

Randomness Concerns When Deploying Differential Privacy

Simson L. Garfinkel

US Census Bureau

Suitland, MD

simson.l.garfinkel@census.gov

Philip Leclerc

US Census Bureau

Suitland, MD

philip.leclerc@census.gov

ABSTRACT

The U.S. Census Bureau is using differential privacy (DP) to protect confidential respondent data collected for the 2020 Decennial Census of Population & Housing. The Census Bureau’s DP system is implemented in the Disclosure Avoidance System (DAS) and requires a source of random numbers. We estimate that the 2020 Census will require roughly 90TB of random bytes to protect the person and household tables. Although there are critical differences between cryptography and DP, they have similar requirements for randomness. We review the history of random number generation on deterministic computers. We also review hardware random number generator schemes, including the use of so-called “Lava Lamps” and the Intel Secure Key RDRAND instruction. We finally present our plan for generating random bits in the Amazon Web Services (AWS) environment using AES-CTR-DRBG seeded by mixing bits from `/dev/urandom` and the Intel Secure Key RDSEED instruction, a compromise of our desire to rely on a trusted hardware implementation, the unease of our external reviewers in trusting a hardware-only implementation, and the need to generate so many random bits.

CCS CONCEPTS

• Security and privacy → Privacy protections; • Theory of computation → Theory of database privacy and security; • Software and its engineering → Software verification;

KEYWORDS

Differential privacy, US Census Bureau, Randomness, RDRAND

ACM Reference Format:

Simson L. Garfinkel and Philip Leclerc. 2020. Randomness Concerns When Deploying Differential Privacy. In *19th Workshop on Privacy in the Electronic Society (WPES’20)*, November 9, 2020, Virtual Event, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3411497.3420211>

1 INTRODUCTION

To date, most of the discussion regarding the use of differential privacy for the US 2020 Census of Population and Housing has focused on the impact of DP on accuracy and the suitability of DP’s privacy guarantee (e.g. [5, 19, 36, 53, 67]), and not on the specific details of the Census Bureau’s DP implementation.

This paper is divided into two sections. In the remainder of this section we present the role of DP in the 2020 Census, discuss DP’s

This paper is authored by an employee(s) of the United States Government and is in the public domain. Non-exclusive copying or redistribution is allowed, provided that the article citation is given and the authors and agency are clearly identified as its source.

WPES’20, November 9, 2020, Virtual Event, USA

2020. ACM ISBN 978-1-4503-8086-7/20/11.

<https://doi.org/10.1145/3411497.3420211>

requirements for randomness, contrast DP’s requirements for randomness with those of cryptography, and present related work. Section 2 provides an overview of the DAS, its randomness requirements, and discusses the engineering challenges we encountered.

1.1 DP and the 2020 Census

As described in the *2020 Census Operational Plan* [82], the 2020 Census uses data collected from households supplemented with data from administrative records to create a dataset known as the Census Unedited File (CUF). This file consists of “[a]ll person and household records for the 50 states, D.C., and Puerto Rico.” [82, p.9] This file is used to produce the Census Edited File (CEF). Following the creation of the CEF, the respondent data travels to a purpose-built application called the Disclosure Avoidance System (DAS).

The output of the DAS consists of two microdata sets: one containing person records, and a second in which each record corresponds to a housing unit or group quarters facility.

Disclosure Avoidance is a term used by the Census Bureau to describe techniques employed to limit the risk of a disclosure of respondent information that would be prohibited by Section 9 of the Census Act (U.S. Code Title 13), as interpreted by the Census Bureau’s Data Stewardship Executive Policy Committee (DSEP), which is the Census Bureau’s executive policy-setting organ [85]. In 2017, the Census Bureau announced that it would use DP [20] as the core privacy-conferring mechanism for the 2020 Census [25].¹

As there was no off-the-shelf mechanism for applying DP to a national census, the Census Bureau developed its own. Although DP was created in part with the protection of a national census in mind, the 2020 Decennial Census will be the first time that a national statistics agency has attempted to use DP for the purpose that it was created.

Increased transparency of disclosure avoidance processes was an important goal in the Census Bureau’s adoption of differential privacy. In 1990, the Census Bureau adopted a rules-based “Confidentiality Edit” termed “data swapping” as a privacy protection mechanism for the original data, and a second technique, called “Blank and Impute,” for sample data [46]. However, no formal proof is available that these techniques can provide meaningful privacy guarantees against broad classes of attackers. Nor is it clear that these privacy guarantees are not undermined by the transparent release of implementation details. Consequently, the Census Bureau has not released details concerning either the previous disclosure avoidance techniques’ implementation, nor their impact on data accuracy.

¹An important subtlety is that, though it uses DP subroutines as its core privacy technology, the DAS is not *end-to-end differentially private*, due to the policy requirement that a modest set of “invariant” statistics not be altered by the infusion of DP noise. The mathematically provable privacy guarantees conferred by the DAS are weakened by this requirement, but the DAS’s privacy guarantee is nevertheless precise and provable, and is similar in form to the guarantees offered by pure DP systems.

By moving to DP, the Census Bureau gained the ability to provide formal, mathematical proof that meaningful, precise privacy guarantees hold against broad classes of attackers, and cannot be undermined by the transparent release of algorithm or implementation details. This allows the Census Bureau to directly and explicitly engage with data users and other external audiences, using publicly available data, concerning the implementation details of the DAS. This option is particularly useful for helping to explore and address the trade-offs between privacy loss and accuracy at various DAS parameter settings.²

In the interests of transparency, and to engage external users for preliminary review of the DAS, the Census Bureau released the source code for the DAS that was used for the 2018 End-to-End Census Test [84], including 62,572 lines of Python source code and 516 lines of configuration files.

Continuing its engagement with the user community, in October 2019 the Census Bureau re-released data from the 2010 Census using a prototype for the 2020 Census DAS system. Called the 2010 Demonstration Data Products (2010DDP), this system was the subject of a December 2019 meeting of the Committee on National Statistics, where attendees compared the statistical accuracy of these data products with previous data publications based on the 2010 Census [38]. The source code used to prototype the 2010DDP was released the following month [81]. This code base included 33,853 lines of Python programs and 1263 lines of configuration files.

1.2 DP and Randomness

This article focuses on the use of randomness in, and randomness requirements for, the 2020 DAS. We believe this is of general interest as a reference implementation of a large-scale DP system.

As traditionally defined, DP is an *information-theoretic* requirement that a disclosure avoidance system must satisfy, by which we mean that DP makes no assumptions about computational limitations of attackers; moreover, analyses of specific DP algorithms are most often carried out in real-valued arithmetic and assume access to truly random variables, ignoring the practical subtleties of floating-point arithmetic and random number generation. The use of pseudo-random number generators is of special concern in this article: using a pseudo-random number generator implies that the information-theoretic privacy-loss budget may be larger than claimed, as pseudo-random iterates are not independent in the strict sense required by information-theoretic definitions. To help overcome this obstacle, computationally aware adaptations of DP have been developed that acknowledge and accommodate the use of cryptographically secure pseudo-random number generators (CSPRNGs) by modeling adversaries as computationally bounded [52]. This is a common assumption when designing practical cryptographic systems. Hence, implementations and interpretations of differentially private algorithms can in principle be made consistent with the use of a pseudo-random number generator.

²Readers familiar with the DP literature may note that many DP systems induce error distributions which are “data-independent,” and can therefore be analyzed even without the use of public data sets as a proxy for sensitive data. Although the DAS generates a large number of estimators with data-independent error distributions, the error distributional properties of the MDF are necessarily data-dependent—a side effect of the policy requirement that the DAS generate nonnegative estimates (which is

a single component of the policy requirement that the DAS generate microdata, that can be readily manipulated by Census systems “downstream” from the DAS).

We shall take up this issue in Section 2.6.

1.3 Comparing DP and Cryptography

DP grew in part from ideas in cryptography, and the application of such ideas to the formalized study of privacy, so it is not surprising that there are many parallels between DP and cryptography:

- Applications in both fields have well-defined secrets that need to be maintained indefinitely.³
- Both require strong (pseudo-)randomness guarantees, and ready access to random numbers.
- Side-channel leakage is a threat to implementations of both kinds of systems.
- Failures are hidden: it’s hard to distinguish working systems from compromised systems.
- Both consider security models in which the attacker has full access to the system’s source code, and, depending on context, may consider attackers with arbitrary prior distributions that can fully Bayes’ update (i.e., who possess unlimited expertise, side information, and computational power), or attackers who face computational bounds, or whose background knowledge is in some way limited.
- Given the complexity of algorithm design and requirements for correct implementations, end-users should generally refrain from creating their systems, and instead use algorithms and implementations developed and vetted by experts.

DP and cryptography also have some important differences. Of particular note, the DP threat model is somewhat different from common cryptography threat models. A common cryptography threat model involves three parties: the message sender (“Alice”), the message receiver (“Bob”), and the eavesdropper (“Eve”). The DP threat model, by contrast, has just two parties: the message sender and the message receiver, *who is also the adversary*.⁴

Indeed, a fundamental insight of the DP literature is that, when guarding against general adversaries,⁵ every novel release based on confidential data leaks some information about that confidential data to the recipient: if too many queries are answered too accurately on a confidential database, this necessarily reveals all of the confidential database’s contents. This observation is sometimes called “the Fundamental Law of Information Recovery.”⁶

³Note, however, that in DP, the meaning of “secret” is more subtle, and it is only the “usual” semantic interpretation of the privacy guarantee that does not weaken over time, while some alternative privacy guarantees smoothly degrade as more external knowledge is accumulated.

⁴Although a reviewer of this manuscript noted that Private Information Retrieval seems to have a similar threat model, in that there are also just two parties, in PIR the goal is for the second party to learn *nothing* about the stored information. DP’s goal is more nuanced.

⁵That is, without making special assumptions about the knowledge or computational capacity of attackers.

⁶As of this writing, there seem to be several different theorems that might qualify for this name: all share the property that if too many queries from a given class can be asked by an attacker, with some pre-specified bounds on the noise infused into the query answers before release, then the attacker will be able, with high probability, to reconstruct exactly all bits in the underlying database. The theorems differ in the structure of queries used by the attacker, and in the convergence rate—i.e., the number of queries required to achieve reconstruction. As a practical matter, the query classes treated in these theorems generally differ from those actually released in practice by national statistical agencies. This may be little comfort, given recent, concrete demonstrations of high-efficacy reconstruction attacks.

