



Microsoft Research

Faculty
Summit

2014 15TH ANNUAL



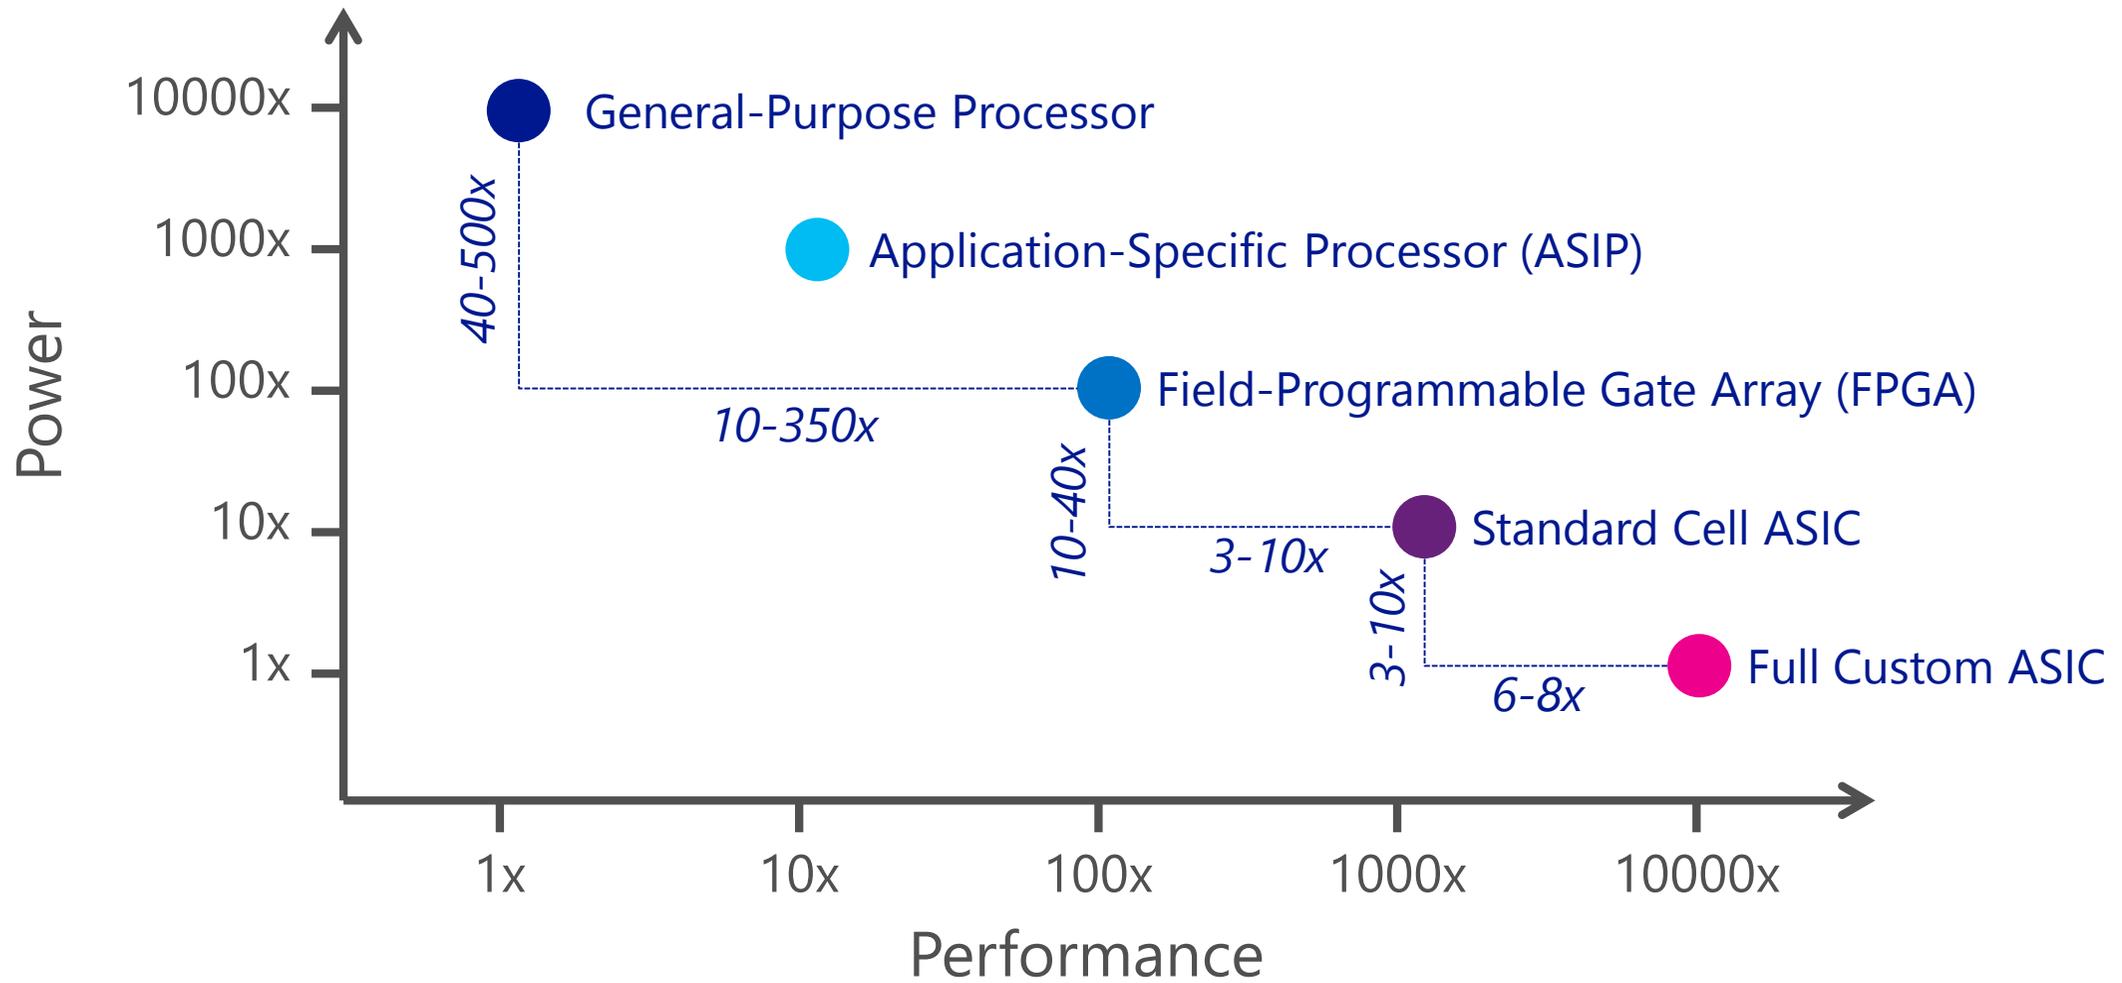
Microsoft Research
Faculty
Summit
2014 15TH ANNUAL

Specialization for Data Analytics

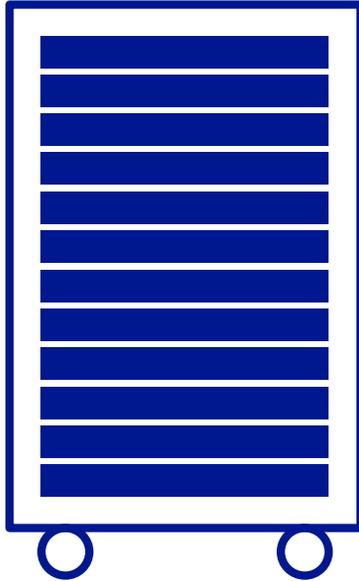
Martha Kim
Columbia University



Why Specialize?



Why Not Specialize?



Homogeneous (design, maintenance)

Uniform software interface

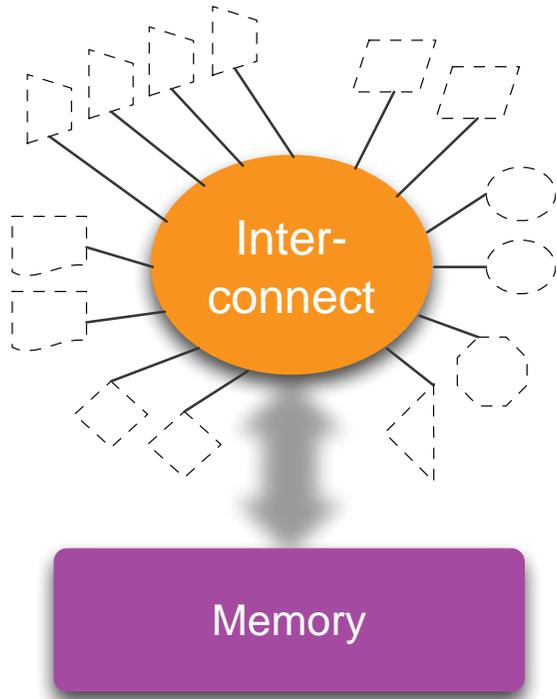
Arbitrary software



Heterogeneous

Complex, arcane

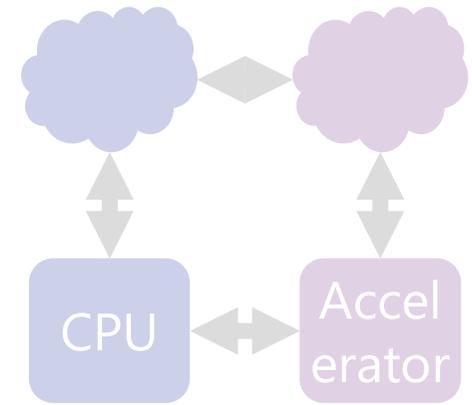
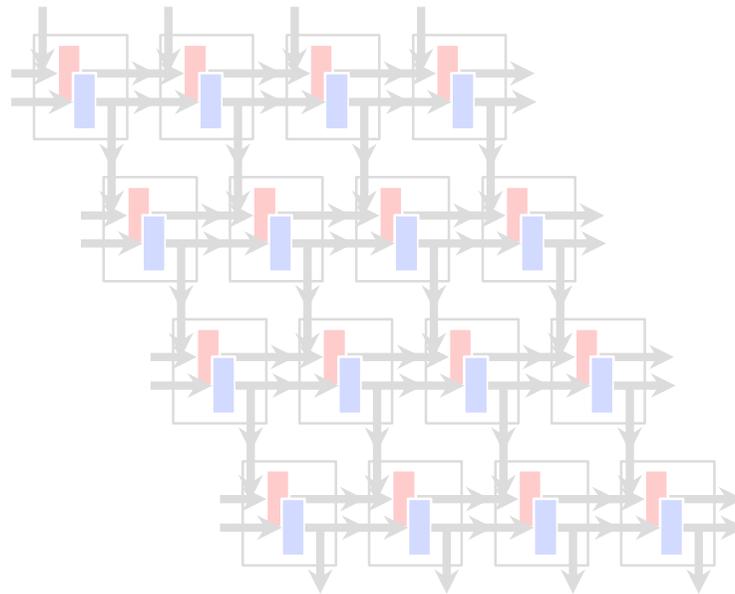
Fixed function(s)



Q100

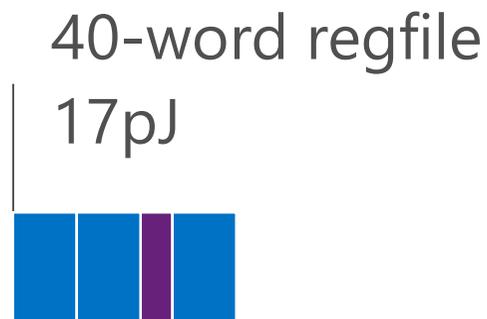
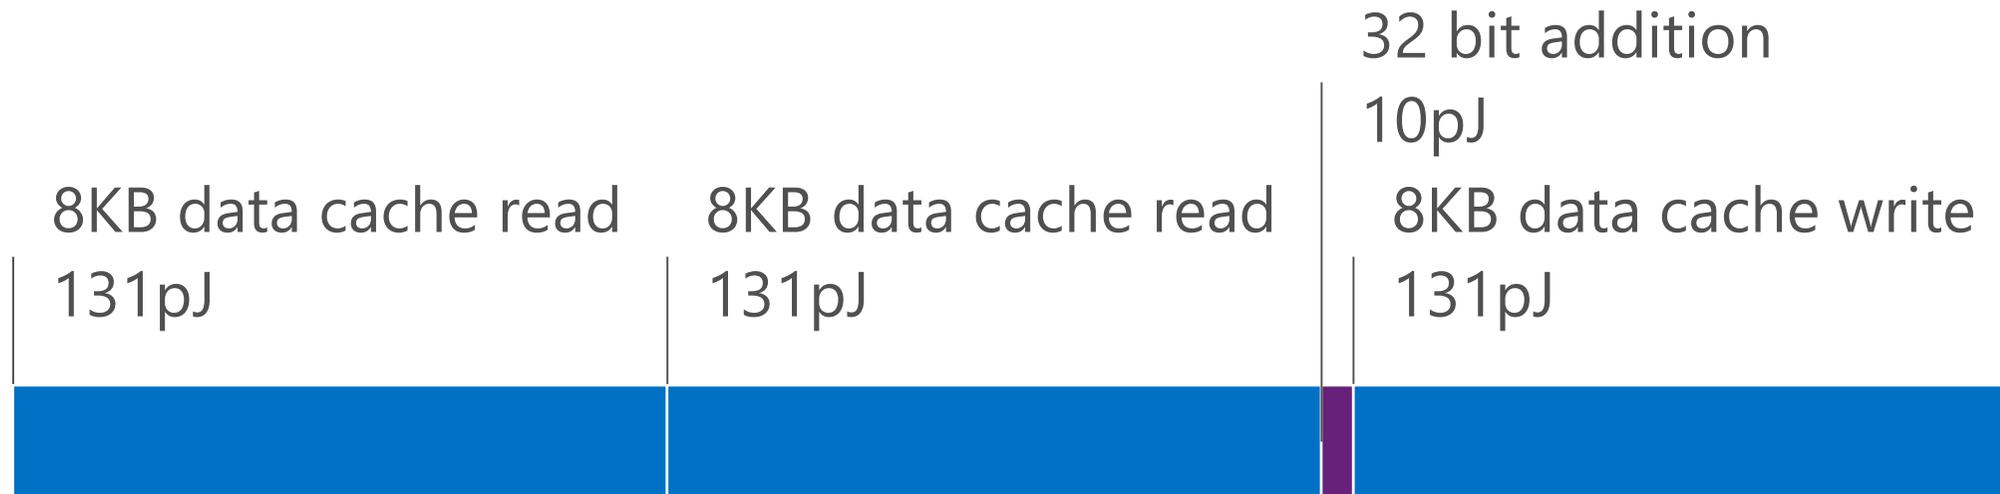
Data Processing Unit

Data Processing Unit
2.0



System Integration
Challenges

Data-Centric Specialization

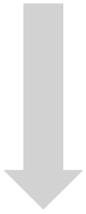


Data supply accounts for 83%-97% of total energy!

