Sandia National Laboratories

Exceptional service in the national interest

From naïve to smart: leveraging offloaded capabilities to enable intelligent NICs

Whit Schonbein
wwschon@sandia.gov

2021-12-02

SAND2021-15240 C

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

U.S. DEPARTMENT OF ENERGY

NNSA

What is a SmartNIC?

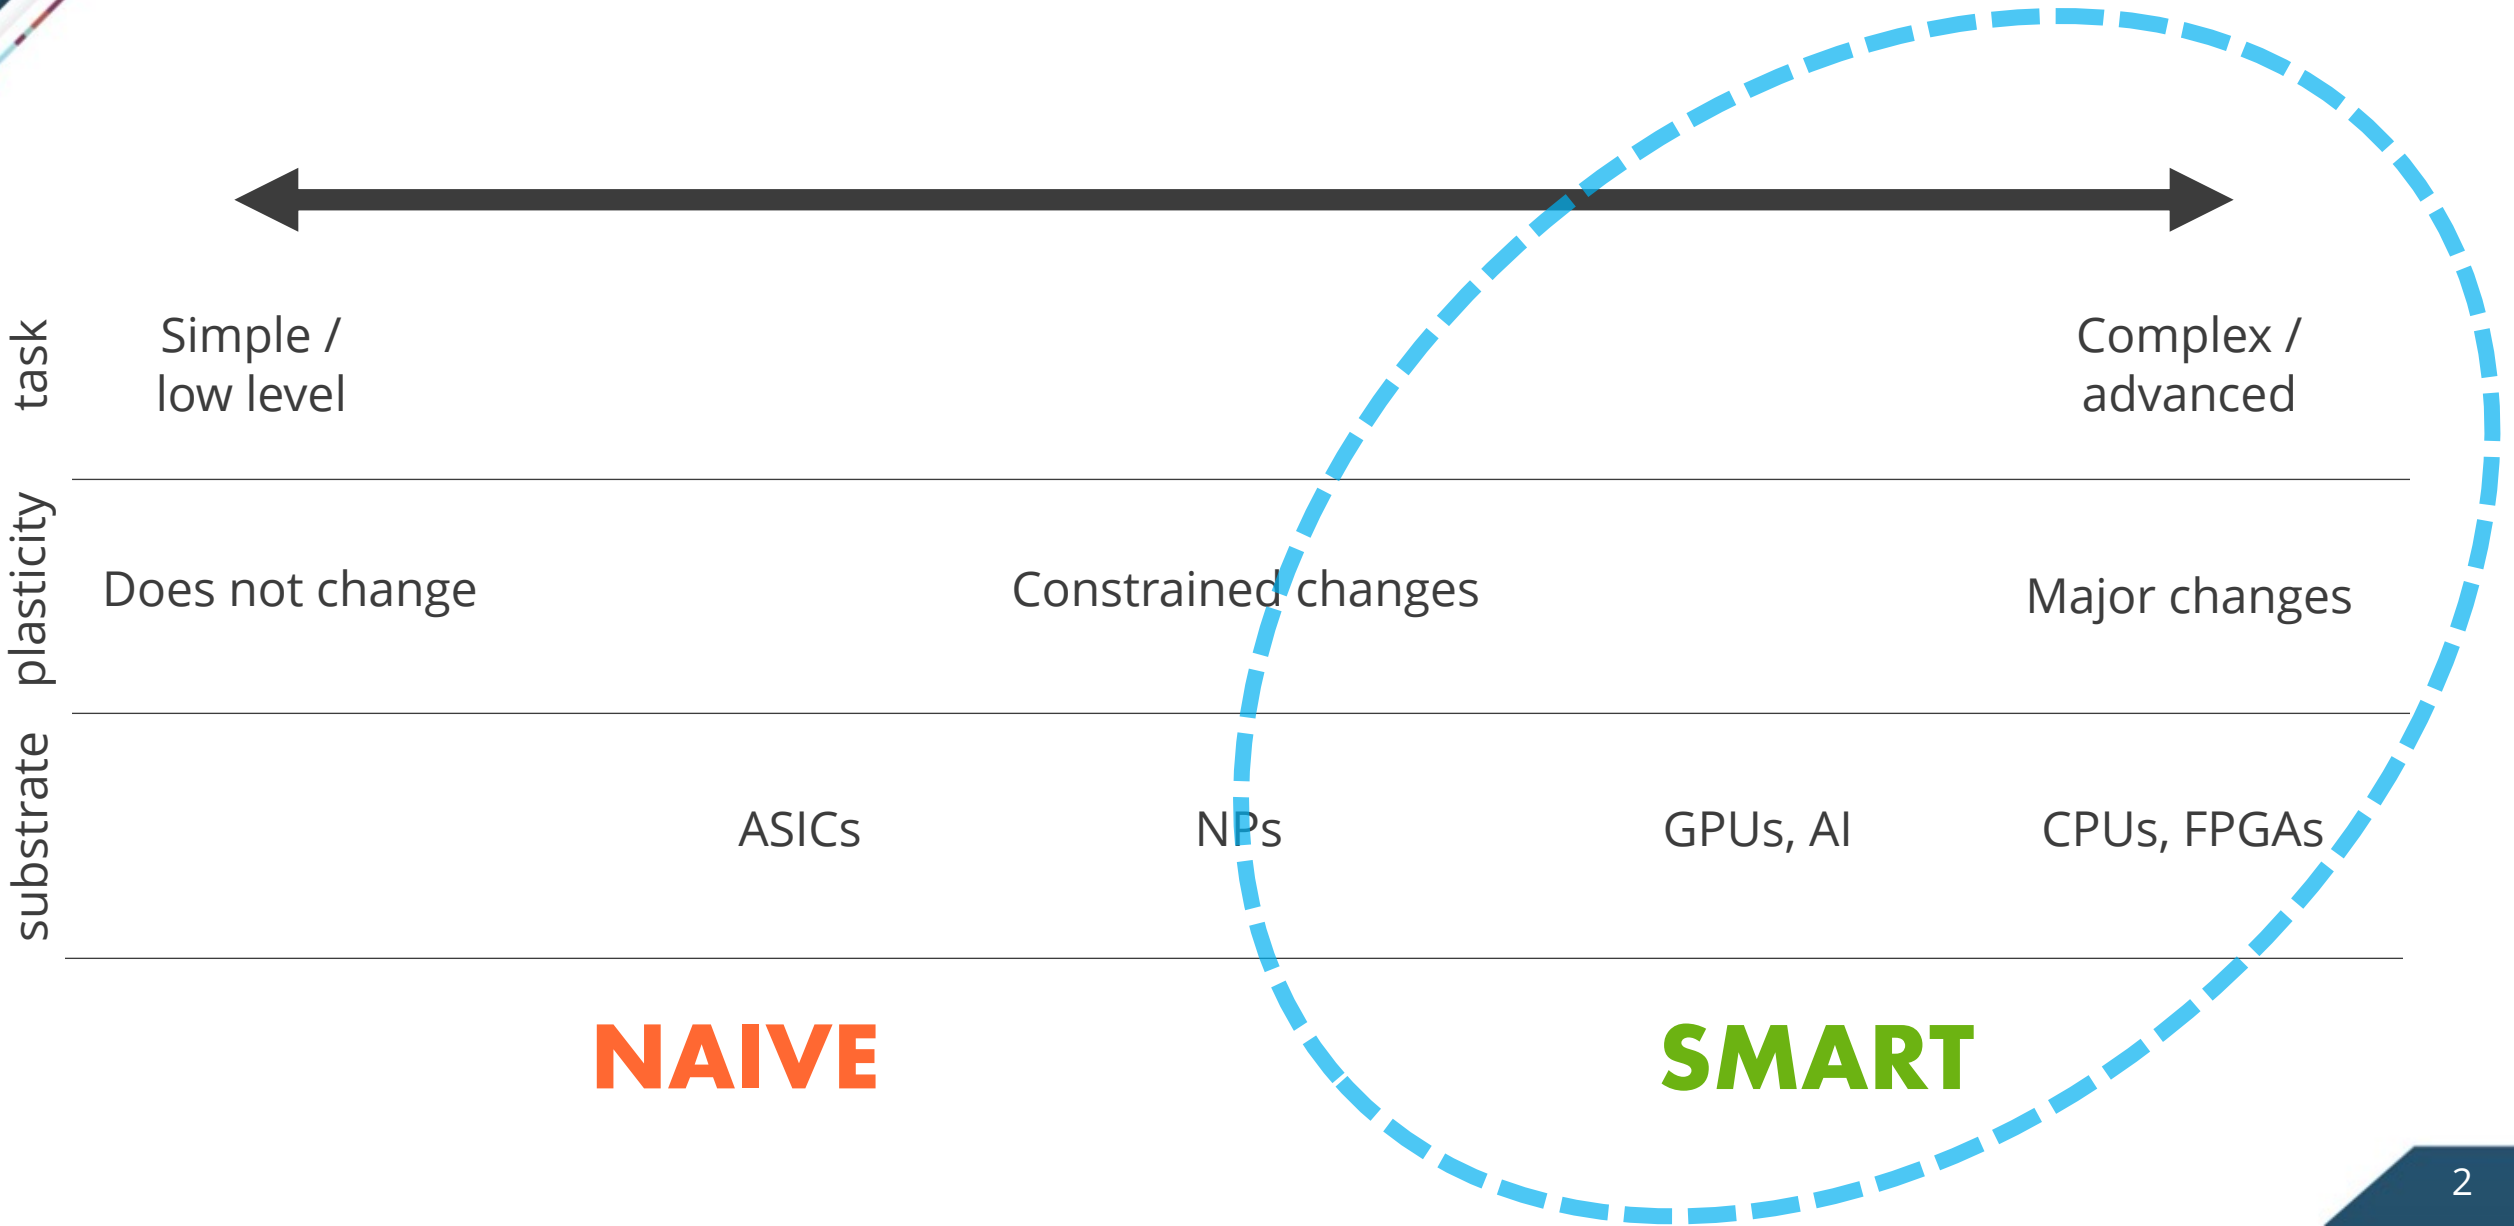

What is a SmartNIC?

- A NIC that can be configured to perform arbitrary tasks.
