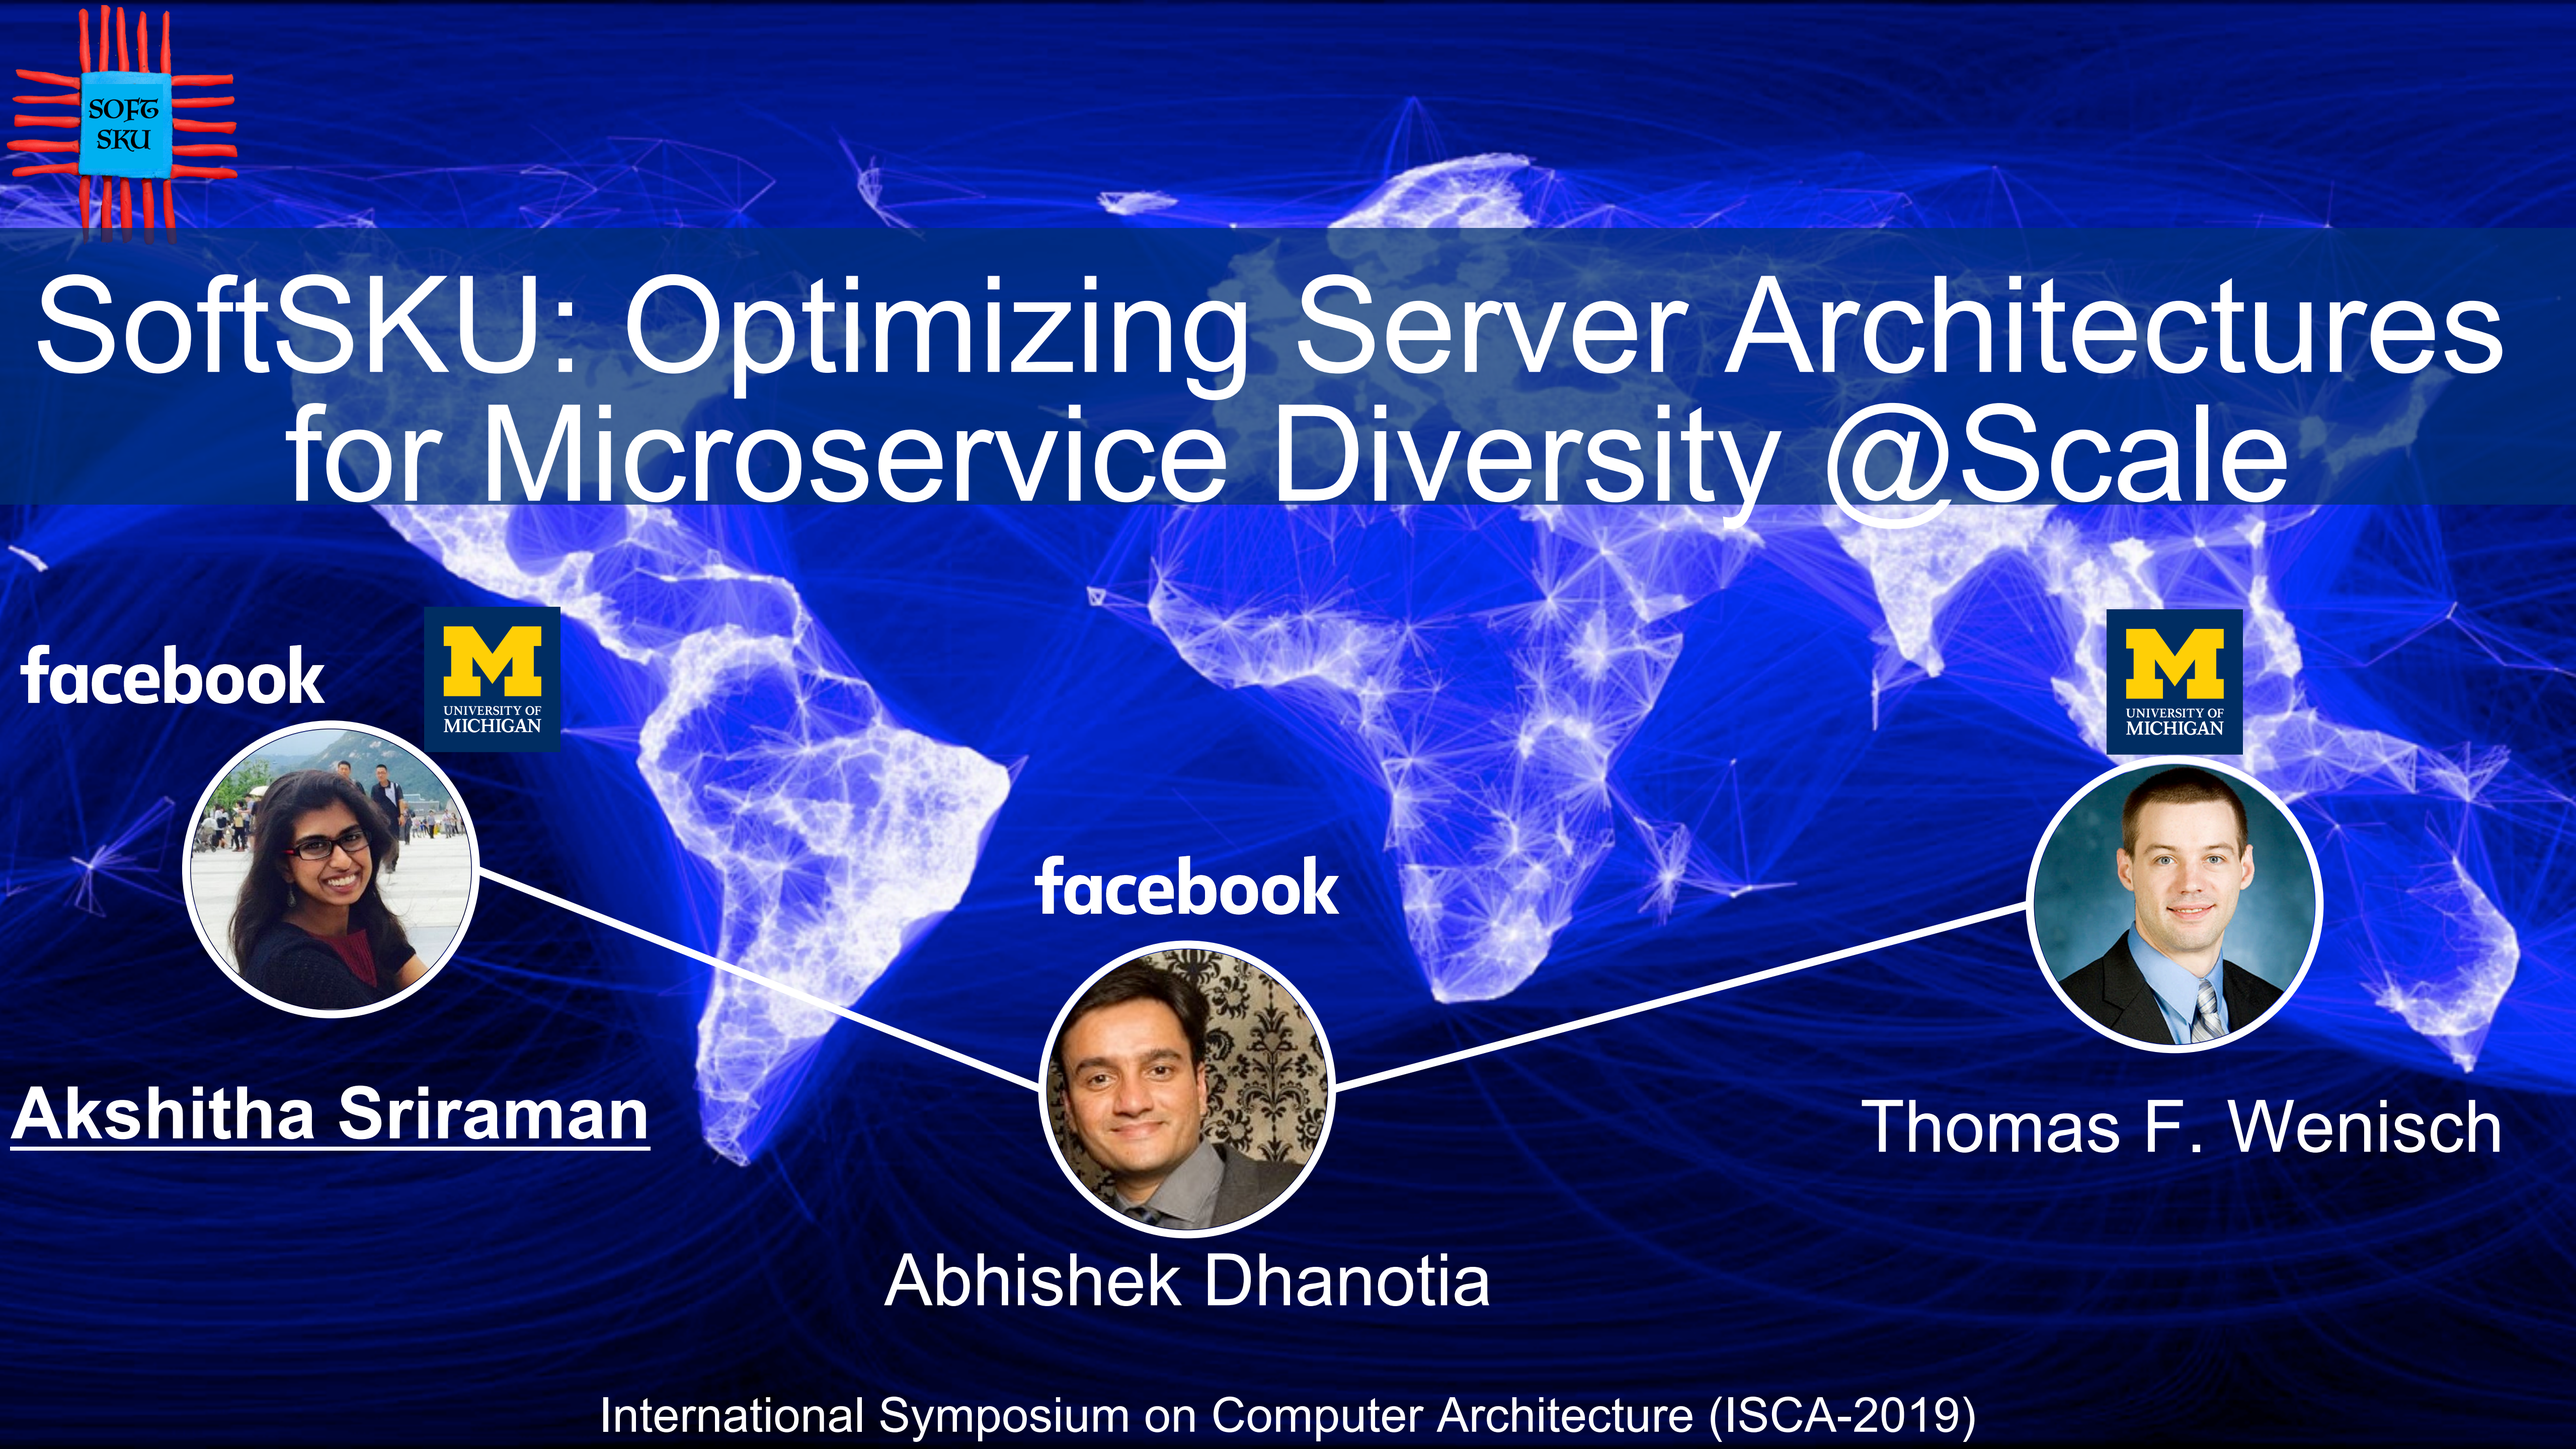

SoftSKU: Optimizing Server Architectures for Microservice Diversity @Scale

facebook

UNIVERSITY OF MICHIGAN

Akshitha Sriraman

facebook

Abhishek Dhanotia

UNIVERSITY OF MICHIGAN

Thomas F. Wenisch

International Symposium on Computer Architecture (ISCA-2019)

Rapid Increase in Modern Web Services

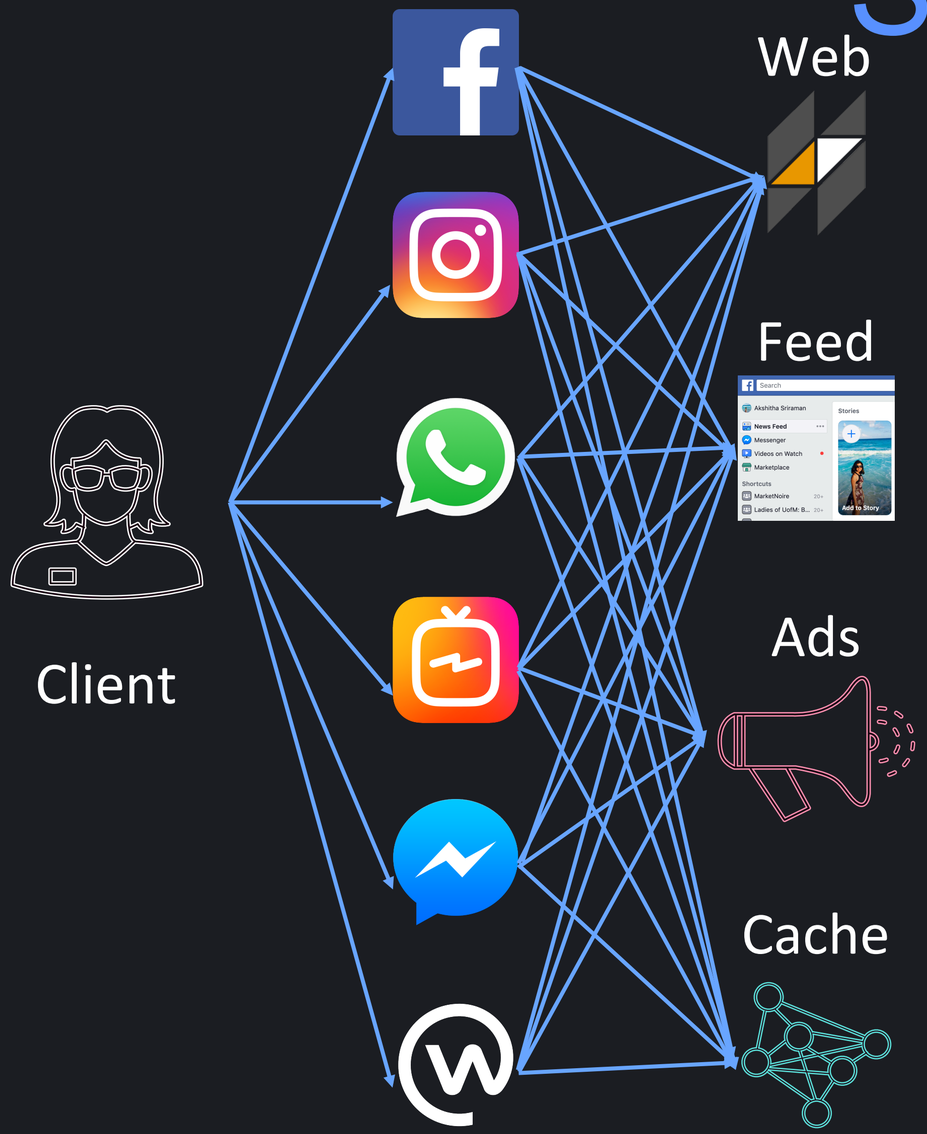

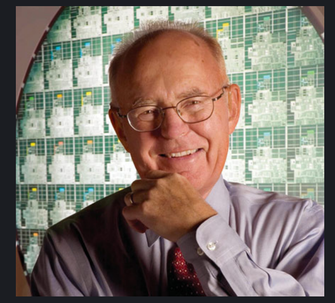

Stringent Service-Level Objectives + Moore's law decline

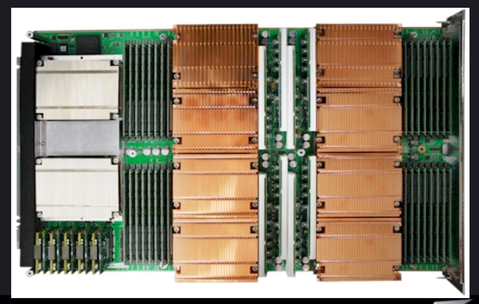

