

# DR-STRaNGe: End-to-End System Design for DRAM-based True Random Number Generators

F. Nisa Bostancı<sup>†§</sup>    Ataberk Olgun<sup>†§</sup>    Lois Orosa<sup>§</sup>    A. Giray Yağlıkcı<sup>§</sup>  
Jeremie S. Kim<sup>§</sup>    Hasan Hassan<sup>§</sup>    Oğuz Ergin<sup>†</sup>    Onur Mutlu<sup>§</sup>  
<sup>†</sup>*TOBB University of Economics and Technology*    <sup>§</sup>*ETH Zürich*

*Random number generation is an important task in a wide variety of critical applications including cryptographic algorithms, scientific simulations, and industrial testing tools. True Random Number Generators (TRNGs) produce cryptographically-secure truly random data by sampling a physical entropy source that typically requires custom hardware and suffers from long latency. To enable high-bandwidth and low-latency TRNGs on widely-available commodity devices, recent works propose hardware TRNGs that generate random numbers using commodity DRAM as an entropy source. Although prior works demonstrate promising TRNG mechanisms using DRAM, practical integration of such mechanisms into real systems poses various challenges.*

*We identify three key challenges for using DRAM-based TRNGs in current systems: (1) generating random numbers with DRAM-based TRNGs can degrade overall system performance by slowing down concurrently-running applications due to the interference between RNG and regular memory operations in the memory controller (i.e., RNG interference), (2) this RNG interference can degrade system fairness by causing unfair prioritization of applications that intensively use random numbers (i.e., RNG applications), and (3) RNG applications can experience significant slowdown due to the high latency of DRAM-based TRNGs.*

*To address these challenges, we propose DR-STRaNGe, an end-to-end system design for DRAM-based TRNGs that (1) reduces the RNG interference by separating RNG requests from regular memory requests in the memory controller, (2) improves fairness across applications with an RNG-aware memory request scheduler, and (3) hides the large TRNG latencies using a random number buffering mechanism combined with a new DRAM idleness predictor that accurately identifies idle DRAM periods.*

*We evaluate DR-STRaNGe using a comprehensive set of 186 multi-programmed workloads. Compared to an RNG-oblivious baseline system, DR-STRaNGe improves the performance of non-RNG and RNG applications on average by 17.9% and 25.1%, respectively. DR-STRaNGe improves system fairness by 32.1% on average when generating random numbers at a 5 Gb/s throughput. DR-STRaNGe reduces energy consumption by 21% compared to the RNG-oblivious baseline design by reducing the time spent for RNG and non-RNG memory accesses by 15.8%.*

## 1. Introduction

Random numbers are used in a wide range of applications, such as cryptographic key generation, authentication, scientific simulations, Monte Carlo methods, and industrial testing [10, 12, 30, 42, 57, 83, 97, 99, 141, 178, 194]. These applications often require a high-throughput random number generator to achieve high performance [82, 138, 182].

Random number generators are categorized into two classes [31, 94, 170, 173]. First, true random number generators (TRNGs) [7, 14, 21–23, 33, 41, 57, 60, 65–67, 78, 93, 100, 118, 122, 141, 146, 149, 166–168, 171, 174, 176, 181, 186, 191] harness the entropy resulting from inherently random physical processes (e.g., electrical noise, clock jitter, Brownian motion, and atmospheric noise) to generate random numbers. Second, pseudo-random number generators (PRNGs) [16, 119, 121, 123, 161] use a seed value to produce a deterministic stream of numbers that appears to be random without the knowledge of the seed.

Random number quality is critical for security applications such as authentication and key generation [20, 30, 34, 38, 42, 57, 71, 83, 95, 97, 99, 116, 117, 178, 194]. Since PRNGs output a deterministic stream of numbers, they are not preferable for security-critical applications [30, 36, 178]. Instead, these applications use TRNGs, because unlike PRNGs, TRNGs do not rely on predictable seed values that can compromise application security.

DRAM is widely available in almost all computer systems and can be easily integrated into mobile and IoT devices as main memory. To enable low-cost and high-throughput true random number generation in almost all commodity systems and emerging processing-in-memory (PIM) architectures, prior works [11, 44, 60, 78, 82, 138, 166, 168] propose various DRAM-based TRNGs that exploit manufacturing process variations in DRAM and the resulting entropy to generate true random numbers. However, no prior work provides an end-to-end system design to enable DRAM-based true random number generation in real systems.

We identify **three key challenges** these proposals face in an end-to-end system design that integrates a DRAM-based TRNG. First, performing random number generation in DRAM can be intrusive in current systems that use DRAM as main memory. It can cause a significant slowdown on concurrently-running applications due to the interference between RNG and non-RNG memory requests in the memory request scheduler (i.e., RNG interference). Second, existing memory schedulers schedule memory requests to achieve high system fairness and performance. However, they are RNG-oblivious and ignore different characteristics of RNG and non-RNG memory requests. RNG requests are received in bursts and served together because interleaving RNG and non-RNG requests induce additional overhead from frequently modifying timing parameters. Existing memory schedulers can prioritize RNG requests to achieve high throughput to serve the high RNG demand. This can increase the memory stall time of non-RNG applications more than RNG applications and degrade system fairness. Third, random number generation in DRAM has a high latency and can cause applications that use random numbers intensively (i.e., *RNG applications*) to experience long memory stall times.

DRAM-based TRNG mechanisms need to perform multiple DRAM reads to gather enough random bits from DRAM. Random number generation can stall the processor’s instruction window if later instructions depend on the generated random number. Therefore, RNG applications can suffer from long memory stall times even when they are the only application running on the system.

**Our goal** is to design an end-to-end system for DRAM-based TRNGs with low cost and high performance that (1) minimizes the slowdown of both RNG and non-RNG applications by reducing and controlling the interference between them, (2) improves system fairness by reducing the memory stall time experienced by non-RNG applications due to random number generation, and (3) mitigates the performance degradation of RNG applications due to the high latency of DRAM TRNG mechanisms. To this end, we propose DR-STRaNGe, a new end-to-end System design for **DRAM-based True Random Number Generators**. DR-STRaNGe has three major components. First, DR-STRaNGe implements a *buffering mechanism* to reduce both the interference between RNG and non-RNG applications and the high latency of a DRAM-based TRNG. Our buffering mechanism comprises (1) a DRAM idleness predictor that leverages a DRAM channel’s idle time to generate random numbers with low interference to the system, and (2) a random number buffer in the memory controller to store the generated values in advance to mitigate the TRNG latency. Second, DR-STRaNGe implements an *RNG-aware scheduler* that improves system fairness and reduces memory stall time by reducing the RNG interference. Our RNG-aware scheduler uses an RNG request queue to separate RNG requests from non-RNG requests, and prioritizes requests based on the priority of the applications set by the operating system (OS). Third, DR-STRaNGe provides an *interface to applications* that enables them to use the system’s DRAM-based TRNG with low latency and high throughput. DR-STRaNGe is independent of the DRAM-based TRNG mechanism used in the system, and is compatible with all previously proposed DRAM-based TRNGs.

**Summary of Results.** Our experimental evaluations across a variety of real workloads show that DR-STRaNGe improves (1) performance of both non-RNG and RNG applications by 17.9% and 25.1% on average, respectively, compared to the RNG-oblivious baseline design and (2) system fairness by 32.1% on average when an RNG application with a 5 Gb/s random number generation throughput requirement runs concurrently with non-RNG applications. DR-STRaNGe provides 21% energy reduction over the RNG-oblivious baseline design by reducing the total time spent for RNG and non-RNG memory accesses by 15.8%. DR-STRaNGe incurs a minor area overhead of  $0.0022\text{mm}^2$  at 22nm technology node (0.00048% of an Intel Cascade Lake CPU Core [187]). This paper makes the following key contributions:

- DR-STRaNGe is the first work that overcomes three key challenges of DRAM-based TRNGs and proposes a viable end-to-end system design for DRAM-based TRNGs.
- We show that an efficient random number buffering mechanism can hide the high TRNG latency and reduce RNG interference on the system. We propose the first random number buffering mechanism with a lightweight DRAM idleness predictor that can predict idle periods with high accuracy (80%). DR-STRaNGe generates random bits during predicted idle periods and fills the buffer, such that, random numbers,

when requested, are served with low latency from the random number buffer.

- We propose the RNG-Aware memory request scheduler, a priority-based scheduler design that improves system fairness and reduces memory stall time by reducing the RNG interference. The scheduler achieves this by differentiating RNG requests from regular requests and employing different scheduling techniques for them to minimize the slowdown of high-priority applications.
- We evaluate the performance, fairness, and energy consumption of DR-STRaNGe, across a variety of real workload mixes, showing significant performance, fairness and energy benefits over the baseline RNG-oblivious system design. DR-STRaNGe improves performance of non-RNG applications, system fairness, and energy by 17.9%, 32.1%, and 21%, respectively, compared to the commonly-used baseline memory request scheduler design.
- We evaluate DR-STRaNGe using two state-of-the-art DRAM-based TRNG mechanisms: D-RaNGe [82] and QUAC-TRNG [138]. We show that DR-STRaNGe is compatible with these mechanisms and improves system performance, fairness and energy with both TRNG mechanisms.

## 2. Background

We provide a brief background on DRAM organization, memory schedulers, and DRAM-based TRNGs.

**DRAM Organization.** DRAM-based main memories are accessed via a memory channel which is internally connected to multiple *banks*. Each bank contains a two dimensional array of DRAM cells, organized as *rows* and *columns*. Upon a request, DRAM cells are fetched at row granularity and the fetched row is temporarily stored in a buffer (*row buffer*). A request is serviced faster if the requested data is already in the row buffer (i.e., *row buffer hit*) [77]. For more information on DRAM, we refer the reader to extensive prior work [15, 25–29, 55, 62, 63, 79–81, 85, 86, 88, 91, 92, 103, 104, 106–108, 113, 114, 130, 142, 147, 152, 154, 155, 193].

**Memory Request Scheduling.** The First Ready First Come First Serve (FR-FCFS) [151, 197] scheduling policy prioritizes 1) row hits over all other requests, 2) then older requests over the younger ones. FR-FCFS policy aims to maximize throughput by leveraging the row buffer. However, it unfairly prioritizes the applications that have high row locality and high memory intensity [128, 135].

Previous works [8, 43, 54, 89, 90, 135–137, 148, 163, 164, 175] propose application-aware memory request schedulers, which improve both performance and fairness across applications. These schedulers monitor different characteristics and schedule based on a maintained ranking [43, 54, 89, 90, 135, 136] or use blacklisting [163, 164] based on application behavior. Ranking based schedulers can unfairly slow down low-ranked applications, and have significant hardware complexity. The blacklisting scheduler (BLISS) [163, 164] is a simple memory scheduler design that categorizes applications into only two groups based on sensitivity to inter-application interference. No prior work considers the high latency of random number generation or requests with different latencies due to non-standard timing parameters.

**DRAM-based True Random Number Generators.** Prior research proposes various DRAM-based TRNGs that leverage the randomness in DRAM timing failures [11, 82, 138], retention failures [60, 78, 166] and start-up values [44, 168].

Retention failure- and start-up value-based TRNGs are limited in the throughput they can provide [82] because random retention failures can take minutes to accumulate at room temperature [58, 113, 114, 142, 147, 177] and random start-up values are formed following expensive DRAM power-cycles [75]. As such, these mechanisms cannot be used in a streaming manner to generate true random numbers.

DRAM TRNG mechanisms based on timing failures can be used in a streaming manner, and they provide random numbers at high throughput ( $> 100 \text{ Mb/s}$ ). These TRNGs operate by deliberately violating DRAM timing parameters that are defined by the DRAM manufacturers to guarantee correct operation in DRAM arrays when obeyed by the memory controller. Random DRAM timing failures can be induced quickly using carefully engineered and timed valid sequences of DRAM commands on commodity DRAM devices [11, 82, 138].

To generate random numbers in DRAM, prior work [82] reserves rows that contain RNG cells, which can be used to extract random numbers by reading with reduced timing parameters. The throughput of the TRNG depends on the number of RNG cells per row. The latency of the TRNG depends on the latency of the DRAM command sequence used to induce random errors. Random number generation can be an intrusive procedure to system operation, especially when the required random number throughput is high, due to two reasons. First, random number generation requires multiple reads to multiple reserved rows. Second, due to non-standard timing parameters, DRAM becomes unavailable for regular memory requests to ensure reliability of data residing in other rows.

### 3. Motivation and Goal

TRNGs are used in a wide range of security applications such as cryptographic key generation, authentication, generation of initial and data padding values, and countermeasures against hardware attacks [30, 42, 57, 83, 88, 97, 99, 178, 194]. The quality of random numbers is important to ensure system security in presence of attacks that aim to obtain confidential user data [30, 178]. Emerging security protocols (e.g., quantum key distribution protocols [34, 116]) provide stronger security guarantees in the presence of attacks targeting weak random numbers. These protocols require very high true random number throughput, in the order of several  $\text{Gb/s}$  [182].

High throughput DRAM-based TRNGs have the main advantage of availability over other TRNGs that typically require dedicated hardware. DRAM devices are widely available in most computer systems and in emerging processing-in-memory architectures [3–6, 9, 17–19, 27, 32, 39, 46, 47, 50–53, 53, 56, 61, 64, 68, 69, 84, 98, 105, 110, 115, 127, 133, 133, 134, 139, 140, 143, 152–159, 165, 180, 192]. DRAM-based TRNGs can provide true random numbers to security-critical applications that run on these systems at high throughput. The impact of integrating DRAM-based TRNGs to these systems is twofold. First, DRAM-based TRNGs can enable security applications on mobile and IoT devices, which becomes more critical with increasing demand for user data privacy. Second, for PIM architectures, DRAM-based TRNGs can improve (1) system efficiency by enabling large contiguous code segments to be executed in memory, and (2) system security by enabling security tasks to run completely in memory.

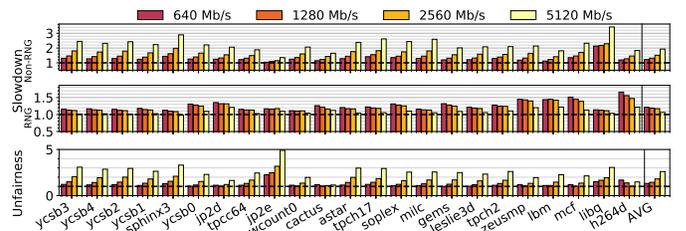
Previous research proposes a set of high-throughput DRAM-based TRNGs. However, no prior work provides an end-to-end system design for integrating DRAM-based TRNGs into

real systems. We identify three key challenges for integrating DRAM-based TRNGs into a baseline *RNG-oblivious* real system: (1) the interference of RNG and non-RNG applications in the memory controller causes unnecessary slowdowns for both types of applications, (2) the unfair prioritization of RNG applications that require high TRNG throughput degrades system fairness, and (3) the high latency of DRAM-based TRNGs degrades the RNG applications’ performance.

