PROBABILISTIC COMMUNICATION COMPLEXITY
(preliminary version)

Ramamohan Paturi and Janos Simon[1]
Department of Computer Science
The Pennsylvania State University

Abstract: We study (unbounded error) probabilistic communication complexity. Our new results include

- one way and two way complexities differ by at most 1
- certain functions like equality and the verification of Hamming distance have upper bounds that are considerably better than their counterparts in deterministic, nondeterministic, or bounded error probabilistic model
- there exists a function which requires $\Omega(\log n)$ information transfer

As an application, we prove that a certain language requires $\Omega(n \log n)$ time to be recognized by a 1-tape (unbounded error) probabilistic Turing machine. This bound is optimal. (Previous lower bound results [Yao 1] require acceptance by bounded error computation. We believe that this is the first nontrivial lower bound on the time required by unrestricted probabilistic Turing machines.)

1. DEFINITIONS.

The essentials of this model are the same as those of Yao [Yao 2] who introduced the notion of communication complexity (see also [PS] and [JKS] for variants of and extensions to the model).

Two processors P_0, and P_1 wish to compute a function of two arguments. (We assume in most of this paper that the function is boolean.) The first argument, x_0, of the boolean function $f:\{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$, is known to P_0, and second argument, x_1, is known to P_1. In order to compute f, P_0, and P_1 communicate with each other in turns by sending messages (sequences of bits) according to some protocol. P_0, and P_1 have unlimited local computing power, and the ability to realize an arbitrary probability distribution over the set of messages they transmit in each turn. The complexity measure is the number of bits transmitted.

Given the input x_i to P_i for $i=0,1$, the computation, according to some protocol φ, will be as follows: P_0 is always the first one to send a message. The processors communicate in turns. The last message is always sent by P_1 and is a single bit. The last bit is the output produced. Each message will be sent with a certain probability, determined by the protocol. A probabilistic computation can be viewed as a stochastic process. An event in this process is a sequence of messages $\beta_1, \beta_2, \ldots, \beta_{2k}$ (where message β_i is sent by processor $P_{i+1 mod 2}$). The probability distribution, given by the protocol, assigns a probability to each event. The *result* of an event is the output produced by the associated set of messages. The

[1]Research partially supported by USARO research contract DAAG29-82-K-0110.

0272-5428/84/0000/0118$01.00 © 1984 IEEE

protocol φ outputs the bit b ($b=0$ or 1) if the probability of events whose result is b is greater than $\frac{1}{2}$.

Formally, a protocol can be specified by a function $\varphi: \{0,1\}^n \times \{0,1\}^* \times \{0,1\}^+ \to [0,1]$. $\varphi(x,\alpha,\beta)$ is the probability with which the message β will be sent by a processor, where x is its input and α is the concatenation of the sequence of messages exchanged so far. φ has the property that the set $\{\beta \mid \exists x\ \varphi(x,\alpha,\beta) \neq 0\}$ is finite and prefix free for each α. Note that $\sum_{\beta}\varphi(x,\alpha,\beta)=1$. Due to the prefix freeness property, a concatenated sequence of messages can again be decomposed into a sequence of messages which is unique for the given protocol.

Let φ be a protocol. Let x_i be the input at P_i for $i=0,1$. Let $(\beta_1,p_1),\ldots,(\beta_{2l},P_{2l})$ be such that

$$\varphi(x_0,\lambda,\beta_1) = p_1$$
where λ is the null string;
$$\varphi(x_0,\beta_1 \cdots \beta_{2j},\beta_{2j+1}) = p_{2j+1} \quad \text{for } j=1,\ldots l-1;$$
$$\varphi(x_1,\beta_1 \cdots \beta_{2j-1},\beta_{2j}) = p_{2j} \quad \text{for } j=1,\ldots,l.$$

The set of all such sequences is the computation $T_\varphi(x_0,x_1)$ under the protocol φ with the input x_i at P_i.

