Amdahl's Law in Big Data Analytics: Alive and Kicking in TPCx-BB (BigBench)

Daniel Richins† Tahrina Ahmed* Russell Clapp※ Vijay Janapa Reddi†‡

†The University of Texas at Austin
※Intel
*Stanford
‡Google

February 28, 2018

TRINITY

What is Big Data Analytics?

What is Big Data Analytics?

Big data analytics is the application of **varied techniques** to very **large and diverse data sets** in order to uncover **hidden patterns** and produce **meaningful insights**.

What is Big Data Analytics?

Big data analytics is the application of **varied techniques** to very **large and diverse data sets** in order to uncover **hidden patterns** and produce **meaningful insights**.

Big data analytics deals with data sets **too large**, problems **too complex**, and patterns **too subtle** to be handled by conventional relational databases.

Big Data Benchmarks

Big Data Benchmarks

There are ample big data-related benchmarks.

Big Data Benchmarks

There are ample big data-related benchmarks.

CloudSuite

BigDataBench

AMPLab Benchmarks

LinkBench

YCSB

DCBench

HiBench

Big Data Benchmarks

There are ample big data-related benchmarks.

CloudSuite

BigDataBench

AMPLab Benchmarks

LinkBench

YCSB

DCBench

HiBench

TPCx-BB BigBench

Big Data Benchmarks

There are ample big data-related benchmarks.

CloudSuite

What is it and why should computer architects care?

LinkBench

YCSB

DCBench

HiBench

TPCx-BB BigBench

TPCx-BB Origin and Features

TPCx-BB Origin and Features

Proposed at SIGMOD 2013, BigBench was developed with input from many industry partners.

TPCx-BB Origin and Features

Proposed at SIGMOD 2013, BigBench was developed with input from many industry partners.

TPCx-BB Origin and Features

Proposed at SIGMOD 2013, BigBench was developed with input from many industry partners.

BigBench was standardized as TPCx-BB in 2014.

TPCx-BB (BigBench) is Uniquely Realistic

TPCx-BB (BigBench) is Uniquely Realistic

It simulates a modern retailer with a physical and online store presence.

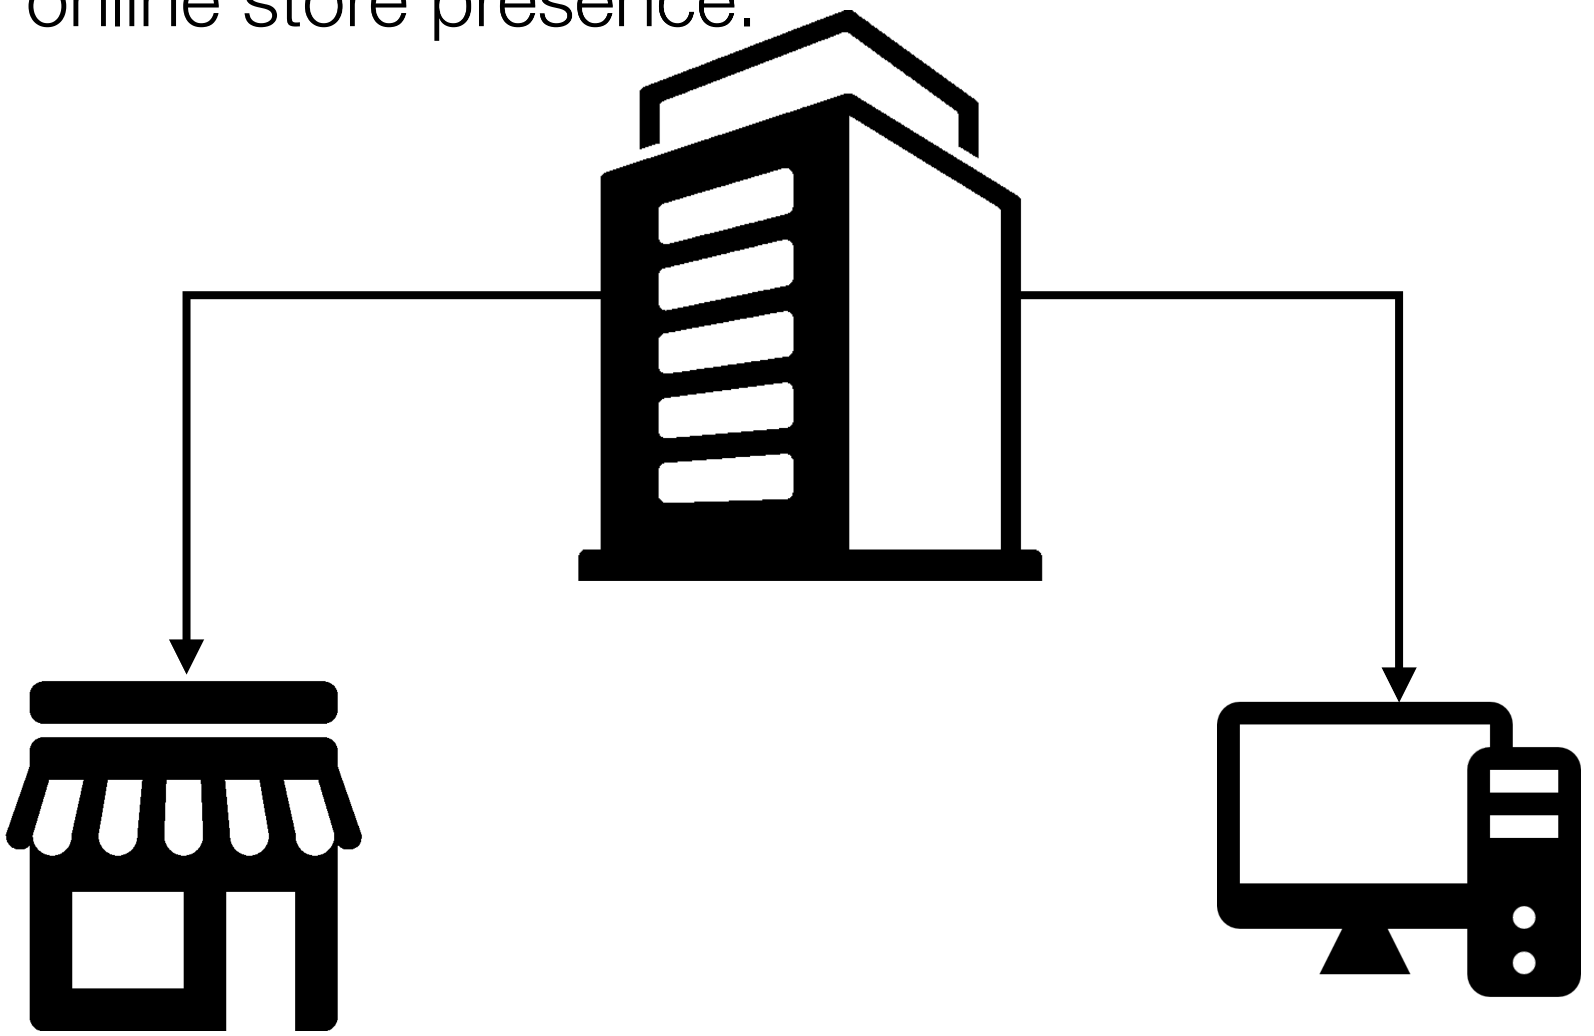

TPCx-BB (BigBench) is Uniquely Realistic

It simulates a modern retailer with a physical and online store presence.

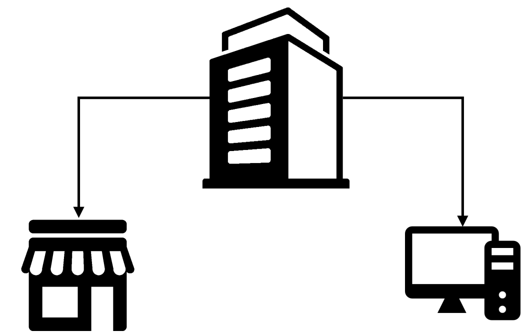

It gathers copious data on products, customers, and competitors.

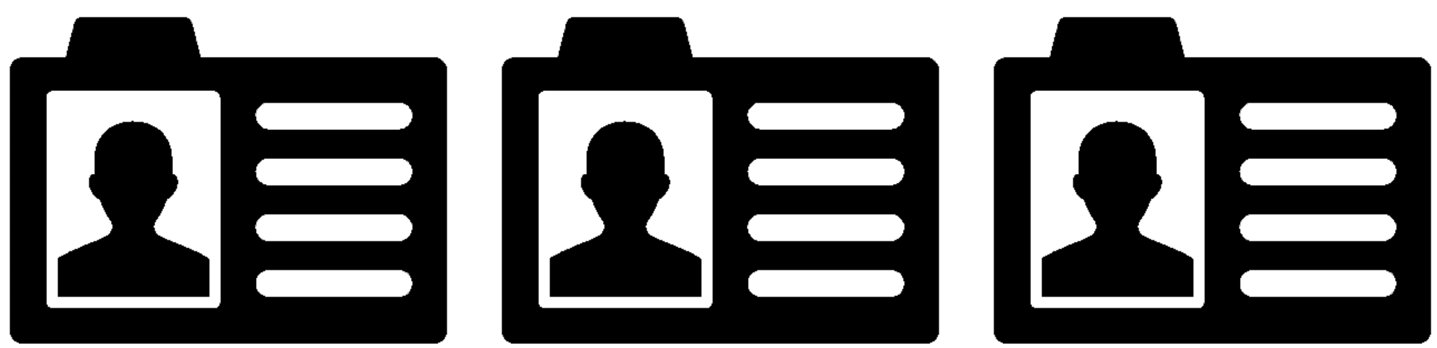

TPCx-BB (BigBench) is Uniquely Realistic

It simulates a modern retailer with a physical and online store presence.

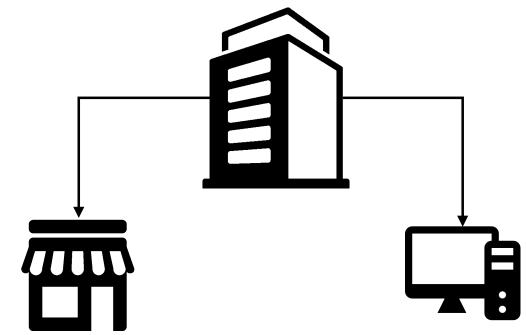

It gathers copious data on products, customers, and competitors.

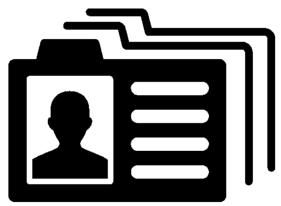

It uses complex queries to extract value from its collected data.

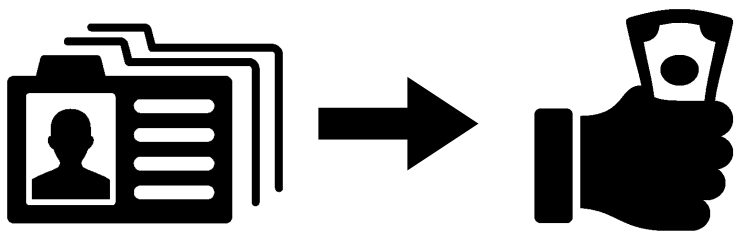

TPCx-BB (BigBench) is Uniquely Realistic

It simulates a modern retailer with a physical and online store presence.

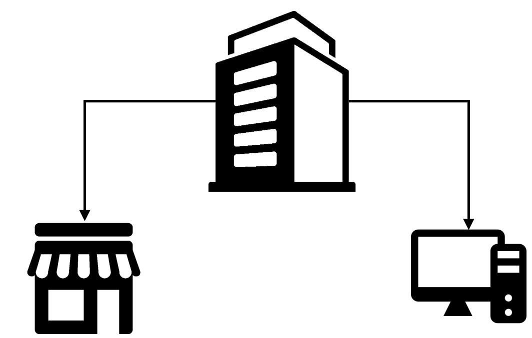

It gathers copious data on products, customers, and competitors.

It uses complex queries to extract value from its collected data.

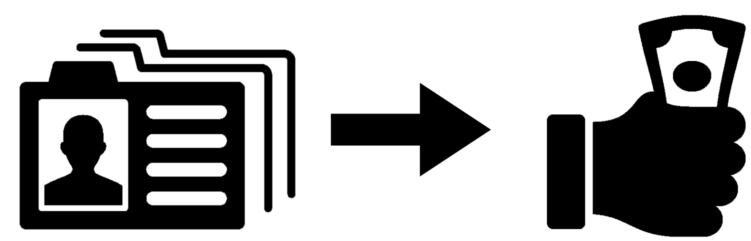

Data size is configurable from 1 TB to 1 PB.

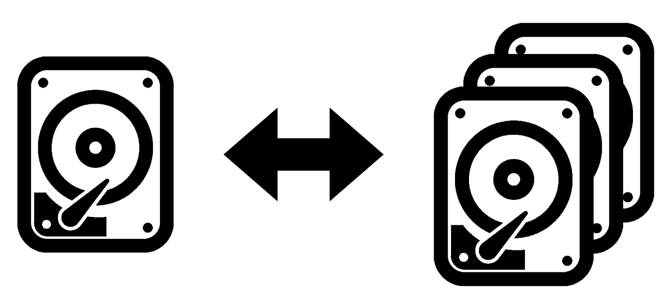

BigBench Queries

BigBench Queries

30 queries operate on collected data to extract useful information.

BigBench Queries

30 queries operate on collected data to extract useful information.

Q01: Find the top 100 products frequently purchased together.

BigBench Queries

30 queries operate on collected data to extract useful information.

Q01: Find the top 100 products frequently purchased together.

Q12: Find customers who viewed certain categories online then made an in-store purchase in the same category.

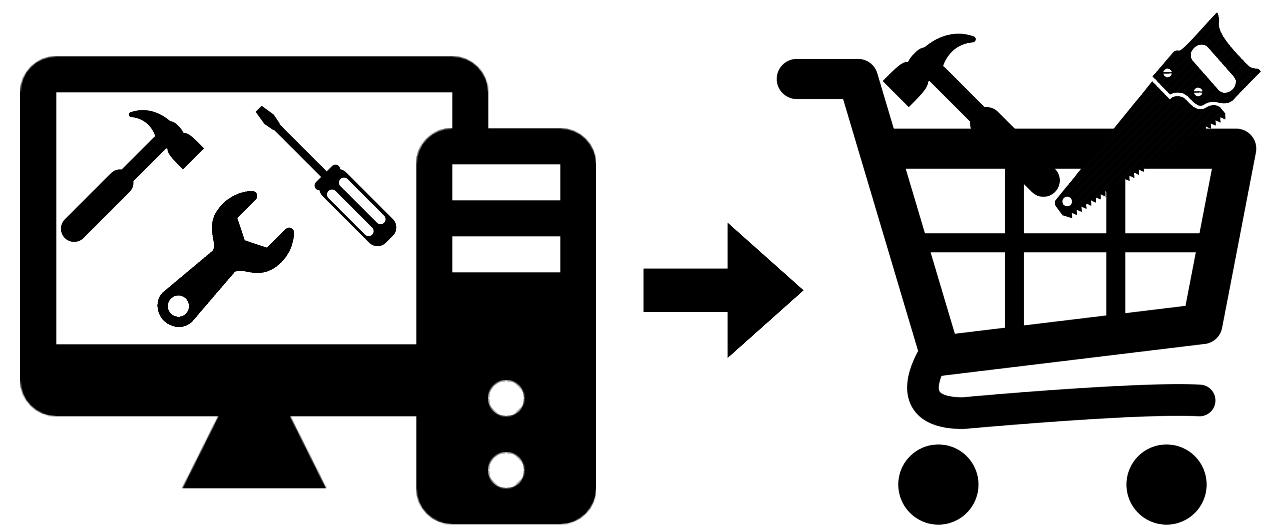

BigBench Queries

30 queries operate on collected data to extract useful information.

Q01: Find the top 100 products frequently purchased together.

Q12: Find customers who viewed certain categories online then made an in-store purchase in the same category.

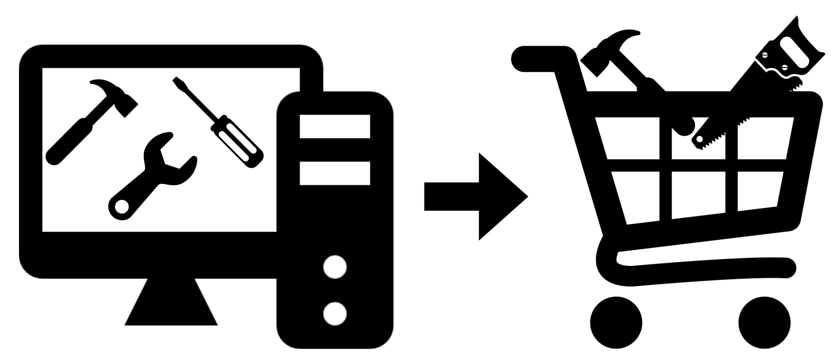

Q27: Extract competitor product names from online reviews.

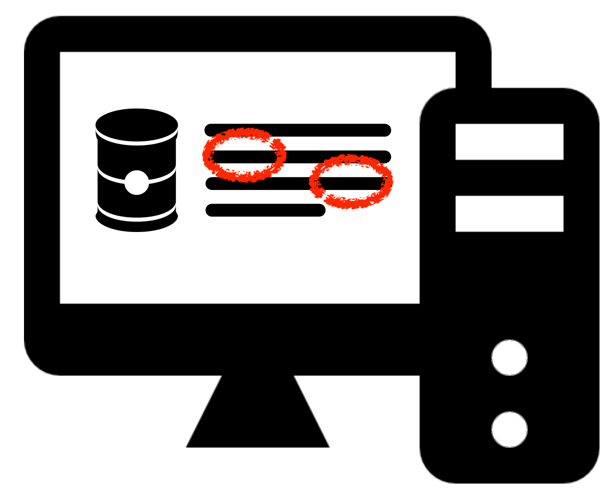

BigBench Queries

30 queries operate on collected data to extract useful information.

Q01: Find the top 100 products frequently purchased together.

Q12: Find customers who viewed certain categories online then made an in-store purchase in the same category.

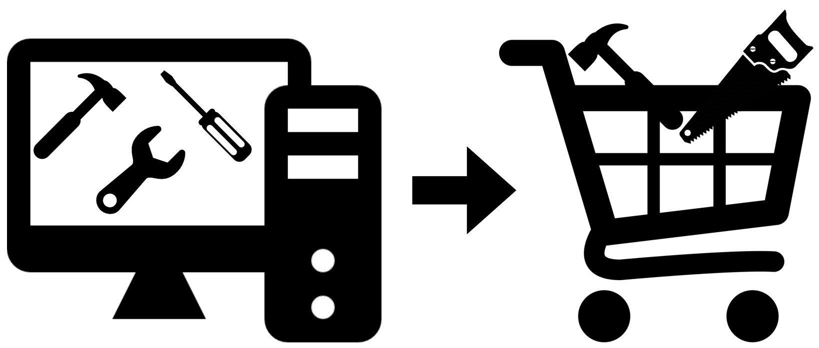

Q27: Extract competitor product names from online reviews.

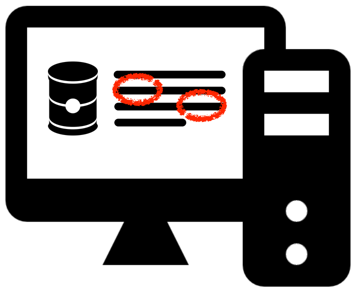

Coverage

Coverage

BigBench data includes

Coverage

BigBench data includes

structured (e.g. customer demographics) data,

Coverage

BigBench data includes

structured (e.g. customer demographics) data,

semi-structured (e.g. web site click streams) data, and

Coverage

BigBench data includes

- **structured** (e.g. customer demographics) data,
- **semi-structured** (e.g. web site click streams) data, and
- **un-structured** (e.g. online product reviews) data.

Coverage

BigBench data includes

- **structured** (e.g. customer demographics) data,
- **semi-structured** (e.g. web site click streams) data, and
- **un-structured** (e.g. online product reviews) data.

BigBench operates on the data using

Coverage

BigBench data includes

structured (e.g. customer demographics) data,
semi-structured (e.g. web site click streams) data, and
un-structured (e.g. online product reviews) data.

BigBench operates on the data using

MapReduce,

Coverage

BigBench data includes

- **structured** (e.g. customer demographics) data,
- **semi-structured** (e.g. web site click streams) data, and
- **un-structured** (e.g. online product reviews) data.

BigBench operates on the data using

- **MapReduce**,
- **machine learning**,

Coverage

BigBench data includes

- **structured** (e.g. customer demographics) data,
- **semi-structured** (e.g. web site click streams) data, and
- **un-structured** (e.g. online product reviews) data.

BigBench operates on the data using

- **MapReduce**,
- **machine learning**,
- **user-defined functions**,

Coverage

BigBench data includes

- **structured** (e.g. customer demographics) data,
- **semi-structured** (e.g. web site click streams) data, and
- **un-structured** (e.g. online product reviews) data.

BigBench operates on the data using

- **MapReduce**,
- **machine learning**,
- **user-defined functions**,
- **query language operations**,

Coverage

BigBench data includes

structured (e.g. customer demographics) data,
semi-structured (e.g. web site click streams) data, and
un-structured (e.g. online product reviews) data.

BigBench operates on the data using

MapReduce,
machine learning,
user-defined functions,
query language operations,
and **natural language processing**.

Coverage

BigBench data includes

structured (e.g. customer demographics) data,
semi-structured (e.g. web site click streams) data, and
un-structured (e.g. online product reviews) data.

BigBench operates on the data using

MapReduce,
machine learning,
user-defined functions,
query language operations,
and **natural language processing**.

Most queries take multiple steps, and many cover multiple data types and operation types.

Why Should We Study One More Benchmark?

Why Should We Study One More Benchmark?

It is essential to understand, emulate, and study industry perspectives in order to produce believable and relevant insights.

Why Should We Study One More Benchmark?

It is essential to understand, emulate, and study industry perspectives in order to produce believable and relevant insights.

There is a lot of heterogeneity in the applications, something that has been oversimplified in past benchmarking efforts.

Why Should We Study One More Benchmark?

It is essential to understand, emulate, and study industry perspectives in order to produce believable and relevant insights.

There is a lot of heterogeneity in the applications, something that has been oversimplified in past benchmarking efforts.

Thread-limited execution is still pervasive even in scale-out big data analytics and demands better scale-up performance (Amdahl's law).

Experimental Setup

Experimental Setup

3+1 setup: 3 workers, 1 master.

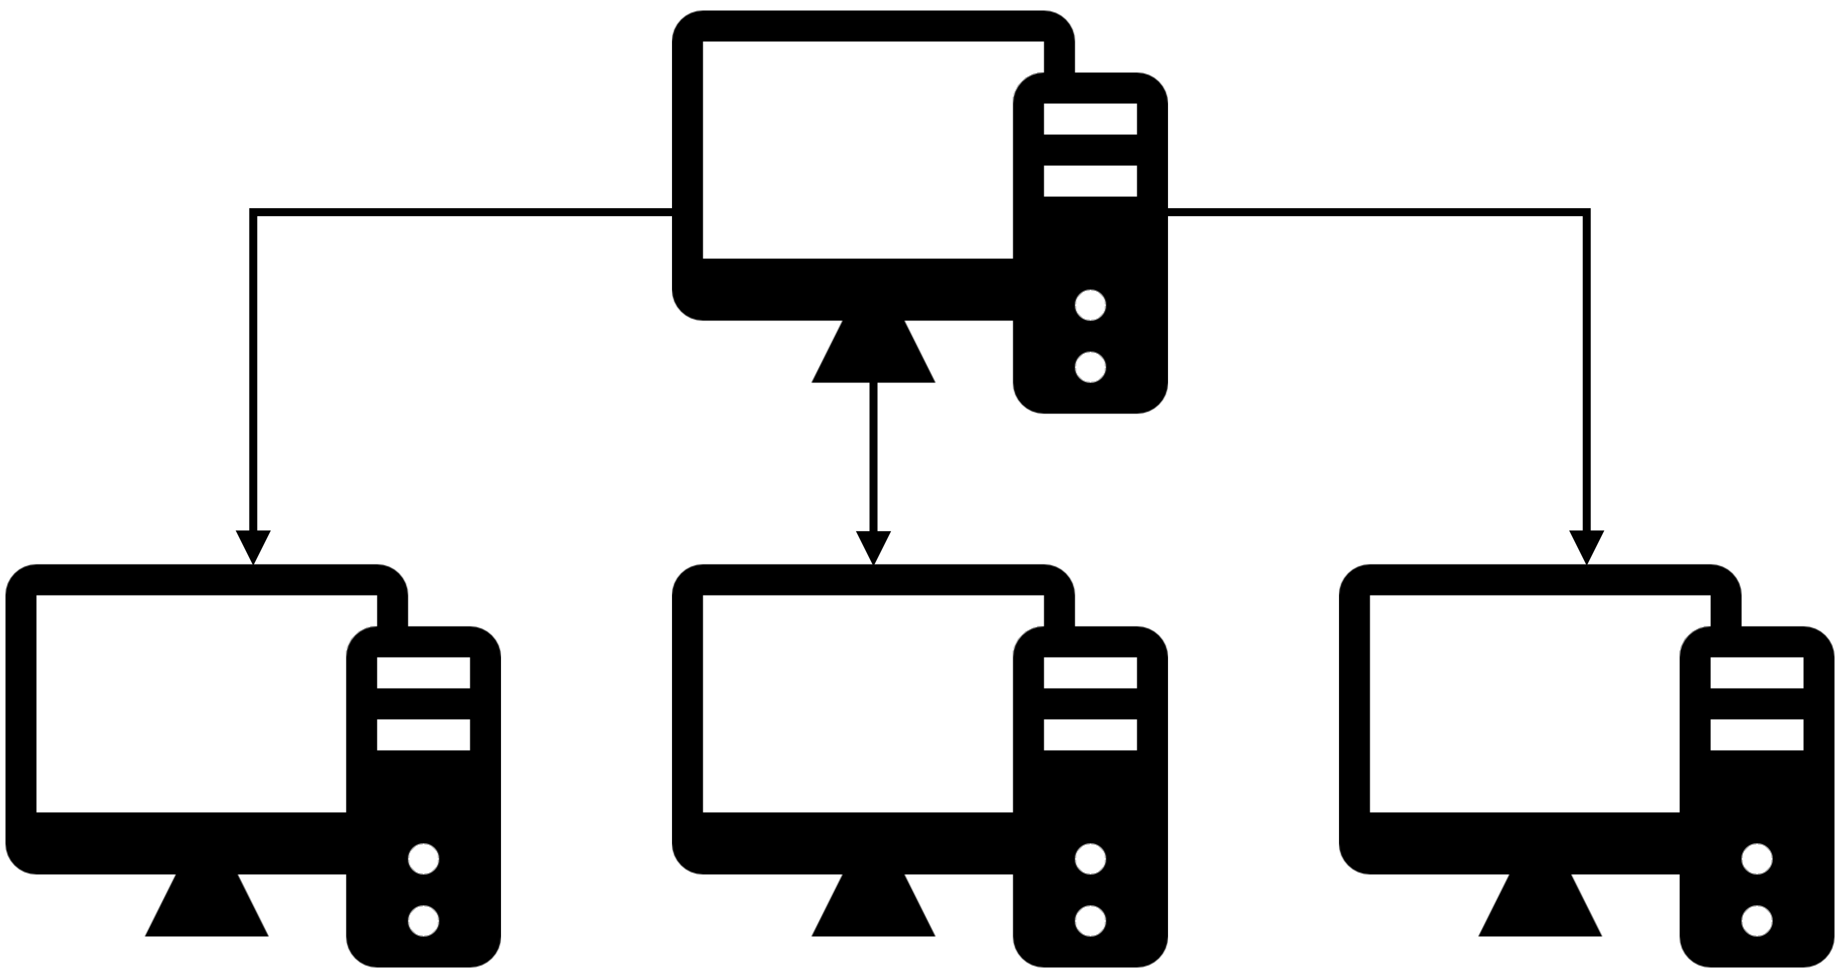

Experimental Setup

3+1 setup: 3 workers, 1 master.

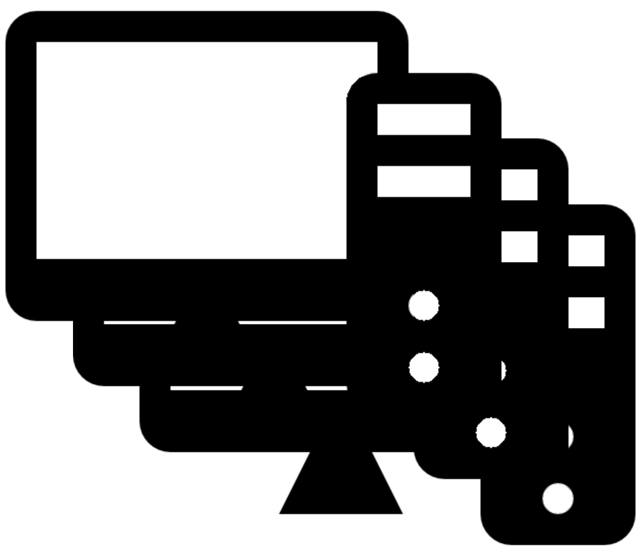

Worker nodes:

- Intel Xeon E5-2699 v3
- 384 GB RAM
- SSD Storage
- 10 Gbps Ethernet

Experimental Setup

3+1 setup: 3 workers, 1 master.

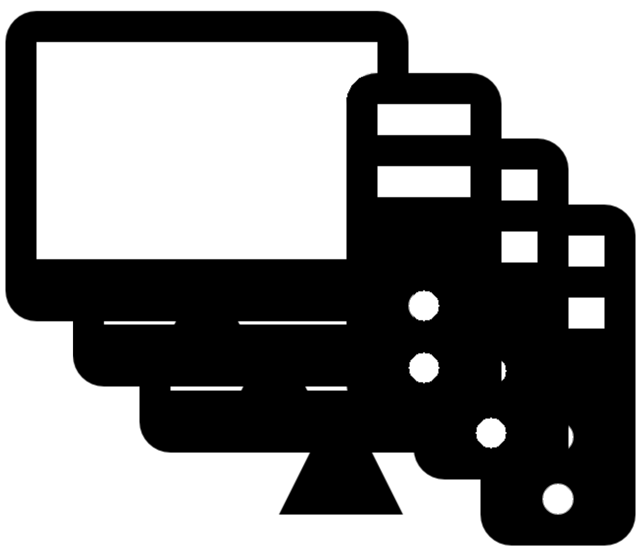

Worker nodes:

- Intel Xeon E5-2699 v3
- 384 GB RAM
- SSD Storage
- 10 Gbps Ethernet

Experimental Setup

3+1 setup: 3 workers, 1 master.

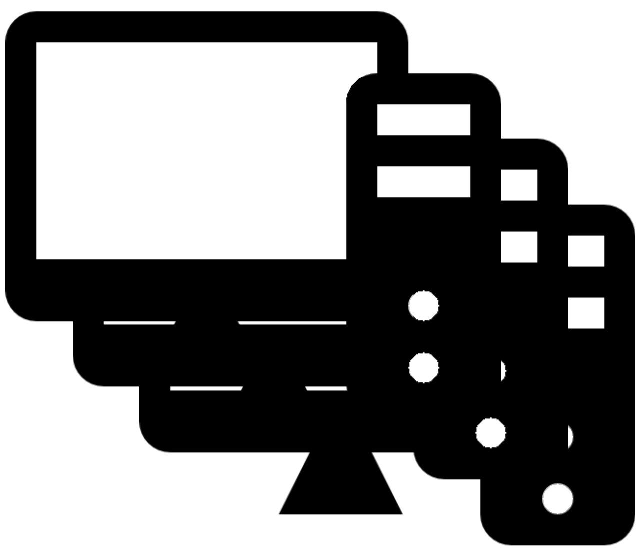

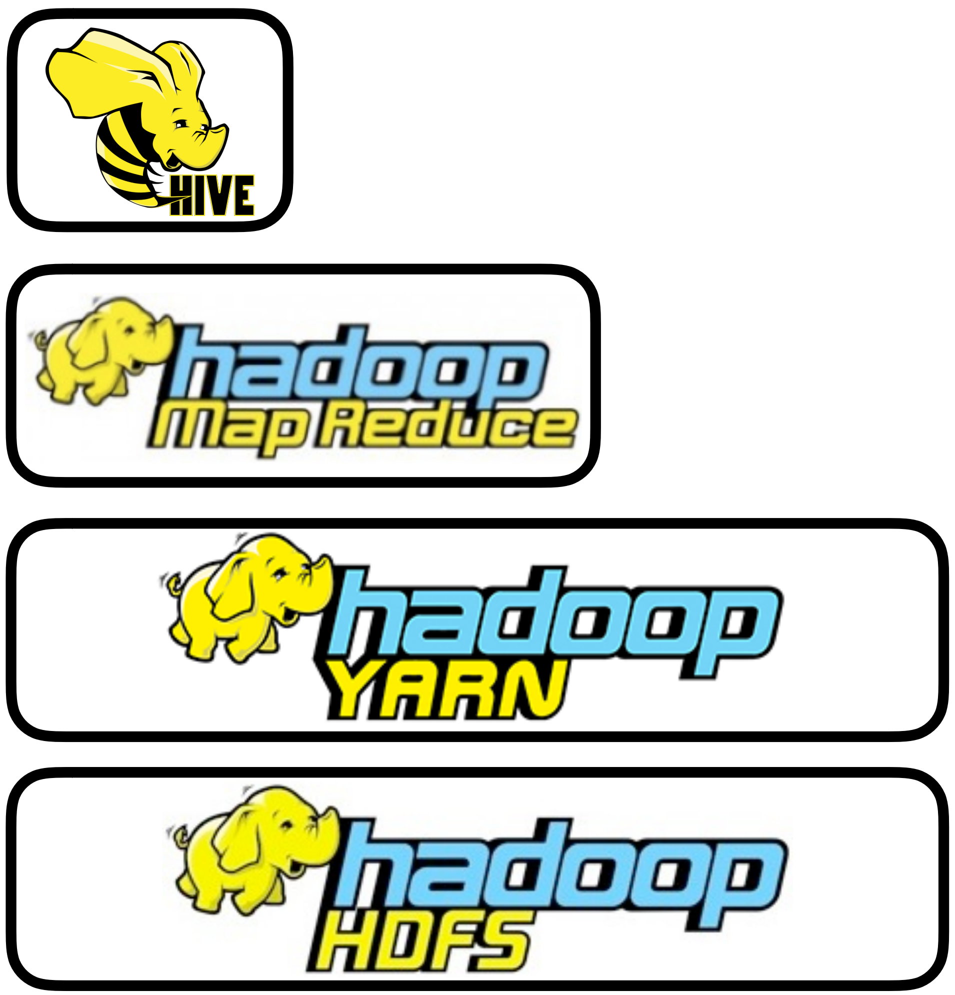

Worker nodes:

- Intel Xeon E5-2699 v3
- 384 GB RAM
- SSD Storage
- 10 Gbps Ethernet

Experimental Setup

3+1 setup: 3 workers, 1 master.

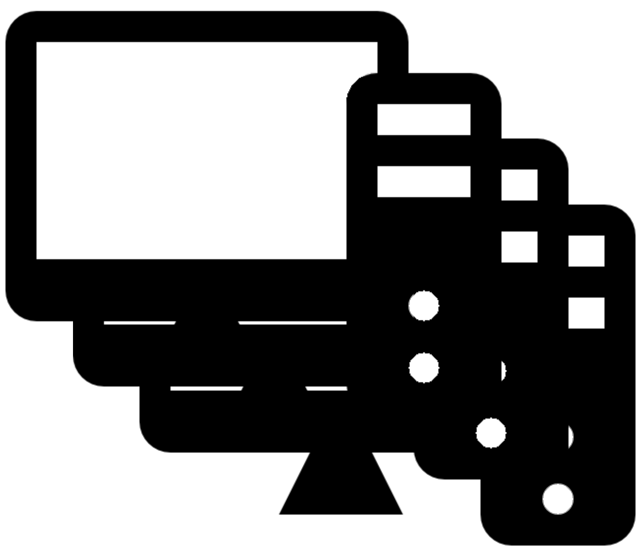

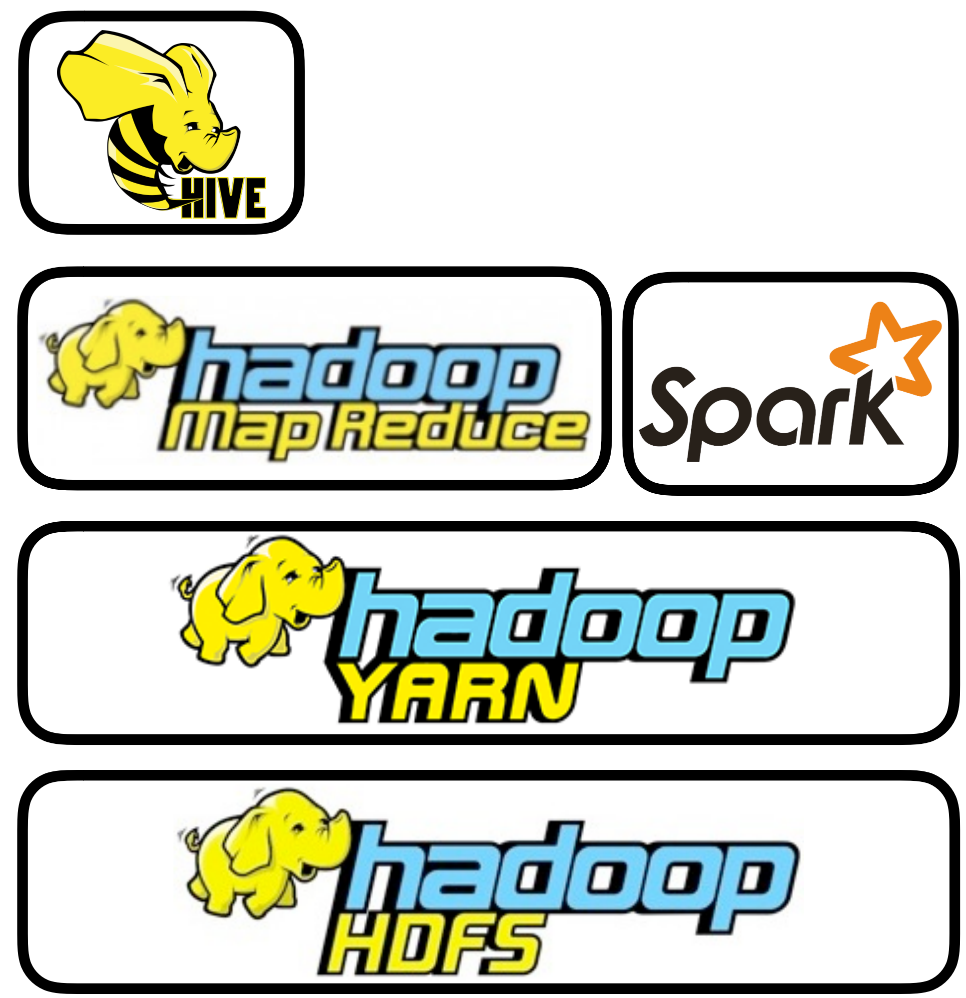

Worker nodes:

- Intel Xeon E5-2699 v3
- 384 GB RAM
- SSD Storage
- 10 Gbps Ethernet

Experimental Setup

3+1 setup: 3 workers, 1 master.

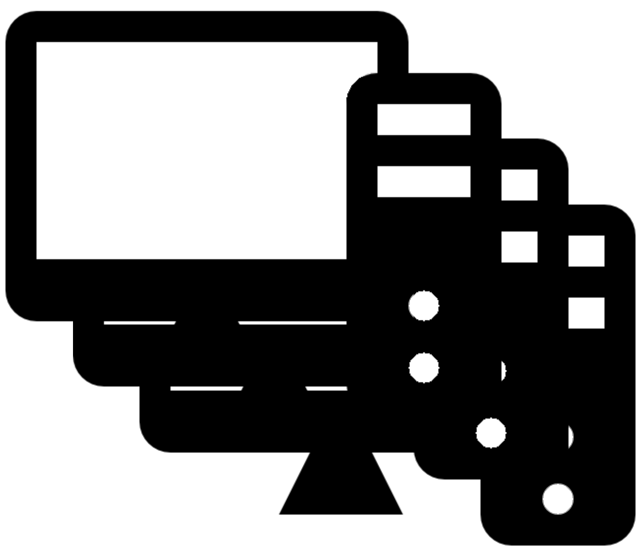

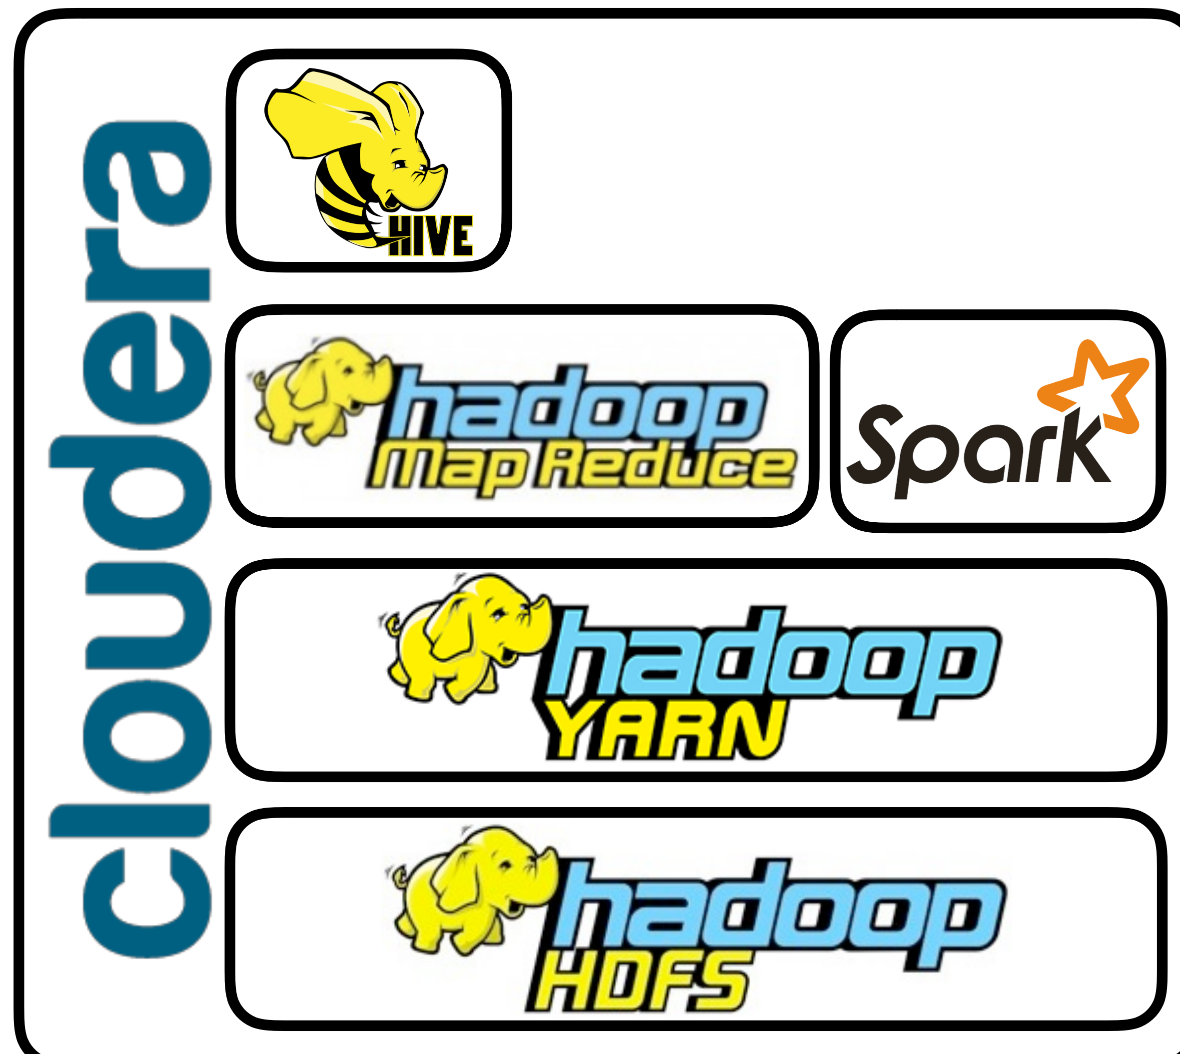

Worker nodes:

- Intel Xeon E5-2699 v3
- 384 GB RAM
- SSD Storage
- 10 Gbps Ethernet

Experimental Setup

3+1 setup: 3 workers, 1 master.

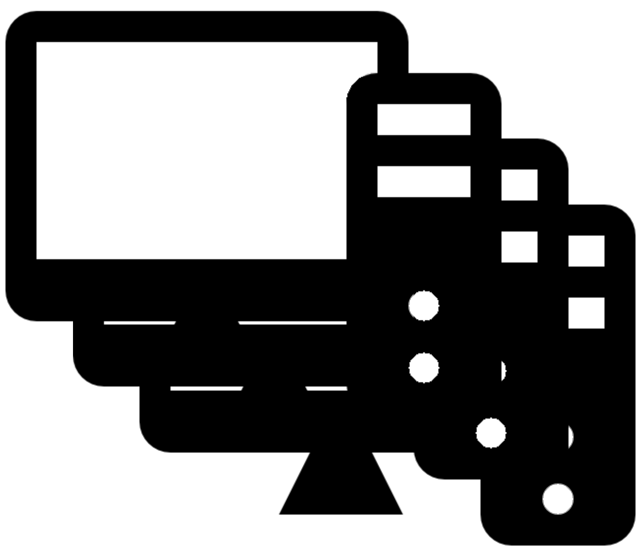

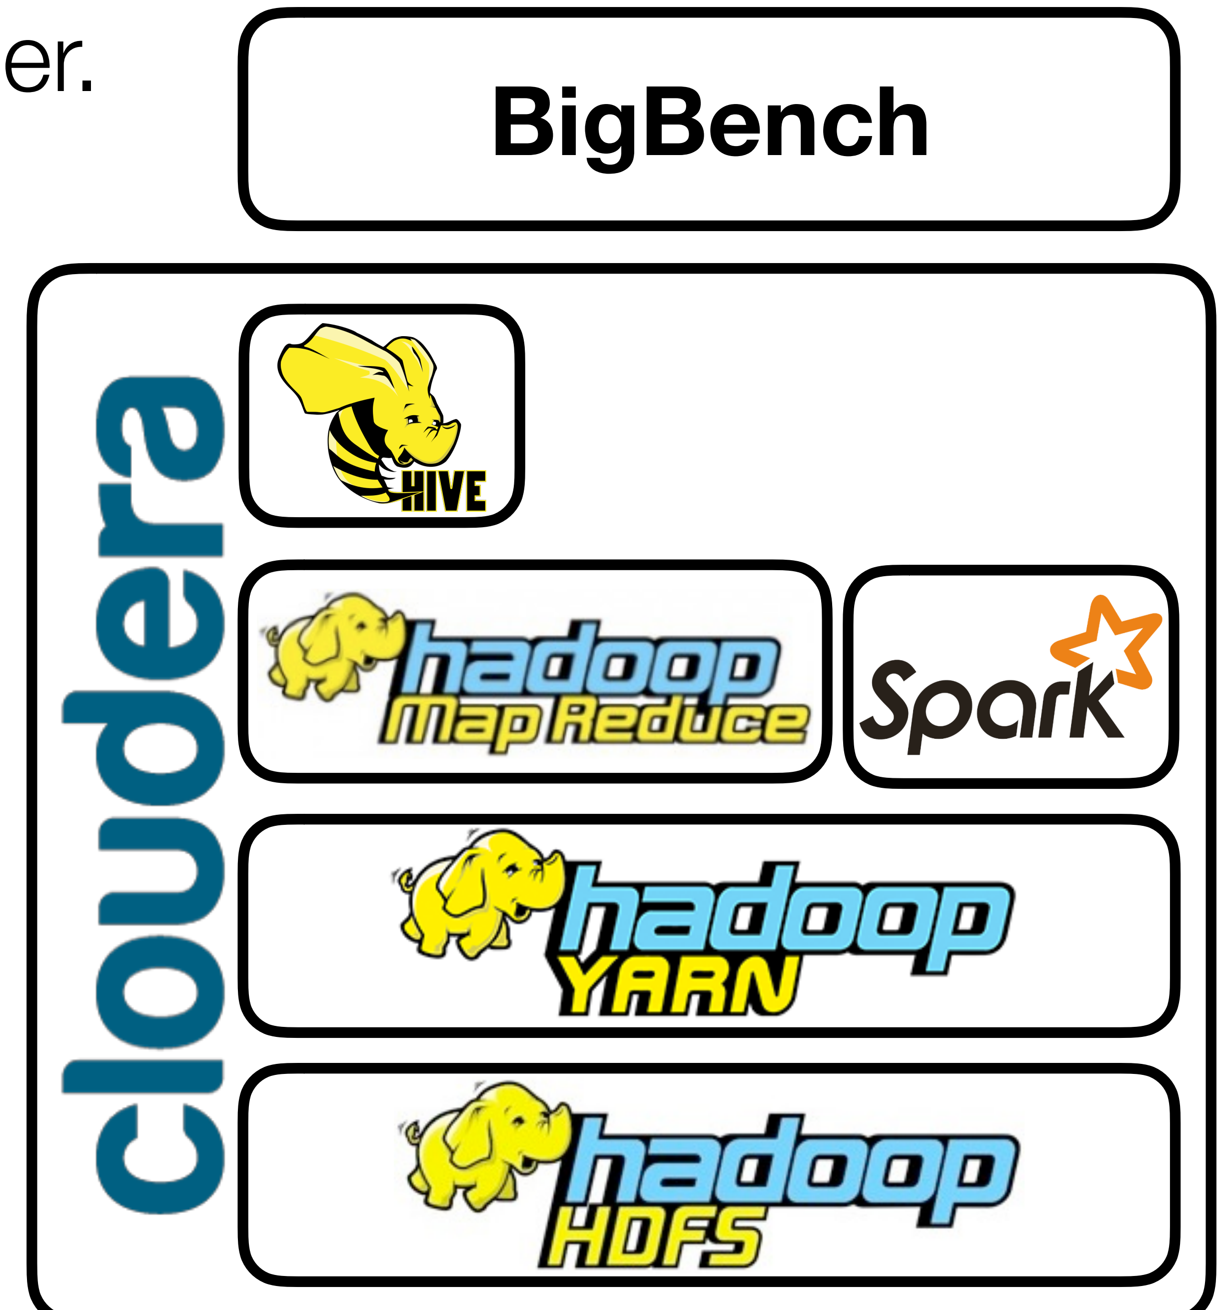

Worker nodes:

- Intel Xeon E5-2699 v3
- 384 GB RAM
- SSD Storage
- 10 Gbps Ethernet

BigBench Stages

BigBench Stages

Data Generation
Creates the data used by the benchmark.

BigBench Stages

Data Generation

Creates the data used by the benchmark.

Load Test

Extracts/transforms/loads the data into a usable form.

BigBench Stages

Data Generation
Creates the data used by the benchmark.

Load Test
Extracts/transforms/loads the data into a usable form.

Power Test
Runs each query sequentially to minimize latency.

BigBench Stages

Data Generation
Creates the data used by the benchmark.

Load Test
Extracts/transforms/loads the data into a usable form.

Power Test
Runs each query sequentially to minimize latency.

Throughput Test
Runs multiple queries at once to maximize throughput.

