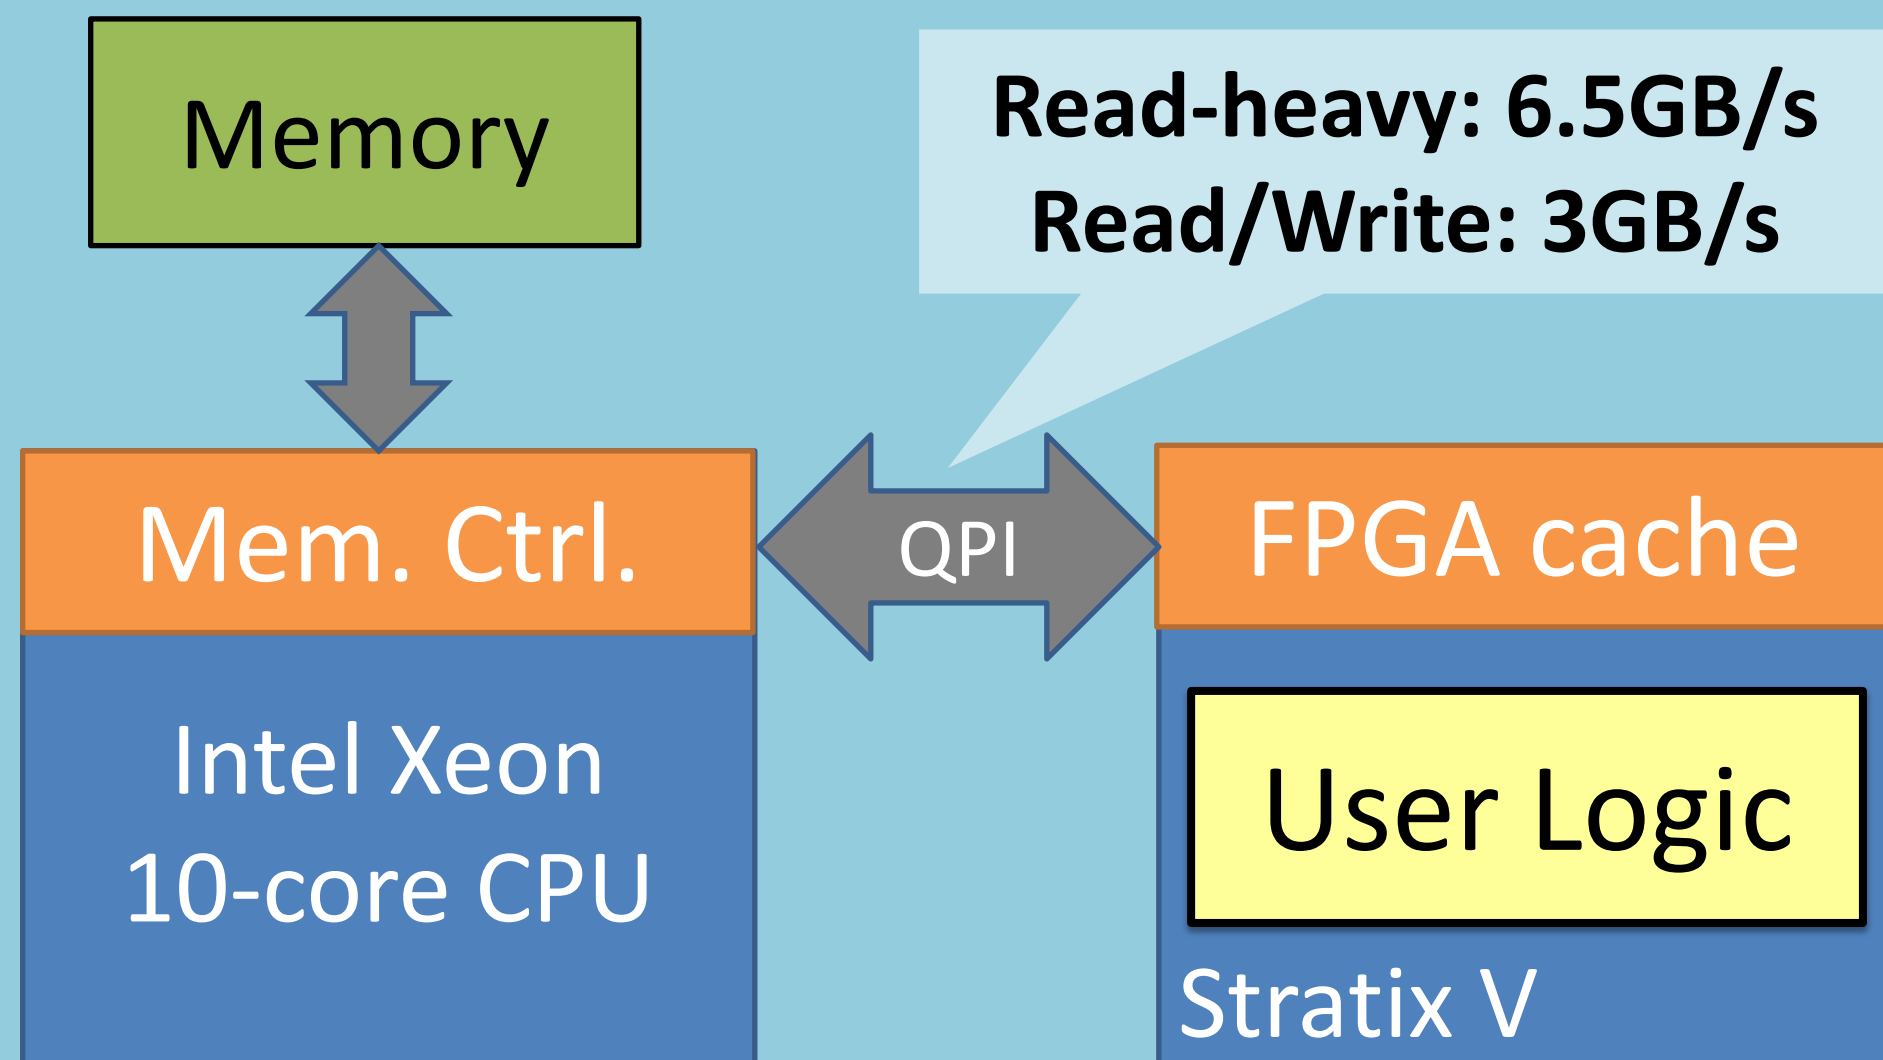


New emerging hybrid architectures

Intel Xeon+FPGA (v1)



Read-heavy: 6.5GB/s
Read/Write: 3GB/s

Hybrid CPU-FPGA Architectures:

- FPGA has cache-coherent access to main memory
- Eliminate the issue of data movement/partitioning
- FPGA is a **specialized** core

Your next CPU might come with an FPGA!

Intel Xeon+FPGA (v2)



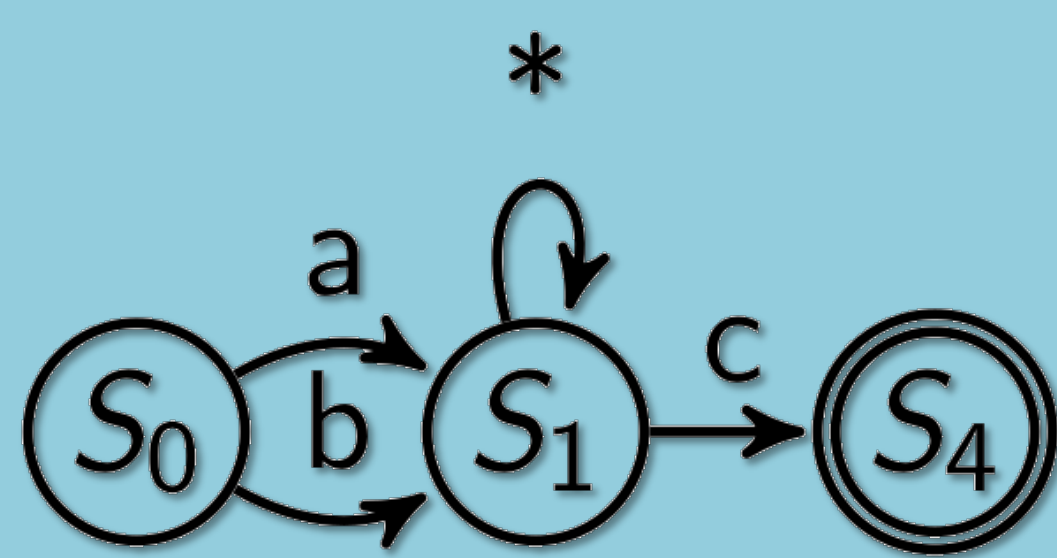
CPU and FPGA in same package

How to use your future FPGA for pattern matching

Regular Expression in Hardware:

- Regular expressions can be mapped to NFAs
- NFAs can be efficiently executed on FPGAs[1,2]

Regular Expression: $(a|b).^*c$



[1] R. Sidhu, V. Prasanna, *Fast regular expression matching using FPGAs*, FCCM'01
[2] L. Woods, J. Teubner, *Complex event detection at wire speed with FPGAs*, VLDB'10

