ETH zürich

ENZIAN Systems@ETH zürich DINFK

Trustworthy Board Management Software

Daniel Schwyn, Ben Fiedler, David Cock, Michael Giardino, Timothy Roscoe, et al.
Systems Group @ ETH Zürich

The Computer in the Computer

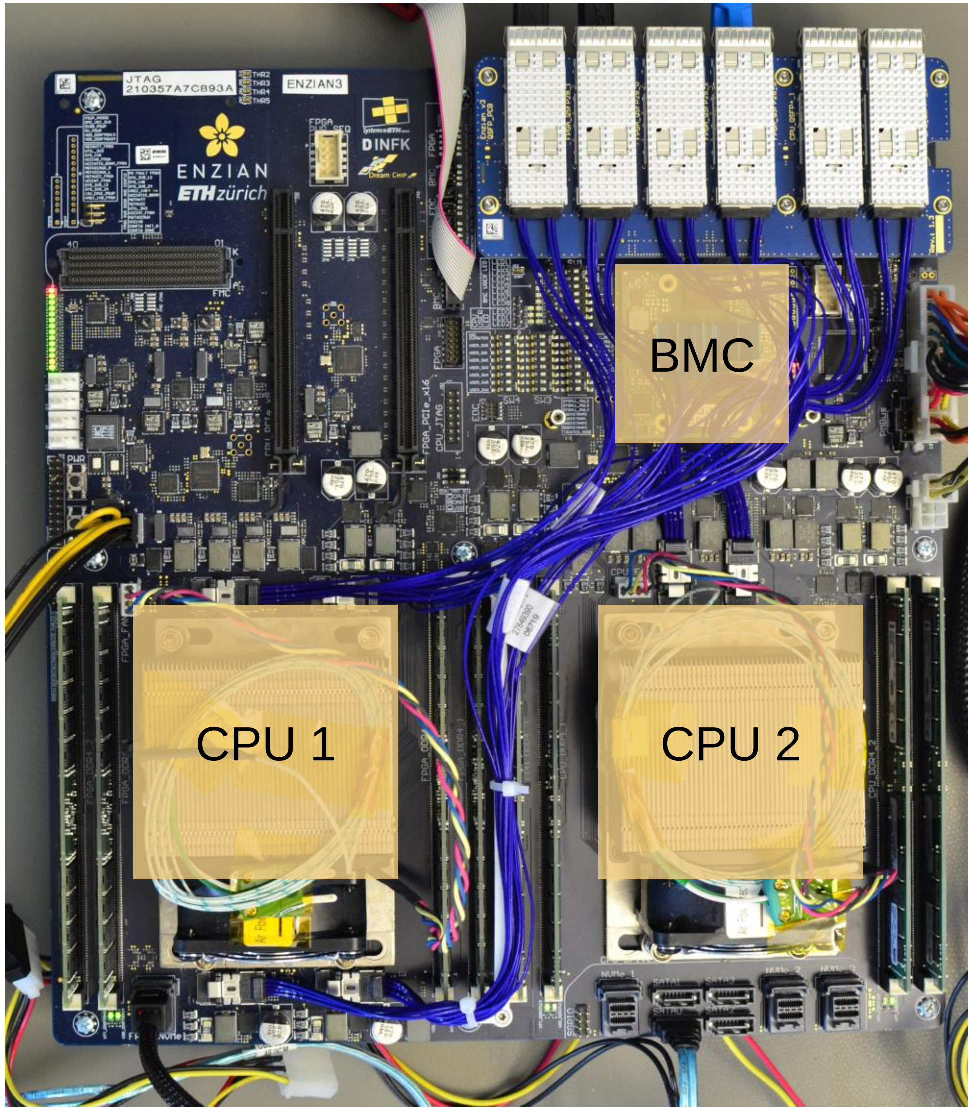

The Computer in the Computer

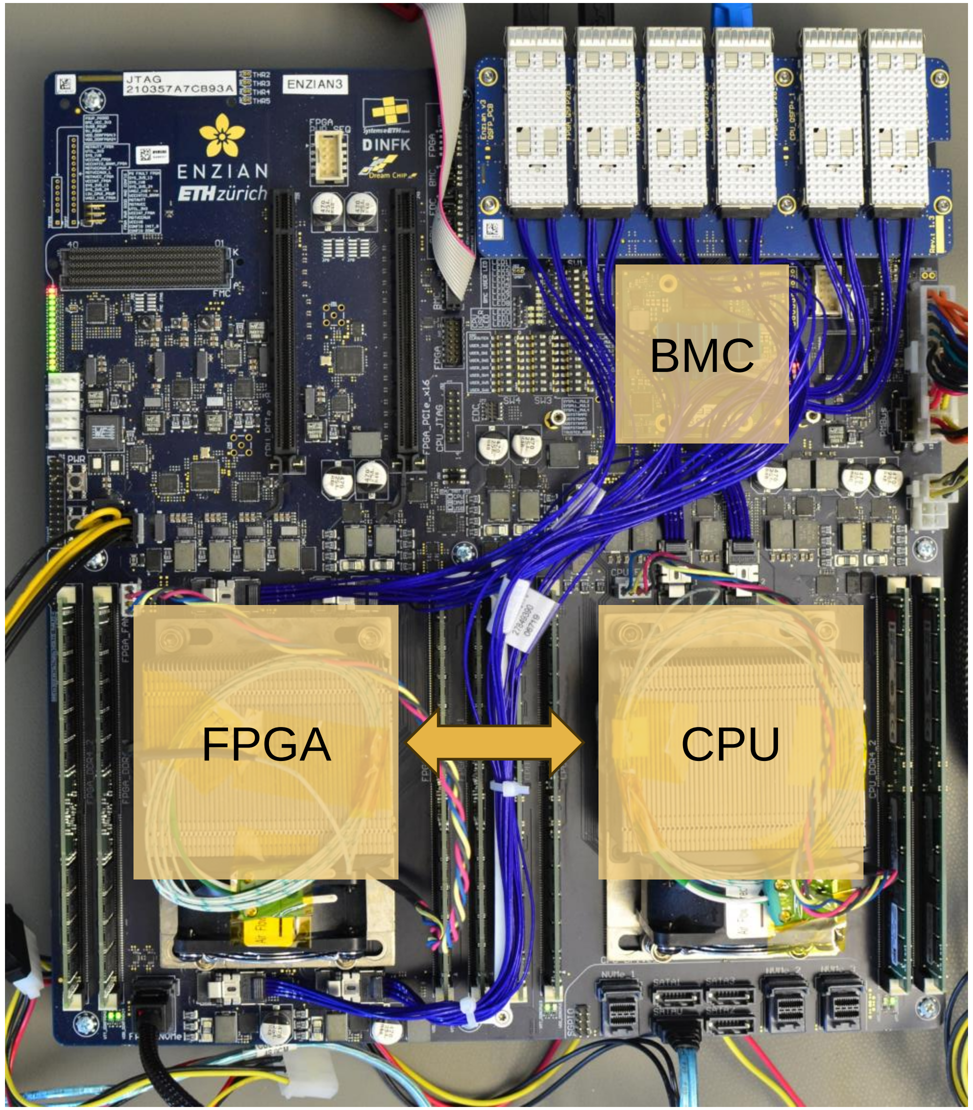

State of the Art – Closed source

ENZIAN Systems@ETH Zürich D INFK

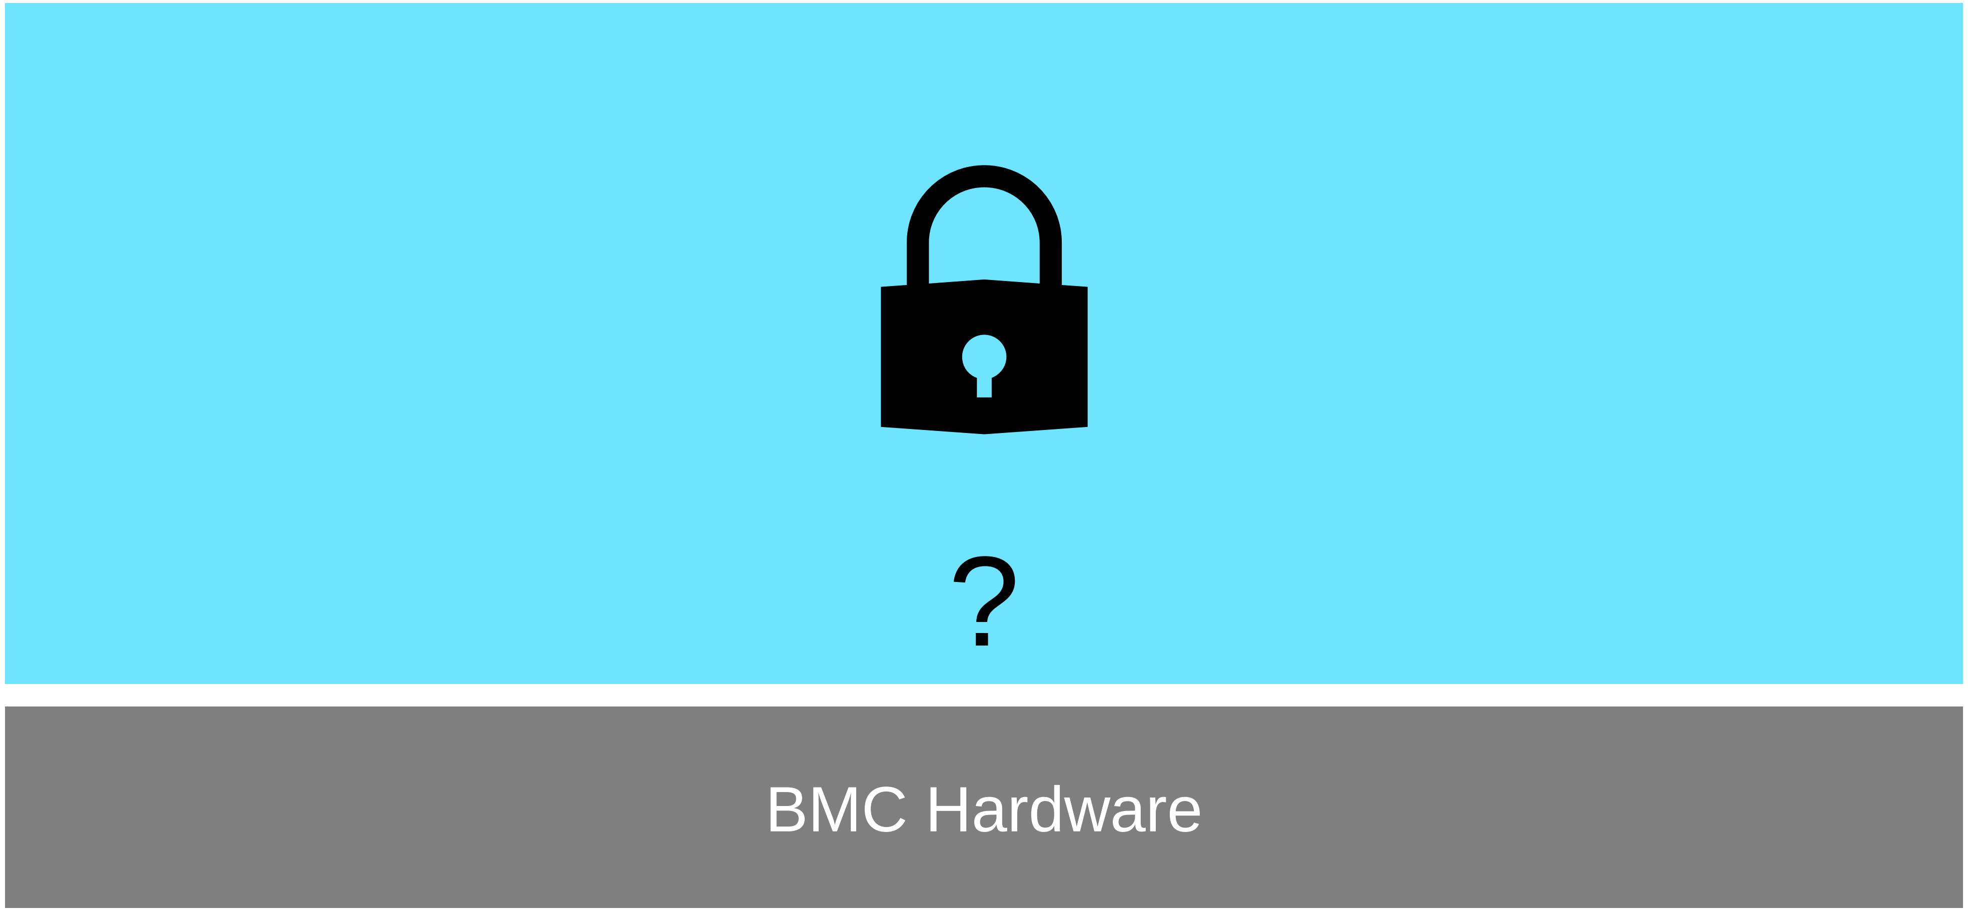

… and buggy

ENZIAN Systems@ETH zürich DINFK

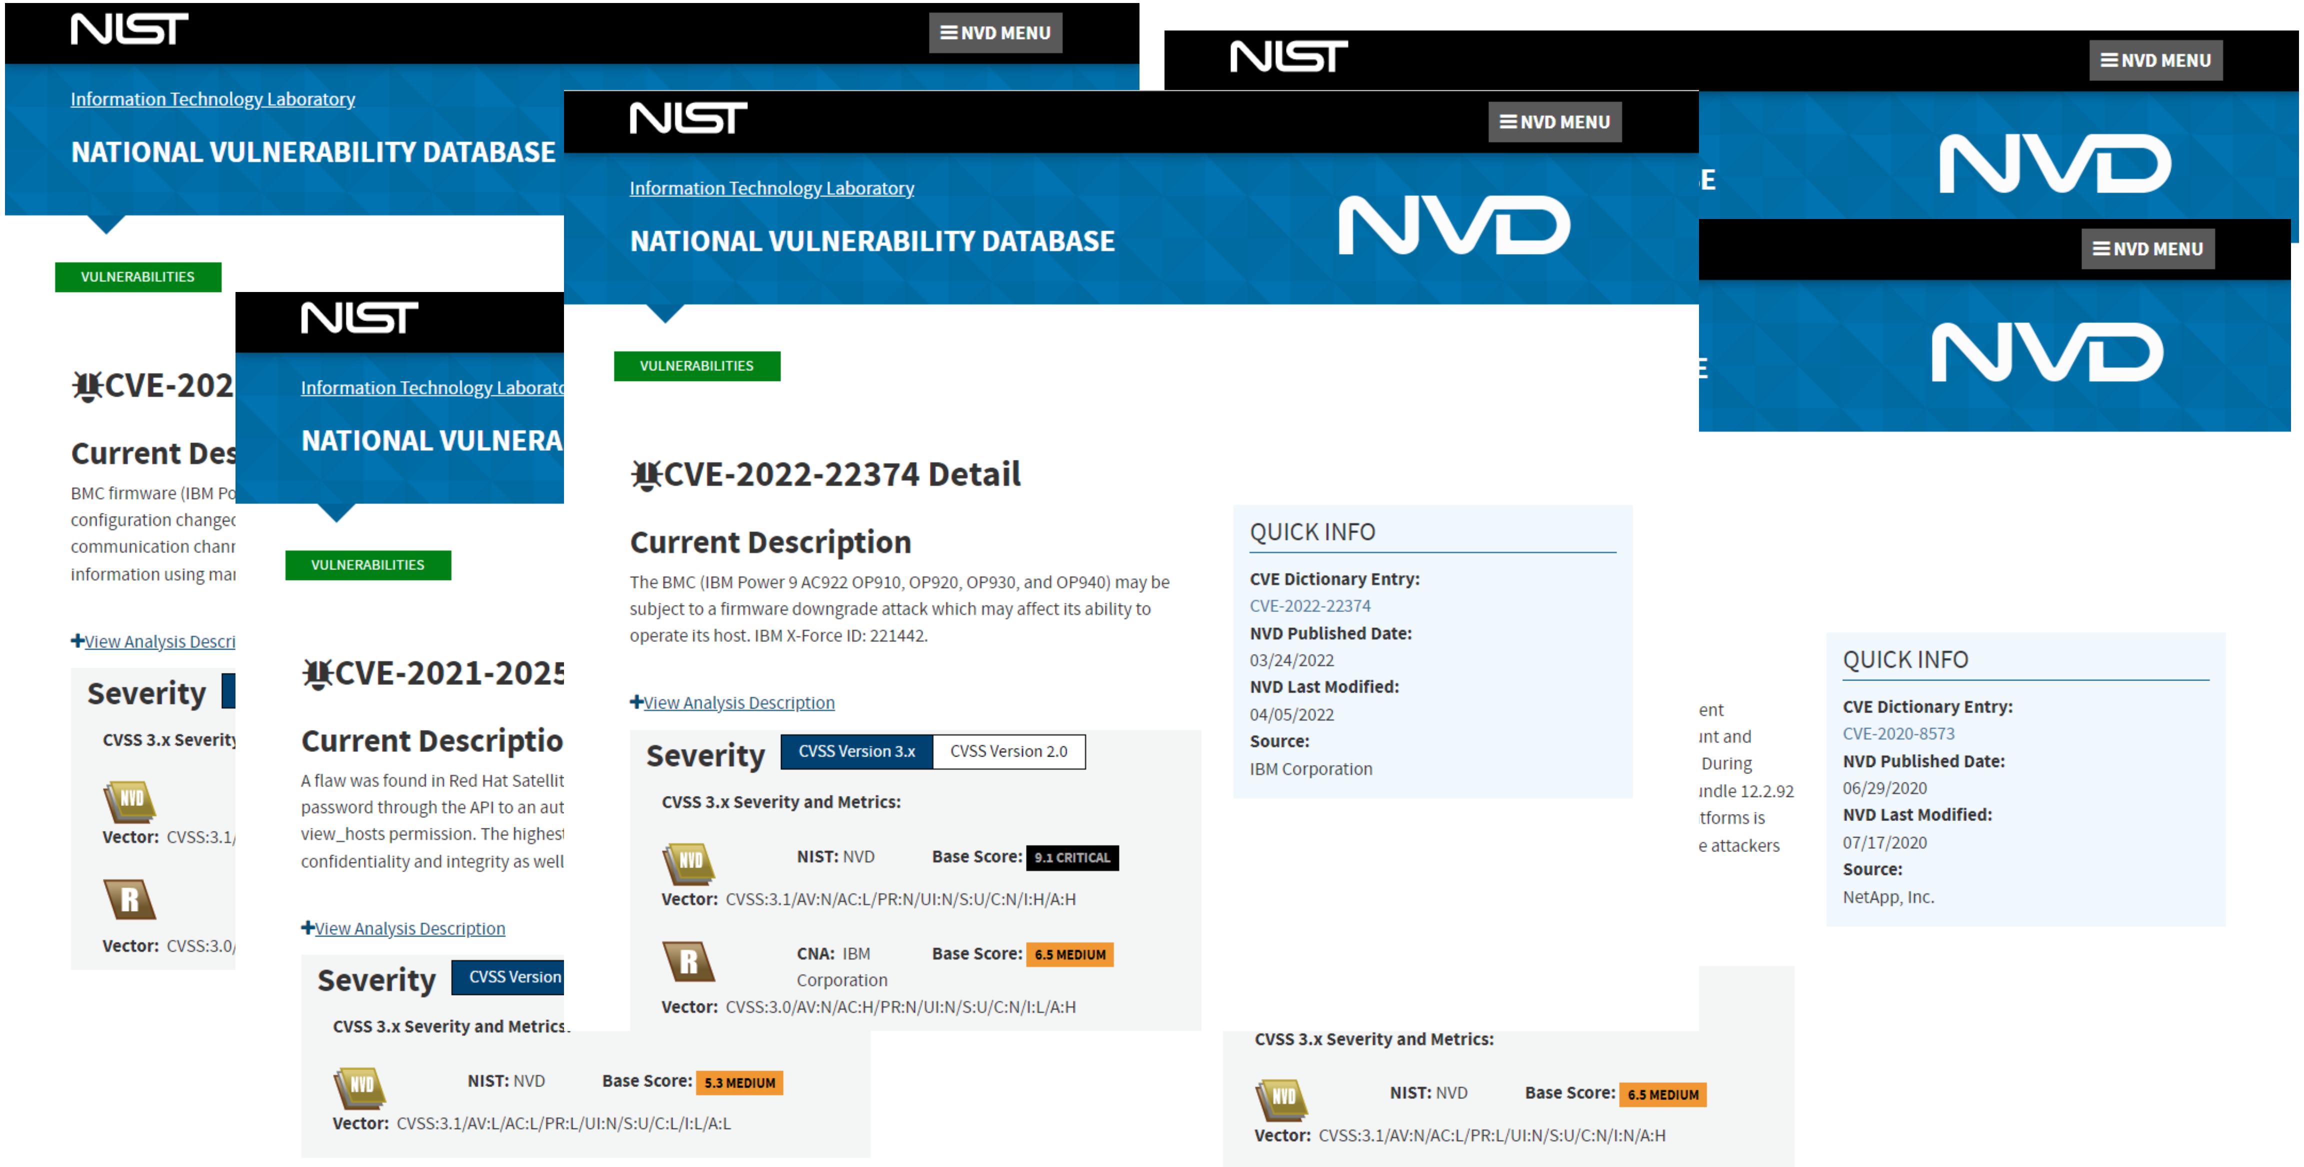

State of the Art – Open Source

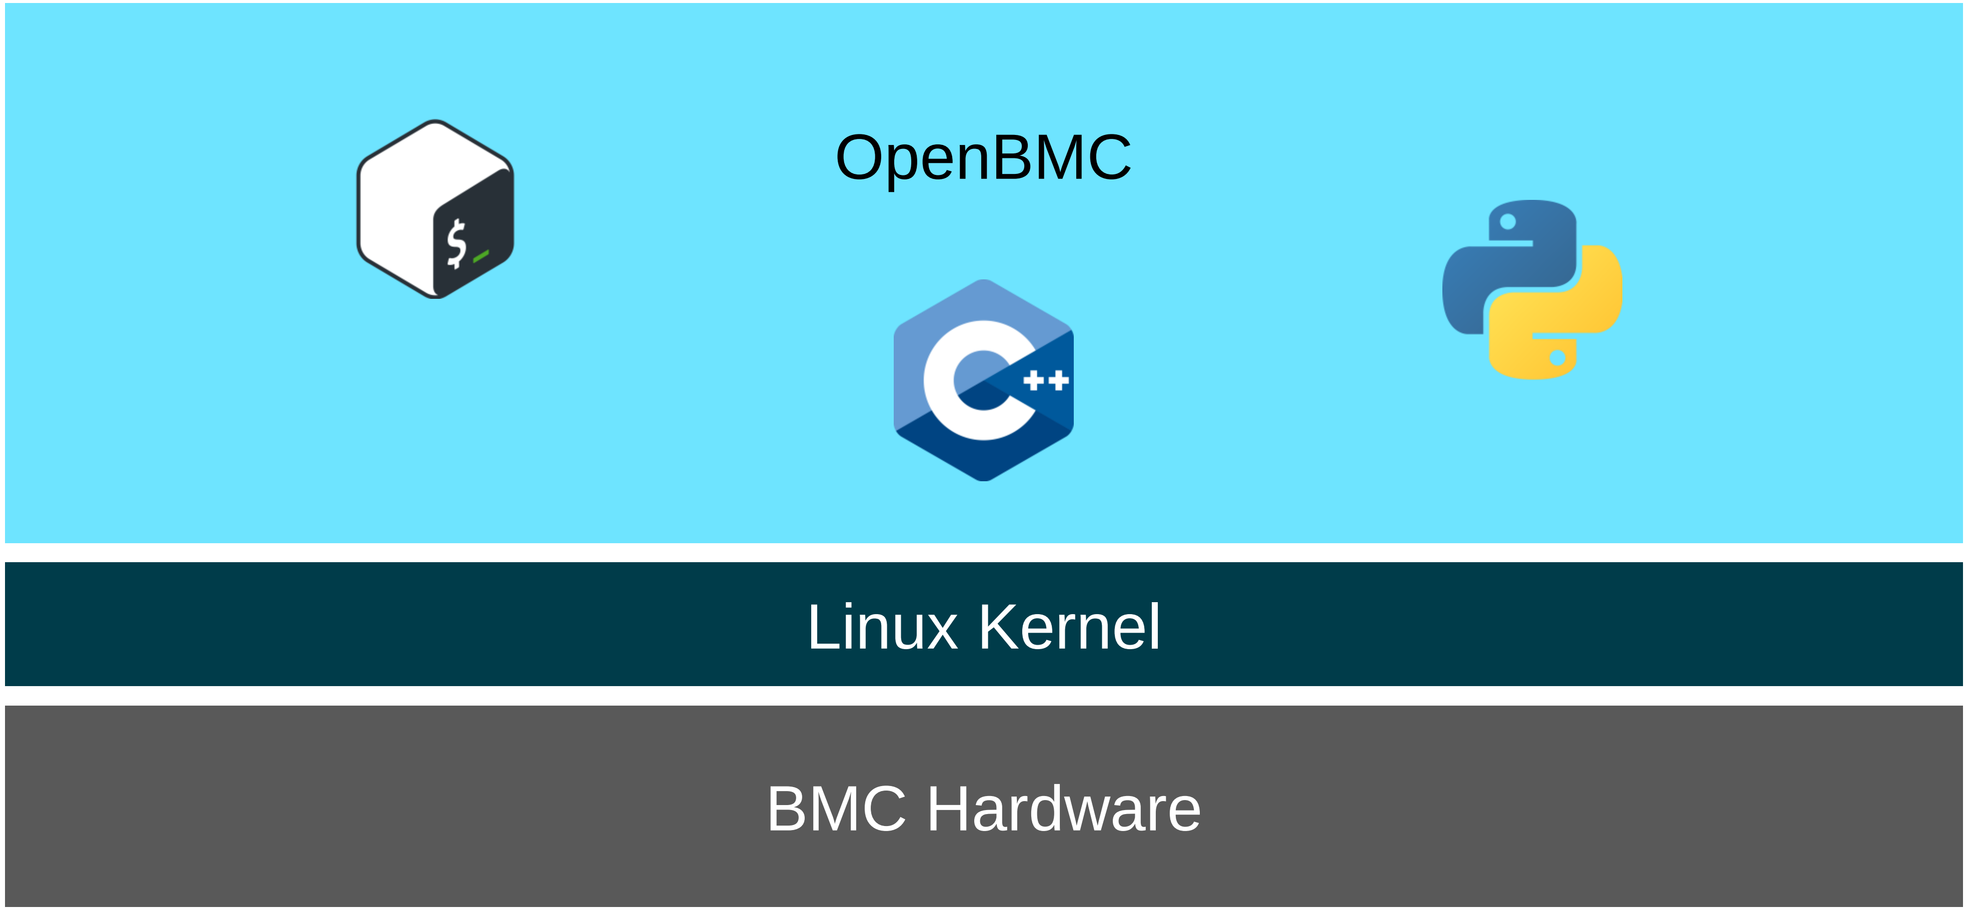

