

ROBUST HEURISTICS: Attacks and Defenses for Job Size Estimation in WSJF Systems

**Erica Chiang, Nirav Atre, Hugo Sadok,
Weina Wang, Justine Sherry**

June 2022
CMU-CS-22-117

Computer Science Department
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Packet scheduling algorithms control the order in which a system serves network packets, which can have significant impact on system performance. Recent work in adversarial scheduling has shown that *Weighted Shortest Job First* (WSJF) – scheduling packets by the ratio of job size to packet size – significantly mitigates a system’s susceptibility to algorithmic complexity attacks (ACAs), a particularly dangerous class of Denial-of-Service (DoS) attacks. WSJF relies on knowledge of a packet’s job size, information that is not available *a priori*. In this work, we explore the theoretical implications of using heuristics for job size estimation. Further, we consider preemption as another technique that may help protect systems when job sizes are unknown. We find that heuristics with certain properties (*e.g.*, estimated job-size-to-packet-size ratios increasing monotonically with the actual ratios, step function-based job categorization) can protect systems against ACAs with varying guarantees, while preemption alone does not.

Keywords: algorithmic complexity attacks, adversarial scheduling, robust heuristics

1 Introduction

Packet scheduling algorithms control the order in which a system serves network packets, which can have significant impact on system performance. Many systems rely on Shortest Job First (SJF), an important packet scheduling algorithm with many desirable properties. Classic results show that SJF provably minimizes the average job completion time [2], and recent work [1] shows that a variant of SJF also protects systems against algorithmic complexity attacks (ACAs), a particularly dangerous class of Denial-of-Service (DoS) attacks [3]. In an ACA, an adversary exploits the worst-case behavior of an algorithm in order to induce a large amount of work in the target system, causing a significant drop in goodput despite using only a small amount of attack bandwidth. SurgeProtector [1] demonstrated that using *Weighted* SJF (WSJF) – scheduling packets by the ratio of job size to packet size – significantly mitigates the impact of ACAs on *any* networked system.

There is just one problem: *how do we determine a packet’s job size without running the job?* A common technique is to estimate job sizes using heuristics. In an adversarial setting, however, inaccuracies in job size estimation may be exploitable, re-opening the door to ACA vulnerabilities. In this work, we explore three strategies for using WSJF in practice, and bound their vulnerability against ACAs. Our key findings are: (1) any heuristic that results in estimated job-size-to-packet-size ratios increasing monotonically with the actual ratios will lead to perfect scheduling, thereby maintaining SurgeProtector’s guarantees; (2) a heuristic that accurately separates jobs into job size categories can also protect a system against ACAs, but the guarantees are not as strong; and (3) preempting jobs that run for longer than their estimates does not maintain bounds on an adversary’s damage if the estimates are inaccurate.

2 Background and Motivation

Atre et al. [1] argue that in the absence of job size information, we can use heuristics to estimate job sizes. In this context, a heuristic h is a mapping from packets to estimated job sizes.¹ However, in an adversarial setting, it is conceivable that incorrect estimates could undermine guarantees of a system’s protection against attacks. Besides heuristics, we also explore *preemption* (i.e., pausing a job while it is in service, then resuming it at a later point), another technique that may be able to help protect systems when job sizes are unknown.

2.1 Mathematical Framework

We first build a mathematical framework for analyzing the impact of adversarial traffic on a system. Each packet can be characterized by a packet size, $s(p)$ (the amount of data sent over the wire, in bits), and a job size, $c(p)$ (the time required to process the packet, in seconds). We define a packet’s z -ratio as the ratio of its job size to packet size, noting that WSJF schedules packets by increasing z -ratio. Finally, we quantify the vulnerability of the system using the *Displacement Factor* (DF) [1], defined as the adversary’s payoff relative to the amount of resources they invest

¹Here, we assume direct correspondence between real and estimated job sizes for simplicity. However, our analysis admits more sophisticated mappings (e.g., probability distributions) as well.

into the attack:

$$DF = \frac{\text{Innocent traffic displaced (Gbps)}}{\text{Attack bandwidth used (Gbps)}}$$

2.2 WSJF and ACAs

In this section, we summarize the results of SurgeProtector [1] in the context of our heuristic-based approach to packet scheduling. SurgeProtector uses the DF to quantify the severity of an ACA, and shows that WSJF scheduling imposes an upper-bound of 1 on the DF. This implies that in order to displace 1 bps of innocent traffic, an adversary must invest *at least* 1 bps of their own bandwidth into the attack. Given the practical limitations of crafting and sending large volumes of data, a bounded DF greatly reduces the harm that an adversary can do to a system. In this paper, we aim to understand how these theoretical findings extend to practical settings where job sizes are not known *a priori*.