Note that the probabilities $p_1,p_3,\ldots,p_{2l-1}$ do not depend on the input at the processor P_1. Similarly $p_2,p_4,\ldots,p_{2l}$ do not depend on the input at P_0. We, therefore, define two functions $\varphi_0,\varphi_1:\{0,1\}^n \times M_\varphi \to [0,1]^+$, where M_φ is the set of all concatenated sequences of messages that are transmitted between P_0 , and P_1 with positive probability for some input. Let $\beta_1,\ldots,\beta_{2l}$ be the decomposition of $\alpha \in M_\varphi$ under the protocol φ. Now, $\varphi_i(x,\alpha) = (p_1,\ldots,p_l)$ where

$$p_j = \varphi(x,\beta_1 \cdots \beta_{2j-2}, \beta_{2j-1}) \quad \text{if } i=0$$
$$p_j = \varphi(x,\beta_1 \cdots \beta_{2j-1}, \beta_{2j}) \quad \text{if } i=1$$
for $j=1,\ldots,l$.

These functions φ_0, and φ_1 together with the decomposition for each $\alpha \in M_\varphi$ capture all the information contained in the protocol φ.

In the computation $T_\varphi(x_0,x_1)$, the probability of outputting the bit b is $\sum_{\alpha b \in M_\varphi} {}^*\varphi_0(x_0,\alpha b)\,{}^*\varphi_1(x_0,\alpha b)$. Here, $*$ is an operator, that applied to a list of real numbers, yields their product.

The communication complexity $\tilde{C}_\varphi$ of the protocol φ is $\max\{|\alpha| \mid \alpha \in M_\varphi\}$. The protocol φ computes a function f if $f(x_0,x_1)=b$ iff the probability fo outputting the bit b in the computation $T_\varphi(x_0,x_1)$ is greater than $\frac{1}{2}$.

The unbounded error probabilistic communication complexity $\tilde{C}_f$ is $\min\{\tilde{C}_\varphi \mid \varphi \text{ computes } f\}$.

A restricted model in which only one processor p_0 is allowed to send messages is also of interest because of its equivalence to the unrestricted two-way model. In this one-way model, P_0 sends the messages $\beta_1,\ldots,\beta_l$ with probabilities $p_1,\ldots,p_l$ respectively. P_1 on the receipt of β_i, outputs 1 with probability q_i and

0 with probability $1-q_i$. The set of messages sent by P_0 and the probability distribution on it is entirely determined by the input at P_0 alone and are not influenced by the input at P_1. Similarly, the probabilities q_i at P_1 depend only on its input and the message received. The one-way protocol φ can therefore be completely specified by two functions $\varphi_0,\varphi_1:\{0,1\}^n \times M_\varphi \to [0,1]$, where M_φ is the set of all messages that are sent by P_0 with positive probability for some input. $\varphi_0(x,\alpha)$ is the probability with which the message α is sent by P_0 with input x. $\varphi_1(x,\alpha)$ is the probability with which P_1 with input x outputs 1 upon receiving the message α. Since the particular set of messages is not relevant, φ_0,φ_1 can be represented as functions from $\{0,1\}^n$ to $[0,1]^K$, where $|M_\varphi| = K$. The communication complexity of the protocol φ is $\lceil \log_2 K \rceil$. K is also called the length of the protocol φ. Other notions for one way protocols are defined in the analogous way.

Equivalence of one-way and two-way complexities Finally, we exhibit a one-way protocol for each two way protocol such that both compute the same function and their communication complexities differ by most 1.

Theorem 1: Let φ be a two-way protocol. Then, there exists a one-way protocol φ' such that

1) φ, and φ' compute the same function
2) $\tilde{C}_{\varphi'} \leq \tilde{C}_\varphi + 1$

Proof: Let φ_0,φ_1, and M_φ be as defined earlier for the two-way protocol φ. Let $M_\varphi = M_\varphi{}^1 \cup M_\varphi{}^0$. $\alpha \in M_\varphi{}^b$ if the last bit of α is b. Let

$$d_x{}^b = \sum_{\alpha \in M_\varphi{}^b} {}^*\varphi_0(x,\alpha) \; ; \; d = \max_x d_x{}^1.$$

