

Low-latency user-level communication for RISC-V clusters



Charisios Loukas^{1,*}, Pantelis Xirouchakis¹, Michalis Gianioudis¹, Aggelos Ioannou^{1,2}, Manolis Katevenis¹ and Nikos Chrysos¹

¹Computer Architecture and VLSI Systems laboratory, Foundation for Research and Technology - Hellas

²Lawrence Berkeley National Laboratory, University of California

*cloukas@ics.forth.gr



FORTH-ICS
COMPUTER ARCHITECTURE &
VLSI SYSTEMS LABORATORY

Motivation

- Tightly couple Ariane cores with caRVnet Network Interface
- Develop and optimize user level library (USL)
- Evaluate and breakdown the end-to-end latency on the hardware testbed

Methodology

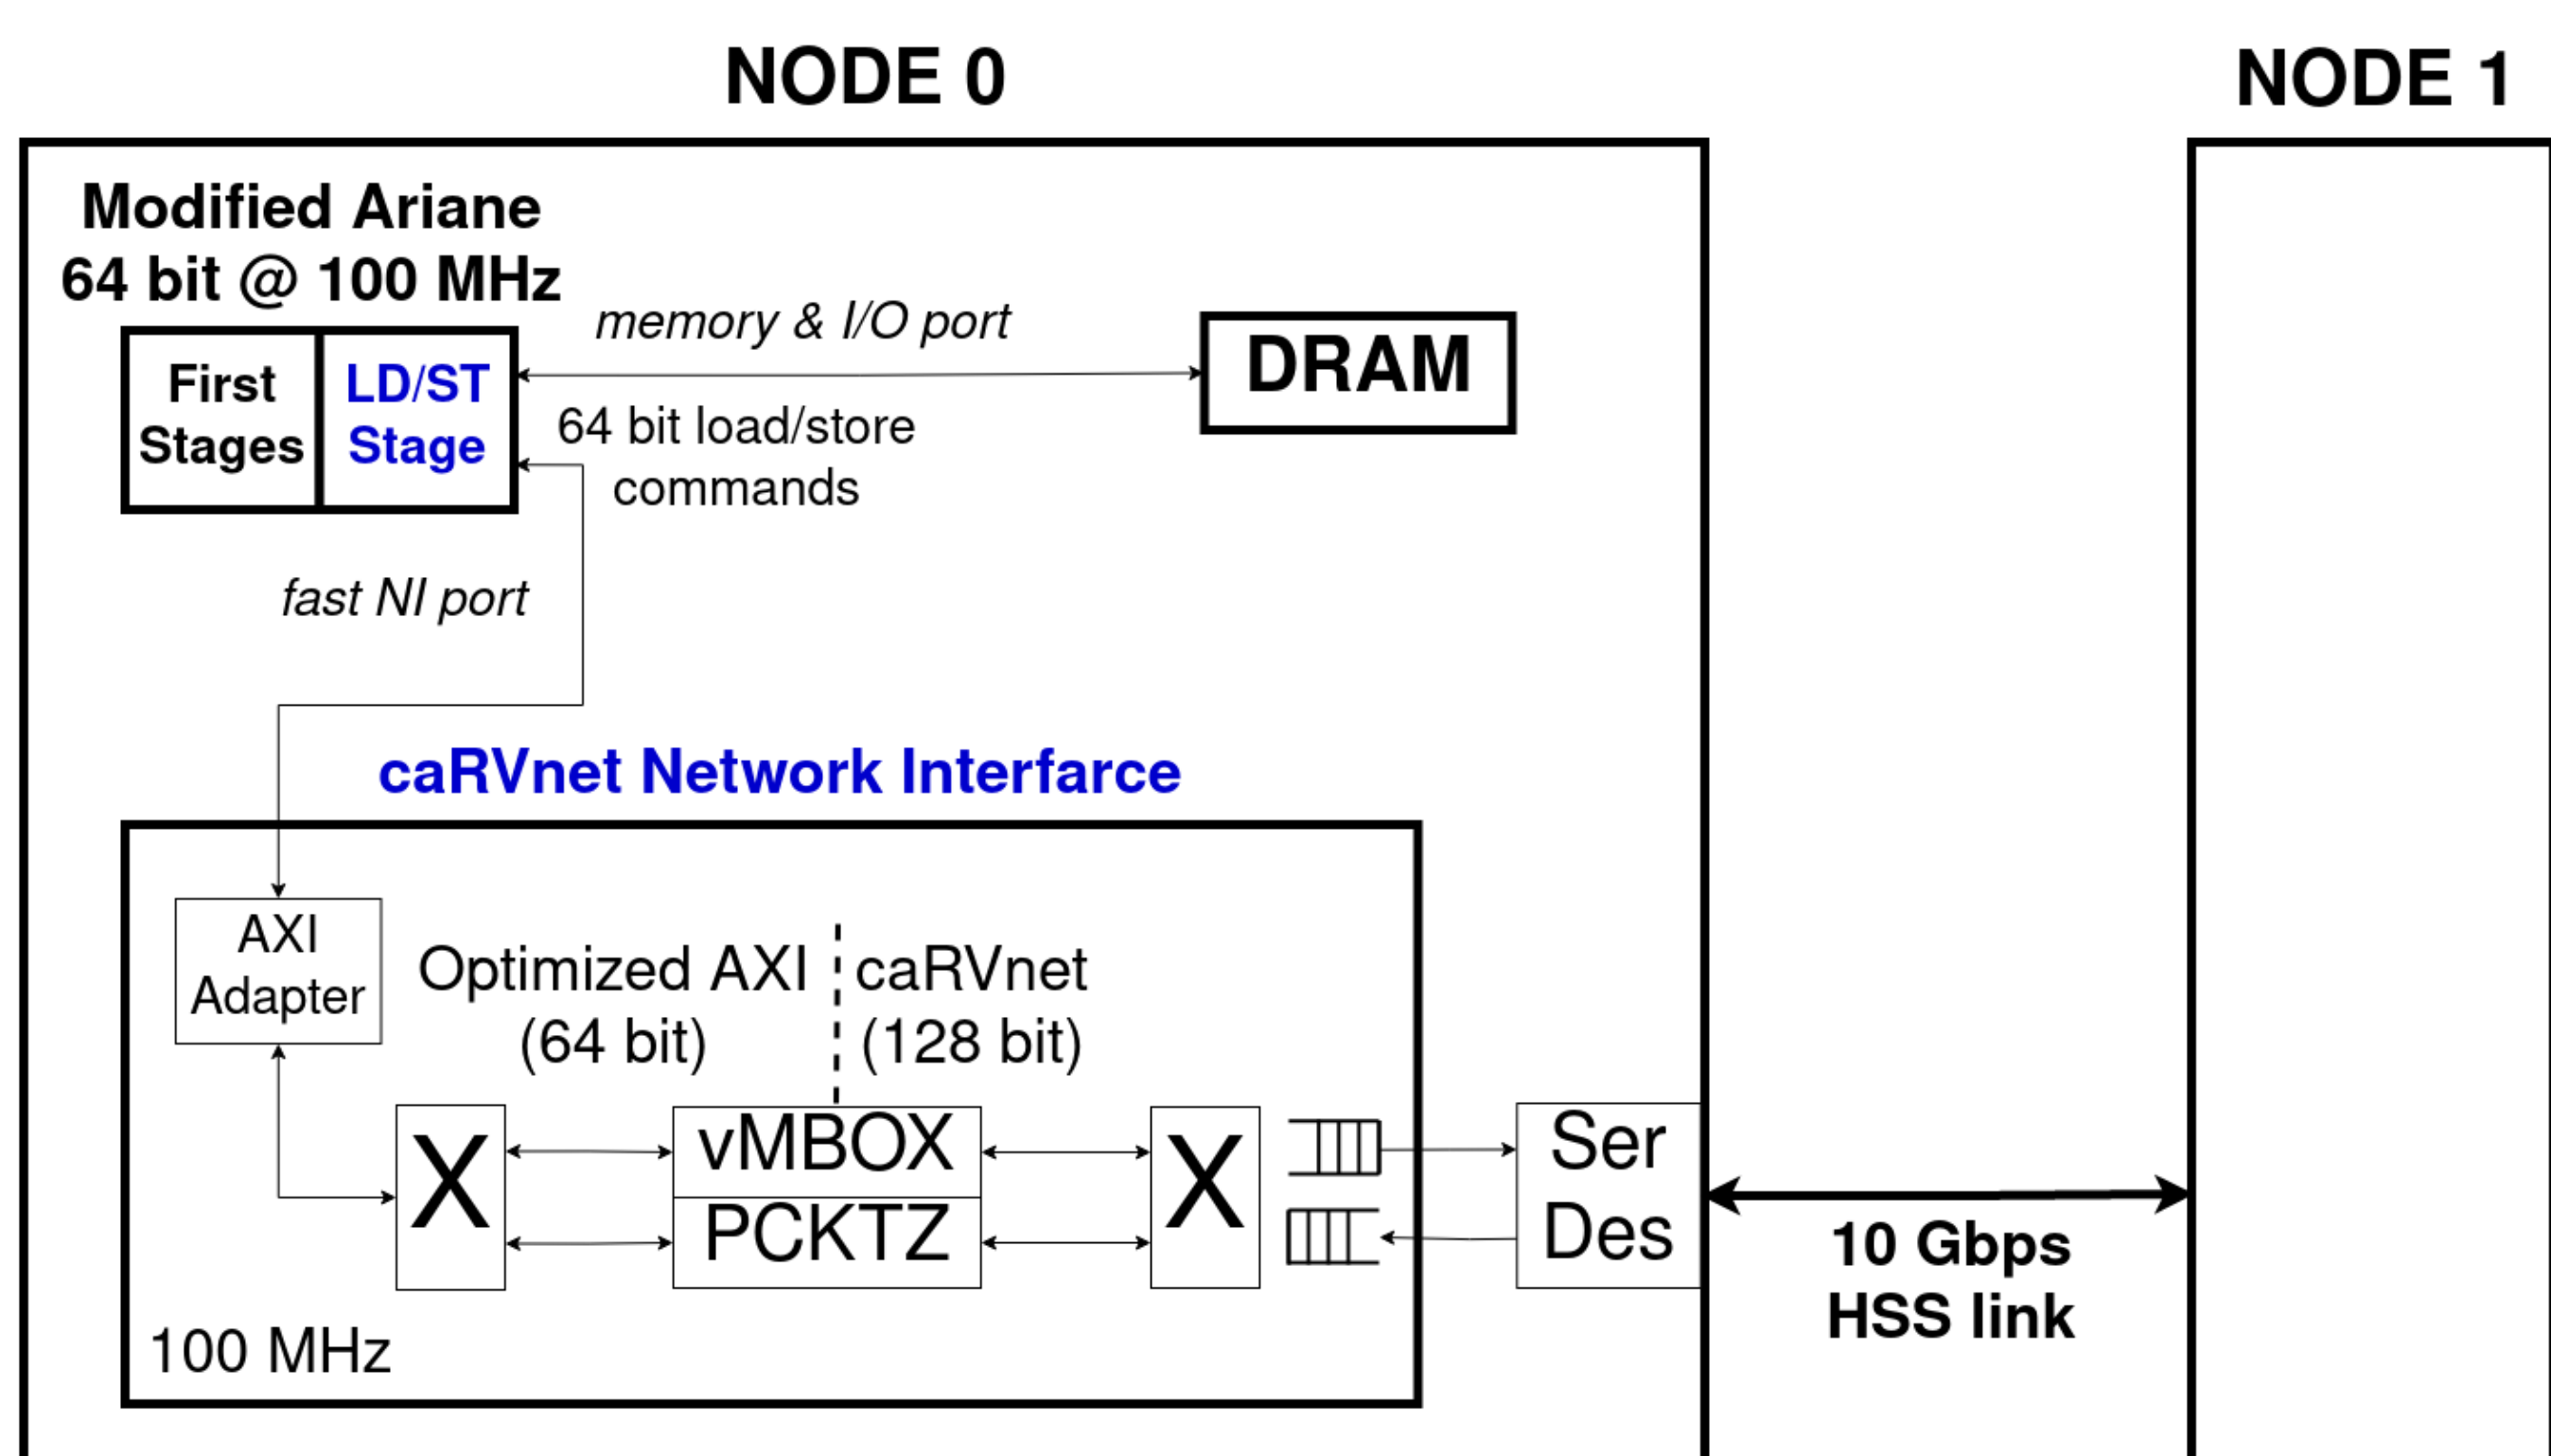


