Technical in-depth

Arithmetic on the Mill™ Architecture

from

Out-of-the-Box Computing

The Mill

- Binary compatible family of cores.
- Wide issue, in-order, statically scheduled.
- 34 to 300+ RISC-equivalent 32-bit operations per cycle per thread, peak.
- Not all members support the full operation set in hardware. Unsupported operations are transparently emulated.
- Lift the hood and a Mill looks like nothing you've ever seen.

Data Types

Supported operand widths, in bits

integer
8, 16, 32, 64, *128*
pointer
64
IEEE 754R floating point
16, 32, 64, *128*
IEEE 754R decimal
32, 64, *128*
C200x fraction
8, 16, 32, 64, *128*
C200x fixed point
16, 32, 64, 128, *256*
IEEE 754R complex
64, 128, *256*

Widths in *italics* are optional, present in hardware only in members intended for certain markets and otherwise emulated in software. Widths in *bold italics* will be present if the type is supported in hardware at all.

Operands

- Up to four arguments per operation
- Double-rate issue if only 2 arguments
- Any number of results per operation
- Operands may be scalar (single value) or slice (short array of same-sized values)
- Slices apply operation SIMD to each element
- 3-4 argument ops (e.g. dot product and complex ops) have only one rounding error

Why decimal?

- Mainframe experience shows an established market for hardware decimal arithmetic
- Significant cycles are spent in decimal emulation in commercial markets (Oracle)
- New decimal standard (IEEE754-2008) should lead to format consolidation in industry

Why us?

- No, we're not competing with IBM – yet
- New design, family model, and missing op emulation lets us leave room for the future
- Why paint ourselves into a corner?

Which decimal format?

- The choice: declets/BCD (IBM) vs. scaled integer (Intel). IEEE754R accepts both (BOO!)
- Technical choice: IBM better for hardware implementation, Intel better for software
- Market choice: unpredictable which will dominate consolidation, maybe both (niche?)

Conclusion: both

- IBM format: full hardware opset support
- Intel: unfold/fold helper operations only