We define the one-way protocol φ' such that

$M_{\varphi'} = M_\varphi \cup \{\gamma\}$
with $\gamma \notin M_\varphi$.

$$\varphi_0'(x,\alpha) = \frac{1}{2d} {}^*\varphi_0(x,\alpha) \quad \text{for } \alpha \in M_\varphi{}^1$$

$$\varphi_0'(x,\gamma) = \frac{1}{2}\left(1-\frac{d_x{}^1}{d}\right)$$

$$\varphi_0'(x,\alpha) = \frac{1}{2d_x{}^0} {}^*\varphi_0(x,\alpha) \quad \text{for } \alpha \in M_\varphi{}^0$$

$$\varphi_1'(x,\alpha) = {}^*\varphi_1(x,\alpha) \quad \text{for } \alpha \in M_\varphi{}^0$$

$$\varphi_1'(x,\gamma) = 0$$

$$\varphi_1'(x,\alpha) = 1-\frac{1}{2d} \quad \text{for } \alpha \in M_\varphi{}^0$$

φ_i' are functions from $\{0,1\}^n \times M_{\varphi'}$ to $[0,1]$. It can be easily verified that φ, and φ' compute the same function. It is also clear that their complexities differ by at most 1. ▪

2. Why Communication Complexity, in particular, Unbounded Error Probabilistic Communication Complexity?

There are well known reasons to study this measure of complexity [Yao 2][PS]:

-Communication is the bottleneck in many parallel algorithms, VLSI implementations, and distributed systems.

-It is closely related to other questions in computational complexity (lower bounds in restricted models of computation, like 1-tape Turing machines, branching programs, and monotone circuits; generaliza-

tion of static measures of complexity, like circuit size and Kolmogorov complexity, etc.)

-It allows us to study, otherwise intractable questions (like the power of nondeterminism, the power of probabilistic choices, etc.) in a favourable environment, where it is possible to settle some of them.

-Perhaps most importantly, it is a rich source of interesting problems and of techniques for solving them. In our study of the unrestricted probabilistic model, we came up with some combinatorial problems related to arrangements of hyperplanes [Za] and oriented matroids [FL]. We hope that these questions will stimulate further research by both mathematicians and computer scientists.

This unrestricted probabilistic model is not intended to serve as the basis for a theory of 'reliable information transfer'. Rather, we are interested in understanding the power of unrestricted probabilistic choice in parallel environments. The facts below show that this power is considerable.

Let $I(x,y) = (x=y)$; $\bar{I}(x,y) = (x \neq y)$; and $G(x,y) = (x \geq y)$, where x and y are interpreted as n-bit integers.

Fact: a[1]) $\tilde{C}_I(1\to2) = \tilde{C}_{\bar{I}}(1\to2) = 2$

b) $\tilde{C}_G(1\to2) = 1$

Recall that any deterministic protocol for $I,\bar{I}$ or G requires n bits of information transfer, and every nondeterministic protocol for $\bar{I}$ or G requires n bits of information transfer [Yao 2] [PS]. Even bounded error probabilistic protocols must exchange $\Omega(\log n)$ bits to compute the functions I, $\bar{I}$, and G [Yao 1]. An optimal protocol for computing I can be found in the appendix.

An immediate question is whether these facts mean that the model is trivial. After all, B could probabilistically guess x, perform the protocol for equality, and compute $f(x,y)$ with just 2 bits of information transfer. Fortunately(?), the strategy does not work since, as one can verify, the computation is not reliable enough. This challenges us to try to prove lower bounds for probabilistic information transfer. The results in this paper partially answer this challenge.

The problem (of proving lower bounds for probabilistic information tranfer) requires new techniques: In the case of deterministic protocols, a counting argument immediately yields a (nonconstructive) proof of the existence of functions with asymptotically linear communication complexity. For example, there are $2^{2^{2n}}$ boolean functions of $2n$ variables, but only $2^{2^{O(l)}}$ different deterministic protocols of length l. There are, on the other hand, nondenumerably many probabilistic protocols of length l, since the probabilities are arbitrary. Although, by a

[1]Fact a) was known to M. Rabin in the context of crossing sequences for Turing machines [B-0].

continuity argument, we can restrict ourselves to rational probabilities with bounded denumerators, the number of resulting protocols still makes the counting argument impossible. In the case of the bounded error probabilistic model, both the logarithmic and linear lower bound arguments make use of the fact that the error in the computation is bounded by a constant.

