

EFFICIENT PARALLEL PSEUDORANDOM NUMBER GENERATION

J. H. REIF† AND J. D. TYGAR‡

Abstract. We present a parallel algorithm for pseudorandom number generation. Given a seed of n^ϵ truly random bits for any $\epsilon > 0$, our algorithm generates n^c pseudorandom bits for any $c > 1$. This takes poly-log time using $n^{\epsilon'}$ processors where $\epsilon' = k\epsilon$ for some fixed small constant $k > 1$. We show that the pseudorandom bits output by our algorithm cannot be distinguished from truly random bits in parallel poly-log time using a polynomial number of processors with probability $\frac{1}{2} + 1/n^{O(1)}$ if the Multiplicative Inverse Problem almost always cannot be solved in RNC. The proof is interesting and is quite different from previous proofs for sequential pseudorandom number generators.

Our generator is fast and its output is provably as effective for RNC algorithms as truly random bits. Our generator passes all the statistical tests in Knuth [14].

Moreover, the existence of our generator has a number of central consequences for complexity theory. Given a randomized parallel algorithm \mathcal{A} (over a wide class of machine models such as parallel RAMs and fixed connection networks) with time bound $T(n)$ and processor bound $P(n)$, we show that \mathcal{A} can be simulated by a parallel algorithm with time bound $T(n) + O((\log n)(\log \log n))$, processor bound $P(n)n^\epsilon$, and only using n^ϵ truly random bits for any $\epsilon > 0$.

Also, we show that if the Multiplicative Inverse Problem is almost always not in RNC, the RNC is within the class of languages accepted by uniform poly-log depth circuits with unbounded fan-in and strictly subexponential size $\bigcap_{\epsilon > 0} 2^{n^\epsilon}$.

Key words. cryptography, parallel algorithms, pseudorandom, number generators

AMS(MOS) subject classifications. 5.24, 5.27.

1. Introduction. A number of parallel randomized algorithms have appeared recently. These algorithms typically use a large number of random bits which must be generated in a small amount of time. Nonetheless, the area of parallel random bit generation remains unexplored.

In reality, our computers are deterministic and unable to generate truly random values. But we can give algorithms which will give pseudorandom bits on input of a random seed s_0 . These pseudorandom bits satisfy conditions which suggest that for algorithmic purposes they are as effective as truly random bits.

What conditions should a pseudorandom bit sequence satisfy?

Improving on an idea by Shamir [16], Blum and Micali [6] argue that the notion of "cryptographic strength" captures the important facets of random sequences. To demonstrate cryptographic strength they follow this schema:

- (1) Upper bound the computational resources by *Resources A*.
- (2) Assume that *Problem B* cannot be solved within the limits of *Resources A*.
- (3) Produce a *Pseudorandom Bit Generator G*.
- (4) Argue that if an opponent sees the first m_0 bits generated by *Pseudorandom Bit Generator G* and can utilize *Resources A* to predict the remaining bits with an

† This work was done at the Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139. Present address, Computer Science Department, Duke University, Durham, North Carolina 27706. The work of this author was supported in part by National Science Foundation grant NSF-MCS-79-21024 and Office of Naval Research contract N0014-80-C-0674.

‡ This work was done at the Aiken Computation Laboratory, Harvard University, Cambridge, Massachusetts 02138. Present address, Department of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213. The work of this author was supported in part by a National Science Foundation graduate fellowship and National Science Foundation grant MCS-81-21431.

accuracy rate of $\frac{1}{2} + \varepsilon(m)$ (where m is the size of the seed and ε is a fixed function satisfying $\lim_{m \rightarrow \infty} \varepsilon(m) = 0$), then the opponent will be able to solve *Problem B* limited to *Resources A* by consulting the bit-guessing oracle, a contradiction.

Several cryptographically-strong pseudorandom bit generators have been proposed (Blum, Blum and Shub [5], Blum and Micali [6],) and many applications have been discussed (Alexi, Chor, Goldreich and Schnorr [3], Goldreich, Goldwasser and Micali [9], Goldwasser, Micali and Tong [10], Vazirani and Vazirani [19], Yao [21].) These generators are all inherently sequential, require polynomial time, and their cryptographic strength relies on some unproven cryptographic assumption.

