Efficient Generation of Shared RSA keys

Dan Boneh*
dabo@cs.stanford.edu

Matthew Franklin
franklin@parc.xerox.com

Abstract

We describe efficient techniques for a number of parties to jointly generate an RSA key. At the end of the protocol an RSA modulus $N = pq$ is publicly known. None of the parties know the factorization of N. In addition a public encryption exponent is publicly known and each party holds a share of the private exponent that enables threshold decryption. Our protocols are efficient in computation and communication.

Keywords: RSA, Threshold Cryptography, Primality testing, Multiparty computation.

1 Introduction

We present efficient protocols for a number of parties to jointly generate an RSA modulus $N = pq$ where p, q are prime. At the end of the computation the parties are convinced that N is indeed a product of two large primes. However, none of the parties know the factorization of N. We then show how the parties can proceed to compute a public exponent e and shares of the corresponding private exponent. Our techniques require a number of steps including a new distributed primality test. The test enables two (or more) parties to test that a random integer N is a product of two large primes without revealing the primes themselves.

Several cryptographic protocols require an RSA modulus $N = pq$ for which none of the participants know the factorization. A good example is the original Fiat-Shamir authentication protocol [19] where all parties use the same modulus N, but none of them know its factorization. For other examples see [18, 23, 28, 30, 31]. Usually a modulus N with an unknown factorization is obtained by asking a dealer to generate it. The dealer must be *trusted* not to reveal the factorization of N. Our results eliminate the need for a trusted dealer since the parties can generate the modulus N themselves.

Threshold cryptography is a concrete example where shared generation of RSA keys is very useful. We give a brief motivating discussion and refer to [24] for a survey. Let $N = pq$ be an RSA modulus and d, e be signing and verification exponents respectively, i.e. $de = 1 \bmod \phi(N)$. A threshold RSA signature scheme involves k parties and enables any subset of t of them to generate an RSA signature of a given message. No subset of $t-1$ parties can generate a signature. Unlike standard secret sharing [35], the signature is generated without having to reconstruct the private decryption exponent d. A simple approach for obtaining a k-out-of-k threshold signature scheme is as follows [20]: pick random d_i's satisfying $d = \sum_{i=1}^{d} d_i \bmod \phi(N)$ and give d_i to party i. Abstractly, to sign a message m each party computes $s_i = m^{d_i} \bmod N$ and sends s_i to a combiner (that holds no secrets). The combiner multiplies the s_i's and obtains the signature. This way the parties are

*Supported by DARPA contract #F30602-97-C-0326

able to generate a standard RSA signature without having to reconstruct the private key d at a single location. This is clearly advantageous for securing a sensitive private RSA key such as the one used by a Certificate Authority. Constructions providing a t-out-of-k RSA threshold signatures schemes can be found in [16, 13, 21, 34].

An important issue left out of the above discussion is the initial generation of the RSA modulus N and the shares d_i. Traditionally the modulus N and the shares of the private key are assumed to be generated by a *trusted dealer*. Clearly, the dealer, or anyone who compromises the dealer, can forge signatures. Our results eliminate the need for a trusted dealer since the k parties can generate N and the private shares themselves. Such results were previously known for the ElGamal public key system [32], but not for RSA.

The paper is organized as follows: we give a high level description of our protocol in Section 2. In Sections 3–5 we explain the various stages for generating a modulus $N = pq$ with an unknown factorization. In Section 6 we describe protocols that given a public exponent e, generate shares of the corresponding private key $d = e^{-1} \bmod \varphi(N)$. To achieve fault tolerance one often shares d using a t-out-of-k scheme. That is, any subset of t users can apply the private key d. We explain how a t-out-of-k sharing of d is obtained in Section 6.3. We discuss various practical improvements of our protocol in Section 7.

We note that generic secure circuit evaluation techniques, e.g. [38, 26, 3, 9] can also be used to generate shared RSA keys. After all, a primality test can be represented as a boolean circuit. However, such general techniques are too inefficient.

1.1 Communication and privacy model

The communication and privacy model assumed by our protocol are as follows:

Full connectivity Any party can communicate with any other party.

Private and authenticated channels Messages sent from party A to party B are private and cannot be tampered with en route. This can be achieved by having A and B share a secret key which they use for encryption, integrity and authentication.

Honest parties We assume all parties are honestly following the protocol. At the end of the protocol no threshold of parties have enough information to factor the generated modulus. This is often called the *honest but curious* scenario. In Section 8 we discuss some results showing how this assumption can be relaxed. Recently, Frankel, MacKenzie and Yung [22] showed how our protocol can be made robust against $\lfloor \frac{k-1}{2} \rfloor$ malicious parties.

Collusion Our protocol is $\lfloor \frac{k-1}{2} \rfloor$ private. That is, any coalition of size at most $\lfloor \frac{k-1}{2} \rfloor$ learns no information about the factorization of $N = pq$. However, a coalition of more than $\lfloor \frac{k-1}{2} \rfloor$ can recover the factorization. The reason for this bound is our reliance on the BGW [3] protocol in one step of our algorithm. The BGW protocol provides information theoretic security and hence is limited to achieving $\lfloor \frac{k-1}{2} \rfloor$ privacy (in the honest parties model). We note that it is possible to achieve $k-1$ privacy by replacing the BGW step with a heuristic protocol due to Cocks [10]. Unfortunately, Cocks' protocol is by far slower than the BGW method. Furthermore, its security is based on a heuristic argument.

To prove the security properties (i.e. privacy) satisfied by our protocol we provide a simulation argument for each of its components.

2 Overview

In this section we give a high level overview of the protocol. The k parties wish to generate a shared RSA key. That is, they wish to generate an RSA modulus $N = pq$ and a public/private pair of exponents e, d where $e \cdot d = 1 \bmod \varphi(N)$. The factors p and q should be at least n bits each. At the end of the computation N and e are public, and d is shared between the k parties in a way that enables threshold decryption. All parties should be convinced that N is indeed a product of two primes, but no coalition of at most $t = \lfloor \frac{k-1}{2} \rfloor$ parties should have any information about the factors of N.

At a high level the protocol works as follows:

(1) pick candidates: The following two steps are repeated twice.

- **(a) secret choice:** Each party i picks a secret n-bit integer p_i and keeps it secret.
- **(b) trial division:** Using a private distributed computation the k parties determine that $p = p_1 + \ldots + p_k$ is not divisible by any prime less than some bound B_1. Details are given in Section 5. If this step fails repeat step (a).

Denote the secret values picked at the first iteration by $p_1, \ldots, p_k$, and at the second iteration by $q_1, \ldots, q_k$.

(2) compute N: Using a private distributed computation the k parties compute

$$N = (p_1 + \ldots + p_k) \cdot (q_1 + \ldots + q_k)$$

Other than the value of N, this step reveals no further information about the secret values $p_1, \ldots, p_k$ and $q_1, \ldots, q_k$. Details are given in Section 4.

Now that N is public, the k parties can perform further trial divisions and test that N is not divisible by small primes in the range $[B_1, B_2]$ for some bound B_2.

(3) primality test: The k parties engage in a private distributed computation to test that N is indeed the product of two primes. If the test fails, then the protocol is restarted from step 1. We note that the primality test protocol is $k-1$ private and applies whenever two (or more) parties are involved. Details are in section 3.

(4) key generation: Given a public encryption exponent e, the parties engage in a private distributed computation to generate a shared secret decryption exponent d. Details are in Section 6.

Notation Throughout the paper we adhere to the following notation: the RSA modulus is denoted by N and is a product of two n bit primes p, q. When $p = \sum p_i$ we denote by p_i the share in possession of party i. Similarly for q_i. When the p_i's themselves are shared among the parties we denote by $p_{i,j}$ the share of p_i that is sent to party j.

Performance issues Our protocol generates two random numbers and tests that $N = pq$ is a product of two primes. By the prime number theorem the probability that both p and q are prime is asymptotically $1/n^2$. Therefore, naively one has to perform n^2 probes on average until a suitable N

is found. This is somewhat worse than the expected $2n$ probes needed in traditional generation of an RSA modulus (one first generates one prime using n probes and then a second prime using another n probes). This $n/2$ degradation in performance is usually unacceptable (typically $n = 512$).

Fortunately, thanks to trial division things aren't so bad. Our trial division (Step 1b) tests each prime individually. Therefore, to analyze our protocol we must analyze the effectiveness of trial division. Suppose a random n-bit number p passes the trial division test where all primes less than B are tested. We take $B = c \cdot n$ for some constant c. How likely is p to be prime? Using a classic result due to Mertens, DeBruijn [12] shows that asymptotically

$$\Pr[p \text{ prime} \mid \text{trial division up to } B] = e^{\gamma} \frac{\ln B}{\ln 2^n}(1 + o(1/n)) = 2.57 \frac{\ln B}{n}(1 + o(1/n))$$

Hence, when $n = 512$ bits and $\ln B = 9$ (i.e. $B = 8103$) the probability that p is prime is approximately $1/22$. Consequently, traditional RSA modulus generation requires 44 probes while our protocol requires 484 probes. This eleven fold degradation in performance is unfortunate, but manageable. We discuss methods to avoid this slowdown in Section 7.

