

Reducing Tail Latencies Through Environment- and Neighbour-aware Thread Management

Andrew Jeffery
University of Cambridge

Chris Jensen
University of Cambridge

Richard Mortier
University of Cambridge

Abstract

Application tail latency is a key metric for many services, with high latencies being linked directly to loss of revenue. Modern deeply-nested micro-service architectures exacerbate tail latencies, increasing the likelihood of users experiencing them. In this work, we show how CPU overcommitment by OS threads leads to high tail latencies when applications are under heavy load. CPU overcommitment can arise from two operational factors: incorrectly determining the number of CPUs available when under a CPU quota, and the ignorance of neighbour applications and their CPU usage. We discuss different languages’ solutions to obtaining the CPUs available, evaluating the impact, and discuss opportunities for a more unified language-independent interface to obtain the number of CPUs available. We then evaluate the impact of neighbour usage on tail latency and introduce a new neighbour-aware threadpool, the *friendlypool*, that dynamically avoids overcommitment. In our evaluation, the *friendlypool* reduces maximum worker latency by up to $6.7\times$ at the cost of decreasing throughput by up to $1.4\times$.

1 Introduction

As shown repeatedly by large companies including Google, Amazon, and independent studies [5, 10, 11, 28, 32]: increased latency reduces client retention and costs money. This is exacerbated by modern micro-service architectures due to their deep nesting [29, 36], leading to a higher likelihood of a tail latency being observed by a user. Bursts of demand may also be correlated [25], leading to resource demand surges.

As Table 1 shows, large CPU counts are now available at many cloud providers. This leads to applications that can scale to the number of CPUs available on the host in order to handle larger volumes of requests [20, 26, 31]. However, they assume that they operate alone on the host, and so can make full use of the CPU resources. This assumption is often false, leading to overcommitment [27]. Each application creates the number of CPUs worth of threads, meaning the total threads

Table 1: Overview of the biggest VMs by CPU core count available at large cloud providers.

Cloud provider	vCPUs	Instance type
AWS	192	c7i.metal-48x1
Azure	96	Standard_D96d_v5
GCP	360	c3d-standard-360

per CPU scales with the number of applications. This leads to contention over the CPUs, leaving each application with far less than the full host it was expecting to have.

CPU quotas are used to isolate applications on the same host from using too much CPU time [34]. However, the interfaces that many applications use for determining the available CPU resources do not encapsulate the CPU quota’s restrictions [17, 18]. This is in part due to the isolation only being lightweight and so a leaky abstraction [4, 12]. This directly leads to CPU overcommitment and high tail latencies.

CPU quotas also prevent an application from bursting its CPU usage in order to handle demand surges. This leads to requests being queued, increasing the latency observed, or rejecting requests outright.

To take advantage of the larger nodes’ size and cost-efficiency applications are packed tightly, aiming for high utilisation [35]. However, when these applications have coinciding demand bursts they can contend over shared CPUs. Since the number of CPUs available to an application is typically treated as a static constant they do not adapt to the reduced CPU usage that they obtain. The CPU overcommitment then leads to increased tail latencies.

In this work we show that CPU overcommitment is still a problem in modern contexts (Section 2). We then draw comparisons and corresponding evaluations from this for the contributions of CPU quotas (Section 3 and 4) and neighbour CPU load (Section 5). Finally, we present a new thread pool design that dynamically avoids overcommitment, improving tail latencies (Section 5.2).




Figure 1: Structure of the applications and the latencies being measured.

2 Overcommitment: OS Threads > CPUs

This section outlines the model of computation that we refer to throughout this work, as well as defining overcommitment and evaluating its impacts on tail latency.


Computation Model A single processor socket is formed of multiple physical processing cores, each presenting potentially multiple CPUs due to hyper-threading [9]. Each CPU is exposed to the operating system (OS), which manages the execution of OS threads over the many CPUs. These OS threads are *scheduled* onto the CPUs based on a scheduling algorithm. OS threads are preemptive, so they may not complete their computation in one scheduling cycle before being taken off the CPU.