We proved that the one way probabilistic model is as powerful as the two way one. In contrast, we have, in the deterministic model, exponential gaps between not only one way and two way protocols, but also between k-turn and $k+1$-turn protocols [DGSch]. We present several equivalent exact characterizations of the probabilistic communication complexity of a function: one in terms of the approximations of a boolean matrix by rank 1 real matrices, and the other, a geometric one, using arrangements of hyperplanes. These characterizations can be used to construct a hierarchy of functions f_i, that require i bits of information transfer for $1 \leq i \leq \log n$.

It is not known whether all functions can be computed using $O(\log n)$ information transfer. This question is equivalent to some combinatorial problems related to oriented matroids that appear interesting on their own. The equivalence follows from our characterization of probabilistic communication complexity in terms of arrangements of hyperplanes.

In the sequel, we present a brief outline of these results.

3. RESULTS:

We consider only one-way protocols. If φ is a one-way protocol of length k, let $\varphi_0,\varphi_1: \{0,1\}^n \to [0,1]^k$ be the associated probability functions.

Arrangements of Hyperplanes and Probabilistic Communication Complexity

We present our first characterization of probabilistic communication complexity in terms of arrangements of hyperplanes.

An arrangement $Arr(H)$ of hyperplanes is a finite set $H = \{h_1, h_2, \ldots, h_m\}$ of hyperplanes in R^d for some d. The regions of an arrangement $Arr(H)$ are the nonempty connected components of R^d, when the hyperplanes in H are deleted. Each region r of the arrangement can be characterized by an m bit string whose ith bit (for $i=1,\ldots,m$) is 1 iff the region r is in the positive half space of the hyperplane h_i. We call this bit string, the *signature* of the region r. We say that the arrangement $Arr(H)$ *realizes* the set $S_H \subseteq \{0,1\}^m$ of signatures if $S_H = \{w \in \{0,1\}^m \mid w$ is a signature of some region r in $Arr(H)\}$.

We call each $w \in \{0,1\}^m$ a *requirement*. A requirement $w \in \{0,1\}^m$ is *satisfied* by an arrangement $Arr(H)$ of m hyperplanes H in R^d for some d, if $w \in S_H$. Similarly, we say that a

boolean valued matrix M of order $k \times m$ is satisfied by an arrangement $Arr(H)$ of m hyperplanes H in R^d if each row of M when viewed as a requirement belongs to S_H.

Theorem 2: Let M be the matrix of a function f. Let d be the smallest dimension in which there is an arrangement $Arr(H)$ of 2^n hyperplanes H that satisfies the matrix M. Then

$$\lceil \log d \rceil \leq \widetilde{C}_f \leq \lceil \log d \rceil + 1.$$

The proof essentially consists of interpreting, for each x_0 and y_0, $\varphi_0(x_0)$ and $\varphi_1(x_1)$ (defined previously for one way protocols of length k) as a hyperplane and a point of R^k respectively, and using a continuity argument.

It is possible to give another equivalent characterization using rank 1 real matrices. We say that a real matrix $\hat{M}$ is an *approximation* of a boolean matrix M of the same order if $\hat{M}[x,y] > 0$, when $M[x,y] = 1$ and, $\hat{M}[x,y] < 0$, when $M[x,y] = 0$.

