COIN FLIPPING BY TELEPHONE: A Protocol for Solving Impossible Problems

Manuel Blum*

Department of Electrical Engineering and Computer Sciences
Computer Science Division
University of California at Berkeley
November 10, 1981

Abstract

Alice and Bob want to flip a coin by telephone. (They have just divorced, live in different cities, want to decide who gets the car.) Bob would not like to tell Alice HEADS and hear Alice (at the other end of the line) say "Here goes... I'm flipping the coin.... You lost!"

Coin-flipping in the SPECIAL way done here has a serious purpose. Indeed, it should prove an INDISPENSABLE TOOL of the protocol designer. Whenever a protocol requires one of two adversaries, say Alice, to pick a sequence of bits at random, and whenever it serves Alice's interests best NOT to pick her sequence of bits at random, then coin-flipping (Bob flipping coins to Alice) as defined here achieves the desired goal:

1. It GUARANTEES to Bob that Alice will pick her sequence of bits at random. Her bit is 1 if Bob flips heads to her, 0 otherwise.
2. It GUARANTEES to Alice that Bob will not know WHAT sequence of bits he flipped to her.

Coin-flipping has already proved useful in solving a number of problems once thought impossible: mental poker, certified mail, and exchange of secrets. It will certainly prove a useful tool in solving other problems as well.

Introduction

Flipping coins by telephone is EASY, as we show below, if one assumes the existence of a COMPLETELY SECURE one-way function. A (NORMALLY SECURE) ONE-WAY FUNCTION is an efficiently computable function of some set into itself whose inverse cannot be computed efficiently except on a negligible fraction of values. A COMPLETELY SECURE ONE-WAY FUNCTION has the additional property that from a knowledge of $f(x)$, one cannot have more than a 50-50 chance to guess efficiently if x has some nontrivial property, e.g., is even (lsb = 0) or odd (lsb = 1).

To flip coins, Alice and Bob should agree on a completely secure 1-1 one-way function f. Alice then selects an integer x unknown to Bob, computes $f(x)$ and sends $f(x)$ to Bob. Bob, who cannot determine some nontrivial property of x from $f(x)$, tells Alice whether he thinks x is even or odd (this is where Bob flips a coin to Alice). At this point, Alice can tell if he guessed right or wrong. To convince Bob, she sends him x.

Such completely secure one-way functions, however, may not only be hard to discover, they may not even exist. We show how a normally secure one-way function can be used to flip coins in a completely secure fashion. Our one-way function is 2-1, i.e., it maps exactly two elements from its domain to each element of its range. A simple property, say even or oddness, will distinguish the two elements x,y that map to the same element $f(x) = f(y)$. If Alice selects x and sends $f(x)$ to Bob, he ABSOLUTELY CANNOT DETERMINE whether she selected the x or the y, $x \neq y$, such that $f(x) = f(y)$. He guesses whether she picked the even or odd number. His guess as told to Alice constitutes his coin flip to her. Since f is one-way, Alice cannot cheat, i.e., cannot tell him y if in fact she selected x.

A coin-flipping protocol with the right properties, of which the one presented here is an example, has numerous applications, eg., to the exchange of (secret) keys [B'81], to the certified mail problem [BR '81] and to the solution of mental poker [SRA'78] WITHOUT commutative locks [MG'81]. The RIGHT PROPERTIES are:

1. If either participant following protocol does not catch the other cheating, he or she can be sure that the coins each have exactly 50-50 independent chance to come up heads (provable under the reasonable assumption that factoring is hard).
2. If either participant catches the other cheating, he or she can prove it to a judge (this assumes that all messages are sent signed).
3. After Bob flips coins to Alice, she knows which coins came up heads, which tails. He should have absolutely NO idea how they came up (not even a good guess).
4. After the sequence of coin flips, Alice should be able to prove to Bob which coins came up heads, which tails.