Characters configuring Character Encoder

Characters	C1	C2	C3	C4
'a'	1	0	0	0
'b'	0	1	0	0
'c'	0	0	1	0

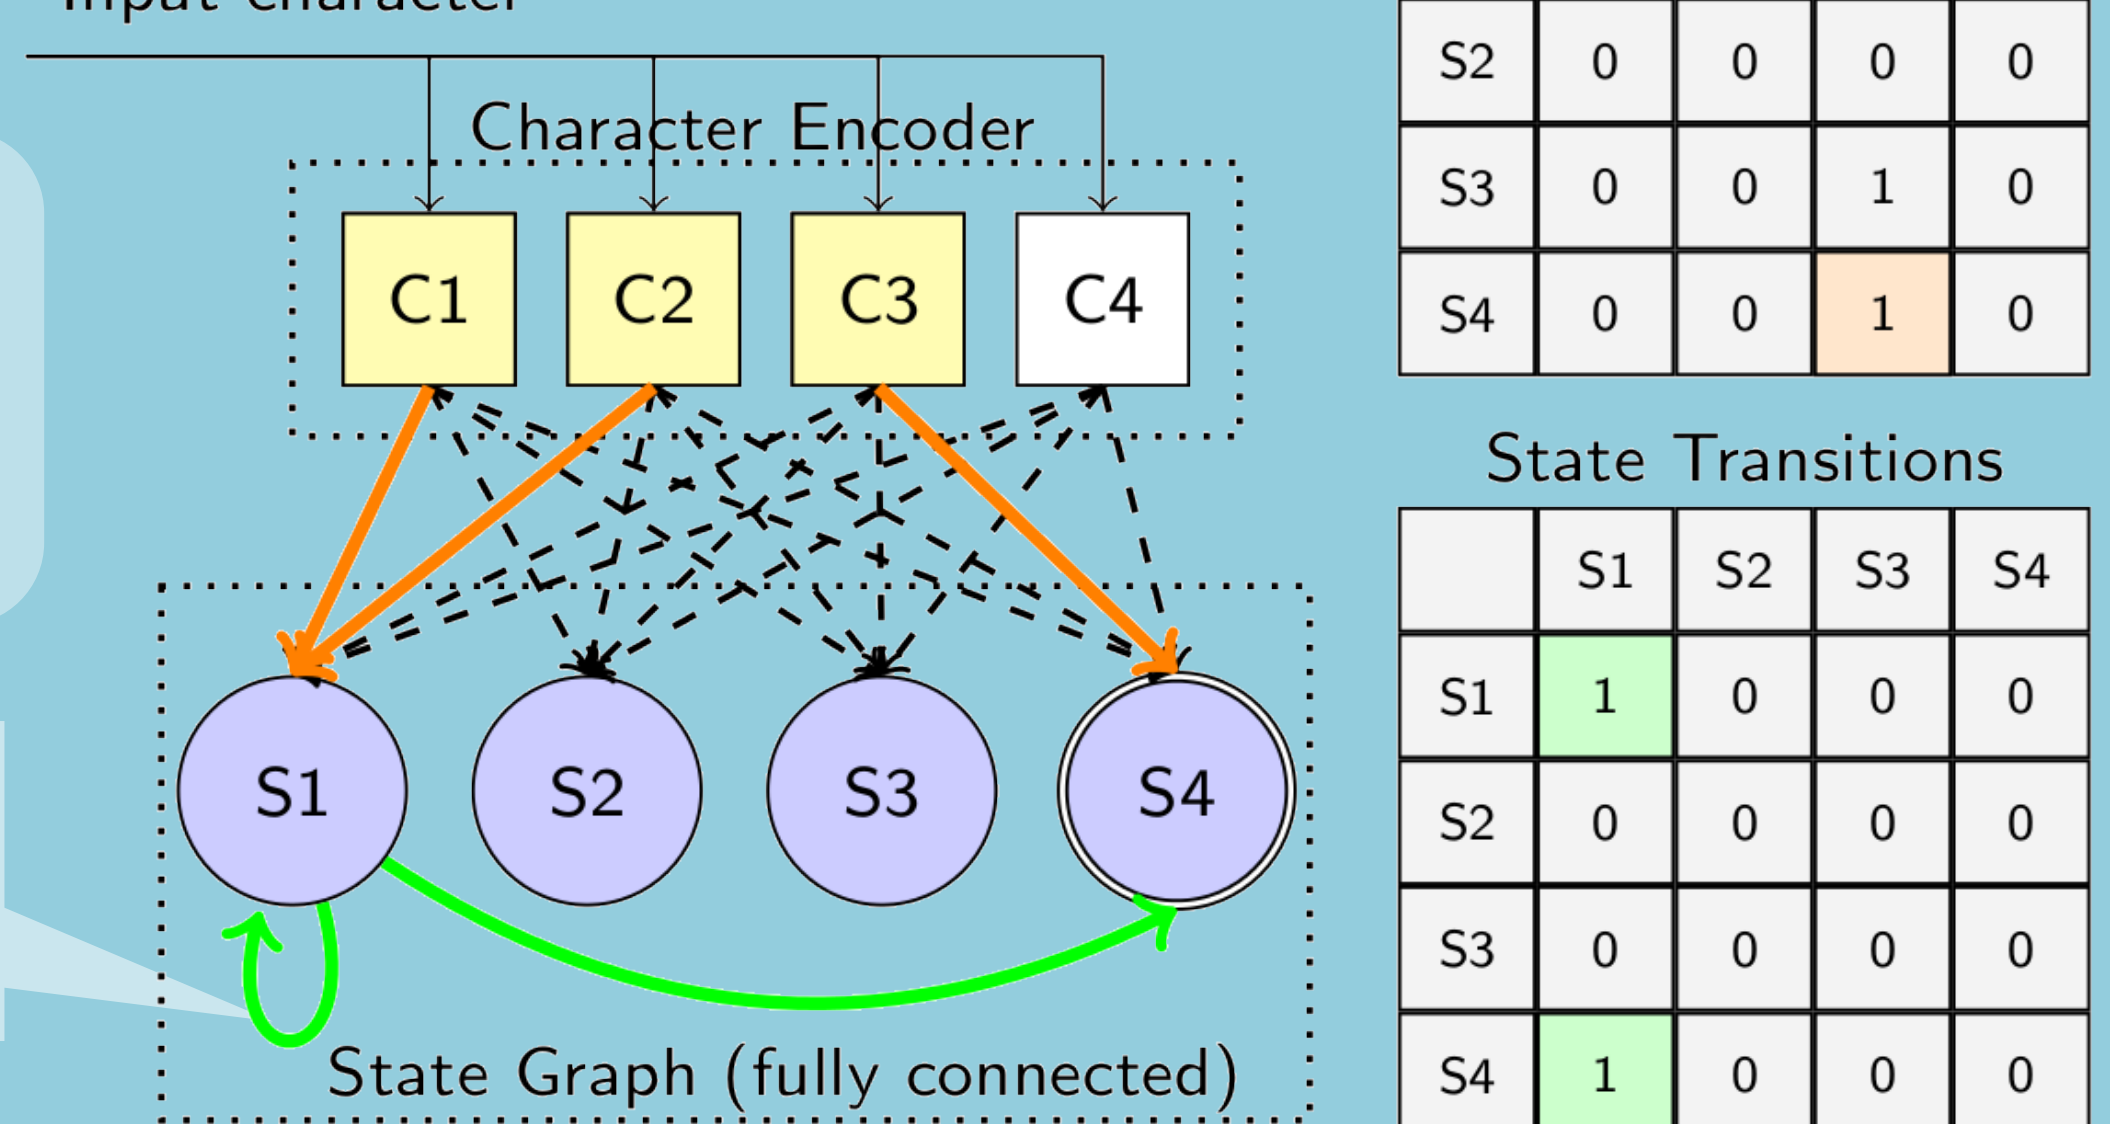
Characters triggering states

Triggers	C1	C2	C3	C4
S1	1	1	0	0