BigBench Stages

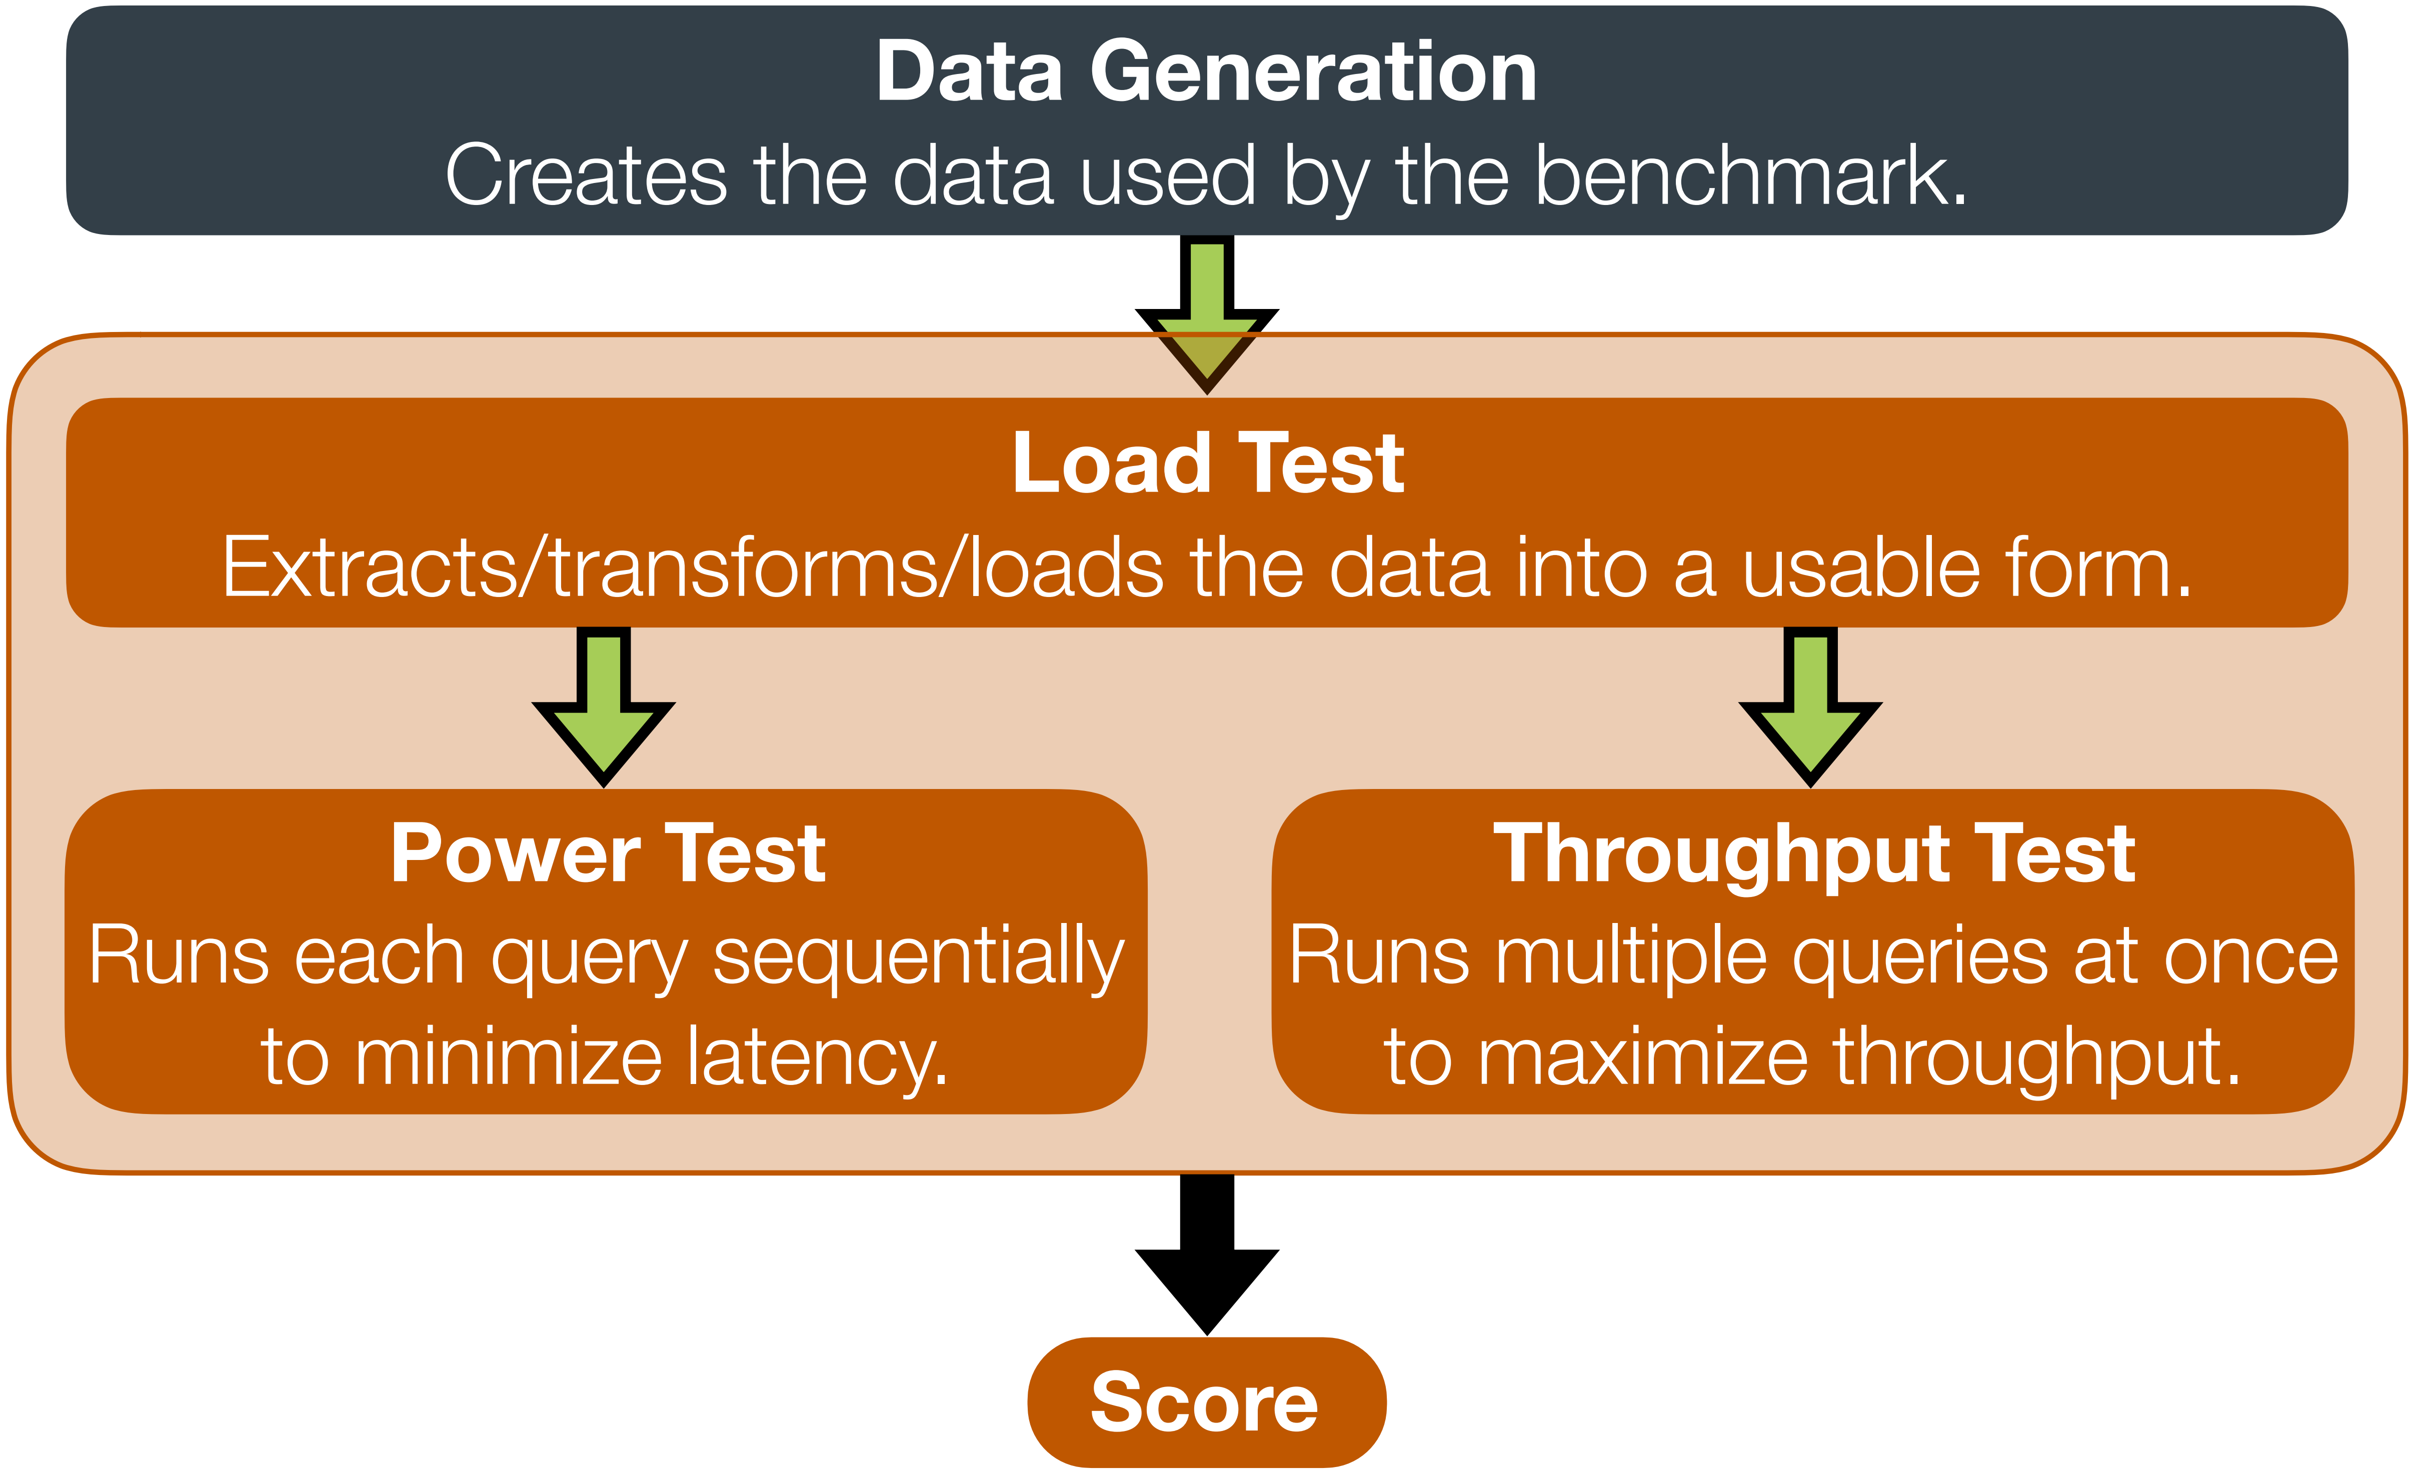

BigBench Stages

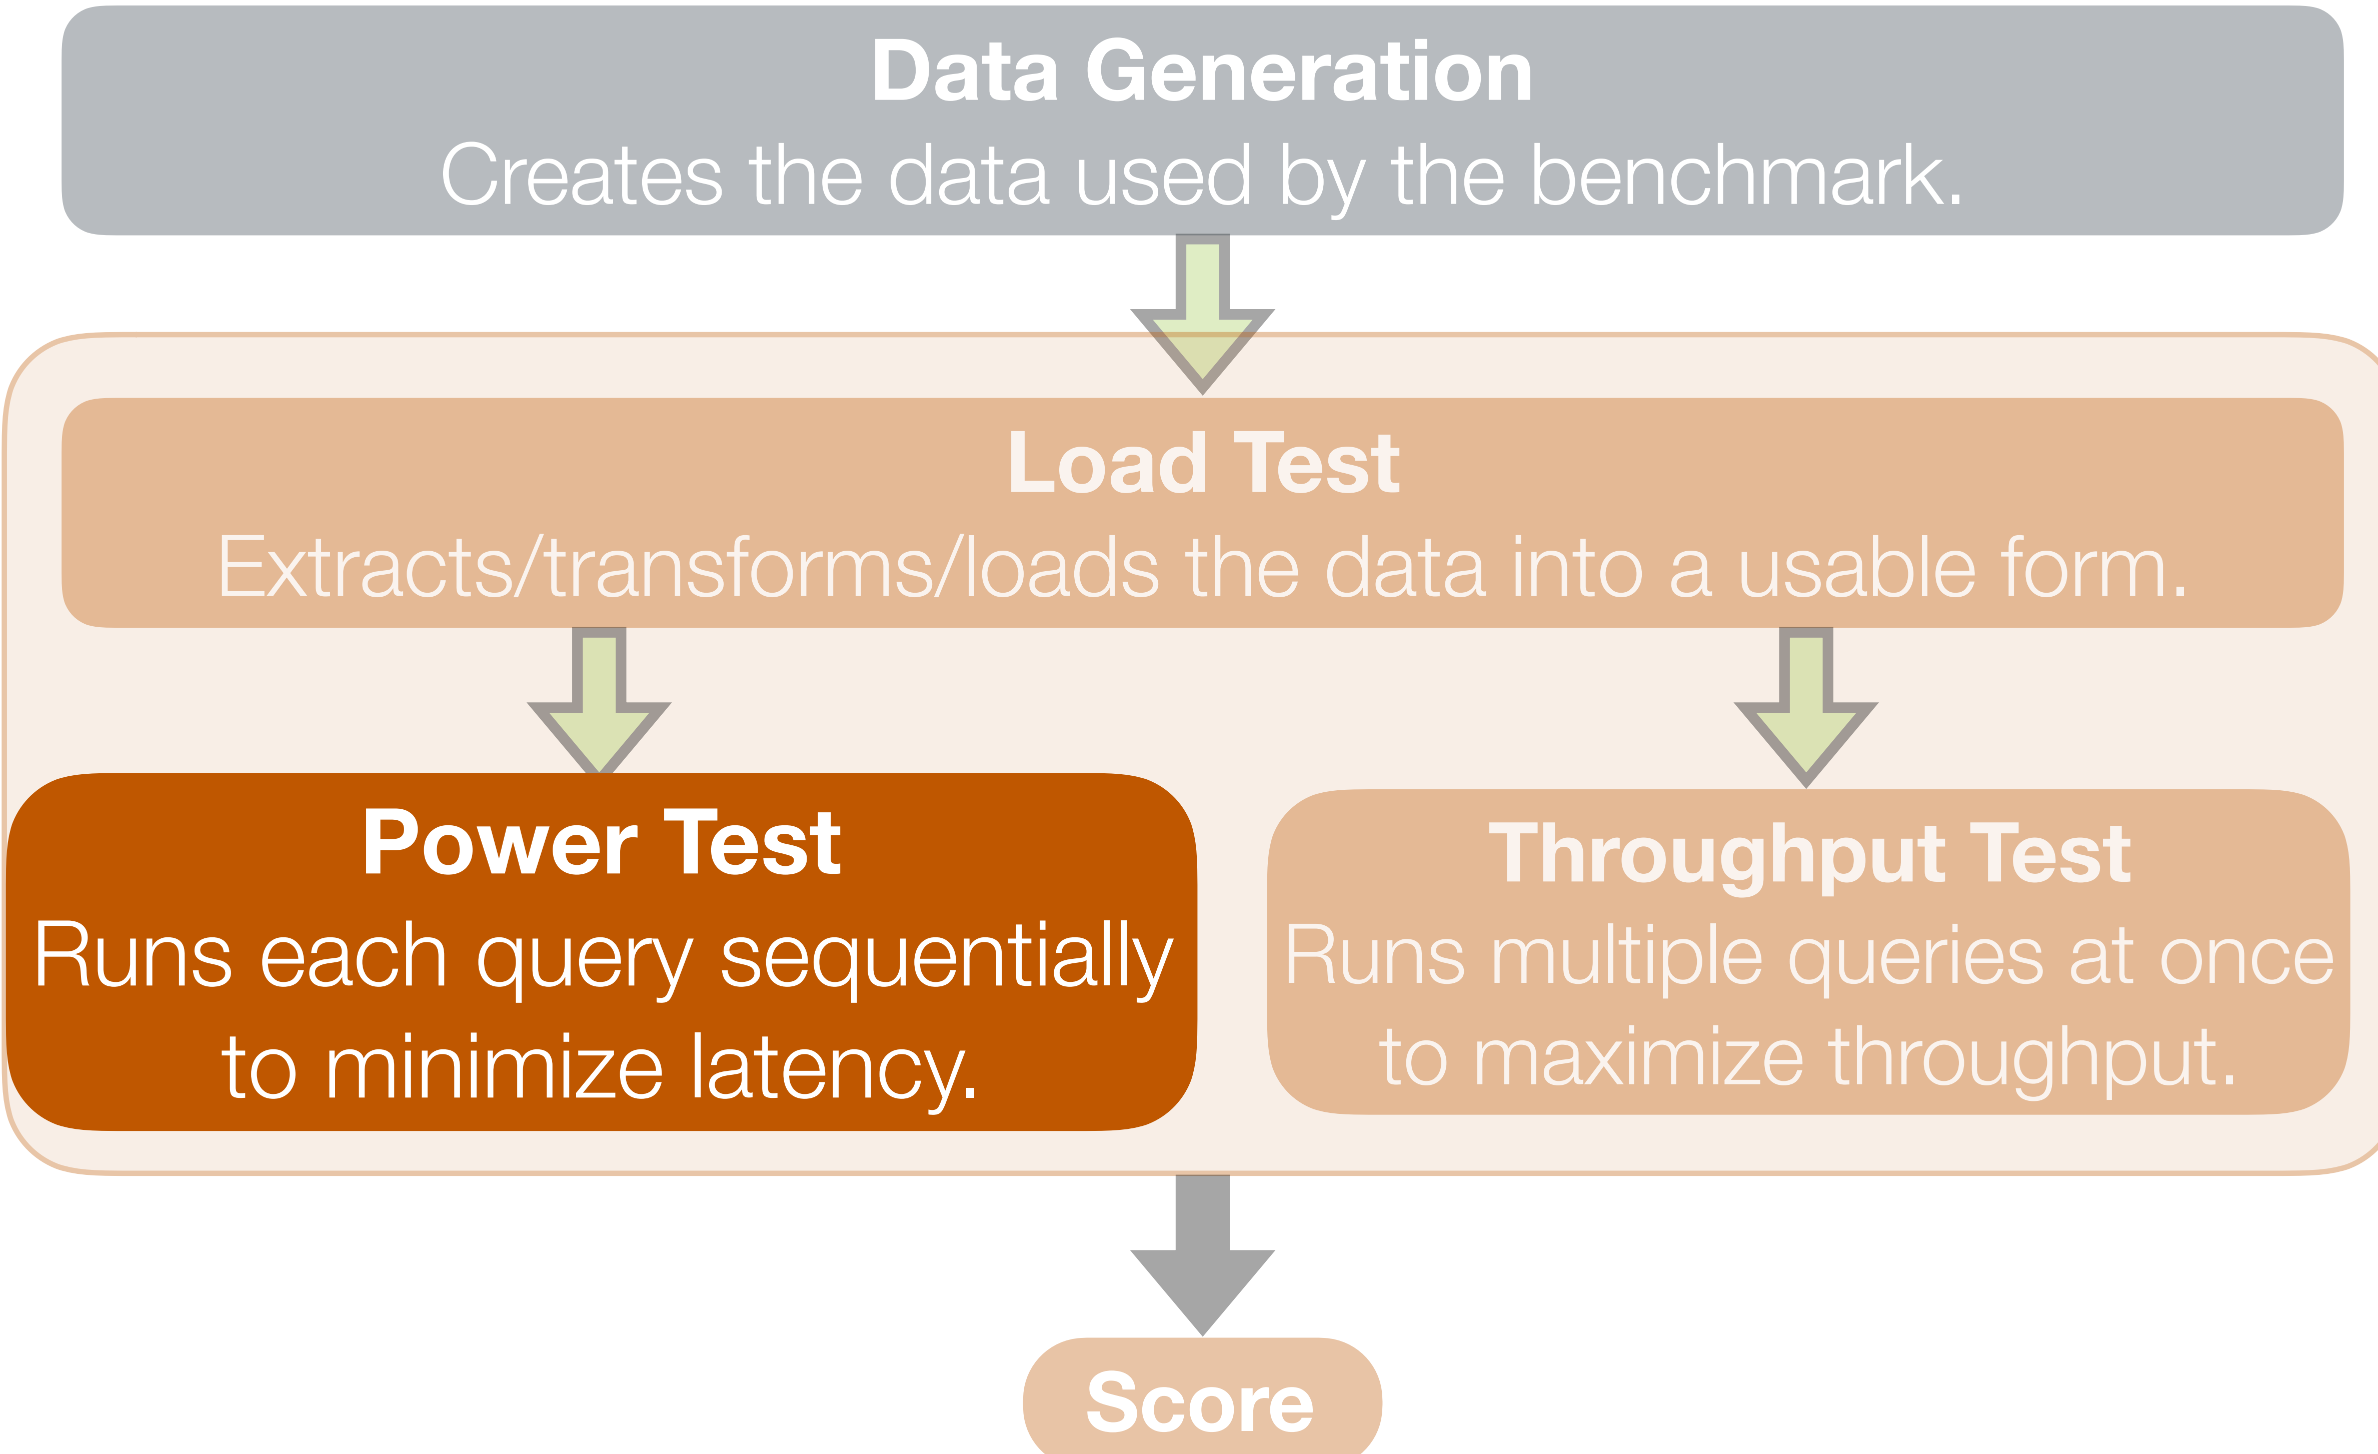

Characterization:
A Caution Against Oversimplification

Characterization:
A Caution Against Oversimplification

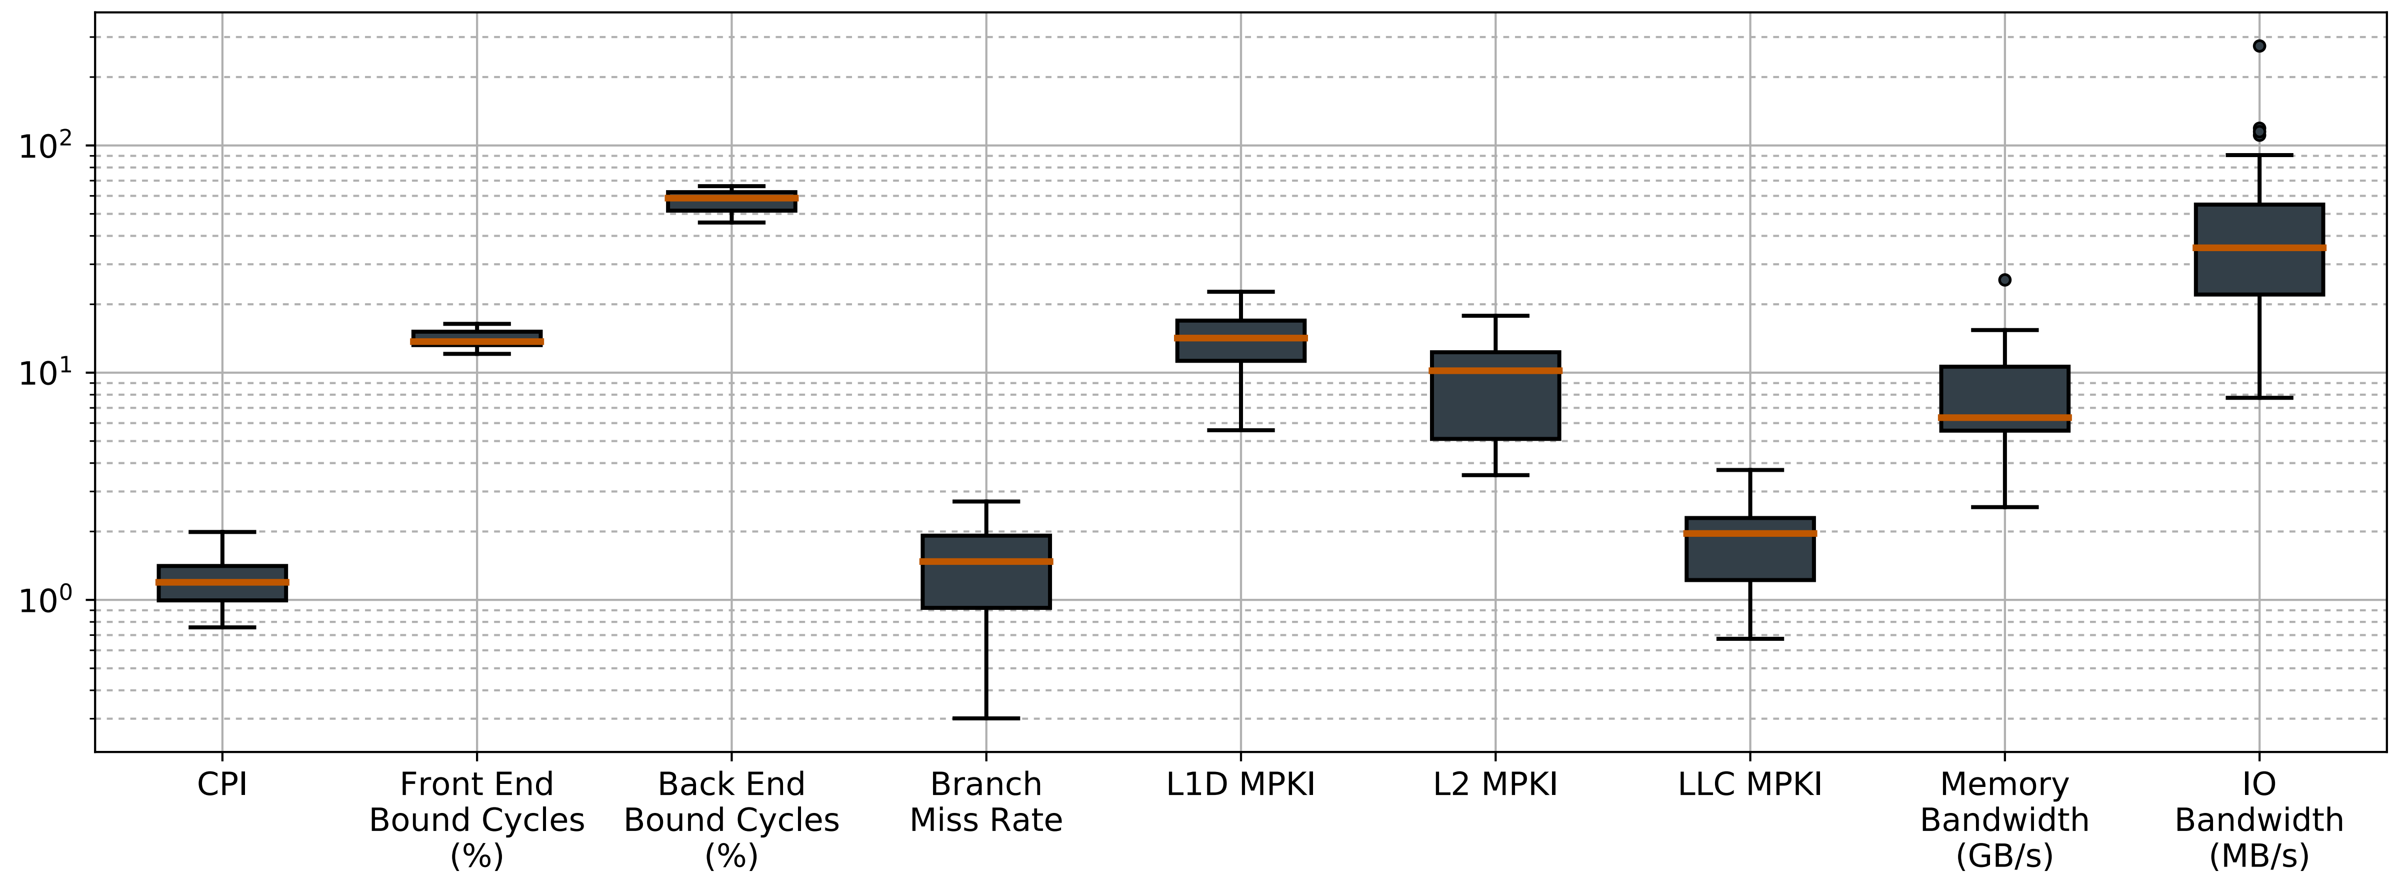

Characterization: A Caution Against Oversimplification

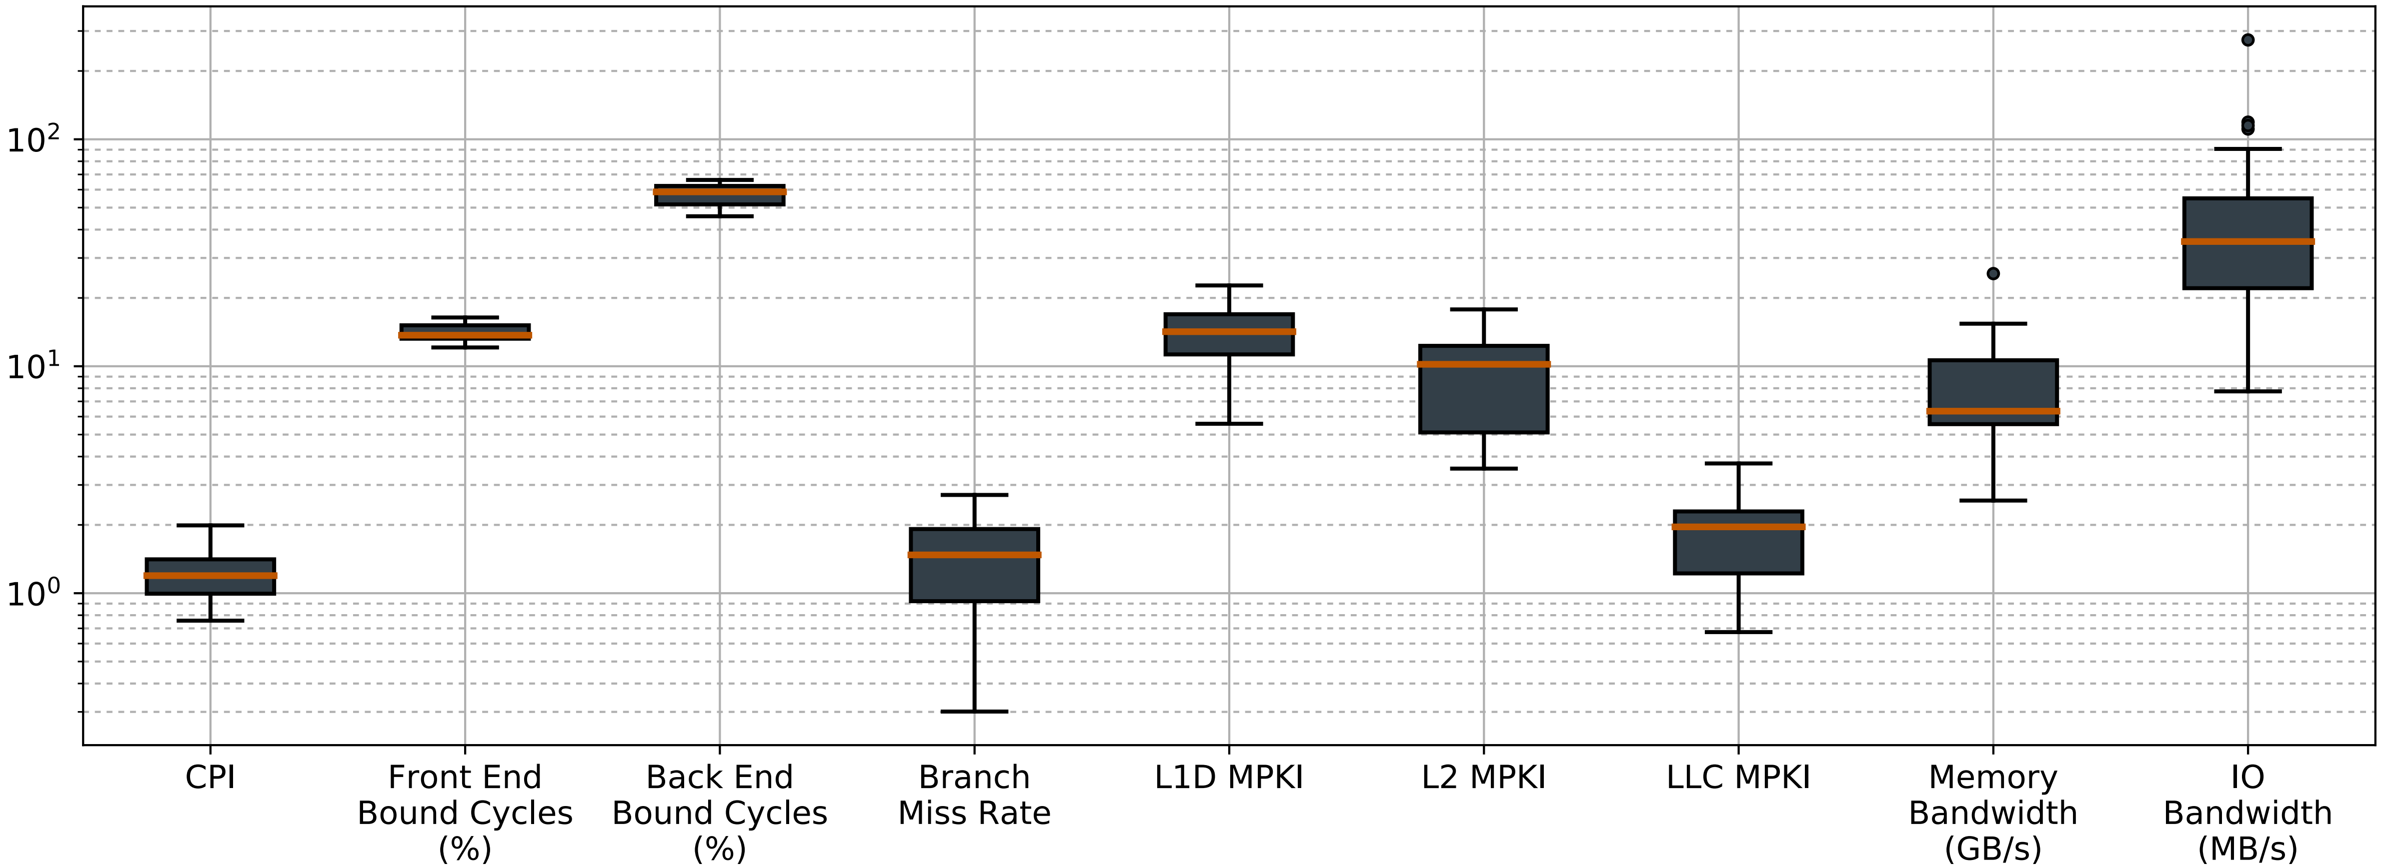

Compute

Characterization:
A Caution Against Oversimplification

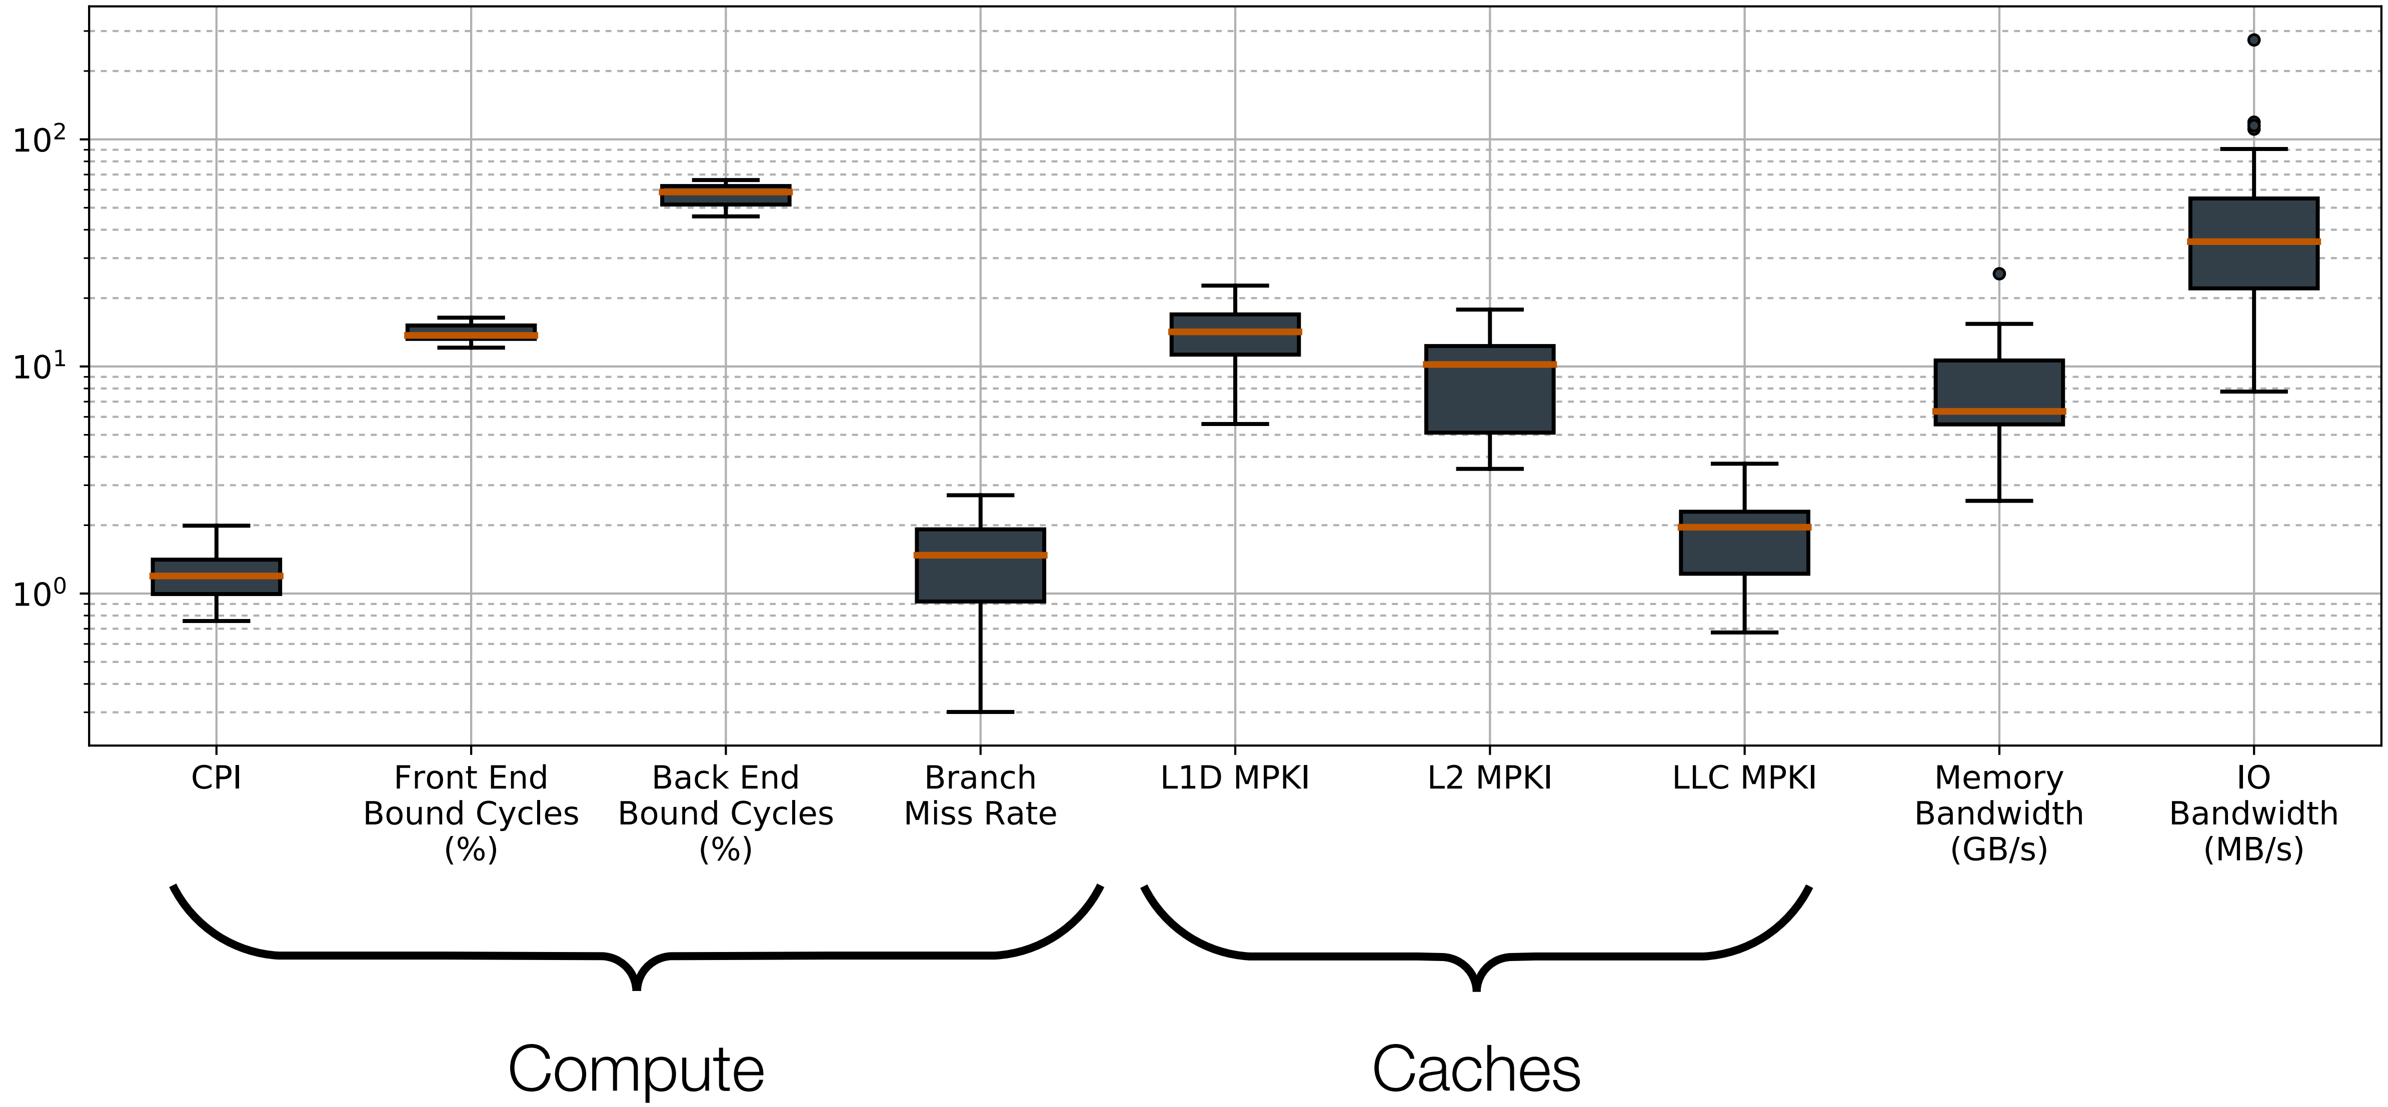

Characterization: A Caution Against Oversimplification

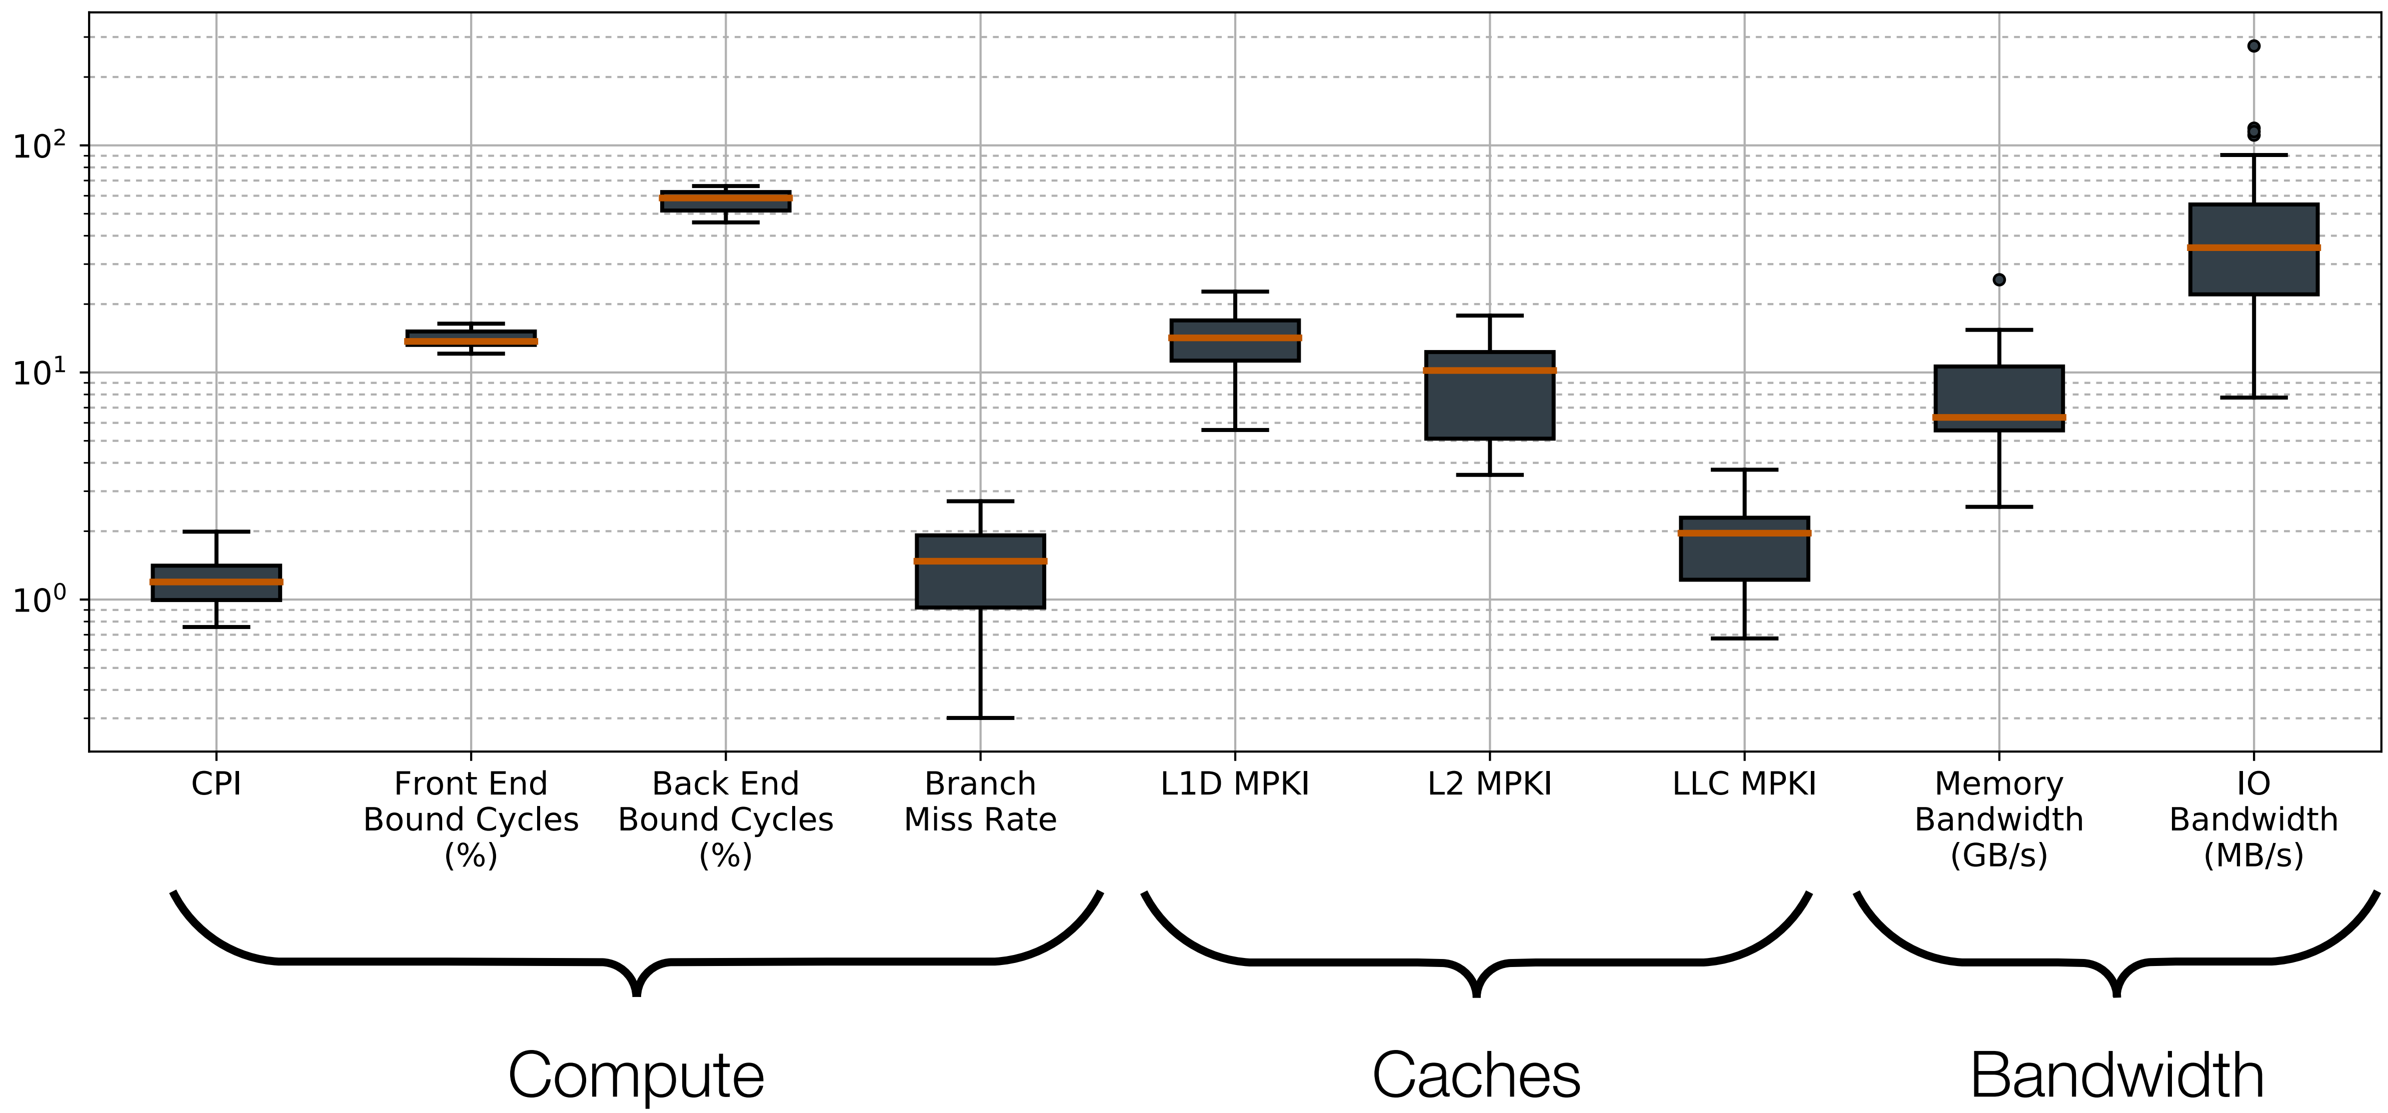

Characterization:
A Caution Against Oversimplification

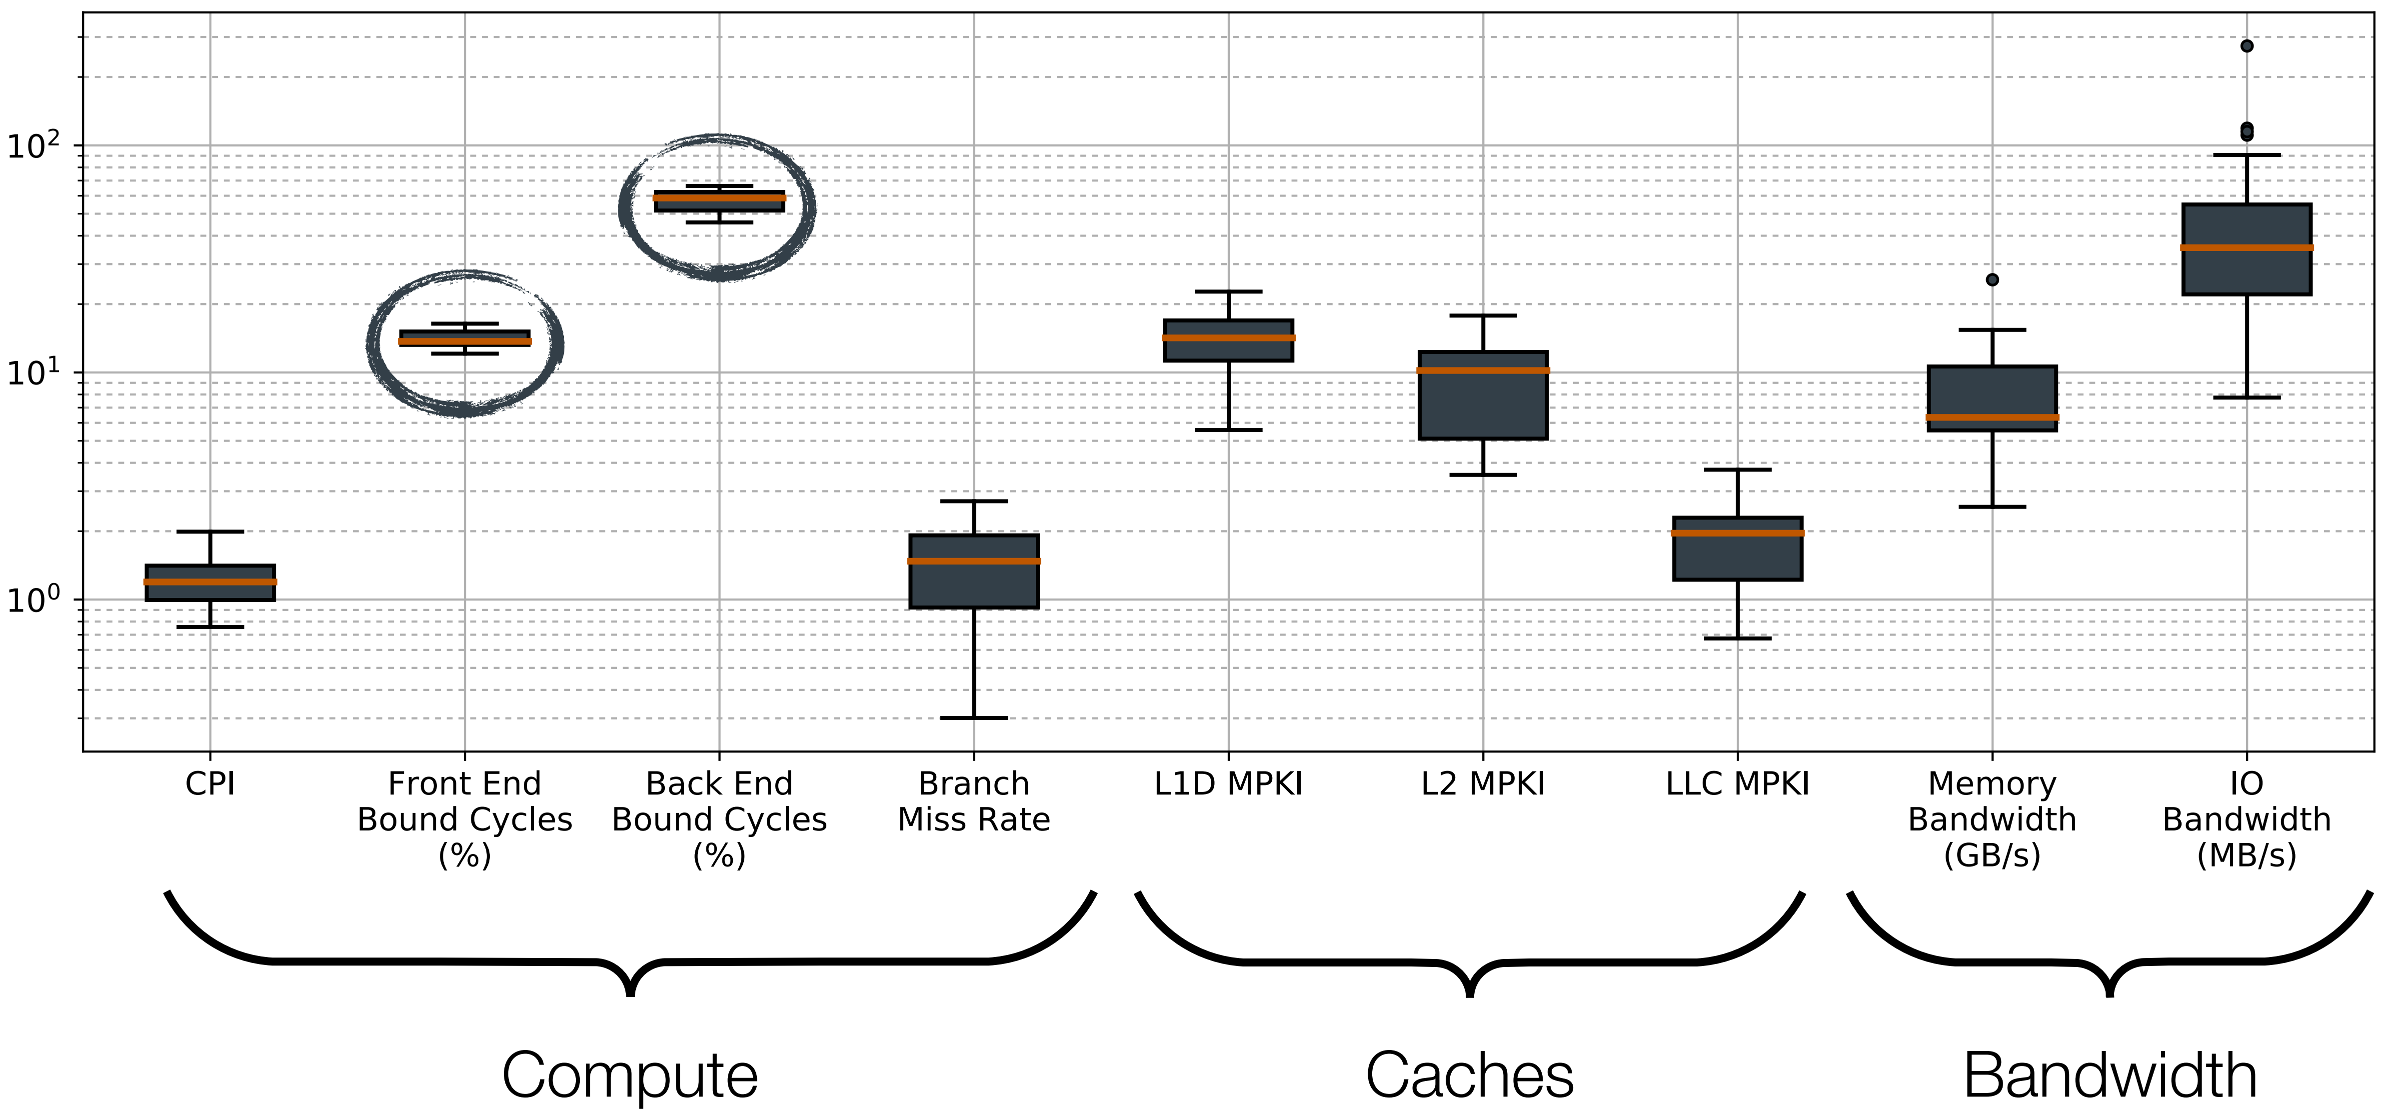

Only the front and back end-bound cycles are stable across queries.

