



**QUEEN'S
UNIVERSITY
BELFAST**

Fast DRAM PUFs on Commodity Devices

Miskelly, J., & O'Neill, M. (2020). Fast DRAM PUFs on Commodity Devices. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 3566-3576. Advance online publication. <https://doi.org/10.1109/TCAD.2020.3012218>

Published in:

IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

© 2020 IEEE.

This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Open Access

This research has been made openly available by Queen's academics and its Open Research team. We would love to hear how access to this research benefits you. – Share your feedback with us: <http://go.qub.ac.uk/oa-feedback>

Fast DRAM PUFs on Commodity Devices

Jack Miskelly and Máire O'Neill, *Senior Member, IEEE*

Abstract—Intrinsic physical unclonable functions (PUFs), which derive hardware identifiers from components already present in a system without modification, are an appealing way to add a layer of hardware rooted security into a system. This is evidenced by the fact that the majority of PUF designs in commercial use today are intrinsic. However, as each intrinsic PUF design is reliant on specific hardware their use is limited to a subset of systems. It is therefore desirable to have practical intrinsic PUF designs for as wide a range of underlying hardware as possible. Most intrinsic PUF designs to date have used memory as the entropy source, with the most well studied type being based on SRAM. More recently designs based on DRAM have been proposed, an appealing prospect considering the ubiquity of that technology. While previous research has demonstrated that entropy can be extracted from DRAM there has not yet been a substantive demonstration of such a PUF operating in real-time on a commodity system. In this article, we present a novel set of algorithms for deriving PUF responses in-runtime from DRAM by altering timing parameters using only software. These algorithms reduce the critical period of system disruption by 96% from 88 ms to 3 ms on average compared to existing designs. We present a large scale dataset derived from 1824 DRAM chips characterized using the proposed design on commodity off-the-shelf desktop hardware running a Linux OS. An analysis of the data shows that in addition to the speed improvements the proposed design shows near ideal (>44%) uniqueness and good (>88%) reliability.

Index Terms—DRAM physical unclonable function (PUF), hardware derived identifier, hardware security, intrinsic PUF, memory PUF, PUF.

I. INTRODUCTION

PHYSICAL unclonable functions (PUFs) are a form of hardware-based security primitive which derive entropy from low level physical variation in components. A PUF typically consists of a circuit that can accept some set of input challenges and derive an output response based on the challenge and entropy extracted from the low level hardware variances in the circuit components. The PUF response acts as a unique “fingerprint” for the circuit as the low-level variations that dictate the response cannot be selected without access to a process orders of magnitude more precise than that used to produce the PUF itself.

Despite close to two decades of research on PUFs there has been minimal adoption of the discrete circuit PUFs which

comprise the majority of research to date. In part, this is due to the difficulty in implementation. PUFs are often touted as being well suited to resource constrained device and in particular IoT devices but introducing new and unfamiliar hardware components into existing designs requires a cost in development resources that on balance does not seem to appeal to many device manufacturers.

It is perhaps because of this that the most widely used sub-category of PUFs at present are Intrinsic PUFs. These PUFs do not use a dedicated circuit to produce the PUF response but rather extract the entropy from components in circuitry already present in the target system. In many senses, this is ideal as it avoids the overheads of discrete circuit PUFs with many Intrinsic PUFs able to be implemented entirely through software on existing devices. However, such PUFs are restricted to devices that contain the necessary base hardware. Due to this, the development of practical Intrinsic PUF designs for as broad a range of base hardware as possible is vital.

A. PUF Fundamentals

The underlying entropy sources for a given PUF design vary substantially. The variance in almost any component property can be used as the entropy source with the right extraction mechanism. For example, in Arbiter PUFs [6], two nominally identical sets of delay stages are used, with an arbiter at the end outputting a 1 or 0 depending on whether the upper or lower path executes faster, which is determined by low-level variance in the delay elements themselves. Similarly, ring-oscillator (RO) PUFs [3], [4] use loops of delay elements to generate frequencies determined by the delay component variances. Comparing two of these frequencies generates a single bit of the response depending on which frequency is higher. Other entropy sources include the startup state of SRAM cells [1], the point-of-failure frequency in overclocked processor cores [5], and the stable state of cross coupled latch cells on FPGA [7].

While the specifics of the hardware implementations are highly varied between PUF designs, at a higher level all PUFs can be conceptualised as a some kind of physically rooted function which cannot be easily cloned. It is sufficient to understand that the PUF function accepts some set of challenges as input and produces a corresponding set of responses as output. Any instance of a PUF should produce a unique response set to the set of challenges. This property is called uniqueness and is derived from the interdevice Hamming distance. However, within a given PUF instance the same challenge should consistently produce the same response under repeated measurements. This property is called reliability and is calculated from the Hamming distance between repeated

Manuscript received April 17, 2020; revised June 17, 2020; accepted July 6, 2020. This article was presented in the International Conference on Embedded Software 2020 and appears as part of the ESWEEK-TCAD special issue. (Corresponding author: Jack Miskelly.)

The authors are with the Centre for Secure Information Technologies, Queen’s University Belfast, Belfast BT7 1NN, U.K. (e-mail: jiskelly08@qub.ac.uk; maire.oneill@qub.ac.uk).

Digital Object Identifier 10.1109/TCAD.2020.3012218

93 measurements of the same PUF. Ideally, a PUF should also
94 show minimal bias toward binary 0 or 1 in the response set.

95 PUFs in practice are most commonly used in two applica-
96 tions. The first and simpler is as a source of secret information
97 for use in cryptographic algorithms. In this case, the PUF,
98 which may have only a single valid challenge and response,
99 provides some secret hardware derived value when prompted
100 which is fed into a standard cryptographic function. In
101 effect, the PUF serves as an alternative to storing the secret
102 information in ROM or similar techniques.

103 The second use case is for device ID and authentication. In
104 this case, it is generally assumed that this process is centrally
105 managed in some way, with a central device controlling the
106 authentication of many other devices. This requires the enrol-
107 ment of each PUF instance on creation. Enrolment involves
108 the characterization of the full set of challenges and responses,
109 such that the controlling device has a record of the expected
110 response of each PUF instance to each possible challenge.
111 When a device needs to be authenticated one of this set of
112 challenges is selected and sent to that device, which passes it
113 to the PUF circuit and returns the response. If the response
114 matches that derived at enrolment then the identity of the target
115 device is verified. This challenge-response verification process
116 is referred to as a PUF query.

117 B. PUF Classification

118 When discussing PUFs it is often useful to classify them as
119 “Weak” or “Strong.” This is not a measure of the attack resis-
120 tance of a PUF but is rather based on the size of the challenge-
121 response space—that is, how many unique challenge-response
122 pairs (CRPs) that are possible for a given PUF design. Weak
123 PUFs have only a single or small amount of CRPs, while
124 strong PUFs have a large CR space with the number of CRPs
125 ideally increasing exponentially as the number of challenge
126 bits is increased.