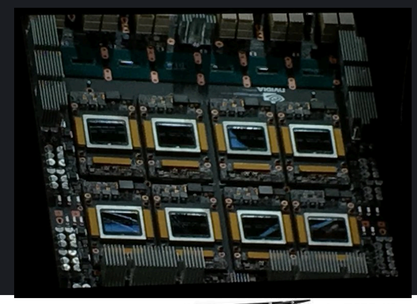

µservices SLOs + Moore's law decline -> need for customized hardware

Are Customized Platforms Always Needed?

Procurement @scale

Customized platforms -> expensive

Hardware resource fungibility

Testing overhead

Urgent need for limited SKUs that support a variety of microservices

Performance of Commodity Servers

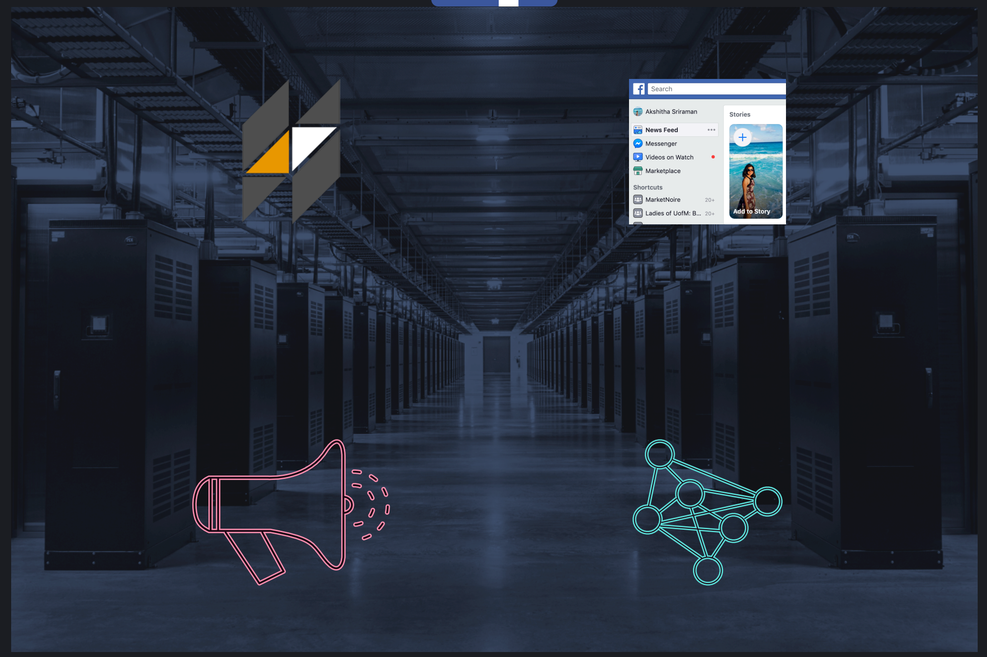

Maintain commodity hardware
for procurement efficiency & scalability

How performant is commodity hardware for these µservices?