Characterization:
A Caution Against Oversimplification

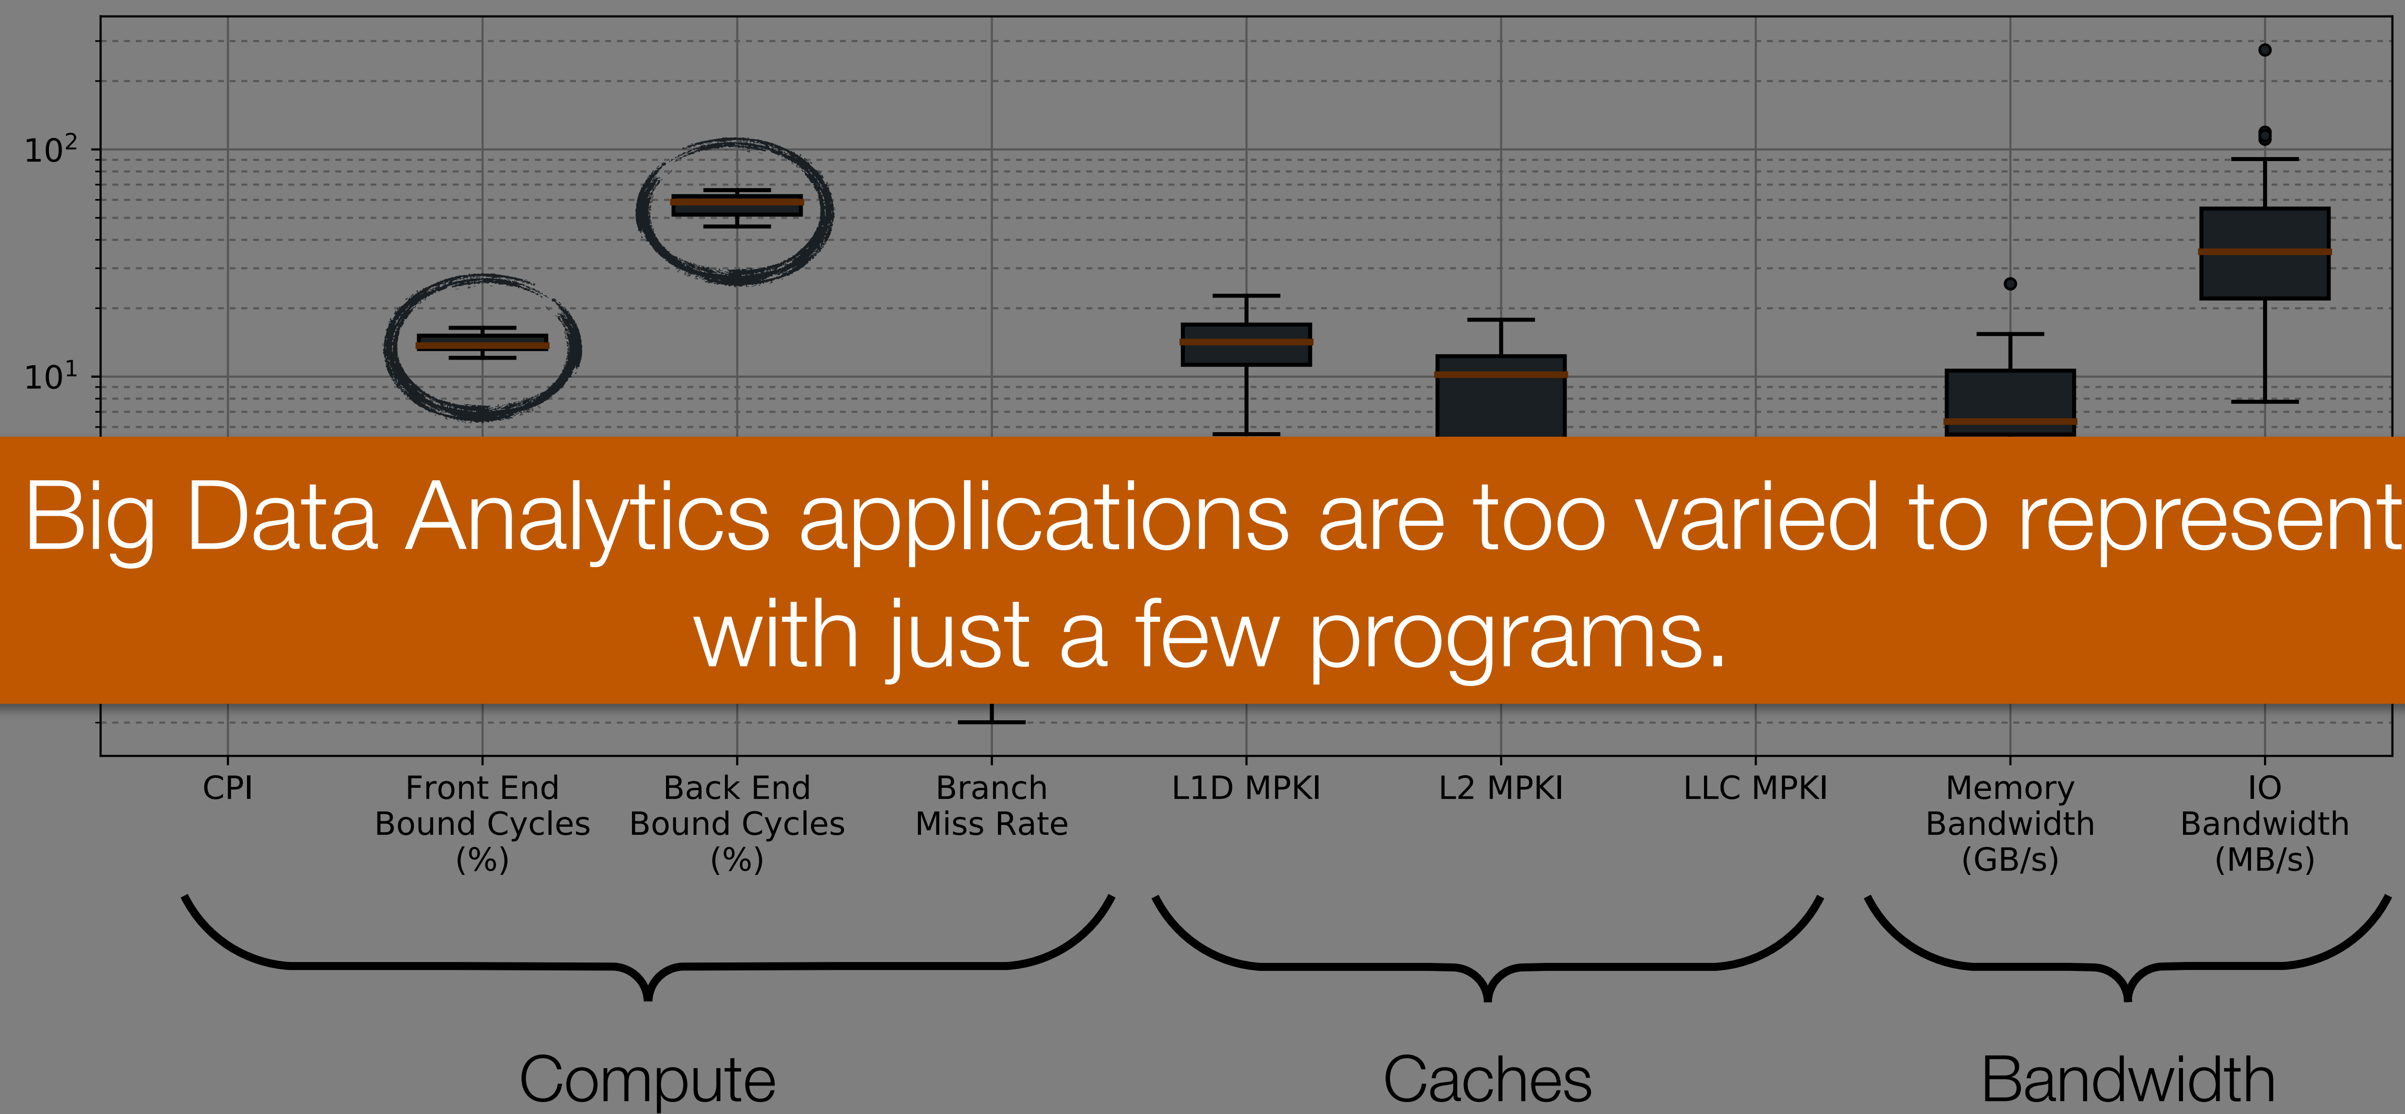

Big Data Analytics applications are too varied to represent with just a few programs.

Only the front and back end-bound cycles are stable across queries.

Big Data Analytics is a Rich and Varied Field

Big Data Analytics is a Rich and Varied Field

Computational Method

Big Data Analytics is a Rich and Varied Field

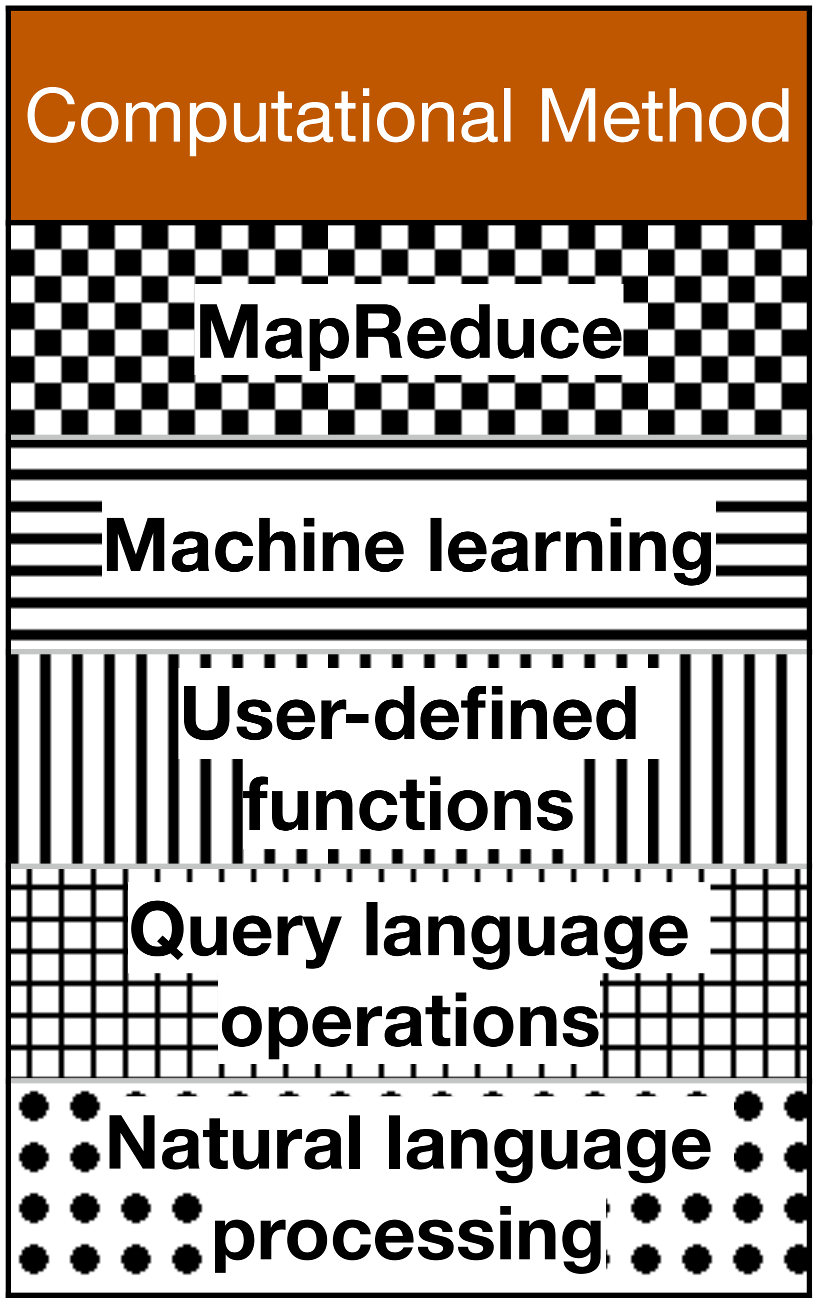

Big Data Analytics is a Rich and Varied Field

Computational Method

- MapReduce
- Machine learning
- User-defined functions
- Query language operations
- Natural language processing

Data Type

Big Data Analytics is a Rich and Varied Field

Computational Method
MapReduce
Machine learning
User-defined functions
Query language operations
Natural language processing

Data Type
Structured
Semi-Structured
Un-Structured

Big Data Analytics is a Rich and Varied Field

Computational Method
MapReduce
Machine learning
User-defined functions
Query language operations
Natural language processing

×

Data Type
Structured
Semi-Structured
Un-Structured

=

Big Data Analytics is a Rich and Varied Field

Computational Method

- MapReduce
- Machine learning
- User-defined functions
- Query language operations
- Natural language processing

×

Data Type

- Structured
- Semi-Structured
- Un-Structured

=

Queries

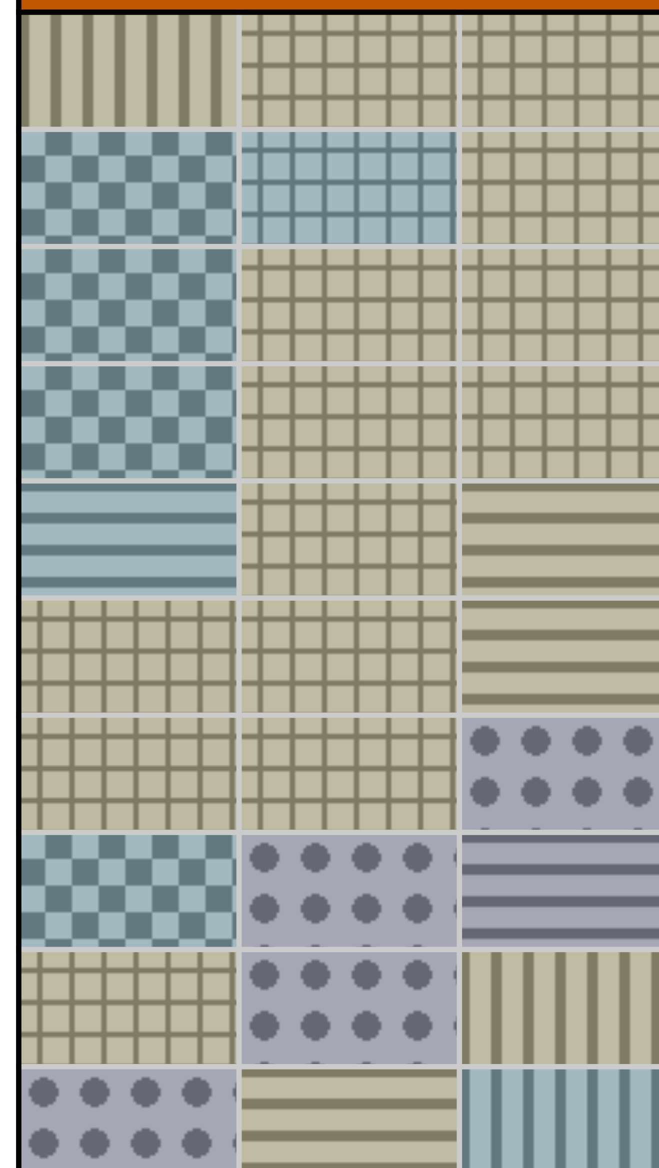

With the Promises of what Big Data Can Do, We Forget About What it Cannot Do

With the Promises of what Big Data Can Do, We Forget About What it Cannot Do

> “Hadoop makes reliable, scalable, distributed computing possible on industry-standard servers…Hadoop is also designed to scale from a single server to thousands of machines.”
>
> **Microsoft Azure’s Hadoop Page**

With the Promises of what Big Data Can Do, We Forget About What it Cannot Do

> "Hadoop makes reliable, scalable, distributed computing possible on industry-standard servers…Hadoop is also designed to scale from a single server to thousands of machines."
>
> **Microsoft Azure's Hadoop Page**

> "Apache Hadoop is a highly scalable storage platform designed to process very large data sets across hundreds to thousands of computing nodes that operate in parallel."
>
> **IBM's Hadoop Page**

With the Promises of what Big Data Can Do, We Forget About What it Cannot Do

> "Hadoop makes reliable, scalable, distributed computing possible on industry-standard servers…Hadoop is also designed to scale from a single server to thousands of machines."
>
> **Microsoft Azure's Hadoop Page**

> "Apache Hadoop is a highly scalable storage platform designed to process very large data sets across hundreds to thousands of computing nodes that operate in parallel."
>
> **IBM's Hadoop Page**

> "Architecturally, the reason you're able to deal with lots of data is because Hadoop spreads it out…You've got all of these processors, working in parallel, harnessed together."
>
> **Cloudera Co-Founder Mike Olson**

With the Promises of what Big Data Can Do, We Forget About What it Cannot Do

"Hadoop makes reliable, scalable, distributed computing possible

We see these claims about Hadoop scaling to thousands of machines and we have a tendency to think we don't need to worry about scale-up.

"Apache Hadoop is a highly scalable storage platform designed to process very large data sets across hundreds to thousands of computing nodes that operate in parallel."

IBM's Hadoop Page

"Architecturally, the reason you're able to deal with lots of data is because Hadoop spreads it out…You've got all of these processors, working in parallel, harnessed together."

Cloudera Co-Founder Mike Olson

With the Promises of what Big Data Can Do, We Forget About What it Cannot Do

“ Hadoop makes reliable, scalable, distributed computing possible

We see these claims about Hadoop scaling to thousands of machines and we have a tendency to think we don’t need to worry about scale-up.

“ Apache Hadoop is a highly scalable storage platform designed to process very large data sets across hundreds to thousands of ”

In reality, scale-up performance is just as important in big data analytics as it has ever been.

“ because Hadoop spreads it out…You’ve got all of these processors, working in parallel, harnessed together. ”

Cloudera Co-Founder Mike Olson

Amdahl’s Law is Alive and Kicking

Amdahl’s Law is Alive and Kicking

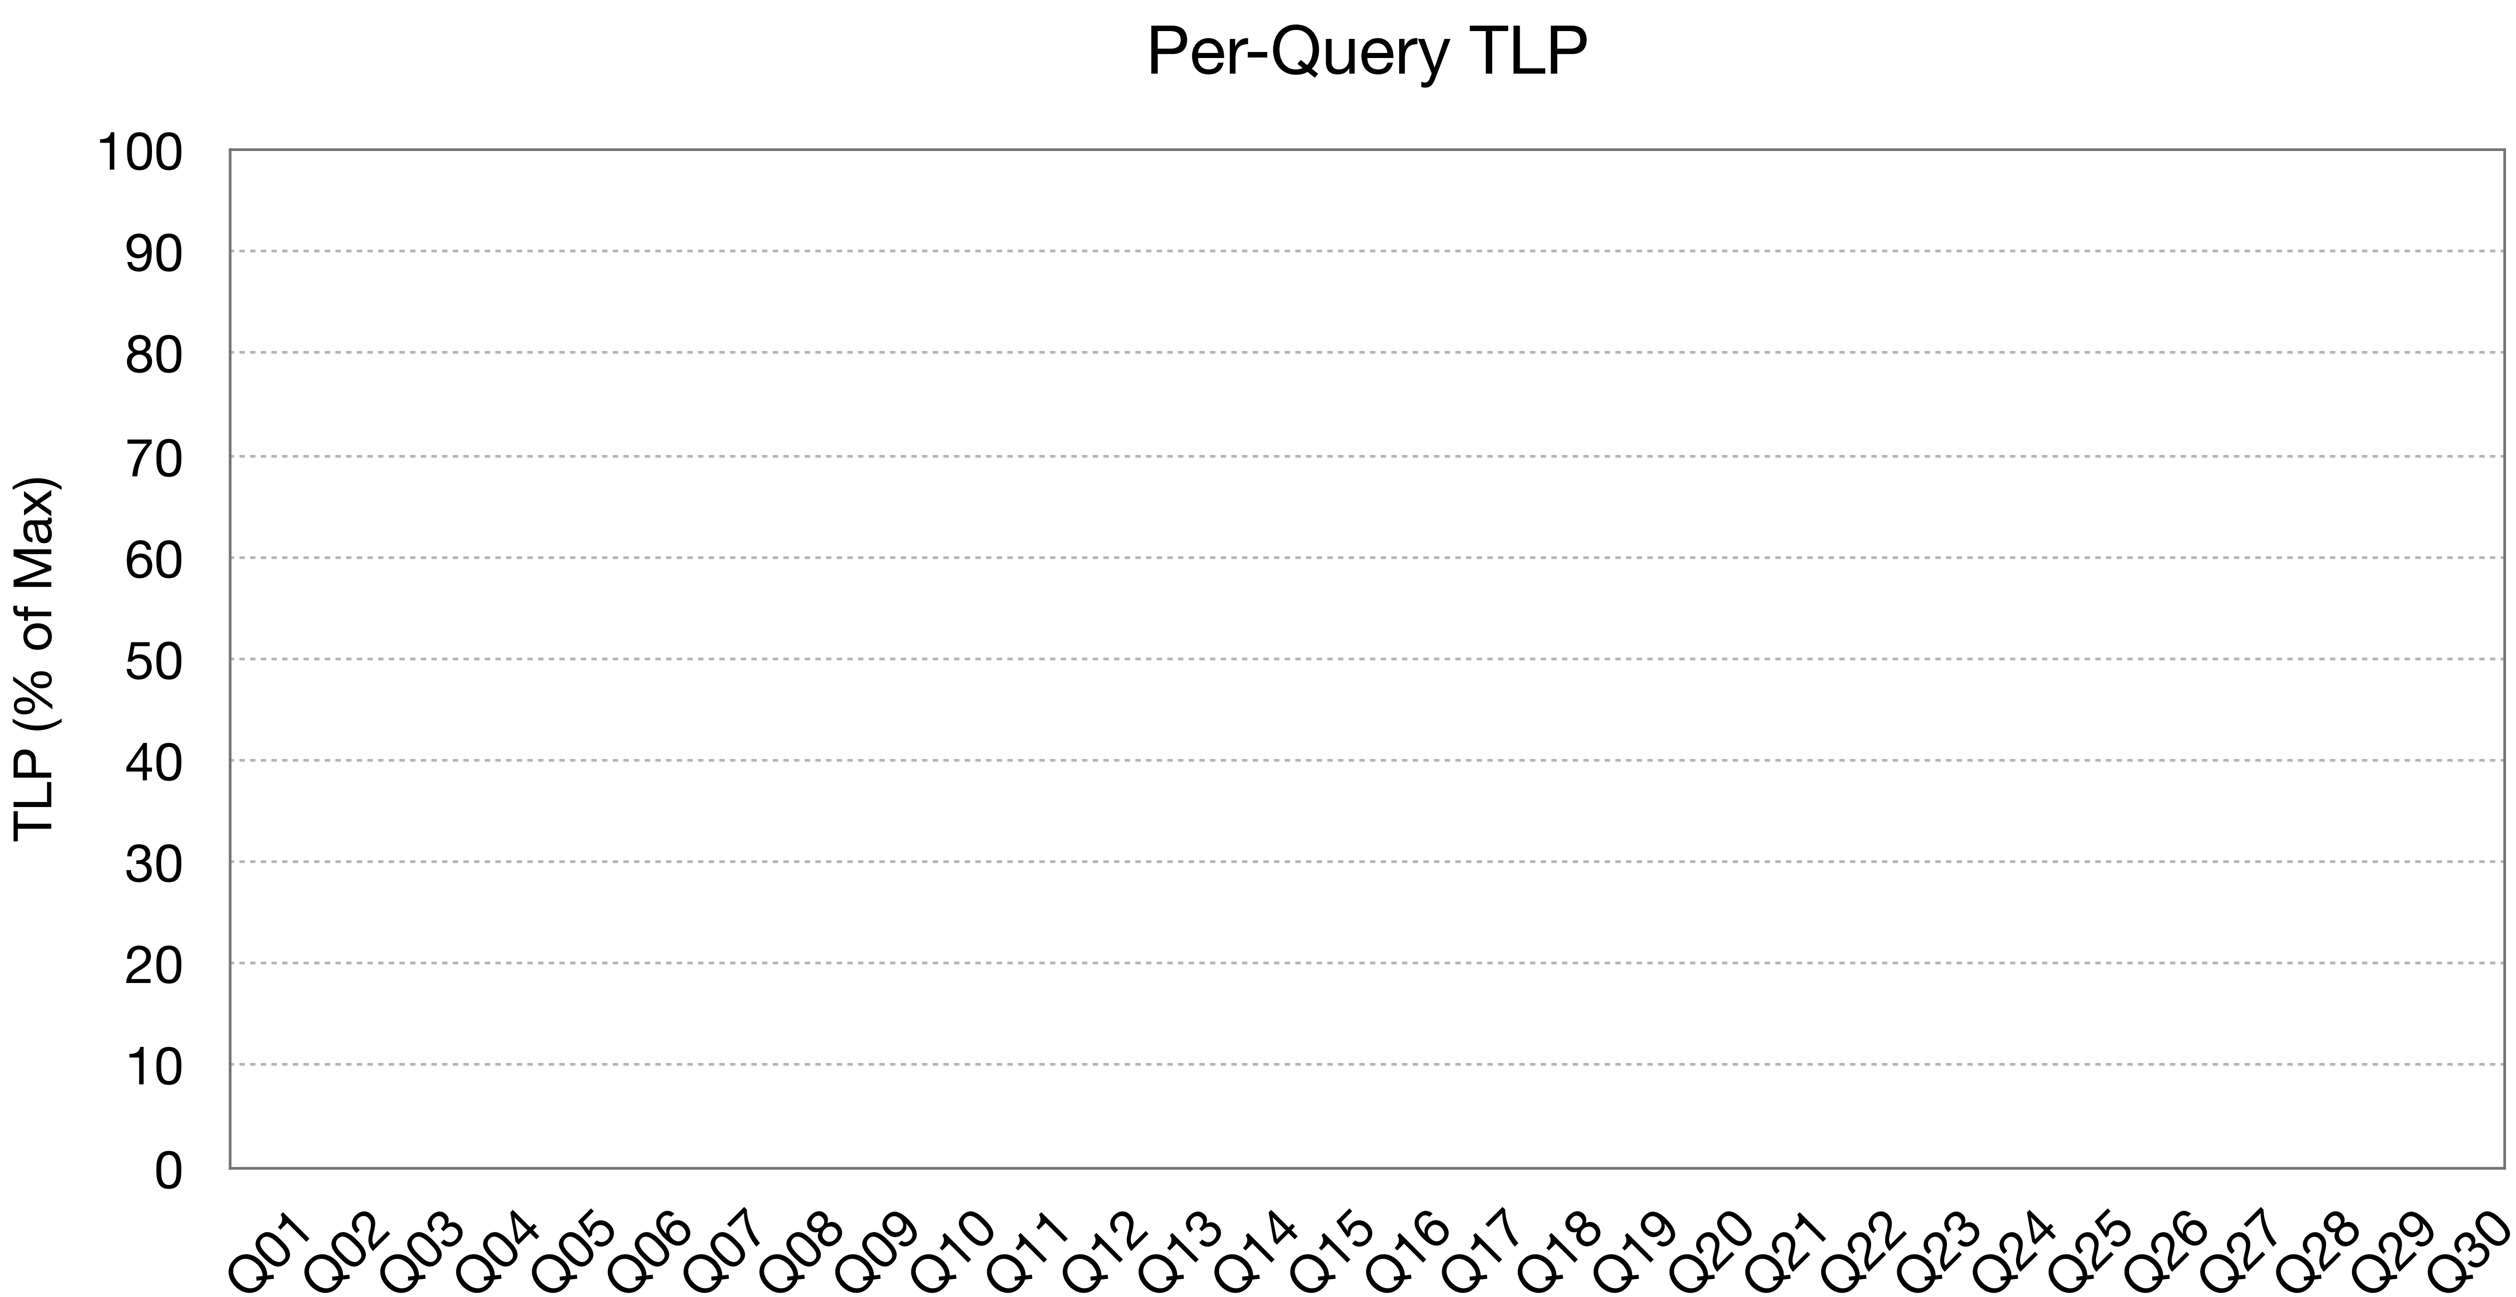

Amdahl’s Law is Alive and Kicking

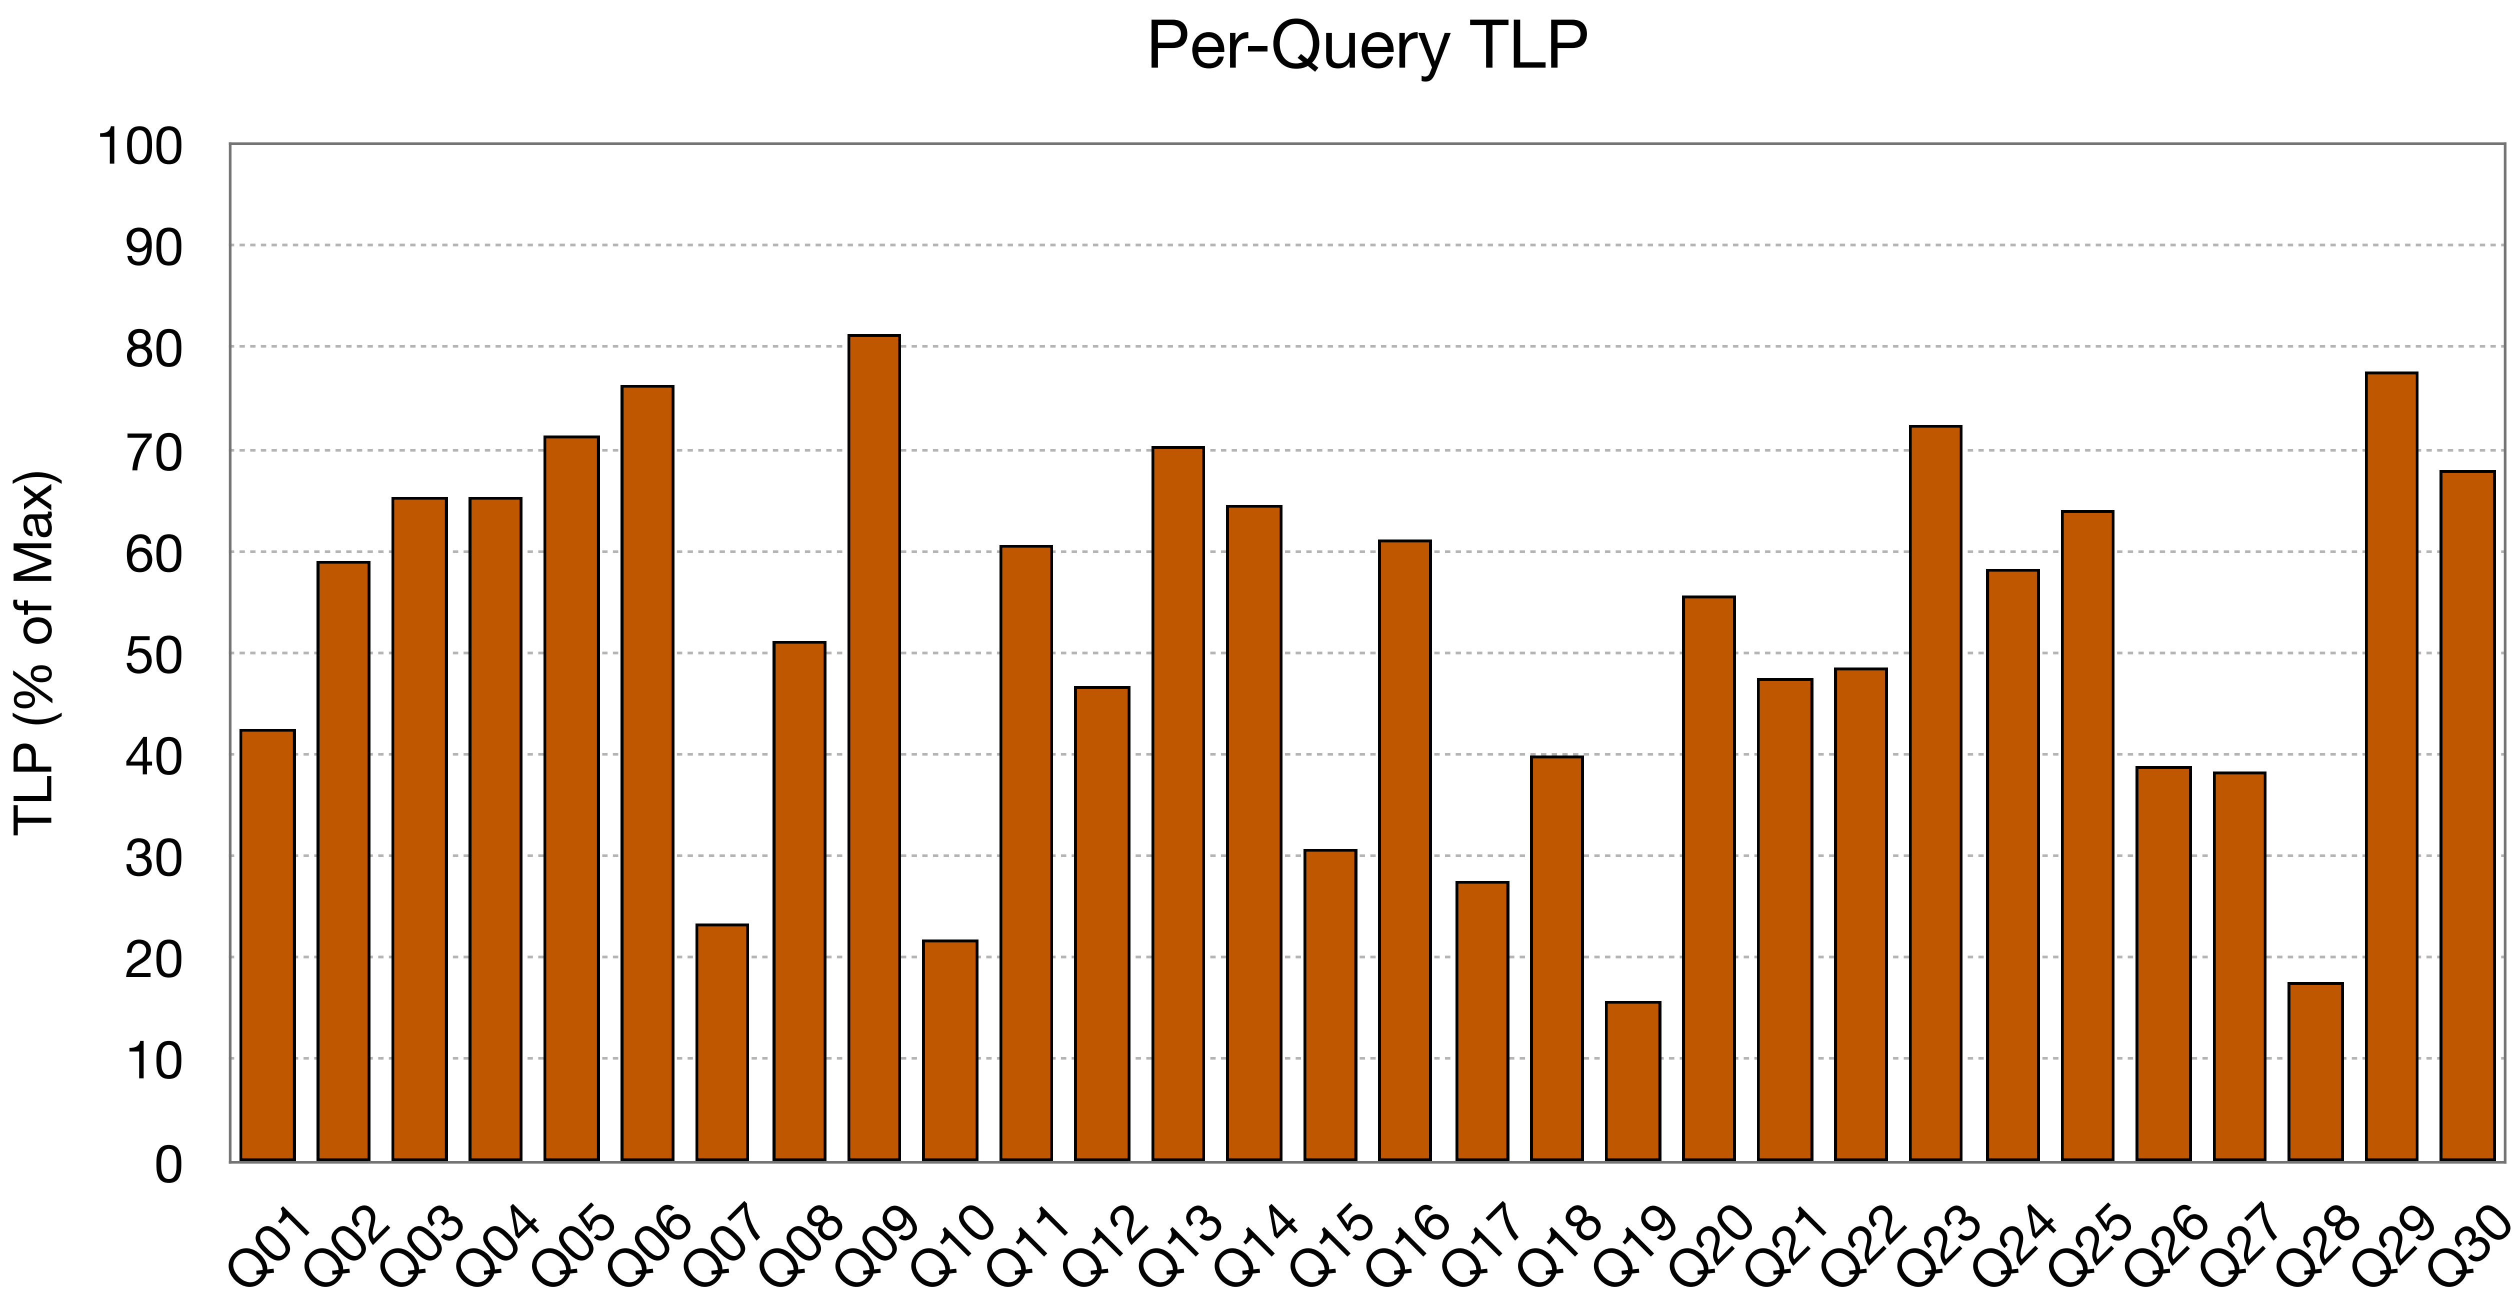

Amdahl's Law is Alive and Kicking

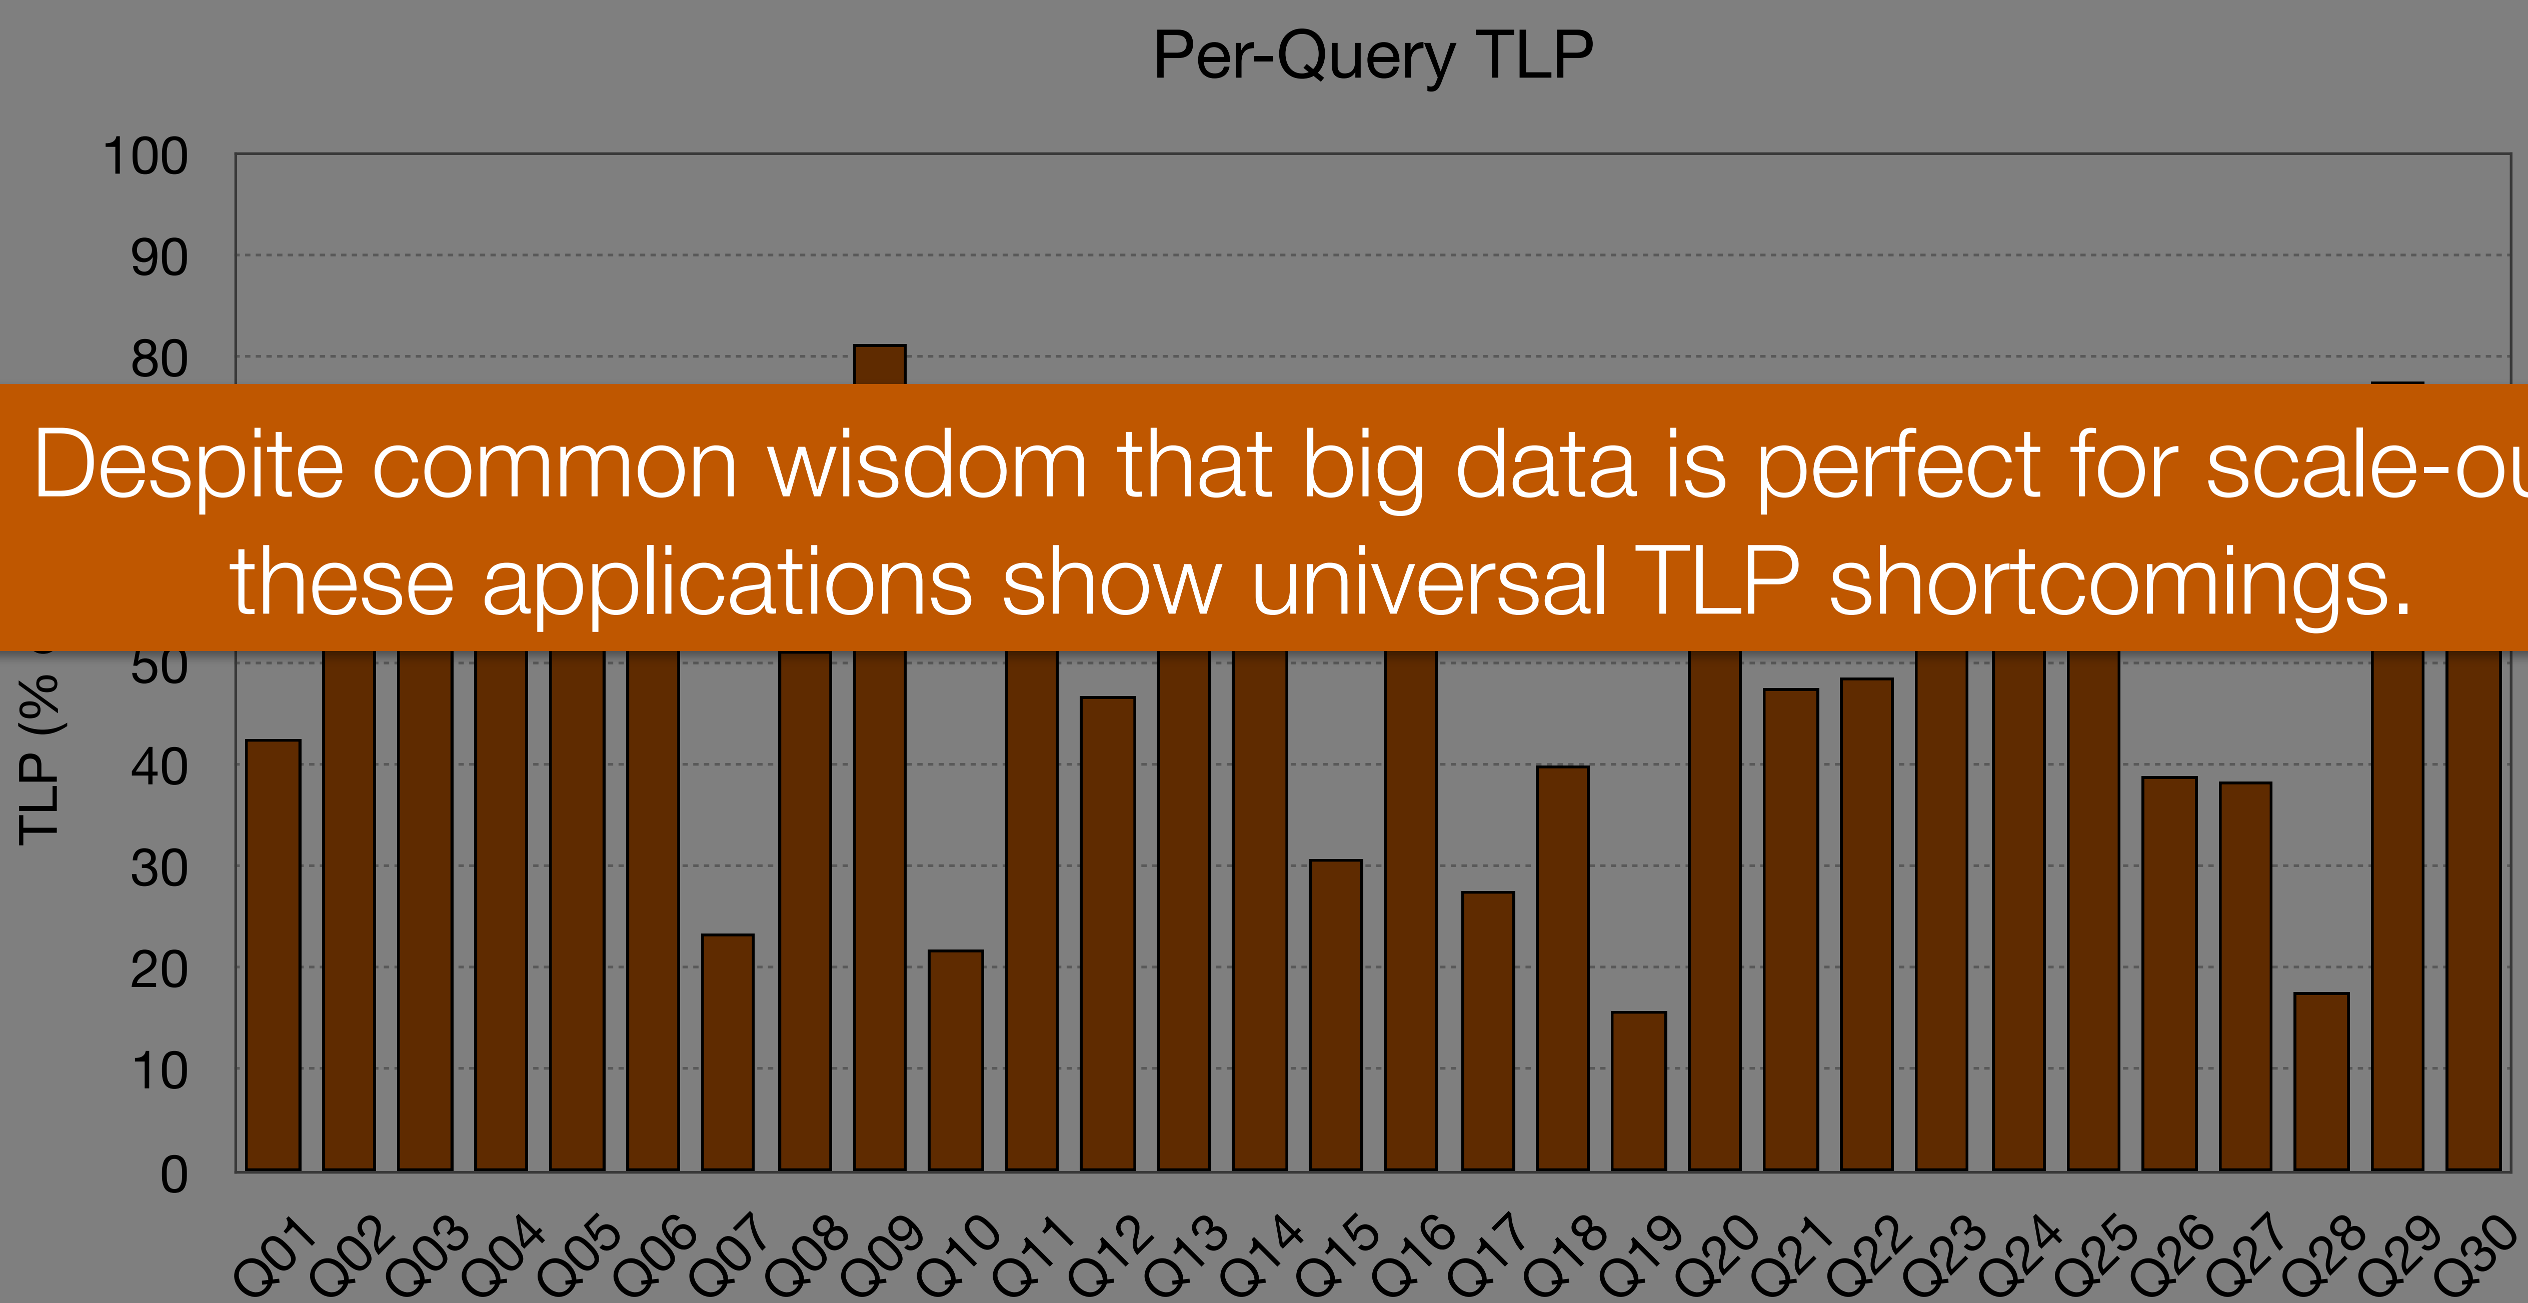

Despite common wisdom that big data is perfect for scale-out, these applications show universal TLP shortcomings.

Contrasting Scale-Out and Scale-Up

Contrasting Scale-Out and Scale-Up

We measure the change in runtime as we scale resources:

Contrasting Scale-Out and Scale-Up

We measure the change in runtime as we scale resources:

- Number of cores

Contrasting Scale-Out and Scale-Up

We measure the change in runtime as we scale resources:

- Number of cores
- Operating frequency

Contrasting Scale-Out and Scale-Up

We measure the change in runtime as we scale resources:

- Number of cores
- Operating frequency

We report the ***efficiency*** of resource scaling.

Contrasting Scale-Out and Scale-Up

We measure the change in runtime as we scale resources:

- Number of cores
- Operating frequency

We report the ***efficiency*** of resource scaling.

$$Efficiency = \frac{RunTime_1 - RunTime_2}{RunTime_1 \cdot (1 - \frac{Resource_1}{Resource_2})}$$

Contrasting Scale-Out and Scale-Up

We measure the change in runtime as we scale resources:

- Number of cores
- Operating frequency

We report the ***efficiency*** of resource scaling.

$$Efficiency = \frac{\text{Change in run time}}{RunTime_1 \cdot (1 - \frac{Resource_1}{Resource_2})}$$

Contrasting Scale-Out and Scale-Up

We measure the change in runtime as we scale resources:

- Number of cores
- Operating frequency

We report the ***efficiency*** of resource scaling.

$$Efficiency = \frac{\text{Change in run time}}{\text{Expected change}}$$

Scaling Efficiency

Scaling Efficiency

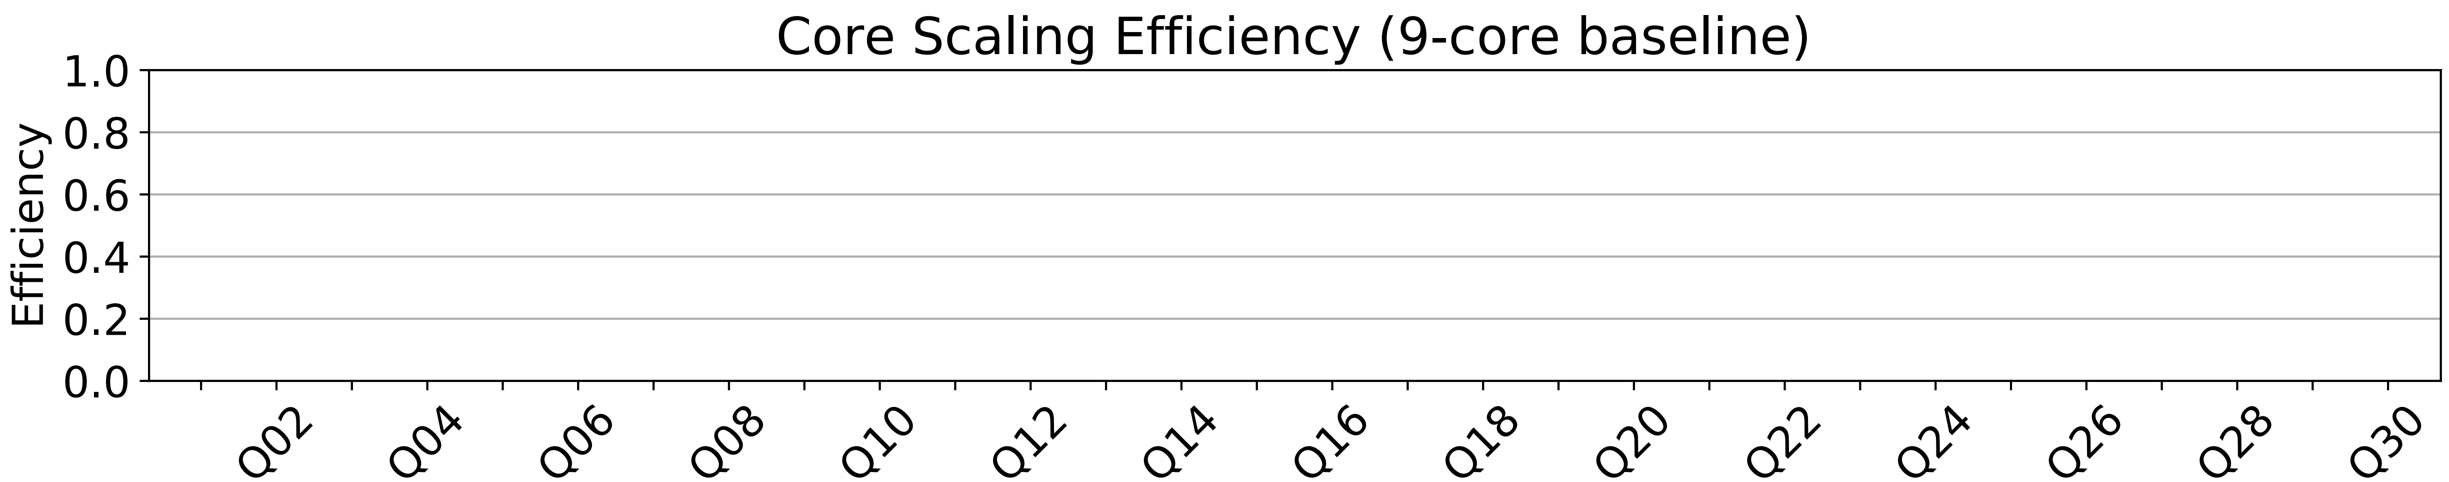

Scaling Efficiency

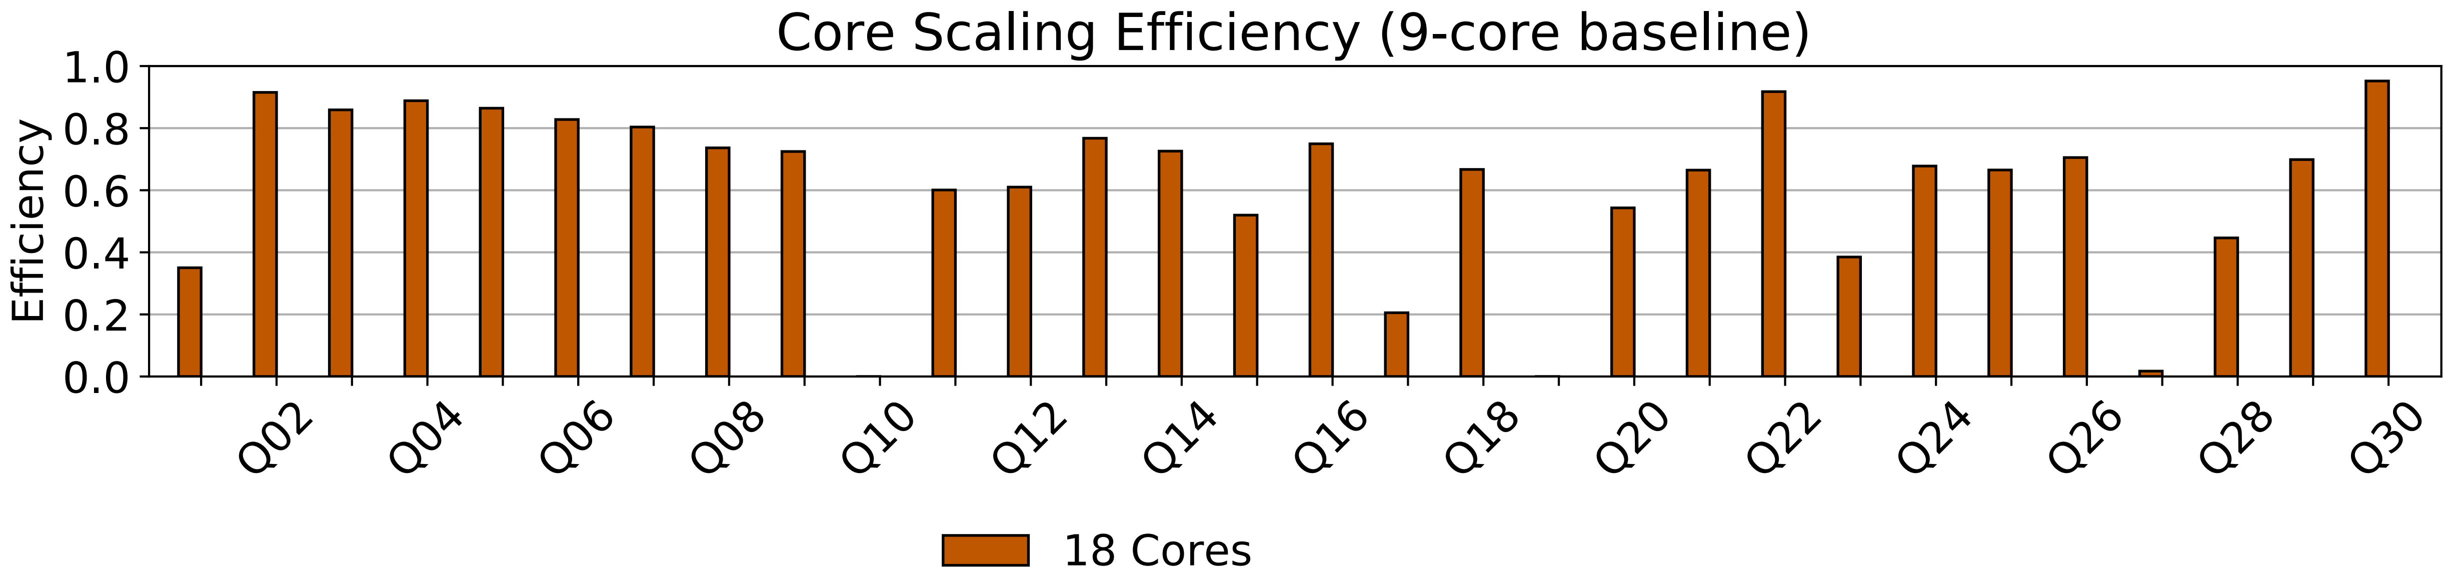

Scaling Efficiency

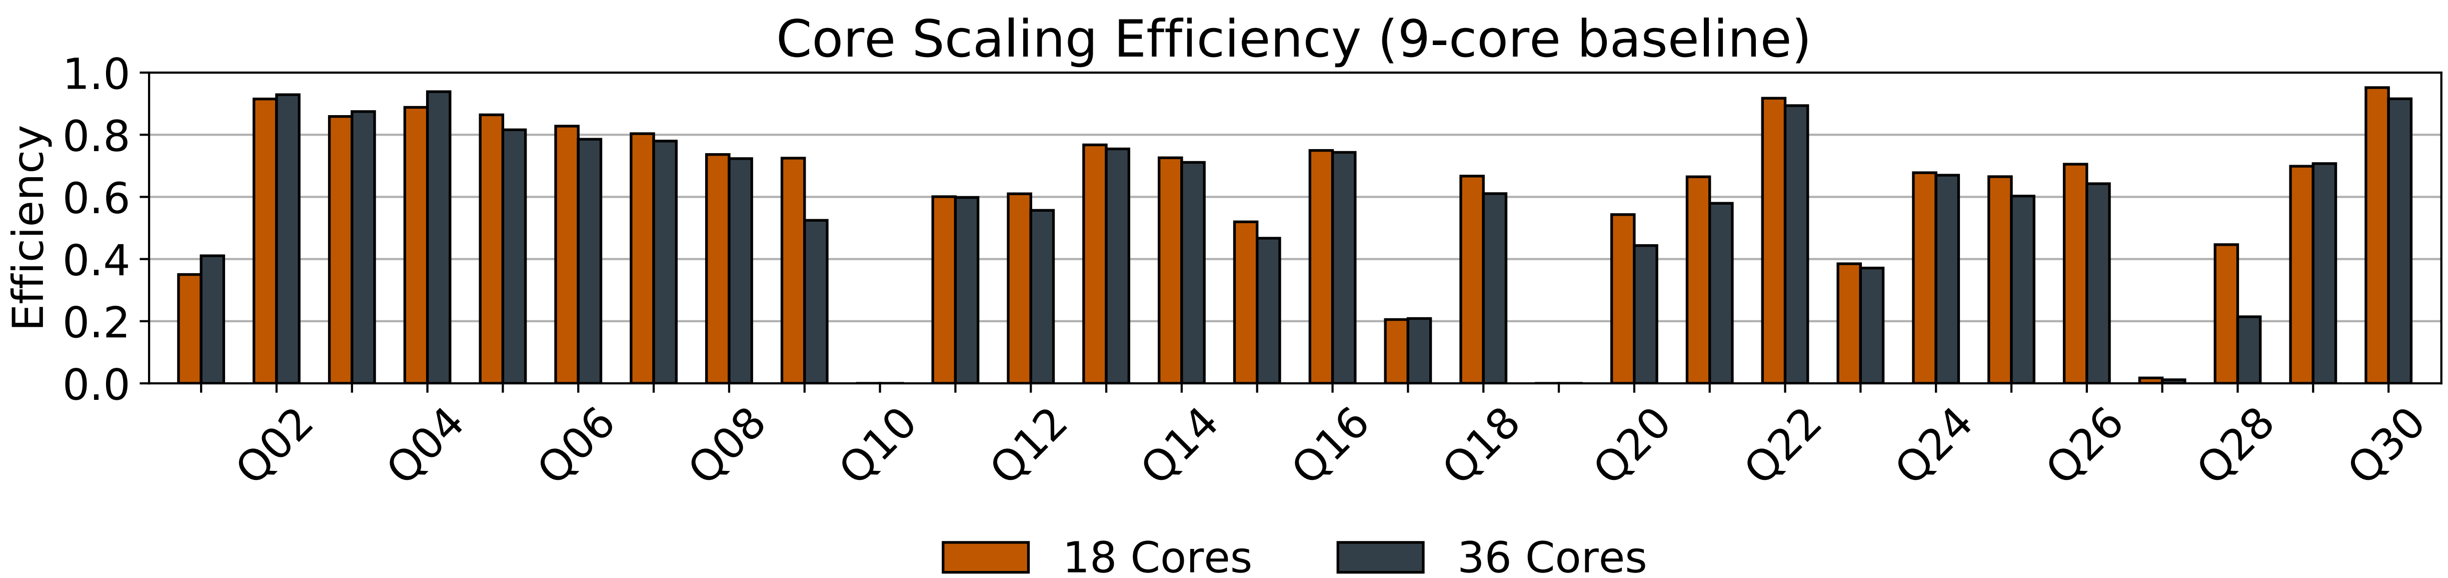

Core scaling shows limited efficiency, sometimes providing no benefit at all.

