



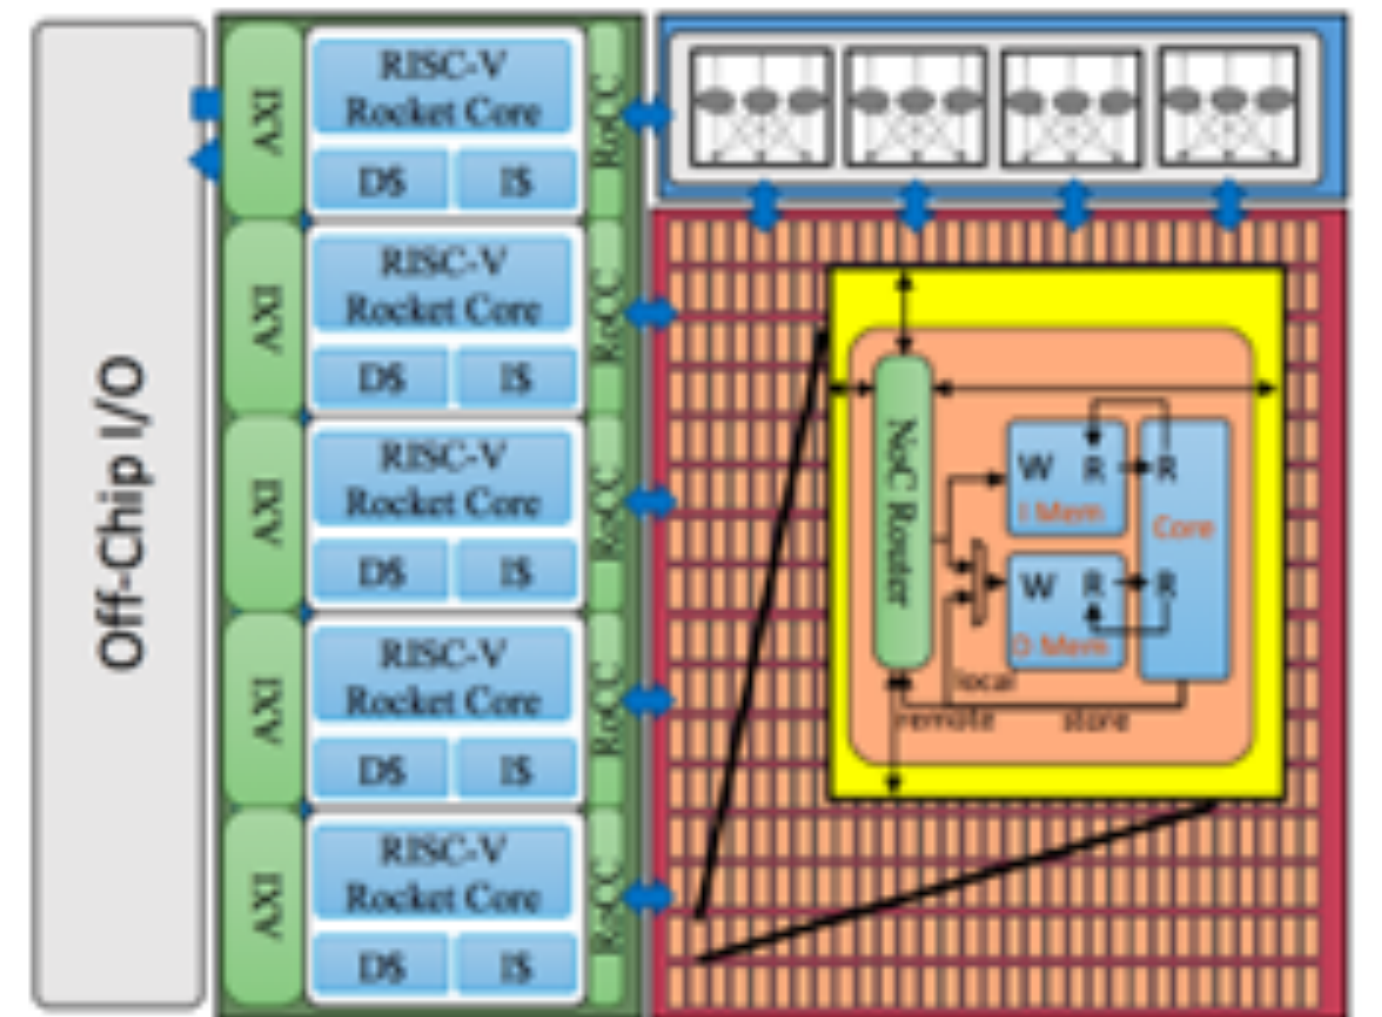



μ TVM

Deep Learning on Bare-Metal Devices

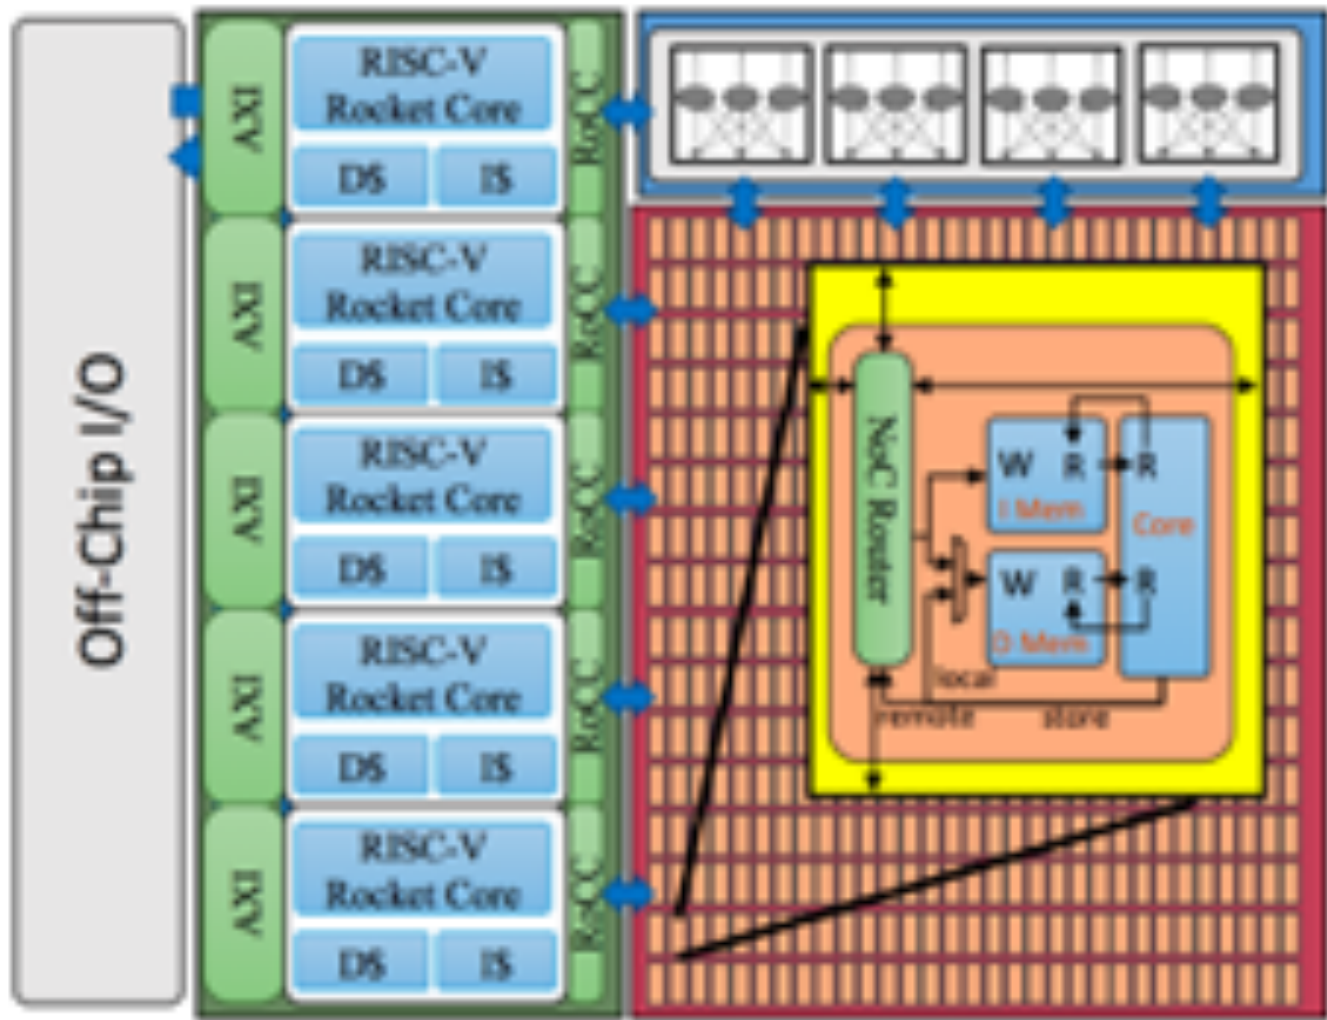



Pratyush Patel



