# Mitigating Memory Controller Side Channels

by

Peter William Deutsch

BASc, University of British Columbia (2020)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2022



Author .....
Department of Electrical Engineering and Computer Science
August 24, 2022

Certified By .....
Mengjia Yan
Homer A. Burnell Career Development Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by .....
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

**Mitigating Memory Controller Side Channels**

by

Peter William Deutsch

Submitted to the Department of Electrical Engineering and Computer Science
on August 24, 2022, in partial fulfillment of the
requirements for the degree of
Master of Science

## Abstract

Memory timing side channels, where attackers utilize contention within DRAM controllers to infer a victim's secrets, pose an important challenge to secure computation in shared memory environments. Attacks utilizing these side channels are broad and highly effective, as memory controllers offer a shared attack surface across all cores on a machine. Attacks have been demonstrated in the wild to leak cryptographic keys and other secret data, emphasizing the importance of employing mitigations to block the ability of an attacker to leak information. Existing state-of-the-art memory timing side channel mitigations have several key performance and security limitations. Prior schemes require onerous static bandwidth partitioning, extensive profiling phases, or simply fail to protect against attacks which exploit fine-grained timing and bank information.

In this thesis we present DAGguise, a defense mechanism which fully protects against memory timing side channels while allowing for dynamic traffic contention in order to achieve good performance. DAGguise utilizes a novel abstract memory access representation, the Directed Acyclic Request Graph ($r$DAG for short), to model memory access patterns which experience contention. DAGguise shapes a victim's access patterns according to a publicly known $r$DAG obtained through a lightweight profiling stage, completely eliminating information leakage.

We formally verify the security of DAGguise, proving that it maintains strong security guarantees. Moreover, by allowing dynamic traffic contention, DAGguise achieves a 12% overall system speedup relative to Fixed Service, which is the state-of-the-art mitigation mechanism, with up to a 20% relative speedup for co-located applications which do not require protection. We further claim that the principles of DAGguise can be generalized to protect against other types of scheduler-based timing side channels, such as those targeting on-chip networks, or functional units in SMT cores.

Thesis Supervisor: Mengjia Yan
Title: Homer A. Burnell Career Development Professor of Electrical Engineering and Computer Science

*"In the beginning the Universe was created.*

*This has made a lot of people very angry and been widely regarded as a bad move."*

DOUGLAS ADAMS

# Acknowledgments

First of all I would like to thank my advisor, Mengjia Yan, in helping me immensely throughout my journey so far. Starting in the middle of a worldwide pandemic was not exactly how I envisioned things would begin, but Mengjia's unwavering support and guidance (over Zoom and otherwise) has propelled me to new heights. The perspectives I've gained over our countless hours of whiteboarding will no doubt help me in my studies and beyond.

To my labmates Jules, Joseph, Weon, and Yuheng: thank you all for being so awesome to work with. Your voices have inspired me to explore our field to new depths, and have gotten me through plenty of research crises over the last two years. A particular thank you to Yuheng, who provided significant contributions to this thesis, particularly in the development of its security analysis.

Thank you to all of the mentors I've been privileged enough to work with, teaching me the ins-and-outs of research. Thomas, thank you for lending your (seemingly limitless) wisdom to my endeavours. Joel, thank you for always pushing me to look at the bigger picture. Mieszko and Prashant, thank you for believing in me and convincing me to take this plunge into academia to begin with.

To my friends: Sacha, thanks for being a fantastic roommate who's always ready for adventure, and never failing to make me laugh. Kyle, thanks for our lunchtime walks along the Charles, and for being a sounding board when I needed it the most. Chris, thanks for all of the laughs, and your endless supply of barbeque tips. Sue, thanks for your unwavering support, and all of the baby animal photos you send my way. To all of those back at home: Ben, Brock, Grisha, Hunter, Keanu, Leah, Matt, Michael, and Nick, thank you for the 3AM D&D games, and for being there for me all of these years. To my dog Elliot, thank you for being a fantastic pup, and for only destroying a single pair of my shoes.

Last but certainly not least, thank you to my family: Stefanie, Lisa, and Thomas. Thank you for supporting me in all of my endeavours, academic and otherwise. I would not be here without you.

# Contents

# List of Figures

# List of Tables

# Chapter 1

# Introduction

Side channel attacks, a class of attacks that exploits micro-architectural vulnerabilities to breach system security, have become a serious security threat in recent years. An attacker can use such vulnerabilities to steal secrets from a victim by monitoring the side effects of the victim's actions on various structures within a processor [31], including caches [21, 19, 36, 7], branch predictors [9, 1], on-chip networks [34], and memory controllers [33].

In this thesis, we focus on studying memory timing side channels which exploit shared memory controllers, a broad attack surface. Memory controllers are responsible for servicing accesses to main memory (DRAM), serving memory requests according to the timing constraints of the DRAM chips. In a shared computing environment, a single memory controller often schedules and buffers requests from *multiple* security domains, leading to visible contention which can be exploited through a side channel attack. These attacks are practical and highly effective. For instance, Wang et al. [33] have demonstrated that contention on memory buses can be used to extract RSA keys. Pessl et al. [24] have further shown that row-buffer contention can be used to monitor keystrokes and recover user passwords.

As first presented in DAWG [15] and CaSA [5], side channel attacks can be described via a telecommunications analogy. In a side channel attack, there is a *transmitter* (the victim) that modulates a *channel*, with that modulation being detected by a *receiver* (the attacker). When the channel is a cache, the receiver is generally *active*,

i.e., it modulates the channel *itself* in order to detect a transmission. In cache-based channels, this involves preconditioning the channel prior to transmission to detect the modulation, e.g., using Prime+Probe [22, 19]. In a memory timing side channel the channel is a memory controller, and the modulation by the transmitter is memory requests based on secret values that make the memory controller busy. In this case, the active receiver must *concurrently* modulate the channel (opposed to preconditioning it) by emitting memory requests to try to contend with the transmitter's requests. Due to memory queuing and scheduling delays, the latency of the receiver's requests can be affected by the transmitter's traffic patterns. Therefore, the receiver can use the timing information of its *own* memory requests to infer the transmitter's secret.

## 1.1 Mitigation Challenges

Prior work has struggled to efficiently mitigate attacks against memory controllers, broadly exploring two directions. The first approach has been to completely block interference between memory requests emitted by different applications using partitioning techniques such as Temporal Partitioning ($TP$) [33] and Fixed Service ($FS$) [28]. While secure, such approaches incur high performance overheads as they statically partition memory bandwidth across applications and allocated bandwidth can often go under-utilized.

The second explored approach has been to shape the transmitter's requests into a predefined pattern that is independent from the secret, such as demonstrated in Camouflage [40]. Camouflage leverages offline profiling to obtain the *distribution* of timing distances between consecutive memory requests, and then shapes request patterns on-the-fly so that the distance between consecutive requests follows this predefined distribution. Although Camouflage can achieve better performance than Temporal Partitioning [33] and Fixed Service [28], it does not offer the same level of security. As described by its authors [40], Camouflage was designed to hide coarse-grained timing information and only provides security when the attacker's timer resolution is low. Camouflage is also severely limited in its flexibility, requiring prior knowledge

of co-running applications' bandwidth requirements during profiling to determine an optimal shaping distribution.

We observe that no prior work can simultaneously meet the following criteria for an *optimal* secure memory controller:

1. ***Security:*** Completely blocking information leakage to an attacker that uses the latency of its own memory requests to infer a secret.
2. ***Limited Performance Overhead:*** Allowing a dynamic allocation of bandwidth across applications.
3. ***Low Profiling Costs:*** Requiring, at most, a simplistic profiling step which does not need prior knowledge of co-located applications.

## 1.2 Our Proposal: $r$DAGs and DAGguise

This thesis introduces *DAGguise*, an effective defense mechanism that simultaneously satisfies the three requirements above. To accomplish this, DAGguise shapes requests according to a novel memory request representation, the *Directed Acyclic Request Graph* (*rDAG* for short).

### 1.2.1 The Directed Acyclic Request Graph Representation

In an $r$DAG, each vertex represents a memory request which experiences an unknown amount of contention in the memory controller. An edge between two vertices indicates the existence of a timing dependency between the two requests, i.e., the destination request can only be emitted by the core after the source request completes. If there exists no path between two vertices, that implies that the corresponding requests can be emitted in parallel.

As a representation of memory request patterns, $r$DAGs have two appealing properties: *generality* and *versatility*. $r$DAGs are *general* enough to fully describe any fine-grained request pattern, describing the distances between consecutive requests,

timing dependencies between requests, and the requests' memory-level parallelism. Moreover, rather than being a constant representation of memory request timing, $r$DAGs are *versatile*, being able to accommodate unknown latencies within the memory controller. Specifically, when a request in an $r$DAG is delayed due to memory contention, its dependent requests are also delayed. We fully describe the structure and properties of $r$DAGs in Section 4.1.

### 1.2.2 Shaping Memory Requests Using $r$DAGs

The key idea of DAGguise is to shape memory requests into a pre-defined pattern described using an $r$DAG, which we call a defense $r$DAG. Specifically, DAGguise introduces a request shaper between the transmitter and the memory controller. The request shaper works as a proxy agent of the transmitter and disguises the transmitter's request patterns. The shaper buffers requests from the transmitter and emits requests following the timing dependencies described by the defense $r$DAG by either delaying some requests or emitting fake requests. At the cycle when the defense $r$DAG prescribes the need to emit a request, the shaper checks whether any request from the transmitter has been buffered. If such a request exists, that request is sent, otherwise a fake request is generated to maintain conformity with the defense $r$DAG and preserve security.

To achieve better performance, DAGguise is assisted with an offline profiling phase which aims to generate a defense $r$DAG that can match the bandwidth requirements of the program to be protected. The profiling is lightweight and is performed on the transmitter in isolation, without requiring the need to account for any co-running applications.

DAGguise satisfies our three secure memory controller requirements, thanks to the generality and versatility of $r$DAGs. First, DAGguise securely hides memory request patterns, as the memory requests emitted by the shaper are fully dependent on the defense $r$DAG, and completely independent from the transmitter's original request patterns. Since the defense $r$DAG is not dependent on any secrets, even if the receiver can fully reconstruct the defense $r$DAG, it cannot glean any information

from the transmitter.

Second, DAGguise can achieve better performance than TP [33] and FS [28]. Rather than statically allocating bandwidth between applications, DAGguise allows the memory controller to *dynamically* adjust the bandwidth allocation between the shaped requests and co-running applications. For instance, when a co-running application emits an increased number of requests, the requests from the defense $r$DAG suffer from more contention. As these requests and the subsequent requests which are dependent on them are delayed in turn, the shaper's bandwidth utilization naturally reduces accordingly.

Lastly, DAGguise's offline profiling cost is low. Since $r$DAGs are versatile, we only need to independently profile the transmitter to derive a defense $r$DAG which achieves good performance, requiring a far smaller profiling cost compared to Camouflage [40].

