

Cryptographic Randomness from Air Turbulence in Disk Drives

Don Davis,¹ Ross Ihaka,² and Philip Fenstermacher³ *

¹ Openvision Technologies, 1 Main St. Cambridge, MA 02142

² University of Auckland, Mathematics Dept, Auckland, NZ

³ 18A Forest St. Cambridge, MA 02138

Abstract. A computer disk drive’s motor speed varies slightly but irregularly, principally because of air turbulence inside the disk’s enclosure. The unpredictability of turbulence is well-understood mathematically; it reduces not to computational complexity, but to information losses. By timing disk accesses, a program can efficiently extract at least 100 independent, unbiased bits per minute, at no hardware cost. This paper has three parts: a mathematical argument tracing our RNG’s randomness to a formal definition of turbulence’s unpredictability, a novel use of the FFT as an unbiasing algorithm, and a “sanity check” data analysis.

1 Introduction

Secure PRNG design commonly rests on computational complexity [2, 5, 6, 13, 24], but none of the underlying problems has been proven to be hard. Specialized hardware can provide naturally random physical noise, but has disadvantages: dedicated devices tend to be expensive; natural noise tends to be biased and correlated; hardware failure can silently suppress randomness; and physical randomness is only an article of faith. Our random number generator, which is based on disk-speed variations,⁴ addresses each of these problems. Timing data are very low-cost, easily whitened, reliable, and mathematically noisy.

I/O randomness is well-known in cryptography [17], and a spinning-disk RNG was used even 50 years ago [11]. Still, our approach is subtly novel, because a disk drive combines three important features most economically. First, the OS detects and reports disk faults, so that silent randomness failures are unlikely. Second, unlike most other I/O devices, the disk can be secured from outside influence and measurement. Last, nonlinear dynamics gives us an *a priori* mathematical argument for our generator’s randomness. This has not been possible for other noise sources, which rely on *a posteriori* statistical measurements.

This paper has three parts. First, we trace the disk’s speed-variations to air turbulence, and we show why these variations can show only short-term correlations. Second, we show that the FFT removes bias and correlations from the disk’s timing-data. Third, we describe our “sanity check” analysis of some timing-data and the resulting random numbers.

* Affiliations during this work: MIT Project Athena, MIT Stat. Ctr., MIT LCS, resp.

⁴ Disk drives use brushless DC motors [10, 18], so these speed variations are independent of the AC line frequency.

2 Turbulence in Disk Drives

In this section, we review studies demonstrating turbulent air flow in disk drives. Oversized mockups have clarified the various turbulent flow regimes inside a disk drive, and have shown how rotational speed, disk spacing, and cooling flow affect the flow [1]. The apparatus was a stack of 1-meter glass disks, spun in a water tank at low speeds (5 – 60 rpm). A close-fitting cylindrical shroud enclosed the disks, and 50 cm.-diameter hubs separated them, to closely model typical modern disk drives of various sizes. Dye and bubbles made the flows visible, and a rotating video camera recorded the results. Our disk’s speed and configuration were similar to those studied.

Turbulence arose at the read/write heads and their support arms, in Coriolis circulation between the disk surfaces, and in Taylor-Couette flow at the disks’ rims. Crucially, the T-C flow pumped turbulence into the Coriolis flow. Numerical simulation of disk flow showed similar turbulence patterns [21], and yielded an estimate of 3% for the consequent fluctuations in the windage torque. This is clearly large enough to influence the disk’s speed.

Spectral measurements of the fluid velocity showed both sharp peaks and broadband features [1], reflecting *weak turbulence*: very noisy motion with a periodic component [3]. The spectra were taken at various rotational speeds, but the peaks always contained only a small proportion of the spectral power, rising only a factor of 2–3 above the white-noise background. This broadband spectral component was maximized at Reynolds numbers near those commonly found in disk drives.