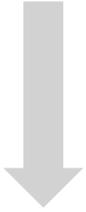
Q100 Database Processing Unit

- Targets analytical queries (not updates or txns)
- Hardware implementations of relational operators
- Processes streams representing columns or tables
- Spatial and temporal ISA

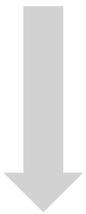
Query



Plan



Program



Result

```
SELECT s_season, SUM(s_qty) as sum_qty  
FROM sales  
WHERE s_shipdate >= '2013-01-01'  
GROUP BY s_season  
ORDER BY s_season
```

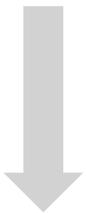
Query



Plan

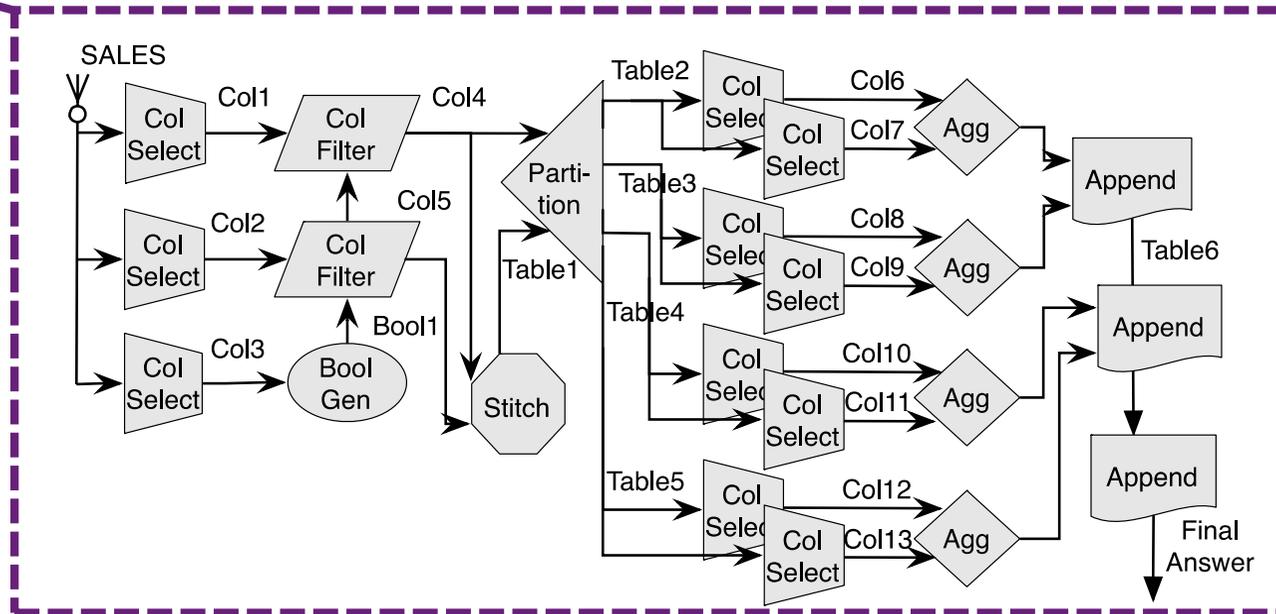


Program



Result

```
SELECT s_season, SUM(s_qty) as sum_qty
FROM sales
WHERE s_shipdate >= '2013-01-01'
GROUP BY s_season
ORDER BY s_season
```



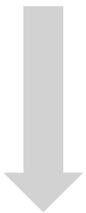
Query



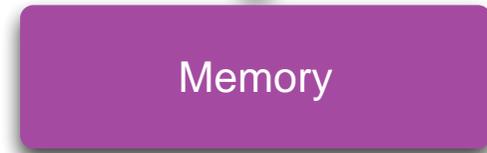
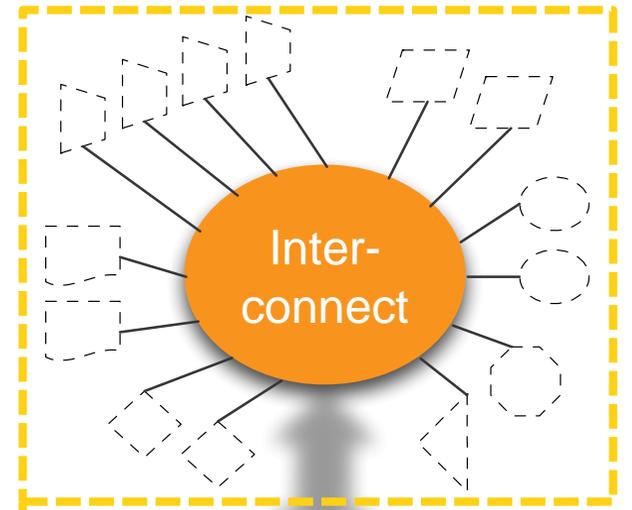
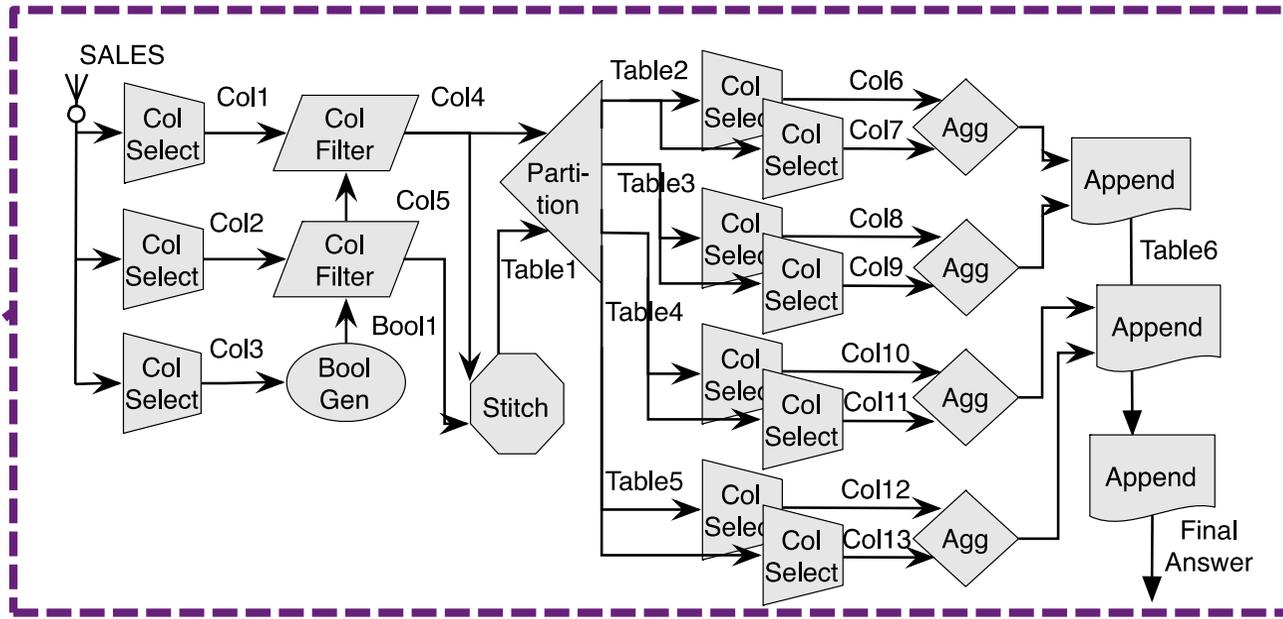
Plan