- Coordination of heterogeneous task-specific accelerators (GPUs, AI engines, etc.)

Portals Network API

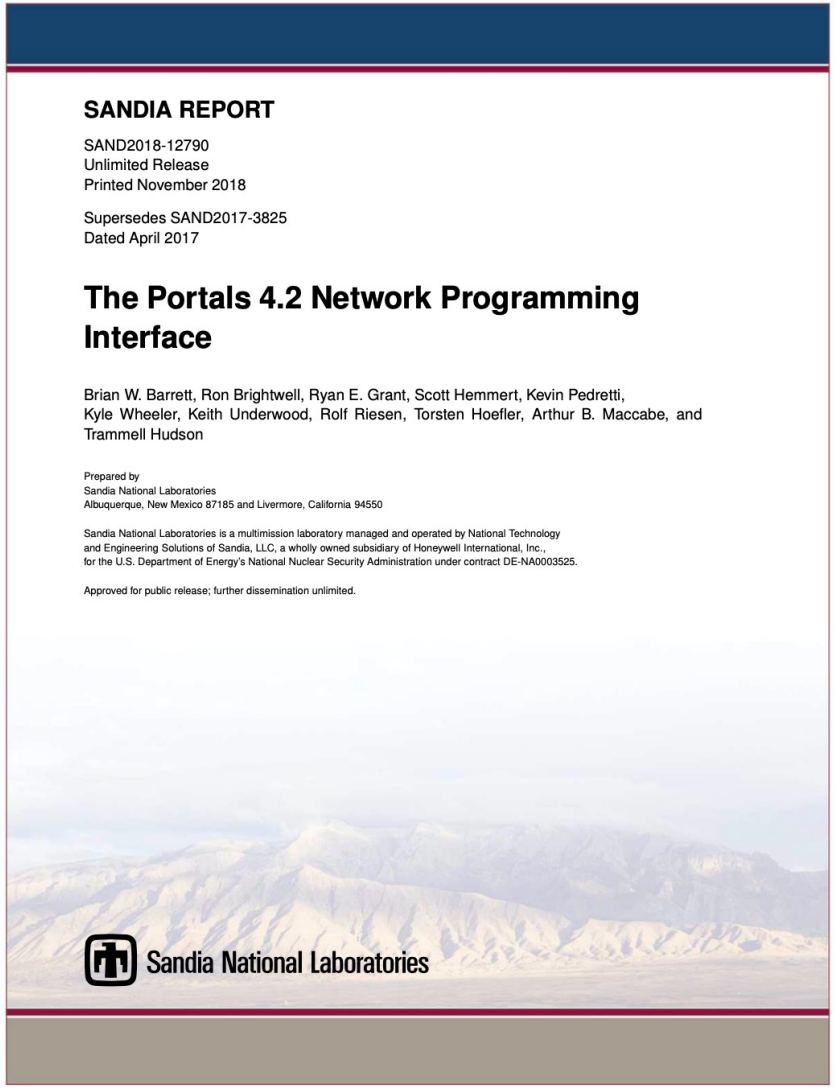

SANDIA REPORT

SAND2018-12790
Unlimited Release
Printed November 2018

Supersedes SAND2017-3825
Dated April 2017

The Portals 4.2 Network Programming Interface

Brian W. Barrett, Ron Brightwell, Ryan E. Grant, Scott Hemmert, Kevin Pedretti, Kyle Wheeler, Keith Underwood, Rolf Riesen, Torsten Hoefler, Arthur B. Maccabe, and Trammell Hudson