The classic Taylor-Couette (T-C) flow experiment models a disk drive’s dominant turbulence pretty well: a tall cylinder spins inside a fluid-filled glass sleeve, which displays the fluid’s toroidal convection. Laser-Döppler velocimetry experiments have precisely measured the fluctuating convection in T-C flow [9]. The flow changed from periodic to quasiperiodic, and then *abruptly* to weakly turbulent, as the fluid’s velocity was gradually increased. This development was consistent with a formal model of weak turbulence in simple quasiperiodic systems:

Theorem 1 (Newhouse, Ruelle, Takens, 1978 [14]). *“Let v be a constant vector field on the torus $T^n = \mathbb{R}^n / \mathbb{Z}^n$. If $n \geq 3$, every C^2 neighborhood of v contains a vector field v' with a strange Axiom A attractor. If $n \geq 4$, we may take C^∞ instead of C^2 .”*

Here, the torus T^n does not represent the toroidal T-C vortices directly, but is a simple dynamical system’s phase space. “Axiom A”⁵ refers to a formal definition of dynamical systems that show a close mixture of periodic and turbulent behavior. (For the definition of Axiom A flows and attractors, see [15].)

⁵ The finite-dimensional Axiom A formalism can’t apply *directly* to Navier-Stokes infinite-dimensional phase flows. The machinery of inertial manifolds, though, has shown that bounded-velocity flows have finite-dimensional attractors [20].

Even weak turbulence is sufficiently random for our purposes, because it shows sensitive dependence on initial conditions (SDIC). (Theorem 1's Axiom A result implies SDIC.) Somewhat formally, a phase-space flow $f_t : S \rightarrow S$ has the SDIC property if \exists an attractor $A \subset S$ s.t. $\forall x \in A, \forall$ small $U \ni x$, the diameter of $f_t(U)$ increases exponentially with time [15]. Informally, to completely forecast a system that shows SDIC, we must specify its parameters and initial conditions with infinite precision; measurement limitations limit the forecast to short-term accuracy [3]. Thus, *it is not computational complexity, but information losses in measurement, that prevent effective prediction in such physical systems.*

3 Converting Access-Times to Random Numbers

In Section 2, we showed that the disk's speed variations show sensitive dependence on initial conditions (SDIC). Even so, the disk's access-times are still strongly structured, biased, and correlated, so they clearly cannot directly simulate a tossed coin. Some solutions to this problem [4, 23] make assumptions that don't fit our noise source. Semi-randomness [16, 22, 8] successfully formalizes imperfect randomness, but is restrictively pessimistic because it requires two sources. This motivates our use of the FFT. In this part of our paper, we show that the FFT is a good unbiasing algorithm, and that SDIC justifies this use of the FFT. In presenting these results we follow Brillinger [7].

Let (X_1, \dots, X_k) be a vector of random variables. The *joint cumulant* of k th order, $cum(X_1, \dots, X_k)$, is defined as the coefficient of $i^k t_1 \dots t_k$ in the Taylor series expansion of the logarithm of the characteristic function of (X_1, \dots, X_k) about the origin. For a stationary time-series X_t with $E|X_t|^k < \infty$, we define the *joint cumulant function* of order k to be

$$c_{X\dots X}(t_1, \dots, t_{k-1}) = cum(X_{t_1}, \dots, X_{t_{k-1}}, X_0) \quad (1)$$

The cumulant is thus the mean of a polynomial function of k staggered copies of the time-series X_t , and its parameters t_i describe the copies' offsets. The cumulants represent the dependencies present in X_t . The requirement that these dependencies fall off over time is known as a *mixing* condition. Our use of the FFT as an unbiasing algorithm rests on the following mixing assumption.

Assumption 2. *The time series X_t possesses moments of all orders and its cumulant functions satisfy*

$$\sum_{t_1} \dots \sum_{t_{k-1}} |c_{X\dots X}(t_1, \dots, t_{k-1})| < \infty. \quad (2)$$

In our case, the disk-timing data X_t are bounded, so the cumulants do exist.