S2	0	0	0	0
S3	0	0	1	0
S4	0	0	1	0

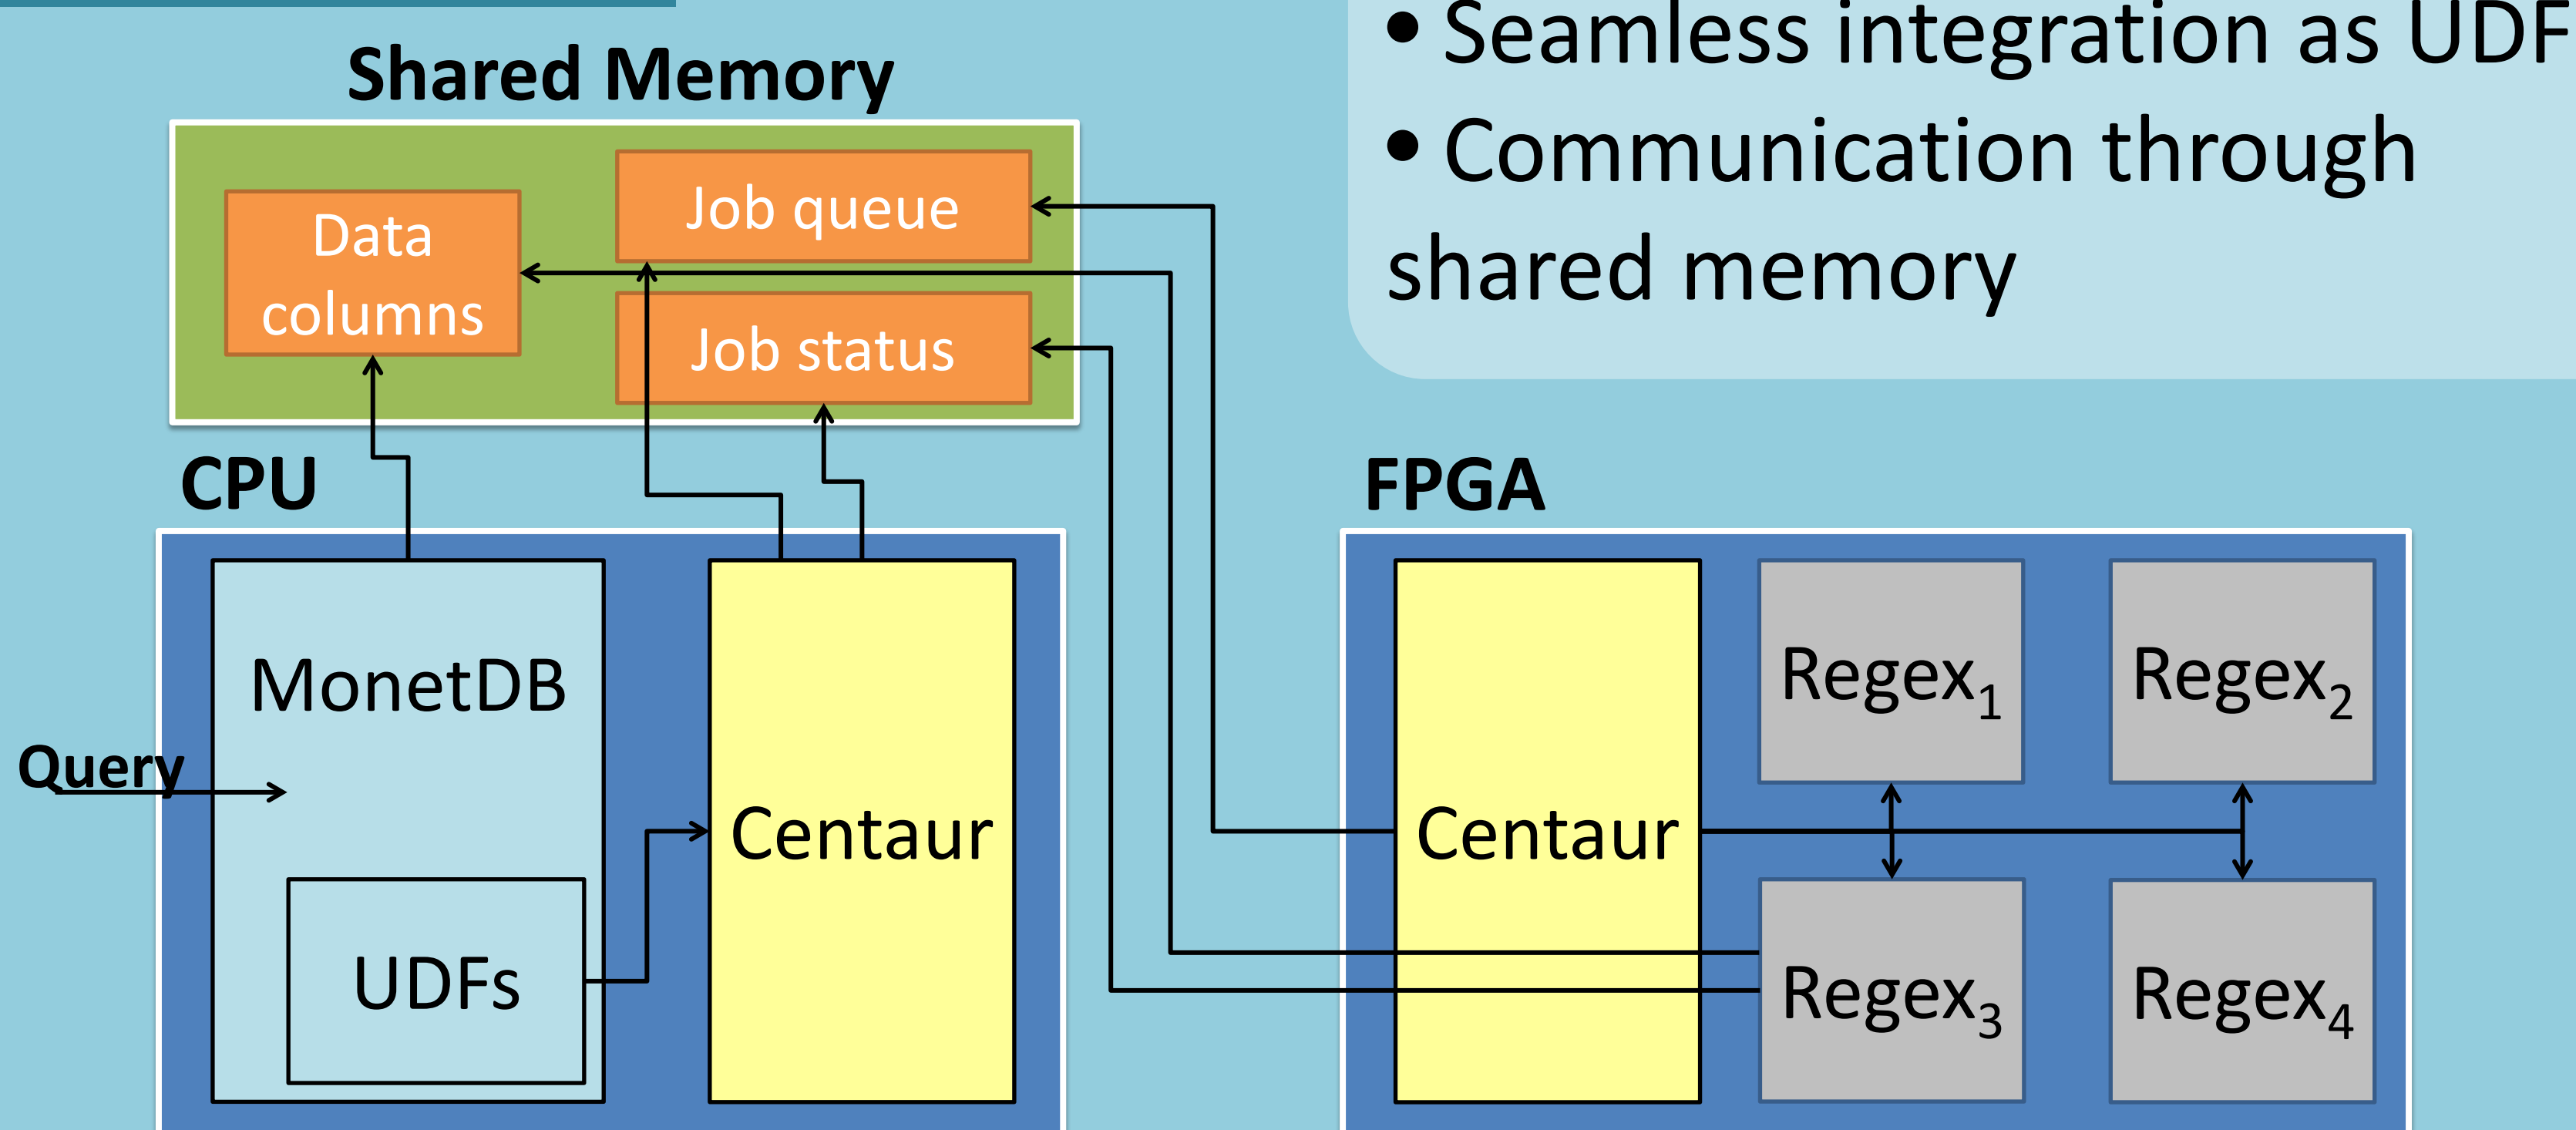
Character Encoder supports:

