

NVIDIA RISC-V Story

4th RISC-V Workshop 7/2016



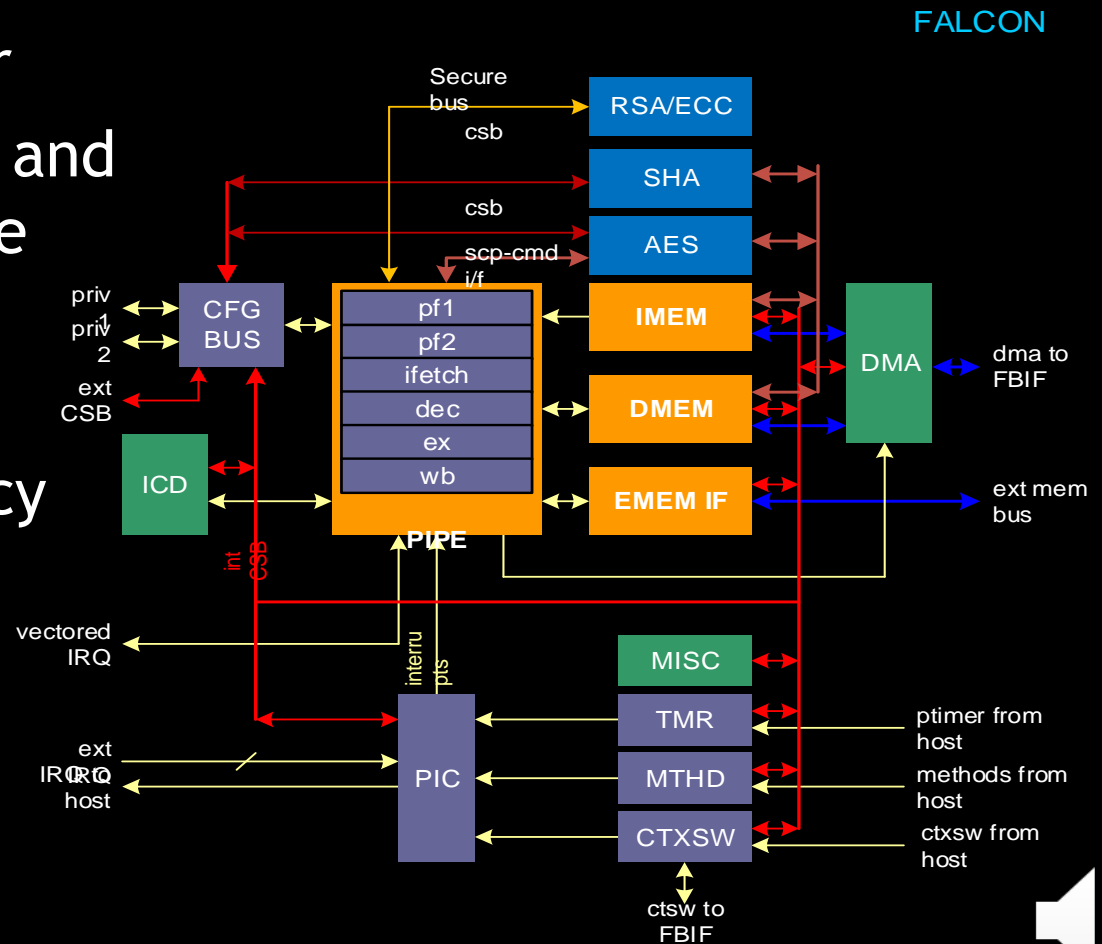
Outlines

- Introduce NVIDIA falcon CPU
- Why a new CPU?
- Introduce NV-RISCV



NVIDIA Falcon overview

- Falcon = **FA**st **L**ogic **CON**troller
- Introduced over 10 years ago, and used in >15 different hardware engines today
- Design for flexibility
- Design for long memory latency
- Design for low area
- Design for security



Why Falcon Next Gen?

- New use cases requiring more horsepower & feature
 - Wide addressing range
 - More performance
 - Not limit to code size
 - Rich OS support
- Falcon has limits
 - Small addressing range
 - Poor performance (0.67DMIPS/Mhz, 1.4Coremark/Mhz)
 - No D\$
 - No rich OS support

Falcon Next Gen - Options

- Buy
 - ARM (A,R family)
 - Synopsys (ARC family)
 - MIPS
 - Cadence
- Build
 - Improve falcon
 - Move to a new ISA (And this is when RISC-V came into the picture..)