A COIN-FLIPPING PROTOCOL WITH THESE PROPERTIES CAN BE USED IN ANY PROTOCOL THAT REQUIRES ONE OF TWO ADVERSARIES, SAY ALICE, TO GENERATE AND USE A RANDOM NUMBER, x, WITHOUT REVEALING IT TO HER OPPONENT, BOB. EVEN IF IT IS TO HER ADVANTAGE TO SELECT A PARTICULAR NONRANDOM x, SHE WILL NOT BE ABLE TO DO SO.

BOB FORCES ALICE TO CHOOSE x AT RANDOM BY FLIPPING COINS TO HER.

THE RESULTING SEQUENCE OF BITS IS COMPLETELY UNKNOWN TO BOB AND, PROVIDED BOB FOLLOWS PROTOCOL, COMPLETELY RANDOM. ALICE USES THE SEQUENCE AS THE REQUIRED RANDOM NUMBER, x. LATER, SHE PROVES TO BOB THAT HE FLIPPED THE SEQUENCE x TO HER, THUS ASSURING HIM IT WAS CHOSEN AT RANDOM.

*Supported in part by NSF grant MCS 79-03767

Reprinted from COMPCON SPRING, February 1982

In addition to the above properties, the coin-flipping protocol presented here has the following useful properties:

1. Each participant KNOWS at each step along the way if the other cheated. He or she does not require later proof (e.g., after Alice has revealed the result of the coin-flips to Bob) to determine this. The court is needed only to provide justice, to give independent proof of wrong (or right) doing, or to force an adversary to complete the protocol.

2. Bob can use his public-key, n, provided it is constructed correctly, to flip coins.

3. Bob does not have to create new primes for each coin-flip. Each coin-flip only requires computation time of the order the time required to compute a Jacobi symbol, (x/n), for $0 < x < n$ and n = a 160-digit number. The Jacobi symbol, (x/n), can be computed quickly, in the same order of time required to compute the greatest common divisor, gcd (x,y).

ASSUMPTIONS:

1. Factorization: We assume that no procedure can efficiently factor a number n that is a product of two large primes, except for a negligible fraction of such numbers n. In 1980 technology, this means that a product of two 80-digit primes cannot be factored in reasonable time (not even 5 years) using the most advanced available technology (1000 CRAY-1's working in parallel) on any but a negligible fraction (one in Avogadro's number) of such numbers. A coin-flipping protocol based on the intractability of the discrete logarithm rather than factorization appears in [BM'81].

2. Random Number Generation: We assume that Bob and Alice each have their own true random number generators. The coin-flipping protocol shows how the random number generators enable Bob to generate and flip random bits to Alice.

3. Signatures: Some of the messages in the protocol below are required to be signed. We assume the existence of a secure signature scheme of the sort first suggested by Diffie and Hellman in 1976 [DH'79], and as implemented by Rivest, Shamir, and Adleman [RSA'78] and Rabin [R'79]. Signed messages are placed in quotes and terminated with the signer's name. It is expected that each participant knows (from the protocol) if the message he or she receives is supposed to be signed and refuses to continue the exchange unless the received message IS properly signed.

THE JACOBI SYMBOL:

The Jacobi symbol (x/n) is defined for odd positive integers n and arbitrary (positive and negative) integers x. It has values 0, +1, or -1. As pointed out earlier, the computation of (x/n) is similar to the computation of gcd(x,n) and takes the same order of time to compute [A'78]. An algorithm for computing (x/n) can easily be constructed from the following of its properties:

x, x1, x2, ... are arbitrary (positive or negative) integers, n, n1, n2, ... are positive odd integers:

1. $(x/n) = 0$ if $\gcd(x,n) \neq 1$
2. $(1/n) = 1$
3. $((x1*x2)/n) = (x1/n)*(x2/n)$
4. $(x/(n1*n2)) = (x/n1)*(x/n2)$
5. $(x1/n) = (x2/n)$ if $x1 = x2 \bmod n$
6. $(-1/n) = +1$ if $n = 1 \bmod 4$
   -1 if $n = 3 \bmod 4$
7. $(2/n) = +1$ if $n = 1$ or $7 \bmod 8$
   -1 if $n = 3$ or $5 \bmod 8$
8. $(n1/n2) = (n2/n1)$ if $\gcd(n1,n2) = 1$ and [n1 or n2 = 1 mod 4]
   $-(n2/n1)$ if $\gcd(n1,n2) = 1$ and [n1 and n2 = 3 od 4]

EXAMPLE: $(23/59) = -(59/23) = -(13/23) = -(23/13) = -(10/13) = -(2/13)*(5/13) = (5/13) = (13/5) = (3/5) = (5/3) = (2/3) = -1$.