Program



Result



Q100 Device

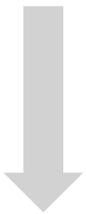
Query



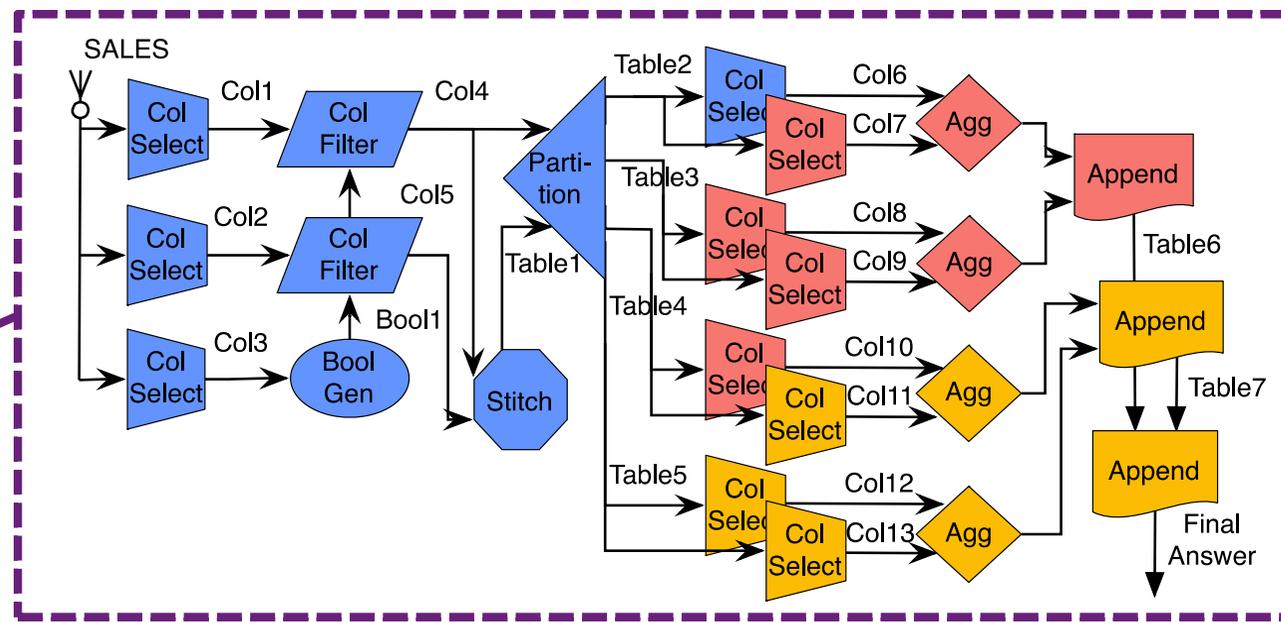
Plan



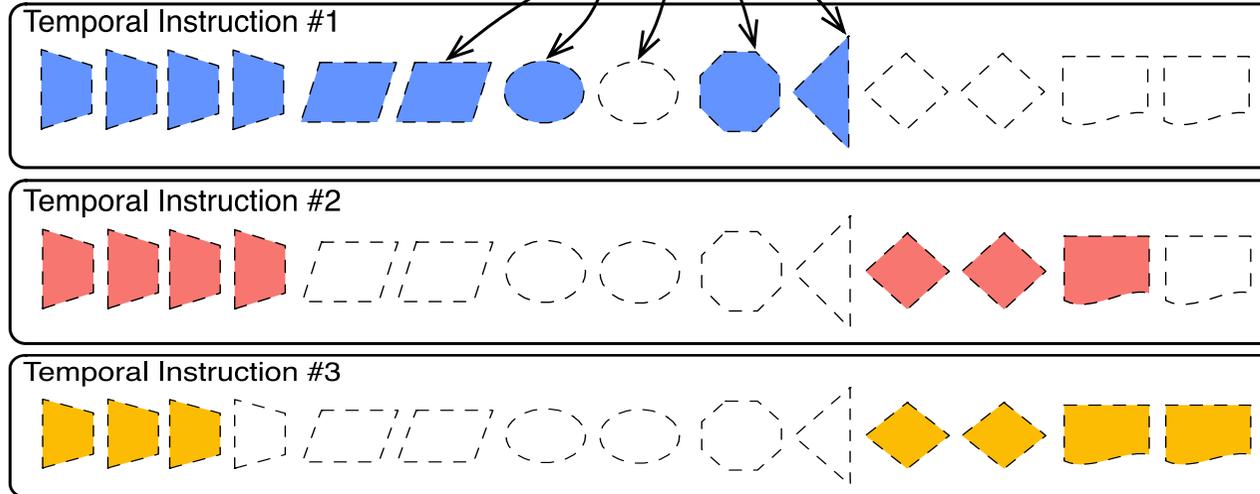
Program



Result



Spatial Instructions



Query



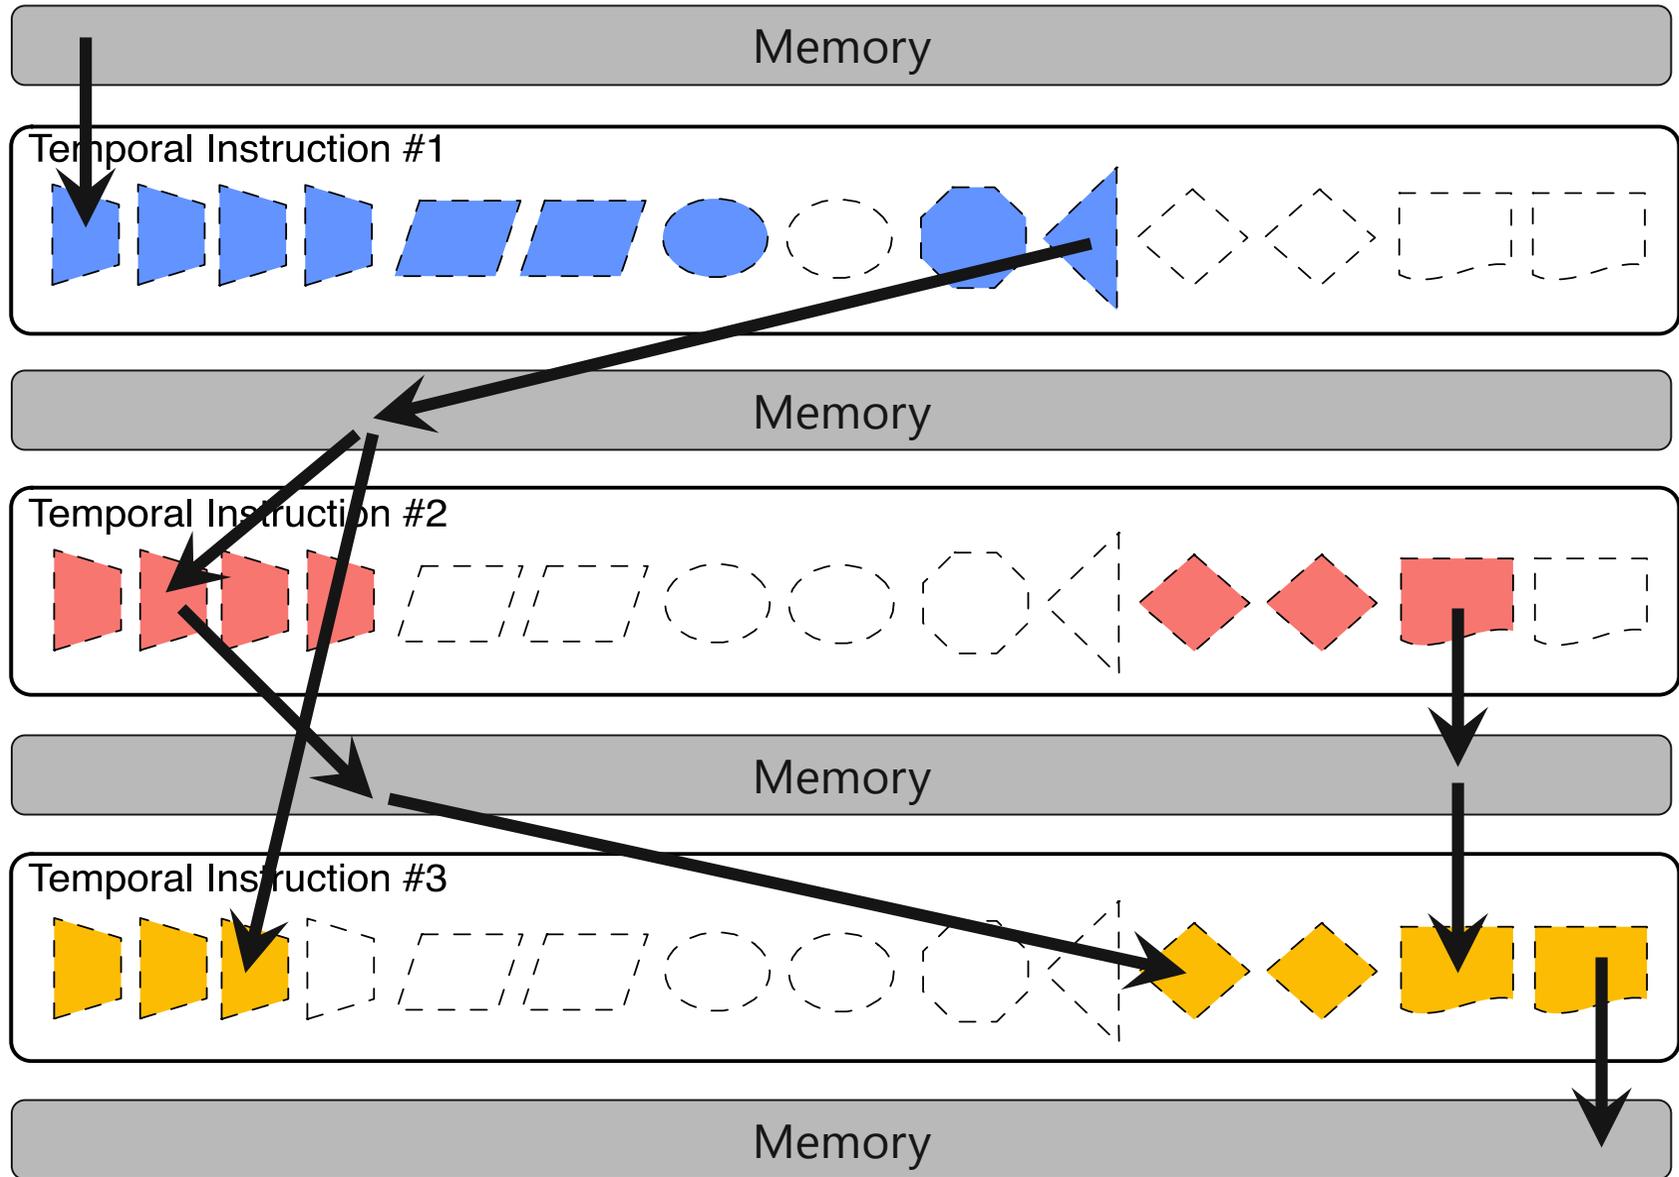
Plan



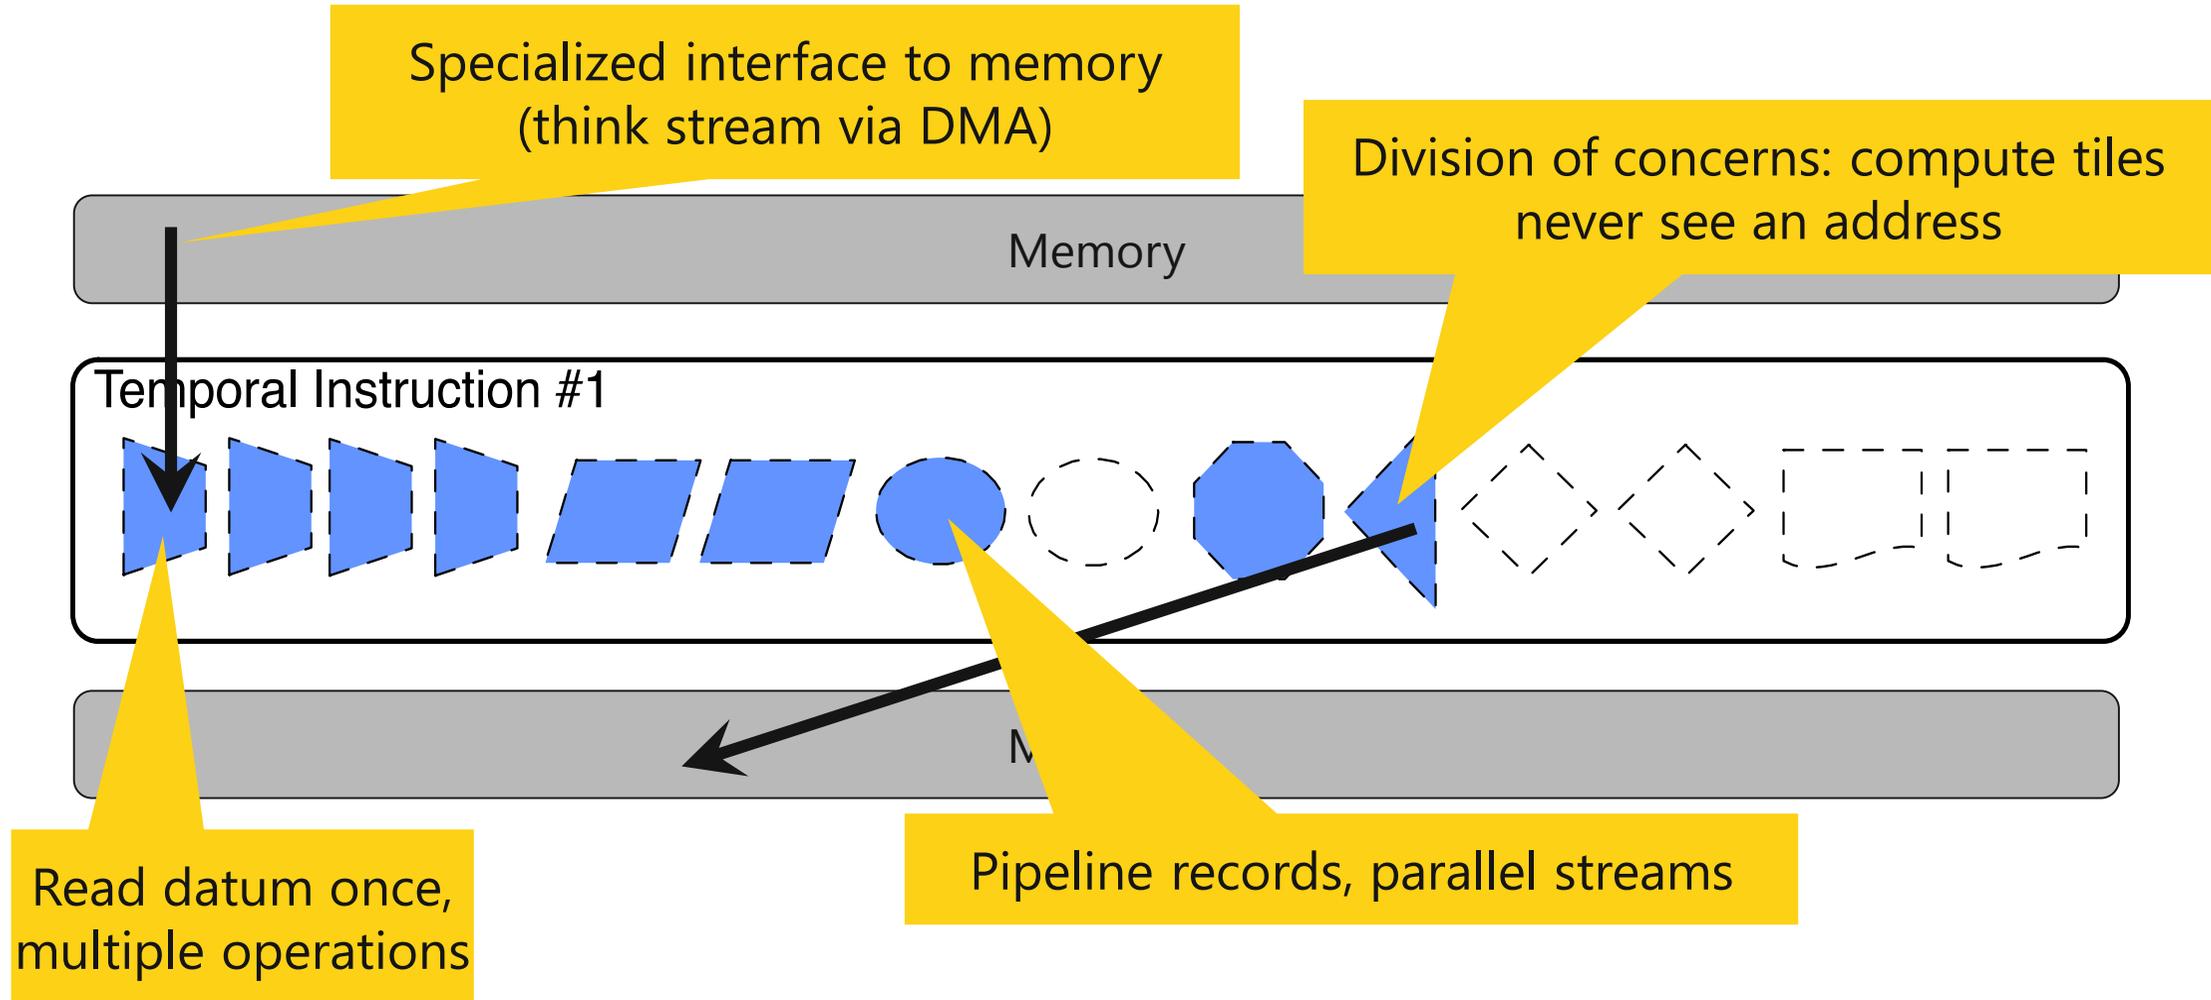
Program