Do µservices have common bottlenecks that can inspire future CPU designs?

Contributions

- Comprehensive characterization of Facebook's microservices
  - System-level and architectural bottlenecks
  - Enormous bottleneck diversity across microservices
- Concept of "soft" server SKUs
  - Tuning coarse-grained OS & hardware configuration knobs
- µSKU
  - Automates soft-SKU search & configuration via production AB-tests
  - Deploys soft SKUs on production microservices

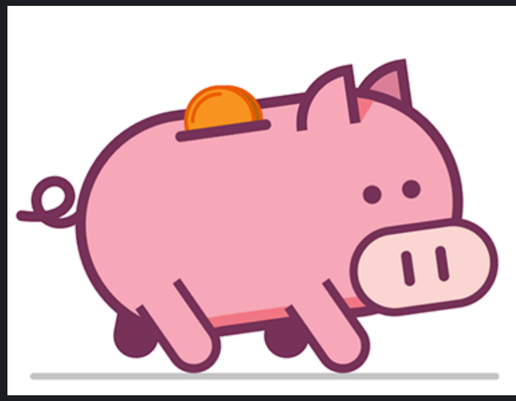

~7.2% perf. boost on **production** µservices + **no extra hardware**

Introduction:

Commodity servers AND performance?

Characterization:

System-level & architectural understanding of key FB µservices

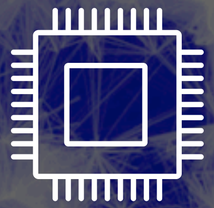

Soft SKU:

Achieving data center performance- & cost-efficiency

µSKU:

Tool to automate Soft SKU creation

Evaluation:

7% better performance on production systems

Introduction:

Commodity servers AND performance?

Characterization:

System-level & architectural understanding of key FB µservices

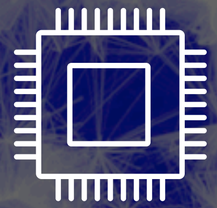

Soft SKU:

Achieving data center performance- & cost-efficiency

µSKU:

Tool to automate Soft SKU creation

Evaluation:

7% better performance on production systems

Facebook's Production Microservices

µService	Throughput (QPS)	Response latency	Pathlength
Web	O(100)	O(ms)	$O(10^6)$
Feed1	O(1000)	O(ms)	$O(10^9)$
Feed2	O(10)	O(s)	$O(10^9)$
Ads1	O(10)	O(ms)	$O(10^9)$
Ads2	O(100)	O(ms)	$O(10^9)$
Cache1	O(100K)	O(µs)	$O(10^3)$
Cache2	O(100K)	O(µs)	$O(10^3)$

