

A Public Randomness Service

Michael J. Fischer¹, Michaela Iorga² and René Peralta²

¹Yale University, Department of Computer Science

²National Institute of Standards and Technology, ITL, Computer Security Division

Keywords: Randomness Server; Electronic Commerce, Electronic Voting

Abstract: We argue that it is time to design, implement, and deploy a trusted public randomness server on the Internet. NIST plans to deploy a prototype during 2011. We discuss some of the engineering choices that have been made as well as some of the issues currently under discussion.

1 INTRODUCTION

The theoretical community has developed many clever cryptographic security protocols over the years for access, authentication, privacy, and authorization in networking and e-commerce applications. However, except for the simplest and most basic protocols, few have been widely deployed. A major reason concerns efficiency. Many of the more sophisticated security protocols, such as Zero Knowledge proof systems, are highly interactive and require too many communication rounds to be feasible in most situations. Other privacy-preserving protocols eliminate the need for many rounds of communication but assume the availability of a trusted source of randomness, an assumption that is not generally valid at present. We argue that it is time to design, implement, and deploy a trusted public randomness server on the Internet. The “NIST Beacon”¹ project aims at doing just that, starting with a prototype in 2011.

Trust is a complex concept, involving technical as well as social components. At the technical level, three fundamental properties will be provided: *unpredictability*, *autonomy*, and *consistency*. Unpredictability means that users cannot predict bits before they are made available by the source. Autonomy means that the source is resistant to attempts to alter the distribution of the random bits. Consistency means that a set of users can access the source in such a way that they are confident that they all receive the same random string. We describe some applications of trusted public randomness servers in sections 2 and 3.

Note that the requirements for some applications

¹Rabin appears to have first proposed this type of service, calling it a “beacon” (Rabin, 1983)

of random numbers are quite different from these. For example, password generation requires confidentiality of the random string and the inability of an adversary to recover the string, even long after the fact. This stands in sharp contrast to our requirement of consistency for the NIST Beacon.

Experience teaches us that it is extremely hard to secure any single point in the Internet. However, it would be hard to simultaneously compromise the integrity of several independent servers. Therefore, we see our role as proposing a format that can be emulated by others so that the end result is several independent servers available to users. These servers can then be incorporated into a suitable multi-party protocol to provide a collective randomness server that achieves unpredictability, autonomy, and consistency, and that is also resistant to the failure or corruption of a limited number of servers. The basic idea is to take the XOR of the shares produced by an agreed-upon set of servers. However, other issues must be addressed as well, including how to control membership in the approved server pool, and how to prevent a corrupted server from learning the other servers’ random shares before choosing its own share.

2 A TRUSTED PUBLIC RANDOMNESS SERVICE

The most basic service a trusted public randomness service provides is to post a random bit string S_L of a fixed length L , time-stamped by its time of generation, and digitally signed by the server. The difference between a random string obtained from the randomness service and one generated locally by flipping coins is that the former is widely trusted to have resulted

from a prescribed random process, was not “cooked” or otherwise falsified, and was not known to anyone before the indicated time.

2.1 A Simple Authentication Mechanism

The following situation occurs frequently in modern cryptographic applications:

- User P claims to have authorization to pass through a security checkpoint.
- The sentinel V says to P , “Passage is restricted to those who are able to invert the function f . Here is a challenge number y . Please tell me what $f^{-1}(y)$ is.”
- P calculates and sends back $x = f^{-1}(y)$ to V . Then V checks that $f(x) = y$ and, if so, allows P to pass.

The terms “sentinel”, “security checkpoint”, and “user” are abstractions. The user can be a program, a person holding a smart card, a client accessing a server, etc. The security checkpoint can be a physical place, an operating system security mechanism, an ATM, etc. The sentinel can be an operating system or one of its subsystems, an algorithm running on a network component, a human guard, etc.

This protocol uses the notion of a trap-door one-way function. Such a function f has the property that, given an arbitrary input x , $f(x)$ can be easily computed. However, only a user who knows the secret trap-door information can invert f in reasonable time.

