# Hybrid Dot-Product Design for FP-Enabled FPGAs 

Bogdan Pasca

Intel
ARITH 2019, June 10-12, 2019

# Hybrid Dot-Product Design for FP-Enabled FPGAs 

Bogdan Pasca

Intel
ARITH 2019, June 10-12, 2019


## Context

- FPGAs intersting neural network training accelerators
- training: mostly dot-products in forward+backward propagation
- current industry-standard: bfloat16 multiplications, SP reduction
- bfloat16 vs SP: 2X bandwidth


## Context

- FPGAs intersting neural network training accelerators
- training: mostly dot-products in forward+backward propagation
- current industry-standard: bfloat16 multiplications, SP reduction
- bfloat16 vs SP: 2X bandwidth

Goal $\rightarrow$ Find a dot-product implementation that:

- maintains an accuracy comparable to bfloat16+SP
- maximizes the dot-product density for a given FPGA


## Density



- FPGA devices have a various mix of resources
- increasing compute density $\rightarrow$ make efficient use of existing mix


## Density



- FPGA devices have a various mix of resources
- increasing compute density $\rightarrow$ make efficient use of existing mix

Focus on "Core Logic Fabric" and VP DSP Blocks

## Density - some current devices



Intel ${ }^{\ominus}$ Agilex ${ }^{\top n}$ FPGAs


Intel ${ }^{6}$ Stratix ${ }^{\oplus}$ Series


Intel ${ }^{\circ}$ Arria ${ }^{\circ}$ Series


Intel ${ }^{\circledR}$ Cyclone ${ }^{\circ}$ Series

## Density - some current devices



Intel ${ }^{\oplus}$ Agilex ${ }^{\top n}$ FPGAs


Intel ${ }^{6}$ Stratix ${ }^{\oplus}$ Series


Intel ${ }^{\circ}$ Arria ${ }^{\circ}$ Series


Intel ${ }^{\ominus}$ Cyclone ${ }^{\circ}$ Series

| PRODUCT LINE |  | $\begin{aligned} & \text { GX } 160 \\ & \text { SX } 160 \end{aligned}$ | $\begin{aligned} & \text { GX } 220 \\ & \text { SX } 220 \end{aligned}$ | $\begin{aligned} & \text { GX } 270 \\ & 5 \times 270 \end{aligned}$ | $\begin{aligned} & \text { GX } 320 \\ & \text { SX } 320 \end{aligned}$ | GX 480 SX 480 | GX SX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & \mathscr{U} \\ & \text { U } \\ & 0 \\ & 0 \\ & 0 \\ & \sim \end{aligned}$ | LEs (K) | 160 | 220 | 270 | 320 | 480 |  |
|  | System logic elements (K) | 210 | 288 | 354 | 419 | 629 |  |
|  | Adaptive logic modules (ALMs) | 61,510 | 83,730 | 101,620 | 118,730 | 181,790 | 217 |
|  | Registers | 246,040 | 334,920 | 406,480 | 474,920 | 727,160 | 86 |
|  | M20K memory blocks | 440 | 588 | 750 | 891 | 1,438 | 1, |
|  | M20K memory (Mb) | 9 | 11 | 15 | 17 | 28 |  |
|  | MLAB memory (Mb) | 1.0 | 1.8 | 2.4 | 2.8 | 4.3 |  |
|  | Hardened single-precision floating-point multiplers/adders | 156/156 | 191/191 | 830/830 | 985/985 | 1,368/1,368 | 1,523 |
|  | $18 \times 19$ multipliers | 312 | 382 | 1,660 | 1,970 | 2,736 | 3, |
|  | Peak fixed-point performance (GMACS) ${ }^{1}$ | 343 | 420 | 1,826 | 2,167 | 3,010 | 3, |
|  | Peak floating-point performance (GFLOPS) | 140 | 172 | 747 | 887 | 1,231 | 1 |