- Sequences
- Case-insensitivity
- Ranges
- Collations

Transitions parameterized by State Transitions matrix



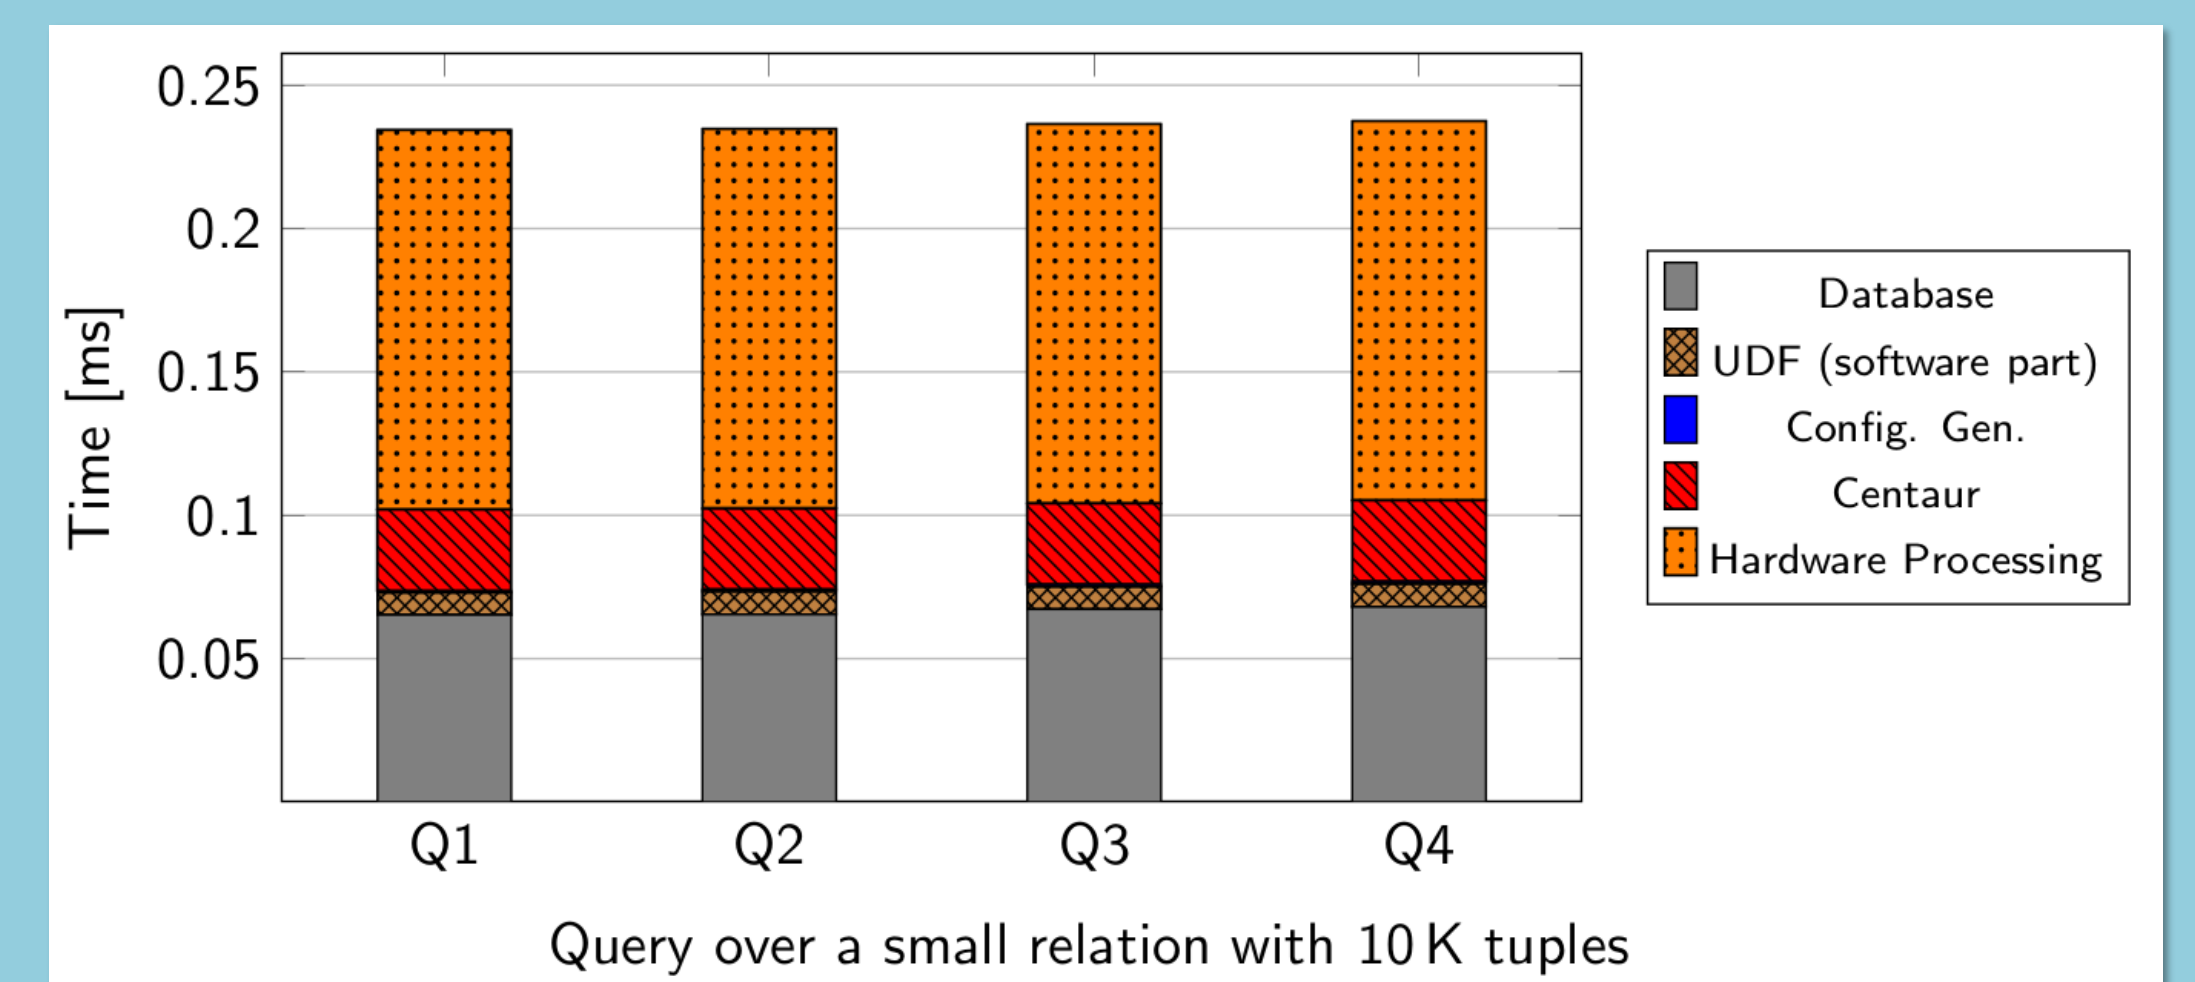
Integration



- Seamless integration as UDF