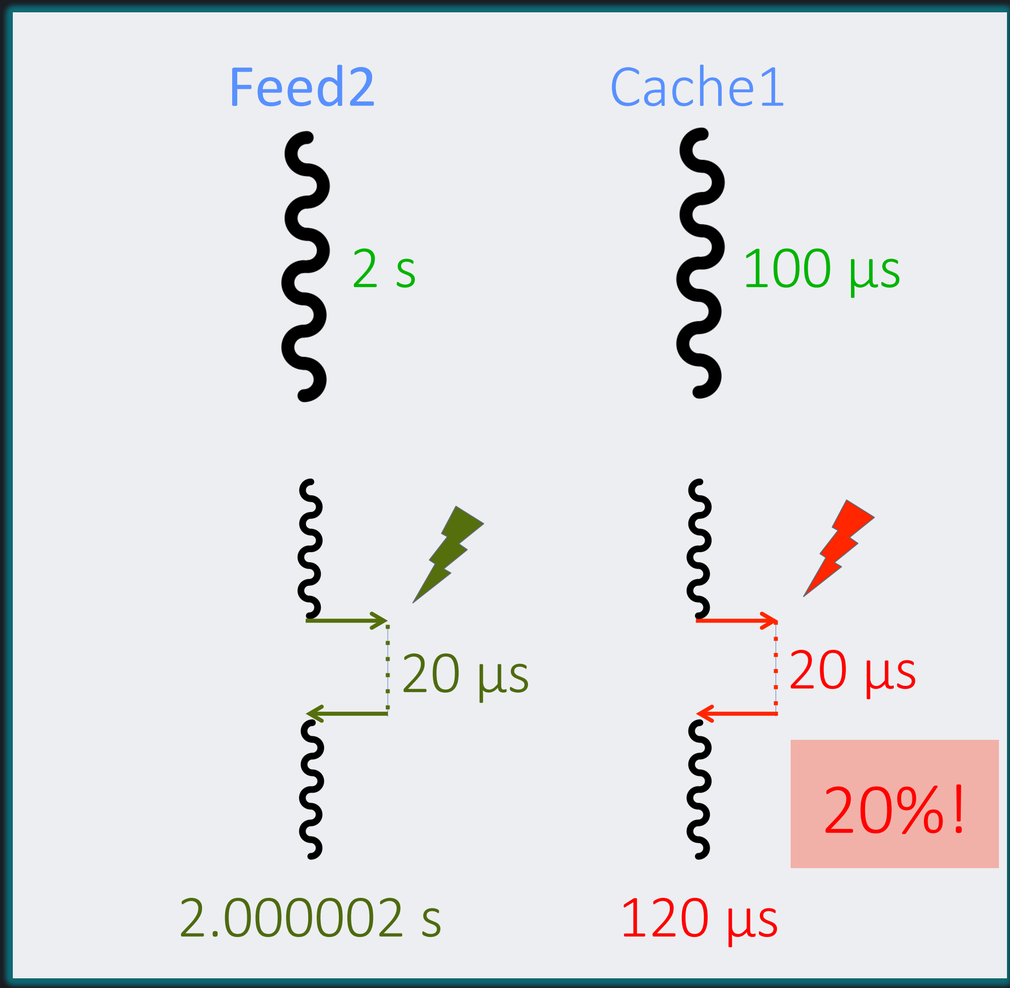

Diversity in service traits causes diversity in bottlenecks faced

Context Switch Penalty

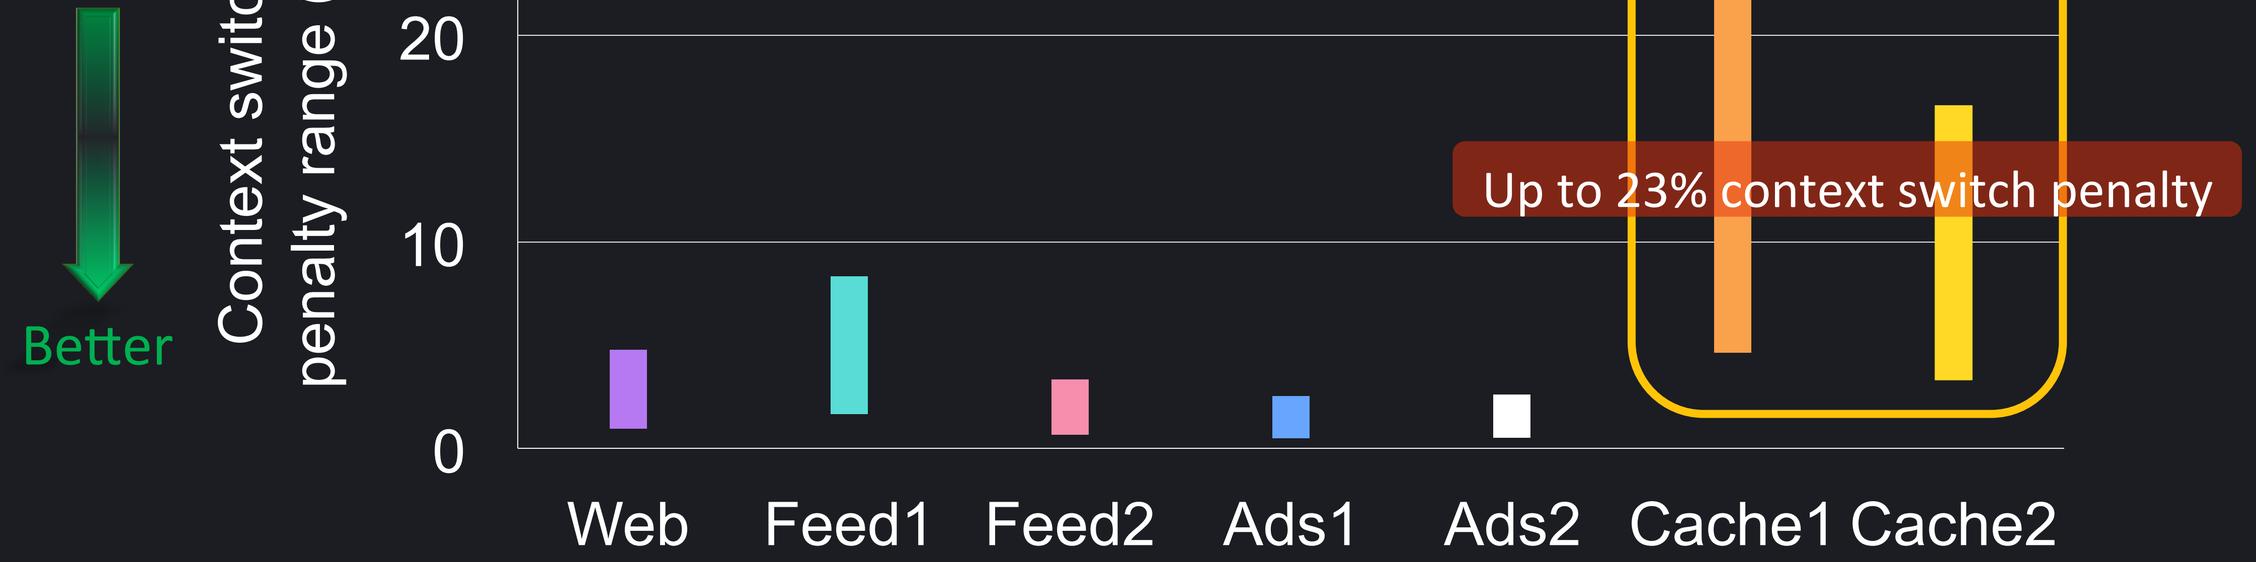

Cache can benefit from context switches/switching delay reductions

Pipeline Slot Breakdown

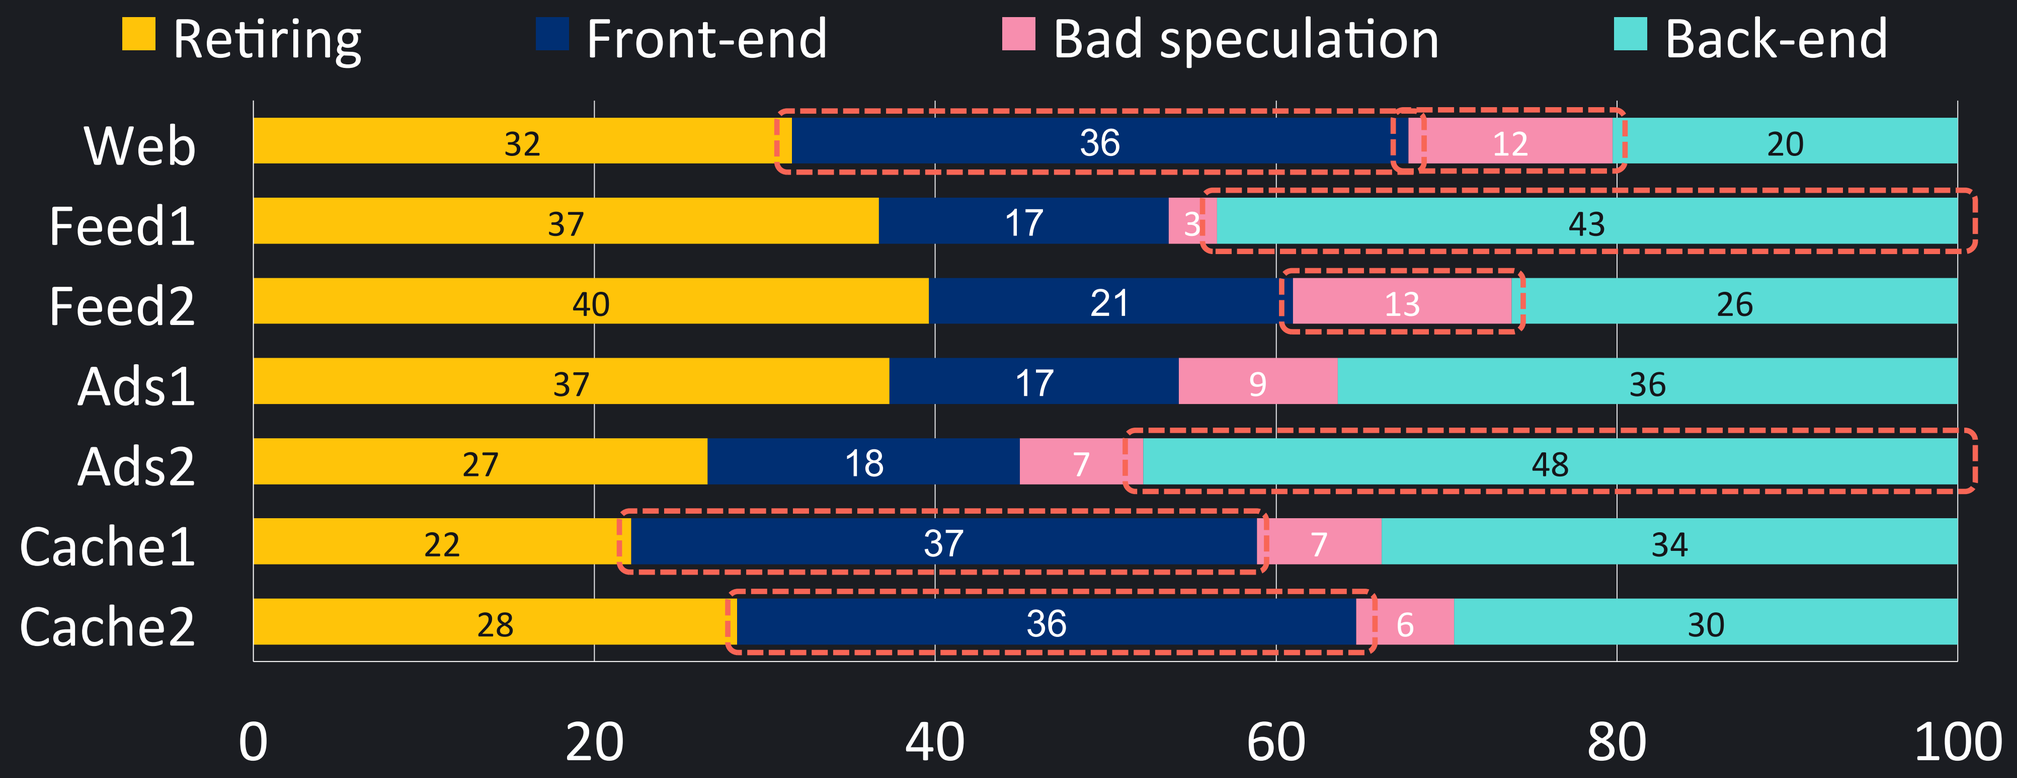

Pipeline slot breakdown (%)

