Randomness and Cryptography
with Yevgeniy Dodis

by April Bacon

©Gent explant rectaque eost laboremoles

© NYU Photo Bureau (Hollenshead)

Across two decades, Yevgeniy Dodis has explored theoretical limits of randomness generation and extraction, and devised ways to illuminate and fortify its foundation, entropy.

“Randomness in cryptography is like the air we breathe. You can’t do anything without it,” says Yevgeniy Dodis, Professor of Computer Science at Courant. “It’s needed for everything: generation of keys, cryptographic protocols, masking—you name it.”

It is fundamental to the field because secrets are fundamental to the field; cryptography is only possible when a secret can be kept safe from a potential attacker, and a secret that isn’t random to that attacker isn’t truly secret. Yevgeniy, a Courant alum (B.S., 1996) who joined the Institute as a faculty member in 2001, has gone deep and wide into the subject.

There are many ways a computer can attempt to find randomness for cryptographic purposes. For example, it can create a sequence of numbers mathematically or collect bits by tracking physical processes such as the processor’s temperature, interrupt timing, or the movement of a computer mouse. While the possibilities for how to generate bits are as wide as a cryptographer’s imagination, these sources are not guaranteed—nor even likely—to produce perfect randomness. Luckily, such imperfect sources can still be sufficient for real-world applications of cryptography. Much of Yevgeniy’s work on randomness consists of characterizing the precise conditions when a source is “good enough” for a given application and, when these conditions are met, devising the most efficient way to use that source.

A source doesn’t have to be perfectly random because randomness is not an on and off switch, it’s a spectrum. Between true randomness and complete predictability (such as the sequence “0000”) lies a mathematical concept called entropy. The entropy of a given source tells us just how unpredictable—and therefore secure—it is. A one-hundred-bit source, for example, can have entropy ranging from zero (totally predictable) to one hundred (truly uniform) bits; twenty bits of entropy guarantees that the source cannot be guessed with probability better than 2^{-20}, which is less than one in a million. Secrets require at least a few hundred bits of entropy, otherwise they can be easily guessed. This measurement alone does not tell the whole story, as 100 truly random bits appears to be much more useful than a million-bit source with 100 bits of entropy “scattered” throughout otherwise predictable bits.

In early work, Yevgeniy investigated formally what degree of entropy is sufficient for different cryptographic tasks. He showed that even very scattered entropy is likely sufficient for authentication tasks such as digital signatures. But in a series of cornerstone papers in the early 2000s, he and various co-authors demonstrated that such is not the case for privacy tasks (such as encryption), which cannot be based on entropy alone. Even more surprisingly—and of great philosophical importance to understanding the role of randomness in cryptography—these privacy tasks require true randomness.

There are ways to meaningfully overcome this fact, such as with privacy amplification, an area first developed in the late 80s. The technique combines two initial sources—one perfect public source and one imperfect secret source—to extract a new, nearly perfect and secret source. In other words, public perfect randomness can be used to “purify” imperfect secret randomness. The success of the process is measured in minimizing its “entropy loss.” Entropy loss is the difference between the entropy of the secret source given as input, and the length of the nearly perfect randomness that is extracted from it. Prior work on privacy amplification achieved entropy loss of 128 bits for “industry-grade” security 2^{-64}. This is a high price to pay because entropy is already scarce in many cases, such as when taken from biometric data (as discussed below).

In a series of recent works, Yevgeniy and co-authors achieved the same level of security with strikingly lower entropy loss: just 10 bits for any authentication and 64 bits for most privacy applications (including encryption). The result has important practical implications; as Yevgeniy puts it, “If you need randomness to produce cryptographic keys, you don’t need as much entropy as for full randomness extraction.”

In work that has over 2500 citations and which this year was selected for a Eurocrypt Test-of-Time award for the year 2004, Yevgeniy and co-authors tackled the question of securely extracting cryptographic keys from biometrics and other noisy data, such as fingerprints and retina scans. Specifically, such data is not only imperfect in terms of its entropy, but also noisy: repeated readings of the same data will likely be close, but not identical. The resulting cryptographic primitive is called a fuzzy extractor. As Yevgeniy explains, “First I measure my fingerprint to derive the key. That’s the true secret. The next time I measure my finger, it’s going to be close but not exactly the same. So how do I reliably extract the same key from close-but-noisy readings?”

Yevgeniy’s approach is to decouple the issue of noise and extraction. With the first reading, helper information is created through another primitive he developed, known as "secure sketch." The original reading maintains most of its entropy, even if the helper information is public, so the helper information can be stored without risk of exposing the key. The next time a noisy reading is taken, the helper information allows the exact initial reading to be reproduced, and so the same key is derived the second time around.