To demonstrate the impact of using DRAM-based TRNGs in an RNG-oblivious system, we simulate a realistic two-core system with a DRAM-based TRNG mechanism. The RNG-oblivious system generates 64-bit random numbers by changing DRAM timing parameters to induce errors in the reserved rows [82] as applications request random values. During random number generation, the system stalls regular memory requests because changed timing parameters can induce errors on real data residing in other rows. The system uses all memory channels in parallel to achieve the minimum RNG latency, and minimizes the time that regular requests are stalled. It is possible to use only one channel at a time. However, since the system does not know which channels will be idle at any given time, this can cause further slowdowns by blocking one busy channel for a longer time period. Using all channels and all banks is important to minimize the interference of RNG while serving the random number requests as quickly as possible.

For our tests, we create 172 two-core workloads, each consisting of one RNG and one non-RNG application. For RNG applications, we use four different synthetic benchmarks that request random numbers with required RNG throughput values of  $640 \text{ Mb/s}$ ,  $1280 \text{ Mb/s}$ ,  $2560 \text{ Mb/s}$ , and  $5120 \text{ Mb/s}$ . For non-RNG applications, we use 43 single-core applications from the following benchmark suites: SPEC CPU2006 [2], TPC [172], STREAM [125], MediaBench [48], and YCSB benchmark suite [35].

Figure 1 (top and middle) shows the execution time of non-RNG and RNG applications running concurrently, normalized to each application running alone. Figure 1 (bottom) shows the unfairness index of the overall system for the same two-core workloads. Unfairness index is calculated as the ratio of the maximum memory-related slowdown experienced by an application in the workload to the minimum memory-related slowdown [49, 128, 135], explained in detail in Section 7. An unfairness index of 1 means all applications experience the same memory-related slowdown. Higher unfairness index values indicate that one or more applications are unfairly prioritized by the memory scheduler.



**Figure 1: Slowdown of non-RNG (top) and RNG (middle) applications running concurrently and the unfairness index of the two-core system (bottom) for various RNG throughput requirements. AVG is across 172 workloads.**

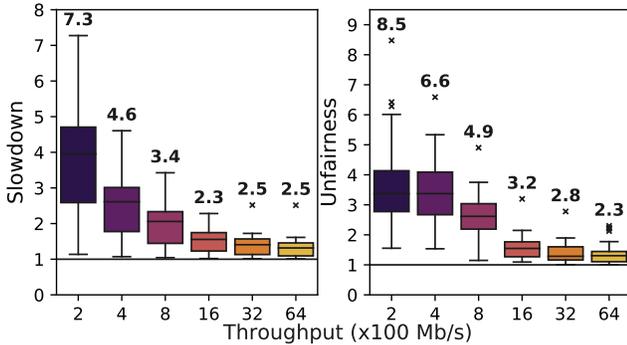
**Impact of the Interference.** Figure 1 (top and middle) shows that both non-RNG and RNG applications experience significant slowdowns due to the interference of both types of applications in the memory scheduler.

The slowdown of non-RNG applications increases as the required RNG throughput increases. On average, non-RNG applications experience 93.1% slowdown when the required RNG throughput is 5 Gb/s. Figure 1 (middle) shows that RNG applications experience slowdowns due to sharing the main memory with another application. We show that the least and the most RNG intensive applications experience 21.4% and 6.2% average slowdown, respectively, compared to when they run alone.

**Unfair Prioritization.** Figure 1 (bottom) shows the unfairness index. We show that the unfairness indices of applications grow significantly with increasing RNG throughput requirement. On average, workloads have an unfairness index of 1.32 at an RNG throughput requirement of 640Mb/s and it gradually increases to 2.61 when the required RNG throughput increases to 5120Mb/s.

**Impact of the RNG Latency.** We observe that the RNG applications spend up to 58.8% of their execution time in random number generation when they require high RNG throughput, due to the high RNG latency.

**Impact of the TRNG Throughput.** We simulate 6 different DRAM-based TRNG mechanisms with different TRNG throughput values ranging from 200 Mb/s to 6.4 Gb/s. Note that two state-of-the-art DRAM-based TRNGs, D-RaNGe [82] and QUAC-TRNG [138], can provide  $\sim 563$  Mb/s and  $\sim 3.44$  Gb/s average random number throughput, respectively, given a state-of-the-art system configuration [138].<sup>1</sup> We evaluate 43 two-core multi-programmed workloads, each consisting of one RNG and one non-RNG application. Figure 2 plots the distribution of (1) the slowdown of non-RNG applications (left), and (2) the system fairness (right) across all 43 workloads.



**Figure 2: The effect of DRAM TRNG throughput on non-RNG application performance (left) and system fairness (right).**

Each box in the figure represents the interquartile range of the observed slowdown ratios and unfairness indices of a dual-core baseline system when the system is provided with the TRNG throughput displayed on the x-axis. The midlines of the boxes indicate the median value of the corresponding interquartile ranges. Cross (×) marks represent the outlier values that fall outside the interquartile range (by more than  $1.5\times$  the length of

<sup>1</sup>All designs assume low latency values based on D-RaNGe’s latency [82] to show only the effect of TRNG throughput. Note that QUAC-TRNG’s latency is higher than what is assumed in this figure. Therefore, the real slowdown and unfairness indices of a system with QUAC-TRNG [138] would be higher.

the box away from the upper quartile). We label the maximum observed slowdown and the unfairness index above each box in the figure.

We make two observations. First, workloads experience significant performance degradation and high system unfairness when executed on systems with state-of-the-art DRAM-based TRNGs. Even the highest-throughput existing TRNG causes a 39.9% average slowdown and degrades average fairness by 28.5%. Second, performance and system fairness do not improve significantly even with future DRAM-based TRNGs that can potentially provide higher throughput. The maximum slowdown and unfairness we observe across all workloads begin to saturate as TRNG throughput reaches 3.2 Gb/s.

**Our goal** in this paper is to develop a low-cost and high-performance end-to-end system design for DRAM-based TRNGs that (1) reduces the interference between RNG and non-RNG applications, (2) improves system fairness across RNG and non-RNG applications, and (3) hides the high latency of DRAM TRNGs that RNG applications suffer from.

#### 4. Design Opportunities

In this section, we discuss the design opportunities of an end-to-end system for DRAM-based TRNGs.

**Random Number Buffering.** Previous work [82, 138] assumes the memory controller can generate random numbers when DRAM utilization is low and uses extra bandwidth to fill a buffer of random numbers. This assumption is not complete without a DRAM idleness predictor since the controller does not know for how many cycles memory channels will be available to generate random numbers with low or no overhead. A buffering mechanism combined with a DRAM idleness predictor is needed to accurately identify the idle periods where the system can generate random numbers with minimal performance impact.

**Memory Request Scheduling.** The latency of RNG is an important factor in system slowdown because it increases the memory stall time of non-RNG requests of both RNG and non-RNG applications. An RNG-aware scheduler is needed to control both types of applications’ memory stall times by managing when to generate random numbers.

**Application Interface.** The system needs to expose an interface so that applications can communicate with the DRAM-based TRNG. The interface can be designed in many ways, including memory-mapped configuration status registers (CSRs) [188], other existing I/O interfaces (e.g., x86 IN and OUT instructions) or new specialized ISA instructions. The interface needs to be exposed to user applications via system calls or static/dynamic library API functions.

#### 5. DR-STRaNGe

DR-STRaNGe is a high performance and low cost end-to-end DRAM-based TRNG system, which (1) reduces the RNG interference that degrades overall system performance, (2) improves the system fairness across RNG and non-RNG applications, and (3) hides TRNG latency to improve the performance of RNG applications. DR-STRaNGe consists of three components. First, its *Random Number Buffering Mechanism* (Section 5.1) aims to hide the high TRNG latency by utilizing idle DRAM cycles to generate random numbers with low interference on other running applications. The buffering mechanism achieves this by predicting and utilizing the idle periods in DRAM channels where it can generate random numbers in small batches. The buffering mechanism then stores the generated random numbers

in a small buffer in the memory controller and serves incoming random number requests from this buffer. Second, the RNG-aware Scheduler (Section 5.2) aims to reduce the memory stall time of regular memory requests by scheduling RNG requests in an order that reduces the RNG interference and improves system performance. The scheduler accumulates RNG and non-RNG memory requests in separate memory request queues and chooses from which queue to schedule a request based on the priority levels of concurrently-running applications. Third, applications use DR-STRaNGe’s Application Interface (Section 5.3) to utilize DRAM-based TRNGs. DR-STRaNGe achieves its design goals by combining these three components.

**DR-STRaNGe overview.** Figure 3 shows an overview of DR-STRaNGe, where (1) the black track shows the flow of a random number request, and (2) the blue track shows how DR-STRaNGe fills the random number buffer. Figure 4 shows the flowcharts of these two tracks.

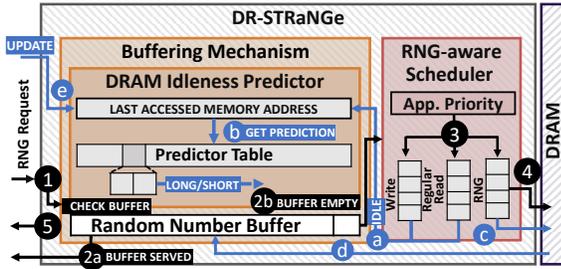


Figure 3: DR-STRaNGe Overview.

The mechanism works as follows. There are two different execution modes in the memory controller: 1) *Regular Execution Mode* processes only regular memory requests, and 2) *RNG Mode* handles RNG read requests and gathers random numbers. The memory controller is initially in Regular Execution Mode.

When the memory controller receives an RNG request, DR-STRaNGe checks the random number buffer (1). If the random number buffer has enough random bits, DR-STRaNGe serves the request from the buffer with low latency (2a). Otherwise, if the buffer does not have a sufficient number of random bits, DR-STRaNGe enqueues an RNG request to the memory request queue (2b). The RNG-aware scheduler determines the memory request scheduling order by checking the concurrently running applications’ priorities (3). Before DR-STRaNGe schedules RNG requests, it first schedules the older regular read requests that belong to high-priority non-RNG applications. After these requests are scheduled, the memory controller switches to the RNG mode, schedules RNG requests and generates random bits in DRAM (4). When enough random bits are gathered, DR-STRaNGe serves the random number request (5) and the memory controller switches back to the Regular Execution Mode.

Every cycle, DR-STRaNGe determines the idleness of each DRAM channel by checking the number of requests in memory request queues (a). If a channel’s memory request queues are empty, DR-STRaNGe predicts the idle period length (b). When the DRAM idleness predictor finds a sufficiently long idle period to generate random bits with no or low overhead, DR-STRaNGe enqueues and schedules RNG requests (c). DR-STRaNGe switches to the RNG Mode and generates random bits by utilizing all banks of the selected channel (d). If the channel remains idle after random number generation, DR-STRaNGe continues to fill the random number buffer. Ran-

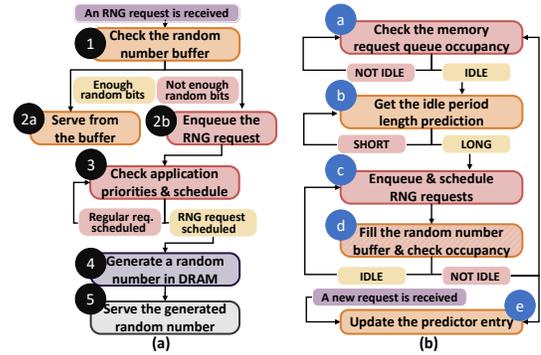


Figure 4: DR-STRaNGe Flowchart.

dom number generation stops when there is a new regular read request to the selected channel or when the buffer is full. DR-STRaNGe switches to the Regular Execution Mode and updates the predictor table corresponding to the selected channel (e). When the random number buffer is full, DR-STRaNGe stops generating random numbers for the random number buffer until a random number is served from the buffer.

### 5.1. Random Number Buffering Mechanism

The goal of the random number buffering mechanism is to hide the high RNG latency by generating random numbers before they are requested. This improves system performance by (1) reducing the memory stall time experienced by RNG applications due to RNG latency and (2) mitigating the RNG interference in the memory controller. The key idea is to use idle DRAM cycles to generate random numbers and store them in a small buffer in the memory controller. When the random number buffer is not empty, DR-STRaNGe quickly serves random number requests from the buffer with low overhead. Many applications often do not fully utilize the DRAM bandwidth [37, 102, 109, 162], and thus DR-STRaNGe can use the idle DRAM bus cycles for RNG. We refer to multiple consecutive idle cycles as *idle periods*. Lengths of idle periods differ between applications based on the observed memory access pattern.

The main challenge of the buffering mechanism is determining the length of an upcoming idle period because the future memory accesses of an application are unknown. However, several heuristics can be employed to predict the idle period length. DR-STRaNGe uses two different approaches: (1) assuming every idle period is sufficiently long to generate 8 random bits, and (2) predicting the length of the next idle period based on idle period length history.

Figure 5 plots the distribution of the idle DRAM period lengths of medium and high memory intensity applications that we evaluate. The distribution is represented as a box-and-whiskers plot where the y-axis shows the length of the observed idle DRAM periods in DRAM cycles. The straight horizontal line represents the time needed to generate a 64-bit random number, which is 198 memory cycles (990 ns) on average in our test setup that we describe in Section 7.

Each box represents the interquartile range of the observed DRAM idle periods and the middle line shows the median. We observe that, for many applications, a significant part of the idle periods do not pass the threshold of 198 cycles. This means that the idle periods are not sufficiently long to generate 64-bit random numbers. Therefore, generating random numbers in smaller batches is preferable to achieve low interference with applications’ memory accesses.

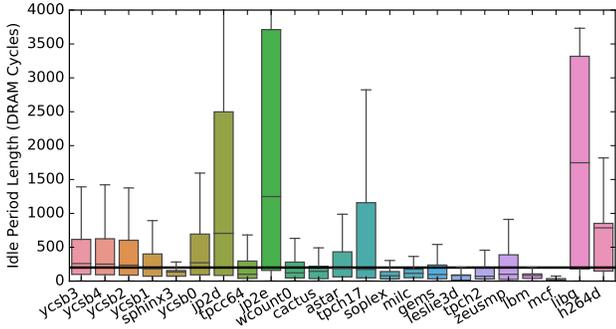


Figure 5: Distribution of DRAM Idle Period Lengths

We design DR-STRaNGe to quickly generate at least eight random bits whenever a channel is idle. To leverage bank-level parallelism, we simultaneously access all DRAM banks and read at least one random bit from each bank during an idle period.