THE GROUP Zn*

For n a positive integer greater than 1, define Zn* to be the group of positive integers less than n that are relatively prime to n, the group operation being multiplication mod n.

LEMMA

Let $n = p1^{e1} * \ldots * pk^{ek}$, where k is an integer greater than 1; p1,...,pk are distinct odd primes; and e1,...,ek are positive integers. Let a in Zn* be a quadratic residue (square of an integer) mod n.

Then every solution in Zn* of $x^2 = a \bmod n$ is of the form

$$x = [\pm (x1*v1*p2^{e2} * \ldots * pk^{ek}) \pm (x2*v2*p1^{e1} * \ldots * pk^{ek}) \pm \ldots \pm (xk*vk*p1^{e1} * \ldots * p(k-1)^{e(k-1)})] \bmod n,$$

where v1 is any integer such that $(v1*p2^{e2} * \ldots * pk^{ek}) \bmod p1^{e1} = 1$ [footnote 1]. v2,...,vk are defined similarly.

PROOF

See LeVeque [L'77], Theorems 3.21 & 5.2.
!!!
!!!

THEOREM 1

If n is any odd positive integer, so $n = p1^{e1} * \ldots * pk^{ek}$ is defined as in the statement of the lemma except that k = 1 is also permitted, then 1, 2, 3 are equivalent:

1. there exist x,y in Zn* such that $x^2 = y^2 \bmod n$ and $(x/n) \neq (y/n)$.

2. $pi^{ei} = 3 \bmod 4$ for some i.

3. Let a in Zn* be a quadratic residue mod n. Then exactly half the roots in Zn* of the equation $a = x^2 \bmod n$ have Jacobi symbol $(x/n) = +1$ (the other half have $(x/n) = -1$).

PROOF

[footnote 1]: v1 is easily generated by the Euclidean algorithm, which, when applied to the entries in $\gcd(p1^{e1}, p2^{e2} * \ldots * pk^{ek}) = 1$, yields integers u1, v1 such that $u1*p1^{e1} + v1*p2^{e2} * \ldots * pk^{ek} = 1$).

1 => 2: Suppose to the contrary that $p_i^{e_i} = 1 \bmod 4$ for all i. Let a in Zn* be any quadratic residue mod n. If k = 1, then $a = x^2 \bmod n$ has exactly two roots, x and -x [footnote 2]. These satisfy

$$(x/n) = (x/n)*(-1/n) \text{ since } (-1/n) = +1 \text{ when } n = 1 \bmod 4$$
$$= (-x/n).$$

This contradicts 1 for k = 1.

If k > 1, then by the lemma, x satisfies equation 1. Therefore,
$(x/p_1^{e_1}) = (x_1/p_1^{e_1})$
$= (-x_1/p_1^{e_1})$ since $(-1/p_1^{e_1}) = +1$.
Therefore, any two solutions x,y satisfy $(x/p_1^{e_1}) = (y/p_1^{e_1})$. Similarly, $(x/p_2^{e_2}) = (y/p_2^{e_2}),\ldots,(x/p_k^{e_k}) = (y/p_k^{e_k})$. Therefore $(x/n) = (y/n)$. This contradicts 1.

2 => 3: Equation 1 gives all roots of $a = x^2 \bmod n$. Let the root y be obtained from x by changing the sign of x_i, where $p_i^{e_i} = 3 \bmod 4$. Then $(x/n) = (x/p_1^{e_1})*\ldots*(x/p_k^{e_k}) = -(y/p_1^{e_1})*\ldots*(y/p_k^{e_k}) = -(y/n)$.