We use Rosette [32] to formally verify the security properties of DAGguise, ensuring that the shaped access patterns of a transmitter are *indistinguishable* to any receiver. We evaluate the performance overhead of DAGguise using gem5 [4] and run SPEC benchmarks alongside two security-sensitive applications, DocDist [12] and DNA sequencing [27]. Our results show that DAGguise introduces considerably less performance overhead compared to FS [28]. DAGguise incurs a 10% system slowdown on a two-core system, improving system-wide performance by 6% compared to Fixed Service [28]. We show that DAGguise also scales well compared to Fixed Service, achieving a 12% performance speedup compared to FS on an eight-core machine. Furthermore, DAGguise is area efficient, requiring only $0.037mm^2$ of area to instantiate eight parallel shaper instances.

Note that, while we focus on utilizing DAGguise to mitigate memory timing side channels, the key insights of DAGguise are generalizable. $r$DAGs are a general representation of request patterns for various microarchitectural structures. The principles of DAGguise can be applied to mitigate other types of timing side channels involving schedulers and queues, such as instruction port contention in SMT cores [2].

### 1.2.3 Thesis Contributions

This thesis (based on the original DAGguise paper published in ASPLOS 2022 [8]) makes the following contributions:

- The introduction of the Directed Acyclic Request Graph ($r$DAG) representation, a generalizable and versatile way to describe memory access patterns.
- The design of an effective and performant defense mechanism, DAGguise, demonstrating that it is possible to exceed the performance of state-of-the-art defenses, i.e., fine-grained static temporal traffic partitioning [28], while preserving the same security guarantees.
- A formal analysis of the security properties of DAGguise using Rosette [32], a solver-aided programming framework.
- A detailed performance and area evaluation of DAGguise, demonstrating a 12% speedup over Fixed Service [28] with an area footprint of only $0.037mm^2$.

# Chapter 2

# Background

## 2.1 Memory Basics

In modern computing systems, processors access the main memory system via one or more memory controllers (MCs). A memory controller manages a memory *channel*, organized hierarchically into *ranks* and *banks* [14]. Each memory channel supports multiple ranks, where a rank is a collection of DRAM chips that work in parallel to handle a memory request, e.g., to fill a cache line. Each rank is partitioned into multiple banks. Banks and ranks help support multiple outstanding requests, thus enabling a high degree of parallelism in the memory system. Each bank contains a row-buffer, which caches the data of the most recent request. Under an open-row policy, temporally adjacent accesses to the same DRAM row can hit in the row-buffer. Conversely, a closed-row policy forbids hits in the row-buffer.

Memory requests can interfere with each other's timing at several points within the memory controller [14]. Upon arrival from the last-level cache, memory requests are first buffered in a *transaction queue*. Then, each memory request is converted to a sequence of DRAM commands, which are placed into a *command queue* based on their addresses. These command queues are arranged such that there is one queue per bank or per rank of memory. Finally, depending on the DRAM command scheduling policy, commands are scheduled to the DRAM devices based on resource availability and timing constraints. Command scheduling can vary in complexity, ranging from a

basic First Come First Served (FCFS) policy, to policies that optimize for row-buffer hits or bus direction switches (i.e. by grouping reads and writes together) [14].

## 2.2 Memory Timing Side Channels

Shared memory controllers expose a large attack surface for timing side channel attacks. Compared to pipeline structures and private caches, memory controllers are shared by all processes and virtual machines running on the same chip. Several attacks have already shown the viability of memory timing side channels [24, 33].

A memory timing side channel involves a victim (transmitter) program and an attacker (receiver) program communicating via contention within a memory controller. The transmitter in the victim's security domain emits a sequence of memory requests based on some secret values. The receiver in the attacker's security domain aims to obtain the secret by monitoring the transmitter's memory access pattern. While the attacker cannot *directly* observe the victim's access pattern, it can emit a sequence of memory requests and observe how the victim's accesses interfere with its own as they contend with each other in the shared memory controller.

### 2.2.1 Attack Example

Fig. 2-1 provides an example of how an attacker can discern a victim's detailed request patterns based on the latency of its own requests. In this example, the attacker always follows the same request pattern, emitting a new request a constant amount of time after the previous request completes. The attacker's requests are always mapped to the same bank and the same row. We consider a simplified memory where each request takes $n$ cycles to service and the DRAM uses an open-row policy.

In Fig. 2-1(a), when the victim does not emit any requests, none of the attacker's requests are delayed. In Fig. 2-1(b), when the victim emits a request targeting a different bank from the attacker's requests, one of the attacker's requests is delayed for $\Delta$ cycles due to contention in the transaction queue and the shared memory bus. In Fig. 2-1(c), when the victim emits a request to the same bank and the same

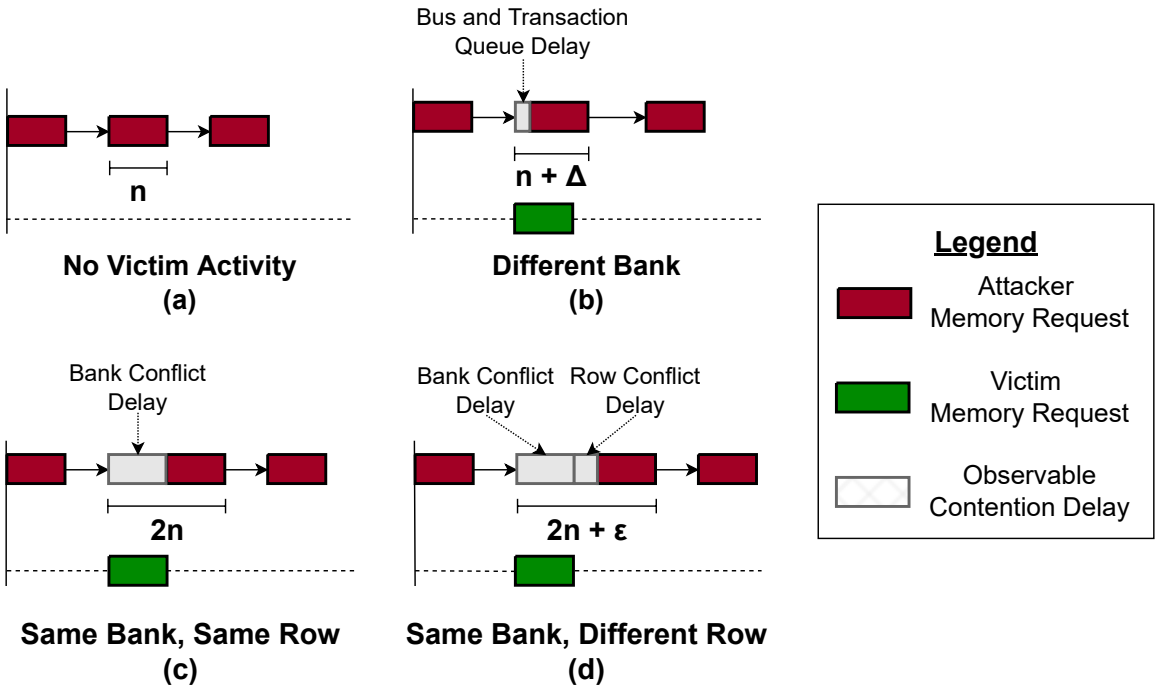


Figure 2-1: Memory timing side channels examples which exploit different types of memory contention. An attacker can discern a victim's detailed memory request patterns based on the latency of its own requests.

row as the attacker's requests, one of the attacker's requests is delayed until the victim's request completes. As a result, the attacker observes a request latency of $2n$ cycles. In Fig. 2-1(d), when the victim's request targets the same bank but a different row compared to the attacker's requests, one of the attacker's requests will suffer an additional $\epsilon$ cycle penalty, i.e., the time for the memory controller to close the current row and open a new row. As demonstrated, an attacker can use its own latency to effectively discern a victim's request patterns, including the number of memory requests, the timing of these requests, as well as their bank and row address information.

## 2.3 Threat Model

We assume the attacker and the victim are in different security domains and the attacker cannot directly access the victim's secret data. The attacker and the victim run on the same machine accessing DRAM via one or more shared memory controllers.

The attacker can be either a user-level application or privileged system software, where the latter case applies to a system with support for enclaves, such as Intel SGX [13], Sanctum [6], or Keystone [18].

The attacker performs a memory timing side channel attack to glean secrets from the victim's domain. As described in Section 2.2, the attacker actively generates requests to interfere with the victim's memory requests and aims to infer the victim's memory request patterns based on its own response latencies. We do not consider physical attacks, that is, where the attacker physically accesses and probes the DRAM bus to *directly* observe the timing, addresses, or even data of memory requests. Such attacks require the attacker to physically possess the attacked device. Similar to other existing defense mechanisms [28, 33, 40], we do not block information leakage due to early termination time. Termination time leakage is intrinsically a program-level issue, and cannot be effectively addressed at the microarchitectural level.

We consider a defense mechanism to be secure if no attacker can distinguish a transmitter's memory request patterns. The memory latencies observed by the attacker should be independent from the victim's actual memory activity. A formal definition of the *indistinguishability* property is provided in Chapter 5.

# Chapter 3

# Motivation

## 3.1 Limitations of Existing Approaches

There exists two directions in mitigating memory timing side channels: partitioning and traffic shaping. We observe that existing mitigation mechanisms [33, 28, 40] along these two directions suffer from several key limitations. In this section, we examine two state-of-the-art defense mechanisms, Fixed Service [28] and Camouflage [40].

### 3.1.1 Fixed Service

Fixed Service (FS) [28] achieves static and fine-grained temporal partitioning by introducing a deterministic schedule for memory requests. Every request is assigned to a certain "slot". Within each slot, a request sequentially passes through the request queues, the command bus, the bank, and the data bus. The slots are pipelined, with a fixed stride inserted between consecutive slots to ensure that each in-flight request uses different resources at any point of time. Therefore, no collisions can occur in any of the shared microarchitectural resources.

The memory controller assigns slots to different security domains using a *round-robin, no-skip* arbitration policy. If a security domain does not have a pending request for its slot, the slot is wasted. This strict partitioning approach completely isolates the memory access patterns of security domains from one another and achieves a

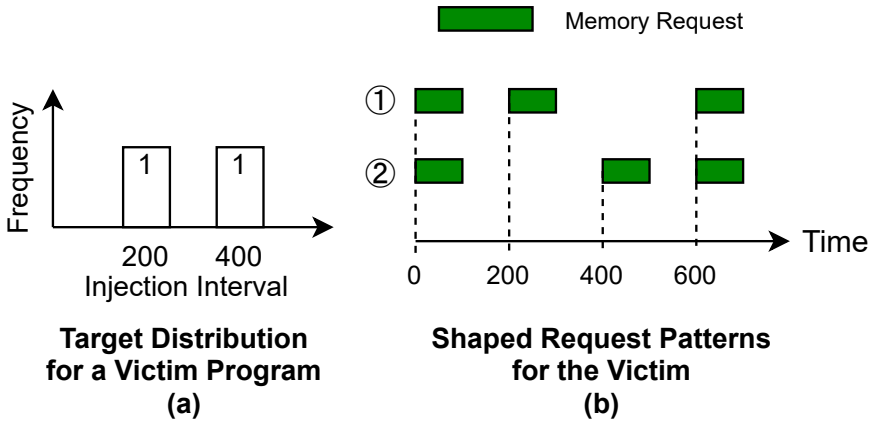


Figure 3-1: A demonstration of how Camouflage cannot hide fine-grained request patterns.

strong non-interference property.

