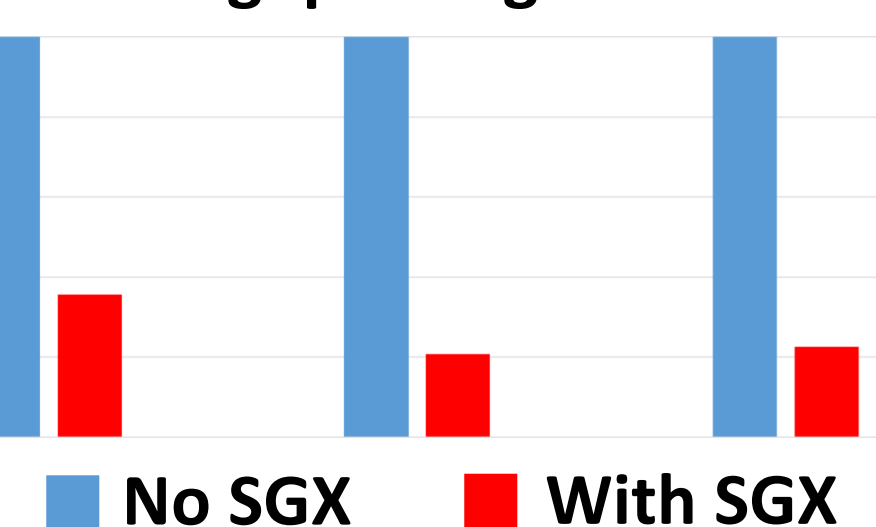


SGX Secure Execution

Authenticated code

Malicious environment

Throughput degradation




Is it **practical**?

- ▶ What is the impact on overall application's performance?
- ▶ What creates the bottlenecks?

Contribution: overcoming the bottlenecks with **HotCalls**

Cloud Attack Surface

- ▶ To lower costs - computation and storage are moved to third party machines
- ▶ This implies trust across the entire software stack



- ▶ And cloud provider employees

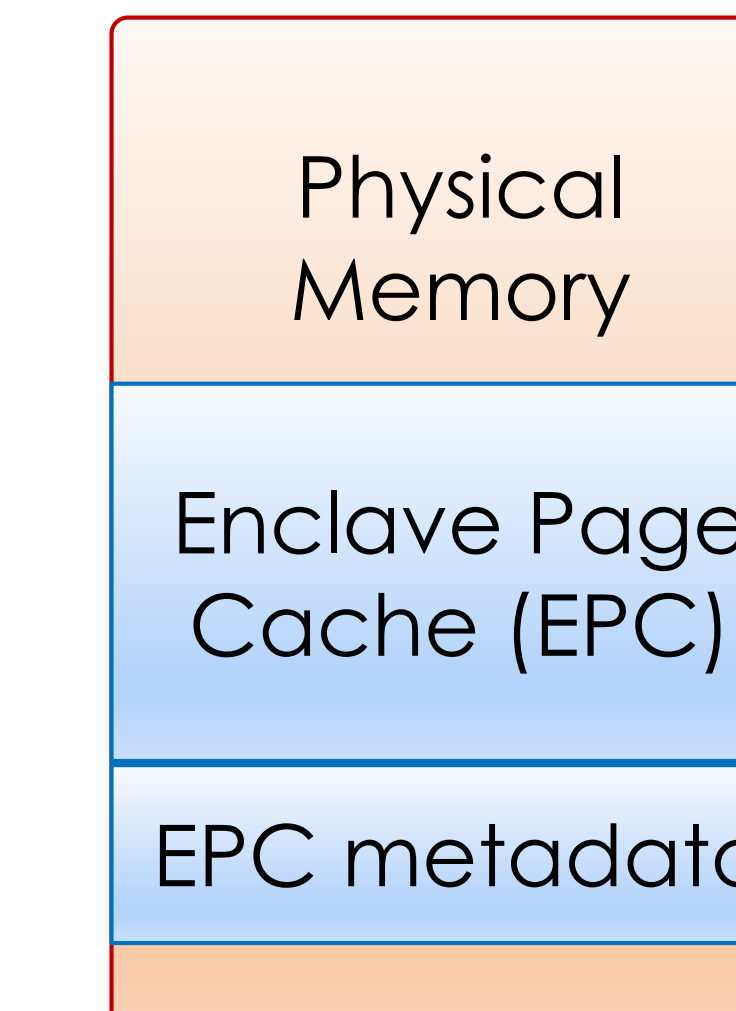
The attack surface is large



SGX in a Nutshell



Memory encrypted on CPU die (via MEE)



Secure enclave is isolated from the system

SGX Life-Cycle

Application - Untrusted Code

ecall

Enclave - Trusted Code

Plaintext Shared Memory

Encrypted Memory

- Can call system API functions (send, fread, etc.)

