

2GRVI Phalanx: W.I.P. Towards Kilocore RISC-V® FPGA Accelerators with HBM2 DRAM

⇒ 1776 RV32I / 1332 RV64I cores, 28 MB SRAM, 30 HBM2 DRAM Channels, PCIe, on Xilinx UltraScale+ VU37P / Alveo U280

Jan Gray | Gray Research LLC | Bellevue, WA | jan@fpga.org | http://fpga.org

Datacenter FPGA accelerators await our apps

- MS Catapult, Amazon AWS F1, Alibaba, Baidu, Nimbix
- Massively parallel, specialized, connected, versatile
- High throughput, low latency, energy efficient

But two hard problems

- Software: Porting & maintaining workload as accelerator
- Hardware: Compose 100s of cores, 100G NICs, many DRAM/HBM channels, with easy timing closure

Mission: GRVI Phalanx FPGA accelerator kit

- **GRVI:** FPGA-efficient RISC-V processing element cores
- **Phalanx:** array of clusters of PEs, SRAMs, accelerators
- **Hoplite NoC:** FPGA-optimal directional 2D torus soft NoC
- Local shared memory, global message passing, PGAS

Software-first, software-mostly accelerators

- Run your C++/OpenCL+ kernels on 100s of soft processors
- Add custom function units/cores/memories to suit
- More 10 sec recompiles, fewer 5 hour synth/place/route
- Complements high level synthesis & OpenCL→FPGA flows

2017: V1: 1680 core GRVI Phalanx in a VU9P

GRVI Phalanx: A Massively Parallel RISC-V® FPGA Accelerator Framework
⇒ A 1680-core, 26 MB SRAM Parallel Processor Overlay on Xilinx UltraScale+ VU9P

Jan Gray | Gray Research LLC | Bellevue, WA | jan@fpga.org | http://fpga.org

Datacenter FPGA accelerators are mainstream

- MS Catapult, Amazon AWS F1, Intel + Altera, Baidu, Alibaba, Nimbix
- Massively parallel, specialized, connected, versatile
- High throughput, low latency, energy efficient

But two hard problems

- Software: Porting & maintaining workload as accelerator

• Hardware: Compose 100s of cores, 25-100G NICs, many DRAM/HBM channels, with easy timing closure

GRVI Phalanx: FPGA accelerator framework

• GRVI: FPGA-efficient RISC-V processing element

• Phalanx: CPU/accelerator/IO fabric: clusters of PEs, SRAMs, accelerators, DRAM/IO controllers on a Hoplite NoC

• Local shared memory, global message passing

Software-first, software-mostly accelerators

• Run your C++/OpenCL+ kernels on 100s of soft processors

• Add custom function units/cores/memories to suit

• More 10 sec recompiles, fewer 5 hour synth/place/route

• Complements high level synthesis & OpenCL→FPGA flows

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)

• 32b pointers, 32b integer, 32b floating point

• 2 cycle loads, 3 cycle taken branches/jumps

• Painstakingly technology mapped and floorplaned

• Typically 320 LUTs @ 375 MHz = 0.7 MIPS/LUT

GRVI: massive RISC-V processing element (PE)

• User mode RV32I, minus all CSR plus -M+u+I (cluster DSP), -A 1/z/c (cluster RAM banks)