`cgroups` [4] enable the assignment of CPU shares and a CPU quota to a process. The OS scheduler, Linux’s CFS in this case, packs OS threads to CPUs based on their share relative to all total shares of OS threads ready to be executed. The scheduler also enforces the quota as the maximum amount of time a process is able to run for in any given period. CPU quotas are typically employed to prevent processes from hogging CPU resources, particularly in multi-tenant environments.


Since OS threads are typically heavy-weight, modern languages use a design of *lightweight threads* [30]. These lightweight threads are language-specific so require a runtime to manage them, multiplexing them over OS threads created by the runtime. The runtime’s OS threads are then commonly partitioned into *workers* and *waiters*. Worker threads are small in number and perform the CPU-bound work, while waiter threads are greater in number and execute blocking operations, spending most of their time in wait queues.

Overcommitment A well known result is that as the number of threads increases, coordination costs also increase, due to contention over shared resources [22]. This also extends to CPUs as once each CPU is fully utilized, adding further threads increases coordination costs without increasing the amount of useful work done. This situation is known as overcommitment. However, considering the total quantity of useful work is not the whole story, instead focusing here on latency.

To study the effect of overcommitment throughout this work we implement a simple program that spawns worker




(a) Queue latency decreases as more threads are used, up to the number of CPUs. Fib latency remains low until overcommitment.

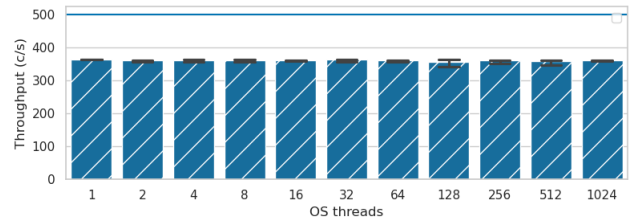


(b) Throughput gradually increases as more threads are used, up to the number of CPUs.

Figure 2: Overall latency and throughput at various amounts of OS threads in Rust. Workers have no contention.



(a) Queue latency decreases as more threads are used. Fib latency increases with overcommitment.



(b) Throughput remains constant, limited by lock contention.

Figure 3: Overall latency and throughput at various amounts of OS thread overcommitment in Rust. Workers have contention over the `fib` computation starting with a lock at `fib(30)`.

threads, each performing CPU-bound work, a naive fibonacci (`fib`) sequence computation, computing `fib(30)`. We drive the workers at a target rate in an open-loop fashion, representing a uniform demand. We collect timing measurements at three points, the time at which the work item is “received” by the service and pushed into the queue (*queue start*), the time the work item is pulled out of the queue by a worker (*fib start*), and the time that the work item has finished being processed (*end*). From these we derive the *queue latency* and the *fib latency*, shown in Figure 1. We also calculate overall throughput of the process as the time taken to complete divided by the total runtime.

The experiments in this paper were conducted on a single machine running Linux 6.2.0 with a dual socket Intel Xeon Silver 4112 Processor featuring 4 physical cores each with hyper-threading, so 16 CPUs total. Experiments are run for 5 seconds with 5 repeats. Latency plots show points for all repeats, throughput plots show the median with an error bar between the minimum and maximum across repeats.

Figure 2 breaks down the total throughput and latency for client requests at various numbers of threads in a Rust implementation. When the total system’s throughput is lower than the request rate (fewer than 16 threads) requests must queue, resulting in higher than optimal latency. However if the system is overcommitted (more than 16 threads), each thread has a lower effective throughput and hence increased fib latency, getting worse with more overcommitment. This occurs because the worker thread pulls an item off the queue, starting its fib latency timer, but due to overcommitment, the OS scheduler may choose to preempt the worker thread and run another before the current one completes working on its item. The original thread must then wait to be scheduled onto a CPU again to continue work and report its result, stopping the timer. The effect of overcommitment on the latency of each request is exacerbated by contention over shared resources, a mutex around the computation of `fib(30)` in this case, shown in Figure 3. Since this is an extreme contention scenario, showing the limit, we proceed only without contention to study the most optimistic case.

In the rest of this work we will not consider queue latency as it arises due to the overload of the service and other strategies such as load shedding or horizontal scaling can be used to mitigate it.