127 Strong PUFs, due to the large set of CRPs, can in theory
128 be used for device authentication without additional crypto-
129 graphic hardware. However, they are also vulnerable to certain
130 attacks, particularly machine learning (ML)-based attacks.
131 Many of the most well known PUF designs, such as the RO
132 PUF [3], [4] or the Arbiter PUF [6] are strong PUFs. While
133 there have been substantial efforts to mitigate the issues with
134 strong PUFs they have yet to see adoption in the industry on
135 the level of weak PUFs.

136 Weak PUFs are more limited in application but less vulner-
137 able to ML attacks due to the small CRP space not allowing
138 an attacker to build a large enough training dataset. Weak
139 PUFs are mainly used for the generation of small amounts
140 of secret information, as True Random Number Generators,
141 and to generate hardware derived identifiers. Weak PUFs have
142 seen substantial adoption in industry, particularly the SRAM
143 PUF [1], a variant of which is deployed in many commercial
144 FPGA accelerator cards.

145 C. Intrinsic PUFs

146 Intrinsic PUFs are a subset of PUF designs that form the
147 PUF response by extracting entropy from hardware already

148 present in a system. Truly intrinsic PUFs require no modifica-
149 tion or addition of hardware whatsoever instead using existing
150 software and hardware functions to extract the entropy from
151 the underlying hardware. This has obvious benefits as it can
152 add a layer of hardware rooted security into a system while
153 requiring no additional hardware resources. Intrinsic PUFs
154 extract additional security functionality from components that
155 will be used in the system regardless. The main drawback to
156 Intrinsic PUFs is that by their very nature they can only be
157 implemented in systems that have specific hardware already
158 present. While it is entirely possible to add hardware to allow
159 for an Intrinsic PUF, this negates most of the benefits of these
160 PUFs over discrete circuit PUFs.

161 The most well studied intrinsic PUF design is the SRAM
162 PUF [1]. When powered on SRAM cells enter into an unsta-
163 ble state which will fall into either a 1 or 0, with the bias
164 of a cell toward either value determined by variances in the
165 cell components. The PUF response is derived from the set
166 of power-on values of the memory cells. In the most basic
167 form this bit pattern is simply used unaltered as the response
168 with more advanced methods also being used to improve the
169 uniqueness and reliability of the PUF response. As the single
170 PUF challenge in this case is to power cycle the memory mod-
171 ule the viable use cases of this design are limited. Designs for
172 memory intrinsic PUFs based on other memory technologies,
173 such as DRAM have also been proposed. These are discussed
174 in greater detail in Section II. As with the SRAM PUF spe-
175 cific conditions are required for the PUF to be viable in a
176 given system.

177 D. Contributions and Outline

178 In this article, we introduce a novel sets of algorithms
179 for extracting PUF responses from commodity-off-the-shelf
180 (COTS) DRAM memory modules by leveraging internal
181 memory timing parameters. These algorithms provide a purely
182 software-based in-runtime intrinsic DRAM PUF. Further, they
183 improve upon existing comparable designs by greatly low-
184 ering the system disruption required to query the PUF. We
185 demonstrate the proposed design on unaltered COTS comput-
186 ing systems running a live Linux OS. From this, we present
187 a dataset based on 1820 discrete DRAM chips, which is the
188 largest of its kind to date to the best knowledge of the authors
189 at the time of writing. This dataset is provided freely for the
190 use of the research community. Finally, we use this dataset to
191 analyse the quality of the proposed PUF design, demonstrat-
192 ing that it exhibits near ideal uniqueness and good reliability.
193 Further, we demonstrate that at the cost of increased memory
194 overheads the reliability can be improved to a near ideal
195 value. The contributions of this article can be summarized as
196 follows.

- 197 1) A novel implementation of the latency-based DRAM
198 PUF concept with a new set of algorithms for PUF
199 enrolment and query which reduce the critical period
200 of system disruption by up to 96% to below 31 ms in
201 the worst case and 3 ms on average.
- 202 2) A demonstration of the proposed design on live desk-
203 top systems, using only COTS hardware and a widely




Fig. 1. DRAM structure for a standard DIMM [8].

B. DRAM Volatility and the Refresh Cycle

DRAM cells are inherently volatile. Over time due to leakage the value stored in a cell will decay and eventually a bit flip may occur depending on the stable state of that particular cell. In order to retain the memory values the entire memory must be refreshed periodically. This is done by reading the value in each cell and writing it back. In modern DRAM this is controlled and carried out by dedicated refresh cycle circuitry.

C. DRAM Operations and Timings

Due to the multiple steps involved in internal memory operations and the inherent latency of some of these steps certain delays are introduced to ensure correct functionality. These delays are dictated by timing parameters imposed by the memory controller. The safe range for each parameter is specified by the device manufacturer. In most systems these parameters can be changed at the BIOS level, and can be adjusted to improve either the performance or the stability of the memory in a system. If timings are lowered past the lower limit given by the manufacturer the stability of memory operations can no longer be guaranteed resulting in undefined behavior, though depending on the exact circumstances the system may still be able to function with “unsafe” values for some timings.

A timing value that is critical in this work is $tRCD$. This is the required delay between the row address strobe (RAS) signal which selects a row to perform the operation on, and the column address strobe (CAS) signal which selects a column within the row. When $tRCD$ is given a value below the recommended limit the memory enters an undefined state in which read operations return incorrect values for some or all of the bits being read. This property is discussed in greater detail in later sections.

D. DRAM-Based PUFs

As a ubiquitous technology found in a wide range of devices the idea of using DRAM as the core hardware for an Intrinsic PUF has been explored in several works. While some designs exploiting more obscure properties of DRAM have been proposed, such as the Rowhammer Effect PUF [10] the majority of DRAM PUF designs proposed to date fall into one of two categories: 1) the more well studied Retention PUFs which use the refresh cycle to extract entropy and 2) the recently proposed Latency PUFs which use the latency of operations to do so.

E. DRAM Retention PUFs

The DRAM Retention PUF proposed by Tehranipoor *et al.* [2] exploited the volatile nature of DRAM cells to extract entropy based on either the drainage rate or the stability of cells. Nominally the charge of a cell containing a 1 and the rate at which that cell discharges should be fixed. In practice this is not the case as due to process variation in the various components the real values vary within an acceptable range centred around the nominal value. This variance is marginal and hence does not impact the normal functions of the memory.

used OS (Ubuntu Linux). Results from six individual machines are presented.

- 3) A large scale dataset using a measurements from over 1800 COTS DRAM chips from three manufacturers and in two form factors [dual in-line memory module (DIMM) and small outline DIMM (SODIMM)]. This is by a substantive amount the largest dataset of its kind to date and is provided for the use of the research community.
- 4) An analysis of the performance of the proposed design taking into account the standard metrics of uniqueness and reliability, the bias of individual responses, the required query time, and the overheads required to implement the proposed design.