CPU comparison

- Conclusion -
 - RISC-V is the right direction to next generation of ISA
 - Build our own implementation of RISC-V core

Item	Requirement	ARM A53	ARM A9	ARM R5	SNPS HS	RISC-V Rocket	Falcon (improved)
Core perf	>2x falcon	Yes	Yes	Yes	Yes	Yes	No
Area (16ff)	<0.1mm ²	No	No	Yes	Yes	Yes	Yes
Security	Yes	TZ	TZ	No	No	Yes	Yes
TCM	Yes	Yes	No	Yes	Yes	No	Yes
L1 I/D \$	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Addressing	64bit	Yes	No	No	No	Yes	No
Extensible ISA	Yes	No	No	No	Yes	Yes	Yes
Safety (ECC/Parity)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Functional Simulation model	Yes	Yes	No	No	Yes	No	Yes

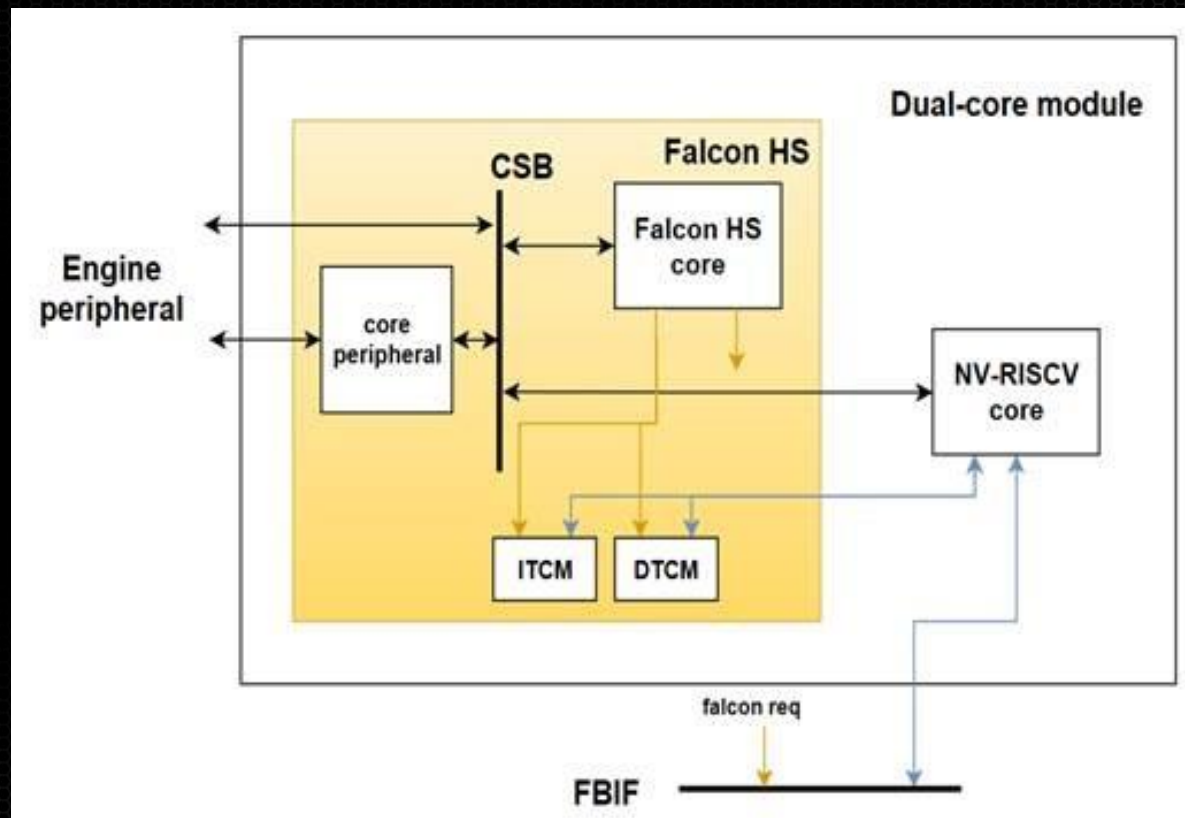
Introducing NV-RISCV

Summary: Falcon -> NV RISC-V

	Falcon	NV-RISCV
	Falcon	NV-RISCV
ISA	Falcon-ISA	RISCV-RV64
Address Width	32/24	64
Data width	32	64
GPR Num.	16	32
Stage	6	5
Micro-arch	In-order issue	In-order issue
	out-of-order exec	out-of-order exec
	out-of-order WB (diff regs)	in-order WB
WAW Hazard	Stall	ROB
Cache	No D cache	I/D configurable
TCM	I/D configurable	I/D configurable
Prediction	static	BTB/BHT/RAS
Load-store	In-order	Load out-of-order
Memory protection	No	MPU
Address mapping	TCM tagging	Base and bound

Falcon Next Gen with RISC-V

- RISC-V plugged-in as 2nd core
 - Back compatibility on interface, easy to integrate
 - Isolation between security and non-security applications