Figure 1: Testbed and hardware architecture

Hardware

- 2 Trenz boards each hosting a Xilinx FPGA (ZU9EG)
- Node hardware architecture: Ariane RISC-V core tightly coupled with caRVnet Network Interface
- Implemented fast NI port inside the Ariane LD/ST stage to allow for b2b stores to NI
- Packetizer and Mailbox: optimized caRVnet endpoints for small messages
- Use caRVnet (128 bit @ 100 MHz) to connect nodes in a Ring 10 Gbps topology
- Developed dedicated hardware mechanism to disable context switch related interrupts in Ariane

Software

- Packetizer/Mailbox drivers (kernel space)

Acknowledgement

This work is part of the RED-SEA project which receives funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955776. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Greece, Germany, Spain, Italy, Switzerland.



- Latency optimized User Space Library for the caRVnet packetizer and mailbox peripherals
- Ping-pong test between adjacent nodes - time measurement with the use of RISC-V timers
- Customized GNU/Linux OS with a shared NFS

Latency Results

- Measurement parameters:
 - 10^6 ping-pong transfer iterations
 - Per iteration user-space to user-space latency measurements
- Baseline average one-way latency: 25 μ s
 - After code optimizations: 930ns
- Average value still does not match 720ns Vivado ILA observations
- Large average value suggests the existence of outliers caused by context switches and pagefaults.