Fixed Service's strict partitioning can often significantly degrade bandwidth utilization, incurring a high performance overhead. For instance, if there are $N$ security domains, one bandwidth-intensive domain and the remaining idling, the memory-intensive domain will only be able to utilize $1/N$ of the total bandwidth, with the remaining bandwidth being wasted.

### 3.1.2 Camouflage

Camouflage [40] is a memory traffic shaping mechanism. It shapes the timing of memory requests to follow a pre-determined distribution which is independent from any victim secrets. Camouflage relies on offline profiling to obtain the distribution that approximately matches the bandwidth utilization of the victim application. Shaping to a distribution is accomplished by selectively delaying existing memory requests, and issuing fake requests when necessary.

Unfortunately, Camouflage does not offer strong security guarantees. Specifically, Camouflage is unable to hide fine-grained memory access patterns, including the *ordering* of memory requests and bank contention. We use an example in Fig. 3-1 to illustrate how distribution-based traffic shaping is insufficient to block information leakage. Assume Camouflage aims to shape the injection time of consecutive victim

requests to the target distribution in Fig. 3-1(a), i.e., one 200-cycle interval and one 400-cycle interval. Under Camouflage, the output of the shaper is not necessarily deterministic. Given different victim request inputs, the shaper can generate two different request sequences as shown in Fig. 3-1(b). Both request sequences ① and ② conform to the distribution in Fig. 3-1(a), but they differ in the *ordering* of the injection intervals. Sequence ① has the 200-cycle interval first and then the 400-cycle interval next, while sequence ② swaps the order of the two intervals. An attacker can use memory timing side channels (Section 2.2) to easily distinguish the two sequences. Moreover, the distribution used by Camouflage does not consider any bank information. Thus, Camouflage is further vulnerable to attacks which exploit bank contention.

Another limitation of Camouflage is that its offline profiling process is both expensive and oftentimes infeasible. The timing distribution of the victim is inherently dependent on co-running applications, as memory contention can slow down the victim program and significantly affect the injection intervals of its memory requests. Therefore, to obtain good performance, the target timing distributions used by Camouflage must not only be tailored to the program being protected, but also to the applications expected to run alongside the victim, significantly increasing the offline profiling cost. Moreover, such a profiling method is completely infeasible if the co-running applications also need protection and the application owners do not want to share any memory bandwidth usage information.

## 3.2 Design Goals

We propose DAGguise to achieve the three design goals as shown in Table 3.1, comparing it with the existing defense mechanisms, Fixed Service [28] and Camouflage [40].

First, our security goal is to block information leakage via fine-grained memory access patterns, including the number of memory requests, the timings of requests, and bank/row information. Second, the defense mechanism should incur a low performance overhead. Different from FS [28], which uses static partitioning which leads

Table 3.1: Design goals of DAGguise and comparison with existing defense mechanisms.

| | Fixed Service [28] | Camouflage [40] | DAGguise (this thesis) |
|---|---|---|---|
| Security | ✓ | x | ✓ |
| Performance Overhead | High | Low | Medium |
| Profiling Cost | – | High | Low |

to significant bandwidth under-utilization, DAGguise can flexibly allocate memory bandwidth among different applications based on each application's actual bandwidth requirements. Finally, we aim to address the substantive profiling issue in Camouflage [40]. DAGguise uses a feasible and lightweight profiling method which only needs to profile the victim application *alone*, without needing knowledge about the bandwidth requirements of potentially co-located programs.

# Chapter 4

# DAGguise: An $r$DAG Request Shaper

We propose DAGguise, an effective defense mechanism to mitigate memory timing side channels. The core idea of DAGguise is to shape memory requests into a predetermined pattern described using a novel graph representation, which we call a *Directed Acyclic Request Graph* or *rDAG* for short. An $r$DAG is general enough to describe any detailed memory request pattern, including those ignored by Camouflage [40], such as the ordering of requests and their bank information. Moreover, $r$DAGs are versatile and can react to contention within the memory controller, helping to adjust memory bandwidth allocation automatically to achieve better performance.

At a high level, DAGguise introduces a request shaper that works as a proxy agent for the transmitter and emits requests following the timing dependencies prescribed by an $r$DAG, which we call a defense $r$DAG. An overview of DAGguise is shown in Fig. 4-1. The shaping operation is achieved by delaying existing requests and emitting fake requests. Note that any secret-independent defense $r$DAG can be used to effectively block information leakage. To achieve better performance, the defense $r$DAG should match the bandwidth utilization of the victim application. We profile the victim application alone to construct a defense $r$DAG by configuring parameters in an $r$DAG template.

In this section, we first introduce the $r$DAG representation in Section 4.1. We then describe the DAGguise scheme through an illustrative example in Section 4.2, demonstrating both the security properties of DAGguise and the versatility of the

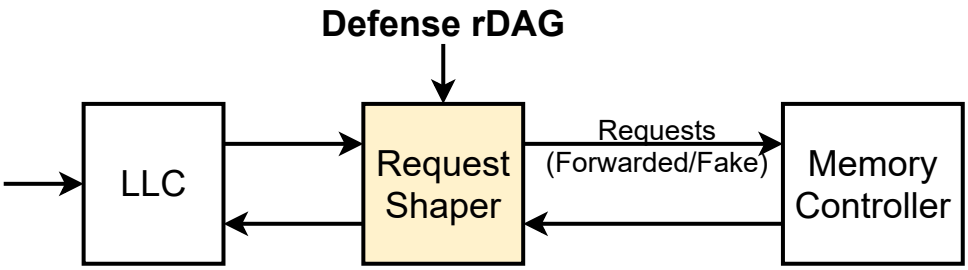


Figure 4-1: DAGguise overview.

$r$DAG representation. We then provide further details of the DAGguise architecture, describing the offline profiling methodology in Section 4.3 and the online shaping mechanism in Section 4.4.

## 4.1 Directed Acyclic Request Graphs ($r$DAGs)

We introduce *Directed Acyclic Request Graphs* (*rDAGs*) to describe memory request patterns. An $r$DAG is a weighted direct acyclic graph that encodes the detailed timing dependencies between memory requests. An example of an $r$DAG is shown in Fig. 4-2. Each vertex represents a memory request. Each edge, connecting two vertices, represents a timing dependency between the two requests. A timing dependency indicates that the destination request can only be emitted after the memory controller finishes serving the source request, e.g., request $v_1$ must be emitted after the response for request $v_0$ leaves the memory controller. If there does not exist a path between two vertices, it means the two corresponding requests can be emitted in parallel, such as request $v_1$ and $v_2$.

To consider possible memory contention, an $r$DAG encodes detailed timing information as follows. First, an $r$DAG encodes two time points for each memory request, its *arrival time* and its *completion time*. The arrival time is the time point when a request arrives at the memory controller and enters the transaction queue; the completion time is the time point when a request has been fully consumed by the memory controller and the response for the request leaves the memory controller. We often conventionally represent the $r$DAG with implicit time flowing from left to right.

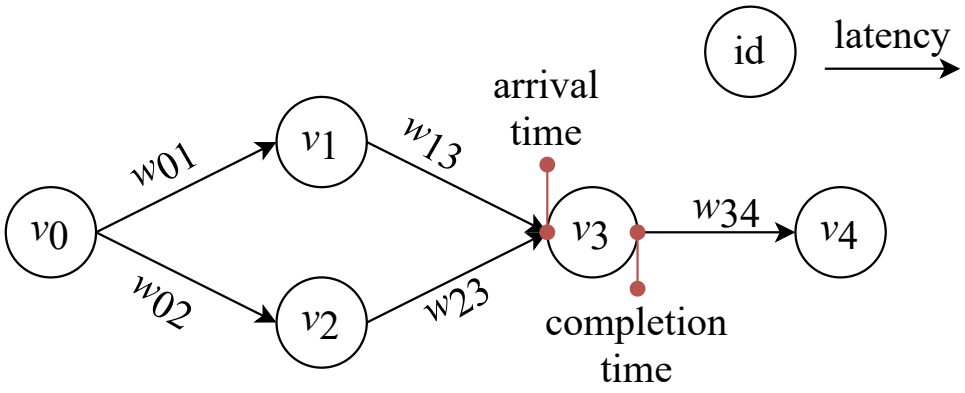


Figure 4-2: An $r$DAG example.

Hence, it is convenient to associate the arrival time with the left side of a vertex, and the completion time with the right side of a vertex, as shown in Fig. 4-2. Note that, due to contention, it can take a *variable amount of time* between when the request arrives at the memory controller and when the memory controller finishes serving the request.

Each vertex is associated with a bank ID and a tag to indicate whether it is a read or write request. This information is included since the memory controller's scheduling policy takes these properties into account in deciding when to serve the request.

Each edge is associated with a weight that measures the latency between the completion time of the source request and the arrival time of the destination request. For example, the weight of the edge connecting vertices $v_0$ and $v_1$ is $w_{01}$, so we have $t_{arrival}(v_1) = t_{completion}(v_0) + w_{01}$.

### 4.1.1 $r$DAG Properties

$r$DAGs have two appealing properties: *generality* and *versatility*. First, an $r$DAG is general enough to represent any fine-grained request pattern, describing the injection intervals between requests (used by Camouflage [40]), timing dependencies between requests, and memory level parallelism. Moreover, an $r$DAG can also describe complex and irregular request patterns generated by real applications.

Second, rather than being a constant representation of memory request patterns, $r$DAGs are versatile, meaning that an $r$DAG can accommodate for unknown memory latencies. Specifically, when a request in an $r$DAG is delayed due to memory contention, its dependent requests will also be delayed. For example, in Fig. 4-2, if the request $v_3$ suffers from bank contention, the completion time of the vertex $v_3$ will be delayed. As a result, the arrival time of the dependent request $v_4$ is also delayed. This versatility property allows an $r$DAG to flexibly represent different memory request injection times.

### 4.1.2 Original $r$DAGs vs. Defense $r$DAGs

A victim's *unshaped* memory request pattern can also be described using an $r$DAG, which we call the *original rDAG*. This original $r$DAG varies with the secret value used by the victim application.

The $r$DAG representation conveniently visualizes the lifetime of memory requests. Specifically, in an original $r$DAG, the time represented by a vertex (i.e. between the vertex's arrival and completion times) corresponds to the time that the request spends within the memory controller, where it may experience contention. The latency indicated by the edge weight represents inter-request timing relationships, corresponding to the time to traverse through the cache hierarchy and perform dependent computations in a core.

A defense $r$DAG is used to describe the memory request patterns that should be emitted by the DAGguise request shaper, which is placed between the LLC and the memory controller (as shown in Fig. 4-1). The request shaper takes a secret-independent defense $r$DAG as input, and shapes the victim's memory requests according to this $r$DAG. Effectively, the shaper *encapsulates* the victim's original $r$DAG inside the defense $r$DAG.

Note that we *do not* need to obtain the original $r$DAG for our defense mechanism to work. We also note that the defense $r$DAG also does not need to closely resemble the original $r$DAG to achieve good performance. We show how to directly obtain a defense $r$DAG via statistical profiling in Section 4.3.

## 4.2 An Illustrative Example

We now describe the DAGguise scheme through an illustrative example, shown in Fig. 4-3, demonstrating both the security properties of DAGguise and the versatility of the $r$DAG representation.