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

- Network programming API
- Hardware oriented

Portals Network API

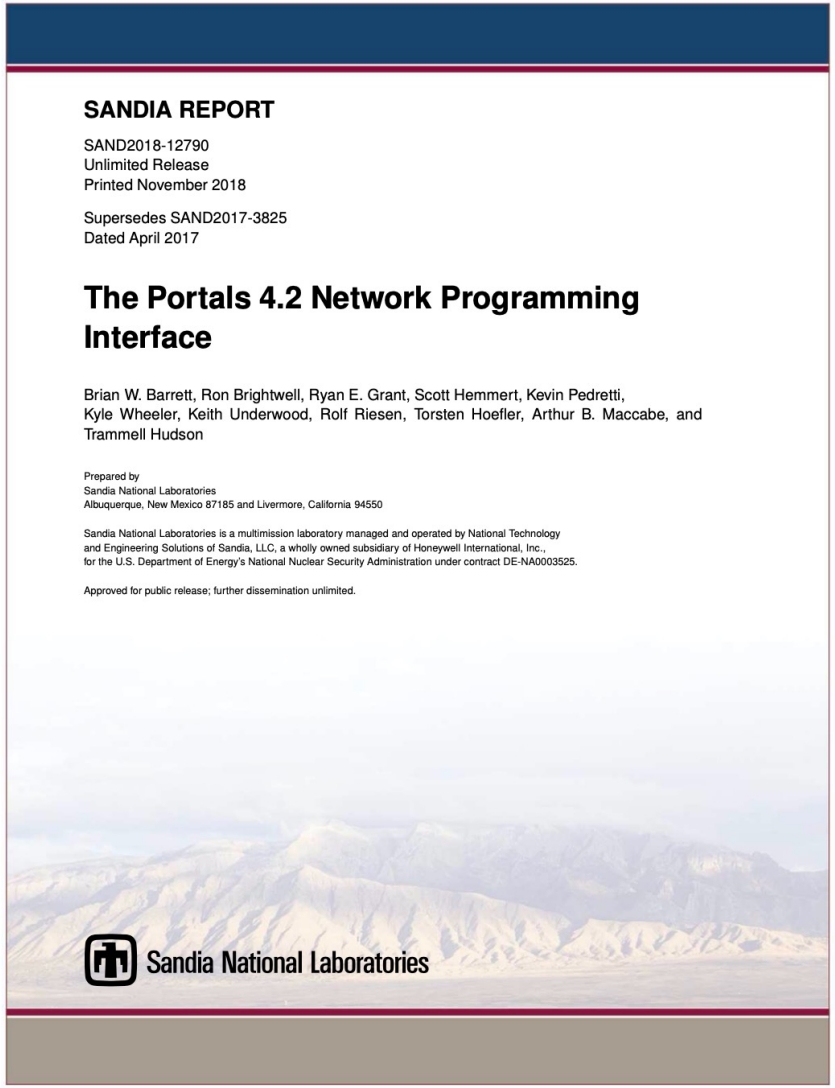

SANDIA REPORT

SAND2018-12790
Unlimited Release
Printed November 2018

Supersedes SAND2017-3825
Dated April 2017

The Portals 4.2 Network Programming Interface

Brian W. Barrett, Ron Brightwell, Ryan E. Grant, Scott Hemmert, Kevin Pedretti, Kyle Wheeler, Keith Underwood, Rolf Riesen, Torsten Hoefler, Arthur B. Maccabe, and Trammell Hudson

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

- How can a Portals NIC support `intelligent' offloads?