## Density - some current devices



Intel ${ }^{\ominus}$ Agilex ${ }^{\top n}$ FPGAs


Intel ${ }^{6}$ Stratix ${ }^{\oplus}$ Series


Intel ${ }^{\circ}$ Arria ${ }^{\circ}$ Series


Intel ${ }^{\ominus}$ Cyclone ${ }^{\circ}$ Series

| PRODUCT LINE |  | $\begin{aligned} & \text { GX } 160 \\ & \text { SX } 160 \end{aligned}$ | $\begin{aligned} & \text { GX } 220 \\ & \text { SX } 220 \end{aligned}$ | $\begin{aligned} & \text { GX } 270 \\ & \text { S } \times 270 \end{aligned}$ | $\begin{aligned} & \text { GX } 320 \\ & \text { SX } 320 \end{aligned}$ | GX 480 SX 480 | GX SX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & \mathscr{U} \\ & \text { U } \\ & 0 \\ & 0 \\ & 0 \\ & \text { O} \end{aligned}$ | LEs (K) | 160 | 220 | 270 | 320 | 480 |  |
|  | System logic elements (K) | 210 | 288 | 354 | 419 | 629 | 7 |
|  | Adaptive logic modules (ALMs) | 61,510 | 83,730 | 101,620 | 118,730 | 181,790 | 21 |
|  | Registers | 246,040 | 334,920 | 406,480 | 474,920 | 727,160 | 868 |
|  | M20K memory blocks | 440 | 588 | 750 | 891 | 1,438 | 1, |
|  | M20K memory (Mb) | 9 | 11 | 15 | 17 | 28 |  |
|  | MLAB memory (Mb) | 1.0 | 1.8 | 2.4 | 2.8 | 4.3 |  |
|  | Hardened single-precision floating-point multiplers/adders | 156/156 | 191/191 | 830/830 | 985/985 | 1,368/1,368 | 1,523 |
|  | $18 \times 19$ multipliers | 312 | 382 | 1,660 | 1,970 | 2,736 | 3, |
|  | Peak fixed-point performance (GMACS) ${ }^{1}$ | 343 | 420 | 1,826 | 2,167 | 3,010 | 3, |
|  | Peak floating-point performance (GFLOPS) | 140 | 172 | 747 | 887 | 1,231 | $\underline{1,}$ |

## Density - some current devices



Intel ${ }^{\ominus}$ Agilex ${ }^{\top m}$ FPGAs


Intel ${ }^{6}$ Stratix ${ }^{\oplus}$ Series


Intel ${ }^{\circ}$ Arria ${ }^{\circ}$ Series


Intel ${ }^{\ominus}$ Cyclone ${ }^{\circ}$ Series

| PRODUCT LINE |  | $\begin{aligned} & \text { GX } 160 \\ & \text { SX } 160 \end{aligned}$ | $\begin{aligned} & \text { GX } 220 \\ & \text { SX } 220 \end{aligned}$ | $\begin{aligned} & \text { GX } 270 \\ & \text { S } \times 270 \end{aligned}$ | $\begin{aligned} & \text { GX } 320 \\ & \text { SX } 320 \end{aligned}$ | GX 480 <br> SX 480 | GX SX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & \mathscr{U} \\ & \text { U } \\ & 0 \\ & 0 \\ & \text { O } \\ & \text { O} \end{aligned}$ | LEs (K) | 160 | 220 | 270 | 320 | 480 |  |
|  | System logic elements (K) | 210 | 288 | 354 | 419 | 629 |  |
|  | Adaptive logic modules (ALMs) | 61,510 | 83,730 | 101,620 | 118,730 | 181,790 | 217 |
|  | Registers | 246,040 | 334,920 | 406,480 | 474,920 | 727,160 | 868 |
|  | M20K memory blocks | 440 | 588 | 750 | 891 | 1,438 | , |
|  | M20K memory (Mb) | 9 | 11 | 15 | 17 | 28 |  |
|  | MLAB memory (Mb) | 1.0 | 1.8 | 2.4 | 2.8 | 4.3 |  |
|  | Hardened single-precision floating-point multiplers/adders | 156/156 | 191/191 | 830/830 | 985/985 | 1,368/1,368 | 1,523 |
|  | $18 \times 19$ multipliers | 312 | 382 | 1,660 | 1,970 | 2,736 | 3, |
|  | Peak fixed-point performance (GMACS) ${ }^{1}$ <br> Peak floating-point performance (GFLOPS) | 394 | 459 | 122 | 120 | 132 | 1 |

## Density - some current devices



Intel ${ }^{\oplus}$ Agilex ${ }^{\top n}$ FPGAs


Intel ${ }^{6}$ Stratix ${ }^{\oplus}$ Series


Intel ${ }^{\ominus}$ Arria ${ }^{\circ}$ Series


Intel ${ }^{\circledR}$ Cyclone ${ }^{\ominus}$ Series

Hardened single-precision floating-point multiplers/adders
$18 \times 19$ multipliers
156/156
312

## Background

- Intel FPGAs: DSP blocks implement SP mult-add
- N-element SP dot-product $=$ N DSPs



## Background

- Intel FPGAs: DSP blocks implement SP mult-add
- N -element SP dot-product $=\mathrm{N}$ DSPs