State of the Art – Open Source

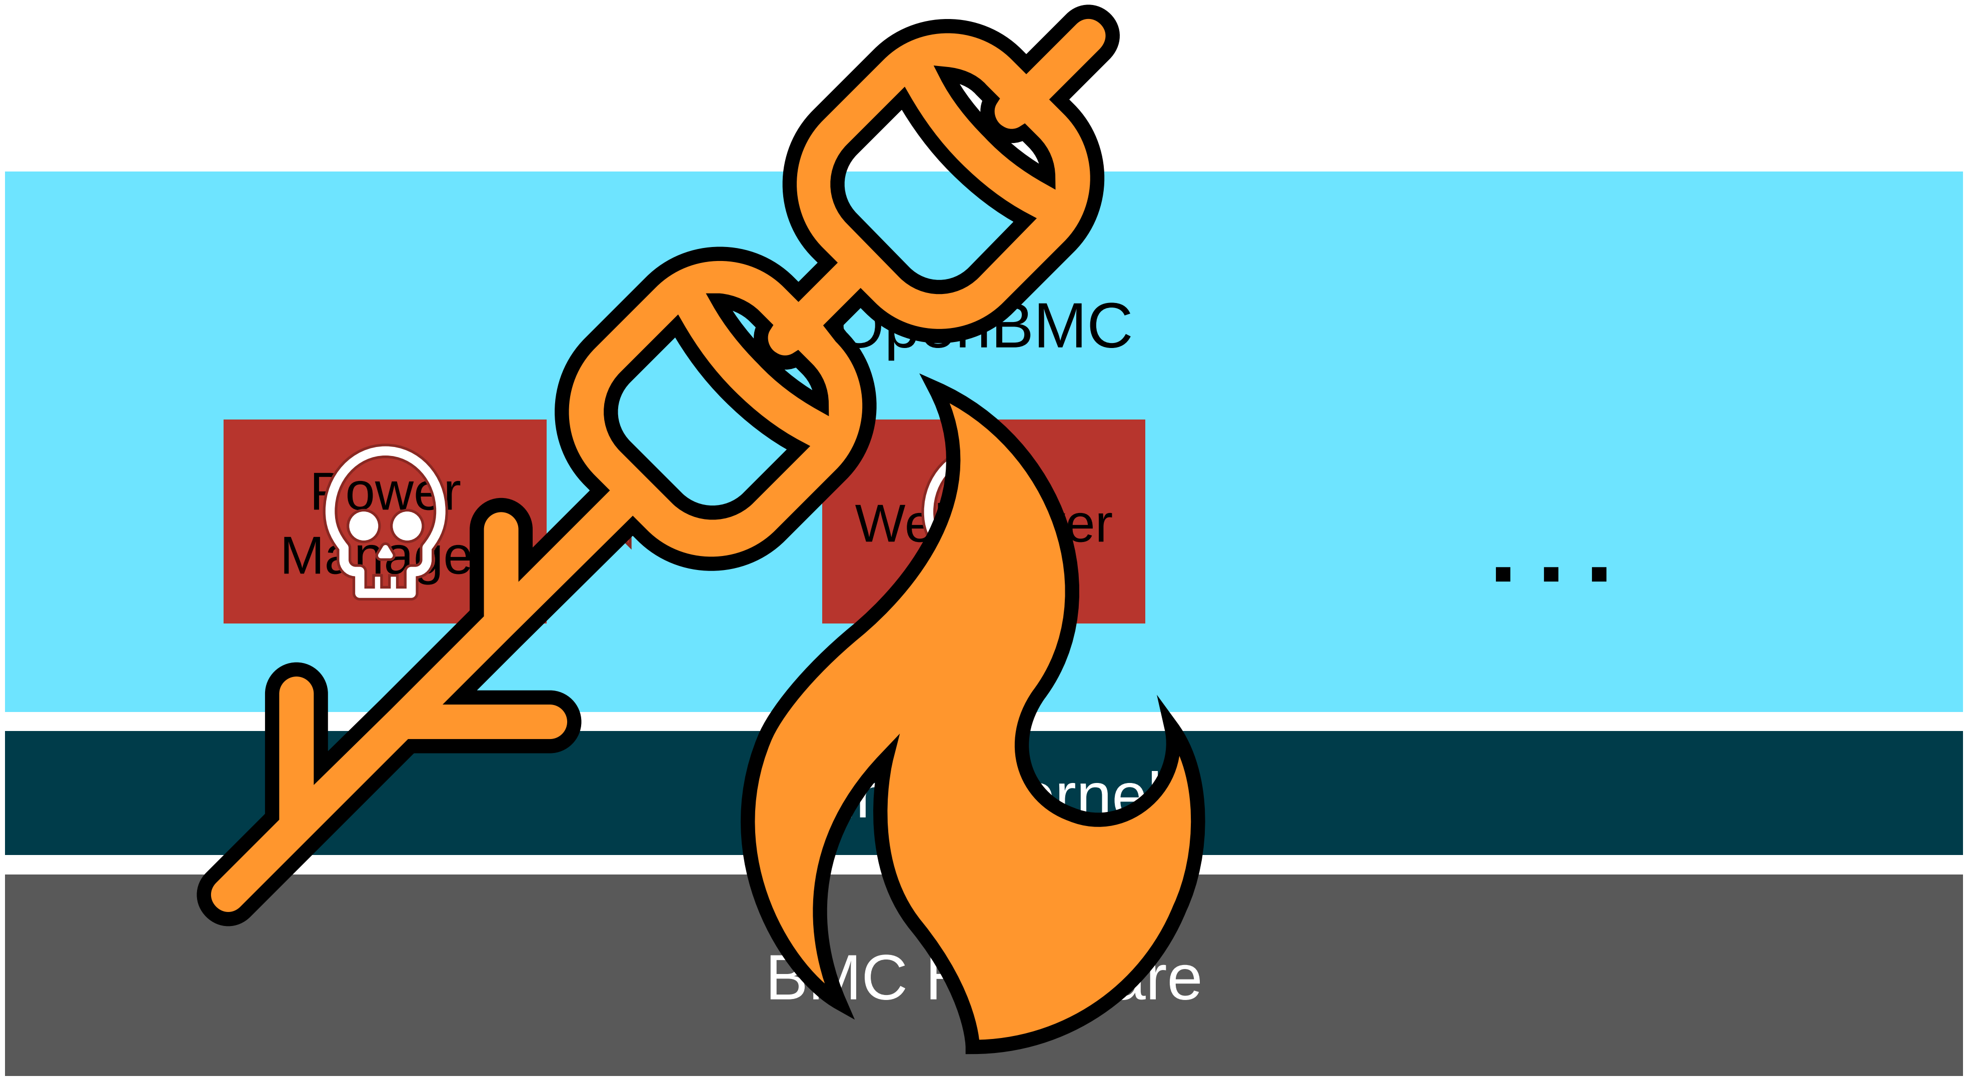

Trust & Threat Model

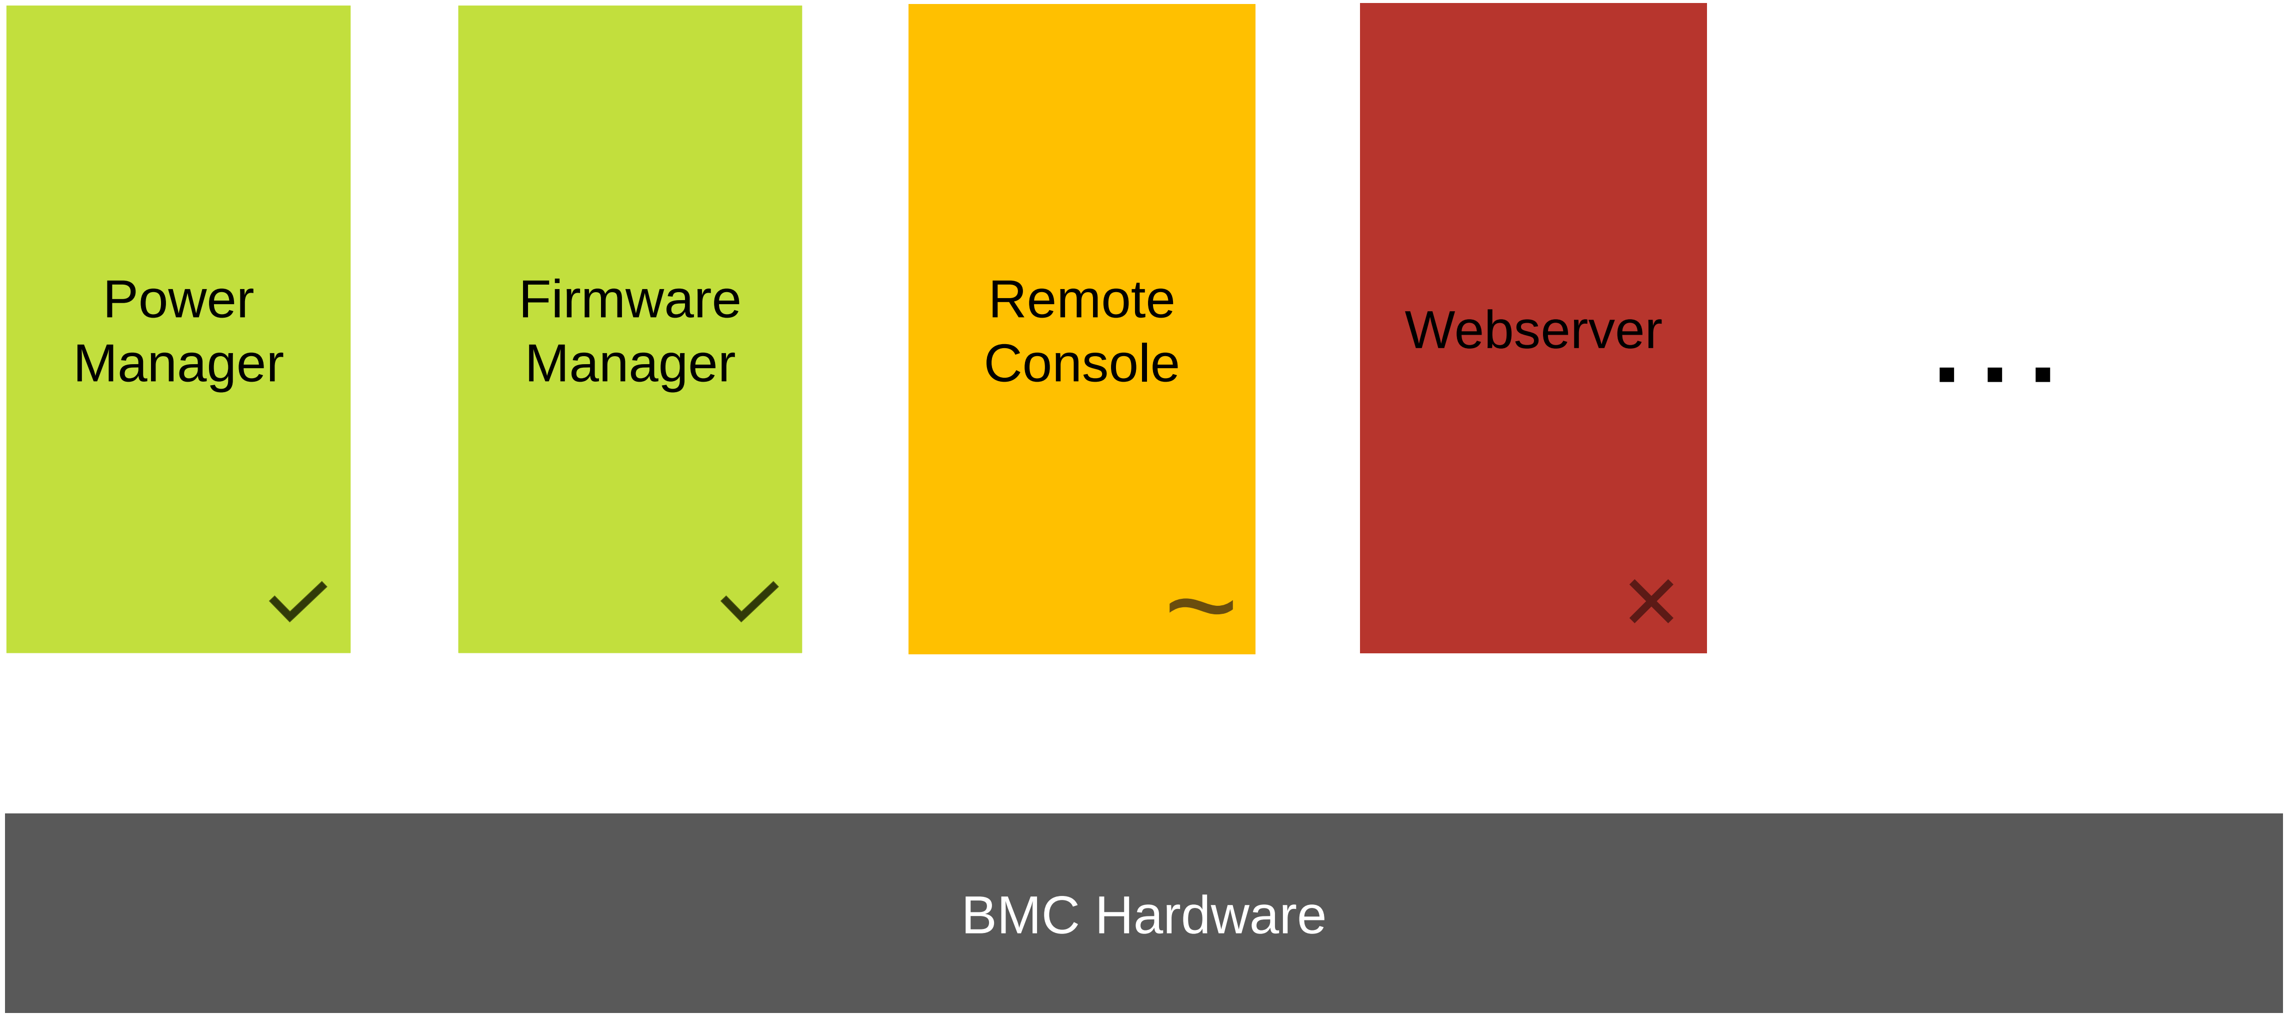

Isolation:
Confidentiality, Integrity, Availability

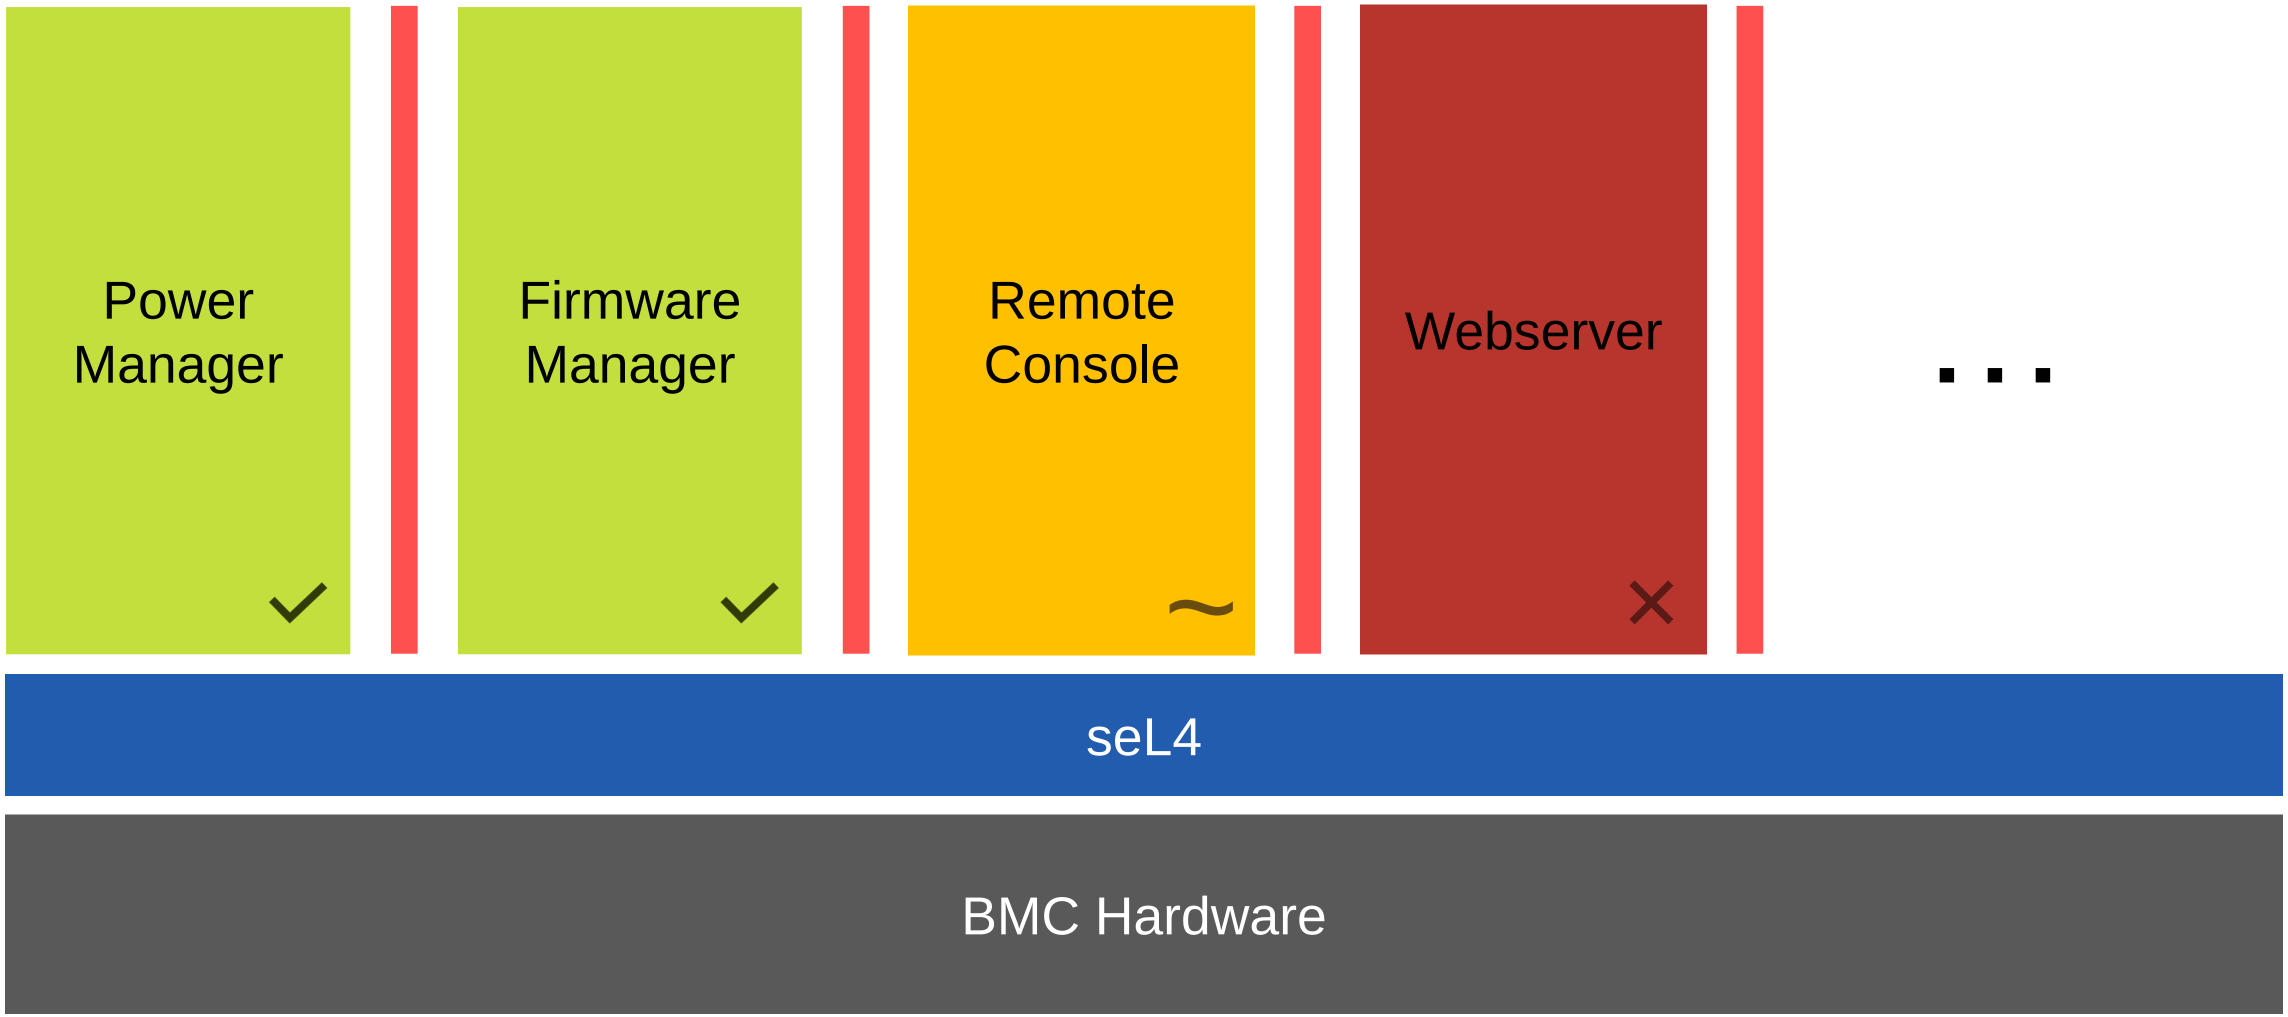

Cyber Retrofit – Step 1

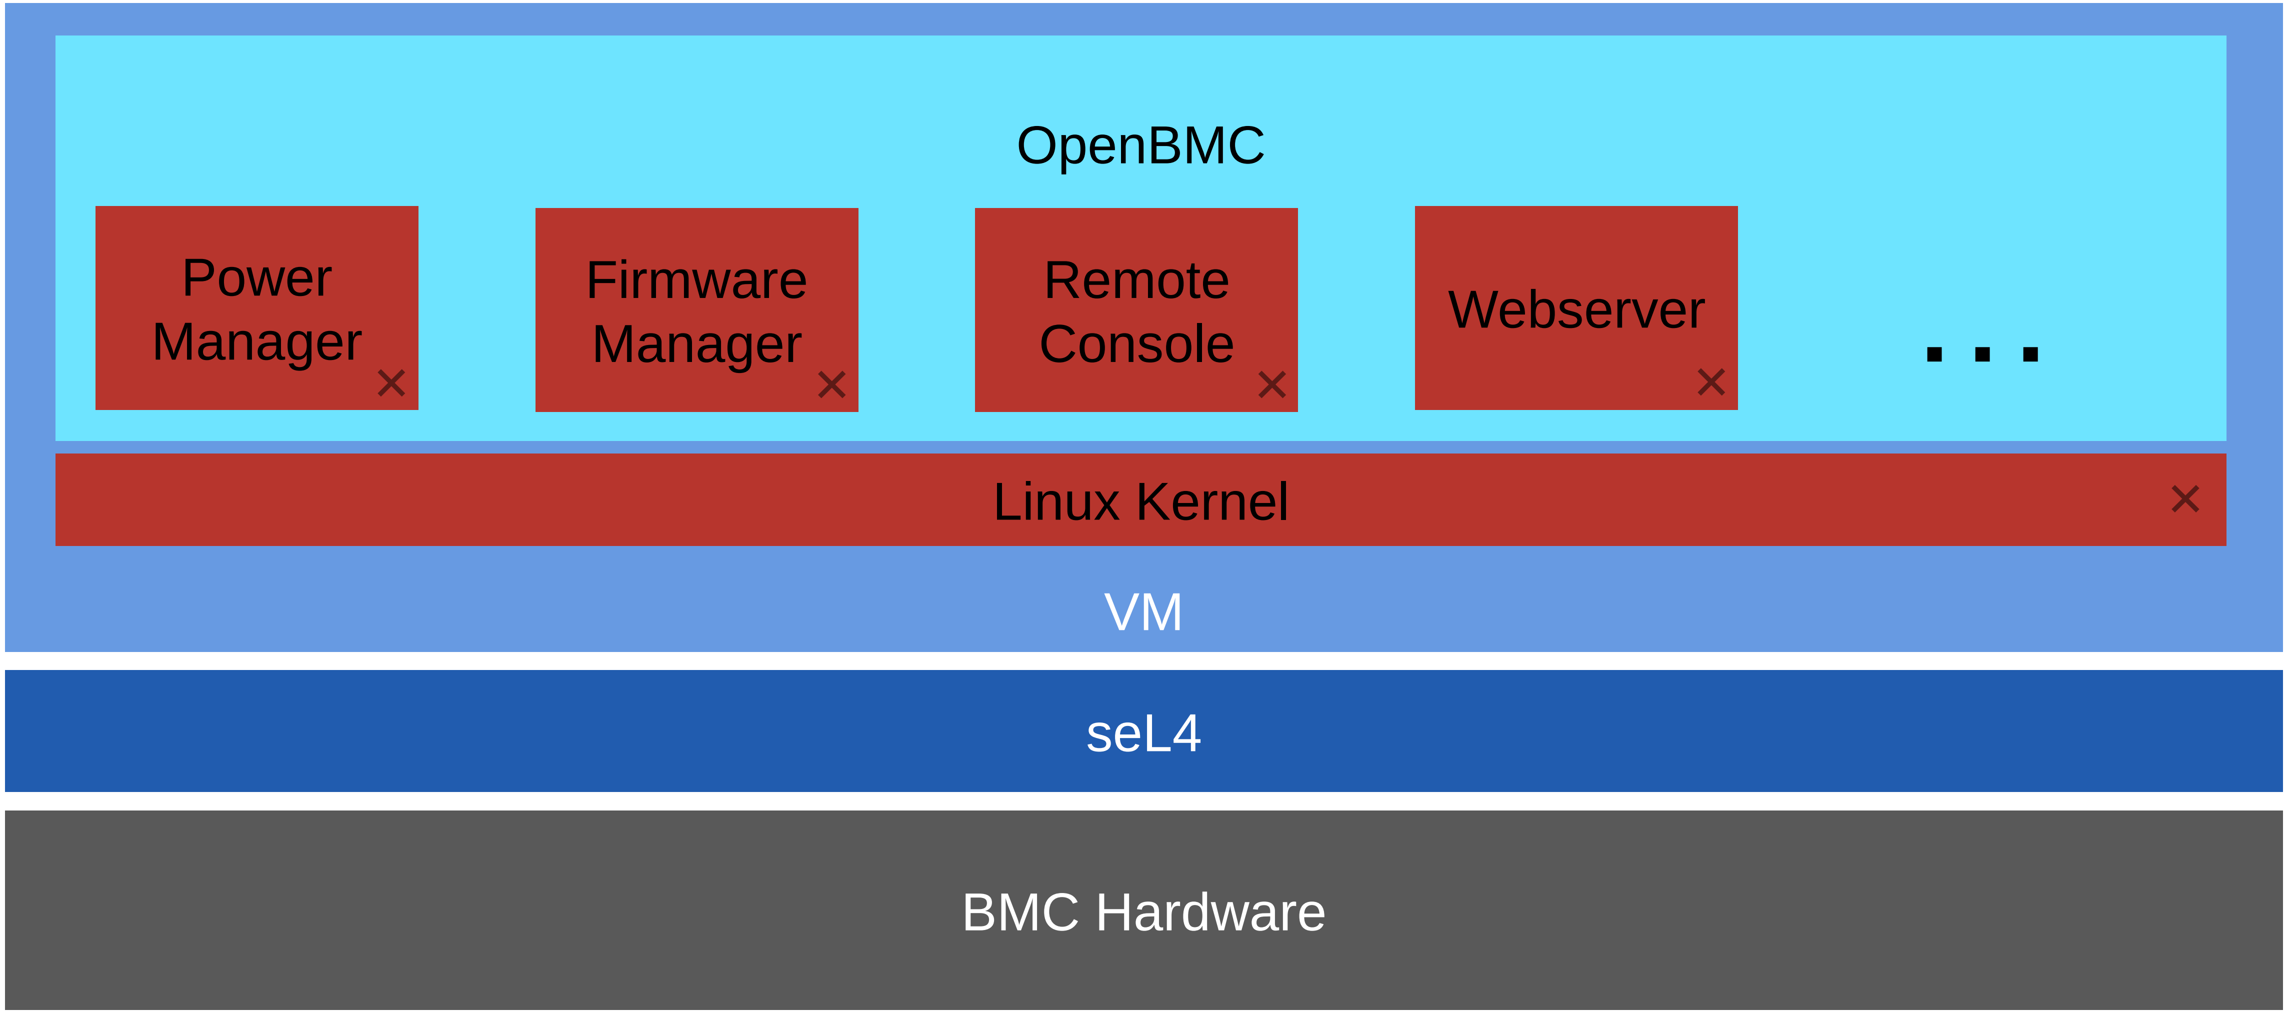

Cyber Retrofit – Step 2

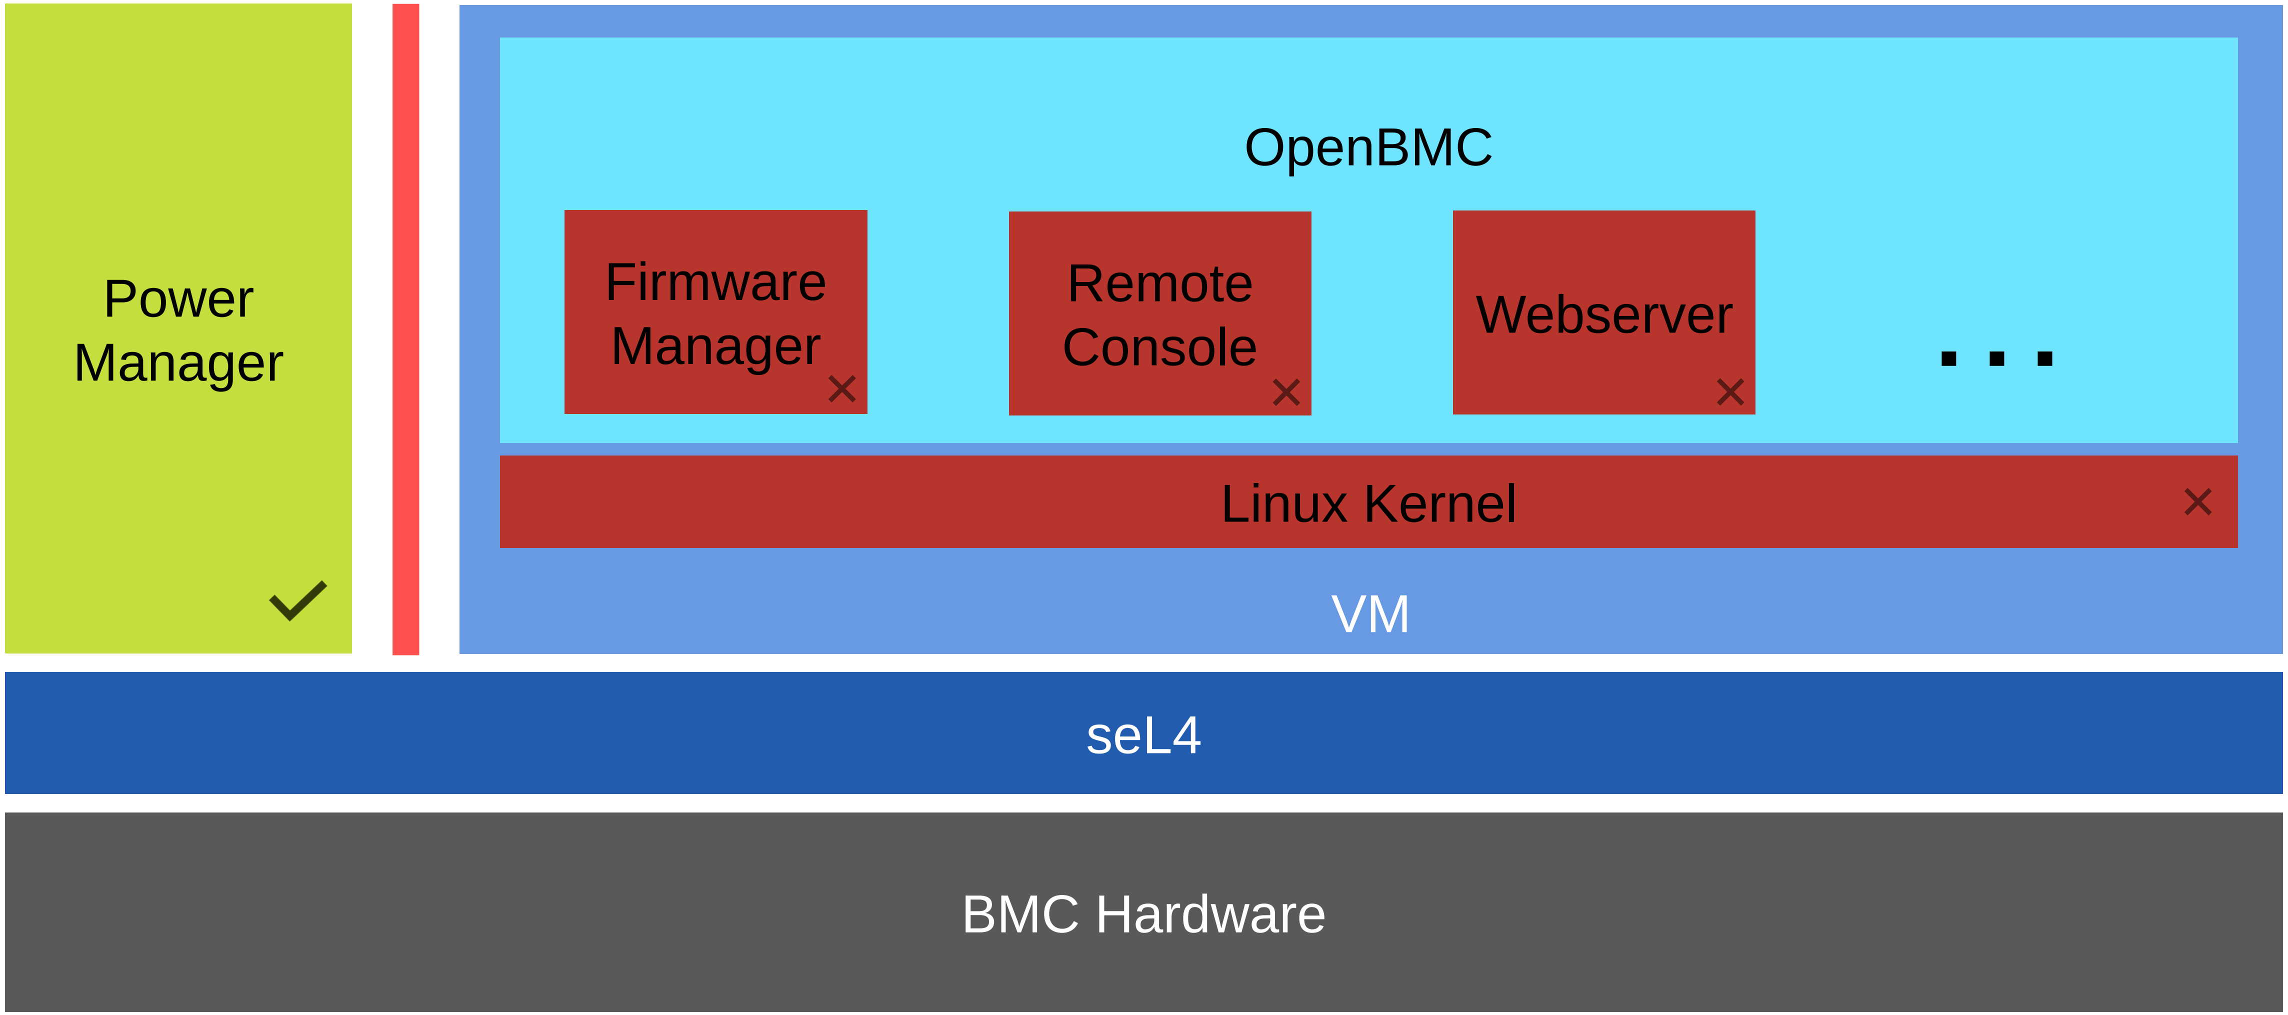

Turning on computers is hard…

ENZIAN

Systems@ETH zürich

D INFK

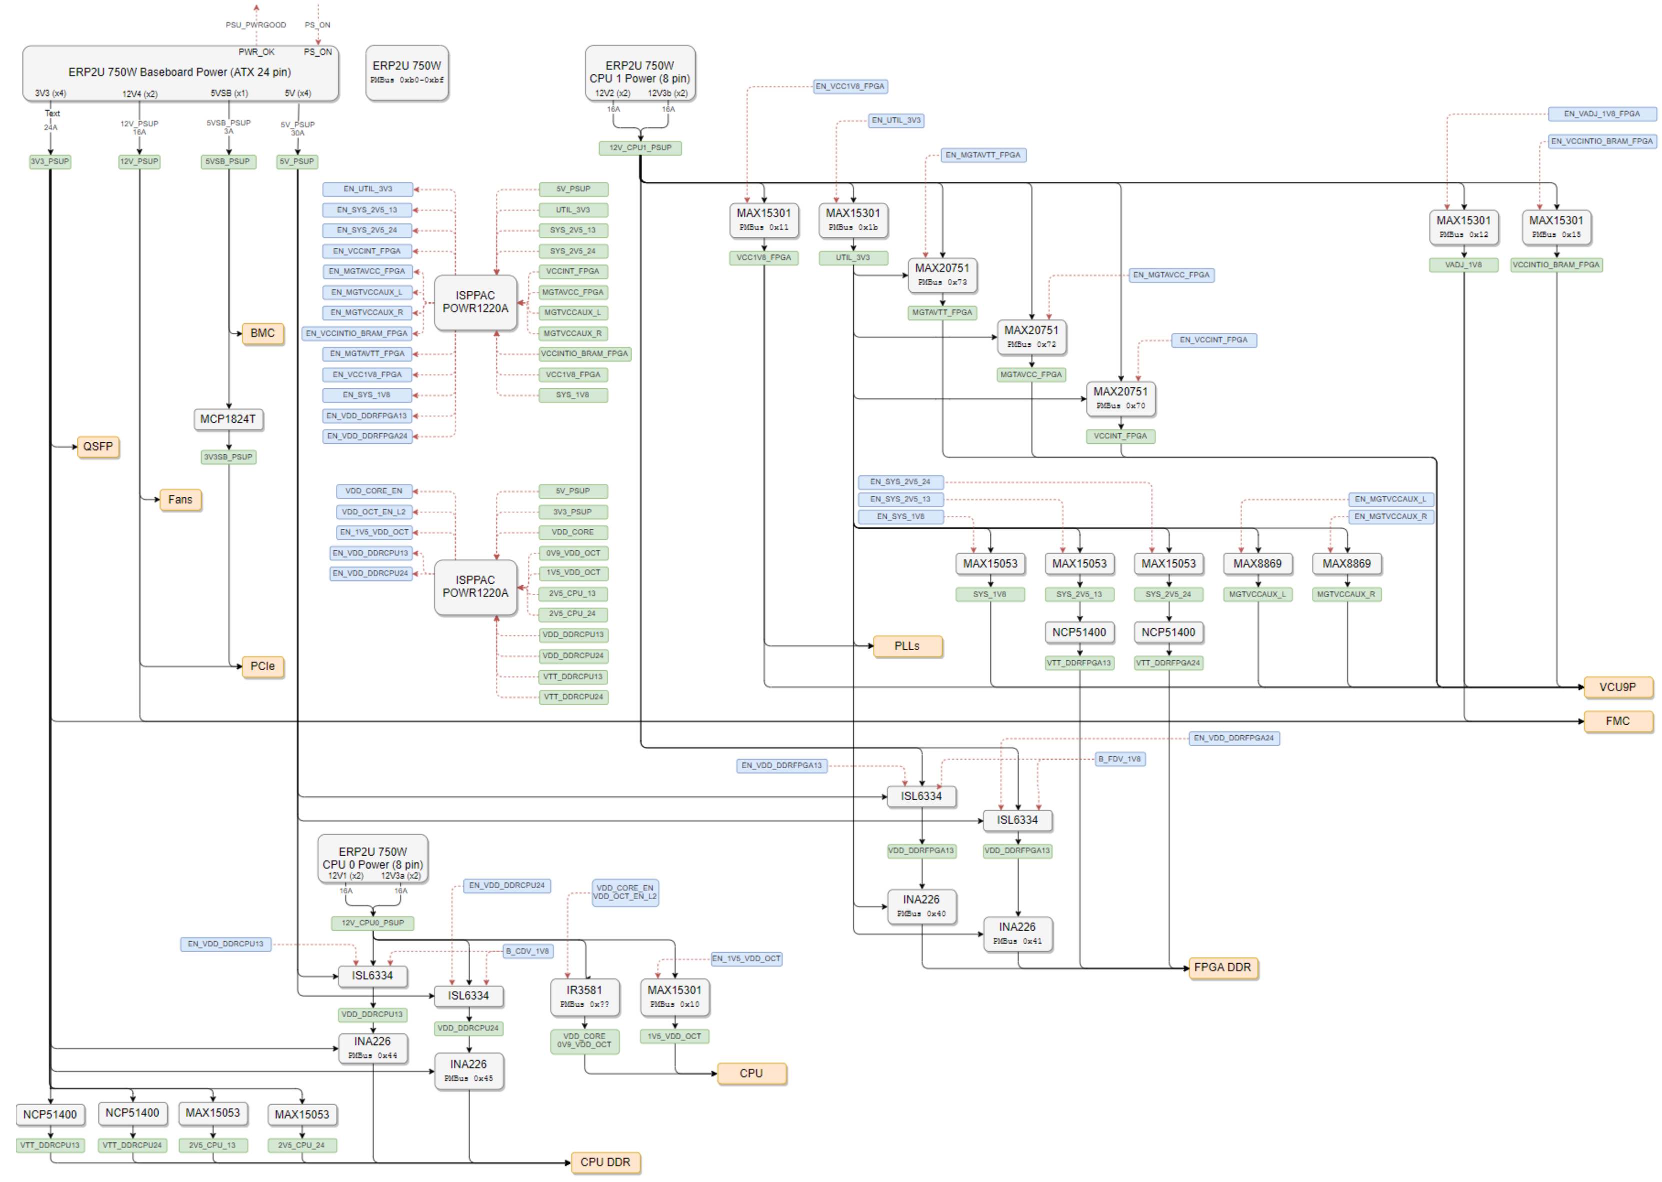

Turning on computers is hard…

ENZIAN Systems@ETH zürich DINFK

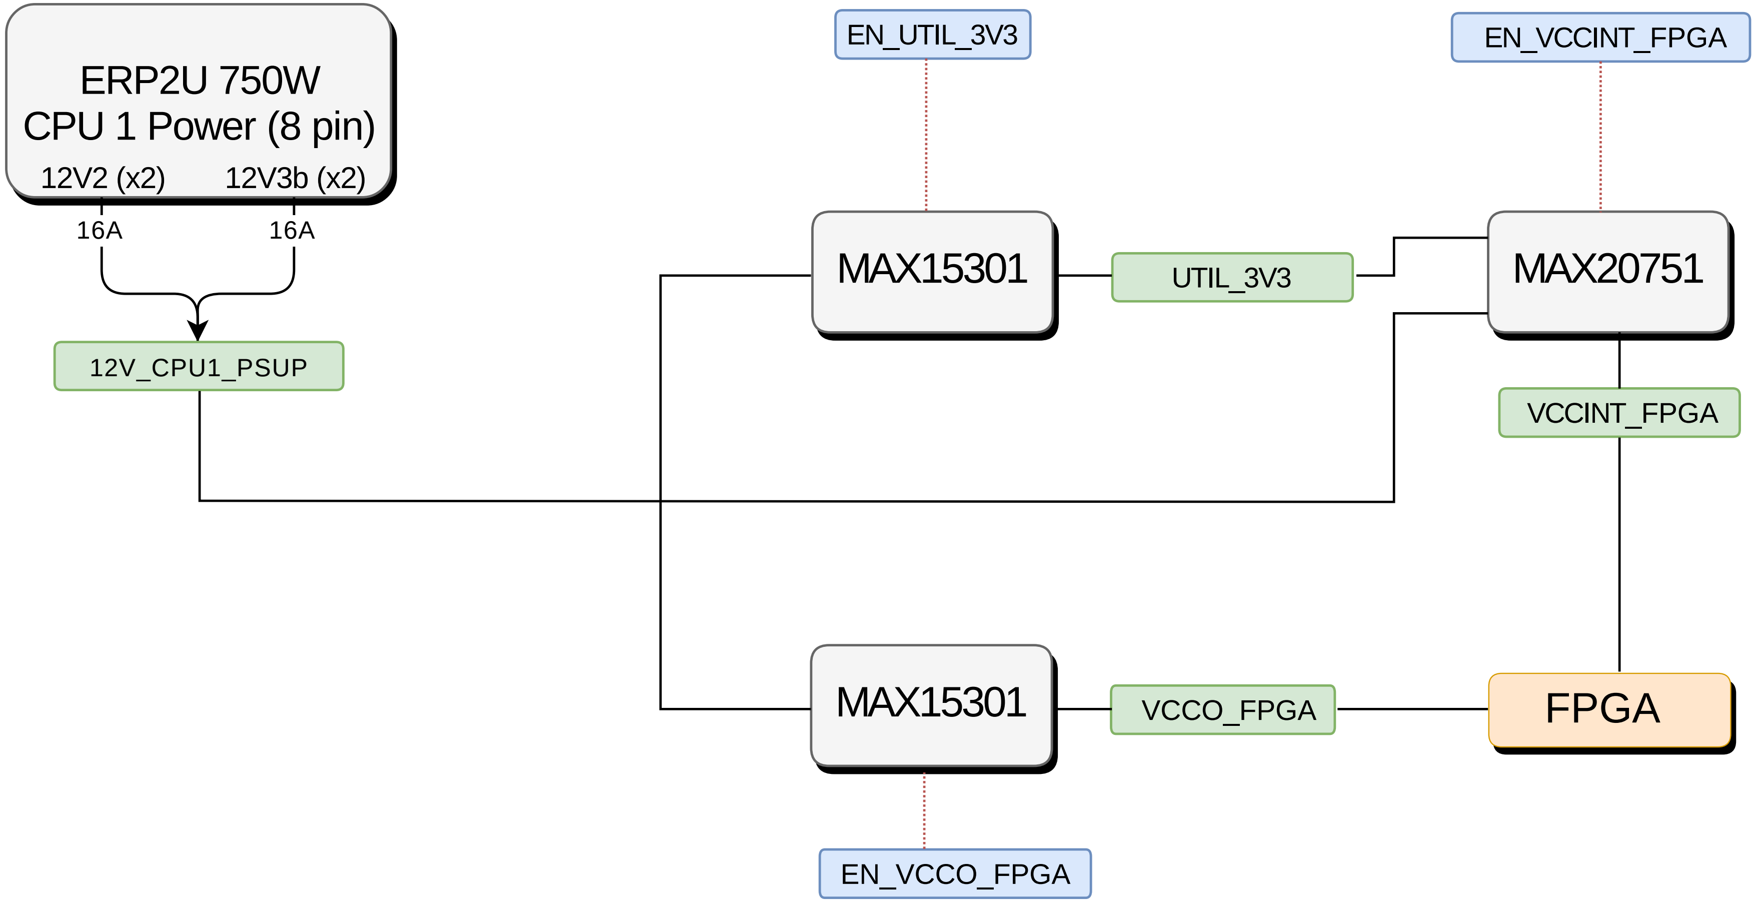

… and we need a more systematic approach!