Theorem 3: Let M be the matrix of a function f. Let d be the smallest number such that there are d rank 1 matrices O_i of order $2^n \times 2^n$, and $O = \sum_{i=1}^{d} O_i$ is an approximation of M. Then

$$\lceil \log d \rceil \leq \widetilde{C}_f \leq \lceil \log d \rceil + 1$$

Proof Sketch: Let O be an approximation of M such that $O = \sum_{i=1}^{d} O_i$, where each O_i is a rank 1 real matrix. Since O_i is a rank 1 matrix, $O_i = a_i \times b_i^T$ for some a_i, and $b_i \in R^{2^n}$. Let $\varphi_0(x) = (a_1(x), a_2(x), \cdots a_d(x))$, and $\varphi_1(x) = (b_1(x), b_2(x), \cdots b_d(x))$. $O[x_0,x_1] = \langle \varphi_0(x), \varphi_1(x) \rangle$ ($\langle s,t \rangle$ is the inner product of vectors s and t). We now have an arrangement $Arr(H)$ in R^d where H consists of the hyperplanes $\varphi_0(x_0)$, and this arrangement satisfies the matrix M.

In a similar way, given an arrangement of hyperplanes in R^d, we can find d rank 1 real matrices whose sum approximates the matrix M. ▪

A Logarithmic Lower Bound

Theorem 4: There exists a function f such that $\lceil \log_2 n \rceil \leq \widetilde{C}_f \leq \lceil \log_2 n \rceil + 1$.

Proof: Consider the function f defined as

$f(x,y) = \text{bin}(x)\text{th bit of } y \quad \text{for } 0 \leq \text{bin}(x) \leq n-1$
$= 0 \text{ otherwise.}$

It can be shown that if $Arr(H)$ is an arrangement of 2^n hyperplanes that satisfies the matrix M, then there is $H' \subseteq H$, such that $|H'| = n$, and $Arr(H')$ has 2^n distinct regions. The number of distinct regions in any arrangement of n hyperplanes in R^d is bounded by $\sum_{i=0}^{d} \binom{n}{i}$ [Bu]. Hence, $d \geq n$. This gives us the required lower bound.

Since any arrangement of d hyperplanes in general position in R^d contains 2^d regions, we also achieve our upper bound. ▪

The theorem can be easily extended to yield a complexity hierarchy for $0 \leq \widetilde{C} \leq \lceil \log n \rceil$.

A Lower Bound for 1-tape Probabilistic Turing Machines:

A 1-tape probabilistic Turing machine M is said to accept (reject) a string x in time t if the probability of the event "M, started in its initial configuration with input x, will enter an accepting (rejecting) configuration after at most t steps" is greater than $\frac{1}{2}$ [Gi]. [Yao 1] has obtained an $\Omega(n \log n)$ lower bound on the time required by certain 1-tape probabilistic Turing machines. However, the definition of acceptance used in [Yao 1] is more restrictive (bounded error), and the proofs use the restriction in an essential way. As an application of our results, we can prove the following.

Theorem 5: Let $L = \{x\#0^n\#y \mid |x|=|y|=n,\ x,y \in \{0,1\}^*,$ bit $\mathrm{bin}(y)$ of x exists and is 1$\}$. Then, any probabilistic 1-tape Turing machine (PTM) that accepts L uses $\Omega(n \log n)$ steps for some input of length n.

sketch of the proof: suppose, by contradiction, that M is a 1-tape PTM that accepts L in time $T(n)=o(n \log n)$ - i.e., for any c, for any input of length n, n sufficiently large, it uses less than $cn \log n$ steps (for any guess string). Then, from the computation of M on input $x\#0^n\#y$, with $x \in \{0,1\}^n$, $y \in \{0,1\}^{\log n}$, we produce a probabilistic protocol for the function f of theorem 4. The protocol uses $o(\log n)$ bits, yielding the contradiction that proves our claim. There is a technical difficulty in doing this: while it is easy to show that at each boundary of the string of n 0's there must be a guess string that causes a crossing sequence of length $O(\log n)$, it is not clear that any guess string will cause many long crossing sequences (there are $2^{T(n)}$ guess strings, but only $O(n)$ boundaries).

We can show, by a 'cut and paste' argument, that, if we assume that $T(n)<cn \log n$, then for any guess sequence g, in the computation of M with input $w=x\#0^n\#y$, using g as a guess string, there must be a crossing sequence in the middle of 0's that is *short* (length $<\varepsilon \log n$) and *close* to the first # (at a distance less than n^δ from the right #), where ε and δ can be chosen to be sufficiently small. Using the short crossing sequences and the (short) address of their position in the string one can construct a probabilistic protocol of complexity less than $\log n$ to compute the function f of Theorem 4. This gives us the desired contradiction.