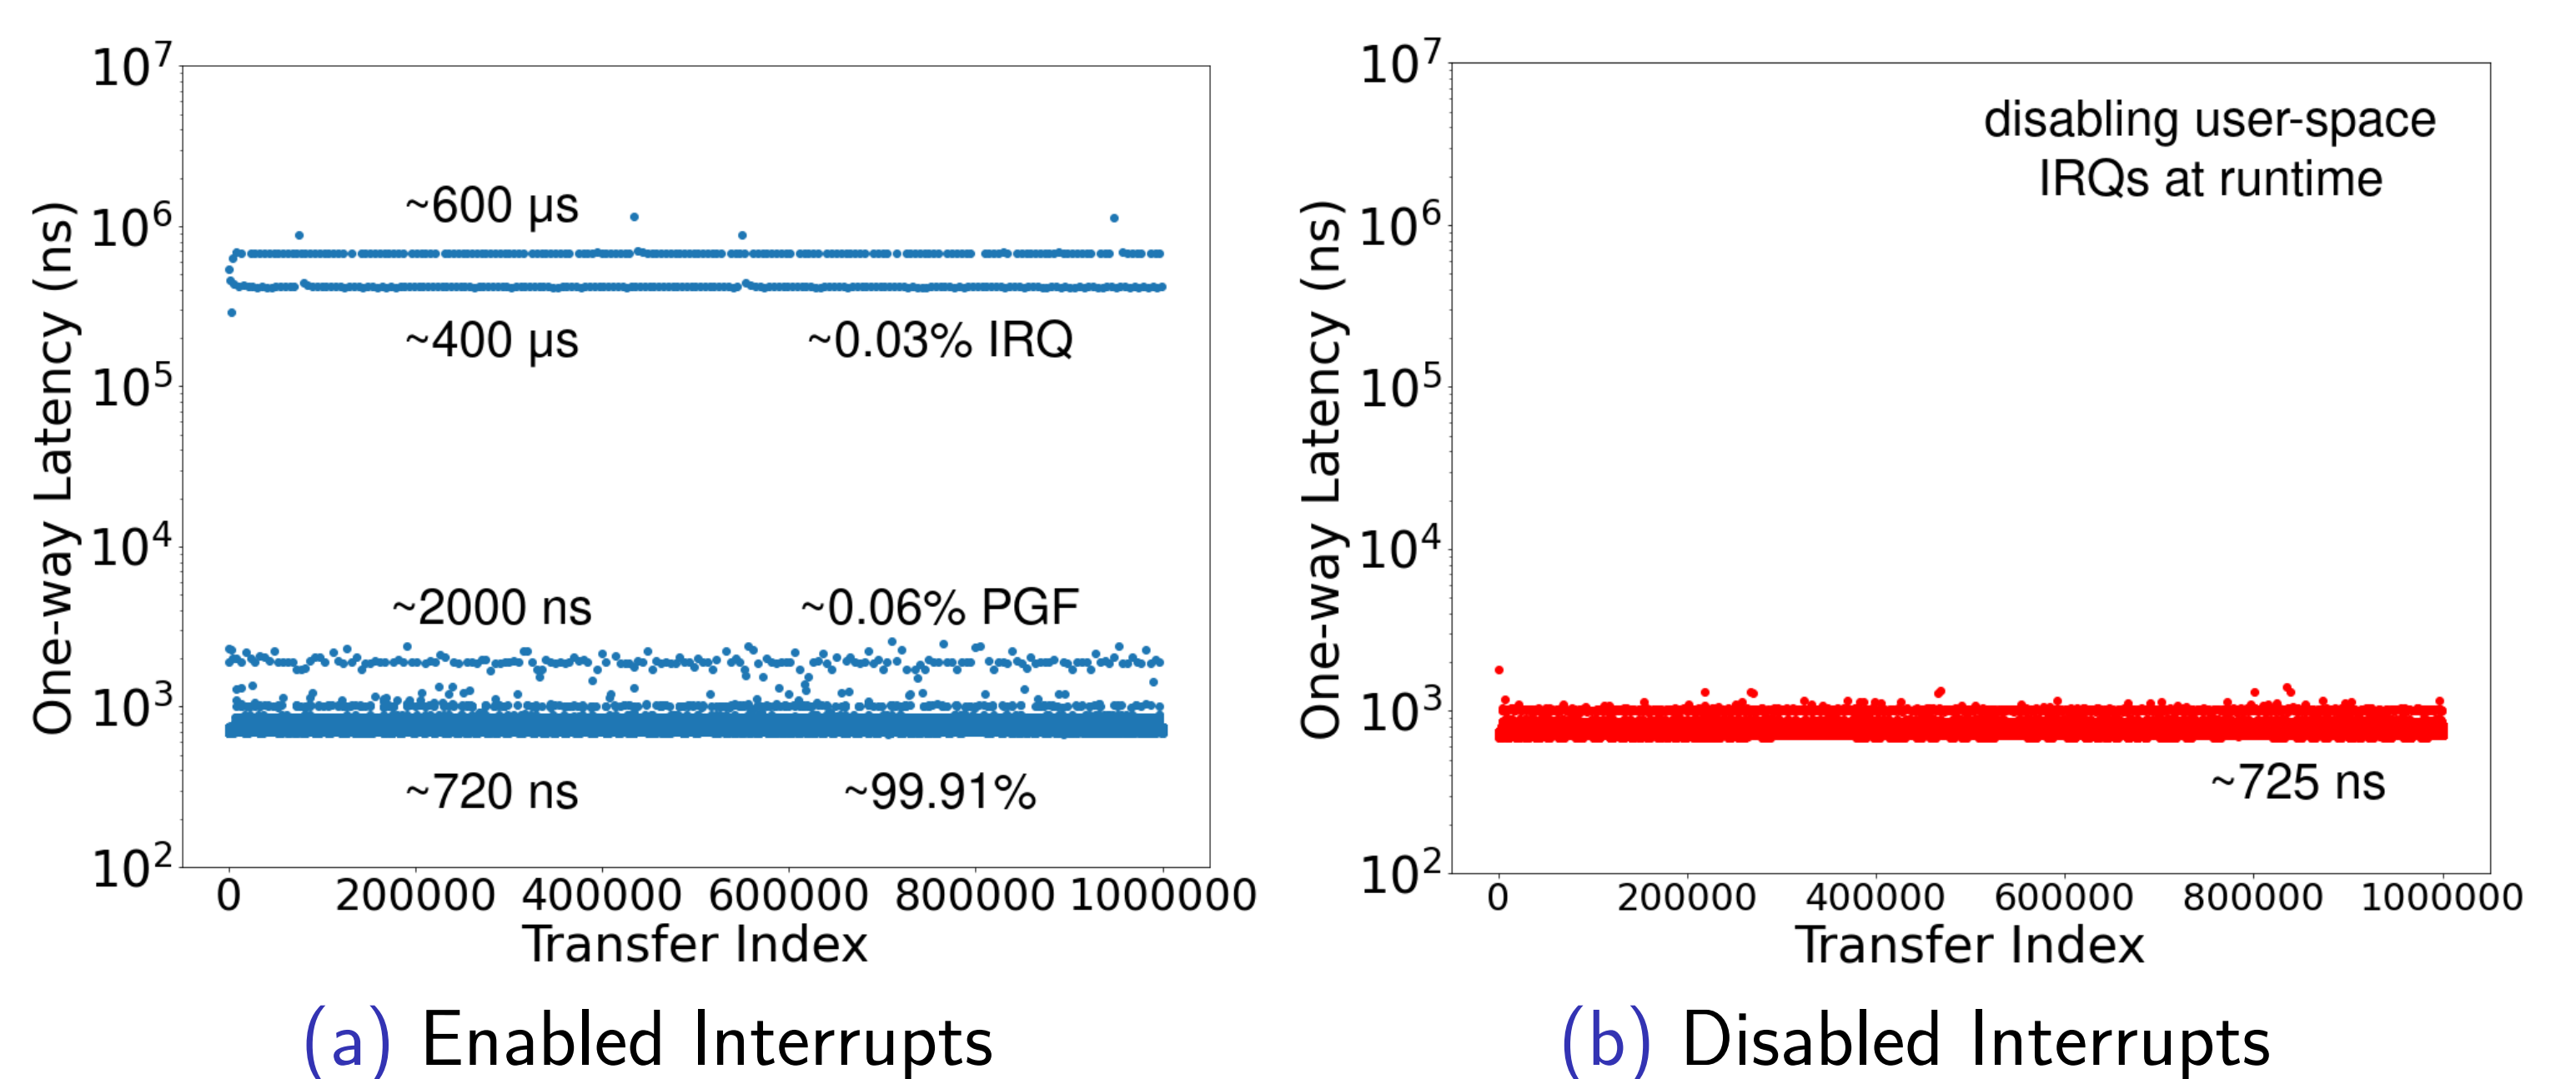


Figure 2: One-way latency measurements

725ns average one-way user-space latency

Breakdown

- 200ns network traverse time
- 180ns hardware time: 6 store (6 CCs), 4 read (8 CCs), network interface 1+1 CC & status poll (2 CCs)
- 330ns software time
- <150ns projected one-way latency on ASIC with RISC-V core @ 1 GHz & caRVnet NI running at 128 Gbps (128 bit, 1 GHz)