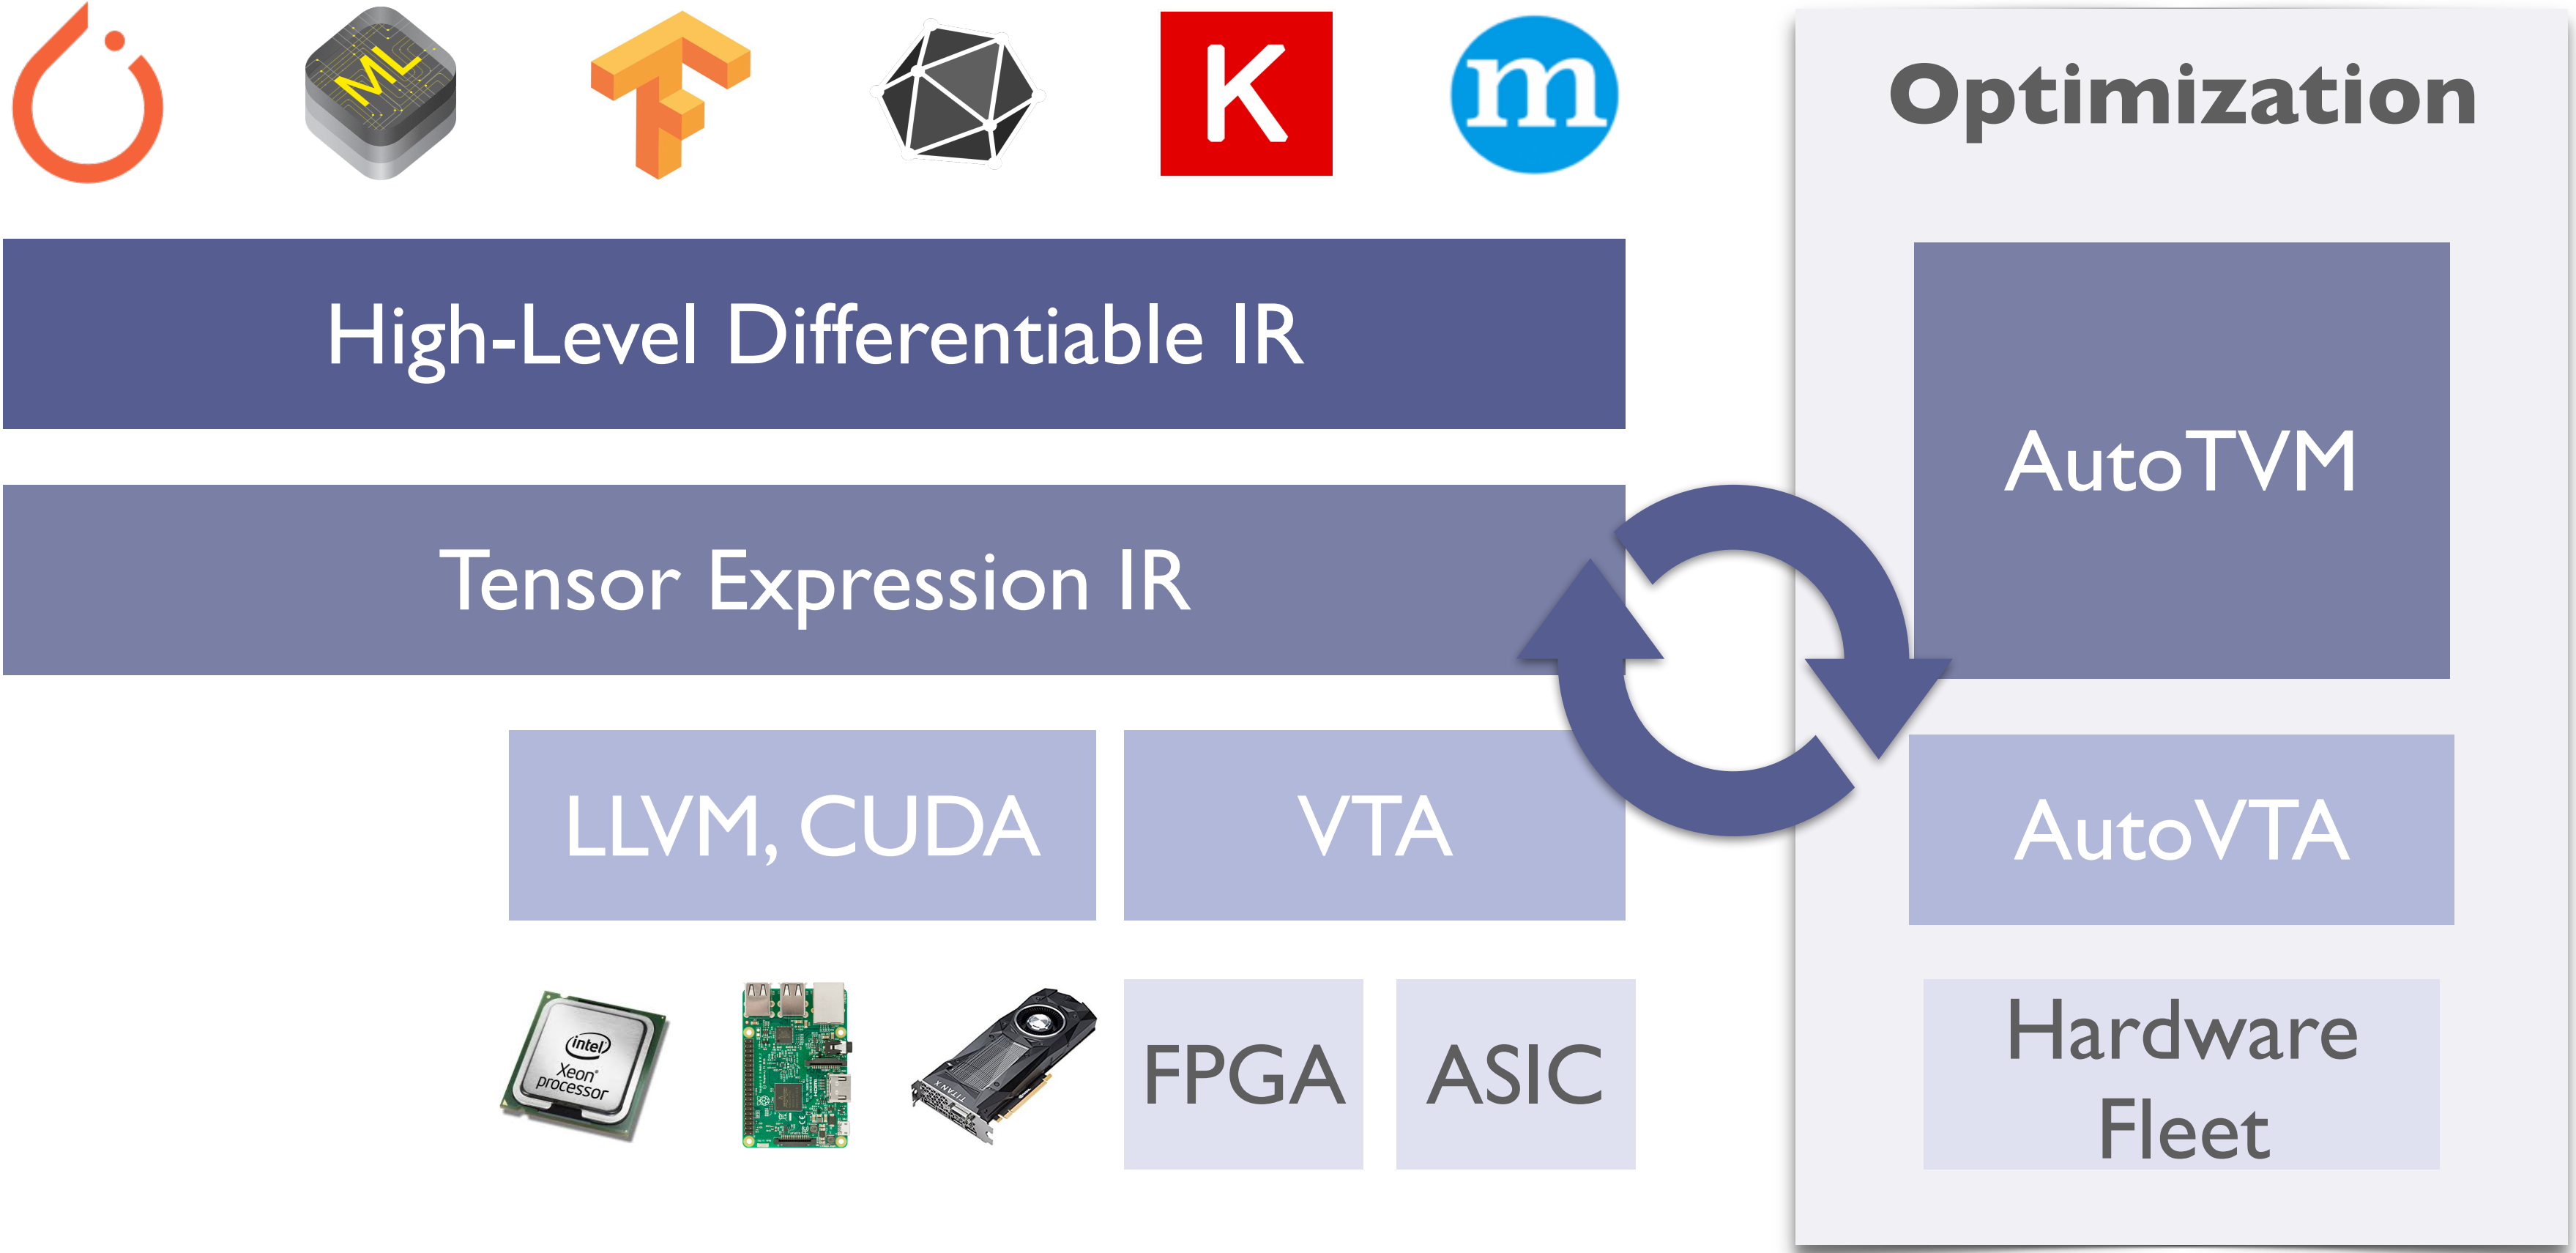




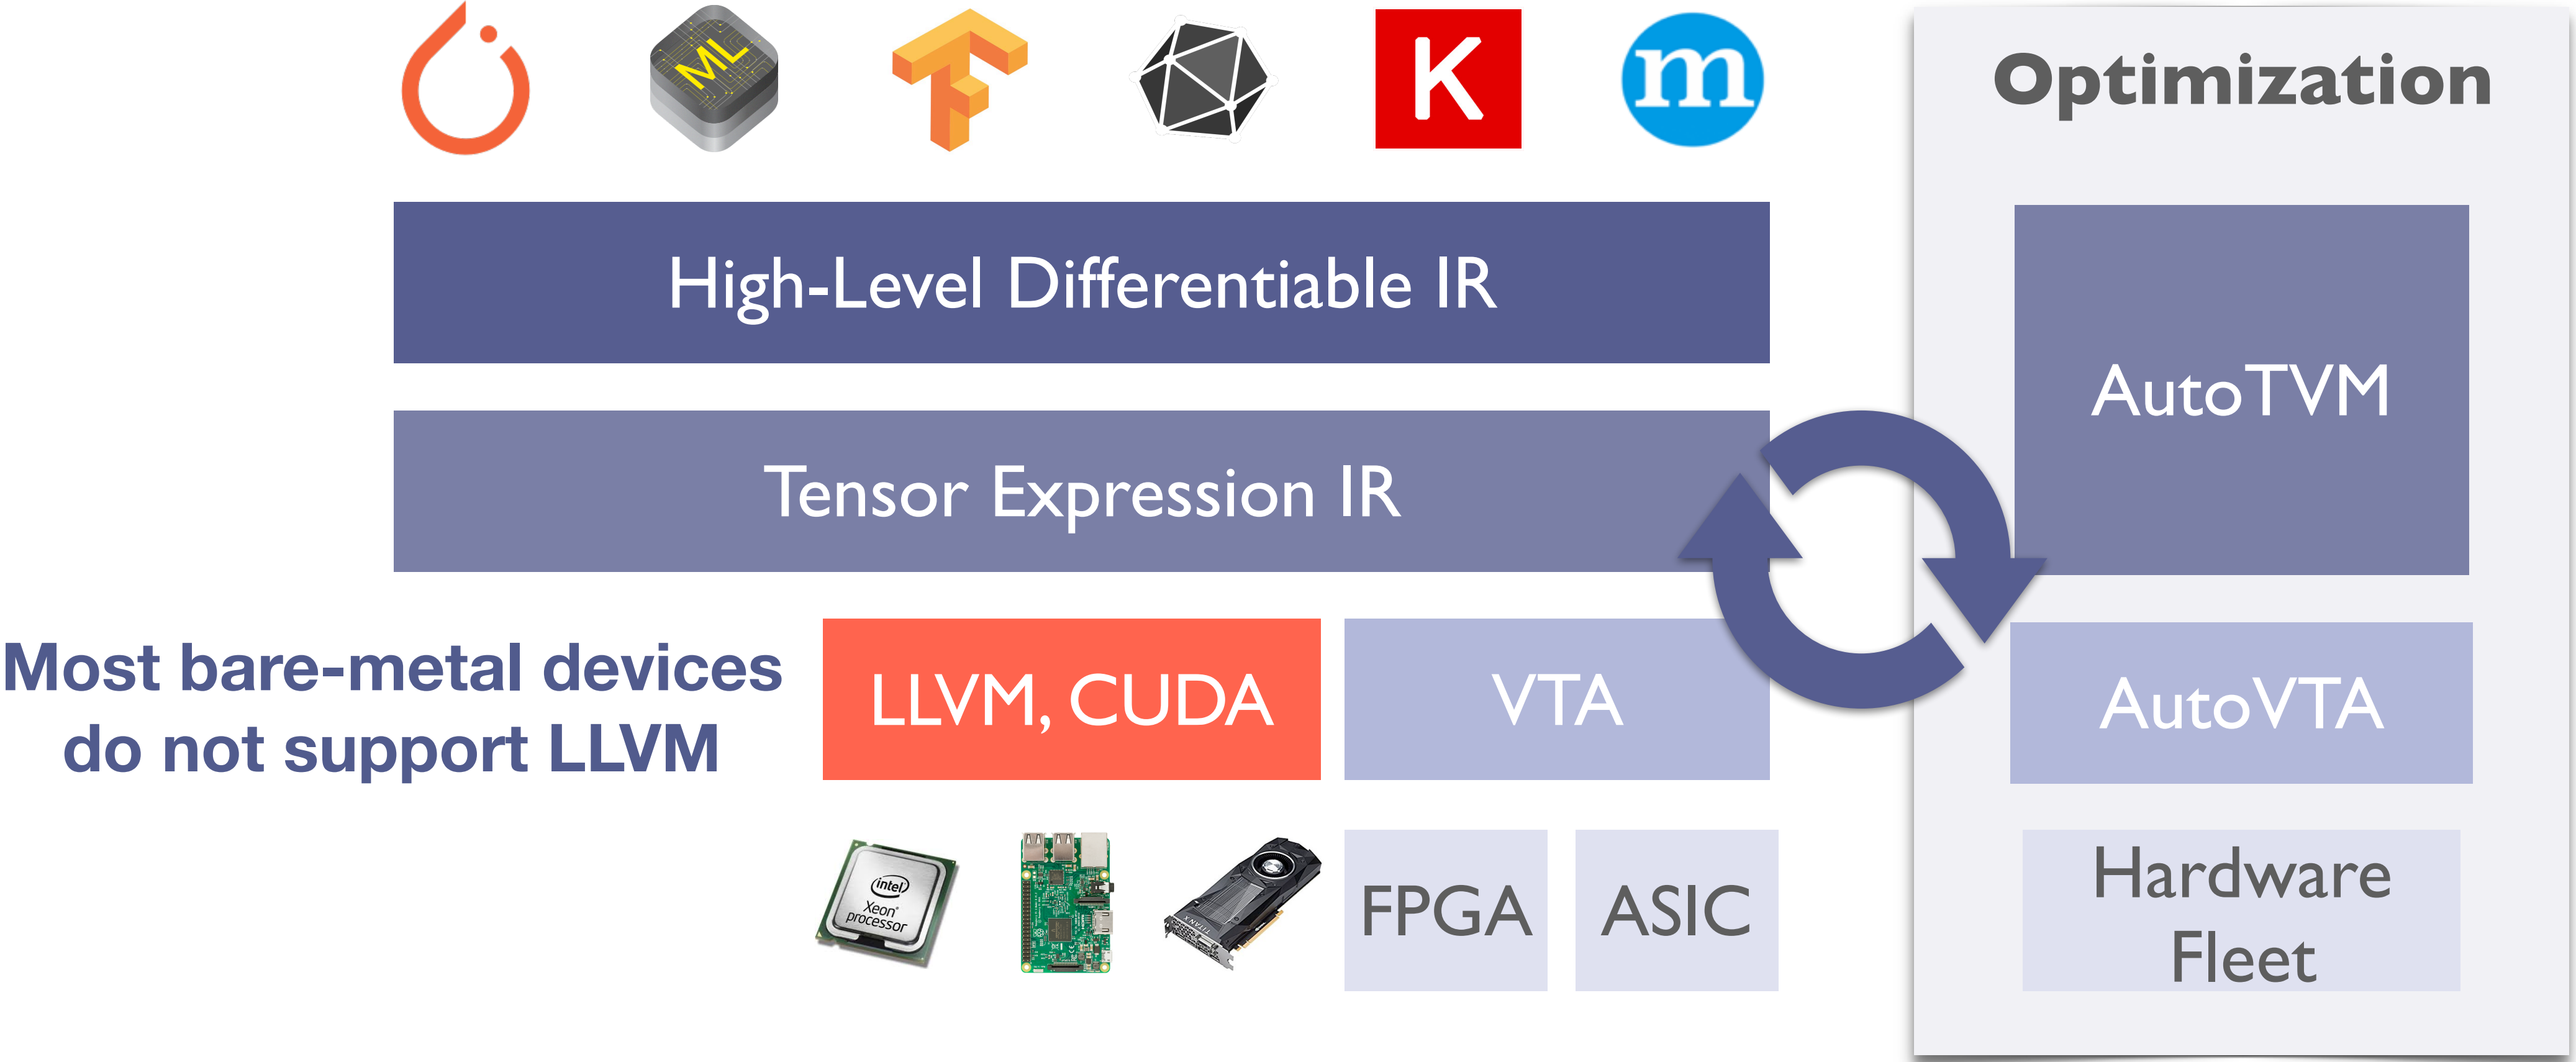
No OS stack



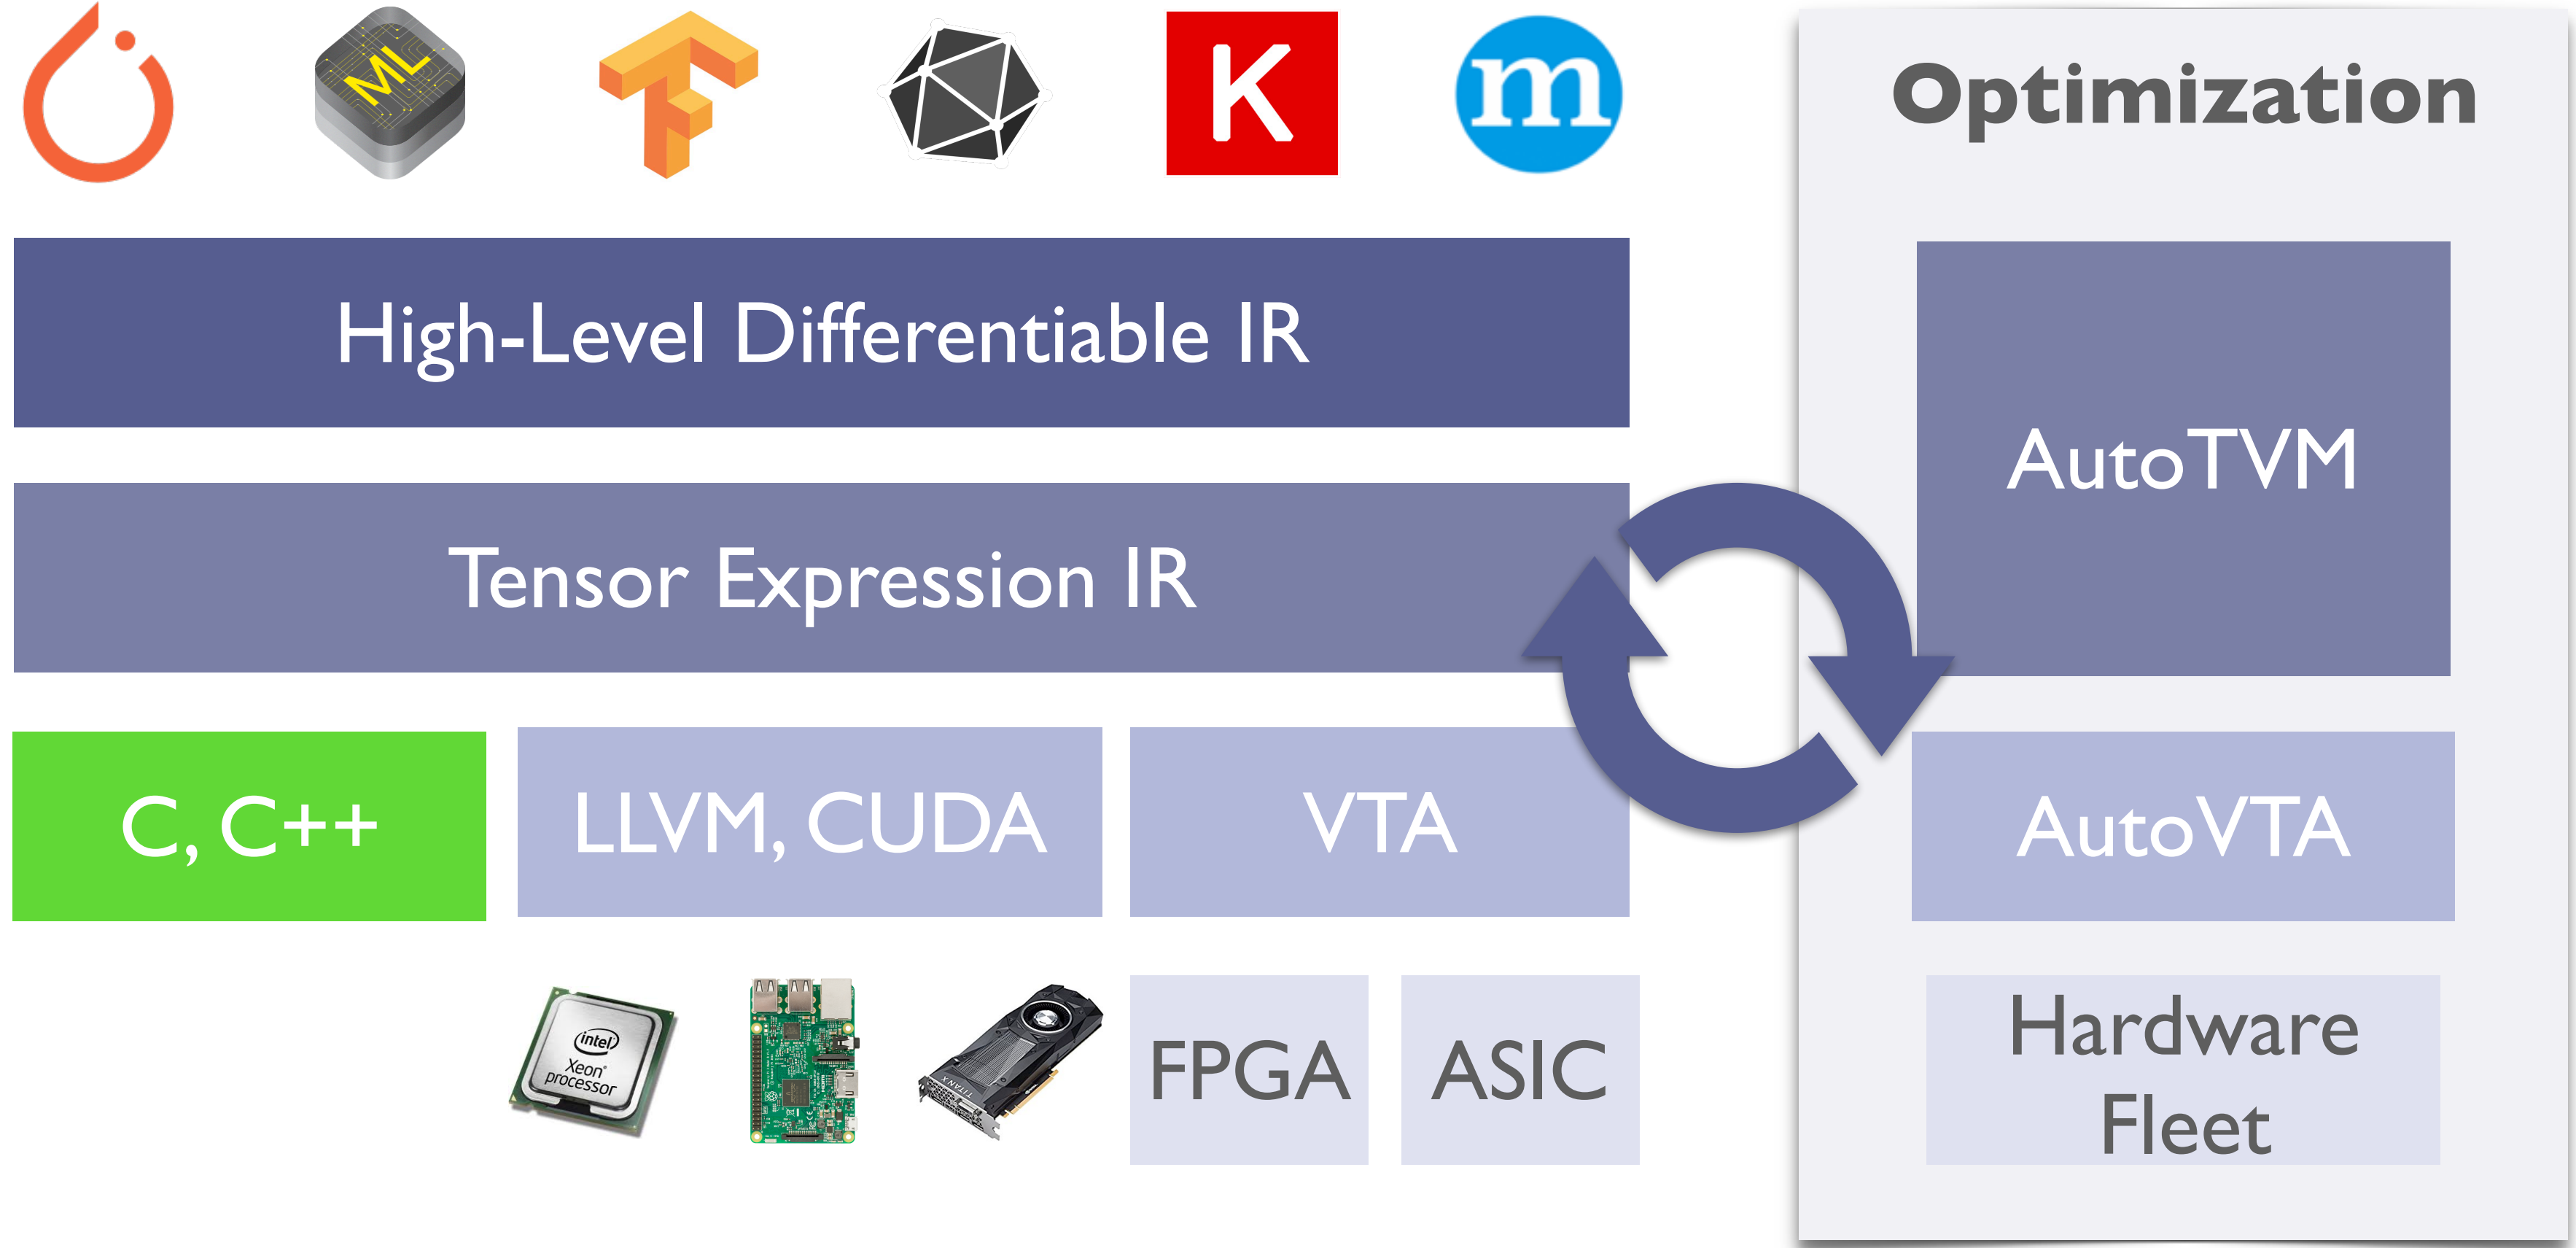