Generating random numbers in small batches does not solve the problem entirely due to existence of idle periods that are even shorter than the time interval required to generate an 8-bit random number (40 cycles). Generating random numbers during these short idle periods can degrade system performance by stalling regular memory requests. Therefore, the mechanism needs to avoid generating random numbers during short idle periods. Our buffering mechanism consists of two parts: (1) a lightweight DRAM idleness predictor that identifies idle periods that are at least 40 cycles long (i.e., *long idle periods*) and (2) a small buffer that stores random numbers. Since the DRAM idleness predictor adds area and time complexity to the design, we evaluate the random number buffering mechanism with and without the idleness predictor in Section 8.

**5.1.1. Simple Buffering Mechanism.** Our simple buffering mechanism does *not* predict for the length of idle DRAM periods and leverages *every* idle cycle of a DRAM channel to generate random numbers and fill the buffer.

With the simple buffering mechanism, DR-STRaNGe selects a channel for RNG when the channel has no regular request in read and write queues. When a channel is selected, DR-STRaNGe puts the channel into the RNG Mode and gathers at least 8 random bits. After generating at least 8 bits, DR-STRaNGe switches back to the Regular Execution Mode if the random number buffer is full or the memory scheduler has new regular memory requests. Regular requests that are at the memory scheduler while the channel is in RNG mode are served after switching to the Regular Execution Mode.

**5.1.2. DRAM Idleness Predictor.** The main goal of the DRAM idleness predictor is to accurately identify sufficiently long idle DRAM periods for RNG. We design two DRAM idle period length predictors that differ in their prediction mechanisms, accuracy and complexity.

**Simple DRAM Idleness Predictor.** We propose a lightweight DRAM channel idleness predictor that uses last accessed memory addresses to predict idle period lengths. Our predictor maintains a table of 2-bit saturating counters, a register for *last accessed address value*, and a counter for *idle period length* initialized as 0 for each channel. A channel’s predictor table is accessed with the *last accessed memory address* when its memory request queues are empty. We group idle periods into two categories: (1) *long* (# of cycles  $\geq$  *PeriodThreshold*) and

(2) *short* (# of cycles  $<$  *PeriodThreshold*). The predictor treats the idle period as *long* if the last accessed address’s counter is 2 or larger and otherwise as *short*.

During idle periods, DR-STRaNGe increments the *idle period length* counter by one every cycle. When a channel receives a regular memory request, DR-STRaNGe updates the predictor table of the corresponding channel as follows. First, DR-STRaNGe retrieves the predictor table entry for the last accessed memory address to access the saturating counter. Second, if the observed idle period length is at least as long as the *Period Threshold* (empirically determined as 40 cycles), DR-STRaNGe increments the saturating counter by one. Otherwise, it decrements the saturating counter. Finally, DR-STRaNGe sets the *idle period length* to zero and updates the *last accessed memory address* with the new request’s address.

The accuracy of the predictor affects DR-STRaNGe’s performance in two ways. First, if the predictor mispredicts a short idle period as a long idle period, which we refer to as a false positive, the interference between RNG and non-RNG applications can increase. Second, if the predictor mispredicts a long idle period as a short idle period, which we refer to as a false negative, it wastes an idle period; hence, future random number requests can experience the full RNG latency if the random number buffer is empty. The predictor cannot starve RNG applications by predicting all idle periods as *short* because DR-STRaNGe serves random number requests by generating random numbers on demand when the random number buffer is empty.

We observe that the predictor often predicts idle periods as short due to the large number of short idle periods. This results in a small number of idle periods used for random number generation. Hence, DR-STRaNGe provides limited performance gain by using the simple DRAM idleness predictor. To increase the RNG opportunities, we propose a method to utilize the periods when a DRAM channel has low utilization (i.e., the memory request queue is largely empty) to generate random numbers for the buffer.

We augment the predictor to determine if a DRAM channel has low utilization based on a threshold, called the *low utilization threshold*. To determine if a DRAM channel has low utilization, the predictor first checks the memory request queue occupancy and determines DRAM utilization to be low if the the number of requests in the memory request queue is less than the low utilization threshold value (empirically determined as 4). Second, when the DRAM utilization is low, the predictor accesses the last accessed memory address’s entry in the predictor table and predicts the length of the low utilization period. If the simple DRAM idleness predictor predicts the period to be long, DR-STRaNGe stalls the regular read queue and generates random numbers for the random number buffer. The predictor stalls only a small number of requests as the predictor does not trigger RNG if there are more requests than the *low utilization threshold* value in the request queue. Some regular memory requests experience higher latency due to RNG but serving the RNG requests with lower TRNG latency improves the overall system performance.

**Reinforcement Learning Agent for DRAM Idleness Prediction.** As a second predictor, we design a reinforcement learning agent to predict the DRAM idle period lengths and create a model by defining the DRAM idleness prediction problem as a reinforcement learning problem.

Reinforcement learning (RL) is a commonly employed technique to solve architectural optimization problems such as branch prediction [196], memory scheduling [74, 131], prefetching [13, 144], dynamic voltage swing control for I/O communication [40], and garbage collection [76]. Among RL techniques, Q-learning [185] is a preferable method due to its simplicity and model-free nature, making it a practical and effective solution. A Q-learning-based RL agent maintains a state machine, where performing a certain action ( $a$ ) at a particular state ( $s$ ) has a value, called Q-value, denoted as  $Q(s, a)$ . These Q-values represent the future cumulative reward of each action for each state and they are stored in a look-up table for fast access. At a given state ( $s$ ), the algorithm chooses the action ( $a$ ) with the largest Q-value,  $Q(s, a)$ . We define two possible actions in our mechanism: 1) to initiate random number generation procedure and 2) to wait. In our implementation, the state is defined as the last accessed address whose least significant 10 bits are XOR'ed with the history of last 10 idle periods, where a logic-1 represents a *long* idle period, and a logic-0 means a *short* idle period.

After the agent takes an action ( $a$ ) based on the state ( $s$ ) of the environment, the Q-value of that state-action pair ( $Q(s, a)$ ) is updated when the reward  $r$  is determined. Reward  $r$  depends on the idle period length and the action taken according to the RL-agent's prediction and it is added to the Q-value of the state-action pair. Positive rewards are used when the agent correctly predicts the idle period length and generates random numbers in long idle periods, or decides to wait in short idle periods. In contrast, negative rewards are used when the agent mispredicts and causes more interference by initiating the RNG procedure in short idle periods (false positives) or waits in long idle periods (false negatives). An action's reward is calculated at the end of the idle period when the agent can observe the idle period's length and determine the correctness of its prediction.

The agent can observe the state only when a DRAM channel is idle. The next state cannot be determined before taking an action because it depends on future memory accesses. Therefore, we omit the next state's part in the update function. The Q-value update function is  $Q(s, a) = (1 - \alpha)Q(s, a) + \alpha * r$  where  $\alpha$  is the learning rate. A high learning rate enables fast adaptation to changes in the access pattern but it is susceptible to noise. Based on our experiments, we observe that the agent achieves peak performance when the learning rate is 0.05.

## 5.2. RNG-Aware Memory Request Scheduler

DR-STRaNGe cannot serve all random number requests from the buffer due to the distribution of idle DRAM periods. Memory request scheduling becomes more critical when a random number is requested, the buffer is empty, and there are RNG and regular read requests waiting in the queues. The goal of the RNG-aware scheduler is to (1) schedule RNG requests without stalling other requests significantly and (2) improve system fairness.

RNG requests can (1) block regular memory requests when a single memory request queue is used for both types of requests and (2) cause the memory controller to switch between the Regular Execution Mode and the RNG Mode frequently. Since RNG and regular memory requests are different types of requests, it is intuitive to have separate queues for each type. DR-STRaNGe uses and maintains an additional queue, *the RNG queue*, for RNG requests. Two queues reduce contention for queue space between the two types of requests.

**5.2.1. Configuring The Scheduler Using Application Priorities.** The operating system (OS) manages hardware resources based on priority levels of applications, and it is possible to use these priority levels for memory request scheduling. The RNG-aware scheduler can use these priority levels to prioritize either the RNG queue or the regular read queue. However, these priority levels are insufficient to employ application-level scheduling decisions because the regular read queue is shared across all applications. The scheduler needs to identify the applications that use both the RNG and regular read queues (i.e., *RNG applications*). DR-STRaNGe marks an application as an RNG application when it requests a random number for the first time. After identifying RNG and non-RNG applications, the scheduler can use the priority bits set by the OS and prioritize one queue over the other. The deprioritized queue can suffer from longer memory stalls, leading to starvation. The RNG-aware scheduler uses a set of rules to enforce priorities while preventing starvation, explained as follows.

**RNG Prioritized.** When an RNG application is prioritized over other applications, DR-STRaNGe prioritizes the RNG application's RNG requests to (1) alleviate the slowdowns incurred due to RNG and (2) reduce the RNG interference. The scheduler chooses the RNG queue over the regular read queue when both queues are non-empty and an RNG application with a request has higher priority than any non-RNG application with a request. The scheduler chooses the RNG queue until the last request in the queue is scheduled. When DR-STRaNGe completes the high-priority random number generation request, the scheduler starts to schedule requests from the regular read queue, preventing the starvation of non-RNG applications with a request. However, as a result of the required RNG throughput, the scheduler can choose the RNG read queue once it receives another high-priority RNG request.

**Non-RNG Prioritized.** The scheduler chooses the regular read queue over the RNG queue and aims to minimize the memory stall time of non-RNG requests when both queues are non-empty and a non-RNG application with a request has higher priority than any RNG application with a request. The scheduler only chooses the RNG queue when the oldest request in the regular read queue is (1) from an RNG application, and (2) received after the oldest RNG request in the RNG queue. In this case, the scheduler chooses the RNG queue until the older RNG requests are scheduled to prevent starvation of RNG applications. After RNG requests are scheduled, it prioritizes the regular read queue again.

**Equal Priorities.** If two RNG applications with an RNG request have the same priority level, the scheduler schedules the older RNG request first. If both RNG applications have regular memory requests, the scheduler schedules them by following the baseline scheduling policy. Similarly, in between two non-RNG applications' requests with the same priority level, the scheduler follows the rules of the baseline scheduler. However, if one RNG application with an RNG request and one non-RNG application with a regular request have the same priority, DR-STRaNGe prioritizes the RNG requests to minimize the RNG interference. Section 8.5 shows that DR-STRaNGe does not degrade the performance of non-RNG applications or system fairness when the RNG queue is prioritized.

**Starvation Prevention Mechanism.** There can be extreme cases that can lead to starvation regardless of the RNG-aware scheduling rules we explained. An RNG application that re-

quires a high RNG throughput can fill the RNG queue frequently. If this application is prioritized, then the regular read requests can experience long memory stalls. Similarly, a high-priority memory-intensive non-RNG application can fill the regular read queue frequently and can cause RNG applications to experience long memory stalls. The RNG-aware scheduling rules might not be sufficient to identify these cases. Therefore, we employ a starvation prevention mechanism.

The prevention mechanism works as follows. DR-STRaNGe observes which memory request queue is deprioritized due to the priority levels of running applications and prevents the scheduler from stalling the deprioritized queue for a long time period. RNG-aware scheduler increments a counter, *the stall time counter*, when it chooses a memory request queue based on the higher priority level of an application, which means the other queue is deprioritized and stalled. It then compares the stall time counter to a threshold value, *the stall limit*, and if the stall time counter reaches the stall limit, the scheduler chooses a request from the deprioritized queue. This prevents high-priority applications from starving other types of applications. In every cycle that a request from the deprioritized queue is scheduled or when the priority levels are changed, the stall time counter is set to 0.

When we set the stall limit as 100 memory cycles in our evaluation, none of the workloads cause the stall time counter to reach the stall limit. We observe that the RNG-aware scheduler's rules are sufficient to prevent starvation even when a memory-intensive non-RNG application or an RNG application with a high RNG throughput requirement is prioritized.

### 5.3. Application Interface

An interface between the software and the DRAM-based TRNG mechanism is needed to integrate the mechanism into the system. In Linux systems, this can be done by changing the existing interface of the kernel's random number generator [57]. The random number generator uses environmental noise from device drivers and gathers bits in an entropy pool. The *getrandom()* system call is used when an application asks for random numbers, which fills a buffer passed to it with a pointer using the random bits in the entropy pool.

Our proposed interface changes the *getrandom()* system call to use DR-STRaNGe. This can be done by setting memory-mapped configuration status registers or using other existing I/O datapaths based on the target system. When a request is made, DR-STRaNGe serves the request either from the random number buffer (see Section 5.1) or by generating random numbers (see Section 5.2).

When the random number buffer is not empty, DR-STRaNGe's system call does not incur any overhead over the baseline system call. However, when the random number buffer is empty, the overhead is related to the difference between the latencies of gathering random bits with the available RNG of the baseline system (e.g., the Linux kernel's built-in RNGs [57]) and DRAM-based TRNGs. Since the first one depends on available devices that can be used as entropy sources, this overhead depends on the target system.

## 6. Security Analysis of DR-STRaNGe

**Secure Random Numbers.** To use random numbers for security-critical applications, a good TRNG should provide two key properties. First, the TRNG should not leak random numbers to applications other than the requesting one. DR-STRaNGe ensures this property by using a random number

buffer that can be accessed by applications only via a system call. The system call returns random bits only to the application that requested them and discards the used bits from the buffer. Second, the TRNG should provide a unique random number to each random number request. DR-STRaNGe ensures this property by discarding each random number after being served. We conclude that the simple and restrictive interface of DR-STRaNGe provides secure true random numbers to security-critical applications.

### The Random Number Buffer as a Timing Side-Channel.

Side-channel attacks [96] are a class of attacks that observe and use side-channel information to infer application behaviour or leak sensitive information and confidential data [96]. Unlike many other shared hardware structures (e.g., caches [87, 113, 184]) that can be exploited to provide timing information about arbitrary data, DR-STRaNGe can be exploited by an attacker to measure the time it takes to get a random number from DR-STRaNGe. This timing information can be used to infer whether or not the buffer is empty, and understand if another application is requesting random numbers. This side channel is difficult to exploit efficiently due to two reasons: (1) DR-STRaNGe continuously fills the buffer with random numbers asynchronously to the applications, (2) this could cause the buffer to be empty very few times, and (3) if two or more applications are using random numbers, it might be challenging to identify which of those applications is emptying the buffer. We conclude that DR-STRaNGe timing side channels are likely to be more difficult to exploit than side-channels from other more general-purpose shared hardware structures (e.g., caches).