Notation. When we say a class of circuit is *uniform*, we mean that it is constructible in logarithmic space by a deterministic Turing Machine.

$\text{NC}(\text{NC}_U)$ is the class of languages accepted by (uniform, respectively) deterministic circuits constructible in log-space with poly-log depth and polynomial size.

$\text{RNC}(\text{RNC}_U)$ is the class of languages accepted by (uniform, respectively) randomized circuits constructible in log-space with two-sided error, poly-log depth, polynomial size, and acceptance probability greater than $\frac{1}{2}$.

We give more precise definitions of these terms in § 4.

Our result. We present a new cryptographically-strong pseudorandom bit generator which runs in NC_U but which is secure against attacks taking parallel poly-log time if the Multiplicative Inverse Problem almost always is not in RNC . While we use the schema described above for demonstrating the cryptographic strength of our random number generator, because of the inherent parallel nature of our generator, the technical details of our proof are quite different from those of previous proofs for sequential pseudorandom number generators. In particular, we prove that if the bits output by our pseudorandom bit generator can be predicted in RNC , then we can solve the multiplicative inverse problem in RNC almost always and this requires that we construct an interesting, nontrivial, parallel algorithm for that problem. (See § 3.)

About the assumption. While our assumption has not been proved, it is quite interesting to observe that it is *testable* in the following sense: If an RNC algorithm takes more than poly-log time using our pseudorandom bits instead of truly random bits then we can observe this event by timing. Thus one of two scenarios is possible: either every application of our generator to an RNC algorithm yields a poly-log algorithm using only a small number of random bits, or some application of our generator is discovered to exceed its poly-log time bounds and we can immediately derive a NC algorithm for multiplicative inverse.

About the measure of randomness. Valiant, Skyum, Berkowitz and Rackoff [18] show that an NC -machine can evaluate any straight-line program which computes a multivariate polynomial which has degree polynomial in the length of the program. Thus if our assumption is correct, our pseudorandom bit generator is secure against any statistical test which can be so formulated as a straight-line program. This includes most standard statistical tests for random number generators (Knuth [14]).

Applications. Our method for parallel pseudorandom bit generation is actually very practical. It requires, for any $\varepsilon > 0$, only $O(\log n(\log \log n))$ added depth and a factor of n^ε for a bounded fan-in circuit. Here is an example: Karp and Wigderson [12] give a deterministic algorithm for the maximal independent set problem in $O((\log n)^4)$ time using $O(n^3/(\log n)^3)$ processors. They also give a uniform randomized algorithm for the same problem running in $O((\log n)^3)$ expected time with $O(n^2)$ processors using $O(n^2)$ random bits. Our results immediately yield a uniform algorithm with $O((\log n)^3)$ running time and $O(n^{2+\varepsilon'})$ processors using only n^ε random bits, where $\varepsilon, \varepsilon' > 0$ can be set arbitrarily small.

Recently, Karp, Upfal and Wigderson [13] have shown that finding a maximum graph matching is in RNC_U , and Anderson and Mayr [2] have shown that finding a maximal path is in RNC_U . Our results also immediately yield efficient randomized uniform algorithms for these problems, using only n^ϵ bits for any $\epsilon > 0$.

Implications. An interesting theoretical application of our result is that RNC_U is contained within the class of languages recognized by uniform deterministic circuits of unbounded fan-in with poly-log depth and 2^{n^ϵ} size for any $\epsilon > 0$. (Adleman [1] proved RNC_U is contained in (nonuniform) NC , but the previous best construction for bounding RNC_U by deterministic uniform circuits of poly-log depth required $2^{\Omega(n)}$ size.) This extends a result of Yao [21] for sequential polynomial time computations to poly-log time parallel computations.

2. Definitions and results.

Notation. We use the following notation throughout the paper:

N A positive composite integer such that each prime factor of N is greater than N^c for a fixed $c > 0$.