Generation of shares In step (1) of the protocol each party uniformly picks a random n bit integer p_i as its secret share. The prime p is taken to be the sum of these shares. Since the sum of uniform independent random variables is *not* uniformly distributed, p is picked from a distribution with slightly less entropy than uniform. We show that this is not a problem. The sum $p = \sum_i p_i$ is at most an $n + \log k$ bit number. One can easily show that p is chosen from a distribution with at least n bits of entropy (since the n least significant bits of p are a uniformly chosen n bit string). These $\log k$ bits of "lost" entropy can not help an adversary, since they can be easily guessed (the number of parties k is small, certainly $k < n$). This is formally stated in the next lemma.

A second issue is the fact that the shares p_i themselves leak some information about the factors of N. For instance, party i knows that $p > p_i$. We argue that this information does not help an adversary either.

The two issues raised above are dealt with in the following lemma. Let $\mathbb{Z}_n^{(2)}$ be the set of RSA moduli $N = pq$ that can be output by our protocol above when k parties are involved. We assume $k < \log N$.

Lemma 2.1 *Suppose there exists a polynomial time algorithm $\mathcal{A}$ that given a random $N \in \mathbb{Z}_n^{(2)}$ chosen from the distribution above and the shares $\langle p_i, q_i \rangle$ of $k-1$ parties, factors N with probability at least $1/n^d$. Then there exists an expected polynomial time algorithm $\mathcal{B}$ that factors $1/4k^3n^d$ of the integers in $\mathbb{Z}_n^{(2)}$.*

Assuming the hardness of factoring, the lemma shows that even an adversary who is given N and the private shares of $k-1$ parties cannot factor the modulus N generated by the protocol. The proof of the lemma is somewhat tedious and is given in Appendix B so as not to distract the reader from the main thrust of the paper.

3 Distributed primality test

We begin the detailed discussion of the protocol with the distributed primality test (step 3). Party i has two secret n-bit integers p_i, q_i. All parties know N where $N = pq = (\sum p_i)(\sum q_i)$. They wish to determine if N is the product of two primes without revealing any information about the factors

of N. We refer to this test as a distributed primality test. Our primality test is a probabilistic test [36, 33] carried out in both $\mathbb{Z}_N^*$ and a quadratic extension of $\mathbb{Z}_N^*$.

Throughout the section we are assuming that $p = q = 3 \bmod 4$ (hence the resulting $N = pq$ is a Blum integer). This can be arranged ahead of time by having party 1 pick shares $p_1 = q_1 = 3 \bmod 4$. All other parties pick shares $p_i = q_i = 0 \bmod 4$.

Before describing the test we briefly discuss the structure of the quadratic extension of $\mathbb{Z}_N^*$ we will be using. We will be working in the twisted group $\mathbb{T}_N = (\mathbb{Z}_N[x]/(x^2+1))^*/\mathbb{Z}_N^*$. Suppose all prime factors of N are 3 mod 4. In this case, x^2+1 is irreducible in $\mathbb{Z}_N[x]$ and $\mathbb{Z}_N[x]/(x^2+1)$ is a quadratic extension of $\mathbb{Z}_N$. A linear polynomial $f(x) = \alpha x + \beta$ in $\mathbb{Z}_N[x]/(x^2+1)$ is invertible if and only if $\gcd(\alpha, \beta, N) = 1$. It follows that elements of $\mathbb{T}_N$ can be viewed as linear polynomials $f(x) = \alpha x + \beta \in \mathbb{Z}_N[x]$ with $\gcd(\alpha, \beta, N) = 1$. Two linear polynomials $f, g \in \mathbb{Z}_N[x]$ represent the same element of $\mathbb{T}_N$ if $f = ag$ for some $a \in \mathbb{Z}_N^*$. We note that elements of $\mathbb{T}_N$ can also be viewed as points on the projective line over $\mathbb{Z}_N$.

Distributed primality test:

Step 1: The parties agree on a random $g \in \mathbb{Z}_N^*$. The value g is known to all k parties.

Step 2: Party 1 computes the Jacobi symbol of g over N. If $\left(\frac{g}{N}\right) \neq 1$ the protocol is restarted at step (1) and a new random g is chosen.

Step 3: Otherwise, party 1 computes $v_1 = g^{(N-p_1-q_1+1)/4} \bmod N$. All other parties compute $v_i = g^{(p_i+q_i)/4} \bmod N$. The parties exchange the v_i values with each other and verify that

$$v_1 = \pm \prod_{i=2}^{k} v_i \pmod{N}$$

If the test fails then the parties declare that N is not a product of two primes.

Step 4: The parties perform a Fermat test in the twisted group $\mathbb{T}_N = (\mathbb{Z}_N[x]/(x^2+1))^*/\mathbb{Z}_N^*$. To carry out the Fermat test in $\mathbb{T}_N$ the parties pick a random $h \in \mathbb{T}_N$. Party 1 computes $u_1 = h^{N+p_1+q_1+1}$. All other parties compute $u_i = h^{p_i+q_i}$. The parties then exchange the u_i values with each other and verify that

$$\prod_{i=1}^{k} u_i = h^{N+p+q+1} = 1$$

If the test fails N is rejected. Otherwise they declare success.

Since $p_1 = q_1 = 3 \bmod 4$ and $p_i = q_i = 0 \bmod 4$ for $i > 1$, the exponents in the computation of the v_i's (Step 3) are guaranteed to be integers after division by 4. The correctness and privacy of the protocol is proved in the next two lemmas.

Lemma 3.1 *Let $N = pq$ be an integer with $p \equiv q \equiv 3 \bmod 4$. If N is a product of two distinct primes then success is declared in all invocations of the protocol. Otherwise, the parties declare that N is not a product of two primes with probability at least $\frac{1}{2}$ (over the random choice of g and h).*

Proof Observe that in step (3) of the protocol we test that $v_1 / \prod v_i = \pm 1 \bmod N$ which amounts to testing $g^{(N-p-q+1)/4} \equiv \pm 1 \bmod N$.

Suppose p and q are distinct primes. In step (2) we verify that $\left(\frac{g}{N}\right) = 1$. This implies $\left(\frac{g}{p}\right) = \left(\frac{g}{q}\right)$. Also, since $\frac{q-1}{2}$ and $\frac{p-1}{2}$ are odd we have,

$$\begin{cases} g^{\phi(N)/4} = \left(g^{\frac{p-1}{2}}\right)^{\frac{q-1}{2}} \equiv \left(\frac{g}{p}\right)^{\frac{q-1}{2}} = \left(\frac{g}{p}\right) \pmod{p} \\ g^{\phi(N)/4} = \left(g^{\frac{q-1}{2}}\right)^{\frac{p-1}{2}} \equiv \left(\frac{g}{q}\right)^{\frac{p-1}{2}} = \left(\frac{g}{q}\right) \pmod{q} \end{cases}$$

Since $\left(\frac{g}{p}\right) = \left(\frac{g}{q}\right)$ it follows that $g^{\phi(N)/4} \equiv \pm 1 \bmod N$. Since $\phi(N) = N - p - q + 1$ when p and q are prime, it follows that the test in Step 3 always succeeds.

Similarly, we show that when p and q are distinct primes the test in Step 4 succeeds. Since $p \equiv q \equiv 3 \bmod 4$ the polynomial $x^2 + 1$ has no root in $\mathbb{F}_p$ and $\mathbb{F}_q$. Therefore, $\mathbb{F}_p[x]/(x^2+1)$ and $\mathbb{F}_q[x]/(x^2+1)$ are quadratic extensions of $\mathbb{F}_p$ and $\mathbb{F}_q$ respectively. It follows that the group $\mathbb{T}_p = (\mathbb{Z}_p[x]/(x^2+1))^*/\mathbb{Z}_p^*$ is of order $p+1$. Similarly, $|\mathbb{T}_q| = q+1$. By the Chinese Remainder Theorem, $\mathbb{T}_N = \mathbb{T}_p \times \mathbb{T}_q$ and hence $|\mathbb{T}_N| = (p+1)(q+1)$. We conclude that all $h \in \mathbb{T}_N$ satisfy $h^{(p+1)(q+1)} = 1$. Consequently, the test of Step 4 always succeeds.

