X^+
0	1	1.0	1.0
0	2	1.0	1.5
1	0	1.0	1.5
1	1	1.5	2.0 ⁻

Table 3. X^- and X^+ indicate the min and max values attainable by X . 2.0⁻ indicates a value close to but less than 2.0.

X^-	X^+	Y^-	Y^+	P^-	P^+	N
1.0	1.0	1.0	1.0	1.0	1.0	0
1.0	1.0	1.0	1.5	1.0	1.5	0
1.0	1.0	1.5	2.0 ⁻	1.5	2.0 ⁺	?
1.0	1.5	1.0	1.0	1.0	1.5	0
1.0	1.5	1.0	1.5	1.0	2.25	?
1.0	1.5	1.5	2.0 ⁻	1.5	3.0 ⁻	?
1.5	2.0 ⁻	1.0	1.0	1.5	2.0 ⁺	?
1.5	2.0 ⁻	1.0	1.5	1.5	3.0 ⁻	?
1.5	2.0 ⁻	1.5	2.0 ⁻	2.25	4.0 ⁻	1

Table 4. The normalization decision based on 2 input digits of each operand.

Module	Conventional	Carry/Save
Recode	1.4	4.8
Radix convert	-	4
Generator + Array	62	73
Sticky	1.3	1.3
Carry	0.7	0.7
SEL	0.1	0.1
pre-compute	6	2.6
MUX	0.9	1.4
R1Shift	0.9	1.4
LADJ	0.1	0.1
Total	73	89

Table 2. Area of blocks given in millions of square Locator Units (LU².) The generator and reduction array are combined because if the reduction array is generated as a separate block there is tremendous overhead due to the large number of inputs

4.1 Refinements

A refinement to the C/S design can be made using the fact that the decision of whether or not to normalize can be made for some fraction of the cases using only the first two digits of the operands $X_0^s, X_1^{c,s}$ and $Y_0^s, Y_1^{c,s}$, thus mitigating the need for ovf_P0 . This only makes sense for the C/S version because for the conventional design the full carry propagation must be made on the upper half of P anyway. This allows the C/S version to work at two different speeds, the finish state being signified by a single bit. In order to determine the normalization state we can determine the *min* and *max* of X and Y , as in table 3, and then determine the range of P based on those ranges as shown in table 4. This leads to knowing the overflow condition for 4 out of 9 cases. For these operands the delay of the multiplier is 58ns, plus some small logic delay for analyzing the two leading digits.

4.2 Conclusion

We have shown that using standard layout tools and Spice analysis one can design a floating-point multiplier which accepts inputs in carry/save form. The delay of this multiplier is comparable with the delay of a multiplier with conventional operands of the same precision, paying only about an extra 15% in area. These savings are in addition to the delay which was not generated in the unit which produced the operands. Also, since the primary difference between the redundant form and the conventional is that of the full length CPA which must be performed, the savings should grow when moving to double and quad precision operands.

References

- [1] A. Glew, "Processor with architecture for improved pipelining of arithmetic instructions by forwarding redundant intermediate data forms," *U. S. Patent no. 5619664*, April 1997.
- [2] T. Nishiyama et al., "High Speed Multiplier Utilizing Signed-Digit and Carry-Save Operands," *U.S. Patent Number 4,868,777*, Sept. 1989.
- [3] M. Matula and A. Nielson *Pipelined packet-forwarding floating point. I. Foundations and a rounder, Proceedings. 13th IEEE Symposium on Computer Arithmetic*, Asilomar, CA, 6-9 July 1997. p.140-7.
- [4] M.I. Ferguson, M. D. Ercegovac, "A Multiplier with Redundant Operands," *Proceedings of the 33rd Annual Conference on Signals, Systems, and Computers*, Asilomar, Pacific Grove, October 24-27,1999 p. 1322-1325.
- [5] *IEEE Standard for Binary Floating-Point Arithmetic*, Inst. Electr. & Electron. Eng., New York, NY, USA, 12 Aug. 1985. 18 pp.
- [6] V. G. Oklobdzija, et al., "A Method for Speed Optimized Partial Product Reduction and Generation of Fast Parallel Multipliers Using an Algorithmic Approach," *IEEE Transactions on Computers*, Vol. 45, No. 3, March 1996.
- [7] M. Santoro, G. Bewick, M. Horowitz, "Rounding Algorithms for IEEE Multipliers". *Proceedings of the 9th Symposium on Computer Arithmetic*, Santa Monica, 6-8 Sept. 1989 p.176-83.
- [8] M. D. Ercegovac and T. Lang, "Digital Arithmetic", Sept. 2000, manuscript in progress.
- [9] Internet address www.mosis.com/Technical/Processes/proc-hp-amos14tb.html

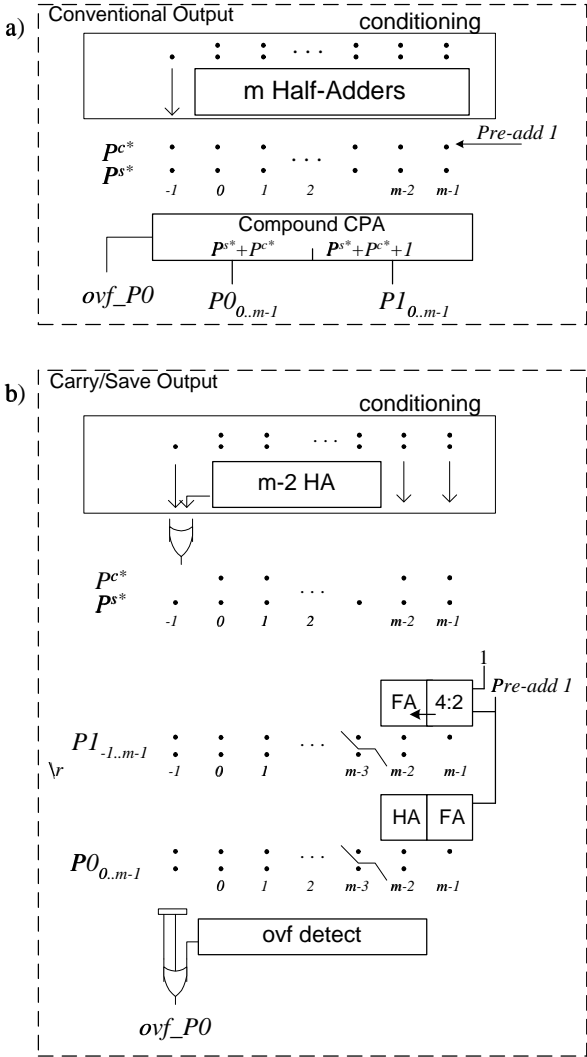


Figure 3. The difference between rounding to a conventional and a carry/save operand. The outputs to the left of the jagged line come directly from Pc^* and Ps^*