\mathbf{Z}_N^* The multiplicative group of positive integers less than and relatively prime to N . (Note that the fact that N has only large factors implies that a random positive integer less than N is an element of \mathbf{Z}_N^* with high probability.)

We will sometimes use $x \bmod N$ to indicate the residue of x modulo N .

Definitions. An NC -machine (Cook [8]) is a deterministic parallel algorithm which runs on $n^{O(1)}$ parallel-RAM (P-RAM) processors in time $(\log n)^{O(1)}$ for input of size n . (Note that NC_U is the class of languages accepted by NC -machines.)

An RNC -machine is a randomized parallel algorithm which runs on $n^{O(1)}$ P-RAM processors in time $(\log n)^{O(1)}$ for input of size n . (Note that RNC_U is the class of languages accepted by RNC -machines.)

Given $s_0 \in \mathbf{Z}_N^*$, the *multiplicative inverse* of s_0 modulo N is the s_0^{-1} such that $s_0 s_0^{-1} = 1 \bmod N$.

For a fixed N , given an arbitrary $k \in \mathbf{Z}_N^*$, the *Multiplicative Inverse Problem* is to find the multiplicative inverse of k modulo N . Note that the input size to the problem is $n = \lceil \log N \rceil$.

The problem of finding multiplicative inverses in poly-log depth has been studied extensively. (Cook [8], Kannan, Miller and Randolph [11], Reif [15] and von zur Gathen [20].) Based on the lack of significant positive results obtained so far we conjecture the following.

Complexity assumption. There exists an infinite sequence of numbers N_1, N_2, \dots constructible in NC_U such that for each $n = 1, 2, \dots$ we have $n = \lceil \log N_n \rceil$ and that no RNC -machine can solve the Multiplicative Inverse Problem for arbitrary elements of $\mathbf{Z}_{N_n}^*$ for almost all values of n .

(Actually we could replace this complexity assumption with the weaker assumption that there exists a k such that for almost all n there exists an n' such that $n < n' < n^k$ and no RNC -machine can solve the Multiplicative Inverse Problem for arbitrary elements of $\mathbf{Z}_{N_{n'}}^*$. All the theorems in this paper would remain true under that weaker assumption.)

Definitions. A set S of bit sequences $\sigma = (b_1, \dots, b_J)$ of length $J = n^{O(1)}$ pseudorandom bits is *RNC -cryptographically strong* if no RNC -machine can, on a random input $b_1, \dots, b_i \in S$ ($i < J$) predict any one bit b_i, \dots, b_J with expected success of $\frac{1}{2} + 1/n^{O(1)}$. Informally, the bit sequences are RNC -cryptographically strong if no RNC -machine can predict untransmitted bits with an expected success rate significantly better than $\frac{1}{2}$.

THEOREM. *If no RNC-machine can solve the Multiplicative Inverse Problem for almost all n , then there exists a deterministic NC-machine \mathcal{G} which on an input seed of n bits outputs an RNC-cryptographically strong sequence of $J = n^{O(1)}$ pseudorandom bits. \mathcal{G} can be computed by a bounded fan-in uniform Boolean circuit of depth $O((\log n)(\log \log n))$ and size $n^{O(1)}$.*

This theorem is proved in § 3.

Definition. An **RNC-statistical test** is an **RNC-machine** which attempts to distinguish truly random bit sequences from pseudorandom bit sequences. A statistical test succeeds if it correctly distinguishes the pseudorandom bit sequences from truly random bit sequences with probability at least $1/n^{O(1)}$.

By a technique due to Yao [21] we can show that no RNC statistical test can succeed on RNC-cryptographically strong bit sequences. Hence we have the following.