- Can access all memory
- No access to system calls

ocall

- Ecall: a secure context switch to enclave code
- Ocall: reverse context switch to request OS services

SGX operations may become a bottleneck

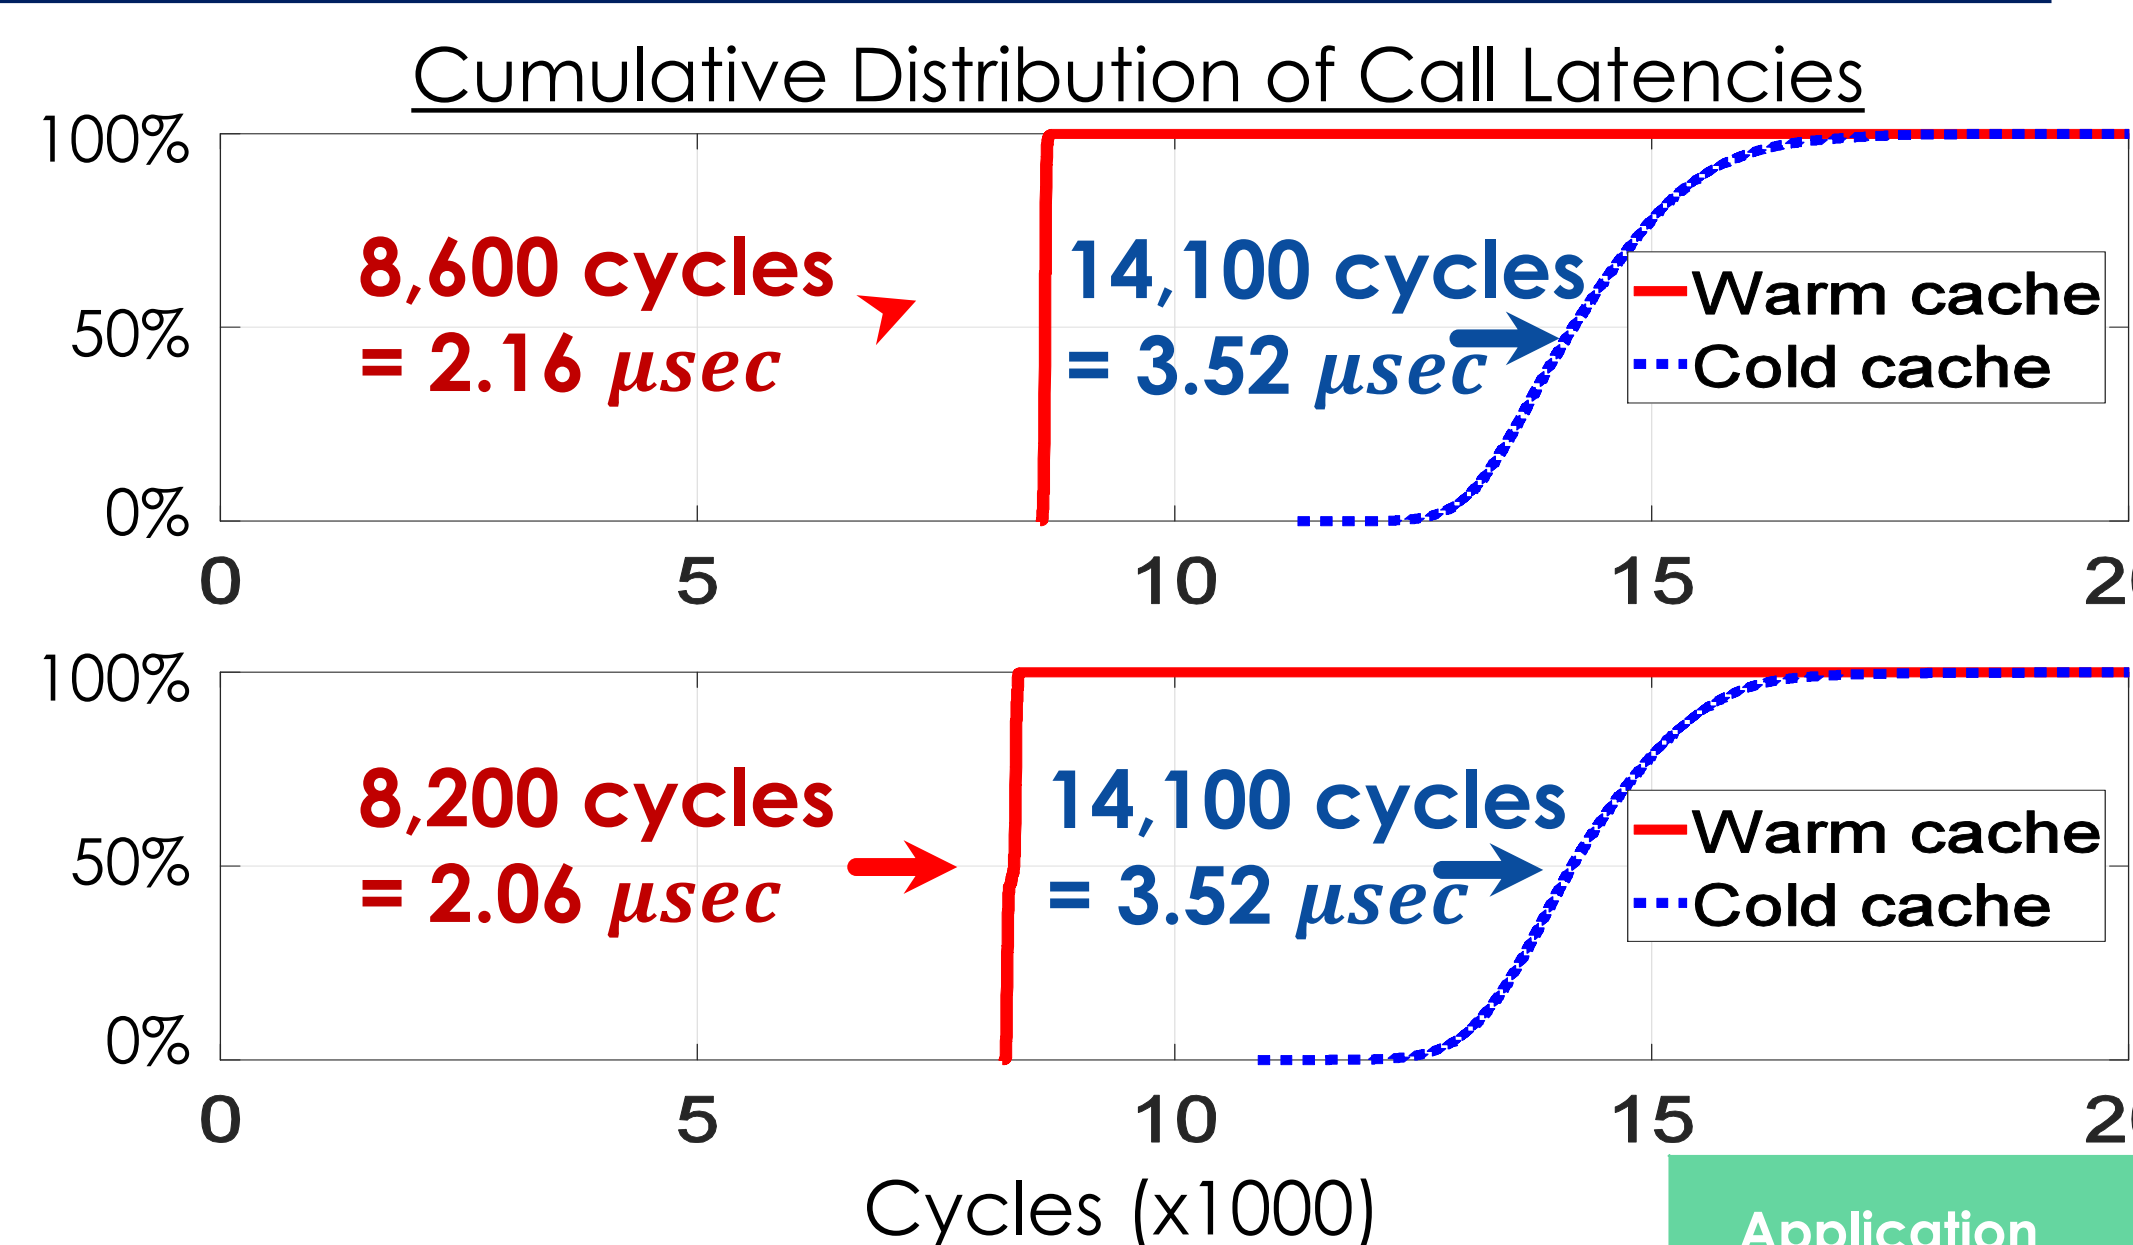
Cost of Ecalls/Ocalls

Ecalls

- SDK code
- EENTER
- EEXIT

Ocalls

- SDK code
- EEXIT