Better

Branch prediction researchers

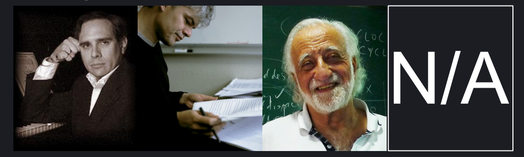

Several microservices face much higher front-end stalls

Last Level Cache: Misses Per Kilo Instruction

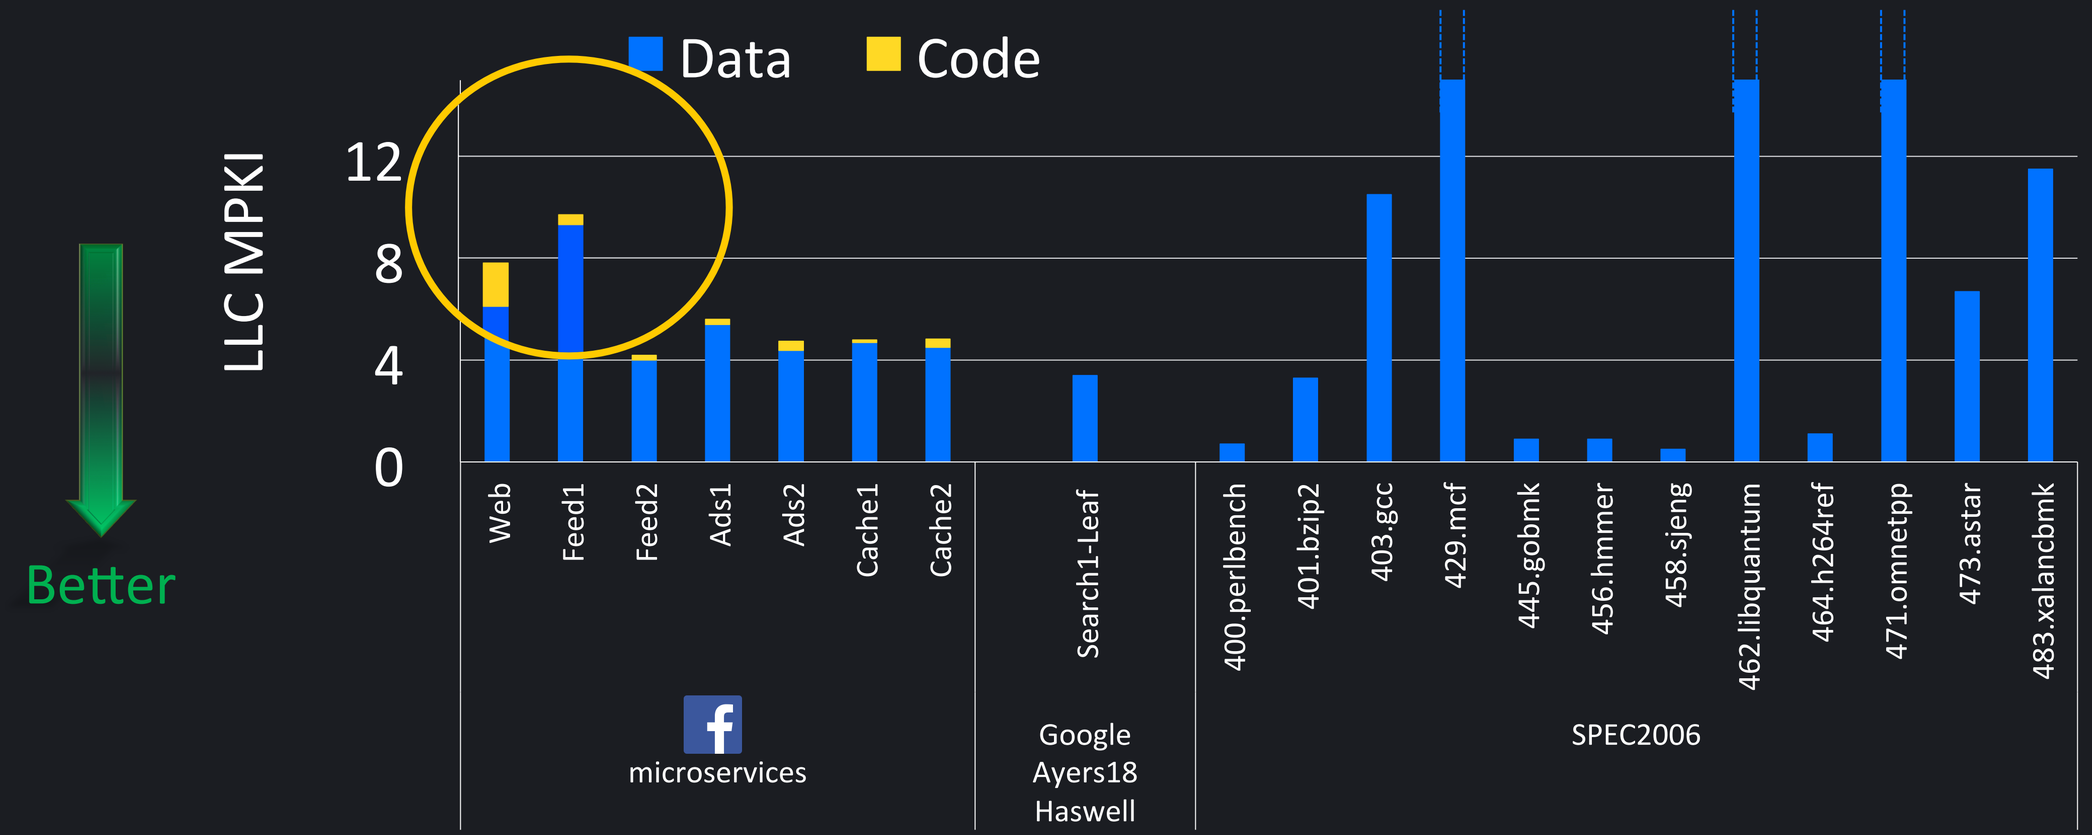

Better

Web incurs an unusually high LLC code MPKI

More Characterizations

Per-core IPC

Instruction mix breakdown

L1 code MPKI

L1 data MPKI

L2 code MPKI

L2 data MPKI

LLC capacity scaling

Memory latency

Memory bandwidth

iTLB load MPKI

iTLB store MPKI

dTLB load MPKI

dTLB store MPKI

More details in the paper

Characterization Takeaways

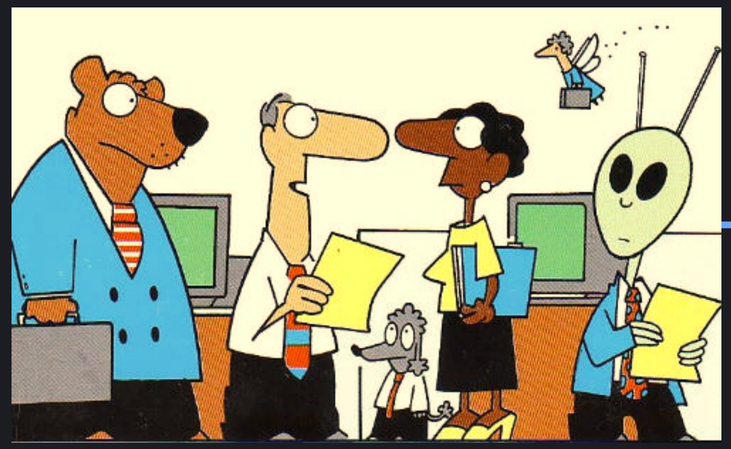

Diversity in bottlenecks across microservices

Custom CPU SKUs

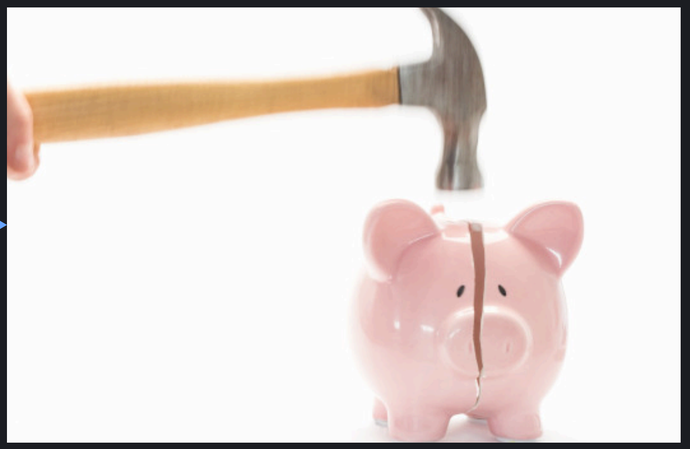

Prohibitively expensive

Can we achieve perf. efficiency without building customized hardware?

Introduction:

Commodity servers AND performance?

Characterization:

System-level & architectural understanding of key FB µservices

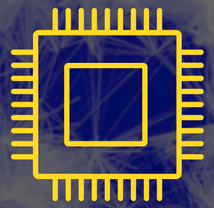

Soft SKU:

Achieving data center performance- & cost-efficiency

µSKU:

Tool to automate Soft SKU creation

Evaluation:

7% better performance on production systems

“Soft” SKUs: Best of Both Worlds

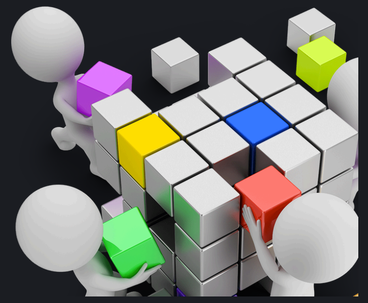

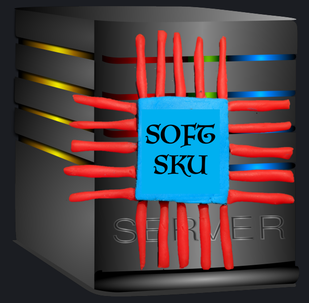

Tune coarse HW & OS knobs on commodity HW

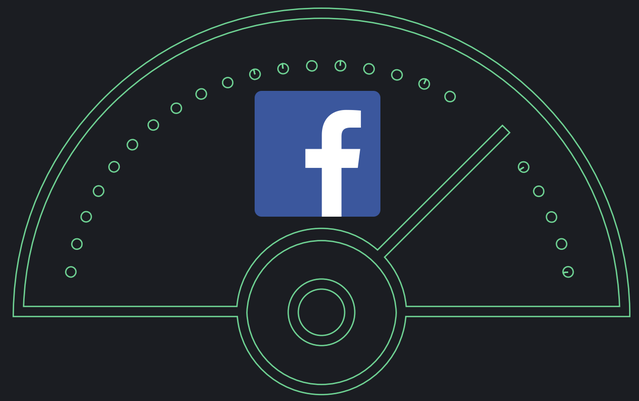

Performance efficiency

Procurement efficiency & scalability

- Core Frequency
- Uncore Frequency
- Core Count
- CDP: LLC
- Prefetcher
- Transparent Huge Pages (THP)
- Statically-Allocated Huge Pages (SHP)

Tuning configurable OS & HW knobs for each microservice may improve perf. efficiency & scalability

Tuning Configurable OS & CPU Knobs

Complex design space

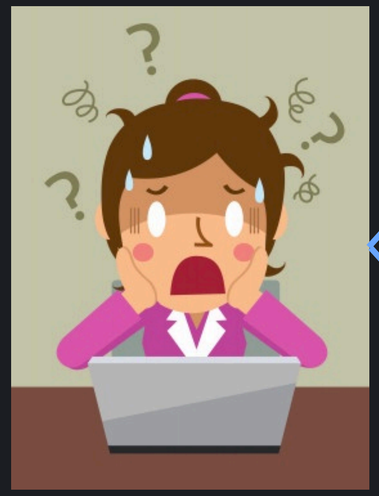

Manual tuning: Not scalable

Quick code evolution

Synthetic load tests don't capture production load

Small knob tuning effects

Need for an automated design & deployment tool to create soft SKUs

Introduction:

Commodity servers AND performance?