COROLLARY 1. *If no RNC-machine can solve the Multiplicative Inverse Problem for almost all n , then no RNC-statistical test can succeed on our pseudorandom bit generator \mathcal{G} .*

COROLLARY 2. *If the N_n are constructible in depth $h(n)$, then given a randomized parallel algorithm \mathcal{A} (over a wide class of machine models such as parallel RAMS and fixed connection networks) with time bound $T(n)$ and processor bound $P(n)$ then \mathcal{A} can be simulated by a parallel algorithm with time bound $T(n) + h(n) + O((\log n)(\log \log n))$, processor bound $P(n)n^\varepsilon$, and only using n^ε truly random bits for any $\varepsilon > 0$, where $\varepsilon' = O(\varepsilon)$.*

CIRCUIT_U ($D(n)$, $S(n)$) is the class of languages accepted by uniform deterministic circuits with unbounded fan-in, depth $D(n)$, and size $S(n)$. (See § 4 for a precise definition of these complexity classes.)

COROLLARY 3. *If for almost all n the Multiplicative Inverse Problem is not in RNC then*

$$\text{RNC}_U \subseteq \bigcup_{c>0} \bigcap_{\varepsilon>0} \text{CIRCUIT}_U((\log n)^c, 2^{n^\varepsilon})$$

This corollary is proved in § 4.

COROLLARY 4. *There exists a cryptosystem where encryption and decryption can be done by an NC-machine on $n^{O(1)}$ bits given a secret shared key exactly n bits long (here n is a security parameter). If no RNC-machine can solve the Multiplicative Inverse Problem, then no RNC-machine can decrypt ciphertext exchanged in this cryptosystem.*

We use the pseudorandom bits as a “one-time pad”—we take the sequential exclusive-or of the plaintext and the pseudorandom bits to produce the ciphertext and take the sequential exclusive-or of the ciphertext and the pseudorandom bits to obtain the plaintext again. Encryption and decryption both take parallel poly-log time but an opponent cannot decrypt the ciphertext with RNC-machine.

3. The proof of the main theorem.

Properties. We recall the following facts which we use implicitly (Beame, Cook and Hoover [4], Reif [15], and Shonhage and Strassen [17]):

- There exists an NC-machine for multiplication of two numbers in \mathbf{Z}_N^* .
- $2 \log p$ multiplications suffice to find the p th power of a number in \mathbf{Z}_N^* .
- If $p < (\log N)^{O(1)}$ there exists an NC-machine for finding the p th power of a number in \mathbf{Z}_N^* .

Fix $m = \lceil \log N \rceil$ throughout this section.

Let \mathcal{G} be the NC-machine which performs the following operations:

Input: random elements $s_0, k \in \mathbf{Z}_N^*$.

Output: b_1, \dots, b_J where $J = m^{O(1)}$.

Method: In parallel each processor P_i ($i = 1, \dots, J$) calculates $s_i = ks_0^i \bmod N$ and $b_{J-i+1} = B(s_i)$ where

$$B(x) = \begin{cases} 0 & \text{if } x \leq N/2, \\ 1 & \text{if } x > N/2. \end{cases}$$

LEMMA. *If there exists an RNC-machine which can determine the value of b_j with probability 1 (i.e., no error) on input b_1, \dots, b_{J-1} , then there exists an RNC-machine which can solve the Multiplicative Inverse Problem for $\mathbf{Z}_{N,n}^*$.*

Proof of the lemma. Suppose that MB (for "magic box") is an oracle which can determine the value of b_j with probability 1. Then given $s_0 \in \mathbf{Z}_N^*$ we can find $s_0^{-1} \bmod N$. We can find this by running in parallel the following algorithm on each processor P_i for ($0 \leq i \leq m$):

Set $k \leftarrow 2^i$. In parallel set $b_i \leftarrow B(ks_0^{J-i-1})$ for $1 \leq i \leq J-1$. Note that $b_J = B(2^i s_0^{-1})$. Feed the sequence (b_1, \dots, b_{J-1}) to MB to get b_J . Set the i th most significant bit of δ to be $B(2^i s_0^{-1})$. Define

$$\phi(\delta) = \left\lfloor \frac{\delta N}{2^m} \right\rfloor.$$

Then $\phi(\delta) = s_0^{-1} \bmod N$. \square

THEOREM. *If there exists an RNC-machine which can determine the value of b_j with probability at least $\frac{1}{2} + 1/m^{O(1)}$ on input b_1, \dots, b_J , then there exists an RNC-machine \mathcal{H} which can solve the Multiplicative Inverse Problem for $\mathbf{Z}_{N,n}^*$. \mathcal{H} can be computed by a bounded fan-in Boolean circuit of depth $O((\log n)(\log \log n))$ and size $n^{O(1)}$.*

Proof of the theorem. Assume that there exists an RNC-machine MB which can predict b_j with probability $\frac{1}{2} + 2/m^c$. Let $H = 2(c+1)\lceil \log m \rceil$. Let δ and ϕ be as in the proof of the lemma.