3 How Do Applications Tune Themselves?

When an application starts up it needs to perform some initialisation. This typically involves getting some information about the environment and using that to set up some internal configuration parameters, such as threadpool sizes. This section investigates how this initialisation works under the presence of CPU quotas and the impacts that it can have.

Table 2: Lightweight threading mechanism, runtime, and CPU source for each language.

Lang	Thread	Runtime	CPU source
C++	Coroutine [3]	Built-in	CPU affinity
Go	Goroutine [6]	Built-in	CPU affinity
Java	Virt. threads [7]	Built-in	cgroup
Julia	Task [8]	Built-in	CPU affinity
OCaml	Eio fiber [14]	Eio [13]	CPU affinity
Rust	Tokio task [16]	Tokio [15]	cgroup




Figure 4: Example schedulings, shown in scheduling periods of 2 apps with CPU quotas equivalent to 1 CPU on a 2 CPU system.

3.1 Working Under CPU Quotas

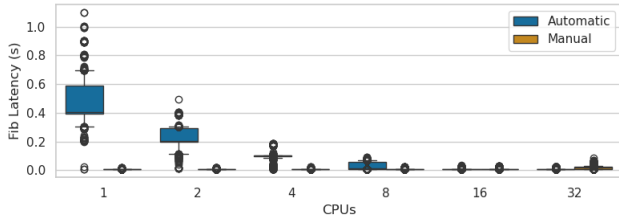
Assuming: Quotas in place.

Problem 1: Applications do not account for their CPU quota when calculating the number of CPUs.

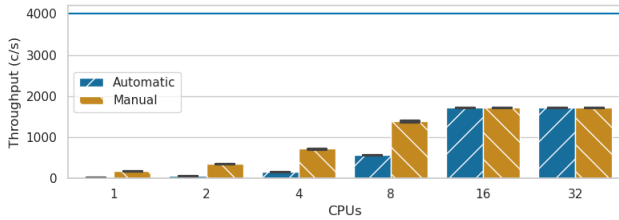
The logical conclusion from the previous section is that applications should limit themselves to OS threads = CPUs, a thread-per-core architecture [21], to avoid overcommitment which should be easy to read from the environment.

Unfortunately, when applications are deployed with CPU quotas, obtaining the correct CPU count becomes more difficult. We would expect that an application should use the formula: OS threads = quota cores, for its self-tuning, where quota cores is quota time/scheduler period. However, despite the seeming ease of obtaining this from the `cgroups` API, multiple languages do not obtain the correct number. Table 2 summarizes the approaches some popular languages use to obtain the number of CPUs available for the application. We observed that a large number of languages are not `cgroup` aware, and so do not look for their allocated CPU quota. Instead, most rely on the OS system calls, such as `sched_getaffinity` [17] and `sysconf(_SC_NPROCESSORS_ONLN)` [18], which do not reflect the CPU quota in place, instead returning the total number of CPUs in the system.

Since the CPU quota is an enforcement on maximum usage, overcommitment of OS threads can easily happen when incorrectly scaled, as shown in Figure 4a by the full run queue. As shown in Section 2, overcommitment of OS threads to CPUs results in increased latency due to preemption. Figure 5



(a) Fib latency remains constantly low when matching threads to CPU quota. Letting the runtime decide the number of threads leads to a large fib latency, decreasing as CPU quota approaches number of CPUs.



(b) Throughput is impacted by overcommitment of CPUs due to incorrect automatic detection of CPU availability.

Figure 5: Impact of using an incorrect OS thread count when under a CPU quota in Go.

presents the fib latency and the throughput from Go when manually setting the optimal number of threads and when letting the runtime determine it automatically. However, in this experiment we place the applications in `cgroups` with CPU quotas corresponding to varying numbers of CPUs. Since Go, the cloud-native language [33], is not `cgroup`-aware it spawns a number of OS threads equal to the number of CPUs on the entire machine (16 in this case). This leads Go to having a large overcommitment for most cases and sees an accordingly high latency on its work. For the Go language, this has been observed widely, including at Uber, a prominent user of Go, having published their library to obtain the CPU quota from the `cgroup` API to contain the issue [1].