Put another way, a factor complicating DP over traditional secret-key cryptography is that the information that is intentionally released (or “leaked”) in a DP system is related to the information that needs to be kept confidential (though learning one from the other may require specific background knowledge). For example, the “aggregate statistics” intentionally released as part of the Decennial Census necessarily leak some information about the values of individual Census responses—indeed, were the DP tabulations released in the Decennial Census not functions of the sensitive data, they would be useless. Thus, the use of DP is similar to the use of property-preserving encryption schemes [54] or functional encryption schemes [9], such as order-revealing encryption [18], in that making the protected data more useful also inherently makes it more revealing of the very private information that the technique seeks to protect.

This difference in threat models is natural, as DP’s central motivation is quite different from the traditional application of secret-key cryptography: rather than seeking to restrict the access of unauthorized parties to confidential information while enabling complete access by authorized parties, the goal of the commonly understood DP semantics⁷ is instead to *propose a formal definition for what kinds of attacker inferences qualify as privacy-eroding*, and then to *quantify how much of this private information is leaked* when using a given algorithm to achieve a targeted level of accuracy in a statistical publication.

In the use-cases considered by DP, the typical expectation is that the amount of information an attacker gains from the data release *should* be non-zero, so that the released tabulations can be useful; the privacy loss incurred by individuals as a result of the release must also be non-zero. DP allows for the privacy loss to be precisely defined, quantified, and sharply bounded.

While practical cryptography deployments assume computationally bounded adversaries, DP research often focuses on semantic interpretations with a strictly stronger adversary. The common DP semantic interpretation assumes an attacker that is computationally unbounded with arbitrarily sophisticated algorithms (specifically, they can fully Bayes’ update), and has access to arbitrary “auxiliary knowledge” (e.g., from external data sets, or from directly knowing a data subject) for use in making inferences about the data subject.

The most common semantic guarantee given in the DP literature is a promised bound on *how much more any fixed attacker can learn about any fixed property of a data subject from a differentially private publication based on confidential data than would have been possible had that data subject’s information never been collected in the first place*. As a simple, concrete example, this interpretation of the privacy guarantee justifies statements such as: “For small ϵ , after a release of data based on a differentially private mechanism, an attacker’s confidence that you are of Voting Age cannot be very much larger than it would have been had the same data release been performed but without your data included.”

Put a bit more tersely, we can summarize this by saying that the risk of an attacker learning anything about you (or, in fact, anyone

⁷“Semantics” is a common term for formal interpretations given of one of the several precise privacy guarantees that can be shown to hold when using a DP disclosure avoidance algorithm.

else) is about the same (for small ϵ), regardless of whether you participate in data collection. This upper bound increases smoothly as the privacy-loss “budget” parameter (ϵ) increases⁸.

In secret-key cryptography there is typically only one bound that is relevant: the attacker should be able to make inferences about an encrypted message’s contents no better than they could have without seeing the encrypted message in the first place. In this restricted sense, information-theoretic cryptographic guarantees correspond to the special case of the guarantee intended by DP systems when $\epsilon = 0$. This statement comes with two caveats: first, it ignores the previously discussed differences in attacker models. Second, standard cryptographic guarantees are better identified with non-standard DP semantic guarantees, rather than the usual DP guarantees. DP guarantees typically compare attacker inference after a release to attacker inference in a world where a data subject did not participate, whereas cryptographic guarantees are more strongly related to DP semantics that compare attacker inference after a release to attacker inference before the release. Guarantees of this kind are more demanding⁹ and necessarily degrade—though smoothly—as more auxiliary knowledge is gathered,¹⁰ unlike the usual DP guarantees (except when $\epsilon = 0$) where they correspond to the usual information-theoretic secrecy promises desired for encrypted messages.

Because $\epsilon = 0$ implies that released tabulations will be useless, deploying and using a DP system *inherently* involves making and understanding social choices and economics. Setting the privacy-loss budget (ϵ) fundamentally requires making a trade-off between the usefulness of the data release and preserving confidentiality. The data custodian must determine the cost of leakage and the benefit of tabulation release at a given accuracy. How best to resolve this trade-off is necessarily a *policy* question; algorithm designers can help to provide more *efficient* algorithms, with higher accuracy for a given ϵ , but the “correct” choice of ϵ is not a question that can be resolved through design of better algorithms. The central research question in much of the DP literature is, therefore, whether there are more efficient mechanisms that have more statistical utility for the same ϵ .

This paper focuses on a property DP shares with cryptography: the need for large amounts of high-quality random numbers. For use of DP in large-scale applications, randomness requirements are driven not by the memory footprint of the underlying microdata, but by the number and scale of output tabulations published. In the DAS’s case, the randomness required is further driven by the need to build intermediate “histograms”—counts of synthetic records of all possible types—in order to readily convert DP statistics into microdata, and to provide some (typically small) expenditure of

⁸It is notable that there is no finite ϵ at which these bounds become meaningless; they always impose some non-degenerate bound on attacker learning. However, they do become loose quite quickly as ϵ grows, as the bounds depend exponentially on ϵ . In cases where a small ϵ cannot be justified, alternative semantic statements can be made that apply with smaller bounds, but only by qualitatively weakening the privacy guarantee—for example, we may have to compare an attacker’s beliefs relative to a world in which only a portion of a person’s data record was not used, rather than their entire record not being used.

⁹These guarantees essentially involve considering the change before collecting any data, not just if a single person’s data had not been collected.

¹⁰The reason for this is that auxiliary knowledge can involve learning about probabilistic dependence (or correlation, colloquially) between distinct persons’ records, which allows for improved inference about a target person, using information concerning other persons’ records.

privacy-loss budget on even arbitrarily complex statistics of the true data. Thus, while the data for the Decennial Census can be stored in a few tens of gigabytes, protecting its output statistics will require the DAS to use roughly 90TB of random data.

1.4 Related Work

1.4.1 Historical Roots. The history of random numbers is littered with the corpses of methods that were known not to be truly random when they were deployed but were incorrectly thought to be good enough for the task at hand.

Tippet published a book of random numbers for use by computers¹¹ in 1927 [77]. Following in this tradition, the RAND Corporation published *A Million Random Digits with 100,000 Normal Deviates* in 1955¹² [66]. During production, RAND discovered that the “electronic roulette wheel” built for generating the numbers exhibited statistical bias and required adjustment prior to publication [1]. Such books of random numbers pose two practical problems: the sequence is available to anyone who has a copy of the book, making the numbers unsuitable for security or privacy applications. And since the book is printed, the numbers are not easily available for use by electronic computers.

Indeed, early *electronic* computing efforts at Los Alamos required large quantities of random numbers implement Stanislaw Ulam’s “Monte Carlo Method” [50]. Early computers did not have an intentional source of usable randomness, so von Neumann invented the “middle-square” method, which generates a stream of digits by starting with an n -digit integer, squaring it, and then extracting the middle n digits [87]. von Neumann recognized that the resulting sequence of digits, while seemingly randomly, were entirely predictable. Reflecting on this contradiction, von Neumann wrote:

“Anyone who considers arithmetical methods of producing random digits is, of course, in a state of sin. For, as has been pointed out several times, there is no such thing as a random number—there are only methods to produce random numbers, and a strict arithmetic procedure of course is not such a method.” [87]

1.4.2 Statistical requirements for randomness. Every modern statistics package provides facilities for generating seemingly random values. Typically these numbers are drawn from a specific distribution such as a uniform distribution, a Gaussian distribution, or some other named and well-studied distribution. Internally, modern statistics packages implement a pseudo-random number generator (PRNG) that takes a single value as a *seed* and emits a sequential series of numbers, qualitatively similar to von Neumann’s method but typically with a longer period. As with von Neumann’s method, if the same seed is provided for two runs, those two runs will present the same sequence of random values. This measure of repeatability is useful when developing statistical programs because it allows for regression testing.

In 1998, Matsumoto *et al.* introduced the Mersenne Twister, a fast pseudo-random number generator based on a linear congruential generator with a period of $2^{19937} - 1$. Known as MT19937,

¹¹That is, human computers [32].

¹²RAND’s book also found uses outside of computing: “a nuclear submarine commander kept a copy of the book with him to chart courses during evasive maneuvers.” [3].

the generator was widely adopted—for example, it was adopted by the popular Python NumPy numeric library.

In 2007 L’Ecuyer and Simard introduced *TestU01* [45], a collection of 144 statistical tests for random number generators, mostly gathered from the literature, and developed three test suites for randomness: *SmallCrush*, *Crush* and *BigCrush*. MT19937 failed the linearity tests of *Crush* and *BigCrush*, which are specifically designed to detect linear generators.

In July 2019, the Python NumPy scientific library 1.17.0 was released, with an upgraded random number system that allows the use of pluggable “bit generators” to produce the random distributions. This allows any source of seemingly random bits to be used for generating a sequence of seemingly random integers in a given range or floating-point numbers that match a desired distribution. In this way, questions of speed, reproducibility, and even suitability for operation in a parallelized environment are pushed from NumPy down to the design and implementation of the bit generator. NumPy version 1.19 includes support for four bit generators: MT19937, PCG 64 [58, 86], Philox [69], and SFC64 [17].

1.4.3 Cryptographic requirements for randomness. Most modern cryptographic systems are based on Kerckhoffs’s principle, that the design should be public and all of the security provided by the system should reside in the secrecy of the cryptographic key. As a result, cryptographic applications require sources of randomness to create unguessable keys.

The underlying rationale for using Kerckhoff’s principle is that strength of a secure systems depends both on the strength of the algorithm *and* on the inability of the attacker to determine the data protection key. Algorithms are hard to design and, once deployed, they are hard to replace. Keys, in contrast, should be relatively easy to change. So it makes sense for most cryptography users to rely on algorithms and implementations that have been publicly developed and vetted, for the simple reason that most users lack the resources to do as good a job.

Eastlake *et al.* [22, 23] extensively documents the requirements for randomness in modern computer systems.

Cryptography researchers have spent considerable effort developing and analyzing random number generators [30, 68]. A common design is to create an *entropy pool*, which is typically implemented as a buffer into which entropy from some source of perceived randomness is added through a bit-mixing function, and from which bits are extracted through the use of a cipher or cryptographic hash function.¹³

The American National Standards Institute (ANSI) adopted the ANSI X9.17/X.931 standard for Financial Institution Key Management (Wholesale) in 1985. The standard presented a design for a secure pseudo-random number generator that combines timestamps with a statically keyed block cipher to produce the pseudo-random output. This design was widely adopted, even though it obviously requires that the static key be kept secret.

¹³For readers unfamiliar with the term *entropy pool*, imagine a jar filled with red and blue marbles that is periodically stirred by the computer’s background tasks. When a program wants random numbers, a mechanical arm reaches into the jar, causing additional mixing, and emerges with a fistful of random bits. These are read in sequence and then returned to the jar.