### 4.2.1 Security Properties of DAGguise

In this example, a victim application emits different memory access patterns based a boolean secret value. Fig. 4-3(a) shows the victim's original request patterns using $r$DAGs. When the victim's secret is 0, the application emits one request at a time with a 100-cycle interval between the completion time of each request and the arrival time of its subsequent request. When the secret is 1, the application emits requests slower with a 200-cycle interval between consecutive requests. In both cases, we assume a fixed DRAM latency of 100 cycles. DAGguise works by shaping the two memory request patterns into the same pattern described by the defense $r$DAG in Fig. 4-3(a), making the interval between consecutive requests 150 cycles.

Fig. 4-3(b) demonstrates how each of the victim's request patterns are shaped in accordance with the defense $r$DAG. The first line for each secret represents the victim's original request pattern, corresponding to the victim's original $r$DAG in Fig. 4-3(a). The shaper delays the victim's requests, as shown on the second line for each secret, to match the timing pattern prescribed by the defense $r$DAG. The final request pattern output by the shaper is shown on the third line.

When the secret is 0, the shaper delays each of the victim's requests by 50 cycles to increase the timing interval between requests to 150 cycles (as required by the defense $r$DAG). When the secret is 1, the shaper needs to both delay the victim's requests and issue fake requests. Since the victim issues requests with 200-cycle intervals, and the defense $r$DAG emits requests faster with a timing interval of 150 cycles, the shaper generates a *fake request* when the victim has no outstanding request pending. For example, the second and the fourth requests output by the shaper (on the third line) are fake requests. The victim's actual requests are further delayed to become

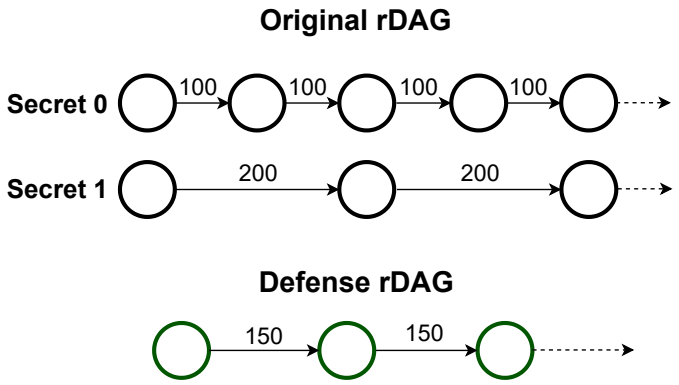


(a) Victim's Request Patterns


(b) Shaping Request Patterns to the Same rDAG

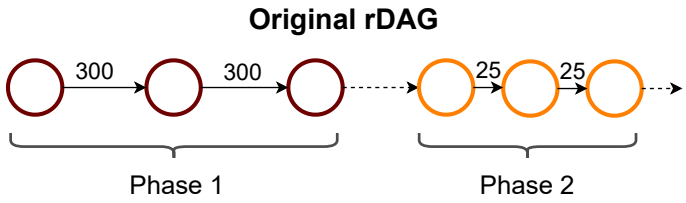


(c) Unprotected Program's Request Pattern

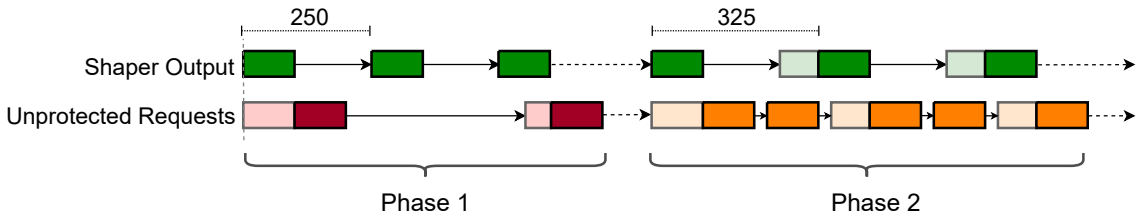


(d) Memory Controller Contention Between Victim and Unprotected Program

Figure 4-3: A running example to demonstrate the security and adaptivity properties of DAGguise.

the third and fifth requests output by the shaper.

Since the requests generated by the shaper have identical timing intervals (shown on the third line) regardless of the victim's secret, and because an attacker cannot differentiate between contention caused by fake and real requests, the shaping scheme employed by DAGguise guarantees that an attacker cannot distinguish between the victim's original request traces. A formal security verification of this property is discussed in Chapter 5.

### 4.2.2 Adaptivity Properties of DAGguise

Figures 4-3(c) and (d) further demonstrate how the versatility property of $r$DAGs helps DAGguise achieve good performance. We continue to use the same victim application and defense $r$DAG as the previous example. In these figures, the victim application protected by DAGguise shares a memory controller with an unprotected co-running application. Fig. 4-3(c) describes the unprotected application with two phases of differing request intervals, modeling a real-world application with varied memory behaviors. Fig. 4-3(d) details the effects of contention between the unprotected application's memory requests and the victim's shaped memory requests.

In phase 1, the unprotected application emits memory requests at a slow interval of 300 cycles, and there is not much contention at the memory controller. As a result, the shaper is able to maintain the victim's ideal injection interval of 250 cycles, i.e., 100 cycles for the memory access latency, and 150 cycles for the timing dependency (i.e., the weighted edge between vertices in the defense $r$DAG).

In phase 2, the unprotected application generates requests at a rapid interval of 25 cycles, causing a large amount of contention at the memory controller and delaying many of the shaper's requests. To maintain the timing dependencies in the defense $r$DAG, the shaper emits the next request 150 cycles after the response of the previous request returns. As a result, due to contention, the injection intervals of the victim's requests in phase 2 are increased from the original 250 cycles to 325 cycles. By slowing down the shaper's emission rate, the scheduler is able to allocate more bandwidth to the unprotected application and achieve better overall memory utilization.

Thanks to the versatility property of $r$DAGs, DAGguise is able to adapt to memory controller contention and adjust its own emission rate, allowing the memory controller to achieve better memory utilization while still maintaining security. Note that, while the example focuses on the case of running a protected application with an unprotected application, DAGguise also works effectively for the case of running multiple protected applications together. In this case, multiple defense $r$DAGs can interact with each other in a similar way as in Fig. 4-3(d), as a "denser" defense $r$DAG can obtain more bandwidth from the memory controller.

## 4.3 Offline Profiling Method

The goal of the offline profiling phase is to find a suitable defense $r$DAG to be used by the memory request shaper. It is important to note that shaping requests to any secret-independent defense $r$DAG will ensure security. The offline profiling step is thus used to optimize for system-wide performance.

DAGguise uses a lightweight two-step profiling method: 1) obtaining an $r$DAG search space by configuring parameters in an $r$DAG template, and 2) profiling the victim application *alone* using different candidate $r$DAGs to select a final defense $r$DAG.

### 4.3.1 Generating an $r$DAG Search Space

Rather than searching the entire space of possible $r$DAGs, we generate an $r$DAG search space by deriving candidate $r$DAGs from an $r$DAG template. The search space can be generated by varying configurable parameters in the $r$DAG template, including the number of parallel sequences, the edge weights, and the write ratio (the frequency of write requests). Note that the template determines the complexity of the $r$DAGs. We intentionally choose templates that follow a regular and repetitive pattern, aiming to simplify the defense $r$DAGs and reduce the hardware overhead of storing and processing the defense $r$DAGs during the online shaping phase (Section 4.4).

Fig. 4-4 shows two examples of $r$DAGs used in DAGguise, as derived from an

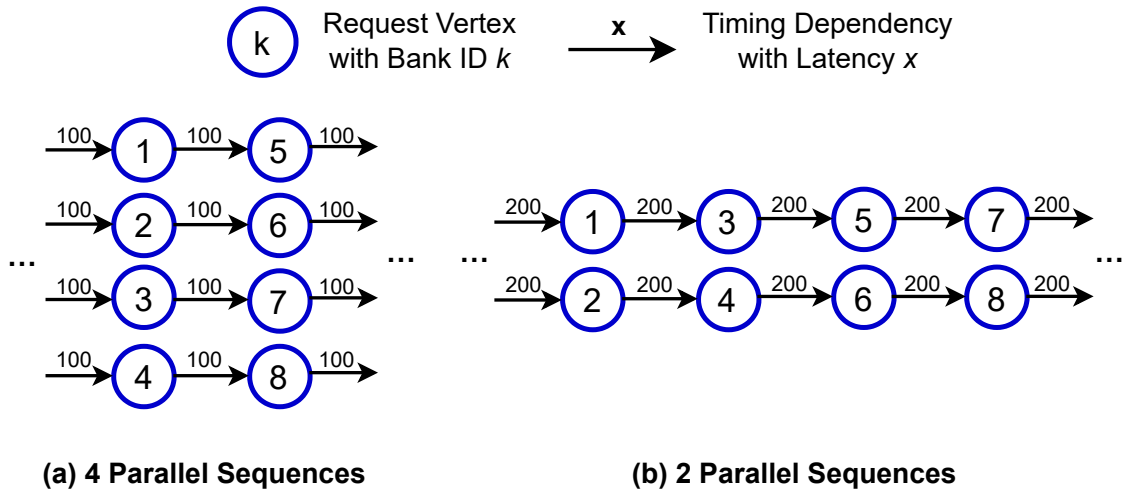


Figure 4-4: Example $r$DAGs used in DAGguise, derived from $r$DAG templates.

$r$DAG template. Fig. 4-4(a) demonstrates an $r$DAG with four parallel sequences (where each sequence contains requests that alternate between two different banks) with uniform edge weights of 100 DRAM cycles. Fig. 4-4(b) demonstrates an $r$DAG derived from the same template when reducing the number of parallel sequences to 2 and increasing the edge weights to 200 DRAM cycles.

### 4.3.2 Selecting a Defense $r$DAG

To select the final $r$DAG from the search space, we test these candidate $r$DAGs on the protected program by feeding each candidate $r$DAG to DAGguise and measuring the impact of DAGguise on the protected program's performance. Intuitively, if we choose a candidate $r$DAG with smaller edge weights and more parallel sequences, the defense $r$DAG becomes denser and thus can request more bandwidth from the memory controller, reducing the amount of bandwidth remaining for co-running applications. In other words, the density of the defense $r$DAG determines the allocated bandwidth to the protected application. To optimize for system-wide performance, we derive the final defense $r$DAG based on the victim program's sensitivity to allocated bandwidth. This presents a trade-off between the protected program's own IPC, and the proportion of memory bandwidth consumed by that program (which is made unavailable to co-running applications).