Portals Network API

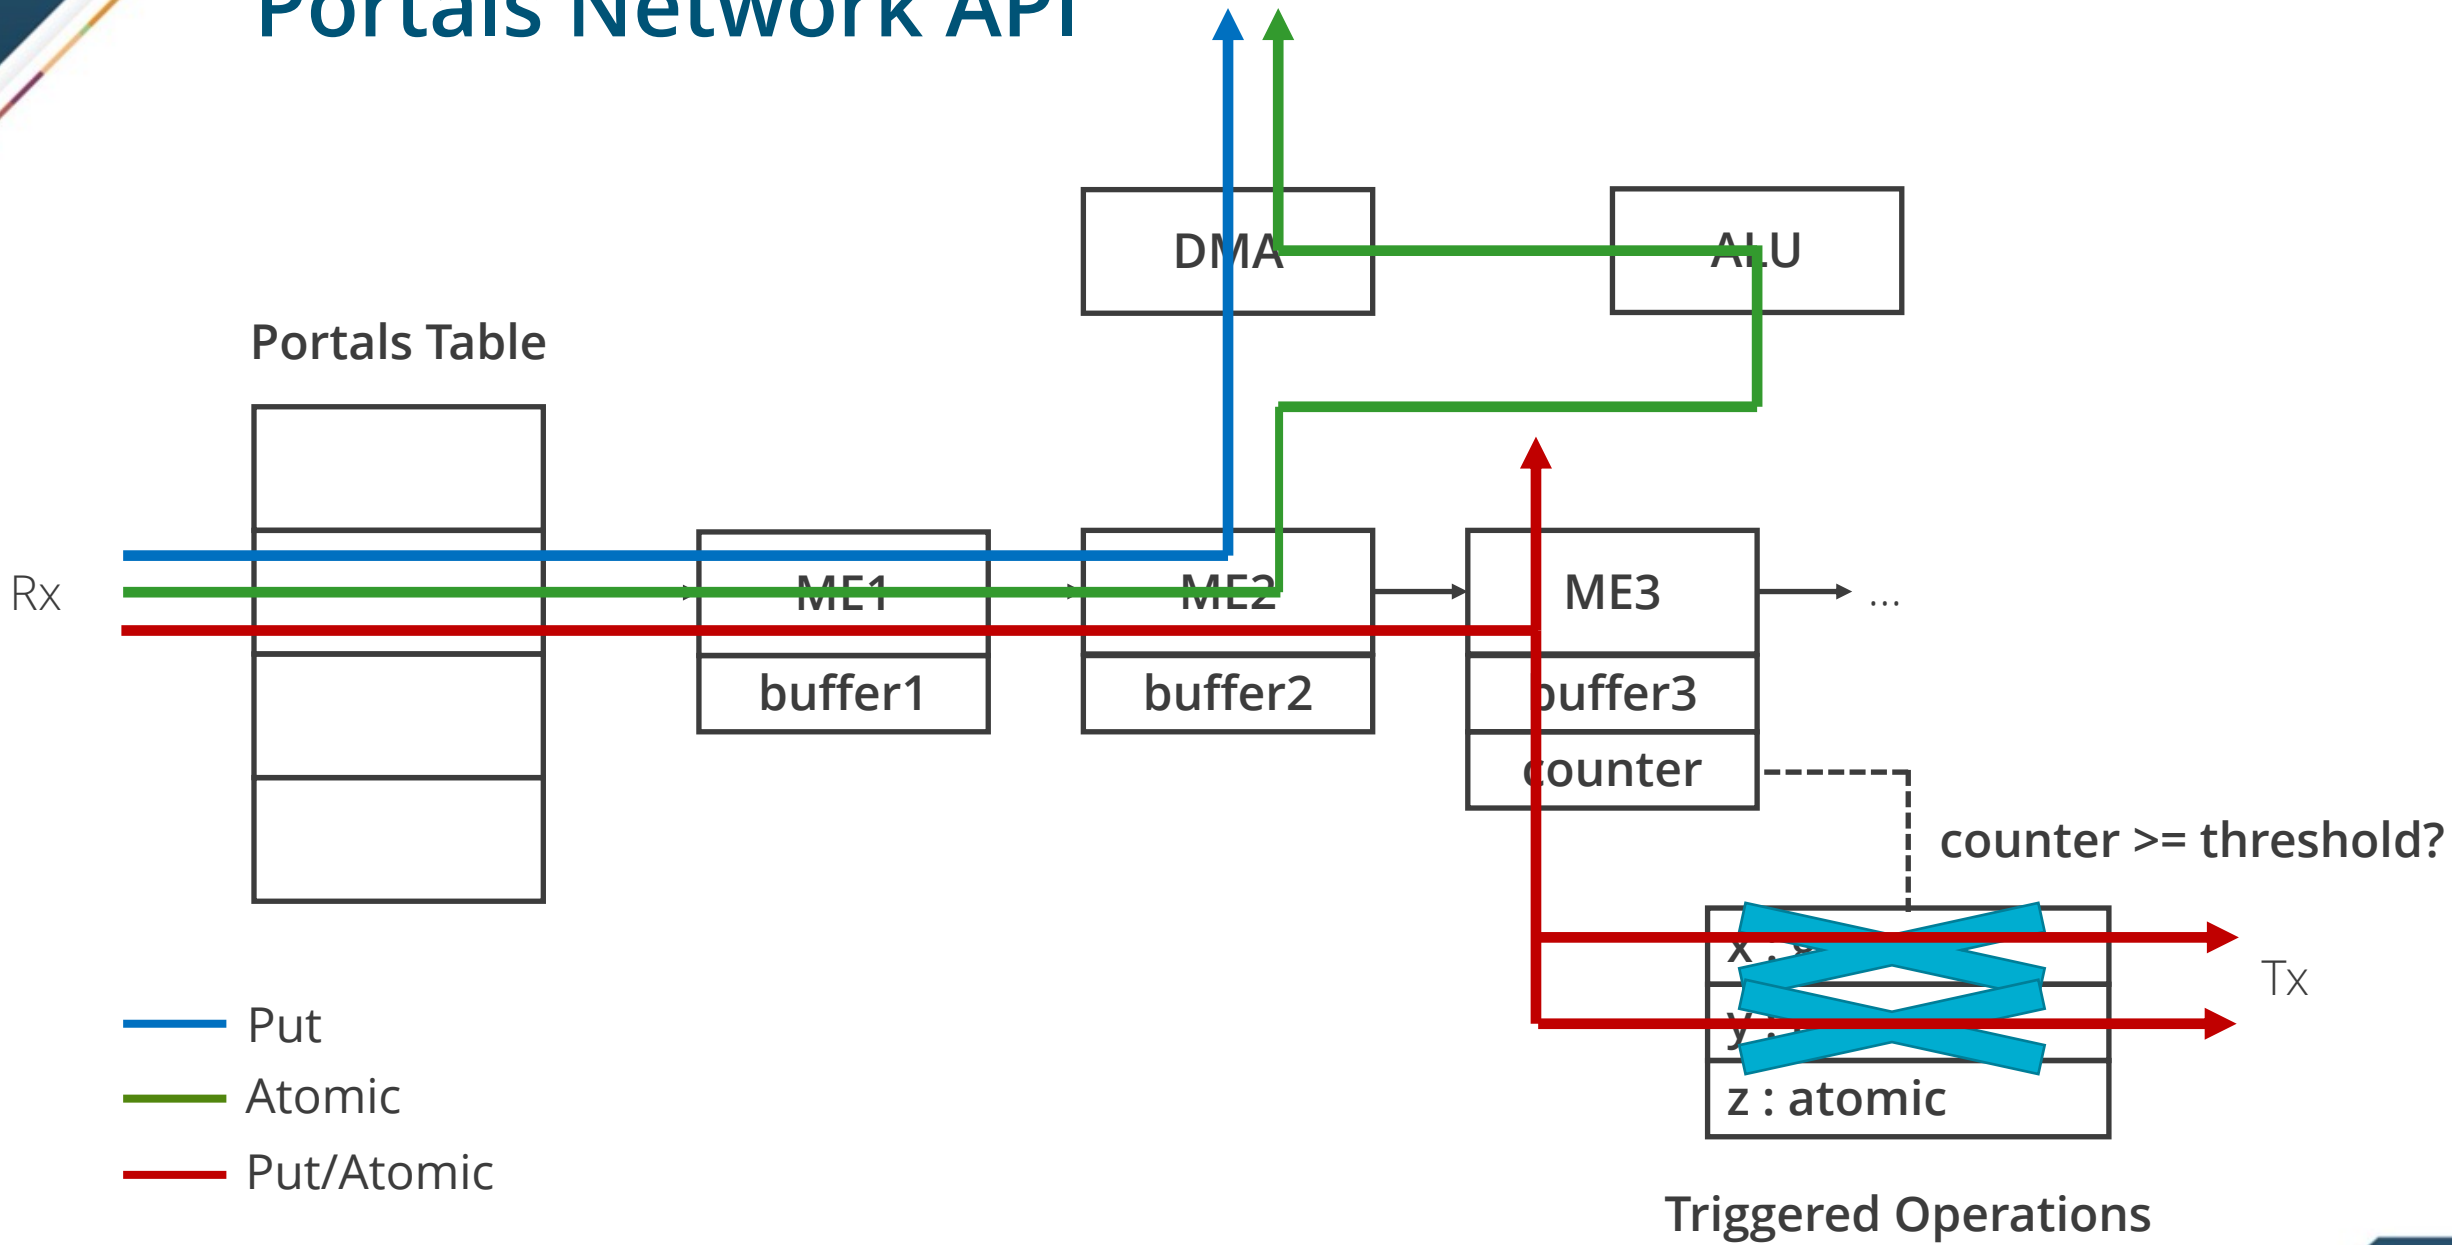

Extending Portals

Current (4.2)		Proposed
Target-side resource management: buffers, offsets	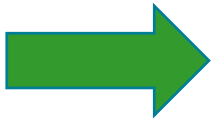	+ ME-specified operations
Transient triggered operations	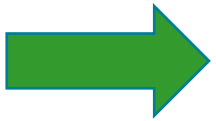	+ persistent triggered operations
Triggering conditions: >=	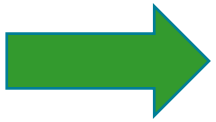	+ strict indexing (==)
Counter updates: bytes, writes	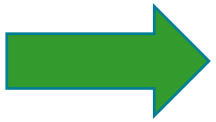	+ conditional counter updates

Extending Portals

+ ME-specified operation:
counter = (cond ? val1 : val2)

```
if [buffer] <= 0:
        counter = x
else:
        counter++
```

Extending Portals

Current (4.2)		Proposed
Target-side resource management: buffers, offsets	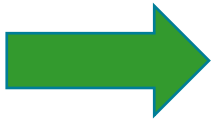	+ ME-specified operations
Transient triggered operations	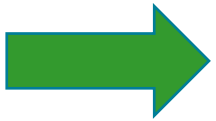	+ persistent triggered operations
Triggering conditions: >=	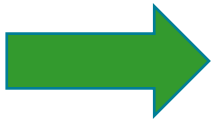	+ strict indexing (==)
Counter updates: bytes, writes	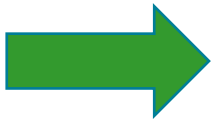	+ conditional counter updates

