

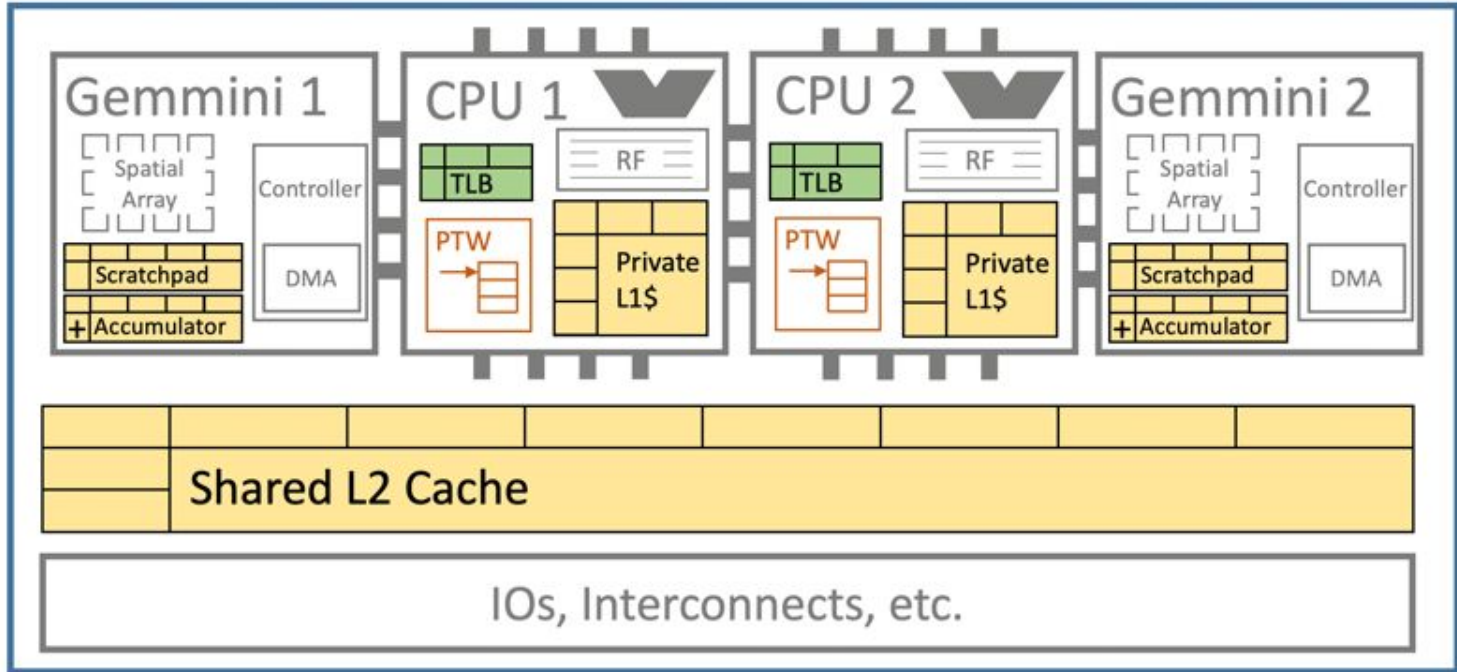
Enabling Systematic Deep-Learning Architecture Evaluation via Full-Stack Integration