3 => 1: Immediate.
!!!

PUBLICATION DATE OF n

In the first message of the following protocol, Bob reveals a number n together with its PUBLICATION DATE defined to be the date on which it was first published. A statute of limitations prohibits either party from bringing the other to court, say, 5 years after that publication date. This limitation is needed for several reasons:

1. Bob cannot expect the prime factorization of his number n to remain hidden for more than 5 years, and once it is out, Alice can fool the judge (but this problem can be avoided without a statute of limitations by having Alice sign an additional message below).
2. It is unreasonable to demand that Alice and Bob keep their correspondence for longer than some fixed length of time (5 years).

We assume that each of the parties is aware of this limitation and do not mention it further.

THE PROTOCOL: BOB FLIPS COINS TO ALICE

I. BOB SELECTS n [footnote 3]:

Bob selects at random two distinct (exactly) 80-digit primes p1, p2, both congruent to 3 mod 4; i.e., he repeatedly selects 80-digit numbers at random and tests each until he obtains two primes both congruent to 3 mod 4 [footnote 4]. He multiplies these together to get n = p1*p2.

B => A: Bob sends Alice "My coin-flipping number is n. Its publication date is May 27, 1980. - signed Bob."

II. ALICE TESTS n:

If Alice trusts that n is a 160-digit product of two primes, both congruent to 3 mod 4, then this part of the protocol may be skipped. Otherwise, Alice checks that n has the following two properties [footnote 5]:

a) Alice checks that n is a 160-digit number and that $n = 1 \bmod 4$. The latter implies that n is odd and $(-1/n) = +1$.
b) Alice checks that for SOME x there is almost surely a y such that $x^2 = y^2 \bmod n$ and $(x/n) \neq (y/n)$, as follows:

B => A: Bob selects 80 (distinct) numbers $x_1,\ldots,x_{80}$ chosen at random from Zn* (it is in HIS interest to select these numbers at random). He sends $x_1^2 \bmod n,\ldots,x_{80}^2 \bmod n$ to Alice [footnote 6].

B => B: Alice sends Bob a sequence of 80 randomly chosen bits, $b_1,\ldots,b_{80}$, where each $b_i = +1$ or -1.

B => A: Let $x_i^2 = y_i^2 \bmod n$ where $(x_i/n) = +1$, $(y_i/n) = -1$. (By Theorem 1, two such roots must exist for every i.) Bob sends Alice a particular sequence of 80 numbers: for each i, from i = 1 to i = 80, he sends Alice x_i if $b_i = +1$, or y_i if $b_i = -1$.

This convinces Alice that condition 1 of Theorem 1 holds (it fails with probability $1/2^{80}$ < 1/Avogadro's number) [footnote 7].

From this point on, the number n has been tested and Alice does not have to retest it.

[footnote 2]: By abuse of notation, we assign to x a dual use as variable (in the equation $a = x^2 \bmod n$) and as constant (a particular solution of $a = x^2 \bmod n$).

[footnote 3]: A TRUSTED INTERMEDIARY may be used to select n. He must choose n at random according to the rules given Bob. If the intermediary is trusted by the entire world (trusted to select n according to the rules and not to give away the primes), then everyone can use n to flip coins to everyone else. It is not necessary for either Bob or Alice to know the prime factors of n in order for Bob to flip coins to Alice. In fact, Alice - or whoever is receiving the flipped coin - should NOT know the prime factorization of n, else she can cheat Bob.

[footnote 4]: PRIMALITY TEST Alice and Bob can test if an 80-digit number is prime using one of the efficient algorithms for primality of Gary Miller [M'76], Strassen and Solovay [SS'77], or Rabin [R'80]. Approximately half of all the 80-digit primes are congruent to 3 mod 4.