Cohen *et al.* performed a “systematic study of publicly available FIPS 140-2 certifications for hundreds of products that implemented the ANSI X9.31 random number generator, and found twelve whose certification documents use of static, hard-coded keys in source code, leaving the implementation vulnerable to an attacker who can learn this key from the source code or binary” [12]. One conclusion of the study is that certified implementations of cryptographic systems are not necessarily secure implementations, and that even professional programmers make mistakes when it comes to implementing random number generators, even when they are following a widely used standard.

Other attacks have been found against the ANSI X9.31 standard (e.g. [40]), demonstrating that the mere fact that an algorithm has been standardized does not imply that it is secure.

Reviewers of the Census Bureau’s initial decision to use Intel’s Secure Key (ISK) as the sole source of randomness cited the experience with ANSI X9.31 as an example of why it is important to build systems that do not have a single point of failure—for example, by using multiple sources of entropy and by frequently reseeding the system’s CSPRNG.

1.4.4 Non-determinism in Linux. Although the EDVAC, the ED-SAC, and other computers designed in von Neumann’s time were intended to be deterministic machines—the occasional moth no longer being of concern given the field’s transition to electronic tube technology—modern computers have many sources of non-determinism. Weaver *et al.* identified “operating system interaction [12], program layout [13], [1], measurement overhead [14], multi-processor variation [15], and hardware implementation details [13], [16]” as potential sources of non-determinism [88]. (Note: references in the Weaver quotation refer to those in Weaver [88], not the references in this paper.) However, Das *et al.* suggested that it is a mistake to use such sources for security purposes, in part because they are either insufficiently unpredictable, or they are susceptible to manipulation [14].

Nevertheless, there has been considerable attention to the use of such non-determinism as a source of entropy for entropy pools, the most prominent example probably being the Linux Random Number Generator (LRNG).

Müller authored a 196-page report analyzing the LRNG in kernels 4.10 through 4.20 and 5.0 through 5.5 [75]. Müller’s assessment states “[t]he goal of the assessment is to determine whether the Linux-RNG is able to provide 100 bits, the threshold defined by [TR021021], of entropy early after a system boot” and concludes:

“Applying the general Linux-RNG entropy heuristics, the Linux-RNG significantly underestimates the available entropy...This allows the conclusion that when the *getrandom* system call unblocks, sufficient entropy has been accumulated to be available for use cases with strong cryptographic requirements. The measurements of the available entropy during boot for virtual environments and native hardware hardly differ. Thus, the conclusion is equally applicable to both environments.

“It is important to note that this conclusion is only applicable to environments with a high-resolution time stamp. Hardware architectures with a low-resolution time stamp will not have significant amounts of entropy after boot.” [75]

The Amazon Elastic Map Reduce Kernel used by the Census Bureau is Linux 4.14.128, and so it is covered by Müller’s report. The source code for the version 4.14.128 kernel’s random device can be downloaded from <https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux.git/tree/drivers/char/random.c?h=v4.14.128>. The Linux 4.14.128 kernel’s random device is ©2017 by Jason A. Donefeld and is 2343 lines long, of which 658 lines are comments, addressing some of the concerns previously raised by Gutterman *et al.* [33].

The Linux kernel maintains two entropy pools: a larger pool that receives entropy from the top-half of the kernel using the *_mix_pool_bytes* function, and a smaller entropy pool that receives bytes mixed in using the *fast_mix* function that is called during system interrupts serviced by the bottom half of the kernel.¹⁴

Pseudorandom bytes are extracting by running the CHACHA20 stream cipher [55] over a portion of the larger pool, which the source code claims creates a CRNG (cryptographically strong random number generator). (The Linux version 3 kernel extracts bytes using the SHA-1 cryptographic hash function; Linux stopped using SHA-1 in version 4.8.)

We tested the r5.24xlarge AWS Linux VMs used by the Census Bureau and found that the systems’s *gettimeofday()* system call had microsecond resolution, in that we were able to observe single microsecond increments in the time returned by the system call as it was repeatedly called from a tight loop in a C program.

RedHat Linux provides two device interfaces to the randomness: */dev/urandom*, which is the output of the ostensible CSPRNG,¹⁵ and */dev/random*, which maintains a counter of the amount of entropy that has been added to the entropy pool and blocks if there is not sufficient entropy remaining until more entropy has been added.

The *random* device exports four interfaces that are meant to be called from other parts of the kernel to add entropy:

- add_device_randomness()*** adds information such as “MAC addresses or serial numbers, or the output of the RTC (real time clock)” that are “likely to differ between two devices.”
- add_input_randomness()*** adds randomness from user input, such as mouse movements or keyboard strikes.
- add_interrupt_randomness()*** adds randomness from the interrupt layer
- add_disk_randomness()*** adds randomness based on the disk seek times.

At this point, the reader may be concerned that a Linux kernel running in a data center may lack a sufficient source of entropy, at least immediately following system start-up:

- (1) MAC addresses and serial numbers are predictable, as is the output of the RTC.
- (2) There is no mouse or keyboard on servers in a data center, so these are not sources of randomness.
- (3) Hardware interrupts during the boot process are predictable.
- (4) Disk seek times are predictable on systems equipped with solid-state drives, a fact noted in the source code.

¹⁴The “top-half of the kernel” refers to the code that is invoked by system calls, while the “bottom-half of the kernel” is the portion that is invoked in response to hardware interrupts. Although the terms are widely used and appear in the kernel source code, we were unable to find a suitable reference to their origin

¹⁵We say “ostensible” not because we doubt whether */dev/urandom*’s use of CHACHA20 constitutes a CSPRNG, but because we are aware of no formal proof of */dev/urandom*’s security properties.

The Linux source code acknowledges this possibility, and provides a small script that saves 512 random bytes from `/dev/urandom` into the file `/var/run/random-seed` at system shutdown, and then copies this file back into `/dev/urandom` at system startup. The source code recommends that these shell scripts be used to preserve the entropy pool between reboots. But this approach does not work with Amazon’s Elastic Map Reduce (EMR): since virtual machines are cloned from a master image, there is no ability for each VM to have its own, unique entropy pool carried in the file `/var/run/random-seed` between system reboots. This is, in fact, the very situation that Müller’s report attempts to address!

Recognizing that the data center needed an improved source of randomness, Linux added support for CPU-based hardware random number generation in 2018 [80], as discussed in §1.4.7.

1.4.5 Hardware Random Number Generators. An alternative approach to using machine randomness is to generate random numbers using an external entropy source. Such an approach based on the movements of the liquid within a Lava Lamp and captured by a digital camera is described by Noll *et al.* in US Patent 5,732,138 [56]. The Internet hosting company Cloudflare famously has a wall of Lava Lamps in its office and uses them to seed the allegedly cryptographically secure pseudo-random number generators that the company reports using in its Internet services [44], and has more recently developed an entropy service that mixes entropy from five sources in different countries, using a combination of Lava Lamps, seismic sources in Chile, environmental noise, extraterrestrial noise, and other sources [10, 15].

We attempted to evaluate whether Lava Lamps in fact provide sufficient randomness for DP after a Census Bureau official suggested using them as a source of randomness.

Lava Lamps were invented in 1963 by Edward Walker, and gained popularity in the 1970s. The lamps use an incandescent light bulb to heat blob of wax inside another material. As the wax heats up, it moves upwards through the liquid, at which point it cools down and descends. The wax blob’s movements are governed by thermodynamically controlled microcurrents and is often regarded to represent a chaotic system, because small variations in energy distribution resulting from micro-fluctuations of the power line and in the air surrounding the lamp are amplified and result in unpredictable movements of the wax. US Patent 5,732,138 conjectures that Lava Lamps are “chaotic systems,” and states that the patent could be implemented “for example, by taking pictures of a moving freeway, clouds, or lava lamps” [56]. However, a conjecture in a patent is not scientific validation. It is also unclear if the patent’s invention requires that the functioning Lava Lamp be a chaotic system in the formal sense of that phrase, and not merely an unpredictable system. In any event, we were informed by a Census Bureau safety officer that Lava Lamps are prohibited from the US Census Bureau’s Suitland, MD, headquarters due to fire-safety concerns. Fortunately, thermal noise is also present at the atomic level, and its use there does not constitute a fire hazard.

Another approach is to rely directly on quantum mechanics as a source of randomness. ID Quantique introduced such a true random number generator in 2001; the device has most recently been reduced to a silicon chip for inclusion in 5G smartphones [39]. Los Alamos National Laboratory partnered with Whitewood Security

to create the Entropy Engine [24, 37], a hardware random number that plugs into a PCI Express slot and can generate 350 Mbit/s of true random numbers.

In January 1999, Intel announced that the forthcoming Pentium III microprocessor would include a hardware true random number generator (TRNG) and a unique processor serial number (PSN) in each chip. The TRNG implementation sampled thermal noise 32 bits at a time into a shift register which was then processed with SHA-1. So long as a few bits of the noise change from moment-to-moment, the output of the SHA-1 function is thought to be unpredictable. Reviewing Intel’s published design, Guttman warned that the generator could fail without detection [34, p.238].

Before the Pentium III’s random number generator could be subject to further analysis, its PSN was attacked by privacy activists and, eventually, a legislative panel of the European Union [13]. Under pressure, Intel removed the PSN from the “Taulatin” (130nm) version of the Pentium III and it was not present in the Pentium IV. The random number generator was on the same section of the chip as the PSN, resulting in the loss of the TRNG as well. However, updated versions of both returned (with little fanfare) in the “Ivy Bridge” (22nm) series of Intel’s Core processors (Core i3, i5 and i7) [62]. The updated PSN is now termed the PPIN (Protected Processor Identification Number). The updated random number generator, code-named Bull Mountain Technology, now has the official name Intel Secure Key [49] (ISK).

1.4.6 Intel Secure Key. In this section we review the extensive documentation regarding Intel’s hardware random number generator, and discuss the controversy surrounding its adoption.

Intel states that the design requirements for a random number generator (RNG) is that each new value be *statistically independent*, that the numbers be *uniformly distributed*, and that the sequence be *unpredictable*, in that “an attacker cannot guess some or all of the values in a generated sequence. Predictability may take the form of *forward* prediction (future values) and *backtracking* (past values)” [49].

The Intel “Digital Random Number Generator” (DRNG) is implemented as a module that is separate from the cores on Intel’s multi-core chips. The hardware entropy source is a noisy asynchronous self-timed circuit that outputs a random stream of bits at 3 GHz. The hardware entropy source feeds into a hardware AES-CBC-MAC based “conditioner” to “spread” the “entropy sample into a large set of random values.” It is not a conventional entropy pool, in that the contents are flushed with every random draw.

