



References 1/6

- <http://pulp-platform.org> has up to date references
 - Papers: <https://pulp-platform.org/publications.html>
 - Talks: <https://pulp-platform.org/conferences.html>
- **32bit cores (RI5CY, Micro, Zero)**
 - M. Gautschi *et al.*, "Near-Threshold RISC-V Core With DSP Extensions for Scalable IoT Endpoint Devices," in *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 10, pp. 2700-2713, Oct. 2017, doi: [10.1109/TVLSI.2017.2654506](https://doi.org/10.1109/TVLSI.2017.2654506).
 - P. Davide Schiavone *et al.*, "Slow and steady wins the race? A comparison of ultra-low-power RISC-V cores for Internet-of-Things applications," *2017 27th International Symposium on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, Thessaloniki, 2017, pp. 1-8, doi: [10.1109/PATMOS.2017.8106976](https://doi.org/10.1109/PATMOS.2017.8106976).



References 2/6

■ ARIANE 64bit core

- F. Zaruba and L. Benini, "The Cost of Application-Class Processing: Energy and Performance Analysis of a Linux-Ready 1.7-GHz 64-Bit RISC-V Core in 22-nm FDSOI Technology," in *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 11, pp. 2629-2640, Nov. 2019, doi: [10.1109/TVLSI.2019.2926114](https://doi.org/10.1109/TVLSI.2019.2926114).

■ ARA Vector Processor

- M. Cavalcante, F. Schuiki, F. Zaruba, M. Schaffner and L. Benini, "Ara: A 1-GHz+ Scalable and Energy-Efficient RISC-V Vector Processor With Multiprecision Floating-Point Support in 22-nm FD-SOI," in *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 28, no. 2, pp. 530-543, Feb. 2020, doi: [10.1109/TVLSI.2019.2950087](https://doi.org/10.1109/TVLSI.2019.2950087).



References 3/6

■ NTX Network training accelerator

- F. Schuiki, M. Schaffner, F. K. Gürkaynak and L. Benini, "A Scalable Near-Memory Architecture for Training Deep Neural Networks on Large In-Memory Datasets," in *IEEE Transactions on Computers*, vol. 68, no. 4, pp. 484-497, 1 April 2019, doi: [10.1109/TC.2018.2876312](https://doi.org/10.1109/TC.2018.2876312)

■ Transprecision Floating Point Unit

- S. Mach, D. Rossi, G. Tagliavini, A. Marongiu and L. Benini, "A Transprecision Floating-Point Architecture for Energy-Efficient Embedded Computing," *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, Florence, 2018, pp. 1-5, doi: [10.1109/ISCAS.2018.8351816](https://doi.org/10.1109/ISCAS.2018.8351816). (extended version: <https://arxiv.org/abs/2007.01530>)



References 4/6

■ PULP Cluster and SoCs

- D. Rossi et al. "A 60 GOPS/W, -1.8V to 0.9V body bias ULP cluster in 28nm UTBB FD-SOI technology", in Solid-State Electronics, vol. 117, pp. 170-184 2016, doi: [10.1016/j.sse.2015.11.015](https://doi.org/10.1016/j.sse.2015.11.015).
- D. Rossi et al., "Energy-Efficient Near-Threshold Parallel Computing: The PULPv2 Cluster," in IEEE Micro, vol. 37, no. 5, pp. 20-31, September/October 2017, doi: [10.1109/MM.2017.3711645](https://doi.org/10.1109/MM.2017.3711645).
- F. Conti et al., "An IoT Endpoint System-on-Chip for Secure and Energy-Efficient Near-Sensor Analytics," in IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 64, no. 9, pp. 2481-2494, Sept. 2017, doi: [10.1109/TCSI.2017.2698019](https://doi.org/10.1109/TCSI.2017.2698019).
- A. Pullini et al., "Mr.Wolf: An Energy-Precision Scalable Parallel Ultra Low Power SoC for IoT Edge Processing," in IEEE Journal of Solid-State Circuits, vol. 54, no. 7, pp. 1970-1981, July 2019, doi: [10.1109/JSSC.2019.2912307](https://doi.org/10.1109/JSSC.2019.2912307).
- F. Zaruba, F. Schuiki, S. Mach and L. Benini, "The Floating Point Trinity: A Multi-modal Approach to Extreme Energy-Efficiency and Performance," *2019 26th IEEE International Conference on Electronics, Circuits and Systems (ICECS)*, Genoa, Italy, 2019, pp. 767-770, doi: [10.1109/ICECS46596.2019.8964820](https://doi.org/10.1109/ICECS46596.2019.8964820).



References 5/6

■ PULP IPs DMA, Synchronization and event unit

- Davide Rossi, Igor Loi, Germain Haugou, and Luca Benini. 2014. Ultra-low-latency lightweight DMA for tightly coupled multi-core clusters. In Proceedings of the 11th ACM Conference on Computing Frontiers (CF '14). Association for Computing Machinery, New York, NY, USA, Article 15, 1–10. DOI: <https://doi.org/10.1145/2597917.2597922>
- F. Glaser, G. Haugou, D. Rossi, Q. Huang and L. Benini, "Hardware-Accelerated Energy-Efficient Synchronization and Communication for Ultra-Low-Power Tightly Coupled Clusters," *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, Florence, Italy, 2019, pp. 552-557, doi: [10.23919/DATE.2019.8715266](https://doi.org/10.23919/DATE.2019.8715266).



References 6/6

■ Integrating accelerators in PULP

- F. Conti and L. Benini, "A ultra-low-energy convolution engine for fast brain-inspired vision in multicore clusters," *2015 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, Grenoble, 2015, pp. 683-688, doi: [10.7873/DATE.2015.0404](https://doi.org/10.7873/DATE.2015.0404).
- F. Conti, P. D. Schiavone and L. Benini, "XNOR Neural Engine: A Hardware Accelerator IP for 21.6-fJ/op Binary Neural Network Inference," in *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 11, pp. 2940-2951, Nov. 2018, doi: [10.1109/TCAD.2018.2857019](https://doi.org/10.1109/TCAD.2018.2857019).

Documentation on how to do it: <https://hwpe-doc.readthedocs.io/en/latest/>