4. CONCLUSIONS AND OPEN PROBLEMS:

Our results start a theory of probabilistic information transfer for unbounded error protocols. We provided interesting characterizations, some surprisingly efficient protocols, and a nontrivial lower bound.

It is pleasing that the basic questions about probabilistic information transfer are

mathematically interesting. Approximations of matrices by matrices of rank 1 (in a different metric) play an important role in numerical analysis [GoVL], and the decomposition of Euclidean space by hyperplanes is a classical geometric problem [Bu] [Za]. Our lower bounds follow from the basic properties of these objects. Strengthening them would be equivalent to settling certain mathematical problems that are interesting on their own.

The main remaining open problem is the optimality of our lower bound: Can one prove superlogarithmic lower bounds, or do all functions have low complexity? We have done little to settle the problem. On the positive side: consider the problem of verifying whether x and y have the Hamming distance d for some $d \in \{0,\ldots,n\}$. Protocol 2 in the appendix achieves $O(\log n)$ information transfer for this problem. Similar techniques yield $O(\log n)$ protocols for other problems. But, the technique fails for the function defined by a Hadamard matrix. We conjecture that this function has maximal (linear) probabilistic communication complexity. Proving this, however, seems to be difficult. Our lower bound proof uses counting of regions in R^d: a linear lower bound results from a choice of 2^n requirements, that would require the existence of $2^{2^{O(n)}}$ other regions in any arrangement of 2^n hyperplanes that satisfies these 2^n requirements. The choice of orthogonal requirements corresponding to the Hadamard matrix of order $2^n \times 2^n$ seems to be a suitable one, and hence the conjecture.

5. APPENDIX

Protocol 1 (Equality)

The following set of 2-dimensional planes p_x, and points q_y in R^3 define a protocol for computing $I(x,y)$. Normalization of the coefficients of these planes and points yield a 2-bit protocol for computing $I(x,y)$.

Let $m = 2^n$, $\varepsilon = \frac{1}{m^{m+3}}$, and $L_k = \sum_{j=0}^{k} \frac{1}{m^j}$.

$p_x(1) = 1;$
$q_y(1) = \mathrm{bin(y)} L_{\mathrm{bin}(y)+1} - (\mathrm{bin}(y)+1) L_{\mathrm{bin}(y)} - -(\mathrm{bin}(y)+1) m\varepsilon + \varepsilon;$
$p_x(2) = \mathrm{bin}(x)+1;$
$q_y(2) = L_{\mathrm{bin}(y)} - L_{\mathrm{bin}(y)+1} + m\varepsilon;$
$p_x(3) = L_{\mathrm{bin}(x)+1};$
$q_y(3) = 1.$

It is easy to verify that $\sum_i p_x(i) q_y(i) > 0$ if $x = y$, and $\sum_i p_x(i) q_y(i) < 0$ if $x \neq y$.

Protocol 2 (Verification of Hamming distance)

Let h_d, for some $d \in \{0,1,\ldots,n\}$, be such that $h_d(x,y) = 1$ iff the Hamming distance between x and y is d. The following protocol computes h_d for $d = \frac{n}{2}$. Protocols for other d can be devised similarly.

A sends two bits of its input x along with their addresses. Each pair of bits is equally likely to be selected. At B, after these two bits are received, one of the two following events Event I or Event II occurs, such that Event I happens with probability $\frac{1}{n}$.

Event I: Output 1 if the Hamming distance between the two bits received

and the corresponding bits of y is 1

Output 0 otherwise.

Event II: Output 1 with probability $(\frac{1}{2} - \frac{(n-2)(n+2)}{2n^2(n-1)} - \frac{1}{n^4})/(1 - \frac{1}{n})$.

It can be verified that this protocol indeed computes the function $h_{\frac{n}{2}}$.