- ERESUME

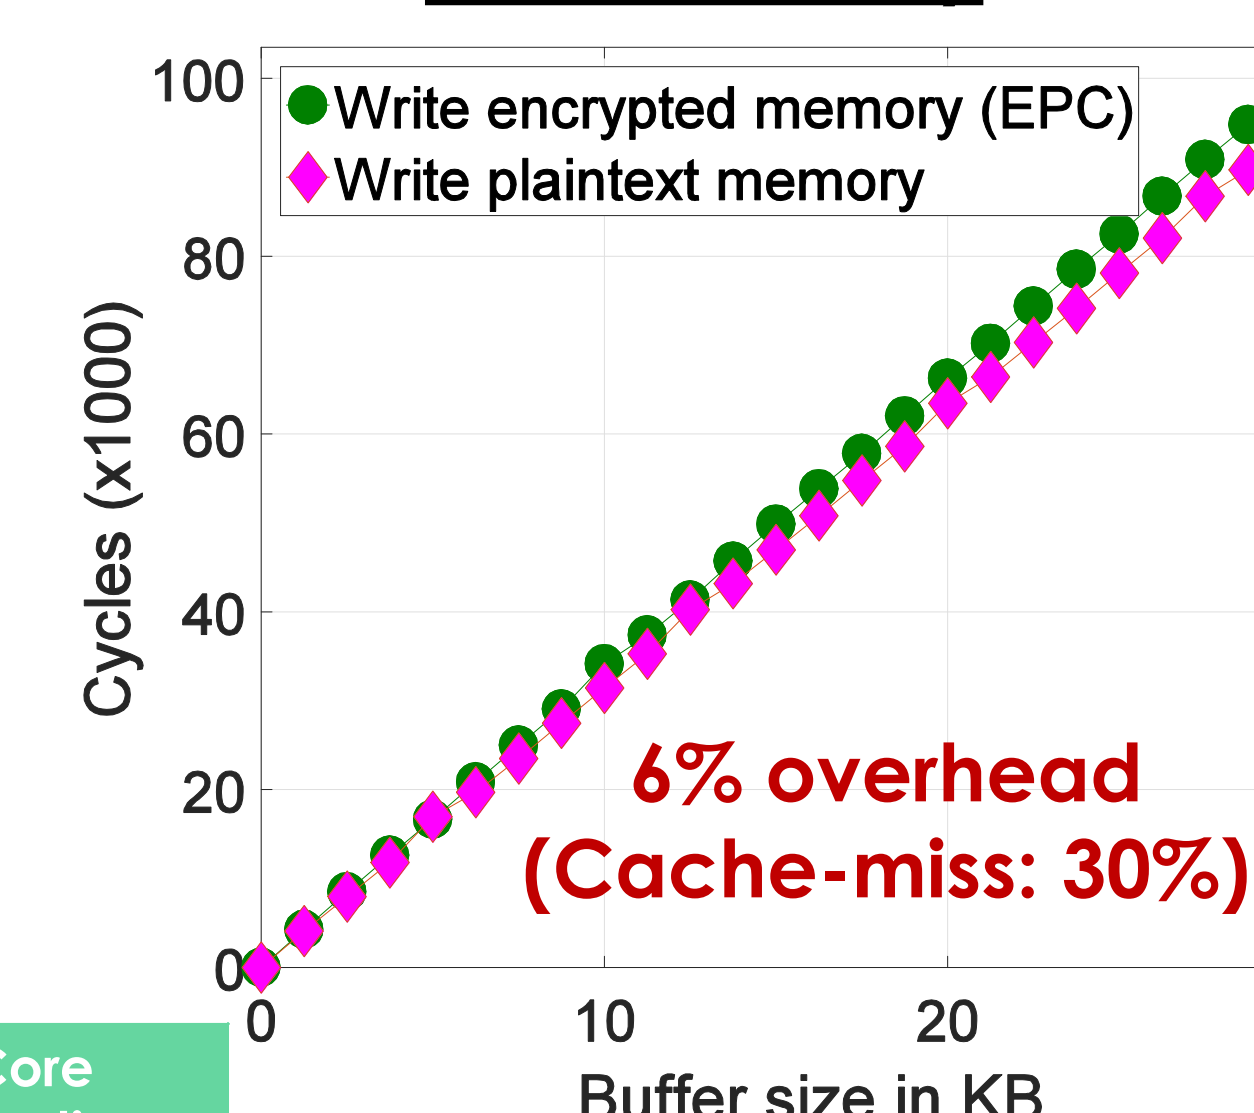


Application	# Calls /second	Core spending
Memcached	200,000	43%
OpenVPN	275,000	57%
Lighttpd	270,000	56%