Extend TVM to support bare-metal devices



Extend TVM to support bare-metal devices



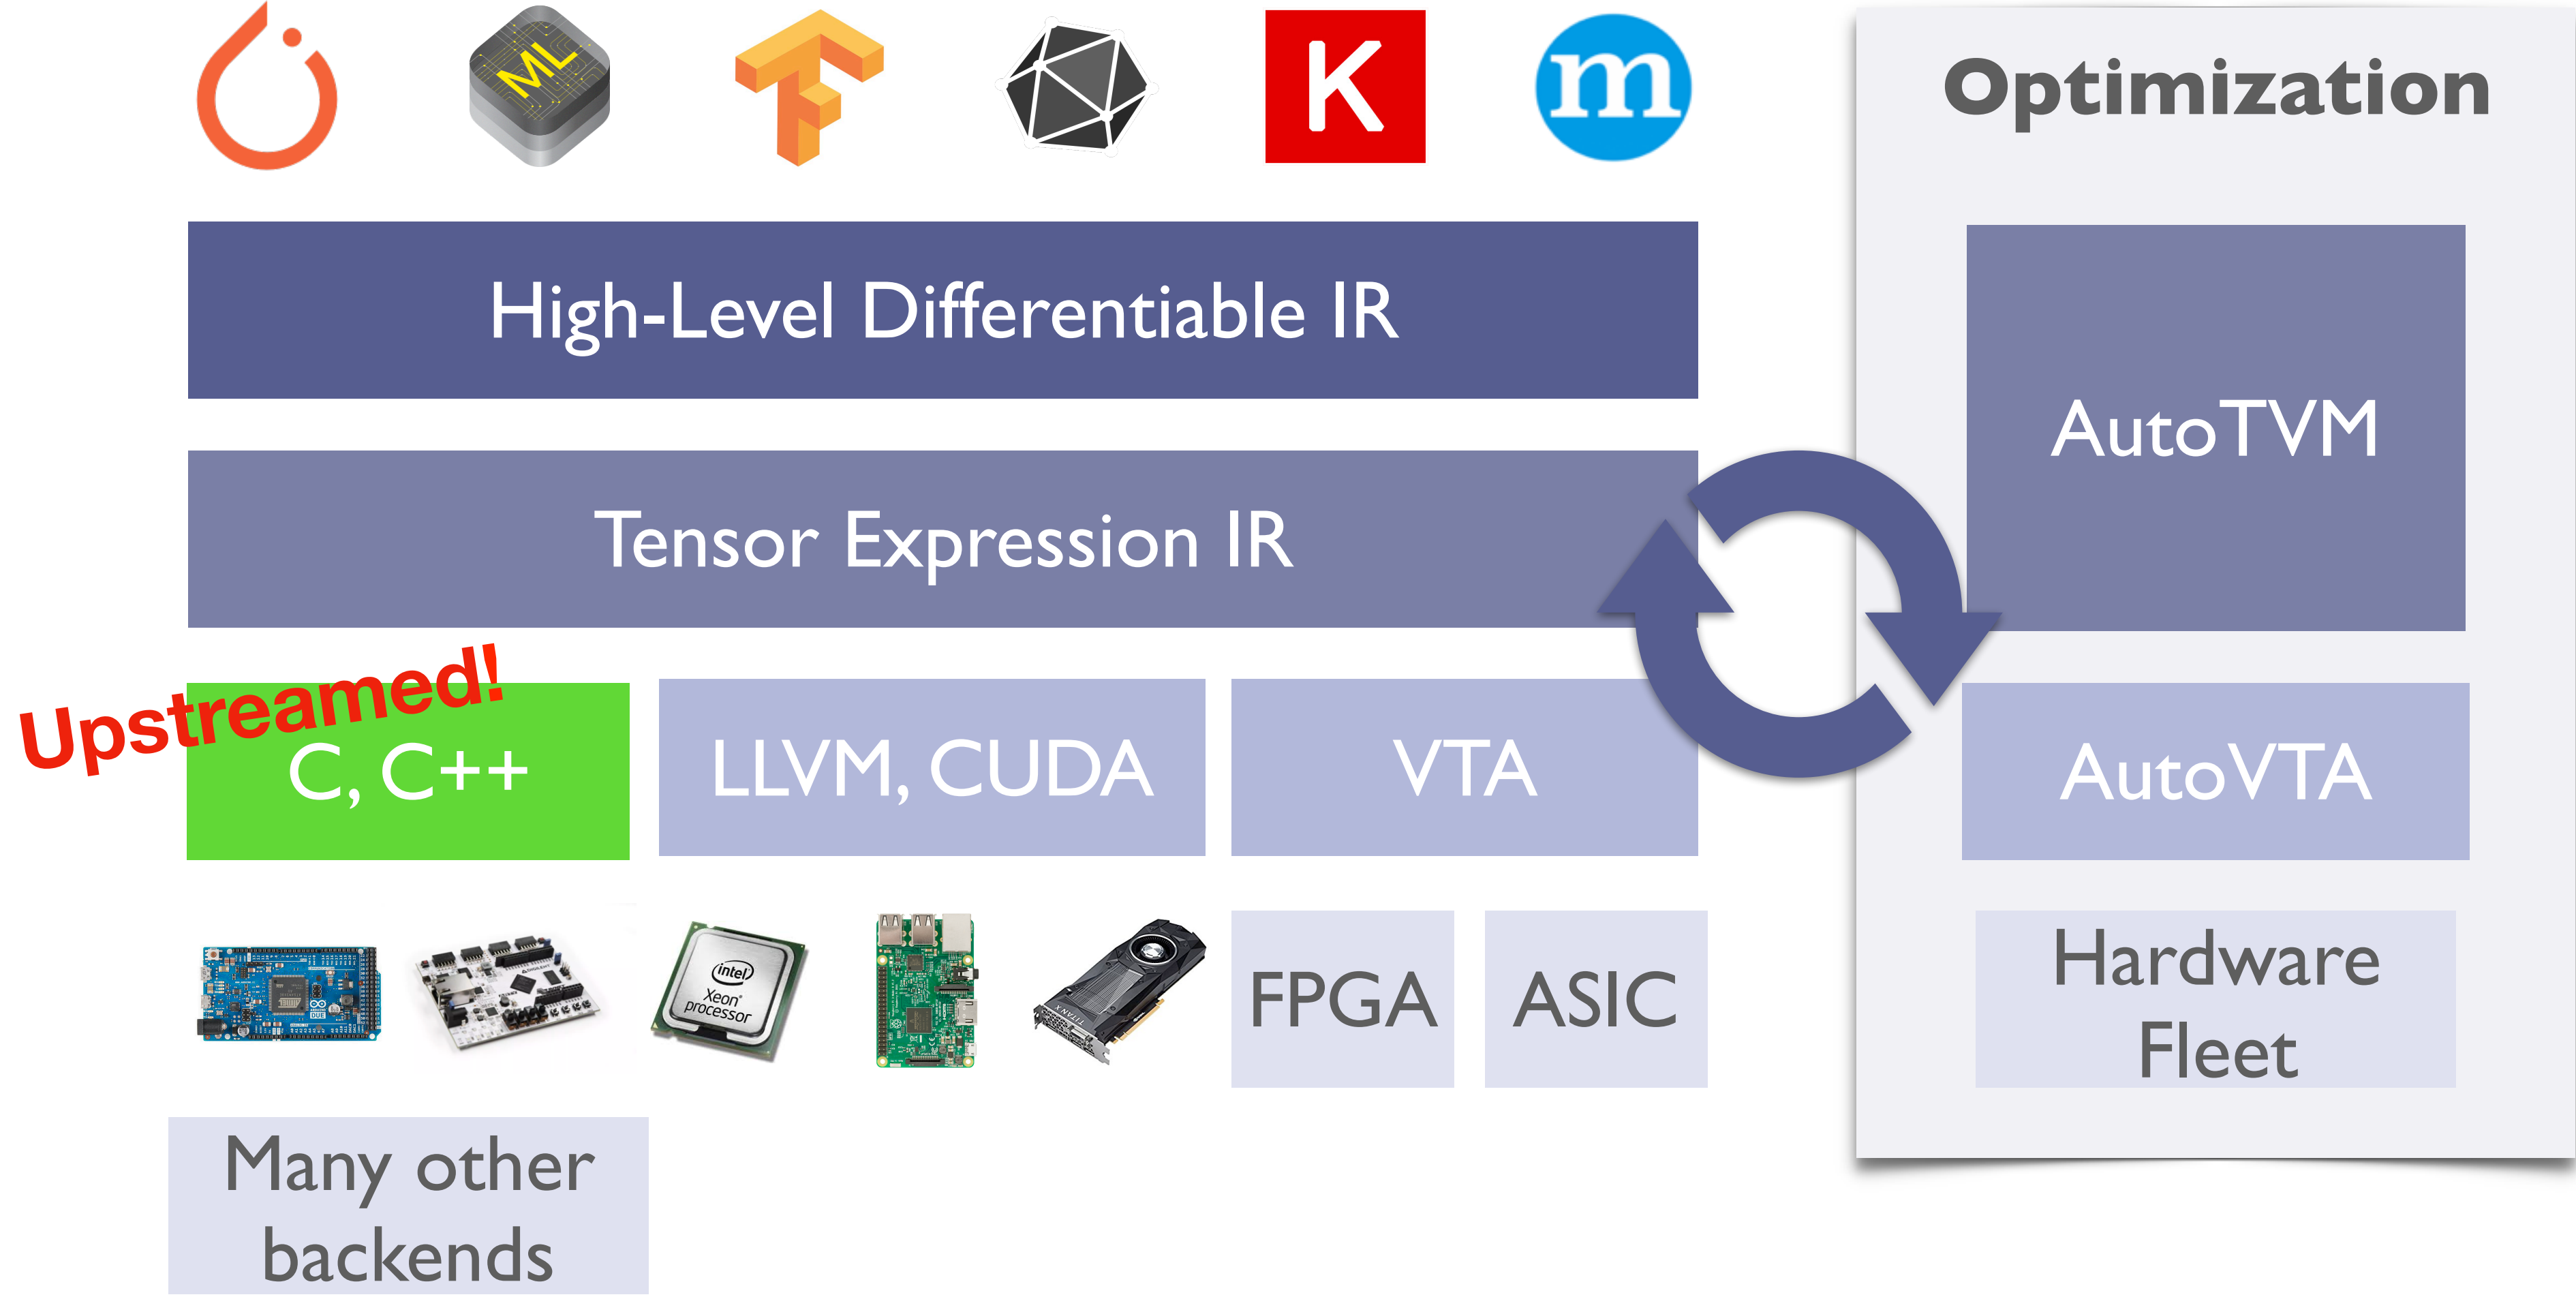
Extend TVM to support bare-metal devices