- soft-logic-only solution
- bfloat16 multiplier $\rightarrow 2 /$ DSP
- SP FP adder $1 \rightarrow$ 1/DSP
- N -element dot product: $C_{D S P}=N / 2+N-1$


## Background

- Intel FPGAs: DSP blocks implement SP mult-add
- N -element SP dot-product $=\mathrm{N}$ DSPs

- soft-logic-only solution
- bfloat16 multiplier $\rightarrow 2 /$ DSP
- SP FP adder $1 \rightarrow$ 1/DSP
- N -element dot product: $C_{D S P}=N / 2+N-1$
- adjust ratio: migrate SP FP adders to logic (300-400 ALMs/add)


## Background

- Intel FPGAs: DSP blocks implement SP mult-add
- N -element SP dot-product $=\mathrm{N}$ DSPs

- soft-logic-only solution
- bfloat16 multiplier $\rightarrow$ 2/DSP
- SP FP adder $1 \rightarrow 1 / D S P$
- N -element dot product: $C_{D S P}=N / 2+N-1$
- adjust ratio: migrate SP FP adders to logic (300-400 ALMs/add)
- solution is too large


## How do we solve this?



## Our implementation



$$
\begin{gathered}
N=\alpha+\beta \\
C_{A L M}=f(\alpha, w) \\
C_{D S P}=\alpha / 4+\beta
\end{gathered}
$$

## Our implementation



$$
\begin{gathered}
N=\alpha+\beta \\
C_{A L M}=f(\alpha, w) \\
C_{D S P}=\alpha / 4+\beta
\end{gathered}
$$

Objective: $C_{A L M} / C_{D S P} \approx$ device ALM/DSP ratio

## Hard FP part



- SP accumulation integrated
- $P_{g}$ will merge with the logic-based dot product
- $P_{/}$recirculated, added with $P_{b}$ using spare adder


## Soft FP part



## Soft FP part



## Soft FP part



## Soft FP part - fused

Multipliers

- $1 D S P=2 \times 18 \times 18=4 \times 8 \times 8$ mantissa multipliers (+ALMs)
- skip multiplier normalization
- $R N \rightarrow R Z_{w}$
- extend exponent to avoid overflow/underlow


## Soft FP part - fused

## Multipliers

- $1 D S P=2 \times 18 \times 18=4 \times 8 \times 8$ mantissa multipliers (+ALMs)
- skip multiplier normalization
- $R N \rightarrow R Z_{w}$
- extend exponent to avoid overflow/underlow


## Adders

- (except first layer inputs) operate on 2's complement mantissas
- mantissa grows by 1 (+1 optional) bit(s) every stage
- mantissa format changes from $(S M, 1, w F) \rightarrow(2 C, 1+1+L, w+L)$
- after final adder, normalization converts to SP


## Soft FP part - fused

## Multipliers

- 1 DSP $=2 \times 18 \times 18=4 \times 8 \times 8$ mantissa multipliers (+ALMs)
- skip multiplier normalization
- $R N \rightarrow R Z_{w}$
- extend exponent to avoid overflow/underlow


## Adders

- (except first layer inputs) operate on 2's complement mantissas
- mantissa grows by 1 (+1 optional) bit(s) every stage
- mantissa format changes from $(S M, 1, w F) \rightarrow(2 C, 1+1+L, w+L)$
- after final adder, normalization converts to SP
- intermediary normalization may be introduced for large $\alpha$


## Accuracy (Average)

- w - knob to control the accuracy
- $e_{c}$ - exponents centered, $e_{s}$ - the exponent span
- $e_{c}=0, e_{s}=10$ - inputs generated in $\left(2^{-10} \cdot 2,2^{10} \cdot 2\right)$

Table: Average relative error comparison between the proposed hybrid dot-product and a typical Al bfloat16+SP implementation for $n=16, \alpha=12$, $\beta=4, \beta_{g}=2, \beta_{b}=2$

| Config | Param | Proposed | AI |
| :--- | :---: | :---: | :---: |
|  | $w=7$ | $1.287601 \mathrm{e}-02$ |  |
| $e_{c}=0, e_{s}=5$ | $w=8$ | $6.172194 \mathrm{e}-03$ | $4.570449 \mathrm{e}-03$ |
|  | $w=9$ | $2.935275 \mathrm{e}-03$ |  |
|  | $w=7$ | $7.934867 \mathrm{e}-03$ |  |
| $e_{c}=0, e_{s}=10$ | $w=8$ | $4.120781 \mathrm{e}-03$ | $3.402314 \mathrm{e}-03$ |
|  | $w=9$ | $1.864206 \mathrm{e}-03$ |  |
| $e_{c}=0, e_{s}=20$ | $w=7$ | $6.672454 \mathrm{e}-03$ |  |
|  | $w=8$ | $3.161355 \mathrm{e}-03$ | $2.996574 \mathrm{e}-03$ |
|  | $w=9$ | $1.588372 \mathrm{e}-03$ |  |