Define the *power spectrum* and the *finite Fourier transform*, respectively, as

$$f_{XX}(\lambda) = \sum_{t=-\infty}^{\infty} |c_{XX}(t)| e^{-i\lambda t} < \infty \quad (3)$$

$$d_V^T(\lambda) = \sum_{t=0}^{T-1} X_t e^{-i\lambda t} \quad (4)$$

where V denotes a T -vector (X_0, \dots, X_{T-1}) . Then

Theorem 3 (Brillinger, 1981 [7]). *Let X_t be a stationary time series which satisfies Assumption 2, and let $\lambda_1, \dots, \lambda_k$ be distinct values in the interval $[0, 2\pi]$ s.t. $\lambda_j \neq 0, \pi, 2\pi$. Then as $T \rightarrow \infty$, the values $d_V^T(\lambda_j)$ converge asymptotically to independent (complex) normal random variables with mean 0 and variances equal to $2\pi T f_{XX}(\lambda_j)$.*

As an immediate consequence it is clear that the phase angles $\phi_j = \arg(d_V^T(\lambda_j))$ are asymptotically independent and uniformly distributed on the interval $[0, 2\pi]$.

Theorem 3 is a generalization of the Law of Large Numbers. Like the L.L.N., it grants perfect normality only in the limit as $T \rightarrow \infty$, so the spectral distributions are only approximately normal. Another price of the theorem's generality is that the output distribution converges only pointwise to the desired joint normal. Lacking convergence-rates, we must measure how well the spectra approach normality, so that we can choose a practical spectrum-length T .

To feed the FFT, we filtered and decimated our raw disk-periods to remove some obvious quantization structure. We discarded the FFT's predictable spectral lines at $\lambda_0, \lambda_{T/2}, \lambda_{T/2+1}, \dots, \lambda_T$. Note that if we take from this algorithm more bits than we feed into it, we run the risk that the mapping $\arg(d_V^T) : V \rightarrow [0, 2\pi)^{T/2}$ can be inverted, even if we only keep part of the spectrum as random numbers.

We claim that disk access-times satisfy Assumption 2. No statistical test can justify such a claim, so we will argue instead that the cumulants' decay follows from SDIC. The N.R.T. theorem ensures exponentially-damped autocorrelation [14], but says nothing about higher-order correlations.

The mixing condition is formally very different from SDIC, but their meanings are similar: loosely, both put limits on how measurements can aid prediction of the system's long-term behavior:

- SDIC means that *measured* initial conditions are insufficient;
- mixing means polynomial functions of past measurements are insufficient.

To see that SDIC implies our mixing condition, suppose that a cumulant of order k fails to decay, so that arbitrarily long-term dependencies can exist among k measurements of the time series. Then $n < k$ early measurements can suffice to predict some function of $k - n$ later measurements' values, for some arbitrarily large separations between the two groups. The n early samples, though, all have limited accuracy, and cannot specify the underlying system's state with infinite precision. Thus, we've contradicted the SDIC property: the first n samples specify a range of initial conditions, whose consequences are recurrently and significantly parallel.

4 Statistical Analysis

In this section we summarize two data analyses, as “sanity checks.” We tested:

- our measured access-times, to ensure that noise was present.
- the RNG’s product, for several deviations from randomness.

We call these measurements sanity checks, because our argument for the disk’s value as a noise-source actually rests on the mathematical properties of the disk’s air turbulence, and not on our observations. These tests’ failure would have disproven our claim that the disk’s motion reflects turbulent flow.

For our measurements, we used an IBM RT/PC desktop workstation and a Micropolis 1320 series 40 Mb hard disk with nonremovable 5.25 inch media. A permanent-magnet brushless DC motor turns the disk spindle at a nominal rate of 3600 r.p.m. The motor’s phase-locked loop stabilizes the rate to $\pm 0.03\%$, which amounts to a positional accuracy of $5 \mu\text{sec}$. [19].