II. BACKGROUND AND LIMITATIONS OF EXISTING DRAM PUFs

A. DRAM Structure

The fundamental storage medium of DRAM is an array of capacitive cells, gated by transistors with each cell storing a single bit of data. The data bit stored depends on the charge of the capacitor—in general a charged capacitor equates to a binary 1 and a discharged capacitor to a binary 0. This is a simplification of the hardware level operations, where in practice there exist both cells and anti-cells which are inverted in terms of what level of charge equates to each value.

These cells are arranged into a grid of rows and columns referred to as a bank. There will usually be many such banks in a single DRAM chip. Often, multiple chips are combined into a rank which share a single chip select. While individual chips are deployed the most recognizable form factor for DRAM is the DIMM which incorporates multiple chips. Another common form factor is the SODIMM which is used in smaller form factor devices.




Fig. 2. Diagram showing the timings used in a burst read operation [9].

291 The Retention PUF exploits this by filling a segment of
 292 DRAM with a bit pattern and then disabling the refresh cycle
 293 of the memory. If the components all had the exact nominal
 294 values this would result in a period of normal operation until
 295 a fixed period had passed, at which point the charge leak-
 296 age would be enough that all the bits would flip. In practice,
 297 however, each cell leaks at a slightly different rate with some
 298 cells flipping after only a short period and others retaining their
 299 value for much longer. Further, the distribution of the fast flip-
 300 ping cells and the slow flipping cells is highly random across
 301 commercial DRAM giving this type of PUF high uniqueness.
 302 By stopping the refresh cycle and reading the memory back
 303 after a fixed period a bit pattern unique to that memory seg-
 304 ment is generated. In a following work it was also shown that
 305 a PUF of this type is fairly resistant to degradation due to
 306 ageing [12].

307 The initial proof-of-concept Retention PUF designs suffered
 308 from the requirement for a lengthy period in the order of
 309 minutes to allow for the decay of cells before the response
 310 could be read. In addition, it was unclear if the PUF response
 311 could be generated in-runtime on a real system without mod-
 312 ification (i.e., as a truly intrinsic PUF). Sutar *et al.* later
 313 proposed means by which the query time could be reduced
 314 to between 20 and 60 s [13] while Schaler, Xiong *et al.* [14]
 315 and Schaller *et al.* [15] demonstrated the viability of using a
 316 Retention PUF on a commodity system, either by generating
 317 the response at boot or by generating the response in-runtime
 318 while manually refreshing key parts of memory to prevent a
 319 system crash.

320 Despite the advances from the original proposal in terms of
 321 practical implementation and query time, there are still seri-
 322 ous drawbacks to the Retention PUF, either the response must
 323 be generated at boot (and may not be regenerated during run-
 324 time) or the response can be generated in-runtime but with

substantial limitations on the system for the query duration
 due to the need to reserve sections of memory and the com-
 putational cost of manually refreshing key areas of memory
 to prevent instability. To date in most designs this period is
 in the order of several minutes and in all designs it is at
 minimum 20 s.

F. DRAM Latency PUFs

331 As a response to the technical challenges inherent to DRAM
 332 Retention PUFs Kim *et al.* [11] proposed the DRAM Latency
 333 PUF. Rather than using cell decay as the entropy source this
 334 design exploits the latency-stability tradeoff of modern DRAM
 335 to extract entropy by placing the memory into a state of unde-
 336 fined behavior through the manipulation of timing parameters.
 337 By lowering specific timing parameters, such as $tRCD$, beyond
 338 the normal lower bound the memory is placed into a state
 339 where read operations do not return the values held in memory
 340 but instead return a pattern derived from both the memory val-
 341 ues and the low-level variances in the DRAM circuitry. The
 342 resultant error patterns are repeatable within a given piece of
 343 hardware and form the PUF response.

344 This PUF has several advantages over the Retention PUF.
 345 Like some of the more recent Retention PUF proposals it
 346 is in theory viable for use in-runtime so long as the system
 347 allows in-runtime changing of timing parameters. This means
 348 it can be used for authentication within a standard challenge-
 349 response framework. Like the Retention PUF it exhibits high
 350 uniqueness. Most importantly it requires only a relatively short
 351 query time in the order of 88 ms to generate an 8 KiB response,
 352 a substantial improvement over the Retention PUF.

353 In this work, we propose a novel PUF based on the latency
 354 PUF concept which works to further minimize the query time.
 355 We will demonstrate that by the careful selection of bit patterns
 356 and a more efficient query scheme that the key period of
 357

358 system disruption can be reduced to just 3 ms while retaining
 359 excellent uniqueness and high reliability.

360 III. PROPOSED PUF DESIGN

361 In this section, we present a set of novel algorithms that
 362 allow for the extraction of PUF responses from COTS DRAM
 363 memory by exploiting the behavior of memory with extremely
 364 reduced timing parameter values. This new approach leads to
 365 drastic reductions in the critical period during which memory
 366 must be reserved for PUF use thus minimizing system dis-
 367 ruption while maintaining near ideal performance in other
 368 metrics.

369 The proposed design is a Weak PUF. The design can pro-
 370 duce a CRP set equal in size to twice the available memory
 371 size. While in many cases this will be relatively large it
 372 increases only linearly with PUF segment size. As such the
 373 operating assumption is that the proposed PUF will be used
 374 as a Weak PUF to generate hardware rooted identifiers. The
 375 enrolment and query algorithms and the threat model have
 376 been chosen with this in mind.

377 A. Challenge-Response Mechanism

378 In the proposed design the PUF challenge consists of a set
 379 of values to set each timing parameter to combined with a
 380 series of offsets used to target sections of the PUF segment.
 381 By default, the timing parameter changes are simply to set
 382 t_{RCD} to the lowest possible value which is sufficient to extract
 383 entropy from the hardware. Supplementary to this is the start
 384 address of the memory segment reserved for the PUF and a bit-
 385 pattern which will be written into it prior to query. As this data
 386 is static and does not reveal any information about the PUF
 387 response to an adversary without root access it can be stored on
 388 the target device in order to reduce the size of the challenge,
 389 or provided as part of the challenge for additional security.
 390 The PUF response consists of a number of bits specified in
 391 the challenge returned as a series of 32 b segments. These
 392 segments are concatenated without further post-processing to
 393 produce the final response.

394 The enrolment process and challenge-response mechanism
 395 used have been designed with the aim of minimizing the query
 396 time (and hence, the amount of system disruption) while max-
 397 imizing the entropy of the PUF. The reasoning for how the
 398 challenge values, memory segment, and bit-pattern are selected
 399 is detailed in the following sections.