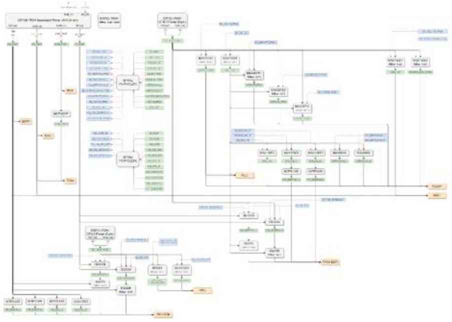

SMT
Solver

1.
2.
3.
4.
5.

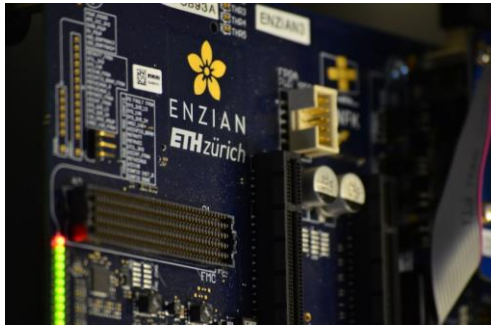

- We can derive sequences efficiently (< 10s)
- Works on real hardware
- More confidence in power-up sequences
- Basis for rigorous specification of correct behavior

Jasmin Schult, Daniel Schwyn, Michael Giardino, David Cock, Reto Achermann, and Timothy Roscoe. 2021. Declarative Power Sequencing. ACM Trans. Embed. Comput. Syst. 20, 5s, Article 84

Trustworthy Power Manager

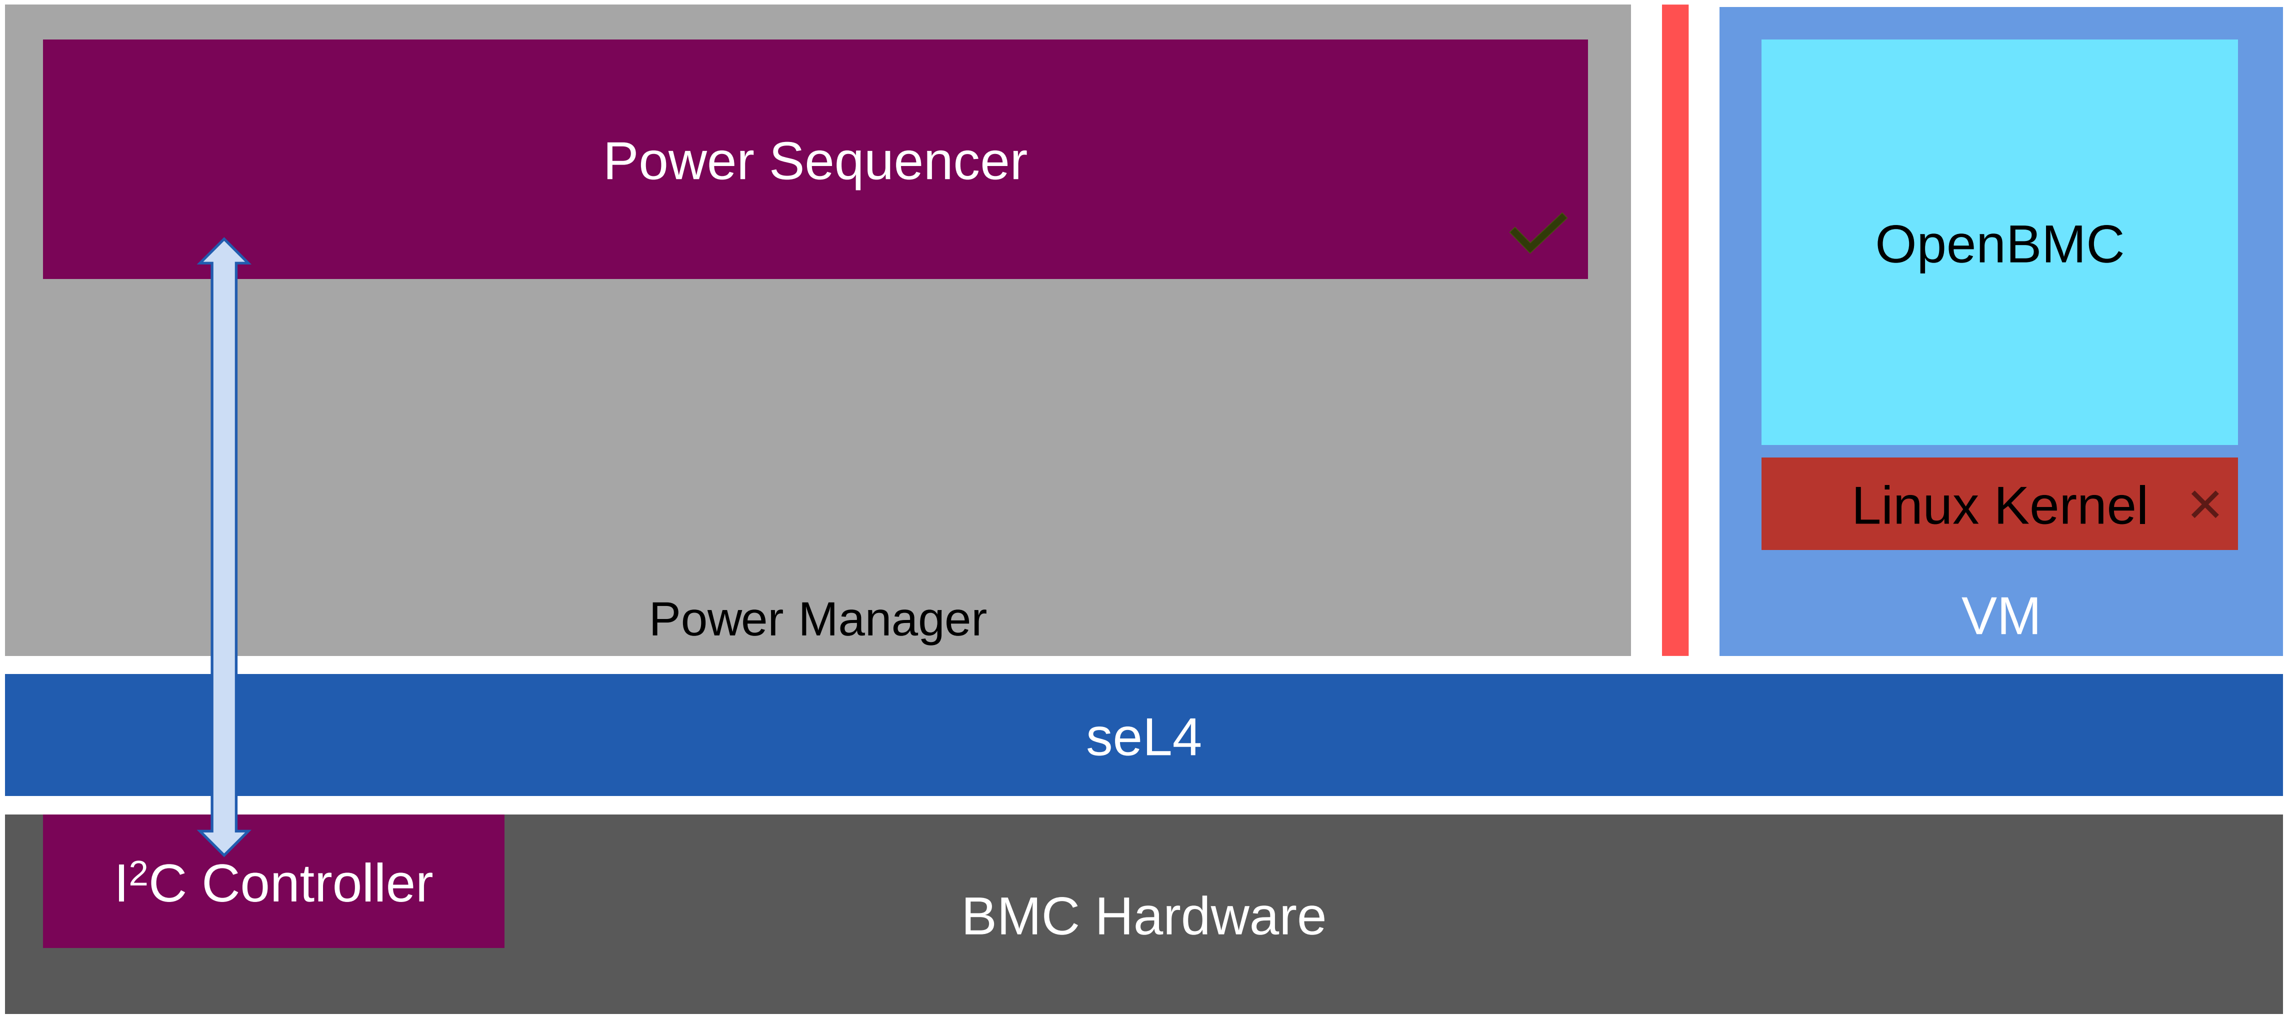

I^{2}C

- Widespread, low-speed configuration bus
- Controls critical hardware components