2.3 Incorrect Estimates

The accuracy of heuristics is crucial to maintaining DF guarantees. To gain intuition as to why poorly designed heuristics can lead to an unbounded DF, we consider a heuristic that incorrectly estimates packets of a certain real job size, while all other packets are estimated correctly. Figure 1 demonstrates why incorrect estimates can be dangerous; in this example, we assume that all packets have unit size, such that WSJF orders packets by job size (represented by the packet width).

More formally, consider a heuristic that estimates the job size for adversarial packets as ϵ , allowing adversarial packets to have an arbitrarily small z -ratio as ϵ goes to 0. This implies that an attacker can push the system into overload using an infinitesimally small amount of their own bandwidth, displacing all innocent traffic in the process, and leading to an unbounded DF. Thus, incorrect job size estimates in a scheduling policy that relies on the job sizes of packets (*e.g.*, WSJF) can lead to arbitrarily bad DFs.




Figure 1: WSJF *should* de-prioritize the attack packet, but with incorrect estimates, innocent packets are displaced.

3 Novel Theoretical Findings

In this section, we present three novel theoretical findings regarding protection against ACAs when job sizes are unknown. Proofs for all theorems can be found in Appendix A.

3.1 Strictly Monotonically Increasing Heuristics Yield Perfect Scheduling

We first develop the concept of a ‘perfect’ heuristic, meaning that all packets are scheduled correctly when using estimated job sizes. Since correctly ordering all packets is equivalent to preserving the relative ordering between any pair of packets, a perfect heuristic must estimate job sizes such that between any two packets, the packet with smaller z -ratio will have a smaller *estimated* z -ratio. We can visualize this as any function mapping actual ratios to estimated ratios that is *strictly monotonically increasing*, as seen in Figure 2.

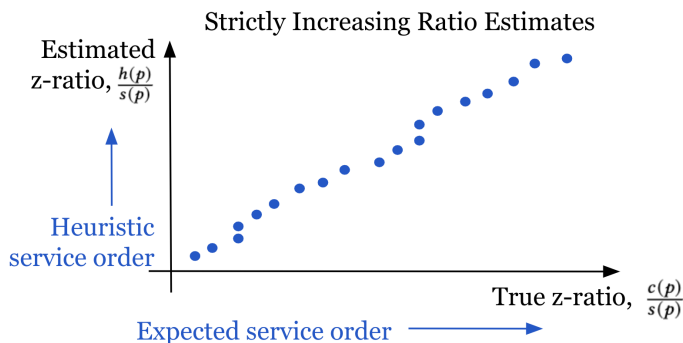


Figure 2: Strictly monotonically increasing ratios lead to perfect scheduling.

Any such heuristic preserves the relative ordering of packets as they are scheduled according to WSJF, which in turn maintains all guarantees from [1] and yields an upper-bound of 1 on the DF. In §A.1, we prove the following:

Theorem 1 (DF of Monotonic Heuristic). *Under WSJF, a heuristic h is perfect if and only if $\frac{h(p)}{s(p)}$ is strictly monotonically increasing relative to $\frac{c(p)}{s(p)}$; such heuristics result in the DF being upper-bounded by 1.*

3.2 Step Functions Guarantee a Constant DF

In this section, we consider ‘step function’ heuristics in which packets are correctly classified into job size categories, but packets within each category (‘step’) are indistinguishable. In particular, we consider heuristics where the range of actual job sizes that each step covers has some upper bound that is a constant multiplicative factor, k , times the lower bound, and estimates of each step increase by the same factor k , as depicted in Figure 3.




Figure 3: Step function heuristic.

Despite categorizing job sizes on a coarse level, the discrete steps still enforce some lower

bound on how small each packet’s job size estimate can be, protecting all jobs below a certain threshold. As we show in §A.2, this yields an upper-bound of k on the DF.

Theorem 2 (DF of Step Function Heuristic). *A heuristic of the form $h(p) = a \cdot k^{\lfloor \log_k c(p) \rfloor}$, where a is some arbitrary constant, results in the DF being upper-bounded by k .*

3.3 Preemption Can’t Guarantee DF Bounds

Finally, we consider preemption as an additional aid to protecting systems against ACAs. The setup is as follows: each incoming job is assigned an estimated job size of J_p ; if the job does not finish running within the allocated J_p time, the system preempts it and reinserts the job (with saved state) back into the scheduling queue, with an increased estimated job size of $2J_p$. The preemption model is depicted in Figure 4.