### The Random Number Buffer as a Covert Channel.

An attacker can use channels that are not designed as communication channels to send data from one process to another. The random number buffer can be used as a covert channel [101] under certain circumstances, such as when no process other than an attacker controlled process is requesting random numbers. The random number buffer provides a covert channel fundamentally similar to existing cache covert channels. Previous work proposes several techniques to use shared caches as covert channels [72, 112, 124, 145, 150, 189, 190] and countermeasures to mitigate such attacks [70, 87, 111, 120, 179, 184]. Countermeasures of cache-based covert channels can be applied to the random number buffer. First, the random number buffer can be partitioned across different threads with small performance overhead, since a small size buffer is effective for many applications as shown in Section 8.3. Second, the system can give access privilege of the buffer only to a small number of applications (possibly to one application at a time). This would hurt the performance of RNG applications that cannot access the random number buffer at a given time, but these applications would still benefit from DR-STRaNGe's RNG-Aware scheduler design.

**Denial of Service Attacks.** An attacker might attempt to occupy the DRAM bandwidth with RNG requests to starve other applications. Our RNG-aware memory request scheduler prevents applications from starving with a set of rules to provide system fairness (see Section 5.2). In addition to these rules, DoS attacks can be easily mitigated at OS-level by adapting fairness-aware process scheduling policies (e.g., [126]) to be aware of RNG requests.

## 7. Methodology

To evaluate DR-STRaNGe’s performance and fairness, we use a modified version of Ramulator [1, 92], a cycle-accurate DRAM simulator, that can simulate two state-of-the-art DRAM-based TRNG mechanisms (i.e., D-RaNGe [82] and QUAC-TRNG [138]). We extend the core model of Ramulator to support random number requests. We simulate a system with configurations given in Table 1. We use the configurations of DR-STRaNGe in Table 1 and simulate the TRNG as proposed in D-RaNGe [82] unless stated otherwise.

**Workloads.** We evaluate 43 single-core applications from five benchmark suites: SPEC CPU2006 [2], TPC [172], STREAM [125], MediaBench [48], and YCSB benchmark suite [35]. Based on the last-level cache misses-per-kilo-instruction (MPKI), we group the applications into three categories: (1) L (low memory-intensity,  $MPKI < 1$ ), (2) M (medium memory-intensity,  $1 \leq MPKI < 10$ ), (3) H (high memory-intensity,  $MPKI \geq 10$ ). To do so, we obtain the MPKI values of the applications by analyzing SimPoint [59] traces (200M instructions) of each application.

We create synthetic RNG benchmarks with varying random number request intensities to test our designs. The random number request intensity is controlled by the number of instructions the benchmark has in between two 64-bit random number requests. The synthetic RNG applications read from all banks across all channels, but they are not memory intensive in terms of non-RNG requests. The least RNG-intensive synthetic benchmark requests random numbers with a throughput of 640Mb/s, which is below the maximum throughput of the current best DRAM TRNG mechanism. The most RNG-intensive benchmark requires 5Gb/s random number throughput. We use the most RNG-intensive synthetic RNG benchmark for all tests unless stated otherwise.

We create 43 two-core workloads, each consisting of one non-RNG and one RNG application. In addition, for the four-core configuration we create four multi-programmed workload groups, each consisting of 10 multi-programmed workloads. Each group has 3 different applications from different memory-intensity categories and one synthetic RNG benchmark. For example, a workload in LLHS group has two randomly selected single-core applications from low memory-intensity category, one randomly selected single core application from high memory-intensity category, and a synthetic RNG benchmark. We also create 30 multi-programmed workloads for 8-core and 16-core configurations consisting of low, medium, and high memory-intensity applications.

**Comparison Points.** We compare DR-STRaNGe to (1) an RNG-oblivious system, and (2) a Greedy Idle Design, for performance and fairness results. We compare our RNG-aware Scheduler design to FR-FCFS [151, 197] with a column cap

**Table 1: Simulated System Configuration**

Processor	1-2-4-8-16 cores, 4GHz clock frequency, 3-wide issue, 128-entry instruction window
DRAM	DDR3-1600 [75], 800Mhz bus frequency, 4 channels, 1 rank/channel, 8 banks/rank, 64K rows/bank
Memory Ctrl.	32-entry read/write queues, FR-FCFS [151, 197] with a column cap of 16 [135]
DR-STRANGE	32-entry random read queue, RNG-aware scheduler, 256-entry predictor table/channel, 16-entry random number buffer

of 16 [135] and BLISS [163, 164] memory schedulers. We compare (1) the simple DRAM idleness predictor with a low utilization threshold of 4 to (2) the RL-based DRAM idleness predictor (see Section 5.1.2).

The Greedy Idle Design is based on the idea of a buffer filling mechanism with no overhead.<sup>2</sup> If an idle period reaches the *Period Threshold*, which is 40 cycles in our tests, we assume we fill the buffer with 8 random bits without any overhead. Greedy Idle Design has a separate RNG read queue similar to DR-STRaNGe. To simulate a fair comparison point, we also employ application-priority-based RNG-aware scheduling in the Greedy Idle Design.

**Metrics.** We measure the performance of the non-RNG and RNG applications running on a dual-core system using normalized execution time of the same number of simulated instructions. For multi-core evaluations, for non-RNG applications we use the weighted speedup metric [160], which prior work shows is a good measure of job throughput [45].

For fairness results, we use the unfairness metric proposed in [49, 128, 135]. Unfairness index is calculated as the ratio of the maximum memory-related slowdown experienced by an application in the workload to the minimum memory-related slowdown. To calculate the memory-related slowdown of an application, we measure the memory stall time when memory is shared normalized to the memory stall time when the application runs alone:

$$MemSlowdown_i = \frac{MCP_i^{shared}}{MCP_i^{alone}}, \text{ Unfairness} = \frac{\max_i MemSlowdown_i}{\min_i MemSlowdown_i}$$

If the unfairness index is equal to one, it means that applications running on the system experience similar slowdowns. A higher unfairness index shows that the system unfairly prioritizes one application, and thus there is a large difference between the slowdowns of the applications.

## 8. Results

We evaluate the performance, fairness, energy efficiency, and area overhead of DR-STRaNGe.

### 8.1. Impact on Performance

**8.1.1. Dual-core Performance.** Figure 6 compares the performance results of three designs: (1) the RNG-Oblivious baseline, (2) the Greedy Idle Design, and (3) DR-STRaNGe. The figure shows the slowdown of non-RNG (top) and RNG (bottom) applications in workloads executed on a dual-core system compared to each application’s performance when executed alone on a single core. We make two observations.

First, DR-STRaNGe improves the performance of both RNG and non-RNG applications. DR-STRaNGe reduces the execution time of non-RNG applications by 17.9% and RNG applications by 25.1% on average compared to the baseline RNG-oblivious system. DR-STRaNGe improves the average performance of RNG applications by 20.6% compared to the performance of RNG applications when executed alone on a single core, due to the lower RNG latency.

Second, the greedy design improves average performance of non-RNG and RNG applications by 7.6% and 10.7%, respectively. DR-STRaNGe outperforms the greedy design in most of the workloads because the greedy design fills the random

<sup>2</sup>The Greedy Idle Design provides an upper bound for performance and fairness results using a greedy algorithm. However, its performance and fairness improvement is limited because the best dynamic memory request scheduling order cannot be determined with a greedy approach in polynomial time.

number buffer only in sufficiently long idle periods. In contrast, DR-STRaNGe leverages the low DRAM utilization to fill the random number buffer with its low utilization prediction mechanism (see Section 5.1.2).

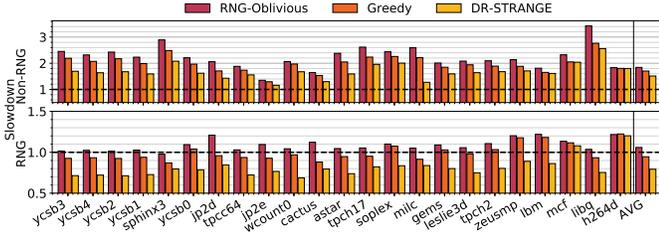


Figure 6: Slowdown over single-core execution of non-RNG (top) and RNG (bottom) applications in dual-core workloads.

### 8.1.2. Multi-core Performance.

**Non-RNG Applications.** Figure 7 shows the normalized weighted speedup of non-RNG applications in four-core workload groups (left) and 4-, 8-, 16-core workload groups (right) normalized to the RNG-oblivious baseline design. We make two observations. First, for four-core workloads, DR-STRaNGe provides 7.6% average performance improvement. As the number of highly memory-intensive applications in the workload increases, DR-STRaNGe’s performance improvement increases. The greedy design has 5.4% higher performance improvement over the RNG-oblivious baseline. DR-STRaNGe outperforms the greedy design because DR-STRaNGe serves a larger number of random number requests from the random number buffer compared to the greedy design. Second, DR-STRaNGe provides 12.1%, 8.2%, and 6.1% average performance improvement for workloads consisting of high, medium, and low memory-intensity applications, respectively. DR-STRaNGe outperforms greedy design in all workload groups with only two exceptions. For highly memory-intensive workloads consisting of 8 and 16 applications, the greedy design provides a slightly higher performance improvement compared to DR-STRaNGe due to the performance overhead of DRAM utilization mispredictions of DR-STRaNGe.

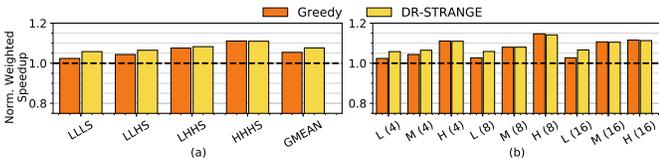


Figure 7: Normalized weighted speedup of non-RNG applications in (a) 4-core workloads, and (b) 4-, 8-, 16-core workloads grouped based on memory-intensity.

**RNG Applications.** Figure 8 compares the performance of RNG applications on (1) the RNG-Oblivious baseline, (2) the Greedy Idle Design, and (3) DR-STRaNGe over the performance of RNG applications when executed on a single-core baseline system. The figure shows the slowdown of RNG applications in four-core workload groups (left) and 4-, 8-, 16-core workload groups (right). For four-core workload groups, DR-STRaNGe improves average performance by 17.8%. For 4-, 8-, 16-core workload groups, DR-STRaNGe improves average performance by 4.5%, 6.7% and 16.9% for high, medium, and low memory-intensity workloads. For all workloads DR-STRaNGe improves performance at least as much as the greedy idle design.

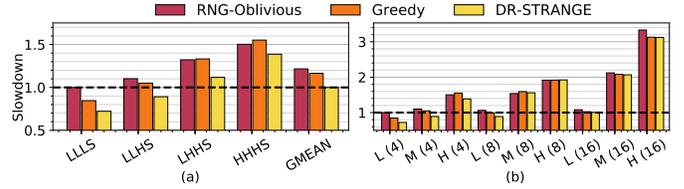


Figure 8: Slowdown of RNG applications in (a) 4-core workloads, and (b) 4-, 8-, 16-core workloads grouped based on non-RNG applications’ memory-intensity.

### 8.2. Impact on System Fairness

We evaluate system fairness impact of three designs: (1) the RNG-Oblivious baseline, (2) the Greedy Idle Design, and (3) DR-STRaNGe. Figure 9 plots system fairness, which we calculate using the *unfairness index* metric [49, 128, 135], for two-core workloads. We make three observations. First, DR-STRaNGe improves average system fairness by 32.1% compared to the RNG-oblivious baseline design. Second, DR-STRaNGe outperforms the greedy design by 15.2% in terms of fairness. Third, some workloads that include non-RNG applications, such as *jp2d* and *cactusADM*, show higher unfairness indices compared to the greedy design because DR-STRaNGe improves performance of RNG applications more than that of such non-RNG applications due to its effective random number buffering. We conclude that DR-STRaNGe outperforms both the RNG-oblivious baseline and the Greedy Idle Design in terms of system fairness by employing an RNG-aware scheduling policy and effectively reducing and controlling the RNG interference.

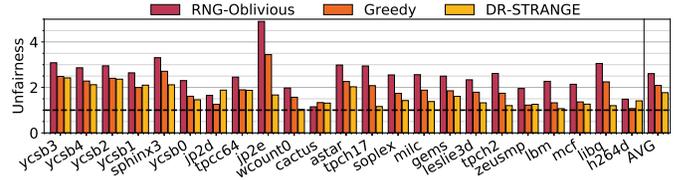


Figure 9: System fairness on a dual-core system.

### 8.3. Impact of the Random Number Buffer

Serving random number requests from the random number buffer reduces the total execution time by reducing the RNG latency. The impact of the random number buffer depends on the ratio of random number requests served from the random number buffer over all random number requests. We call this ratio *the buffer serve rate*. The buffer serve rate depends on (1) DRAM channel utilization, (2) required TRNG throughput, and (3) the random number buffer size. In this section, we evaluate DR-STRaNGe with different random number buffer sizes maintained with the simple buffering mechanism.

Figure 10 shows the slowdown of non-RNG and RNG applications executed on a dual-core system compared to when each application is executed alone on a single-core system (top and middle) and the buffer serve rate of workloads (bottom). We make two observations.

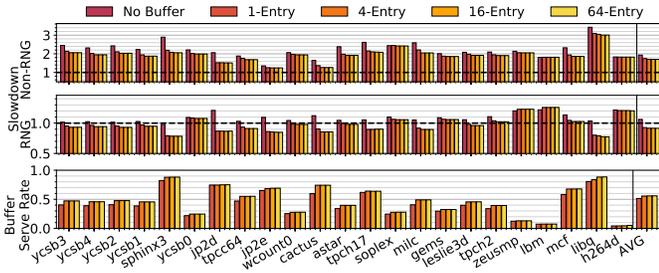
First, a 16-entry random number buffer improves the average performance of non-RNG and RNG applications by 11.7% and 13.8%, respectively. The performance improvement of non-RNG and RNG applications (top and middle) increases with larger buffer sizes until 16 entries and a significant performance improvement is achieved with a 16-entry random number buffer.

Second, a 16-entry random number buffer achieves an average buffer serve rate of 0.55. Increasing the buffer size improves

the buffer serve rate of only a few workloads, such as the workloads including *jp2e*, *cactus*, and *libquantum*.

Third, RNG applications in some workloads consisting of memory-intensive non-RNG applications, such as *zeusmp* and *lbm*, experience slowdown due to the low number of sufficiently long idle periods and increased RNG interference as a result of random number generation for the buffer. These workloads do not benefit much from the random number buffer and have low buffer serve rates as shown in Figure 10 (bottom).

We conclude that the buffering mechanism improves the average performance of non-RNG and RNG applications by reducing the RNG interference and TRNG latency.



**Figure 10: Impact of the buffer size on slowdown of non-RNG (top, lower is better) and RNG applications (middle, lower is better) and buffer serve rate (bottom, higher is better).**

#### 8.4. Impact of RNG-Aware Scheduling

We evaluate (1) DR-STRaNGe, (2) BLISS [163, 164]<sup>3</sup>, and (3) the RNG-oblivious baseline design with the FR-FCFS scheduler [151, 197] with a column cap of 16, called FR-FCFS+Cap [135]. In these tests, we evaluate DR-STRaNGe designs with no random number buffer.