Extending Portals

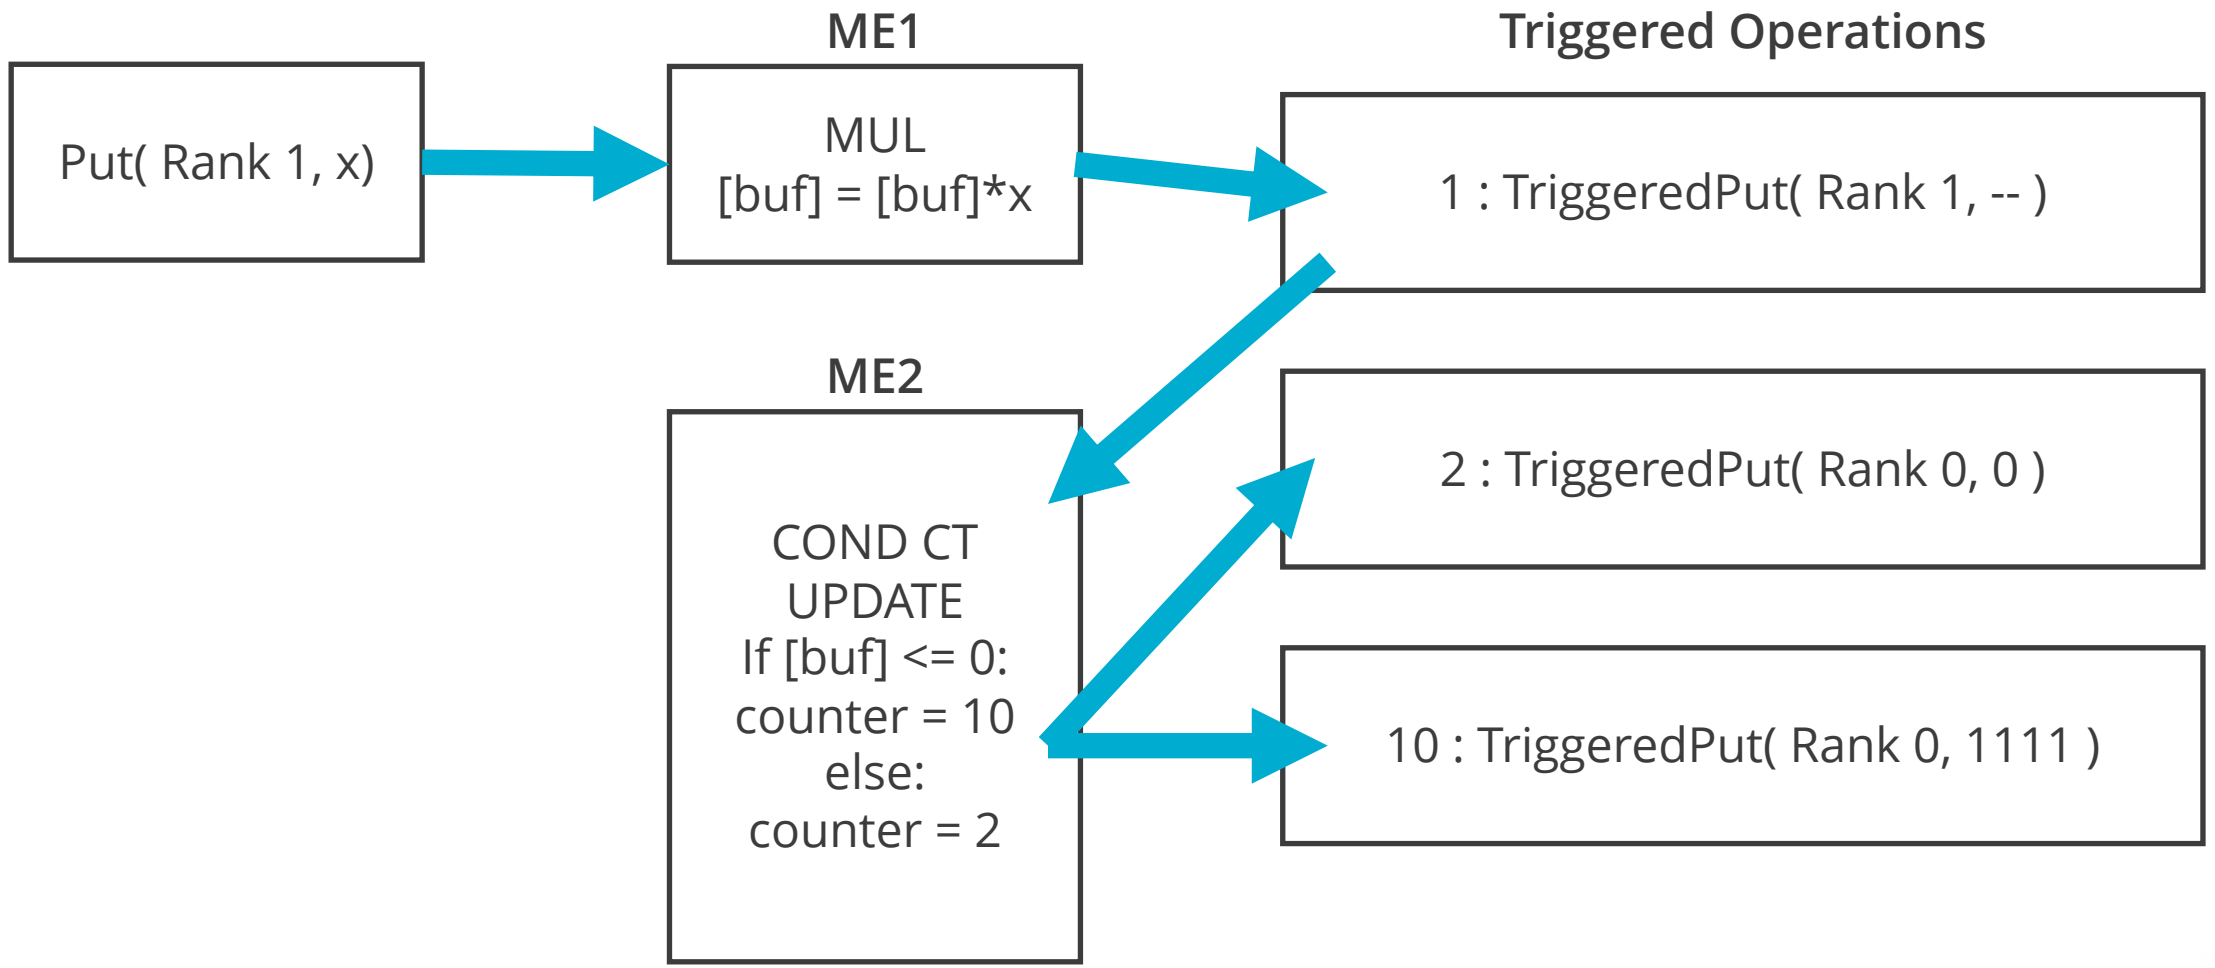

Enabling Intelligence

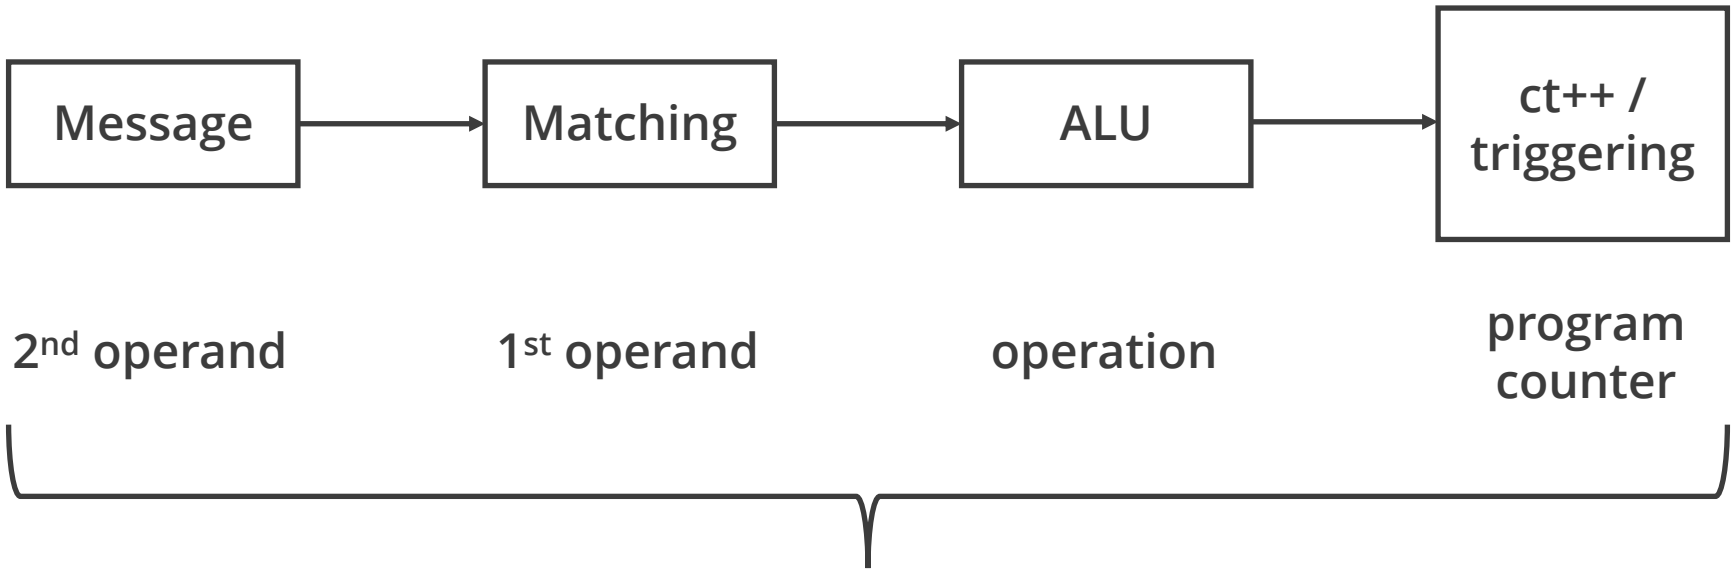

Program: Sequence of instructions sharing the same counter

Turing Equivalent

INCA: In-Network Compute Assistance