To prove the converse, suppose at least one of p, q is not prime. That is, $N = r_1^{d_1} \cdots r_s^{d_s}$ is a non-trivial factorization of N with $\sum d_i \geq 3$ and $s \geq 1$. Set $e = (N-p-q+1)/4 = (p-1)(q-1)/4$ to be the exponent used in step (3). Note that e is odd since $p \equiv q \equiv 3 \bmod 4$. Define the following two subgroups of $\mathbb{Z}_N^*$:

$$G = \{g \in \mathbb{Z}_N \text{ s.t. } \left(\frac{g}{n}\right) = 1\} \quad \text{and} \quad H = \{g \in G \text{ s.t. } g^e = \pm 1 \bmod N\}$$

To prove the lemma we show that $|H| \leq \frac{1}{2}|G|$. Since H is a subgroup of G it suffices to prove proper containment of H in G, i.e. prove the existence of $g \in G \setminus H$. There are four cases to consider.

Case 1. Suppose $s \geq 3$. Let r_2, r_3 be distinct prime factors of N such that $d_2 = d_3 \bmod 2$ (such a pair must exist by the pigeonhole principle). Let r_1 be a prime factor of N distinct from both r_2 and r_3 and let a be a quadratic non-residue modulo r_3. Define $g \in \mathbb{Z}_N$ to be an element satisfying

$$g \equiv 1 \bmod r_1 \quad \text{and} \quad g \equiv -1 \bmod r_2 \quad \text{and} \quad g \equiv \begin{cases} 1 \bmod r_3 & \text{if } \left(\frac{-1}{r_2}\right) = 1 \\ a \bmod r_3 & \text{if } \left(\frac{-1}{r_2}\right) = -1 \end{cases}$$

and $g \equiv 1 \bmod r_i$ for $i > 3$. Observe that $g \in G$. Since e is odd $g^e = g = 1 \bmod r_1$ and $g^e = g = -1 \bmod r_2$. Consequently, $g^e \neq \pm 1 \bmod N$ i.e. $g \notin H$.

Case 2. Suppose $\gcd(p, q) > 1$. Then there exists an odd prime r such that r divides both p and q. Then r^2 divides N implying that r divides $\phi(N)$. It follows that in $\mathbb{Z}_N^*$ there exists an element g of order r. Since r is odd we have $\left(\frac{g}{N}\right) = \left(\frac{g^r}{N}\right) = \left(\frac{1}{N}\right) = 1$, i.e. $g \in G$. Since r divides both p and q we know that r does not divide $N - p - q + 1 = 4e$. Consequently $g^{4e} \neq 1 \bmod N$ implying that $g^e \neq \pm 1 \bmod N$. Hence, $g \notin H$.

Case 3. The only way $N = pq$ does not fall into both cases above is if $p = r_1^{d_1}$ and $q = r_2^{d_2}$ where r_1, r_2 are distinct primes and at least one of d_1, d_2 is bigger than 1 (Case 2 handles N that are a prime power $N = r^d$). By symmetry we may assume $d_1 > 1$. Since $\mathbb{Z}_p^*$ is a cyclic

group of order $r_1^{d_1-1}(r_1-1)$ it contains an element of order $r_1^{d_1-1}$. It follows that $\mathbb{Z}_N^*$ also contains an element g of order $r_1^{d_1-1}$. As before, $\left(\frac{g}{N}\right) = 1$, i.e. $g \in G$. If $q \neq 1 \bmod r_1^{d_1-1}$ then $4e = N - p - q + 1$ is not divisible by $r_1^{d_1-1}$. Consequently, $g^{4e} \neq 1 \bmod N$, i.e. $g \notin H$.

Case 4. We are left with the case $N = pq$ with $p = r_1^{d_1}$, $q = r_2^{d_2}$, $d_1 > 1$ as above and $q \equiv 1 \bmod r_1^{d_1-1}$. Since $p = q = 3 \bmod 4$ we know that $r_1 = r_2 = 3 \bmod 4$. In this case it may indeed happen that $H = G$. We show that in this case, Step (4) of the primality test will fail with probability at least half (over the choice of $h \in \mathbb{T}_N$).

Define the group $H' = \{h \in \mathbb{T}_N \text{ s.t. } h^{(p+1)(q+1)} = 1\}$. We show that $|H'| \leq \frac{1}{2}|\mathbb{T}_N|$. Since H' is a subgroup of $\mathbb{T}_N$ it suffices to prove proper containment, i.e. we must exhibit an element $w \in \mathbb{T}_N \setminus H'$. Since $p = r_1^{d_1}$ the group $\mathbb{T}_p$ has order $r_1^{d_1-1}(r_1+1)$. Therefore it contains an element h of order r_1. It follows that there exists an element $w \in \mathbb{T}_N$ of order r_1. Since by assumption $q \equiv 1 \bmod r_1$ we know that r_1 does not divide $q+1$. Hence r_1 does not divide $N+p+q+1$ and therefore $w^{N+p+q+1} = w^{(p+1)(q+1)} \neq 1$. This proves $w \in \mathbb{T}_N \setminus H'$ completing the proof of the lemma. □

We note that Step (4) of the protocol is needed to filter out integers that fall into Case 4 above. Indeed, such integers will pass steps (1)-(3). For example, consider integers $N = pq$ with $p = 3^s$ and $q = 2 \cdot 3^{s-1} + 1$ with s odd and q prime. For these integers we have $H = G$ where H, G are as defined in the body of the proof above. Consequently, such integers always pass steps (1)-(3) even though they are not a product of two distinct primes. However, they will fail step (4). In Section 3.1 we give an alternate approach for filtering out integers that fall into Case 4.

The following lemma shows that when N is the product of two distinct primes, the primality test protocol reveals no other information about the private shares of the participants.

Lemma 3.2 *Suppose p, q are prime. Then any coalition of $k-1$ parties can simulate their view of the primality testing protocol. Consequently, the protocol is $k-1$ private.*

Proof Since p, q are prime we know that $v_1 = \pm \prod_{i=2}^k v_i \bmod N$ where the v_i's are defined as in step (3) of the protocol. Also, $\prod u_i = 1$ where the u_i's are defined in step (4). Let U be a coalition of $k-1$ parties. Say party m is not a member of the coalition. The coalition's view consists of $\langle u_i, v_i, N, g, h\rangle$ for $i = 1, \ldots, k$ and $\langle p_i, q_i\rangle$ for all $i \neq m$.

We construct a simulator of the coalition's view as follows: the simulator is given N, p_i, q_i for $i \neq m$ as input (this is the coalition's input when the protocol is started). To simulate the coalition's view, the simulator first picks a random $g \in \mathbb{Z}_N^*$ with $\left(\frac{g}{N}\right) = 1$ and a random $h \in \mathbb{T}_N$. It then computes the u_i, v_i values for all $i \neq m$ using the p_i, q_i values it is given as input. Next the simulator must generate u_m, v_m. Generating u_m is trivial since $u_m = (\prod_{i \neq m} u_i)^{-1}$. Generating v_m is a bit harder. From the proof of Lemma 3.1 it follows that

$$\begin{cases} v_1 = \prod_{i=2}^k v_i \pmod N & \text{implies} \quad \left(\frac{g}{p}\right) = \left(\frac{g}{q}\right) = 1 \\ v_1 = -\prod_{i=2}^k v_i \pmod N & \text{implies} \quad \left(\frac{g}{p}\right) = \left(\frac{g}{q}\right) = -1 \end{cases}$$

That is, the sign of $v_1 / \prod_{i=2}^k v_i$ leaks whether g is a quadratic residue modulo N. To simulate v_m the simulator computes $\hat{v}_m = v_1 / \prod_{i \neq m} v_i$. It then flips an unbiased coin and sets $v_m = \hat{v}_m$ or $v_m = -\hat{v}_m$ accordingly. The resulting distribution on v_m is computationally indistinguishable from the true distribution, assuming the hardness of quadratic residuosity modulo a Blum integer. To

conclude, given that N is the product of two distinct primes, the set of values $\langle u_i, v_i, N, g, h\rangle$ produced by the simulator is computationally indistinguishable from a real transcript. Consequently, the coalition learns no information other than what it had at the beginning of the protocol. □

We note that step (2) of the protocol is crucial. Without it the condition of step (3) might fail (and reveal the factorization) even when p and q are prime. We also note that in practice the probability that a non RSA modulus passes even one iteration of this test is actually far less than a half.

3.1 An alternative to Step (4)

Step (4) of the distributed primality protocol is necessary to filter out integers that fall into Case 4 in the proof of Lemma 3.1. We describe an alternative simpler approach to filter out such integers. It requires less computation, although there is more communication between the parties.