NV-RISCV Core perf/area

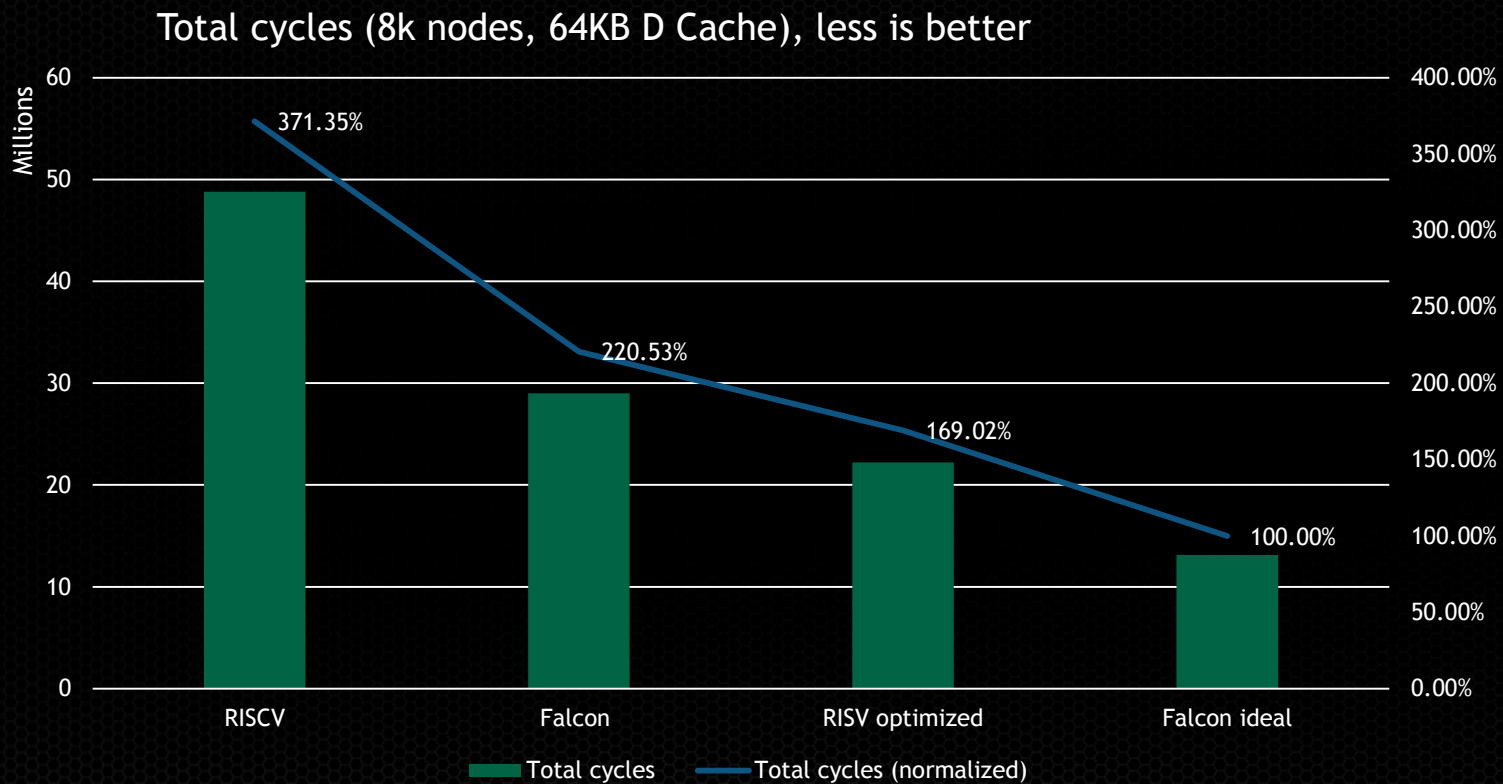
- Area data under 16ff

Core (TSMC)	Falcon (Today)	RISCV rocket chip	NV-RISCV	BOOM
Dhrystone (no inline)	0.67	1.72	1.9~2.0	N/A
EEMBC Core Mark	1.4	2.3	~2.5	3.91
Core Area(mm ²)	0.03	0.055 (no FPU, no L2)	0.05~0.06	N/A
Frequency(GHz)	1.5	~1	>1.5	N/A

Cache design to tolerate large latency

- Configurable cache size/line size/associativity/write policy
- Cache Optimizations
 - Store Buffer
 - Write merging
 - Line-fill Buffer
 - Victim Buffer
 - Stream buffer
 - SW pre-fetch (future)
 - L2 (future)
 - Banked cache (future)
- I/DTCM

D\$ perf - btree (8k nodes)



RISCV - Area of Interests to NV

- Tool chain
 - Tool for automotive - To meet ISO26262/ASIL-D or SIL3 requirements
 - Tool for debug - To debug ucode on silicon/FPGA/emulator
 - Tool for performance tuning - For ucode profiling and tuning
 - Tool for flexibility - That users can easily customize ISA
 - Other compiler features - ILP32/LP64
- Security
 - Crypto instruction & extensions
- More instructions
 - Cache instruction - pre-fetch/invalidate/flush...

Conclusion

- Falcon is NVIDIA proprietary control processor
- New use-cases require more feature and performance from falcon
- It is hard to improve the current CPU/ISA to meet all new requirements
- We evaluated different options in the market, result showed that RISC-V is overall best choice as next generation of falcon
- We will build a new core from RISC-V ISA

Thank You!