Intel provides two unprivileged user-level instructions for accessing the DRNG: RDSEED and RDRAND [48]. Both instructions are available in 16, 32 and 64-bit versions that allow seemingly random bits to be stored in the designated destination register.

The RDSEED instruction passes the output of the AES-CBC-MAC based conditioner through a “non-deterministic random bit generator” that is compliant with NIST SP 800-90B and C (drafts) (as of November 17, 2012) and provides the bits directly to the caller [49]. According to Intel, “RDSEED is intended for seeding a software PRNG of arbitrary width.”

The RDRAND instruction causes the output of the AES-CBC-MAC based conditioner to be used as an input for a AES-256 circuit operating in CTR mode. The RDRAND instruction draws from

the output of the AES circuit, with an upper-bound of 511 128-bit samples being used for each random seed. Intel states that this is a “cryptographically secure pseudo-random number generator” (CSPRNG) that is compliant with NIST SP 800-90A.¹⁶

Intel uses the terms *multiplicative prediction resistance* and *additive prediction resistance* to describe the difference in security between the RDSEED and the RDRAND instructions. We could find no other reference for these terms, so we provide Intel’s:

“The numbers returned by RDSEED have multiplicative prediction resistance. If you use two 64-bit samples with multiplicative prediction resistance to build a 128-bit value, you end up with a random number with 128 bits of prediction resistance ($2^{128} \times 2^{128} = 2^{256}$). Combine two of those 128-bit values together, and you get a 256-bit number with 256 bits of prediction resistance. You can continue in this fashion to build a random value of arbitrary width and the prediction resistance will always scale with it. Because its values have multiplicative prediction resistance, RDSEED is intended for seeding other PRNGs.

“In contrast, RDRAND is the output of a 128-bit PRNG that is compliant to NIST SP 800-90A. It is intended for applications that simply need high-quality random numbers. The numbers returned by RDRAND have additive prediction resistance because they are the output of a pseudo-random number generator. If you put two 64-bit values with additive prediction resistance together, the prediction resistance of the resulting value is only 65 bits ($2^{64} + 2^{64} = 2^{65}$). To ensure that RDRAND values are fully prediction-resistant when combined together to build larger values you can follow the procedures in the DRNG Software Implementation Guide on generating seed values from RDRAND, but it’s generally best and simplest to just use RDSEED for PRNG seeding.” [48]

The DRNG monitors its output using Online Health Tests (OHTs) and Built-In Self Tests (BISTs) and shuts the system down if the output of the DRNG fails to meet statistical quality tests. As of 2014, the system could produce a maximum of 800 MB/sec of random data, an upper bound for all threads and all cores on the CPU. Because of this limit, if a random value is not available, the RDRAND and RDSEED instructions set the CPU carry flag (CF) if the returned value is actually random; otherwise CF is set to zero and the destination register is cleared. Intel recommends that applications calling RDRAND attempt 10 retries when in a tight loop if either instruction returns and CF is not set. For RDSEED, Intel recommends that a PAUSE instruction be inserted in the retry loop, and a maximum of 100 retries be performed. It is not clear what a program should do when the retry limit is exceeded.

Intel notes that RDSEED is not available on all processors, and that it can be simulated on such processors that have RDRAND by using RDRAND to generate 512 128-bit samples and cryptographically mixing the results to assure that one of the values was a fresh value from the DRNG and not the result of AES counter mode.

¹⁶As with the Linux random number generator’s use of CHACHA20, the claim that Intel’s use of AES constitutes a CSPRNG is not, as far as we are aware, supported by formal proof. Instead, the security of AES is justified by the fact that decades of study has failed to discover significant vulnerabilities in it.

1.4.7 ISK Adoption and Controversy. In 2012, Cryptography Research (a consulting firm that was highly respected in the cryptography community) conducted an independent review of the DRNG and concluded “the Ivy Bridge RNG is a robust design with a large margin of safety that ensures good random data is generated even if the ES [Entropy Source] is not operating as well as predicted.” [35]

In 2015, Shrimpton and Terashima presented “A Provable-Security Analysis of Intel’s Secure Key RNG” at EUROCRYPT [74] and concluded that the security guarantees offered by Intel Secure Key “tell a mixed story.”

“We find that ISK-RNG lacks backward-security altogether, and that the forward-security bound for the “truly random” bits fetched by the RDSEED instruction is potentially worrisome. On the other hand, we are able to prove stronger forward-security bounds for the pseudorandom bits fetched by the RDRAND instruction.” [74]

ISK generated considerable controversy in the Linux community. Hardware random number generators would seem ideal for providing a source of randomness in a data center, especially if one boots virtual machine snapshots as is the case with Amazon Web Services. For this reason, the Linux kernel includes the functions `arch_get_random_seed_long()` and `arch_get_random_long()` to return hardware-generated random numbers. On Intel-based systems, these functions gateway to the RDRAND instruction, and the output of that instruction is added to the Linux entropy pool.

Nevertheless, by 2013 there was already broad knowledge of ISK within the Linux community and a growing desire on the part of some developers to use it, while a reluctance on the part of others to do so. After all, ISK could not be readily audited, because it was implemented in silicon. In 2013, a petition on *change.org* requested that Linux remove the use of the RDRAND instruction from the `/dev/random` device. The fear was that a hardware backdoor might give Intel, and perhaps other organizations, the ability to predict its random output.

The concern over ISK is similar to concerns that were raised regarding the adoption of the Dual Elliptic Curve Deterministic Random Bit Generator (Dual_EC_DRBG) in NIST SP800-90, “Recommendation for Random Number Generation Using Deterministic Random Bit Generators (Revised)” [8]. The Dual_EC_DRBG algorithm depends on several pre-specified constants, and the way that those constants were created can make the algorithm vulnerable to attack. The summer that Dual_EC_DRBG was proposed, concerns were raised that there might be a “secret backdoor” in the standard [70]. In 2013, those concerns about Dual_EC_DRBG were confirmed [60]. NIST responded by issuing guidance stating “NIST strongly recommends that, pending the resolution of the security concerns and the re-issuance of SP 800-90A, the Dual_EC_DRBG, as specified in the January 2012 version of SP 800-90A, no longer be used.” [41]¹⁷ A 2015 article by the Director of Research at the National Security Agency described the agency’s “failure to drop support for the Dual_EC_DRBG” after vulnerabilities were identified in 2007 as “regrettable.” [89].

¹⁷To avoid this sort of problem, other algorithms that require constants sometimes rely on a so-called “nothing-up-my-sleeve number,” in which the number is chosen from a well-known mathematical sequence, such as the decimal expansion of π . For example, the SHA-2 Secure Hash Algorithm uses the square roots and cube roots of small primes for its constants.

Linux’s inventor Linus Torvalds refused to remove support for RDRAND from the kernel, because RDRAND’s entropy is mixed into the Linux entropy pool, and not used to replace the Linux entropy pool. Torvalds responded to the *change.org* petition:

“Short answer: we actually know what we are doing. You don’t.

“Long answer: we use `rand` as *one* of many inputs into the random pool, and we use it as a way to *improve* that random pool. So even if `rand` were to be back-doored by the NSA, our use of `rand` actually improves the quality of the random numbers you get from `/dev/random` [78].

Torvalds’s comments were supported by the comments of Linux engineer Theodore Y. Ts’o, who noted: “I am so glad I resisted pressure from engineers working at Intel to let `/dev/random` in Linux rely blindly on the output of the RDRAND infrastructure. Relying solely on an implementation sealed inside a chip and which is impossible to audit is a BAD idea.” [79].

Nevertheless, the petition’s request *was* ultimately implemented in the Linux kernel. In August 2018, the Linux v1.19-rc1 release candidate kernel included a flag that allowed the kernel to be compiled without support for RDRAND [11, 80]. (We note that even if kernel support for RDRAND is disabled, the instruction can still be accessed from user-level programs, since use of the instruction does not require privilege.) And in 2019, the Linux kernel added “sanity checking” to the output of the hardware random number generator, after the discovery that the hardware random number on some AMD-based systems stopped providing random values after a suspend/resume cycle [43, 61].

In June 2020, Ragab *et al.* published an attack called CrossTalk that allows one user-level process running on an Intel-based computer to eavesdrop on the output of the RDRAND instruction run in another process by observing the Intel CPU’s shared “staging” buffer [65]. This attack is not relevant to the DAS, since we assume that no unauthorized software is running in the Census Bureau’s secure computing environment. However, some systems running ISK have now been patched to address this issue. After the patch is applied, RDRAND reportedly generates random values with only 3% of its unpatched performance.

1.4.8 Use of RDRAND in Statistical Software. The Python programming language adopted MT19937 as its default random number generator in Python 2.3 and still uses it for version 3.8, the 3.9.0b1 release candidate, and the Python 3.10 development tree. The documentation notes “Warning: The pseudo-random generators of this module should not be used for security purposes. For security or cryptographic uses, see the `secrets` module.” [64] The Python `secrets` module includes a function `SystemRandom` which calls the Python function `os.urandom()` as a source of randomness; on Linux systems this reads from `/dev/urandom`, which (as noted above) may incorporate entropy mixed-in from RDRAND. (Note: because it can block and will supply only a small number of random values, `/dev/random` may not be appropriate for use in production statistical software.)

As noted above, the Python `numpy` numeric package now provides for user-supplied random bit generators. Thus, it is now straightforward to combine modern versions of `Numpy` with an RDRAND-based bit-generator such as Sheppard’s *randomgen* [72]. This is a non-standard mode of operation, and requires that *randomgen*

be separately downloaded. We analyzed *randomgen.RDRAND*’s behavior and source code. We verified that `NumPy` run in this configuration is in fact using *randomgen.RDRAND* by running on an 2011 MacBook Air laptop and observing that the `NumPy` random number generator raised an exception, as the MacBook Air’s processor lacked the ISK. However, our analysis of the *randomgen.RDRAND* source code revealed that, as of July 9, 2020, it did *not* check the Carry Flag (CF) as recommended in Intel’s software implementation guide [26]. This implementation error was reported to Sheppard and it was promptly corrected.

We also analyzed the source code for the Python Package Index (pypi) `rand` model version 1.5.0 [76] and found that it did properly implement the CF check.

1.4.9 Randomness Requirements for DP. As noted above, DP’s requirements for the quality of randomness sources are similar to the corresponding requirements in cryptography.