Context switches may dominate performance

Cost of Encrypted Memory

Write Latency



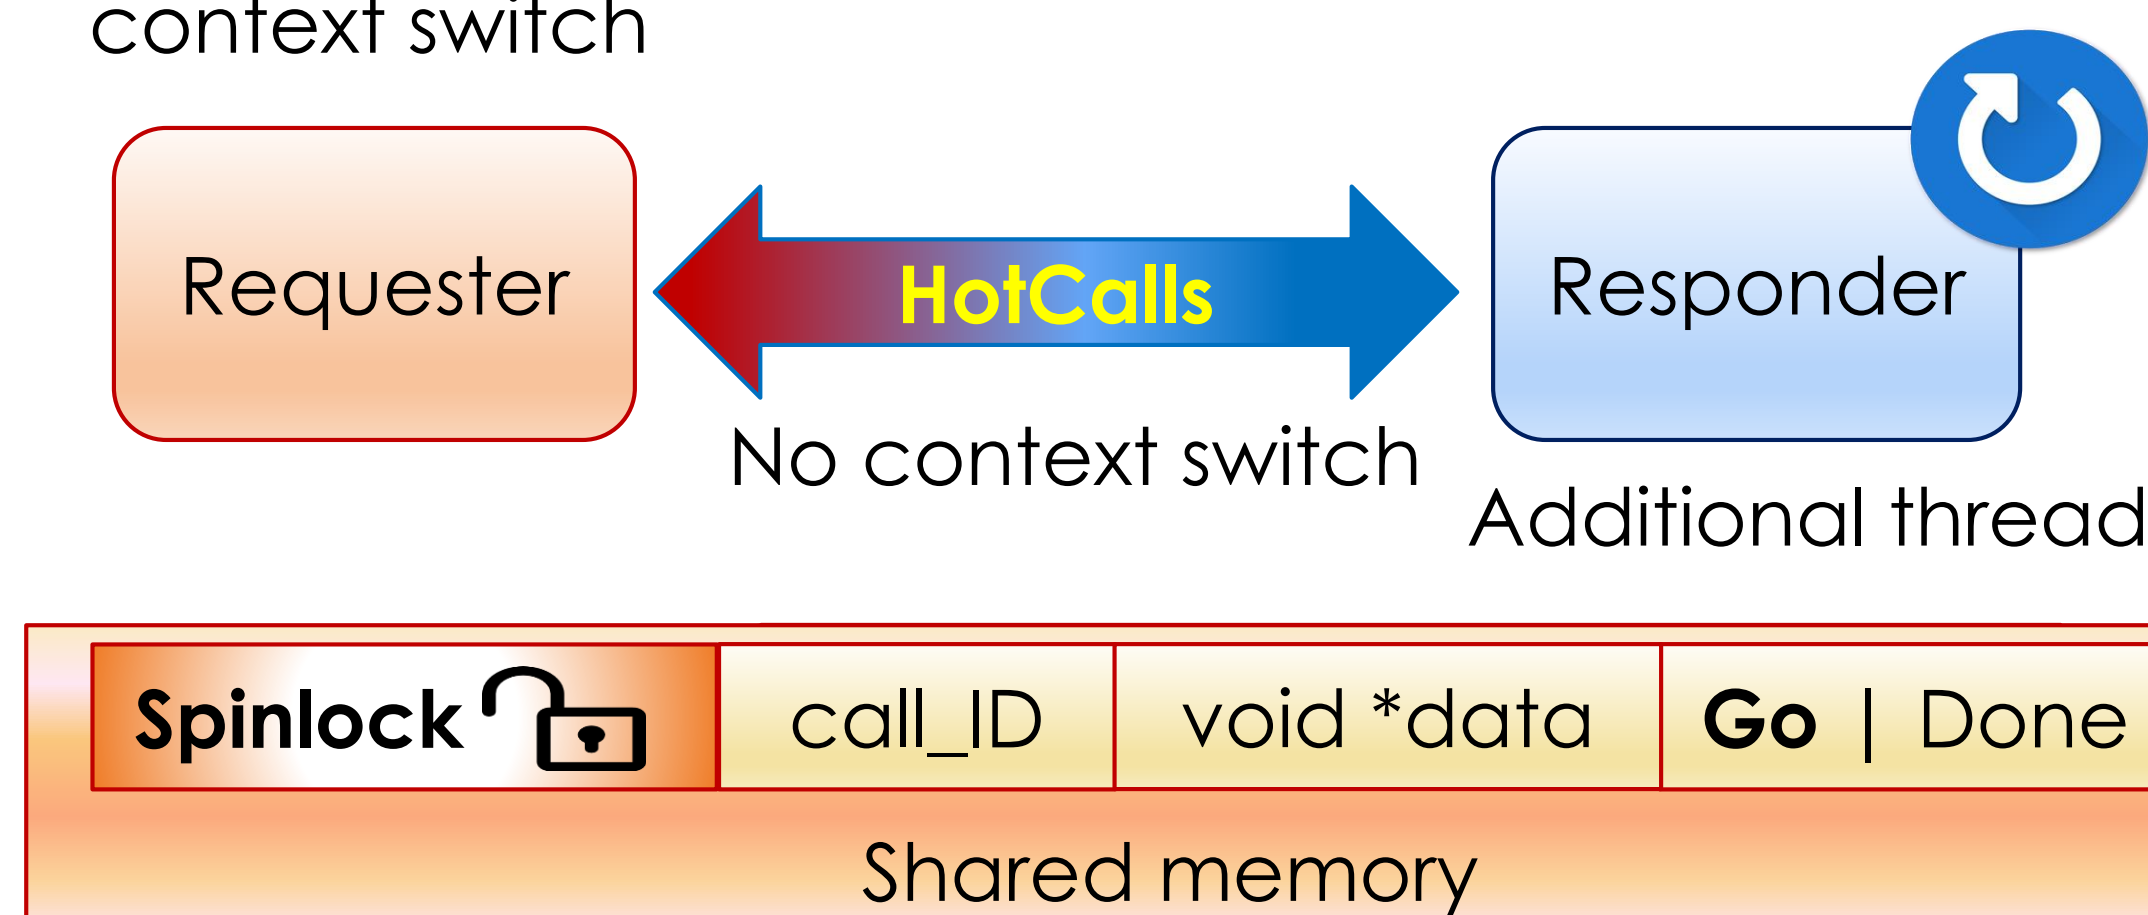