Notably, other CPU containment mechanisms do not feature these limitations. CPU shares, which govern a weighting for the time share a thread gets, are only a lower bound. CPU affinity allows pinning a thread to a CPU but, working on the thread-level, does not introduce a direct limitation on how many threads can run.

3.2 Getting the Correct CPU Count

Assuming: Quotas in place.

Solution 1: Fix the syscall for number of CPUs to be quota-aware, or add a new one.

Given that we know the interfaces languages use to deter-

mine the number of CPUs to use we see two complementary solutions. Firstly, languages and their runtimes could be made `cgroup`-aware. This is likely the fastest method to fix this situation but requires that each language independently work out their fix and correctly interpret the `cgroup` API. It also places maintenance burden on the developers as new versions of the `cgroup` API are added, such as with `cgroups` v1 and v2. Secondly, as most languages already see fit to offload the calculation to the OS, a new `syscall` could be added to obtain the correct number of CPUs that the process should be using. This would need to take into account the CPU affinity mask for the process as well as implications of a `cgroups` CPU quota. This solution would be portable between languages and reduce the complexity required in them. Importantly, this would encapsulate being `cgroup`-aware but also aim to incorporate future changes too. We leave the implementation of these fixes out of the scope of this work, but note that the results of their changes would follow the pattern shown by the “Manual” thread count in Figure 5.

4 Capitalising On Spare Resources

Assuming that the issue with CPU quotas has been fixed, we now look at limitations that are still imposed. To do this, we extend our view broader than our single application, out to the host and all the processes it may be running.

4.1 CPU Quotas Are Wasteful

Assuming: Quotas in place, OS threads = quota cores.

Problem 2: Applications without sufficiently busy neighbours can see wasted CPUs that they can’t use.

When all applications are fully using their quota and the sum of the quota cores is equal to the number of CPUs then the system behaves well, each OS thread gets matched to a CPU. However, such a case with constant work is unlikely, even for batch jobs. When an application does not fully utilise its quota then some portion of the CPU will be left under-utilised, shown in Figure 4b. This wastage is clearly suboptimal from a utilisation perspective and would ideally be avoided. Particularly this involves the application being burstable, often in response to its demand. This is more common in modern deployments with dense packing of applications to hosts.

4.2 Enabling Bursts

Assuming: Quotas in place, OS threads = quota cores.

Solution 2: Do not assign CPU quotas.

In order to solve the issue of applications with quotas being unable to burst up their usage we simply recommend




Figure 6: Example schedulings, shown in scheduling periods of 2 apps with no CPU quotas on a 2 CPU system.

removing the CPU quota unless strictly required. They have been shown to have negative performance implications in real-world contexts [2]. This can also be used as a stopgap solution until better support for detecting the number of CPUs based off CPU quotas is added (Section 3).

5 The Cost Of Ignoring Your Neighbours


This section analyses issues arising when other applications (neighbours) are also present on the host. We assume that CPU quotas are not in place for the applications, but the scheduler is fair, as is Linux’s CFS.

5.1 Noisy Neighbourhood


Assuming: No quotas in place, OS Threads = CPUs.
Problem 3: Applications ignore their neighbours’ usage when calculating the number of threads to use.

Having removed the CPU quotas on applications, each application is free to create a number of threads equal to the number of CPUs as shown in Figure 6a. Provided that the sum of busy threads at any instant in time does not exceed the number of CPUs the scheduler should be able to map each thread to its own distinct CPU. However, in reality the overcommitment of OS threads to CPUs is severely higher than 1, by a factor of the number of applications, as each will spawn one thread per CPU, shown in Figure 6b. Unknowingly to each application, there is contention over the CPU cores.

CPU shares should be used to ensure that each application, when it is busy, has a minimum portion of CPU time. Assuming that all applications on a host are busy at the same time, then they will each be given their CPU shares. In this scenario each application is then directly subject to its own overcommitment. This effectively partitions the applications from each other, pushing the impact of overcommitment directly, and solely, onto them. This means that applications that do not overcommit their share will not observe the associated increased latencies. However, those applications that do overcommit will only impede the performance of their own other threads, and so observe the associated increase in latencies seen previously in Figures 2 and 3.