Figure 4: Preemption system model.

Since job sizes are unknown, this allows us to systematically allocate resources to each packet and ensures that packets finish according to job size order. However, even if there is no preemption cost – an overly optimistic assumption – this setup can result in an unbounded DF. As we show in §A.3, preemption alone cannot guarantee any bound on the DF:

Theorem 3 (DF of Preemptive Model). *Under WSJF with preemption but without heuristics, there exist some regimes of system parameters for which the DF is lower bounded by $\frac{\rho}{1-\rho}$, where $\rho \leq 1$ is the load on the system due to innocent traffic.*

4 Next Steps

Having identified desirable properties for heuristics and a framework for reasoning about their vulnerability, the main unanswered question is: how to design data structures and corresponding heuristics such that we see these properties in practice? In addition, while we do not see theoretical bounds on the DF as a result of preemption alone, is it possible that some level of preemption could still be beneficial in practice?

References

- [1] Nirav Atre, Hugo Sadok, Erica Chiang, Weina Wang, and Justine Sherry. SurgeProtector: Mitigating temporal algorithmic complexity attacks using adversarial scheduling. In *Proceedings of the 2022 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, New York, NY, USA, 2022. Association for Computing Machinery.
- [2] Alan Cobham. Priority Assignment in Waiting Line Problems. *Journal of the Operations Research Society of America*, 2(1):70–76, 1954.

- [3] Scott A. Crosby and Dan S. Wallach. Denial of Service via Algorithmic Complexity Attacks. In *12th USENIX Security Symposium (USENIX Security 03)*, Washington, D.C., August 2003. USENIX Association.

A Proofs for DF Analysis

A.1 Proof of Theorem 1 (Perfect Heuristic)

Proof. We first prove that a heuristic with strictly monotonically increasing ratios must be a perfect heuristic. Consider two arbitrary packets p_A and p_B , and a heuristic h with strictly monotonically increasing ratios. $\frac{c(p_A)}{s(p_A)} < \frac{c(p_B)}{s(p_B)} \implies \frac{h(p_A)}{s(p_A)} < \frac{h(p_B)}{s(p_B)}$, which means that if p_A has a smaller true z -ratio then it must be served first in a WSJF system, and thus the heuristic is optimal.

Next, we show that if a heuristic does not estimate z -ratios to be strictly monotonically increasing, then the heuristic is *not* optimal, meaning that it is not guaranteed to schedule all packets correctly. If a function h is not strictly monotonically increasing in ratio estimates, there must exist packets p_A, p_B such that $\frac{c(p_A)}{s(p_A)} < \frac{c(p_B)}{s(p_B)}$ and $\frac{h(p_A)}{s(p_A)} \geq \frac{h(p_B)}{s(p_B)}$. Then it is possible for the system to serve p_B before p_A , an exploitable point that makes the heuristic imperfect. \square

We also provide an example of a class of perfect heuristic: Any heuristic of the form $h(p) = k \cdot c(p)$ for some $k \in \mathbb{R}^+$ will ensure perfect scheduling under WSJF.

We first prove that a heuristic of this form is, in fact, perfect. Given that the expected service order (priority under WSJF scheduling) is in increasing value of $\frac{c(p)}{s(p)}$, while the heuristic's service order will be in increasing value of $\frac{h(p)}{s(p)}$, we want to show that these orderings will always be the same. Consider 2 packets p_A and p_B . It holds that

$$\begin{aligned} \frac{c(p_A)}{s(p_A)} < \frac{c(p_B)}{s(p_B)} &\iff k \cdot \frac{c(p_A)}{s(p_A)} < k \cdot \frac{c(p_B)}{s(p_B)} \\ &\iff \frac{h(p_A)}{s(p_A)} < \frac{h(p_B)}{s(p_B)} \end{aligned}$$

A perfect scheduling order ensures that all of the assumptions from SurgeProtector are maintained, allowing WSJF to provide the same DF upper bound of 1.

A.2 Proof of Theorem 2 (Step Function Heuristic)

Proof. In order for an adversary to achieve a DF greater than k , it must be able to displace innocent packets such that for every byte of data transmitted by the adversary, k times the amount of innocent data is displaced. Equivalently, the innocent packet job size to packet size ratio is a factor of k less than the adversary's. We will prove that the step function guarantees this property by showing the contrapositive to be true.