Read Latency




Encrypted memory is a potential bottleneck

HotCalls Mechanism

Key insight: requesting services does not mandate a context switch

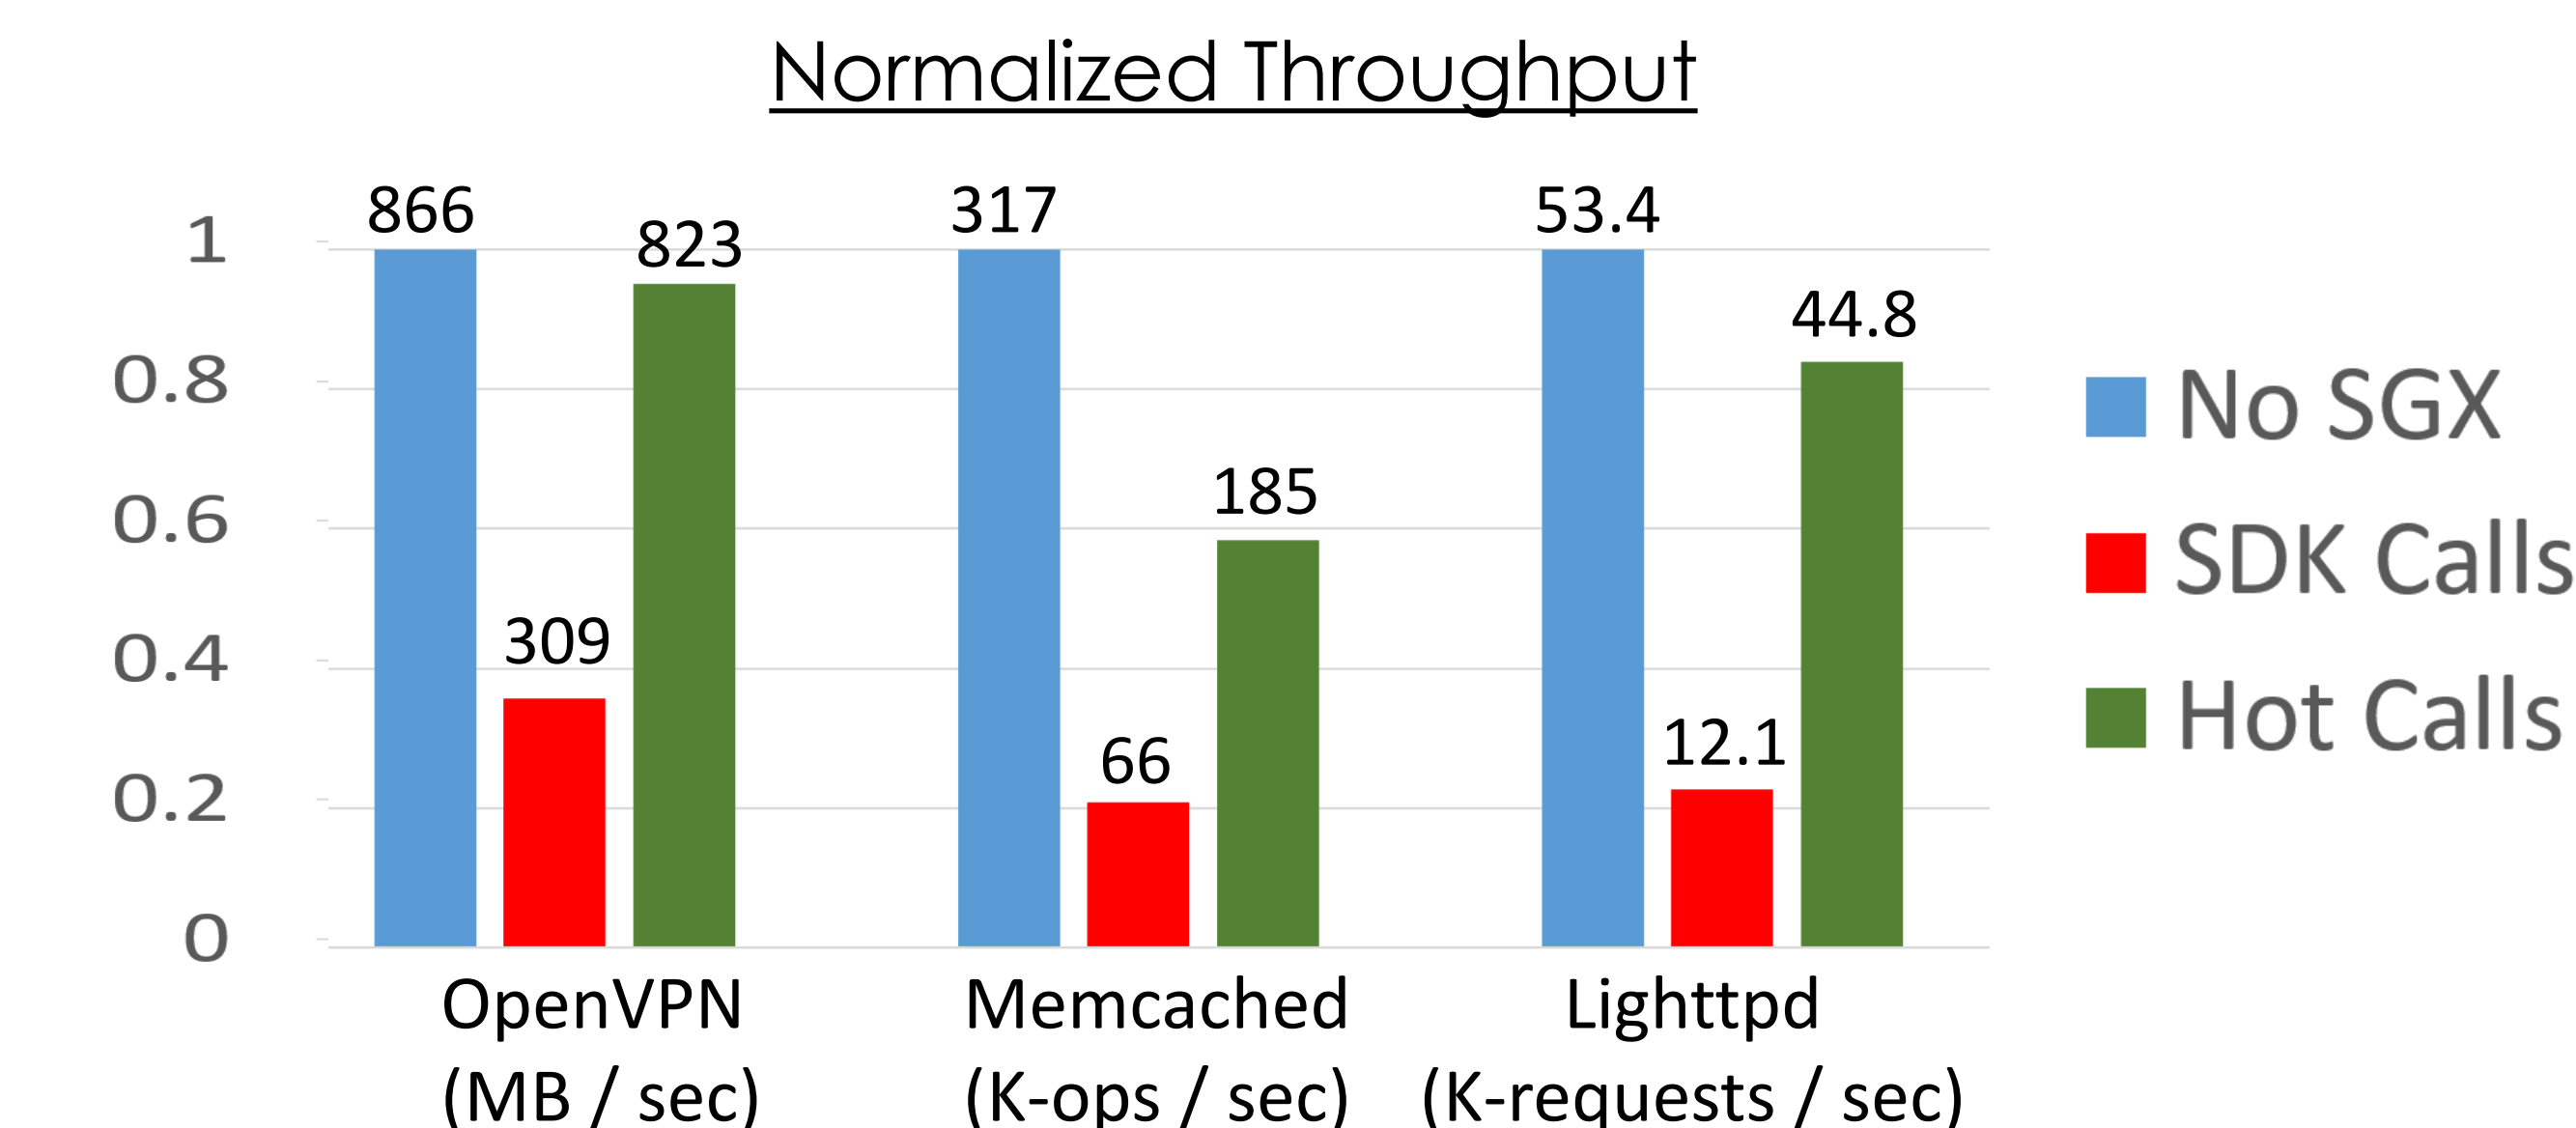


HotCalls vs. SDK Calls



HotCalls are order of magnitude faster

HotCalls in Action



Making SGX great again