400 B. Enrolment

401 The bit pattern held in memory when the PUF is queried
 402 has a significant impact on the generated response. In previous
 403 works fixed bit patterns have been used in which the memory
 404 segment is filled with all 1's, all 0's, or some other regular pat-
 405 tern, such as alternating 1's and 0's. This approach is relatively
 406 simple but does not extract the maximum available entropy
 407 from each memory cell, as we observed that the responses
 408 generated using the same memory segment but inverted bit
 409 patterns were distinct and could exhibit markedly different
 410 characteristics in terms of the probability of bit flip and the

reliability of bit flips. This observation and how it was reached
 is discussed in further detail in Section V.

As such the key goal and contribution of the novel enrol-
 ment algorithms presented here is in the derivation of nonreg-
 ular mixed bit patterns which optimize the PUF response in
 terms of reliability such that fewer repeated queries are needed
 to acquire error-corrected responses, with a corresponding
 improvement in the critical period of system disruption when
 querying the PUF. This new approach also has the advantage
 of exhibiting similar near ideal uniqueness values to compar-
 able works, despite optimizing for reliability and speed rather
 than other metrics.

1) *Changes to the Method of Entropy Extraction:* Unlike
 the approach in [11] where a count of the number of read
 failures is used to determine the PUF response bits, here the
 output of the read operations is used directly as the PUF
 response and all error correction and post-processing is per-
 formed after the query has finished. While the end result is
 fairly similar, our approach allows for a further reduction in
 the amount of system disruption by performing any process-
 ing on the data after reserved memory has been released and
 normal operation restored. In Section V—PUF analysis it will
 be shown that this approach can produce a sufficiently robust
 PUF with the only tradeoff being an increase to memory
 overheads.

2) *Initial Characterisation:* The enrolment process starts
 with characterizing the response of the PUF segment for a bit
 pattern of all 1's and also for an all 0 b pattern over a large
 number of repeated measurements, n . The result of this will
 be a set of n responses for each solid bit pattern.

3) *Granularity of Characterisation:* It may seem intuitive
 that in deriving a mixed bit pattern that the responses should
 be analysed on a bitwise basis. We found that this approach
 while certainly possible produced no better results in prac-
 tice than analysing each 32-b subsegment while requiring a
 substantially more complex enrolment process. Further, if this
 bitwise derivation is used, then a bit pattern equal in size to the
 PUF memory segment must be stored on device, or transmit-
 ted as part of the PUF challenge. As such the 32-b response
 from each single read operation is taken as a discrete unit and
 kept or discarded as a whole. This means the mixed bit pattern
 will consist of alternating blocks of 32-b solid patterns of 1
 or 0 and the mixed pattern can be stored using only 1-b per
 32-b of the output response.

4) *Calculating the Reliability of Subsegments:* The key
 metric to select for in the mixed bit pattern is reliability as this
 allows for minimal repetition during query and corresponding
 reductions in query time. In addition, the reliability of sub-
 segments is sufficiently randomly distributed that selecting for
 reliability does not lower the uniqueness of PUF responses
 significantly from what can be achieved using a regularly
 alternating pattern of 0 and 1. This is detailed further in
 Section V. To calculate reliability for each subsegment first
 simple error correction is performed on the n repeat mea-
 surements to derive a reference value. Reliability for each
 subsegment is then calculated in reference to this value using
 the equation in Section V-B. This operation is performed on
 every characterized sub segment for both solid bit patterns.

469 5) *Filtering of Zero Entropy Responses:* We observed that
 470 in certain devices there would be segments of memory that
 471 produce no bit-flips under any of the possible PUF setups for
 472 one or both of the inputs. Including these segments would
 473 lower the entropy of the response. Further, if selecting for
 474 reliability a disproportionate number of these segments will
 475 be included as they will be near or at 100% reliability. In
 476 cases where one input produces bit flips and the other does
 477 not one value can be discarded immediately and the reliability
 478 calculation skipped for that segment. If neither input produces
 479 any bit flips then the location must be flagged and excluded
 480 from the final CRP set. Again this improves the quality of the
 481 response at the expense of further memory overheads.

482 6) *Deriving the Optimal Bit Pattern:* For all subsegments
 483 which are not yet determined the value of that 32 b subseg-
 484 ment in the mixed bit pattern is simply the value of whichever
 485 bit pattern produced the more reliable response. Thus, the final
 486 pattern should consist of only subsegments in which at least
 487 one pattern exhibited entropy. If only one pattern exhibited
 488 entropy the value of the subsegment in the mixed pattern is
 489 whichever solid pattern produced entropy. If both solid patterns
 490 exhibited entropy then the value in the mixed bit pattern sub-
 491 segment is whichever solid pattern produced the most reliable
 492 response.

493 Optionally a further step can be taken to filter out unreliable
 494 responses by repeating some of the previous step using the
 495 derived bit pattern. This allows a certain degree of control
 496 over the reliability of the overall CRP set but at the expense
 497 of lowering the size of the set. This can be compensated for
 498 by increasing the PUF segment size if the increased overheads
 499 can be tolerated. In Section V—PUF analysis it is shown that
 500 it is possible to use this technique to produce a PUF with a
 501 minimum reliability of 99% for all CRPs.

502 Through burst reads of memory offsets found to exhibit high
 503 reliability and the use of a derived bit pattern which further
 504 enhances reliability the throughput of the PUF is maximized
 505 and the critical period of system disruption minimized. The
 506 result is a PUF which has a greatly reduced query time while
 507 maintaining high performance on key metrics.

508 C. Query

509 Much of the complexity in the operation of the proposed
 510 design is in the initial enrolment phase. Once the PUF has been
 511 enrolled the PUF query that will be performed on the target
 512 system is much simpler and lightweight. Once the bit pattern
 513 and map of memory locations is enrolled to query the PUF all
 514 that is needed is to reserve the PUF memory segment, write the
 515 known bit-pattern into the segment, lower the timing param-
 516 eters to the specified values, and read the memory location
 517 at each specified offset. As the timing parameters associated
 518 with write operations remain unchanged the PUF response can
 519 be stored in any part of memory, including the same area of
 520 memory as the PUF segment if desired.

521 D. Selection of Parameter Values

522 The selection of the timing values is a key factor in getting
 523 a useful CRP set from the proposed design. In all cases, the

Algorithm 1: Characterise PUF for 0 and 1

```

Reserve Memory Channel  $B$ ;
Reserve Segment of size (PUF Segment)*2 in Memory
Channel  $A$ ;
Fill PUF segment  $S$  in  $B$  with 0xFFFFFFFF;
Set timings to  $LOW$ ;
for <large number of reads> do
  while NOT end of segment do
    Read 32 bit double word from next location;
    Store value in array in Memory Channel  $A$ ;
  end
end
Restore original timings;
Fill PUF segment  $S$  in  $B$  with 0x00000000;
Set timings to  $LOW$ ;
for <large number of reads> do
  while NOT end of segment do
    Read 32 bit double word from next location;
    Store value in array in Memory Channel  $A$ ;
  end
end
Restore original timings;