Observe that if N falls into Case 4 then $\gcd(N, p+q-1) > 1$. The alternative to Step (4) is to directly test this condition. To test this condition with no information leakage the parties do the following: each party picks a random $r_i \in \mathbb{Z}_N$ and keeps it secret. Using the protocol of the next section they compute

$$z = \left(\sum_{i=1}^{k} r_i\right)\left(-1 + \sum_{i=1}^{k}(p_i + q_i)\right) \pmod N$$

without leaking any information about the private shares p_i, q_i or r_i. Finally, the parties check if $\gcd(z, N) > 1$. If so, N is rejected. Using the BGW method (Section 4.1) this approach is $\lfloor \frac{k-1}{2} \rfloor$ private. All N that fall into case 4 are rejected. All that are a product of two distinct primes pass the test with overwhelming probability. We note that this alternate test eliminates a few valid RSA moduli, i.e. moduli $N = pq$ with p, q prime and $q = 1 \bmod p$.

4 Distributed computation of N

Next we describe the computation of N. Each party has a secret p_i, q_i. They wish to make the product $N = (\sum p_i)(\sum q_i)$ public without revealing any information about their private shares beyond what is revealed by the knowledge of N.

4.1 The BGW method

BenOr, Goldwasser and Wigderson [3] describe an elegant protocol for private evaluation of general functions for three or more parties. Their full technique is an overkill for the simple function we have in mind. We adapt their protocol to the computation at hand so as to minimize the amount of computation and communication between the parties. From here on, let $P > N$ be some prime. Unless otherwise stated, all arithmetic operations are done modulo P. The protocol works as follows:

Step 1: Let $l = \lfloor \frac{k-1}{2} \rfloor$. For all $i = 1, \ldots, k$ party i picks two random degree l polynomials $f_i, g_i \in \mathbb{Z}_P[x]$ satisfying $f_i(0) = p_i$ and $g_i(0) = q_i$. In other words, the constant term of f_i, g_i are set to p_i, q_i and all other coefficients are chosen at random in $\mathbb{Z}_P$. Similarly, party i picks a random degree $2l$ polynomial $h_i \in \mathbb{Z}_p[x]$ satisfying $h_i(0) = 0$.

Step 2: For all $i = 1, \ldots, k$ party i computes the $3k$ values:

$$\forall j = 1, \ldots, k: \quad p_{i,j} = f_i(j) \quad ; \quad q_{i,j} = g_i(j) \quad ; \quad h_{i,j} = h_i(j)$$

Party i then privately sends the triple $\langle p_{i,j}, q_{i,j}, h_{i,j} \rangle$ to party j for all $j \neq i$. Note that the $p_{i,j}$ for $j = 1, \ldots, k$ are standard l-out-of-k Shamir secret sharings of p_i. The same holds for q_i.

Step 3: At this point, each party i has all of $\langle p_{j,i}, q_{j,i}, h_{j,i} \rangle$ for $j = 1, \ldots, k$. Party i computes:

$$N_i = \left(\sum_{j=1}^{k} p_{j,i} \right) \cdot \left(\sum_{j=1}^{k} q_{j,i} \right) + \sum_{j=1}^{k} h_{j,i} \pmod{P}$$

Party i broadcasts N_i to all other parties.

Step 4: At this point each of the parties has all values N_i for $i = 1, \ldots, k$. Let $\alpha(x)$ be the polynomial

$$\alpha(x) = \left(\sum_{j} f_j(x) \right) \left(\sum_{j} g_j(x) \right) + \sum_{j} h_j(x) \pmod{P}$$

Observe that $\alpha(i) = N_i$ and by definition of f_i, g_i, h_i we have $\alpha(0) = N$. Furthermore, $\alpha(x)$ is a polynomial of degree $2l$. We note that l is defined so that $k \geq 2l+1$. Consequently, since all parties have at least $2l+1$ points on $\alpha(x)$ they can interpolate it and discover its coefficients. Finally, each party evaluates $\alpha(0)$ and obtains $N \bmod P$. Since $N < P$ the parties learn the correct value of N.

From the description of the protocol it is clear that all parties learn the value N. We note that the protocol requires that at least three parties be involved. In the case of exactly three parties, linear polynomials are used and the protocol is 1-private.

The following lemma shows that no coalition of $\lfloor \frac{k-1}{2} \rfloor$ parties learns any further information about the private shares. This statement holds in the information theoretic sense – no complexity assumptions are needed. For completeness we sketch the proof of the lemma and refer to [3] for the complete details.

Lemma 4.1 *Given N, any coalition of $\lfloor \frac{k-1}{2} \rfloor$ parties can simulate the transcript of the protocol. Consequently, the protocol is $\lfloor \frac{k-1}{2} \rfloor$ private.*

Proof Sketch Set $l = \lfloor \frac{k-1}{2} \rfloor$. By symmetry we may assume the coalition is made up of parties $1 \ldots l$. In what follows we consistently let the index i vary in $i = 1, \ldots, l$, the index j vary in $j = 1, \ldots, k$ and the index r vary in $r = l+1, \ldots, k$. Then the coalition's view consists of $\langle p_i, q_i, f_i, g_i, p_{i,j}, q_{i,j}, h_{i,j}, p_{r,i}, q_{r,i}, h_{r,i}, N_j, N \rangle$.

To simulate the coalition's view, the simulator is given p_i, q_i, N as input. It first picks random $f_i, g_i, h_i \in \mathbb{Z}_P[x]$ with $f_i(0) = p_i$, $g_i(0) = q_i$ and $h_i(0) = 0$. Next, it computes $p_{i,j} = f_i(j)$, $q_{i,j} = g_i(j)$ and $h_{i,j} = h_i(j)$. It picks $p_{r,i}, q_{r,i}, h_{r,i}$ as random independent elements of $\mathbb{Z}_P$ and computes $N_i = (\sum_j p_{i,j})(\sum_j q_{i,j}) + \sum_j h_{i,j}$. Next, it picks a random degree $2l$ polynomial $\alpha(x) \in \mathbb{Z}_P[x]$ satisfying $\alpha(i) = N_i$ and $\alpha(0) = N$. It completes the simulation by setting $N_r = \alpha(r)$. These values are a perfect simulation of the coalition's view. $\square$

The above protocol consists of one phase of the full BGW method. It differs from the BGW protocol in that there is no need for a truncation step. Also, we combine the addition and multiplication stages into one phase. The resulting computation is surprisingly efficient. Essentially, there is only *one* multi-precision multiplication performed by each party (the one in Step 3). We note that communication between the parties can be reduced by a factor of two using a variant of the protocol as described in [6, Section 4].

4.2 BGW Modulo non primes

In our description of the BGW protocol all arithmetic operations were carried out modulo a prime $P > N$. In the coming sections it will be useful to run the BGW protocol while working modulo a non prime P. That is, the k parties wish to compute $N = (\sum p_i)(\sum q_i) \bmod M$ for some integer M, not necessarily prime. One can easily show that if M has no prime divisors smaller than k then the protocol can be used *as is*. Indeed, all Lagrange coefficients used during the interpolation at Step (4) exist. Lemma 4.1 remains correct.

Running the protocol modulo M containing small factors require a slight modification. Write $M = M_1 \cdot M_2$ where M_1 has no prime factors smaller than k and M_2 has only prime factors smaller than k. As mentioned above the protocol immediately works modulo M_1. The problem when working modulo M_2 is that Shamir secret sharing (which is the basis of BGW) is not possible. For instance, consider the case when $M_2 = 3$. It is not possible to use Shamir secret sharing in $\mathbb{F}_3$ among $k > 3$ parties since $\mathbb{F}_3$ does not contain enough points (in Shamir secret sharing each party must be given its own unique point). A simple solution is to run the entire protocol in an algebraic extension of $\mathbb{F}_3$ that contains more than k points. The simulation argument immediately extends to this case. In general, one can factor M_2 into its prime factors and run the protocol in a large enough extension for each factor. Using the Chinese Remainder Theorem one can recover the value of $N \bmod M$.

4.3 Sharing the final outcome

In some cases (as in Section 6.2) we wish to have the parties evaluate the function $N = (\sum p_i)(\sum q_i)$; however the result should be additively shared among the parties rather than become publicly available. That is, at the end of the computation each party should have an M_i such that

$$\sum_{i=1}^{k} M_i = (\sum_{i=1}^{k} p_i)(\sum_{i=1}^{k} q_i) \pmod{P}$$

and no information is revealed about the private shares or the final result.

The modification to BGW in order to achieve the above goal is immediate. The parties do not perform Step 4 of the protocol and do not perform the broadcast described at the end of Step 3. Consequently, they each end up with a point on a polynomial $\alpha(x)$ of degree $2l$ that evaluates to N at $x = 0$. Using Lagrange interpolation we know that

$$N = \alpha(0) = \sum_{i=1}^{k} \lambda_i(0) N_i \pmod{P}$$

where $\lambda_i(x) = \prod_{j \neq i} \frac{(x-j)}{(i-j)}$ is the appropriate Lagrange coefficient. Therefore, rather than broadcast N_i at the end of Step 3, party i simply sets $M_i = \lambda_i(0) N_i$. The resulting M_i's are an additive sharing

of N as required. As before, there is a simple simulation argument showing that any minority of parties obtains no information from this protocol.