One of the rewarding experiences from this work on fuzzy extractors was that it found so many unexpected applications beyond biometrics, such as differential privacy and physically unclonable functions. As Yevgeniy says, "If you do something clean and elegant, science will be kind to you."

Another important area of Yevgeniy's research on randomness is his influential work on random number generators (RNGs). Random number generators are tools built into computer operating systems to produce, as Yevgeniy says "randomness on steroids." From a small amount of randomness in their secret state, RNGs repeatedly produce plentiful amounts of "pseudo-randomness" in the foreground for any process that requires it. Although this pseudo-randomness is not perfect, no efficient attacker can tell it apart from true randomness. The foreground part of this process has been well understood since the late 80s for cases in which the source in the secret state is random to begin with. A far less understood process happens in the background, where an RNG repeatedly incorporates fresh entropy from various imperfect entropy sources (e.g., timing of computer interrupts, etc.) into their small state. This background process should "work like a sponge," says Yevgeniy, looking for entropy everywhere and absorbing it like water. Like a sponge, the generator will "mix up" the entropy that it takes in, without necessarily knowing how much it has or where it might be located. This rapid entropy accumulation safeguards the RNG in face of a computer reboot or potential state compromise—without it, the foreground process of pseudo-randomness generation will lack enough initial entropy and will provably fail.

Yevgeniy was the first to formalize the process of entropy accumulation, which is at the heart of all existing RNG designs. Formerly, "random number generators inside computers were all ad hoc," he explains. RNGS are "complex and hard to understand; as such, they're hard to attack. And because they're hard to attack, the theory behind them was lacking. I wanted to change that—to bring this important area of cryptography on par with encryption and authentication." In particular, Yevgeniy reduced part of the problem of sound entropy accumulation to an online randomness extractor and then made several constructions of such online extractors.

Yevgeniy has applied his theory to real-world RNGs, revealing theoretical weaknesses in the RNG used by the Linux operating system. By comparison, Windows has a very secure random number generator, and macOS is somewhere in between. His work attracted several high-profile discussions on the subject and ongoing interest from Microsoft and Apple, which Yevgeniy hopes will influence their future RNG releases.

Randomness extraction—applying methods to an imperfect source to "extract" a much better one—appears in all of the above examples as a powerful tool to deal with imperfect randomness. Yevgeniy first utilized such extractors for his doctoral dissertation at M.I.T. in 2000. With randomness extractors as one important component, Yevgeniy developed solutions for "Exposure-Resilient Cryptography"—i.e. maintaining the viability of a key even when that key has been partially exposed. For example, hardware may be physically stolen and halfway hacked, or malware may extract bits of secret information. Yevgeniy's dissertation shows that an attacker can uncover quite a bit about the actual secret without the application being compromised, by carefully extracting a shorter, "virtual secret" inside the actual secret. This virtual secret will be perfectly secure, even if the actual secret is partially compromised.

"A lot of things you can do in cryptography are seemingly impossible," says Yevgeniy. "I can prove to you that a statement is true without telling you anything else about the statement, beyond its validity. You're convinced, have no doubts, but you don't know why. This is zero-knowledge. I can do electronic currency—I can give you a string of bits which is money. You can see that it is money and, somehow, can spend it only once. These things are counterintuitive—they are like puzzles."

"Cryptography is really all about puzzles, and I love puzzles," he says.

Yevgeniy's first experience with cryptography was as a graduate student at M.I.T., in a class with Shafi Goldwasser.

"It really intimidated me," he says. "She went full speed into research, and I was used to just taking classes and doing homework." At the time, Yevgeniy's primary research area was in lower bounds. He did well in cryptography, but didn't think it was for him. The following year another cryptographer, Silvio Micali, was the head of his qualifying committee for candidacy into the doctoral program.

"Instead of just saying, 'You passed,' he said, 'You know, why don't I take you for lunch? Let's talk.' It was luck—he was looking for students because he had been on sabbatical. He said, 'You seem to be a talented guy, here is a cool problem.'" Micali had just picked up the problem while visiting another professor at the University of Montreal. It was about lower bounds in cryptography and didn't require much knowledge in the field.

"It was just complete serendipity," says Yevgeniy. Not only was the problem related to lower bounds, then his primary area of study, but it was also solvable using techniques he had learned while taking an elective outside of the computer science department, in electrical engineering.