Figure 11 shows the performance impact of the RNG-Aware scheduler on non-RNG applications (top) and RNG applications (middle) normalized to the performance of applications executed on a single core. We make two observations.

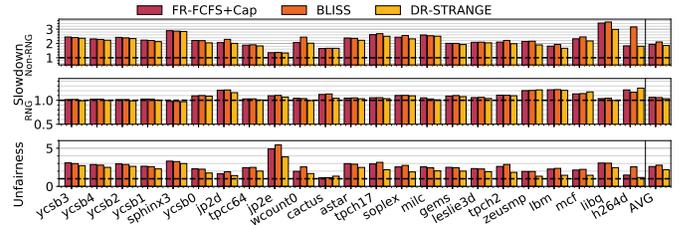
First, the RNG-Aware scheduler outperforms both FR-FCFS+Cap and BLISS mechanisms and improves average fairness by 16.1%. It improves average performance of non-RNG and RNG applications by 5.6% and 1.6%, respectively. The performance of the RNG-aware scheduler is greater than the performance of FR-FCFS+Cap and BLISS for almost all workloads.

Second, on average, BLISS has a higher unfairness index compared to FR-FCFS+Cap and increases average unfairness by 6.6%. Some workloads consisting of memory-intensive non-RNG applications, such as *jp2e*, *wcount0*, *tpch17*, *soplex*, *tpch2*, *lbm*, *mcf*, and *h264d*, have significantly higher unfairness indices compared to the FR-FCFS+Cap and the RNG-Aware schedulers because BLISS frequently blacklists these highly memory-intensive non-RNG applications and prioritizes RNG applications. Due to the frequently blacklisted non-RNG applications, BLISS degrades average performance of non-RNG applications over FR-FCFS+Cap by 8.9%.

#### 8.5. Impact of Priority-Based Scheduling

To show the impact of priority-based RNG-aware scheduling, we compare three designs: (1) the RNG-Oblivious baseline, (2) DR-STRaNGe with high-priority RNG applications, and (3) DR-STRaNGe with high-priority non-RNG applications.

<sup>3</sup>We use a value of four for *Blacklisting Threshold* and a value of 10000 cycles for *Clearing Interval* as proposed in [164].



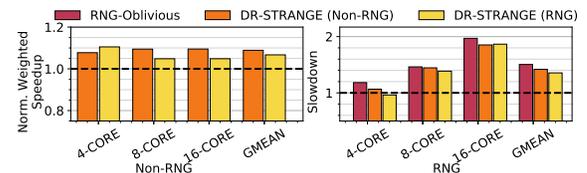
**Figure 11: Impact of the memory request scheduler on the performance of non-RNG (top) and RNG (middle) applications and system fairness (bottom).**

Figure 12 shows the normalized weighted speedup of non-RNG applications (left) and the slowdown of RNG applications normalized to their performance when executed on a single core (right). We make three observations.

First, priority-based RNG-aware scheduling provides significant performance improvement when workloads fully utilize the DRAM bandwidth. Both types of applications benefit from RNG-aware scheduling and performance improvements increase when they have high priority levels.

Second, Figure 12 (left) shows that prioritizing the non-RNG applications using DR-STRaNGe improves the average weighted speedup of the non-RNG applications by 8.9% compared to the RNG-oblivious baseline. Figure 12 (right) shows that prioritizing the RNG applications using DR-STRaNGe improves average performance of the RNG applications by 9.9% compared to the RNG-oblivious baseline.

Third, prioritizing RNG applications using DR-STRaNGe improves the average performance of both non-RNG and RNG applications in 4-core workloads. Some workloads consisting of low and medium memory-intensity non-RNG applications switch between RNG and regular memory request queues frequently when non-RNG applications are prioritized due to low memory-intensity applications creating requests at lower rates. These workloads benefit from lower RNG interference when RNG applications are prioritized using DR-STRaNGe.

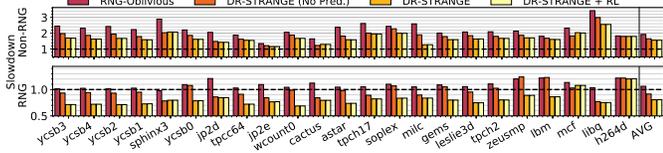


**Figure 12: Impact of the priority-based RNG-aware scheduler on performance of non-RNG (left) and RNG (right) applications.**

#### 8.6. Impact of the DRAM Idleness Predictor

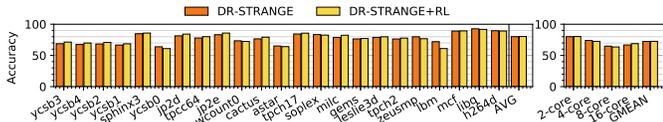
Figure 13 evaluates the performance of four designs: (1) the RNG-oblivious baseline, (2) DR-STRaNGe with no DRAM idleness predictor, (3) DR-STRaNGe with the simple DRAM idleness predictor, and (4) DR-STRaNGe with the RL-based DRAM idleness predictor. Figure 13 plots the slowdown of non-RNG (top) and RNG applications (bottom). We make three observations. First, DR-STRaNGe with the simple DRAM idleness predictor improves average performance of non-RNG and RNG applications by 17.9% and 25.1%, respectively, compared to the RNG-oblivious baseline. Second, the simple DRAM idleness predictor improves DR-STRaNGe's performance by 12.4% and 13.8% for non-RNG and RNG applications, respectively. Third, DR-STRaNGe with the simple predictor achieves similar performance improvement at lower hardware complexity compared to DR-STRaNGe with the RL-based idleness predic-

tor, which improves non-RNG and RNG applications’ average performance by 19.3% and 23.9%, respectively.



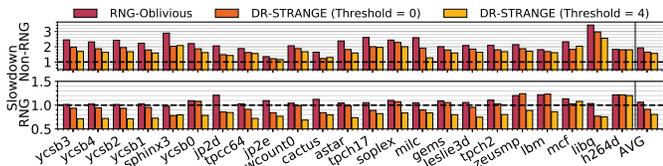
**Figure 13: Impact of the DRAM Idleness Predictor on non-RNG (top) and RNG (bottom) applications’ performance.**

Figure 14 shows the predictor accuracy of the 2-core workloads (left) and 2-, 4-, 8-, 16-core workloads (right). We make the following observations. First, with two-core workloads, the simple DRAM idleness predictor and the RL-based idleness predictor have an accuracy of 80.0% and 80.3%, respectively. Second, the simple DRAM idleness predictor’s accuracy is slightly higher than the RL-based predictor’s accuracy when the workload’s memory-intensity is high and the number of long idle periods is low. In workloads containing low memory-intensity applications, the simple DRAM idleness predictor mispredicts long idle periods to be short and has a higher false negative rate. Third, in 4-, 8-, and 16-core workloads, both predictors have lower accuracy due to less idleness and more complex interference patterns.



**Figure 14: DRAM idleness predictor accuracy in two-core (left) and multicore workloads (right).**

**8.6.1. Impact of the Low Utilization Prediction.** We evaluate three designs: (1) the RNG-oblivious baseline, (2) the simple DRAM idleness predictor without low utilization prediction, and (3) the simple DRAM idleness predictor with low utilization prediction using the threshold value of 4. Figure 15 shows the performance impact of DR-STRaNGe with and without the low utilization prediction mechanism on non-RNG (top) and RNG (bottom) applications. We conclude that the simple DRAM idleness predictor with a low utilization threshold of 4 improves the average performance of non-RNG and RNG applications by 5.5% and 11.7%, respectively, compared to the DRAM idleness predictor without low utilization prediction.



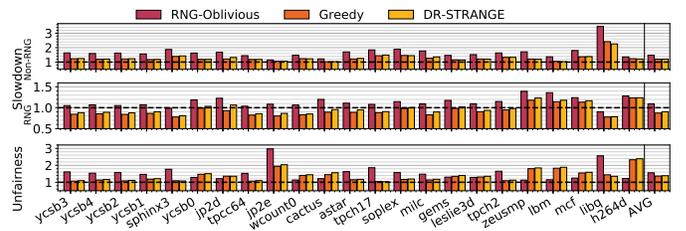
**Figure 15: Impact of the low utilization prediction on non-RNG (top) and RNG (bottom) applications’ performance.**

## 8.7. Experiments Using QUAC-TRNG

To show that DR-STRaNGe is compatible with different DRAM-based TRNGs, we evaluate DR-STRaNGe with QUAC-TRNG [138]. QUAC-TRNG provides a higher RNG throughput compared to D-RaNGe [82]. However, it also has higher latency for 64-bit random number generation. We compare DR-STRaNGe’s impact on performance and system fairness compared to the RNG-oblivious baseline when both systems use QUAC-TRNG to generate the random numbers.

Figure 16 plots the performance of non-RNG and RNG applications (top and middle) and system fairness (bottom) compared to the RNG-Oblivious baseline. We make three observations. First, DR-STRaNGe improves average performance of non-RNG and RNG applications by 18.2% and 17.2%, respectively. Second, DR-STRaNGe improves average system fairness by 10.9%. Some workloads with high memory-intensity (*zeusmp*, *lbm*, *mcf*, *h264d*) suffer from higher unfairness indices because DR-STRaNGe improves performance of non-RNG applications more than the performance of RNG applications.

We conclude that DR-STRaNGe improves both performance and system fairness regardless of the implemented DRAM TRNG. DR-STRaNGe can also potentially leverage hybrid DRAM TRNGs that can generate random numbers using (1) a low-latency DRAM-based TRNG to fill the random number buffer, and (2) a high-throughput DRAM-based TRNG to generate random numbers after receiving a random number request when the random number buffer is empty. We leave the evaluation of such design to future work.



**Figure 16: Performance and fairness in dual-core workloads consisting of RNG and non-RNG applications in a system that uses QUAC-TRNG [138].**

## 8.8. Results of Low-Intensity Applications

We evaluate DR-STRaNGe with workloads consisting of (1) one low-intensity RNG application (i.e., 640 *Mb/s* RNG throughput), and (2) one non-RNG application. We observe that DR-STRaNGe improves the average performance of the RNG and non-RNG applications by 3.2% and 4.6% respectively. Due to the low RNG interference caused by the low required TRNG throughput, DR-STRaNGe does not significantly improve system fairness over the baseline.

## 8.9. Area and Energy Consumption Analysis

To evaluate DR-STRaNGe’s energy consumption, we use DRAMPower [24] with Ramulator’s output traces. We observe that DR-STRaNGe reduces energy consumption and total memory cycles by 21% and 15.8%, respectively, compared to the RNG-oblivious baseline system.

We use CACTI [132] at 22nm process technology node to model the DR-STRaNGe configuration shown in Table 1, including the random number buffer, RNG request queue, and the DRAM idleness predictor. We find that DR-STRaNGe incurs minor area overhead: 0.0022mm<sup>2</sup> (0.00048% of an Intel Cascade Lake CPU Core [187]). These evaluations include the simple DRAM idleness predictor that consists of 256 entries and a 2-bit saturating counter in each entry, with a total of 0.0625 KB area overhead. If we use the RL-based predictor, DR-STRaNGe has an area cost of 0.012mm<sup>2</sup> (0.0033% of an Intel Cascade Lake CPU Core [187]) and the RL agent requires 8KB storage assuming 4-byte Q-values and 10-bit state values.

## 9. Other Related Work

To our knowledge, DR-STRaNGe is the first work to propose an end-to-end system design for DRAM TRNGs that (1) reduces the interference between RNG and non-RNG applications in the memory controller, (2) improves system fairness, and (3) reduces high TRNG latency. We have already discussed closely related work on memory request schedulers (Section 8.4) and DRAM-based TRNG mechanisms and compared DR-STRaNGe to prior proposals (Sections 8.2, 8.1, and 8.7). In this section, we discuss other related works.

**Memory Request Scheduling.** Previously proposed memory request scheduler designs [8, 43, 54, 73, 89, 90, 129, 135, 135–137, 163, 164, 175, 195] aim to reduce inter-application interference to improve the system performance and fairness. These proposals are RNG-oblivious. We already compare DR-STRaNGe to BLISS [163, 164] and FR-FCFS [151, 197] with a column cap [135] in Section 8.4.

**Low-Throughput DRAM-based TRNGs.** Prior work proposes DRAM-based TRNGs that generate random numbers using different entropy sources, such as DRAM start-up values [44], and retention failures [60, 78, 166, 168]. These DRAM-based TRNGs are limited in the throughput they can provide and cannot be used in a streaming manner. As such, they are less practical than the high-throughput DRAM-based TRNGs we evaluate [82, 138].

**DRAM Idleness Predictors.** Prior work proposes DRAM idleness predictors that predict the idle period lengths to reduce DRAM energy consumption [169] and efficiently schedule last level cache writebacks [183]. Compared to these techniques, DR-STRaNGe’s idleness predictor can be implemented using simple hardware at low cost (Section 8.9) and requires no modifications to the interface between the processor and the memory controller.

## 10. Conclusion

We propose DR-STRaNGe, the first end-to-end system design for DRAM true random number generators (TRNGs) that reduces the RNG interference in the memory controller, provides system fairness among different types of applications, and successfully hides the latency of DRAM-based TRNGs. Our system design consists of three main parts: (1) a random number buffering mechanism combined with a DRAM idleness predictor, (2) an RNG-aware memory request scheduler, and (3) an application interface. Our evaluations show that DR-STRaNGe improves the performance of RNG and non-RNG applications by 17.9% and 25.1% respectively, and the overall system fairness by 32.1% while reducing average system energy consumption by 21%. We conclude that with an end-to-end system design, DRAM-based TRNGs can be seamlessly integrated into today’s systems with low overheads and provide true random numbers at high throughput and low latency.

## Acknowledgments

We thank the anonymous reviewers of MICRO 2021 and HPCA 2022 for feedback. We thank the SAFARI group members for feedback and the stimulating intellectual environment. We acknowledge the generous gifts provided by our industrial partners: Google, Huawei, Intel, Microsoft, VMware.

## References

- [1] “Ramulator Source Code,” <https://github.com/CMU-SAFARI/ramulator>.
- [2] “Standard Performance Evaluation Corporation,” <http://www.spec.org/cpu2006>.
- [3] S. Aga, S. Jeloka, A. Subramanian, S. Narayanasamy, D. Blaauw, and R. Das, “Compute Caches,” in *HPCA*, 2017.
- [4] J. Ahn, S. Hong, S. Yoo, O. Mutlu, and K. Choi, “A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing,” in *ISCA*, 2015.
- [5] J. Ahn, S. Yoo, O. Mutlu, and K. Choi, “PIM-Enabled Instructions: a Low-overhead, Locality-aware Processing-in-Memory Architecture,” in *ISCA*, 2015.