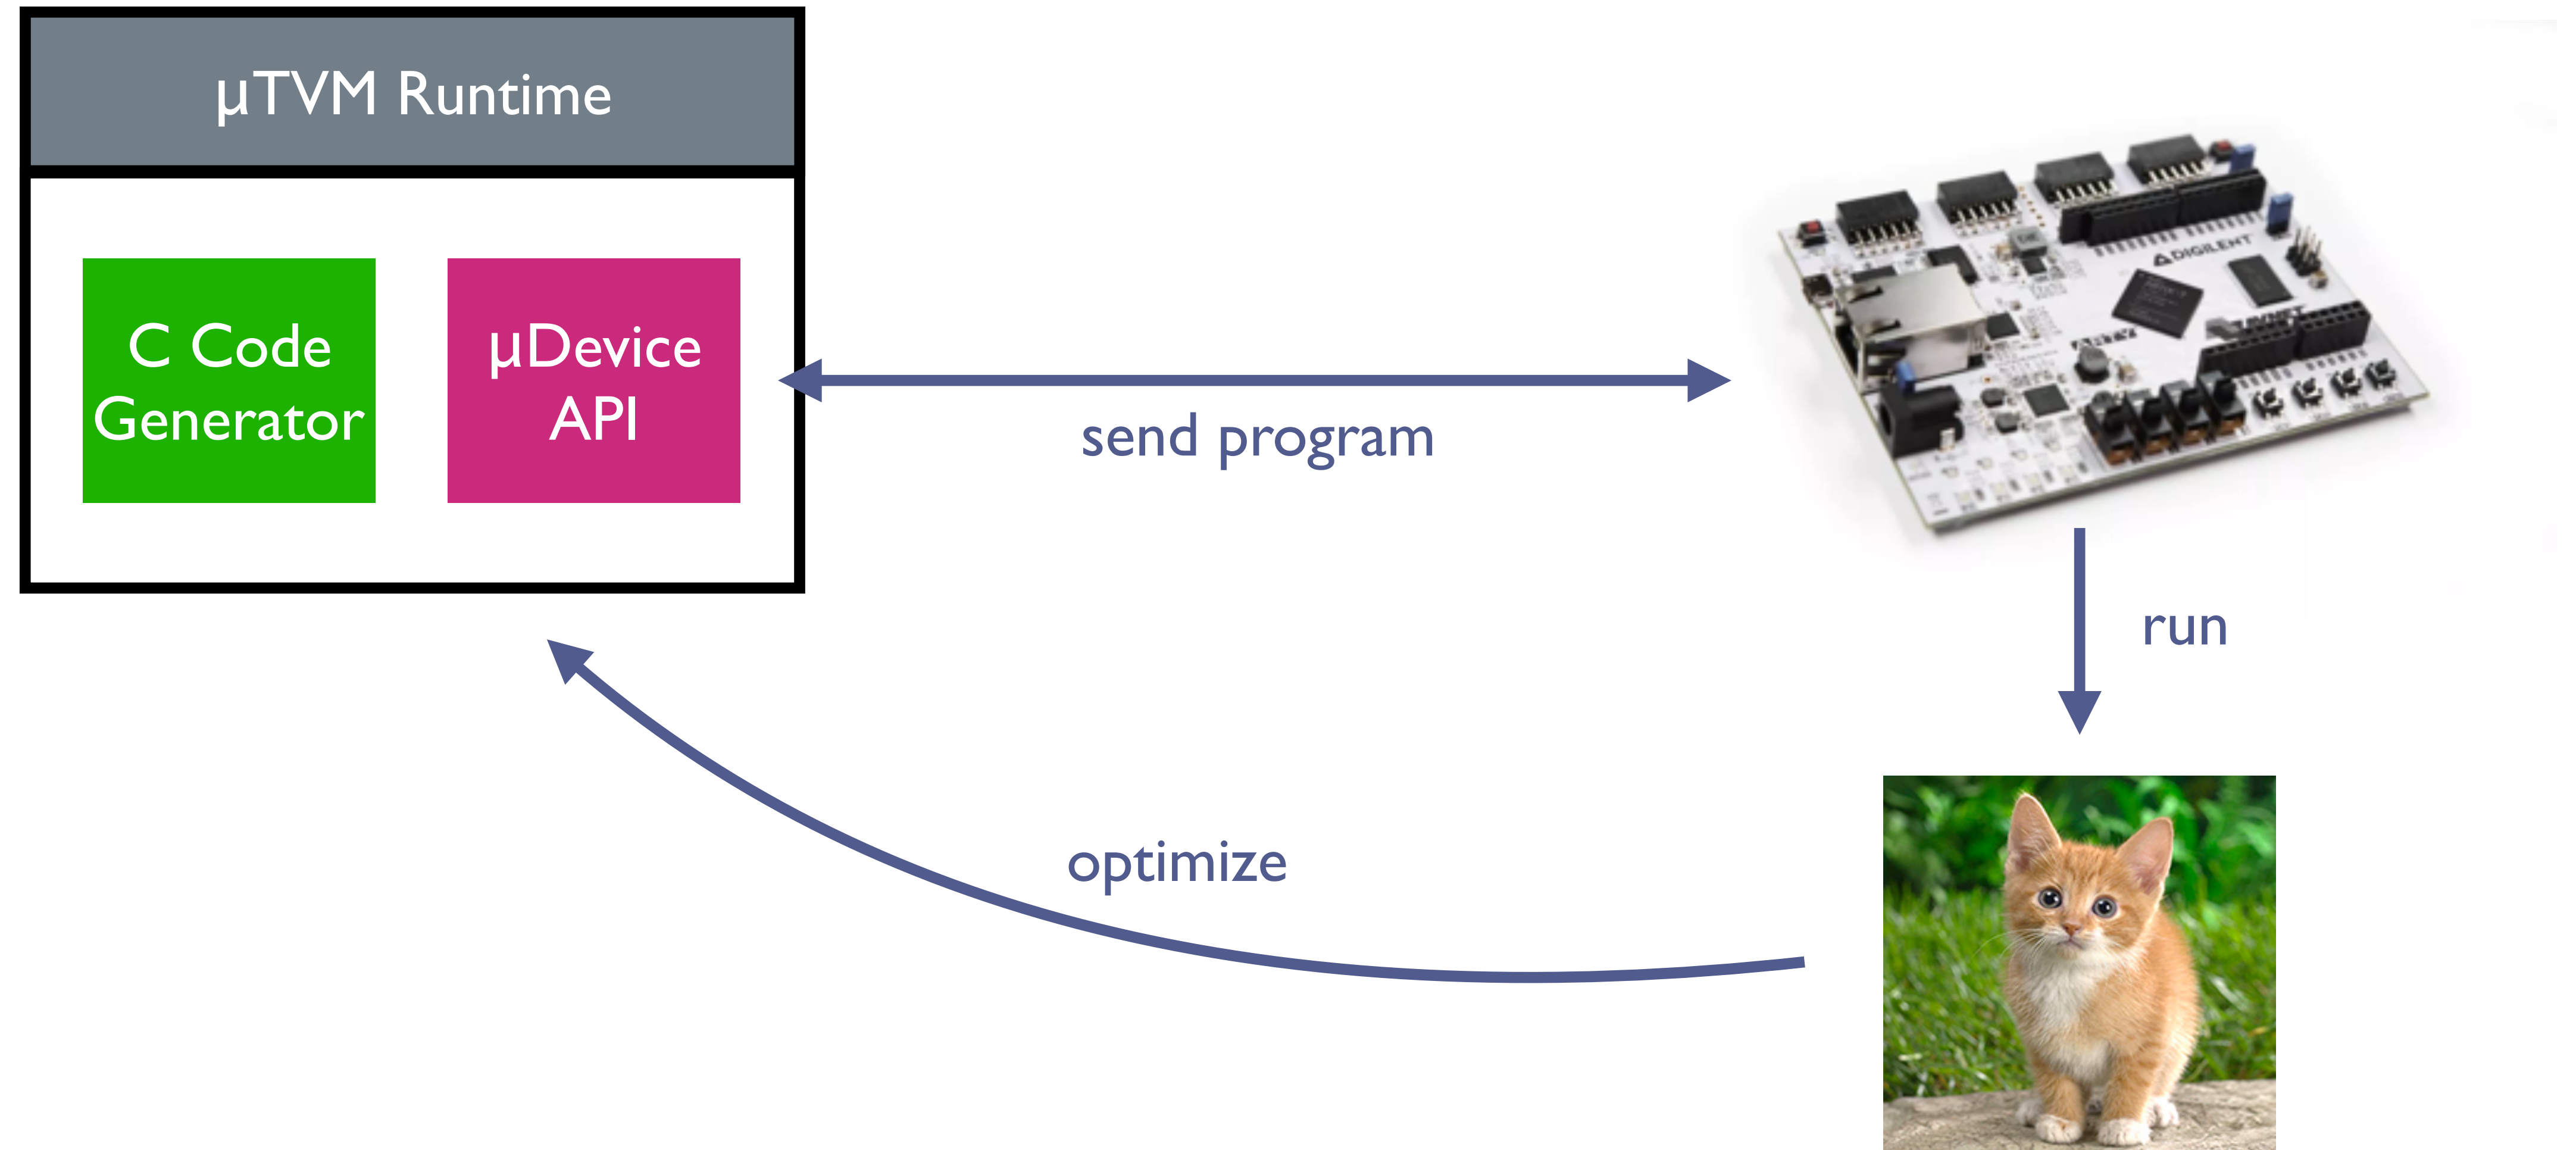
Extend TVM to support bare-metal devices