```

Algorithm 2: Generate Optimal Bit Pattern

```

for <all segments> do
  if bit flips in both responses then
    Value = input of
       $\max\{\text{rel}(0x00000000), \text{rel}(0xFFFFFFFF)\}$ ;
  else
    if bit flips in one response then
      Value = input of bit-flip response;
    else
      Exclude segment from CRP set;
    end
  end
end

```

$tRCD$ timing parameter must be lowered enough to initiate 524
 an unstable state. The point at which this state is initiated is 525
 heavily dependant on the specifics of the memory controller 526
 and the DRAM itself. In cases where this point is unknown 527
 or uncertain it is viable to simply select the lowest possible 528
 value for $tRCD$. Other timing parameters may be lowered in 529
 conjunction with $tRCD$ to varying effect but the exact results 530
 of this and the means of selecting these parameters is outside 531
 the scope of this work. 532

During the enrolment stage the number of repeated mea- 533
 surements used to characterize the PUF can be varied, with 534
 smaller numbers of repetitions speeding up the enrolment pro- 535
 cess but higher numbers providing more accurate values for 536
 reliability. In the experiments used in this work 100 rep- 537
 etitions were used. Likewise, when querying the PUF the 538
 number of repetitions presents a similar tradeoff between query 539
 speed and the accuracy of the error corrected response. In all 540

Algorithm 3: Filter CRP Set

```

Reserve Memory Channel  $B$ ;
Reserve Segment of size (PUF Segment)*2 in Memory
Channel  $A$ ;
Fill PUF segment  $S$  in  $B$  with derived bit pattern;
Set timings to  $LOW$ ;
for  $\langle$ large number of reads $\rangle$  do
  while  $NOT$  end of segment do
    | Generate PUF response;
  end
end
Restore original timings;
for  $\langle$ all response segments $\rangle$  do
  Calculate reliability;
  if reliability is above threshold  $T$  then
    | Include segment in final CRP set;
  else
    | Exclude segment from final CRP set;
  end
end

```

541 measurements taken in this work ten repeated measurements
542 were used.

543 IV. EXPERIMENTAL SETUP AND METHODOLOGY

544 The following section details the hardware and process used
545 to derive the dataset used in the analysis in Section V.

546 A. Testing Platforms

547 The platforms used to perform the experiments consisted
548 of a set of three COTS desktop systems for testing DRAM
549 in the DIMM form factor and a further three laptop systems
550 for testing SODIMM form factor memory. Each device in the
551 testbed contained an AMD A-series processor and was running
552 a live instance of Ubuntu Linux 16.04.

553 The PUF itself was implemented as a set of Linux drivers
554 which could be activated to perform individual queries or to
555 perform a full enrolment of the PUF. As all test devices con-
556 tained related models of AMD processors using the same
557 memory controller architecture this driver was able to be
558 ported to each device with only minimal changes. A more
559 detailed description of the technical challenges of implement-
560 ing the PUF is provided in Section V.

561 The datasets were taken from a set of 1824 discrete DRAM
562 chips from three different manufacturers and in two form
563 factors. Of these 120 chips were from Apacer DRAM in
564 SODIMM form factor, 904 were from QUMOX DRAM in
565 DIMM form factor, and the remaining 800 were from Kingston
566 DRAM in DIMM form factor.

567 B. Experimental Methodology and Dataset

568 When extracting data from the test devices the follow-
569 ing methodology was used. From each DRAM chip a

representative 4 KB key was generated for the solid bit patterns 570
0xFFFFFFFF and 0x00000000, respectively. This was taken 571
from a contiguous 4KB segment with the same starting offset 572
relative to each chip. This segment was queried under eleven 573
parameter sets and repeated ten times within each set, produc- 574
ing an array of $11 \times 10 \times 8$ KB for each chip. All tests were 575
performed at nominal supply voltage and at room temperature. 576

In addition to more accurately measure the timing of the 577
critical period of system disruption a single set of ten 40 KB 578
segments was generated from each chip starting at the same 579
base offset under the default timing parameters of $tRCD = 0$. 580

As a key part of this work is the reduction in system interrup- 581
tion in order to measure this accurately the experiments were 582
performed while the system was live and running the same set 583
of background processes. 584

The full dataset is made openly available for the use 585
of the research community and can be downloaded at 586
10.6084/m9.figshare.12149799. 587

588 V. PUF ANALYSIS

In this section, we present an analysis of the gathered data, 589
focusing on the key metrics of reliability, uniqueness, and 590
query time. Further, we give an analysis of the bias of response 591
bits across the gathered responses. Finally, an analysis of the 592
overheads required in implementing the design is provided. 593

594 A. Uniqueness

Uniqueness is a metric that measures how distinct the PUF 595
response of a given instance is likely to be from all the other 596
instances in a population. It is calculated by measuring the 597
interchip Hamming distance between each the devices in a 598
population of m devices. The ideal value of this metric is 50%. 599

The calculation used in this article is as follows, where a key 600
of n bits has been generated from a population of m devices: 601

$$U = \frac{2}{m(m-1)} \sum_{i=1}^{m-1} \sum_{j=i+1}^m \frac{HD(R_i, R_j)}{n} \times 100\%. \quad (1) \quad 602$$

While the memory used in these experiments is all stan- 603
dards compliant DDR memory, the precise specifications and 604
layout of memory components contributing to the entropy of 605
the PUF will not be identical across chips made by different 606
manufacturers. As such we have calculated both the overall 607
uniqueness for the entire population of devices from all man- 608
ufacturers, and on a per manufacturer basis. As can be seen 609
in Fig. 3 and Table I there is a substantial variance in PUF 610
uniqueness from manufacturer to manufacturer. 611

It can be observed from Fig. 3 that when using a solid 612
bit-pattern in memory (0x00000000 or 0xFFFFFFFF) the 613
uniqueness is generally poor and varies substantially between 614
manufacturers, with the median value of Kingston being above 615
25% while the median value for Apacer is below 10%. 616
However, when using a mixed bit-pattern (see Algorithm 4) 617
the uniqueness increases substantially, becoming close to the 618
ideal value and generally more consistent, as can be seen in 619
Fig. 4. 620

This can be attributed to several factors. First the solid 621
bit-pattern responses are heavily biased toward the pattern 622

Algorithm 4: Query PUF

```

Receive challenge;
From challenge extract timing values;
From challenge extract offset list;
Reserve Memory Channel B;
Reserve Segment of size (PUF Segment)*2 in Memory
Channel A;
Write bit pattern into B;
for <x repetitions> do
  Set timings to LOW;
  for <offset 1:n> do
    Read 32 bit dword;
  end
  Restore default timings;
end
Perform error correction;
Concatenate segments into response;
Send response;

```




Fig. 3. Uniqueness of solid bit-pattern 32-b responses.