Acknowledgements: We have benefitted from the work of P. Berman, G. Schnitger, and V. K. Prasanna Kumar, and from extensive discussions with the first two. In particular, the idea of using rank 1 matrices is due to Prasanna Kumar and Schnitger, who obtained an (essentially equivalent) lower bound. The efficient protocol for the verification of Hamming distance is suggested by P. Berman, who also helped in clarifying some of our ideas. We are grateful for their help. We also bothered a large number of mathematicians, believing that they must have solved our combinatorial problems long ago. A partial list of these helpful mathematicians is P. Rejto, R.M. Hardt, B. Grunbaum (who suggested interesting ways to try to solve it), J. Edmonds, S. Friedland, R. Bland, L. Vaserstein. We appreciate their patience, and the time they spent on our questions.

6. REFERENCES:

[Bu] R. C. Buck, "Partition of Space", Amer. Math. Monthly, 50 (1943), 541-544.

[B-O] M. Ben-Or, private communication, May 1984

[DGSch] P. Duris, Z. Galil, and G. Schnitger, "Lower Bounds on Communication Complexity", 16th ACM STOC, 81-91, 1984.

[FL] J. Folkman, and J. Lawrence, "Oriented Matroids", J. of Combinatorial Theory, Series B 25, 199-236 (1978).

[Gi] J. T. Gill, III, " Computational Complexity of Probabilistic Turing Machines", SIAM J. Computing, 6 (1977), 675-695.

[GoVL] G. H. Golub, and C. F. Van Loan, Matrix Computations, The Johns Hopkins University Press, 1983.

[JKS] J. Ja'Ja', V. K. Prasanna Kumar and J. Simon: "Information Transfer under Different Sets of Protocols", SIAM J. on Computing, to appear.

[PS] C. H. Papidimitriou, and M. Sipser, "Communication Complexity", Proc. 14th ACM STOC, 330-337, 1982.

[Yao 1] A. C.-C. Yao, "A Lower Bound to Palindrome Recognition by Probabilistic Turing Machines", Computer Science Department, Stanford University, Dec. 1977.

[Yao 2] A. C.-C. Yao, "Some Complexity Questions related to Distributed Computations", Proc. 11th ACM STOC, 209-213, 1979.

[Yao 3] A. C.-C. Yao, "Lower Bounds by Probabilistic Arguments", Proc. 24th IEEE FOCS, 420-428, 1983.

[Za] T. Zaslavsky, Facing up to arrangements: face-count formulas for partitions of space by hyperplanes, Memoirs Amer. Math. Soc. 154, Providence, 1975.

25th Annual

Symposium on Foundations of Computer Science

(Formerly called the Annual Symposium on Switching and Automata Theory)