- Devices implement the standard only partially or violate it

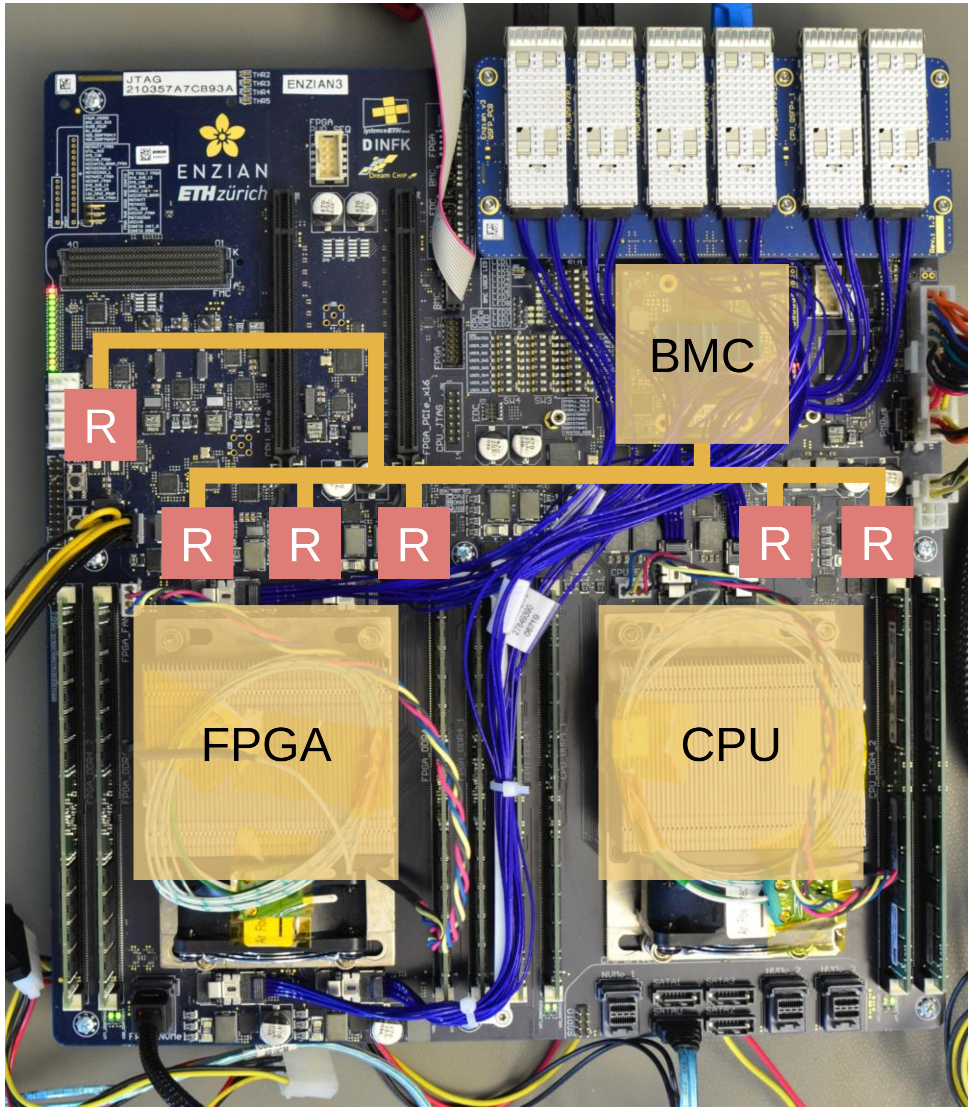

I^{2}C bus

R regulator / sensor

I^2C Modelling Framework

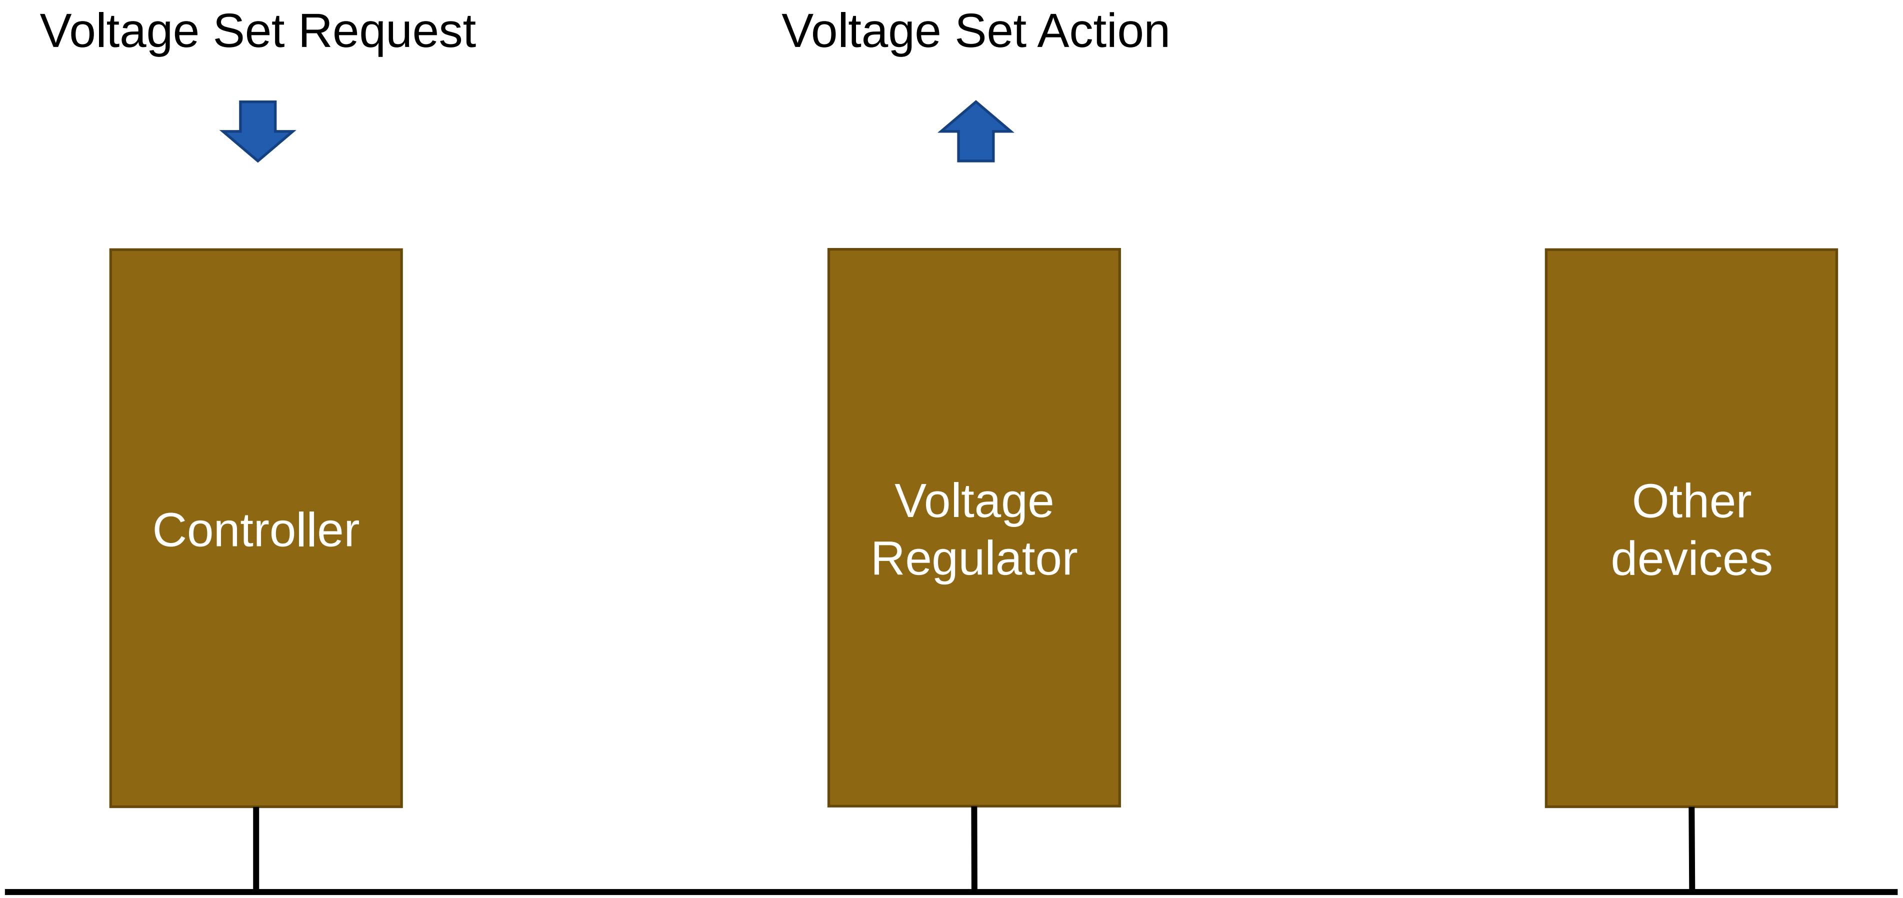

Lukas Humbel, Daniel Schwyn, Nora Hossle, Roni Haecki, Melissa Licciardello, Jan Schaer, David Cock, Michael Giardino, Timothy Roscoe. 2021. A Model Checked I^2C Specification. 27th International Symposium on Model Checking Software (SPIN 2021)

I²C Modelling Framework

ENZIAN Systems@ETH Zürich D INFK

Spin checked

Voltage Set Request

Voltage Set Action

Generated C Code running on actual hardware

Driver
Transaction
Byte
Symbol
Electrical

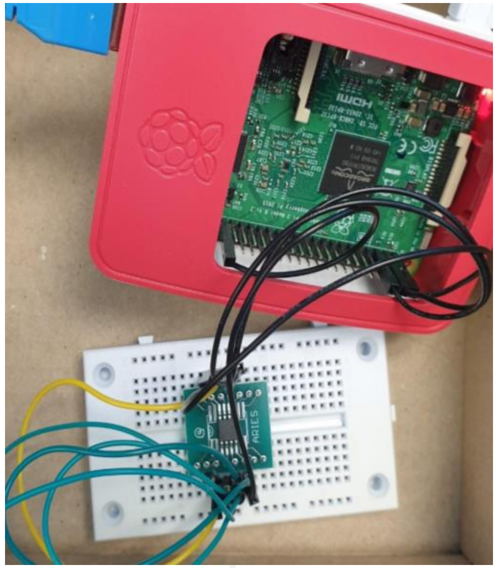

Driver
Transaction
Byte
Symbol
Electrical

Lukas Humbel, Daniel Schwyn, Nora Hossle, Roni Haecki, Melissa Licciardello, Jan Schaer, David Cock, Michael Giardino, Timothy Roscoe. 2021. A Model Checked I²C Specification. 27th International Symposium on Model Checking Software (SPIN 2021)