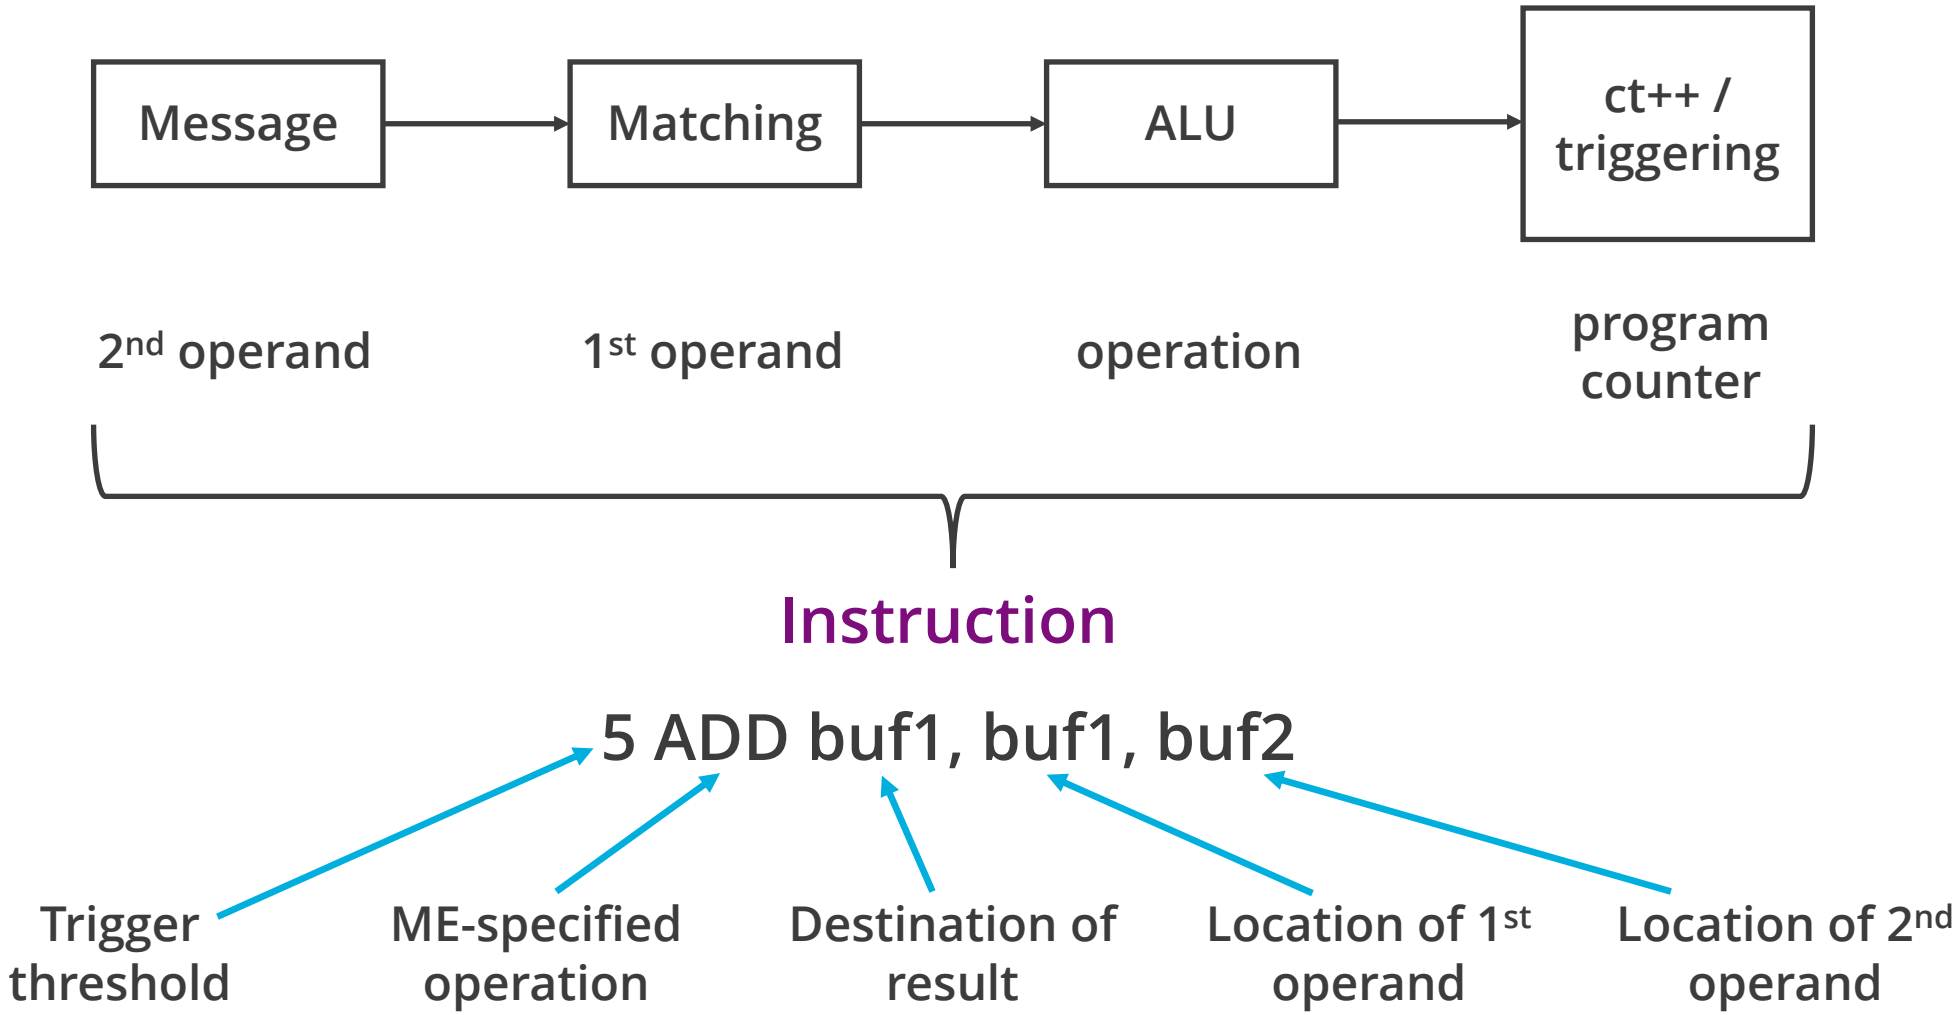

INCA: In-Network Compute Assistance

```
i = 0
while i < 16 do {
    c = c + (A[i] * B[i])
    i = i + 1
}
end
```

Dot product

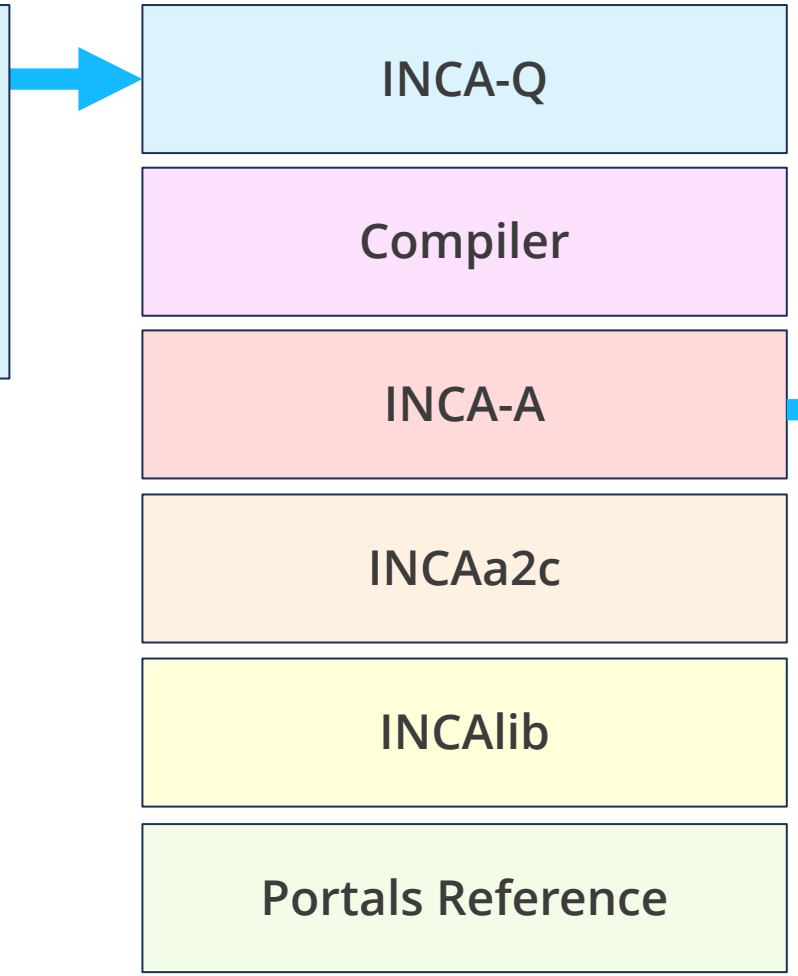