[footnote 5]: The two properties do not prove that n is a product of two distinct primes both congruent to 3 mod 4 (it might for example be a product of three distinct primes, two of them congruent to 3 mod 4), but they suffice to prove that (with extremely high probability) the protocol is trustworthy.

[footnote 6]: Alice does not have to check if the numbers she receives are distinct or quadratic residues mod n relatively prime to n. This check is automatically part of the following B => A message.

[footnote 7]: Alice has NOT proved that n is a product of just two primes - she has just "proved" that n satisfies statement 1 of Theorem 1, which is all she needs to know. It will actually be in Bob's interest to make n a product of just two primes to prevent Alice from factoring n within the 5-year statute of limitations.

III. BOB FLIPS COINS TO ALICE:

For this protocol to hold in court, messages exchanged in this part (as in I but not II) are signed before delivery. Alice checks that the publication date is within an acceptable tolerance, e.g., within 1 year of the current date, else aborts. If she accepts the publication date, then Bob and Alice can flip 80 (or any number of) coins as follows:

A => B: Alice selects 80 numbers $x1,\ldots,x80$ in Zn^* at random. She sends Bob "n, publication date of n, $x1^2$ mod n,..., $x80^2$ mod n - signed Alice."

The purpose in sending n is for the court to know which of perhaps several transactions this is.

This is a delicate point in the negotiations for Alice. The delicacy has to do with the publication date. If Bob does not respond to the above message, he could nevertheless take Alice to court, saying he sent her his guess ($\pm$ 1) and she didn't pursue the matter (he claims she probably lost the coin-toss). With the judge's protocol given here, the court would then enforce completion of the protocol. If Bob does not wish to continue, it would be a courtesy for him to tell her so in a signed message. The protocol could then stop here. It can stop here anyway provided Alice understands that she may be forced to continue the protocol in court. If this is inconvenient, she should take Bob to court and get a letter from the judge terminating this protocol.

At this point, Bob should check that Alice sent him the correct n and the correct publication date (if not, he should ask her to stop fooling around). Bob cannot tell if $(xi/n) = +1$ or -1 since by 3 of Theorem 1, xi^2 mod n has as many roots with $(xi/n) = +1$ as it has with $(xi/n) = -1$ (this is true even if Bob did not choose n to be a product of two primes).

B => A: Bob sends Alice "n, $x1^2$ mod n,..., $x80^2$ mod n, b1,...,b80 - signed Bob" [footnote 8].

Alice should check that n, $x1^2$ mod n,..., $x80^2$ mod n have not been altered (if so, she asks Bob to stop fooling around). Alice now determines her sequence of random bits: her ith random bit is $ri = +1$ if Bob guessed right about xi; -1 if he guessed wrong.

At this point in the protocol, Alice knows what Bob flipped to her; Bob has absolutely no idea. Whenever she wants to prove to Bob what sequence r1,...,r80 of random bits was flipped to her, she sends the confirmation message:

A => B: Alice sends x1,...,x80 to Bob.

This message need not be signed, but in that case Bob must make sure his factorization of n will not be given away during the next 5 years. If Alice DOES sign this message, then the statute of limitations need NOT apply.

To guard against Alice cheating, Bob computes $x1^2$ mod n,..., $x80^2$ mod n and compares them with what Alice sent him: if they do not agree, then Alice cheated. If they do agree, he computes (x1/n),...,(x80/n) and thereby determines r1,...,r80 [footnote 9].

If Bob should need to flip more coins to Alice, the two can continue to use the same n, skipping parts I and II of the above protocol.

END OF PROTOCOL

THE JUDGE'S PROTOCOL

An iron-clad judge's protocol is one that can be programmed and thereby save needless expense. In what follows, the judge should be viewed as a computer.

In case of a dispute, the judge proceeds as follows:

1. Subpoena all signed messages that have been exchanged. Check that the statute of limitations has not been exceeded (5 years past the publication date on the number n). If so, throw the case out of court.

2. If one of the participants, say Alice, produced a signed-by-Bob message that refers back to a previously signed-by-Alice message (she sent Bob), then Bob must produce the message he received or be found guilty of cheating.