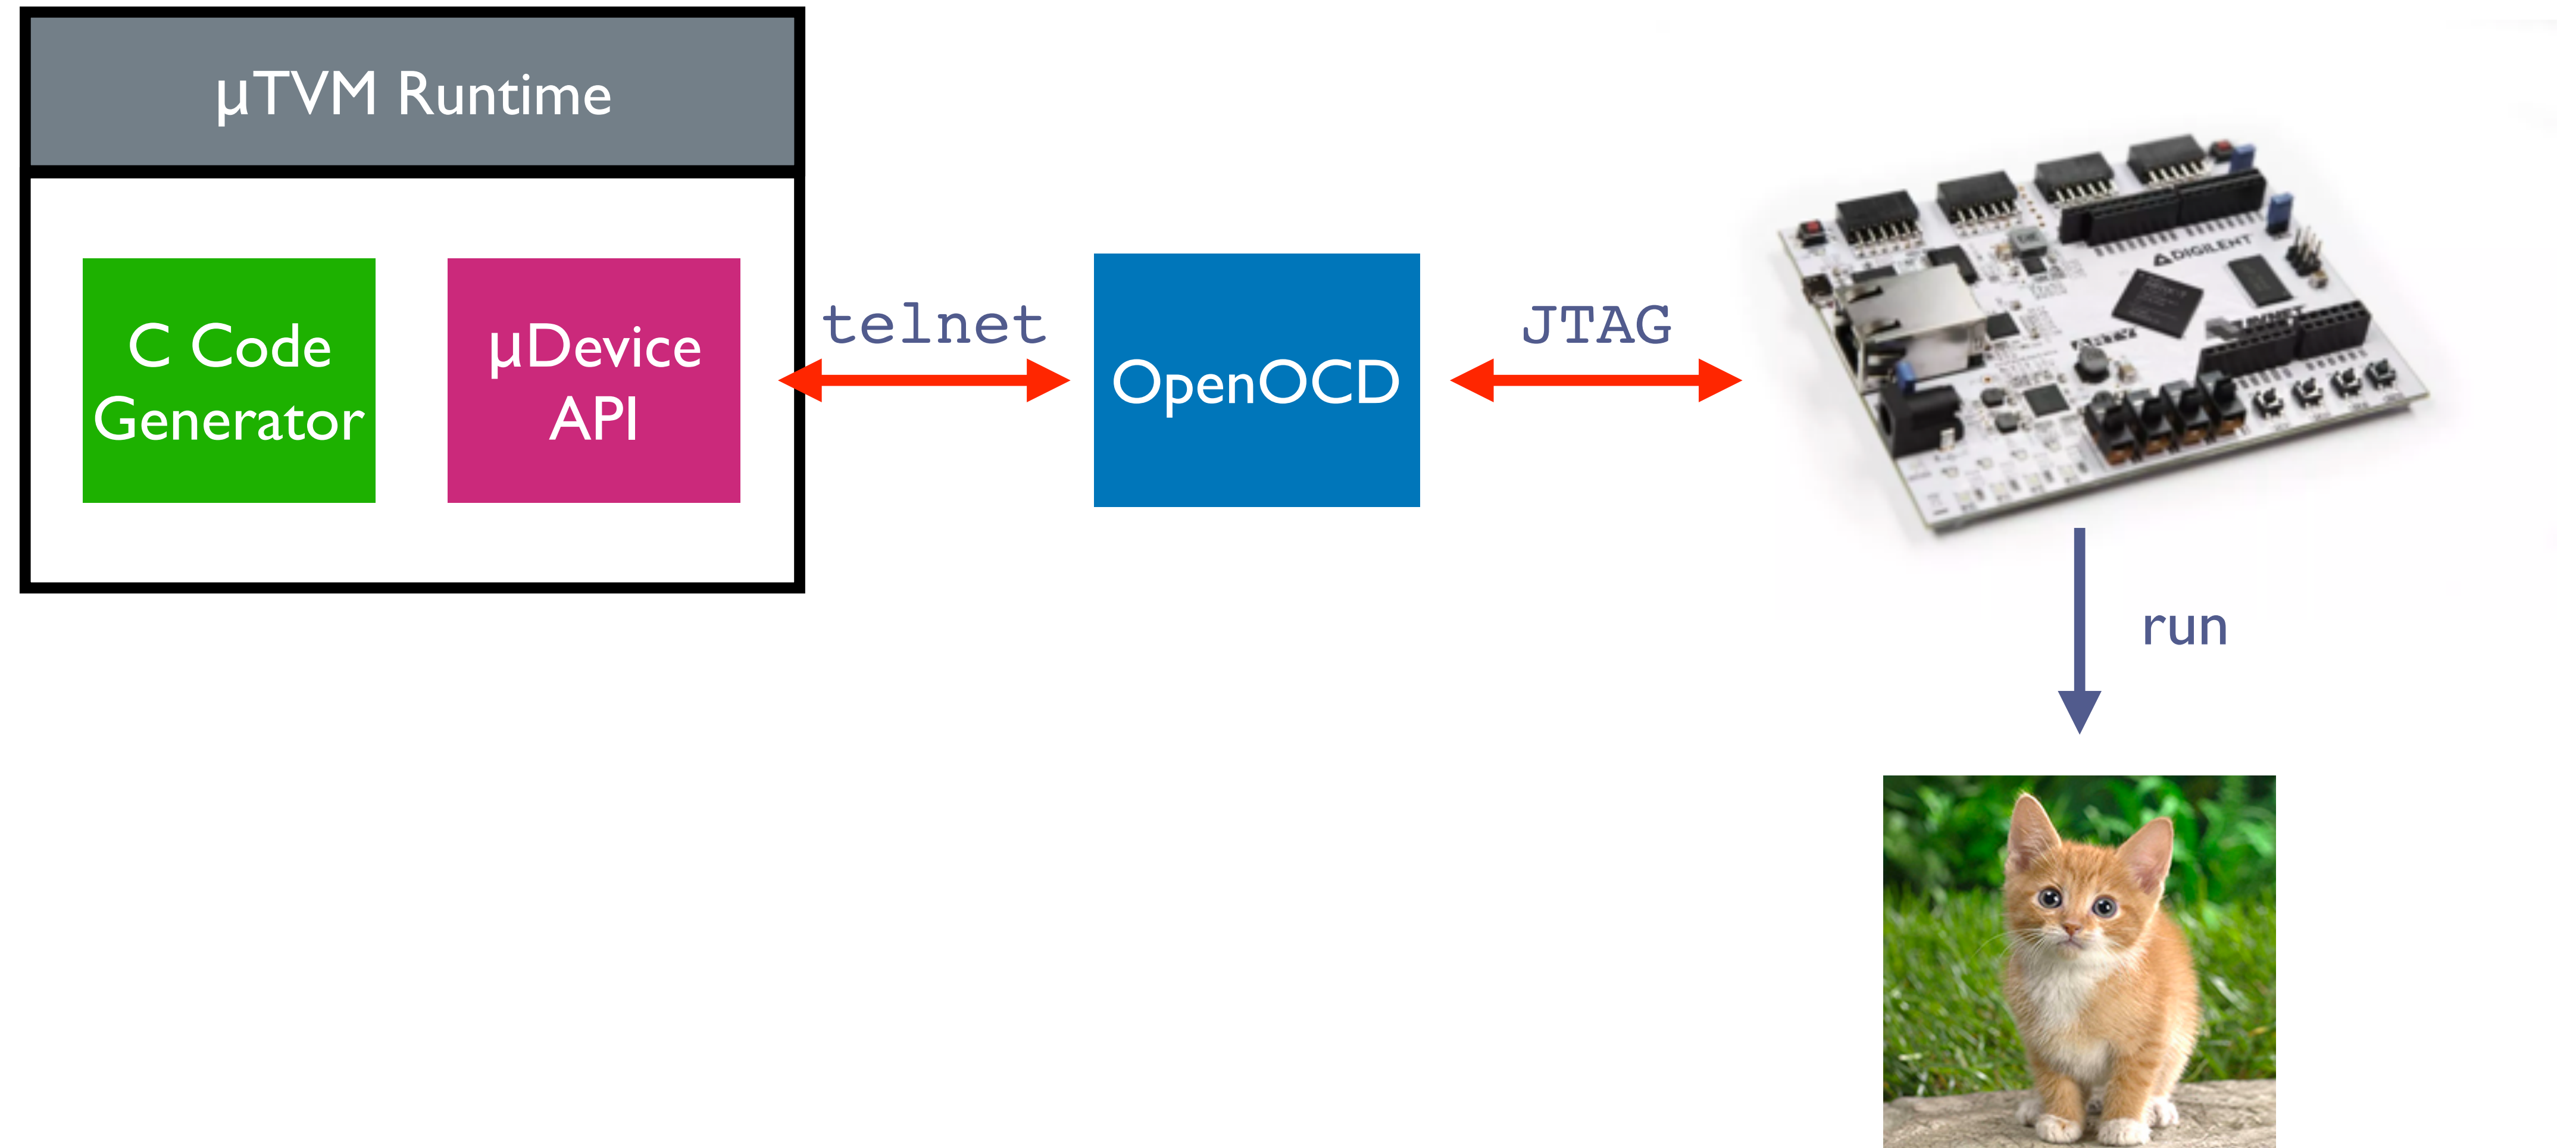
Extend TVM to support bare-metal devices