Scaling Efficiency

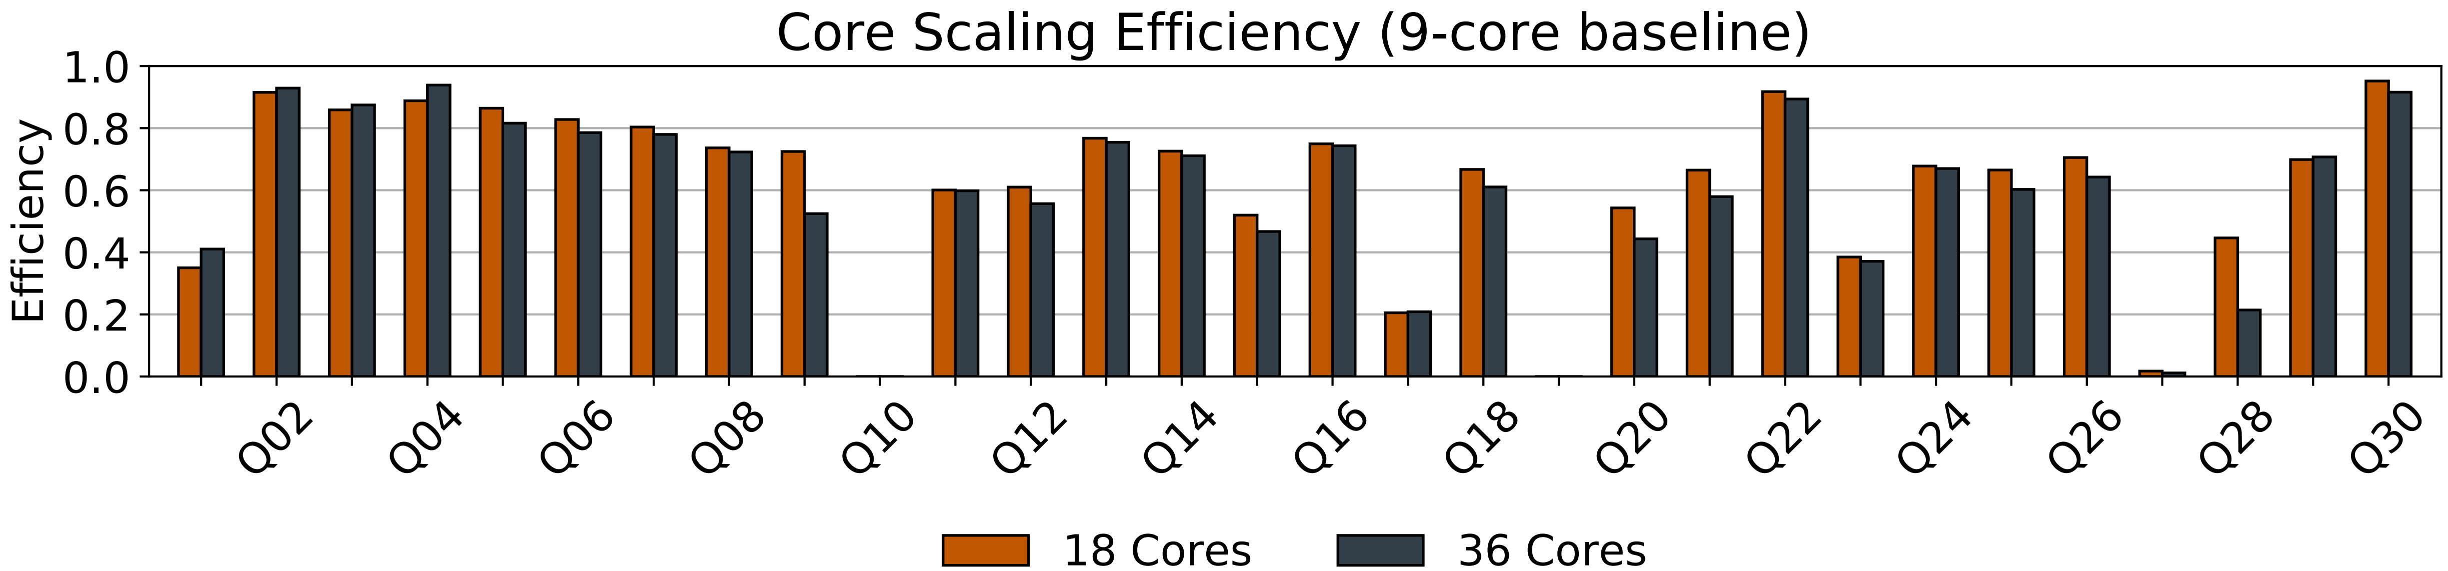

Core scaling shows limited efficiency, sometimes providing no benefit at all.

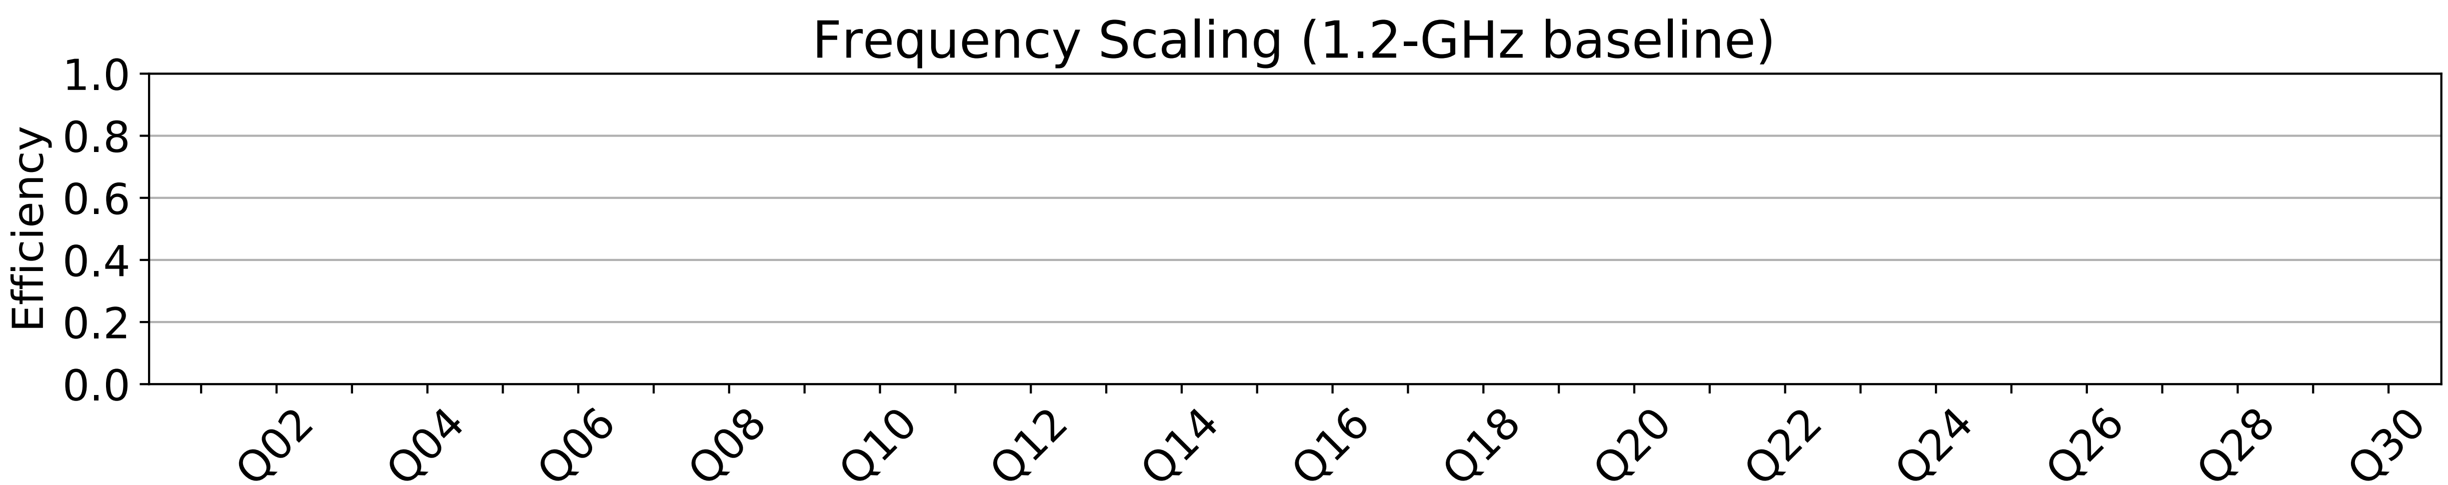

Scaling Efficiency

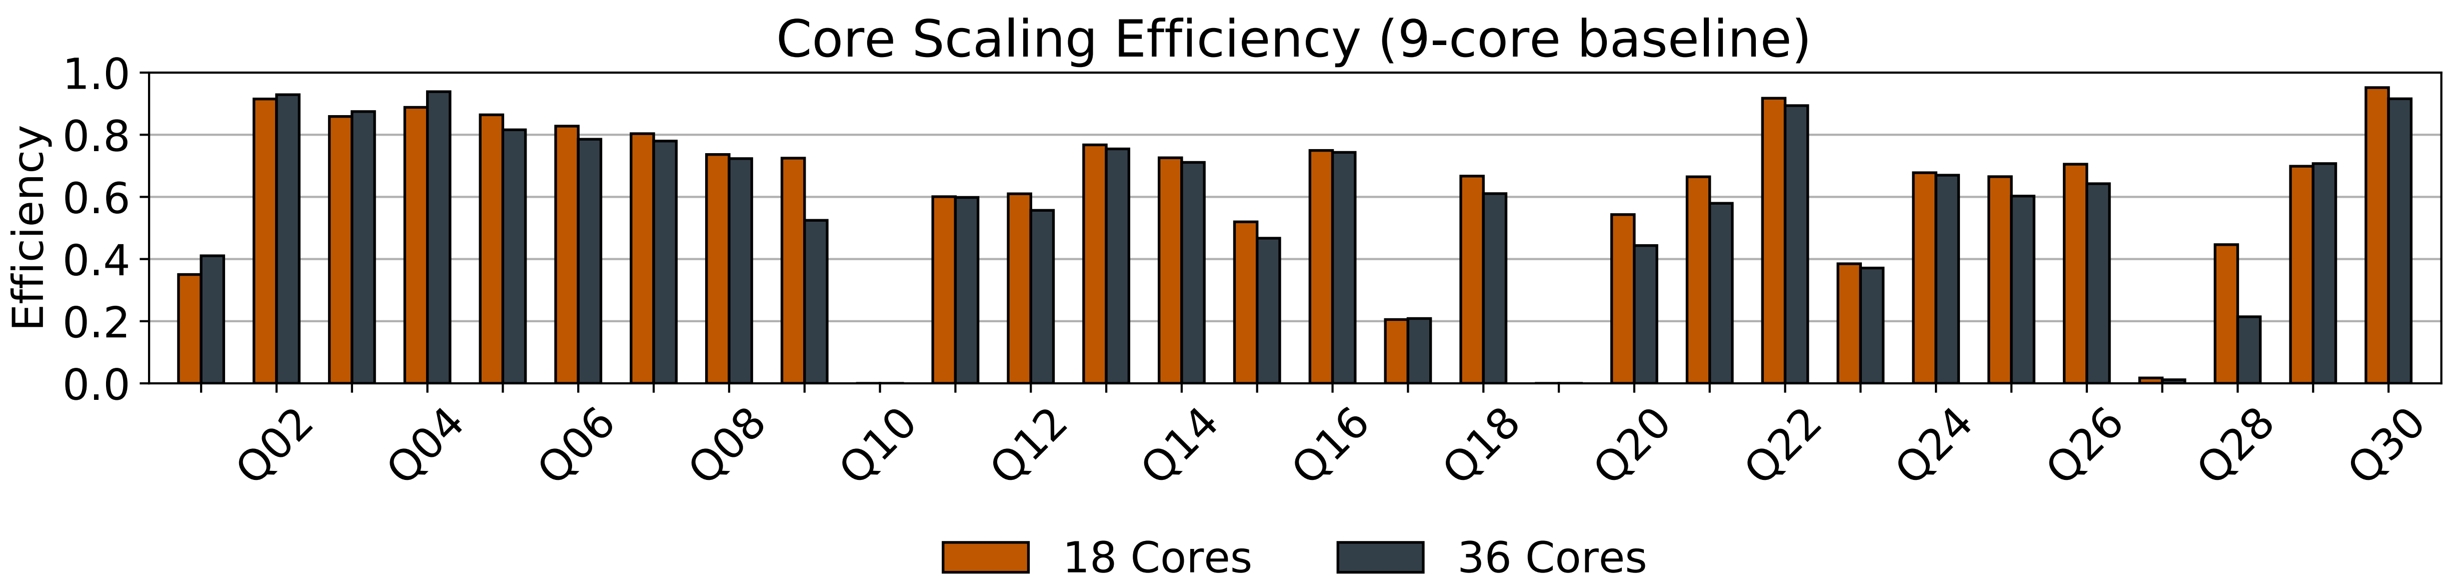

Core scaling shows limited efficiency, sometimes providing no benefit at all.

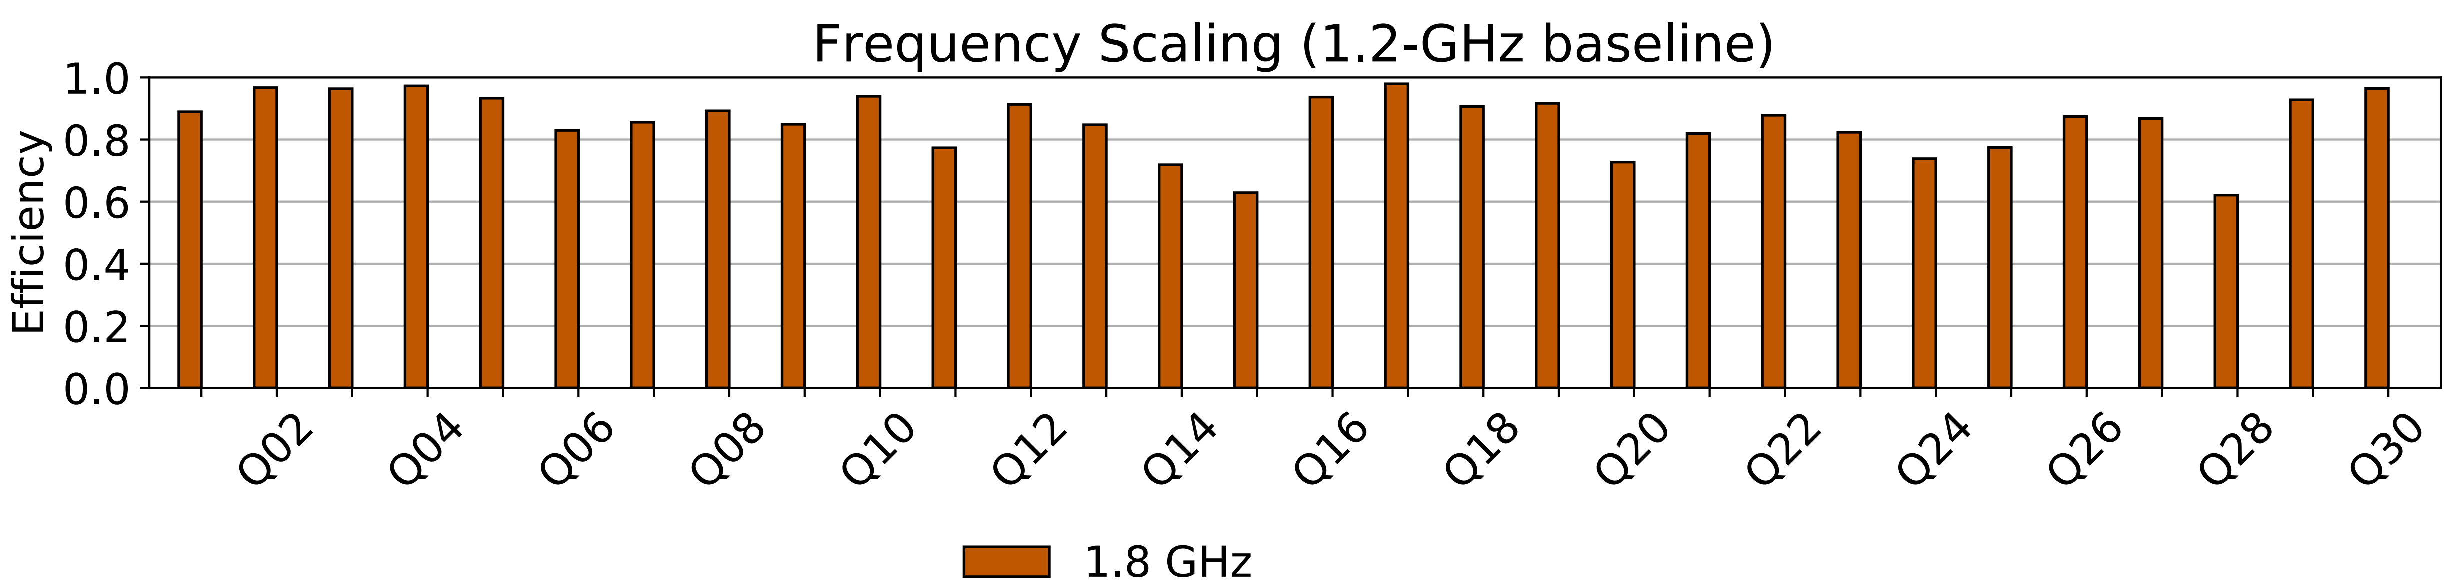

Scaling Efficiency

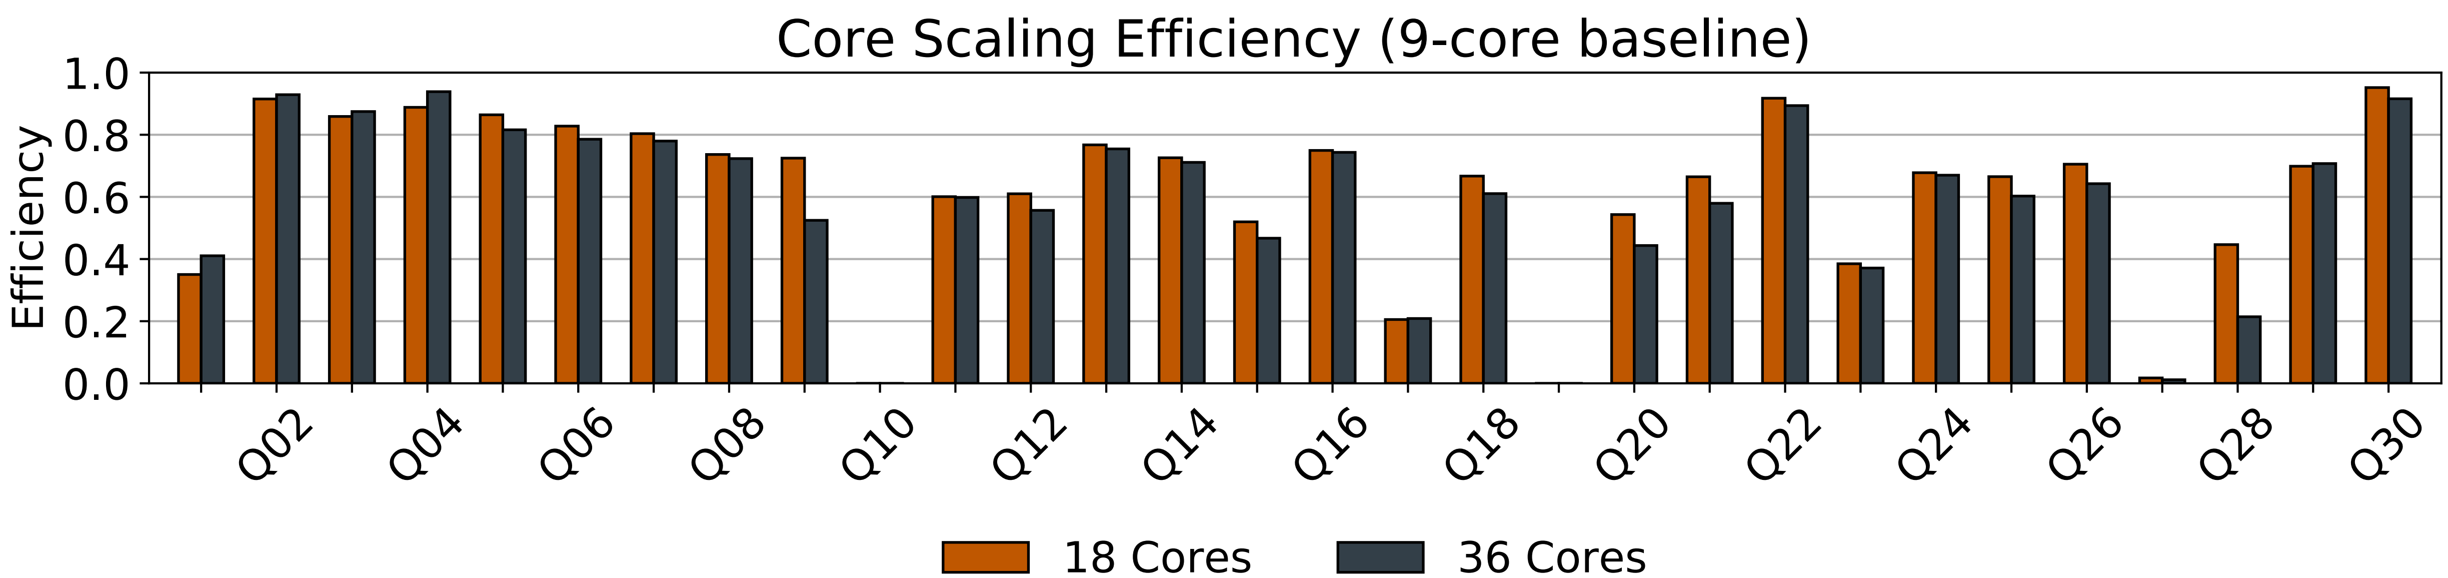

Core scaling shows limited efficiency, sometimes providing no benefit at all.

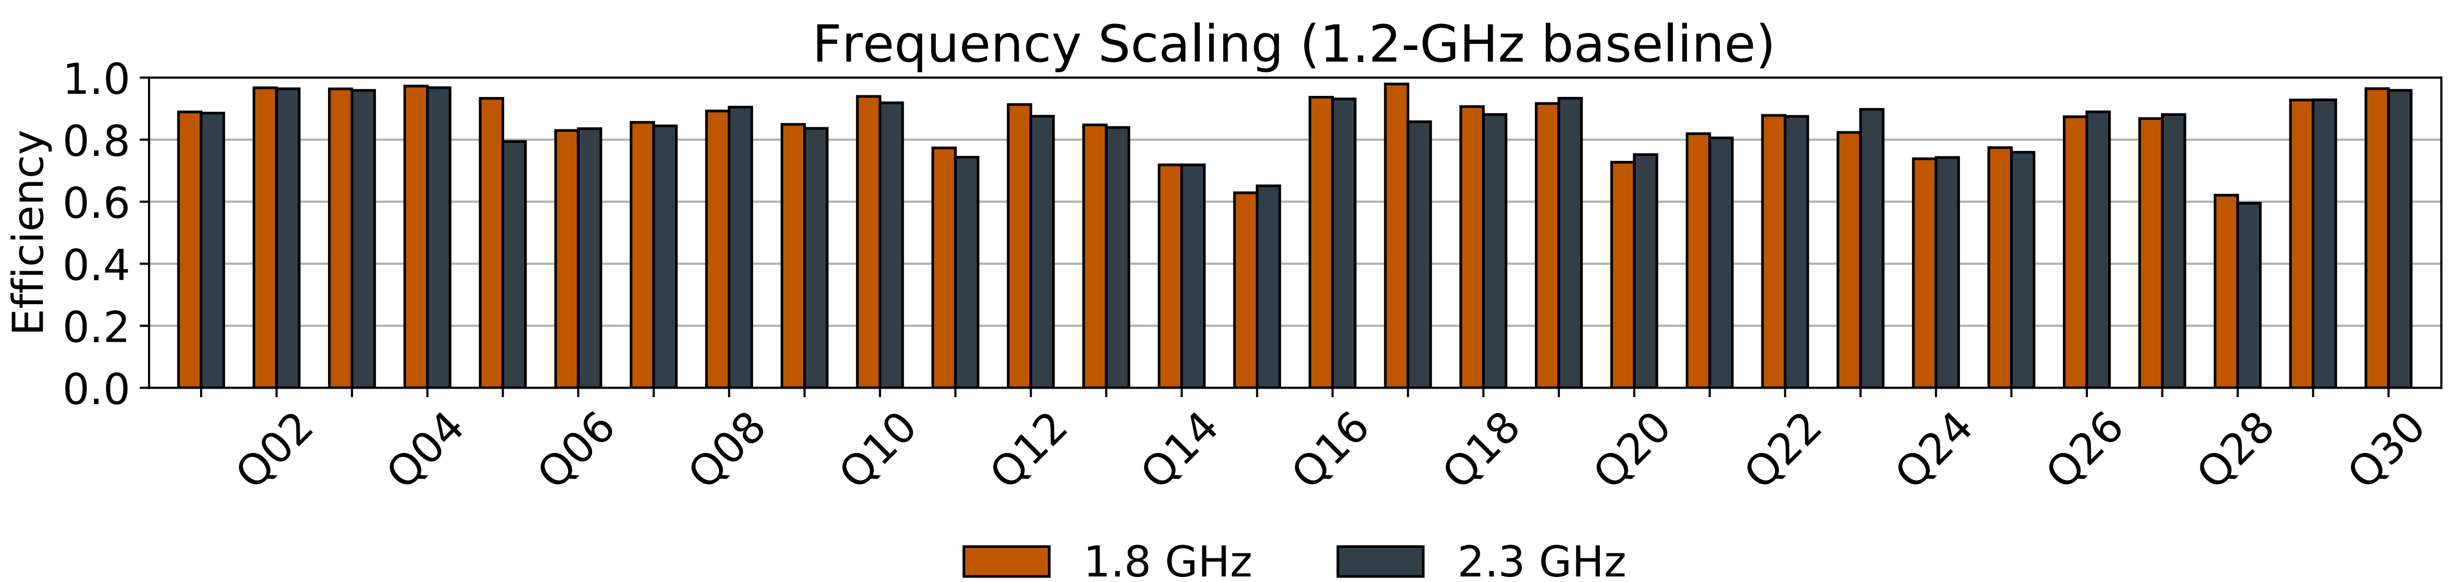

Scaling Efficiency

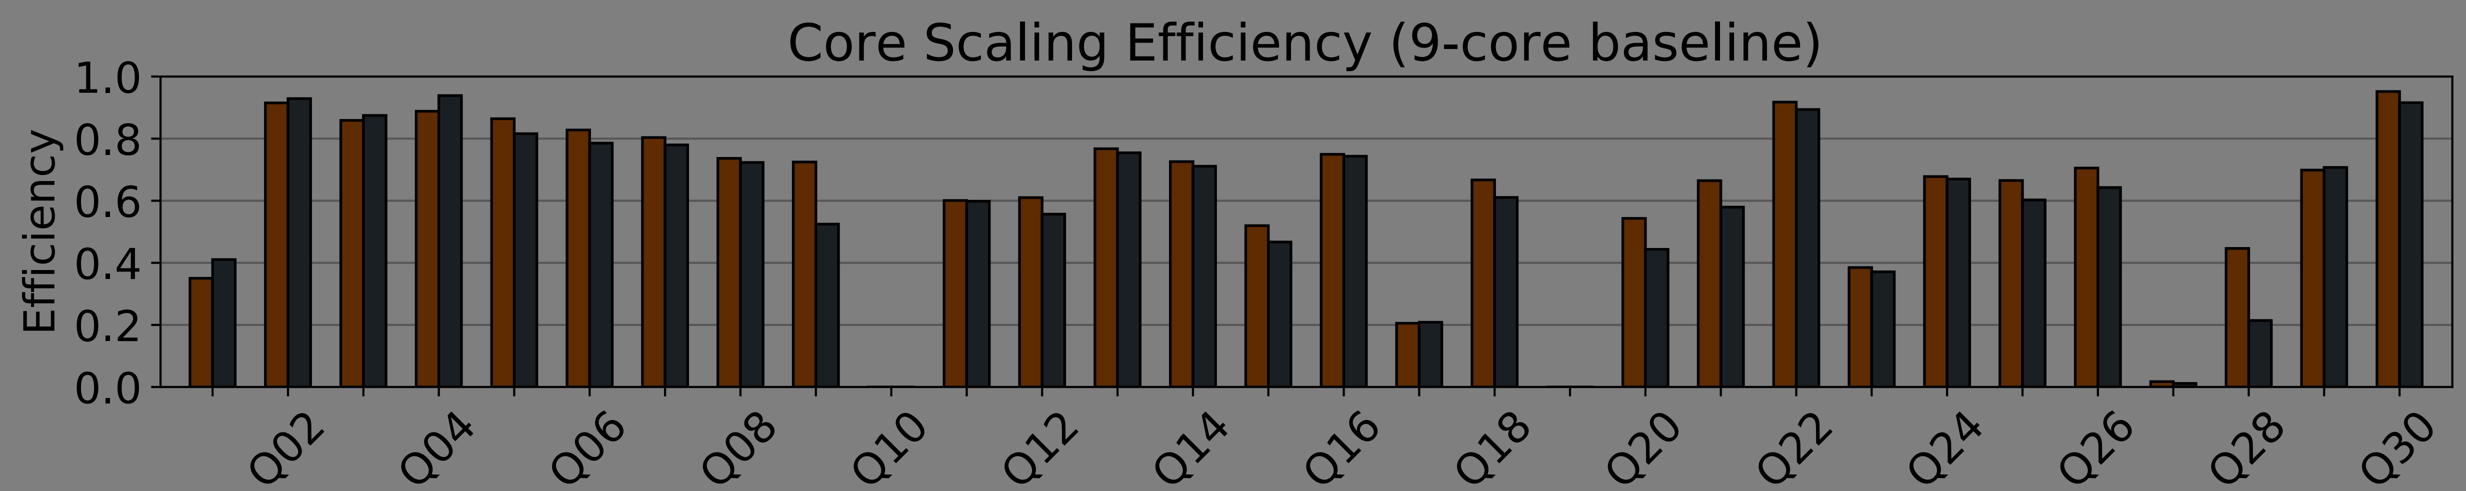

Frequency scaling is more efficient than core scaling.

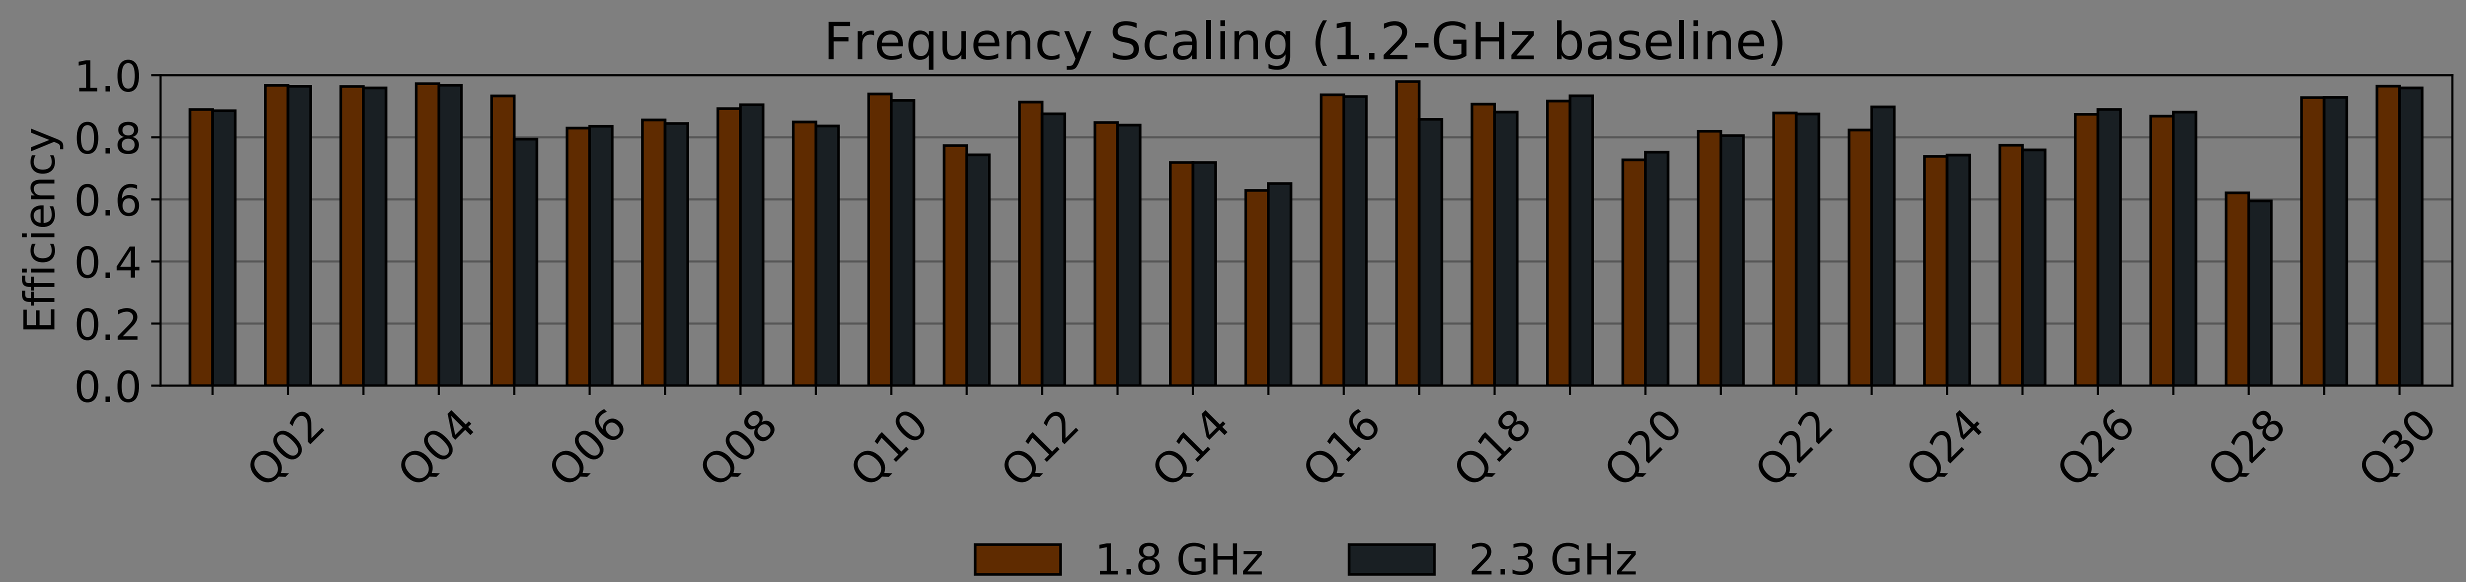

What About Turbo Boost?

What About Turbo Boost?

Turbo Boost uses slack in thermal and current margins to increase the operating frequency of CPUs.

What About Turbo Boost?

Turbo Boost uses slack in thermal and current margins to increase the operating frequency of CPUs.

The Turbo Boost ceiling is a function of the number of active cores. The fewer the cores, the higher the ceiling.

What About Turbo Boost?

Turbo Boost uses slack in thermal and current margins to increase the operating frequency of CPUs.

The Turbo Boost ceiling is a function of the number of active cores. The fewer the cores, the higher the ceiling.

With so many halted cycles, Turbo Boost should be perfect for BigBench, right?

What About Turbo Boost?

Turbo Boost uses slack in thermal and current margins to increase the operating frequency of CPUs.

The Turbo Boost ceiling is a function of the number of active cores. The fewer the cores, the higher the ceiling.

With so many halted cycles, Turbo Boost should be perfect for BigBench, right?