5 Trial division

In this section, we consider the trial division step (Step 1b in Section 2). Let $q = q_1 + \ldots + q_k$ be an integer shared among k parties. Let p be a small prime. To test if p divides q each party picks a random $r_i \in \mathbb{Z}_p$. Using the BGW protocol (as described in Section 4.1) they compute $qr = (\sum q_i)(\sum r_i) \bmod p$. If $qr \neq 0$ then p does not divide q. Furthermore, since r is unknown to any minority of parties, qr provides no other information about q.

Note that in the above approach, a bad candidate q is always rejected. Unfortunately, a good candidate might also be rejected. Indeed, even if p does not divide q it is still possible that $qr = 0 \bmod p$. To alleviate this problem one can repeat the above test twice for each small prime p. Then the probability that a good candidate is rejected is at most $1 - \prod_{p<B}(1 - \frac{1}{p^2}) < \frac{1}{2}$.

One caveat in the above approach is that the BGW protocol as described in Section 4.1 cannot be applied to test divisibility of q by very small p, namely $p \leq k$. The reason is that for such small p the field $\mathbb{F}_p$ does not contain enough points to do Shamir secret sharing among k parties. For such small p one must run the BGW protocol in an extension field of $\mathbb{F}_p$ that contains at least $k+1$ elements, as explained in Section 4.2.

6 Shared generation of public/private keys

Once the parties successfully construct an RSA modulus $N = pq = (\sum p_i)(\sum q_i)$ they may wish to compute shares of $d = e^{-1} \bmod \phi(N)$ for a given encryption exponent e. We have two approaches for computing shares of d. The first only works for small e (say $e < 1000$) but is very efficient requiring very little communication between the parties. The second works for any e and is still efficient, but requires more communication.

Throughout the section we set $\phi = \phi(N)$. Recall that the public modulus $N = (\sum p_i)(\sum q_i)$ satisfies $\phi(N) = \phi = \sum_{i=1}^{k} \phi_i$ where $\phi_1 = N - p_1 - q_1 + 1$ and $\phi_i = -p_i - q_i$ for $i = 2, \ldots, k$. For all $i = 1, \ldots, k$ party i can locally compute ϕ_i.

To compute shares of d the parties must invert e modulo $\sum \phi_i$ without exposing their ϕ_i's. Unfortunately, traditional inversion algorithms, e.g. extended gcd, involve computations modulo $\sum \phi_i$. We do not know how to efficiently perform modular arithmetic when the modulus is shared among the participants. Fortunately, there is a trick for computing $e^{-1} \bmod \phi$ with no modulo ϕ reductions. When only a single user is involved the inversion is done in three steps: (1) Compute $\zeta = \phi^{-1} \bmod e$. (2) Set $T = -\zeta \cdot \phi + 1$. Observe that $T \equiv 0 \bmod e$. (3) Set $d = T/e$. Indeed $d = e^{-1} \bmod \phi$ since $d \cdot e \equiv 1 \bmod \phi$. There is no need for modulo ϕ reductions since $e^{-1} \bmod \phi$ can be immediately deduced from $\phi^{-1} \bmod e$. Both methods rely on this observation.

6.1 Small public exponent

We begin by describing an efficient technique for generating shares of d when the public exponent e is small. The method leaks the value of $\phi(N) \bmod e$ (hence it can only be applied when e is small). On the plus side, it is $k-1$ private.

Step 1: The parties jointly determine the value of $l = \phi \bmod e$. Since $l = \sum \phi_i \bmod e$ it is possible to compute l without revealing any other information about the private shares. To do so we use a simple protocol due to Benaloh [4] which is $k-1$ private.

Step 2: Let $\zeta = l^{-1} \bmod e$. Then, as explained above, $d = (-\zeta \cdot \phi + 1)/e$. Each party i locally computes:

$$d_i = \left\lfloor \frac{-\zeta \cdot \phi_i}{e} \right\rfloor$$

As a result we have $d = \sum d_i + r \bmod \phi(N)$ where $0 \leq r < k$.

Step 3: The above sharing of d enables shared decryption [20] using the equality $c^d \equiv c^r \prod c^{d_i} \bmod N$. Party 1 can determine the value of r by trying all possible values of $0 \leq r \leq k$ during a trial decryption.

The above approach leaks $\phi(N) \bmod e$ and r. This is a total of $\log e + \log k$ bits.

6.2 Arbitrary public exponent

Unlike the previous technique, our second method for generating shares of d works for arbitrary public exponent e. It leaks less than $2 \log k$ bits. This information cannot help an opponent since it can be easily guessed. Rather than exposing $\phi \bmod e$ and then inverting it (as we did above) we show how to invert $\phi \bmod e$ while it is shared among the parties. As a result, no information about ϕ is revealed. The protocol is $\frac{k-1}{2}$ private.

Step 1. Each party picks a random $r_i \in \mathbb{Z}_e$.

Step 2. Using the protocol of Section 4.2 they compute $\psi = (\sum r_i) \cdot (\sum \phi_i) \bmod e$. At the end of the computation ψ is known to all parties. If ψ is not invertible modulo e the protocol is restarted at Step 1.

Step 3. Each party locally computes $\zeta_i = r_i \psi^{-1} \bmod e$. Observe $\sum \zeta_i = (\sum r_i)\psi^{-1} = \phi^{-1} \bmod e$. Hence, the parties are able to share the inverse of $\phi \bmod e$ without revealing any information.

Step 4. Next, the parties agree on a prime $P > 2Ne$. They view the shares $0 \leq \zeta_i < e$ as elements of $\mathbb{Z}_P$. Using the modified BGW protocol of Section 4.3 they compute an additive sharing

$$\sum_{i=1}^{k} T_i = -(\sum \zeta_i)(\sum \phi_i) + 1 \pmod P$$

Each party has a T_i and any minority of parties learns no other information.

Step 5. From here on we regard the T_i's as integers $0 \leq T_i < P$. Our objective is to ensure that over the integers

$$\sum_{i=1}^{k} T_i = -(\sum_{i=1}^{k} \zeta_i)(\sum_{i=1}^{k} \phi_i) + 1 \tag{1}$$

We know that at the end of Step 4 we have $\sum T_i < kP$. Therefore, $\sum T_i = sP + u$ where $s < k$ and $0 \leq u < P$. Given a candidate value of $s \in [0, k)$ party 1 can set $d_1 \leftarrow d_1 - sP$. If the given s is correct then $0 \leq \sum T_i < P$ and the above equality holds over the integers. To determine the correct s the protocol proceeds to Step 6 with each possible value of s until the trial decryption in Step 6 succeeds.

Step 6. Assuming equality (1) holds over the integers we know that e divides $\sum T_i$. To see this observe that:

$$\sum T_i = -(\sum \zeta_i)(\sum \phi_i) + 1 = -(\sum \phi_i)^{-1}(\sum \phi_i) + 1 \equiv 0 \pmod{e}$$

Therefore, $d = (\sum T_i)/e$. Each party i now sets $d_i = \lfloor T_i/e \rfloor$. As a result we have $d = \sum d_i + r \bmod \phi(N)$ where $0 \leq r \leq k$. Party 1 can determine the value of r by trying all possible values of $0 \leq r \leq k$ during a trial decryption.

The protocol leaks the value of r and s. Hence, a total of at most $2 \log k$ bits is exposed. The value of r is found using only one trial decryption: a gateway first picks a random message $m \in \mathbb{Z}_N$ and computes $c = m^e \bmod N$. It then asks the parties to decrypt c. Each party computes $m_i = c^{d_i} \bmod N$ and sends the result to the gateway. By performing at most k multiplications the gateway finds an $r \in [0, k)$ satisfying $m = (\prod m_i)c^r \bmod N$. It then relates the value of r back to the parties who fix their shares of d accordingly. To determine s this procedure is repeated for each of the k candidate values of s. Recall that k is typically small (e.g. less than 10).

Note 6.1: Observe that Step 5 of the above protocol is needed due to the fact that the BGW method is carried out modulo P. However, it is possible to carry out the BGW protocol of Step 4 directly over the integers, avoiding Step 5 altogether. This can be done using a variant of Shamir secret sharing over the integers (see [21, 22]). Let $L = k!$. Then, each party i picks two polynomials $f_i, g_i \in \mathbb{Z}[x]$ of degree $\lfloor \frac{k-1}{2} \rfloor$ as follows: the constant terms are set to $f_0 = f(0) = L\zeta_i$ and $g_0 = g(0) = L\phi_i$. All other coefficients are chosen at random from $\{0, L, 2L, \ldots, L^3N^2\}$. By interacting with the other parties, each party i computes a point on the polynomial $(\sum f_i)(\sum g_i)$ over the integers. These points, multiplied by the appropriate Lagrange coefficients, become the additive sharing $\sum T_i$ of $(\sum \zeta_i)(\sum \phi_i)$ over the integers. Unfortunately, in this approach, the resulting shares d_i of the private key are of order N^2 rather than N as in the above protocol. This results in a factor of two slowdown during threshold signature generation. On the positive side, there is no need to leak the value of s.