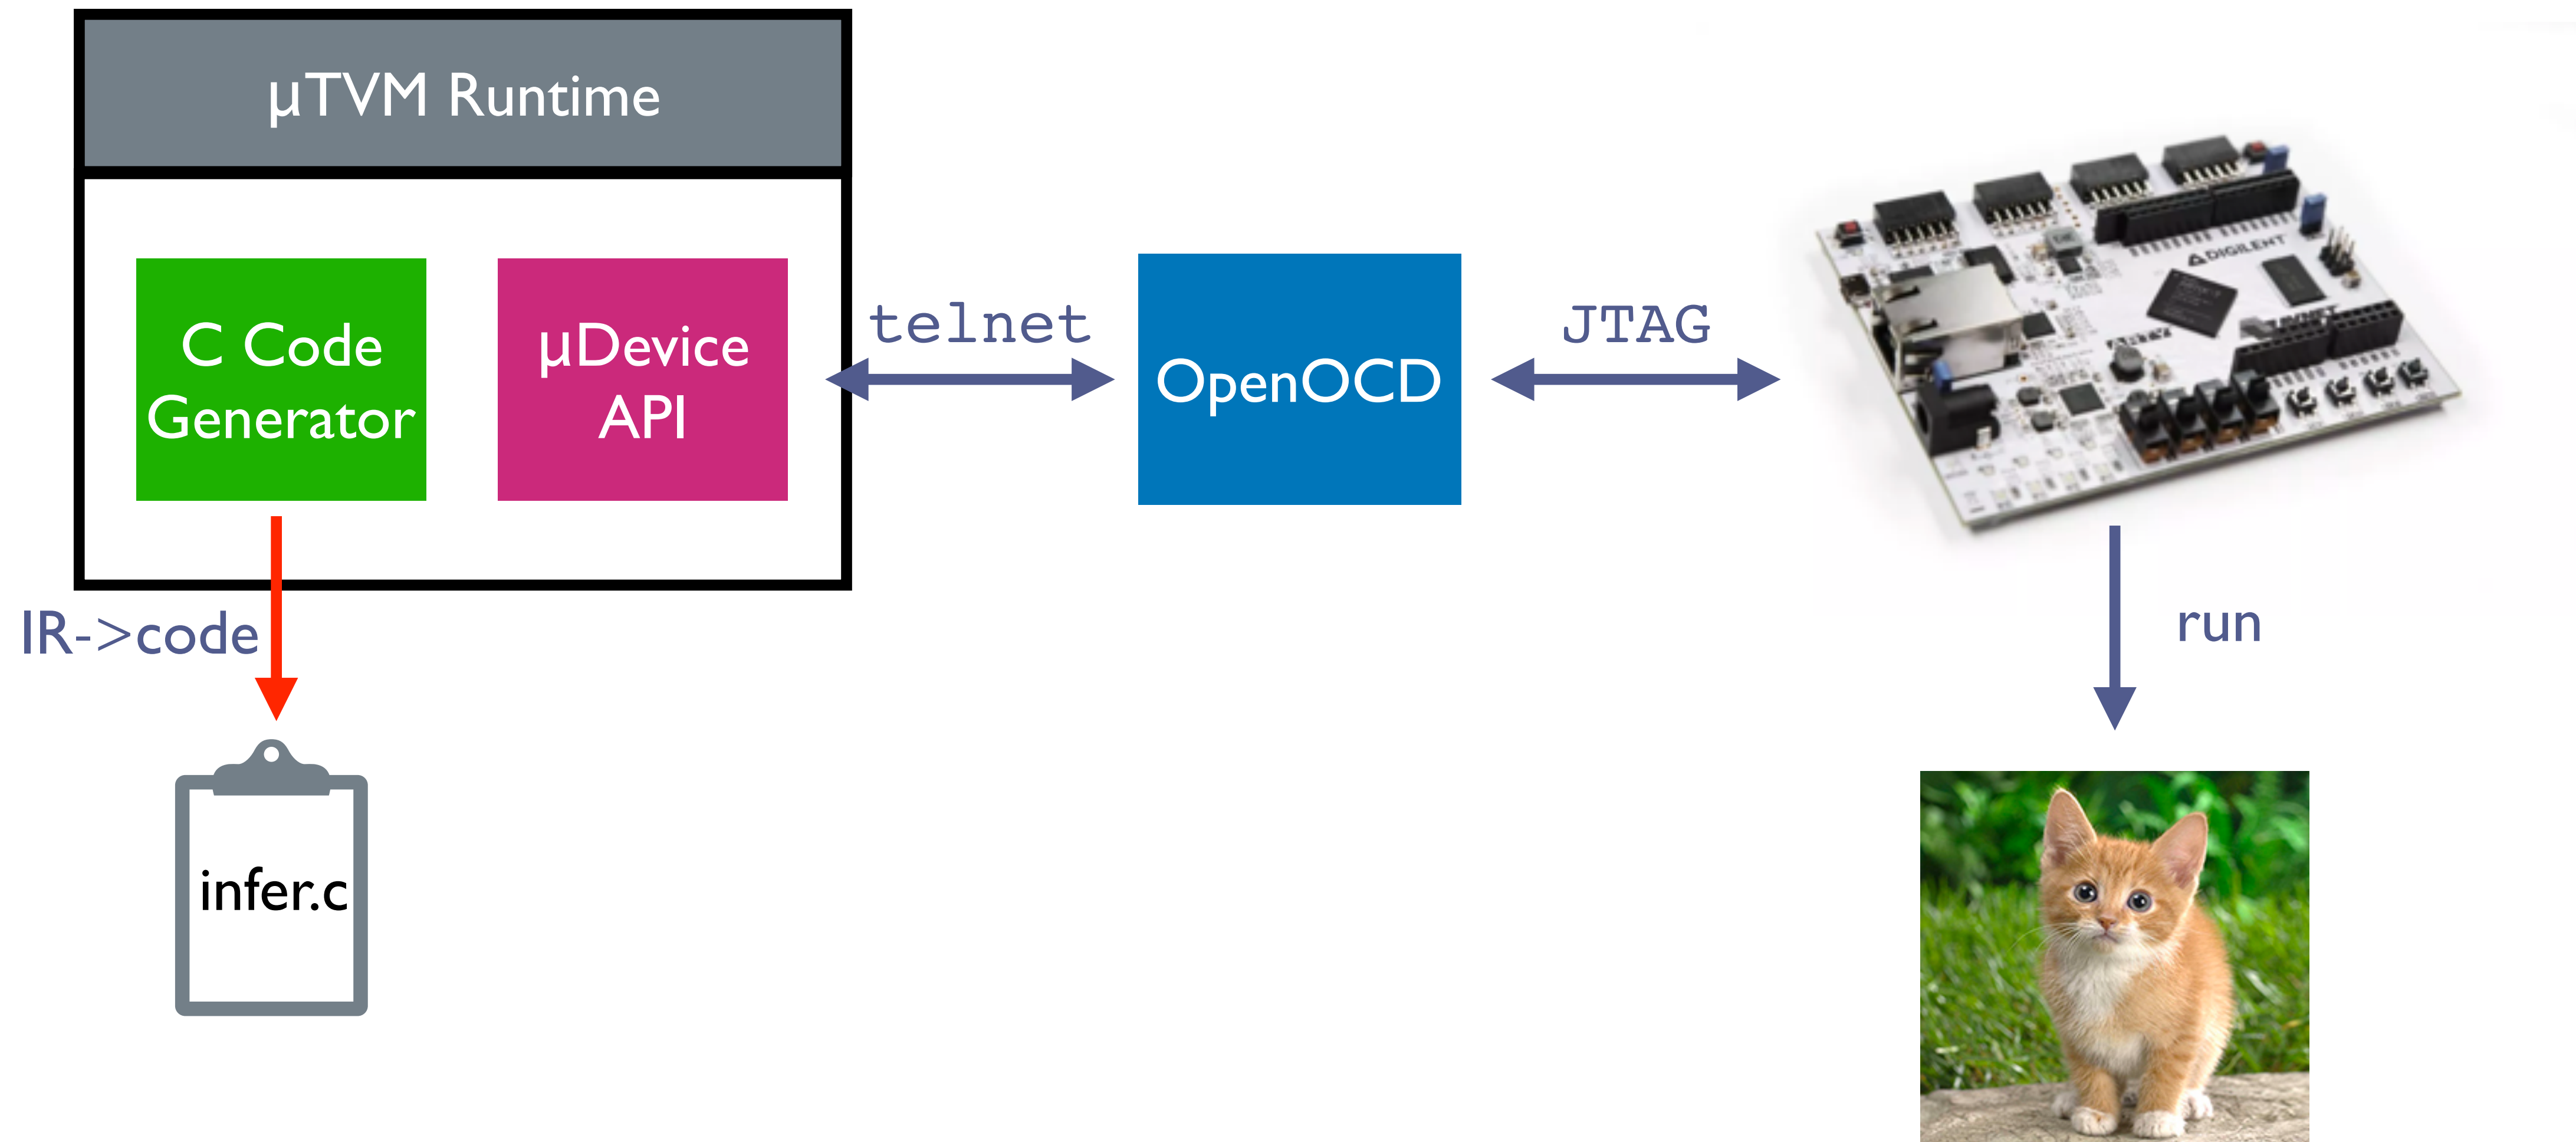
μ TVM builds upon AutoTVM



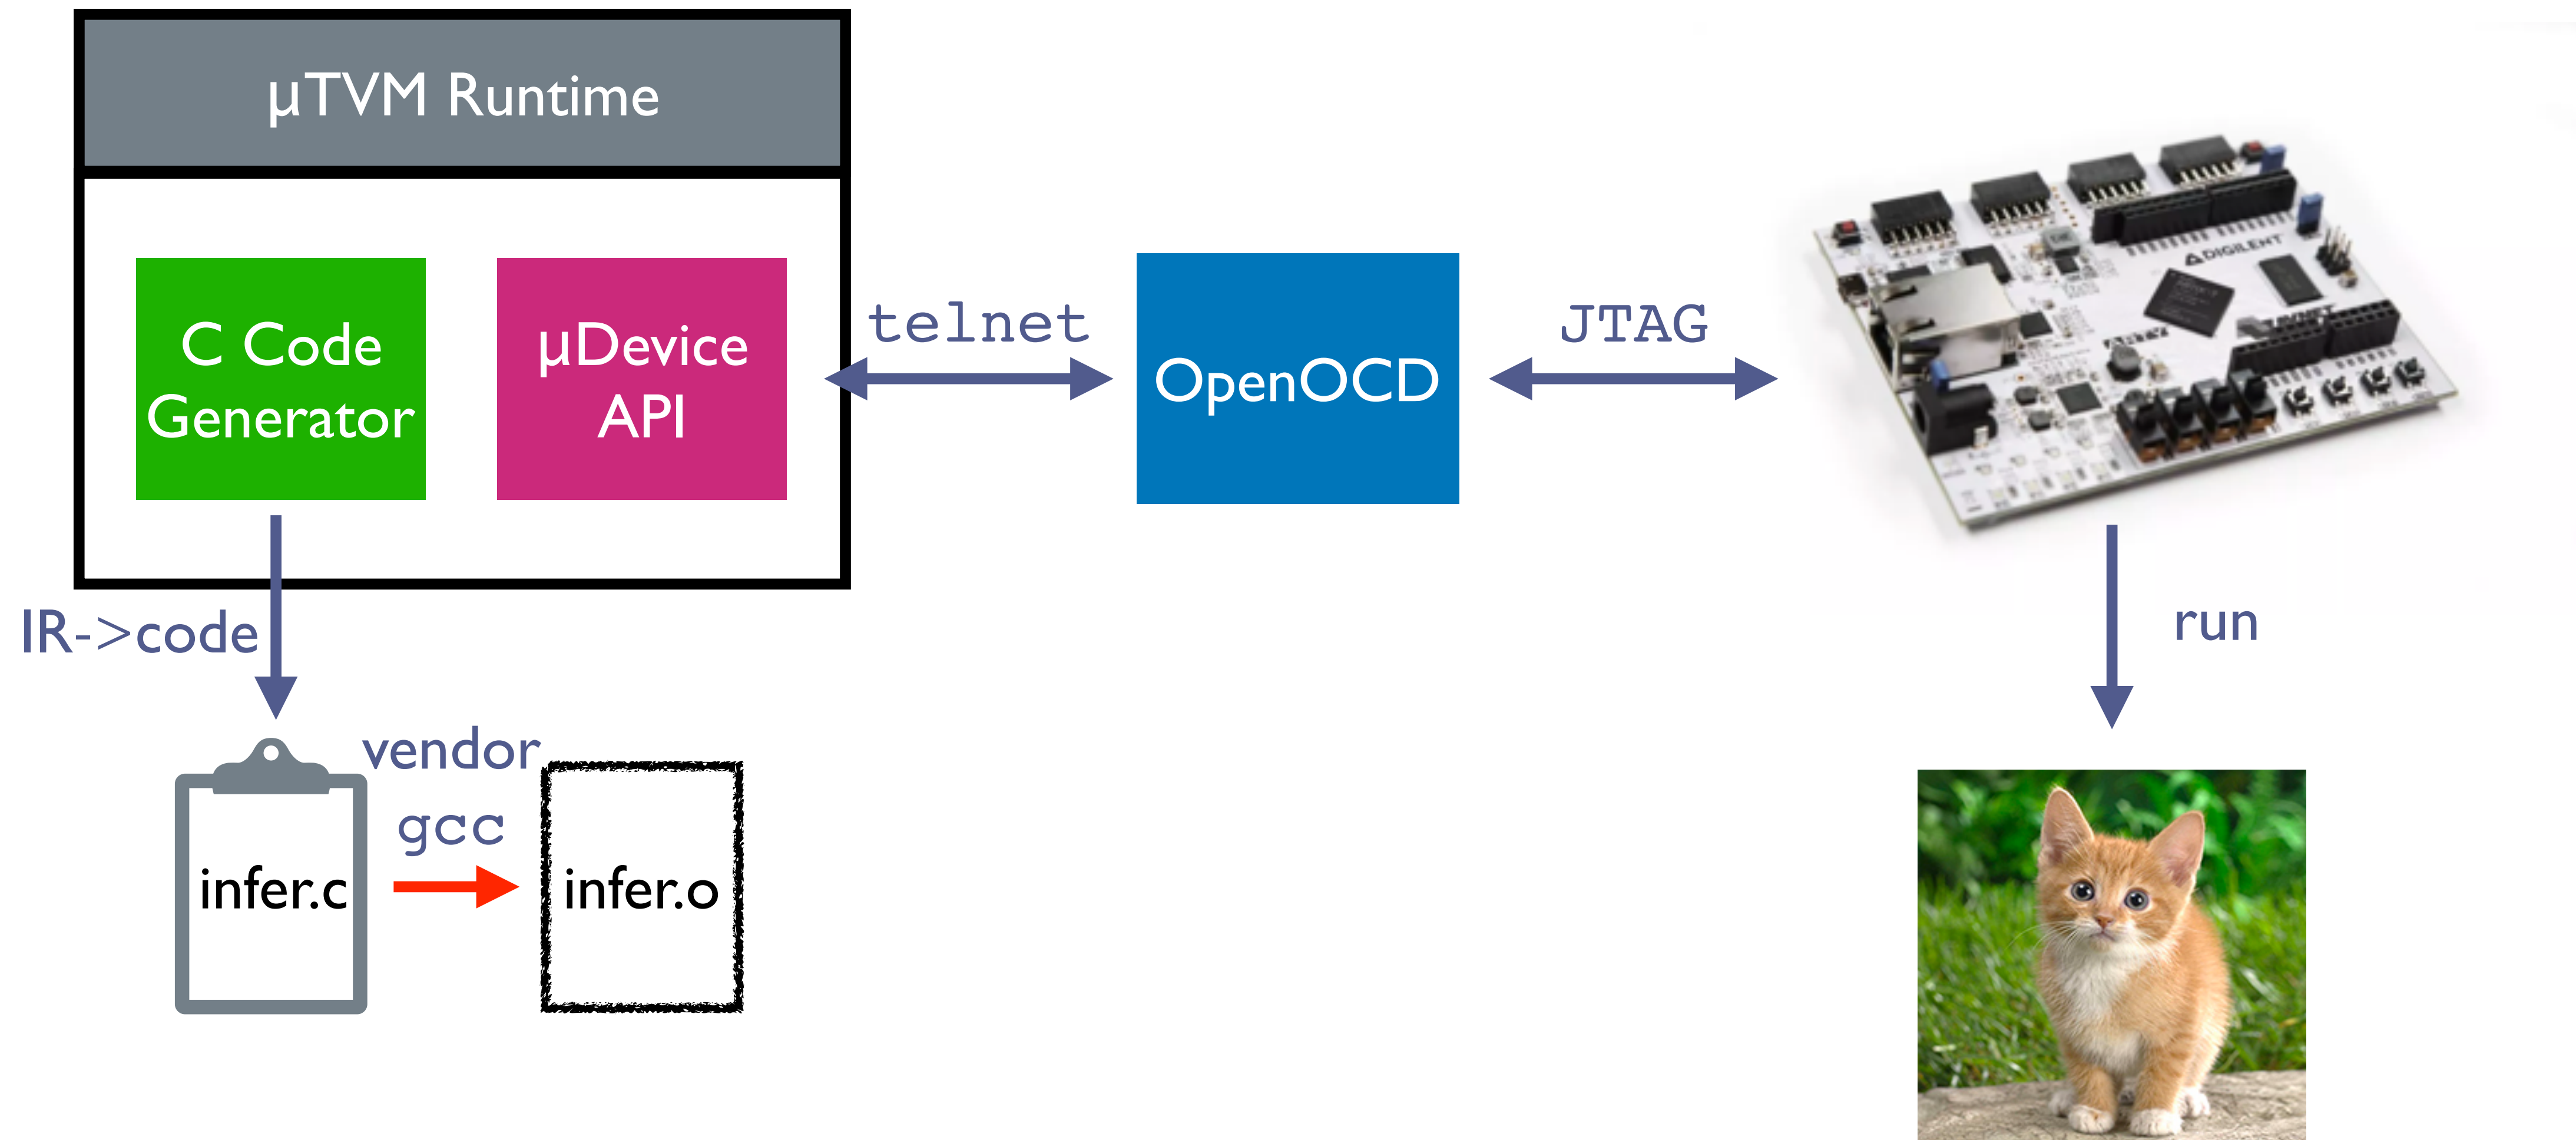
A closer look at μ TVM



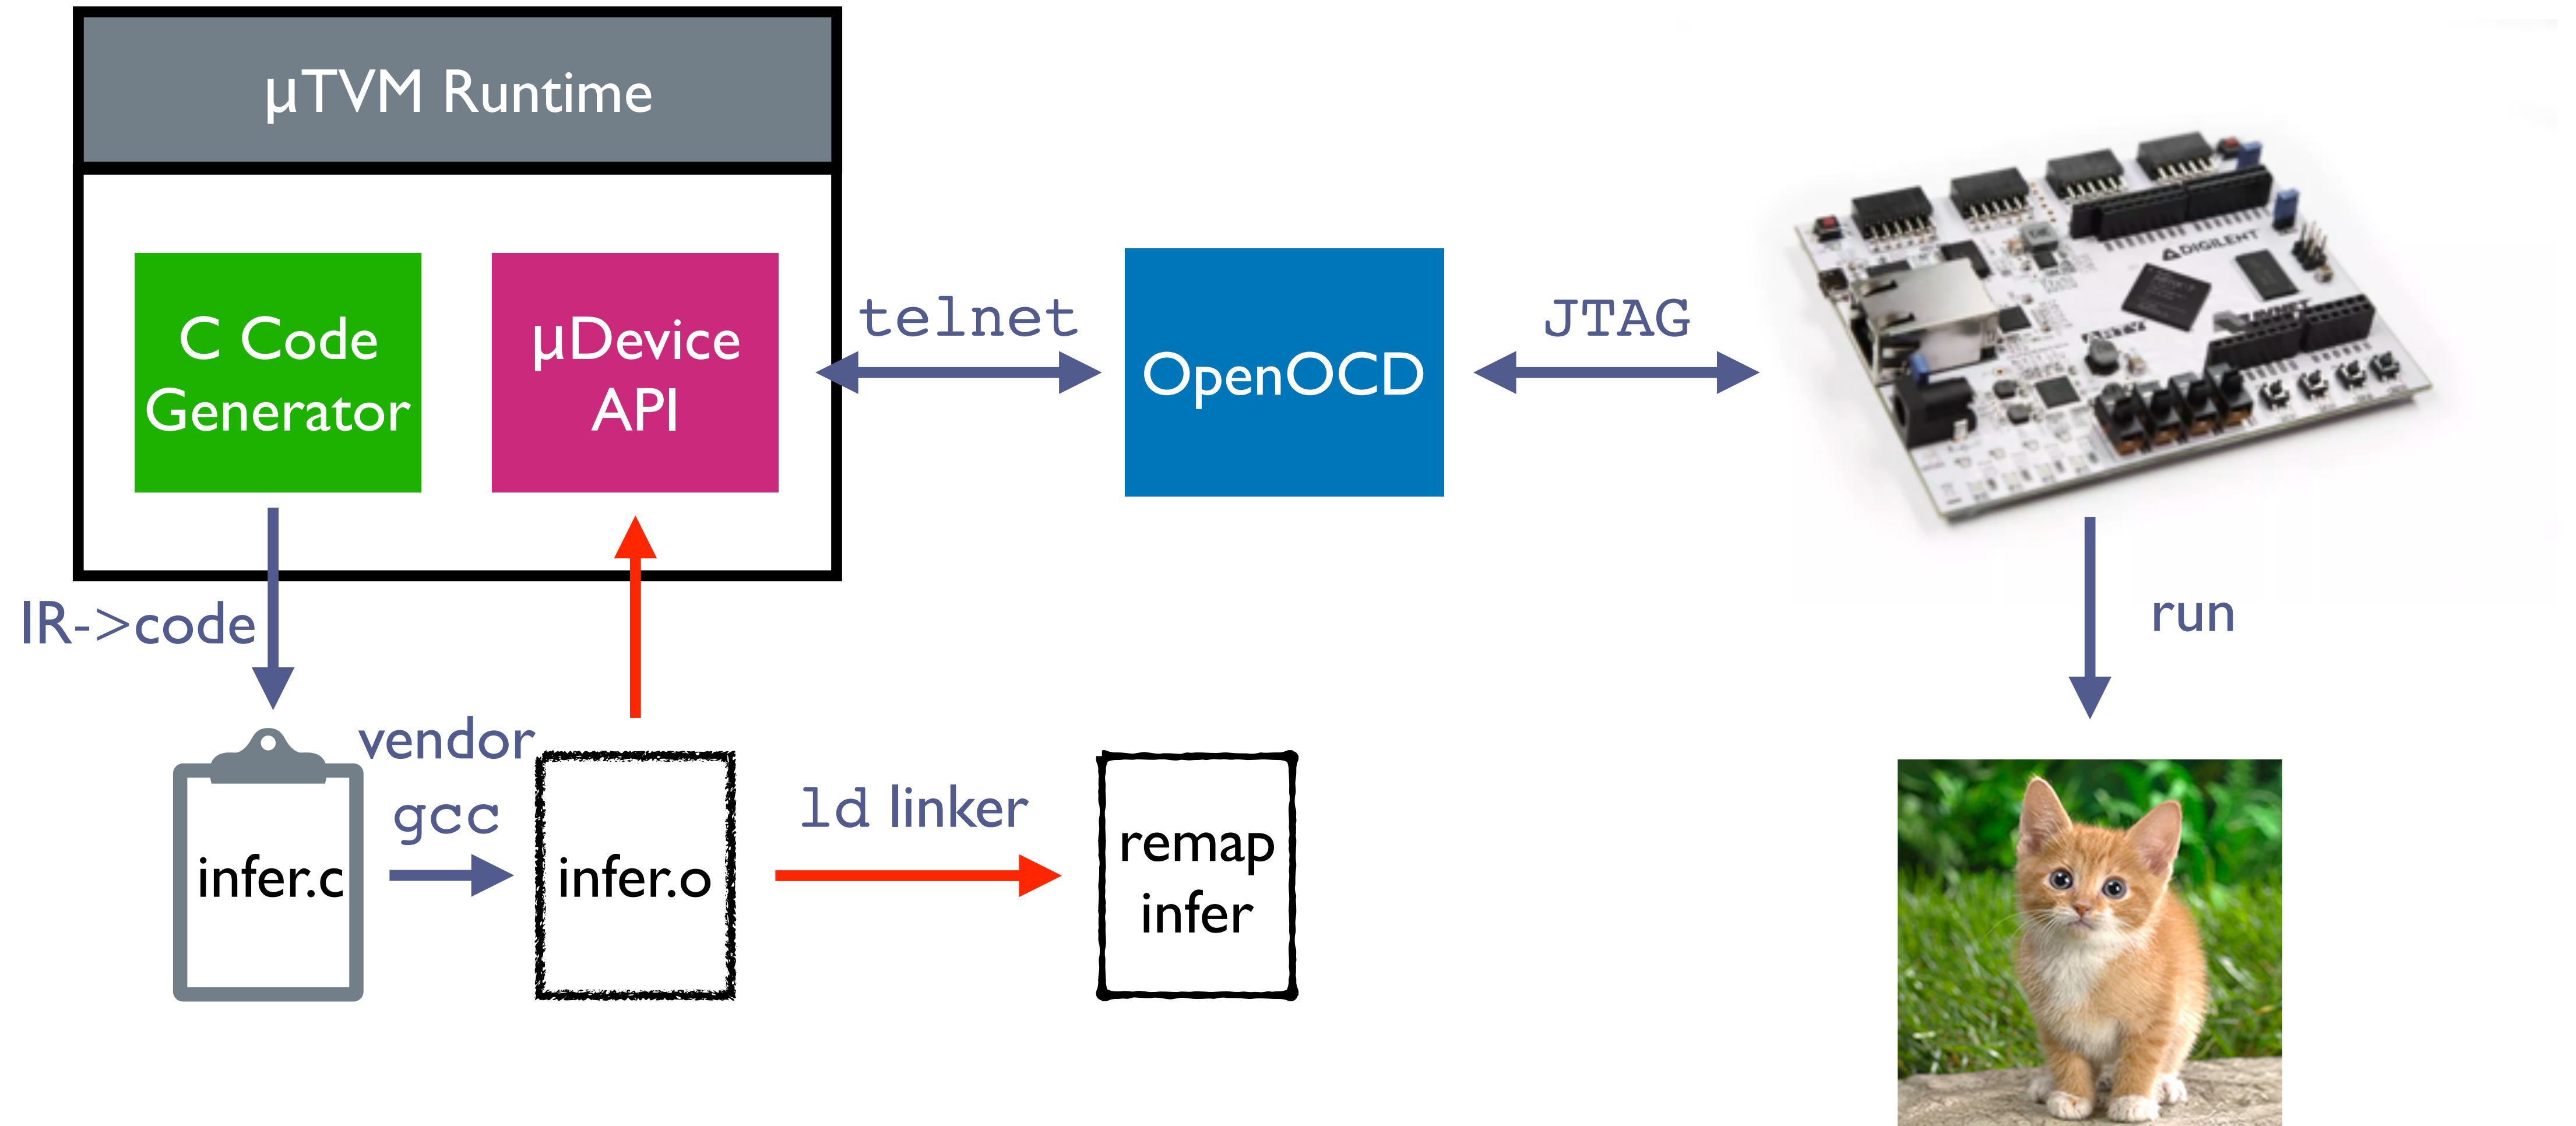
A closer look at μ TVM



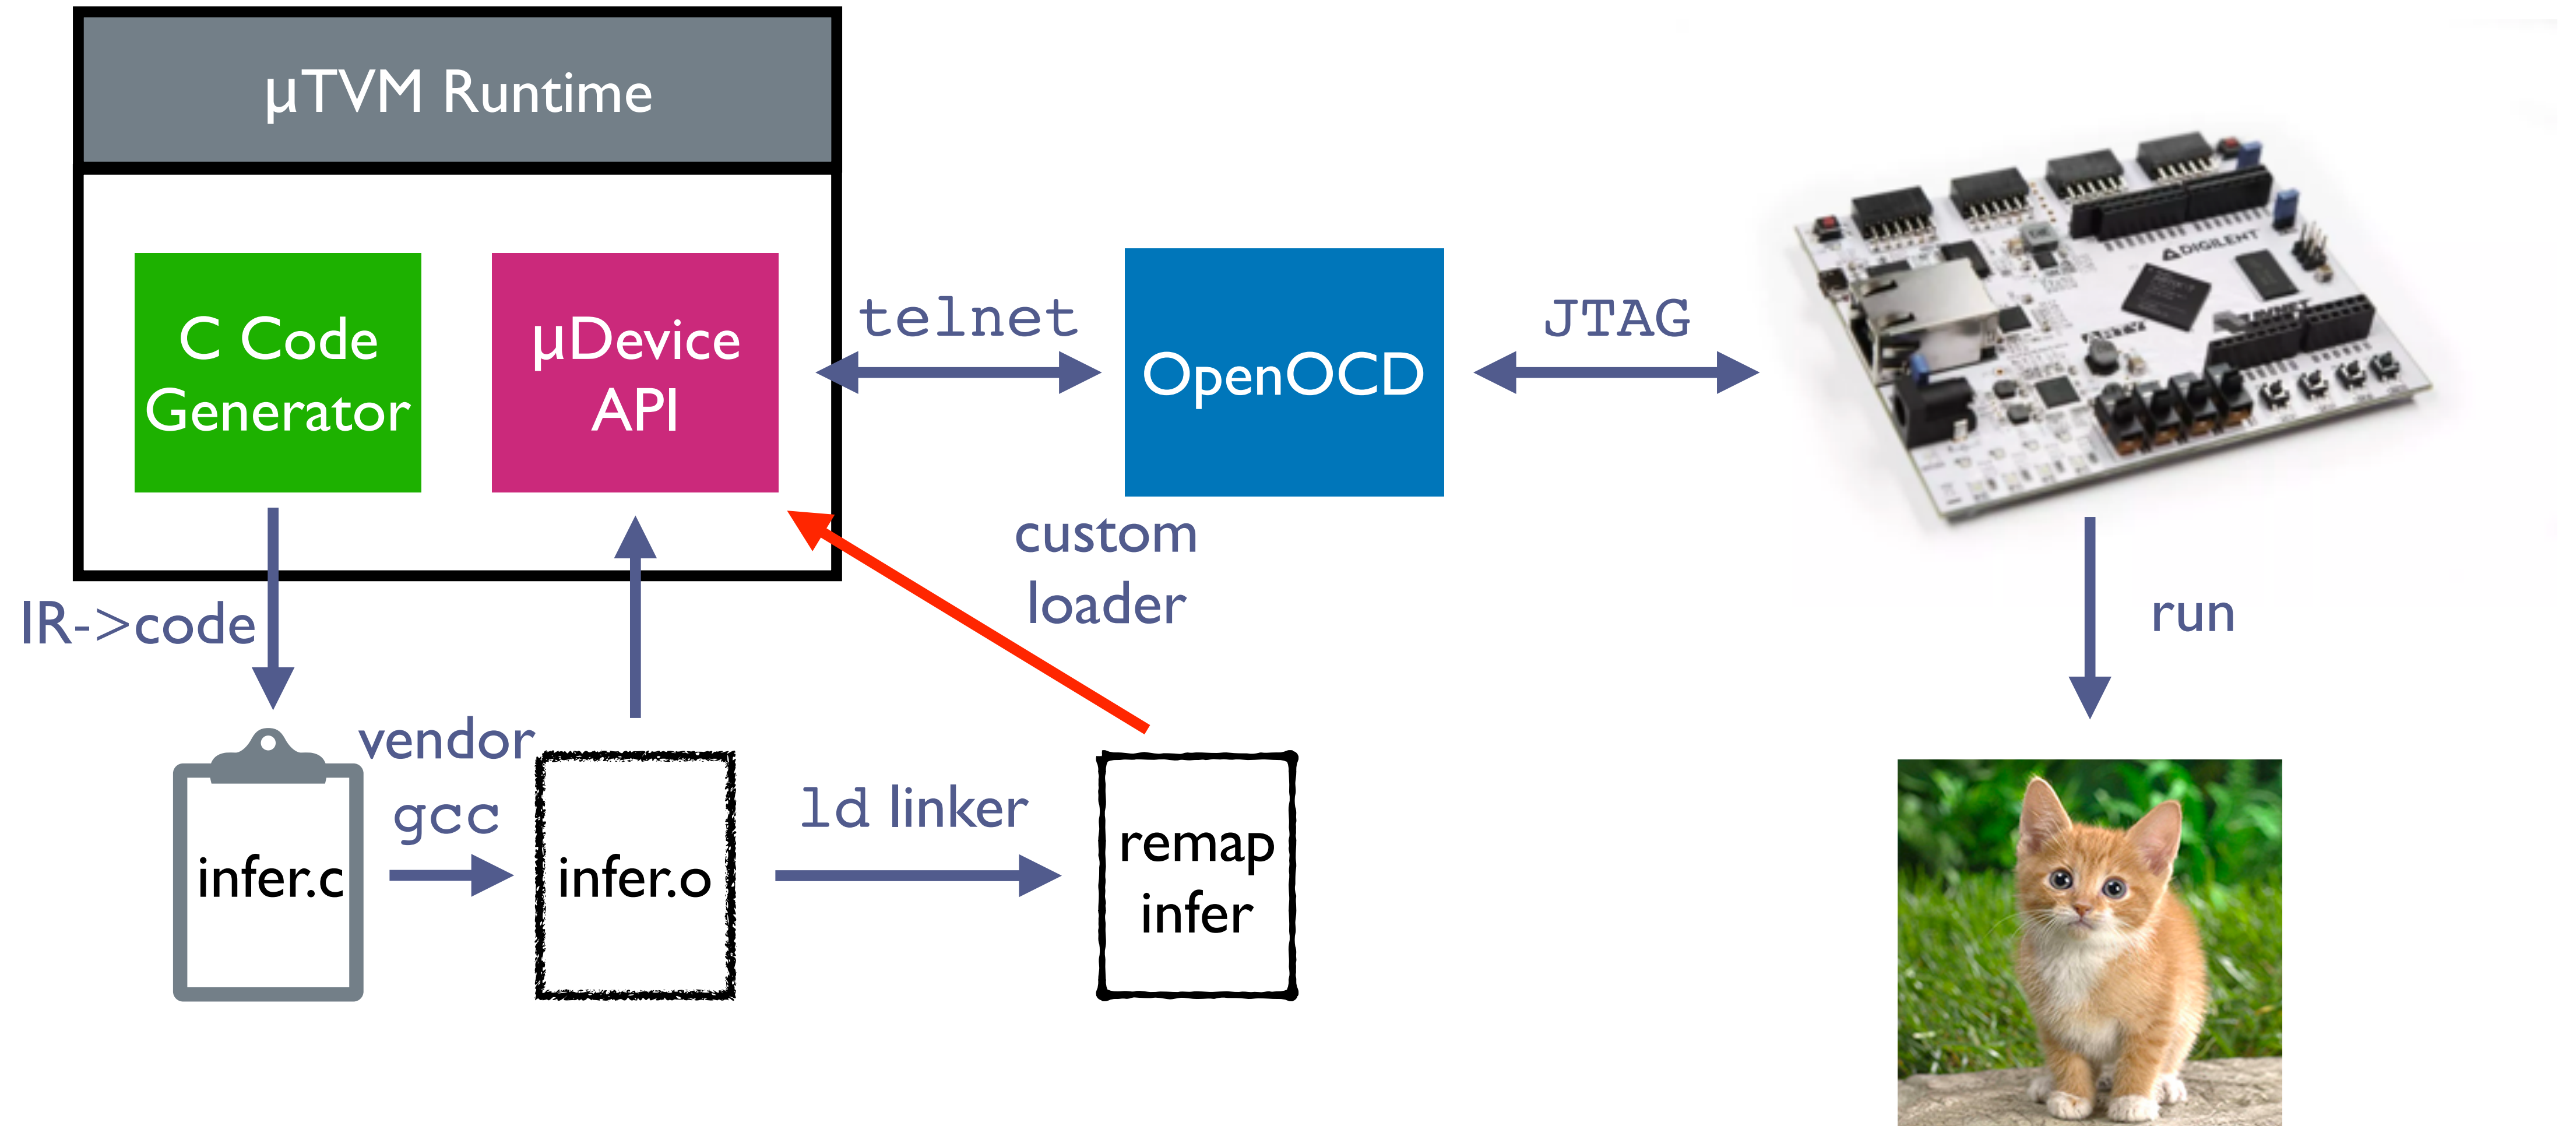