TABLE I
INTERCHIP UNIQUENESS FOR EACH MANUFACTURER WITH SOLID AND MIXED BIT-PATTERNS

	Solid Bit-Pattern	Mixed Bit-Pattern
OVERALL	18.05%	47.95%
QUMOX	18.91%	49.55%
KINGSTON	25.46%	49.79%
APACER	9.8%	44.53%

value. This reduces uniqueness as all of the response bits are weighted toward a certain value meaning the odds of two chips having similar responses similarly increases. However, as can be seen in Fig. 5 the bias for each bit-patterns is proportional to the other but inverted.

This means so long as the selection criteria used to generate the mixed bit-pattern results in roughly equal amounts of segments with each bit-pattern and that there is no correlation between which segments use each bit-pattern in a given PUF




Fig. 4. Uniqueness of mixed bit-pattern 32-b responses.

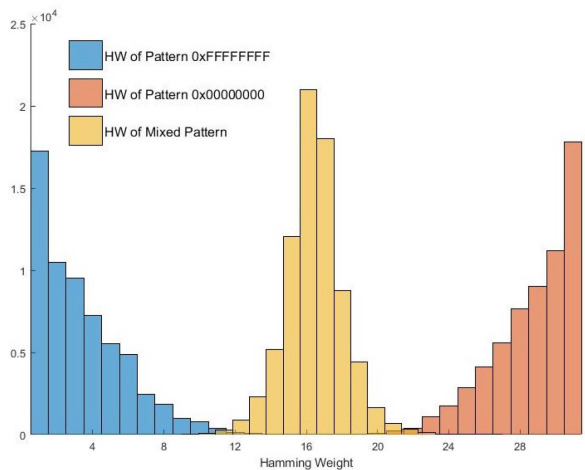


Fig. 5. Hamming Weights of 32-b responses within Apacer chips using various input values.

instance then the biases are canceled out, resulting in minimal bias in the overall PUF responses with mixed bit-patterns. It can be seen in Fig. 5 that this does occur in practice.

The variances which dictate the nature of the PUF response seem to be distributed across the memory structure with a high degree of randomness. This is consistent with the findings of previous works albeit with a more complex set of properties contributing to the overall result. As the PUF is exploiting latency in the internal memory commands some part of the entropy likely comes from the circuitry relating to these operations. Further, some variance was observed between the response of the same chips on different test devices under the same test conditions, implying that the memory controller itself may act as a partial entropy source.

Despite these factors playing some role we propose that the primary entropy source is from the DRAM cells, more specifically in the charge and discharge rates of the cell capacitors. From the data we observed no obvious correlation between the probability of bit-flips in cells when they contain a 1 or 0. That is, a cell which reliably flips when containing one value is no more or less likely to exhibit the same behavior when

TABLE II
RELIABILITY AND PROPORTION OF HIGH RELIABILITY (ABOVE 99%)
32-B SEGMENTS FOR EACH MANUFACTURER

	Reliability	Ideal Segments
OVERALL	93.59%	16.09%
QUMOX	88.66%	2.67%
KINGSTON	95.37%	12.26%
APACER	96.75%	33.35%

653 containing the inverse value. From this the most likely expla-
654 nation is that the cell charge and discharge rate are the main
655 entropy source for each value, respectively, and that these two
656 variables are independent of one another.

657 B. Reliability

658 The ability of the PUF to generate keys consistently is of
659 paramount importance. Reliability is a metric which measures
660 the ability of a PUF design to generate the same response to
661 a given challenge over repeated measurements. In this arti-
662 cle, reliability has been calculated as follows, for m repeat
663 measurements of n bit keys:

$$664 \text{Rel} = \frac{1}{m} \sum_{i=1}^m \frac{HD(R_{\text{ref}}, R_i)}{n} \times 100\%. \quad (2)$$

665 We observed that the manner in which the memory is read
666 during a PUF query has a significant impact on the stability of
667 the PUF. The larger the contiguous segment queried the more
668 unstable the responses became. This is not a product of the
669 DRAM cells as querying the same cells in smaller sections
670 produced stable results. In practice, we found that querying
671 more than a few KB continuously led to a degradation of
672 results. Where larger keys are desired they can be broken down
673 into smaller queries of a sustainable size.

674 Reliability is generally high with values ranging from
675 80%–100%. The worst reliability is found in the QUMOX
676 chips with a mean value of 88%. As an unreliable response
677 requires more error correction when optimizing for minimal
678 query time it is best to use only the most reliable segments. As
679 this excludes some percentage of the CRP set a larger PUF seg-
680 ment is needed to generate a response of the desired size. As
681 can be seen in Table II if a cutoff of 99% is chosen the propor-
682 tion of suitable segments varies highly between manufacturers
683 but remains above 2.5%. These values have been used when
684 calculating the required memory overheads in Section V-C.

685 C. Speed and Overheads

686 In order to minimize the disruption to the overall system the
687 critical period during which memory must be reserved for PUF
688 use must be made as short as possible. In the design proposed
689 in [11] the average time for generating an 8 KiB key reliably is
690 88 ms. As can be seen in Table III our proposed design reduces
691 this critical period to an average of 3 ms per query, a reduction
692 of 96%. Even if ten repeated measurements are taken for the
693 purposes of error correction as in the experiments in this work
694 the query time is only on average 30.3 ms, a reduction of 65%.

TABLE III
MEAN QUERY TIME AND OVERHEAD FOR AN 8 KiB SEGMENT

	Query Time	Overhead
OVERALL	3.03ms	129KiB
QUMOX	2.81ms	299KiB
KINGSTON	3.19ms	65KiB
APACER	3.1ms	24KiB

This improvement comes mainly from two factors. First, 695
by using the results of the read operations directly to form the 696
response and by performing the processing for this formation 697
after the PUF memory has been released and timings restored 698
the period in which the PUF disrupts the system is kept to 699
a minimum. Additionally, using the new proposed enrolment 700
and query scheme to create a set of CRPs with near ideal reli- 701
ability allows for a reduction in the number of read operations 702
required to produce a response. 703

The main drawback of this approach is that it substan- 704
tially increases the memory needed for the PUF. In the worst 705
case only 2.5% of segments provide sufficient reliability to be 706
considered ideal meaning that, as can be seen in Table III, 707
to generate an 8 KiB response requires almost 300 KiB of 708
memory for PUF use. However, in comparison to the capaci- 709
ty of most DRAM chips these overheads are still manageable 710
and the balance of overhead to query time can be controlled 711
by adjusting the threshold at which a segment is considered 712
ideal. Furthermore, these overheads are required only while 713
the PUF is operating and do not need to be reserved exclu- 714
sively for PUF use. It is entirely possible to move data out of 715
the PUF segment, perform the query, and restore the previous 716
values, though this does incur some computational costs. It 717
is also important to note that while a segment of 8 KiB has 718
been evaluated for the purposes of timing analysis in prac- 719
tice the response size needed for device authentication and 720
the corresponding overheads will be much smaller. 721

D. Potential Vulnerabilities and Security Concerns 722

The threat model for fully intrinsic PUFs, such as the 723
proposed design is not identical to that models used with more 724
conventional PUF designs. As such it is useful to discuss some 725
of the contingencies which would need to be considered if 726
using a PUF of this type in practice. 727