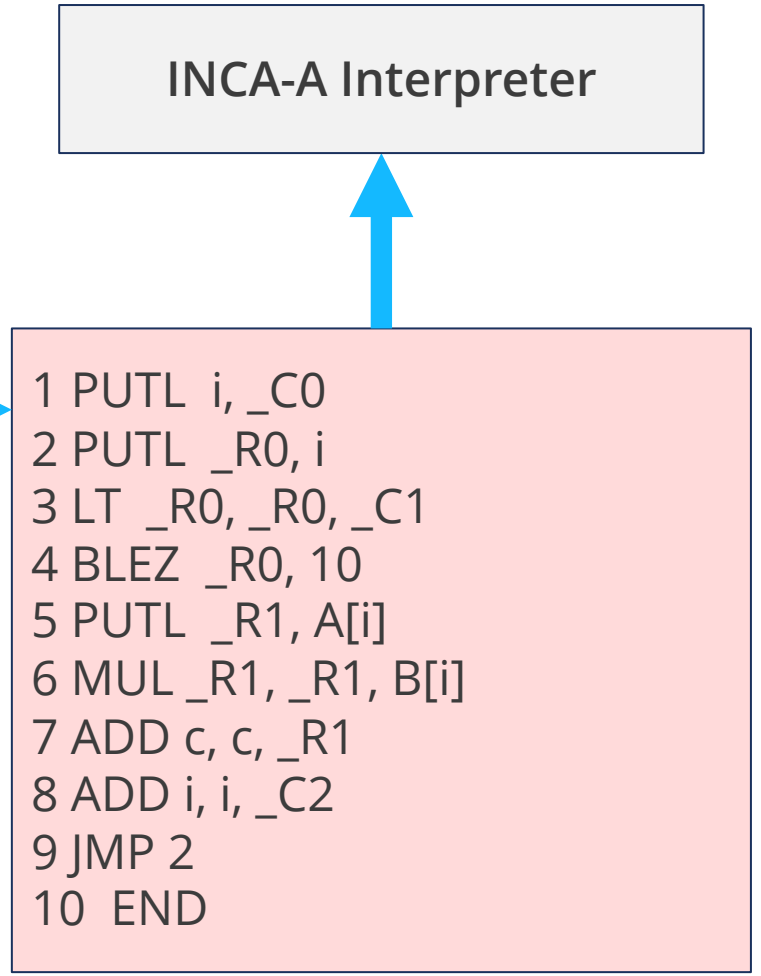

INCA: In-Network Compute Assistance

Kernel	usecs
convolution	190.49
dot-product	31.76
hadamard-product	31.81
linear-interpolation	212.85
matrix-multiplication	1051.20
matrix-transpose	23.62
unpack	59.93