Let $S = \{0, 1, \dots, 2^{H-1} - 1\}$. For each $0 \leq y < x \leq m$, we will create, by randomized methods, two functions $F_{x,y}: S \rightarrow \{0, 1\}^{x-y}$ and $G_{x,y}: S \rightarrow S$. Informally, values in S are guesses; $F_{x,y}$ is a rule for transforming a guess $j_x \in S$ into the x th to y th most significant bits of δ ; and $G_{x,y}$ is a rule for transforming the guess $j_x \in S$ into the guess $j_{x-1} \in S$.

If an RNC-machine could find δ for arbitrary s_0 , we could solve the Multiplicative Inverse Problem. It will turn out that for some $j_m \in S$, $F_{m,0}(j_m) = \delta$ with probability $\frac{1}{2}$. We can verify this occurrence simply by checking whether $s_0 \phi(\delta) = 1 \bmod N$. If we do not immediately find $s_0^{-1} \bmod N$, we simply form a new $F_{m,0}$ by randomized methods and continue testing until we do find $s_0^{-1} \bmod N$.

Suppose we can determine j_x such that we know $(2^x s_0^{-1} \bmod N)$ belongs to one of the two intervals

$$\left\{ \left[\left\lfloor \frac{j_x N}{2^H} \right\rfloor, \left\lfloor \frac{(j_x + 1)N}{2^H} \right\rfloor \right], \left[\left\lfloor \frac{(2^{H-1} + j_x)N}{2^H} \right\rfloor, \left\lfloor \frac{(2^{H-1} + j_x + 1)N}{2^H} \right\rfloor \right] \right\}.$$

We can pick 2^H random values $\beta \in \mathbf{Z}_N^*$ and let v be MB's prediction for

$$B \left(2^x s_0^{-1} + \left\lfloor \frac{N j_x}{2^H} \right\rfloor + \beta \bmod N \right).$$

When β lies in the interval

$$\left[0, \left\lfloor \frac{(2^{H-1} - 1)N}{2^H} \right\rfloor \right],$$

mark a vote for v , when β lies in the interval

$$\left[\left\lfloor \frac{N}{2} \right\rfloor, \left\lfloor \frac{(2^H - 1)N}{2^H} \right\rfloor \right],$$

mark a vote the complement of v , and mark a *null* vote when β lies in other intervals. By assumption, MB predicts correctly with probability at least $\frac{1}{2} + 2/m^c$.

We can assign a processor to calculate MB's prediction for each of the 2^H randomly chosen values of $\beta \in \mathbf{Z}_N^*$. This computation can be done in poly-log time for each β . The expected fraction of *null* votes is $2^{1-H} < 1/m^c$. Thus we have a bias of at least $2/m^c - 1/m^c = 1/m^c$ between 0 and 1 votes. Set $F_{x,x-1}(j_x)$ (our guess for $B(2^x s_0^{-1} \bmod N)$) to be a value which got the most votes. If our guess for $B(2^x s_0^{-1} \bmod N)$ is right, this immediately identifies which of the two intervals $(2^x s_0^{-1} \bmod N)$ belongs to. By the argument in the Appendix, 2^H tests are sufficient to make our guess correct with probability at least $1 - 1/2^m$. If our guess is right, that immediately determines the value of j_{x-1} ; that is, we can determine that $(2^{x-1} s_0^{-1} \bmod N)$ lies in one of the two intervals