- [6] B. Akin, F. Franchetti, and J. C. Hoe, “Data Reorganization in Memory Using 3D-Stacked DRAM,” in *ISCA*, 2015.
- [7] T. Amaki, M. Hashimoto, and T. Onoye, “An Oscillator-based True Random Number Generator with Process and Temperature Tolerance,” in *DAC*, 2015.
- [8] R. Ausavarungnirun, K. K.-W. Chang, L. Subramanian, G. H. Loh, and O. Mutlu, “Staged Memory Scheduling: Achieving High Performance and Scalability in Heterogeneous Systems,” in *ISCA*, 2012.
- [9] O. O. Babarinsa and S. Idreos, “JAFAR: Near-Data Processing for Databases,” in *SIGMOD*, 2015.
- [10] V. Bagini and M. Bucci, “A Design of Reliable True Random Number Generator for Cryptographic Applications,” in *CHES*, 1999.
- [11] B. M. S. Bahar Talukder, J. Kerns, B. Ray, T. Morris, and M. T. Rahman, “Exploiting DRAM Latency Variations for Generating True Random Numbers,” in *ICCE*, 2019.
- [12] M. Barangi, J. S. Chang, and P. Mazumder, “Straintronics-Based True Random Number Generator for High-Speed and Energy-Limited Applications,” in *IEEE Trans. Magn.*, 2016.
- [13] R. Bera, K. Kanellopoulos, A. Nori, T. Shahroodi, S. Subramoney, and O. Mutlu, “Pythia: A Customizable Hardware Prefetching Framework Using Online Reinforcement Learning,” in *MICRO*, 2021.
- [14] M. Bhargava, K. Sheikh, and K. Mai, “Robust True Random Number Generator using Hot-carrier Injection Balanced Metastable Sense Amplifiers,” in *HOST*, 2015.
- [15] I. Bhati, Z. Chishti, S.-L. Lu, and B. Jacob, “Flexible Auto-Refresh: Enabling Scalable and Energy-Efficient DRAM Refresh Reductions,” in *ISCA*, 2015.
- [16] L. Blum, M. Blum, and M. Shub, “A Simple Unpredictable pseudo-random Number Generator,” in *SIAM Journal on Computing*, 1986.
- [17] A. Boroumand, S. Ghose, B. Akin, R. Narayanaswami, G. F. Oliveira, X. Ma, E. Shiu, and O. Mutlu, “Mitigating Edge Machine Learning Inference Bottlenecks: An Empirical Study on Accelerating Google Edge Models,” arXiv:2103.00768, 2021.
- [18] A. Boroumand, S. Ghose, Y. Kim, R. Ausavarungnirun, E. Shiu, R. Thakur, D. Kim, A. Kausela, A. Knies, P. Ranganathan, and Ö. Mutlu, “Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks,” in *ASPLOS*, 2018.
- [19] A. Boroumand, S. Ghose, B. Lucia, K. Hsieh, K. Malladi, H. Zheng, and O. Mutlu, “LazyPIM: An Efficient Cache Coherence Mechanism for Processing-in-Memory,” in *CAL*, 2017.
- [20] R. Botha, “The Development of a Hardware Random Number Generator for Gamma-ray Astronomy,” Ph.D. dissertation, North-West University, 2005.
- [21] R. Brederlow, R. Prakash, C. Paulus, and R. Thewes, “A Low-power True Random Number Generator using Random Telegraph Noise of Single Oxide-traps,” in *ISSCC*, 2006.
- [22] M. Bucci, L. Germani, R. Luzzi, A. Trifiletti, and M. Varanouoovo, “A High-speed Oscillator-based Truly Random Number Source for Cryptographic Applications on a Smart Card IC,” in *TC*, 2003.
- [23] J. J. M. Chan, B. Sharma, J. Lv, G. Thomas, R. Thulasiram, and P. Thulasiraman, “True Random Number Generator using GPUs and Histogram Equalization Techniques,” in *HPCC*, 2011.
- [24] K. Chandrasekar, C. Weis, Y. Li, S. Goossens, M. Jung, O. Naji, B. Akesson, N. Wehn, and K. Goossens, “DRAMPower: Open-Source DRAM Power & Energy Estimation Tool,” <http://www.drampower.info/>.
- [25] K. K. Chang, A. Kashyap, H. Hassan, S. Ghose, K. Hsieh, D. Lee, T. Li, G. Pekhimenko, S. Khan, and O. Mutlu, “Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization,” in *SIGMETRICS*, 2016.
- [26] K. K. Chang, D. Lee, Z. Chishti, A. R. Alameldeen, C. Wilkerson, Y. Kim, and O. Mutlu, “Improving DRAM Performance by Parallelizing Refreshes with Accesses,” in *HPCA*, 2014.
- [27] K. K. Chang, P. J. Nair, D. Lee, S. Ghose, M. K. Qureshi, and O. Mutlu, “Low-cost Inter-linked Subarrays (LISA): Enabling Fast Inter-subarray Data Movement in DRAM,” in *HPCA*, 2016.
- [28] K. K. Chang, A. G. Yaglikci, S. Ghose, A. Agrawal, N. Chatterjee, A. Kashyap, D. Lee, M. O’Connor, H. Hassan, and O. Mutlu, “Understanding Reduced-Voltage Operation in Modern DRAM Devices: Experimental Characterization, Analysis, and Mechanisms,” in *SIGMETRICS*, 2017.
- [29] K. K. Chang, A. Yaglikci, S. Ghose, A. Agrawal, N. Chatterjee, A. Kashyap, D. Lee, M. O’Connor, H. Hassan, and O. Mutlu, “Understanding Reduced-Voltage Operation in Modern DRAM Devices: Experimental Characterization, Analysis, and Mechanisms,” in *SIGMETRICS*, 2017.
- [30] A. Cherkauoi, V. Fischer, L. Fesquet, and A. Aubert, “A Very High Speed True Random Number Generator with Entropy Assessment,” in *CHES*, 2013.
- [31] P. Chevalier, C. Menard, and B. Dorval, “Random Number Generator,” US Patent 3,790,768, 1974.
- [32] P. Chi, S. Li, C. Xu, T. Zhang, J. Zhao, Y. Liu, Y. Wang, and Y. Xie, “PRIME: A Novel Processing-in-Memory Architecture for Neural Network Computation in ReRAM-Based Main Memory,” in *ISCA*, 2016.
- [33] P. P. Chu and R. E. Jones, “Design Techniques of FPGA Based Random Number Generator,” in *MAPLD*, 1999.
- [34] P. J. Clarke, R. J. Collins, P. A. Hiskett, P. D. Townsend, and G. S. Buller, “Robust Gigahertz Fiber Quantum Key Distribution,” *Applied Physics Letters*, 2011.
- [35] B. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears, “Benchmarking Cloud Serving Systems with YCSB,” in *SoCC*, 2010.
- [36] H. Corrigan-Gibbs, W. Mu, D. Boneh, and B. Ford, “Ensuring High-quality Randomness in Cryptographic Key Generation,” in *CCS*, 2013.
- [37] H. David, C. Fallin, E. Gorbato, U. R. Hanebutte, and O. Mutlu, “Memory Power Management via Dynamic Voltage/Frequency Scaling,” in *ICAC*, 2011.
- [38] P. Davis and P. Rabinowitz, “Some Monte Carlo Experiments in Computing Multiple Integrals,” *Mathematical Tables and Other Aids to Computation*, 1956.
- [39] F. Devaux, “The True Processing in Memory Accelerator,” in *HCS*, 2019.
- [40] S. M. P. Dinakarrao, H. Yu, H. Huang, and D. Xu, “A Q-Learning Based Self-Adaptive I/O Communication for 2.5D Integrated Many-Core Microprocessor and Memory,” *IEEE Transactions on Computers*, vol. 65, pp. 1185–1196, 2016.
- [41] L. Dorrendorf, Z. Gutterman, and B. Pinkas, “Cryptanalysis of the Windows Random Number Generator,” in *CCS*, 2007.
- [42] M. Drutarovsky and P. Galajda, “A Robust Chaos-based True Random Number Generator Embedded in Reconfigurable Switched-Capacitor Hardware,” in *Radiotelektronika*, 2007.
- [43] E. Ebrahimi, R. Miftakhutdinov, C. Fallin, C. J. Lee, J. A. Joao, O. Mutlu, and Y. N. Patt, “Parallel application memory scheduling,” in *MICRO*, 2011, pp. 362–373.

