



Power Matters.™

Efficient INT8 Dot Product using Microsemi Math Block

WP0216 White Paper

Introduction

Recent breakthroughs in Deep Learning algorithms have enabled applications in a breadth of end markets. Efficient neural network architectures along with quantization methods have led to significant reductions in both compute and memory requirements for several applications including image and speech recognition.

The Dot Product is a basic computation requirement for both fully connected and convolutional layers in a neural network. Microsemi FPGA families (PolarFire, RTG4, and SmartFusion2) have a built-in Dot Product mode that delivers higher computational efficiency compared to the competition. This paper highlights the architectural advantage of the Dot Product mode of the Math block in Microsemi FPGAs to perform INT8 matrix operations.

Compute Requirement




Figure 1: (a) Fully Connected and (b) Convolutional layers in a neural network

- The basic computation requirement for both the fully connected and convolutional layers is a multiply accumulate of the form [4]:

$$o_j = f\left(\sum_i a_i w_{i,j}\right) \quad \text{where } (i \in [1, n])$$

where $f(x)$ is a non-linear activation function like ReLU.

The equation in expanded form is:

$$o_i = f(a_1 w_{11} + a_2 w_{21} + a_3 w_{31} \dots)$$

- Recent work [6, 7] has shown that 8-bit integer quantization is sufficient to produce acceptable accuracy for image classification applications.

Dot Product Mode in Microsemi FPGA Math Block

All three families of Microsemi FPGAs (PolarFire, RTG4, and SmartFusion2) have a built-in Dot Product (DOTP) mode that is especially suited for 8-bit arithmetic operations [1, 2, 3].

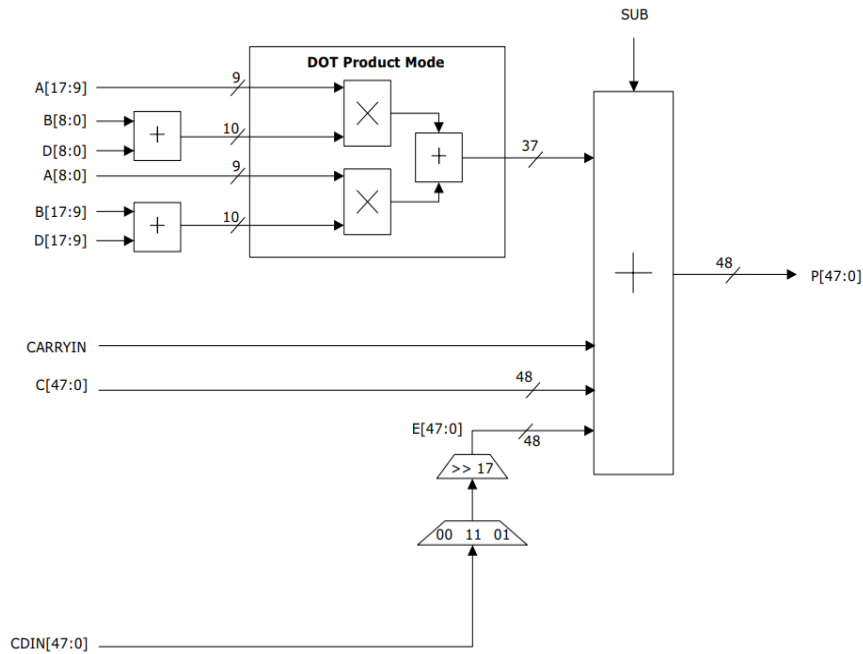


Figure 2: DOTP Mode available in PolarFire Math block

In the DOTP mode, a single math block accommodates the following four useful operations:

- Two multiplication and one addition:

$$a_1w_{11} + a_2w_{21}$$

- The accumulator blocks performs one add needed to accumulate the result from the previous block in cascade.

Competing FPGAs

Unlike the Microsemi Math block, competing FPGAs do not have a built-in DOTP mode for INT8 arithmetic. Due to this limitation, each DSP block can only accommodate two useful operations (one multiply, one add).

To alleviate this inefficiency, a weight sharing architecture has been proposed that computes the DOTP for two different inputs using the same weights [4, 5]. However, since the computation is for two different inputs, the results need to be separated in the accumulator. This limits the length of the DSP cascade before the upper and lower words result become unrecoverable. An additional DSP must be used to handle the separation and summing of the lower and upper words after each cascade.

With the workaround, each DSP block can perform four useful operations (two multiply, two add). The addition performed in the pre-adder is not useful since its goal is to generate a bit shifted output to feed in to the multiplier. However, the requirement of an additional DSP following the cascade reduces the effective useful operations per DSP.

Also, additional control logic will be needed to shift input data operands and unpack results while tracking multiple data flows with separate inputs. The increased complexity in the data flow could lead to additional limitations, which reduce operational efficiency.

Microsemi vs Competing FPGA Comparison

Table 1 • Microsemi vs Competing FPGA Comparison

Family	INT8 Operations per DSP / Math Block
Microsemi FPGA ¹	4
Competing FPGA	2
Competing FPGA (Weight Sharing)	up to 3.5 ²

1. *PolarFire, SmartFusion2, and RTG4.*
2. *Requires overhead logic to shift input data and unpack results.*

Conclusion

The presence of the in-built DOTP mode in Microsemi FPGAs enables all three families (PolarFire, RTG4, and SmartFusion2) to deliver higher efficiency for INT8 computations used by the fully connected and convolutional layers:

- 100% higher compared to Competing FPGA.
- 14% higher compared to Competing FPGA (Weight Sharing).

The higher efficiency with the Microsemi FPGAs is achieved without increasing data flow complexity to:

- Shift input data operands and unpack results.
- Include control logic to track multiple data flows with separate inputs.

References

1. UG0574: RTG4 FPGA Fabric User Guide
2. UG0680: PolarFire FPGA Fabric User Guide
3. Microsemi Digital Signal Processing Reference Guide
4. Y. Fu et al. Deep Learning with INT8 Optimization on Xilinx Devices, 2017
5. Y. Fu et al. 8-bit Dot Product Acceleration, 2017
6. P. Gysel et al. Hardware-oriented Approximation of CNNs, 2016
7. S. Han et al. Deep Compression: Compressing DNNs with Pruning, Trained Quantization and Huffman Coding



Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Sales: +1 (949) 380-6136
Fax: +1 (949) 215-4996

E-mail: sales.support@microsemi.com

© 2018 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi Corporation (MSCC) offers a comprehensive portfolio of semiconductor and system solutions for communications, defense & security, aerospace and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, Calif., and has approximately 4,800 employees globally. Learn more at www.microsemi.com.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.