We first consider an adversarial packet p_A and innocent packet p_I , and assume that $\frac{c(p_A)}{s(p_A)} > k \cdot \frac{c(p_I)}{s(p_I)}$, and we want to show that then the packets cannot be swapped by the scheduler, *i.e.* $\frac{h(p_A)}{s(p_A)} > \frac{h(p_I)}{s(p_I)}$.

$$\begin{aligned} h(p_A) &= a \cdot k^{\lfloor \log_k c(p) \rfloor} \\ &\geq a \cdot k^{\lfloor \log_k (k \cdot c(p_I) \cdot \frac{s(p_A)}{s(p_I)}) \rfloor} \quad (\text{By assumption}) \end{aligned}$$

$$\begin{aligned}
&= a \cdot k^{\lceil 1 + \log_k c(p_I) + \log_k(\frac{s(p_A)}{s(p_I)}) \rceil} \\
&\geq a \cdot k^{1 + \lceil \log_k c(p_I) \rceil + \log_k(\frac{s(p_A)}{s(p_I)})} \\
&= a \cdot k \cdot k^{\lceil \log_k c(p_I) \rceil} \cdot k^{\log_k(\frac{s(p_A)}{s(p_I)})} \\
&= k \cdot h(p_I) \cdot \frac{s(p_A)}{s(p_I)} \\
\implies \frac{h(p_A)}{s(p_A)} &\geq k \cdot \frac{h(p_I)}{s(p_I)} \implies \frac{h(p_A)}{s(p_A)} > \frac{h(p_I)}{s(p_I)}
\end{aligned}$$

Given that an adversarial packet cannot displace innocent packets with a job size to packet size ratio more than a factor of k smaller than the adversarial packet's, we see that the step function heuristic upper bounds the DF at the fixed value of k . □

A.3 Proof of Theorem 3 (Preemption)

Proof. We consider a system that starts with estimated job size $J_p = \epsilon$ for all packets, in order to complete analysis with no assumptions about the quality of job size estimates. We also assume no preemption cost. We first look at a period of T seconds, during which N innocent packets arrive. We can represent the true job sizes j_i of incoming packets in scheduled order, as $S = [j_1, j_2, \dots, j_N]$ where $j_i \leq j_{i+1} \forall i$. The main insight here is that an attacker can exploit the system by injecting adversarial packets such that all innocent packets are *partially* served but displaced (preempted and then never fully served). An adversary's goal thus becomes to "weaponize" innocent packet work by causing the system to serve each innocent packet for some amount of time that is *less* than its true job size.

Since job estimates for all packets increase by factors of 2 of the initial estimate ϵ , for a packet of job size c , the maximum work the adversary can weaponize is $w = \max_{k \in \mathbb{Z}^+} (\epsilon \cdot 2^k)$ such that $w < c$. We see worst-case behavior when the value of $c - w$ is minimized for all innocent packets, such that practically all of the innocent work can be weaponized. So, for a fixed $k \in \mathbb{Z}^+$ (and thus a fixed $w = \epsilon \cdot 2^k$), we consider a scenario where all innocent packets have the same job size $c = w + \delta$, with $\delta \rightarrow 0$. For simplicity, we assume that all packets have a packet size of 1.

Assume that the attacker pushes the system to capacity by using l adversarial packets, each of packet size 1 and a true job size of J_A . In order to minimize l , it is in the attacker's interest to encode as much work as possible in each attack packet. Observe that, as long as service capacity is available, each adversarial packet is guaranteed at least w service time (equivalent to the work served in each innocent packet). Thus, we choose $J_A \geq w$. Since the system must be at capacity in order to displace any traffic, we have:

$$\begin{aligned}
&\overbrace{w \cdot N}^{\text{Weaponized work}} + \overbrace{w \cdot l}^{\text{Adversarial work}} = T \\
l = \frac{T - w \cdot N}{w} &= \lim_{\delta \rightarrow 0} \frac{T - (c - \delta)N}{c - \delta} = \frac{T - cN}{c} = \frac{(T - t)N}{t}
\end{aligned}$$

where $t = cN$ is the cumulative true service time for innocent traffic alone. We also note that given the uniform-sized packets in innocent traffic, we can express the load due to innocent traffic as $\rho = \frac{t}{T}$. We are now able to bound the DF for preemptive WSJF, the innocent traffic displaced relative to the adversarial traffic sent:

$$DF = \lim_{T \rightarrow \infty} \frac{N}{\frac{(T-t)N}{t}} = \lim_{T \rightarrow \infty} \frac{t}{T-t} = \lim_{T \rightarrow \infty} \frac{\frac{t}{T}}{1 - \frac{t}{T}} = \frac{\rho}{1 - \rho},$$

which becomes unbounded as $\rho \rightarrow 1$.

□