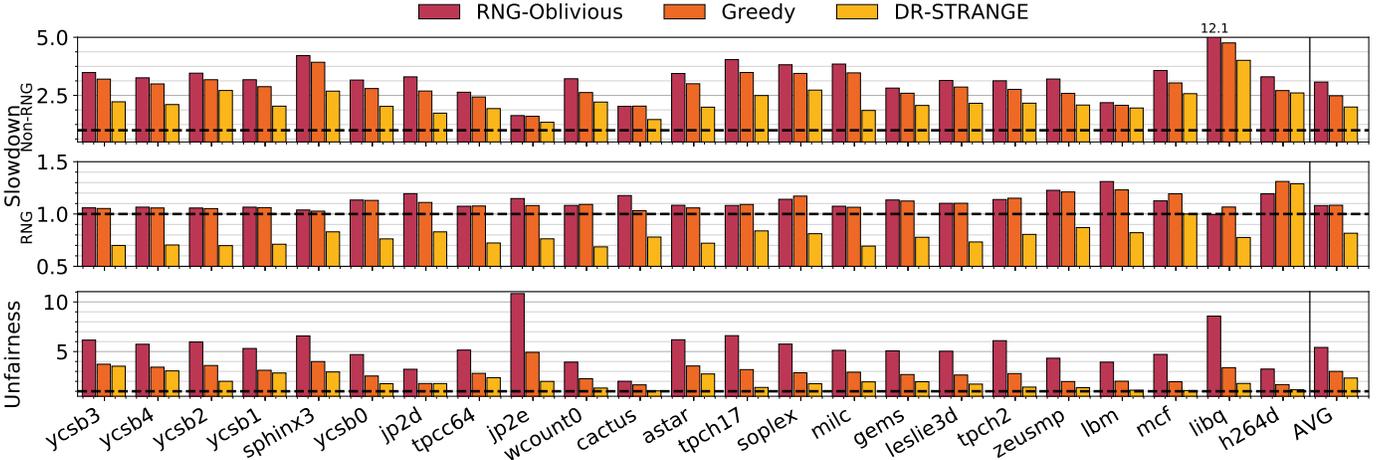
- [44] C. Eckert, F. Tehranipoor, and J. A. Chandy, "DRNG: DRAM-based Random Number Generation Using its Startup Value Behavior," in *MWSCAS*, 2017.
- [45] S. Eyerman and L. Eeckhout, "System-level Performance Metrics for Multiprogram Workloads," in *IEEE Micro*, 2008.
- [46] A. Farnahini-Farahani, J. H. Ahn, K. Morrow, and N. S. Kim, "NDA: Near-DRAM Acceleration Architecture Leveraging Commodity DRAM Devices and Standard Memory Modules," in *HPCA*, 2015.
- [47] I. Fernandez, R. Quisilant, C. Giannoula, M. Alser, J. Gómez-Luna, E. Gutiérrez, O. Plata, and O. Mutlu, "NATSA: A Near-Data Processing Accelerator for Time Series Analysis," 2020.
- [48] J. E. Fritts, F. W. Steiling, J. A. Tucek, and W. Wolf, "MediaBench II Video: Expediting the next Generation of Video Systems Research," *Microprocess. Microsyst.*, 2009.
- [49] R. Gabor, S. Weiss, and A. Mendelson, "Fairness and Throughput in Switch on Event Multithreading," in *MICRO*, 2006.
- [50] M. Gao, G. Ayers, and C. Kozyrakis, "Practical Near-Data Processing for In-Memory Analytics Frameworks," in *PACT*, 2015.
- [51] M. Gao and C. Kozyrakis, "HRL: Efficient and Flexible Reconfigurable Logic for Near-Data Processing," in *HPCA*, 2016.
- [52] S. Ghose, A. Boroumand, J. S. Kim, J. Gómez-Luna, and O. Mutlu, "Processing-in-memory: A workload-driven perspective," *IBM JRD*, 2019.
- [53] S. Ghose, K. Hsieh, A. Boroumand, R. Ausavarungnirun, and O. Mutlu, "Enabling the Adoption of Processing-in-Memory: Challenges, Mechanisms, Future Research Directions," *arXiv preprint arXiv:1802.00320*, 2018.
- [54] S. Ghose, H. Lee, and J. F. Martínez, "Improving Memory Scheduling via Processor-Side Load Criticality Information," in *ISCA*, 2013.
- [55] S. Ghose, G. Yaglikci, R. Gupta, D. Lee, K. Kudrolli, W. Liu, H. Hassan, K. Chang, N. Chatterjee, A. Agrawal, M. O'Connor, and O. Mutlu, "What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study," in *SIGMETRICS*, 2018.
- [56] C. Giannoula, N. Vijaykumar, N. Papadopoulou, V. Karakostas, I. Fernandez, J. Gómez-Luna, L. Orosa, N. Koziris, G. Goumas, and O. Mutlu, "SynCron: Efficient Synchronization Support for Near-Data-Processing Architectures," in *HPCA*, 2021.
- [57] Z. Gutterman, B. Pinkas, and T. Reinman, "Analysis of the Linux Random Number Generator," in *SP*, 2006.
- [58] J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten, "Lest we remember: Cold-boot attacks on encryption keys," *Commun. ACM*, may 2009.
- [59] G. Hamerly, E. Perelman, J. Lau, and B. Calder, "SimPoint 3.0: Faster and More Flexible Program Phase Analysis," *Journal of Instruction-Level Parallelism*, 2005.
- [60] M. S. Hashemian, B. Singh, F. Wolff, D. Weyer, S. Clay, and C. Papachristou, "A Robust Authentication Methodology Using Physically Unclonable Functions in DRAM Arrays," in *DATE*, 2015.
- [61] H. Hassan, M. Patel, J. S. Kim, A. G. Yaglikci, N. Vijaykumar, N. M. Ghiasi, S. Ghose, and O. Mutlu, "CROW: A Low-Cost Substrate for Improving DRAM Performance, Energy Efficiency, and Reliability," in *ISCA*, 2019.
- [62] H. Hassan, G. Pekhimenko, N. Vijaykumar, V. Seshadri, D. Lee, O. Ergin, and O. Mutlu, "ChargeCache: Reducing DRAM Latency by Exploiting Row Access Locality," in *HPCA*, 2016.
- [63] H. Hassan, N. Vijaykumar, S. Khan, S. Ghose, K. Chang, G. Pekhimenko, D. Lee, O. Ergin, and O. Mutlu, "SoftMC: A Flexible and Practical Open-source Infrastructure for Enabling Experimental DRAM Studies," in *HPCA*, 2017.
- [64] S. M. Hassan, S. Yalamanchili, and S. Mukhopadhyay, "Near Data Processing: Impact and Optimization of 3D Memory System Architecture on the Uncore," in *MEMSYS*, 2015.
- [65] D. E. Holcomb, W. P. Burleson, and K. Fu, "Initial SRAM State as a Fingerprint and Source of True Random Numbers for RFID Tags," in *RFID*, 2007.
- [66] D. E. Holcomb, W. P. Burleson, and K. Fu, "Power-Up SRAM State as an Identifying Fingerprint and Source of True Random Numbers," in *TC*, 2009.
- [67] J. Holleman, S. Bridges, B. P. Otis, and C. Diorio, "A 3mu W CMOS True Random Number Generator with Adaptive Floating-Gate Offset Cancellation," *JSSC*, 2008.
- [68] K. Hsieh, E. Ebrahimi, G. Kim, N. Chatterjee, M. O'Connor, N. Vijaykumar, O. Mutlu, and S. W. Keckler, "Transparent Offloading and Mapping (TOM): Enabling Programmer-Transparent Near-Data Processing in GPU Systems," in *ISCA*, 2016.
- [69] K. Hsieh, S. Khan, N. Vijaykumar, K. K. Chang, A. Boroumand, S. Ghose, and O. Mutlu, "Accelerating Pointer Chasing in 3D-Stacked Memory: Challenges, Mechanisms, Evaluation," in *ICCD*, 2016.
- [70] W.-M. Hu, "Reducing timing channels with fuzzy time," *Journal of Computer Security*, 1992.
- [71] T. E. Hull and A. R. Dobell, "Random Number Generators," *SIAM Review*, 1962.
- [72] C. Hunger, M. Kazdagli, A. Rawat, A. Dimakis, S. Vishwanath, and M. Tiwari, "Understanding Contention-based Channels and Using Them for Defense," in *HPCA*, 2015.
- [73] I. Hur and C. Lin, "Adaptive History-Based Memory Schedulers," in *MICRO*, 2004.
- [74] E. Ipek, O. Mutlu, J. F. Martínez, and R. Caruana, "Self-Optimizing Memory Controllers: A Reinforcement Learning Approach," in *ISCA*, 2008.
- [75] JEDEC, *Double Data Rate 3 (DDR3) SDRAM Specification*, 2012.
- [76] W. Kang and S. Yoo, "Dynamic Management of Key States for Reinforcement Learning-assisted Garbage Collection to Reduce Long Tail Latency in SSD," in *DAC*, 2018.
- [77] B. Keeth and R. Baker, *DRAM Circuit Design: A Tutorial*, 2001.
- [78] C. Keller, F. Gurkaynak, H. Kaeslin, and N. Felber, "Dynamic Memory-based Physically Unclonable Function for the Generation of Unique Identifiers and True Random Numbers," in *ISCAS*, 2014.
- [79] S. Khan, D. Lee, Y. Kim, A. R. Alameldeen, C. Wilkerson, and O. Mutlu, "The Efficacy of Error Mitigation Techniques for DRAM Retention Failures: A Comparative Experimental Study," in *SIGMETRICS*, 2014.
- [80] S. Khan, D. Lee, and O. Mutlu, "PARBOR: An Efficient System-Level Technique to Detect Data-Dependent Failures in DRAM," in *DSN*, 2016.
- [81] S. Khan, C. Wilkerson, Z. Wang, A. R. Alameldeen, D. Lee, and O. Mutlu, "Detecting and Mitigating Data-Dependent DRAM Failures by Exploiting Current Memory Content," in *MICRO*, 2017.
- [82] J. S. Kim, M. Patel, H. Hassan, L. Orosa, and O. Mutlu, "D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput," in *HPCA*, 2019.
- [83] J. Kim, T. Ahmed, H. Nili, N. D. Truong, J. Yang, D. S. Jeong, S. Sriram, D. C. Ranasinghe, and O. Kavehei, "Nano-Intrinsic True Random Number Generation," arXiv:1701.06020, 2017.
- [84] J. S. Kim, D. S. Cali, H. Xin, D. Lee, S. Ghose, M. Alser, H. Hassan, O. Ergin, C. Alkan, and O. Mutlu, "GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-memory Technologies," *BMC Genomics*, 2018.
- [85] J. S. Kim, M. Patel, H. Hassan, and O. Mutlu, "Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines," in *ICCD*, 2018.
- [86] J. S. Kim, M. Patel, H. Hassan, and O. Mutlu, "The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern Commodity DRAM Devices," in *HPCA*, 2018.
- [87] T. Kim, M. Peinado, and G. Mainar-Ruiz, "STEALTHMEM: System-level protection against cache-based side channel attacks in the cloud," in *USENIX Security*, 2012.
- [88] Y. Kim, R. Daly, J. Kim, C. Fallin, J. H. Lee, D. Lee, C. Wilkerson, K. Lai, and O. Mutlu, "Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors," in *ISCA*, 2014.
- [89] Y. Kim, D. Han, O. Mutlu, and M. Harchol-Balter, "ATLAS: A Scalable and High-Performance Scheduling Algorithm for Multiple Memory Controllers," in *HPCA*, 2010.
- [90] Y. Kim, M. Papamichael, O. Mutlu, and M. Harchol-Balter, "Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior," in *MICRO*, 2010.
- [91] Y. Kim, V. Seshadri, D. Lee, J. Liu, and O. Mutlu, "A Case for Exploiting Subarray-level Parallelism (SALP) in DRAM," in *ISCA*, 2012.
- [92] Y. Kim, W. Yang, and O. Mutlu, "Ramulator: A Fast and Extensible DRAM Simulator," in *CAL*, 2016.
- [93] D. Kinniment and E. Chester, "Design of an On-chip Random Number Generator using Metastability," in *ESSCIRC*, 2002.
- [94] D. E. Knuth, "The Art of Computer Programming, 2: Seminumerical Algorithms, Addition Wesley," 1998.
- [95] C. K. Koç, "About Cryptographic Engineering," in *Cryptographic Engineering*, 2009.
- [96] P. C. Kocher, "Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems," in *Annual International Cryptology Conference*, 1996.
- [97] S. H. Kwok and E. Y. Lam, "FPGA-based High-speed True Random Number Generator for Cryptographic Applications," in *TENCON*, 2006.
- [98] Y.-C. Kwon, S. H. Lee, J. Lee, S.-H. Kwon, J. M. Ryu, J.-P. Son, O. Seongil, H.-S. Yu, H. Lee, S. Y. Kim et al., "25.4 A 20nm 6GB Function-In-Memory DRAM, Based on HBM2 with a 1.2 TFLOPS Programmable Computing Unit Using Bank-Level Parallelism, for Machine Learning Applications," in *ISSCC*, 2021.
- [99] Q. Labs, "Random Number Generators," White Paper, 2015.
- [100] P. Lacharme, A. Rock, V. Strubel, and M. Videau, "The Linux Pseudorandom Number Generator Revisited," *Cryptography ePrint Archive*, Report 2012/251, 2012.
- [101] B. W. Lampson, "A note on the confinement problem," *CACM*, 1973.
- [102] C. J. Lee, V. Narasiman, E. Ebrahimi, O. Mutlu, and Y. N. Patt, "DRAM-Aware Last-Level Cache Writeback: Reducing Write-Caused Interference in Memory Systems," *HPS Technical Report*, 2010.
- [103] D. Lee, Y. Kim, G. Pekhimenko, S. Khan, V. Seshadri, K. Chang, and O. Mutlu, "Adaptive-latency DRAM: Optimizing DRAM Timing for the Common-case," in *HPCA*, 2015.
- [104] D. Lee, "Reducing DRAM Latency at Low Cost by Exploiting Heterogeneity," Ph.D. dissertation, Carnegie Mellon University, 2016.
- [105] D. Lee, S. Ghose, G. Pekhimenko, S. Khan, and O. Mutlu, "Simultaneous Multi-Layer Access: Improving 3D-Stacked Memory Bandwidth at Low Cost," in *TACO*, 2016.
- [106] D. Lee, S. Khan, L. Subramanian, S. Ghose, R. Ausavarungnirun, G. Pekhimenko, V. Seshadri, and O. Mutlu, "Design-Induced Latency Variation in Modern DRAM Chips: Characterization, Analysis, and Latency Reduction Mechanisms," in *SIGMETRICS*, 2017.
- [107] D. Lee, Y. Kim, V. Seshadri, J. Liu, L. Subramanian, and O. Mutlu, "Tiered-Latency DRAM: A Low Latency and Low Cost DRAM Architecture," in *HPCA*, 2013.
- [108] D. Lee, L. Subramanian, R. Ausavarungnirun, J. Choi, and O. Mutlu, "Decoupled Direct Memory Access: Isolating CPU and IO Traffic by Leveraging a Dual-Data-Port DRAM," in *PACT*, 2015.
- [109] H. Lee, G. Tyson, and M. Farrens, "Eager Writeback - A Technique for Improving Bandwidth Utilization," in *MICRO*, 2000.
- [110] S. Li, C. Xu, Q. Zou, J. Zhao, Y. Lu, and Y. Xie, "Pinatubo: A Processing-in-Memory Architecture for Bulk Bitwise Operations in Emerging Non-Volatile Memories," in *DAC*, 2016.
- [111] F. Liu, Q. Ge, Y. Yarom, F. Mckeen, C. Rozas, G. Heiser, and R. B. Lee, "Catalyst: Defeating Last-level Cache Side Channel Attacks in Cloud Computing," in *HPCA*, 2016.
- [112] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, "Last-Level Cache Side-Channel Attacks are Practical," in *SP*, 2015.
- [113] J. Liu, B. Jaiyen, Y. Kim, C. Wilkerson, and O. Mutlu, "An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms," in *ISCA*, 2013.
- [114] J. Liu, B. Jaiyen, R. Veras, and O. Mutlu, "RAIDR: Retention-Aware Intelligent DRAM Refresh," in *ISCA*, 2012.
- [115] Z. Liu, I. Calciu, M. Herlihy, and O. Mutlu, "Concurrent Data Structures for Near-Memory Computing," in *SPAA*, 2017.
- [116] X. Lu, L. Zhang, Y. Wang, W. Chen, D. Huang, D. Li, S. Wang, D. He, Z. Yin, Y. Zhou, C. Hui, and Z. Han, "FPGA Based Digital Phase-coding Quantum Key Distribution System," *Science China Physics, Mechanics & Astronomy*, 2015.
- [117] X. Ma, X. Yuan, Z. Cao, B. Qi, and Z. Zhang, "Quantum Random Number Generation," *Quantum Inf.*, 2016.
- [118] M. Majzoobi, F. Koushanfar, and S. Devadas, "FPGA-based True Random Number Generation using Circuit Metastability with Adaptive Feedback Control," in *CHES*, 2011.
- [119] G. Marsaglia, "Xorshift RNGs," *Journal of Statistical Software*, 2003.
- [120] R. Martin, J. Demme, and S. Sethumadhavan, "Timewarp: Rethinking Timekeeping and Performance Monitoring Mechanisms to Mitigate Side-Channel Attacks," in *ISCA*, 2012.
- [121] M. Mascagni and A. Srinivasan, "Algorithm 806: SPRNG: A Scalable Library for Pseudorandom Number Generation," in *TOMS*, 2000.
- [122] S. K. Mathew, S. Srinivasan, M. A. Anders, H. Kaul, S. K. Hsu, F. Sheikh, A. Agarwal, S. Satpathy, and R. K. Krishnamurthy, "2.4 Gbps, 7 mW All-digital PVT-variation Tolerant True Random Number Generator for 45 nm CMOS High-performance Microprocessors," in *JSSC*, 2012.