Scalar vs. Slice operation

Scalar operation – only low element

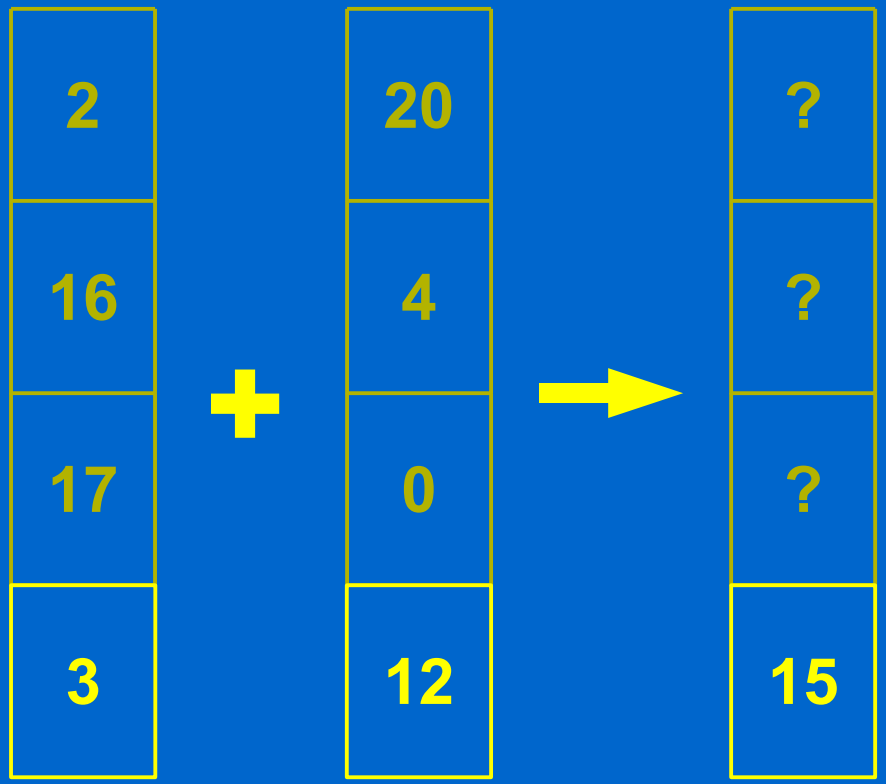

Slice operation – all elements in parallel

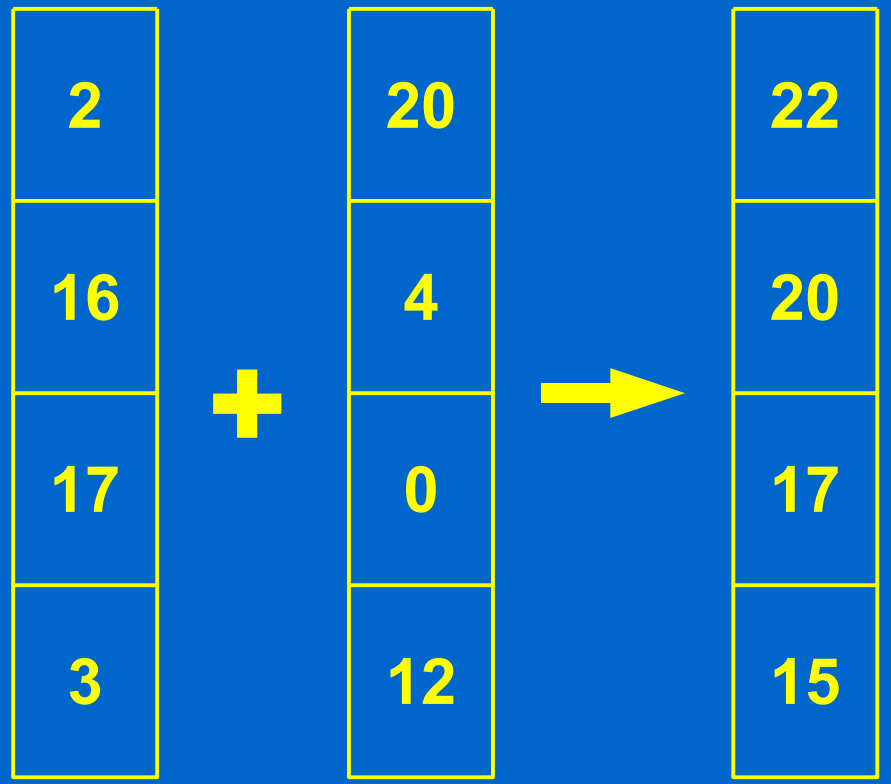

Overflow Semantics

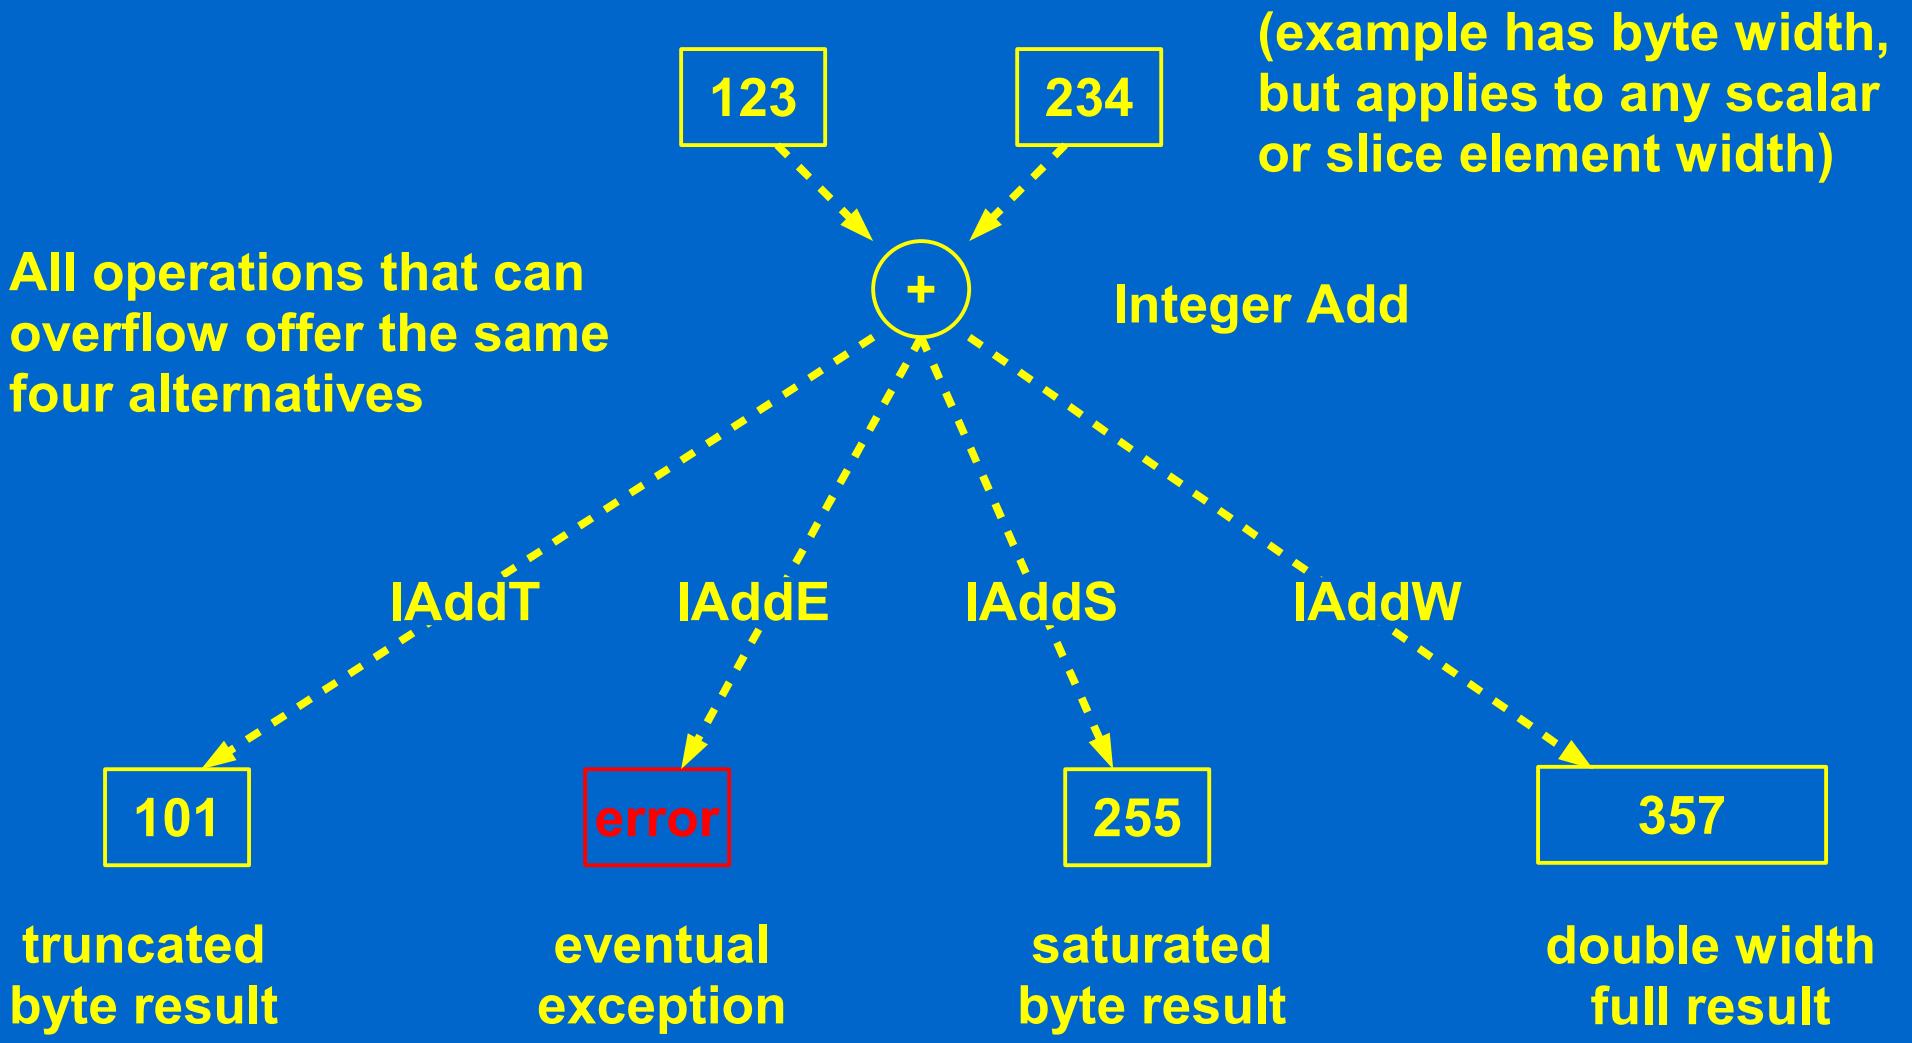

Mill status

- OOTBC is a five year old bootstrap startup
- Tool chain works, no vectorization yet
- Cycle-accurate simulator; porting Linux
- Key hardware in Verilog; FPGA soon
- Standard cell ASIC process
- 10X per thread vs. conventional CPUs on general purpose mixed workloads (ops/second/Watt/mm^2)
- 60+ patents in process

Contact:

Out-of-the-Box Computing

Ivan Godard

ivan@ootbcomp.com

Additional technical details under NDA

We seek strategic partners and Bay Area software and hardware technical contributors