The computation of $\phi^{-1} \bmod e$ (steps 1-3 above) is based on a technique due to Beaver [2].

6.3 t-out-of-k sharing

The previous two subsections explain how one can obtain a k-out-of-k sharing of d. However, to provide fault-tolerance it is often desirable to have a t-out-of-k sharing enabling any subset of t parties to apply the private key. The simplest solution, due to T. Rabin [34], makes use of a generic technique for converting k-out-of-k sharing of a private RSA key into a t-out-of-k sharing scheme. T. Rabin's approach immediately applies once our k-out-of-k sharing of d is obtained.

7 Optimizations

We describe several practical techniques to improve the performance of our distributed protocol. These optimizations are incorporated in an implementation of our protocol [29].

Sieving In Step (1) of our protocol (Section 2) the parties repeatedly pick a random shared integer p until they find one that is not divisible by small primes. For this they engage in an interactive trial division protocol. We briefly outline a more efficient approach. Let M be the product of all odd primes less than some bound B_1. Suppose the k parties could generate an additive sharing $a = \sum a_i$ of a random integer a in $\mathbb{Z}_M^*$. Party i could then set its share of p to be $p_i = rM + a_i$ where r is some random number of the appropriate length to make p_i be an n-bit integer. The resulting candidate prime $p = \sum p_i$ is a random n-bit integer that is relatively prime to M. There is no need to run the trial division protocol on p with primes less than B_1. The only question that remains is how do the parties generate an additive sharing of a random integer $a \in \mathbb{Z}_M^*$? To do so, each party i generates a random element $b_i \in \mathbb{Z}_M^*$. Then $a = \prod b_i \bmod M$ is a random element of $\mathbb{Z}_M^*$. The parties then use the BGW method of Section 4.3 to convert the multiplicative sharing of a into an additive sharing. One caveat is that for this to work the BGW method must be made to work in $\mathbb{Z}_M$ (which is not a field). This is done as explained in Section 4.2.

Load balancing In Step (3) of the primality test party 1 computes $v_1 = g^{(N+1-p_1-q_1)/4} \bmod N$ while all other parties only had to compute $v_i = g^{(p_i+q_i)/4} \bmod N$. Notice that $N+1$ is roughly $2n$ bits while $p_i + q_i$ is only about n bits. Consequently, party 1 works twice as hard as the other parties. To even things out it makes sense to test k candidate N in parallel. In each of these tests the role of party 1 is played by a different party. This way $g^{N+1} \bmod N$ is computed by a different party for each N. This results in better load balancing improving over all performance by up to a factor of two.

Parallel trial division Recall that once N is computed the parties perform trial division on it before invoking the distributed primality test (Step 2 in Section 2). The k parties can perform this trial division in parallel – each party is in charge of verifying that N is not divisible by some set of small primes. This can be efficiently done by hard-coding all primes $p_1, p_2, \ldots, p_l$ in the range $[B_1, B_2]$ in a list. Party i is in charge of testing that N is not divisible by any primes p_j for all $j \equiv i \bmod k$. This factor of k speedup enables us to use a large bound B_2, increasing the effectiveness of trial division.

Fermat filter There is no need to run the full primality test of Section 3 on every candidate modulus N. Instead, one can do a Fermat test first and proceed to perform the full test only if the Fermat test succeeds. We refer to the test

$$g^{N-p-q+1} = g^{N+1-\sum(p_i+q_i)} = 1 \bmod N$$

as a Fermat test. If equality does not hold, then N is not a product of two primes. To carry out this test in zero-knowledge all k parties compute $v_i = g^{p_i+q_i} \bmod N$. They exchange v_i's and verify that $\prod v_i = g^{N+1} \bmod N$. This protocol is easily shown to be $k-1$ private. It saves the computation of the Jacobi symbol of g for most integers N.

Avoiding quadratic slowdown In Section 2 we noted that our protocols suffer from a quadratic slowdown in comparison to single user generation of an RSA key. The main reason is that

both primes p, q are generated at once. We briefly outline a potential solution. Consider the case of three parties, Alice, Bob and Carol. They could generate N as

$$N = p_a q_b (r_a + r_b + r_c)$$

where Alice has p_a, r_a and p_a is an n-bit prime. Bob has q_b, r_b and q_b is an n-bit prime. Carol has r_c. The number of probes until $r_a + r_b + r_c$ is found to be prime is just as in single user generation of N. Hence, we are able to avoid the quadratic slowdown. Furthermore, no single party knows the *complete* factorization of the resulting N. Unfortunately, this approach doesn't scale well. To enable t-privacy N must be a product of $t + 2$ primes. Also note that in the example above, the parties must perform a distributed primality test to verify that N is the product of *three* primes. The techniques of this paper do not easily generalize to enable such a test. We note that one can use recent results of [7] to do just that: test that a shared modulus is the product of three primes without revealing its factorization.

8 Robustness

Throughout the paper we use a model in which parties honestly follow the protocol. This is fine when parties are honestly trying to generate a shared RSA key. For some applications it is desirable to make the protocol robust against active adversaries that cheat during the protocol. Since the RSA function is verifiable (the parties can simply check that they correctly decrypt encrypted messages) active adversaries are limited in the amount of damage they can cause. However, it may still be possible that a party cheat during the protocol and consequently be able to factor the resulting N. Similarly, a party can cheat and cause a non-RSA modulus to be incorrectly accepted.

Recently, Frankel, MacKenzie and Yung [22] showed how our protocol can be made to withstand $\lfloor \frac{k-1}{2} \rfloor$ malicious parties. The protocol enables the parties to detect and exclude the malicious party. In practice, one could run our non-robust protocol until a modulus N is found which is believed to be a product of two primes. Then, the robust Frankel-MacKenzie-Yung protocol can be used to determine that no majority of parties cheated during the non-robust phase. For more results on robust generation of shared RSA keys see [5].

We describe a simple method for making our non-robust protocol robust when the number of participants is small (e.g. less than ten). Consider the case of four parties where at most one of them is malicious. One can run our non-robust protocol until a candidate modulus N is found. At this point the protocol is run four more times, once for each triplet of users. In the first run, party 1 shares her values p_1, q_1 with the other three parties by writing $p_1 = p'_2 + p'_3 + p'_4$ and $q_1 = q'_2 + q'_3 + q'_4$ where p'_j, q'_j are random integers in the range $[0, N]$. Party 1 then sends p'_i, q'_i to party i for $i = 2, 3, 4$. Party i adds these values to its own p_i, q_i. Next, parties $2, 3, 4$ run our non-robust protocol among the three of them (ignoring party 1). If the resulting N does not match the N computed when all four parties were involved, or if N turns out to not be an RSA modulus, the N is rejected and the parties announce that one of them is misbehaving. This experiment is repeated with all four triplets – each time exactly one party is excluded from the computation. Assuming at most one party is malicious, the resulting N must be a product of two large primes. Furthermore, the malicious party cannot know the factorization of N since at no point in the protocol does an honest party reveal any information about its share to another single party. We note that this approach enables the parties to detect cheating, but it does not help in detecting who the malicious party is.

In general, when k parties are engaged in our non-robust protocol, and c of them are malicious, the protocol can be made robust at the cost of $\binom{k}{c}$ extra invocations. The resulting computation is $\lfloor \frac{k-c-1}{2} \rfloor$ private. Clearly this approach can only be applied as long as both k and c are very small.

9 Summary and open problems

We presented techniques that allow three or more parties to generate an RSA modulus $N = pq$ such that all parties are convinced that N is indeed a product of two primes; however none of them can factor N. Our methods achieves $\lfloor \frac{k-1}{2} \rfloor$ privacy. We also show show how the parties can generate shares of a private decryption exponent to allow threshold decryption. To test that N is the product of two primes we presented a distributed double-primality test. We note that our test was recently extended to a triple-primality test [7] enabling k parties to test that N is the product of *three* primes, without revealing any information about the factors.

To demonstrate the effectiveness of our protocols we implemented them. Generating a 1024 bits shared RSA key among three 300MhZ Pentium machines takes 90 seconds. See [29] for a description of the implementation as well as detailed timing measurements.

An important open problem is the generation of shared keys of special form. For example, a modulus which is a product of "safe primes" (i.e., where both $\frac{p-1}{2}$ and $\frac{q-1}{2}$ are prime) has been considered for security purposes [27] as well as for technical reasons related to threshold cryptography [16, 25]. Currently, our techniques do not enable shared generation of moduli of special form. Progress in this directions would be very helpful.