$$\left\{ \left[\left[\frac{j_{x-1}N}{2^H} \right], \left[\frac{(j_{x-1}+1)N}{2^H} \right] \right], \left[\left[\frac{(2^{H-1}+j_{x-1})N}{2^H} \right], \left[\frac{(2^{H-1}+j_{x-1}+1)N}{2^H} \right] \right] \right\},$$

namely

$$j_{x-1} = G_{x,x-1}(j_x) = \lfloor j_x/2 \rfloor + 2^{H-2}(F_{x,x-1}(j_x)).$$

We can calculate in parallel, for each $m \geq x \geq 1$, the functions $F_{x,x-1}$ and $G_{x,x-1}$, since the domain is finite and of polynomial size. If $x - y > 1$, then $F_{x,y}$ and $G_{x,y}$ can be recursively defined as

$$F_{x,y}(j_x) = F_{z,y}(G_{x,z}(j_x))2^{x-z} + F_{x,z}(j_x)$$

and

$$G_{x,y}(j_x) = G_{z,y}(G_{x,z}(j_x))$$

where $z = \lceil (x+y)/2 \rceil$. For each x, y pair ($0 \leq y < x \leq m$) and each $j_x \in S$ we repeatedly calculate the appropriate compositions of these functions for all j_x in the domain of the functions. Thus we can compute $F_{m,0}$ in $\lceil \log m \rceil$ stages.

Some guess j_m is correct. Suppose that for all $1 \leq i \leq m$, that (1) $G_{i,i-1}(j_i)$ is the correct value of j_{i-1} . Then (2) $F_{m,0}(j_m)$ would be the correct value of δ . For each i , the probability that (1) is true for a particular j_i is $(1 - 2^{-m})$, so the probability that (2) is true is $(1 - 2^{-m})^{m-1} > 1 - (m-1)2^{-m} > \frac{1}{2}$.

For some $j_m \in S$, it will be true that $F_{m,0}(j_m) = \delta$ with probability $\frac{1}{2}$. We can try all possible j_m in parallel, and find out if we have a correct value by checking whether $\phi(F_{m,0}(j_m))s_0 = 1 \bmod N$. (Of course, it might happen that an incorrect guess for j_m might give a correct value for δ but this can only speed the calculation.) In the event that we do not get the correct value for $s_0^{-1} \bmod N$, we simply form new $F_{x,y}$ and $G_{x,y}$ functions and continue until we do get the correct value. \square

4. Randomized and deterministic parallel complexity. Let \mathcal{C} be a list of circuits (C_1, C_2, \dots) of unbounded fan-in where C_n and n inputs and size $S(n)$. We consider \mathcal{C} to be *uniform* if there exists a $(\log S(n))$ space deterministic Turing machine which, given any n , outputs the circuit C_n . Let $\text{CIRCUIT}(D(n), S(n))$ be the class of all languages accepted by deterministic Boolean circuits with unbounded fan-in, depth $D(n)$, and size $S(n)$. As usual we define

$$\text{NC} = \bigcup_{k_1 > 0, k_2 > 0} \text{CIRCUIT}((\log n)^{k_1}, n^{k_2}).$$

We allow a *randomized Boolean circuit* \mathbf{C} to have r special nodes, each of which are assigned independent random bits chosen from $\{0, 1\}$ with equal probability. \mathbf{C} *accepts* an input $\omega \in \{0, 1\}^n$ if \mathbf{C} outputs 1 with probability $> \frac{1}{2}$; otherwise \mathbf{C} *rejects* the input. For simplicity, we consider only one-sided error randomized circuits which

never output a 1 on an input they have rejected. (The construction below can easily be extended to two-sided error randomized circuits which have an acceptance probability of at least $\frac{1}{2} + 1/n^k$ for some $k > 1$.) Let $\text{RCIRCUIT}(D(n), S(n))$ be the class of languages accepted by randomized circuits with unbounded fan-in, depth $D(n)$, and size $S(n)$. We define

$$\text{RNC} = \bigcup_{k_1 > 0, k_2 > 0} \text{RCIRCUIT}((\log n)^{k_1}, n^{k_2}).$$

We define CIRCUIT_U , NC_U , RCIRCUIT_U , and RNC_U analogously—restricting the circuits to be uniform.