Not quite…

A Deeper Problem

A Deeper Problem

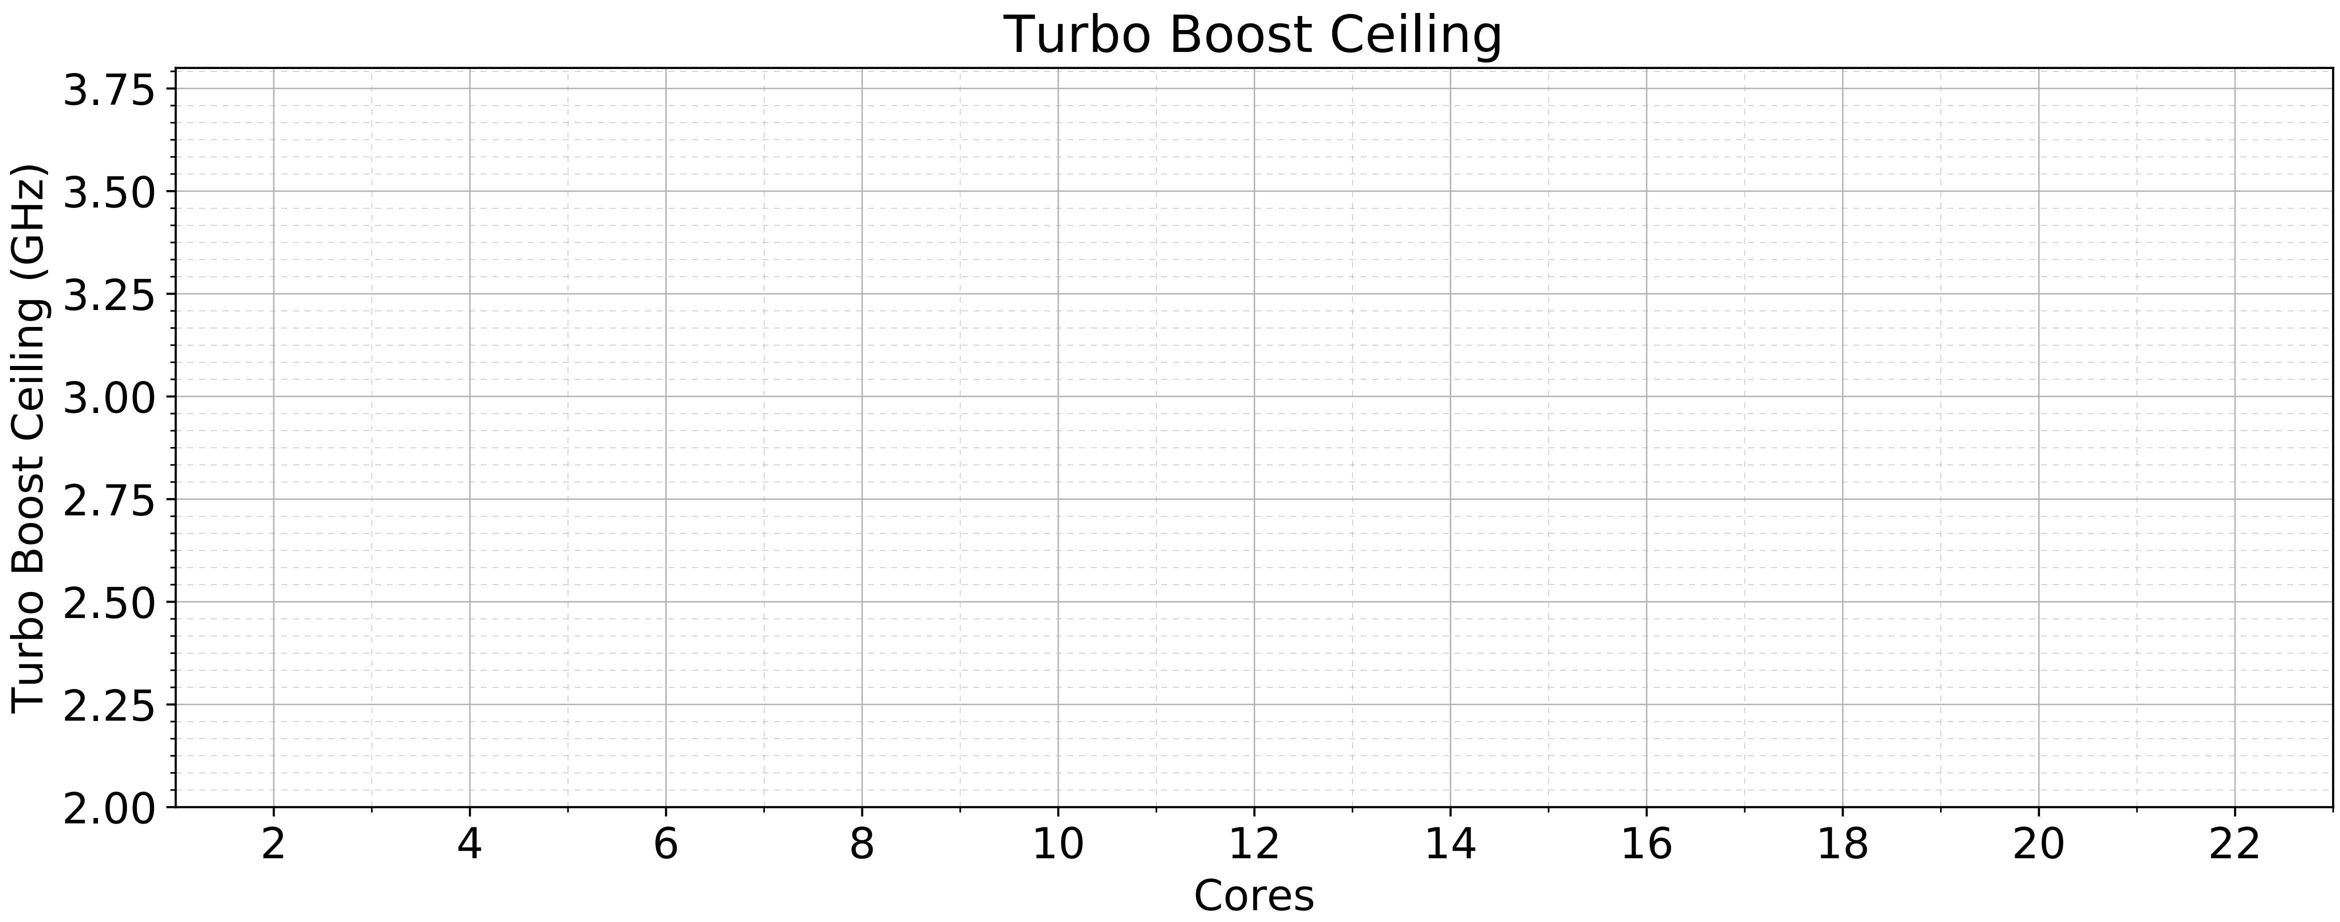

A Deeper Problem

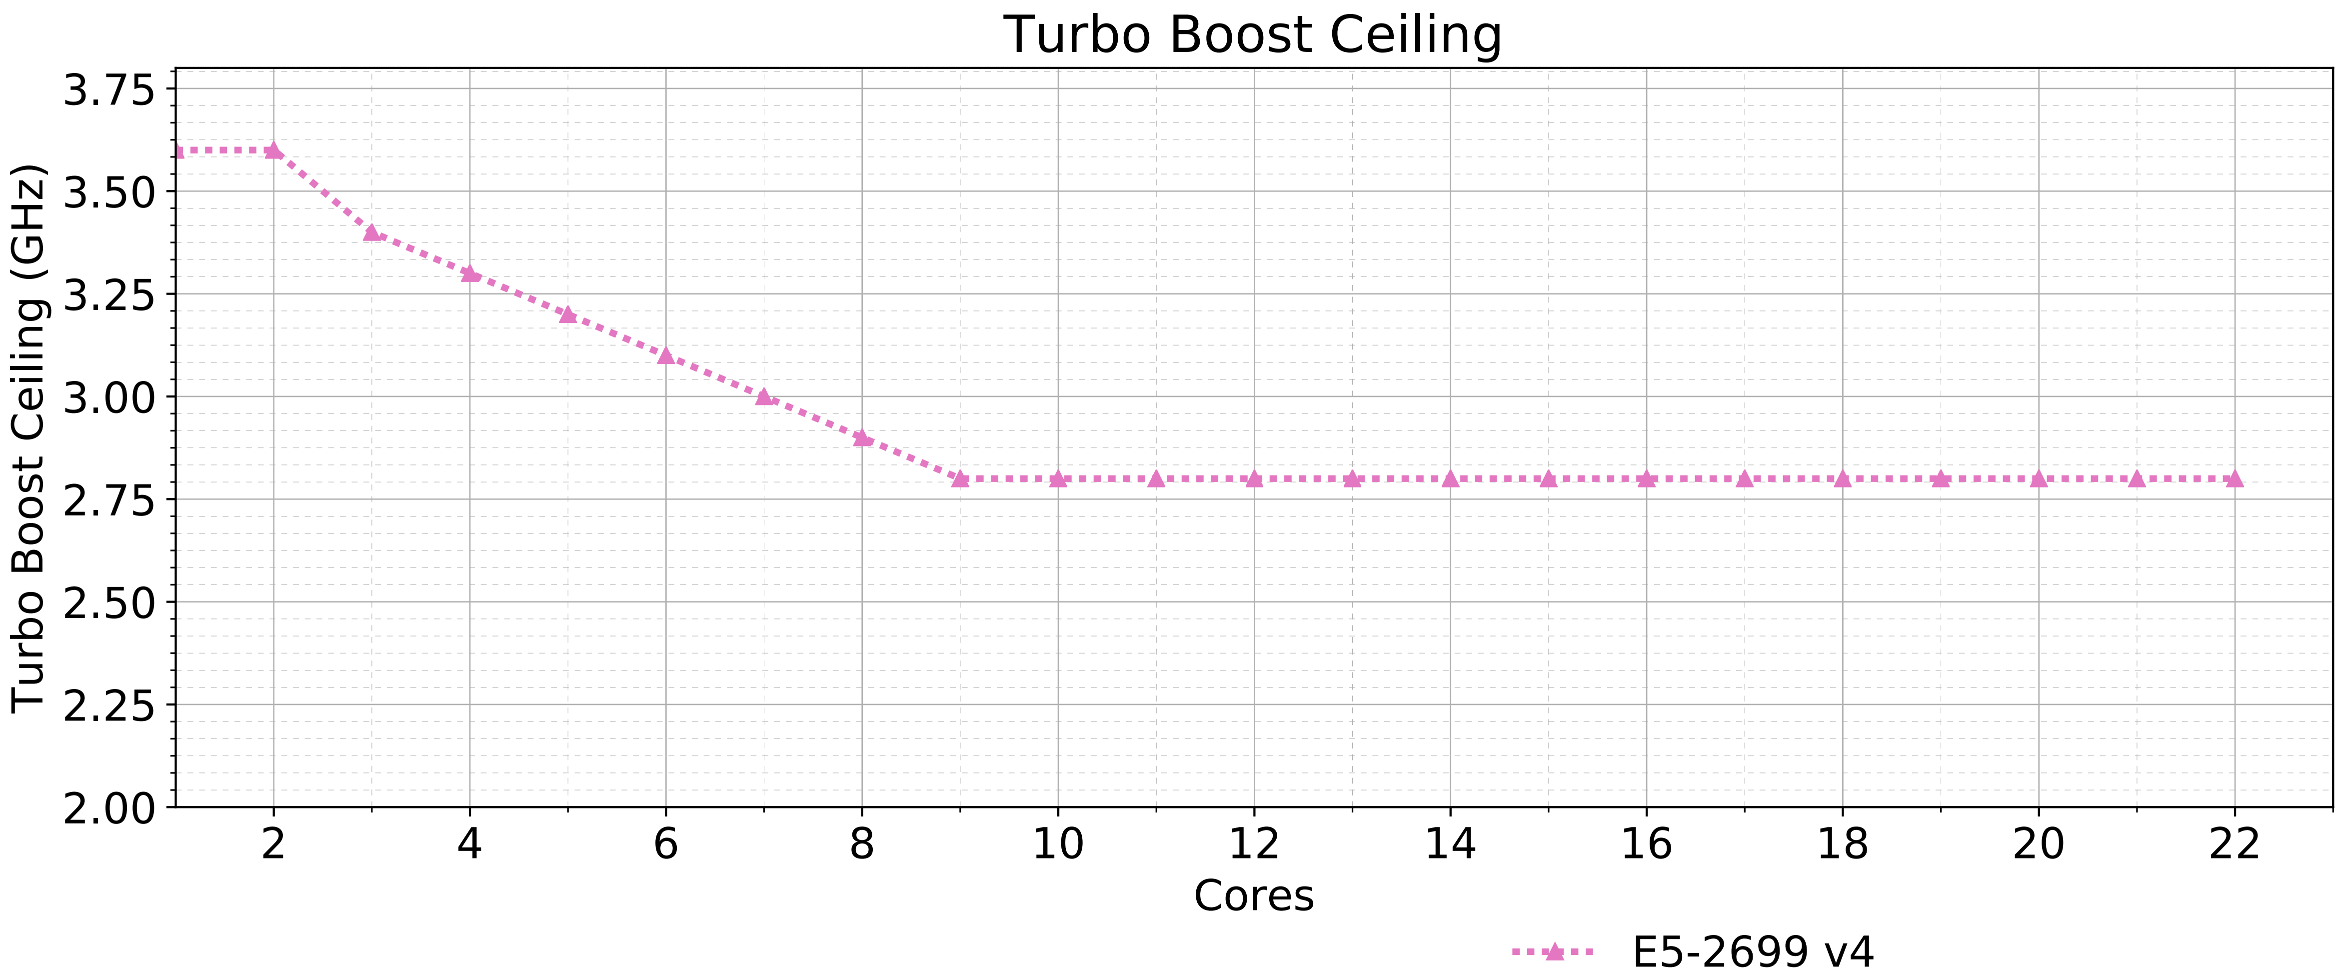

A Deeper Problem

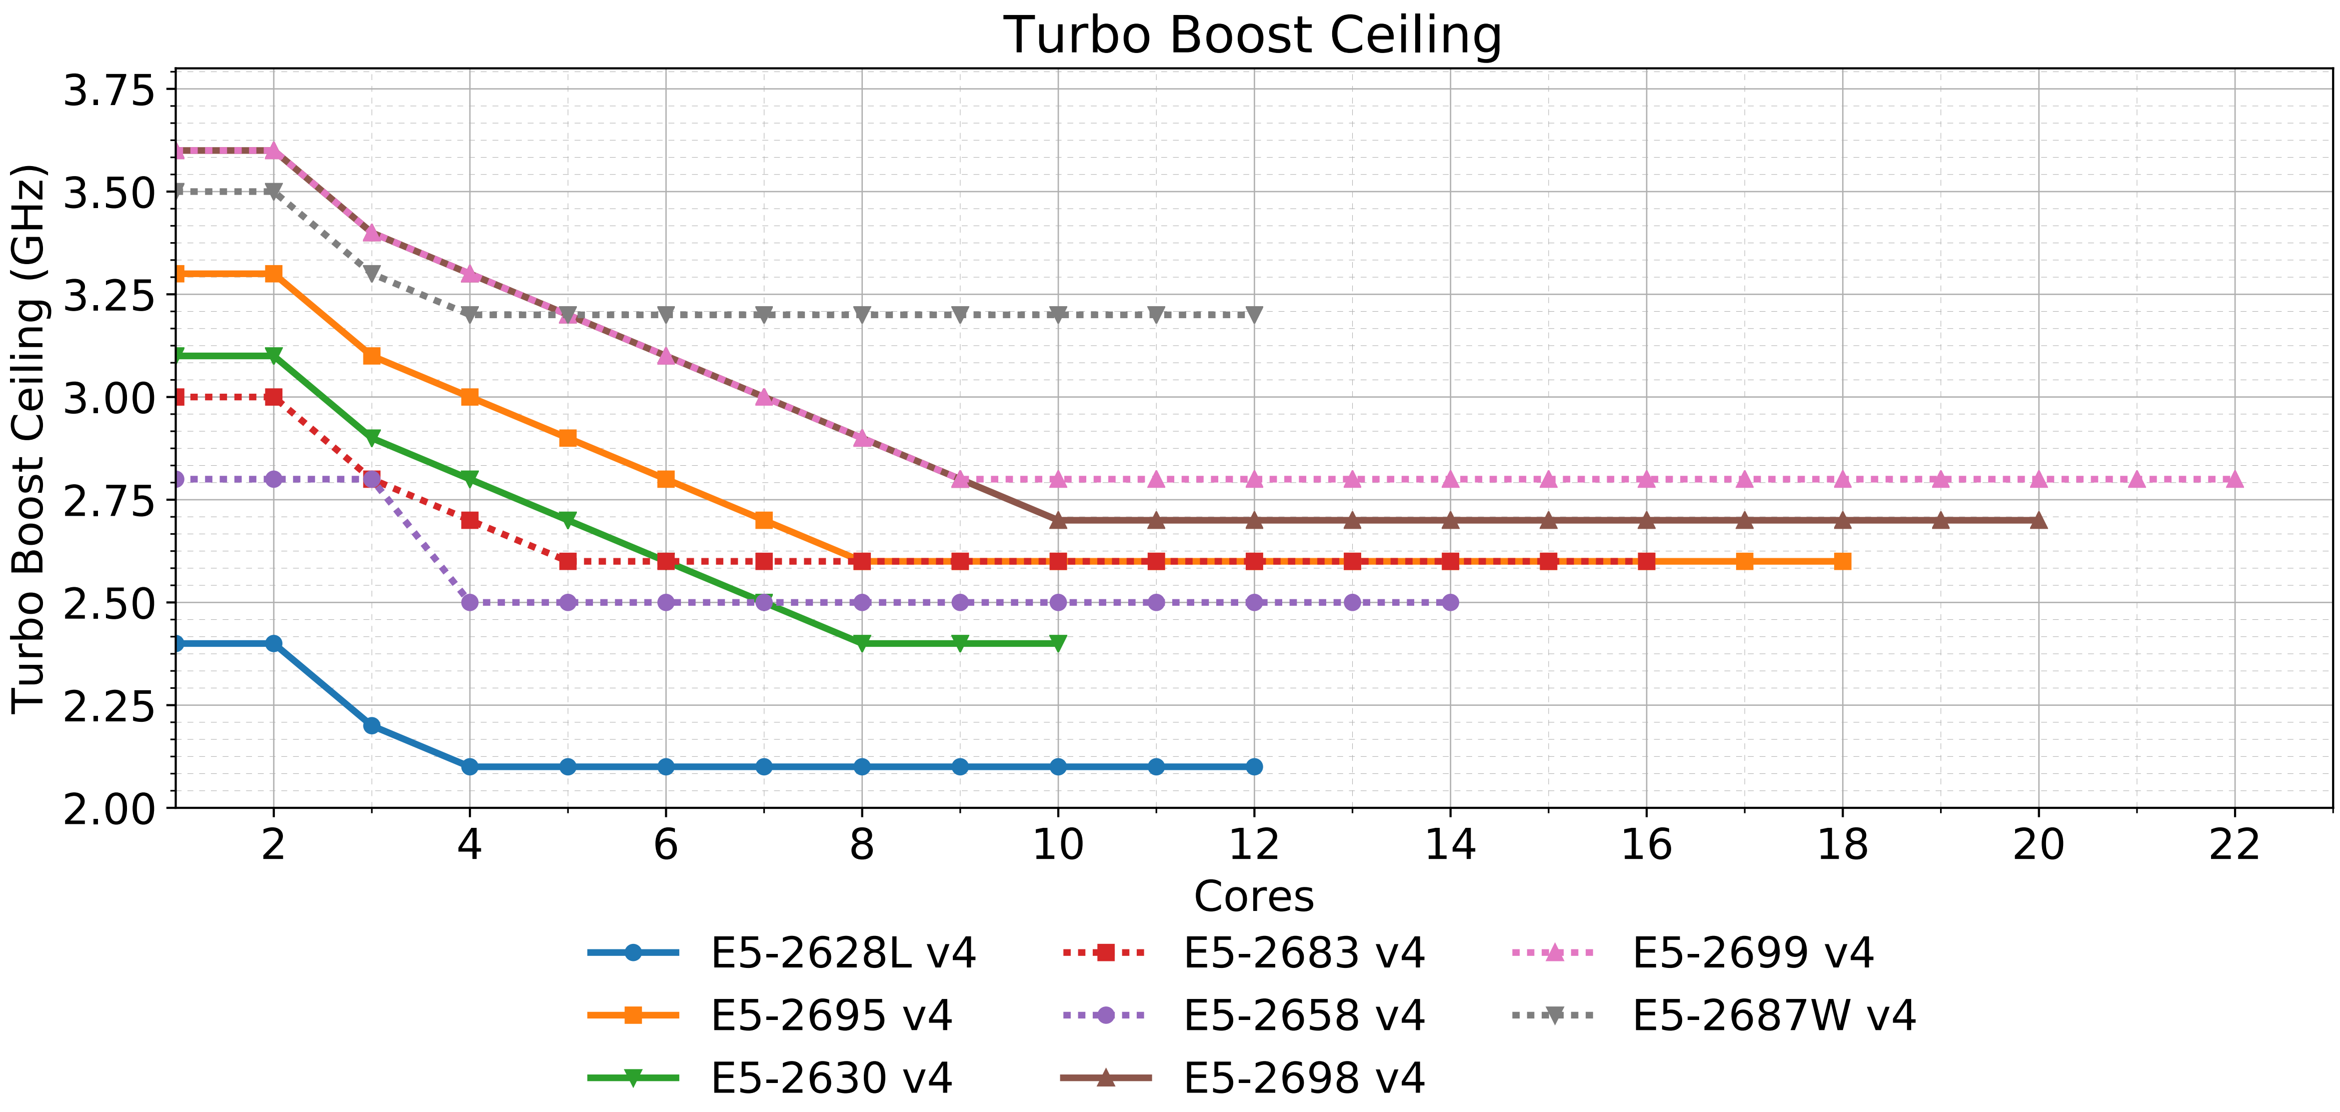

A Deeper Problem

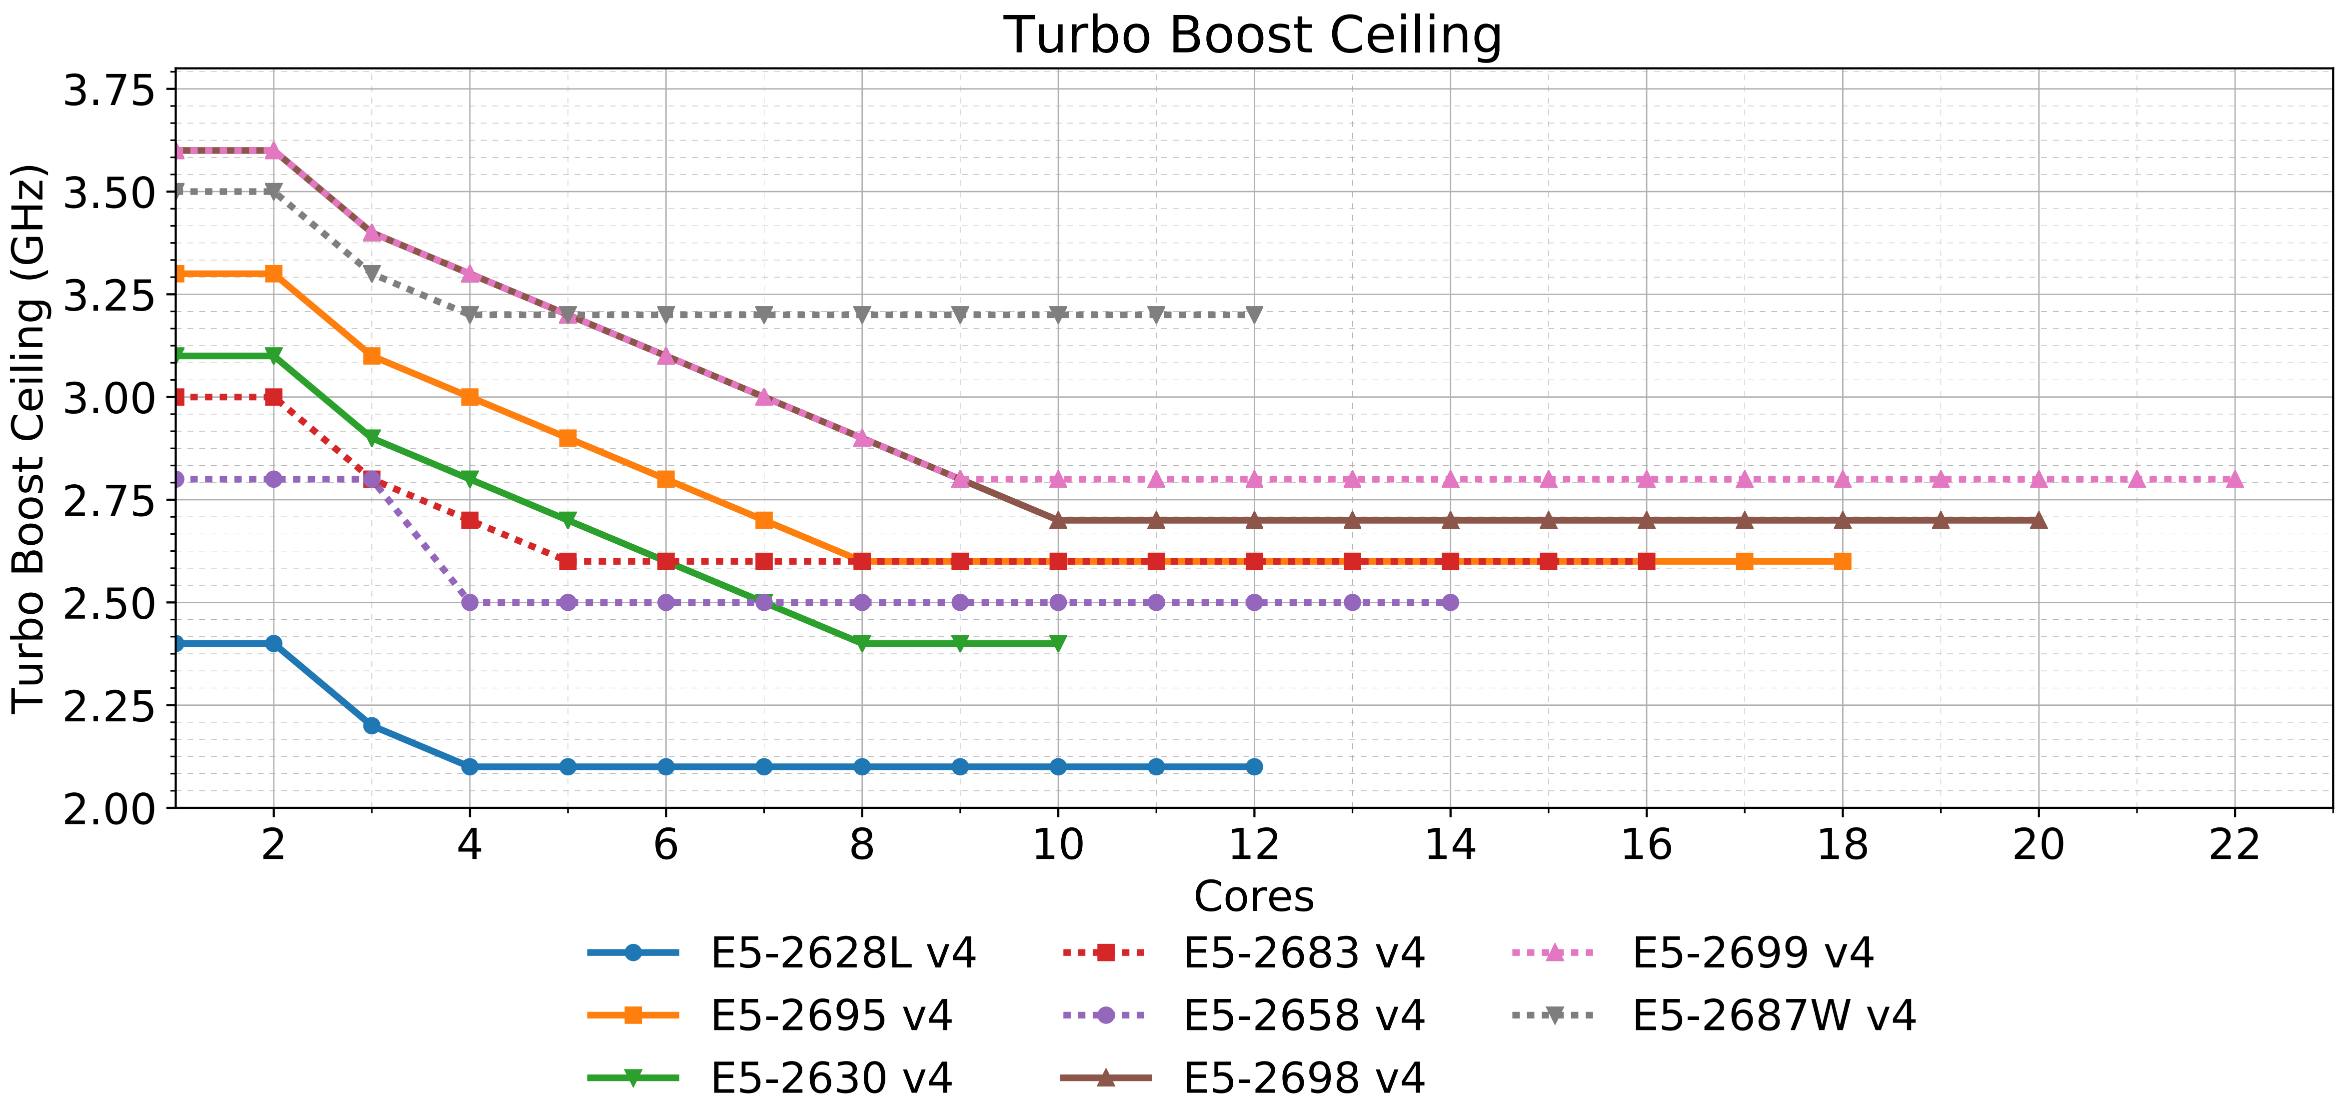

The Turbo Boost ceiling does not increase until most cores have been disabled.

A Deeper Problem

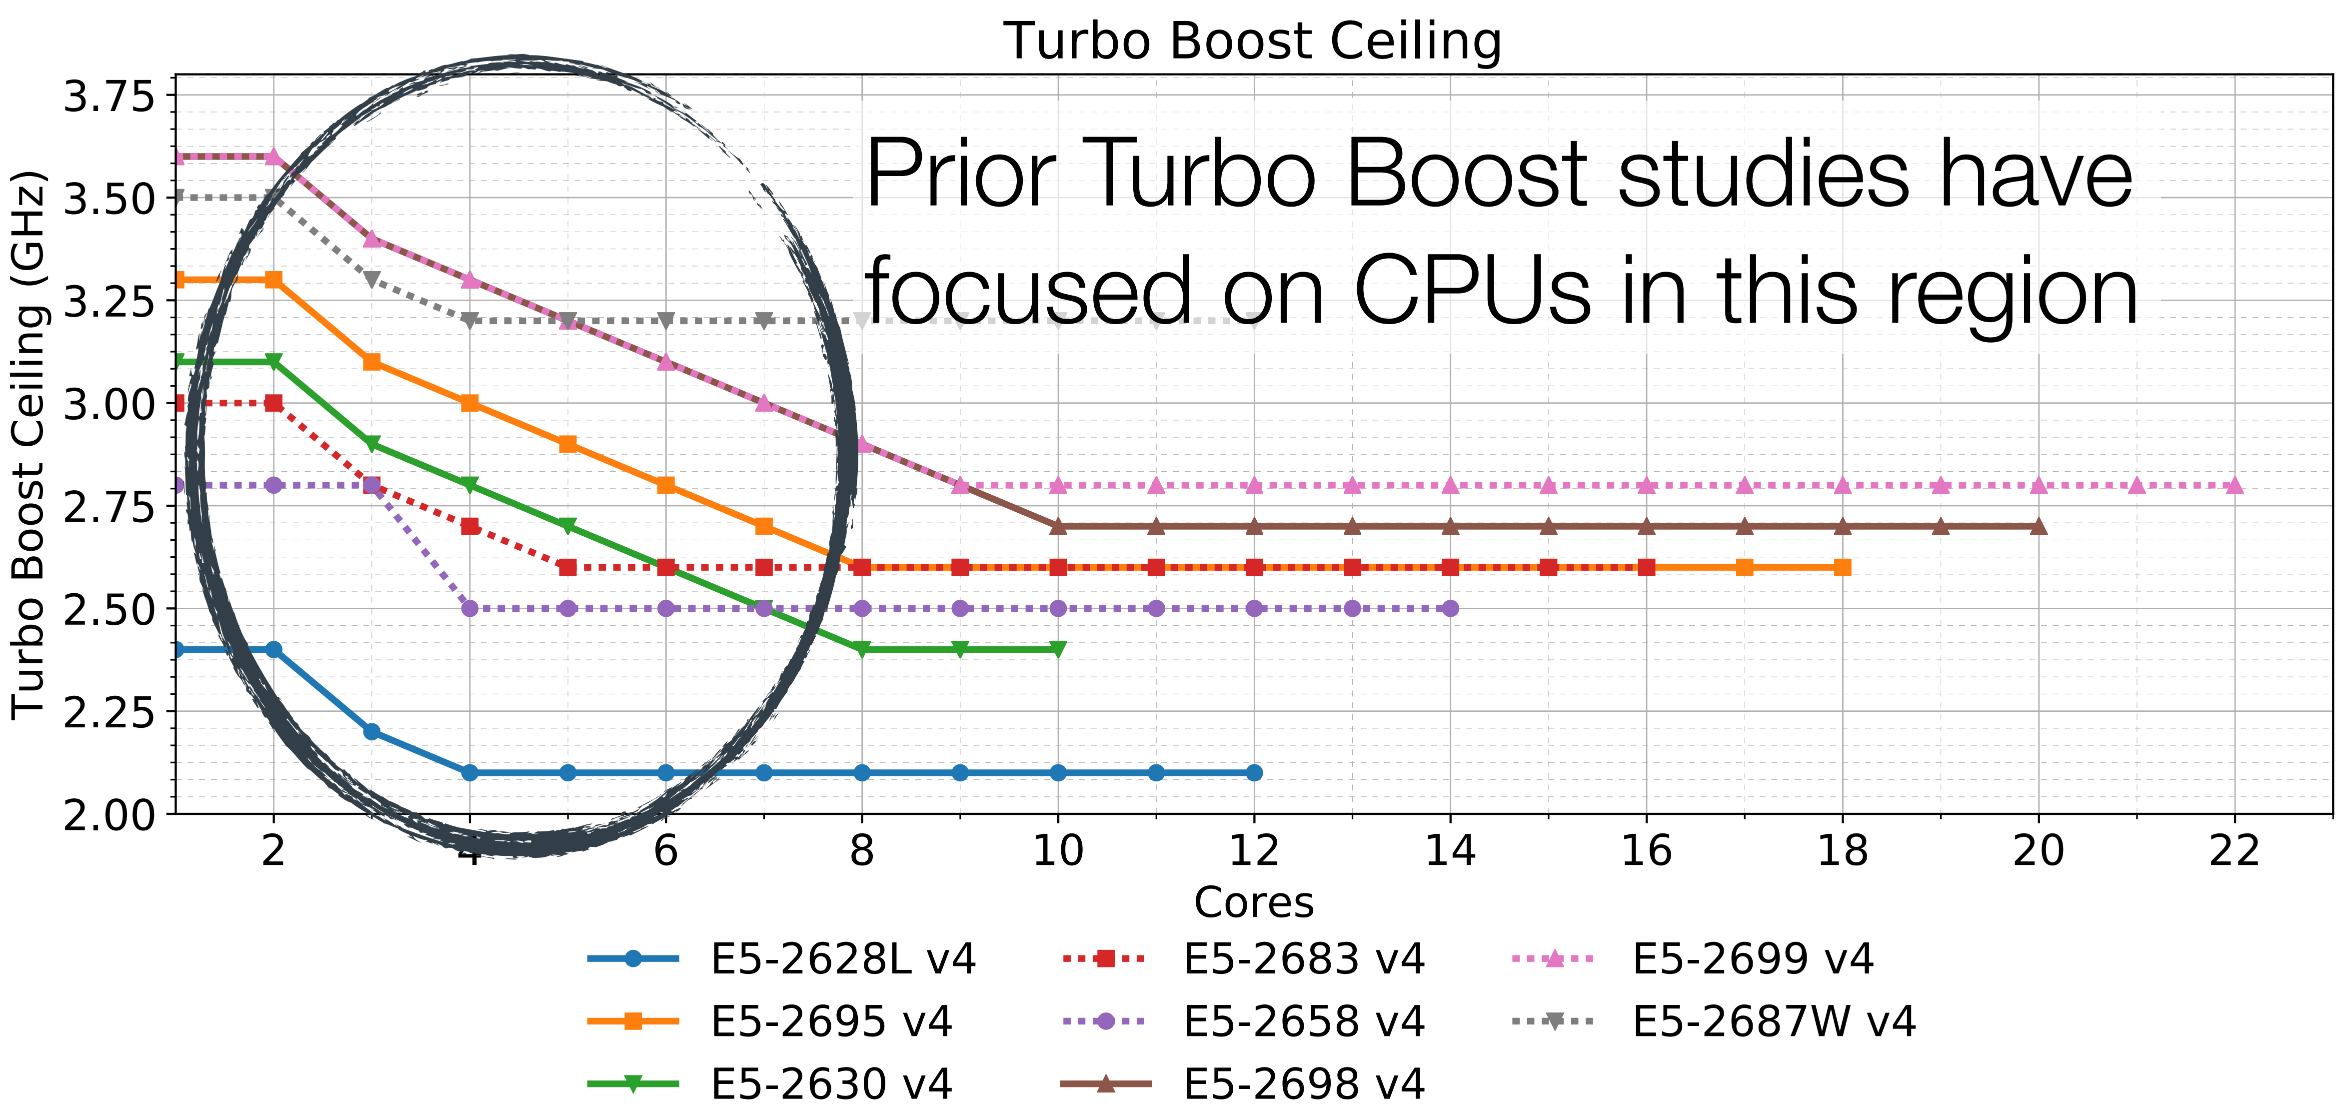

The Turbo Boost ceiling does not increase until most cores have been disabled.

A Deeper Problem…Compounded By Software

A Deeper Problem…Compounded By Software

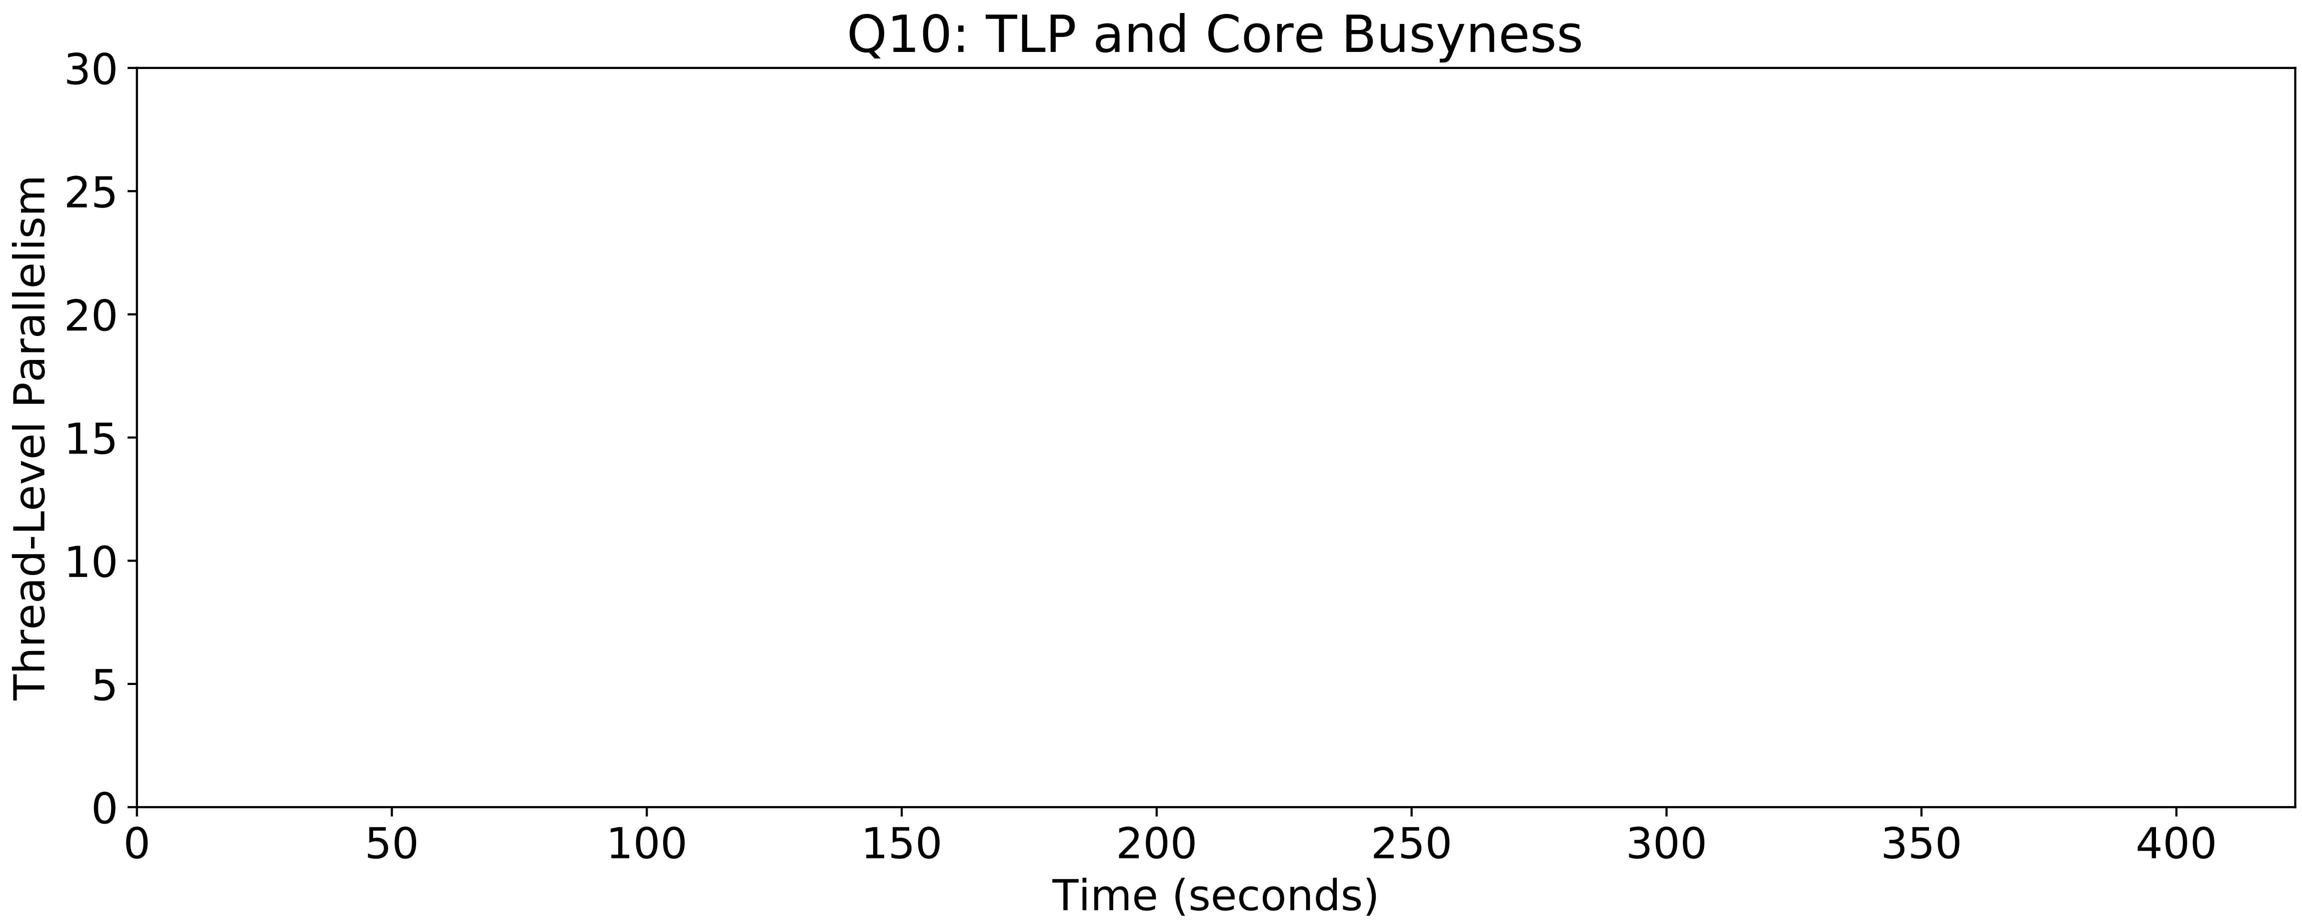

A Deeper Problem…Compounded By Software

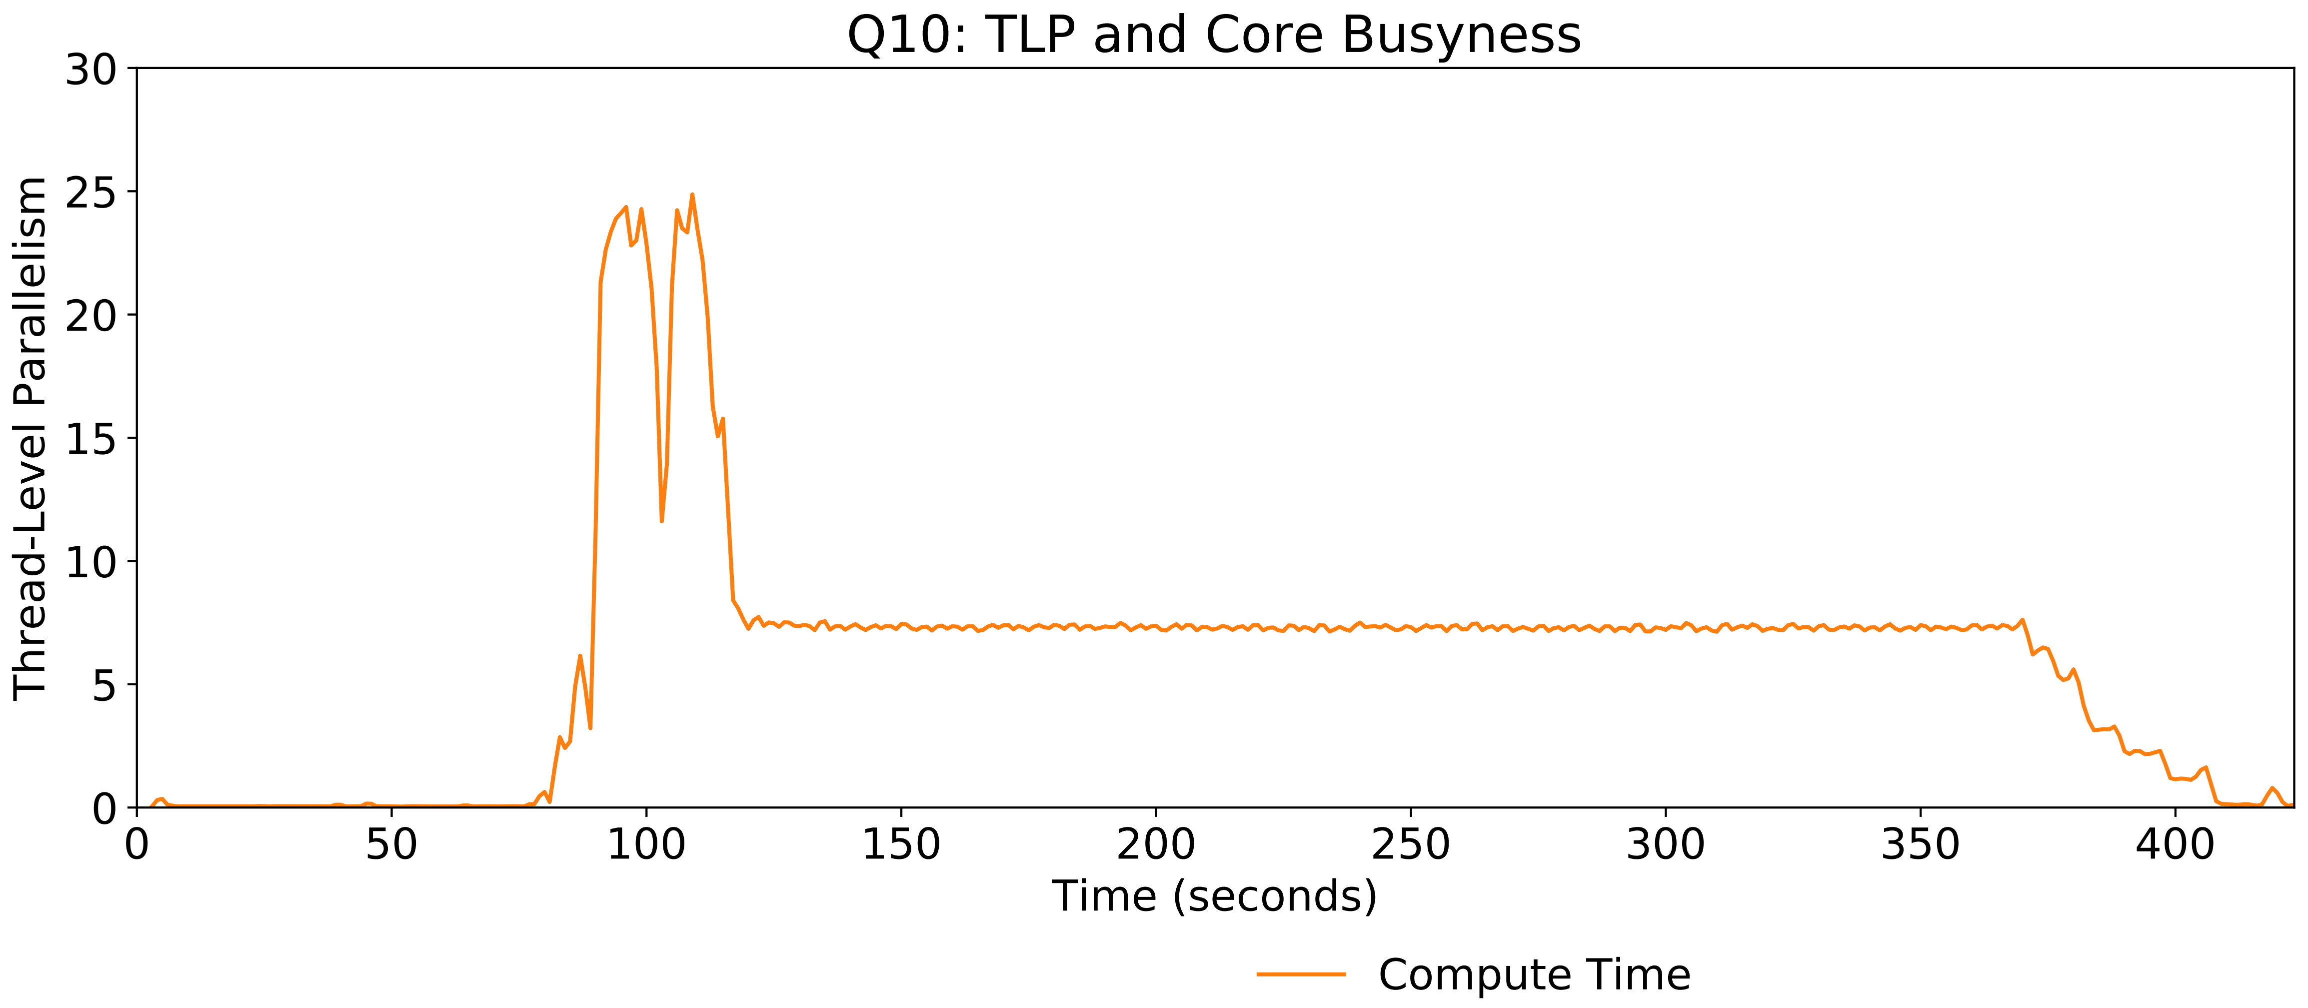

A Deeper Problem…Compounded By Software

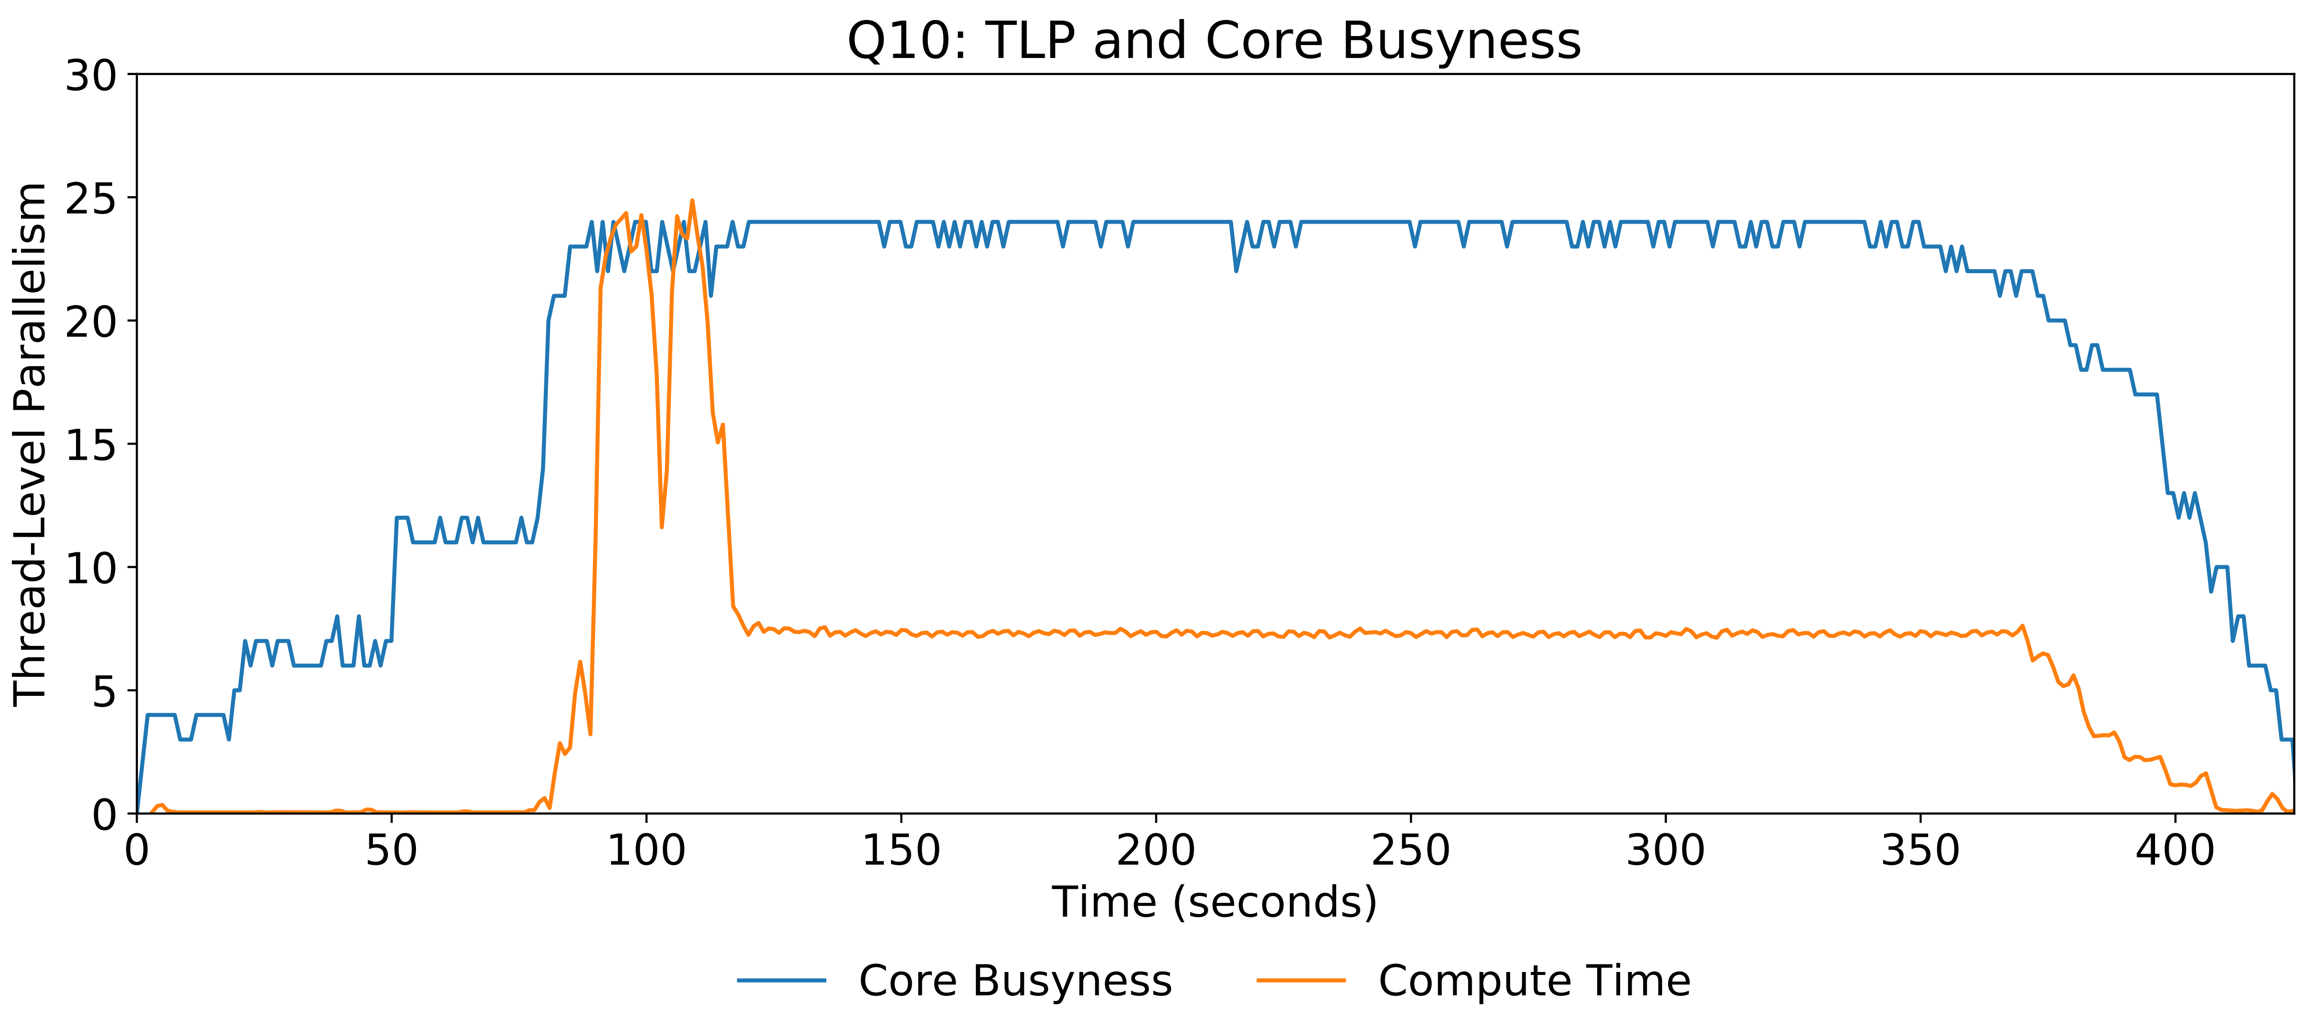

A Deeper Problem…Compounded By Software

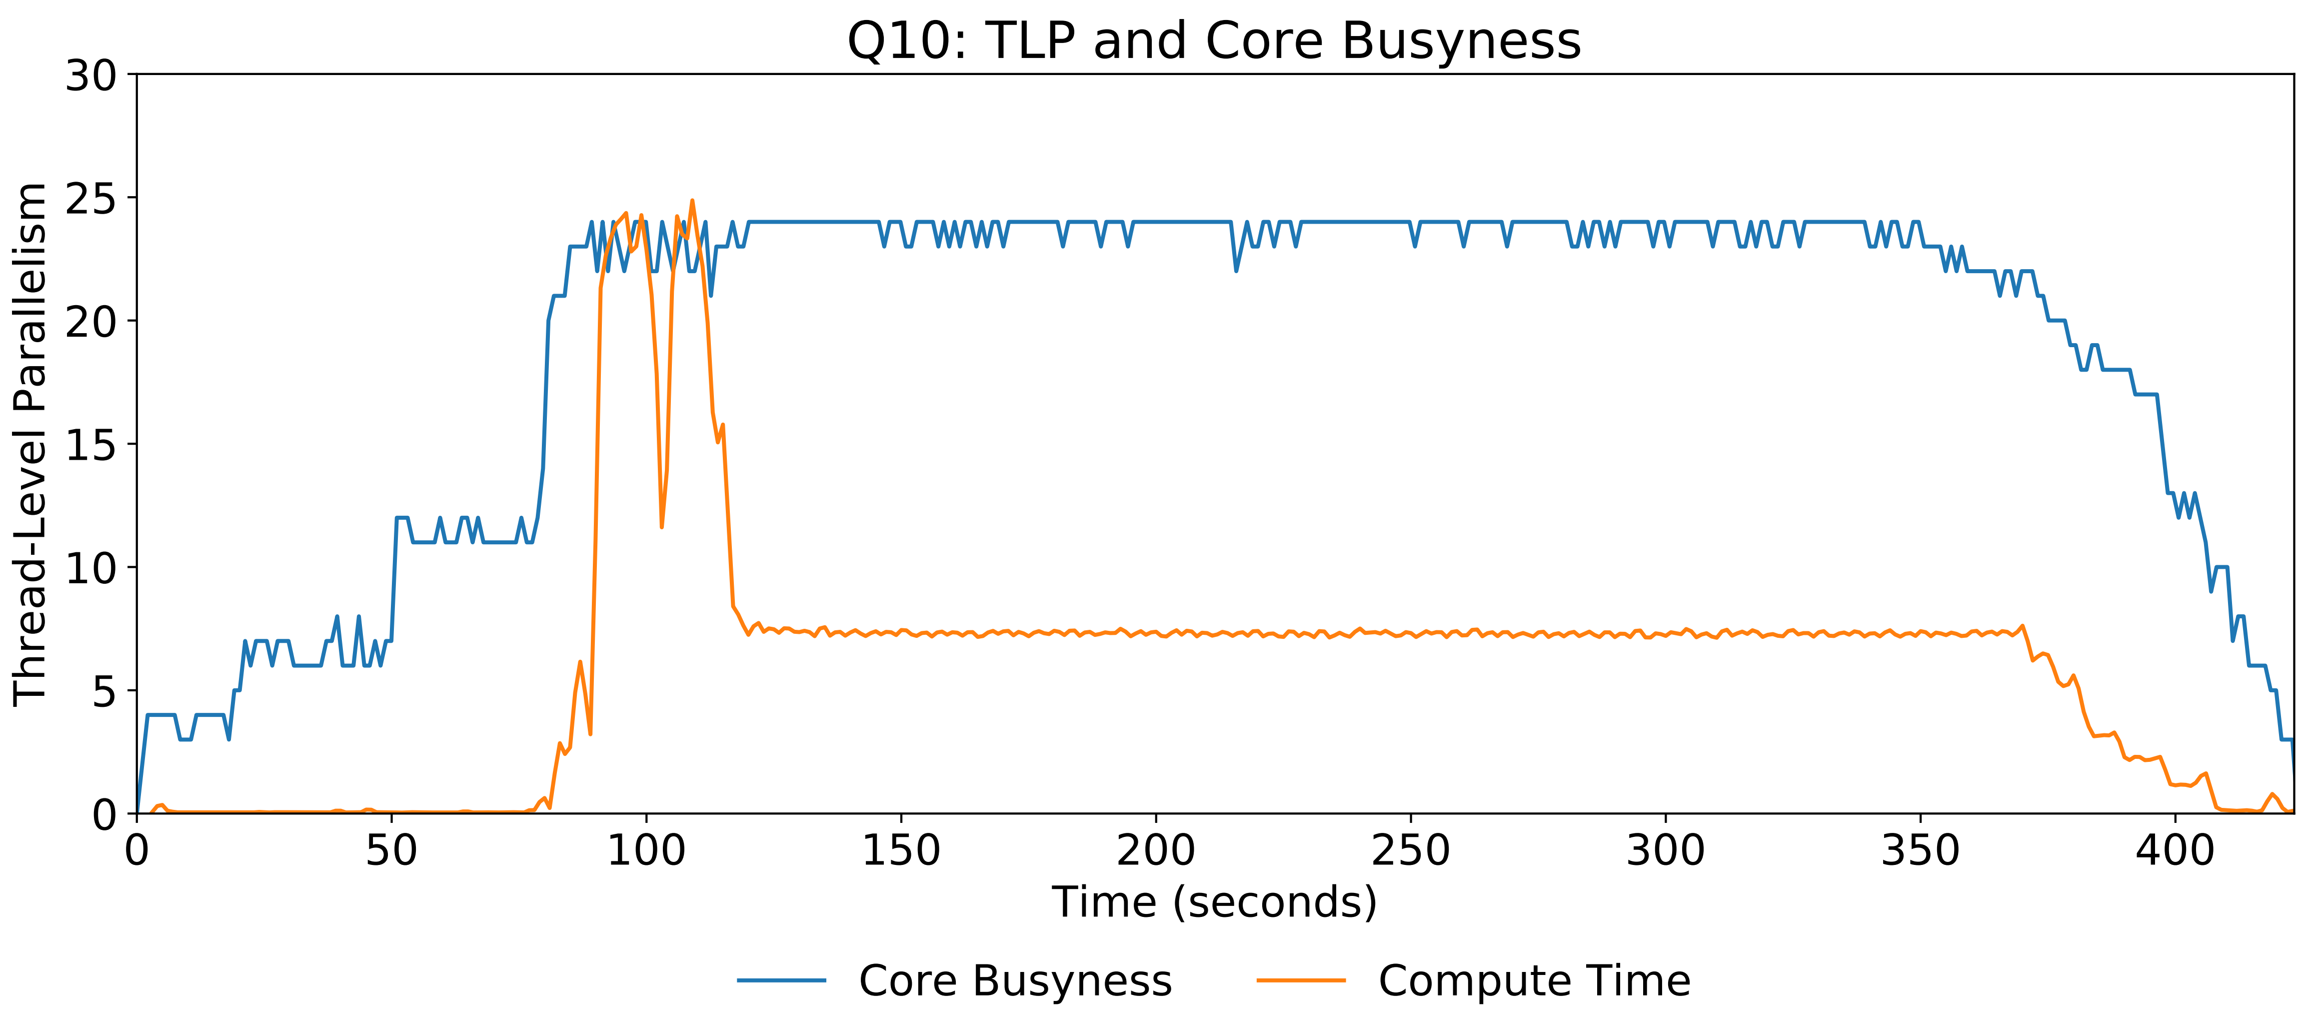

Despite having very little TLP, nearly all the cores are kept busy most of the time.

Double Whammy

Double Whammy

Hardware: Slack in thermal and current margins does not translate to a higher Turbo Boost ceiling until most cores are disabled.

Double Whammy

Hardware: Slack in thermal and current margins does not translate to a higher Turbo Boost ceiling until most cores are disabled.

Software: The abundant software threads are being scheduled onto virtually all available cores, whether they are actually needed or not.

Double Whammy

Hardware: Slack in thermal and current margins does not translate to a higher Turbo Boost ceiling until most cores

Hardware and software cooperatively prevent
Turbo Boost from ever exceeding its baseline ceiling.

Software: The abundant software threads are being scheduled onto virtually all available cores, whether they are actually needed or not.

Core Packing

Core Packing

Core packing is our proposal for proactive enhancement of Turbo Boost through hardware and software cooperation.

Core Packing

Core packing is our proposal for proactive enhancement of Turbo Boost through hardware and software cooperation.

Core Packing

Core packing is our proposal for proactive enhancement of Turbo Boost through hardware and software cooperation.

Status Quo

Core Packing

Core packing is our proposal for proactive enhancement of Turbo Boost through hardware and software cooperation.

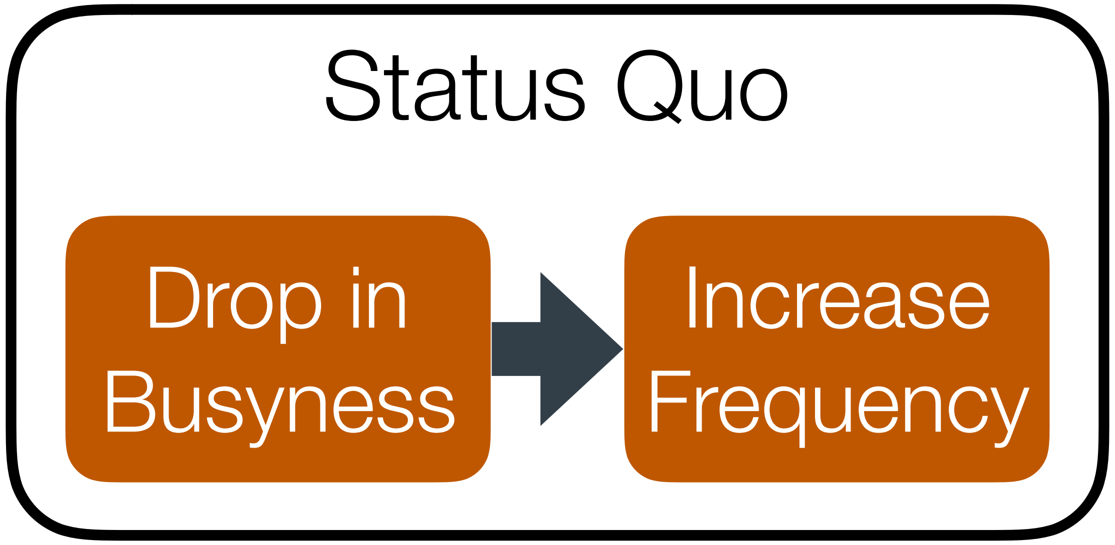

Core Packing

Core packing is our proposal for proactive enhancement of Turbo Boost through hardware and software cooperation.

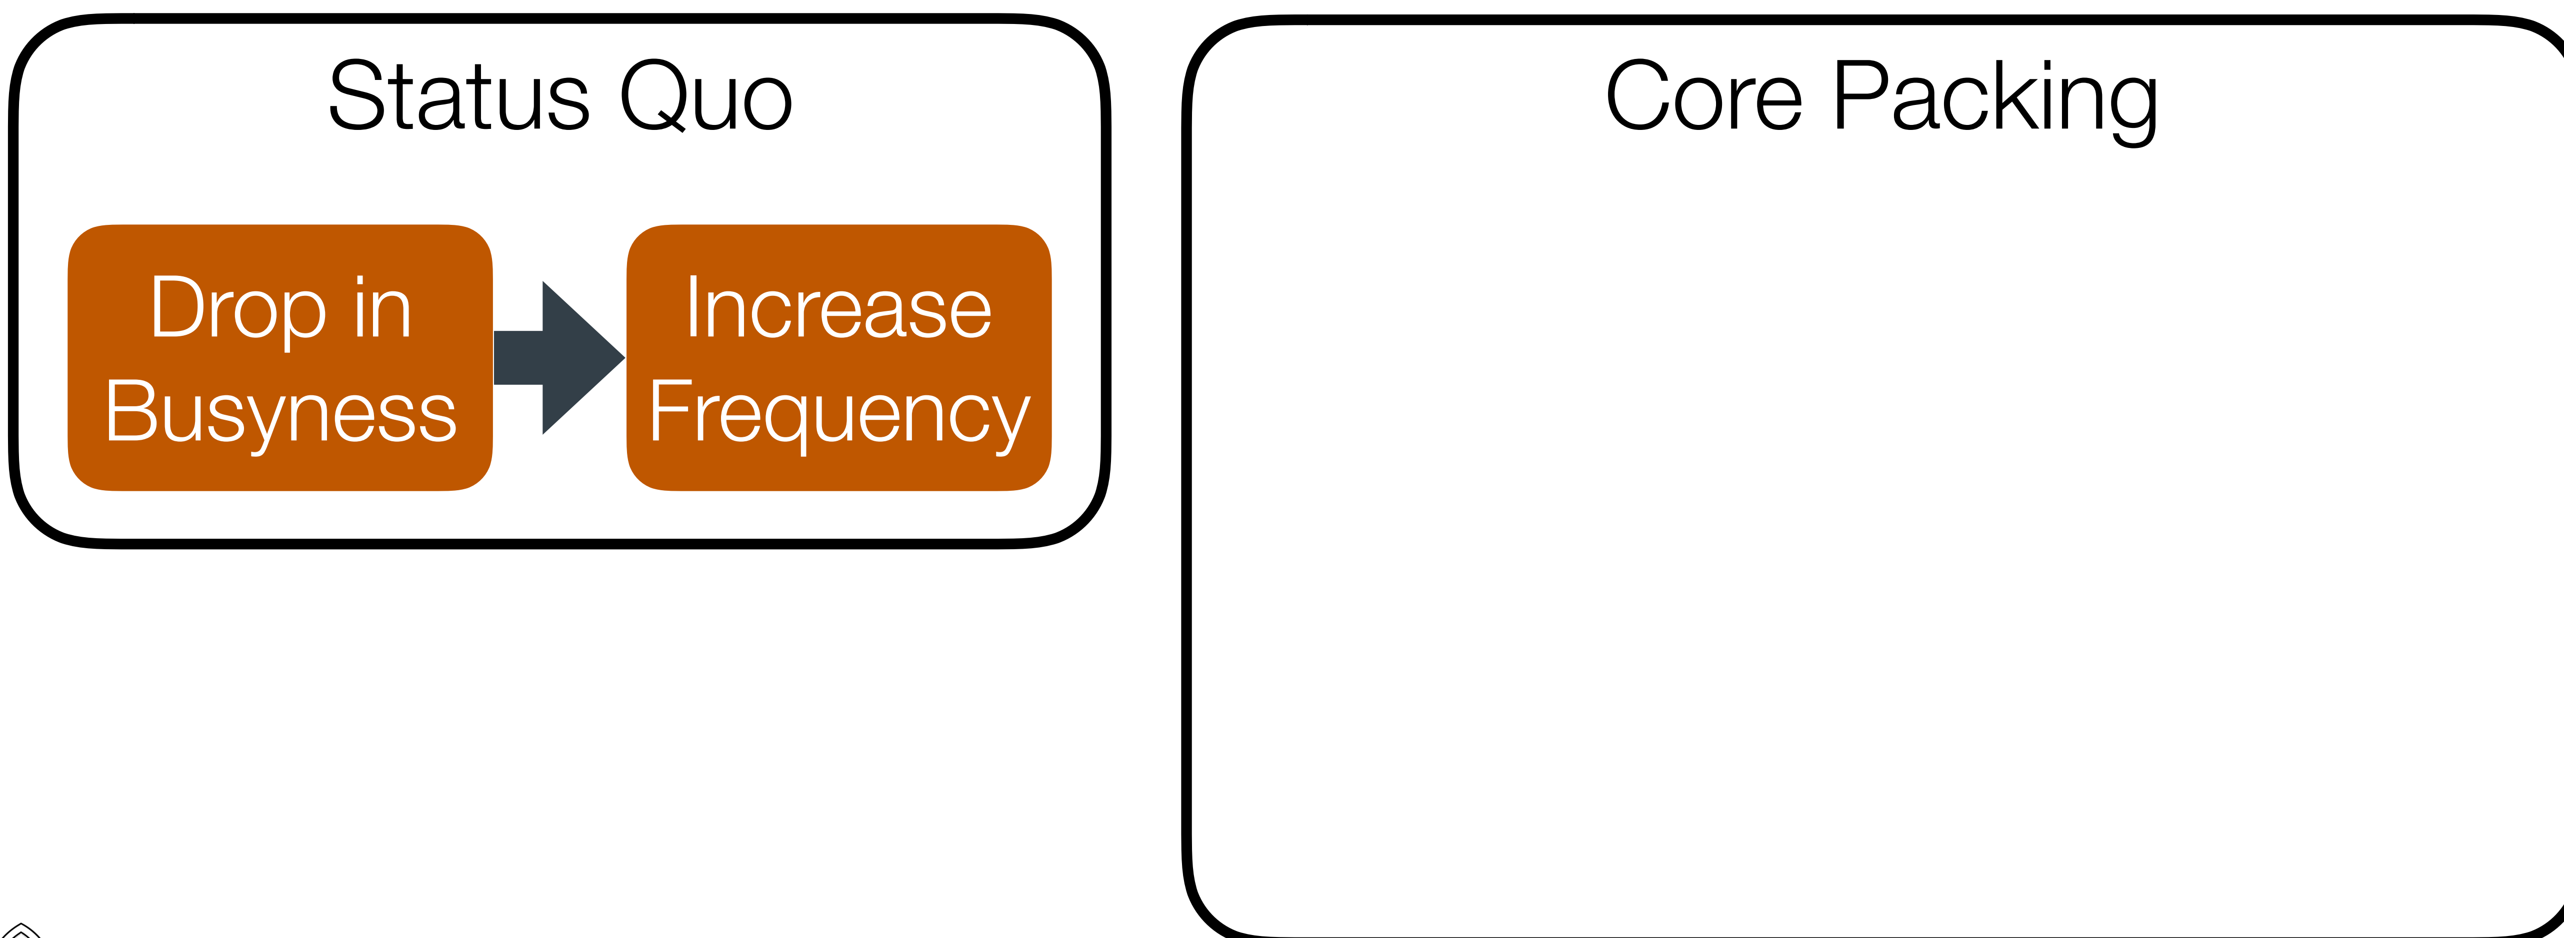

Core Packing

Core packing is our proposal for proactive enhancement of Turbo Boost through hardware and software cooperation.