Acknowledgments

We thank Don Beaver for helpful discussions on our results.

References

[1] N. Alon, Z. Galil and M. Yung, "Dynamic-resharing verifiable secret sharing," ESA 1995.

[2] D. Beaver, "Security, fault tolerance, and communication complexity in distributed systems," Ph.D. thesis, Harvard University, May 1990.

[3] M. Ben-Or, S. Goldwasser, A. Wigderson, "Completeness theorems for non-cryptographic fault tolerant distributed computation", STOC 1988, pp. 1–10.

[4] J. Benaloh (Cohen), "Secret sharing homomorphisms: keeping shares of a secret secret," Crypto '86, 251-260.

[5] S. Blackburn, S. Blake-Wilson, M. Burmester, S. Galbraith, "Shared generation of shared RSA keys", tech. report CORR 98-19, available at `http://www.cacr.math.uwaterloo.ca/`

[6] D. Boneh, M. Franklin, "Efficient generation of shared RSA keys", extended abstract in Proceedings Crypto' 97, pp. 425–439.

[7] D. Boneh, J. Horwitz, "Generating a product of three primes with an unknown factorization", Proc. of the third Algorithmic Number Theory Symposium (ANTS), 1998, pp. 237–251.

[8] J. Carter and M. Wegman, "Universal classes of hash functions", *J. Comput. Syst. Sci.* 18 (1979), 143–154.

[9] D. Chaum, C. Crépeau, and I. Damgård, "Multiparty unconditionally secure protocols," ACM STOC 1988, 11-19.

[10] C. Cocks, "Split knowledge generation of RSA parameters", Available from the author `cliff_cocks@cesg.gov.uk`.

[11] D. Coppersmith, "Private Communications".

[12] N. De Bruijn, "On the number of uncanceled elements in the sieve of Eratosthenes", Proc. Neder. Akad. Wetensch, vol. 53, 1950, pp. 803–812. Reviewed in LeVeque Reviews in Number Theory, Vol. 4, Section N-28, p. 221.

[13] A. DeSantis, Y. Desmedt, Y. Frankel, M. Yung, "How to share a function securely", STOC 1994, pp. 522–533.

[14] Y. Desmedt, "Threshold cryptography," European Transactions on Telecommunications and Related Technologies, Vol. 5, No. 4, July-August 1994, pp. 35–43.

[15] Y. Desmedt, G. Di Crescenzo, M. Burmester, "Multiplicative non-abelian sharing schemes and their applications to threshold cryptography", Asiacrypt '94, pp. 21–32.

[16] Y. Desmedt and Y. Frankel, "Shared generation of authenticators and signatures", Crypto '91, 457–469.

[17] T. ElGamal, "A public key cryptosystem and a signature scheme based on the discrete logarithm", IEEE Transactions on Information Theory, 31(4):469–472, 1985.

[18] U. Feige, A. Fiat, and A. Shamir, "Zero-knowledge proofs of identity," Journal of Cryptology 1 (1988), 77-94.

[19] A. Fiat and A. Shamir, "How to prove yourself: Practical solutions to identification and signature problems," Crypto '86, 186-194.

[20] Y. Frankel, "A practical protocol for large group oriented networks", Eurocrypt 89, pp. 56–61.

[21] Y. Frankel, P. MacKenzie, M. Yung, "Optimal resilience proactive proactive public key cryptosystems", FOCS 1997.

[22] Y. Frankel, P. MacKenzie, M. Yung, "Robust efficient distributed RSA key generation", STOC 1998, pp. 663–672.

[23] M. Franklin and S. Haber, "Joint encryption and message-efficient secure computation," Journal of Cryptology, 9 (1996), 217-232.

[24] P. Gemmel, "An introduction to threshold cryptography", in CryptoBytes, a technical newsletter of RSA Laboratories, Vol. 2, No. 7, 1997.

[25] R. Gennaro, S. Jarecki, H. Krawczyk, T. Rabin, "Robust and efficient sharing of RSA functions", Crypto 96, pp. 157–172.

[26] O. Goldreich, S. Micali, A. Wigderson, "How to play any mental game", STOC 1987, pp. 218–229.

[27] J. Gordon, "Strong primes are easy to find", Eurocrypt 84, pp. 216–223.

[28] L. Guillou and J. Quisquater, "A practical zero-knowledge protocol fitted to security microprocessor minimizing both transmission and memory," Eurocrypt '88, 123-128.

[29] M. Malkin, T. Wu, D. Boneh, "Experimenting with shared RSA key generation", In proceedings of the Internet Society's 1999 Symposium on Network and Distributed System Security (SNDSS).

[30] K. Ohta and T. Okamoto, "A modification of the Fiat-Shamir scheme," Crpto '88, 232-243.

[31] H. Ong and C. Schnorr, "Fast signature generation with a Fiat Shamir-like scheme," Eurocrypt '90, pp. 432-440.

[32] T. Pederson, "A threshold cryptosystem without a trusted party," Proceedings of Eurocrypt 91, pp. 522–526.

[33] M. Rabin, "Probabilistic algorithm for testing primality", J. of Number Theory, vol. 12, pp. 128–138, 1980.

[34] T. Rabin, "A simplified approach to threshold and proactive RSA", Proceedings of Crypto' 98.

[35] A. Shamir, "How to share a secret", Comm. of the ACM, Vol. 22, 1979, pp. 612–613.

[36] R. Solovay, V. Strassen, "A fast monte carlo test for primality", SIAM journal of computing, vol. 6, pp. 84–85, 1977.

[37] M. Wegman and J. Carter, "New hash functions and their use in authentication and set equality", *J. Comput. Syst. Sci.* 22 (1981), 265–279.

[38] A. Yao, "How to generate and exchange secrets", FOCS 1986, pp. 162–167.

A. Cocks' multiplication method

The BGW protocol described in Section 4.1 achieves $\lfloor \frac{k-1}{2} \rfloor$ privacy and is information theoretically secure. Cocks [10] describes a multiplication protocol that heuristically appears to provide $k-1$ privacy. Unfortunately, Cocks' protocol is far slower than the BGW method. Furthermore, it does not seem possible to prove its security using "natural" complexity assumptions. For completeness we describe a generalization of Cocks' protocol. As in the settings of Section 4.1 each of the k parties has a private n-bit p_i, q_i. The parties wish to compute $N = (\sum p_i)(\sum q_i)$ without revealing any other information about their private shares.

Let E, D be a weakly-multiplicative homomorphic public key encryption scheme. $E_p[m]$ encrypts a message m using the public key p, while $D_s(c)$ decrypts a ciphertext c using the secret key s. The homomorphic property implies that given a ciphertext $c_1 = E_p[m_1]$ and a plaintext m_2 there is an efficient algorithm to compute $c = E_p[m_1 \cdot m_2]$. Here we are assuming the plaintexts m_1, m_2 lie in a ring R where $R = \mathbb{Z}/(M\mathbb{Z})$ for some integer M. Examples of multiplicatively homomorphic public key encryption schemes are RSA and ElGamal encryption [17]. In what follows we are assuming each of the participants has its own public/private key pair. We denote by E_i, D_i the pair belonging to party i. All parties can encrypt using E_i while only party i can decrypt using D_i. Also, we assume the plaintext ring R is larger than $2N$.

Step 1: Each party i picks $k-1$ random integers $\langle a_{i,1}, \ldots, a_{i,i-1}, a_{i,i+1}, \ldots, a_{i,k} \rangle$ in $[0, \frac{|R|}{2}]$ and sets $b_i = \sum_{j \neq i} a_{i,j}$.

Step 2: Each party i sends $\langle E_i[p_i], E_i[q_i] \rangle$ to all other parties.

Step 3: For all $j \neq i$ party i does the following:

a. Using the homomorphic property it computes the three values

$$\alpha_{i,j} = E_j[p_i q_j] \ , \quad \beta_{i,j} = E_j[q_i p_j] \ , \quad \gamma_{i,j} = E_j[a_{i,j}]$$

b. Each of $\alpha_{i,j}, \beta_{i,j}, \gamma i, j$ is then broken into l shares, where l is a security parameter. The shares of $\alpha_{i,j}$ are constructed by picking random $u_1, \ldots, u_l \in R$ such that $\sum u_i = 1$. Then $\alpha_{i,j,1}, \ldots, \alpha_{i,j,l}$ are defined as: $\alpha_{i,j,m} = E_j[u_m p_i q_j]$. The $\beta_{i,j}$ and $\gamma_{i,j}$ are broken into l pieces in a similar way.

c. Party i sends the $3l$ values $\langle \alpha_{i,j,m}, \beta_{i,j,m}, \gamma_{i,j,m} \rangle$ to party j in a *random order*. The value l should be chosen large enough so that the random order cannot be undone. This is the main reason this protocol is so much slower than the BGW method.