Sophia Shao
ysshao@berkeley.edu
Electrical Engineering and Computer Sciences

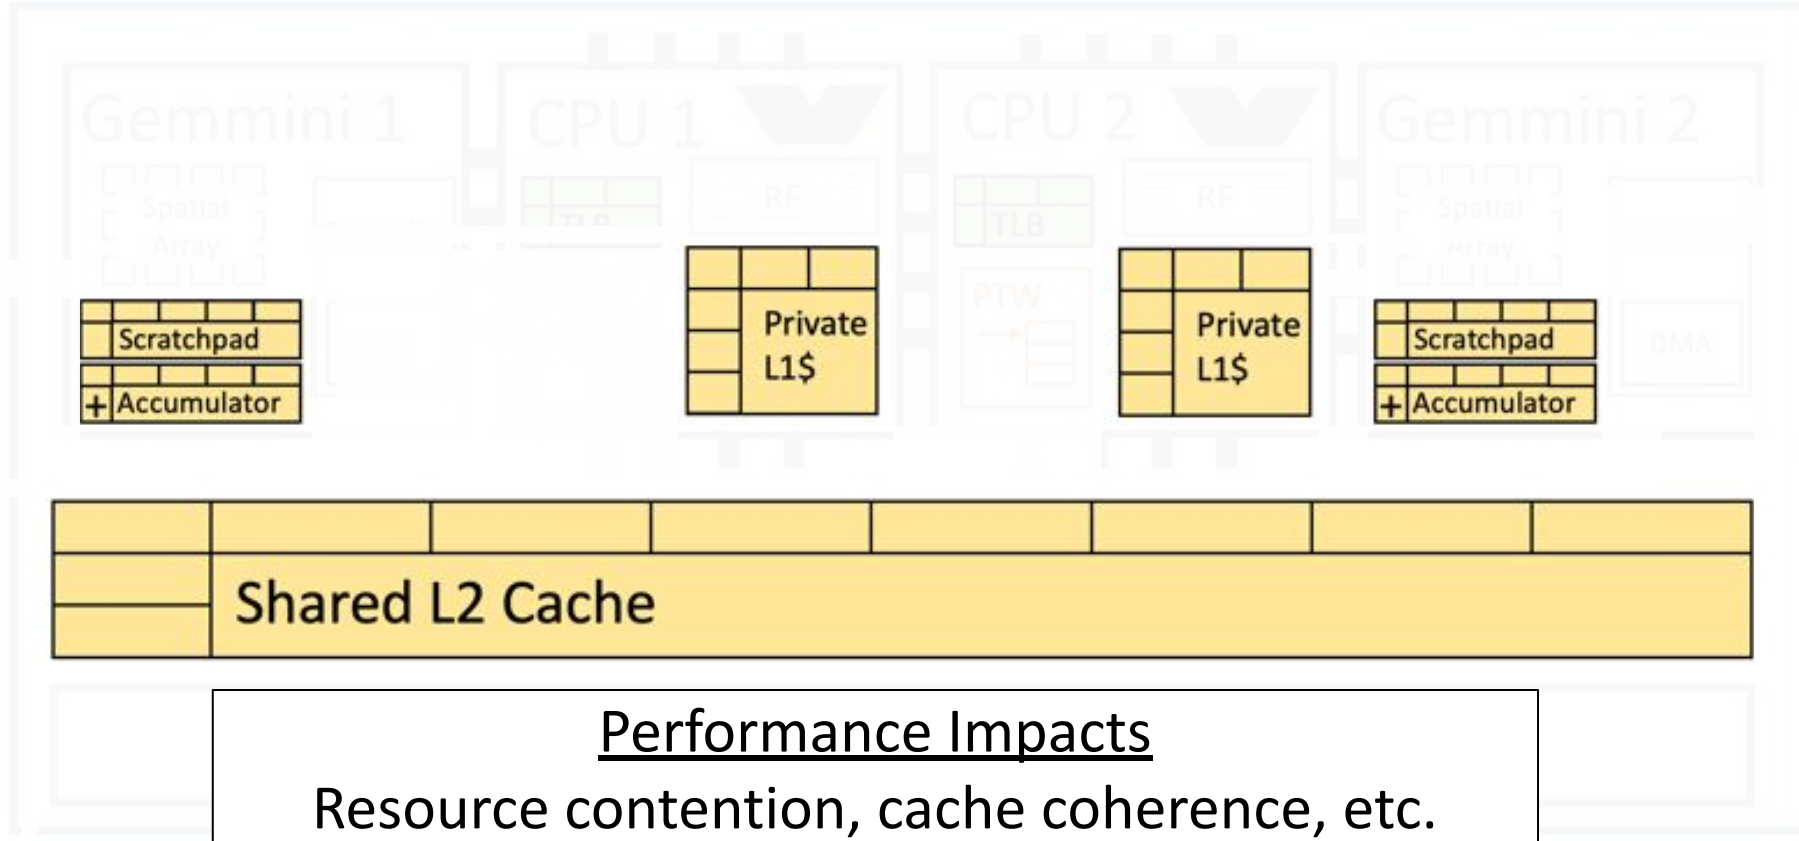


DL accelerators don't exist in isolation.