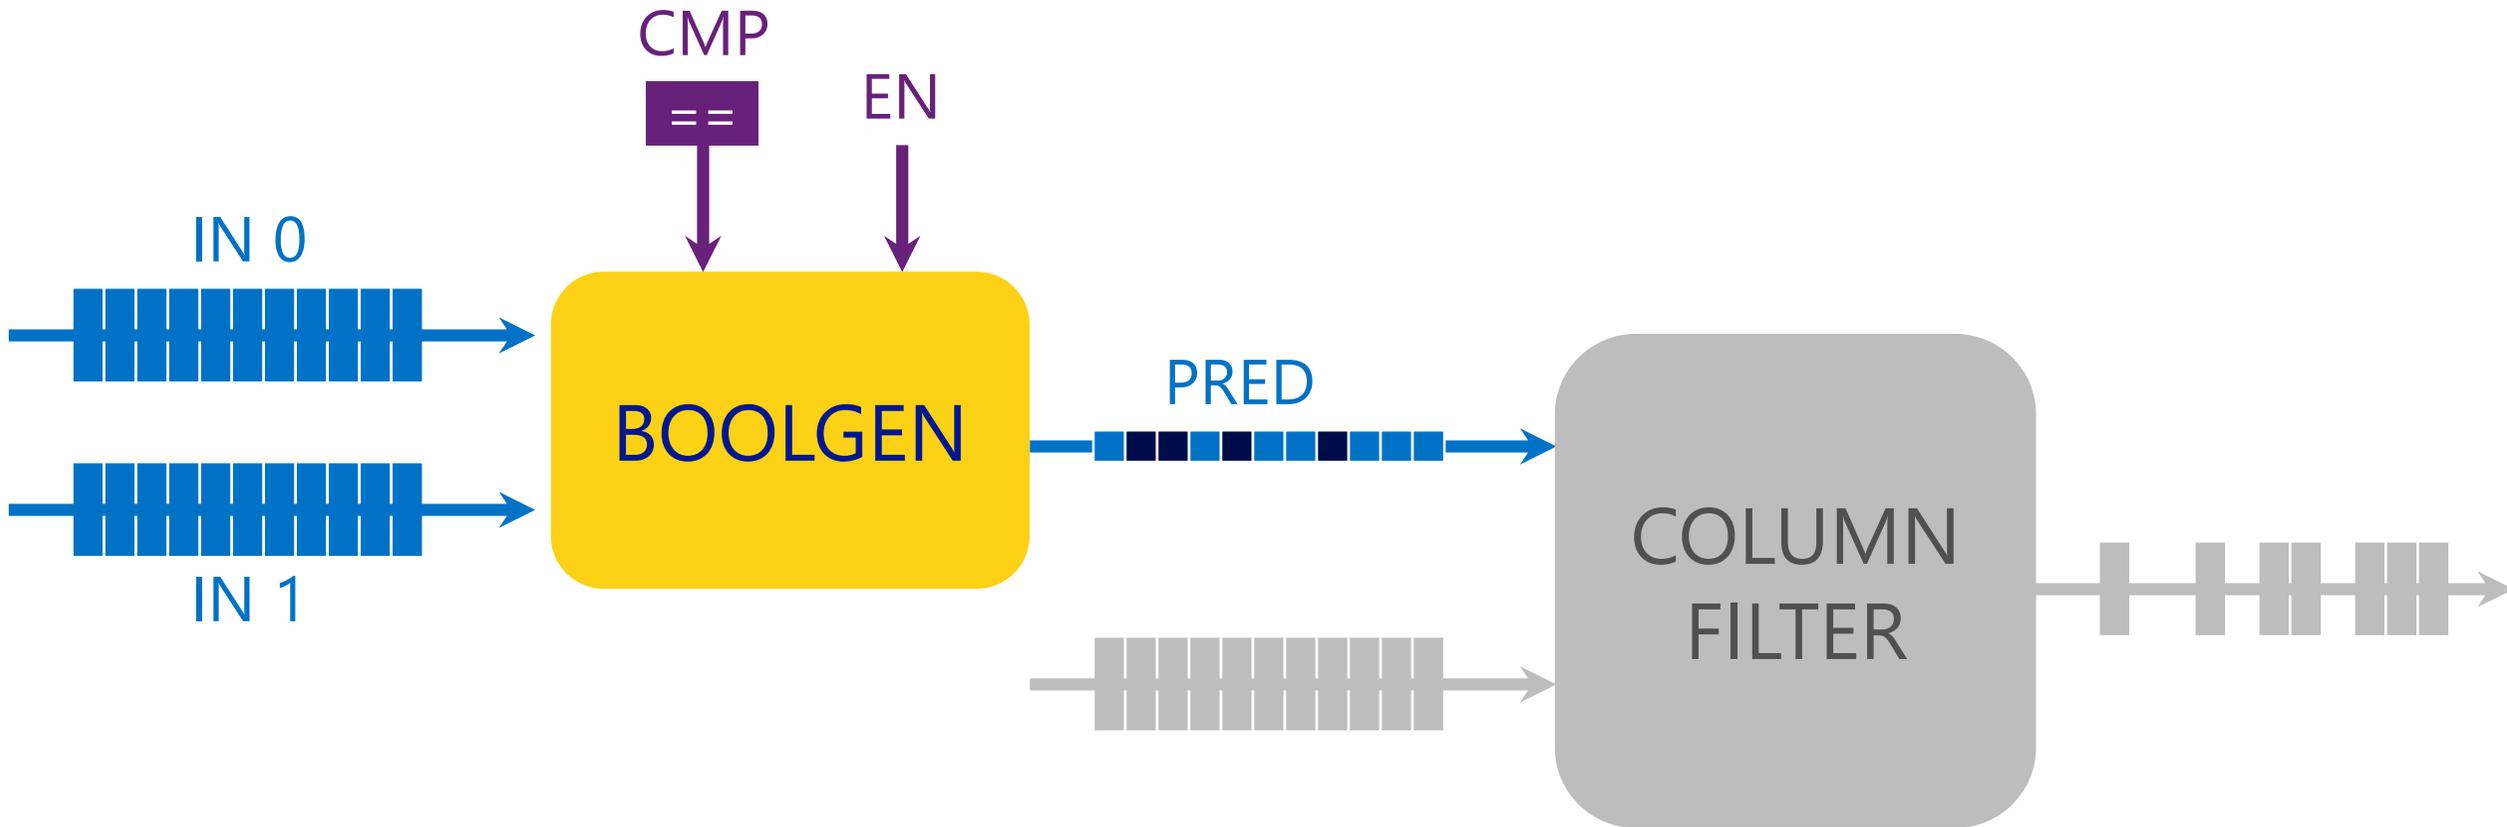
Result



Key Sources of Efficiency



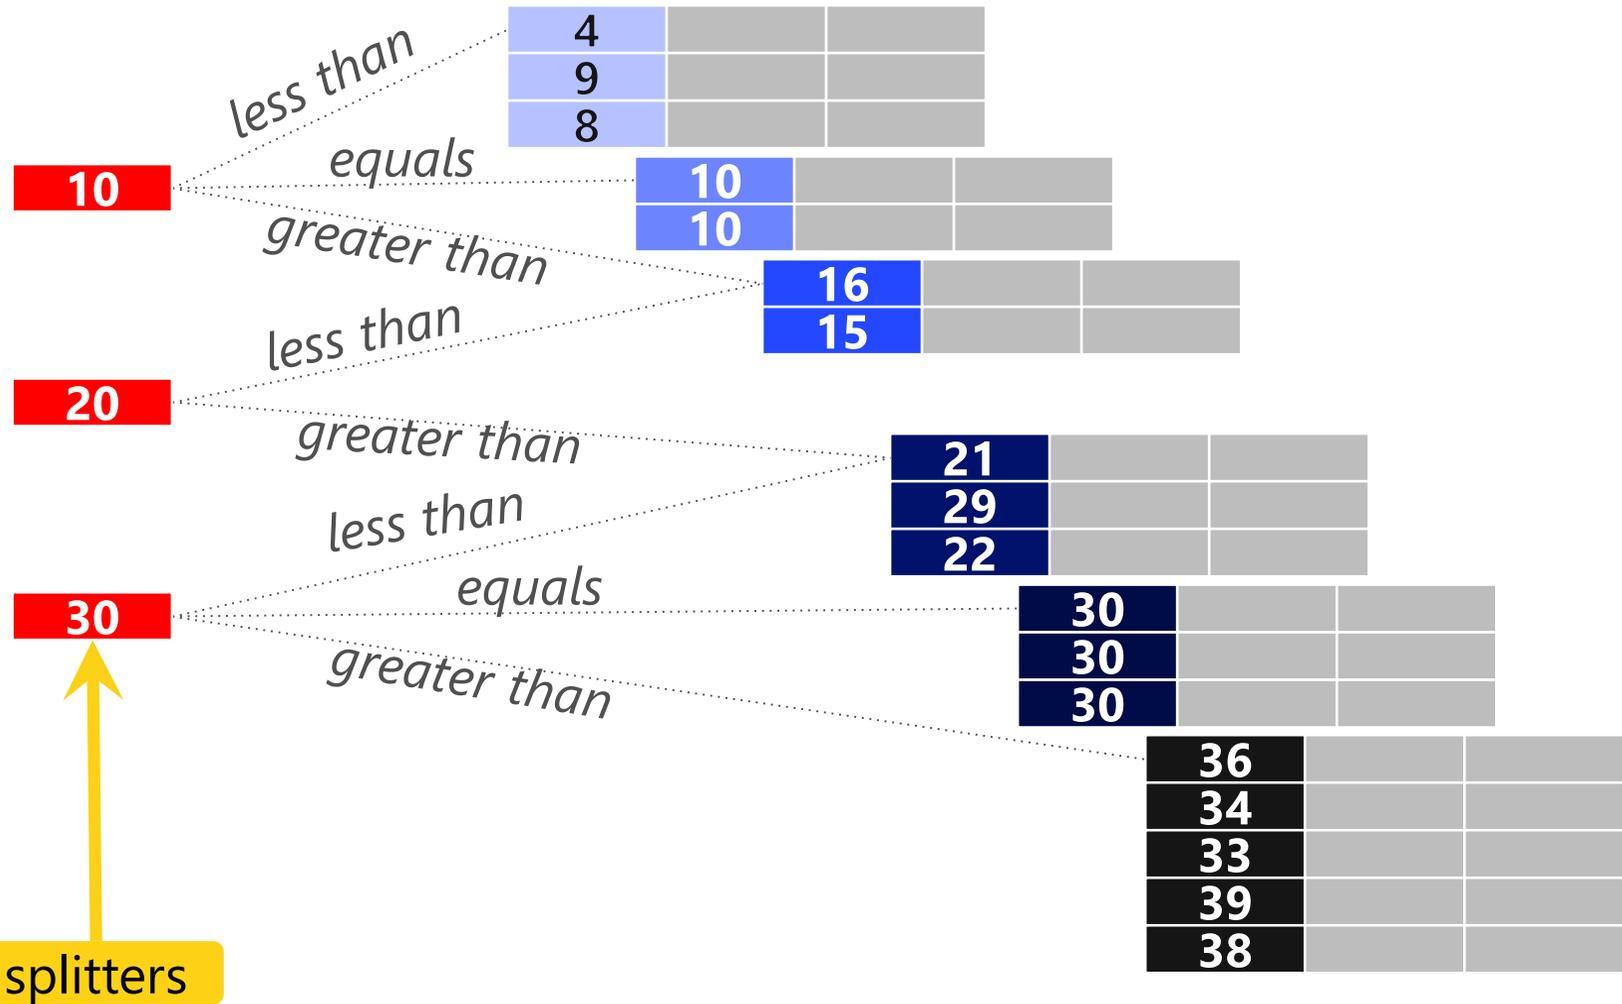
Example Tile: BoolGen



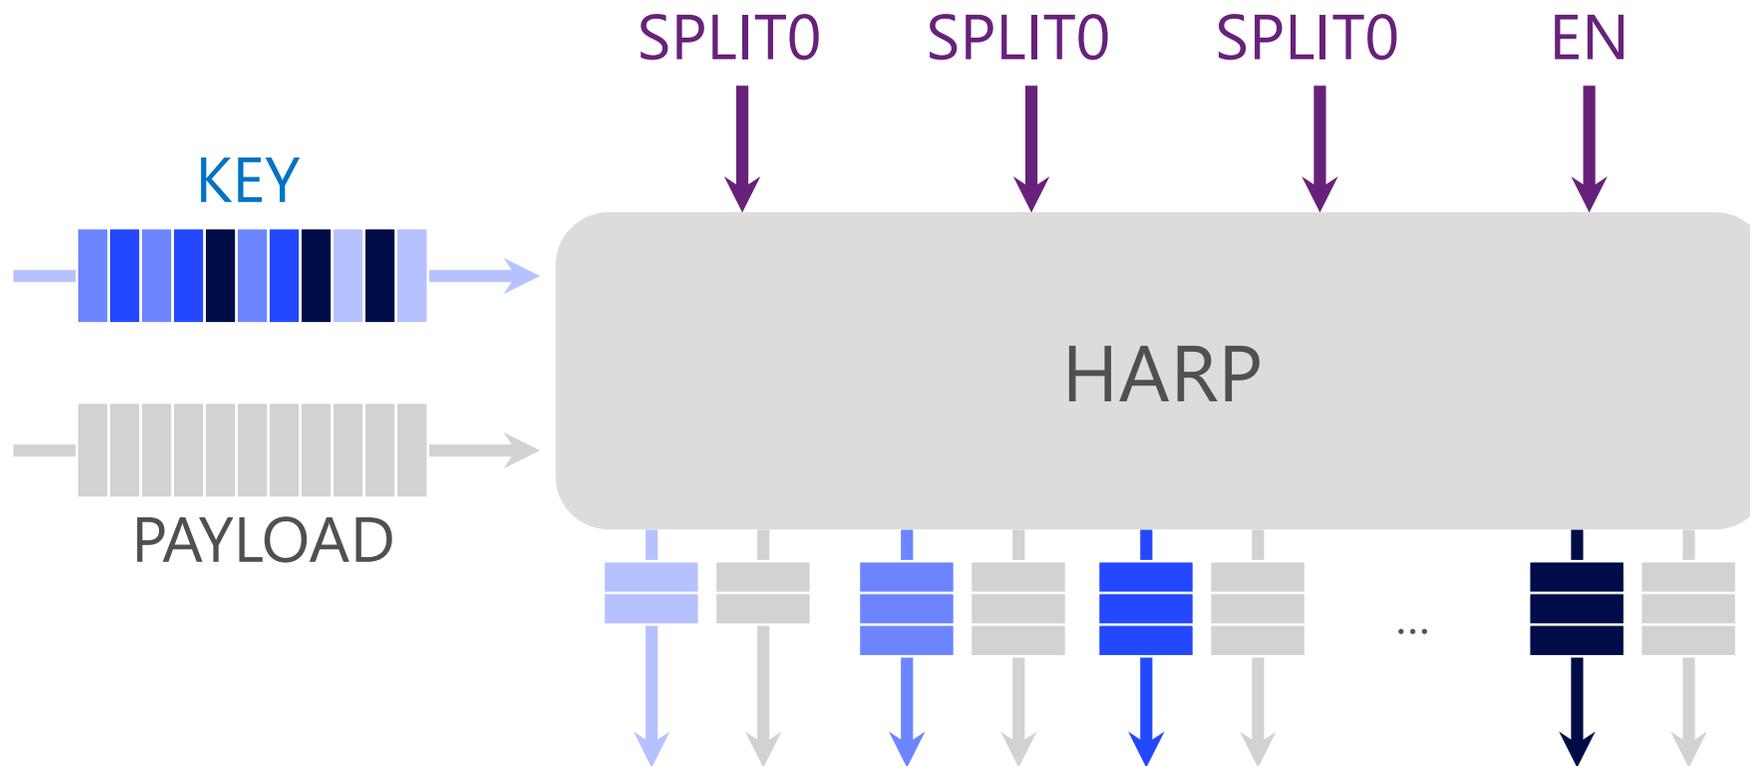
Range Partitioning Operation

36		
4		
21		
16		
30		
10		
34		
29		
30		
9		
33		
15		
10		
39		
22		
30		
8		
38		

↑
keys



Hardware Accelerated Range Partitioner (HARP)