- Communication through shared memory

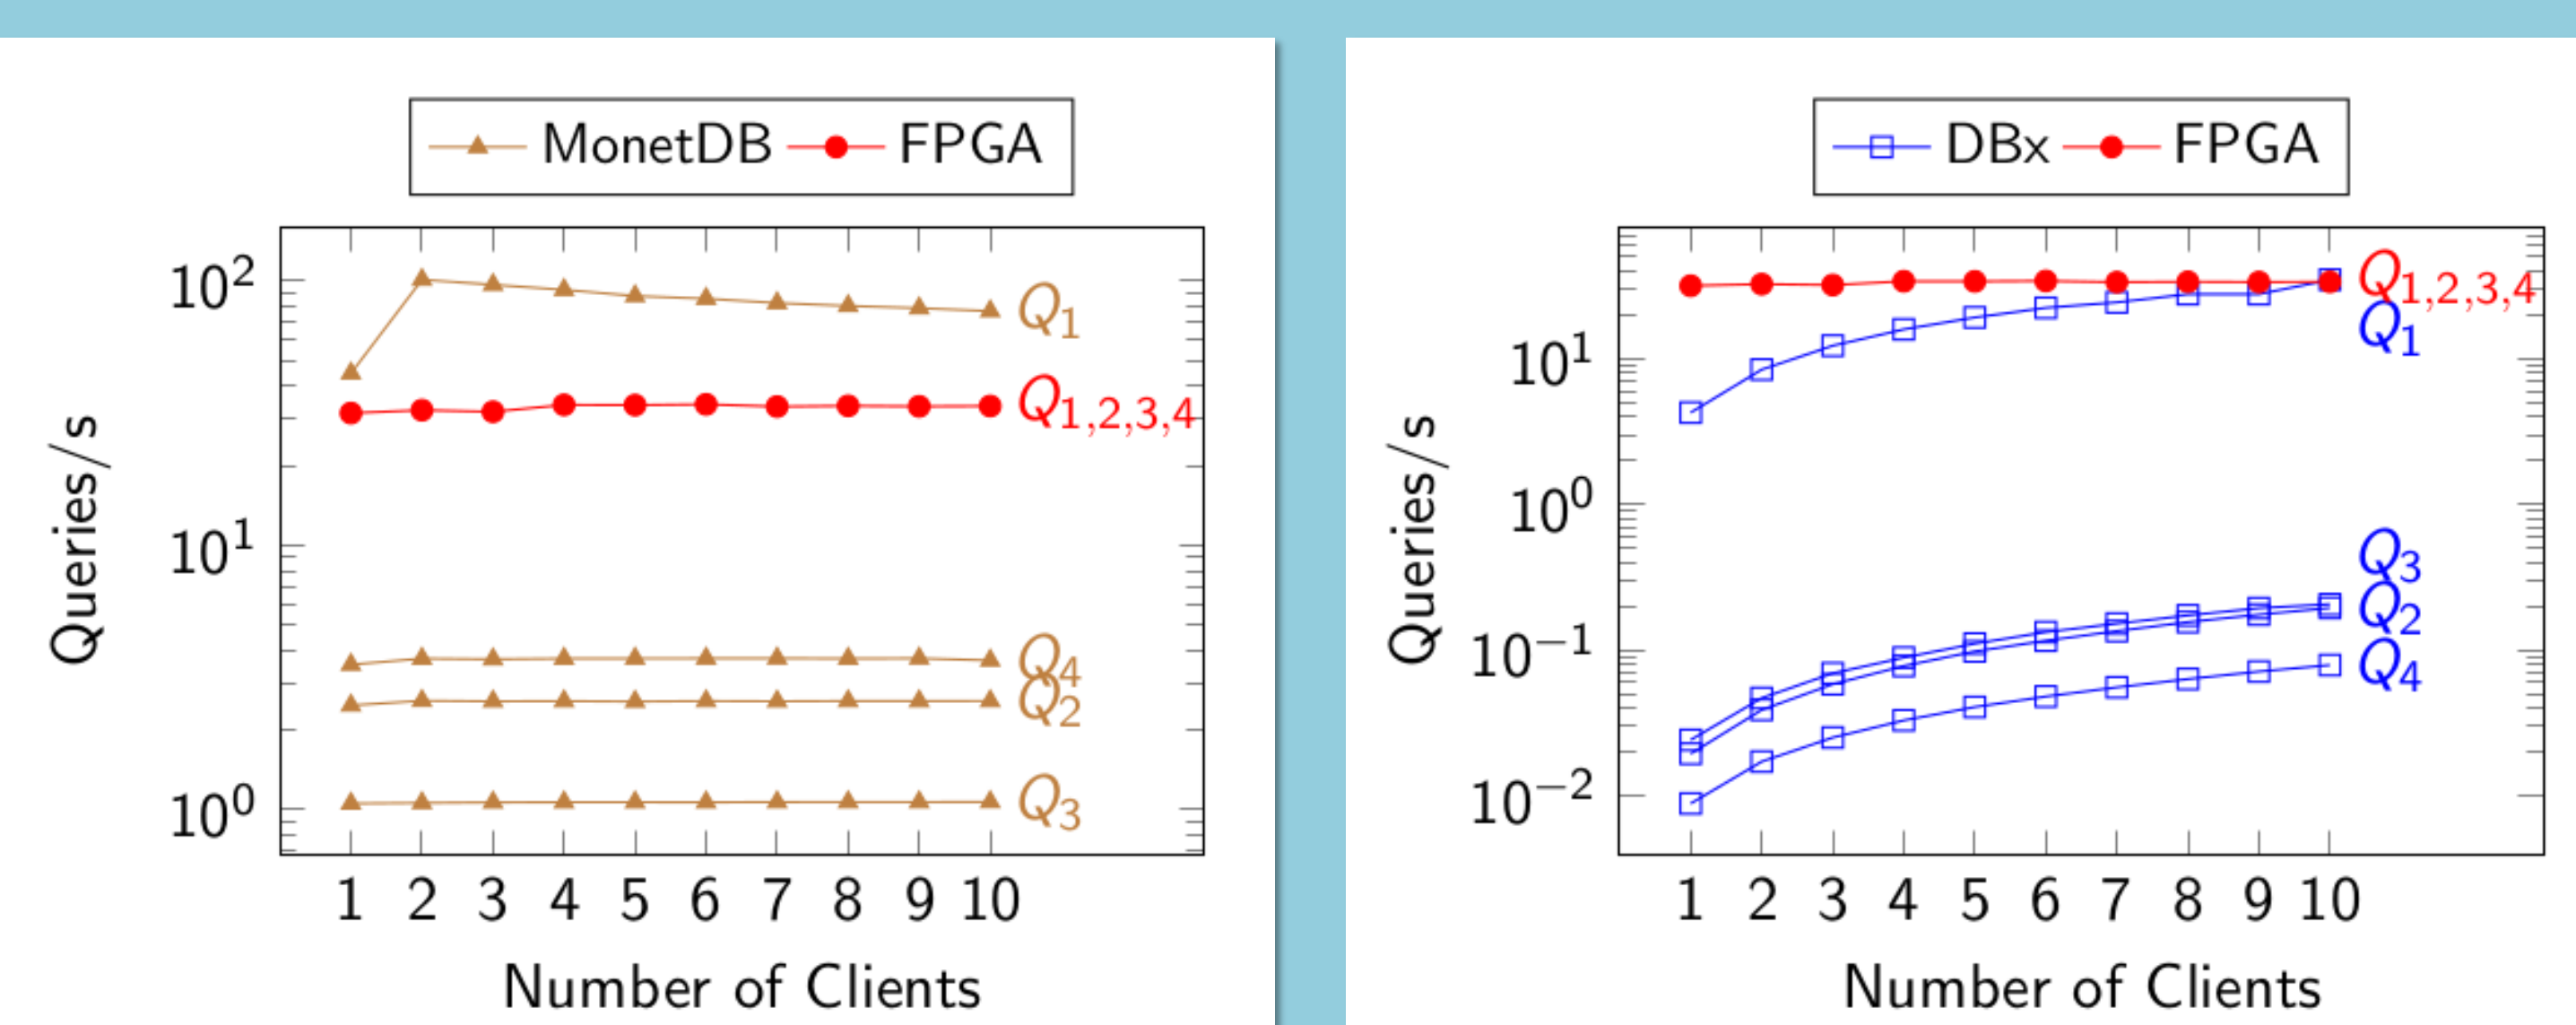
[3] M. Owaida, D. Sidler, et al., *Centaur: A Framework for Hybrid CPU-FPGA Databases*, FCCM'17

Overhead



Minimal overhead which decreases proportionally for larger data sets

Throughput



On a par with a 10-core CPU!

Hybrid Execution

- Complex expressions might not fit into deployed NFA
- Divide regex evaluation between CPU and FPGA