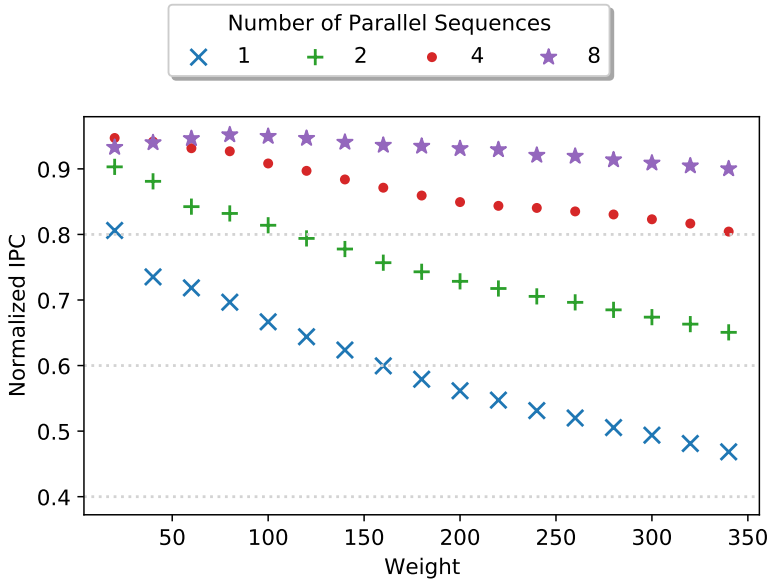


(a) Effect of $r$DAG Edge Weight on IPC

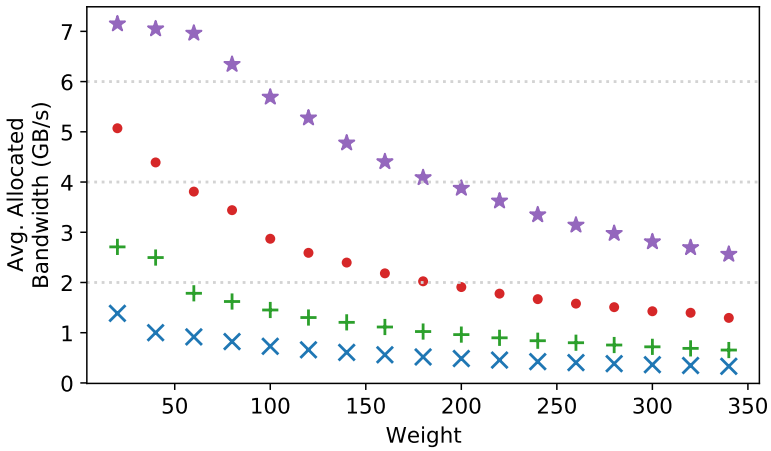


(b) Effect of $r$DAG Edge Weight on Bandwidth Consumption

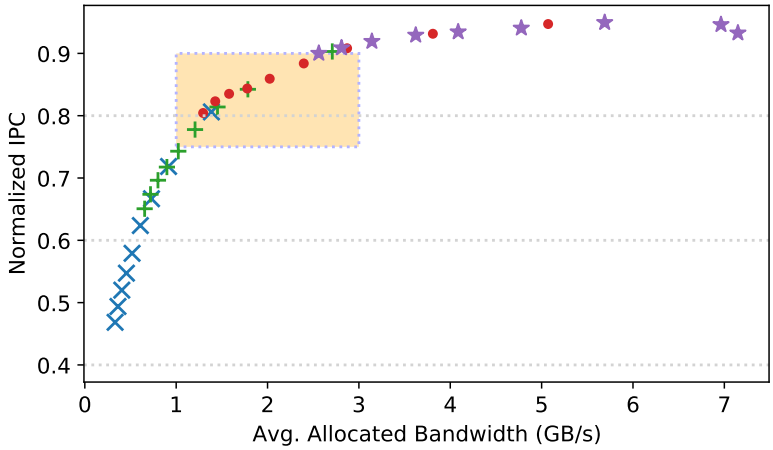


(c) Normalized IPC versus Allocated Bandwidth

Figure 4-5: Selecting a defense $r$DAG for DocDist based on sensitivity to allocated bandwidth.

Fig. 4-5 shows an example of selecting a defense $r$DAG for DocDist, a security sensitive application (see Section 6.1 for details about the experimental setup and DocDist). In this example, a candidate $r$DAG can have 1, 2, 4, or 8 parallel sequences, and a uniform edge weight varying from 0 to 400 DRAM cycles. We run the victim alone, recording the victim's IPC (Fig. 4-5(a)) and memory bandwidth utilization (Fig. 4-5(b)) for each candidate defense $r$DAG. Fig. 4-5(c) combines the IPC and bandwidth results of (a) and (b), demonstrating how the victim program's IPC changes according to the bandwidth allocated to it.

From Fig. 4-5(a) and (b), we observe that as the edge weight decreases and the number of parallel sequences increases, the normalized IPC of the protected program increases, and the allocated bandwidth also increases. From Fig. 4-5(c), we observe that the IPC of DocDist increases quickly when increasing the allocated bandwidth from 0 to 3 GB/s, with a diminishing return after the allocated bandwidth exceeds 3 GB/s. A cost-effective selection of defense $r$DAGs should lie within the highlighted region where the allocated bandwidth is around 1-3 GB/s. Thus, for our evaluation in Chapter 6, we select a defense $r$DAG for DocDist identical to Fig. 4-4(a), comprising of 4 parallel sequences and a uniform edge weight of 100 DRAM cycles. As DocDist is a streaming application which performs very few writes, we set the write ratio (the proportion of vertices in the defense $r$DAG marked as writes) to be small (i.e. $\frac{1}{1000}$). For applications with more varied access patterns, further profiling can be performed to derive an appropriate write ratio which maximizes IPC and minimizes allocated bandwidth.

Note that the actual system-wide performance can vary as co-running applications can have varying bandwidth demands. We heavily rely on the versatility property of $r$DAGs to dynamically adjust the bandwidth utilization for the protected application when running with other applications.

### 4.3.3 Profiling Cost

The profiling cost required by DAGguise is low for two reasons. First, as the $r$DAG representation is versatile, we only need to profile the victim application *alone*, with-

out requiring any information about co-located applications. This makes our approach much more practical than the approach used by Camouflage [40]. Second, we generate candidate $r$DAGs from $r$DAG templates, significantly reducing the search space of defense $r$DAGs.

### 4.3.4 Multithreaded Applications

So far, we have considered how to profile a single-threaded application to generate an optimal defense $r$DAG. For programs with multiple threads belonging to the same security domain, it is possible to utilize a single defense $r$DAG shared across all threads of a program, or multiple defense $r$DAGs with each one being exclusively used by one thread. These two approaches have different trade-offs in profiling cost and system-wide performance. Using a single defense $r$DAG for all threads allows for different threads to share vertices in the defense $r$DAG, reducing the number of fake requests issued and thus increasing system-wide performance. However, this approach may increase the offline profiling cost, since using a single defense $r$DAG may require re-profiling the application for each possible number of concurrent threads. Using one $r$DAG per thread reduces the overall offline profiling cost, in exchange for an increased online performance overhead.

## 4.4 Online Shaping Mechanism

DAGguise uses a request shaper to disguise the transmitter's request pattern. Specifically, the shaper works as a proxy agent for the transmitter and emits requests following the defense $r$DAG by either delaying some of the transmitter's requests or emitting fake requests. Fig. 4-6 shows the memory controller with DAGguise's hardware highlighted.

The baseline memory controller has a transaction queue and multiple command queues which are arranged so that there is one queue per bank. We have discussed how the memory controller manages these queues in Section 2.1.

DAGguise's hardware needs to include the following three components for *each*

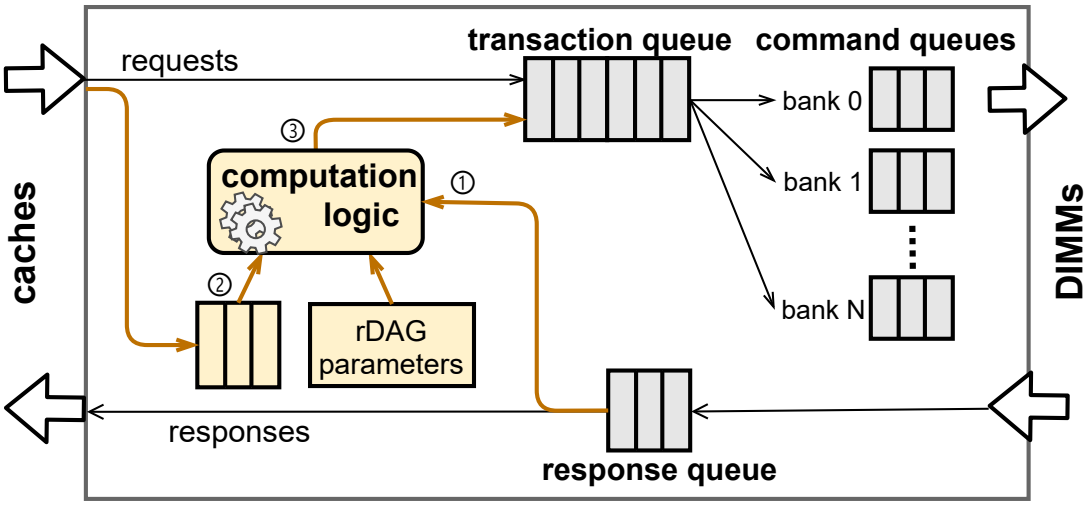


Figure 4-6: DAGguise memory controller architecture.

security domain that needs protection: a private transaction queue, $r$DAG parameter registers, and the shaper logic. To make our mechanism work, every memory request is tagged with a security domain ID. If a memory request is emitted by a domain that is under protection, the request will be inserted into the corresponding domain's private queue, while other requests are directly inserted into the global transaction queue.

We describe the shaper's operations using a single-bank configuration as an example. The shaper logic keeps track of the responses and decides when to emit the next request. When the shaper receives a response for its own security domain (①), the shaper logic computes whether the next request is ready to be emitted. If the next request is ready to be emitted, the shaper checks the private transaction queue to see whether there exists a pending request (②) matching the type of request to be issued (i.e. read or write). In the case that such a matching pending request exists in the private queue, the request will be transferred to the global transaction queue; otherwise, a fake request will be generated and inserted into the transaction queue (③).

In a multi-bank scenario, the access pattern to different memory banks could leak information. Thus, DAGguise needs to ensure that the bank access pattern is the same no matter what secret value is used. Recall that each vertex in an $r$DAG

is associated with a bank ID (Section 4.1). Each step in Fig. 4-6 takes the bank information into account. For example, when the shaper is ready to emit the next request, in addition to checking whether a pending request exists in the private queue, it needs to search the private queue to look for a request with the same bank ID. Similarly, when generating fake requests, the shaper needs to generate the request with the bank ID as prescribed by the defense $r$DAG.

As discussed in Section 2.2, row-buffer hits and misses can also leak information [24]. To hide row-buffer access patterns when using DAGguise, the memory controller must use a closed-row policy, ensuring that DRAM rows are closed immediately after every read or write. It also is possible to encode row-buffer activity in the defense $r$DAG to avoid the overhead of using a closed-row policy. Each vertex could additionally specify whether the memory access is a row-hit, and $r$DAGs with varying row-buffer hit ratios could be explored during the profiling stage. Such a scheme would save costs related to always closing a row, but when a vertex is marked as a row-hit in the $r$DAG and the program actually accesses a different (closed) row, DAGguise would need to emit a fake request to maintain security (negatively impacting performance). We leave further exploration of this direction as future work.

### 4.4.1 $r$DAG Computation Logic

The $r$DAG computation logic is responsible for tracking the execution status of the defense $r$DAG and determining whether a request needs to be emitted by the shaper on a given cycle (and the bank ID/write status of that request if one is required). The complexity of this logic is fully determined by the defense $r$DAG. Recall that the defense $r$DAG is derived from an $r$DAG template following a regular and repetitive pattern (Section 4.3). Consequently, the corresponding computation logic is fairly simple.