Trustworthy Power Manager

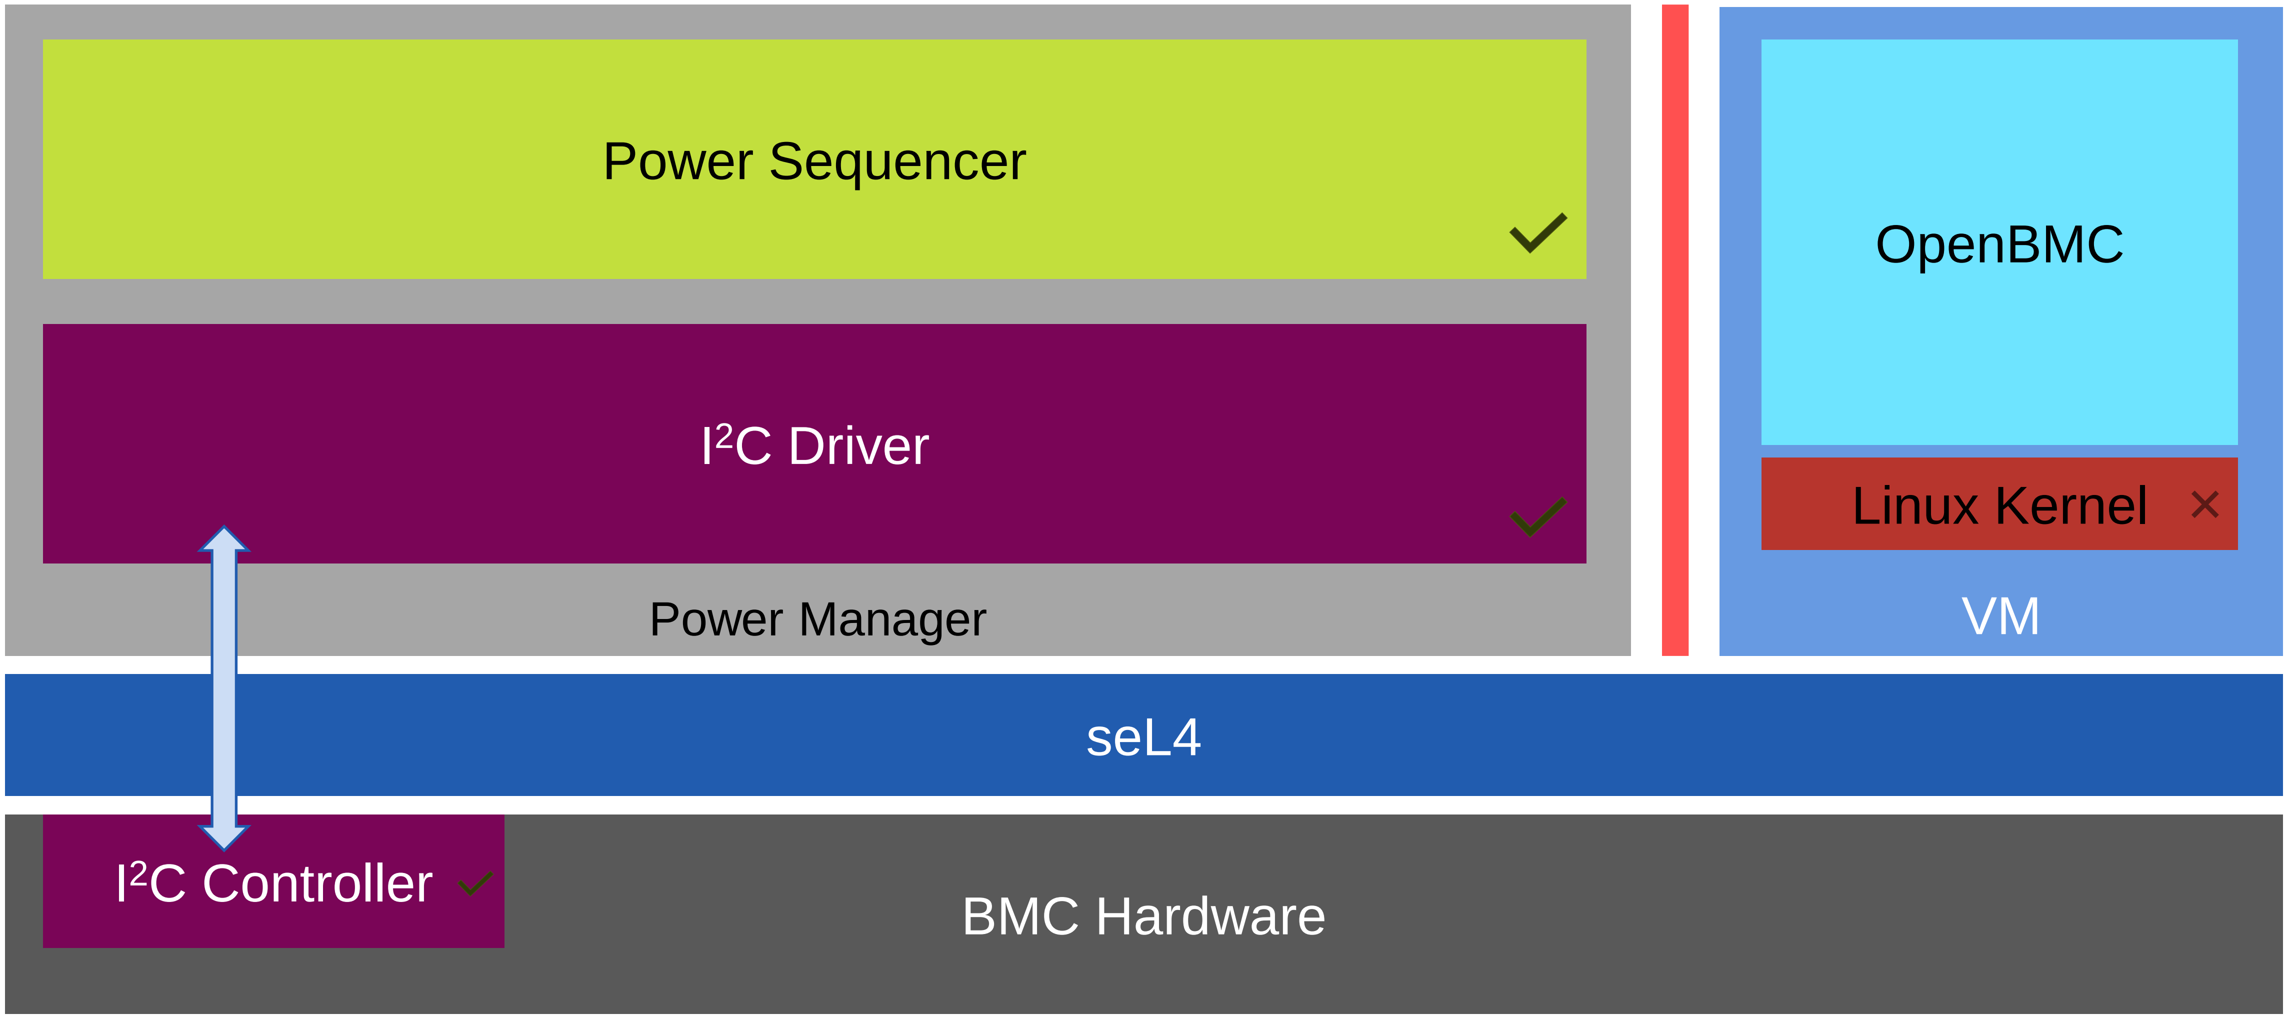

Beyond Power Management

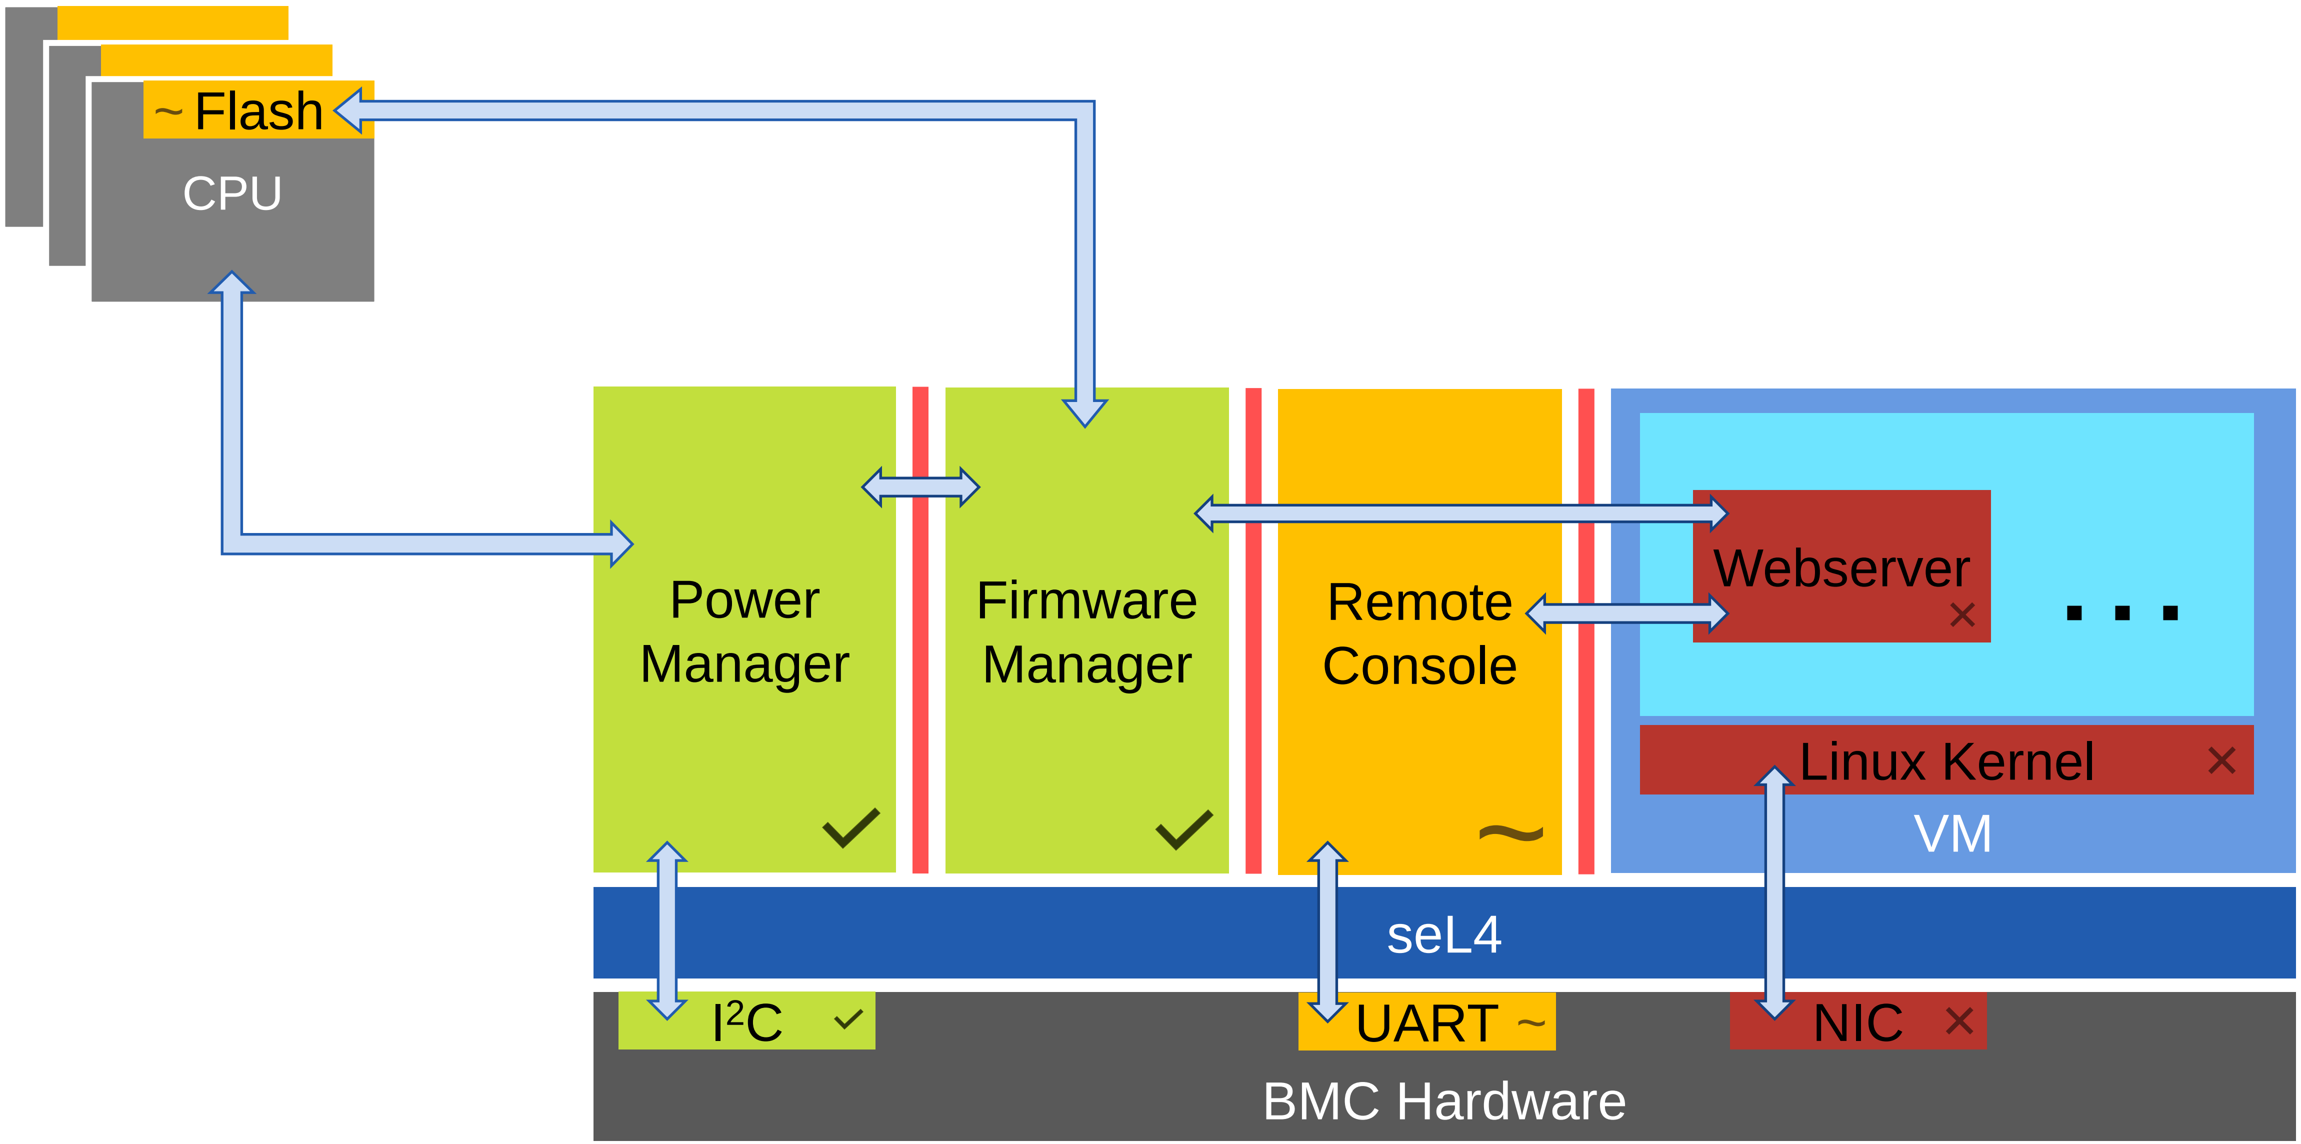

Trustworthy BMCs are only the beginning…

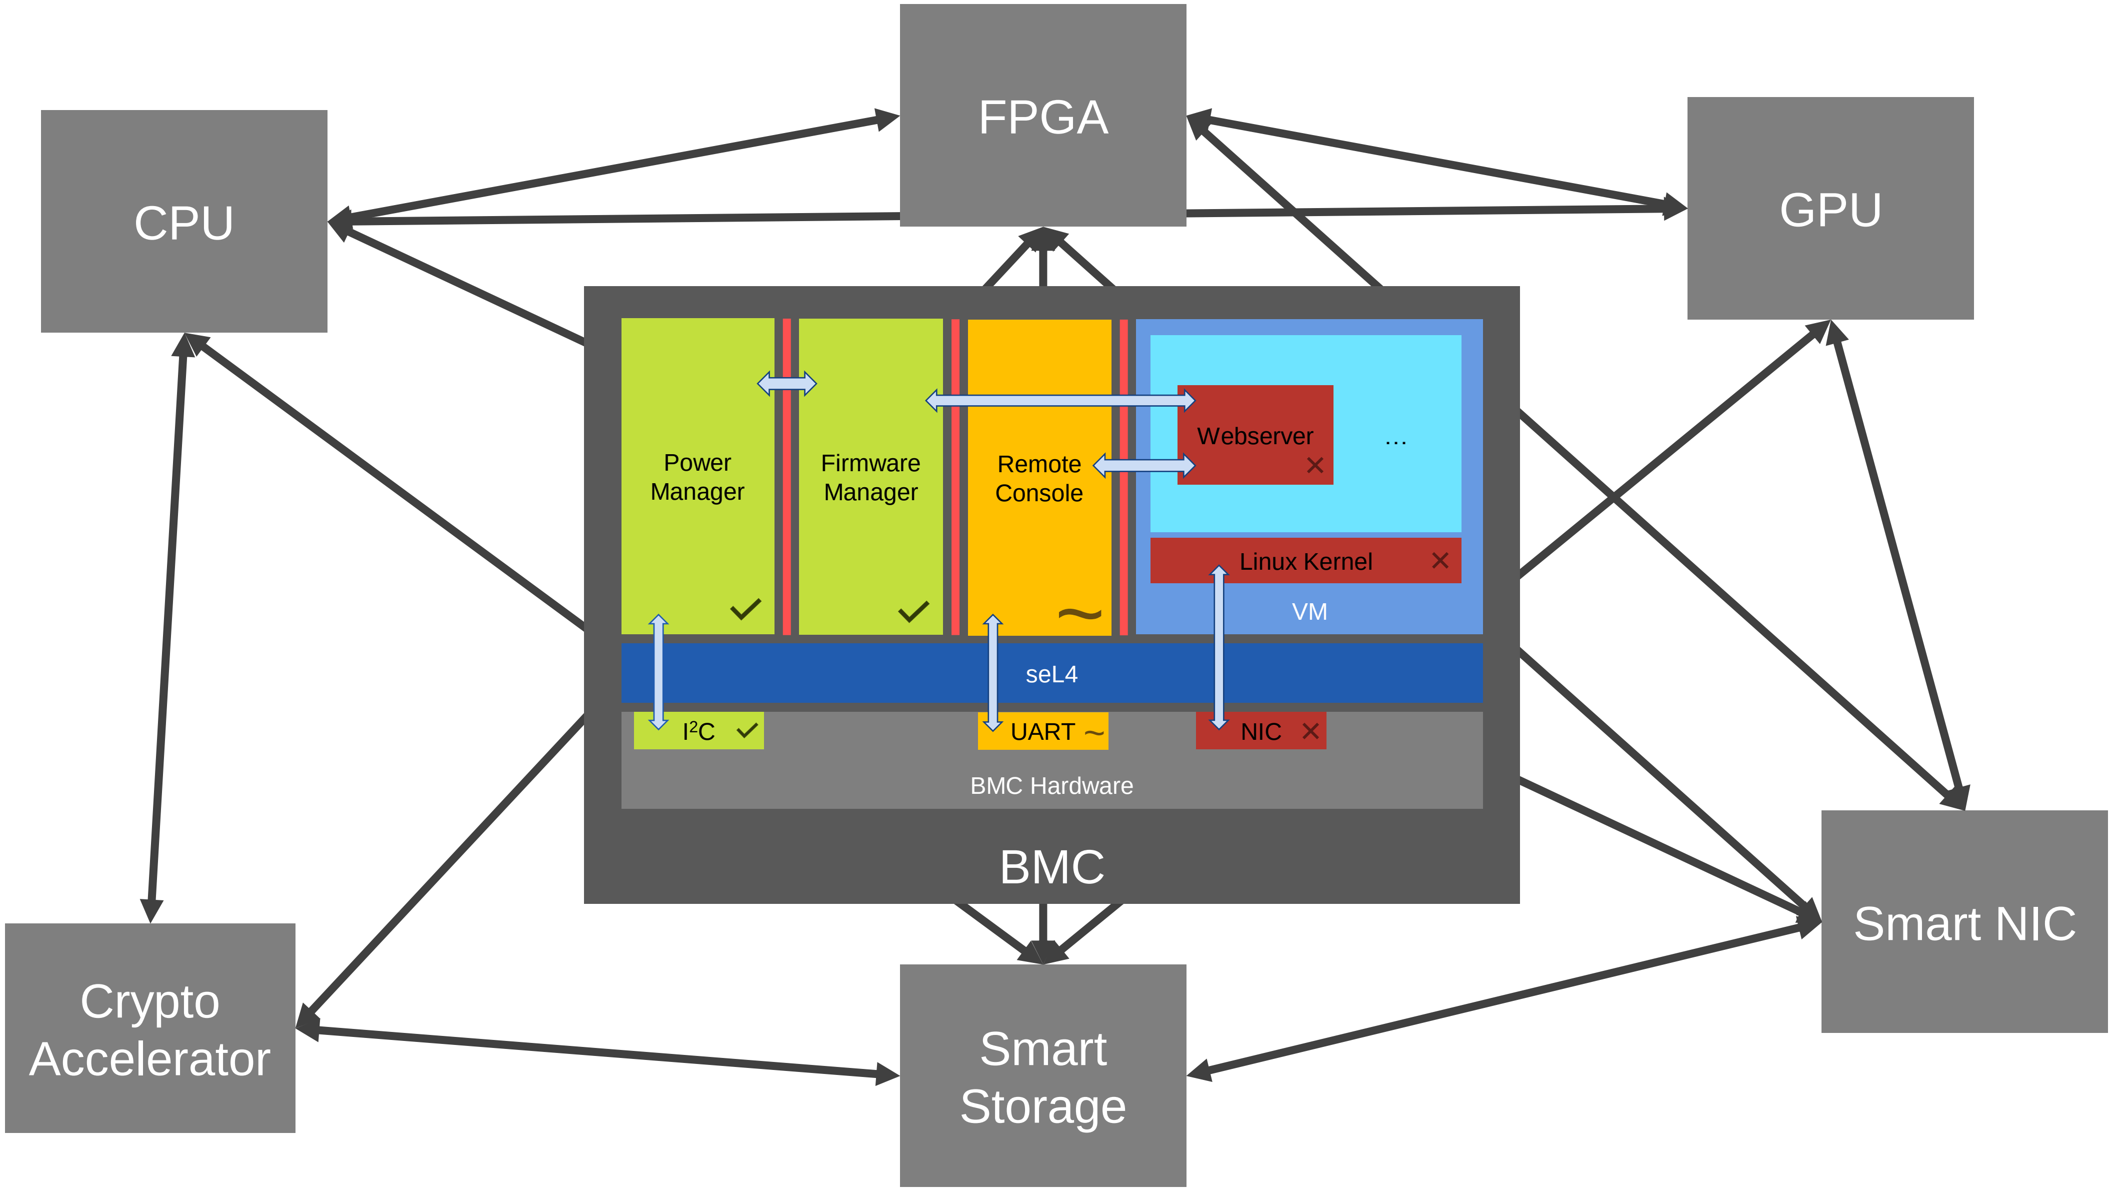

The Enzian Research Computer

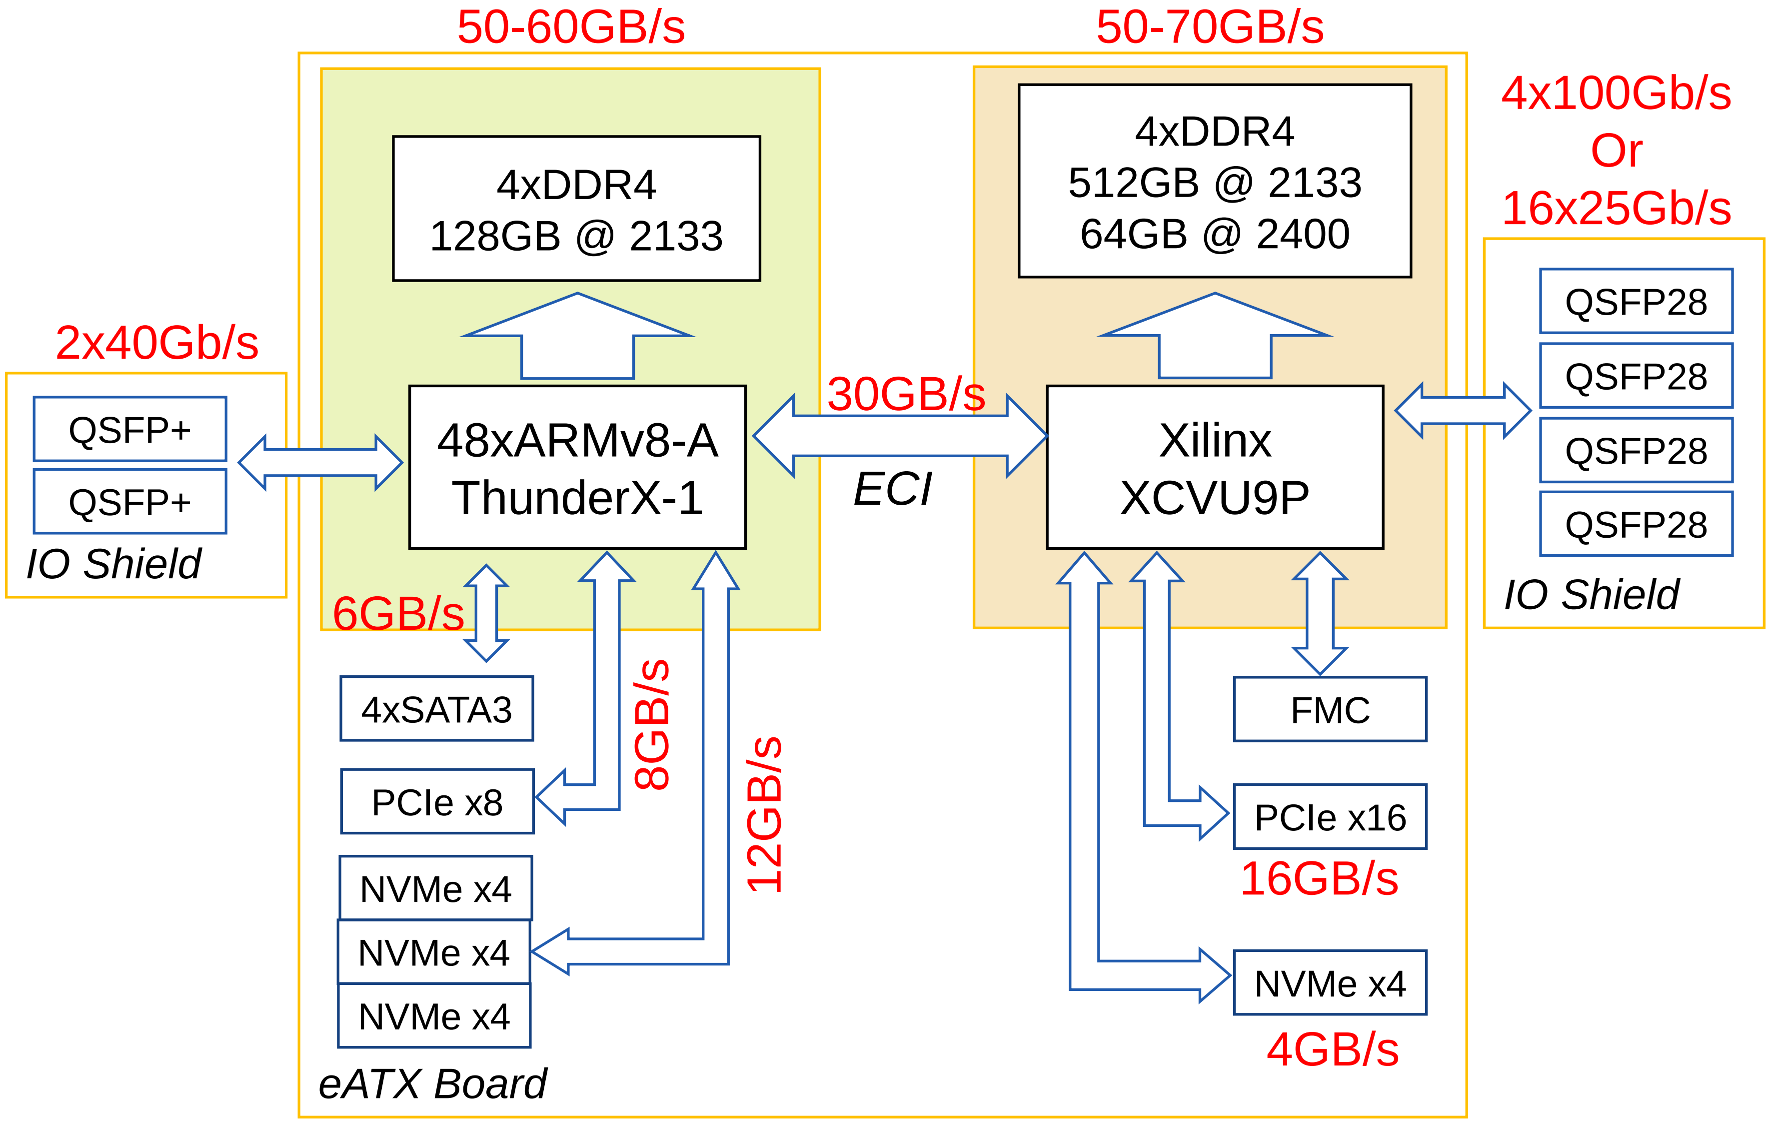

https://enzian.systems

Ongoing Work – Hardware Model Applications

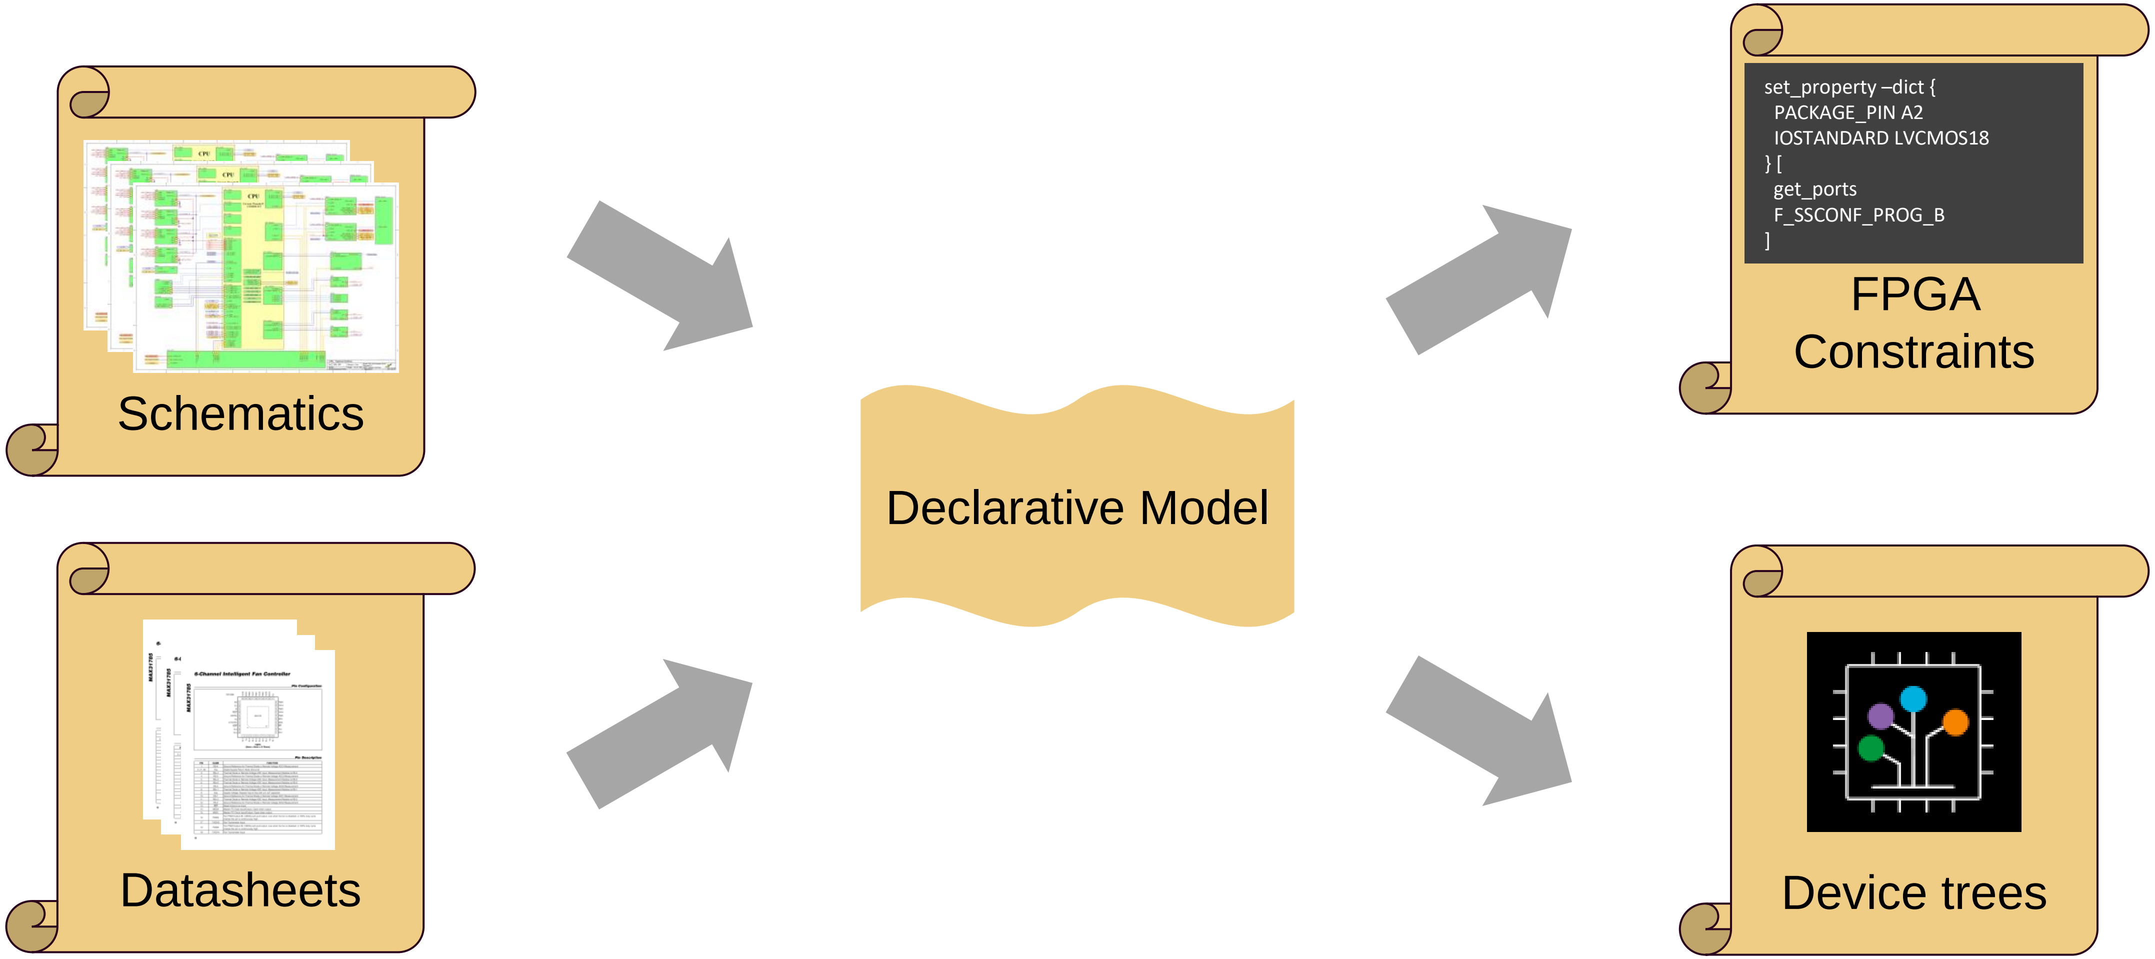

Ongoing Work – Hardware Model Applications

ENZIAN Systems@ETH Zürich DINFK

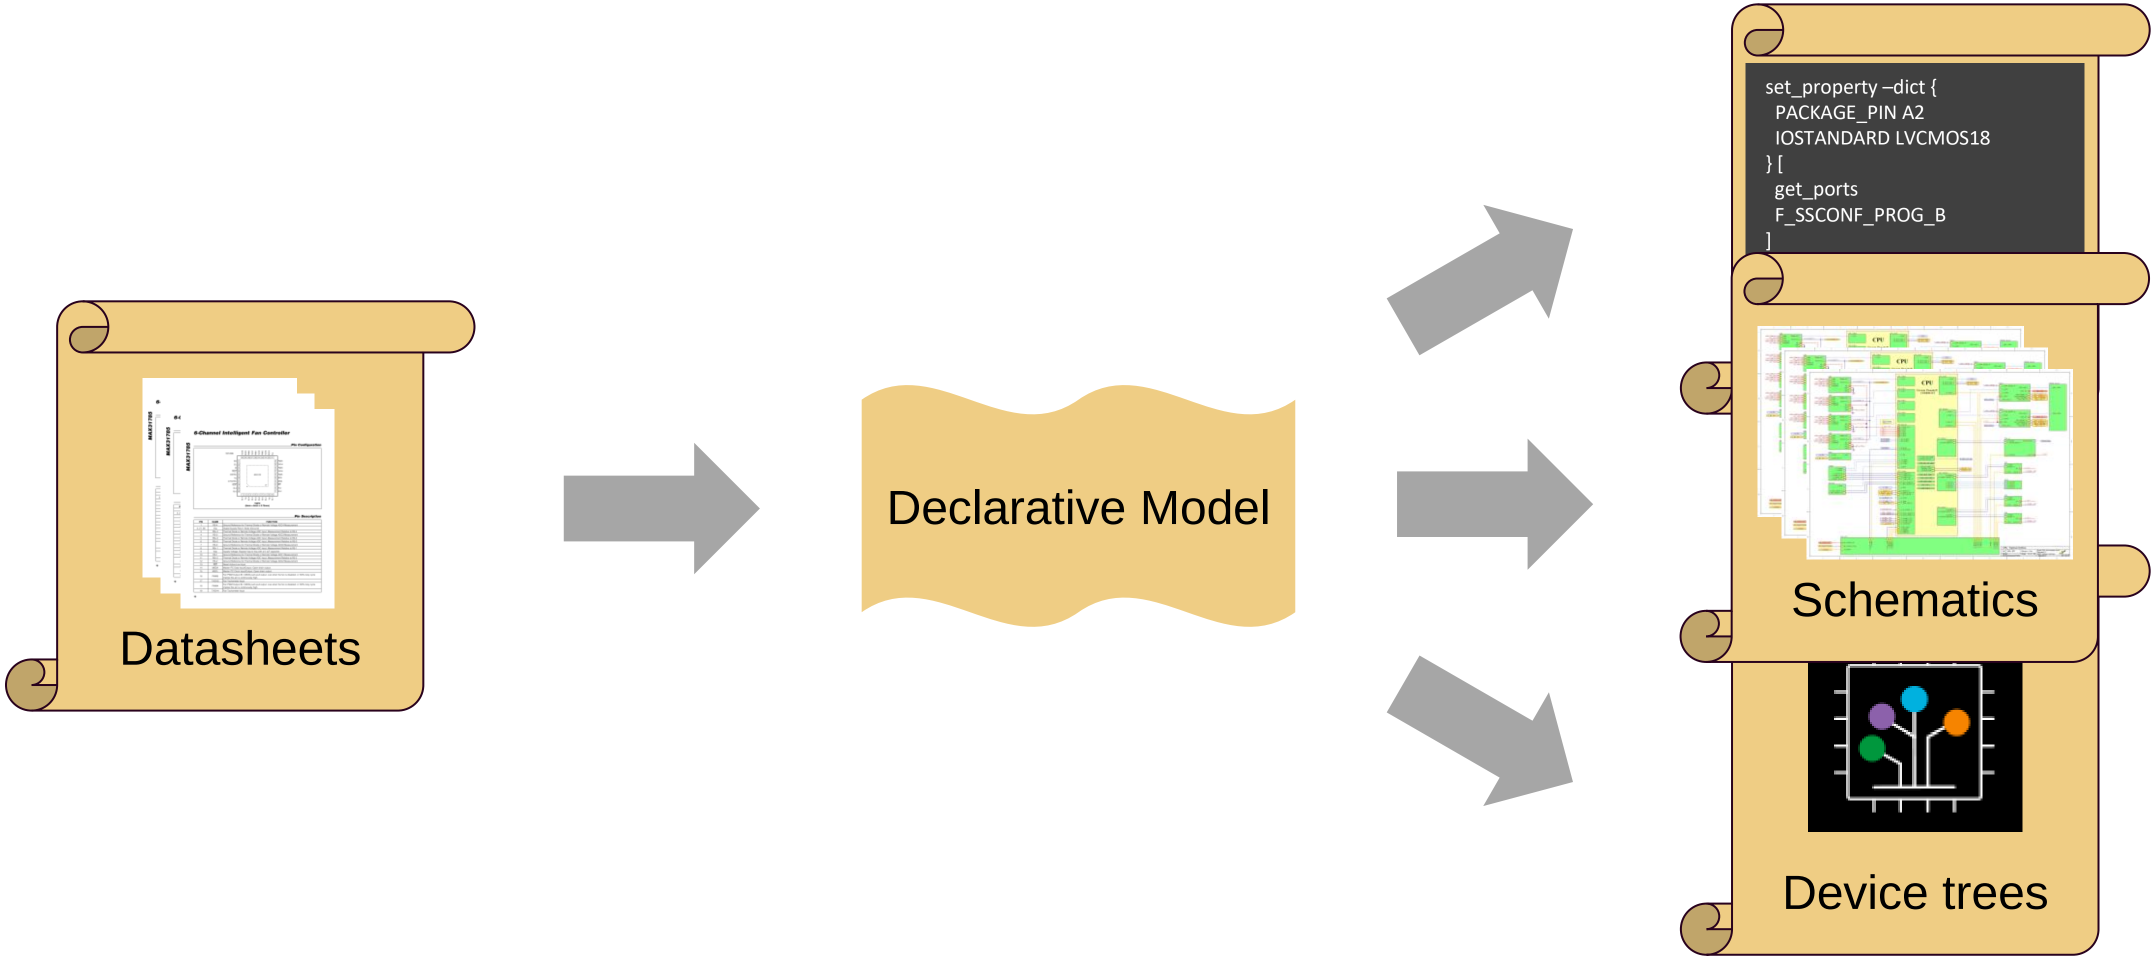

Modelling the Topology: Directed Graph

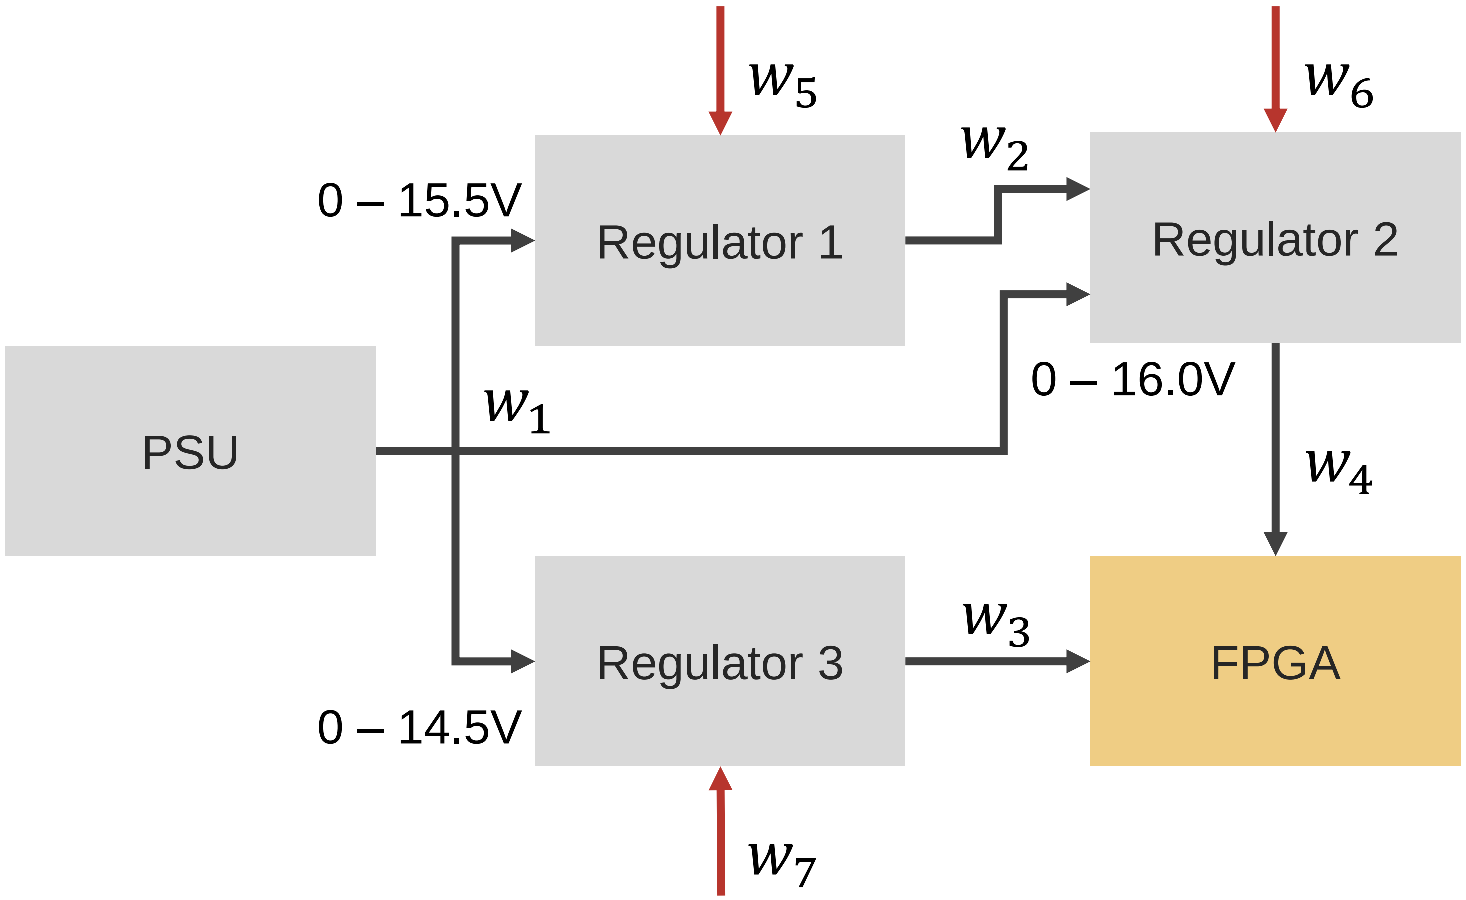

$$0 \leq w_1 \leq 15.5$$
$$0 \leq w_1 \leq 16.0$$
$$0 \leq w_1 \leq 14.5$$

Modelling the Components: State Diagram

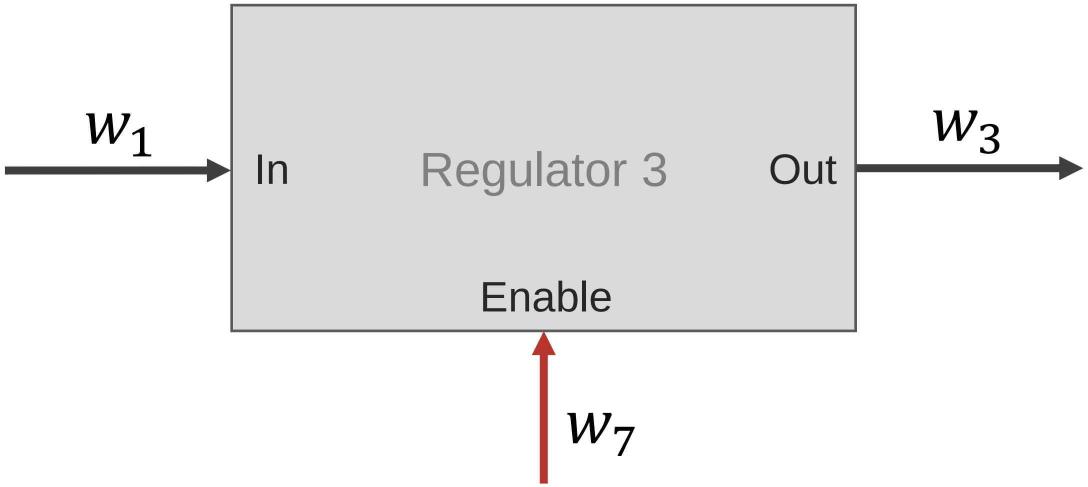

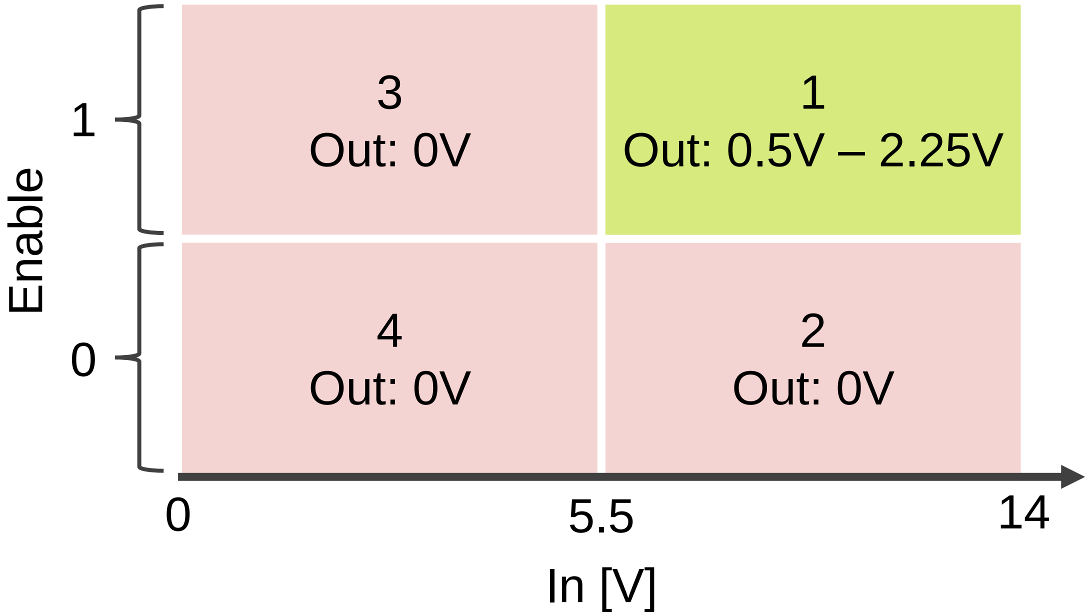

$$
\begin{aligned}
&(0.5 \leq w_3 \leq 2.25 \wedge 5.5 \leq w_1 \leq 14 \wedge w_7 = 1) \\