Because pure DP’s definition is stated information theoretically, pure DP is inconsistent with the use of a PRNG, even a CSPRNG. Mironov *et al.* introduced several computationally-aware DP variants, in which attackers are assumed to face clearly defined computational bounds (and so can no longer fully Bayes’ update); these attacker restrictions are similar to those used in the definition of a CSPRNG. “The good news is that a DP mechanism coupled with a [cryptographically secure] PRNG will satisfy the stronger definition of the two. This is the theoretical underpinning for conveniently ignoring the issue of (information-theoretic) DP vs computational DP” [52].

Dodis *et al.* considered the impact of an imperfect randomness source on the privacy guarantees offered by DP by comparing them to the privacy guarantees associated with using such a randomness source to generate cryptographic keys [16]. The authors also discuss the impact of using “infinite-precision” mechanisms that rely on an infinitely long random tape in $\{0, 1\}^*$ and discuss how to approximate it with a tape that offers randomness of finite precision. However, as the randomness sources modeled in Dodis *et al.* are incomparable to CSPRNGs, they are not directly relevant to use of CSPRNGs in the DAS.

The impact of mathematical precision on the privacy guarantees of DP was taken up by Mironov’s 2012 paper [51], which presented an attack that allowed the compromise of underlying confidential data due to “irregularities of floating-point implementations of the privacy-preserving Laplacian mechanism.” The attack is effective because, in some settings, the differences between IEEE floating-point representations and arithmetic can cause the least significant bits of certain queries to leak much more information about individuals in the confidential database than the information-theoretic definition of DP implies—i.e., these differences cause the actually achieved ϵ to be much larger than the ϵ claimed on the basis of interpreting floating-point implementations of probability distributions as equivalent to their real-valued descriptions. This theme is further explored by Gazeau *et al.* [28].

2 RANDOMNESS IN THE 2020 DAS

This section provides an overview of the 2020 DAS, discusses its requirements for randomness, and then discusses how those requirements are achieved.

2.1 Overview of the DAS

The 2020 DAS is a Python-Spark (pyspark) application deployed on an Amazon Elastic Map Reduce (EMR) cloud-computing cluster. EMR runs on top of Amazon Linux, with nodes being configured to run either the pyspark driver program, or as pyspark workers. Currently DAS is running on EMR version 5.25.

The *Amazon EMR: Amazon EMR Release Guide* [7] details all of the version numbers of the open source Apache software used by each EMR release, but it does not mention the version of Amazon Linux on which the release is based. Although it is possible to run EMR with a custom Amazon Machine Image (AMI) [2], the Census Bureau is using a standard AMI, specifically Amazon Linux AMI release 2018.03, which identifies itself as “amzn” and is based on RedHat Fedora.

2.2 Operation of the DAS

The operation of the 2020 DAS has been described by the Census Bureau in other publications, including an overview presented to the Census Bureau’s Scientific Advisory Committee [25], the published design specification for the DAS [83], and a draft academic article describing the so-called TopDown Algorithm (TDA) [4].

Briefly, the algorithm runs twice to produce privacy protected microdata: once, to produce the table containing person-level microdata, and a second time to produce the table containing microdata for housing units and group quarters facilities. For each table, a histogram of counts is computed at each geographic level for the United States that is described by the Census Bureau’s geographical “spine” (currently the US as a whole, the states and D.C., the counties and county equivalents, the census tracts, and the blocks). Each of the cells of each of these histograms, as well as a set of queries chosen to reflect the to-be-published Census tabulations, is then protected using an existing DP mechanism “using the Laplace mechanism [21], Geometric Mechanism [31], or more advanced techniques such as the high dimensional matrix mechanism [47].” [4] These protected values are called the “noisy measurements.”

Next, the algorithm performs two optimizations on the US-level histogram¹⁸ subject to external knowledge constraints (e.g., that state-level population totals must not be noisily perturbed): the first optimization minimizes squared error to the DP estimates of the histogram cell and query values, while requiring that the output be a non-negative, floating-point-valued estimate of the true histogram.¹⁹ The second optimization pass forces all counts to be integers, performing a variant of controlled rounding while maintaining external knowledge constraints.

Following the US-level optimization, the algorithm performs analogous optimizations in which the counts in the US-level histogram are allocated to the state-level histograms, while minimizing differences between the counts assigned to the states and the counts computed from the noisy measurements within each state.

¹⁸Data for Puerto Rico is processed using a separate pass of the same algorithm.

¹⁹That is, this optimization imposes non-negativity and *self-consistency* as requirements: after this optimization is complete, all queries on the microdata can be calculated based off of it, and the estimates attained will be unique. This circumstance does not hold for the initial DP measurements, where we may, for example, have an estimated total population that is not the sum of the estimated white alone and not-white-alone populations.

Once again, this is a two-step optimization process. This process then repeats for every geographic node in each geographic level of the geographical hierarchy until all counts are distributed to individual Census blocks. Finally, each block’s histogram is expanded into the microdata that the histogram specifies. This process is referred to as *post-processing* in some Census Bureau presentations.²⁰ These optimizations provably do not violate the guarantee provided by TDA’s use of differentially private mechanisms, although the use of invariants in this post-processing does, as we previously commented, imply that the achieved privacy guarantee is qualitatively weaker than it would be under pure DP.

The application of DP is a relatively small but essential part of the TDA: the DP subroutines in use by the DAS tend to be much simpler and faster than those used by the DAS for optimization-based post-processing. However, it is essential for the DP mechanism to be correct, as these are the principal source of the DAS’s privacy guarantees. Among other considerations, this means that acquiring suitable high-quality randomness is necessary for the DAS to achieve its purpose.

2.3 Randomness Requirements for the DAS

We estimate the randomness requirements for the DAS by observing that the privacy mechanism starts by computing a histogram of counts for every geographical unit at every geographical level. (In practice, each “run” of the DAS actually consists of two distinct runs: one for the persons table and one for the housing unit/group quarters facility table. For each of these runs the DAS usually manipulates two histograms simultaneously—a principal histogram concerned with variables that heavily interact with one another in published tabulations, and a much smaller histogram featuring a small set of variables that mostly do not interact with those in the main histogram. For simplicity we ignore this much smaller histogram in our descriptions and calculations here. Additionally, our calculations focus on the histograms rather than individual queries, as the histograms dominate the randomness requirements by orders of magnitude, and the queries can change from run to run, depending on configuration file specifications.) We compute the total number of bits for both runs of the top-down algorithm for the United States (but not for Puerto Rico, which is run separately) in Table 1 and find it to be a minimum of 90TB of random data. Table 1 modestly underestimates the randomness requirements because the final DP workload includes not just each histogram cell, but additional queries of summary statistics (some of which have not yet been determined).

2.4 Threat Models

In deciding upon the source of randomness for the DAS, the development team engaged in several threat modeling exercises. We assumed that an attacker would have access to the entire DAS software and hardware stack, including the actual implementation of the TDA used to generate the published statistics, the Python runtime environment, the Linux operating system, the same hardware

²⁰This use of the term post-processing is a reference to the technical use of the same term in the DP literature, where the “post-processing property” refers to a straightforward but important result: that the DP guarantee cannot be undermined by performing further processing on the output of a DP algorithm, so long as the confidential data is not directly accessed in doing so.

$$\begin{aligned}
\text{Total \# Random Bits} &= 64 \times (\text{Total \# protected histogram cells}) \\
&= 64 \times (|H_p| + |H_u|)(\#\text{geounits}) \\
&= 64 \times (|H_p| + |H_u|) \sum G_{\text{geolevels}} \\
&= 64 \times (|H_p| + |H_u|)(G_{\text{nat}} + G_{\text{state equivs}} + G_{\text{county equivs}} + G_{\text{tracts}} + G_{\text{block groups}} + G_{\text{blocks}}) \\
&= 64 \times ((42 \times 2 \times 116 \times 2 \times 63) + (2 \times 9 \times 2 \times 7 \times 4 \times 2 \times 522))(1 + 51 + 3,143 + 73,782 + 217,550 + (\sim 8,000,000)) \\
&\approx 7 \times 10^{14} \text{bits} \quad (\approx 90\text{TB})
\end{aligned}$$

Where $|H_p|$ = Size of the person-level histogram
 $|H_u|$ = Size of the unit-level histogram
 G_{level} = The number of geounits at geolevel *level*

Figure 1: Randomness requirements for the US run of the 2020 DAS, including 50 states and the District of Columbia. Geography counts for county equivalent and tracts taken from 2010 Census. It is estimated that there will be 8 million habitual blocks for the 2020 census.

on which the DAS had been run, and detailed information of the network configuration. We assumed that the attacker would have all data publications produced from the confidential data, and that the attacker could combine these publications with significant external knowledge. For example, we assumed that an attacker was likely to know the rough age, race and sex distributions of every community in the U.S., since the Census Bureau already publishes this information as part of the American Community Survey. We also assumed that the attacker has unbounded mathematical skills and computational capabilities, although we did assume that the attacker could not crack AES-256.

Except as otherwise noted, however, we exclude most risks pertaining to the correctness of software and hardware implementations. That is, we assume that attackers do not have access to the systems on which the TDA is running, as this would give the attacker access to the underlying confidential data, rendering moot DP’s privacy protections. We likewise assume that Linux kernel on which the TDA executes matches the source code for the Linux kernel that we reviewed.

2.5 MT19937 and DP

Because MT19937 is the default PRNG in Python and was the default Numpy RNG prior to Numpy 1.17, MT19937 has been widely used in DP demonstrations. Indeed, MT19937 was used in the initial prototype implementation of the DAS.

It is inadvisable to use MT19937 in a production system that is intended to protect confidential data: though MT19937 has a large period of $2^{19937} - 1$, MT19937 has substantial security vulnerabilities.

MT19937 maintains its internal state as a vector of 624 32-bit unsigned integers, and its output can, as a byproduct of this, be predicted after observing just 624 output iterations. Indeed, there is publicly available source code that implements this attack (e.g. [57]), and there are several user-friendly blog posts outlining how to perform a variant of this attack [29, 63, 71].

In an application like the DAS, which requires protecting large, sparse histograms, the vast majority of the protected histogram counts are the result of taking true counts of zero and adding noise derived from sequential draws of the RNG. Thus, if an attacker

knows that certain populations are not present in a geographic area, the attacker can immediately infer the noise iterates from viewing the DP query estimates. This situation is problematic if MT19937 is used, as it reduces the security of the implementation to the difficulty of inverting Laplace or Geometric distribution random draws: if these can be inverted, the attack on MT19937 can be carried out, and the DP guarantees unravel entirely for geographic units where the necessary auxiliary knowledge is available.