(a) Ignoring neighbour usage leads to a large fib latency, increasing with number of neighbours. A collaborative approach reduces this.



(b) Ignoring neighbour usage maintains a high throughput, whereas partitioning CPUs by reducing thread count limits it.

Figure 7: Performance when sharing a node with neighbours using the Rust app. Ignorant: threads = CPUs for each process. Collaborative: the dynamic friendly threadpool. Optimal: threads = CPUs / processes. This highlights the trade-off space between latency and throughput from overcommitment.


There is now a conflict, between having enough OS threads available to be able to handle bursts when neighbours are not fully utilising their share, and not impacting latency when neighbours are busy. Figure 7 shows the effect on latency of overcommitting OS threads to CPUs when neighbours are busy, looking at the “Ignorant” kind.

5.2 Dynamically Adapting To Neighbours


Assuming: No quotas in place, OS Threads = CPUs.
Solution 3: Dynamically alter the number of OS threads that perform work for an application.

To handle the adaptive nature of neighbour usage the pool of worker threads needs to dynamically scale up and down. For this, we introduce the *friendlypool*, a threadpool that is friendly to its neighbours and is aware of their usage, marked as *Collaborative* in Figures. It works by simply spawning another thread that is in charge of periodically obtaining the number of the CPUs that the current process should use and scaling the threadpool to that number. When scaling, we do not destroy or create OS threads, but rather they each watch the number of desired active threads and pull work or not, accordingly. Initially one OS thread is created per CPU.

The number of OS threads the current application should be using is equivalent to its proportion of the usage of the



(a) The overcommit factor is directly proportional to the fib latency.



(b) The overcommit factor is directly proportional to the throughput.

Figure 8: Performance of the collaborative friendly pool with variable overcommitment factor.

CPUs in the last period. For instance, if the application is alone on the host, then it will be consuming all of the CPU, and so the usage will be 100%. However, say another identical application is run, then the usages may be evenly split, each 50%. This is a key property obtained from the fairness of the OS thread scheduler. This usage can then be scaled by the number of CPUs in the system to count how many OS threads should be active. The calculation is:

$$\text{active threads} = \left\lceil O \cdot \frac{\text{cpu time}_{\text{self}}}{\text{cpu time}_{\text{all}}} \cdot \text{CPUs} \right\rceil$$

A *control thread* polls the “cpu time” values periodically, limited by how often they are updated by the Kernel, the clock tick frequency. This defaults to 100Hz in our setup, and so this is evaluated every 10ms. When a worker thread has no work to perform, it *park*s itself, waiting for the control thread to *unpark* it when it should execute.

Figure 7 presents the results of different thread pool division strategies in comparison to an “Optimal” allocation. In Figure 7a we observe that the collaborative friendly pool does improve latency of the program compared to a simple static thread pool. However, Figure 7b highlights that the throughput of the friendly pool does not match that of the ignorant pool. This drop in throughput is due to the overhead of the “frontend” thread generating the work, which is masked when more worker threads are active, as they enable concurrent processing of the requests. This decreased throughput is also present for the optimal allocation of threads (one per CPU).

To traverse the trade-off space between latency and throughput we can augment the formula for the number of active threads by a variable O . This is an *overcommitment fac-*

tor, specifying by what amount the number of active threads should overcommit the CPUs and is intended to be tuned for different applications. Figure 8 shows the impact of choosing some sample *overcommitment factor* values, highlighting the trade-off between tail latency and throughput and the space that this parameter enables the user to traverse.

We believe that this approach of making the thread pool dynamic is scalable to current language runtimes for lightweight threading. They follow the typical pattern of having worker threads that are CPU-bound and blocking threads that consume negligible CPU, waiting for operations to complete. The lightweight threads can then be multiplexed over the currently active worker threads, enabling improved latency.

6 Related Work

Janssen [24] provides a motivation for moving towards dynamic threadpools that take into account multiple dimensions in order to maximise performance, particularly in the presence of neighbours. However, the performance in focus is throughput, not tail latency as we argue for. Additionally, our work focuses purely on tuning the threadpool for CPU usage.