3.2 GHz Haswell:
139.5 -- 10.5 usecs

400 MMsgs/s, 1ns scratchpad
8KiB payloads

Enabling Intelligence

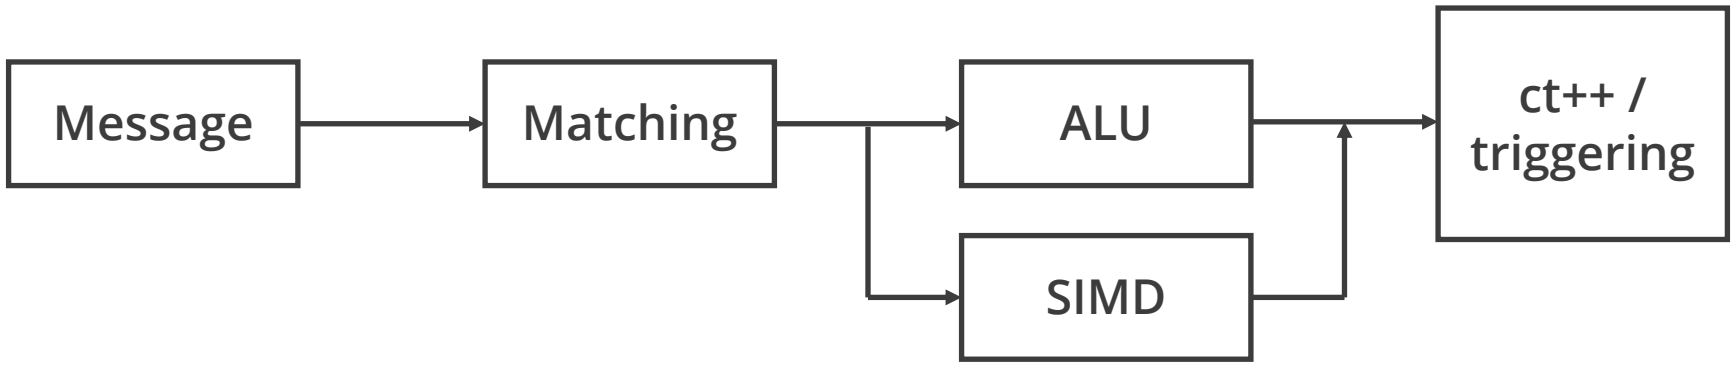

2 MULM A[_Z0], A[_Z0], B[_Z0], 256

Kernel	usecs	usecs
dot-product	31.76	23.74
hadamard-product-pc	31.81	0.176
matrix-multiplication-p	1051.20	819.61
matrix-multiplication-pc	1051.20	153.60

Enabling Intelligence

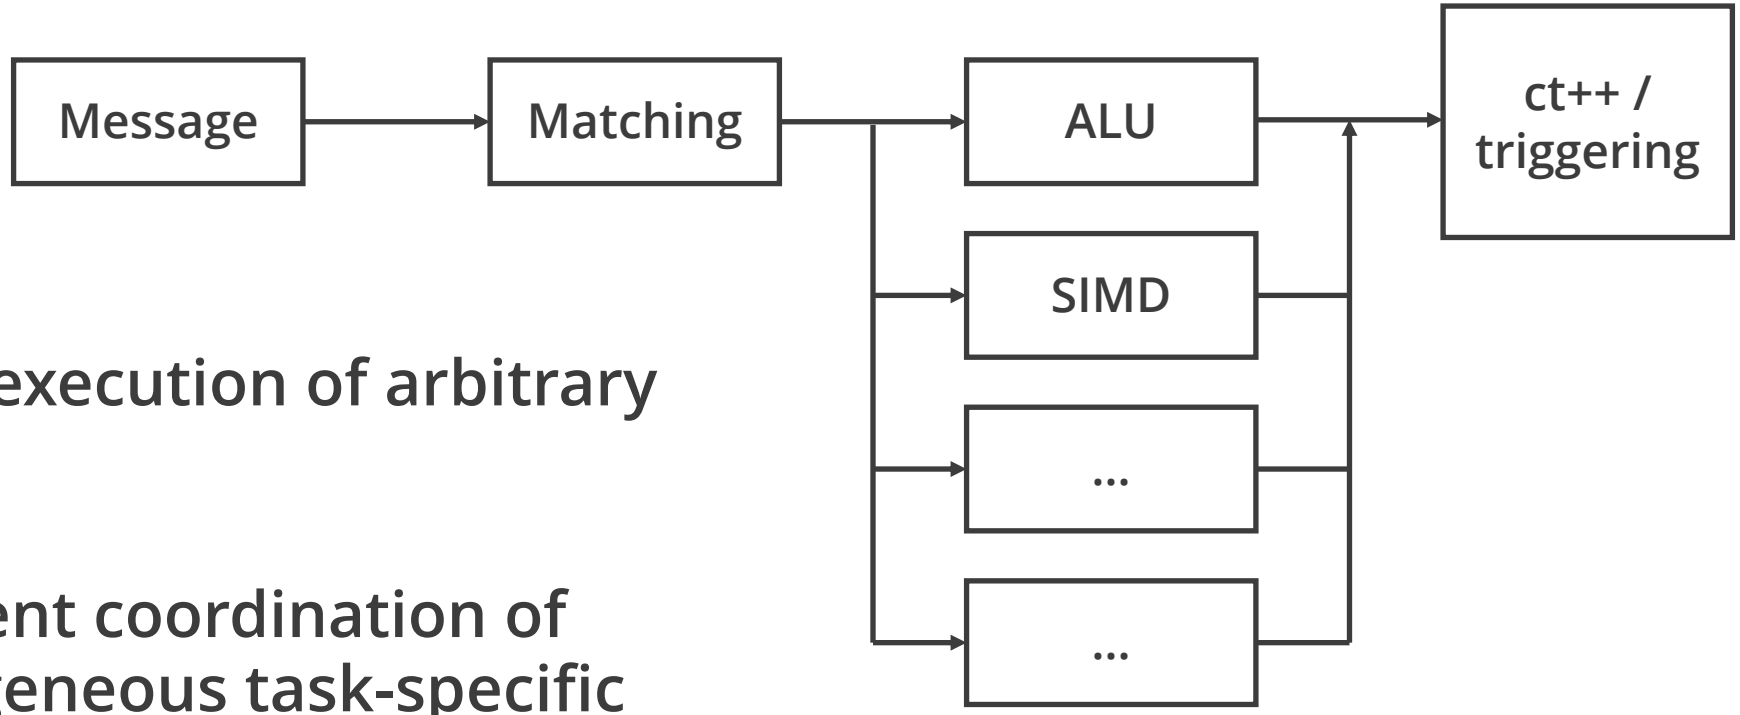

- Enable execution of arbitrary tasks
- Intelligent coordination of heterogeneous task-specific accelerators

Thank You

- Whit Schonbein
  wwschon@sandia.gov

- Portals specification
  https://cs.sandia.gov/Portals/ [BROKEN; but see researchgate]

- Portals reference implementation
  https://github.com/Portals4