Characterization:

System-level & architectural understanding of key FB µservices

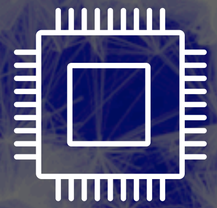

Soft SKU:

Achieving data center performance- & cost-efficiency

µSKU:

Tool to automate Soft SKU creation

Evaluation:

7% better performance on production systems

µSKU: Tool to Design & Deploy Soft SKUs

Design Phase

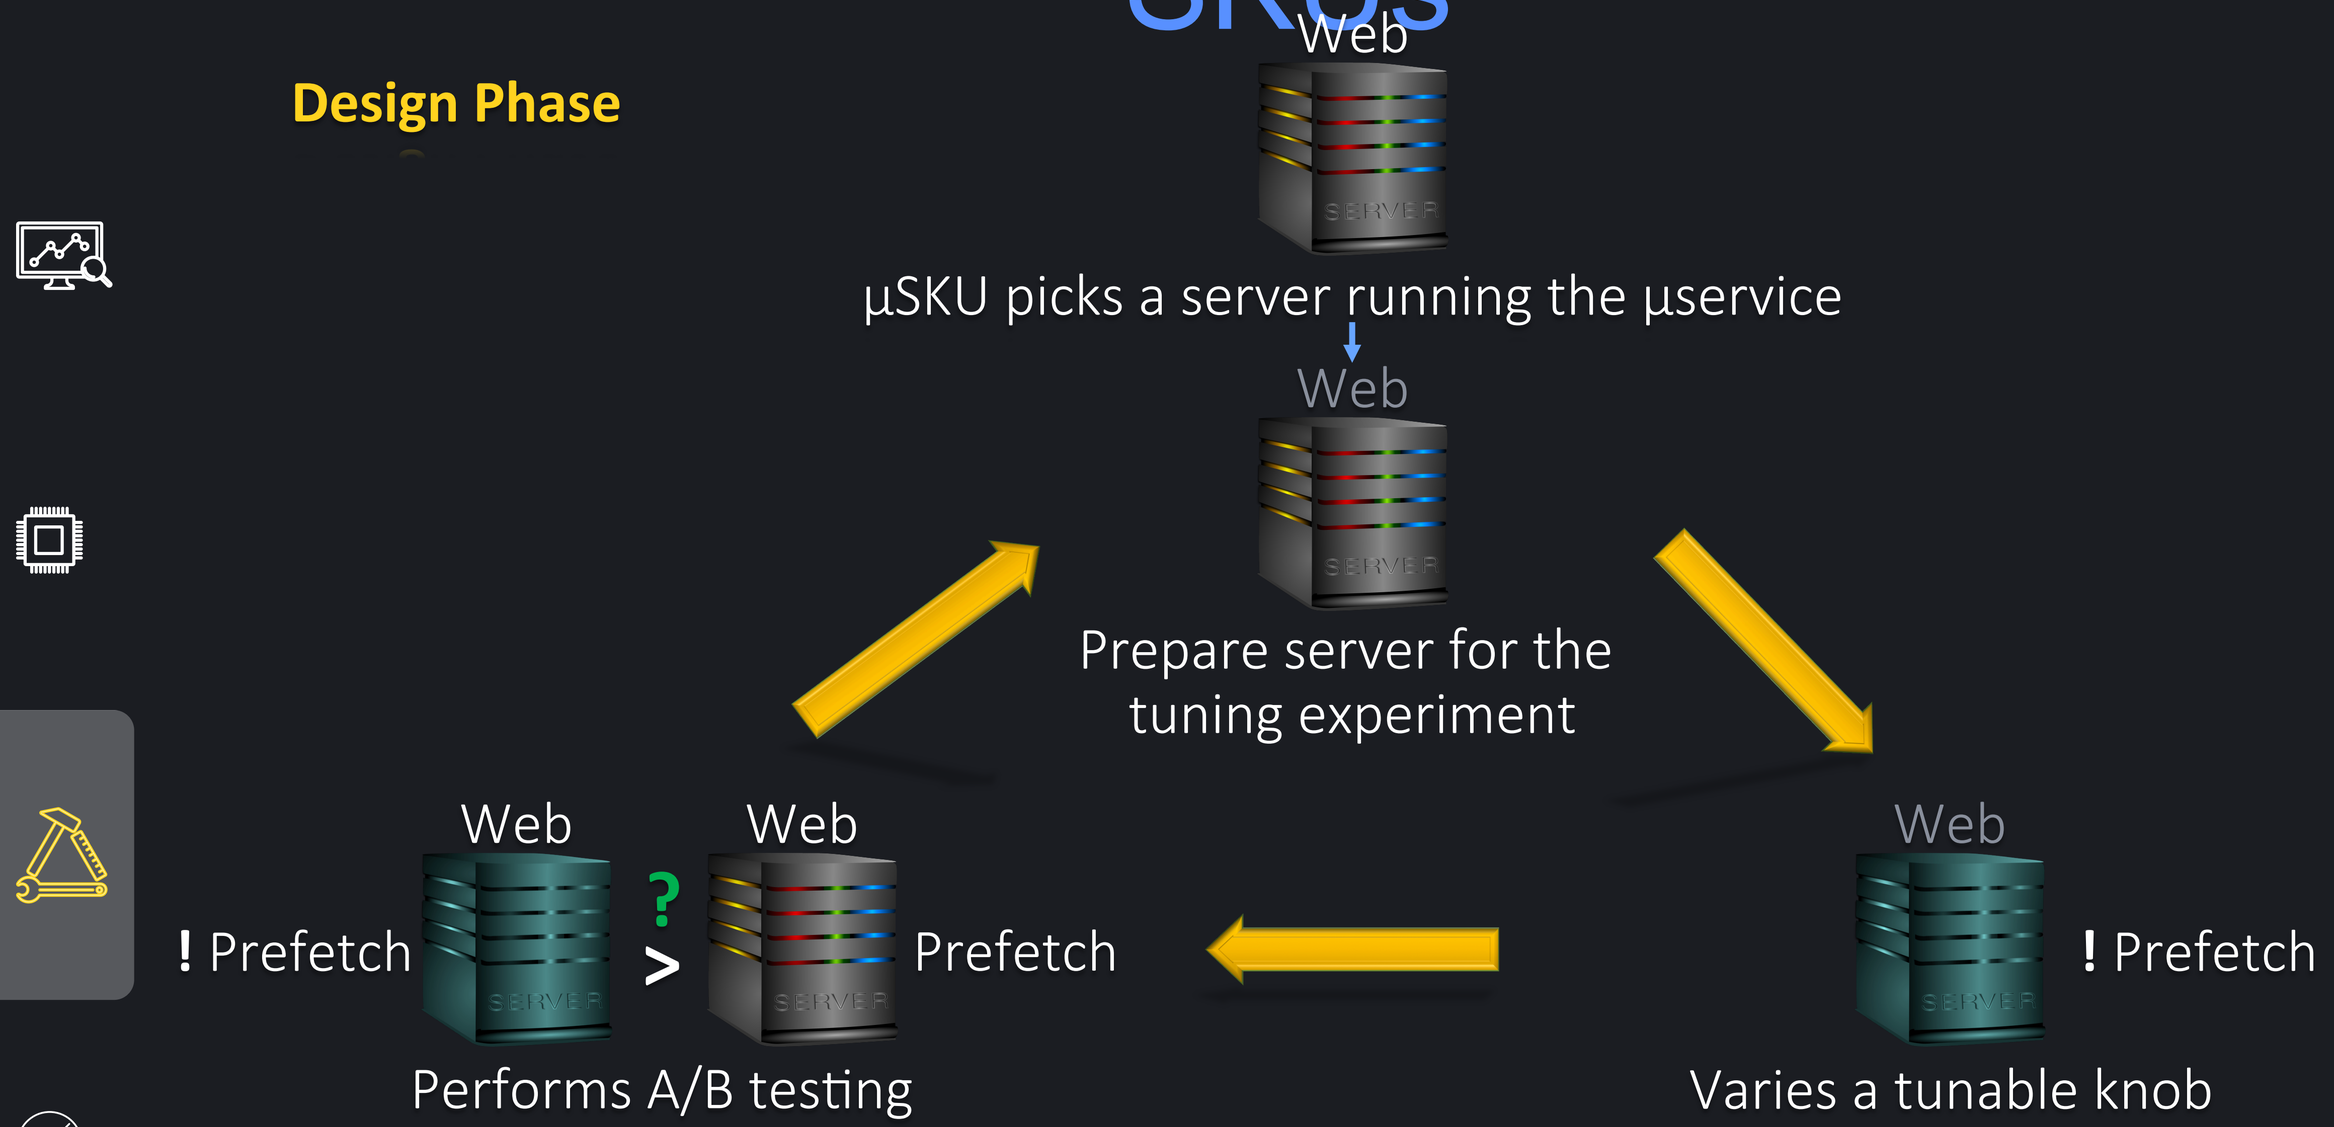

µSKU: Tool to Design & Deploy Soft SKUs

Deployment Phase

A/B tester's design space map

Knob	Best config
Core freq.	2.2 GHz
.	
.	
Prefetcher	on

Selects performant knob configs.