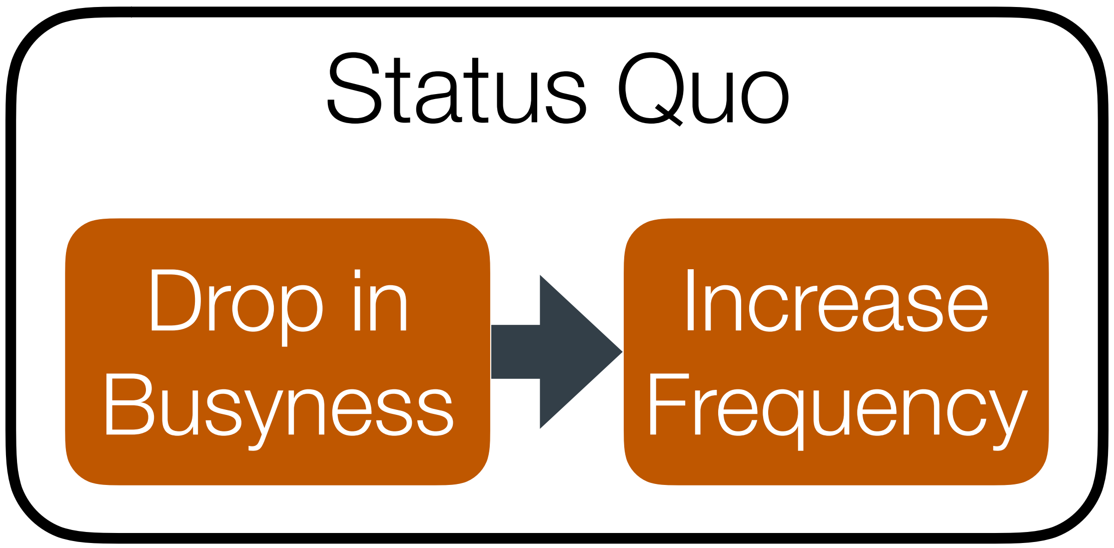

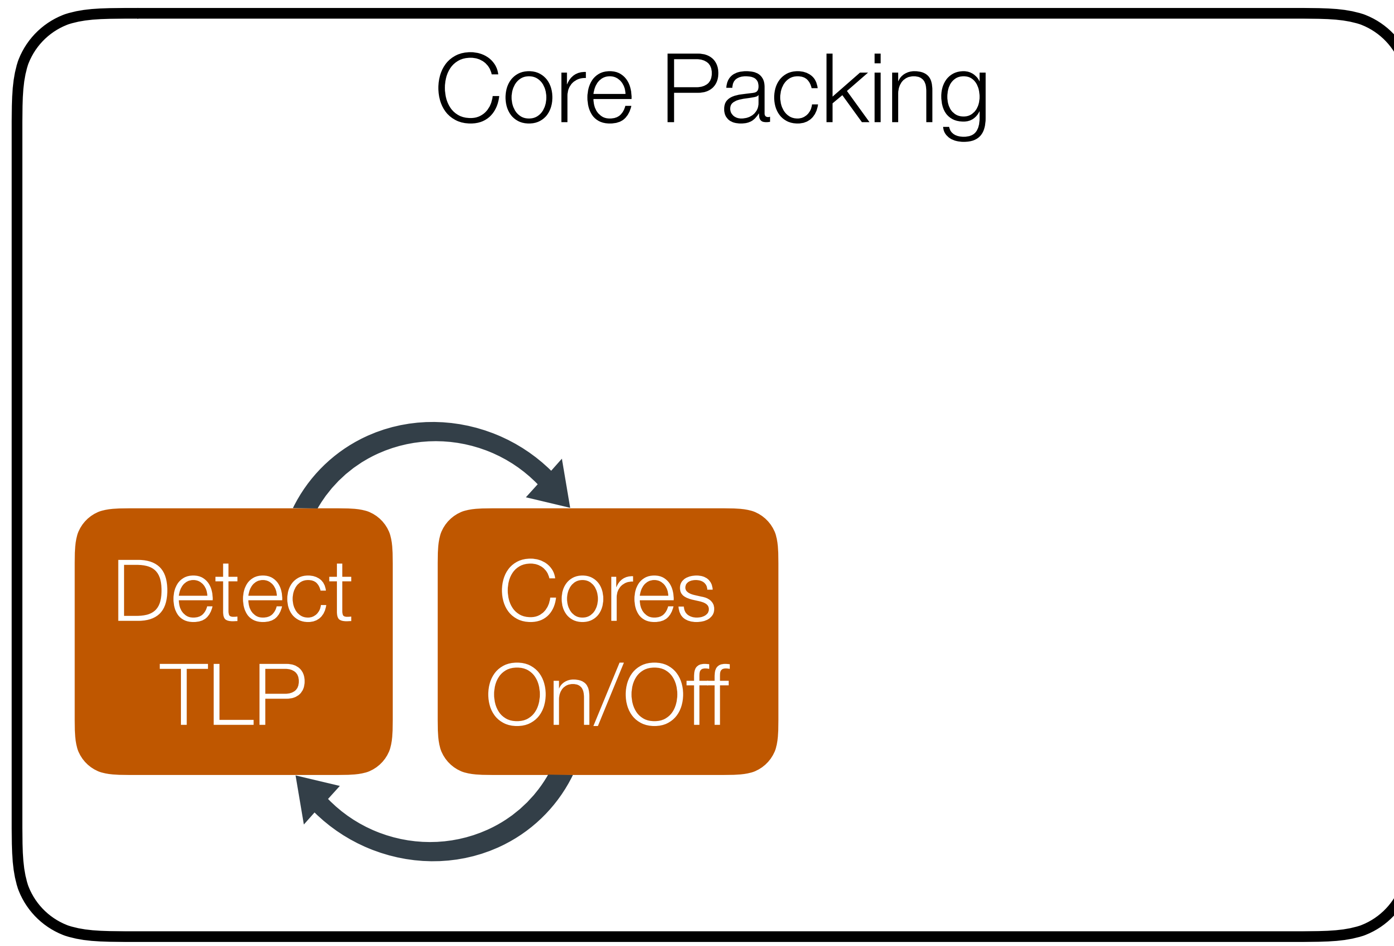

Core Packing

Core packing is our proposal for proactive enhancement of Turbo Boost through hardware and software cooperation.

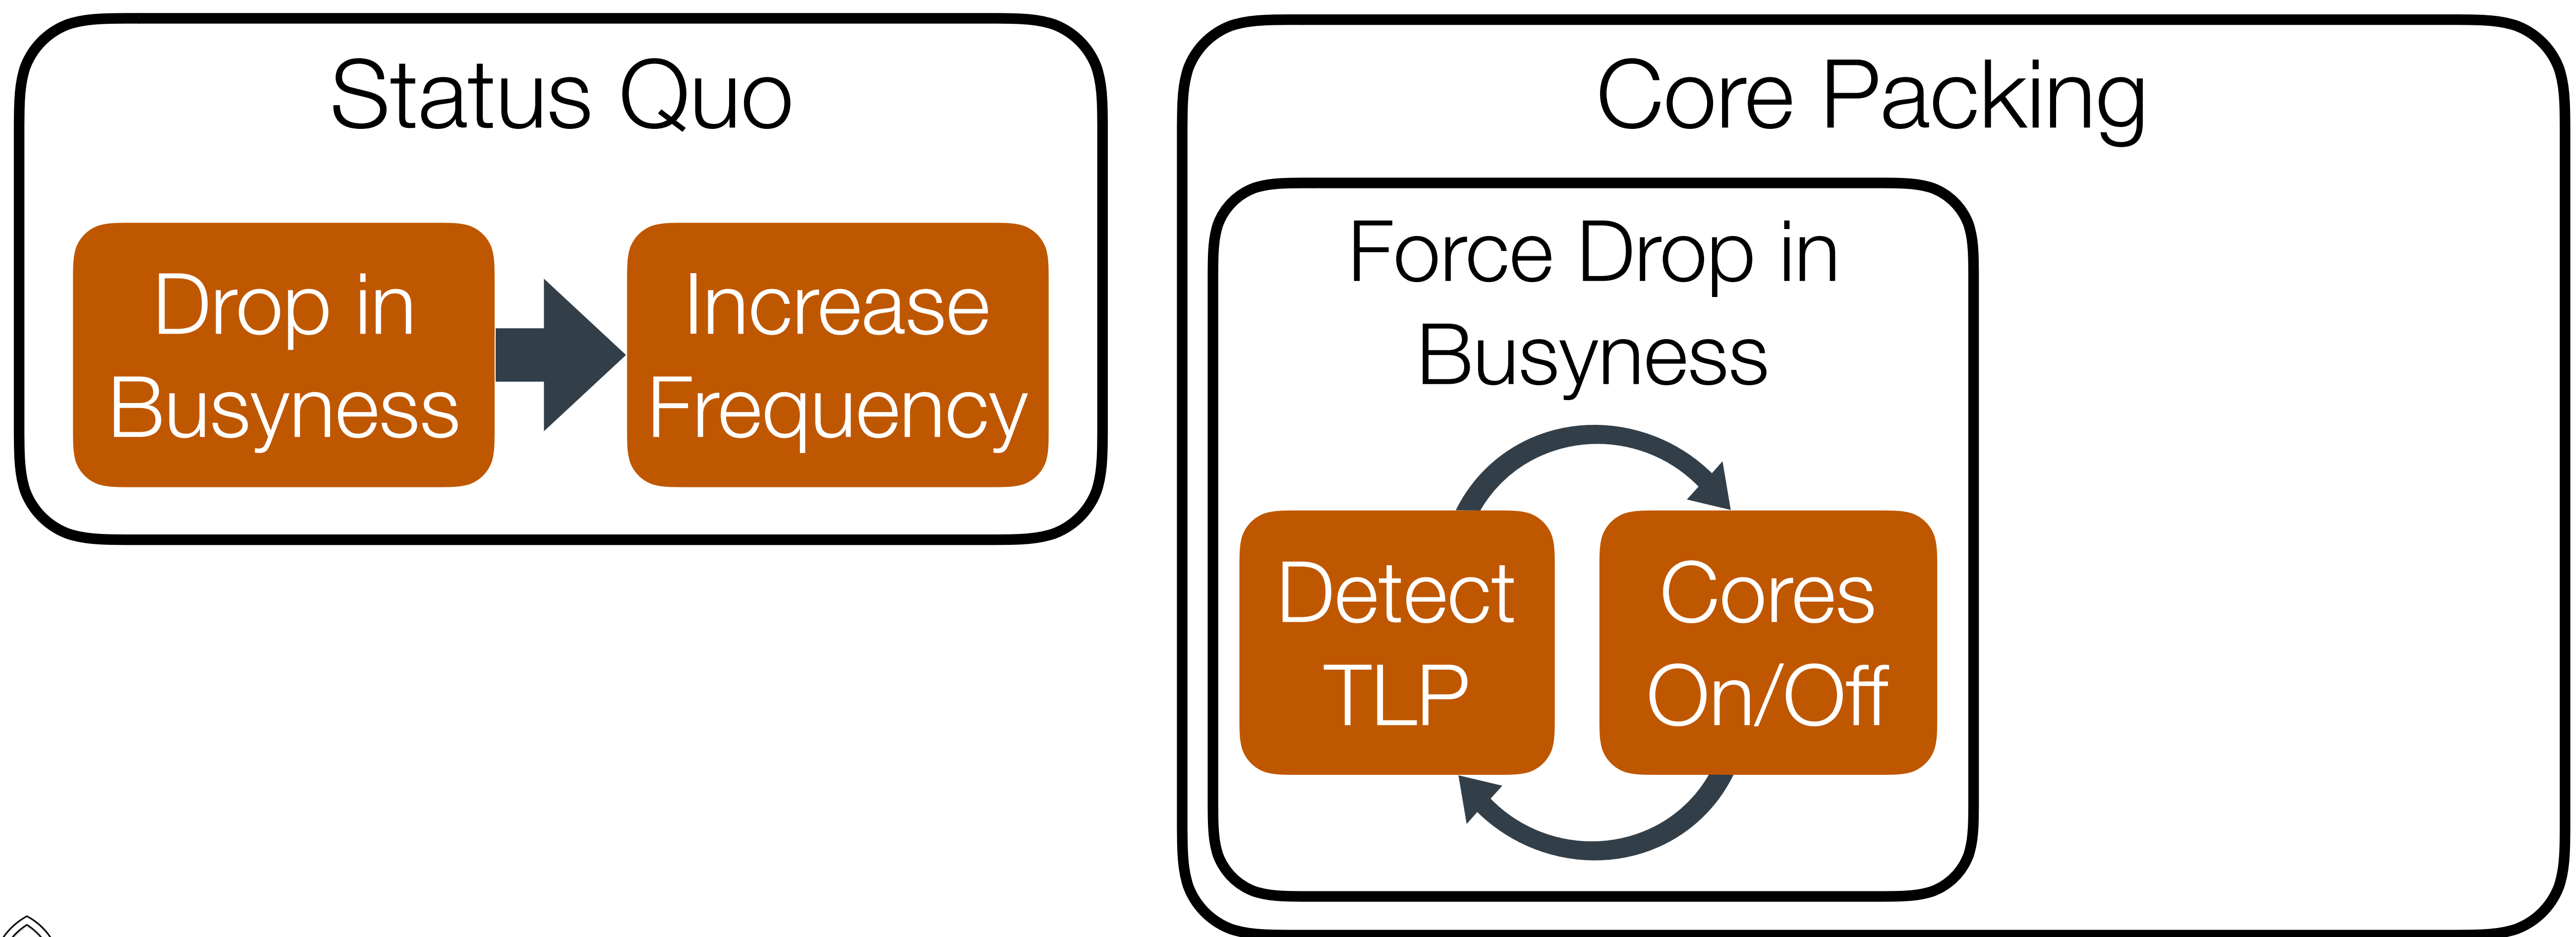

Core Packing

Core packing is our proposal for proactive enhancement of Turbo Boost through hardware and software cooperation.

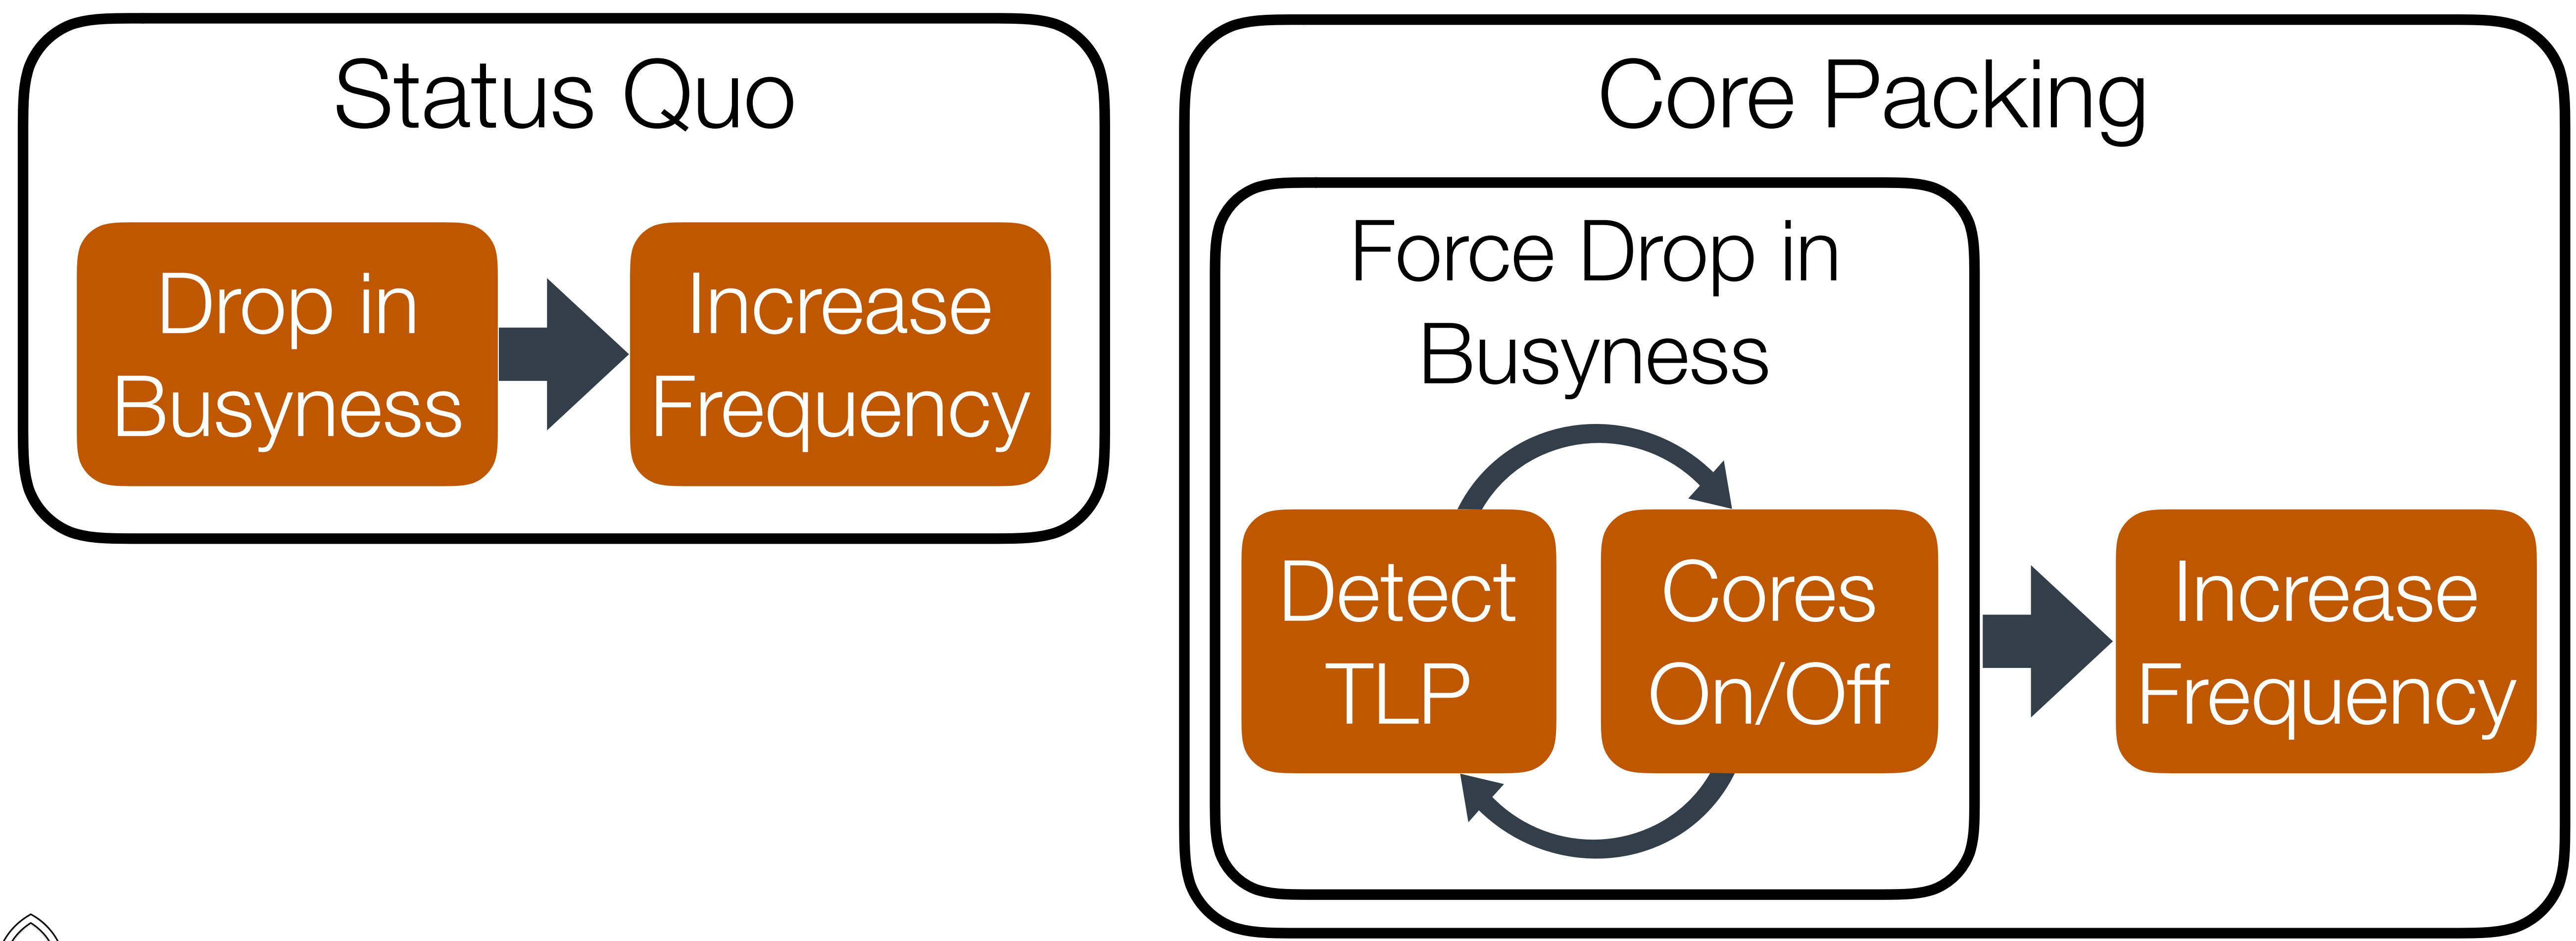

Core Packing—Hardware Proposal

Core Packing—Hardware Proposal

Allow for linear increase in Turbo Boost ceiling with the number of deactivated cores.

Core Packing—Hardware Proposal

Allow for linear increase in Turbo Boost ceiling with the number of deactivated cores.

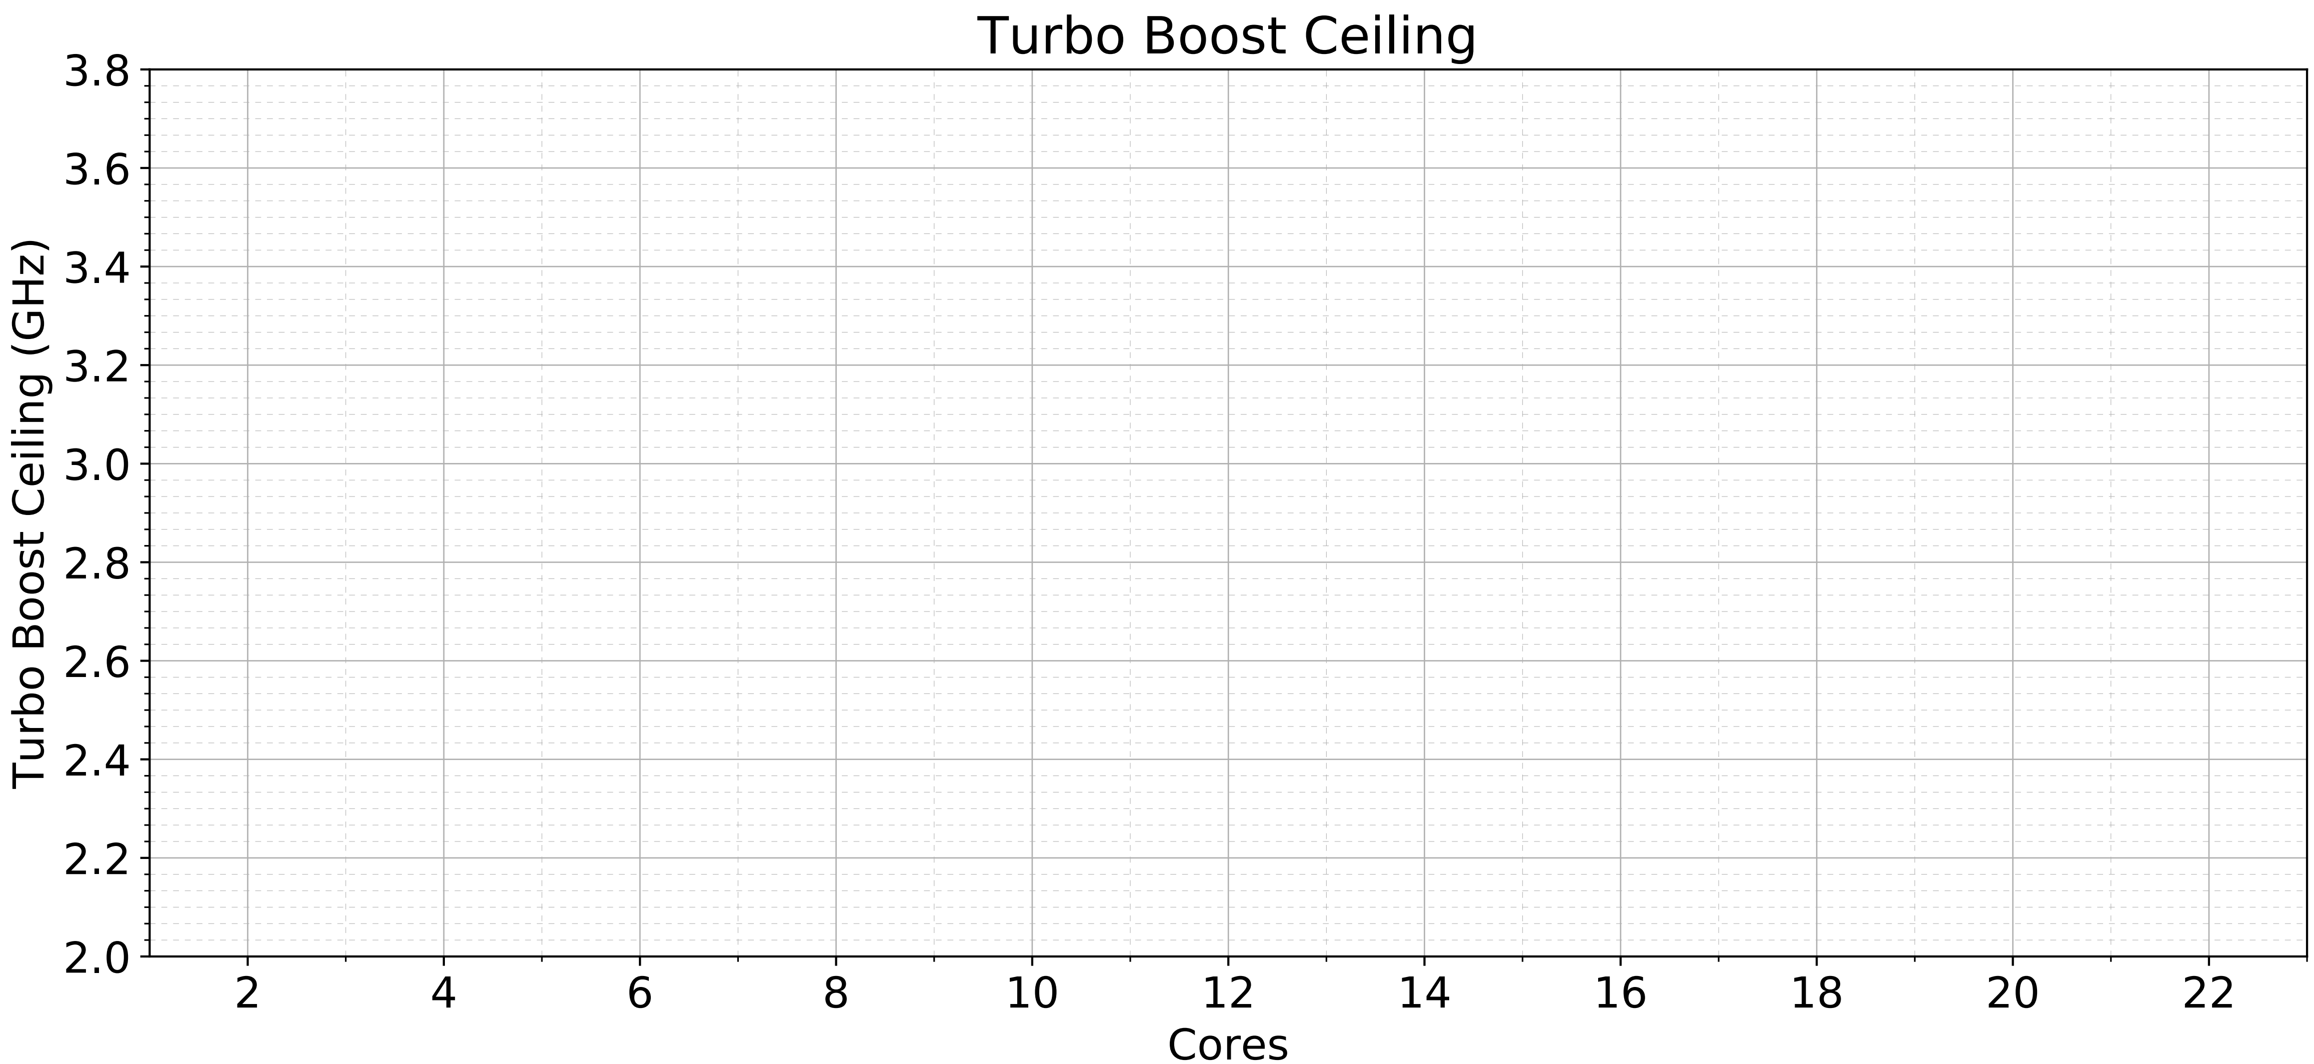

Core Packing—Hardware Proposal

Allow for linear increase in Turbo Boost ceiling with the number of deactivated cores.

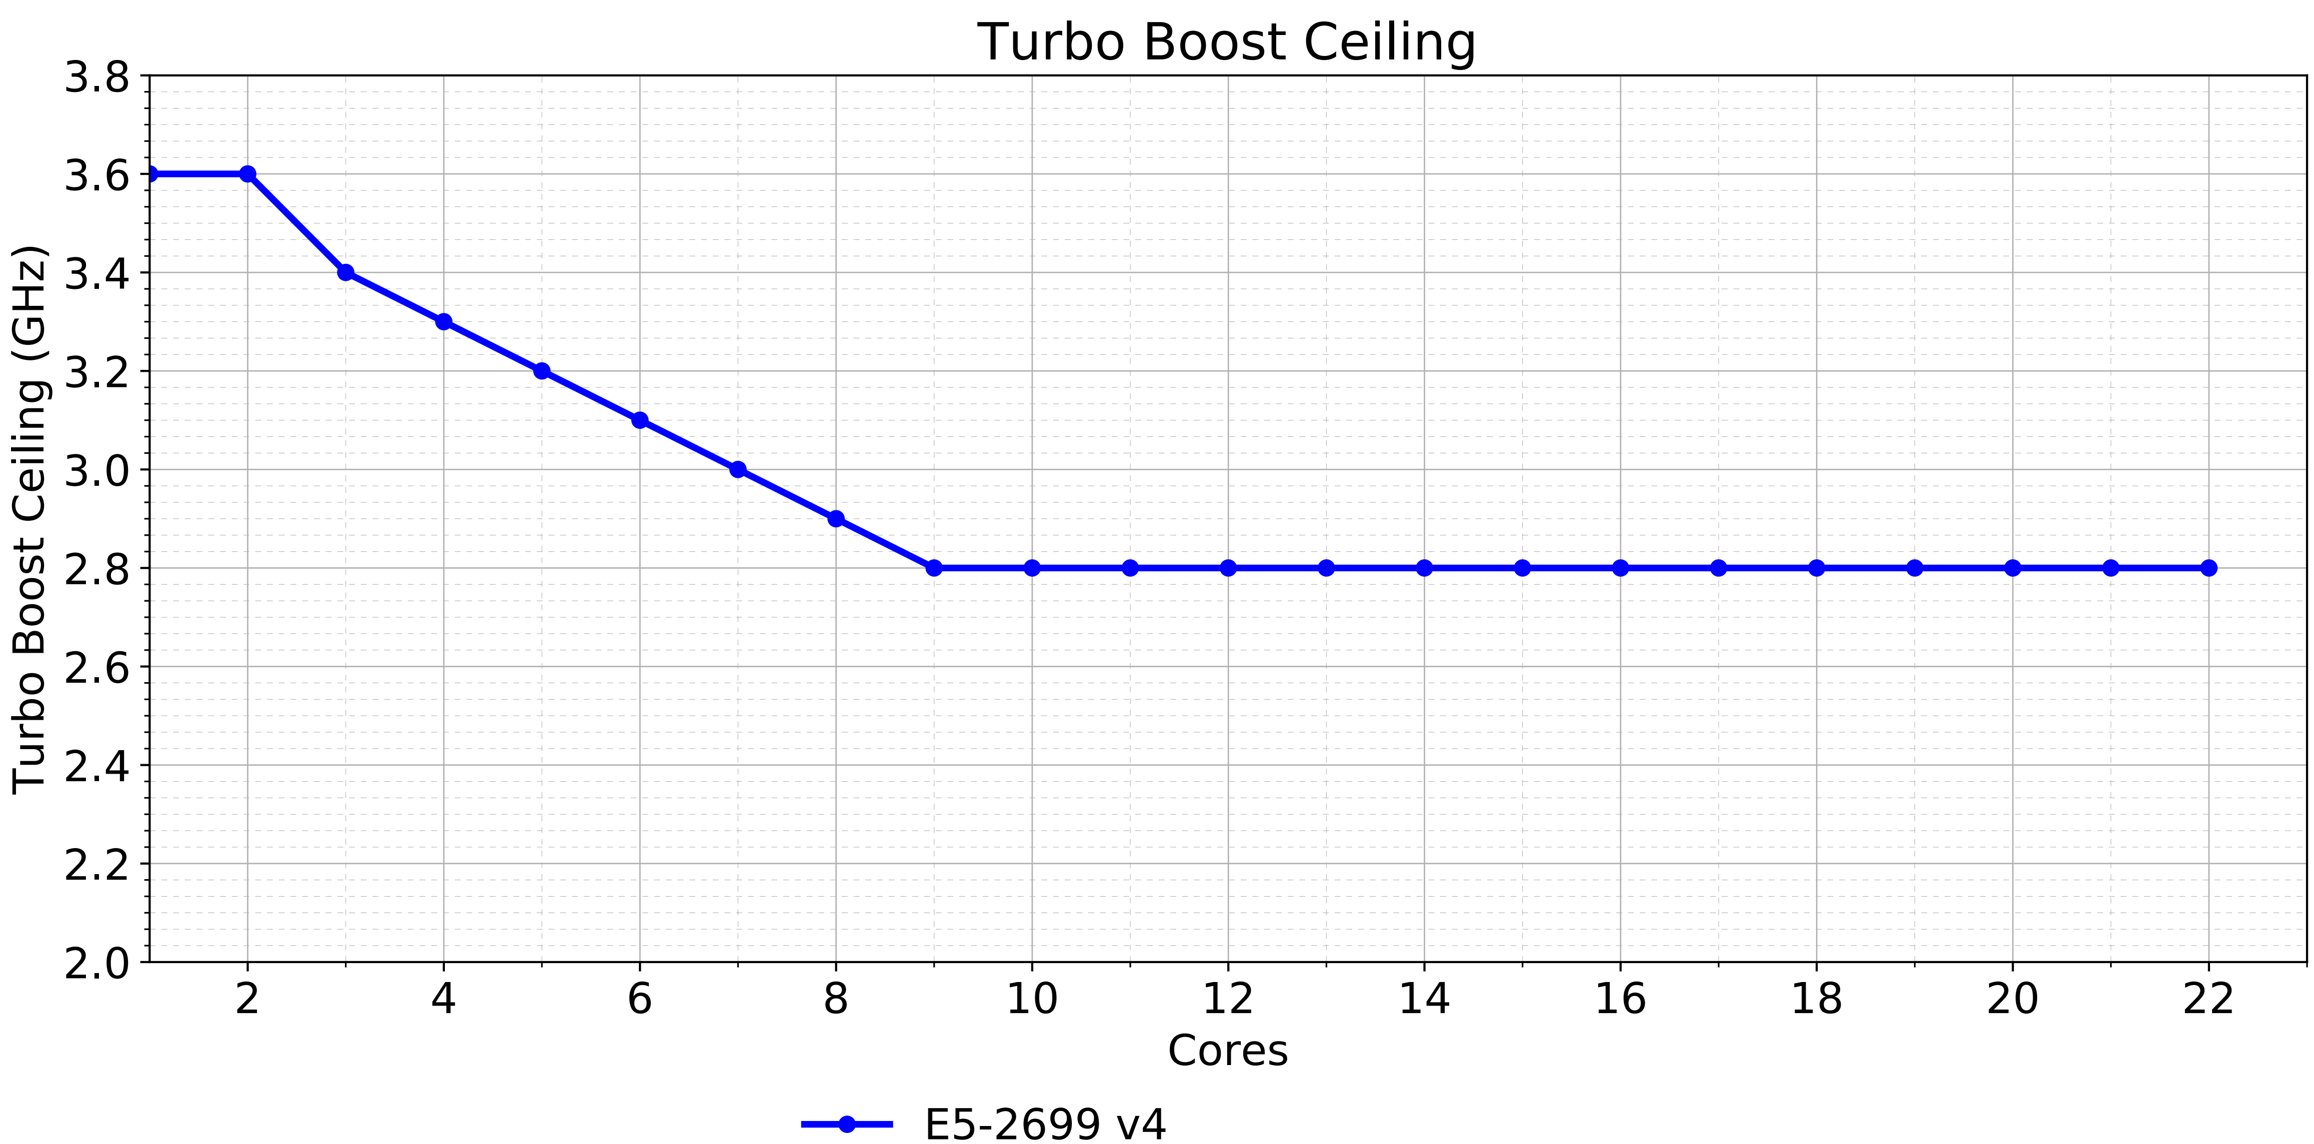

Core Packing—Hardware Proposal

Allow for linear increase in Turbo Boost ceiling with the number of deactivated cores.

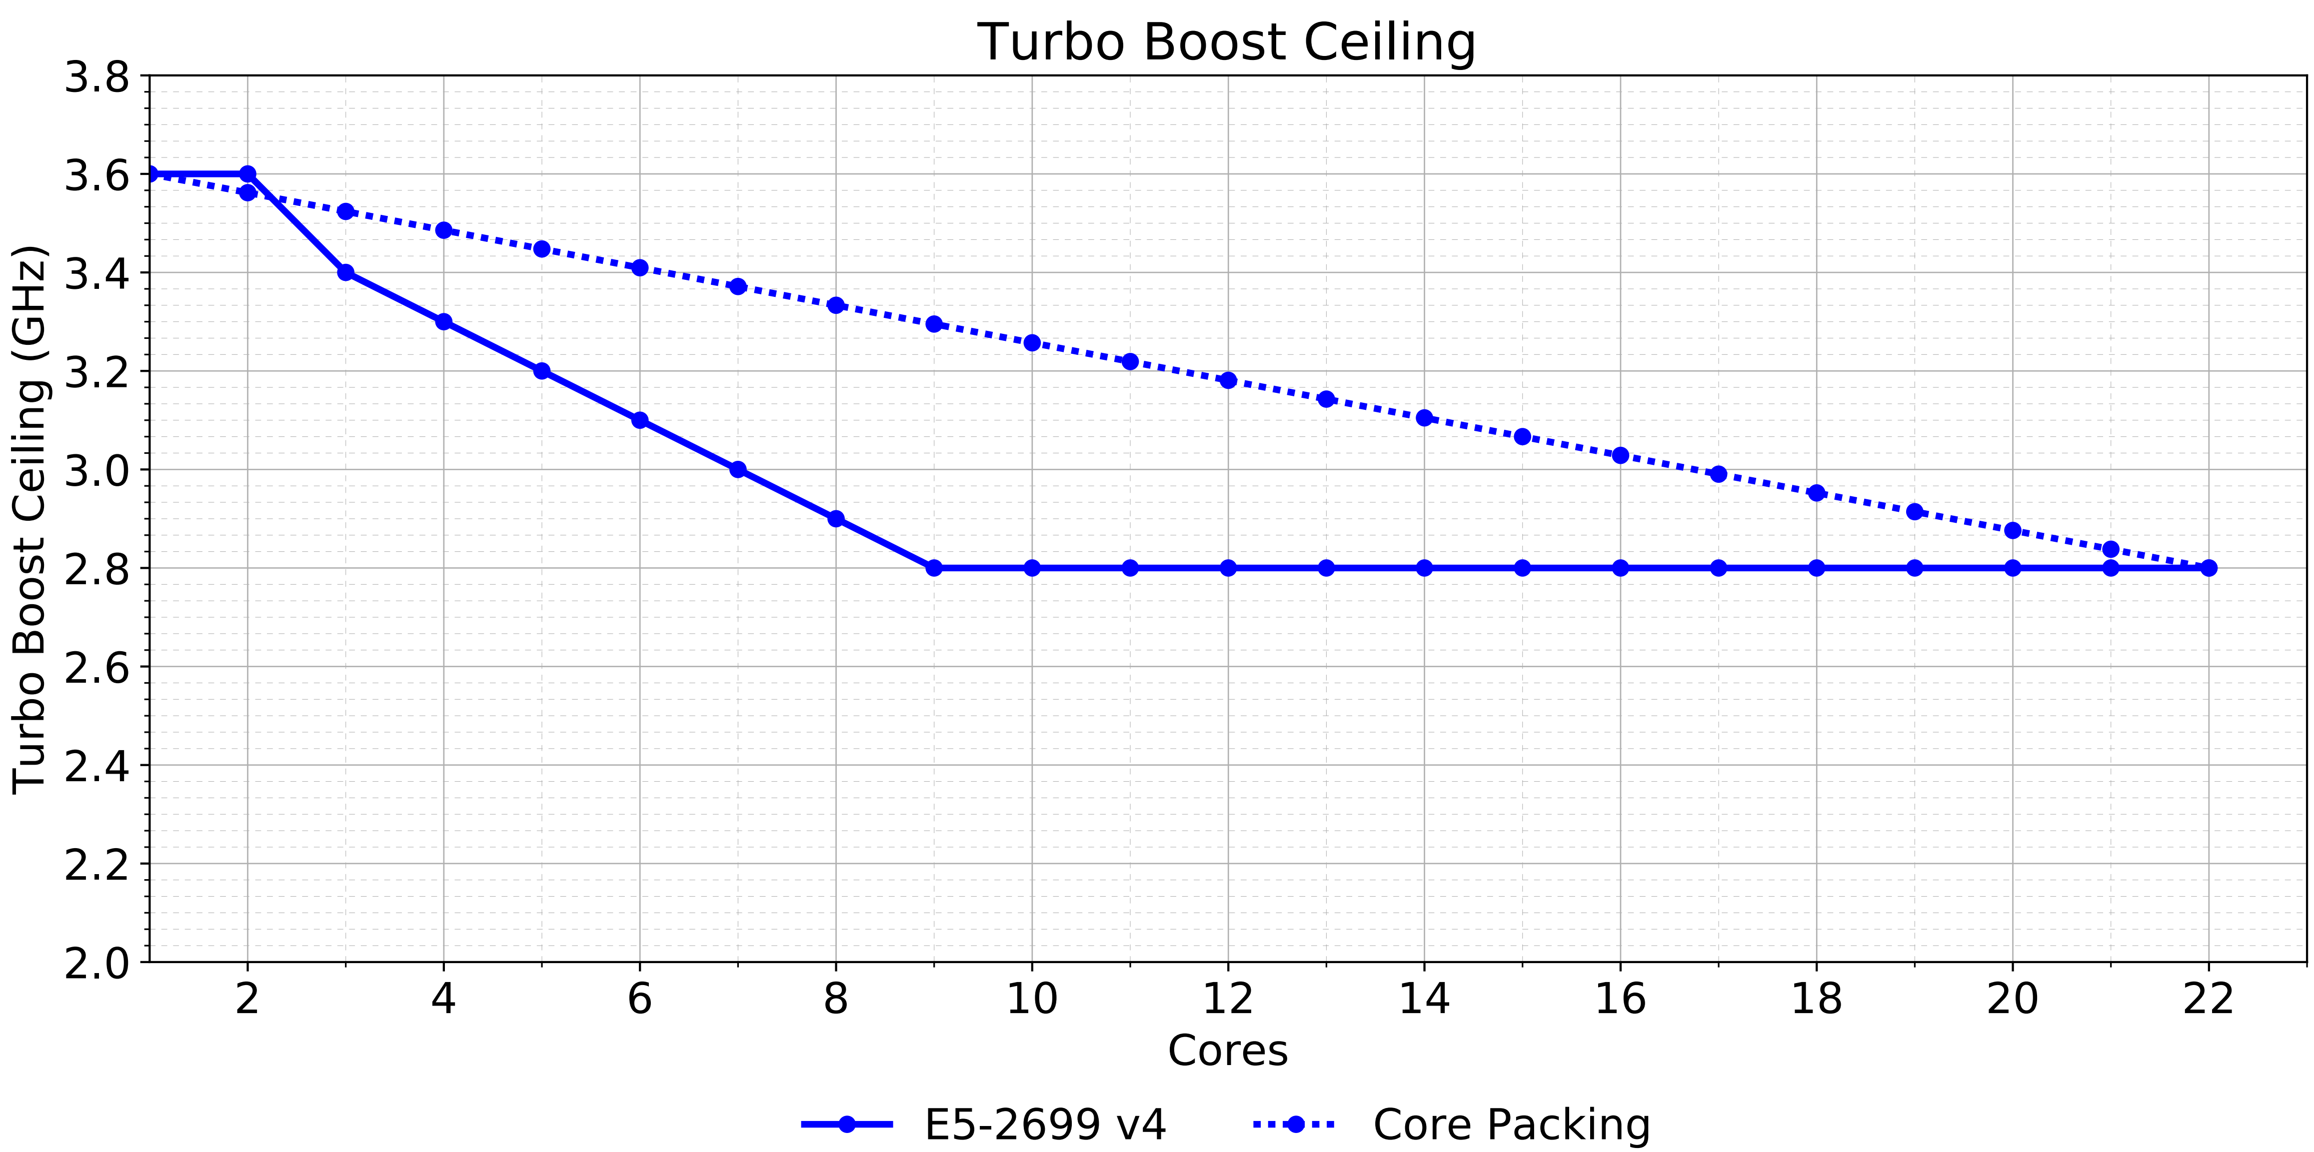

Core Packing—Software Proposal

Core Packing—Software Proposal

Proactively restrict the number of active cores to just meet the workload.

Core Packing—Software Proposal

Proactively restrict the number of active cores to just meet the workload.

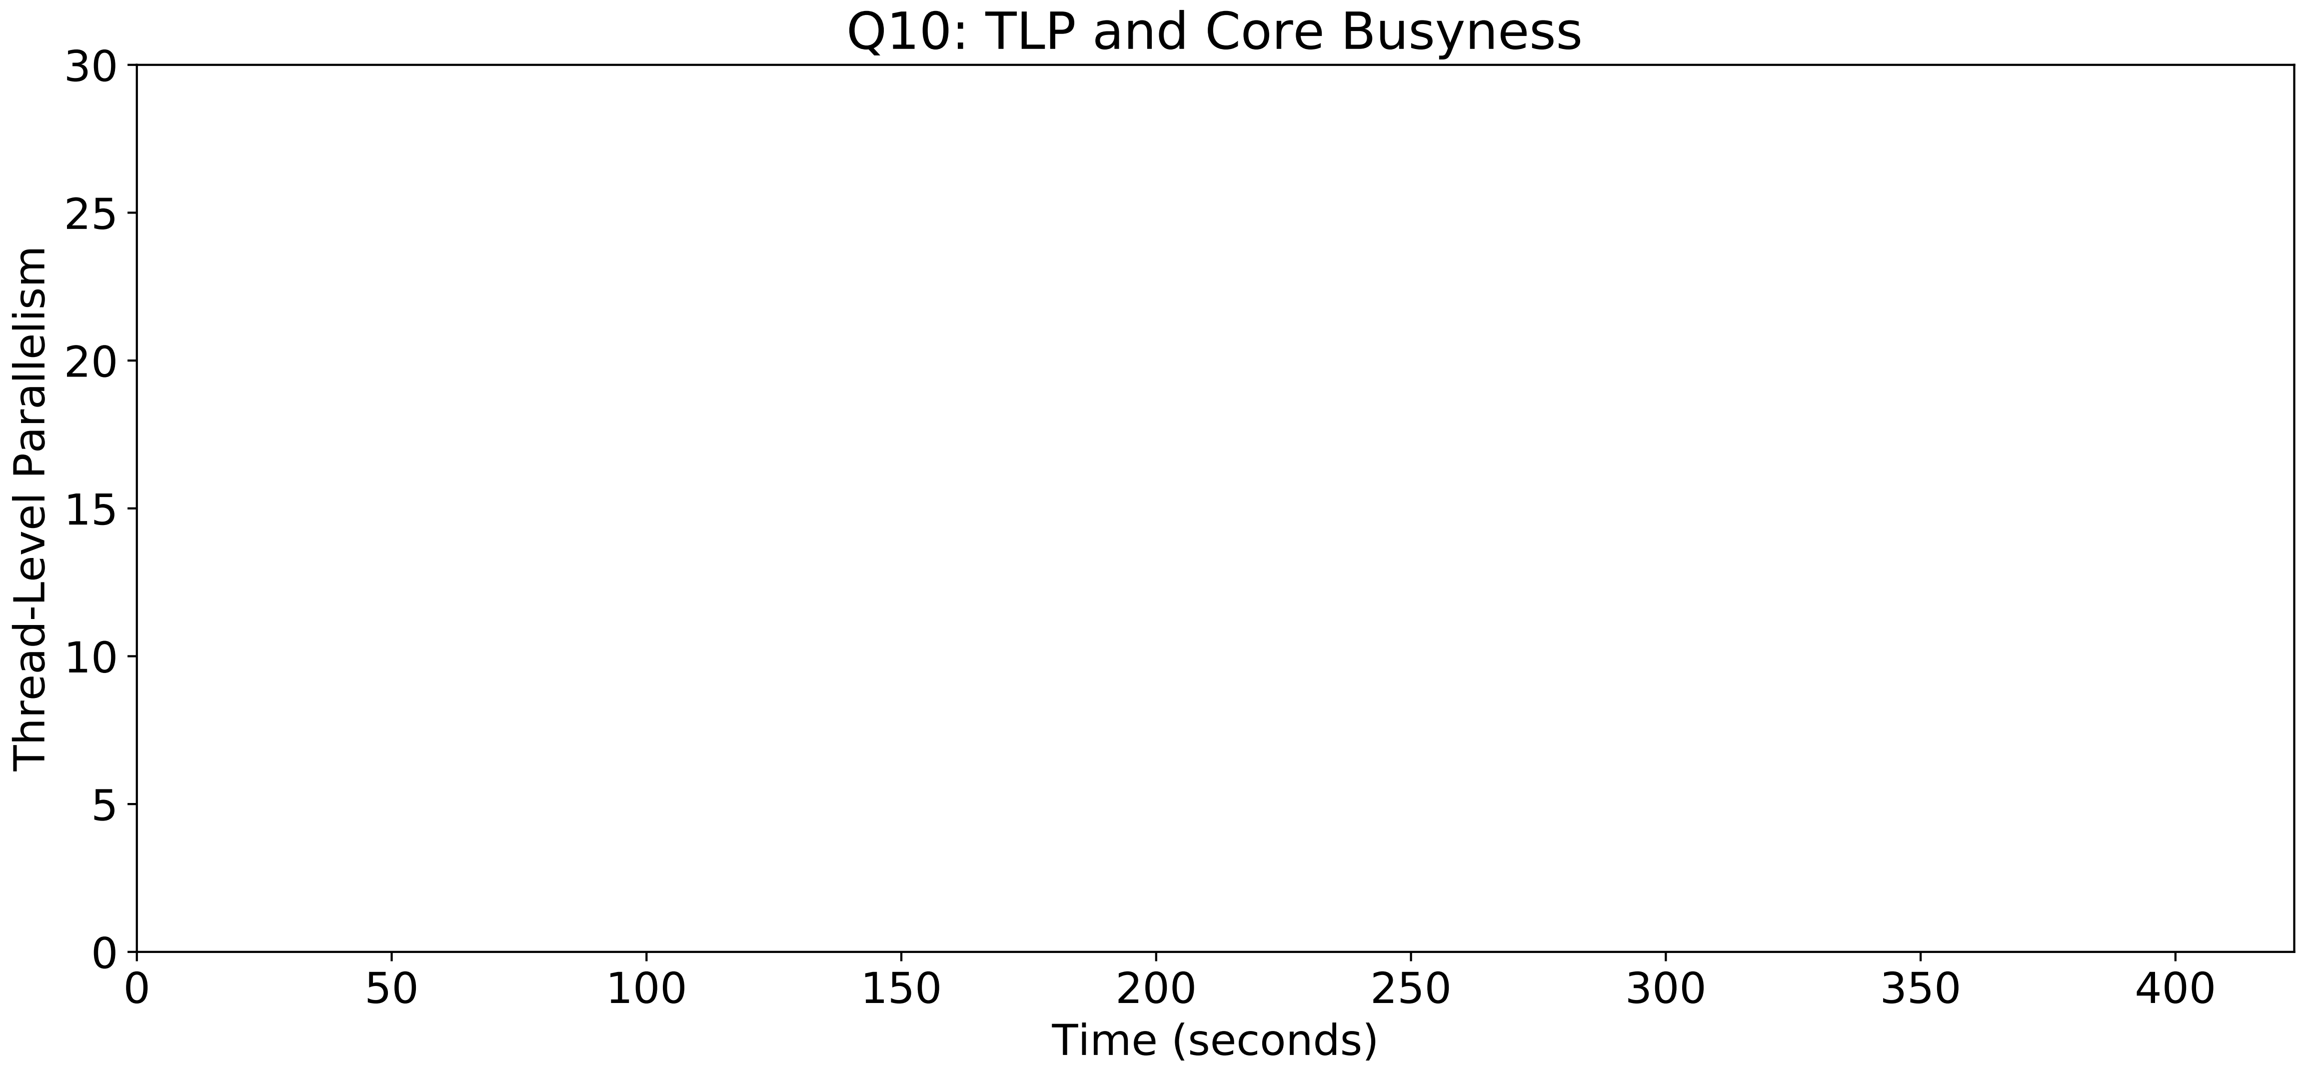

Core Packing—Software Proposal

Proactively restrict the number of active cores to just meet the workload.