Grand central dispatch [19] for Darwin platforms acts as a per-application threadpool manager. It presents interfaces such as queues for the application to schedule work in multiple classes, from which the dispatch runtime can order execution on OS threads. The runtime monitors the CPU configuration to balance work, dynamically scaling the thread pool. This solution shares similarities in design with our threadpool but is more heavily tied into the language runtime with its work dispatch design.

Huang et al [23] make similar arguments to us focusing on burstable containers, they however use CPU sets to dynamically scale the number of threads. Instead, we simply park a thread, or send it to sleep, awaiting its next runnable time. This lets the OS scheduler treat the thread as not runnable, rather than requiring more user-level interventions.

7 Conclusion

We have presented an analysis of CPU overcommitment from OS threads and the various ways that this situation can arise, focusing on CPU quotas and ignorance of neighbour CPU usage. We presented a new threadpool design, the *friendlypool*, that dynamically scales the number of active worker threads to match the number of available CPUs and showed that it significantly reduced latencies, by up to $6.7\times$. However, it comes at the cost of some throughput, up to $1.4\times$, which we mitigate through an overcommitment factor, giving the user an option for traversing the trade-off of latency and throughput. This threadpool design is applicable to modern lightweight thread runtimes, enabling the potential for reducing tail latencies from the core of applications in a scalable manner.

References

- [1] automaxprocs. <https://github.com/uber-go/automaxprocs>. Accessed: 2024-01-08.
- [2] Avoiding cpu throttling in a containerized environment. <https://www.uber.com/en-GB/blog/avoiding-cpu-throttling-in-a-containerized-environment/>. Accessed: 2024-01-08.
- [3] C++ coroutines. <https://en.cppreference.com/w/cpp/language/coroutines>. Accessed: 2024-01-16.
- [4] cgroups. <https://www.man7.org/linux/man-pages/man7/cgroups.7.html>. Accessed: 2024-01-16.
- [5] The cost of latency. <https://perspectives.mvdirona.com/2009/10/the-cost-of-latency/>. Accessed: 2024-01-08.
- [6] Go's goroutines. https://go.dev/doc/effective_go#goroutines. Accessed: 2024-01-16.
- [7] Java virtual threads. <https://docs.oracle.com/en/java/javase/21/core/virtual-threads.html>. Accessed: 2024-01-16.
- [8] Julia tasks. <https://docs.julialang.org/en/v1/base/parallel/>. Accessed: 2024-01-16.
- [9] lscpu. <https://linux.die.net/man/1/lscpu>. Accessed: 2024-01-16.
- [10] Make data useful. <https://web.archive.org/web/20090619174403/http://home.blarg.net/~glinden/StanfordDataMining.2006-11-29.ppt>. Accessed: 2024-01-08.
- [11] Marissa mayer at web 2.0. <http://glinden.blogspot.com/2006/11/marissa-mayer-at-web-20.html>. Accessed: 2024-01-08.
- [12] namespaces. <https://man7.org/linux/man-pages/man7/namespaces.7.html>. Accessed: 2024-01-16.
- [13] Ocaml eio. <https://github.com/ocaml-multicore/eio>. Accessed: 2024-01-16.
- [14] Ocaml eio fiber. <https://ocaml-multicore.github.io/eio/eio/Eio/Fiber/index.html>. Accessed: 2024-01-16.
- [15] Rust tokio. <https://tokio.rs/>. Accessed: 2024-01-16.
- [16] Rust tokio task. <https://docs.rs/tokio/latest/tokio/task/>. Accessed: 2024-01-16.
- [17] sched_getaffinity. https://linux.die.net/man/2/sched_getaffinity. Accessed: 2024-01-11.
- [18] _SC_NPROCESSORS_ONLN. https://www.gnu.org/software/libc/manual/html_node/Processor-Resources.html. Accessed: 2024-01-11.
- [19] Using grand central dispatch on macos. <https://apple.github.io/swift-corelibs-libdispatch/tutorial/>. Accessed: 2024-01-08.
- [20] Raja Appuswamy, Christos Gkantsidis, Dushyanth Narayanan, Orion Hodson, and Antony Rowstron. Scale-up vs scale-out for hadoop: Time to rethink? In *Proceedings of the 4th Annual Symposium on Cloud Computing, SOCC '13*, New York, NY, USA, 2013. Association for Computing Machinery.
- [21] Pekka Enberg, Ashwin Rao, and Sasu Tarkoma. The impact of thread-per-core architecture on application tail latency. In *2019 ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*, pages 1–8. IEEE, 2019.
- [22] Hugo Guiroux, Renaud Lachaize, and Vivien Quéma. Multicore locks: The case is not closed yet. In *Proceedings of the 2016 USENIX Conference on Usenix Annual Technical Conference, USENIX ATC '16*, page 649–662, USA, 2016. USENIX Association.
- [23] Hang Huang, Yuqing Zhao, Jia Rao, Song Wu, Hai Jin, Duoqiang Wang, Suo Kun, and Lisong Pan. Adapt burstable containers to variable cpu resources. *IEEE Transactions on Computers*, 72(3):614–626, 2022.
- [24] Martijn Janssen. Adaptive thread pool scaling and real-time thread monitoring. 2022.
- [25] Hao Jiang and Constantinos Dovrolis. Why is the internet traffic bursty in short time scales? In *Proceedings of the 2005 ACM SIGMETRICS international Conference on Measurement and Modeling of Computer Systems*, pages 241–252, 2005.
- [26] Mijung Kim, Jun Li, Haris Volos, Manish Marwah, Alexander Ulanov, Kimberly Keeton, Joseph Tucek, Lucy Cherkasova, Le Xu, and Pradeep Fernando. Sparkle: Optimizing spark for large memory machines and analytics, 2017.
- [27] Li Liu, Haoliang Wang, An Wang, Mengbai Xiao, Yue Cheng, and Songqing Chen. Mind the gap: Broken promises of cpu reservations in containerized multi-tenant clouds. In *Proceedings of the ACM Symposium on Cloud Computing, SoCC '21*, page 243–257, New York, NY, USA, 2021. Association for Computing Machinery.
- [28] Ciamac C. Moallemi and Mehmet Sağlam. Or forum—the cost of latency in high-frequency trading. *Operations Research*, 61(5):1070–1086, 2013.