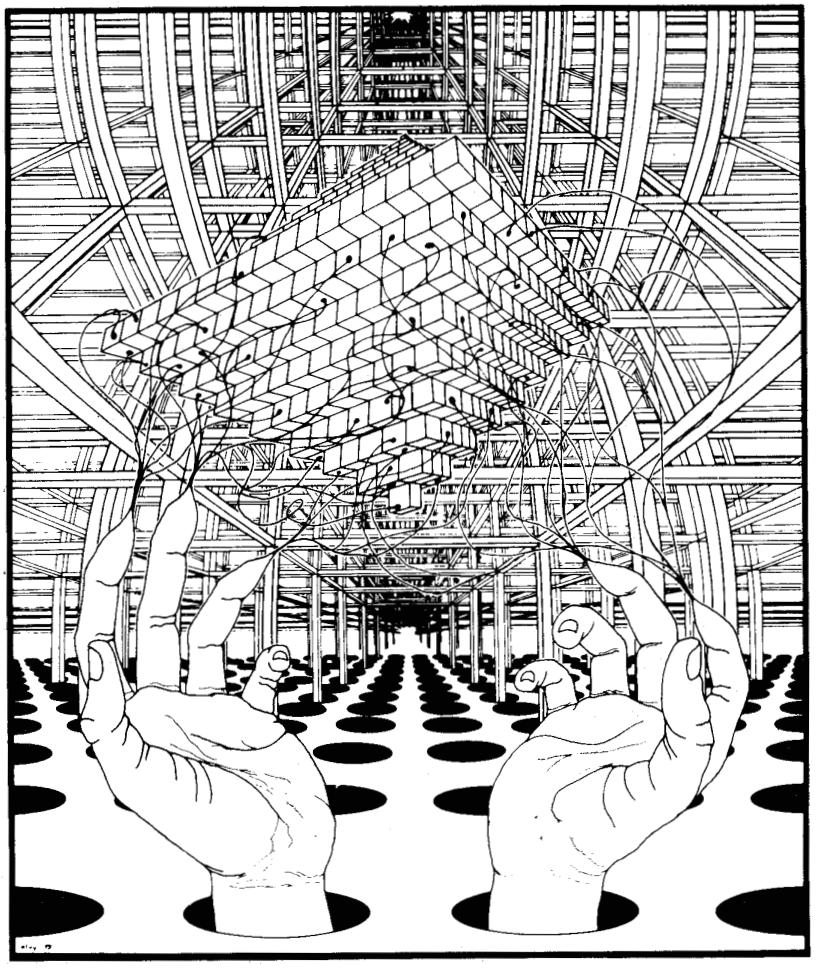

OCTOBER 24–26, 1984 IEEE 84CH2085-9

sponsored by
the IEEE Computer Society's Technical Committee on
Mathematical Foundations of Computing

ISSN 0272-5428
IEEE CATALOG NUMBER 84CH2085-9
LIBRARY OF CONGRESS NUMBER 80-646634
IEEE COMPUTER SOCIETY ORDER NUMBER 591
ISBN 0-8186-0591-X

IEEE COMPUTER SOCIETY PRESS

Symposium on Foundations of Computer Science.
Annual Symposium on Foundations of Computer Science [papers]. 16th- 1975-
[New York, Institute of Electrical and Electronics Engineers]

v. ill. 28 cm.

Annual.
Vols. for 1975- sponsored by IEEE Computer Society, Technical Committee on Mathematical Foundations of Computing and the ACM Special Interest Group for Automata and Computability Theory, and various universities.

Symposium on Foundations of Computer Science. Annual Symposium on Foundations of Computer Science ... (Card 2)

Continues: Annual Symposium on Switching & Automata Theory, ISSN 0272-4847.
Key title: Annual Symposium on Foundations of Computer Science, ISSN 0272-5428.

1. Switching theory—Congresses. 2. Machine theory—Congresses. 3. Electronic data processing—Congresses. I. IEEE Computer Society. Technical Committee on Mathematical Foundations of Computing. II. ACM Special Interest Group for Automata and Computability Theory. III. California. University. Dept. of Electrical Engineering and Computer Sciences. IV. Title.

QA268.5.S9a 519.4 80-646634
MARC-S

Library of Congress

ISSN 0272-5428
IEEE Catalog Number 84CH2085- 9
Library of Congress Number 80-646634
IEEE Computer Society Order Number 591
ISBN 0-8186-0591-X (paper)
ISBN 0-8186-4591-1 (microfiche)
ISBN 0-8186-8591-3 (casebound)

Order from: IEEE Computer Society
Post Office Box 80452
Worldway Postal Center
Los Angeles, CA 90080

IEEE Service Center
445 Hoes Lane
Piscataway, NJ 08854

1884 1984 A CENTURY OF ELECTRICAL PROGRESS THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Twenty-Fifth Annual Symposium on Foundations of Computer Science Singer Island, Florida October 24-26, 1984

Table of Contents

Wednesday Morning, October 24, 1984

Session 1: Leslie Valiant, Chair

Wednesday Afternoon, October 24, 1984

Session 2: Larry Stockmeyer, Chair

Thursday Morning, October 25, 1984

Session 3: Richard M. Karp, Chair

Thursday Afternoon, October 25, 1984

Session 4: Michael O'Donnell, Chair

Friday Morning, October 26, 1984

Session 5: Leo Guibas, Chair

Friday Afternoon, October 26, 1984

Session 6: Martin Tompa, Chair