The major difference is the method of accessing PUF 728
entropy—namely, the lowering of timing parameters in 729
memory. In systems where this functionality exists it is a normal, 730
if somewhat obscure, function of the memory controller. 731
This means that if an attacker gains root access they can easily 732
trigger the PUF. This is largely consistent with any PUF inte- 733
grated into a complex system. If an adversary has control of 734
the controlling system, they can access the PUF. The PUF, at 735
least in the case of the proposed design, is a mean to authen- 736
ticate devices, not a means in itself to secure the device from 737
illicit access. What is different in fully intrinsic PUFs is that 738
we cannot add any hardware which might limit the rate of this 739
access without rendering them nonintrinsic. 740

741 Should an attacker gain full root access they can perform the
 742 same enrolment procedure as the device owner did to derive
 743 the full set of CRPs. This will take a considerable amount of
 744 time (in the order of 100 or 1000 times longer than generating
 745 the keys from a known challenge) and if they do not have
 746 helper data like the area of memory used for the PUF they
 747 may have to also perform this process for the entire memory
 748 of the device. While this is being done device operation will
 749 be limited by necessity and this may allow the attack to be
 750 detected if it disrupts device operations to the point that other
 751 nodes can detect the discrepancy. However, a clever adversary
 752 will perform this attack piecemeal so as to avoid detection.
 753 Once the full CRP set is collected the attacker has a digital
 754 model of that particular device that can be used to impersonate
 755 it, but this would not be transferable to another device of the
 756 same type as PUF CRP sets are unique to a specific item of
 757 hardware.

758 Each cell provides a distinct entropy source, meaning that
 759 acquiring the CRP data for a given memory segment does not
 760 allow the prediction through modeling of other segments in
 761 the same chip. We can conceptualize the proposed design as
 762 a weak PUF in the vein of SRAM PUF, with two CRPs that
 763 have very large responses and can be divided and recombined
 764 to form a response of the desired size. ML modeling attacks
 765 are therefore not applicable to the proposed design.

766 If an attacker were to gain root access to the system it is
 767 entirely possible to use the PUF query mechanism without the
 768 careful segregation of memory channels in order to destabi-
 769 lize the system and force a crash. It is worth noting this is
 770 not actually a weakness of the PUF design specifically, and
 771 could be done in any system with the right memory controller
 772 regardless of whether it was being used for PUF response gen-
 773 eration. Likewise, a sufficiently knowledgeable adversary with
 774 root access has many options for crashing a system, many of
 775 which are easier to perform than this method.

776 If we assume that the attacker only has root access for a lim-
 777 ited window then acquiring the CRP set becomes a nontrivial
 778 challenge. The enrolment process takes considerable time, and
 779 the attacker must possess helper data to know which memory
 780 segment to target. It would be possible for an attacker to gather
 781 the CRP set relatively quickly if the challenge set is known, but
 782 as this set is not held on device they attacker would have to first
 783 compromise the authentication server or perform some kind of
 784 man-in-the-middle attack to gradually gather the challenge set
 785 by listening in on normal operations.

786 In general, as is the case with any individual security prim-
 787 itive, there are means by which a sufficiently knowledgeable
 788 and resourced adversary can compromise a PUF of this type.
 789 It is important, therefore, to keep in mind that this PUF
 790 provides only an element of hardware derived authentica-
 791 tion and cannot be relied on as a means of device security
 792 in and of itself, but rather should be integrated into exist-
 793 ing security mechanism as an additional assurance of device
 794 authenticity. Hardware rooted security is not a silver bullet. It
 795 is simply another hurdle for adversaries to overcome before
 796 they can compromise a system, and critically which requires
 797 different knowledge to attack than other aspects of system
 798 security.

E. Technical Challenges of Implementation 799

There are several difficulties inherent in implementing a 800
 DRAM latency PUF on commodity hardware. In this section 801
 we discuss the technical challenges which must be overcome 802
 when integrating a PUF of this type into a COTS system. 803

1) *Memory Controller Requirements:* A dynamic memory 804
 controller (DMC) which has in-runtime access to tim- 805
 ing parameters is needed. This is not insurmountable as 806
 many existing architectures use such a memory controller. 807
 Snapdragon processors for mobile phones have this capability, 808
 as do the majority of server and HPC processor architectures. 809
 At the level of desktop computing AMD processors use a 810
 DMC on all multichannel systems. The testbeds used in this 811
 article all use AMD processors for this reason. Nonetheless, 812
 it limits the viable platforms for implementing a PUF of this 813
 type. In architectures, where this functionality is present it is 814
 a low level operation of the memory controller and so must 815
 be implemented at a level where direct control of the memory 816
 controller is possible. A major factor in how well a PUF of 817
 this type can be implemented is the granularity of the memory 818
 controller. Ideally the timing parameters would be modified 819
 for only a single bank of memory, so that disruption to the 820
 system is minimal. In practice depending on the specifics of 821
 the memory controller the granularity may in fact only be to 822
 the level of a chip, a rank, or an entire memory channel. In 823
 the case of our experimental testbeds this function was only 824
 possible at the level of memory channels, meaning that during 825
 PUF operation an entire memory channel had to be rendered 826
 temporarily inaccessible. 827

2) *Maintaining System Stability:* In the case where the 828
 smallest section of memory that can have timing values altered 829
 is relatively large there is a real technical challenge in lowering 830
 these parameters without destabilizing the entire system. This 831
 is because when the timings are lowered to the required levels, 832
 while the data in memory remains unaltered, read operations 833
 during this period will return junk data. If a critical process 834
 attempts to do this it will cause severe problems. This means 835
 that before the PUF is triggered it must be ensured that there 836
 is no data in the same channel being used by the PUF which 837
 is required by any critical process. It must also be ensured 838
 that noncritical processes either do not have any necessary 839
 data held in this part of memory, or that those processes are 840
 paused for the duration of the PUF query. This is not a trivial 841
 task, though it is worth noting that the query time as listed in 842
 Section V-C is for a much larger response than would be used 843
 in practice. The query time required to generate an authentica- 844
 tion ID in a real system would likely be no more than 100 ms 845
 (for a 2048 b ID). 846

Similarly, there are challenges in regards to the management 847
 of the data held in the PUF segment and of the response data 848
 as it is read. The PUF has to write a specific bit pattern into a 849
 specific segment of memory as part of the query procedure. If 850
 data is already present it must be shuffled out of this part of 851
 memory and restored after the PUF query is complete. This 852
 could be avoided by reserving a small amount of memory 853
 exclusively for PUF usage depending on how much memory 854
 is available. Some area of memory must also be used to store 855

856 the response as it is generated. Here, there is less of a problem,
 857 as the alteration of timing parameters used in this design does
 858 not affect the stability of write operations. Hence, the memory
 859 used for this can be in the same memory being used by the
 860 PUF already.