- [29] Irakli Nadareishvili, Ronnie Mitra, Matt McLarty, and Mike Amundsen. *Microservice Architecture: Aligning Principles, Practices, and Culture*. O'Reilly Media, Inc., 1st edition, 2016.
- [30] Dan Stein and Devang Shah. Implementing lightweight threads. In *USENIX Summer*, volume 575, 1992.
- [31] Disha Talreja, Kanishka Lahiri, Subramaniam Kalambar, and Prakash Raghavendra. Performance scaling of cassandra on high-thread count servers. In *Proceedings of the 2019 ACM/SPEC International Conference on Performance Engineering, ICPE '19*, page 179–187, New York, NY, USA, 2019. Association for Computing Machinery.
- [32] Parth Thakkar, Rohan Saxena, and Venkata N Padmanabhan. Autosens: inferring latency sensitivity of user activity through natural experiments. In *Proceedings of the 21st ACM Internet Measurement Conference*, pages 15–21, 2021.
- [33] Matthew A Titmus. *Cloud Native Go*. " O'Reilly Media, Inc.", 2021.
- [34] Paul Turner, Bharata B Rao, and Nikhil Rao. Cpu bandwidth control for cfs. 2010.
- [35] Abhishek Verma, Luis Pedrosa, Madhukar Korupolu, David Oppenheimer, Eric Tune, and John Wilkes. Large-scale cluster management at google with borg. In *Proceedings of the Tenth European Conference on Computer Systems, EuroSys '15*, New York, NY, USA, 2015. Association for Computing Machinery.
- [36] Zhizhou Zhang, Murali Krishna Ramanathan, Prithvi Raj, Abhishek Parwal, Timothy Sherwood, and Milind Chabbi. CRISP: Critical path analysis of Large-Scale microservice architectures. In *2022 USENIX Annual Technical Conference (USENIX ATC 22)*, pages 655–672, Carlsbad, CA, July 2022. USENIX Association.