For example, to track the status of the $r$DAG template in Fig. 4-4, this logic only needs to track the following states for each bank: a bit to indicate whether the shaper is waiting for a response, a bit to indicate whether the next request is a read or write, and a counter to track the remaining cycles until the next request

is required. We evaluate the area overhead of this computation logic for multiple parallel security domains, in addition to the required supplementary private queue storage in Section 6.4.

### 4.4.2 Fake Requests

When the shaper needs to emit a request but there does not exist a matching pending request, the shaper must insert a fake request into the global transaction queue. The fake request accesses a random address in the targeted bank. Issuing fake requests, however, can incur high energy consumption. Prior work [28] has introduced multiple approaches to address these energy concerns. One possible approach is to "suppress" fake requests. Rather than issuing these requests to the DIMMs, we can update the timing parameters and DRAM states as if the request was actually performed, as the data of these fake requests is irrelevant. An alternative approach is to use the fake requests to do useful work, e.g., issuing prefetching requests. For simplicity, we use the suppression approach in this thesis.

### 4.4.3 Shaper Management

The DAGguise hardware structures, particularly the $r$DAG parameter registers for each security domain, need to be securely managed by privileged software that is part of the system's trusted computing base (TCB). Such software could be a security monitor [6] or microcode [13] in a system with support for enclaves, or the operating system or hypervisor when protecting user-level applications/virtual machines. Specifically, the privileged software is responsible for initializing and clearing the $r$DAG parameter registers when requested. During context switches, the privileged software is required to save and restore the $r$DAG registers, private queue state, and computation logic states.

# Chapter 5

# Security Verification

In this section, we start by describing a formally modeled system of DAGguise, followed by an explicitly defined indistinguishability property. We then provide details on the verification process. We perform our verification using k-induction [29] (a common technique for the verification of transition systems) in Rosette [32], a solver-aided programming language that integrates SMT solvers.

## 5.1 System Modeling

We model the DAGguise system as a state machine, consisting of an $r$DAG request shaper and a memory controller. We denote the simulation state of the DAGguise system at the beginning of cycle $i$ as $S_i$, which includes the states of the shaper, the memory controller, and the buffers between them. We use $S_{\text{reset}}$ to denote the state after a reset operation.

The inputs to the state machine are two memory request traces from a transmitter and a receiver, denoted as $\text{Req}_{Tx}$ and $\text{Req}_{Rx}$ respectively. The transmitter's request trace is passed to the request shaper, while the receiver's request trace is directly passed to the memory controller. The outputs of the state machine are two memory response traces to the transmitter and the receiver, which are denoted as $\text{Resp}_{Tx}$ and $\text{Resp}_{Rx}$. A request/response trace is a vector, with the element at index $i$ describing whether the trace contains a request/response at cycle $i$ and the bank ID of the

request/response. For example, $\mathrm{Req}_{Tx}[i] = (\mathrm{valid}_i, \mathrm{bankID}_i)$.

Given a state $S_i$ at cycle $i$ and request traces for $j$ cycles, $\mathrm{Req}_{Tx}$ and $\mathrm{Req}_{Rx}$, we use the notation $S_i \xrightarrow[\mathrm{Req}_{Tx},\mathrm{Req}_{Rx}]{\mathrm{Resp}_{Tx},\mathrm{Resp}_{Rx}} S_{i+j}$ to denote simulation of the system for $j$ cycles from state $S_i$ to state $S_{i+j}$ that outputs the response traces $\mathrm{Resp}_{Tx}$ to the transmitter and $\mathrm{Resp}_{Rx}$ to the receiver.[1] The transition function is determined by the configuration of the DAGguise system, including the chosen defense $r$DAG and the scheduling policy used by the memory controller.

For demonstration purposes, in our Rosette implementation we model a simplified memory controller that uses a FCFS scheduling policy and a constant memory latency of 2 cycles. The modeled request shaper uses a defense $r$DAG with a sequence of strictly dependent requests. It is feasible to extend our tool to verify whether the security property holds for different $r$DAGs and complex memory controllers.

## 5.2 Security Property

Recall that in Section 2.3, we consider a system as secure if an adversary (the receiver) cannot distinguish between different request traces of the victim (the transmitter) based on its own response latencies. By *indistinguishability* of request traces, we mean that the receiver's response trace $\mathrm{Resp}_{Rx}$ is independent from the transmitter's request trace $\mathrm{Req}_{Tx}$.

We formally define $\mathrm{P}(S_0, n)$, meaning the system achieves indistinguishability when running the system from the state $S_0$ for $n$ cycles.

$$
\begin{aligned}
\mathrm{P}(S_0, n) := \quad & \forall \ \mathrm{Req}_{Tx}, \mathrm{Req}'_{Tx}, \quad \forall \ \mathrm{Req}_{Rx} \\
& \textbf{if} \quad S_0 \xrightarrow[\mathrm{Req}_{Tx},\mathrm{Req}_{Rx}]{\mathrm{Resp}_{Tx},\mathrm{Resp}_{Rx}} S_n \quad \text{and} \quad S_0 \xrightarrow[\mathrm{Req}'_{Tx},\mathrm{Req}_{Rx}]{\mathrm{Resp}'_{Tx},\mathrm{Resp}'_{Rx}} S'_n \\
& \textbf{then} \quad \mathrm{Resp}_{Rx} = \mathrm{Resp}'_{Rx}
\end{aligned}
$$

We verify that $\mathrm{P}(S_0, n)$ holds for an arbitrary $n$ when setting $S_0 = S_{\mathrm{reset}}$. Note that in

[1] This standard notation of state machine transition places the input below the arrow and the output above the arrow.

practice, a request may depend on previous responses, and so may depend on previous requests. This is not a problem as we prove the property for all possible sequences of requests, independently of how they came to be.

## 5.3 Verifying the Security Property Using K-Induction

We use k-induction [29] to verify the security property above. Our verification involves the following two high-level steps:

1) *Base step*: Perform bounded model checking to verify that the security property $\mathrm{P}(S_{\mathrm{reset}}, k)$ holds for a small integer $k$;

2) *Induction step*: Simulate the system from two arbitrary starting states $S$ and $S'$, taking two arbitrary request traces $\mathrm{Req}_{Tx}$ and $\mathrm{Req}'_{Tx}$ for $k$+1 cycles. Assuming the receiver cannot distinguish between the two cases in the first $k$ cycles, check whether the receiver can distinguish them in the $(k+1)$-th cycle.

**Base Step.** To perform the bounded model checking of $\mathrm{P}(S_{\mathrm{reset}}, k)$, we model arbitrary inputs to the system by defining three symbolic vectors to represent $\mathrm{Req}_{Tx}$, $\mathrm{Req}'_{Tx}$, and $\mathrm{Req}_{Rx}$. We then simulate the system symbolically for $k$ cycles and obtain the response traces for the receiver, $\mathrm{Resp}_{Rx}$ and $\mathrm{Resp}'_{Rx}$. We implement the symbolic simulation process in Rosette and call the SMT solver to search for a binding of symbolic vectors to concrete values that violates the assertion $\mathrm{Resp}_{Rx} = \mathrm{Resp}'_{Rx}$.

**Induction Step.** The induction step is equivalent to searching for a violation of the following assertion.

$$
\begin{aligned}
&\forall \;\; Req_{Tx}, Req'_{Tx}, \quad \forall \;\; Req_{Rx}, \quad \forall \;\; S, S' \\
&\textbf{\textit{if}} \;\; S \xrightarrow[Req_{Tx}, Req_{Rx}]{Resp_{Tx}, Resp_{Rx}} S_{k+1} \;\; \textit{and} \;\; S' \xrightarrow[Req'_{Tx}, Req_{Rx}]{Resp'_{Tx}, Resp'_{Rx}} S'_{k+1} \\
&\qquad \textit{and} \;\; Resp_{Rx}[0\!:\!k) = Resp'_{Rx}[0\!:\!k) \\
&\textbf{\textit{then}} \;\; Resp_{Rx}[k] = Resp'_{Rx}[k]
\end{aligned}
$$

Similarly, we use symbolic vectors to represent the request traces and the starting states, $S$ and $S'$. Again we call the SMT solver to search for a binding of symbolic vectors to concrete values that satisfies the assumption $\text{Resp}_{Rx}[0:k) = \text{Resp}'_{Rx}[0:k)$ and violates the assertion $\text{Resp}_{Rx}[k] = \text{Resp}'_{Rx}[k]$.

We follow standard methodology by incrementing the value of $k$ until the induction step succeeds. The minimal $k$ is related to the system's configuration, proportional to the number of cycles needed for a request to traverse the whole system. The time complexity of the verification process increases significantly with the value of $k$. For our specific implementation, the induction step works with $k = 6$, thanks to the simplified configuration of our model. While the verification is performed on a simplified model, the verification process itself is sound, with potential to extend to more complex configurations.

# Chapter 6

# Evaluation

## 6.1 Experimental Setup

To evaluate the performance overhead of DAGguise, we use gem5 [4], a cycle-accurate simulator. We use DRAMSim2 [25] to model the memory controller and DIMMs. Table 6.1 shows the details of the simulated architectures.

### 6.1.1 Experiment Configurations

We compare DAGguise and FS-BTA, a state-of-the-art protection scheme, against an insecure baseline. *FS-BTA* is short for Fixed Service Bank Triple Alternation, a performance optimized variant of Fixed Service [28]. FS-BTA aggressively pipelines requests such that parallel bank accesses can occur under narrow circumstances, while still maintaining non-interference.

The insecure baseline uses an open-row policy, while FS-BTA and DAGguise utilize a closed-row policy to mitigate row-buffer attacks. Both schemes can mitigate the memory timing side channels described in Section 2.2. Note that we do not provide a performance comparison to Camouflage [40] as it does not fully hide bank contention, while most of the performance penalties associated with FS-BTA and DAGguise stem from protecting bank access patterns.

Table 6.1: Baseline architecture configurations.

| **Parameter** | **Value** |
|---|---|
| Multicore | 2 and 8 out-of-order cores at 2.4GHz |
| Core | 8-issue, out-of-order, 192-entry ROB |
| Private L1<br>I-Cache/D-Cache | 32KB each, 64B line, 8-way<br>4-cycle round-trip (RT) latency |
| Private L2 Cache | 256kB, 64B line, 16-way, 13-cycle RT latency |
| Shared L3 Cache | 1MB per core, 64B line, 16-way<br>42-cycle RT latency |
| DRAM Configuration | 4GB (2-core) and 8GB (8-core)<br>1 Channel, 1 Rank/Channel, 8 Banks/Rank<br>Frequency: 1600Mbps |
| DRAM Timing Parameters | $t_{RC} = 39, t_{RCD} = 11, t_{RAS} = 28, t_{FAW} = 24,$<br>$t_{WR} = 12, t_{RP} = 11, t_{RTRS} = 2, t_{CAS} = 11,$<br>$t_{RTP} = 6, t_{BURST} = 4, t_{CCD} = 4, t_{WTR} = 6,$<br>$t_{RRD} = 5, t_{REFI} = 7.8\mu s, t_{RFC} = 260ns$ |

### 6.1.2 Benchmarks

To evaluate the impact of DAGguise on overall system performance, we co-locate victim programs with benchmark applications on separate cores. The sample set of fifteen co-running benchmark applications are selected from the SPEC2017rate suite [17]. For each SPEC application, we utilize the SimPoint methodology [23] to run up to 10 representative intervals of 50 million instructions each to accurately reflect the application's performance [39]. The caches are populated prior to interval data collection using 1 million warm-up instructions, and all simulations are run using gem5's system call emulation mode.