Step 4: For all $j \neq i$ party i receives $3l$ ciphertexts from party j. It decrypts all of them and computes their sum. The sum comes out to be

$$S_{i,j} = p_i q_j + q_j p_i + a_{j,i}$$

Next, party i computes:

$$N_i = 2 p_i q_i - b_i + \sum_{j \neq i} S_{i,j}$$

A routine calculation shows that

$$\sum_{i=1}^{k} N_i = 2 \sum_{i,j} p_i q_j = 2 \left(\sum p_i \right) \left(\sum q_i \right) = 2N$$

Hence, an additive sharing of $2N$ is obtained. To complete the protocol, the parties may exchange the values N_i with each other and make the final value of N public.

When generating a 1024 bit modulus, Cocks states that one should use l on the order of 400. This implies that each party must perform $1200 \cdot k$ public key decryptions at Step 4 (!) This is significantly slower than the BGW method (over 1000 times slower) where no encryption is used.

Although the protocol may provide $k-1$ privacy, we cannot give a simulation proof. Note that a coalition of $k-1$ parties learns both the value of N and the encryptions $E_i[p_i], E_i[q_i]$ of the party not in the coalition. Providing a simulation for such a transcript seems unlikely under "natural" assumptions. Indeed, Coppersmith [11] shows that if one uses RSA for E_i (i.e. $E_i(m) = m^{e_i} \bmod M_i$) then there is an attack that factors N in time polynomial in e. Consequently, one cannot use low public exponent RSA as E_i. Similarly, it is not clear how to simulate the random permutation used in Step 3.

Throughout the paper, Cocks' method can be used instead of BGW to compute the function $N = (\sum p_i)(\sum q_i)$. This substitution will make all our protocols heuristically $k-1$ private. However, we stick with the BGW protocol due to its superior performance and rigorous proof of security. Improving the performance of Cocks' method and providing a *rigorous* protocol achieving $k-1$ privacy are excellent open problems.

B. Proof of Lemma 2.1

Let $\mathbb{Z}_n^{(2)}$ be the set of RSA moduli $N = pq$ that can be output by our protocol when k parties are involved and each party picks two n bit shares p_i and q_i. We assume $k < \log N$. Lemma 2.1 shows that even an adversary who is given the private shares of $k-1$ of the parties cannot factor N (assuming the hardness of factoring a non-negligible fraction of the RSA moduli in $\mathbb{Z}_n^{(2)}$).

Lemma 2.1 *Suppose there exists a polynomial time algorithm $\mathcal{A}$ that given: (1) a random $N \in \mathbb{Z}_n^{(2)}$ chosen from the distribution on $\mathbb{Z}_n^{(2)}$ induced by the protocol, and (2) the shares $\langle p_i, q_i \rangle$ of $k-1$ parties, factors N with probability at least $1/n^d$. Then there exists an expected polynomial time algorithm $\mathcal{B}$ that factors $1/4k^3n^d$ of the integers in $\mathbb{Z}_n^{(2)}$.*

Proof Given $N \in \mathbb{Z}_n^{(2)}$ algorithm $\mathcal{B}$ works by repeating the following two steps until N is factored: (1) pick $2(k-1)$ random independent n-bit integers $\langle p_i, q_i \rangle_{i=1}^{k-1}$ (2) Run algorithm $\mathcal{A}$ on the input N and $\langle p_i, q_i \rangle_{i=1}^{k-1}$. We show that for $1/4k^3n^d$ of the integers in $\mathbb{Z}_n^{(2)}$, after an expected polynomial number of iterations algorithm $\mathcal{B}$ will output a factorization of N.

Our protocol for generating an RSA modulus $N = pq$ induces a distribution on $\mathbb{Z}_n^{(2)}$: an element N is generated by picking random n-bit numbers $\langle p_1, \ldots, p_k \rangle$ and $\langle q_1, \ldots, q_k \rangle$ and outputting the product of their sums. First, we slightly restrict the choice of p and q by disallowing values close to the boundaries, i.e. the largest and smallest p's and q's. The total weight of the elements we disallow is sufficiently small so as not to hurt $\mathcal{A}$'s probability of factoring the remaining N's. Then we show that on the remaining N's algorithm $\mathcal{B}$ succeeds in factoring essentially the same N's as algorithm $\mathcal{A}$.

Let $S_p = [a, b]$ be the shortest interval such that $\Pr[\sum_{i=1}^{k} p_i \in S_p] \geq 1 - \frac{1}{4n^d}$ where the distribution is over the choice of random n-bit integers $\langle p_1, \ldots, p_k \rangle$ for which $\sum p_i$ is prime. The set S_p contains all possible prime factors produced by the protocol except those close to the boundaries. Define $S \subseteq \mathbb{Z}_n^{(2)}$ as

$$S = \{p \cdot q \mid p, q \in S_p\}$$

Then when a random $N \in \mathbb{Z}_n^{(2)}$ is chosen by the protocol we have that

$$\Pr[N \in S] > (1 - \frac{1}{4n^d})^2 > 1 - \frac{1}{2n^d}$$

Hence, given a random $N \in S$ and $k-1$ of the shares used to generate p and q algorithm $\mathcal{A}$ will factor N with probability at least $\frac{1}{n^d} - \frac{1}{2n^d} = \frac{1}{2n^d}$. It follows that there exists a subset $S' \subseteq S$ with $\Pr[N \in S'] > 1/4n^d$ such that algorithm $\mathcal{A}$ succeeds in factoring all $N \in S'$ with probability at least $1/4n^d$ (over the choice of the p_i's and q_i's and the random bits used by $\mathcal{A}$). We show that algorithm $\mathcal{B}$ will factor all $N \in S'$ in expected polynomial time. Since the size of S is at least $1/k^2$ of $\mathbb{Z}_2^{(n)}$ and S' is at least $1/4n^d k$ of S we have that S' contains at least $1/4n^d k^3$ of the elements in $\mathbb{Z}_2^{(n)}$.

Consider a hypothetical algorithm $\mathcal{B}'$ that takes an integer $N = pq \in S'$ as input. Algorithm $\mathcal{B}'$ is identical to algorithm $\mathcal{B}$ except that it invokes algorithm $\mathcal{A}$ only if $\sum_{i=1}^{k-1} p_i$ is in the range $[p - 2^n, p - 2^{n-1}]$ and $\sum_{i=1}^{k-1} q_i$ is in the range $[q - 2^n, q - 2^{n-1}]$. Otherwise, algorithm $\mathcal{B}'$ restarts the iteration picking new random $\langle p_i, q_i \rangle$. Note that we are assuming algorithm $\mathcal{B}'$ knows p and q. Clearly if $\mathcal{B}'$ succeeds in causing algorithm $\mathcal{A}$ to factor N then so will $\mathcal{B}$ (the only difference between $\mathcal{B}$ and $\mathcal{B}'$ is that algorithm $\mathcal{B}'$ applies a filter prior to running $\mathcal{A}$ while $\mathcal{B}$ always runs $\mathcal{A}$). We show that for $N \in S'$ algorithm $\mathcal{B}'$ will cause $\mathcal{A}$ to succeed in expected polynomial time.

Given $N = pq \in S'$ define a distribution $\mathcal{P}_N$ on tuples of n-bit integers $\langle p_1, \ldots, p_{k-1} \rangle$ as follows: the distribution $\mathcal{P}_N$ is obtained by uniformly picking a k tuple $\langle p_1, \ldots, p_k \rangle$ among all the k-tuples that sum up to p and outputting $\langle p_1, \ldots, p_{k-1} \rangle$. Define $\mathcal{Q}_N$ similarly. For $N \in S'$ we know that

$$\Pr_{\mathcal{P}_N, \mathcal{Q}_N}[\mathcal{A} \text{ factors } N] \geq 1/4n^d$$

Next, observe that the distribution $(\mathcal{P}_N, \mathcal{Q}_N)$ is identical to the distribution that algorithm $\mathcal{B}'$ feeds $\mathcal{A}$. Hence, algorithm $\mathcal{B}'$ succeeds in factoring an $N \in S'$ in polynomial time if the filtering step does not take too long. It remains to show that for $N \in S'$ the filtering step does not take too long. Let p be an integer in the range $[k \cdot 2^{n-1}, k \cdot 2^n]$ (this is the range of all possible values of p produced by the protocol). The question is, what is the probability that a random tuple $\langle p_1, \ldots, p_{k-1} \rangle$ of n-bit integers satisfies that $\sum_{i=1}^{k-1} p_i$ is in the range $[p - 2^n, p - 2^{n-1}]$. The definition of S_p implies that for all $p \in S_p$ this probability is at least $1/k^2 n^d$. Hence, since the filtering step applies to both p and q, the expected number of tries until the filtering step is satisfied is $(k^2 n^d)^2$. Consequently, algorithm $\mathcal{B}$ runs in polynomial time (in k and n). $\square$