861 VI. CONCLUSION

862 In this article, we have presented a novel method for extract-
 863 ing PUF responses from DRAM memory modules in-runtime
 864 on live COTS systems. We propose a new enrolment and query
 865 system for latency-based DRAM PUFs and have used these
 866 algorithms to generate a large scale dataset of PUF challenge-
 867 responses under various timing parameter adjustments which
 868 is provided for the use of the research community. Our analysis
 869 of this data shows that the proposed design exhibits near ideal
 870 uniqueness and high reliability. Furthermore, the proposed
 871 design reduces the critical period of system disruption by up
 872 to 96% relative to comparable designs though at the cost of
 873 increased memory overheads.

874 ACKNOWLEDGMENT

875 The data from the experiments in this article is available for
 876 the use of the research community and can be downloaded at
 877 [10.6084/m9.figshare.12149799](https://doi.org/10.6084/m9.figshare.12149799).

878 REFERENCES

- 879 [1] J. Guajardo, S. S. Kumar, G.-J. Schrijen, and P. Tuyls, "FPGA intrinsic PUFs and their use for IP protection," in *Proc. Int. Workshop Cryptograph. Hardw. Embedded Syst. (CHES)*, 2007, pp. 63–80. doi: [10.1007/978-3-540-74735-2_5](https://doi.org/10.1007/978-3-540-74735-2_5).
- 880 [2] F. Tehranipoor, N. Karimian, W. Yan, and J. A. Chandy, "DRAM-based intrinsic physically unclonable functions for system-level security and authentication," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 25, no. 3, pp. 1085–1097, Mar. 2017.
- 881 [3] G. E. Suh and S. Devadas, "Physical unclonable functions for device authentication and secret key generation," in *Proc. 44th ACM/IEEE Des. Autom. Conf.*, 2007, pp. 9–14.
- 882 [4] A. Maiti and P. Schaumont, "Improved ring oscillator PUF: An FPGA-friendly secure primitive," *J. Cryptol.*, vol. 24, pp. 375–397, Oct. 2010.
- 883 [5] A. Maiti and P. Schaumont, "A novel microprocessor-intrinsic physical unclonable function," in *Proc. FPL*, 2012, pp. 380–387. [Online]. Available: <https://doi.org/10.1109/FPL.2012.6339208>
- 884 [6] B. Gassend, D. Clarke, M. van Dijk, and S. Devadas, "Silicon physical random functions," in *Proc. 9th ACM Conf. Comput. Commun. Security (CCS)*, 2002, p. 148.
- 885 [7] S. Kumar, J. Guajardo, R. Maes, G.-J. Schrijen, and P. Tuyls, "Extended abstract: The butterfly PUF protecting IP on every FPGA," in *Proc. IEEE Int. Workshop Hardw. Orient. Security Trust (HOST)*, 2008, pp. 67–70, doi: [10.1109/HST.2008.4559053](https://doi.org/10.1109/HST.2008.4559053).
- 886 [8] L. A. Bathen, M. Gottscho, N. Dutt, P. Gupta, and A. Nicolau, "ViPZonE: OS-level memory variability-driven physical address zoning for energy savings," in *Proc. CODES+ISSS*, 2012, pp. 33–42, doi: [10.1145/2380445.2380457](https://doi.org/10.1145/2380445.2380457).

- 906 [9] "256Mb: x4, x8, x16 SDRAM" *MT48LC64M4A2/MT4*
 907 *8LC32M8A2/MT48LC16M16A2 Datasheet*, Micron, Boise, ID,
 908 USA, Nov. 1999.
- 909 [10] A. Schaller *et al.*, "Intrinsic Rowhammer PUFs: Leveraging the rowhammer effect for improved security," in *Proc. HOST*, 2017, pp. 1–7, doi: [10.1109/HST.2017.7951729](https://doi.org/10.1109/HST.2017.7951729).
- 910 [11] J. 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 *Proc. HPCA*, 2018, pp. 194–207, doi: [10.1109/HPCA.2018.00026](https://doi.org/10.1109/HPCA.2018.00026).
- 911 [12] F. Tehranipoor, N. Karimian, W. Yan, and J. A. Chandy, "Investigation of DRAM PUFs reliability under device accelerated aging effects," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, Baltimore, MD, USA, 2017, pp. 1–4, doi: [10.1109/ISCAS.2017.8050629](https://doi.org/10.1109/ISCAS.2017.8050629).
- 912 [13] S. Sutar, A. Raha, and V. Raghunathan, "D-PUF: An intrinsically reconfigurable DRAM PUF for device authentication in embedded systems," in *Proc. Int. Conf. Compilers Archit. Synth. Embedded Syst. (CASES)*, Pittsburgh, PA, USA, 2016, pp. 1–10, doi: [10.1145/2968455.2968519](https://doi.org/10.1145/2968455.2968519).
- 913 [14] W. Xiong *et al.*, *Run-Time Accessible DRAM PUFs in Commodity Devices* (Lecture Notes in Computer Science), pp. 432–453, 2016, doi: [10.1007/978-3-662-53140-2_21](https://doi.org/10.1007/978-3-662-53140-2_21).
- 914 [15] A. Schaller *et al.*, "Decay-based DRAM PUFs in commodity devices," *IEEE Trans. Depend. Secure Comput.*, vol. 16, no. 3, pp. 462–475, May/Jun. 2019, doi: [10.1109/TDSC.2018.2822298](https://doi.org/10.1109/TDSC.2018.2822298).



930 **Jack Miskelly** received the B.Eng. (First-Class
 931 Hons.) and M.Eng. degrees in software and elec-
 932 tronic systems engineering from Queen's University
 933 Belfast, Belfast, U.K., where he is currently pur-
 934 suing the Ph.D. degree with the Centre for Secure
 935 Information Technologies.

936 His research interests include hardware security,
 937 true random number generators, IoT security, and
 938 physical unclonable functions with a focus on intrinsic PUFs.
 939



940 **Máire O'Neill** (Senior Member, IEEE) received
 941 the M.Eng. (with Distinction) and Ph.D. degrees in
 942 electrical and electronic engineering from Queen's
 943 University Belfast, Belfast, U.K., in 1999 and 2002,
 944 respectively.

945 She is the Chair of information security with
 946 Queen's University Belfast and received an EPSRC
 947 leadership fellowship to conduct research into
 948 next generation data security architectures. She
 949 previously held a United Kingdom Royal Academy
 950 of Engineering research fellowship from 2003 to
 951 2008. She has authored two research books and has more than 100 peer-
 952 reviewed conference and journal publications. Her research interests include
 953 hardware cryptographic architectures, side channel analysis, physical unclon-
 954 able functions, and post-quantum cryptography.

955 Dr. O'Neill was awarded a Royal Academy of Engineering Silver Medal
 956 in 2014, which recognizes outstanding personal contribution by an early or
 957 mid-career engineer that has resulted in successful market exploitation.