We use two victim programs: Document Distance (*DocDist*) and DNA sequence matching (*DNA*), which process unstructured data and can leak information via memory accesses [38].

DocDist [12] compares documents for similarity, computing the distance between a private input document and a public reference document. DocDist precomputes a feature vector counting the frequency of each word in the reference document. Upon receiving an input document, it first computes a feature vector for that document, then computes the euclidean distance between the input and the reference feature vectors. The access pattern to the feature vectors can leak information.

DNA sequence matching [27] takes a private DNA sequence as input and aligns it with a public DNA sequence. Specifically, the public DNA sequence is divided into substrings and stored in a hash table. To do the alignment, the hash table is searched for common substrings with the private DNA sequence. The access pattern to the hash table can leak information.

We perform system-wide performance evaluations across two system environments. In Section 6.2 we evaluate a two-core system running one protected application and one SPEC benchmark. We extend our analysis in Section 6.3 to evaluate DAGguise's scalability on an eight-core system running four protected domains alongside four co-running SPEC benchmarks.

## 6.2 Performance Overhead

To measure the impact that DAGguise has on overall system performance within a two-core system, we measure the IPCs of each application (DocDist protected by DAGguise, and one SPEC benchmark), and then normalize each IPC to its baseline performance under the insecure configuration (under the same co-location). We then take the *average* of these values to arrive at an average normalized IPC, representing the overall performance of the system, shown in Fig. 6-1.

We observe that in a two-core environment DAGguise has a 10% system slowdown compared to the insecure baseline, while enjoying a modest performance increase over FS-BTA, achieving a relative 6% performance increase.

As a general trend, we note that DAGguise is particularly good at maintaining the performance of co-running applications at the expense of the protected program's performance. In most observed cases, the SPEC program performs better using DAGguise than FS-BTA (20% better, on average), while DocDist does worse (7% worse, on average), resulting in an overall system speedup. For some non-memory-bound benchmarks (such as *leela*) we observe an overall *decrease* in performance, as the additional bandwidth made available to unprotected applications by DAGguise is not used by the benchmark, while the protected program still pays costs for the shaper.

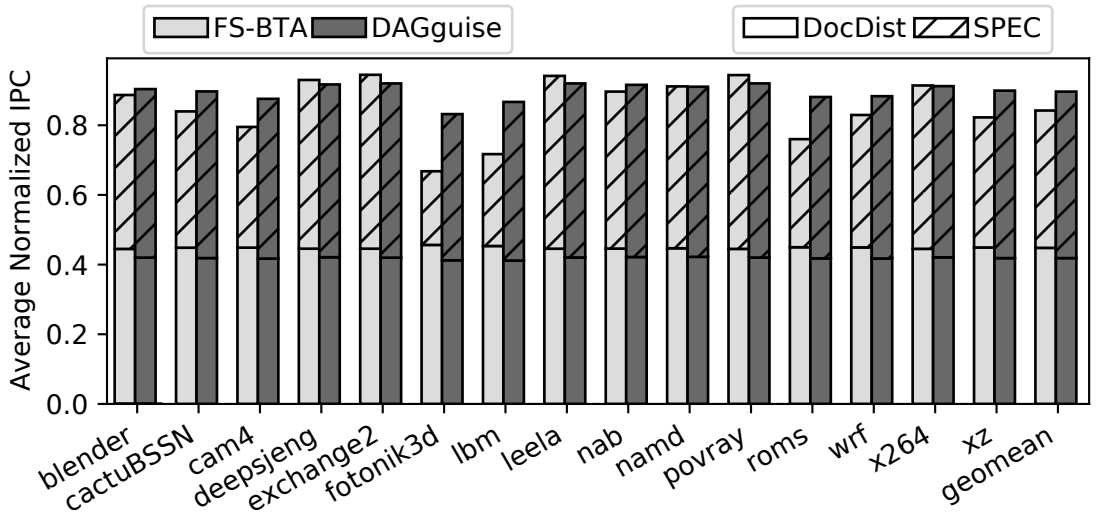


Figure 6-1: Average Normalized IPC running DocDist with one SPEC application on a two-core system.

We hypothesize that additional performance for *protected* programs may be achieved by expanding the $r$DAG search space to consider more complex defense $r$DAGs.

## 6.3 Scalability

In order to demonstrate our system's scalability compared to FS-BTA, we expand our simulation to encompass multiple co-running victim programs alongside unprotected applications. On an eight-core system we use four DAGguise shapers to protect four programs, two copies of DocDist and two of DNA, co-located with four identical unprotected copies of SPEC benchmarks. Under FS-BTA, each individual victim receives $\frac{1}{8}$ of the total number of slots, while the remaining $\frac{4}{8}$ slots are shared amongst the SPEC applications.

The average normalized IPC results for our eight-core experiment are shown in Fig. 6-2. DAGguise encounters a 34% system-wide slowdown compared to the insecure baseline, with an improved 12% average system-wide performance gain relative to FS-BTA. In a heavily provisioned system protected by DAGguise we observe that most applications, not just unprotected ones, achieve a relative speed-up compared to their performance under FS-BTA. This suggests that DAGguise is indeed

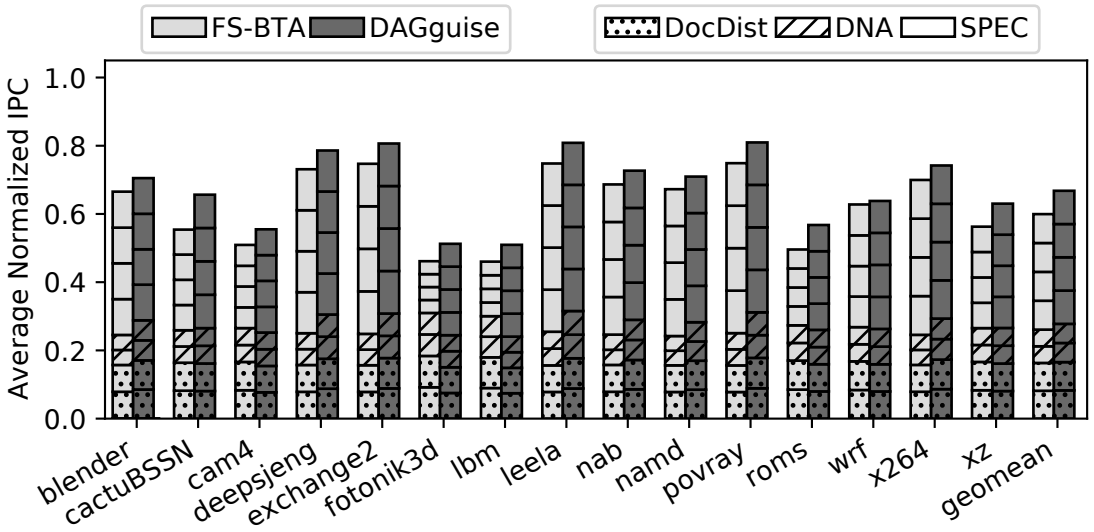


Figure 6-2: Average Normalized IPC of two DocDist, two DNA, and four SPEC processes on an eight-core system.

versatile to complex bandwidth patterns, allowing for better bandwidth utilization compared to FS-BTA.

## 6.4 Area Overhead

To evaluate the area overhead of DAGguise, we compute the combined area of the computation logic (described in Section 4.4.1) and the private transaction queues.

We implement the computation logic in RTL, synthesizing it using the YoSys suite [35] and the 45nm FreePDK45 cell library [16]. We evaluate an eight shaper configuration (allowing for eight independent security domains), each supporting eight banks, 16-bit $r$DAG weights, and eight private queue entries. The implementation supports template $r$DAGs like in Fig. 4-4, allowing 1, 2, 4, and 8-parallel patterns.

To evaluate the SRAM area overhead of the private transaction queues, we use Cacti [3]. Each private queue is sized to match the expected maximum number of parallel memory accesses of a protected program. Each queue entry contains a request's 64-bit address and, if the request is a write, 64B of data. Even if the queue is full, we do not leak any information, as each queue is private to a security domain.

The area of the computation logic and private transaction queues is reported

Table 6.2: Area overhead of DAGguise for 8 protected domains.

| **Component** | **Resources** | **Area ($mm^2$)** |
|---|---|---|
| Computation Logic | 13424 Gates | 0.02022 |
| Private Queue ($8 \times 8$ entries) | 4608 B (72B×64) SRAM | 0.01705 |
| Total | – | 0.03727 |

in Table 6.2, with an eight shaper configuration ultimately requiring a footprint of only $0.037mm^2$.

# Chapter 7

# Related Work

We have discussed most related work in Chapter 3. A few other pertinent works are as follows.

*Temporal partitioning (TP)* [33] divides time into fixed-length periods during which only requests from a given security domain are scheduled, rather than interleaving requests, as in Fixed Service [28]. TP still guarantees non-interference but performs worse than FS, suffering from a static bandwidth allocation.

In addition to FS-BTA, which we compare against in Chapter 6, *Fixed Service* [28] has other variants which perform spatial partitioning at the bank, rank, or channel-level. While these variants can improve performance, they severely limit the number of simultaneous programs and the allowable memory usage of each. DAGguise has no such spatial partitioning requirements.

*Ascend* [10] and its follow-up paper [11] examine a different threat model and consider a passive (non-interfering) attacker probing memory buses to observe request patterns, and tries to obfuscate them using traffic shaping. They rely on a fixed request rate that is changed periodically, resulting in bounded amounts of leakage.

Shaping traffic to secret-independent patterns has also been explored in the area of network timing side channels. Issuing network packets at a fixed rate has been examined to protect SSH traffic [30] and multimedia traffic [26]. Pacer [20] offers a *cloaked tunnel* abstraction which protects against network timing side channels, shaping packets to profiled traffic schedules via software support in the hypervisor/kernel.

# Chapter 8

# Conclusion

In this thesis we introduced DAGguise, an effective defense mechanism against memory timing side channels. DAGguise utilizes Directed Acyclic Request Graphs ($r$DAGs), a novel memory request pattern representation, to shape memory access patterns into secret-independent ones. DAGguise is able to attain formally verified security guarantees while allowing for dynamic traffic contention to achieve good performance, only requiring a lightweight and feasible offline profiling process. DAGguise's insights can further be applied to other scheduler-based side channels which exploit contention in other microarchitectural structures, such as SMT cores and on-chip networks.

## 8.1 Future Work

### 8.1.1 Mitigating Other Types of Side Channels

While this thesis focuses on mitigating memory timing side channels, its key insights (to shape request patterns using an $r$DAG) are even more general. DAGguise can be applied to address a broader range of side channels, such as those that exploit contention in SMT cores [2], on-chip networks [34], cache banks [37], etc. Similar to main memory, these resources are all associated with schedulers which decide the order in which requests are served, such as a pipeline scheduler deciding which instruction will use a functional unit, and a Network-on-Chip scheduler deciding which packet

will use a network link. The scheduler introduces extra latency to some requests due to contention, which can leak information.