Proof of Corollary 3. Let C be a (one-sided error) uniform randomized Boolean circuit with n inputs, depth $D(n) = (\log n)^{k_1}$, and size $S(n) = n^{k_2}$. Fix any $\varepsilon > 0$.

First suppose we had a source of $b = \lceil n^{\varepsilon/2} \rceil$ truly random bits. Observe that C uses at most $S(n) = n^{k_2}$ random bits on each execution. Since $S(n) \leq b^{\varepsilon'}$ where $\varepsilon' = \lceil \varepsilon/k_2 \rceil$ is constant, we can apply our parallel pseudorandom bit generator \mathcal{G} to produce $S(n)$ pseudorandom bits in $(\log n)^{O(1)}$ parallel time using $n^{O(1)}$ processors and using the b truly random bits as the seed. We can view the execution of C on the given input ω as a statistical test. By Corollary 2, given an input $\omega \in \{0, 1\}^n$, we need only execute C on ω for each of the 2^b possible pseudorandom bit sequences. We accept ω if C ever outputs 1.

Furthermore, we can avoid the use of a truly random seed by simply (1) enumerating all b -bit numbers in parallel; (2) executing the parallel pseudorandom bit generator using each of the b -bit numbers as a seed; and (3) executing C in parallel on ω on each of the resulting pseudorandom bit sequences. If C ever outputs 1 we accept ω . The resulting uniform circuit requires size $2^{b^2} \leq 2^{n^\varepsilon}$ and depth $(\log n)^{O(1)} + O(D(n)) = (\log n)^{O(1)}$. \square

Note that if we require that our simulation circuit have bounded fan-in, then to simulate a circuit accepting a language in RNC_U , we require $n^{O(1)}$ (rather than $(\log n)^{O(1)}$ depth) and 2^{n^ε} size. This is an improvement over previous size bounds for RNC_U .

Appendix. Let X be the binomial variable which is the sum of τ independent Bernoulli trials each of which has a probability $p = \frac{1}{2} + 1/m^c$ of giving a value 1 and probability $1 - p = \frac{1}{2} - 1/m^c$ of giving a value of 0. We need to find τ large enough so that

$$\Pr[X < \tau/2] < 1/2^n.$$

Using Chernoff bounds (Chernoff [7]) we recall that

$$\Pr[X < (1 - \delta)\tau p] < e^{-\delta^2 \tau p / 2}.$$

Substituting $p = \frac{1}{2} + 1/m^c$ and $(1 - \delta)p = \frac{1}{2}$ we get

$$\Pr[X < \tau/2] < e^{-(1/m^{2c})\tau(1/2 - 1/m^c)(1/2)} < e^{-\tau/8m^{2c}}.$$

If we set $\tau = m^{2c+1}$, our initial condition is satisfied. Since $2^H = \tau$, setting

$$H = 2(c+1)\lceil \log m \rceil = O(\log m)$$

will give us conditions sufficient to prove the main theorem.

Acknowledgments. We would like to thank Michael Rabin for being an inspiration to us in the fields of randomized algorithms and cryptography.

We are indebted to Silvio Micali for his many helpful and insightful comments on this manuscript.

Also, thanks to Benny Chor, Shafi Goldwasser, Johan Hastad, Brian O'Toole, Charles Rackoff, Les Valiant and Vijay Vazirani for their comments.