Tile Characterization



AGG

ALU

BOOLGEN COLFILTER

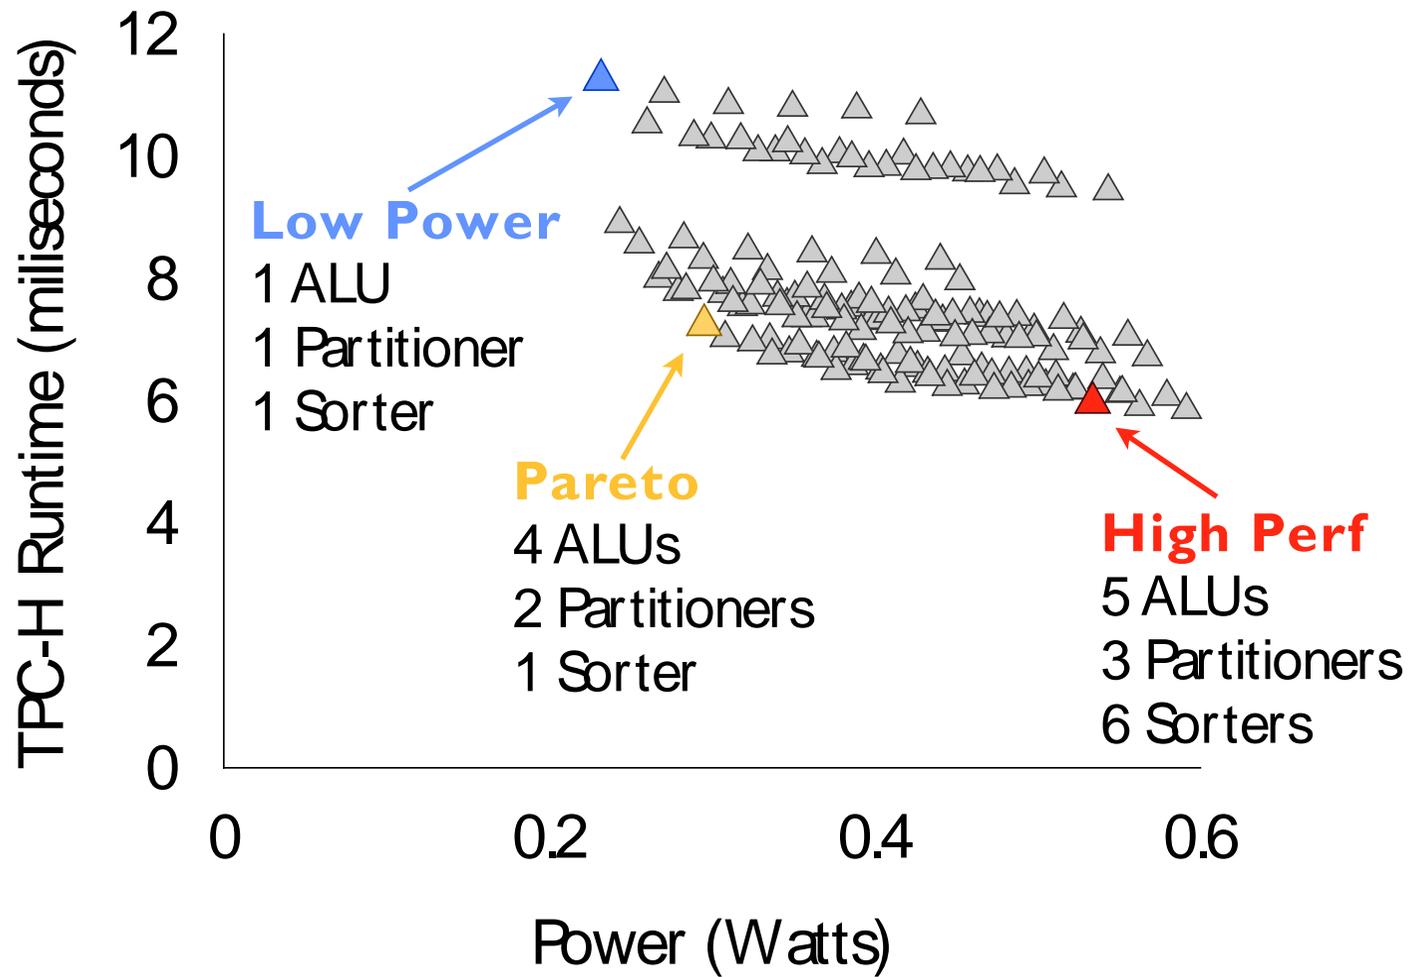
JOIN

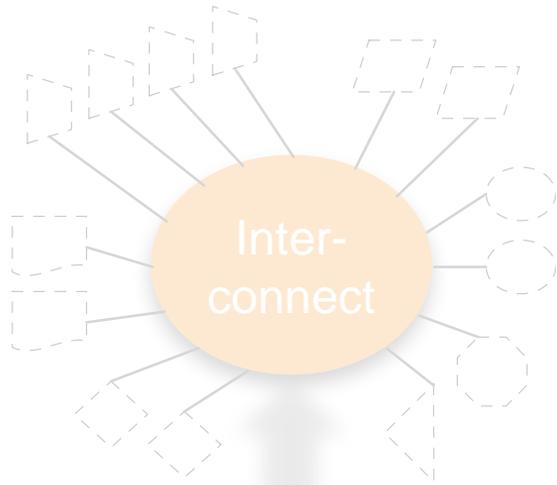
PART

SORT

APPEND COLSELECT CONCAT

STITCH

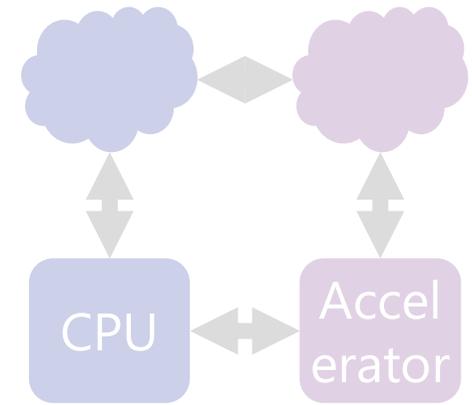
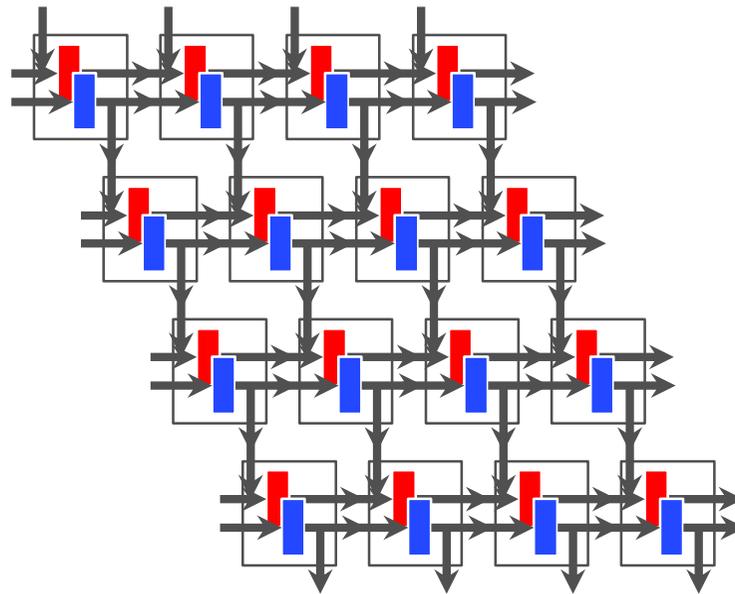




Q100

Data Processing Unit

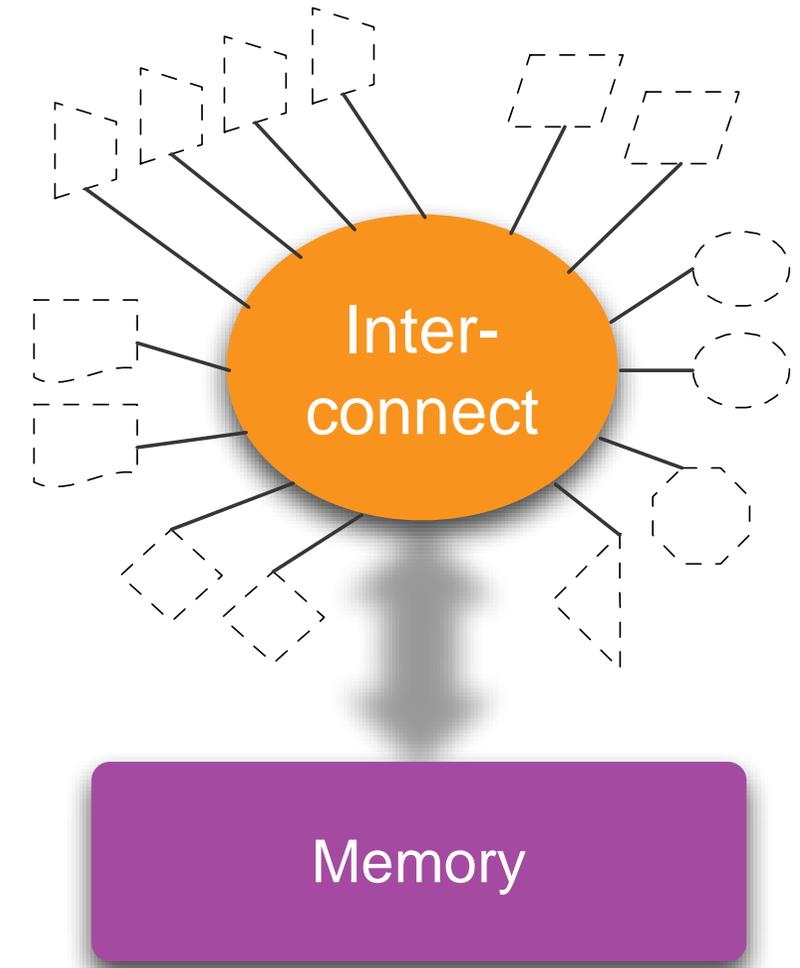
Data Processing Unit 2.0



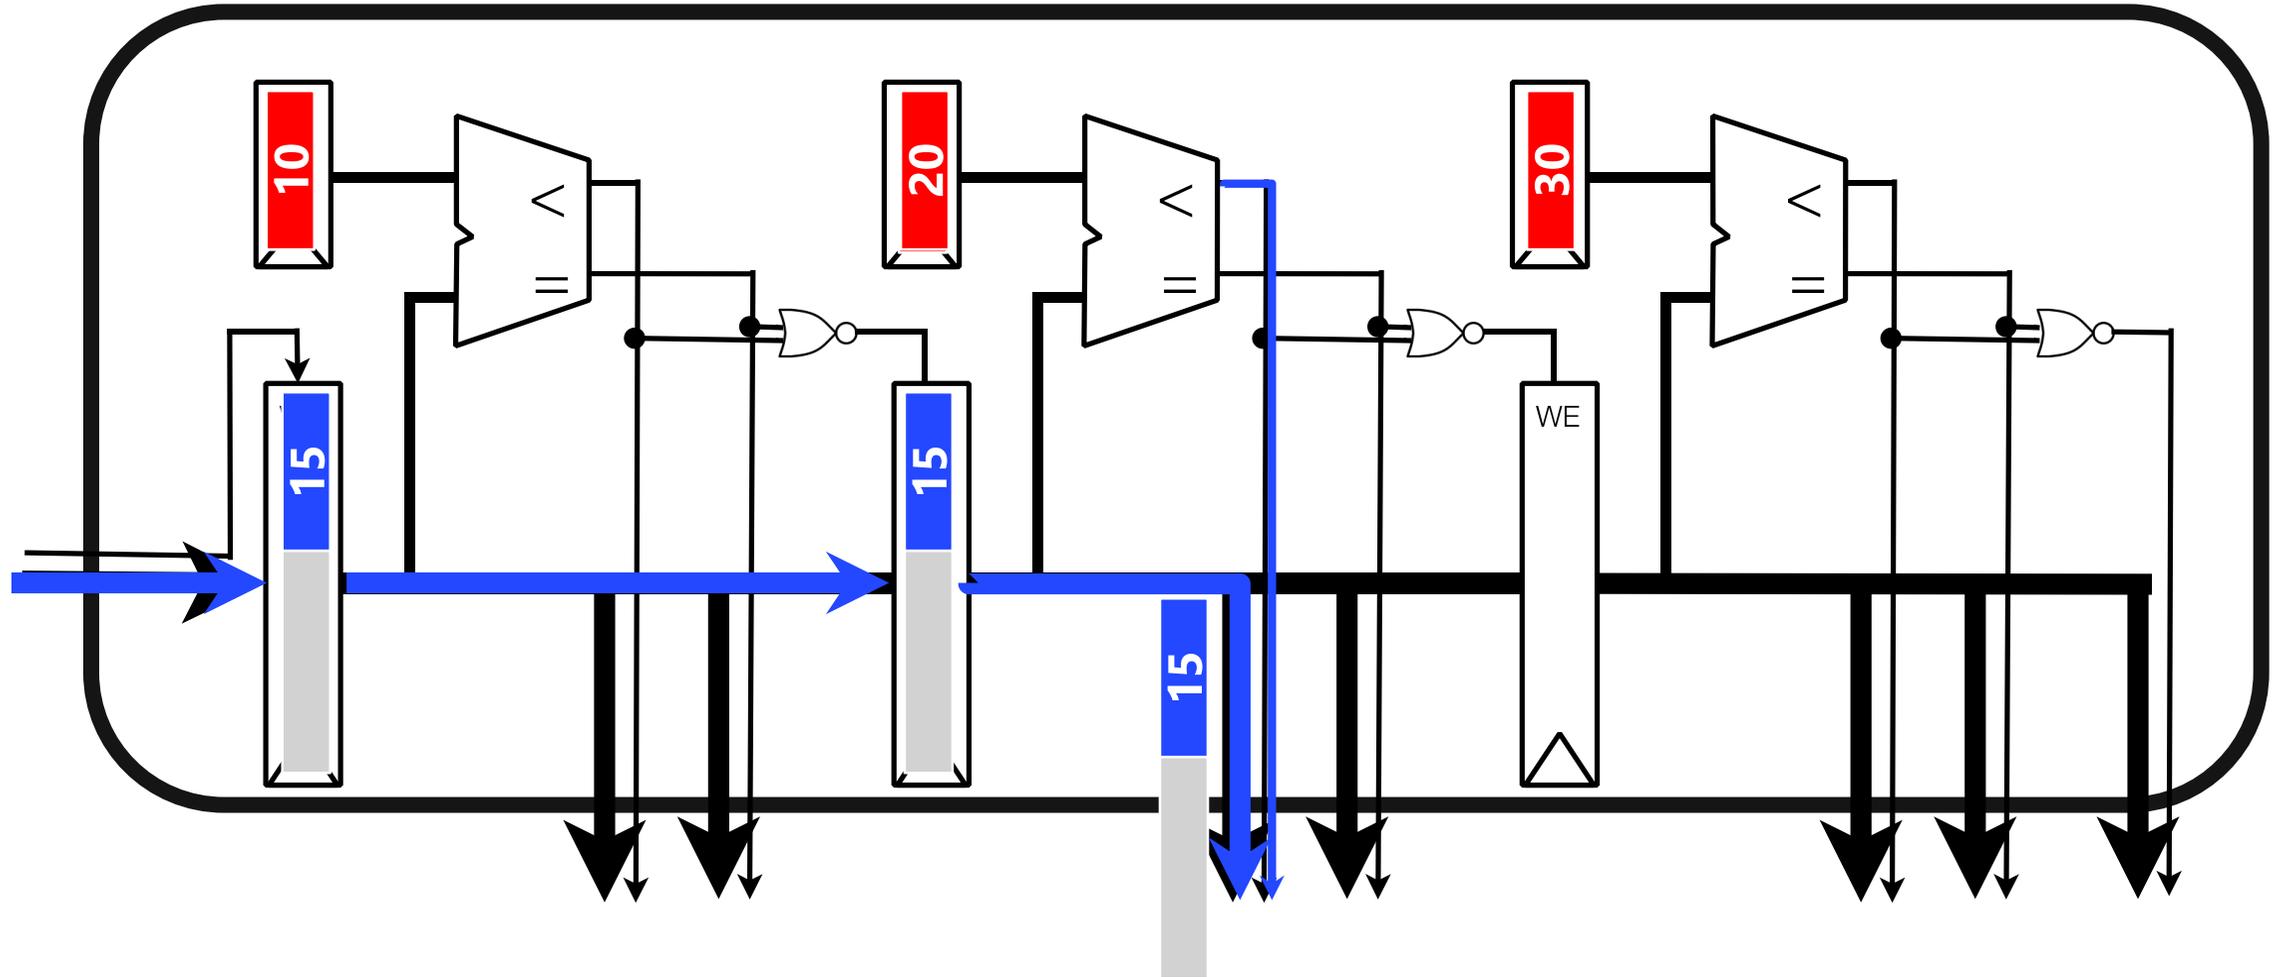
System Integration
Challenges

Is a heterogeneous design best?