The principles of DAGguise can be used to mitigate these *scheduler-based* channels. For example, consider using DAGguise to block leakage via functional unit contention in SMT cores. We can profile the transmitter program and construct a defense $r$DAG, where each vertex represents an instruction's request to use a specific type of functional unit. We then place a request shaper between the decode and the dispatch stages. The shaper emits requests following the defense $r$DAG by delaying instructions and emitting fake instructions. It is promising to improve performance by pairing defense $r$DAGs with complementary functional unit bandwidth requirements on the same core.

### 8.1.2 Using $r$DAGs for Security Analysis

While this thesis focuses on using the $r$DAG representation to describe *shaped* request patterns, $r$DAGs can also be used to describe the *original* request pattern of a given program (as described in Section 4.1.2). Recall that $r$DAGs are general enough to describe any request pattern, and are a versatile representation that is agnostic to contention in a shared microarchitectural resource. These properties make it possible to use $r$DAGs to model a program's detailed request behavior, which can be useful in performing security analyses. In this subsection we explore two examples of using $r$DAGs for security analysis.

First, $r$DAGs can be used to locate the point in a victim program where information leaks via memory timing side channels. We first start by constructing the original memory request $r$DAGs when running a given program with differing secret values. If the derived $r$DAGs for different secret values are identical, the program's memory request patterns are independent from the secret and thus the program is not vulnerable to memory controller side channels. Otherwise, we can locate the position in these $r$DAGs where the program's access patterns begin to diverge, pointing to the information leak.

Second, $r$DAGs can be used to evaluate the security properties of defense mech-

anisms, particularly those which do not completely eliminate leakage (such as Camouflage). Recall that in a scheduler-based side channel attack, the attacker attempts to distinguish between the victim's secret values based on the latencies of its own requests. The attacker's own request patterns can also be represented as an $r$DAG. By modelling interactions between an attacker and victim's original $r$DAGs under a given microarchitectural defense scheme, it may be possible to model what an attacker can observe given a particular attacker $r$DAG, and subsequently determine *how much* this observation leaks about the victim's secret. A space exploration could then be performed, deriving an optimized attacker $r$DAG which *maximizes* the amount of leakage the attacker observes.

# Bibliography


[1] Onur Acıiçmez, Çetin Kaya Koç, and Jean-Pierre Seifert. Predicting Secret Keys via Branch Prediction. In *Cryptographers' Track at the RSA Conference (CT-RSA)*. Springer, 2007.

[2] Alejandro Cabrera Aldaya, Billy Bob Brumley, Sohaib ul Hassan, Cesar Pereida García, and Nicola Tuveri. Port Contention for Fun and Profit. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019.

[3] Rajeev Balasubramonian, Andrew B Kahng, Naveen Muralimanohar, Ali Shafiee, and Vaishnav Srinivas. CACTI 7: New Tools for Interconnect Exploration in Innovative Off-chip Memories. *ACM Transactions on Architecture and Code Optimization*, 2017.

[4] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D. Hill, and David A. Wood. The Gem5 Simulator. *ACM SIGARCH Computer Architecture News*, 2011.

[5] Thomas Bourgeat, Jules Drean, Yuheng Yang, Lillian Tsai, Joel Emer, and Mengjia Yan. CaSA: End-to-end Quantitative Security Analysis of Randomly Mapped Caches. In *Proceedings of the 53th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2020.

[6] Victor Costan, Ilia Lebedev, and Srinivas Devadas. Sanctum: Minimal Hardware Extensions for Strong Software Isolation. In *25th USENIX Security Symposium (USENIX Security)*. USENIX Association, 2016.

[7] Shuwen Deng, Wenjie Xiong, and Jakub Szefer. A Benchmark Suite for Evaluating Caches' Vulnerability to Timing Attacks. In *Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, 2020.

[8] Peter W. Deutsch, Yuheng Yang, Thomas Bourgeat, Jules Drean, Joel S. Emer, and Mengjia Yan. DAGguise: Mitigating Memory Timing Side Channels. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '22, page 329–343, New York, NY, USA, 2022. Association for Computing Machinery.

[9] Dmitry Evtyushkin, Ryan Riley, Nael CSE Abu-Ghazaleh, ECE, and Dmitry Ponomarev. BranchScope: A New Side-Channel Attack on Directional Branch Predictor. In *Proceedings of the 23rd International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, 2018.

[10] Christopher W Fletcher, Marten van Dijk, and Srinivas Devadas. A Secure Processor Architecture for Encrypted Computation on Untrusted Programs. In *Proceedings of the seventh ACM workshop on Scalable trusted computing*. ACM, 2012.

[11] Christopher W Fletchery, Ling Ren, Xiangyao Yu, Marten Van Dijk, Omer Khan, and Srinivas Devadas. Suppressing the Oblivious RAM Timing Channel while Making Information Leakage and Program Efficiency Trade-offs. In *2014 IEEE 20th International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2014.

[12] Hazem Gamal. Document Distance. https://github.com/Hazem-Gamall/document-distance, 2020.

[13] Intel. Intel Software Guard Extensions Programming Reference. https://software.intel.com/en-us/sgx/sdk, 2013.

[14] Bruce Jacob, David Wang, and Spencer Ng. *Memory Systems: Cache, DRAM, Disk*. Morgan Kaufmann, 2010.

[15] Vladimir Kiriansky, Ilia Lebedev, Saman Amarasinghe, Srinivas Devadas, and Joel Emer. DAWG: A Defense Against Cache Timing Attacks in Speculative Execution Processors. In *51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2018.

[16] Jesper Knudsen. Nangate 45nm open cell library. *CDNLive, EMEA*, 2008.

[17] Samuel Kounev, Klaus-Dieter Lange, and Jóakim von Kistowski. The SPEC CPU Benchmark Suite. In *Systems Benchmarking*. Springer, 2020.

[18] Dayeol Lee, David Kohlbrenner, Shweta Shinde, Krste Asanović, and Dawn Song. Keystone: An Open Framework for Architecting Trusted Execution Environments. In *Proceedings of the Fifteenth European Conference on Computer Systems*. ACM, 2020.

[19] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B Lee. Last-Level Cache Side-Channel Attacks are Practical. In *2015 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2015.

[20] Aastha Mehta, Mohamed Alzayat, Roberta De Viti, Björn B. Brandenburg, Peter Druschel, and Deepak Garg. Pacer: Comprehensive network Side-Channel mitigation in the cloud. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 2819–2838, Boston, MA, August 2022. USENIX Association.

[21] Michael Neve and Jean-Pierre Seifert. Advances on Access-driven Cache Attacks on AES. In *Selected Areas in Cryptography*. Springer, 2006.

[22] Colin Percival. Cache Missing for Fun and Profit. http://www.daemonology.net/papers/htt.pdf, 2005.

[23] Erez Perelman, Greg Hamerly, Michael Van Biesbrouck, Timothy Sherwood, and Brad Calder. Using Simpoint for Accurate and Efficient Simulation. *ACM SIGMETRICS Performance Evaluation Review*, 2003.

[24] Peter Pessl, Daniel Gruss, Clémentine Maurice, Michael Schwarz, and Stefan Mangard. DRAMA: Exploiting DRAM Addressing for Cross-CPU Attacks. In *25th USENIX Security Symposium (USENIX Security)*. USENIX Association, 2016.

[25] Paul Rosenfeld, Elliott Cooper-Balis, and Bruce Jacob. DRAMSim2: A Cycle Accurate Memory System Simulator. *IEEE Computer Architecture Letters*, 2011.

[26] T. Scott Saponas, Jonathan Lester, Carl Hartung, Sameer Agarwal, and Tadayoshi Kohno. Devices that tell on you: Privacy trends in consumer ubiquitous computing. In *16th USENIX Security Symposium (USENIX Security 07)*, Boston, MA, August 2007. USENIX Association.

[27] sfu compbio. DNA Sequence Matching. https://github.com/sfu-compbio/mrsfast, 2020.

[28] Ali Shafiee, Akhila Gundu, Manjunath Shevgoor, Rajeev Balasubramonian, and Mohit Tiwari. Avoiding Information Leakage in the Memory Controller with Fixed Service Policies. In *Proceedings of the 48th International Symposium on Microarchitecture (MICRO)*. ACM, 2015.

[29] Mary Sheeran, Satnam Singh, and Gunnar Stålmarck. Checking Safety Properties Using Induction and a SAT-solver. In *Formal Methods in Computer-Aided Design*. Springer, 2000.

[30] Dawn Xiaodong Song, David Wagner, and Xuqing Tian. Timing analysis of keystrokes and timing attacks on SSH. In *10th USENIX Security Symposium (USENIX Security 01)*, Washington, D.C., August 2001. USENIX Association.

[31] Jakub Szefer. Survey of Microarchitectural Side and Covert Channels, Attacks, and Defenses. *Journal of Hardware and Systems Security*, 2016.

[32] Emina Torlak and Rastislav Bodik. Growing Solver-Aided Languages with Rosette. In *Proceedings of the 2013 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software*. ACM, 2013.

[33] Yao Wang, Andrew Ferraiuolo, and G. Edward Suh. Timing Channel Protection for a Shared Memory Controller. In *2014 IEEE 20th International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2014.

[34] Hassan M. G. Wassel, Ying Gao, Jason K. Oberg, Ted Huffmire, Ryan Kastner, Frederic T. Chong, and Timothy Sherwood. SurfNoC: A Low Latency and Provably Non-interfering Approach to Secure Networks-on-chip. In *Proceedings of the 40th Annual International Symposium on Computer Architecture (ISCA)*. ACM, 2013.

[35] Clifford Wolf. Yosys open synthesis suite. https://yosyshq.net/yosys/, 2016.

[36] Yuval Yarom and Katrina Falkner. Flush+Reload: A High Resolution, Low Noise, L3 Cache Side-channel Attack. In *23rd USENIX Security Symposium (USENIX Security)*. USENIX Association, 2014.

[37] Yuval Yarom, Daniel Genkin, and Nadia Heninger. CacheBleed: A Timing Attack on OpenSSL Constant Time RSA. *Journal of Cryptographic Engineering*, 2017.

[38] Xiangyao Yu, Christopher W Fletcher, Ling Ren, Marten van Dijk, and Srinivas Devadas. Generalized External Interaction with Tamper-Resistant Hardware with Bounded Information Leakage. In *Proceedings of the 2013 ACM Workshop on Cloud Computing Security Workshop*. ACM, 2013.

[39] Zirui Zhao, Houxiang Ji, Mengjia Yan, Jiyong Yu, Christopher W. Fletcher, Adam Fletcher, Darko Marinov, and Josep Torrellas. Speculation Invariance (InvarSpec): Faster Safe Execution Through Program Analysis. In *Proceedings of the 53th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2020.

[40] Yanqi Zhou, Sameer Wagh, Prateek Mittal, and David Wentzlaff. Camouflage: Memory Traffic Shaping to Mitigate Timing Attacks. In *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2017.