For the 2020 Census, the Census Bureau is in fact considering the release of the so-called “noisy measurements,” or the raw values, of each cell in each histogram after the noise has been added, as these enjoy in principle the same privacy guarantees as the final microdata. The histograms being considered for the DAS are currently 217,550 and approximately 8 million cells. Since the vast majority of these cells are likely to be zero, if MT19937 is used, an attack of this kind may be quite practical. Here is a sketch:

- This attack relies on finding a run of 313 cells known to correspond to true counts of zero, and being able to invert the noise algorithm such that the output of the MT19937 algorithm can be determined. With these output values, the internal state of the MT19937 engine is then determined.
- The internal state is validated by running the MT19937 algorithm two steps forward to determine the next 64 output bits. The 64 output bits predict the contents of the 313th cell. If the 313th cell matches the prediction, then the internal state of the MT19937 algorithm is validated.
- Once the state is validated, the amount of noise that was added can be inferred for every other cell in the histogram, and the privacy protection mechanism is undone.

This attack requires being able to transform a 64-bit noise value into two successive 32-bit draws, which may require some additional thought and computation to achieve, so this attack is not immediately practical. Nevertheless, it is not desirable for the security of a DP implementation to rely on the difficulty of inverting noise-distribution sampling functions, which are not designed with security in mind. Thus, it is clear that MT19937 is inappropriate for production DP applications. We believe that only CSPRNGs should be used for production privacy applications.

2.6 Developing the DAS

During initial efforts to develop the DAS, the development team discussed whether or not it would be desirable to have a “repeatable” source of random numbers to support regression tests [27]. Initial efforts to create a repeatable sequence failed: the research team discovered that Apache Spark’s scheduler was non-deterministic, and that the same workload might be scheduled simultaneously on multiple nodes if there were available resources. Although there was discussion that the random seed could be made dependent upon the geographical unit, ultimately it was decided that this was unnecessary; instead, approximate repeatability could be achieved by caching the DP measurements and building functionality to re-load from them, as this functionality would be required for other purposes regardless.²¹

The DAS development team turned its attention to the Python random number generator in the early part of 2018. Learning of the problems with MT19937, but unwilling to create its own implementation of the Laplace or exponential distribution random samplers, the initial work-around employed was to reseed the MT19937 generator with a read of `/dev/urandom` on every scalar draw. This proved to be unacceptably slow as the scale of problems the DAS needed to address increased beyond what was required for early testing.

The DAS development team was familiar with the Intel RDRAND instruction and learned that Intel had developed its own Python distribution which made use of hardware acceleration on the Intel platform. The team assumed that the Intel Python distribution adopted RDRAND for use in the Numpy RNG, but it had not.

Intel also created `mkl_random`, “a NumPy-based Python interface to Intel (R) MKL [Math Kernel Library] Random Number Generation functionality” [42, 59]. This software was installed on the DAS development clusters and used for the 2018 Decennial Census End-to-end test. However, additional testing revealed several implementation flaws,²² causing the 2020 DAS development team to stop using the software.

2.7 Randomness and the 2020 DAS

Although Müller would seem to be the final word on the subject, the DAS development team had concerns about using `/dev/urandom` as the sole source of randomness for the privacy mechanism of the 2020 Census. Our first concern was that it was surprisingly difficult to verify that the version of `/dev/urandom` that Müller had reviewed was the same version that the DAS team was running, as there is no obvious mechanism for verify the integrity of modules that make up a kernel running in the AWS environment. Our second concern validating whether or not the entropy seed of the AWS device might be inadvertently replicated as part of the cluster boot process. Our third concern was performance: there is only a single `/dev/urandom` device and corresponding entropy pool for each AWS kernel, creating a single-threaded bottleneck in our otherwise parallelized implementation. With 96 cores per cluster, such a bottleneck might adversely impact performance.

²¹Caching the DP measurements in fact ensures exact repeatability, except for a small degree of non-auditable non-determinism that can occur in the Gurobi Optimizer (e.g., when using it in a mode where multiple optimization algorithms are deployed in parallel, terminating when the first such algorithm terminates), on which the DAS relies.

²²See https://github.com/IntelPython/mkl_random/issues

The alternative to using `/dev/urandom` is to implement a user-level CSPRNG. Such a CSPRNG could be seeded from `/dev/urandom` or from a suitable source of hardware entropy.

Two sources of hardware entropy in the AWS environment are the ISK (on appropriately equipped systems) and a network-accessible hardware security module for generating cryptographically strong random numbers that Amazon provides [6]. Amazon’s service is part of the AWS Key Management Service (KMS), which is an integrated system for managing and using both symmetric and asymmetric keys and encryption algorithms. We tested the KMS `GenerateRandom` and determined that it generates 1024 bytes at a time, and has a round-trip-time of approximately 0.5 seconds per invocation, for a maximum single-threaded, non-pipelined performance of 2KiB/sec. Although we could find no documented concerns regarding the operation of the KMS and have no reason to not to trust the quality of its entropy, we could also find no reason to trust it other than an appeal to Amazon’s authority.

Intel’s documentation states that the DRNG can produce random data at the rate of 800 MB/sec per CPU chip. Current CPU chips have faster clock rates and more cores, but it does not appear that the DRNG has been improved significantly or that it is no longer shared between cores. Assuming that today’s chips have the same DRNG unit and the CrossTalk patches have not been applied, a single Intel chip could produce the required randomness to protect the 2020 Census using the Census Bureau’s DP mechanism in $90 \times 10^{12} \div 800 \times 10^6 \approx 112,000$ seconds, or 31 hours.

Of course, the Census Bureau is *not* running the TDA on a single CPU. Currently, the TDA runs on AWS r5.24xlarge virtual machines. These systems report 96 cores each arranged in 8 12-core Intel Xeon chips. For an EMR cluster with 20 workers, there will be 160 DRNGs, allowing the required randomness to be computed in about 700 seconds, or about 12 minutes. (Sheppard observed that faster times might be achieved with the same level of security by using the RDSEED instruction to seed a software random number generator based on AES run in counter mode—essentially a software version of RDRAND without the periodic reseeding [73].)

Given the extensive analysis that the ISK has received, and given the simplicity of accessing it with a single user-level machine instruction, the DAS team originally planned to simply use RDRAND as the source of *all* randomness for the TDA. This approach was discouraged by outside reviewers, who stressed that a silicon-only implementation could not be audited. Some reviewers suggested that the Census Bureau rely on the Linux `/dev/urandom` device exclusively as the source of randomness, arguing that it already mixes in entropy from RDRAND, as well as from other sources such as hardware interrupt timing.

The Census Bureau performed a speed test of `/dev/urandom` on the AWS r5.24xlarge server in AWS GovCloud US-West region and found that a single-threaded `dd` process could retrieve pseudo-random bytes at the rate of roughly 200 MB/sec, or one fourth the rate of RDRAND. However, the `/dev/urandom` device is effectively single-threaded, as there is only one entropy pool and it is protected with a lock. Four concurrent `dd` processes on the same server retrieved bytes from `/dev/urandom` with data transfer rates between 52 MB/sec and 54 MB/sec each, the slight performance boost likely coming from parallelization of the user-level `dd` code.

On a 96-core machine with 8 CPUs, the 8 ISK devices are able to provide a maximum throughput of 6400 MB/sec—32 times the bandwidth of random data than can be provided by the `/dev/urandom` device. On an EMR cluster with 20 worker nodes, this raises the amount of time required to produce the randomness to 8 hours, which was deemed to be unacceptable.

The Census Bureau’s current solution is to operate a user-level CSPRNG based on AES-CTR-DRBG that will be seeded from user-level calls to RDSEED mixed with output from `/dev/urandom`. The output of this CSPRNG will be used as a bit generator for NumPy. Although this effectively runs 96 CSPRNGs per `r5.24xlarge` system, resulting in a significant performance boost, higher rates of reseeding will lower performance to the maximum rate at which RDSEED can provide such seeds. Thus, the performance of this arrangement is tunable, and those tunings have not yet been decided.

The DAS development team has determined that AES side channel and timing attacks are not of concern in this circumstance, since all software running on the virtual machine is trusted, the use of `r5.24xlarge` VMs assures that there are no other tenants on the physical hardware, and precise timing information will not be available to potential attackers. Nevertheless, the DAS will use an AES implementation that employs the native AES acceleration instructions provided by Intel’s microprocessors to prevent side channel attacks.

3 CONCLUSION

The need to generate a large number of high-quality random numbers is a largely unrecognized requirement of a production differential privacy system. Many DP tutorials and texts assume the availability of high-quality floating point random numbers taken from the Laplace, geometric, or exponential distribution. In practice, these examples use MT19937 or PCG64. These algorithms are not CSPRNGs and should not be used to protect confidential information. Because this important detail may not be obvious to developers looking for a DP implementation, we recommend that DP tutorials and texts discuss this issue and use CSPRNGs as their randomness sources.

The prototype implementation of the 2020 DAS used MT19937 as seeded by `/dev/urandom` on a cluster of 96-core Linux servers. After the DAS development team learned that MT19937 is not secure, the team changed the DP primitives to use Intel’s RDRAND instruction as the source of randomness, as accessed through the Python `mkl_random` library. To avoid relying on the ISK as a single source of randomness, and after discovering that `mkl_random` is based on the closed-source Intel Math Kernel Library, the DAS team pivoted to using a user-level random bit generator based on AES-256 and seeded by the Linux `/dev/urandom` mixed with bits from the RDSEED instruction. Throughout the process, the DAS team found it necessary to review multiple randomness implementations down to the level of assembler code, and found several software quality issues and implementation errors, some of which are discussed in this paper.

The Census Bureau is now four years into the process of modernizing its disclosure avoidance systems to incorporate formal privacy protection techniques. This process has proven to be challenging across disciplines. Beyond the 2020 Census, the Census Bureau intends to use DP or related formal privacy systems to protect all of its future statistical publications.

ACKNOWLEDGMENTS

At the US Census Bureau, we wish to thank John Abowd for supporting our work on this project and Donald E Badrak II for his technical input. We had invaluable assistance and technical input from Galois Inc.’s audit team (Jose Calderon, Josh Heumann, Scott Moore and Marc Rosen) and Kevin Sheppard (Oxford). We also had useful input from Paul Bartholomew, Andrew Hong and Drew Lipman at MITRE, Jim Hughes (UC Santa Cruz), John M. Kelsey (NIST), and Sebastiano Vigna (Università degli Studi di Milano). Additional helpful comments were provided by Abraham D. Flaxman (University of Washington), Rich Salz (Akami), and Ted Ts’o (Google). Our summer intern Joshua Schmidt provided valuable editorial assistance. We also wish to thank the anonymous reviewers, who provided useful guidance.

The views in this paper are those of the authors, and do not represent those of the US Census Bureau.