Applies config. to production servers

Further A/B testing

?
>

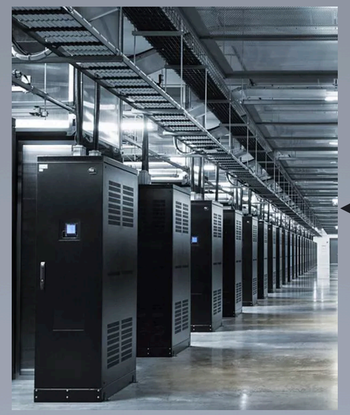

Soft SKU generator

Introduction:

Commodity servers AND performance?

Characterization:

System-level & architectural understanding of key FB µservices

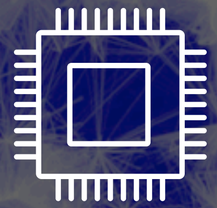

Soft SKU:

Achieving data center performance- & cost-efficiency

µSKU:

Tool to automate Soft SKU creation

Evaluation:

7% better performance on production systems

Methodology

- µSKU evaluated on Web & Ads1
  - Web evaluated on two platforms: Skylake & Broadwell

- Soft SKUs compared against:
  - Stock and hand-tuned production server configurations

µSKU Design Phase: Prefetcher Tuning

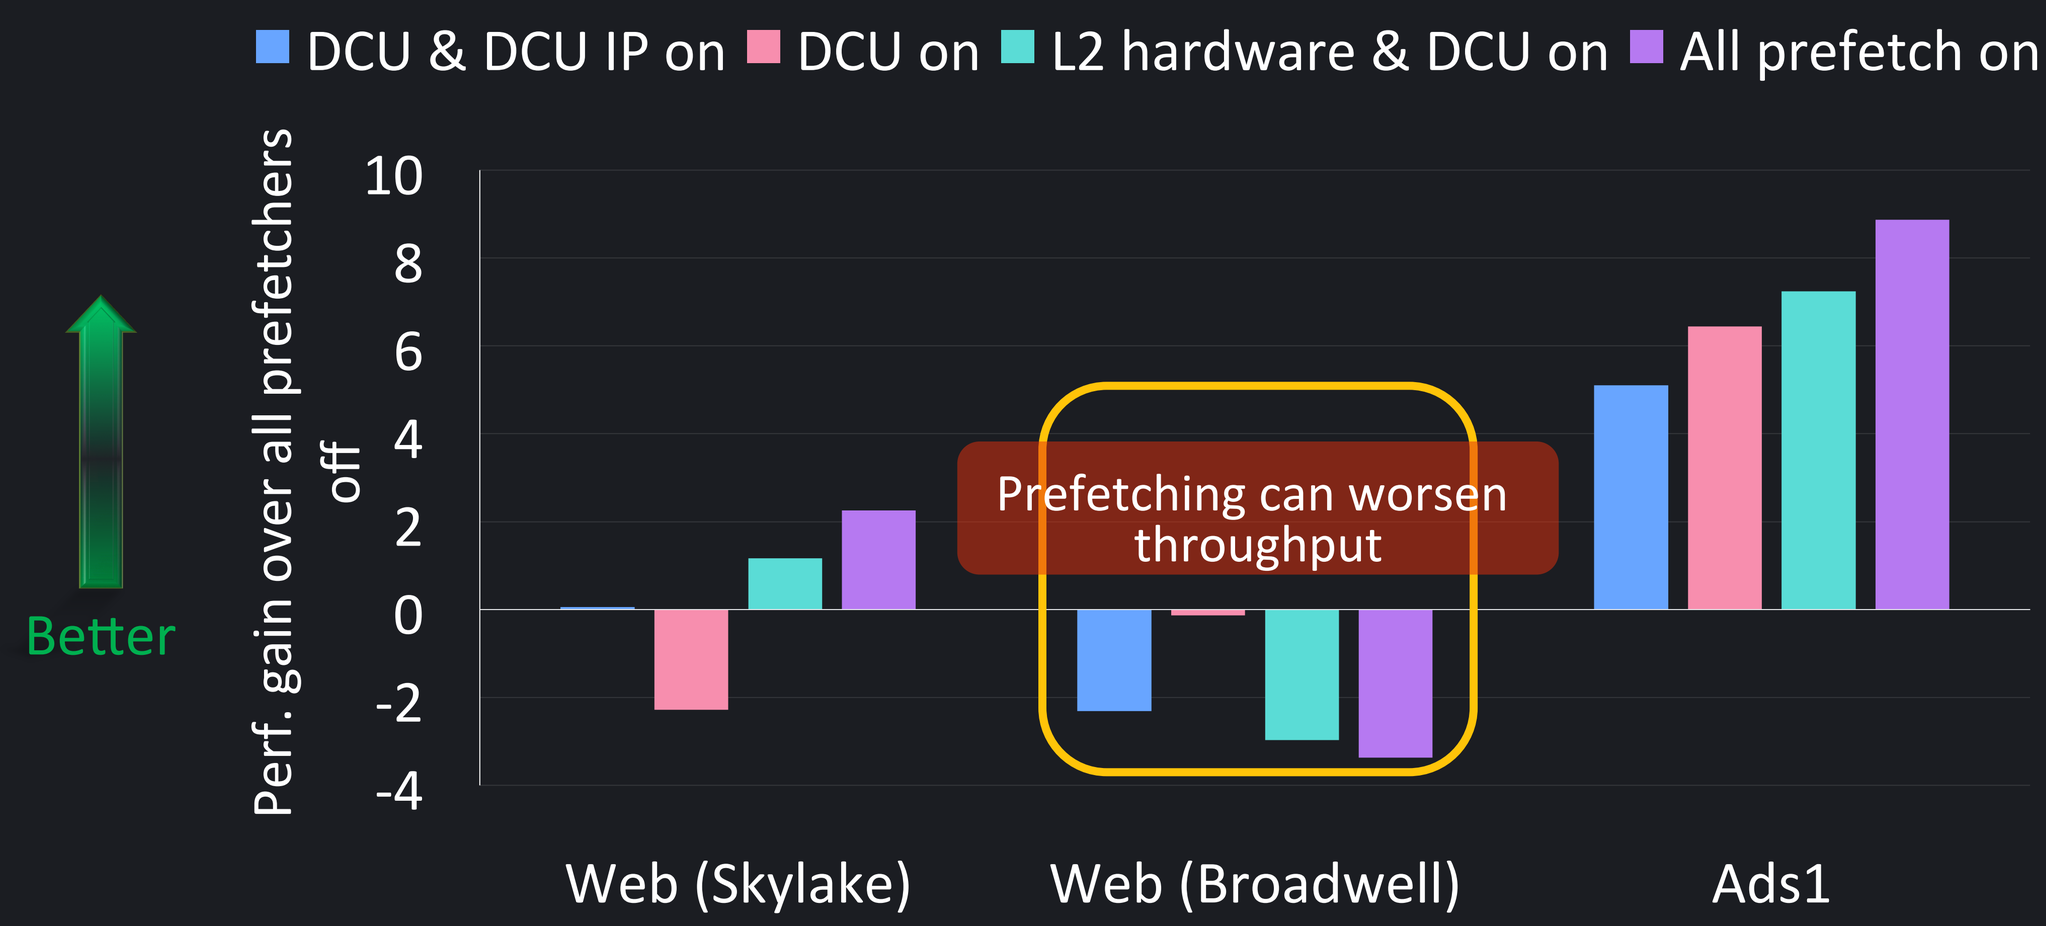

Web (Broadwell) improves when prefetching is off: Better BW util.

µSKU Deployment Phase: Soft SKU Perf.