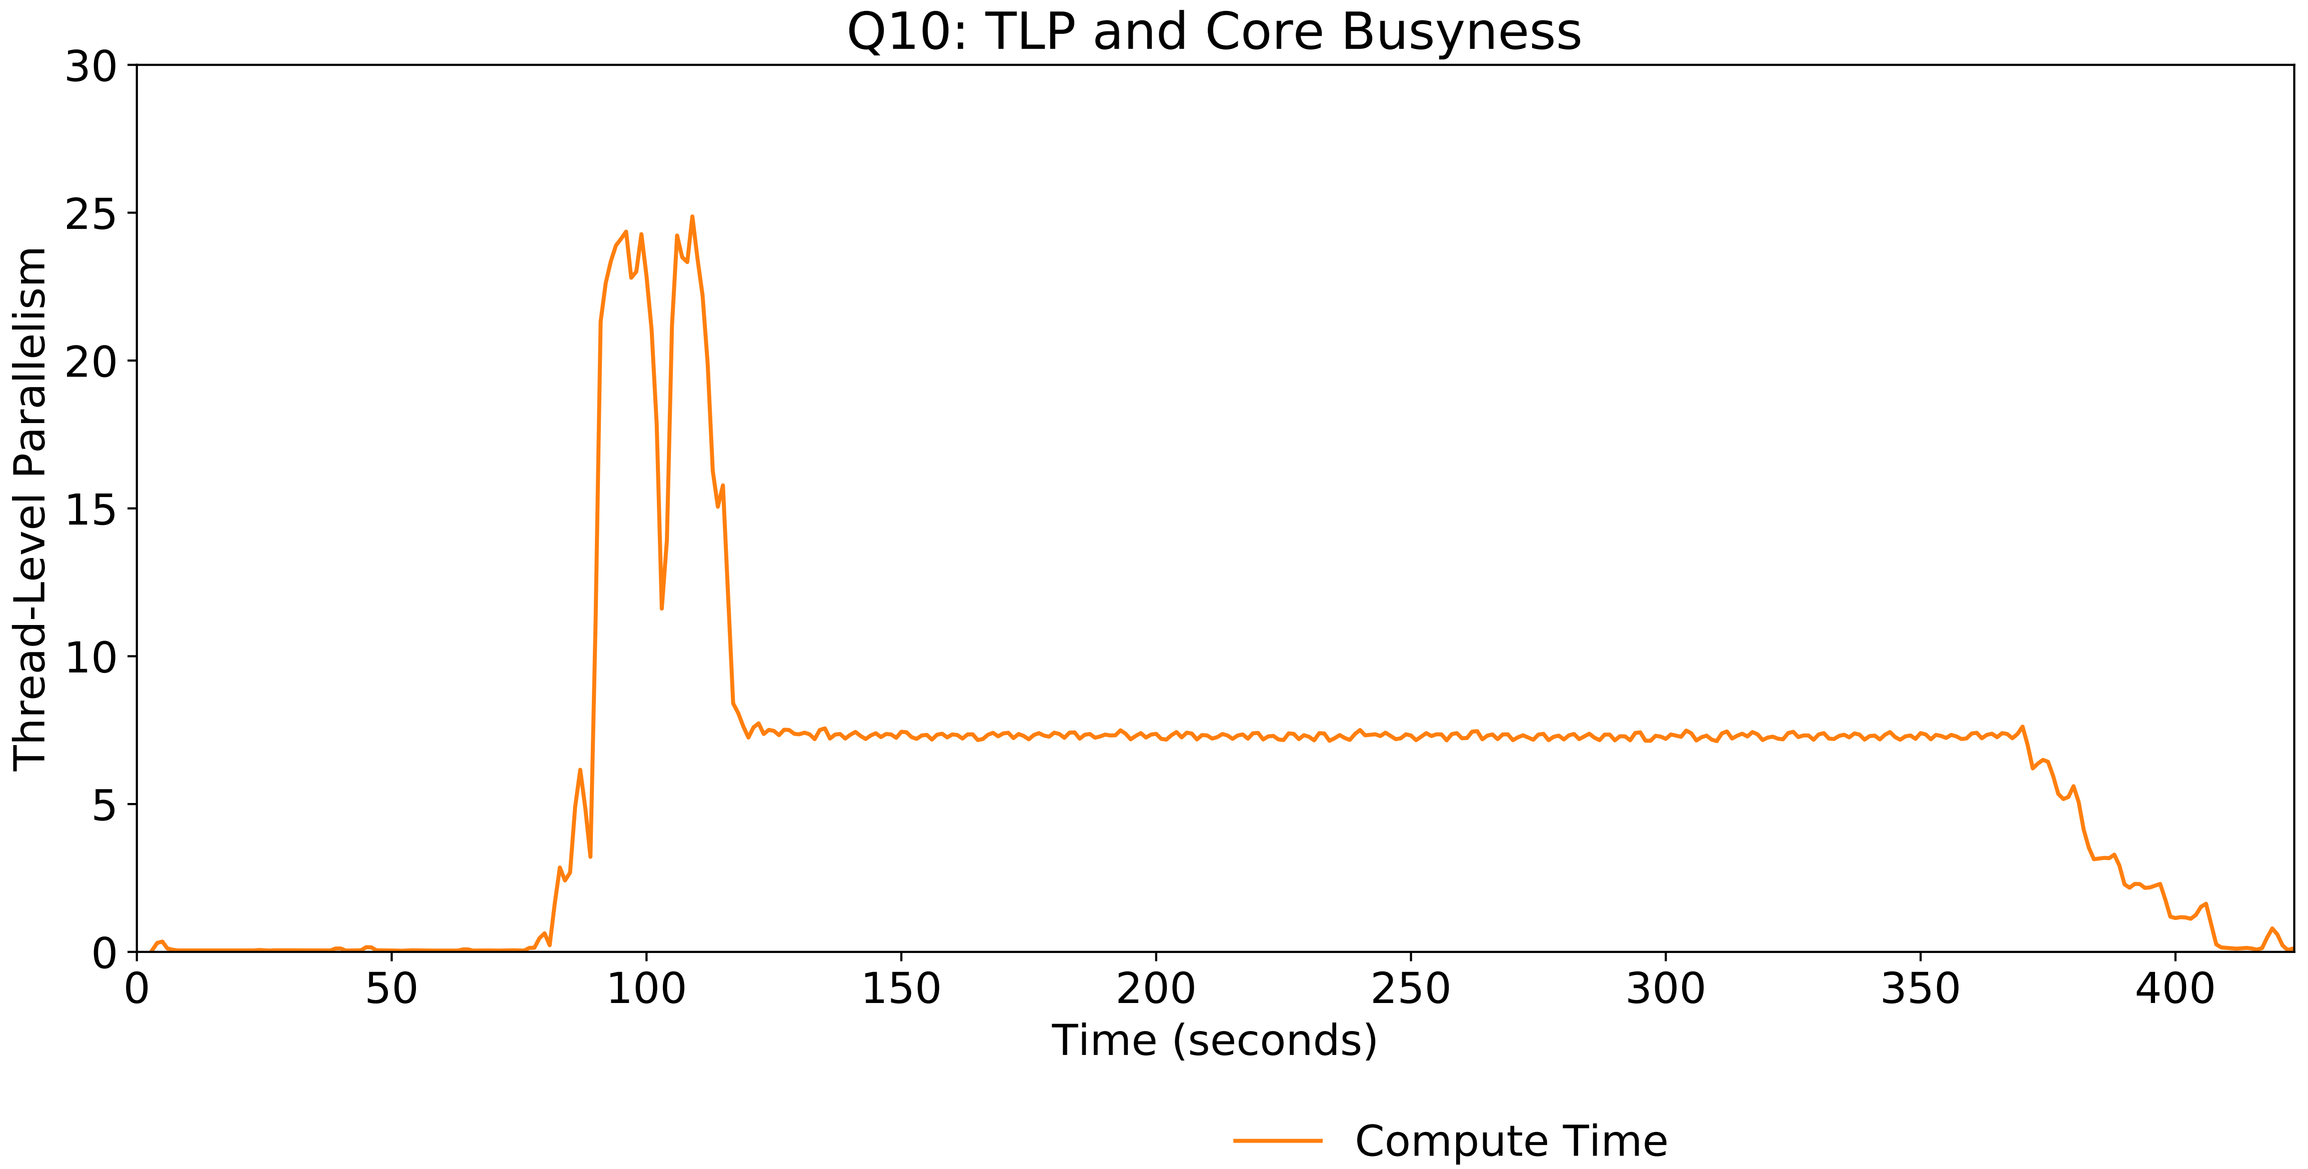

Core Packing—Software Proposal

Proactively restrict the number of active cores to just meet the workload.

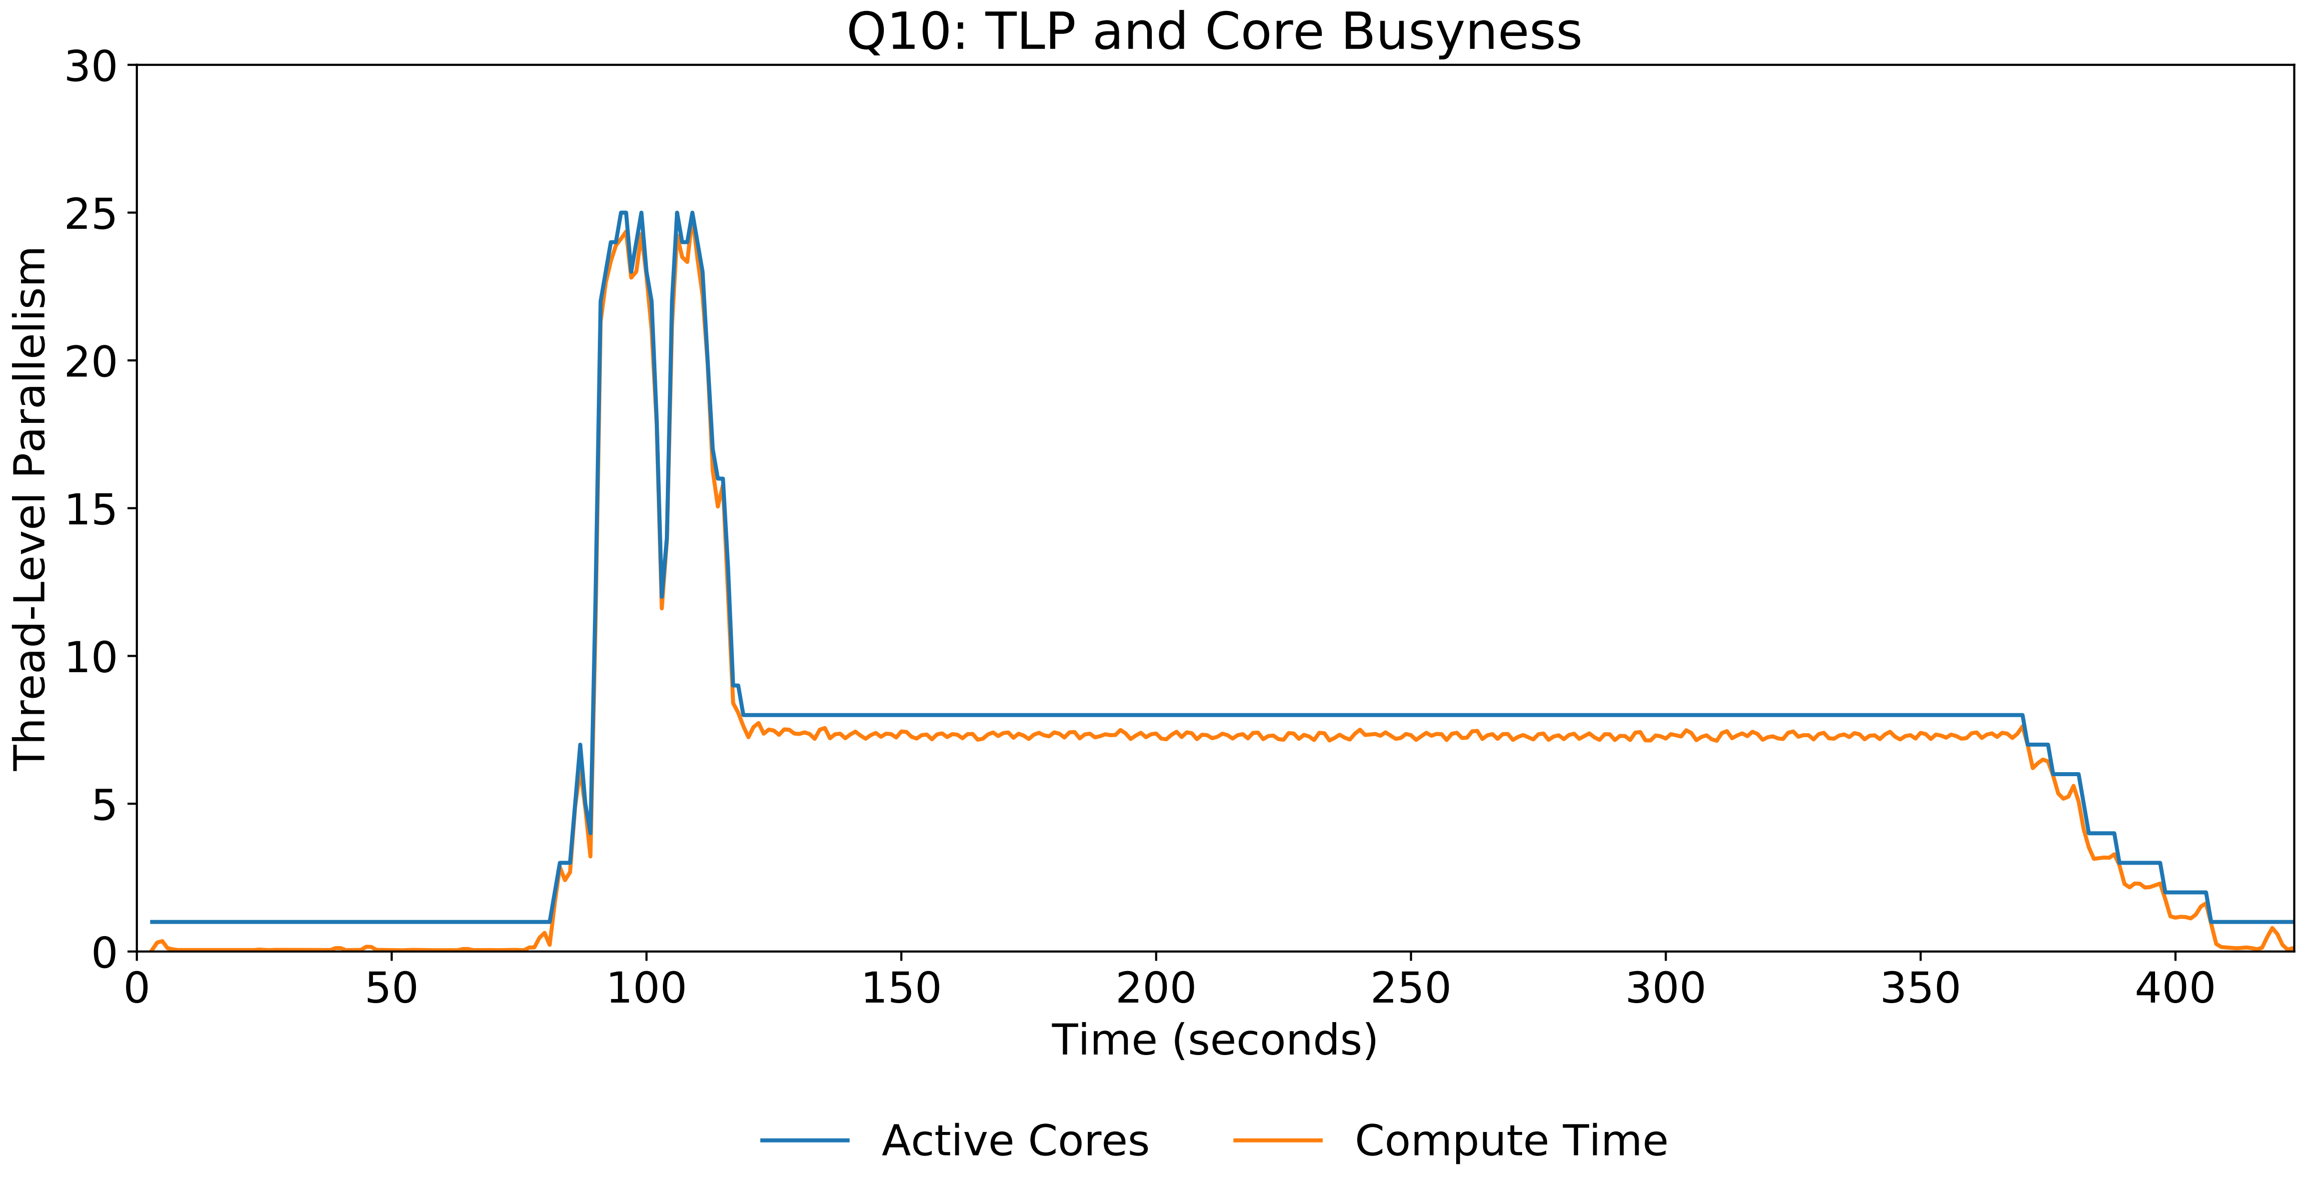

Core Packing—Approximation

Core Packing—Approximation

Q10 has such consistently low TLP, we can use it to approximate core packing behavior.

Core Packing—Approximation

Q10 has such consistently low TLP, we can use it to approximate core packing behavior.

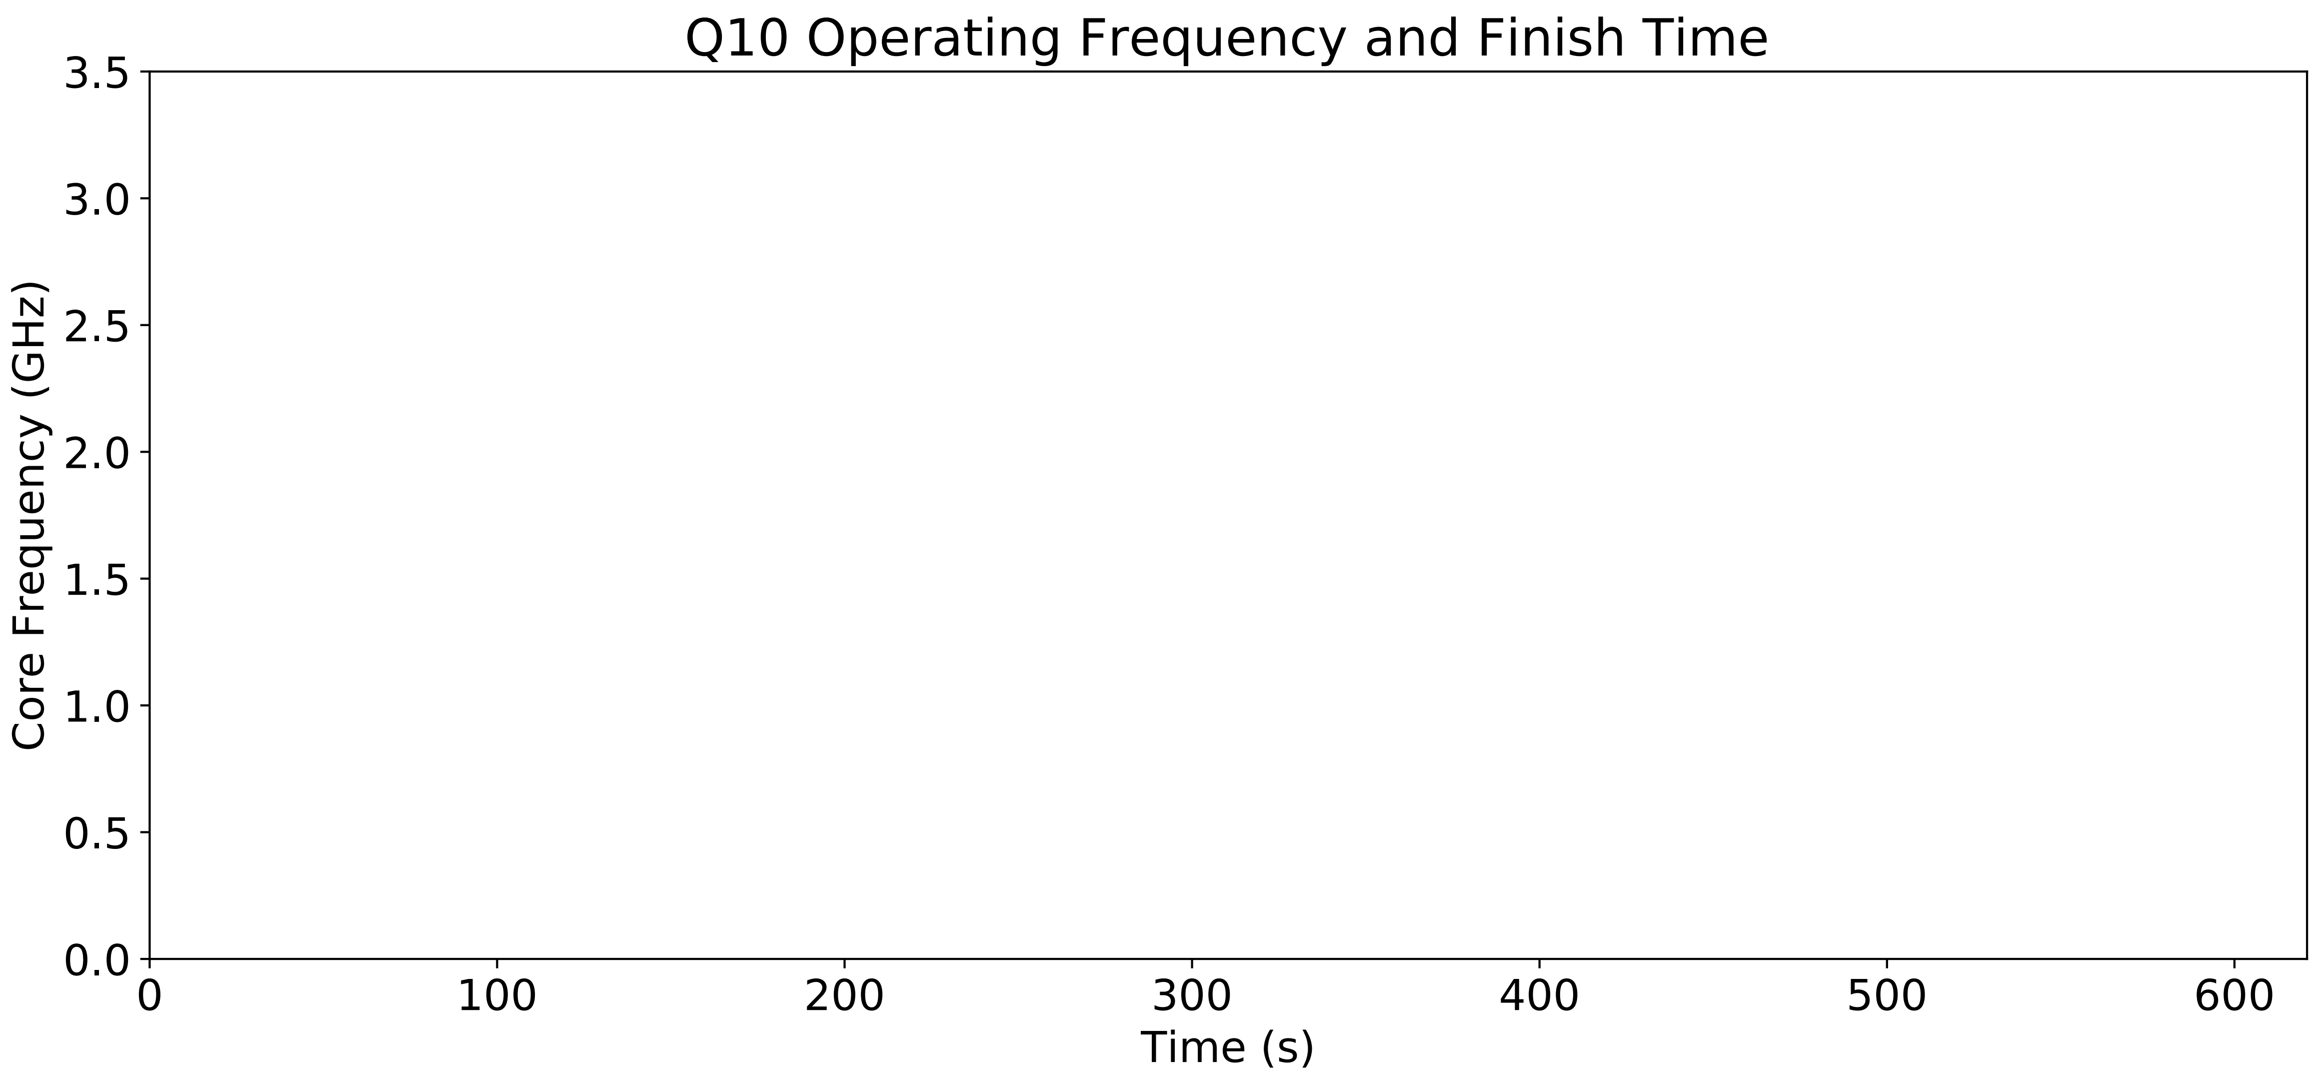

Core Packing—Approximation

Q10 has such consistently low TLP, we can use it to approximate core packing behavior.

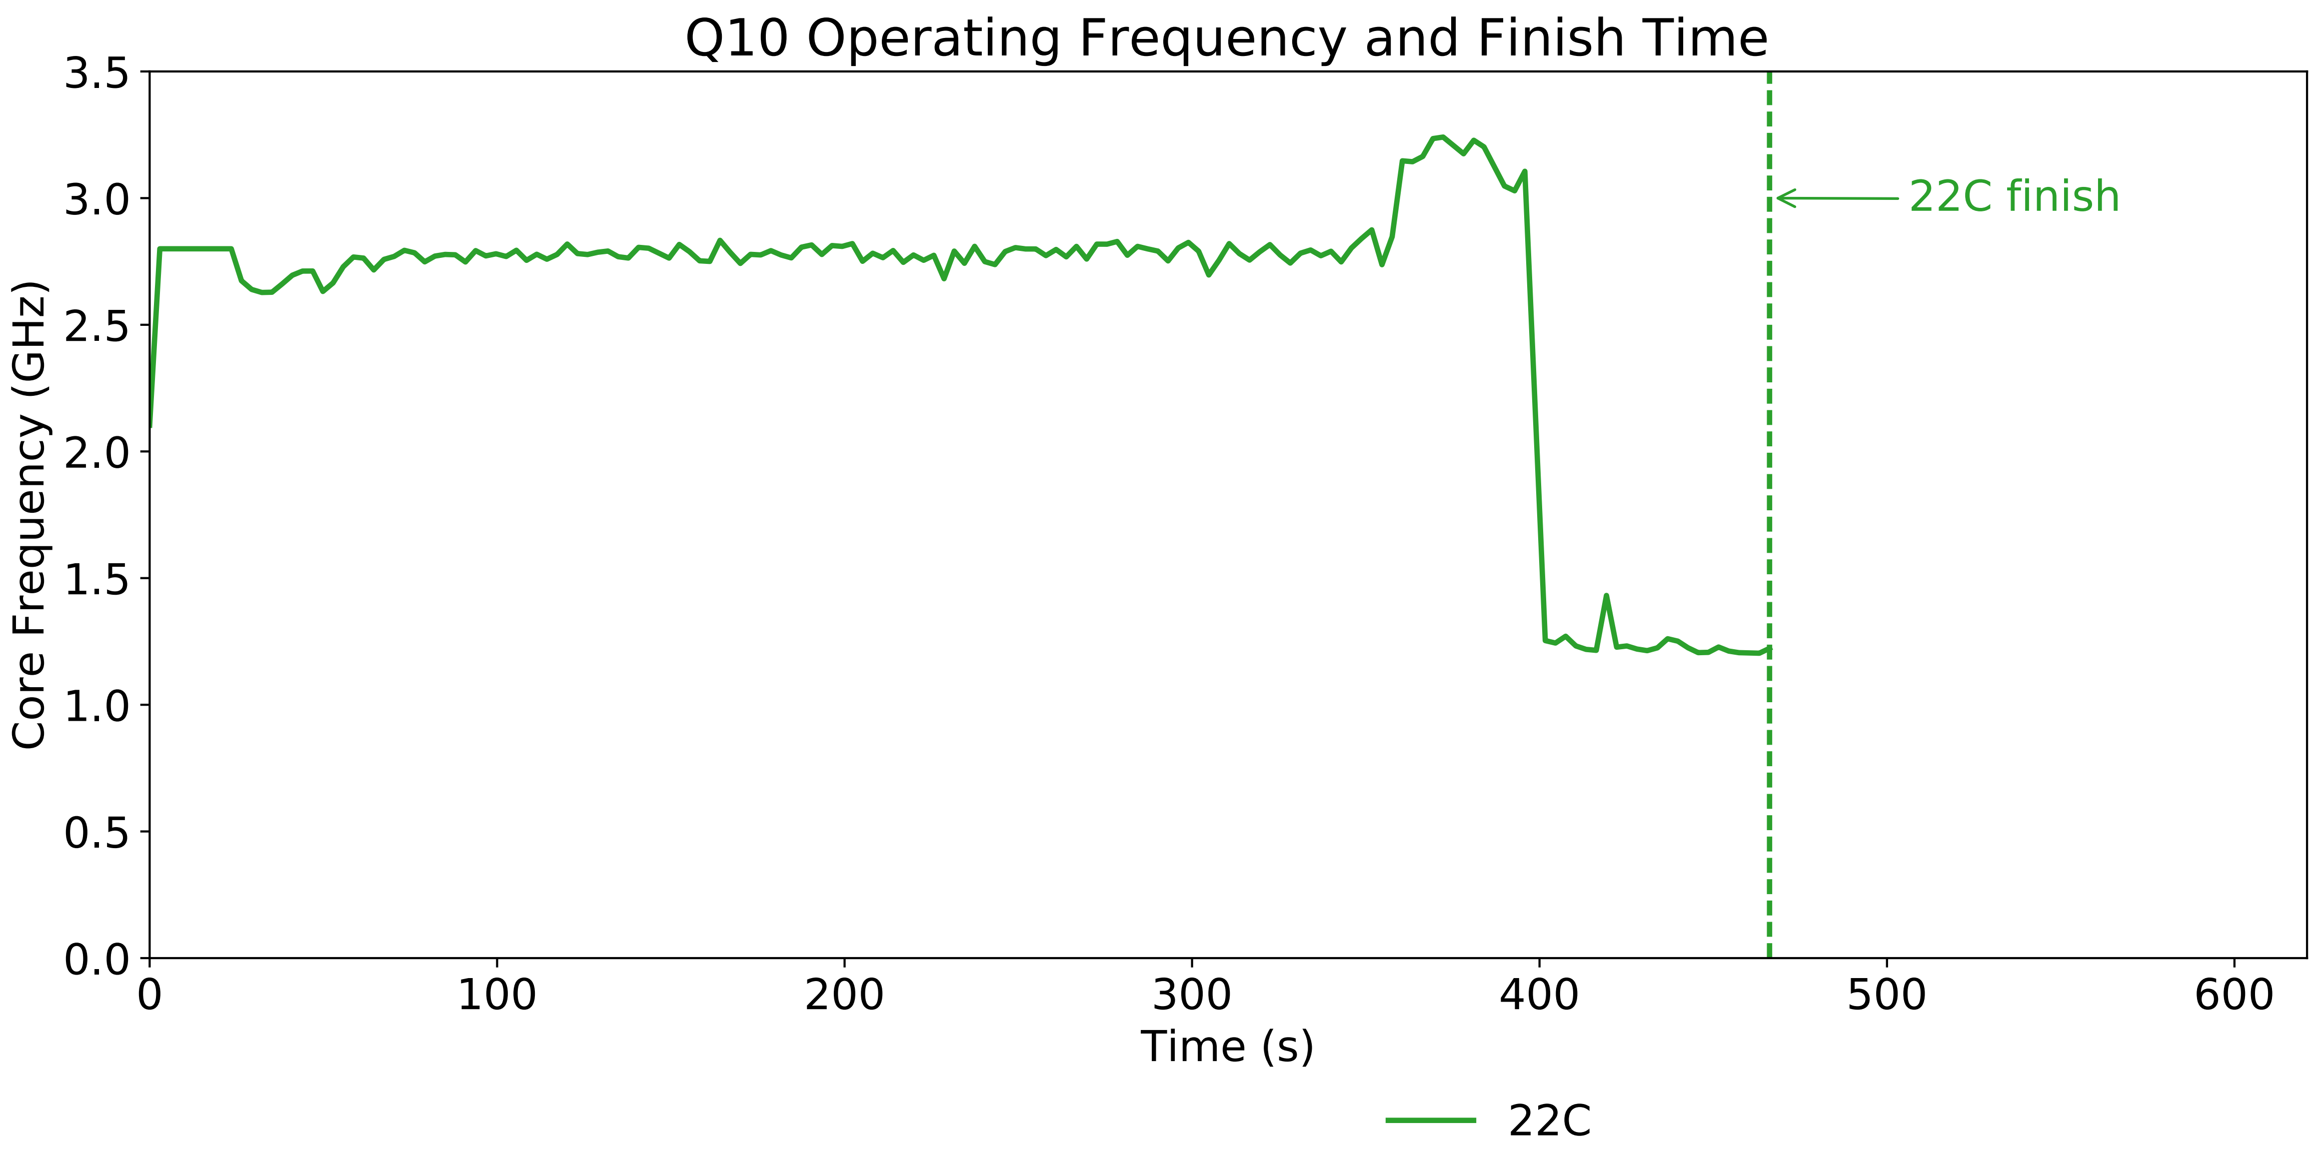

Core Packing—Approximation

Q10 has such consistently low TLP, we can use it to approximate core packing behavior.

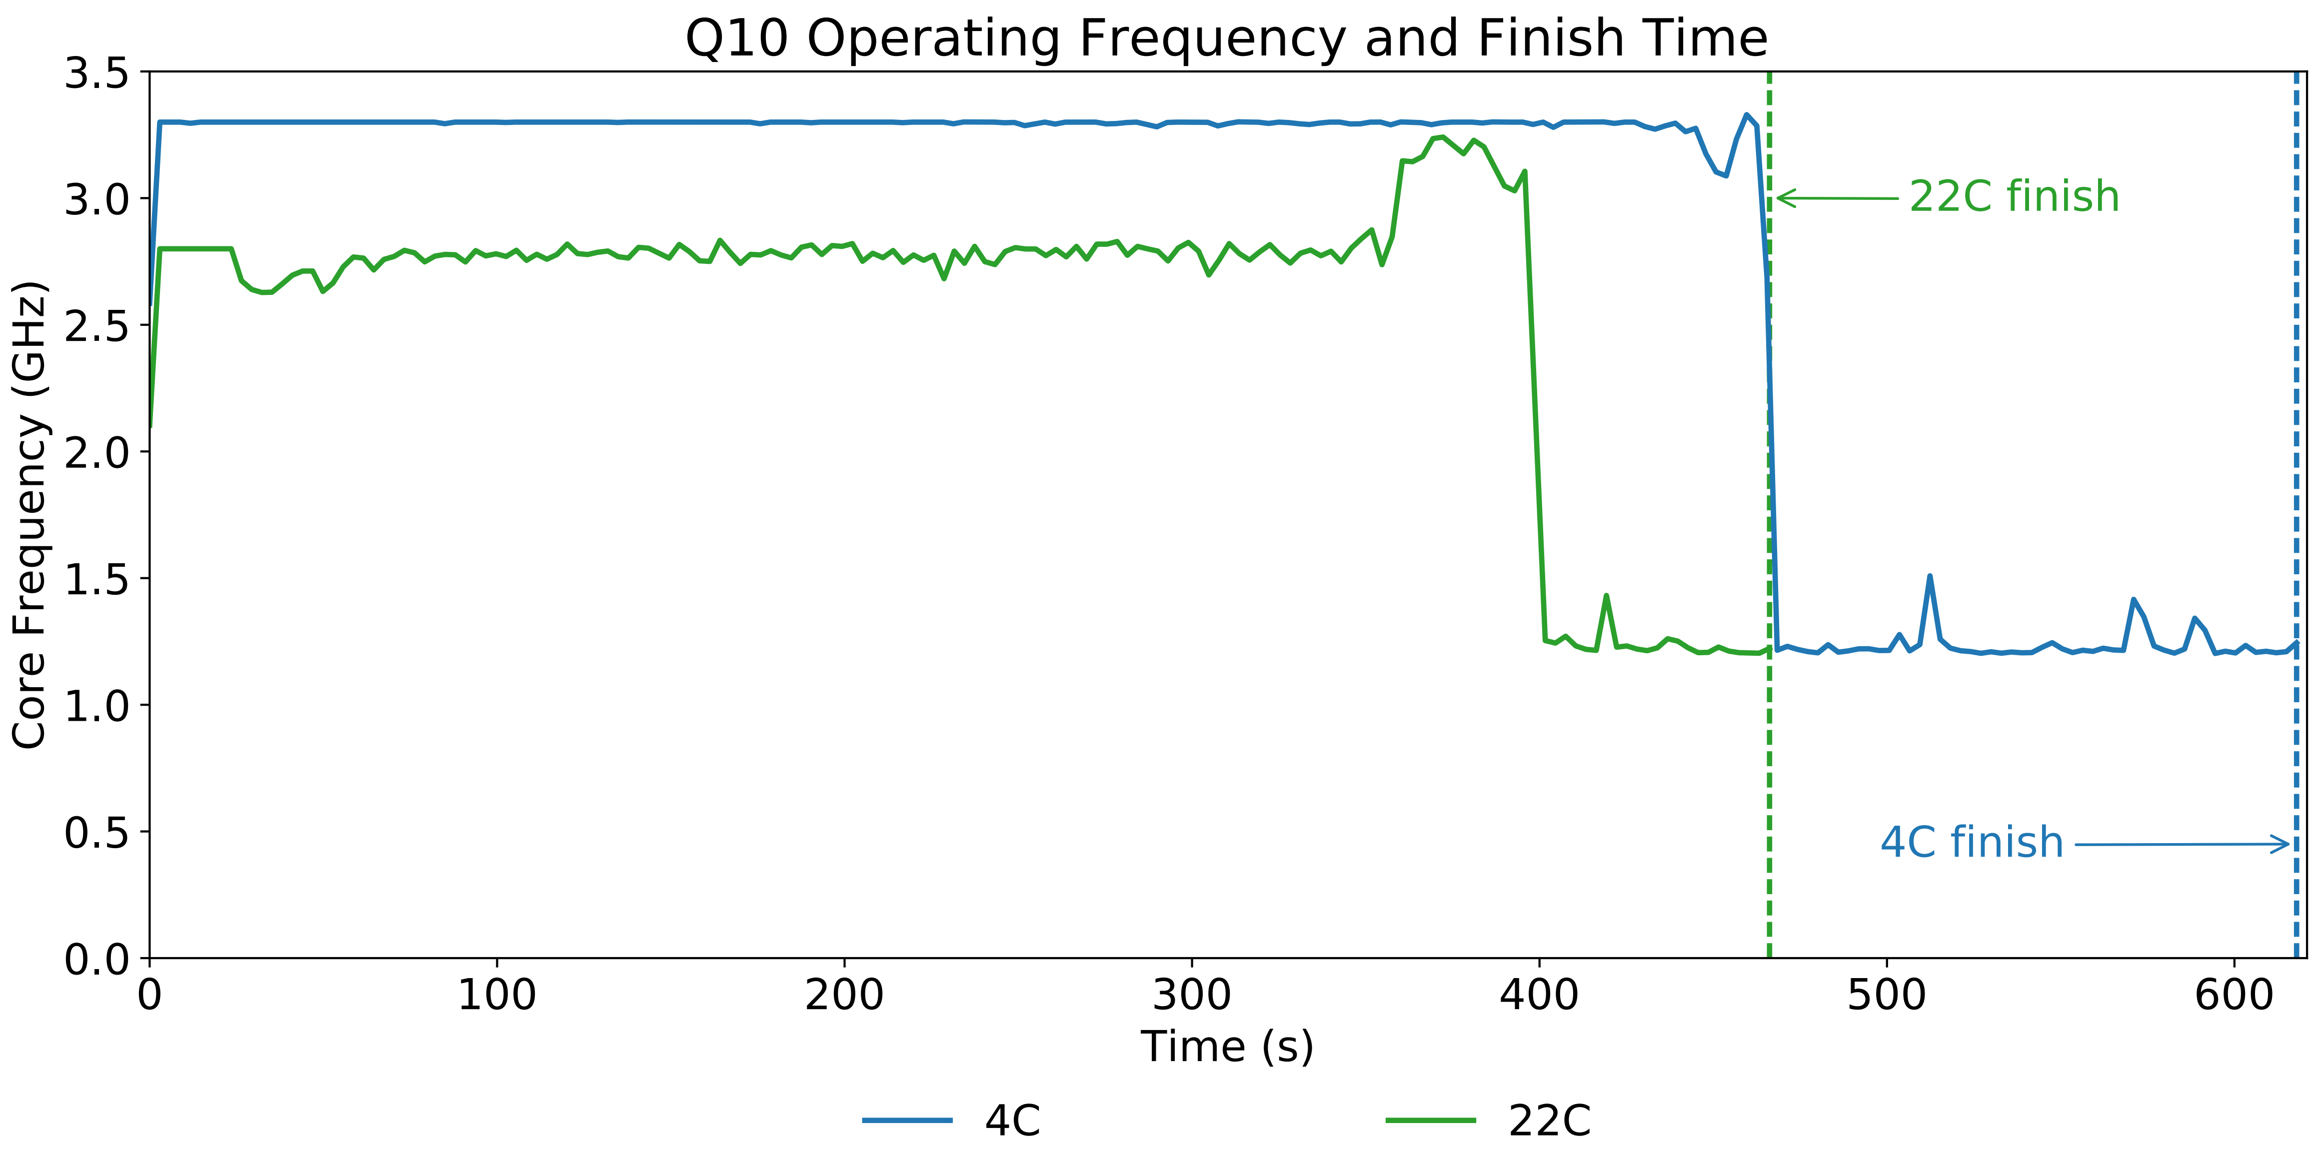

Core Packing—Approximation

Q10 has such consistently low TLP, we can use it to approximate core packing behavior.