REFERENCES

- [1] 1955. *A Million Random Digits with 100,000 Normal Deviates: Foreword to the Online Edition*. Technical Report MR-1418. RAND Corporation. https://www.rand.org/pubs/monograph_reports/MR1418/index2.html Last accessed May 30, 2020.
- [2] 2018. Restricted-Use Microdata. https://www.census.gov/research/data/restricted_use_microdata.html#CRE1 Last Accessed July 14, 2018.
- [3] Alex Abella. 2009. *Soldiers of Reason: The RAND Corporation and the Rise of the American Empire*. Mariner Books.
- [4] John Abowd, Robert Ashmead, Simson Garfinkel, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, Brett Moran, and Pavel Zhuravlev. 2019. Census TopDown Algorithm: Differentially Private Data, Incremental Schemas, and Consistency with Public Knowledge. (Oct. 2019). https://github.com/usensusbureau/census2020-das-2010ddp/blob/master/doc/20191020_1843_Consistency_for_Large_Scale_Differentially_Private_Histograms.pdf
- [5] John M. Abowd. 2018. The U.S. Census Bureau Adopts Differential Privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '18)*. Association for Computing Machinery, New York, NY, USA, 2867. <https://doi.org/10.1145/3219819.3226070>
- [6] Amazon Web Services. 2020. GenerateRandom. In *AWS Key Management Service*. https://docs.aws.amazon.com/kms/latest/APIReference/API_GenerateRandom.html Last accessed May 18, 2020.
- [7] Amazon Web Services, Inc. and/or its affiliates. 2020. *Amazon EMR: Amazon EMR Release Guide*. <https://docs.aws.amazon.com/emr/latest/ReleaseGuide/emr-whatsnew-history.html>
- [8] Elaine Barker and John Kelsey. 2007. *Recommendation for Random Number Generation Using Deterministic Random Bit Generators (Revised)*. National Institute of Standards and Technology. <https://csrc.nist.gov/publications/detail/sp/800-90/revise/archiv/2007-03-14> Last accessed May 30, 2020.
- [9] Dan Boneh, Amit Sahai, and Brent Waters. 2011. Functional Encryption: Definitions and Challenges. In *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011 (Lecture Notes in Computer Science)*, Vol. 6597. Springer, 253. https://doi.org/10.1007/978-3-642-19571-6_16
- [10] Catalin Cimpanu. 2019. Cloudflare launches decentralized service for generating random numbers. *ZDNet* (June 17 2019). <https://www.zdnet.com/article/cloudflare-launches-decentralized-service-for-generating-random-numbers/>
- [11] Thomas Claburn. 2018. Linux 4.19 lets you declare your trust in AMD, IBM and Intel. (Aug. 28 2018). https://www.theregister.com/2018/08/28/linux_419_trust/
- [12] Shaanan N. Cohny, Matthew D. Green, and Nadia Heninger. 2018. Practical State Recovery Attacks against Legacy RNG Implementations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*. Association for Computing Machinery, New York, NY, USA, 265–280. <https://doi.org/10.1145/3243734.3243756>
- [13] Mary Lisbeth D'Amico. 1999. Advisory group asks EU to consider Pentium III ban. *CNN.com* (Nov. 29 1999). <http://www.cnn.com/TECH/computing/9911/29/eu.p3.ban.idg/index.html>
- [14] S. Das, J. Werner, M. Antonakakis, M. Polychronakis, and F. Monrose. 2019. SoK: The Challenges, Pitfalls, and Perils of Using Hardware Performance Counters for Security. In *2019 IEEE Symposium on Security and Privacy (SP)*. 20–38.
- [15] Alex Davidson. 2019. *The Cloudflare Blog* (June 17 2019). <https://blog.cloudflare.com/inside-the-entropy/>
- [16] Yevgeniy Dodis, Adriana López-Alt, Ilya Mironov, and Salil Vadhan. 2012. Differential Privacy with Imperfect Randomness. In *Proceedings of the 32nd Annual Cryptology Conference on Advances in Cryptology – CRYPTO 2012 - Volume 7417*. Springer-Verlag, Berlin, Heidelberg, 497–516. https://doi.org/10.1007/978-3-642-32009-5_29
- [17] Chris Doty-Humphrey. 2016. Small Fast Chaotic (SFC) 64. <http://prcrand.sourceforge.net/> Last accessed July 3, 2020.
- [18] F. Betül Durak, Thomas M. DuBuisson, and David Cash. 2016. What Else is Revealed by Order-Revealing Encryption?. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*. Association for Computing Machinery, New York, NY, USA, 1155–1166. <https://doi.org/10.1145/2976749.2978379>
- [19] Cynthia Dwork. 2019. Differential Privacy and the US Census. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS '19)*. Association for Computing Machinery, New York, NY, USA, 1. <https://doi.org/10.1145/3294052.3322188>
- [20] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the Third Conference on Theory of Cryptography (TCC'06)*. Springer-Verlag, Berlin, Heidelberg, 265–284. https://doi.org/10.1007/11681878_14
- [21] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the Third Conference on Theory of Cryptography (TCC'06)*. Springer-Verlag, Berlin, Heidelberg, 265–284. https://doi.org/10.1007/11681878_14
- [22] D. Eastlake 3rd, S. Crocker, and J. Schiller. 1994. Randomness Recommendations for Security. RFC 1750 (Informational). <http://www.ietf.org/rfc/rfc1750.txt> Obsoleted by RFC 4086.
- [23] D. Eastlake 3rd, J. Schiller, and S. Crocker. 2005. Randomness Requirements for Security. RFC 4086 (Best Current Practice). <http://www.ietf.org/rfc/rfc4086.txt>
- [24] Feynman Center. 2016. Five Los Alamos innovations win R&D 100 Awards. <https://www.lanl.gov/projects/feynman-center/about/news/2016-11-15-r-and-d-awards.php>
- [25] Simson L. Garfinkel. 2018. Modernizing Disclosure Avoidance: Report on the 2020 Disclosure Avoidance System as Implemented for the 2018 End-to-End Test. <https://www.census.gov/about/cac/sac/meetings/2017-09-meeting.html>
- [26] Simson L. Garfinkel. 2020. RDRAND implementation does not appear to check CF as required in Intel software implementation guide. *GitHub.com* (July 9 2020). <https://github.com/bashtage/randomgen/issues/246>
- [27] Simson L. Garfinkel, John M. Abowd, and Sarah Powazek. 2018. Issues Encountered Deploying Differential Privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society (WPES'18)*. Association for Computing Machinery, New York, NY, USA, 133–137. <https://doi.org/10.1145/3267323.3268949>
- [28] Ivan Gazeau, Dale Miller, and Catuscia Palamidessi. 2016. Preserving differential privacy under finite-precision semantics. *Theoretical Computer Science* 655 (2016), 92 – 108. <https://doi.org/10.1016/j.tcs.2016.01.015> Quantitative Aspects of Programming Languages and Systems (2013-14).
- [29] Gehn. 2018. Breaking MT19937 Crypto. <https://eldipa.github.io/book-of-gehn/articles/2018/12/23/Mersenne-Twister-PRNG.html>
- [30] R. Gennaro. 2006. Randomness in Cryptography. *IEEE Security & Privacy* 4, 02 (mar 2006), 64–67. <https://doi.org/10.1109/MSP.2006.49>
- [31] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. 2009. Universally Utility-Maximizing Privacy Mechanisms. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing (STOC '09)*. Association for Computing Machinery, New York, NY, USA, 351–360. <https://doi.org/10.1145/1536414.1536464>
- [32] David Alan Grier. 2007. *When Computers Were Human*. Princeton University Press, USA.
- [33] Z. Gutterman, B. Pinkas, and T. Reinman. 2006. Analysis of the Linux random number generator. In *2006 IEEE Symposium on Security and Privacy (S'06)*. 15 pp–385.
- [34] Peter Guttman. 2004. *Cryptographic Security Architecture: Design and Verification*. Springer.
- [35] Mike Hamburg, Paul Kocher, and Mark E. Marson. 2012. *Analysis of Intel's Ivy Bridge Digital Random Number Generator*. Cryptography Research, Inc. https://web.archive.org/web/20141230024150/http://www.cryptography.com/public/pdf/Intel_TRNG_Report_20120312.pdf
- [36] Michael B. Hawes. 2020. Implementing Differential Privacy: Seven Lessons From the 2020 United States Census. *Harvard Data Science Review* (30 4 2020). <https://doi.org/10.1162/99608f92.3536cf99> <https://hdsr.mitpress.mit.edu/pub/dgg03v06>.
- [37] Sandra Henry-Stocker. 2017. True random numbers are here — what that means for data centers. *Network World* (July 27 2017). <https://www.networkworld.com/article/3211529/true-random-numbers-are-here-what-this-can-mean-for-all-of-us.html>
- [38] V. Joseph Hotz, Joseph Salvo, Catherine Fitch, Daniel Goroff, Eddie Hunsinger, Linda Jacobsen, Michael McDonald, and Matthew Snipp. 2019. Workshop on 2020 Census Data Products: Data Needs and Privacy Considerations. https://sites.nationalacademies.org/dbasse/cnstat/dbasse_196518
- [39] ID Quantique. 2020. ID Quantique and SK Telecom announce the world's first 5G smartphone equipped with a Quantum Random Number Generator (QRNG) chipset. <https://www.idquantique.com/random-number-generation/overview/>
- [40] K. Inayah, B. E. Sukmono, R. Purwoko, and S. Indarjani. 2013. Insertion attack effects on standard PRNGs ANSI X9.17 and ANSI X9.31 based on statistical distance tests and entropy difference tests. In *2013 International Conference on Computer, Control, Informatics and Its Applications (IC3INA)*. 219–224.
- [41] Information Technology Laboratory. 2013. Supplemental ITL Bulletin for September 2013. <https://csrc.nist.gov/csrc/media/publications/shared/documents/itl-bulletin/itlbul2013-09-supplemental.pdf>
- [42] Intel. 2020. Intel Math Kernel Library. <https://software.intel.com/content/www/us/en/develop/tools/math-kernel-library.html> Last accessed July 15, 2020.
- [43] Michael Larabel. 2019. Following Buggy AMD RdRand, The Linux Kernel Will Begin Sanity Checking Randomness At Boot Time. *phoronix* (Oct. 2 2019). https://www.phoronix.com/scan.php?page=news_item&px=Linux-RdRand-Sanity-Check Last accessed May 30, 2020.
- [44] Joshua Liebow-Feeser. 2017. Randomness 101: LavaRand in Production. *The Cloudflare Blog* (Nov. 6 2017). <https://blog.cloudflare.com/randomness-101-lavarand-in-production/>
- [45] Pierre L'Ecuyer and Richard Simard. 2007. TestU01: A C Library for Empirical Testing of Random Number Generators. *ACM Trans. Math. Softw.* 33, 4, Article 22 (Aug. 2007), 40 pages. <https://doi.org/10.1145/1268776.1268777>
- [46] Laura McKenna. 2018. *Disclosure Avoidance Techniques Used for the 1970 through 2010 Decennial Censuses of Population and Housing*. Technical Report CDAR2018-01. US Census Bureau.
- [47] Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. 2018. Optimizing Error of High-Dimensional Statistical Queries under Differential Privacy. *Proc. VLDB Endow.* 11, 10 (June 2018), 1206–1219. <https://doi.org/10.14778/3231751.3231769>