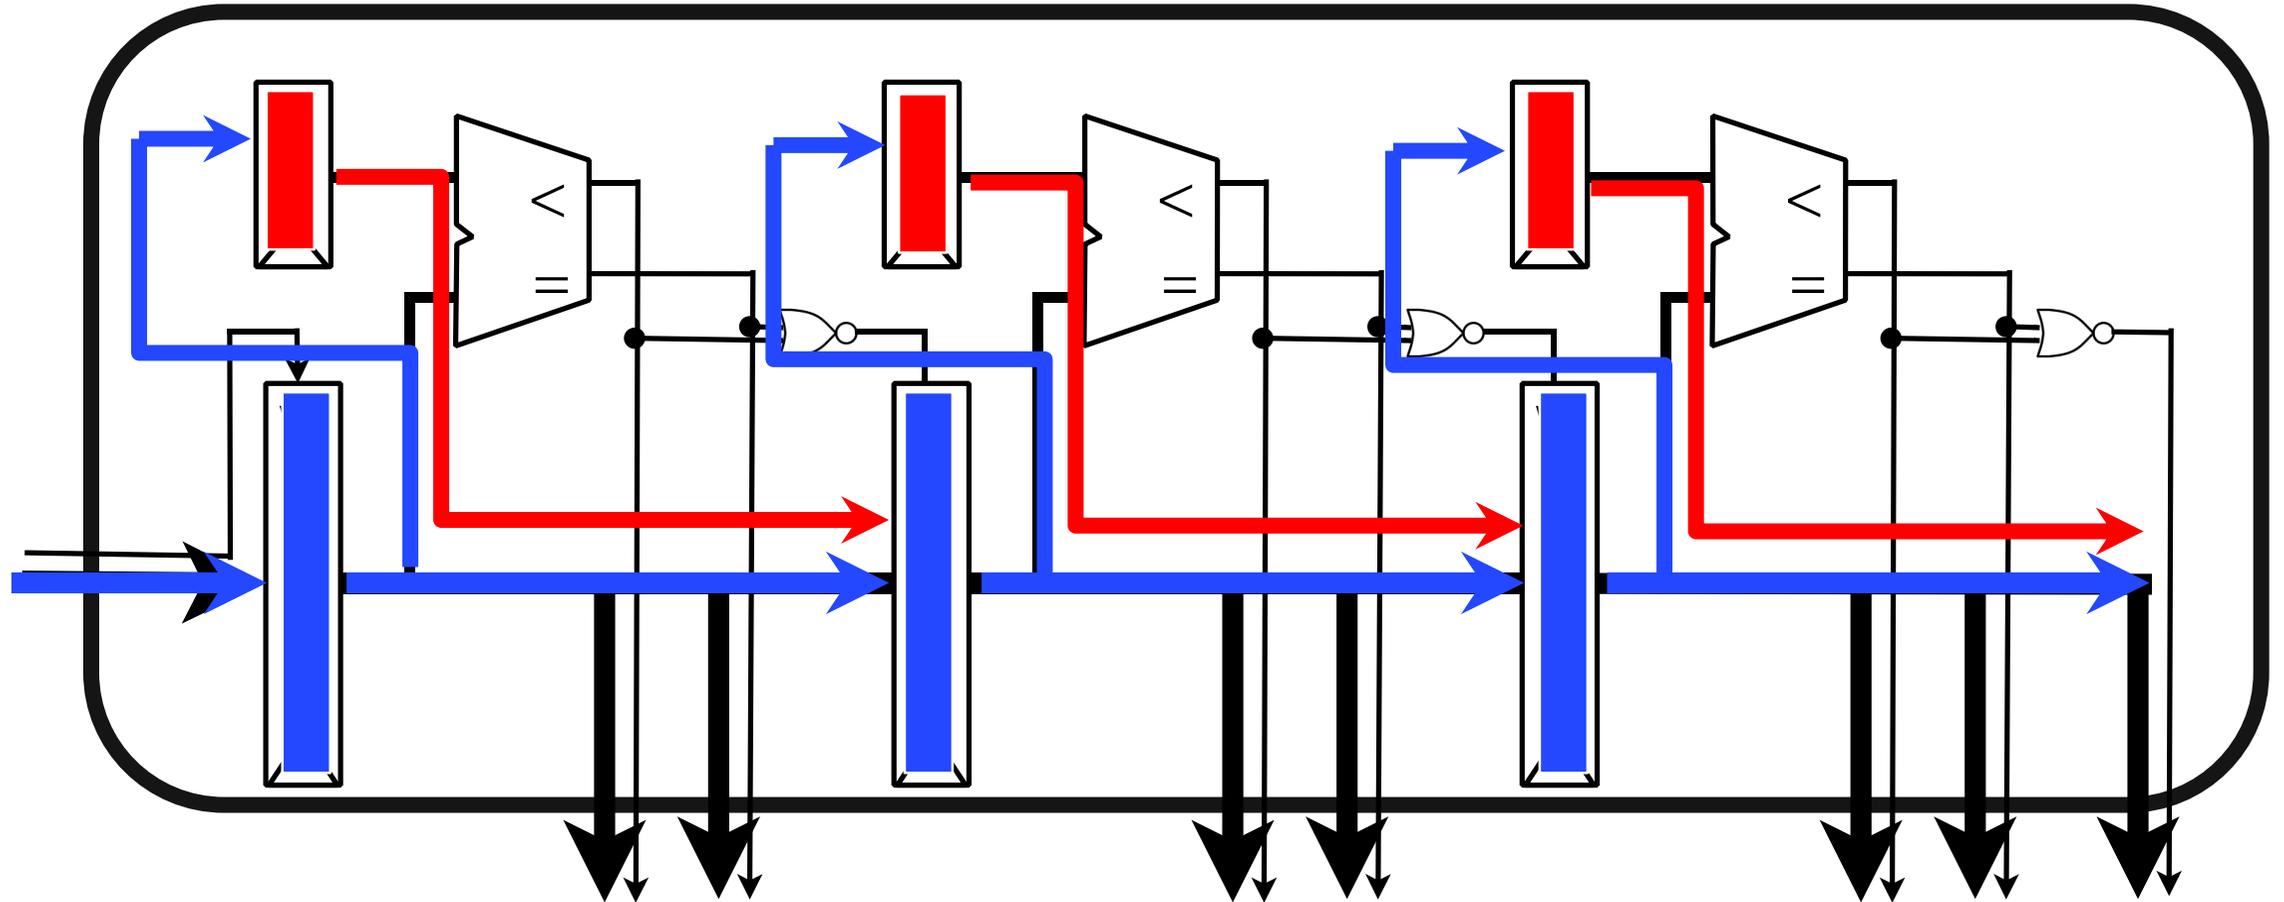
- Tiles are unitaskers
- Yet, they have a lot in common
- Irregular connectivity, tricky NoC
- Ad-hoc support for non-streaming operators