&\vee (w_3 = 0 \wedge 5.5 \leq w_1 \leq 14 \wedge w_7 = 0) \\
&\vee (w_3 = 0 \wedge w_1 = 0 \wedge w_7 = 1) \\
&\vee (w_3 = 0 \wedge w_1 = 0 \wedge w_7 = 0)
\end{aligned}
$$

Full Platform State: Propagate Constraints to the Root

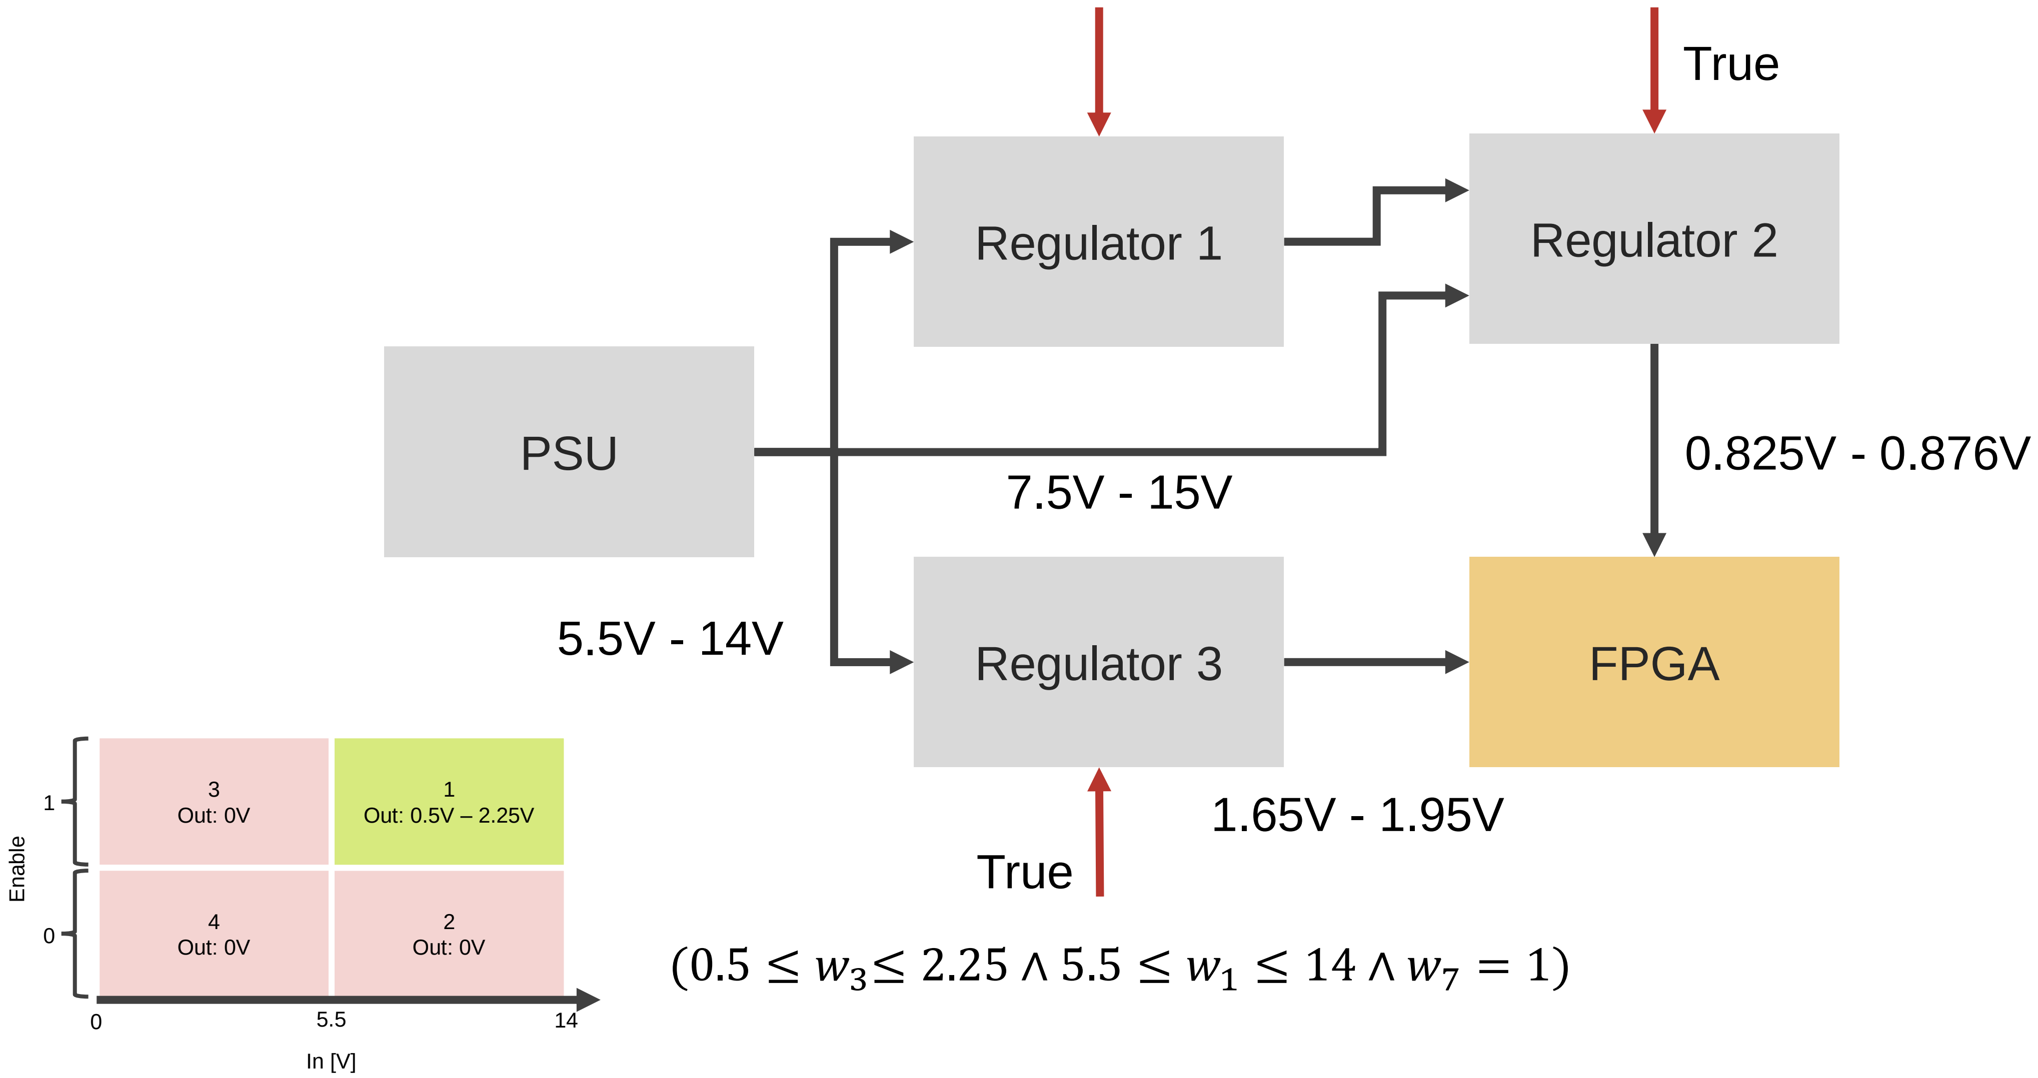

$$(0.5 \leq w_3 \leq 2.25 \land 5.5 \leq w_1 \leq 14 \land w_7 = 1)$$

Synthesizing Configurations from the Model is Practicable

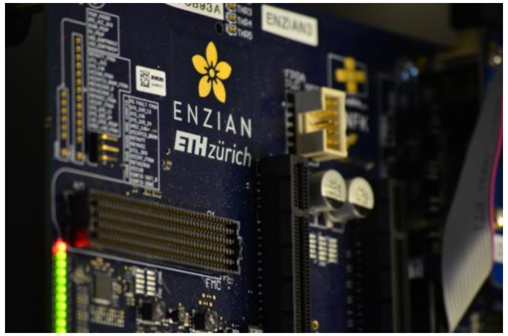

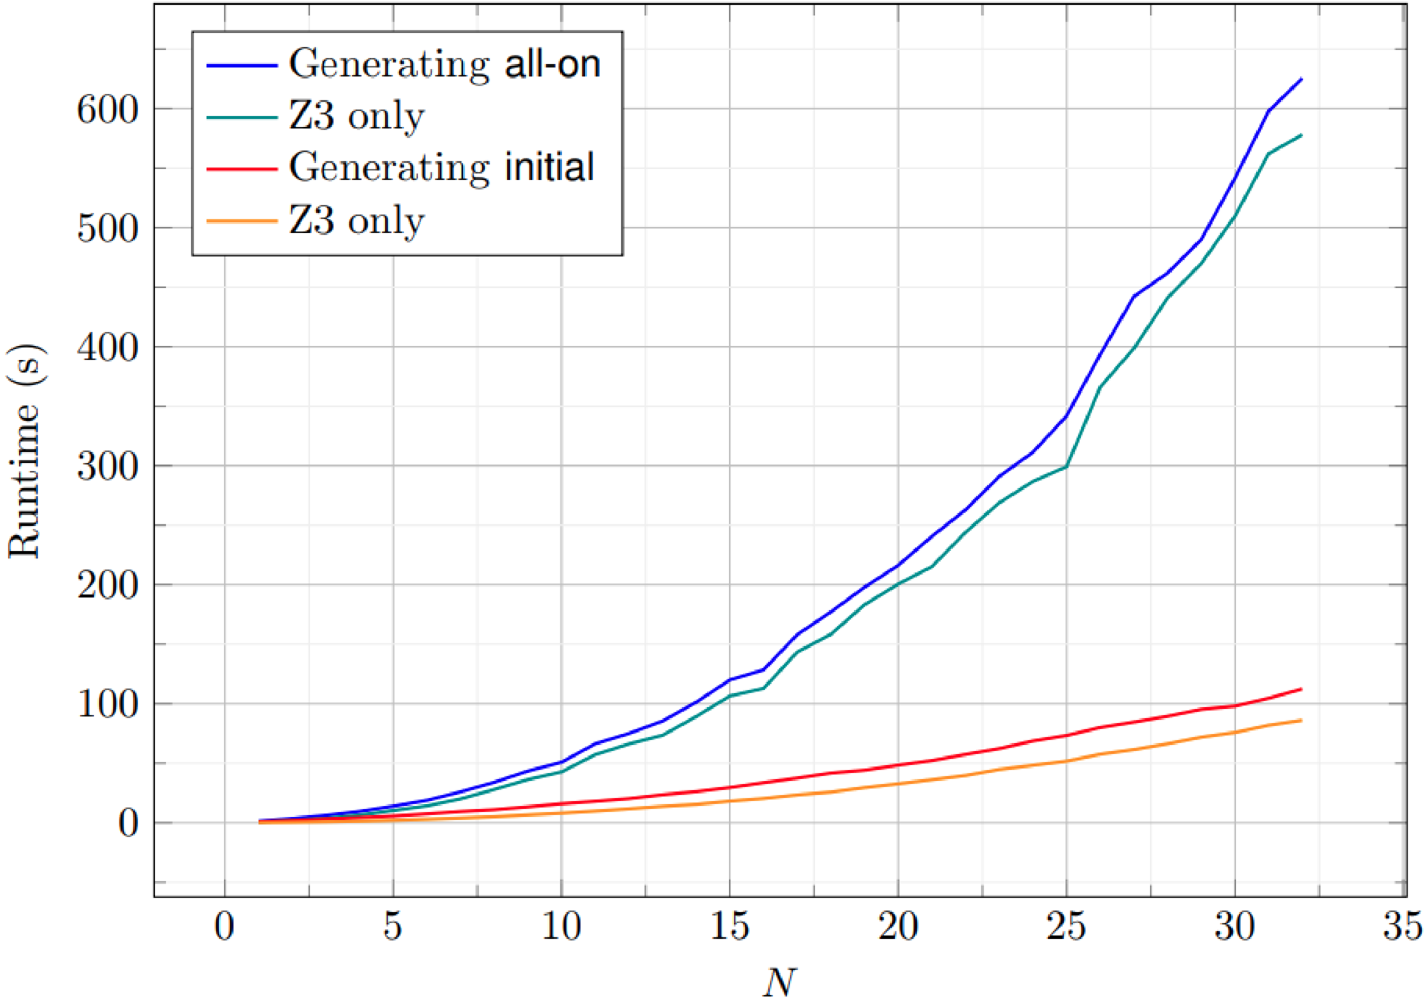

Moritz Knüsel. 2021. Optimizing Declarative Power Sequencing. Master’s Thesis, ETH Zürich

Our work so far

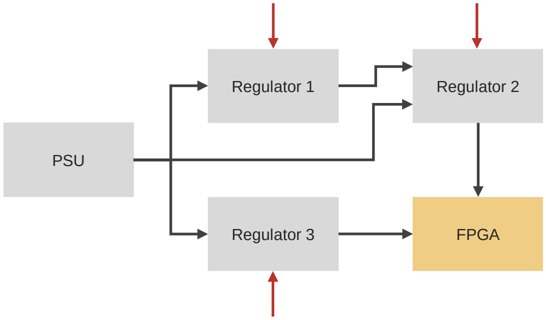

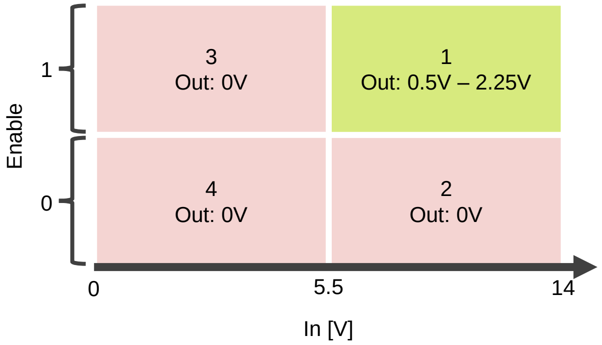

CSP Solver

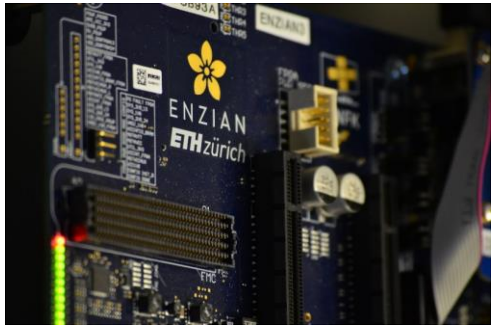

- More confidence in power sequences
- Better automation & maintainability
- Basis for rigorous specification