- [123] M. Matsumoto and T. Nishimura, "Mersenne Twister: A 623-dimensionally Equidistributed Uniform Pseudo-random Number Generator," in *TOMACS*, 1998.
- [124] C. Maurice, C. Neumann, O. Heen, and A. Francillon, "C5: Cross-Cores Cache Covert Channel," in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, 2015.
- [125] J. D. McCalpin, "STREAM: Sustainable Memory Bandwidth in High Performance Computers," Technical Report, 2007.
- [126] I. Molnar, "Modular Scheduler Core and Completely Fair Scheduler," <https://lwn.net/Articles/230501>, 2007.
- [127] A. Morad, L. Yavits, and R. Ginosar, "GP-SIMD Processing-in-Memory," in *TACO*, 2015.
- [128] T. Moscibroda and O. Mutlu, "Memory Performance Attacks: Denial of Memory Service in Multi-Core Systems," in *USENIX Security*, 2007.
- [129] T. Moscibroda and O. Mutlu, "Distributed Order Scheduling and Its Application to Multi-Core DRAM Controllers," in *PODC*, ser. PODC '08. New York, NY, USA: Association for Computing Machinery, 2008.
- [130] J. Mukundan, H. Hunter, K.-h. Kim, J. Stuecheli, and J. F. Martínez, "Understanding and Mitigating Refresh Overheads in High-density DDR4 DRAM Systems," in *ISCA*, 2013.
- [131] J. Mukundan and J. F. Martínez, "MORSE: Multi-objective reconfigurable self-optimizing memory scheduler," in *HPCA*, 2012.
- [132] N. Muralimanohar, R. Balasubramanian, and N. P. Jouppi, "CACTI 6.0: A Tool to Model Large Caches," HP Laboratories, Tech. Rep. HPL-2009-85, 2009.
- [133] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "Processing Data Where it Makes Sense: Enabling In-Memory Computation," *Microprocessors and Microsystems*, 2019.
- [134] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "A modern primer on processing in memory," 2020.
- [135] O. Mutlu and T. Moscibroda, "Stall-Time Fair Memory Access Scheduling for Chip Multiprocessors," in *MICRO*, 2007.
- [136] O. Mutlu and T. Moscibroda, "Parallelism-Aware Batch Scheduling: Enabling High-performance And Fair Shared Memory Controllers," in *ISCA*, 2008.
- [137] K. Nesbit, N. Aggarwal, J. Laudon, and J. Smith, "Fair Queuing Memory Systems," in *MICRO*, 2006.
- [138] A. Olgun, M. Patel, A. G. Yağlıkcı, H. Luo, J. S. Kim, N. Bostancı, N. Vijaykumar, O. Ergin, and O. Mutlu, "QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips," in *ISCA*, 2021.
- [139] L. Orosa, Y. Wang, I. Puddu, M. Sadrosadati, K. Razavi, J. Gómez-Luna, H. Hassan, N. Mansouri-Ghiassi, A. Tavakkol, M. Patel, J. Kim, V. Seshadri, U. Kang, S. Ghose, R. Azevedo, and O. Mutlu, "Dataplant: Enhancing system security with low-cost-in-dram value generation primitives," arXiv:1902.07344, 2019.
- [140] L. Orosa, Y. Wang, I. Puddu, M. Sadrosadati, K. Razavi, J. Gómez-Luna, H. Hassan, N. Mansouri-Ghiassi, A. Tavakkol, M. Patel, J. Kim, V. Seshadri, U. Kang, S. Ghose, R. Azevedo, and O. Mutlu, "CODIC: A Low-cost Substrate for Enabling Custom In-DRAM Functionalities and Optimizations," in *ISCA*, 2021.
- [141] F. Pareschi, G. Setti, and R. Rovatti, "A Fast Chaos-based True Random Number Generator for Cryptographic Applications," in *ESSCIRC*, 2006.
- [142] M. Patel, J. S. Kim, and O. Mutlu, "The Reach Profiler (REAPER): Enabling the Mitigation of DRAM Retention Failures via Profiling at Aggressive Conditions," in *ISCA*, 2017.
- [143] A. Pattnaik, X. Tang, A. Jog, O. Kayiran, A. K. Mishra, M. T. Kandemir, O. Mutlu, and C. R. Das, "Scheduling Techniques for GPU Architectures with Processing-in-Memory Capabilities," in *PACT*, 2016.
- [144] L. Peled, S. Mannor, U. Weiser, and Y. Etsion, "Semantic Locality and Context-Based Prefetching Using Reinforcement Learning," in *ISCA*, 2015.
- [145] C. Percival, "Cache missing for fun and profit," in *BSDCan Ottawa*, 2005.
- [146] C. Pyo, S. Pae, and G. Lee, "DRAM as Source of Randomness," in *IET*, 2009.
- [147] M. K. Qureshi, D. Kim, S. Khan, P. J. Nair, and O. Mutlu, "AVATAR: A Variable-Retention-Time (VRT) Aware Refresh for DRAM Systems," in *DSN*, 2015.
- [148] N. Rafique, W.-T. Lim, and M. Thottethodi, "Effective Management of DRAM Bandwidth in Multicore Processors," in *PACT*, 2007.
- [149] B. Ray and A. Milenković, "True Random Number Generation Using Read Noise of Flash Memory Cells," in *IEEE Transactions on Electron Devices*, 2018.
- [150] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, You, Get off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds," in *CCS*, 2009.
- [151] S. Rixner, W. J. Dally, U. J. Kapasi, P. Mattson, and J. D. Owens, "Memory Access Scheduling," in *ISCA*, 2000.
- [152] V. Seshadri, "Simple DRAM and Virtual Memory Abstractions to Enable Highly Efficient Memory Systems," Ph.D. dissertation, Carnegie Mellon University, 2016.
- [153] V. Seshadri, K. Hsieh, A. Boroum, D. Lee, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry, "Fast Bulk Bitwise AND and OR in DRAM," *IEEE CAL*, 2015.
- [154] V. Seshadri, Y. Kim, C. Fallin, D. Lee, R. Ausavarungnirun, G. Pekhimenko, Y. Luo, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. Mowry, "RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization," in *MICRO*, 2013.
- [155] V. Seshadri, D. Lee, T. Mullins, H. Hassan, A. Boroumand, J. Kim, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry, "Ambit: In-memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology," in *MICRO*, 2017.
- [156] V. Seshadri, T. Mullins, A. Boroumand, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "Gather-scatter DRAM: In-DRAM Address Translation to Improve the Spatial Locality of Non-unit Strided Accesses," in *MICRO*, 2015.
- [157] V. Seshadri and O. Mutlu, "Simple Operations in Memory to Reduce Data Movement," in *Advances in Computers*, 2017.
- [158] G. Singh, D. Diamantopoulos, C. Hagleitner, J. Gomez-Luna, S. Stuijk, O. Mutlu, and H. Corporaal, "NERO: A Near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling," in *FPL*, 2020.
- [159] G. Singh, M. Alser, D. S. Cali, D. Diamantopoulos, J. Gomez-Luna, H. Corporaal, and O. Mutlu, "FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications," *IEEE Micro*, Jul 2021.
- [160] A. Snaveley and D. M. Tullsen, "Symbiotic Jobscheduling for a Simultaneous Multithreading Processor," in *ASPLOS*, 2000.
- [161] G. L. Steele Jr, D. Lea, and C. H. Flood, "Fast Splittable Pseudorandom Number Generators," in *OOPSLA*, 2014.
- [162] J. Stuecheli, D. Kaseridis, D. Daly, H. C. Hunter, and L. K. John, "The Virtual Write Queue: Coordinating DRAM and Last-Level Cache Policies," *SIGARCH Comput. Archit. News*, 2010.
- [163] L. Subramanian, D. Lee, V. Seshadri, H. Rastogi, and O. Mutlu, "The Blacklisting Memory Scheduler: Achieving High Performance And Fairness At Low Cost," in *ICCD*, 2014.
- [164] L. Subramanian, D. Lee, V. Seshadri, H. Rastogi, and O. Mutlu, "BLISS: Balancing Performance, Fairness and Complexity in Memory Access Scheduling," in *TPDS*, 2016.
- [165] Z. Sura, A. Jacob, T. Chen, B. Rosenburg, O. Sallenave, C. Bertolli, S. Antao, J. Brunheroto, Y. Park, K. O'Brien et al., "Data Access Optimization in a Processing-in-Memory System," in *CF*, 2015.
- [166] S. Sutar, A. Raha, D. Kulkarni, R. Shorey, J. Tew, and V. Raghunathan, "D-PUF: An Intrinsically Reconfigurable DRAM PUF for Device Authentication and Random Number Generation," in *TECS*, 2018.
- [167] J. S. Teh, A. Samsudin, M. Al-Mazrooe, and A. Akhavan, "GPUs and Chaos: A New True Random Number Generator," in *Nonlinear Dynamics*, 2015.
- [168] F. Tehranipoor, W. Yan, and J. A. Chandy, "Robust Hardware True Random Number Generators using DRAM Remanence Effects," in *HOST*, 2016.
- [169] G. Thomas, K. Chandrasekar, B. Åkesson, B. Juurlink, and K. Goossens, "A Predictor-Based Power-Saving Policy for DRAM Memories," in *2012 15th Euromicro Conference on Digital System Design*, 2012.
- [170] H. C. v. Tilborg and S. Jajodia, "Encyclopedia of Cryptography and Security," 2011.
- [171] C. Tokunaga, D. Blaauw, and T. Mudge, "True Random Number Generator with a Metastability-based Quality Control," in *JSSC*, 2008.
- [172] Transaction Processing Performance Council, "TPC Benchmarks," <http://tpc.org/>.
- [173] K. H. Tsoi, K. Leung, and P. H. W. Leong, "Compact FPGA-based True and Pseudo Random Number Generators," in *FCCM*, 2003.
- [174] S. Tzeng and L. Wei, "Parallel White Noise Generation on a GPU via Cryptographic Hash," in *ISD*, 2008.
- [175] H. Usui, L. Subramanian, K. K. Chang, and O. Mutlu, "DASH: Deadline-Aware High-performance Memory Scheduler for Heterogeneous Systems with Hardware Accelerators," in *TACO*, 2016.
- [176] V. van der Leest, E. van der Sluis, G.-J. Schrijen, P. Tuyls, and H. Handschuh, "Efficient Implementation of True Random Number Generator Based on SRAM PUFs," in *Cryptography and Security: From Theory to Applications*, 2012.
- [177] R. K. Venkatesan, S. Herr, and E. Rothenberg, "Retention-aware Placement in DRAM (RAPID): Software Methods for Quasi-non-volatile DRAM," in *HPCA*, 2006.
- [178] V. von Kaenel and T. Takayanagi, "Dual True Random Number Generators for Cryptographic Applications Embedded on a 200 Million Device Dual CPU SOC," in *CICC*, 2007.
- [179] Y. Wang, A. Ferraiuolo, D. Zhang, A. C. Myers, and G. E. Suh, "SecDCP: Secure Dynamic Cache Partitioning for Efficient Timing Channel Protection," in *DAC*, 2016.
- [180] Y. Wang, L. Orosa, X. Peng, Y. Guo, S. Ghose, M. Patel, J. S. Kim, J. G. Luna, M. Sadrosadati, N. M. Ghiassi, and O. Mutlu, "FIGARO: Improving System Performance via Fine-Grained In-DRAM Data Relocation and Caching," in *MICRO*, 2020.
- [181] Y. Wang, W. Yu, S. Wu, G. Malysa, G. E. Suh, and E. C. Kan, "Flash Memory for Ubiquitous Hardware Security Functions: True Random Number Generation and Device Fingerprints," in *SP*, 2012.
- [182] Y. Wang, C. Hui, C. Liu, and C. Xu, "Theory and Implementation of a Very High Throughput True Random Number Generator in Field Programmable Gate Array," *Review of Scientific Instruments*, 2016.
- [183] Z. Wang, S. M. Khan, and D. A. Jiménez, "Rank Idle Time Prediction Driven Last-Level Cache Writeback," in *MSPC*, ser. MSPC '12. New York, NY, USA: Association for Computing Machinery, 2012.
- [184] Z. Wang and R. B. Lee, "New Cache Designs for Thwarting Software Cache-based Side Channel Attacks," in *ISCA*, 2007.
- [185] C. J. C. H. Watkins and P. Dayan, "Q-learning," in *Machine Learning*, 1992.
- [186] P. Z. Wiczorek, "An FPGA Implementation of the Resolve Time-Based True Random Number Generator With Quality Control," *IEEE Transactions on Circuits and Systems*, 2014.
- [187] WikiChip, "Cascade Lake SP - Intel," [https://en.wikichip.org/wiki/intel/cores/cascade\\_lake\\_sp](https://en.wikichip.org/wiki/intel/cores/cascade_lake_sp).
- [188] G. Wolrich, D. Bernstein, D. Cutter, C. Dolan, and M. J. Adiletta, "Mapping Requests from a Processing Unit That Uses Memory-Mapped Input-Output Space," US Patent 6,694,380, 2004.
- [189] Z. Wu, Z. Xu, and H. Wang, "Whispers in the Hyper-Space: High-Bandwidth and Reliable Covert Channel Attacks Inside the Cloud," *Transactions on Networking*, 2014.
- [190] Y. Xu, M. Bailey, F. Jahani, K. Joshi, M. Hiltunen, and R. Schlichting, "An Exploration of L2 Cache Covert Channels in Virtualized Environments," in *CCSW*, 2011.
- [191] K. Yang, D. Blaauw, and D. Sylvester, "An All-digital Edge Racing True Random Number Generator Robust Against PVT Variations," in *JSSC*, 2016.
- [192] D. Zhang, N. Jayasena, A. Lyashevsky, J. L. Greathouse, L. Xu, and M. Ignatowski, "TOP-PIM: Throughput-Oriented Programmable Processing in Memory," in *HPDC*, 2014.
- [193] T. Zhang, K. Chen, C. Xu, G. Sun, T. Wang, and Y. Xie, "Half-DRAM: A High-bandwidth and Low-power DRAM Architecture from the Rethinking of Fine-grained Activation," in *ISCA*, 2014.
- [194] T. Zhang, M. Yin, C. Xu, X. Lu, X. Sun, Y. Yang, and R. Huang, "High-speed True Random Number Generation Based on Paired Memristors for Security Electronics," *Nanotechnology*, 2017.
- [195] J. Zhao, O. Mutlu, and Y. Xie, "FIRM: Fair and High-Performance Memory Control for Persistent Memory Systems," in *MICRO*, 2014.
- [196] A. Zouzias, K. Kalaitzidis, and B. Grot, "Branch Prediction as a Reinforcement Learning Problem: Why, How and Case Studies," arXiv:2106.13429, 2021.
- [197] W. K. Zuravleff and T. Robinson, "Controller for a Synchronous DRAM that Maximizes Throughput by Allowing Memory Requests and Commands to be Issued Out of Order," US Patent 5,630,096, 1997.

## A. Appendix

### A.1. Analysis of RNG Applications with Higher RNG Throughput Requirements

Figure 17 compares the performance and fairness results of three designs: (1) the RNG-oblivious baseline, (2) the Greedy Idle Design, and (3) DR-STRaNGe. The figure shows the slowdown of non-RNG applications (top) and RNG applications that require 10 Gb/s RNG throughput (middle) executed on a dual-core system compared to each application’s performance when executed alone on a single core. Figure 17 (bottom) plots system fairness, calculated using the unfairness index metric [49, 128, 135]. We make two observations. First, DR-STRaNGe improves the average performance of non-RNG and RNG applications by 34.9% and 24.5%, respectively. Second, DR-STRaNGe improves the average system fairness by 56.9%, compared to the RNG-oblivious baseline design.



**Figure 17: Performance and fairness in dual-core workloads consisting of non-RNG applications and RNG applications that require 10 Gb/s RNG throughput.**

### A.2. Multicore Workloads

Table 2 shows the dual-core workloads that we use in Section 3. We create 172 workloads, each workload consisting of one non-RNG and one RNG application. We use 43 non-RNG applications across different benchmark suites and create 4 synthetic RNG benchmarks with required RNG throughputs ranging from 640 Mb/s to 5120 Mb/s.

**Table 2: RNG applications**

Multicore Workloads	Non-RNG Applications	RNG Applications
2-core	43 applications	x1 RNG application with 640 Mb/s RNG throughput requirement x1 RNG application with 1280 Mb/s RNG throughput requirement x1 RNG application with 2560 Mb/s RNG throughput requirement x1 RNG application with 5120 Mb/s RNG throughput requirement

Table 3 shows the multicore workloads that we use to evaluate DR-STRaNGe in Section 8. We create 186 workloads for 2-, 4-, 8-, and 16-core systems. We create 43 two-core workloads, each consisting of one non-RNG and one RNG application. In addition, for the four-core configuration we create four workload groups, each consisting of 10 multi-programmed workloads. Each group has 3 different applications from different memory-intensity categories and one synthetic RNG benchmark. We also create 30 multi-programmed workloads for 8-core and 16-core configurations consisting of low, medium, and high memory-intensity applications.

**Table 3: Multicore workloads in Section 8**

<b>Number of Cores</b>	<b>Non-RNG Applications/Workloads</b>	<b>RNG Applications</b>
2-core	43 non-RNG applications	x1 RNG application with 5120 <i>Mb/s</i> RNG throughput requirement x1 RNG application with 640 <i>Mb/s</i> RNG throughput requirement
4-core	40 workloads consisting of non-RNG applications (4 memory-intensity groups, 10 workloads each.)	x1 RNG application with 5120 <i>Mb/s</i> RNG throughput requirement
8-core	30 workloads consisting of non-RNG applications (3 memory-intensity groups, 10 workloads each.)	x1 RNG application with 5120 <i>Mb/s</i> RNG throughput requirement
16-core	30 workloads consisting of non-RNG applications (3 memory-intensity groups, 10 workloads each.)	x1 RNG application with 5120 <i>Mb/s</i> RNG throughput requirement