The workstation’s operating system was MIT Project Athena’s port of 4.3bsd UNIX, with machine-dependent routines from the IBM Academic Information Systems release. For our tests, we debugged the UNIX kernel’s `microtime()` subroutine, and we modified the disk-scheduler software to record the time at each disk-access’ initiation and completion. We were particularly careful to avoid disturbing the spindle’s speed with head motion. During experimental sessions, the workstation ran “standalone,” isolated from the MIT network, with no time-synchronization software or other inessential processes running. Sessions lasted from 30 minutes to 8 hours. To measure disk-speed fluctuations, we repeatedly read a chosen disk block, and recorded each access-completion time. This entails so little software overhead that we could read the block on every rotation, so the completion-time differences gave a running account of the disk’s period. The RT’s 1024 Hz. hardware clock limited our measurement precision to ~ 1 msec.

Our measurements were consistent with the $5 \mu\text{sec}$ variation. We considered a variety of influences whose timing effects might resemble rotational latency:

- delays within the disk controller,
- bus arbitration,
- instruction and I/O caching effects.

We believe that the RT’s very simple disk controller and interrupt mechanism make these effects negligible. Our analysis of 1.7 million disk-periods showed that some noise was present in the variation, its auto-correlation fell off within 5 seconds, and its entropy amounted to about 100 bits/minute [12], enough for 2,600 highly random DES keys/day.

From 100,000 access-times, and using an FFT vector-length $T = 1000$, we gathered $\sim 50,000$ complex-valued spectra, and found that

- the real and imaginary parts passed various Q-Q plot normality tests;
- the angles passed various runs tests, and their a.c.f showed no correlation;
- the angles’ 64 bits were pairwise uncorrelated.

We by no means intend these tests to be definitive, because we take it for granted that *a posteriori* arguments for randomness are inconclusive.

5 Conclusion

Experiments by ourselves and others show that the disk's speed fluctuates measurably because of air turbulence and other factors. Our random number generator uses the FFT algorithm to convert the measured variations into uniformly-distributed and independent variables. In a "worst-case" experimental scenario, we have measured 100 bits/min of entropy in a quiescent disk's speed variation.

We have also sketched a mathematical justification for our claim that our generator's product is truly random. In summary,

1. Disk drives have Taylor-Couette turbulence [1];
2. The N.R.T. theorem applies to Taylor-Couette turbulence [9];
3. N.R.T. theorem \Rightarrow Sensitive dependence on initial conditions [14];
4. SDIC \Rightarrow cumulants decay [Section 3];
5. Cumulants decay \Rightarrow independent, normal spectra [7].

Turbulence's unpredictability is formally and experimentally well-founded in nonlinear dynamics. The SDIC criterion ensures that disk access-times satisfy a statistical mixing condition, which in turn ensures that the time-series' spectra are nearly independent, nearly normal variables. Sanity check statistical analyses, of disk periods and their spectra, are consistent with our argument.

Our experimental scenario is unrealistically constrained, but yields enough random bits to meet a large installation's key-service needs. We believe we could amplify this high-quality entropy by allowing the head's motion, the disk scheduler, and spindle-speed variations to influence each other synergistically. Other hardware noise-sources offer more bandwidth, but this one costs nothing, so our "per bit" price is very competitive.

Acknowledgments We thank John Carr for his help with `microtime()`; Fred Kurzweil and Hugh Sierra (formerly of IBM), and Brian Tanner of Micropolis, for their patient explanations of disk drive design and manufacture; and Peter Constantine, John Eaton, Ciprian Foiaş, Shafi Goldwasser, Bernardo Huberman, Mark Lillibridge, Zbigniew Nitecki, Olin Sibert, Ralph Swick, Harry Swinney, Roger Temam, and Jim White, for helpful discussions, references, and advice.

References

1. S.D. Abrahamson, C. Chiang, and J.K. Eaton, "Flow structure in head-disk assemblies and implications for design," *Adv. Info. Storage Syst.*, **1** (1991). pp. 111–132.
2. W. Alexi, B. Chor, O. Goldreich, and C.P. Schnorr, "RSA and Rabin functions: certain parts are as hard as the whole," *Proc. 25th IEEE Symp. on Foundations of Computer Science*, 1984, pp. 449–457; see also *SIAM J. on Comput.*, **17**(2) (1988).
3. P. Bergé, Y. Pomeau, and C. Vidal, *Order Within Chaos: Towards a Deterministic Approach to Turbulence*, Wiley, New York, 1984.
4. M. Blum, "Independent unbiased coin flips from a correlated biased source: a finite state Markov chain," *Proc. 25th Ann. Symp. on Foundations of Computer Science*, 1984. pp. 425–33.