"That very evening I solved the problem," he says. "Silvio was excited. Because I didn't have any experience in cryptography, he sat with me, and we wrote the entire paper together. He had to translate my technique to the proper notation because I had never written cryptography papers." Previously, Yevgeniy had been struggling to get papers on lower bounds accepted to conferences, but this new paper was accepted to Eurocrypt, the most prestigious conference in cryptography.

"I can summarize what I learned from [Goldwasser and Micali, now Turing Award winners] in one word: aesthetics," he says. "This is something I try to teach to my students. There are proofs which are beautiful; there are proofs which are ugly. I'm a deep believer that aesthetics governs the world, at least in science. There are counterexamples—complex papers which require lengthy and tedious calculations. Some of my papers are like that as well, they require you to just roll up your sleeves and dive in. But my favorite work is elegant: clever work that can be explained to an expert in five minutes. I don't write it on my grant applications, but for me, one of the main values in a paper is what is beautiful." ■

Spring 2019

Published by the Courant Institute of Mathematical Sciences at New York University

Courant Newsletter

Volume 14, Issue 2

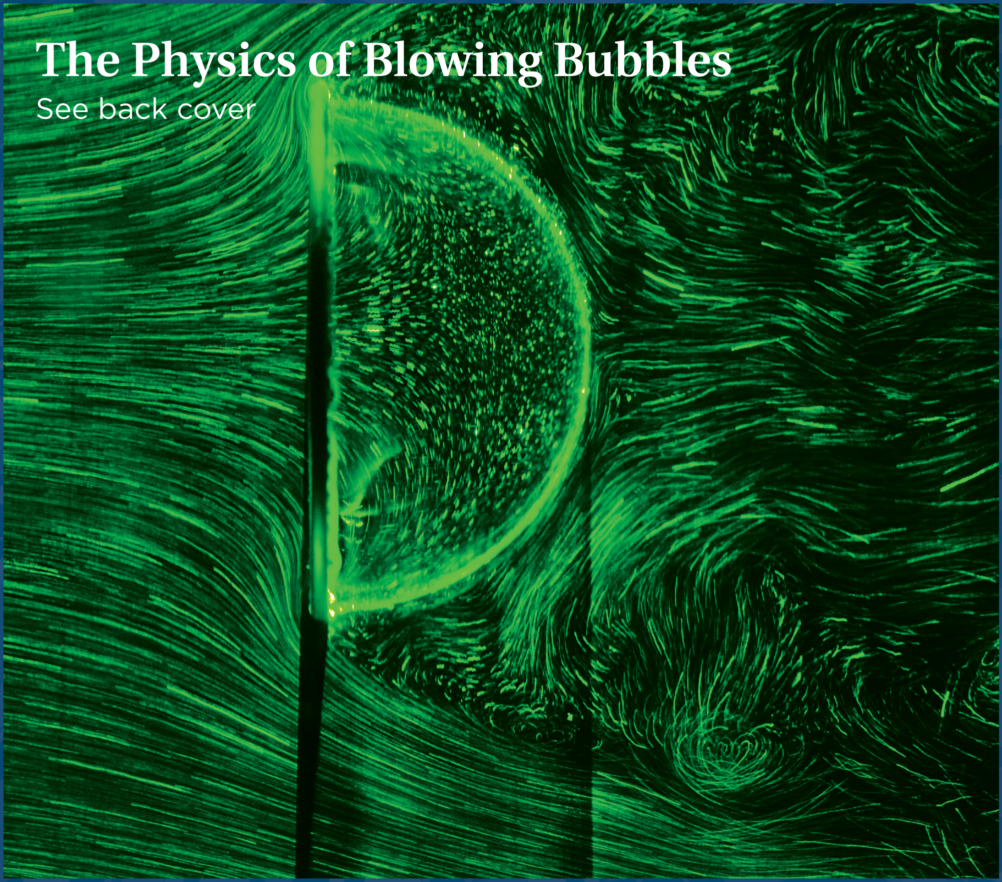

The Physics of Blowing Bubbles

See back cover

p04

YEVGENIY DODIS ON RANDOMNESS AND CRYPTOGRAPHY

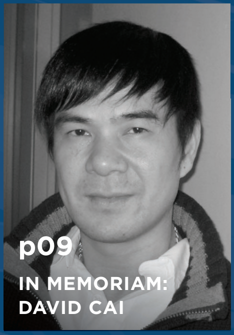

p09

IN MEMORIAM: DAVID CAI

p02

MIRANDA HOLMES-CERFON ON SELF-ASSEMBLY

p07

HENRY MCKEAN RETIRES

ALSO IN THIS ISSUE:

NYU | COURANT INSTITUTE OF MATHEMATICAL SCIENCES