3. Check that n is a 160-digit number that passes the test in II [footnote 10]. If not, Bob is found guilty of cheating.

4. If no messages have provably been exchanged in III, i.e., neither party has produced a message of III signed by the other, then give each of the participants a dated letter asserting that this protocol is declared terminated. (even if Alice sent $x1^2$ mod n to Bob, who decided not to pursue the matter and threw out her signed message). Otherwise, enforce completion of the protocol (if it has not already been completed) [footnote 11].

5. Check that the $x1^2$ mod n,... in the signed-by-Alice message received by Bob is the square mod n of the numbers x1,... in the signed-by-Alice message later received by Bob. If not, Alice is found guilty of cheating.

6. Determine the sequence of random bits that Bob flipped to Alice. Give each of the participants a (dated) paper asserting the court's findings (this is necessary because the case cannot be brought to

[footnote 8]: The extra information, i.e., n, $x1^2$ mod n,..., $x80^2$ mod n, enables the courts to know which of perhaps several transactions this is.

[footnote 9]: Alice cannot cheat since, first, $(xi/n) = (-xi/n)$ so Alice cannot change the Jacobi symbol of xi by changing the sign of xi, and second, Alice cannot compute yi such that $yi \neq xi \bmod n$ and $yi^2 = xi^2 \bmod n$ (under the assumption that she cannot factor n) since gcd $(xi \pm yi) = p1$ or p2.

[footnote 10]: The judge might, instead of testing n, subpoena the prime factors of n, check there are just two 80-digit primes congruent to 3 mod 4, else find Bob guilty of cheating. This requires, however, that the judge be trusted not to divulge the primes, which Bob might be using in other transactions.

[footnote 11]: This is necessary to guard against the possibility that Alice sent $x1^2$ mod n,... to Bob, that Bob responded with +1, -1,... and that Alice decided she did not like the result (she might argue that Bob never sent her the required bits). Or perhaps Alice liked the result, responded to Bob by sending him x1,... and Bob decided he did not like the result (he might argue that Alice never sent him x1,...).

court again once the statute of limitations is exceeded - not to mention the waste of duplicated proceedings).

END OF PROTOCOL

References

[A'78] D. Angluin, "The Complexity of Some Problems in Number Theory," Lecture Notes available from author, Dept. of Comp. Science, Yale U. (1978).

[B'81] M. Blum, "How to Exchange (Secret) Keys," to appear.

[BM'81] M. Blum and S. Micali, "Coin-Flipping into a Well," to appear.

[BR'81] M. Blum and M.O. Rabin, "Mail Certification by Randomization," to appear.

[DH'79] W. Diffie and M.E. Hellman, "Privacy and Authentication: An Introduction to Cryptography," Proc. IEEE, vol. 67 no. 3 (March 1979), 397-427.

[L'77] W.J. LeVeque, "Fundamentals of Number Theory," Addison-Wesley Pub., 1977.

[M'76] G.L. Miller, "Riemann's Hypothesis and a Test for Primality," J. Comput. and System Sci. vol. 13 (1976), 300-317.

[MG'81] S. Micali and S. Goldwasser, "Mental Poker Without Commutative Locks," to appear.

[R'79] M.O. Rabin, "Digital Signatures and Public Key Systems as Intractable as Factorization," MIT LCS TM 1979.

[R'80] M.O. Rabin, "Probabilistic Algorithm for Testing Primality," J. Number Theory, vol. 12 (1980), 128-138.

[RSA'78] R.L. Rivest, A. Shamir, and L.L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," Commun. ACM, vol. 21 (1978), 120-126.

[SRA'78] A. Shamir, R.L. Rivest, and L.M. Adleman, "Mental Poker," in The Mathematical Gardner, ed. D.A. Klarner, pub. by Wadsworth Intrntl (1981), 37-43.

[SS'77] R. Solovay and V. Strassen, "A Fast Monte-Carlo Test for Primality," SIAM J. Comput. vol. 6 (1977), 84-85.