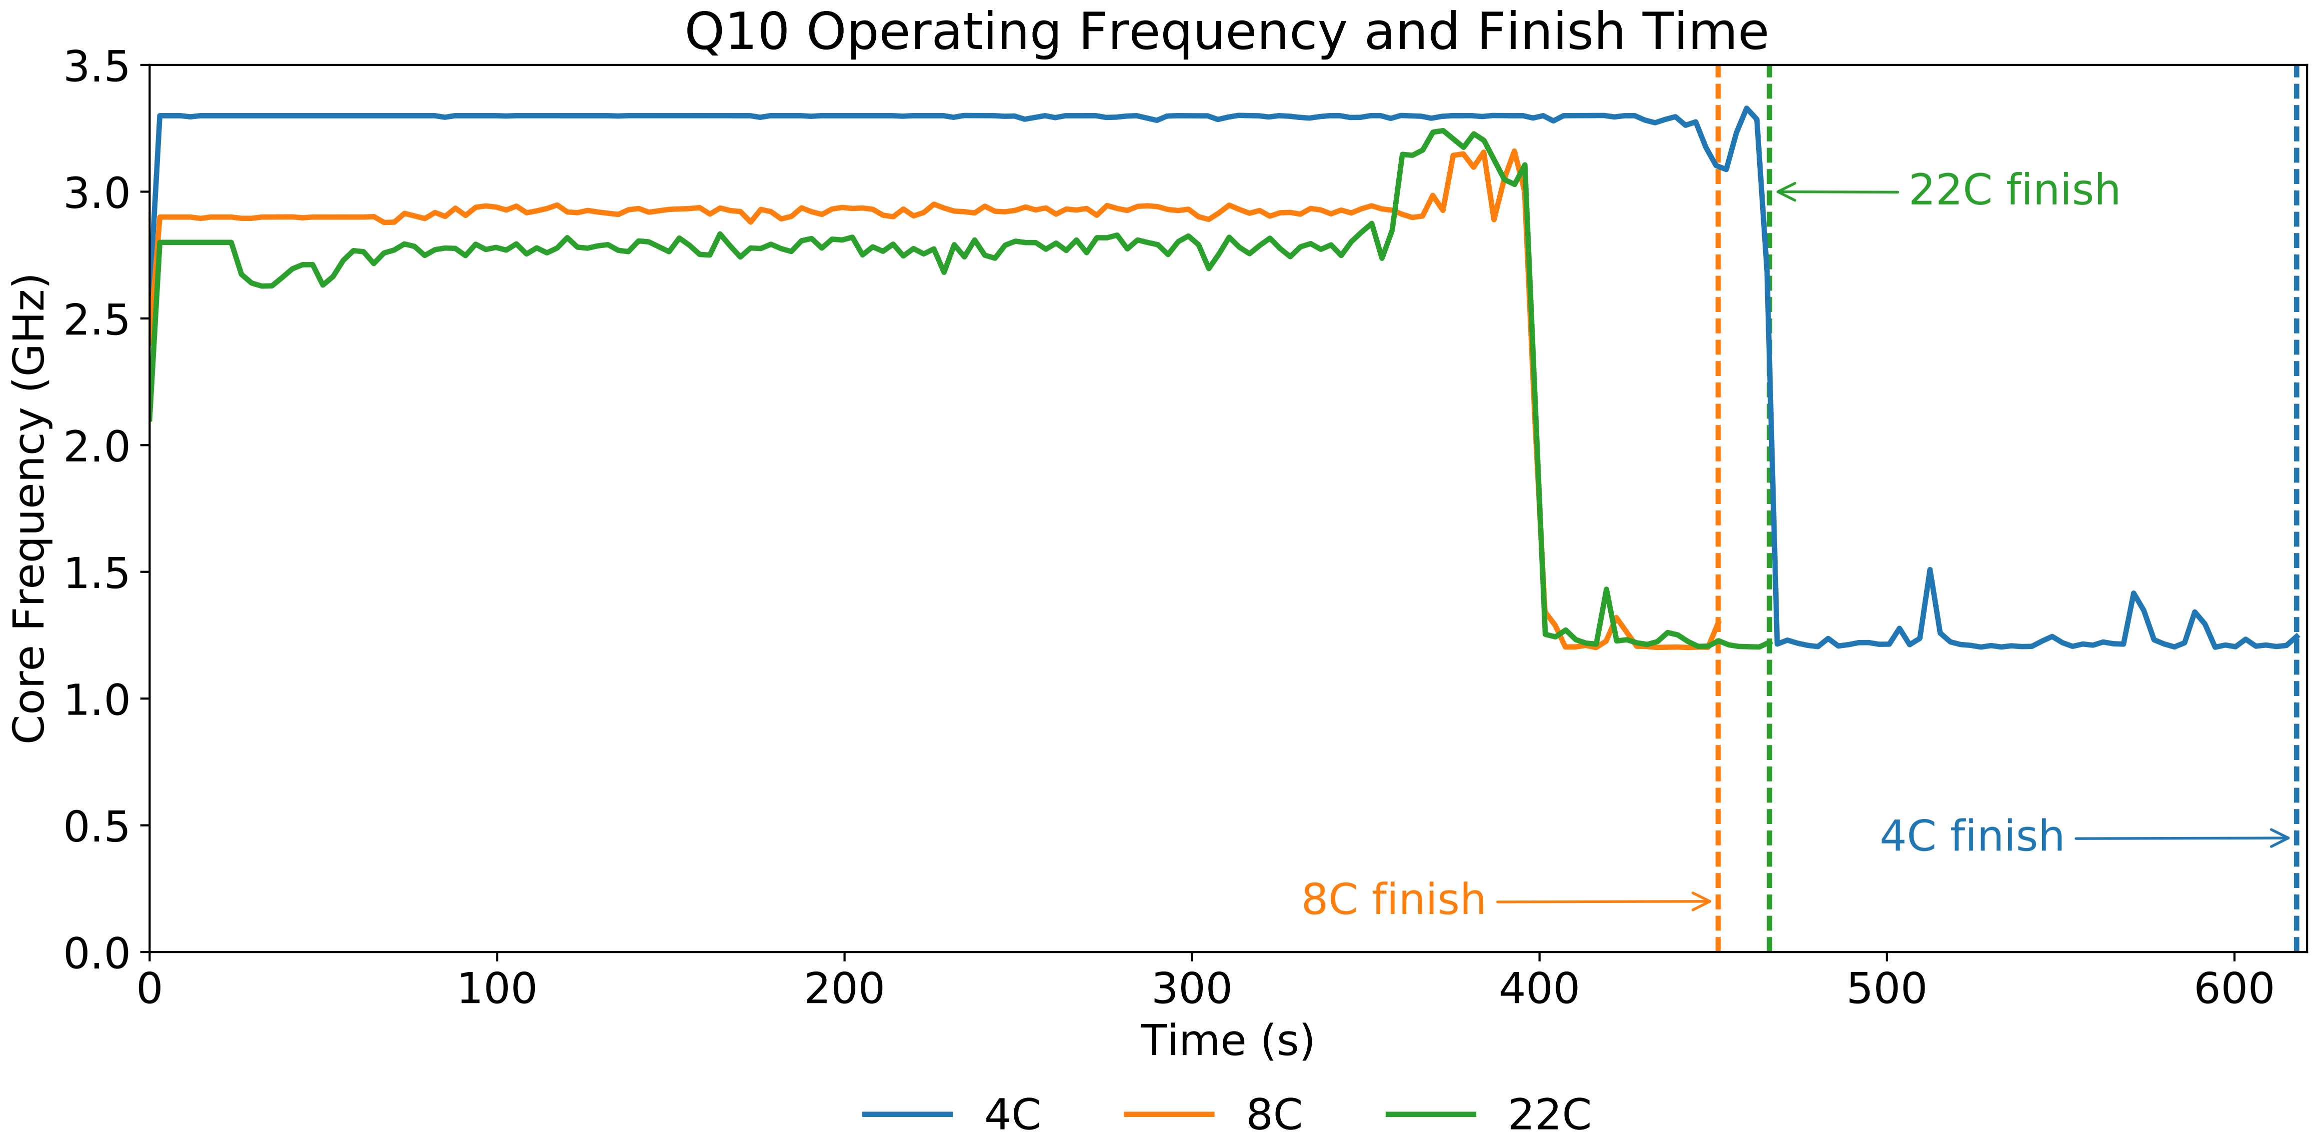

Core Packing—Better Approximation

Core Packing—Better Approximation

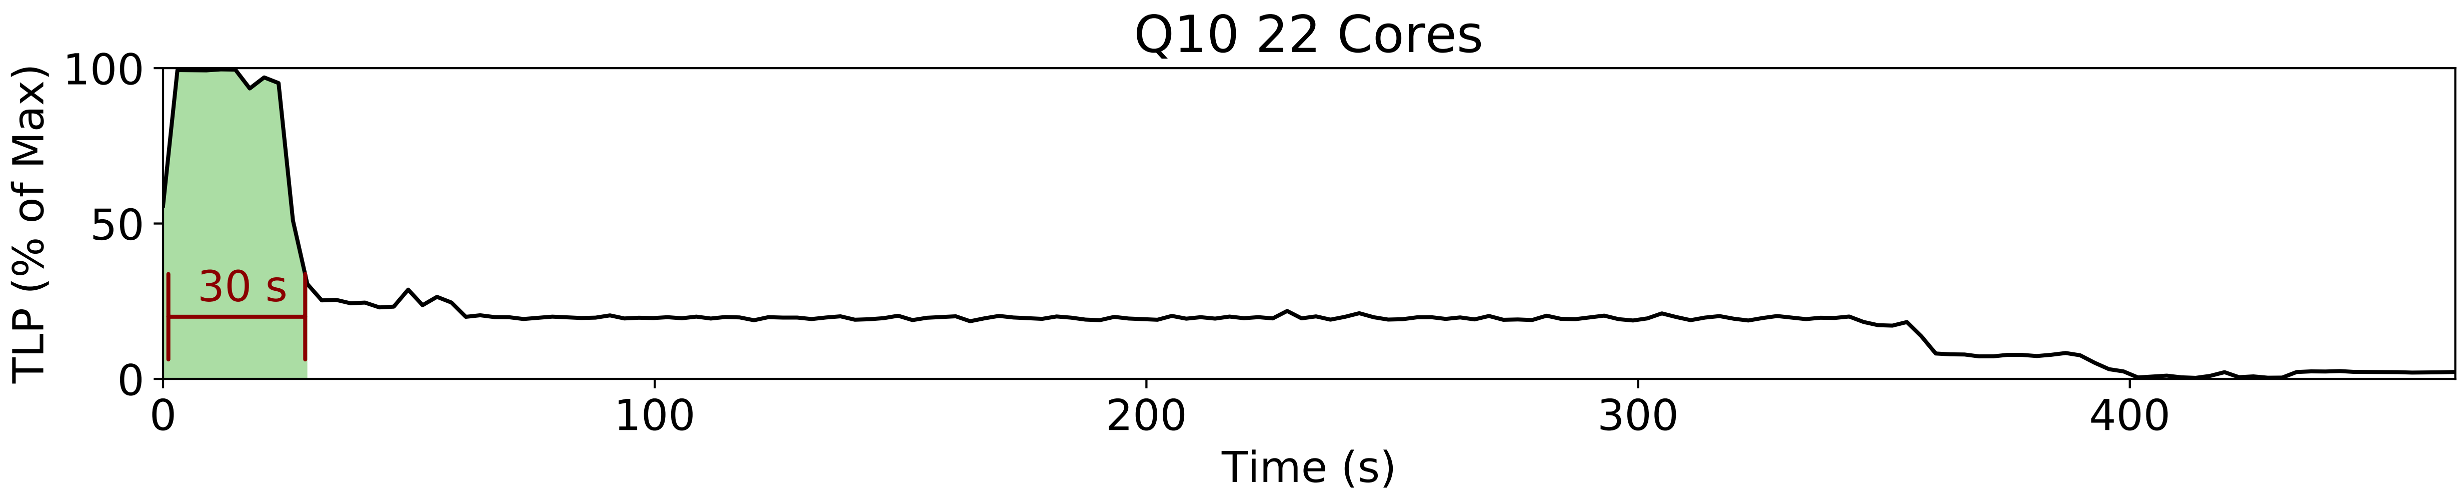

Core Packing—Better Approximation

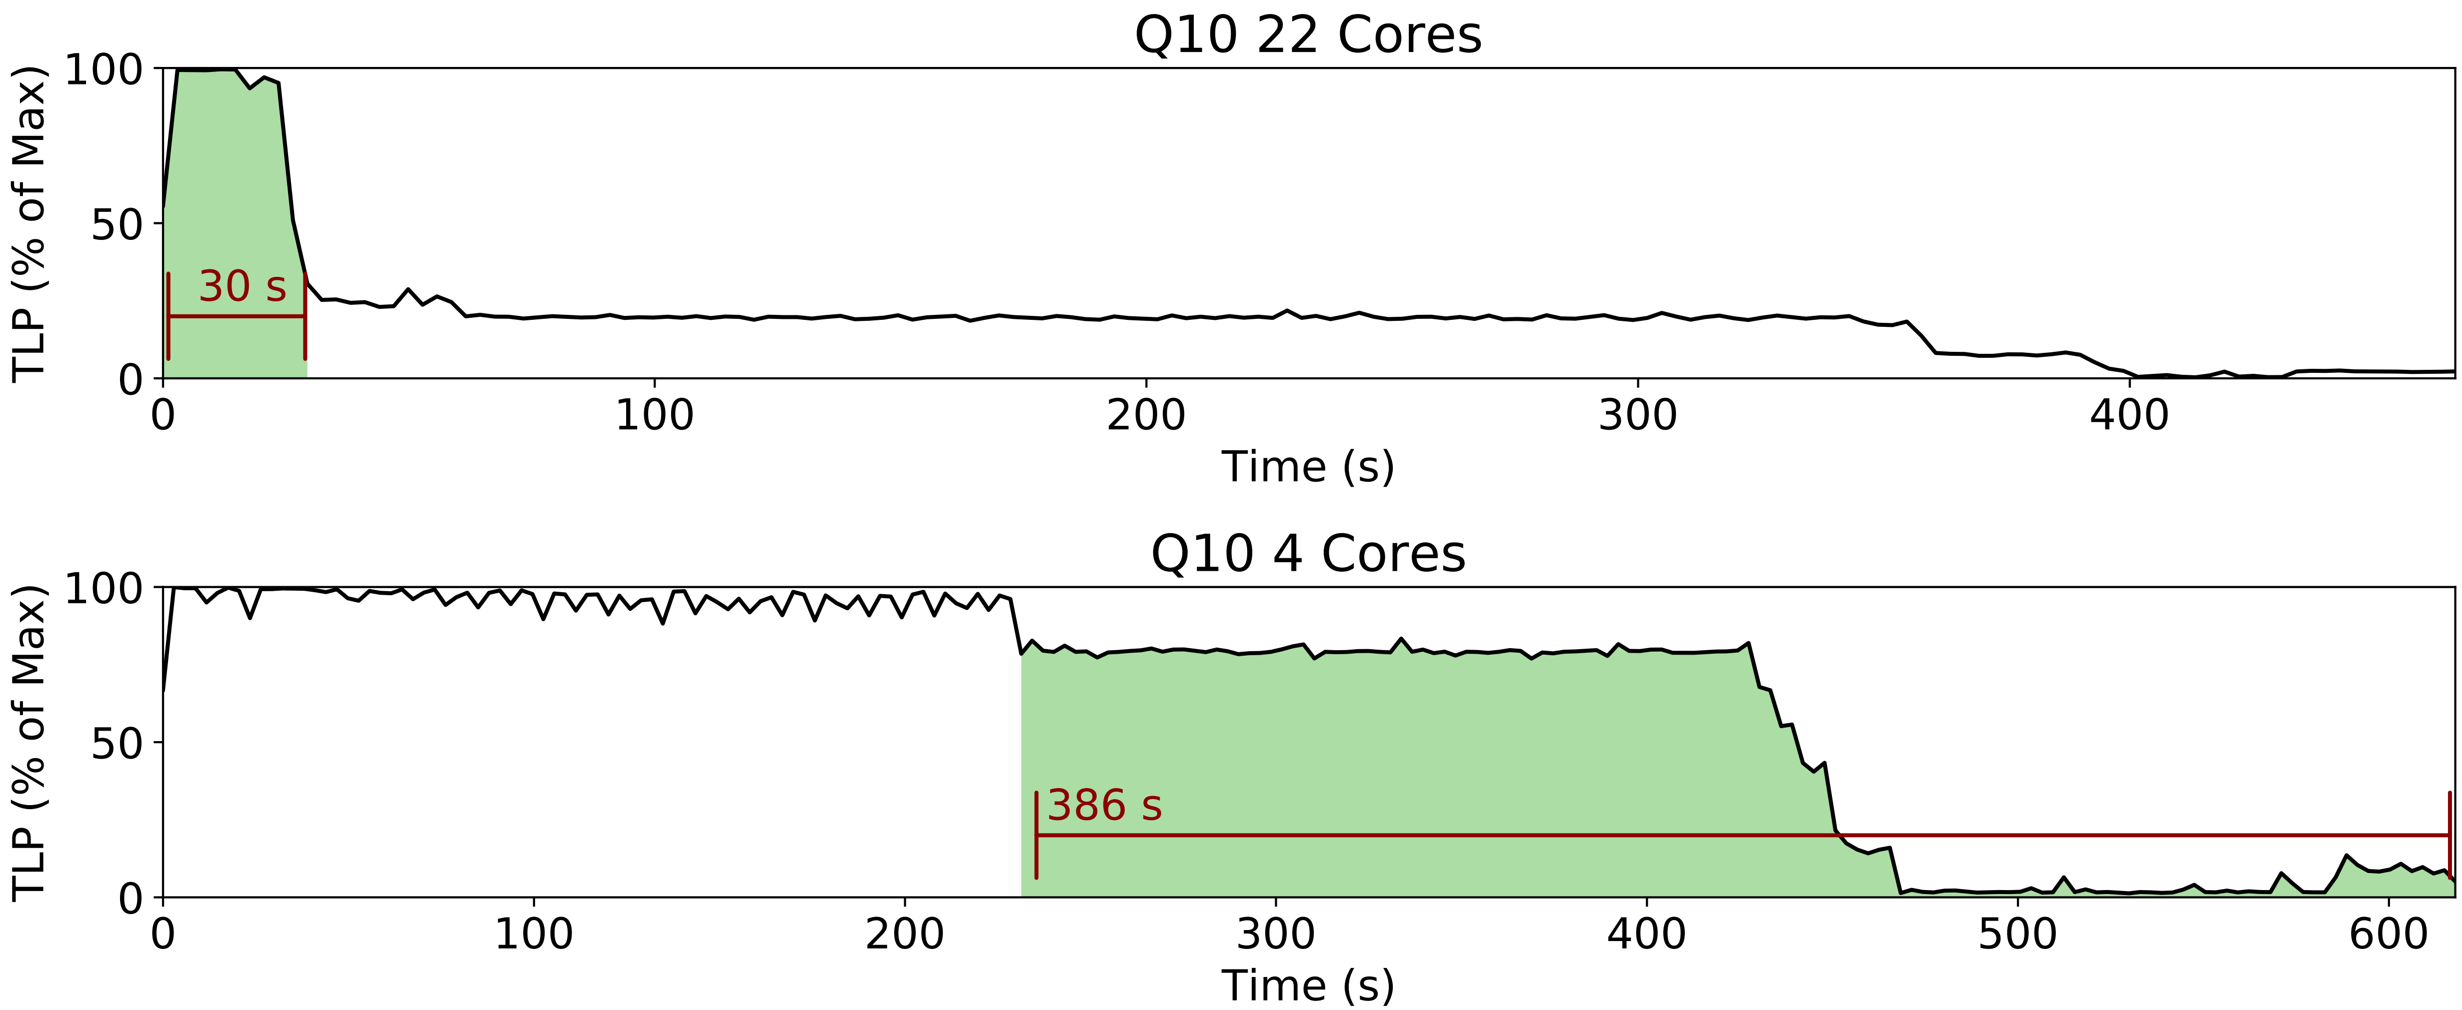

Core Packing—Better Approximation

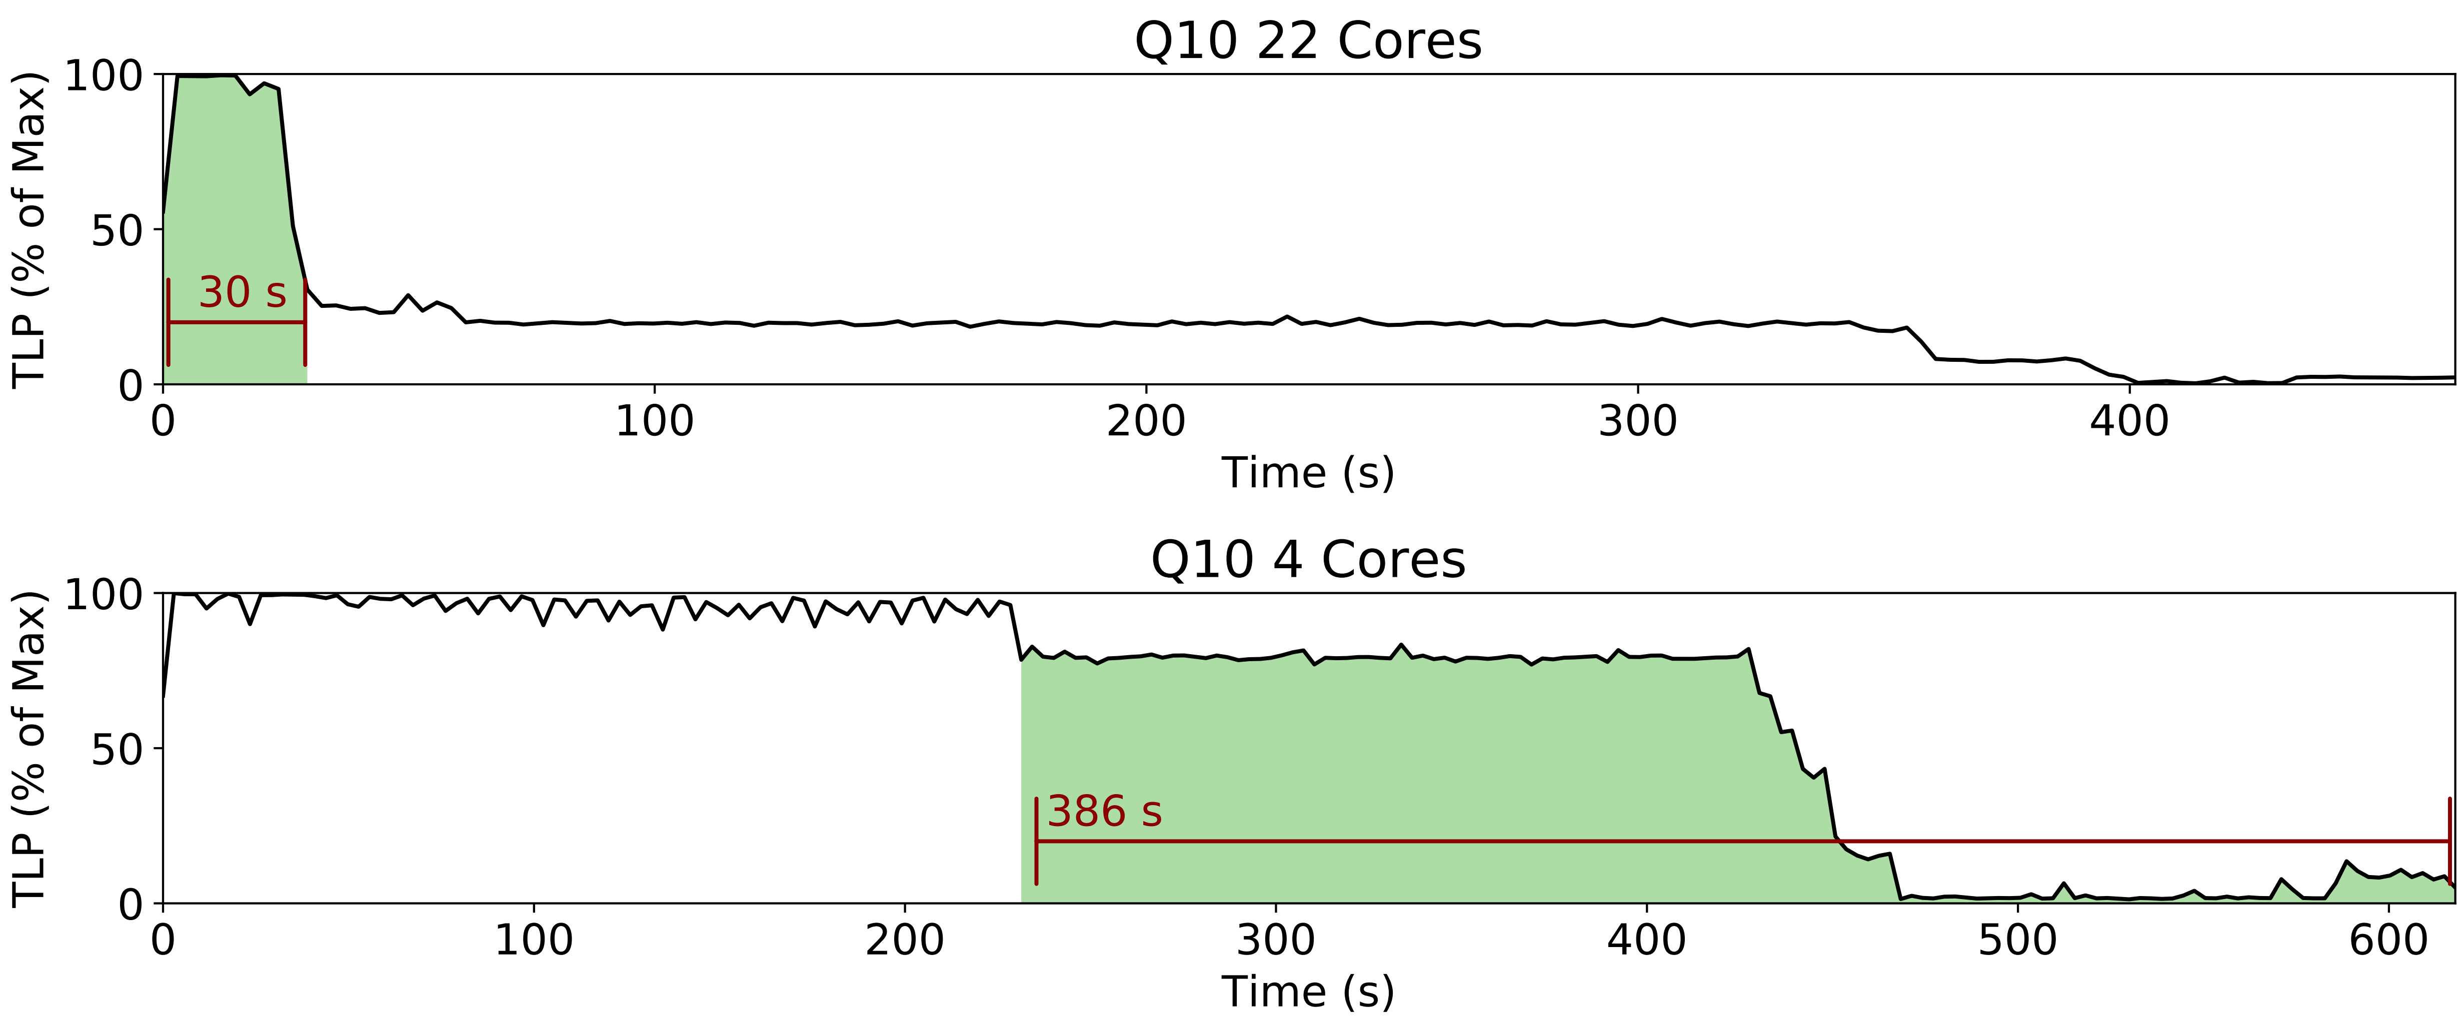

Query 10 could finish in 416 seconds: a 13.5% speedup.

Core Packing—Better Approximation

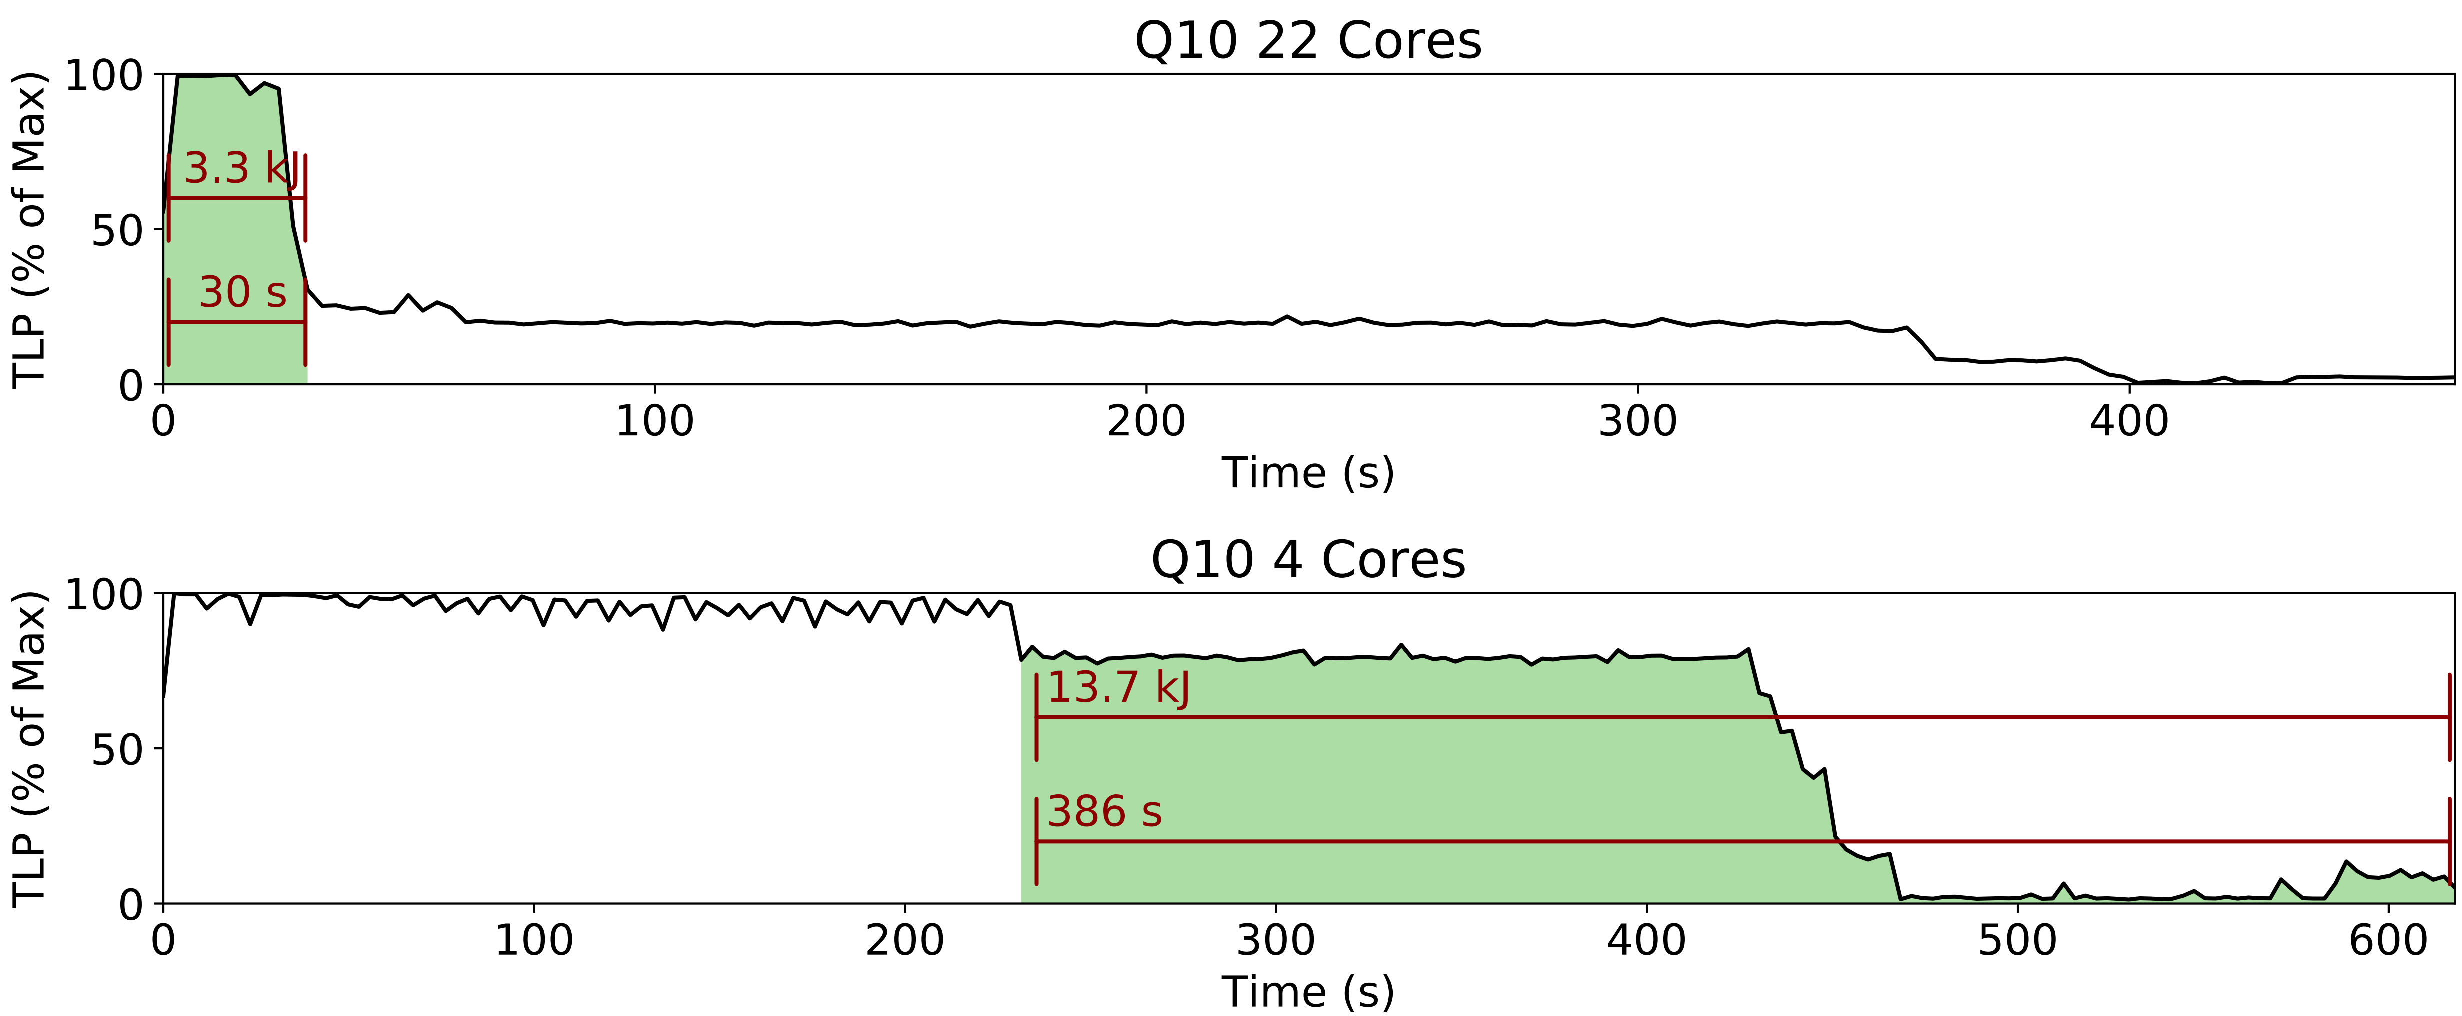

Query 10 could finish in 416 seconds: a 13.5% speedup.

Furthermore, energy consumption decreases by 30%.

Analysis

Analysis

Most queries are far too complex to enable such simple core packing approximation.

Analysis

Most queries are far too complex to enable such simple core packing approximation.

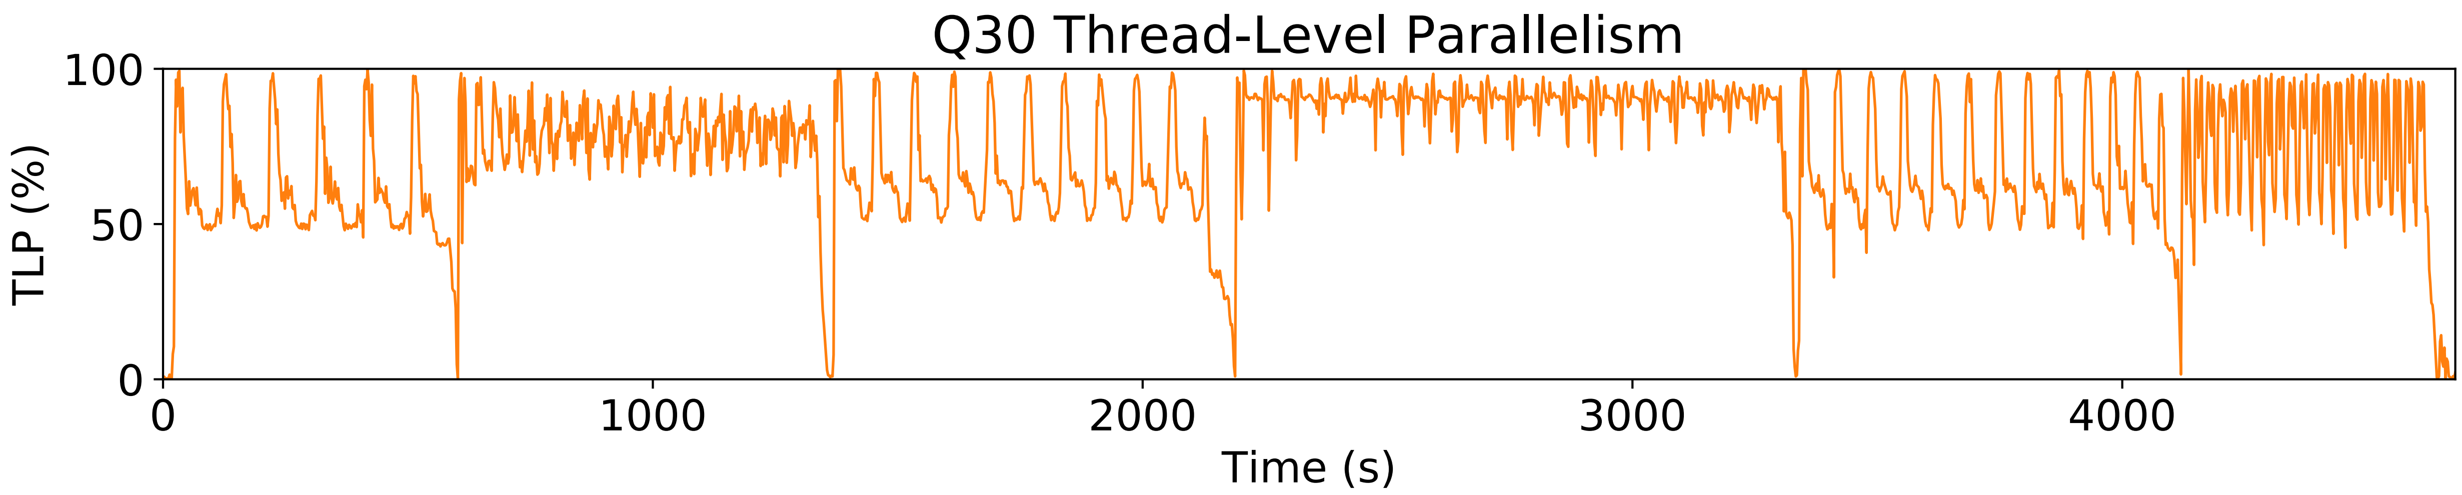

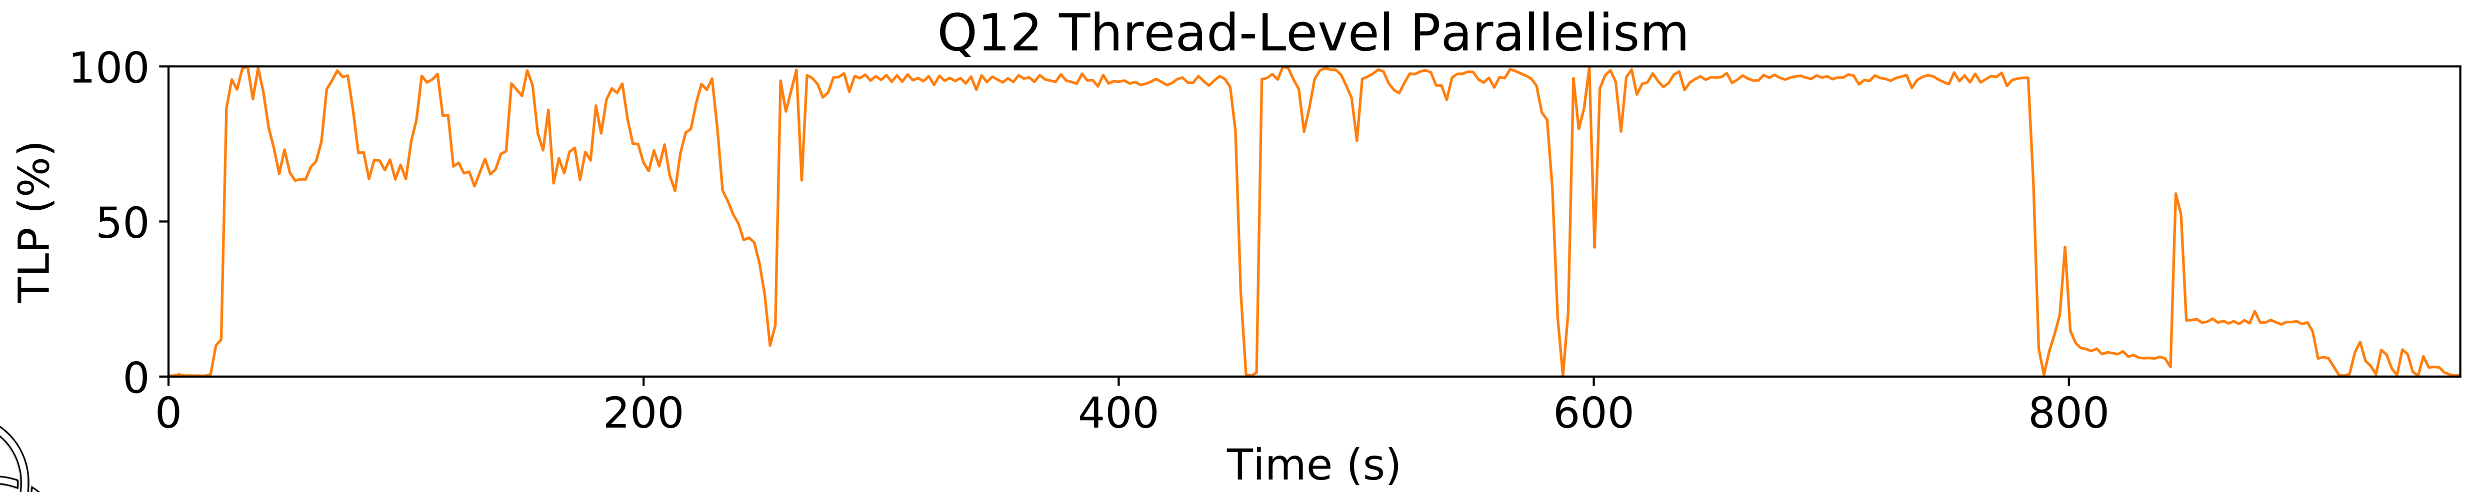

Analytical Model—Turbo Boost Frequency

Analytical Model—Turbo Boost Frequency

Assuming a linear increase in Turbo Boost ceiling for each deactivated core, at what frequency can each query run?

Analytical Model—Turbo Boost Frequency

Assuming a linear increase in Turbo Boost ceiling for each deactivated core, at what frequency can each query run?

$$freq_{TB} = freq_0 + (1 - TLP) \cdot (freq_{Max} - freq_0)$$

Analytical Model—Turbo Boost Frequency

Assuming a linear increase in Turbo Boost ceiling for each deactivated core, at what frequency can each query run?

$$freq_{TB} = \boxed{freq_0} + (1 - TLP) \cdot (freq_{Max} - freq_0)$$

$\boxed{freq_0}$ Base (current) frequency

Analytical Model—Turbo Boost Frequency

Assuming a linear increase in Turbo Boost ceiling for each deactivated core, at what frequency can each query run?

$$freq_{TB} = freq_0 + (1 - TLP) \cdot (freq_{Max} - freq_0)$$

$freq_0$ Base (current) frequency

$(1 - TLP)$ Available slack in cores

Analytical Model—Turbo Boost Frequency

Assuming a linear increase in Turbo Boost ceiling for each deactivated core, at what frequency can each query run?

$$freq_{TB} = freq_0 + (1 - TLP) \cdot (freq_{Max} - freq_0)$$

$freq_0$	Base (current) frequency
$(1 - TLP)$	Available slack in cores
$(freq_{Max} - freq_0)$	Turbo Boost ceiling room

Analytical Model—Packed Cores

Analytical Model—Packed Cores

Given a potentially higher operating frequency, how much speedup can we expect?

Analytical Model—Packed Cores

Given a potentially higher operating frequency, how much speedup can we expect?

$$speedup = \frac{freq_{TB} - freq_0}{freq_0} \cdot efficiency_{freq}$$

Analytical Model—Packed Cores

Given a potentially higher operating frequency, how much speedup can we expect?

$$speedup = \boxed{\frac{freq_{TB} - freq_0}{freq_0}} \cdot efficiency_{freq}$$

$\boxed{\dfrac{freq_{TB} - freq_0}{freq_0}}$ Fractional increase in frequency

Analytical Model—Packed Cores

Given a potentially higher operating frequency, how much speedup can we expect?

$$speedup = \boxed{\frac{freq_{TB} - freq_0}{freq_0}} \cdot \boxed{efficiency_{freq}}$$

$\boxed{\frac{freq_{TB} - freq_0}{freq_0}}$ Fractional increase in frequency

$\boxed{efficiency_{freq}}$ Efficiency of frequency scaling

Analytical Model—Predictions

Analytical Model—Predictions

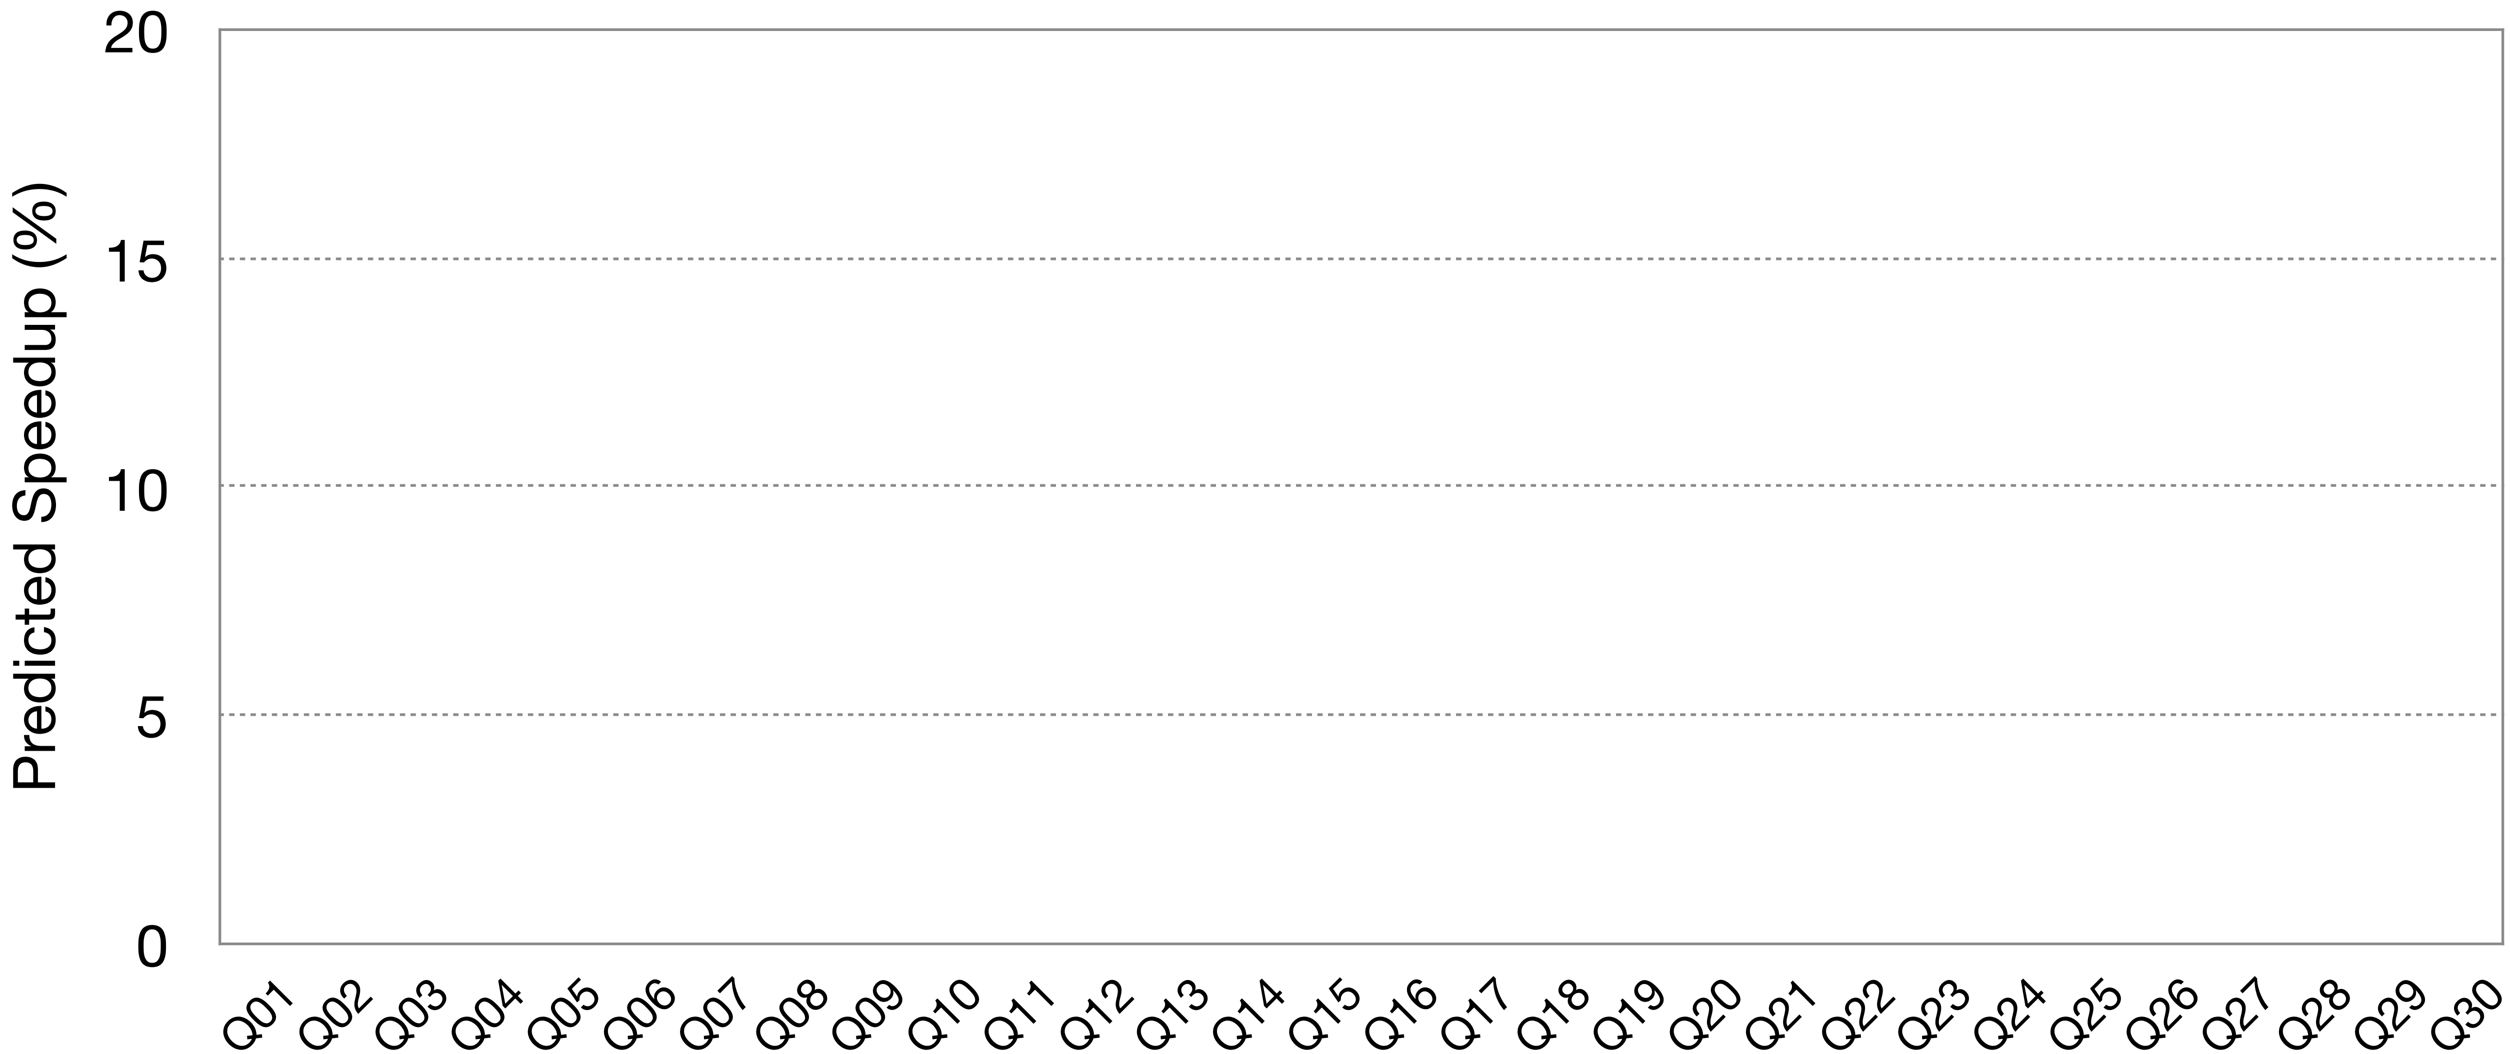

Analytical Model—Predictions

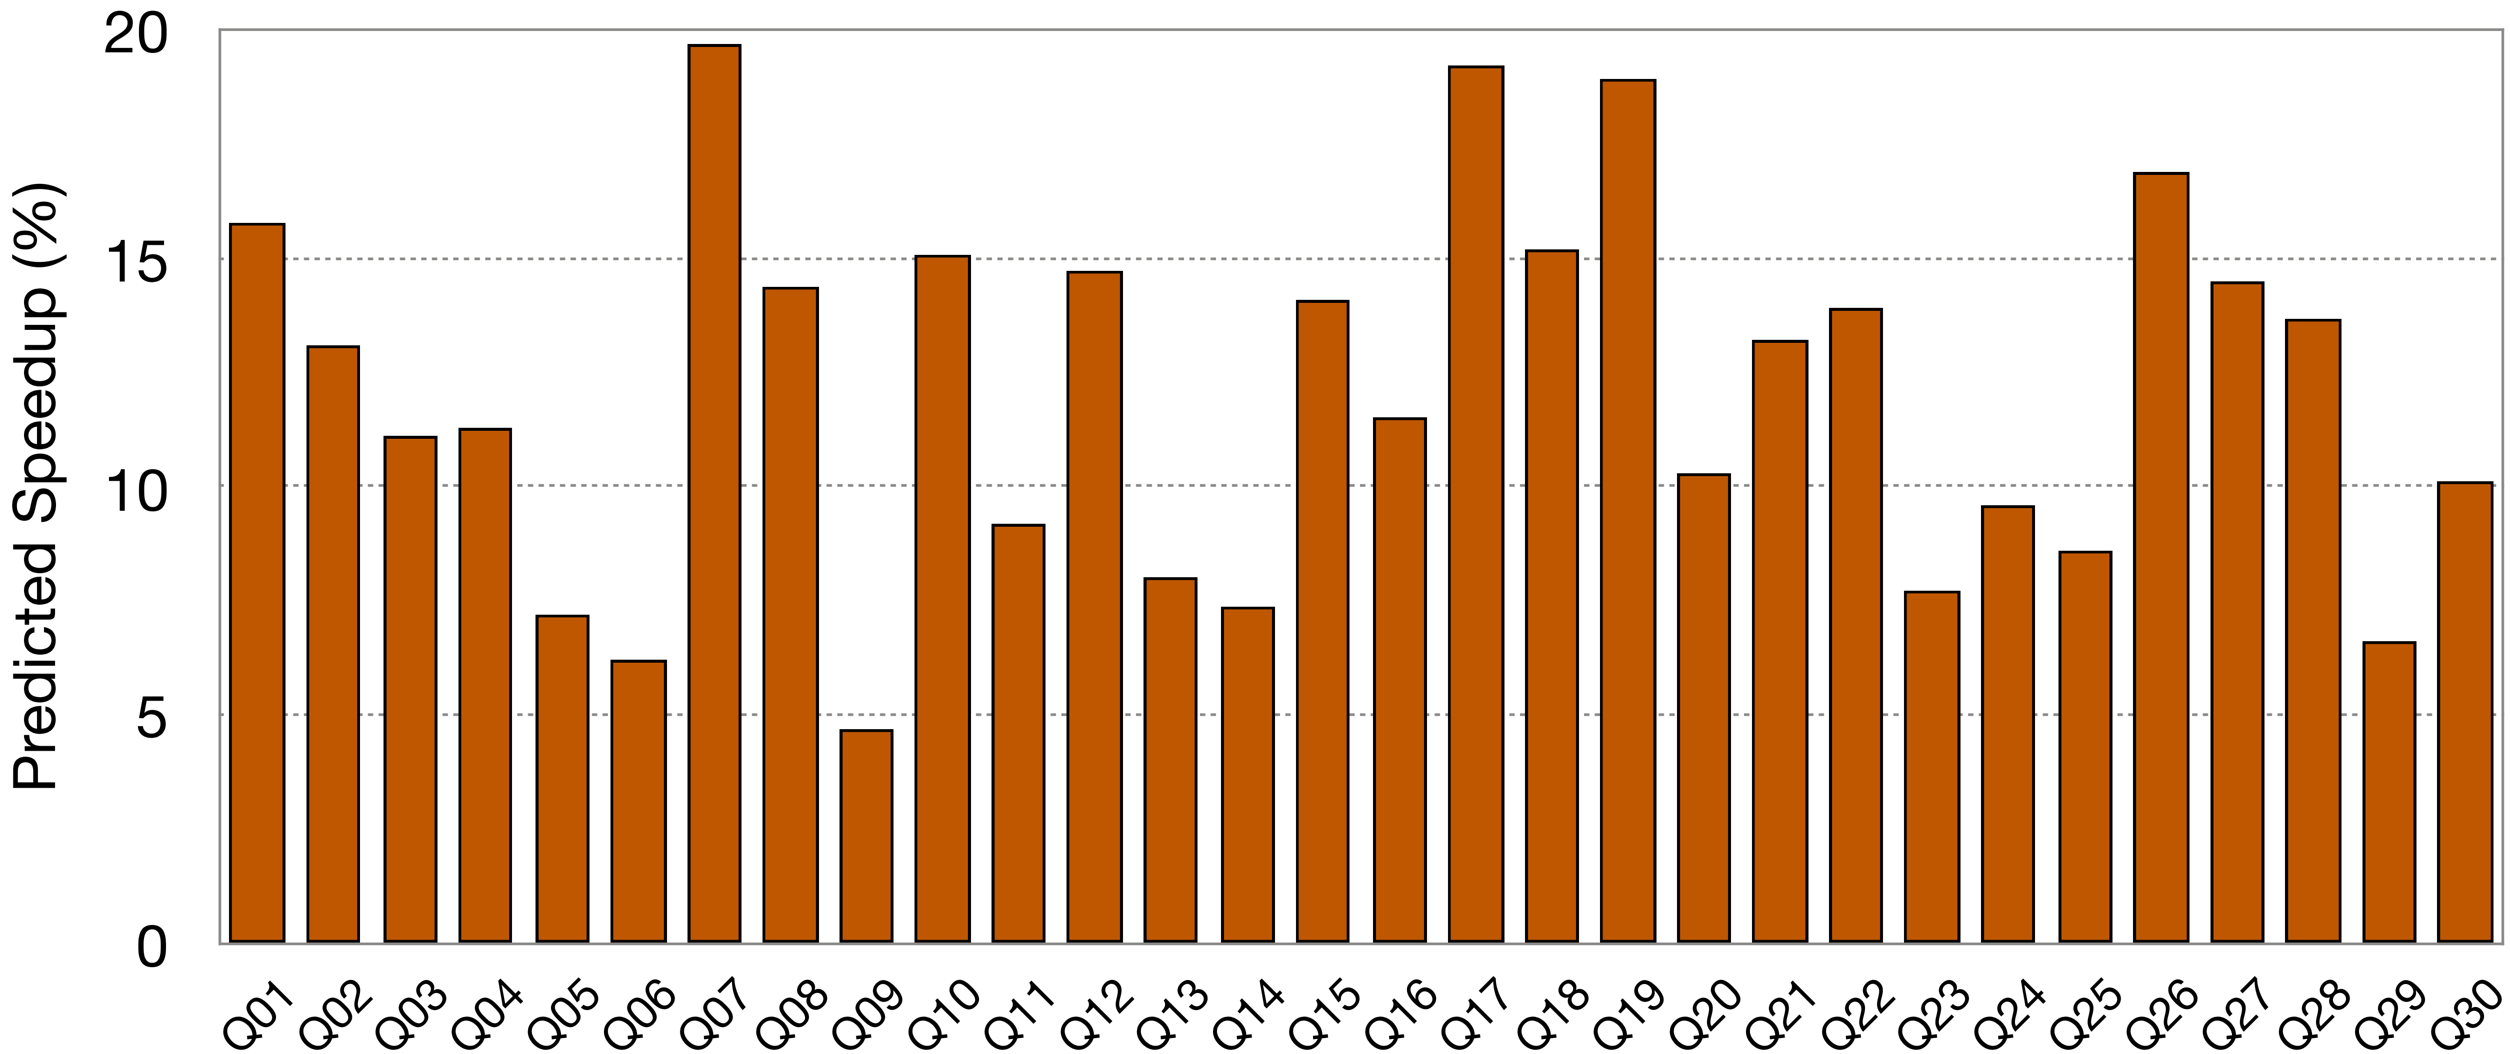

Analytical Model—Predictions

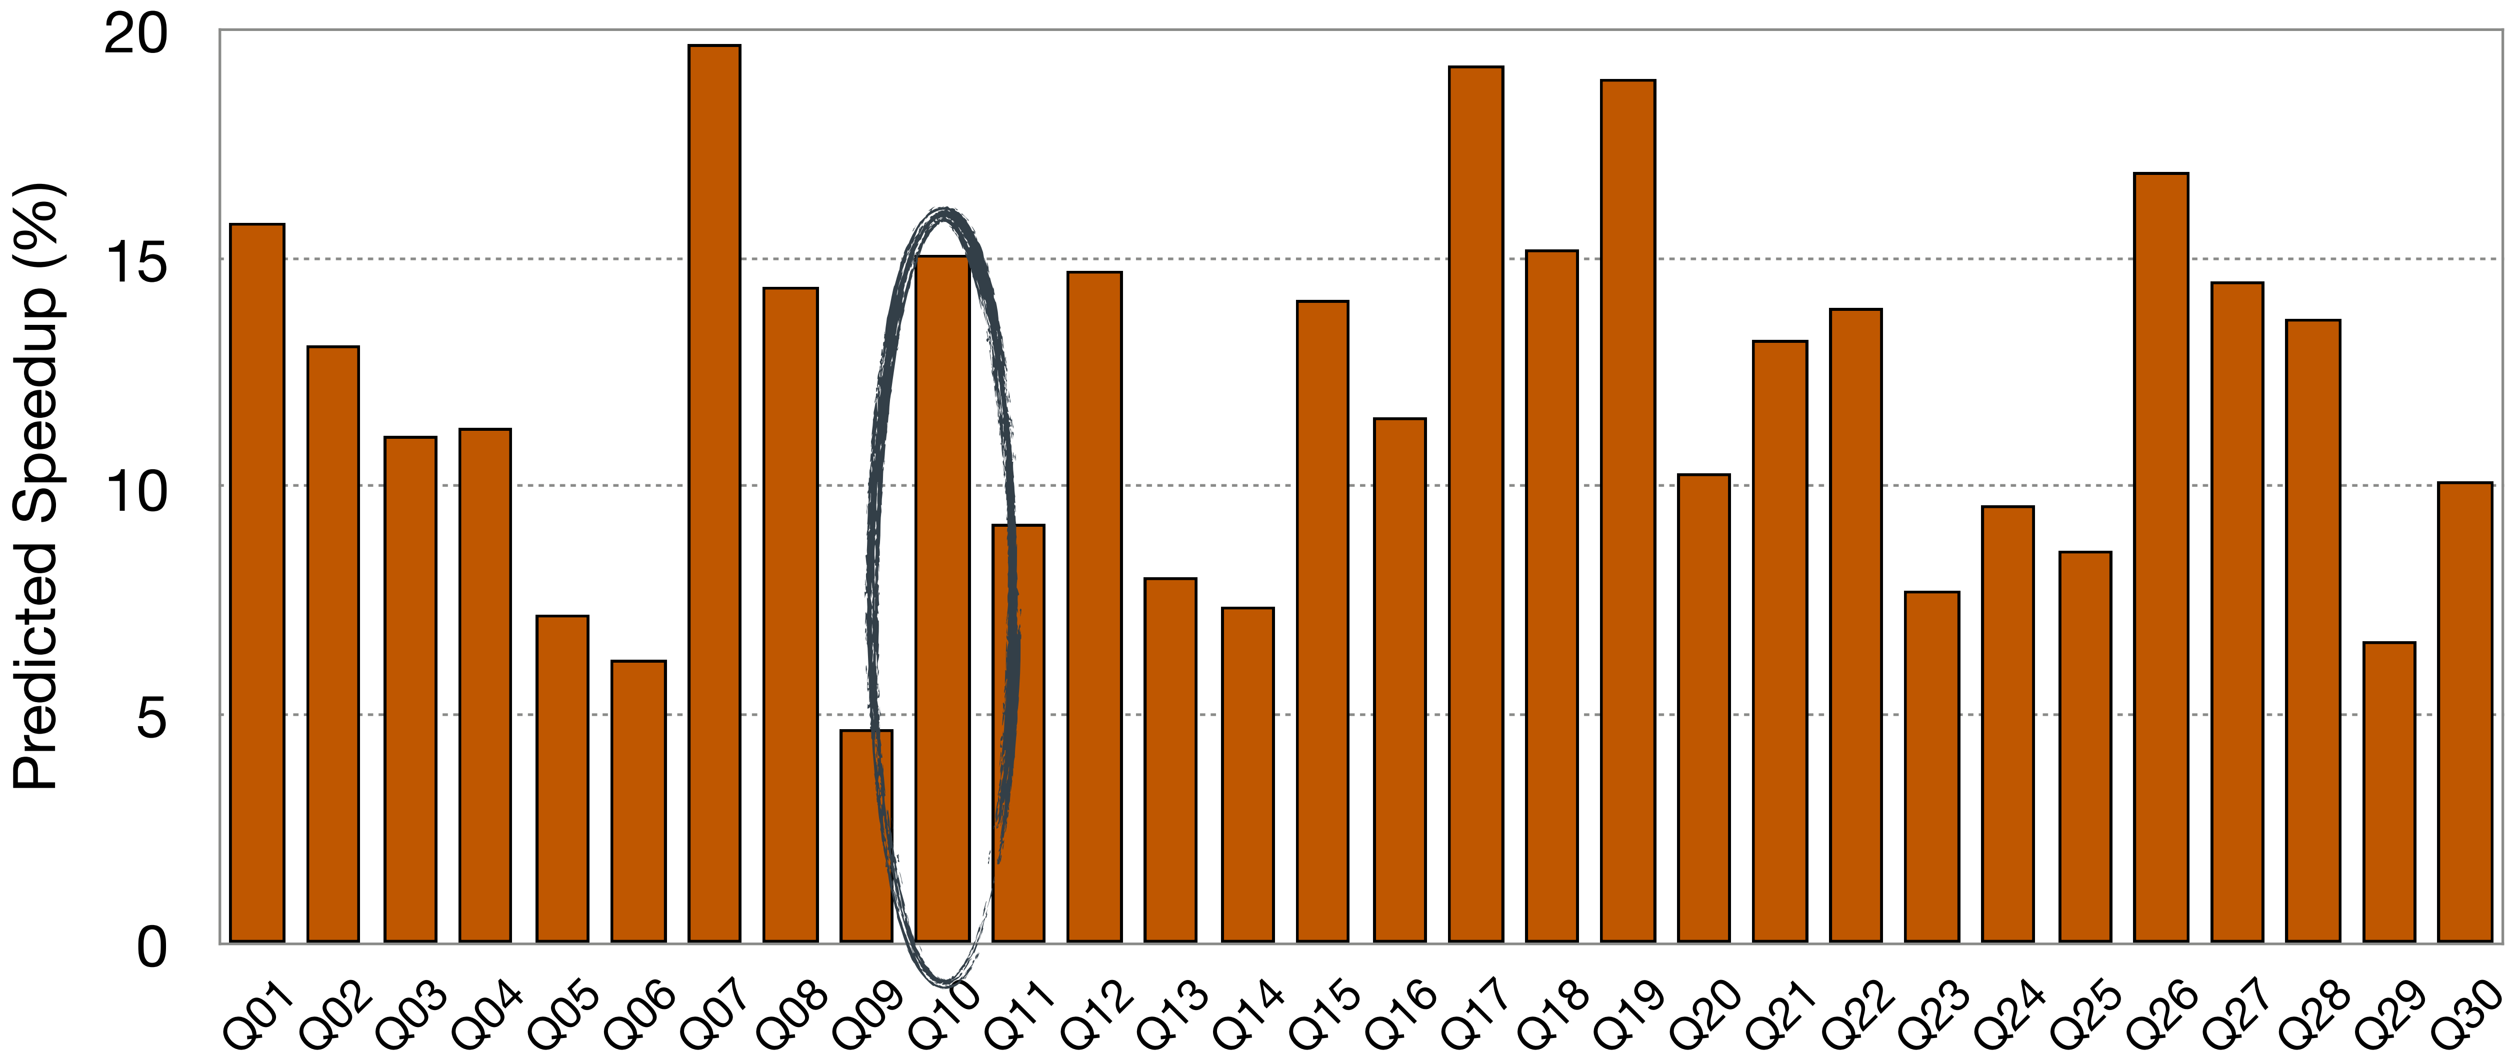

Notably, Q10's modeled speedup of 15.1% is remarkably close to our predicted speedup of 13.5%.

Takeaways

Takeaways

TPCx-BB (BigBench) is the only benchmark to faithfully capture the realism and diversity of industry workloads.

Takeaways

TPCx-BB (BigBench) is the only benchmark to faithfully capture the realism and diversity of industry workloads.

Thread-level parallelism is not so abundant in big data workloads as common wisdom would have us believe.

Takeaways

TPCx-BB (BigBench) is the only benchmark to faithfully capture the realism and diversity of industry workloads.

Thread-level parallelism is not so abundant in big data workloads as common wisdom would have us believe.

Current hardware and software cooperatively undermine Turbo Boost's ability to accelerate execution.

Takeaways

TPCx-BB (BigBench) is the only benchmark to faithfully capture the realism and diversity of industry workloads.

Thread-level parallelism is not so abundant in big data workloads as common wisdom would have us believe.

Current hardware and software cooperatively undermine Turbo Boost's ability to accelerate execution.

We propose core packing to aid Turbo Boost and predict 4-20% speedup.

Amdahl's Law is Alive and Well In Big Data Analytics

Questions

http://www.tpc.org/tpcx-bb/default.asp