REFERENCES

- [1] L. ADLEMAN, *Two theorems on random polynomial time*, Proc. 19th IEEE Symposium on Foundations of Computer Science, Ann Arbor, MI, October 1978, pp. 75-83.
- [2] R. ANDERSON, *A parallel algorithm for the maximal path problem*, Proc. 17th ACM Symposium on Theory of Computing, Providence, RI, May 1985, pp. 33-37.
- [3] W. ALEXI, B. CHOR, O. GOLDREICH AND C. SCHNORR, *RSA/Rabin bits are $\frac{1}{2} + 1/\text{poly}(\log N)$ secure*, Proc. 25th IEEE Symposium on Foundations of Computer Science, Singer Island, FL, October 1984, pp. 449-457.
- [4] P. BEAME, S. COOK AND H. HOOVER, *Small depth circuits for integer products, powers, and division*, Proc. 25th IEEE Symposium on Foundations of Computer Science, Singer Island, FL, October 1984, pp. 1-6.
- [5] L. BLUM, M. BLUM AND M. SHUB, *A simple secure pseudo-random number generator*, Proc. CRYPTO-82, Santa Barbara, CA, September 1982, pp. 112-117.
- [6] M. BLUM AND S. MICALI, *How to generate cryptographically strong sequences of pseudo-random bits*, this Journal, 13 (1984), pp. 850-864.
- [7] H. CHERNOFF, *A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations*, Ann. Math. Statist., 23 (1952), pp. 493-507.
- [8] S. A. COOK, *The classification of problems which have fast parallel algorithms*, Lecture Notes in Computer Science, Vol. 158, M. Karpenski, ed., Springer-Verlag, New York, Berlin, 1983.
- [9] O. GOLDREICH, S. GOLDWASSER AND S. MICALI, *How to construct random functions*, Proc. 25th IEEE Symposium on Foundations of Computer Science, Singer Island, FL, October 1984, pp. 464-479.
- [10] S. GOLDWASSER, S. MICALI AND P. TONG, *Why and how to establish a private code on a public network*, Proc. 23rd IEEE Symposium on Foundations of Computer Science, Chicago, IL, October 1982, pp. 134-144.
- [11] R. KANNAN, G. MILLER AND L. RUDOLF, *Sublinear parallel algorithms for the greatest common divisor of two integers*, Proc. 25th IEEE Symposium on Foundations of Computer Science, Singer Island, FL, October 1984, pp. 7-11.
- [12] R. KARP AND A. WIGDERSON, *A fast parallel algorithm for the maximal independent set problem*, Proc. 16th ACM Symposium on Theory of Computation, Washington, DC, May 1984, pp. 266-272.
- [13] R. KARP, E. UPFAL AND A. WIGDERSON, *Constructing a perfect graph matching in Random NC*, Proc. 17th ACM Symposium on Theory of Computing, Providence, RI, May 1987, pp. 22-32.
- [14] D. KNUTH, *The Art of Computer Programming, Vol. 2: Seminumerical Algorithms*, 2nd ed., Addison-Wesley, Reading, MA, 1981.
- [15] J. REIF, *Logarithmic depth circuits for algebraic functions*, Proc. 24th IEEE Symposium on Foundations of Computer Science, Tucson, AZ, October 1983, pp. 138-145. Revised in Technical Report TR-84-18, Center for Research in Computing Technology, Harvard University, Cambridge, MA, this Journal, to appear.
- [16] A. SHAMIR, *On the generation of cryptographically strong pseudo-random sequences*, ACM Trans. Comput. Sys., 1 (1983), pp. 38-44.
- [17] A. SHONHAGE AND V. STRASSEN, *Schnelle Multiplikation grosser Zahlen*, Computing, 7 (1974), pp. 281-292.
- [18] L. VALIANT, S. SKYUM, S. BERKOWITZ AND C. RACKOFF, *Fast parallel computation of polynomials using few processors*, this Journal, 12 (1983), pp. 641-644.
- [19] U. VAZIRANI AND V. VAZIRANI, *Trapdoor pseudorandom number generators with applications to protocol design*, Proc. 24th IEEE Symposium on Foundations of Computer Science, Tucson, AZ, October 1983, pp. 23-30.
- [20] J. VON ZUR GATHEN, *Parallel algorithms for algebraic problems*, Proc. 15th ACM Symposium on Theory of Computing, Boston, MA, April 1983, pp. 13-18.
- [21] A. YAO, *Theory and applications of trapdoor functions*, Proc. 23rd IEEE Symposium on Foundations of Computer Science, Chicago, IL, October 1982, pp. 80-91.