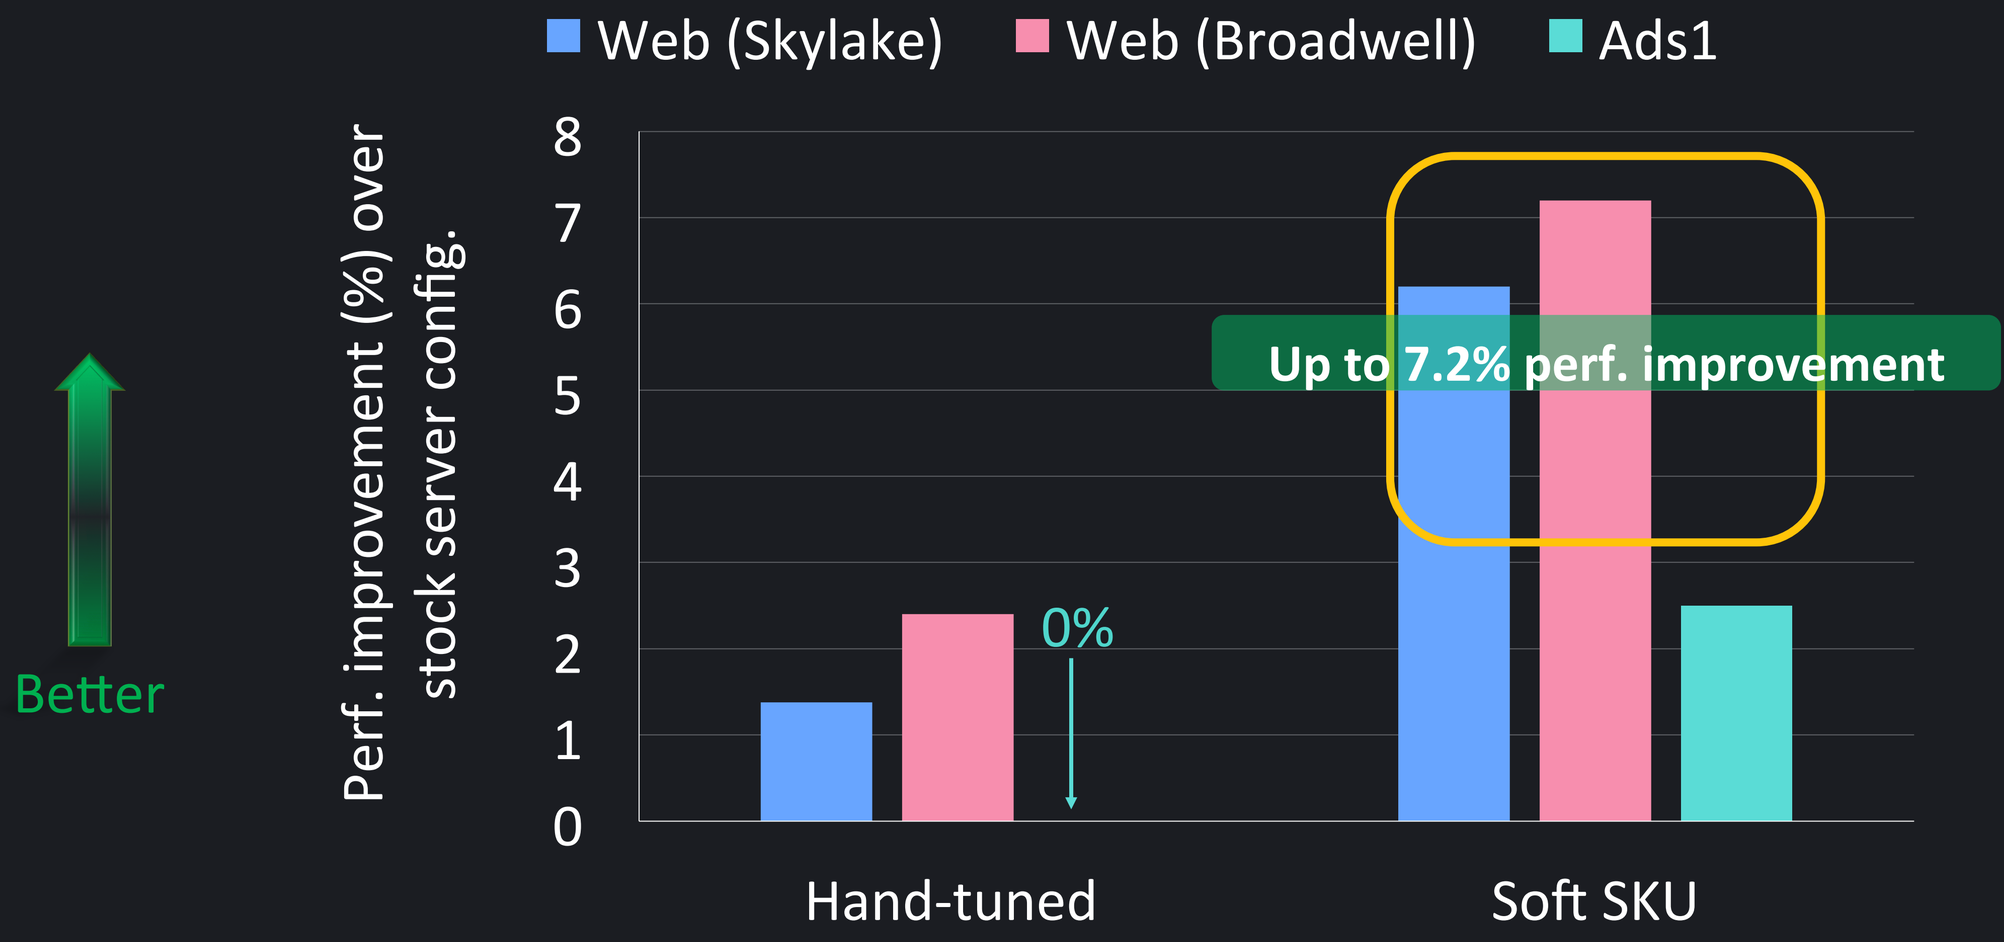

Soft SKU can achieve **~7.2% throughput improvement** on **production systems** with **no extra hardware requirement**

Conclusion

Hardware resource fungibility
+ performance effciency

Soft SKUs can achieve perf.
improvements on existing hardware

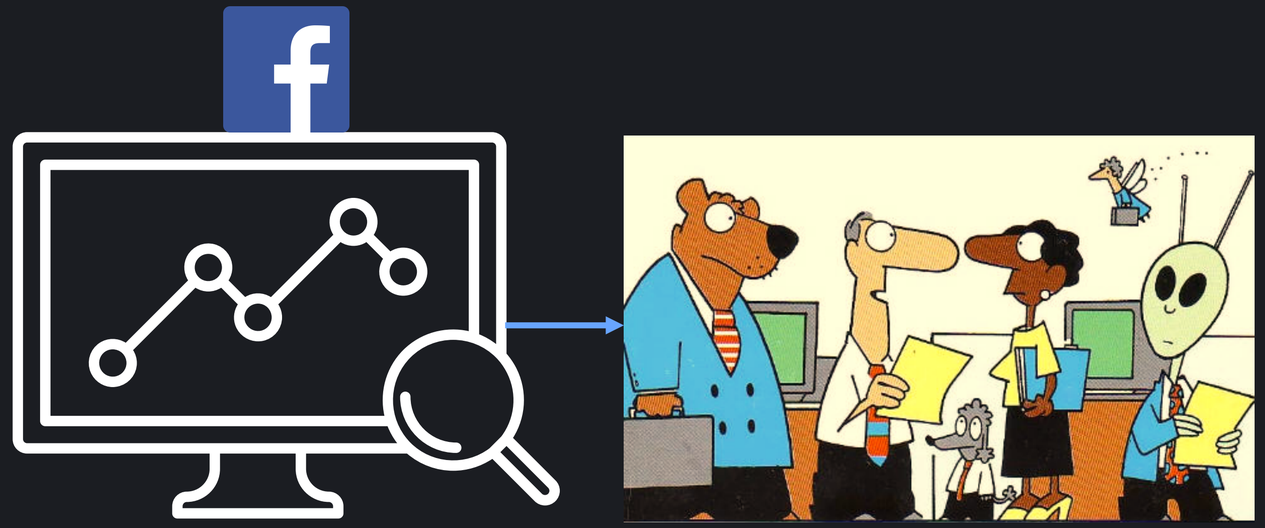

FB’s microservices are significantly diverse

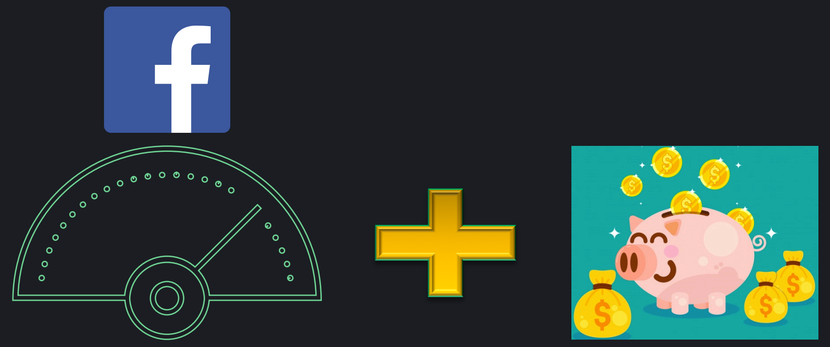

µSKU achieves ~7.2% throughput gains
on production systems

SoftSKU: Optimizing Server Architectures for Microservice Diversity @Scale

facebook

UNIVERSITY OF MICHIGAN

Akshitha Sriraman

facebook

Abhishek Dhanotia

UNIVERSITY OF MICHIGAN

Thomas F. Wenisch

International Symposium on Computer Architecture (ISCA-2019)