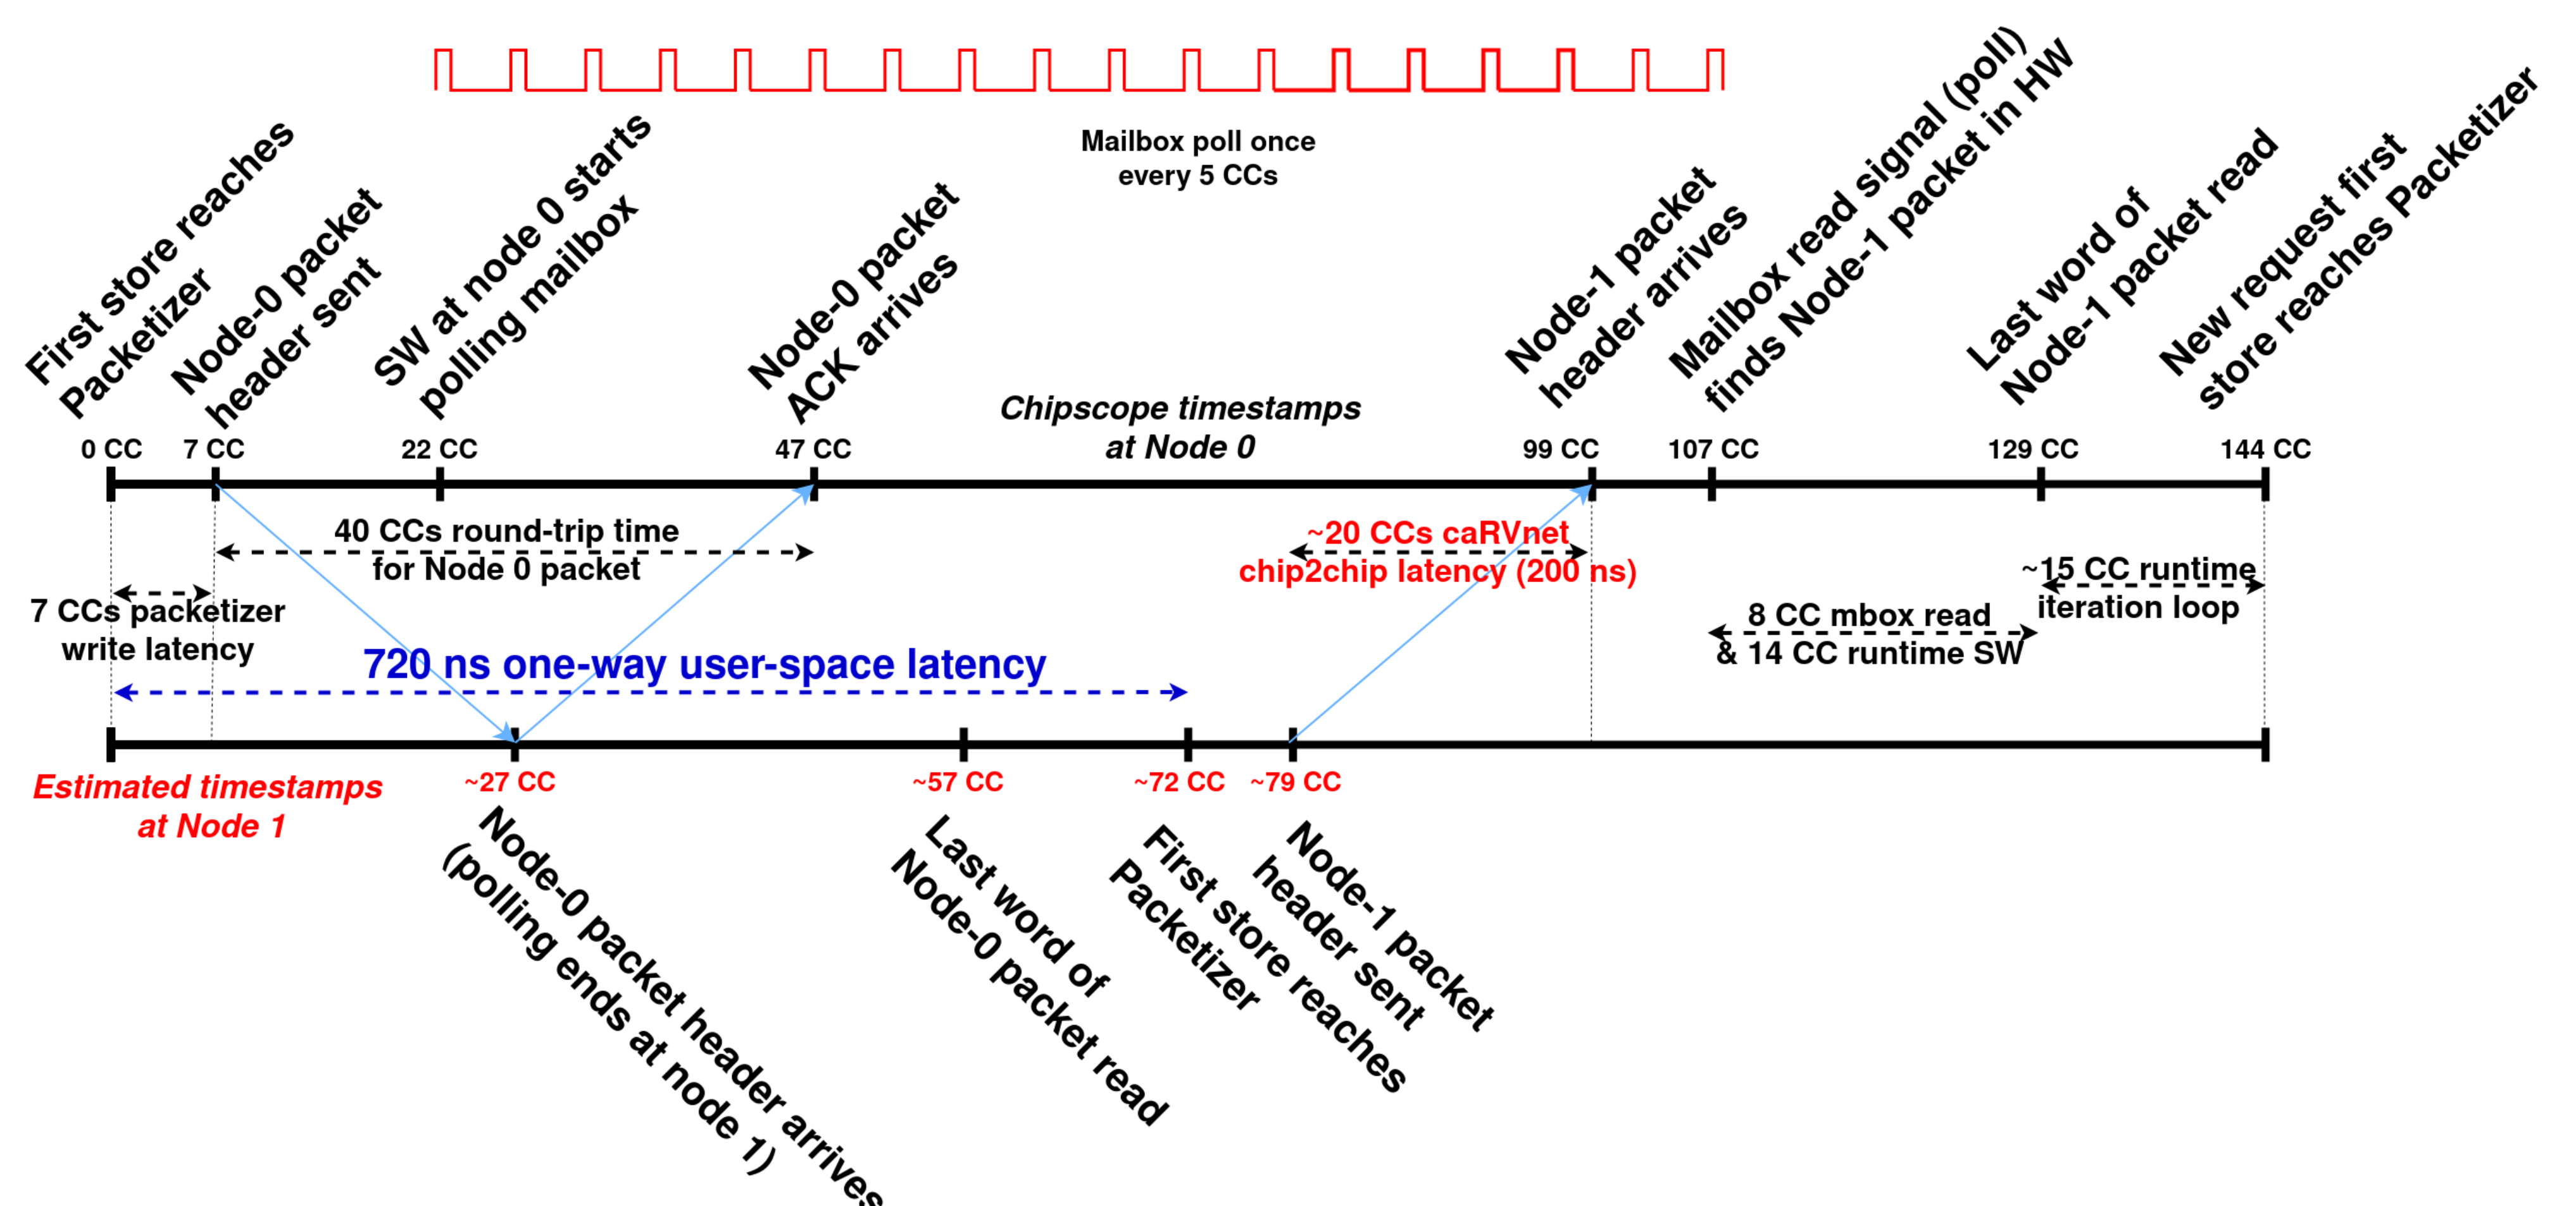


Figure 3: Latency measurement breakdown