$$
r_{d o t}=C_{D S P} / A L M s
$$

| Config | Param | ALMs | DSPs | $r_{\text {dot }}$ |
| :---: | :---: | :---: | :---: | :---: |
| $n=16$ | $w=7$ | 1030 | 7 | 147 |
| $\alpha=12, \beta=4$ | $w=8$ | 1075 |  | 153 |
| $\beta_{g}=2, \beta_{b}=2$ | $w=9$ | 1141 |  | 163 |
| $\begin{gathered} n=16 \\ \alpha=10, \beta=6 \\ \beta_{g}=4, \beta_{b}=2 \end{gathered}$ | $w=7$ | 863 | 8.5 | 102 |
|  | $w=8$ | 894 |  | 106 |
|  | $w=9$ | 948 |  | 112 |

$$
r_{d o t}=C_{D S P} / A L M s
$$

| Config | Param | ALMs | DSPs | $r_{\text {dot }}$ |
| :---: | :---: | :---: | :---: | :---: |
| $\begin{gathered} n=16 \\ \alpha=12, \beta=4 \\ \beta_{g}=2, \beta_{b}=2 \\ \hline \end{gathered}$ | $w=7$ | 1030 | 7 | 147 |
|  | $w=8$ | 1075 |  | 153 |
|  | $w=9$ | 1141 |  | 163 |
| $\begin{gathered} n=16 \\ \alpha=10, \beta=6 \\ \beta_{g}=4, \beta_{b}=2 \\ \hline \end{gathered}$ | $w=7$ | 863 | 8.5 | 102 |
|  | $w=8$ | 894 |  | 106 |
|  | $w=9$ | 948 |  | 112 |


| PRODUCT LINE |  | $\begin{aligned} & \text { GX } 160 \\ & \text { SX } 160 \end{aligned}$ | $\begin{aligned} & \text { GX } 220 \\ & \text { SX } 220 \end{aligned}$ | $\begin{aligned} & \text { GX } 270 \\ & \text { S× } 270 \end{aligned}$ | $\begin{aligned} & \text { GX } 320 \\ & \text { SX } 320 \end{aligned}$ | GX 480 SX 480 | GX SX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & \check{u} \\ & \text { U } \\ & 0 \\ & 0 \\ & \text { U } \end{aligned}$ | LEs (K) | 160 | 220 | 270 | 320 | 480 |  |
|  | System logic elements (K) | 210 | 288 | 354 | 419 | 629 |  |
|  | Adaptive logic modules (ALMs) | 61,510 | 83,730 | 101,620 | 118,730 | 181,790 | 21 |
|  | Registers | 246,040 | 334,920 | 406,480 | 474,920 | 727,160 | 86 |
|  | M20K memory blocks | 440 | 588 | 750 | 891 | 1,438 | 1, |
|  | M20K memory (Mb) | 9 | 11 | 15 | 17 | 28 |  |
|  | MLAB memory (Mb) | 1.0 | 1.8 | 2.4 | 2.8 | 4.3 |  |
|  | Hardened single-precision floating-point multiplers/adders <br> $18 \times 19$ multipliers | $\begin{gathered} 156 / 156 \\ 312 \end{gathered}$ | $191 / 191$ 382 | $830 / 830$ 1,660 | $985 / 985$ 1,970 | $\begin{gathered} 1,368 / 1,368 \\ 2,736 \end{gathered}$ | 1,523 |
|  | Peak fixed-point performance (GMACS) ${ }^{1}$ <br> Peak floating-point performance (GFLOPS) | 394 | 459 | 122 | 120 | 132 | 1 |

## Open questions

- reduction tree topology?


## Open questions

- reduction tree topology?
- accuracy?


## Open questions

- reduction tree topology?
- accuracy?
- resource ratio-account for plumbing?


## Open questions

- reduction tree topology?
- accuracy?
- resource ratio-account for plumbing?
- integration/syncronization?


## Open questions

- reduction tree topology?
- accuracy?
- resource ratio-account for plumbing?
- integration/syncronization?
- design/portability?


## Open questions

- reduction tree topology?
- accuracy?
- resource ratio-account for plumbing?
- integration/syncronization?
- design/portability?
- routability?