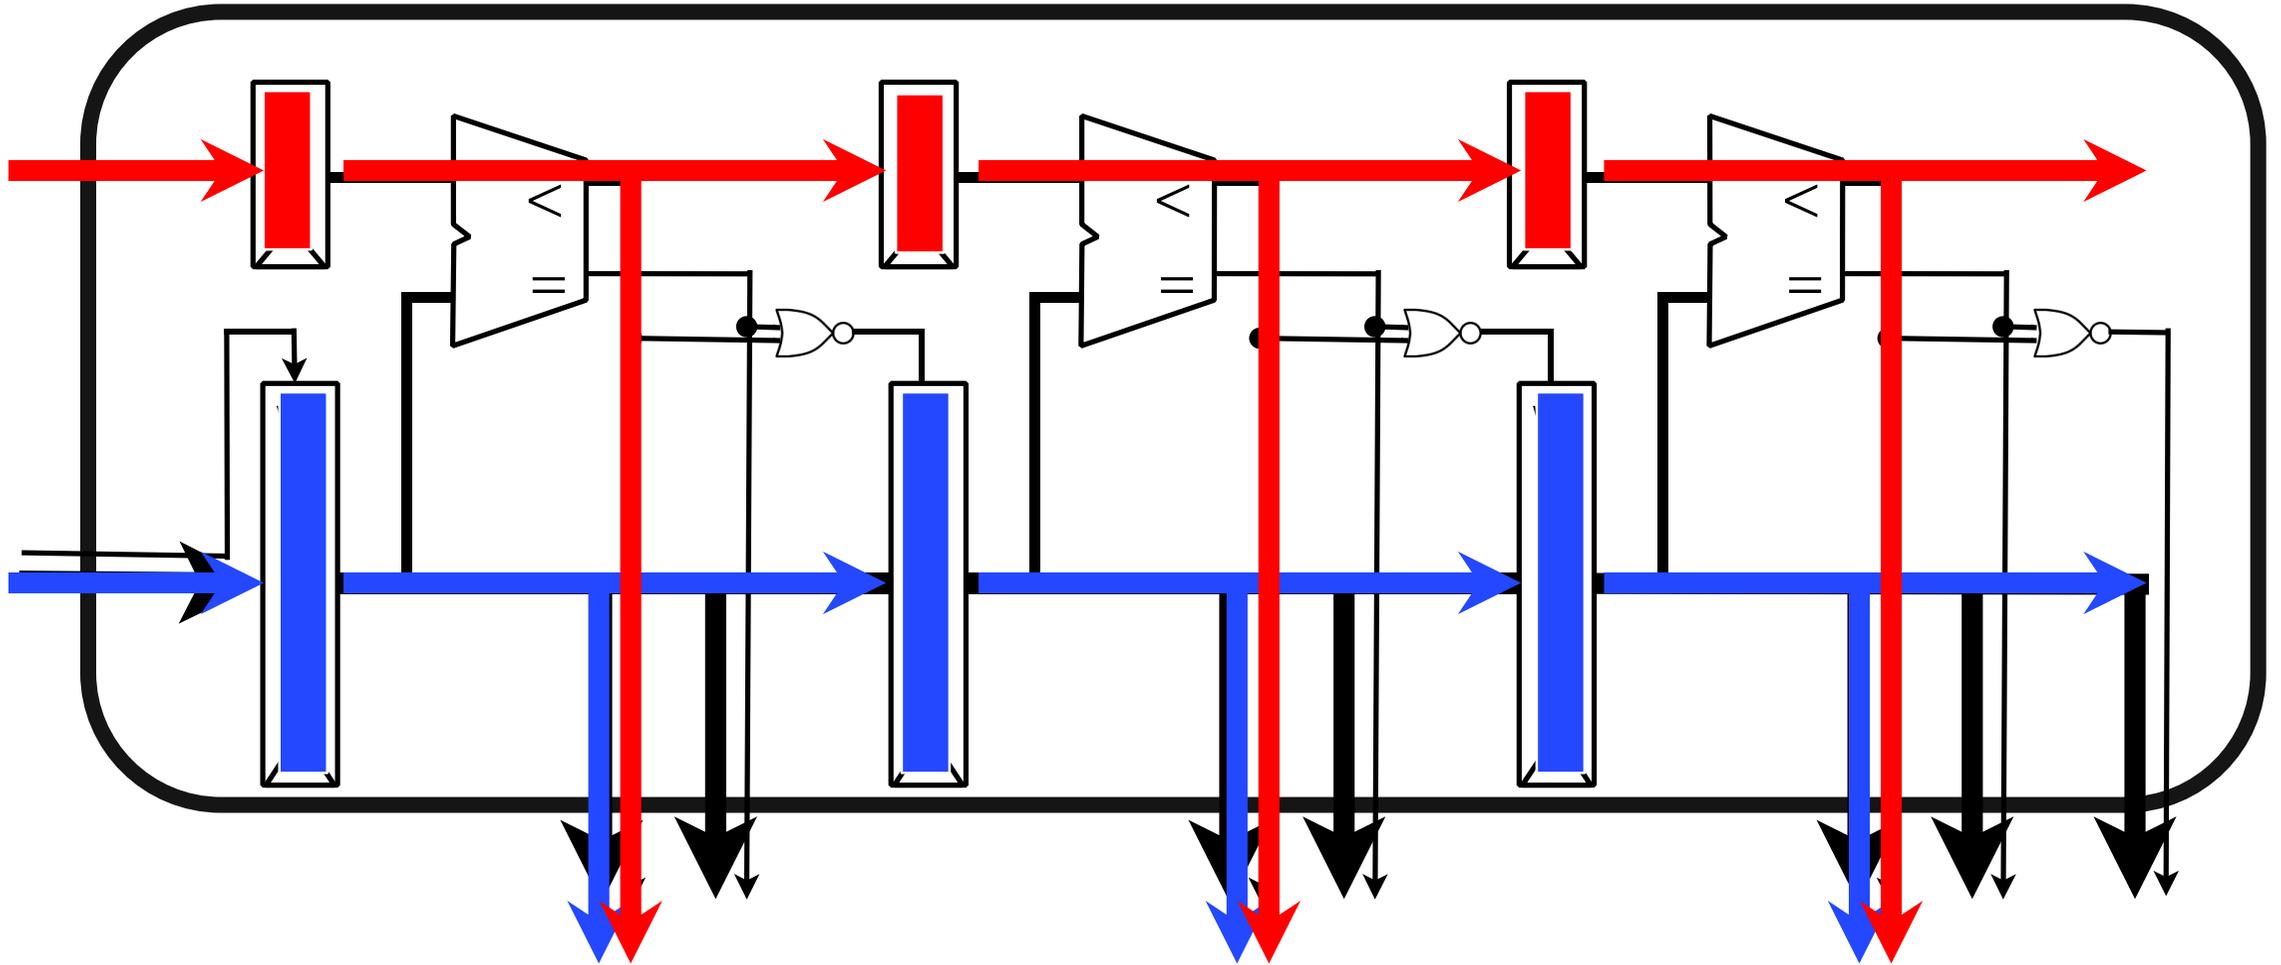
Recall the Partitioner



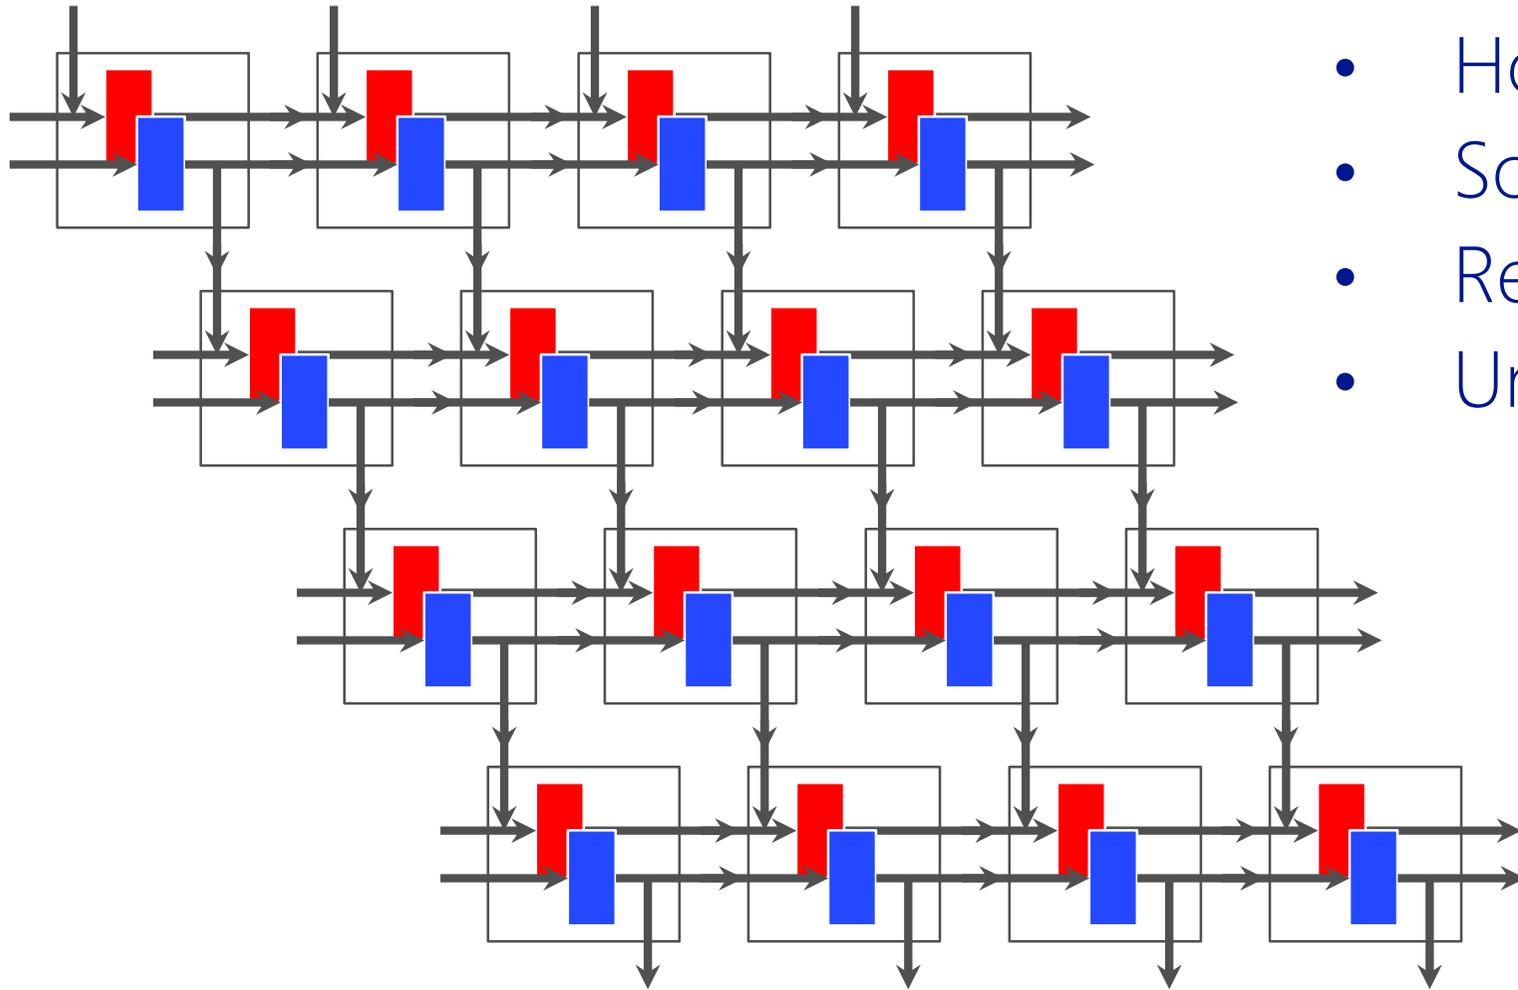
Just a Small Step to Merge...



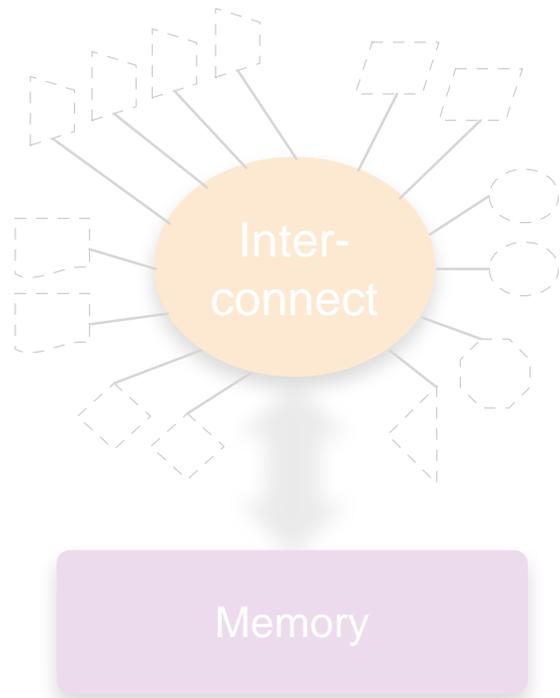
Or to Nested Block Loop Join.



Programmable, Homogeneous DPU



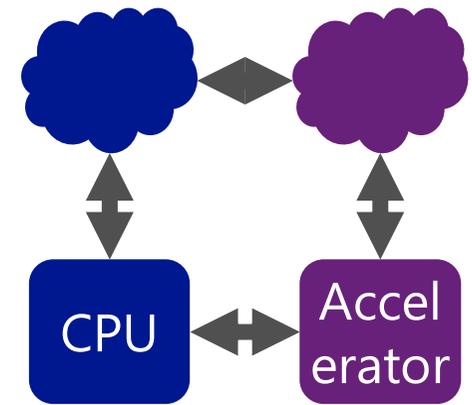
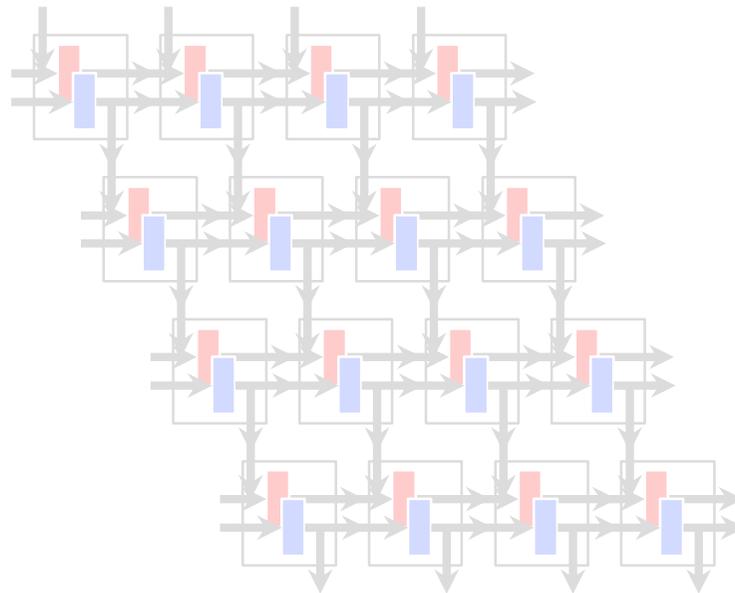
- Homogeneous resources
- So, better usage
- Regular, point-to-point NoC
- Unified buffering and routing



Q100

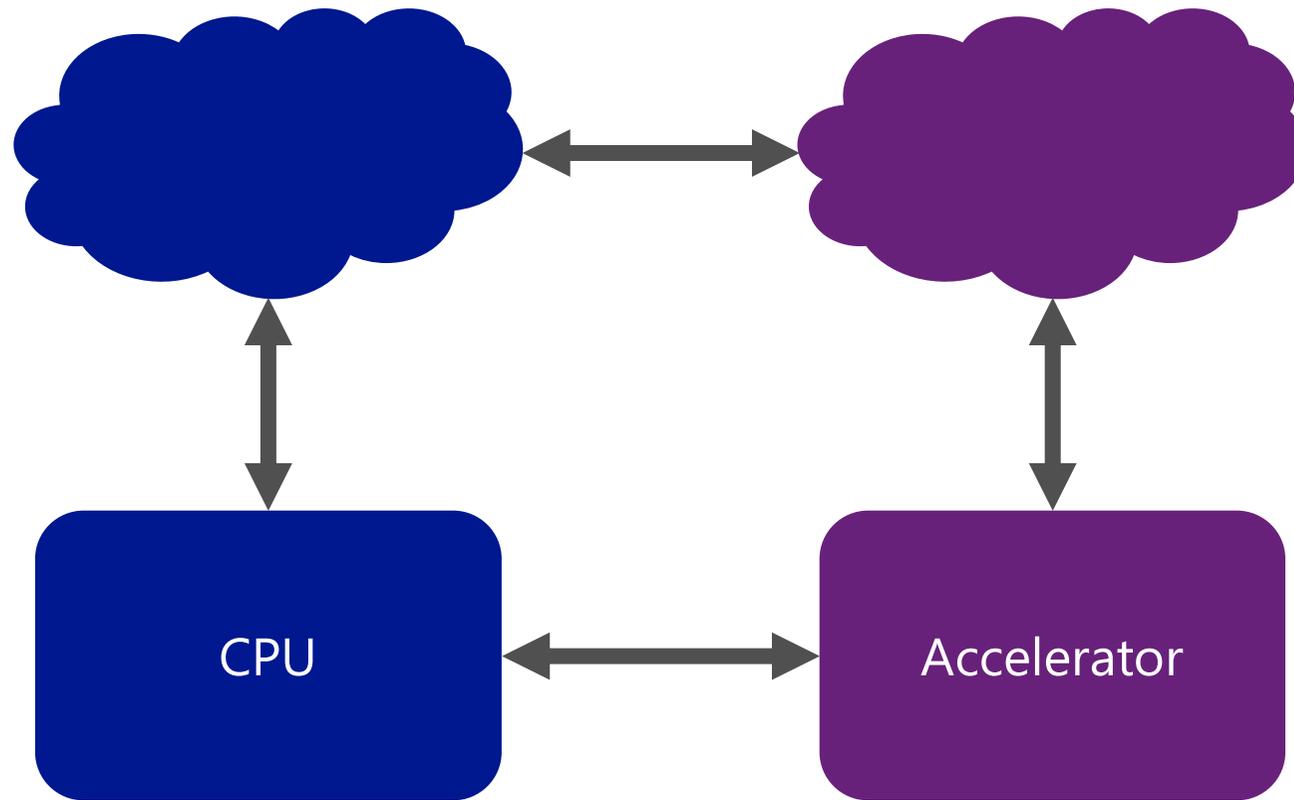
Data Processing Unit

Data Processing Unit
2.0



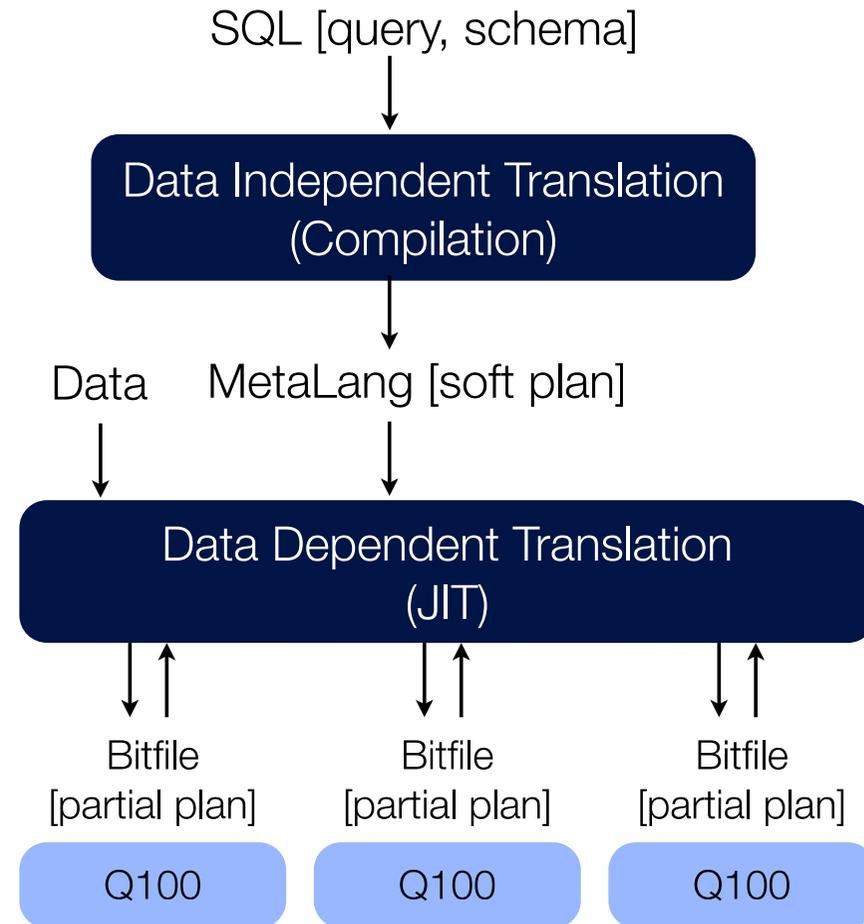
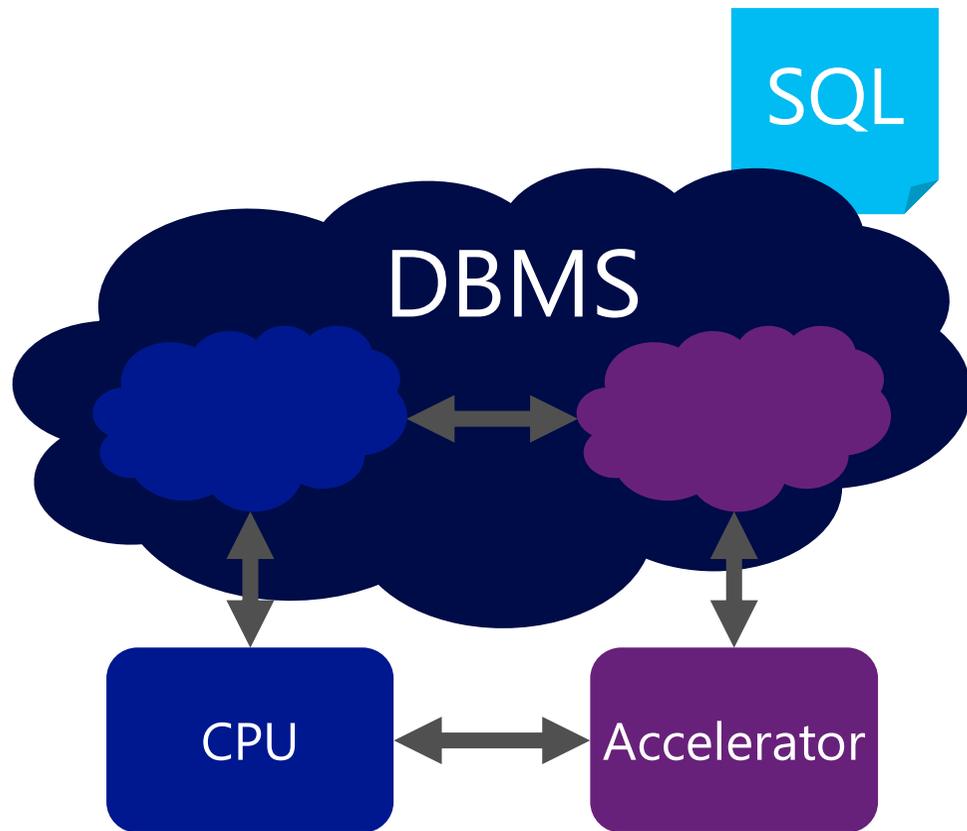
System Integration
Challenges