- [48] John P. Mechalas. 2012. The Difference Between RDRAND and RDSEED. (Nov. 17 2012). last accessed May 29, 2020.
- [49] John P. Mechalas. 2018. Intel (R) Digital Random Number Generator (DRNG) Software Implementation Guide. (Oct. 17 2018). <https://software.intel.com/content/www/us/en/develop/articles/intel-digital-random-number-generator-drng-software-implementation-guide.html> Revision 2.1; originally published on May 14, 2014. Last accessed May 29, 2020.
- [50] N. Metropolis. 1987. The Beginning of the Monte Carlo Method. *Los Alamos Science Special Issue* (1987), 125–130. <https://library.lanl.gov/cgi-bin/getfile?00326866.pdf>
- [51] Ilya Mironov. 2012. On Significance of the Least Significant Bits For Differential Privacy. ACM. <https://www.microsoft.com/en-us/research/publication/on-significance-of-the-least-significant-bits-for-differential-privacy/>
- [52] Ilya Mironov, Omkant Pandey, Omer Reingold, and Salil P. Vadhan. 2009. Computational Differential Privacy. In *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings (Lecture Notes in Computer Science)*, Shai Halevi (Ed.), Vol. 5677. Springer, 126–142. https://doi.org/10.1007/978-3-642-03356-8_8
- [53] National Conference of State Legislatures. 2020. Differential Privacy for Census Data Explained. <https://www.ncsl.org/research/redirecting/differential-privacy-for-census-data-explained.aspx> Last accessed May 13, 2020.
- [54] Muhammad Naveed, Seny Kamara, and Charles V. Wright. 2015. Inference Attacks on Property-Preserving Encrypted Databases. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*. Association for Computing Machinery, New York, NY, USA, 644–655. <https://doi.org/10.1145/2810103.2813651>
- [55] Y. Nir and A. Langley. 2018. ChaCha20 and Poly1305 for IETF Protocols. RFC 8439 (Informational). <http://www.ietf.org/rfc/rfc8439.txt>
- [56] Landon Curt Noll, Robert Mende, and Sanjeev Sisodiya. 1998. Method for seeding a pseudo-random number generator with a cryptographic hash of a digitization of a chaotic system. Filed January 29, 1996.
- [57] Kimiyuki Onaka. 2018. Mersenne Twister Predictor. <https://github.com/kmyk/mersenne-twister-predictor> Last Accessed May 29, 2020.
- [58] Melissa E. O'Neill. 2014. *PCG: A Family of Simple Fast Space-Efficient Statistically Good Algorithms for Random Number Generation*. Technical Report HMC-CS-2014-0905. Harvey Mudd College Computer Science Department, Claremont, CA. <https://www.cs.hmc.edu/tr/hmc-cs-2014-0905.pdf>
- [59] Oleksandr Pavlyk, Mike Sarahan, Andres Guzman-Ballen, and Todd Tomashek. 2019. IntelPython/mkl_random. https://github.com/IntelPython/mkl_random
- [60] Nicole Perlroth. 2013. Government Announces Steps to Restore Confidence on Encryption Standards. (Sept. 10 2013). <https://bits.blogs.nytimes.com/2013/09/10/government-announces-steps-to-restore-confidence-on-encryption-standards/>
- [61] Borislav Petkov. 2019. x86/rdrand: Sanity-check RDRAND output. <https://git.kernel.org/pub/scm/linux/kernel/git/tip/tip.git/commit/?h=x86/cpu&id=7879fc4bdc7506d37bd67b6fc29442c53c066fda> Last accessed May 30, 2020.
- [62] Gilles Pokam, Cristiano Pereira, Klaus Danne, Lynda Yang, Sam King, and Jesep Torrellas. 2009. Hardware And Software Approaches For Deterministic Multi-Processor Replay Of Concurrent Programs. *Intel Technology Journal* 13 (2009). Issue 4.
- [63] Thomas Ptacek, Sean Devlin, Alex Balducci, and Marcin Wielgoszewski. 2020. Challenge 23: Clone an MT19937 RNG from its output. <https://cedricvanrompay.gitlab.io/cryptopals/challenges/23.html> Last accessed July 8, 2020.
- [64] Python Software Foundation. 2020. random — Generate pseudo-random numbers. <https://docs.python.org/3.8/library/random.html> Last accessed May 30, 2020.
- [65] Hany Ragab, Alyssa Milburn, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. 2021. CrossTalk: Speculative Data Leaks Across Cores Are Real. In *IEEE Security and Privacy*. Paper=https://download.vusec.net/papers/crosstalk_sp21.pdf Web=<https://www.vusec.net/projects/crosstalkCode=https://github.com/vusec/ridl>
- [66] RAND. 1955. *A Million Random Digits with 100,00 Normal Deviates*. Rand Corporation.
- [67] Steven Ruggles, Catherine Fitch, Diana Magnuson, and Jonathan Schroeder. 2019. Differential Privacy and Census Data: Implications for Social and Economic Research. In *AES Papers and Proceedings*, Vol. 109. 403–408. <https://doi.org/10.1257/pandp.20191107>
- [68] Sylvain Ruhault. 2017. SoK: Security Models for Pseudo-Random Number Generators. *LACR Transactions on Symmetric Cryptology* 2017, 1 (Mar. 2017), 506–544. <https://doi.org/10.13154/tosc.v2017.i1.506-544>
- [69] J. K. Salmon, M. A. Moraes, R. O. Dror, and D. E. Shaw. 2011. Parallel random numbers: As easy as 1, 2, 3. In *SC '11: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–12.
- [70] Bruce Schneier. 2007. Did NSA Put a Secret Backdoor in New Encryption Standard? *Wired* (Nov. 15 2007). <https://www.wired.com/2007/11/security-matters-1115/> Last accessed May 30, 2020.
- [71] Mike Shema. 2012. Chapter 7 - Leveraging Platform Weaknesses. In *Hacking Web Apps*. Elsevier, Chapter 7, 209–238. <https://doi.org/10.1016/B978-1-59-749951-4.00007-2>
- [72] Kevin Sheppard. 2020. randomgen. <https://pypi.org/project/randomgen/> Last accessed May 31, 2020.
- [73] Kevin Sheppard, Nathaniel J. Smith, Ralf Gommers, Robert Kern, Victor Rodriguez, and Simson Garfinkel. 2020. ENH: RDRAND in numpy to improve performance. <https://github.com/numpy/numpy/issues/9365>
- [74] Thomas Shrimpton and R. Seth Terashima. 2015. A Provable-Security Analysis of Intel's Secure Key RNG. In *EUROCRYPT (1)*, Springer, 77–100. https://doi.org/10.1007/978-3-662-46800-5_4
- [75] Stephan Müller. 2020. *Documentation and Analysis of the Linux Random Number Generator*. Technical Report 3.6. https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/Studies/LinuxRNG/LinuxRNG_EN.pdf?__blob=publicationFile&v=20
- [76] Christopher Stillson. [n. d.]. rdrand 1.5.0. *Python Package Index* (Dec. 9 [n. d.]). <https://pypi.org/project/rdrand>
- [77] L. H. C. Tippett. 1927. *Random Sampling Numbers: Tracts for Computers*, No. XV. University Press, Cambridge.
- [78] Linux Torvalds. 2013. Linux Torvalds's response. <https://www.change.org/p/linux-torvalds-remove-rdrand-from-dev-random-4/responses/9066>
- [79] Theodore Y. Ts'o. 2013. on: N.S.A. Foils Much Internet Encryption. <https://news.ycombinator.com/item?id=6336505> Last accessed May 30, 2020.
- [80] Theodore Y. Ts'o. 2018. random: add a config option to trust the CPU's hwrng 10531149 diffmbx. <https://patchwork.kernel.org/patch/10531149/> Last accessed May 30, 2020.
- [81] US Census Bureau. 2019. 2020 Census 2010 Demonstration Data Products Disclosure Avoidance System. <https://github.com/uscensusbureau/census2020-das-2010ddp> Last accessed May 11, 2020.
- [82] US Census Bureau. 2019. 2020 Census Detailed Operational Plan for: 19. Response Processing Operation (RPO). US Department of Commerce. https://www2.census.gov/programs-surveys/decennial/2020/program-management/planning-docs/RPO_detailed_operational_plan-v2.pdf
- [83] US Census Bureau. 2019. *Disclosure Avoidance System for the 2010 Demonstration Data Products: Design Specification, Version 1.4.1*. Department of Commerce. <https://github.com/uscensusbureau/census2020-das-2010ddp/blob/master/doc/2010-Demonstration-Data-Products-Disclosure-Avoidance-System-Design-Specification%20FINAL.pdf>
- [84] US Census Bureau. 2019. Disclosure Avoidance System for the 2020 Census, End-to-End release. <https://github.com/uscensusbureau/census2020-das-e2e> Last accessed May 11, 2020.
- [85] US Census Bureau. 2020. Disclosure Avoidance and the 2020 Census. https://www.census.gov/about/policies/privacy/statistical_safeguards/disclosure-avoidance-2020-census.html Last accessed May 11, 2020.
- [86] Sebastiano Vigna. 2020. The wrap-up on PCG generators. <http://pcg.di.unimi.it/pcg.php> Last accessed June 13, 2020.
- [87] Jon von Neumann. 1951. Various Techniques Used in Connection With Random Digits. *Journal of Research, Applied Math Series* 3 (1951), 36–38. https://mcnp.lanl.gov/pdf_files/nbs_vonneumann.pdf Summary written by George E. Forsythe.
- [88] V. M. Weaver, D. Terpstra, and S. Moore. 2013. Non-determinism and overcount on modern hardware performance counter implementations. In *2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 215–224.
- [89] Michael Wertheimer. 2015. Encryption and the NSA Role in International Standards. (2015). <http://www.ams.org/notices/201502/rnoti-p165.pdf> Note: At the time of publication, Michael Wertheimer was the Director of Research at the US National Security Agency. Last Accessed May 30, 2020.