Desired Properties User P is assumed not to be trusted. The purpose of this protocol is to prevent a corrupted P from getting past the sentinel without actually knowing the secret trap-door information. In many potential applications, the sentinel V also cannot be trusted. A corrupted V might try to obtain the secret trap-door information in order to enable her to bypass security mechanisms not only at her station, but at any other station using the same one-way function. Even with a trusted sentinel, if the sentinel knows the secret trap-door information, then the sentinel’s data must be protected from the outside world. Depending on the application, this may be hard, expensive, or even impossible to do. Thus, one would like to implement this protocol in such a way that the sentinel does not need to know the trap-door information. Furthermore, we would like the trap-door to remain secret after the execution of the protocol so that the mechanism can be safely reused.

The beautiful insight, that perhaps it is not necessary for the sentinel to know the trap-door, is attributed to von Neumann.² Clearly, only P needs to know the trap-door secret in order to carry out this protocol, for only P inverts f . The sentinel V only needs the ability to generate suitable challenge numbers and to compute f in order to verify the challenge response. But before we can be assured of the safety of this protocol, we must ensure that V doesn’t inadvertently learn the trap-door secret during the protocol’s execution. Moreover, this should remain true even if V is corrupt and is maliciously trying to compromise P ’s secret.

Use of Randomness A question that comes to mind when considering the security of this protocol is, “How does the sentinel V choose the number y ?” Clearly, if y were a fixed value, then anybody who witnessed an execution of the protocol could later pass through the checkpoint, as he would know $f^{-1}(y)$. Therefore, the sentinel must change y in each iteration of the protocol. But a simple change of y (such as adding 1 to the previously-used challenge) might allow special properties of f to be used to compute $f^{-1}(y)$ from previously known values of $f^{-1}(y)$. Ideally, from the point of view of preventing a corrupt P from passing the checkpoint, the number y should be chosen uniformly at random from a large set of “hard to invert” numbers. Thus, we must provide the sentinel with a good random (or pseudo-random) number generator. Fortunately, this is not hard to do. But neither is it trivial. For example, linear congruential generators turn out to be predictable by polynomial-time algorithms (Boyar, 1989).

Another, less obvious, question that arises is, “How does the user P know that the sentinel chooses y at random, and why does he care?” In general, the user has no way of knowing, but it is important that the protocol be designed in such a way that the secret trap-door information remains secret regardless of how a corrupt V chooses y .

An Insecure Implementation: Consider the following implementation of a trap-door function

- N is a product of two large primes p and q , each congruent to 3 modulo 4. Such a number is called a “Blum Integer”. The trap-door is the prime p . It is known by authorized users but not by sentinels. N is publicly known.
- $f(x) = x^2 \pmod N$.

²He, of course, did not pose the problem in these terms, as the notion of trap-door one-way function is quite recent.

To pass the checkpoint, users must be able to demonstrate the ability to compute square roots modulo N . Since not all numbers have modular square roots, we must first solve the problem of how a sentinel should produce a random challenge. For Blum integers it is easy to find a number α modulo N such that for all $y \in Z_N^*$, exactly one of the four numbers in the set $C_y = \{\pm y \bmod N, \pm \alpha y \bmod N\}$ has a modular square root.^{3,4} The challenge issued by the sentinel to the party who wants to prove its knowledge of the factors of N is simply a number y chosen uniformly at random from Z_N^* . The response should be a modular square root x of the element in C_y that is guaranteed to have a modular square root. To verify that x is a valid response to the challenge, the sentinel only needs to compute $x^2 \bmod N$ and test for membership in the set C_y .

It can be formally shown that users reveal no information about the trap-door by responding to a challenge generated according to this protocol. On the other hand, factorization of a composite N is (probabilistic) polynomial-time reducible to the problem of computing square roots modulo N . Thus, the usual assumption of intractability of factoring implies that a user who does not know the factorization of N will not be able to compute the necessary square root.