Hardware and Software Integration

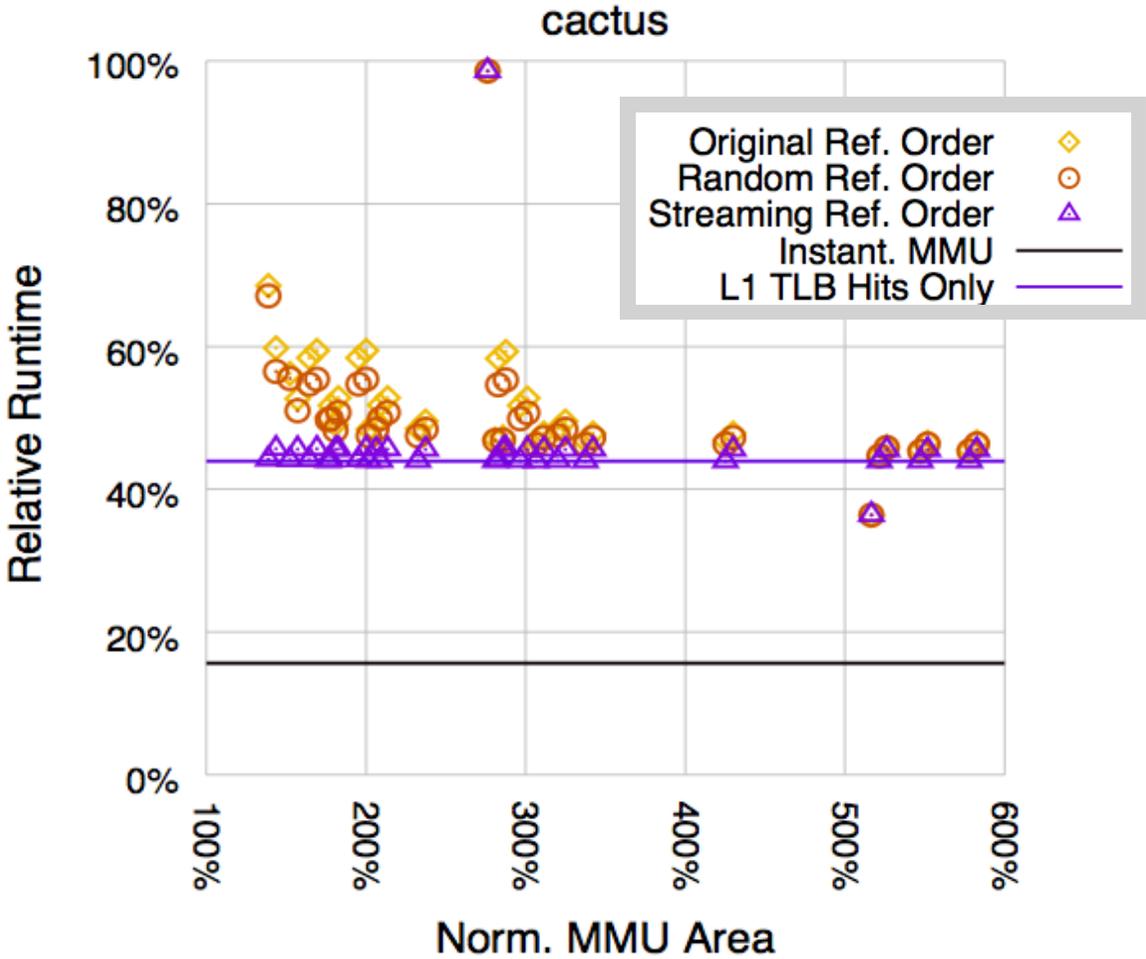
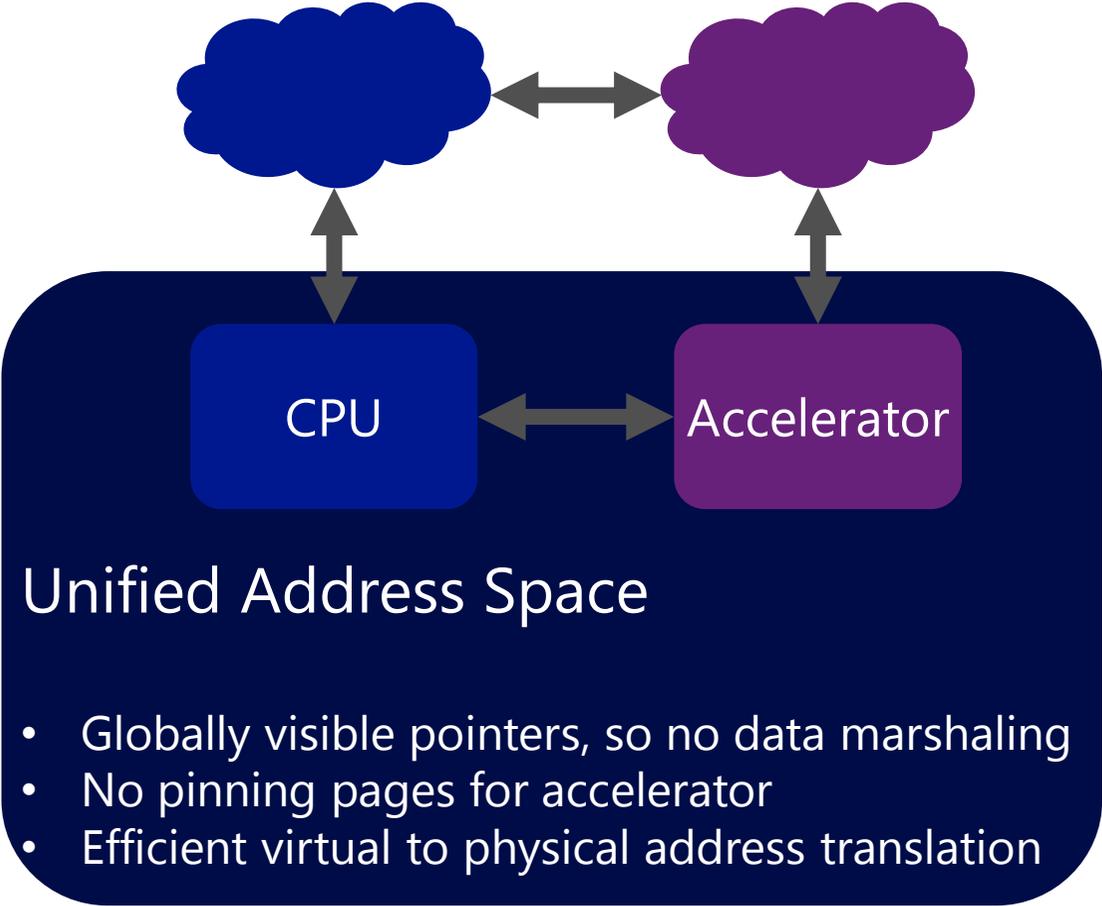


It is all one problem!

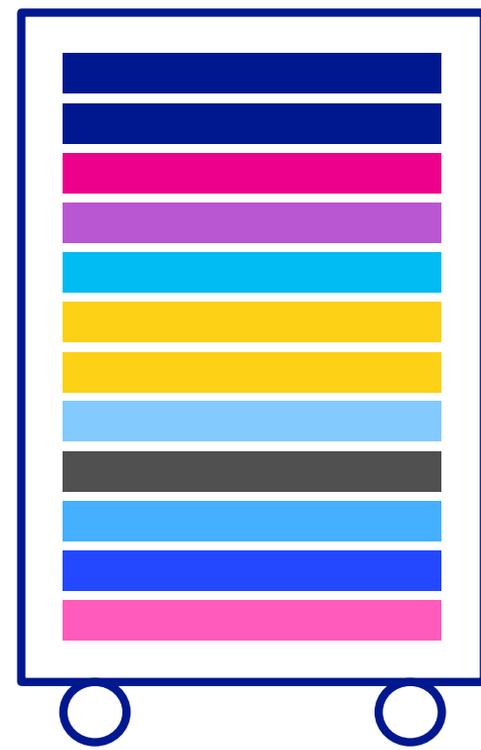
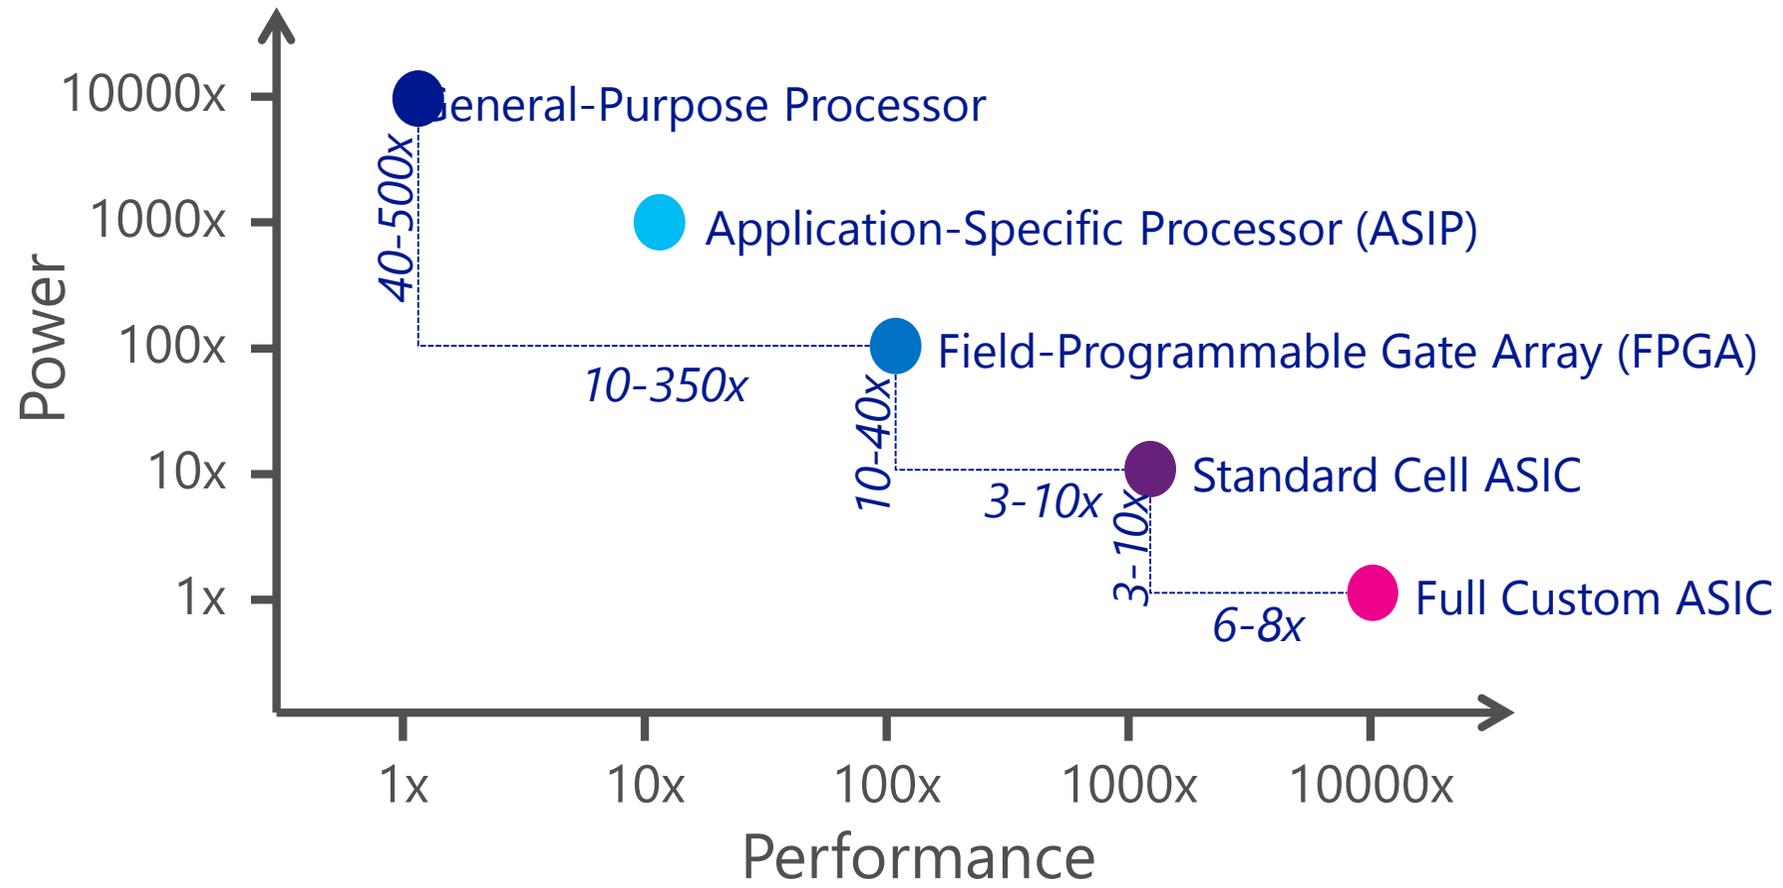
Exploit established software interfaces



Develop software-friendly microarchitectures



How To ~~Why~~ Specialize?



Thank You

<http://arcade.cs.columbia.edu>



Save the planet and return
your name badge before you
leave (on Tuesday)