Sequence Constraints

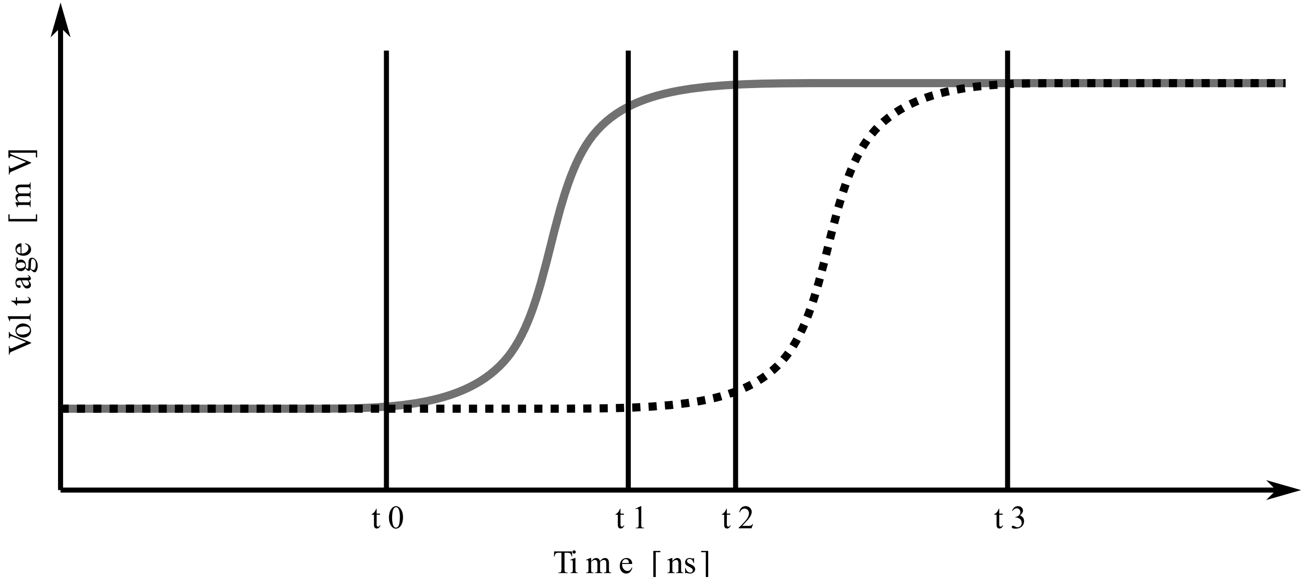

Initiate event:
Action that triggers change

Complete event:
Measurement that confirms change

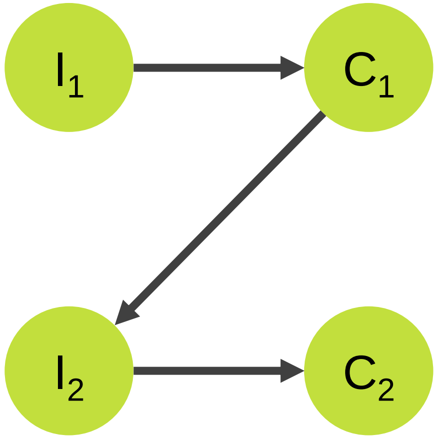

Sequence Constraints

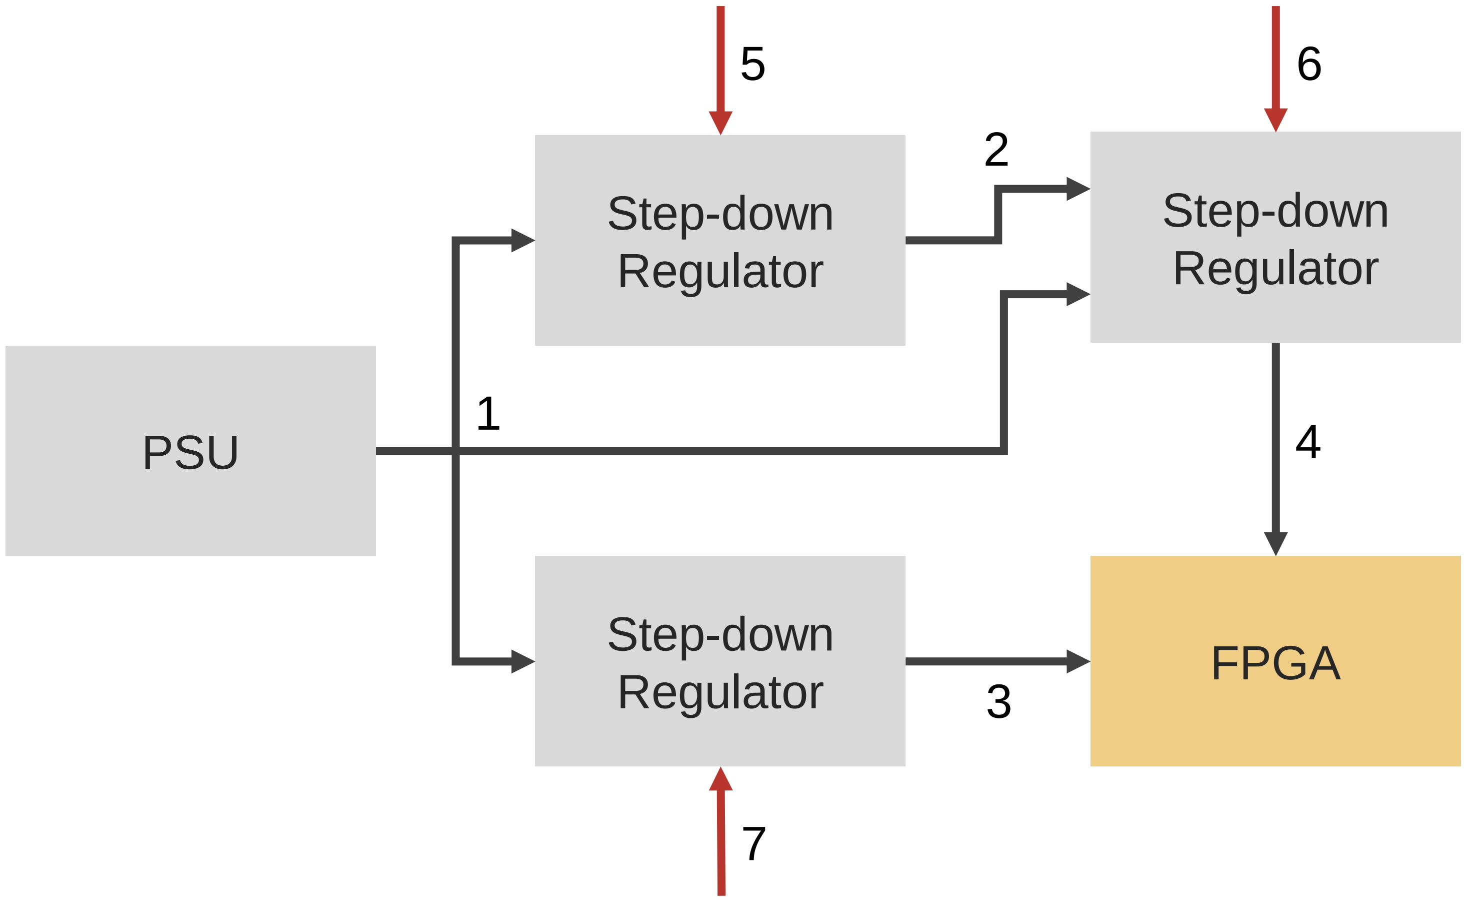

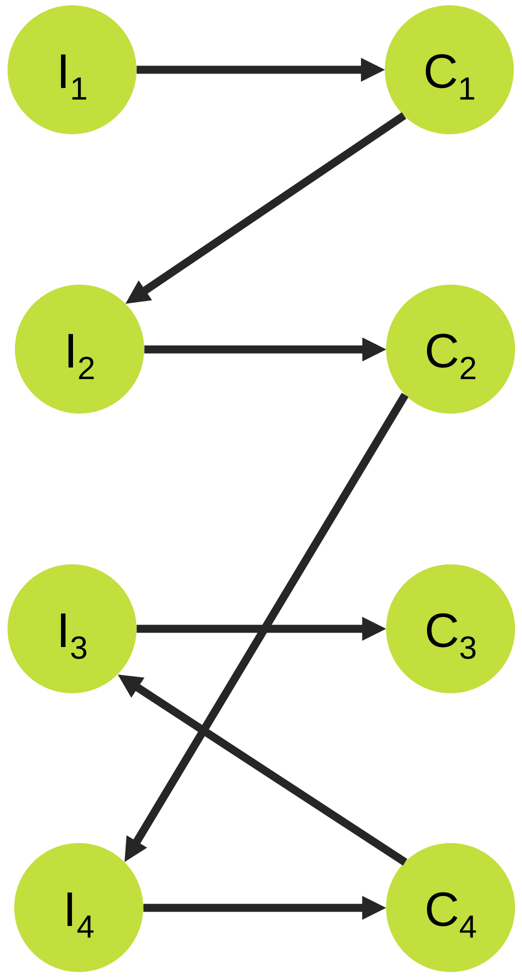

Power Distribution networks became complex

Power Distribution networks became complex

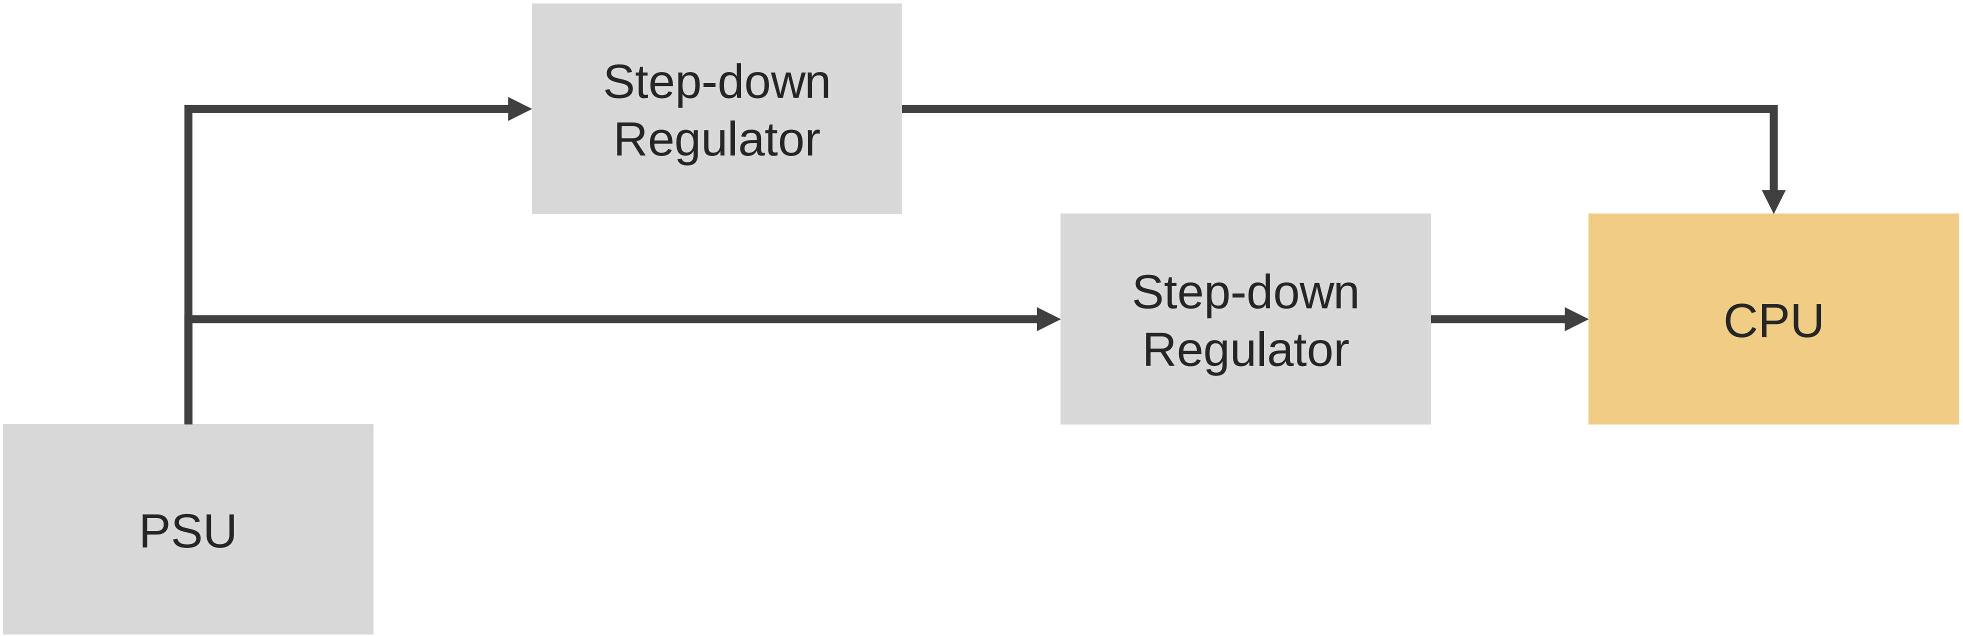

Power Distribution networks became complex

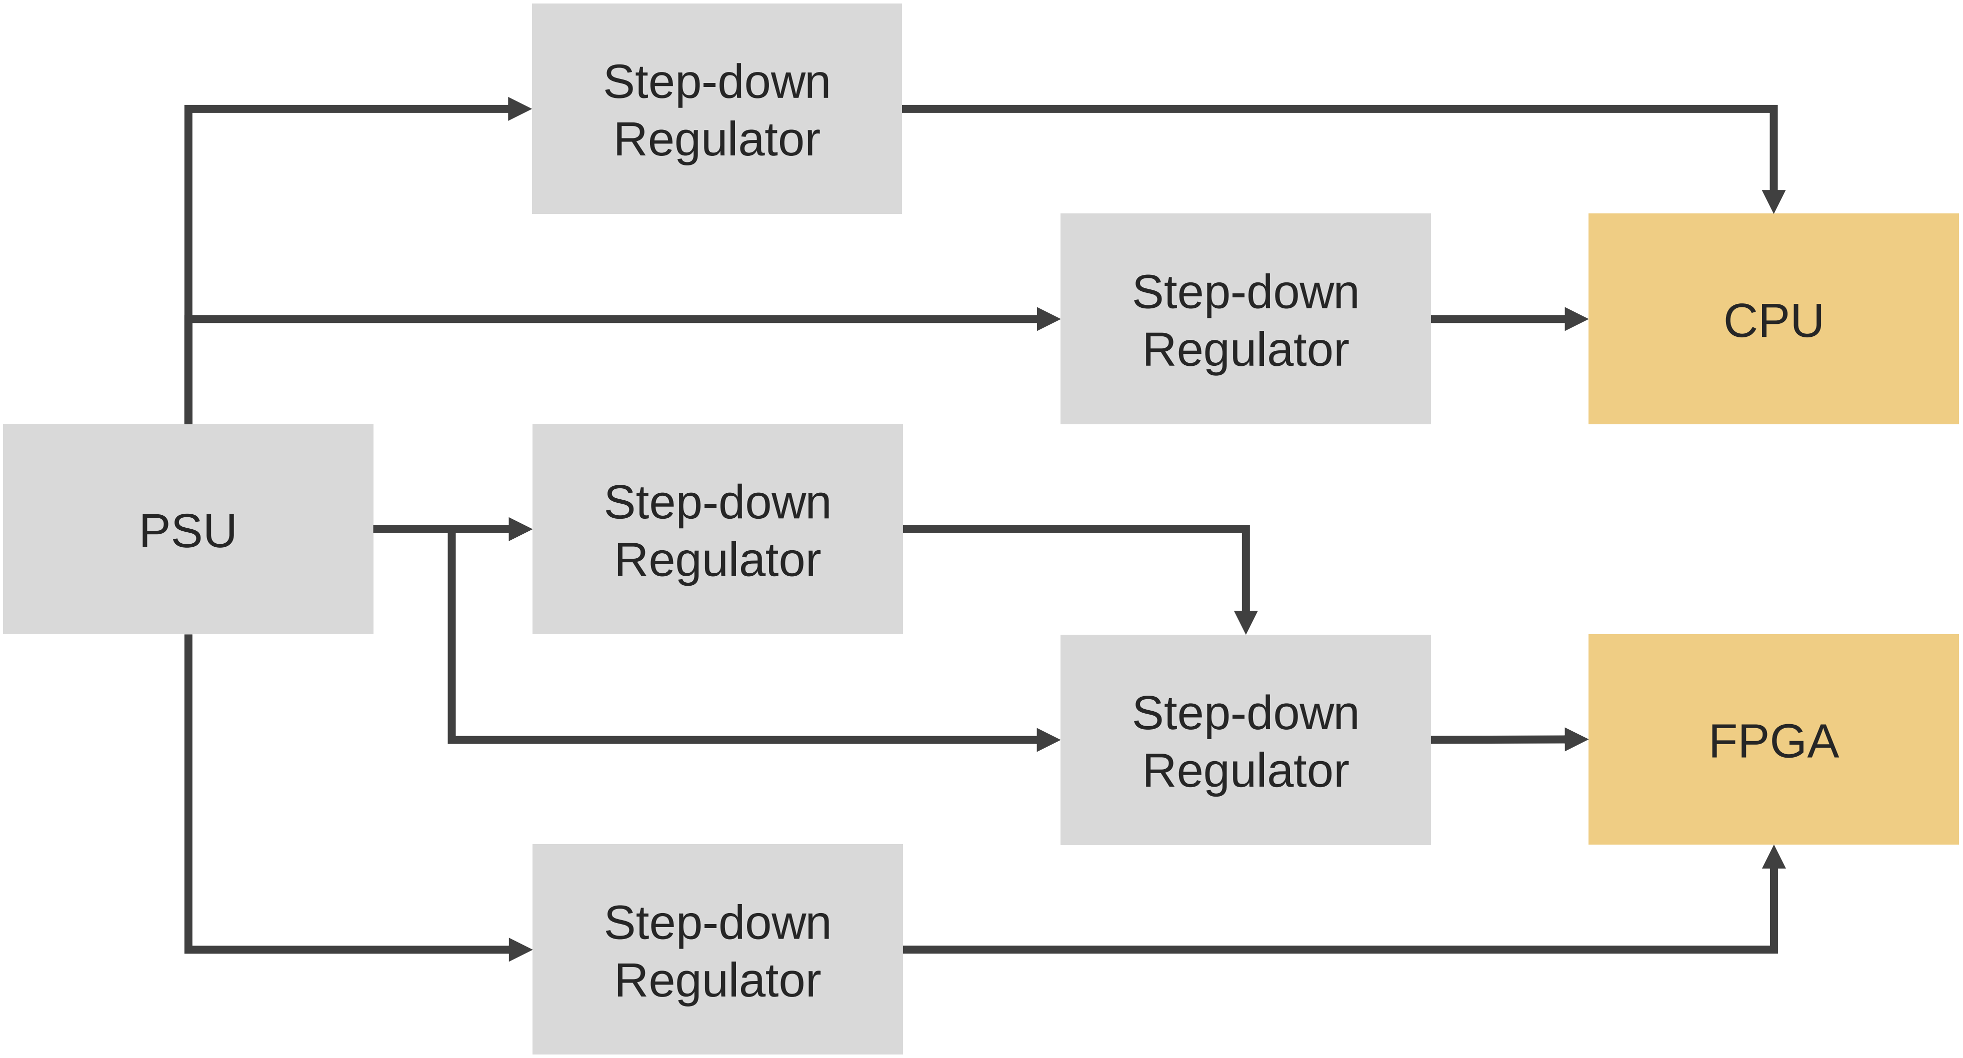

Maximum Ratings

ENZIAN Systems@ETH Zürich DINFK

Table 1: **Absolute Maximum Ratings** *(cont'd)*

Symbol	Description[1]	Min	Max	Units
V_{CCBRAM}	Supply voltage for the block RAM memories	–0.500	1.000	V
V_{CCO}	Output drivers supply voltage for HD I/O banks (VU19P and VU23P only)	–0.500	3.400	V
	Output drivers supply voltage for HP I/O banks	–0.500	2.000	V
V_{CCAUX_IO}[3]	Auxiliary supply voltage for the I/O banks	–0.500	2.000	V
V_{REF}	Input reference voltage	–0.500	2.000	V
V_{IN}[4, 5, 6]	I/O input voltage for HD I/O banks (VU19P and VU23P only)	–0.550	V_{CCO} + 0.550	V
	I/O input voltage for HP I/O banks	–0.550	V_{CCO} + 0.550	V
V_{BATT}	Key memory battery backup supply	–0.500	2.000	V
I_{DC}	Available output current at the pad	–20	20	mA
I_{RMS}	Available RMS output current at the pad	–20	20	mA
High Bandwidth Memory (HBM)				
V_{CC_HBM}	Supply voltage for the high-bandwidth memory	–0.300	1.500	V
$V_{CC\_IO\_HBM}$	I/O supply voltage for the high-bandwidth memory	–0.300	1.500	V
V_{CCAUX_HBM}	Auxiliary supply voltage for the high-bandwidth memory	–0.300	3.000	V
GTY or GTM Transceiver[7]				
V_{CCINT_GT}	Digital supply voltage for select modules in the GTM transceivers	–0.500	1.000	V

Source: Xilinx, Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics

Sequencing Constraints

Power-On/Off Power Supply Sequencing

The recommended power-on sequence is V_{CCINT}, V_{CCINT_IO}/V_{CCBRAM}, V_{CCAUX}/V_{CCAUX_IO}, and V_{CCO} to achieve minimum current draw and ensure that the I/Os are 3-stated at power-on. The recommended power-off sequence is the reverse of the power-on sequence. If V_{CCINT} and V_{CCINT_IO}/V_{CCBRAM} have the same recommended voltage levels, they can be powered by the same supply and ramped simultaneously. V_{CCINT_IO} must be connected to V_{CCBRAM}. If V_{CCAUX}/V_{CCAUX_IO} and V_{CCO} have the same recommended voltage levels, they can be powered by the same supply and ramped simultaneously. V_{CCAUX} and V_{CCAUX_IO} must be connected together. V_{CCADC} and V_{REF} can be powered at any time and have no power-up sequencing requirements.

power-off sequence is the reverse of the power-on sequence to achieve minimum current draw. If these recommended sequences are not met, current drawn from $V_{MGTAVTT}$ can be higher than specifications during power-up and power-down.

Source: Xilinx, Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics

seL4: Full functional correctness proof & practical

DARPA HACMS Project

Source: https://trustworthy.systems/projects/TS/SMACCM/ulb.png

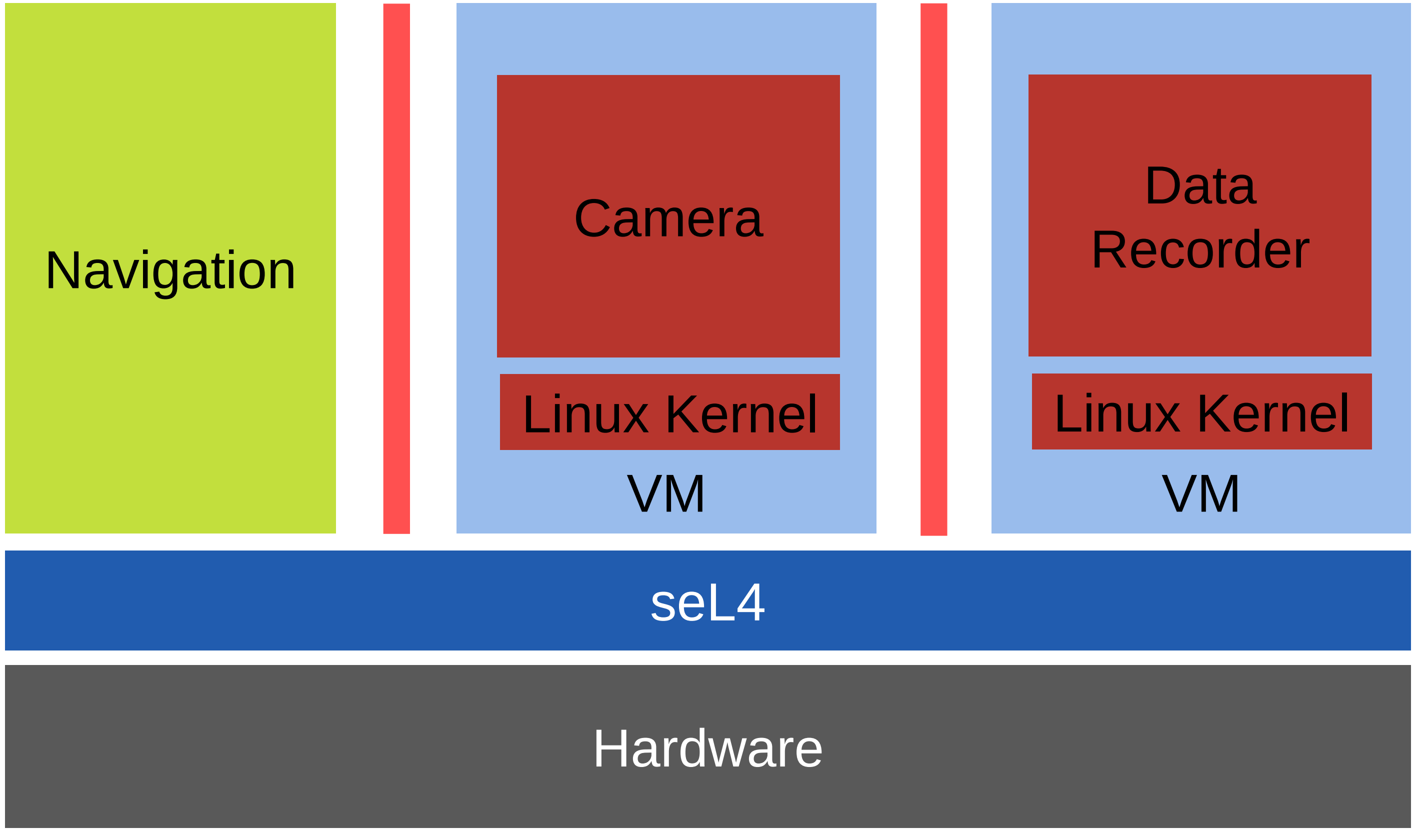

How to make correctness statement at system level?