We have just described an authentication mechanism that works provided all parties follow the protocol. The problem, in many applications, is that neither party can assume the other is acting honestly. For example, the sentinel may be after the trap-door secret and may not follow the protocol in generating y . Such a sentinel can discover the factorization of N with high probability, rendering this protocol absolutely insecure. Here's how. The dishonest sentinel generates y by choosing u at random from Z_N^* and then choosing y at random from the set C_{u^2} . It is easily shown that the numbers y chosen in this way are uniformly distributed over Z_N^* ; hence, it is undetectable by P that V is not following the protocol. The value x that P returns satisfies $x^2 = u^2 \bmod N$. With probability 0.5, $x \neq \pm u \bmod N$, in which case $\gcd(x + u, N)$ is a proper factor of N . The gcd is an easy computation; therefore the trap-door will not remain secret for long from this cheating sentinel.

To summarize, the failure of this protocol came about because P could not detect that V picked the number y in a special way that gave her additional information about y . However, the protocol does work

³ Z_N^* is the group of invertible elements modulo N .

⁴The reader acquainted with Number Theory will recognize that α can be any number with Jacobi symbol -1 modulo N . Such a number is easily found in probabilistic polynomial time without having to factor N .

correctly if both parties can trust that the challenge y is an exogenously generated random number. This is precisely what a public trusted randomness server provides. If we simply change the protocol so that in the second step, P and V obtain y from the trusted randomness server, then the protocol indeed becomes secure.⁵

The above example, aside from whatever merits it may have as an authentication scheme, was constructed to introduce the main practical issue addressed by this paper:

Access to a common trusted source of randomness can make simple protocols secure that otherwise would not be.

3 Applications of a Trusted Public Randomness Server

A trusted randomness beacon has many different uses. Here are some examples.

3.1 Cryptographic Primitives

Providing network security and reliability requires the use of *cryptographic primitives*. Examples of such primitives are encryption, decryption, and digital signatures. Over the last three decades, cryptographers have identified a number of other primitives as being powerful tools for developing secure network applications. Some early examples of these are bit-commitment (see (Boyar et al., 1990; Boyar et al., 1993; Brassard et al., 1988; Brassard and Crépeau, 1987; Goldwasser and Micali, 1984)), oblivious transfer (see (Halpern and Rabin, 1983; Fischer et al., 1985; Berger et al., 1985; Fischer et al., 1996)), digital coin-flipping (Blum, 1982), cryptographically secure pseudo-random number generators (Blum and Micali, 1984), and zero-knowledge (ZK) proofs (Brassard and Crépeau, 1987; Goldreich et al., 1991). The latter primitive implies the ability to prove to a third party that a Boolean function $f(x)$ is satisfiable without revealing a satisfying assignment. Furthermore, some instantiations of this primitive allow proving *knowledge* of a satisfying assignment x_0 for f without revealing x_0 .

ZK proofs are interactive: the prover engages the other party (the “verifier”) in a conversation. After the

⁵In practice, we would probably use a one-way version of this protocol in which P presented both y and $f^{-1}(y)$ to V , and V checked that $f(x) = y$ and that y was a recent and previously-unused value from the beacon. Time-stamps, signatures and possibly other features can be used in order to guard against replay attacks.

conversation is over, the verifier is convinced that $f(x)$ is satisfiable but has not obtained any information besides this fact. Unfortunately, ZK proofs are usually impractical: they require too much interaction and involve too much communication and computation.

There are a number of variants of ZK proofs in which the interaction is minimized, both in total number of bits communicated and in number of rounds. Among these, the most practical protocols assume, in one way or another, access to a common random string.

3.2 Voting

Voting technology is currently in a state of flux. There are various ways in which new technologies are being used. Ensuring security and promoting trust in these new applications is a difficult challenge. A common source of randomness will be useful in at least two ways: i) in random auditing of machines and ballots (see, for example, (Norden et al., 2007)); ii) in facilitating so-called end-to-end voting systems (see, for example, (Adida et al., 2009)).