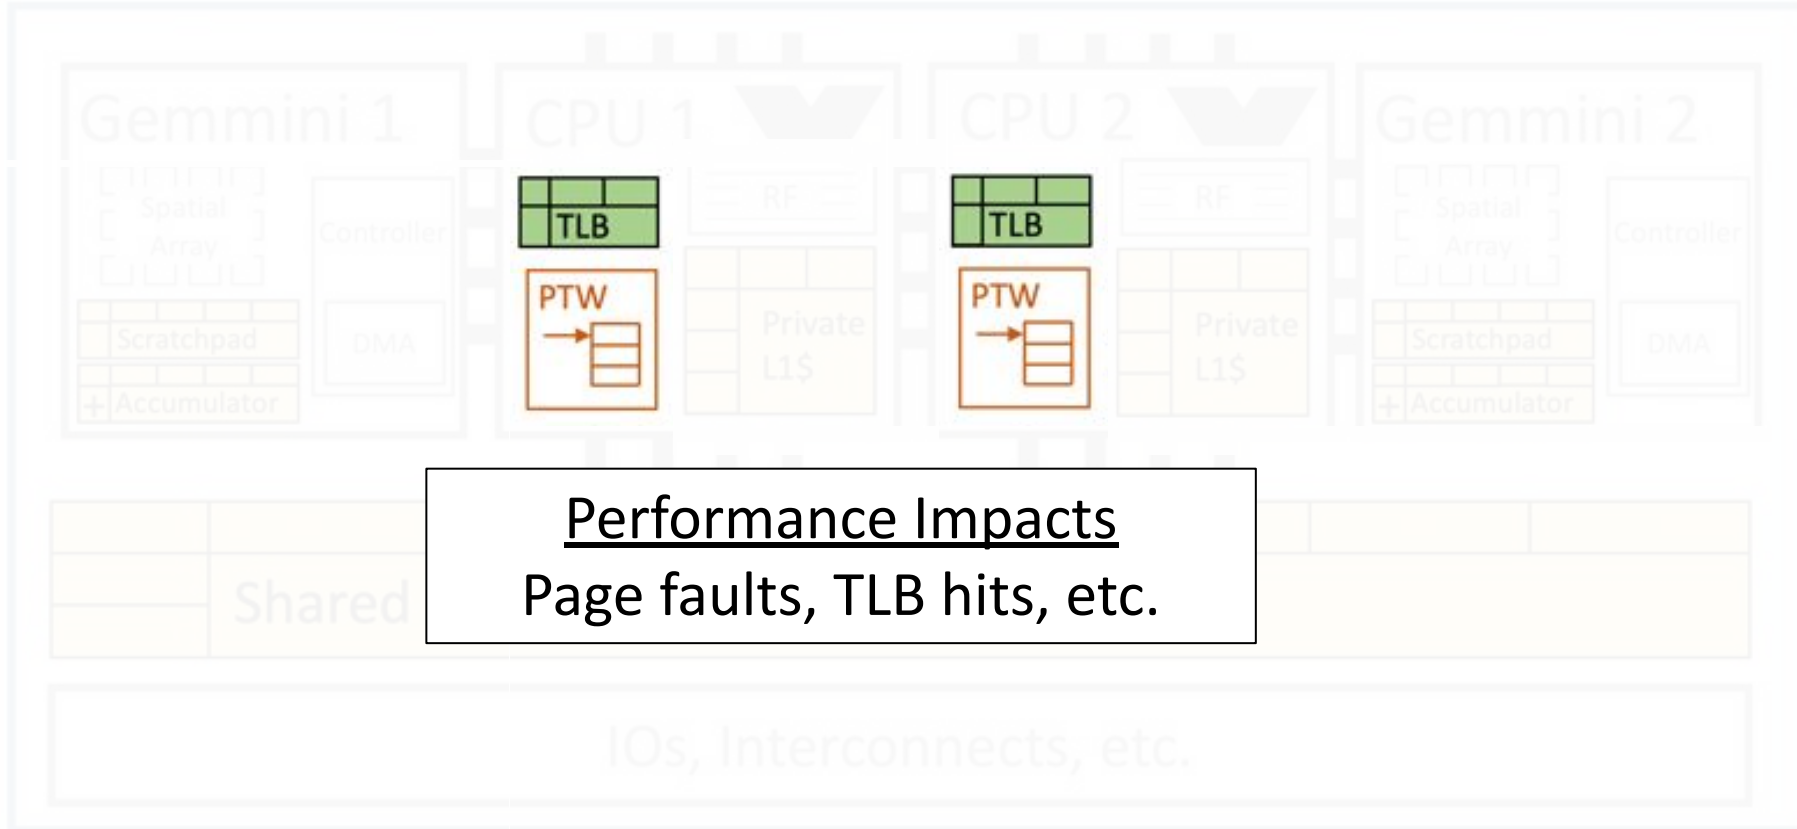
SoC



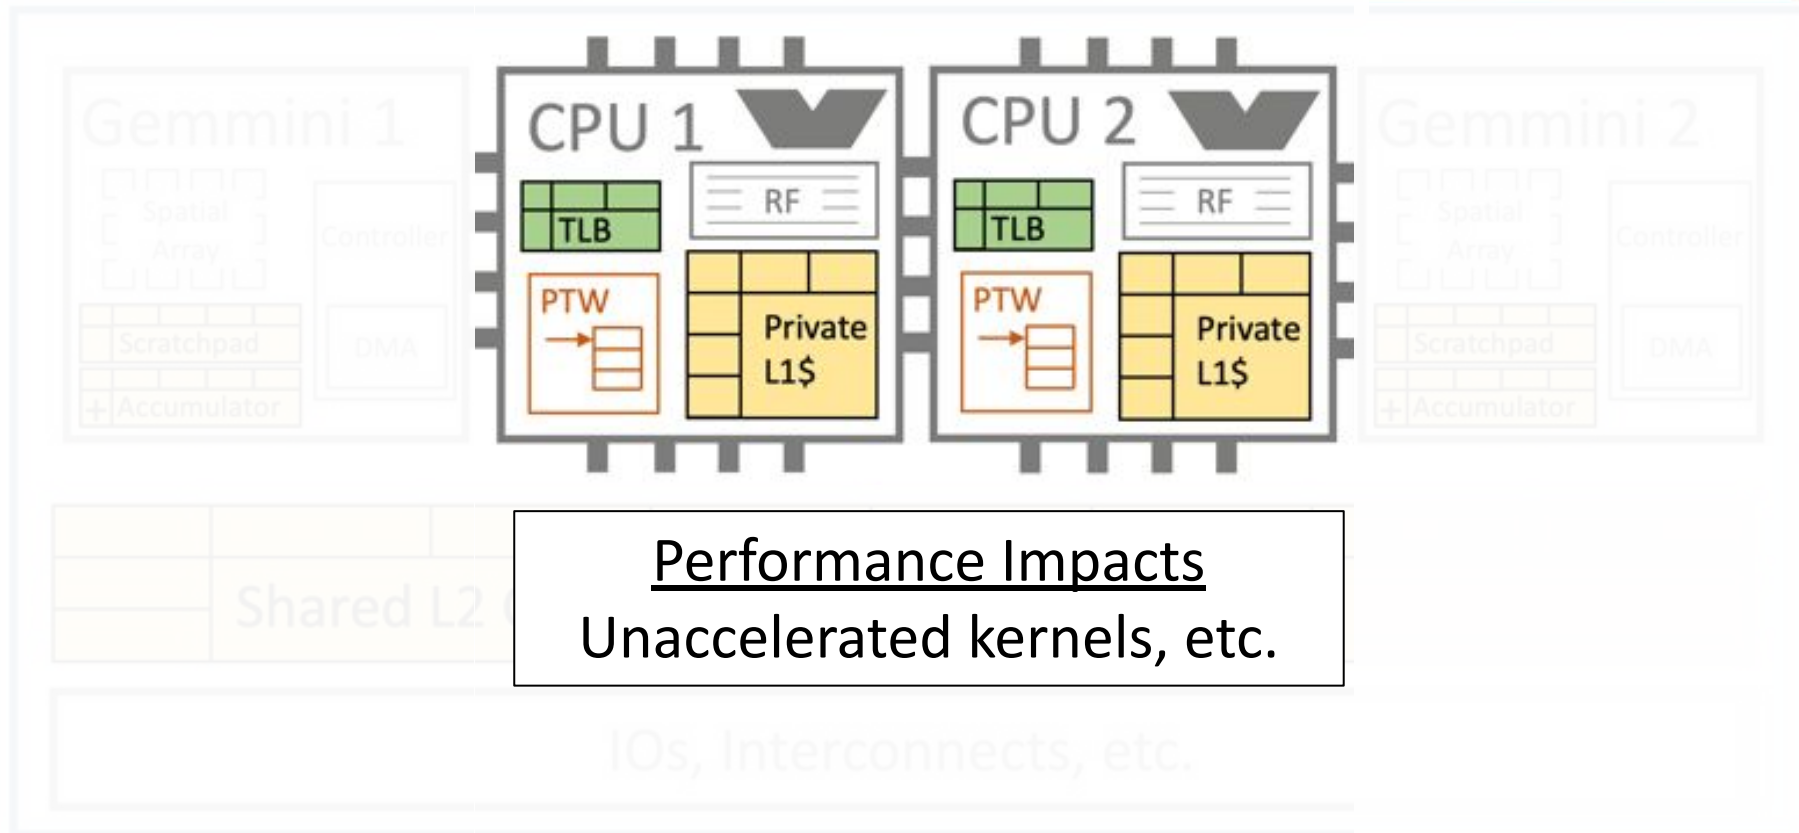
Full-System Visibility: Memory Hierarchy



Full-System Visibility: Virtual Addresses



Full-System Visibility: Host CPUs





Gemmini

- DNN accelerator generator
- Flexible hardware template
- Full-system
- Open-source!
 - <https://github.com/ucb-bar/gemmini>