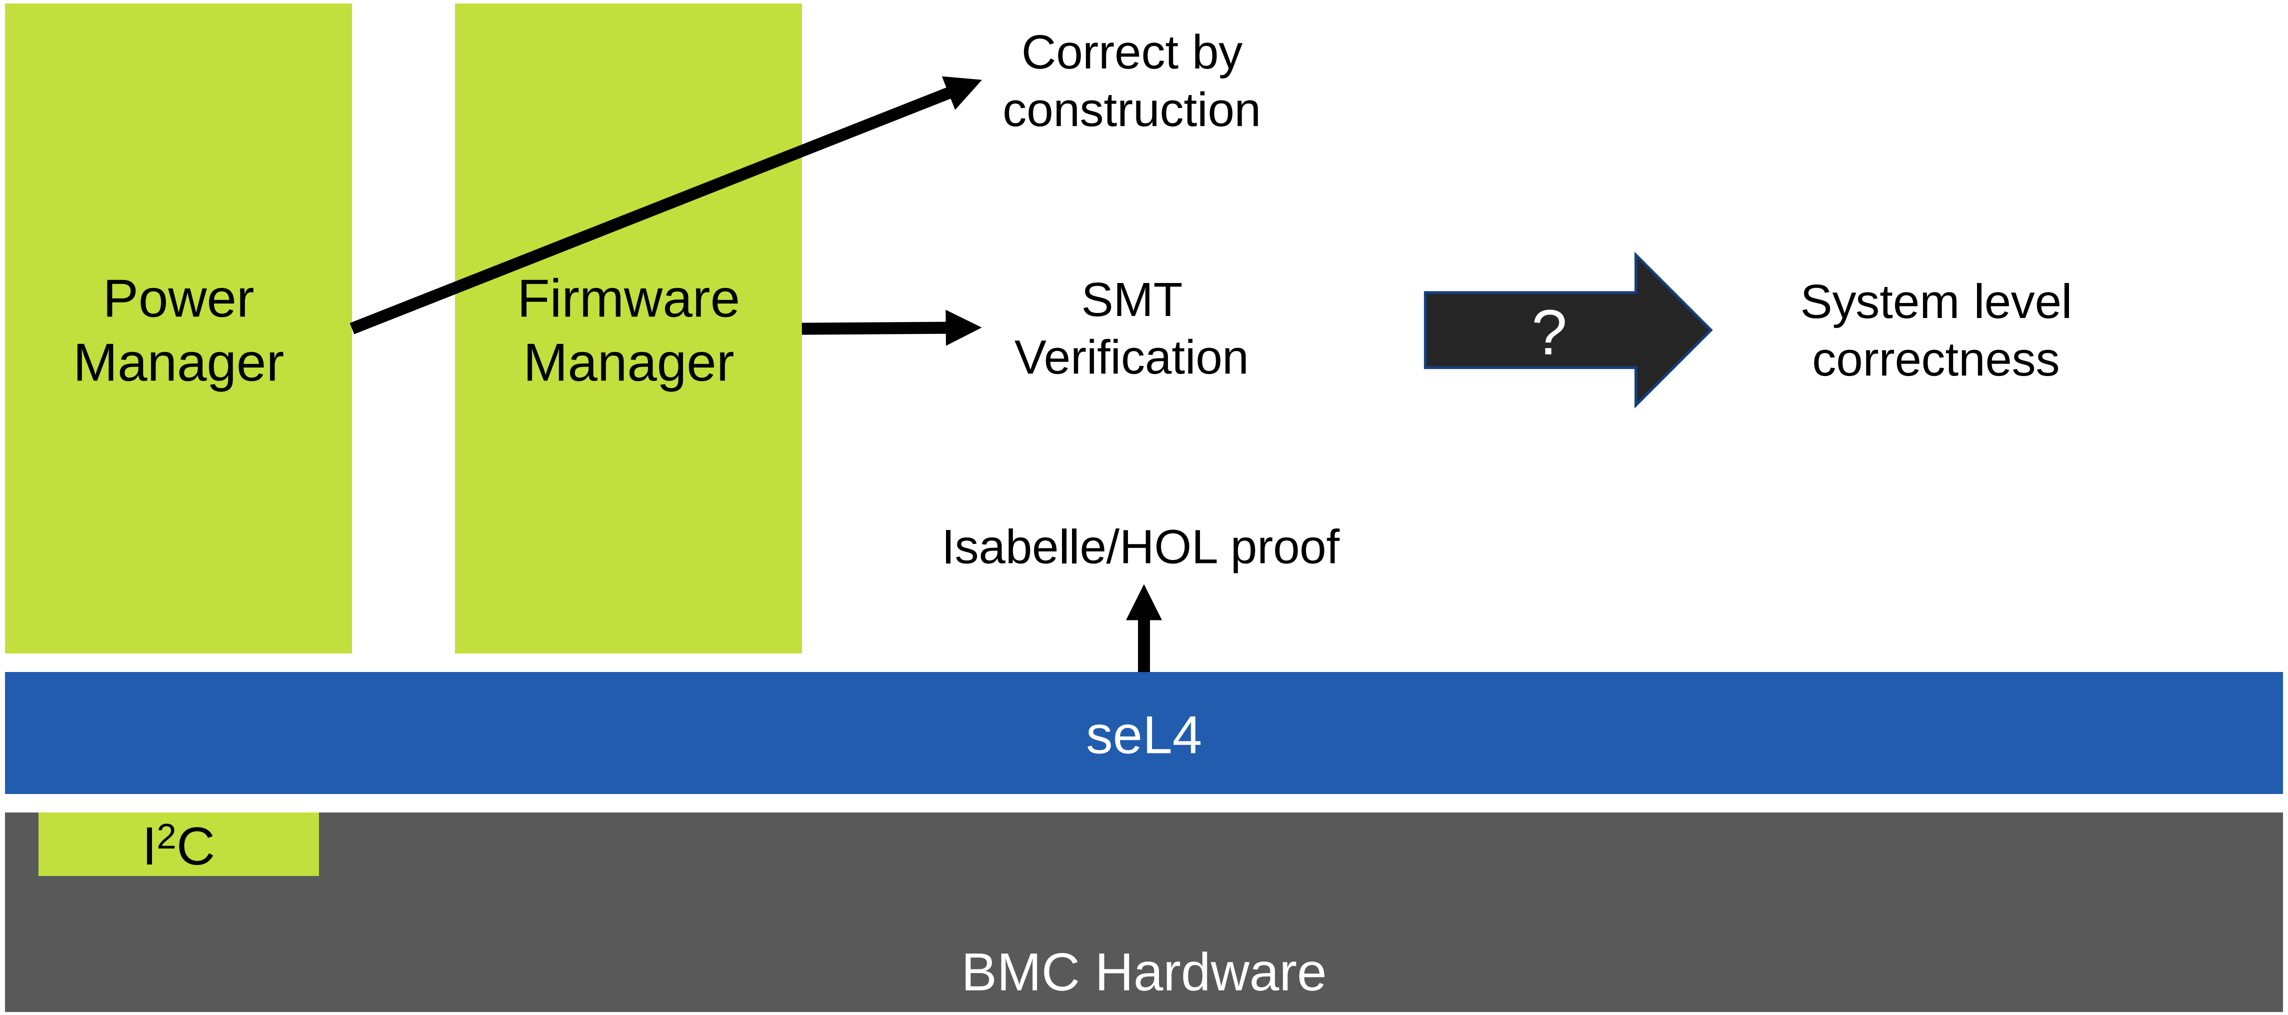

How to make correctness statements about hardware interaction?

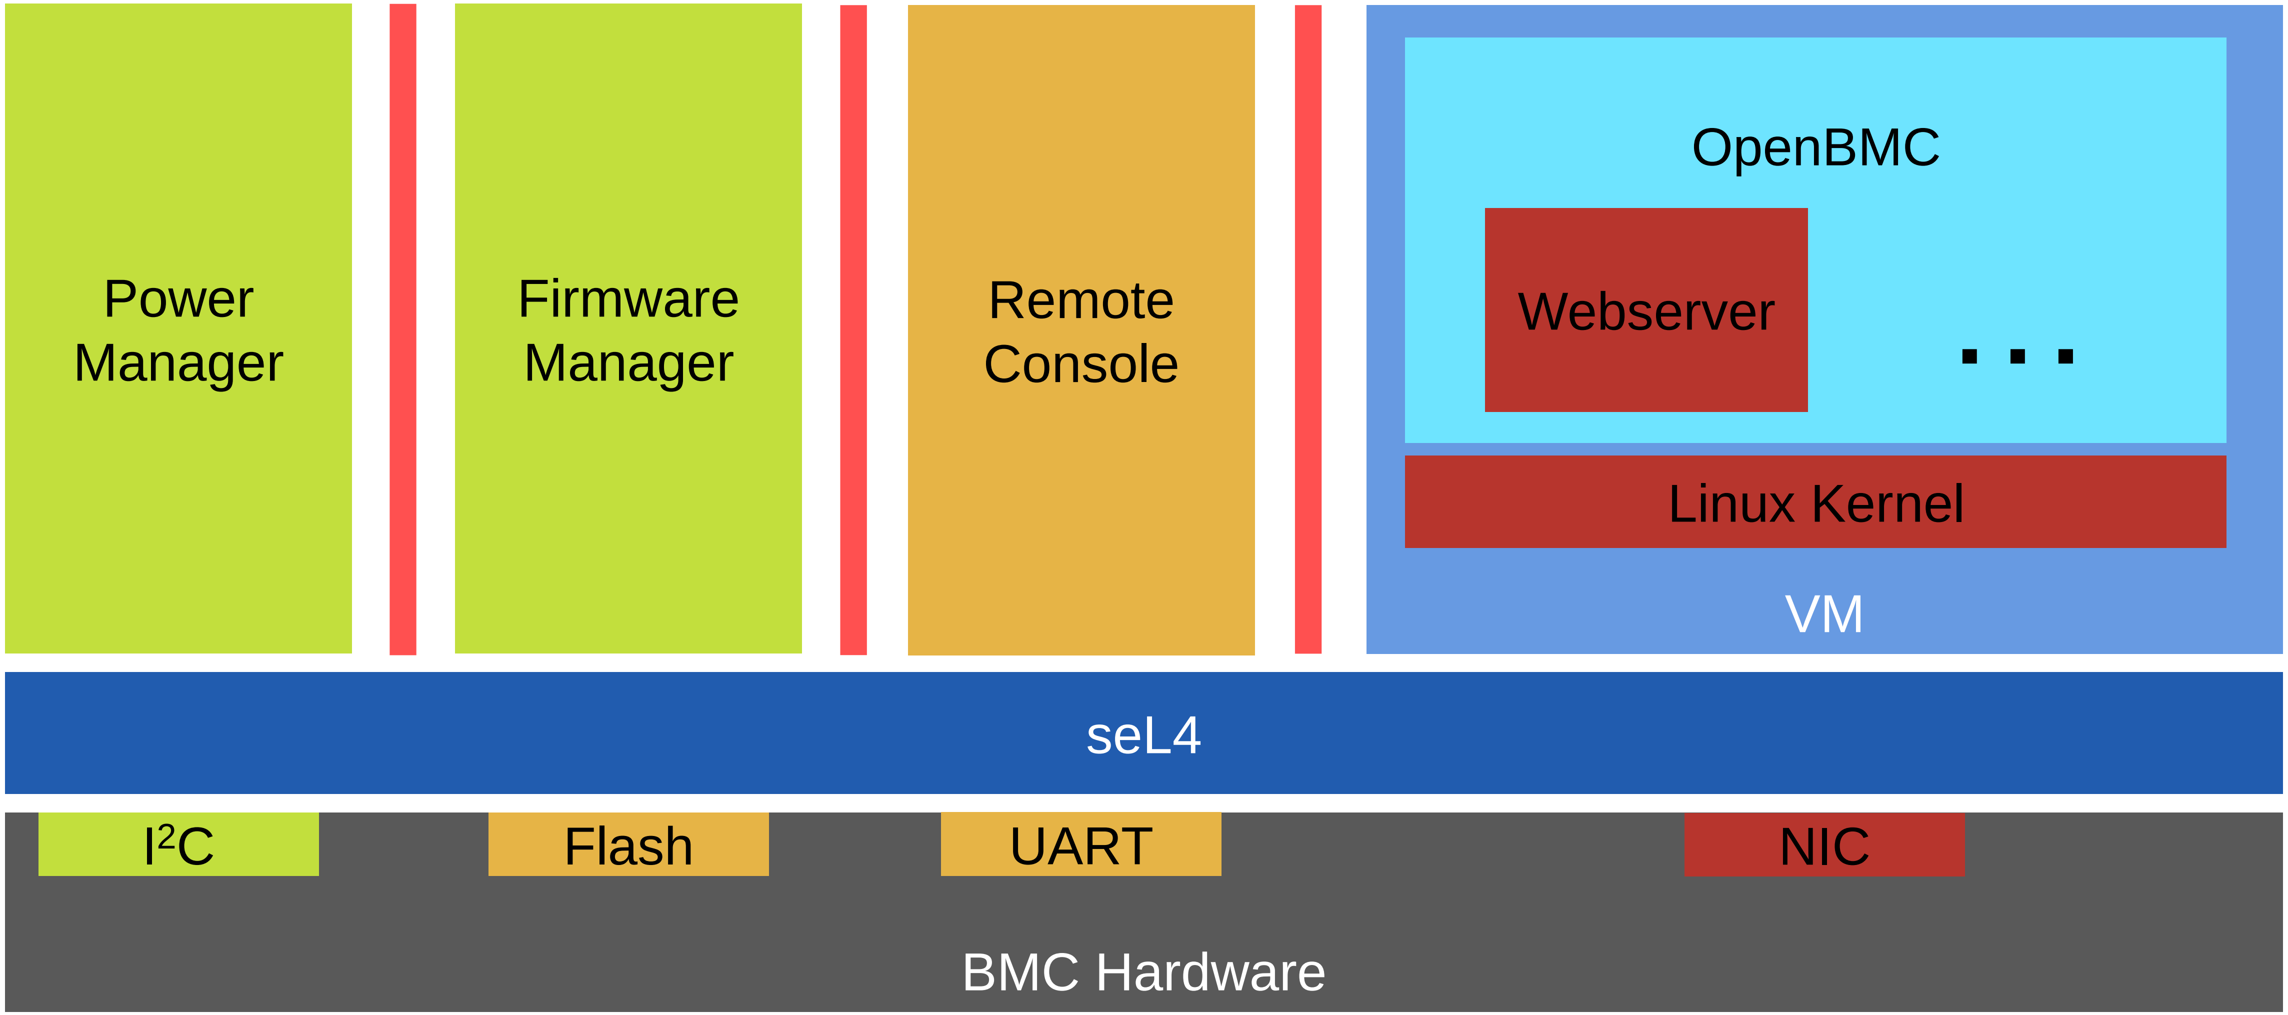

How do components of different trust levels communicate?

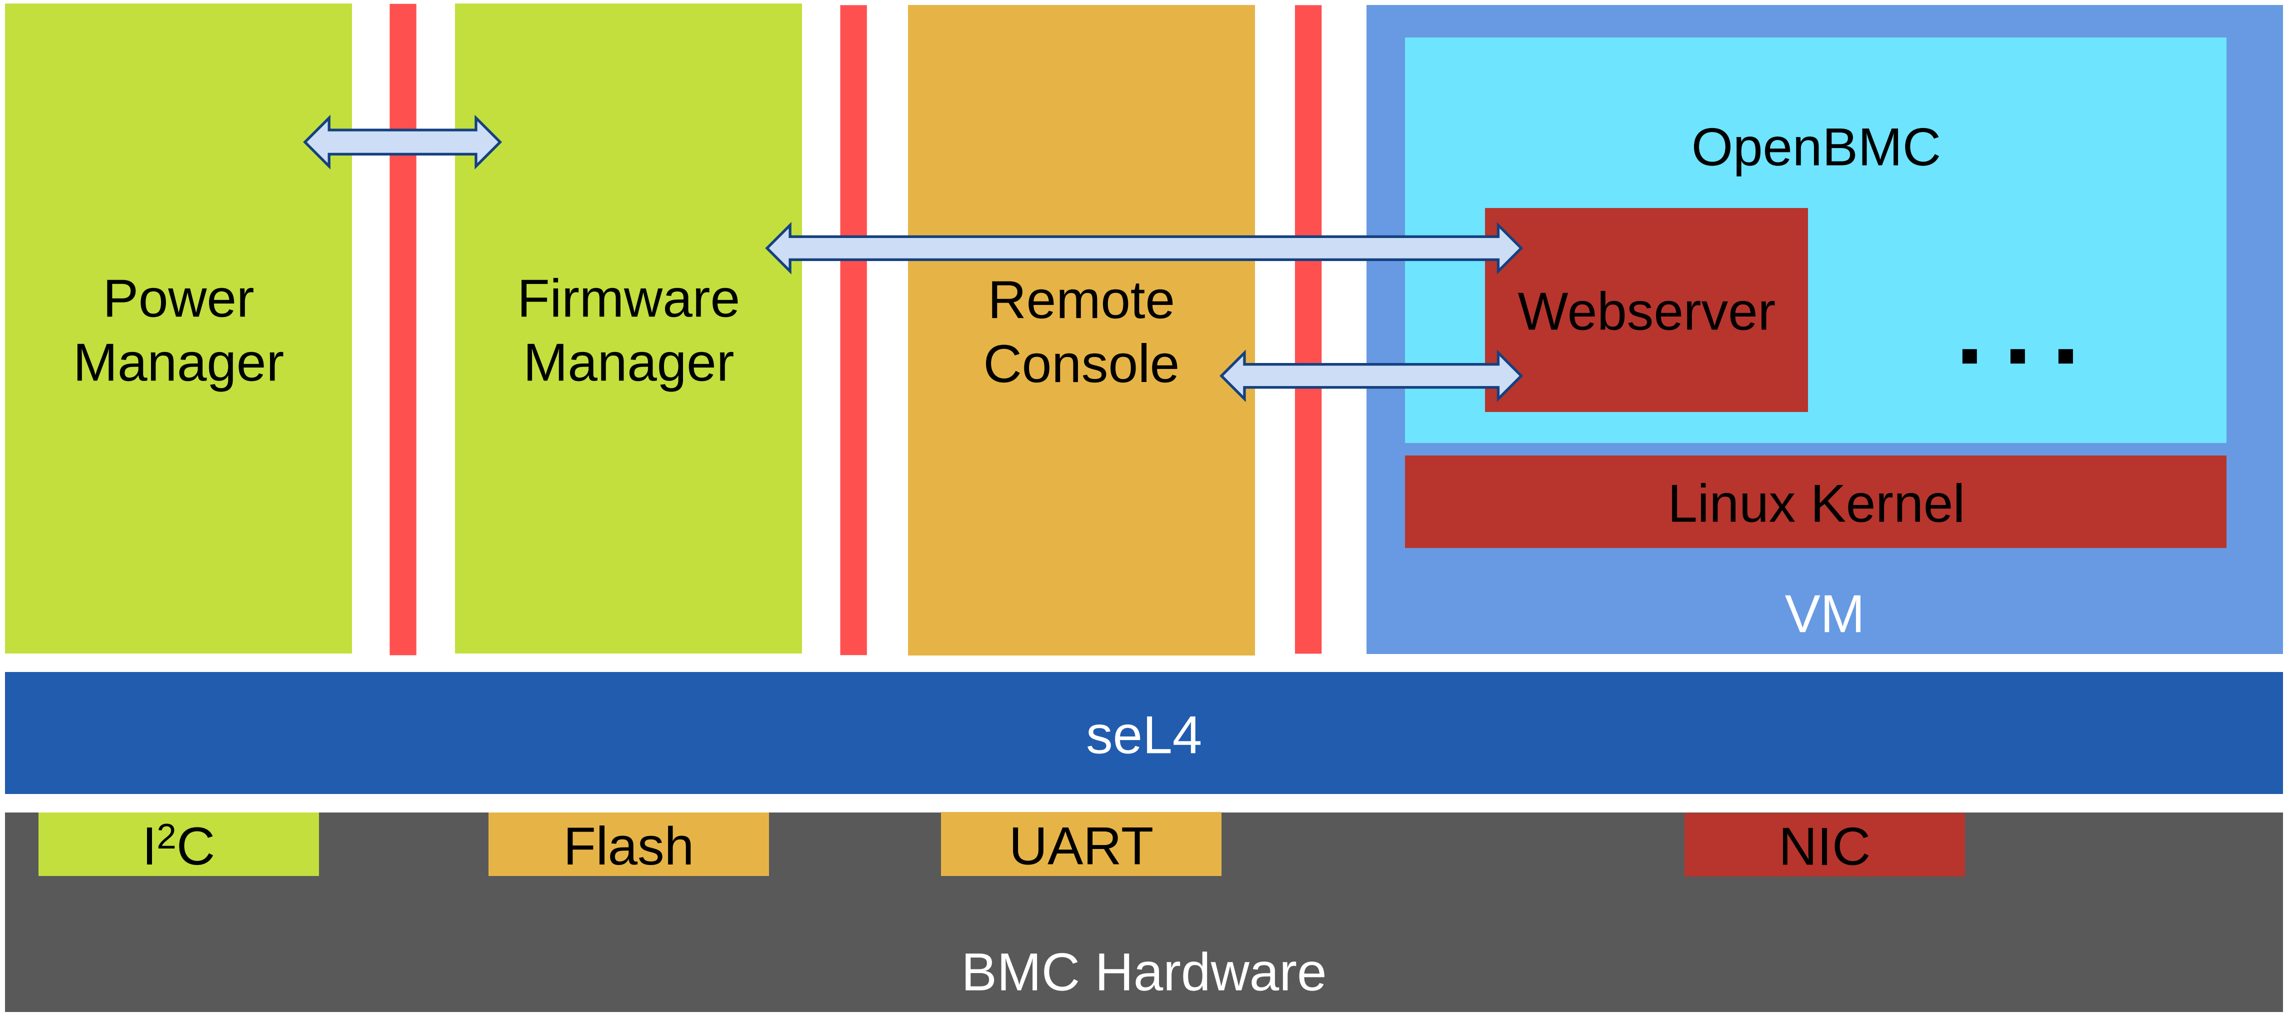

How do we componentize the system?

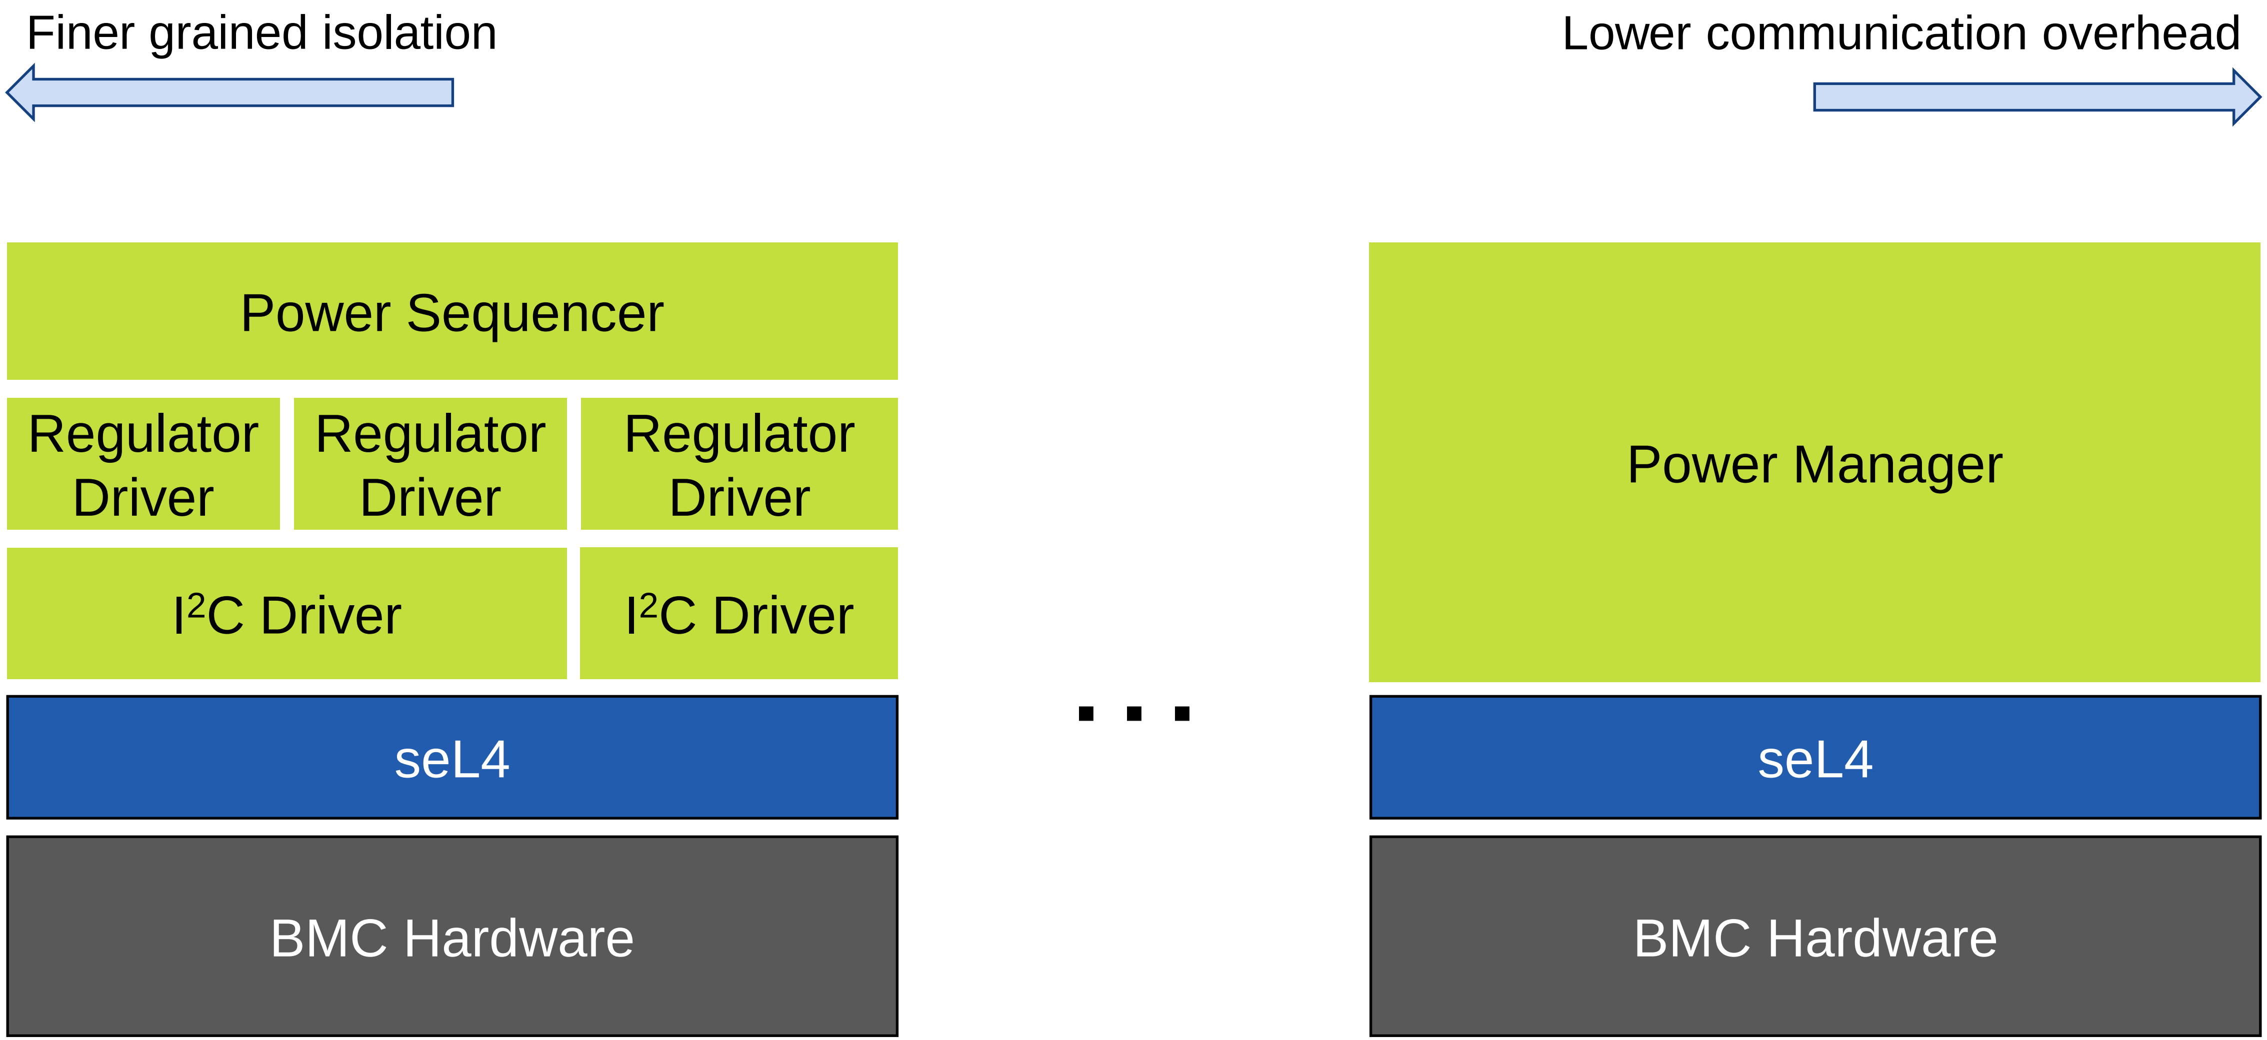

What tooling do we need to push the complexity of verifiable systems?