The Work of Peter W. Shor

Ronald Graham

Much of the work of Peter Shor has a strong geometrical flavor, typically coupled with deep ideas from probability, complexity theory or combinatorics, and always woven together with brilliance and insight of the first magnitude. Due to the space limitations of this note, I will restrict myself to brief descriptions of just four of his remarkable achievements, (unfortunately) omitting discussions of his seminal work [8] on randomized incremental algorithms (of fundamental importance in computational geometry) and his provocative results in computational biology on self-assembling virus shells.

1 Two-dimensional discrepancy, minimax grid matchings and online bin packing

The minimax grid matching problem is a fundamental combinatorial problem arising the the average case analysis of algorithms. To state it, we consider a square S of area N in the plane, and a regularly spaced $\sqrt{N} \times \sqrt{N}$ array G (=grid) of points in S. Let P be a set of N points selected independently and uniformly in S. By a perfect matching of P to G we mean a 1 - 1 map $\lambda : P \to G$. For each selection P, define $L(P) = \min_\lambda \max_{p \subset P} \mathbf{d}(p, \lambda(p))$, where λ ranges over all perfect matchings of P to G, and $\mathbf{d}$ denotes Euclidean distance.

Theorem [Shor [24], Leighton/Shor [21]]
With very high probability,

$$\mathbf{E}(L(P)) = \Theta((logN)^{3/4})$$

The proof is very intricate and ingenious, and contains a wealth in new ideas which have spawned a variety of extensions and generalizations, notably in the work of M. Talagrand [30] on majorizing measures and discrepancy.

A classical paradigm in the analysis of algorithms is the so-called bin packing problem [10], in which a list $W = (w_1, w_2, \dots, w_n)$ of "weights" is given, and we are to required to pack all the w_i into "bins" with the constraint that no bin can contain a weight total of more than 1. Since it is NP-hard to determine the minimum number of bins which W requires for a successful packing
(or even to decide if this minimum number is 2!), extensive efforts have been made for finding good approximation algorithms for producing near-optimal packings.

In the Best Fit algorithm, after the first i weights are packed, the next weight w_{i+1} is placed into the bin in which it fits best, i.e., so that the unused space in that bin is less than it would be in any other bin. (This is actually an online algorithm). In his thesis [23] , Shor proved the very surprising (and deep) result that when the w_i are chosen uniformly at random from $[0, 1]$, then with very high probability, the amount of wasted space has size $\Theta(n^{1/2}(\log n)^{3/4})$.

An "up-right" region $R = R(f)$ of the square S is defined as the region in S lying above some continuous monotonically non-increasing function f (e.g., S is itself up-right). If P is a set of N points chosen uniformly and independently at random in S, we can define the *discrepancy* $\Delta(R) = || R \cap P | - area(R) |$. An old problem in mathematical statistics (from the 1950's; see [5]) was the estimation of $\sup_R \Delta(R)$ over all up-right regions of S. This was finally answered by Leighton and Shor in [24, 21], and it is now known that

$$\sup_R \Delta(R) = \Theta(N^{1/2}(\log N)^{3/4})).$$

The preceding results give just a hint of the numerous applications these fertile techniques have found to such diverse areas as pseudo-random number generation, dynamic storage allocation, wafer-scale integration and two-dimensional bin packing (see [9, 20, 17]).

2 Davenport-Schinzel Sequences

A Davenport-Schinzel sequence $DS(n, s)$ is a sequence $U = (u_1, u_2, ..., u_m)$ composed of n distinct symbols such that $u_i \neq u_{i+1}$ for all i, and such that U contains no alternating subsequence of length $s + 2$, i.e., there do not exist indices $i_1 < i_2 < \ldots < i_{s+2}$ such that $u_{i_1} = u_{i_3} = u_{i_5} = \ldots = a \neq b = u_{i_2} = u_{i_4} = \ldots$. We define

$$\lambda_S(n) = \max\{m : (u_1, \ldots , u_m) \text{ is a } DS(n, s) - \text{sequence}\}.$$

Davenport-Schinzel sequences have turned out to be of central importance in computational and combinatorial geometry, and have found many applications in such areas as motion planning, visibility, Voronoi diagrams and shortest path algorithms. It is known that $DS(n, s)$-sequences provide a combinatorial characterization of the lower envelope of n continuous univariate functions, each pair of which intersect in at most s points. Hence, $\lambda_s(n)$ is just the maximum number of connected components of the graphs of such functions, and accurate estimates of $\lambda_s(n)$ can often be translated into sharp bounds for algorithms which depend on function minimization. It is trivial to show that $\lambda_1(n) = n$ and $\lambda_2(n) = 2n - 1$. The first surprise came when it was shown [15] that $\lambda_3(n) = \Theta(n\alpha(n))$ where $\alpha(n)$ is defined to be the functional inverse of the Ackermann function $A(t)$, i.e., $\alpha(n) := \min\{t : A(t) \geq n\}$. Note that $\alpha(n)$ is an extremely slowly growing function of n since A is defined as follows:

$$A_1(t) = 2t,\ t \geq 1\,, \quad \text{and} \quad A_k(t) = A_{k-1}(A_k(t-1))\,, \quad k \geq 2,\ t \geq 2\,.$$