4 Design issues

An online source of randomness is not a new idea. Implementations date to the 1980s (George Davida, at the Univ. of Wisconsin, deployed a system that provides on-demand random strings using white noise from radio waves as the source of entropy.) A currently functioning source of randomness can be found at <http://www.random.org/>. There are many adequate technologies for entropy extraction. There are also published guidelines for randomness generation by standards organizations (see, for example http://csrc.nist.gov/groups/ST/toolkit/random_number.html). This position paper simply argues that it is time to design, standardize, and deploy a service tailored to electronic commerce applications. There are a number of design and implementation issues that need to be addressed. Some of them are the following:

- source of entropy;
- rate: how many bits per second;
- user interface;
- full-entropy strings or cryptographically secure pseudo-random strings;
- authentication method;
- time-stamping method;

- archival properties (e.g. can old strings be authenticated?);
- trust model: what, exactly, can the consumer assume?
- securing the source from cyber attacks;
- using multiple sources to provide tolerance against failed or corrupted sources.

At this moment we are thinking of broadcasting full-entropy bit-strings. We plan to post them in blocks of 256 bits per second. We intend to sign and time-stamp the bit-strings. We also plan to link the sequence of blocks with a secure hash so that it will not be possible, even for the source itself, to retroactively change a block without detection. As for source of entropy, we are talking to NIST physicists. We see no reason not to use the most sophisticated entropy source we can afford.

REFERENCES

- Adida, B., Pereira, O., Marnette, O. D., and Quisquater, J. (2009). Electing a university president using open-audit voting: Analysis of real-world use of helios. In *Electronic Voting Technology/Workshop on Trustworthy Elections (EVT/WOTE)*.
- Berger, R., Peralta, R., and Tedrick, T. (1985). A provably secure oblivious transfer protocol. In *Advances in Cryptology - Proceedings of EUROCRYPT 84*, volume 209 of *Lecture Notes in Computer Science*, pages 379–386. Springer-Verlag.
- Blum, M. (1982). Coin flipping by telephone. In *IEEE COMPCON*, pages 133–137.
- Blum, M. and Micali, S. (1984). How to generate cryptographically strong sequences of pseudo-random bits. *SIAM Journal on Computing*, 13:850–864.
- Boyar, J. (1989). Inferring sequences produced by pseudo-random number generators. *J. ACM*, 36(1):129–141.
- Boyar, J., Krentel, M., and Kurtz, S. (1990). A discrete logarithm implementation of zero-knowledge blobs. *Journal of Cryptology*, 2(2):63–76.
- Boyar, J., Lund, C., and Peralta, R. (1993). On the communication complexity of zero-knowledge proofs. *Journal of Cryptology*, 6(2):65–85.
- Brassard, G., Chaum, D., and Crépeau, C. (1988). Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37:156–189.
- Brassard, G. and Crépeau, C. (1987). Zero-knowledge simulation of boolean circuits. In *Advances in Cryptology - Proceedings of CRYPTO 86*, volume 263 of *Lecture Notes in Computer Science*, pages 223–233. Springer-Verlag.
- Fischer, M. J., Micali, S., and Rackoff, C. (1996). A secure protocol for the oblivious transfer (extended abstract). *J. Cryptology*, 9(3):191–195. This work was originally presented at EuroCrypt 84.
- Fischer, M. J., Micali, S., Rackoff, C., and Wittenberg, K. D. (1985). An oblivious transfer protocol equivalent to factoring. Presented at the NSF Workshop on the Mathematical Theory of Security, MIT Endicott House, Dedham, Massachusetts, 1985.
- Goldreich, O., Micali, S., and Wigderson, A. (1991). Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *JACM*, 38:691–729.
- Goldwasser, S. and Micali, S. (1984). Probabilistic encryption. *Journal of Computer and System Sciences*, 28:270–299.
- Halpern, J. and Rabin, M. (1983). A logic to reason about likelihood. In *Proceedings of the 15th Annual ACM Symposium on the Theory of Computing*, pages 310–319.
- Norden, L., Burstein, A., Hall, J., and Chen, M. (2007). Post-election audits: restoring trust in elections. Technical report, Brennan Center for Justice at New York University School of Law.
- Rabin, M. (1983). Transaction protection by beacons. *J. Comput. Syst. Sci.*, 27(2):256–267.