A closer look at μ TVM



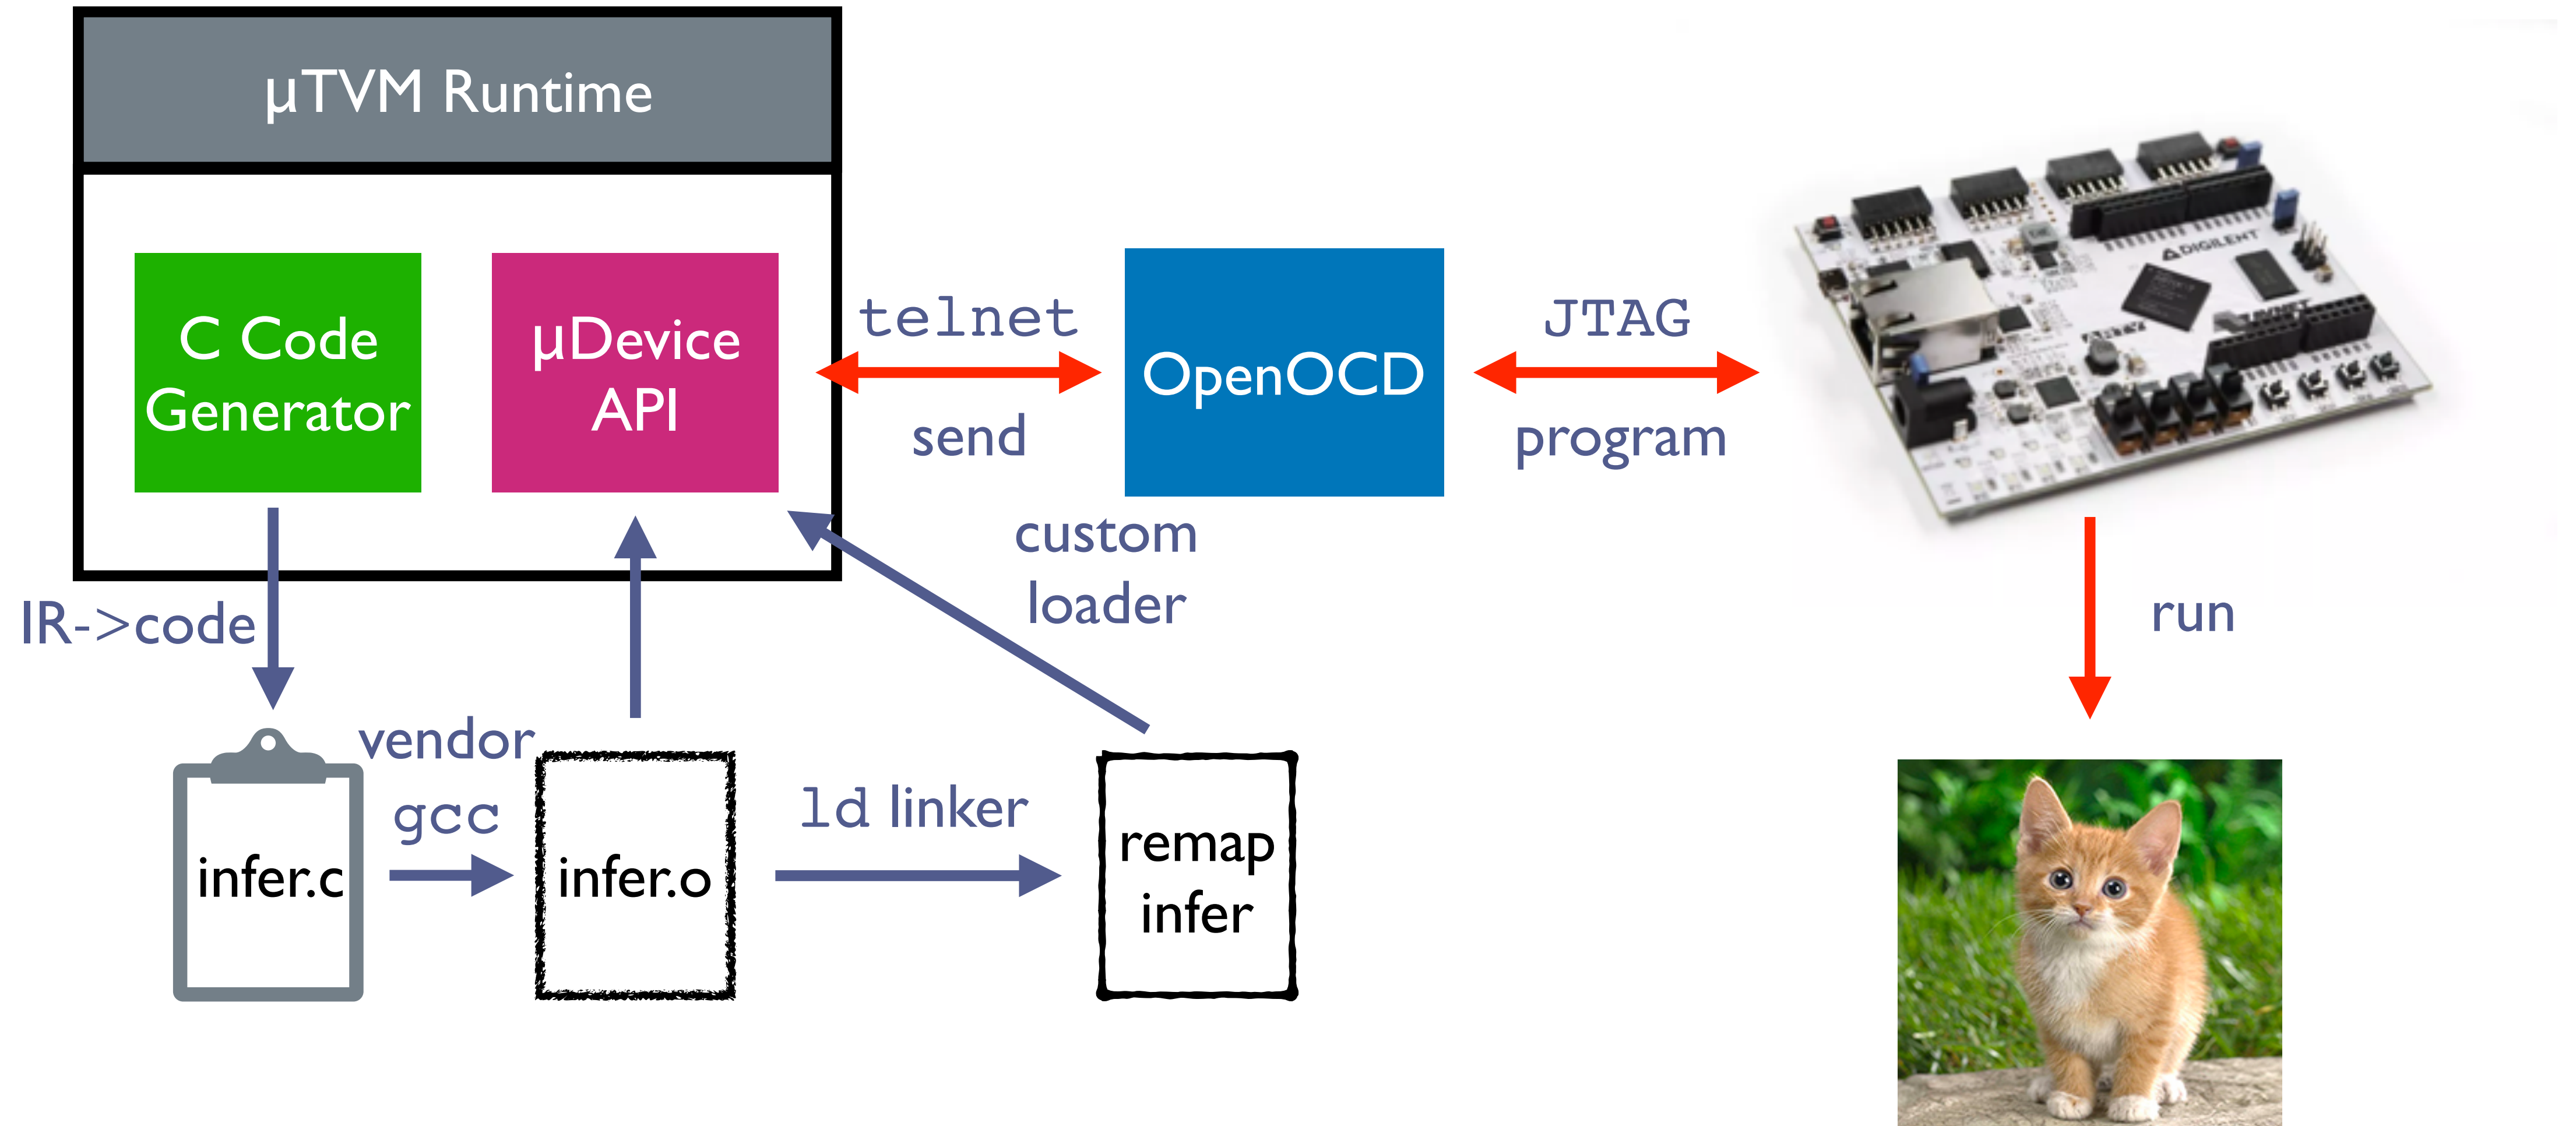
A closer look at μ TVM



A closer look at μ TVM

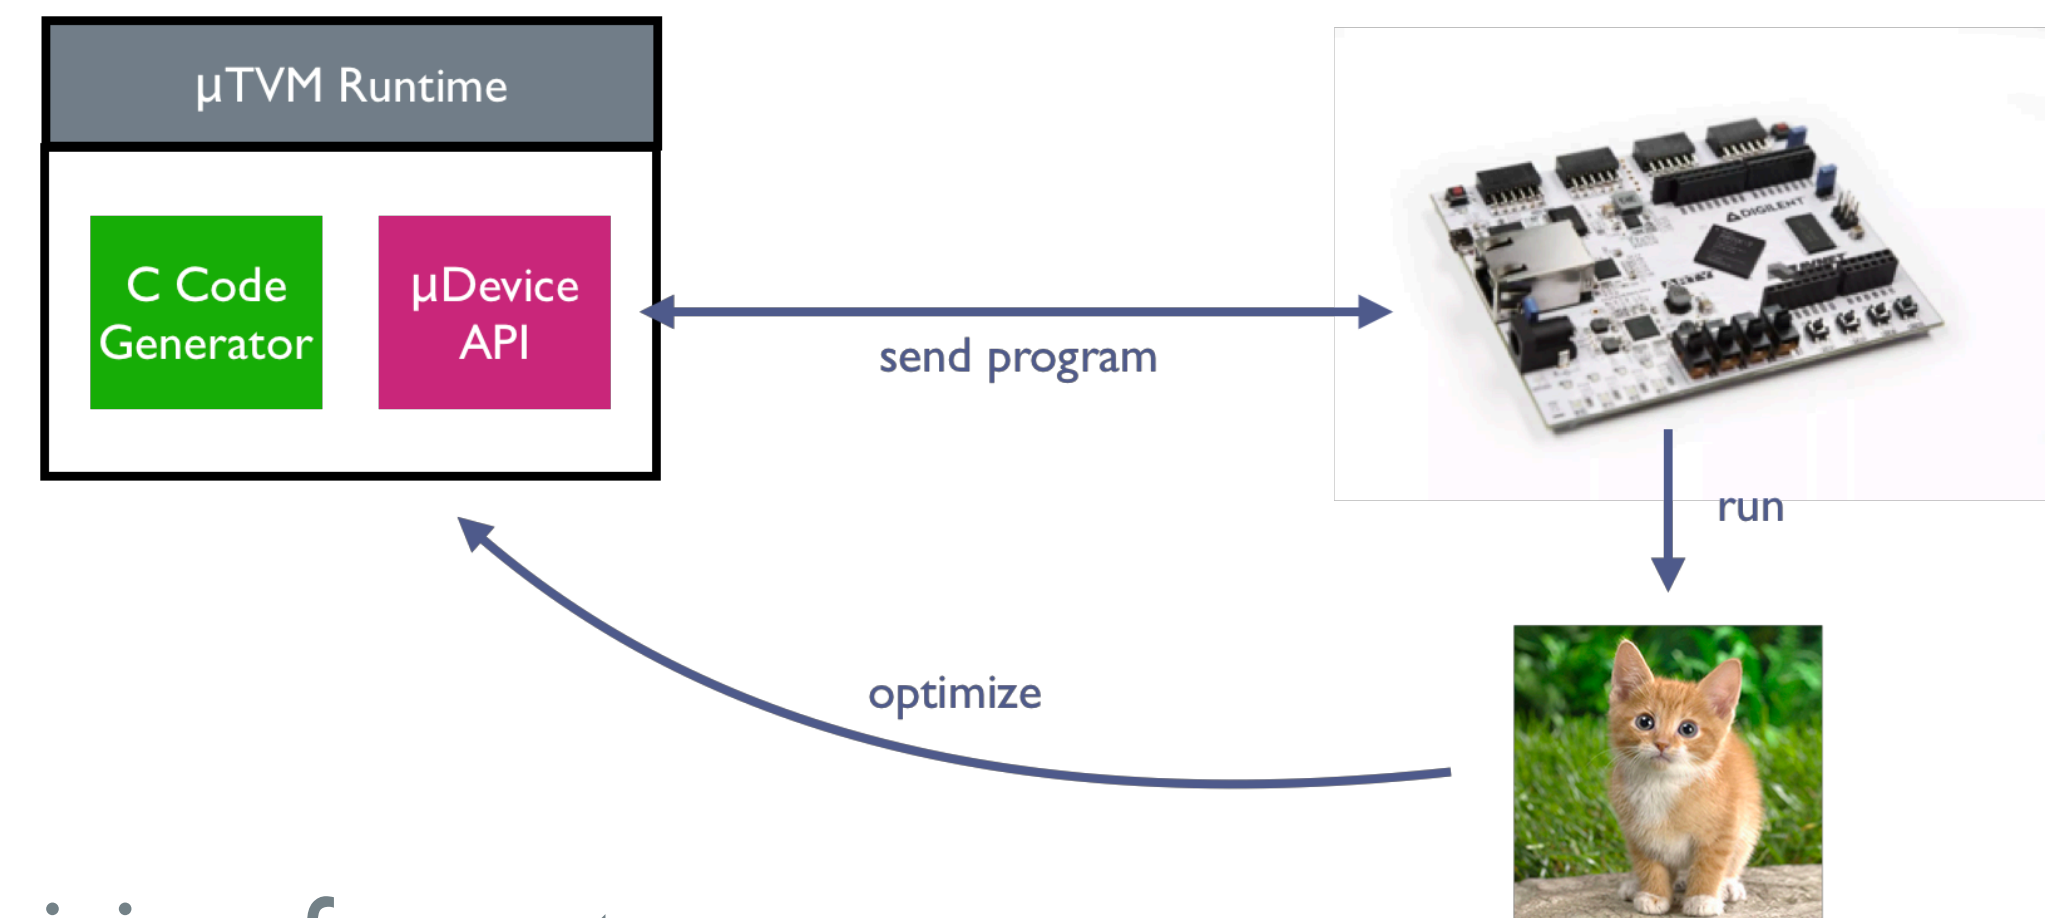


A closer look at μ TVM



Next steps

- Iron out interfaces with actual hardware.
- Optimize with AutoTVM.
- Support restricted and configurable model sizes.
- Enable custom data types such as fixed-point precision formats.



Get in touch!

Pratyush Patel – pate1p1@cs.uw.edu