5. L. Blum, M. Blum, and M. Shub, "A simple unpredictable pseudo-random number generator," *SIAM J. Comput.*, **15**(2) (1986). pp. 364–83.
6. M. Blum and S. Micali, "How to generate cryptographically strong sequences of pseudo-random bits," *SIAM J. Comput.*, **13** (4) (Nov. 1984). pp. 850–864.
7. D. Brillinger, *Time Series: Data Analysis and Theory*, Holden-Day, San Francisco, 1981. Addendum.
8. B. Chor and O. Goldreich, "Unbiased bits from sources of weak randomness and probabilistic communication complexity," *Proc. 26th Ann. Symp. on Foundations of Computer Science*, 1985. pp. 429–42.
9. P.R. Fenstermacher, H.L. Swinney, and J.P. Gollub, "Dynamical instabilities and the transition to chaotic Taylor vortex flow," *J. Fluid Mech.* **94**(1) (1979). pp. 103–128.
10. T. Kenjo and S. Nagamori, *Permanent-Magnet and Brushless DC Motors*, Monographs in Electrical and Electronic Engineering No. 18, Clarendon Press, Oxford, UK, 1985.
11. T.G. Lewis, *Distribution Sampling for Computer Simulation*, Lexington Books, Lexington, Mass., 1975. p.3.
12. U. Maurer, "A universal statistical test for random bit generators," *Crypto '90 Conference Proceedings*, Springer-Verlag Lecture Notes in Computer Science **537**, New York, 1991. pp. 408–420.
13. S. Micali, and C.P. Schnorr, "Efficient, perfect random number generators," *Crypto '88 Conference Proceedings*, Springer-Verlag Lecture Notes in Computer Science **403**, New York, 1990. pp. 173–198.
14. S. Newhouse, D. Ruelle, and F. Takens, "Occurrence of strange Axiom A attractors near quasi-periodic flows of T^m $m \geq 3$," *Commun. Math. Phys.* **64** (1978), pp. 35–40.
15. D. Ruelle, *Elements of Differentiable Dynamics and Bifurcation Theory*, Academic Press, San Diego, 1989.
16. M. Santha and U.V. Vazirani, "Generating quasi-random sequences from semi-random sources," *J. Comput. System Sci.*, **33** (1986). pp. 75–87.
17. B. Schneier, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*; Wiley, New York, 1994. p. 370.
18. H. Sierra, *An Introduction to Direct-Access Storage Devices*, Academic Press, Boston, Mass., 1990. pp. 100–106.
19. Brian Tanner, personal communication.
20. R. Temam, *Infinite-dimensional dynamical systems in mechanics and physics*, Springer-Verlag Applied Mathematical Sciences **68**, 1988, pp. 389–92.
21. D.F. Torok and R. Gronseth, "Flow and thermal fields in channels between corotating disks," *IEEE Trans. on Components, Hybrids, and Manuf. Tech.*, **11**(4) (Dec. 1988). pp. 585–593.
22. U. Vazirani, "Towards a strong communication complexity theory, or generating quasi-random sequences from two communicating slightly-random sources," (extended abstract, undated).
23. J. von Neumann, "Various techniques used in connection with random digits," Notes by G.E. Forsythe, National Bureau of Standards, Applied Math Series, Vol. 12, pp. 36–38, Reprinted in von Neumann's *Collected Works*, Vol. 5, Pergamon Press (1963). pp. 768–770.
24. A.C. Yao, "Theory and applications of trapdoor functions," *Proc. 23rd IEEE Symp. on Foundations of Computer Science*, 1982. pp. 80–91.

This article was processed using the \LaTeX macro package with LLNCS style