Thus, $A_2(t) = 2^t$, $A_3(t)$ is an exponential tower of n 2's, and so on. Then $A(t)$ is defined to be $A_t(t)$. The best bounds for $\lambda_s(n)$, $s > 3$ in [15] were rather weak. This was remedied in [1] where Shor and his coauthors managed to show by extremely delicate and clever techniques that $\lambda_4(n) = \Theta(n2^{\alpha(n)})$. Thus, $DS(n,4)$-sequences can be *much* longer than $DS(n,3)$-sequences (but are still only *slightly* non-linear). In addition, they also obtained almost tight bounds on all other $\lambda_s(n), s > 4$.

3 TILING $\mathbb{R}^n$ WITH CUBES

In 1907, Minkowski made the conjecture (in connection with his work on extremal lattices) that in any *lattice* tiling of $\mathbb{R}^n$ with unit n-cubes, there must be two cubes having a complete facet ($= (n-1)$-face) in common. This was generalized by O. Keller [18] in 1930 to the conjecture that *any* tiling of $\mathbb{R}^n$ by unit n-cubes must have this property. This was confirmed by Perron [22] in 1940 for $n \leq 6$, and shortly thereafter, Hajós [14] proved Minkowski's original conjecture for all n. However, in spite of repeated efforts, no further progress was made in proving Keller's conjecture for the next 50 years. Then in 1992, Shor struck. He showed (with his colleague J. Lagarias) that in fact Keller's conjecture is *false* for all dimensions $n \geq 10$. They managed to do this with an very ingenious argument showing that certain special graphs suggested by Corrádi and Szabó [11] of size 4^n, must always have cliques of size 2^n (contrary to the prevailing opinions then), from which it followed that Keller's conjecture must fail for $\mathbb{R}^n$. The reader is referred to [19] for the details of this combinatorial gem, and to [29] for a fascinating history of this problem. I might point out that this is another example of an old conjecture in geometry being shattered by a subtle combinatorial construction, an earlier one being the recent disproof of the Borsuk conjecture by Kahn and Kalai [16]. It is still not known what the truth for Keller's conjecture is when $n = 7, 8$, or 9.

4 QUANTUM COMPUTATION

It has been generally believed that a digital computer (or more abstractly, a Turing machine) can simulate any physically realizable computational device. This, in fact is the thrust of the celebrated Church - Turing thesis. Moreover, it was also assumed that this could always be done in an efficient way, i.e., involving at most a polynomial expansion in the time required. However, it was first pointed out by Feynman [13] that certain quantum mechanical systems seemed to be extremely difficult (in fact, impossible) to simulate efficiently on a standard (von Neumann) computer. This led him to suggest that it might be possible to take advantage of the quantum mechanical behavior of nature itself in designing a computer which overcame these difficulties. In fact, in doing so, such a "quantum" computer might be able to solve some of the classical difficult problems much more efficiently as well. These ideas were pursued by Benioff [4], Deutsch [12], Bennett [2] and others, and slowly, a model of quantum computation began to evolve. However, the first bombshell in this embryonic field occurred when Peter Shor [25, 26] in 1994 announced the first *significant* algorithm for such a hypothetical quantum computer, namely a method for factoring an arbitrary composite integer N in

$$c(\log N)^2 \log\log N \log\log\log N$$

steps. This should be contrasted with the best current algorithm on (classical) digital computers whose best running time estimates grow like

$$\exp(cN^{1/3}(\log N)^{2/3}).$$

Of course, no one has yet ruled out the possibility that a polynomial-time factoring algorithm exists for classical computers (cf. the infamous P vs. NP problem), but it is felt by most knowledgeable people that this is extremely unlikely. In the same paper, Shor also gives a polynomial-time algorithm for a quantum computer for computing discrete logarithms, another (apparently) intractable problem for classical computers.

There is not space here to describe these algorithms in any detail, but a few remarks may be in order. In a classical computer, information is represented by binary symbols 0 and 1 (bits). An n-bit memory can exist in any of 2^n logical states. Such computers also manipulate this binary data using functions like the Boolean AND and NOT. By contrast, a quantum bit or "qubit" is typically a microscopic system such as an electron (with its spin) or a polarized photon. The Boolean states 0 and 1 are represented by (reliably) distinguishable states of the qubit, e.g., $|0\rangle \leftrightarrow$ spin $\frac{1}{2}$ and $|1\rangle \leftrightarrow$ spin $-\frac{1}{2}$. However, according to the laws of quantum mechanics, the qubit can also exist in a continuum of intermediate states, or "superpositions", $\alpha|0\rangle + \beta|1\rangle$ where α and β are complex numbers satisfying $|\alpha|^2 + |\beta|^2 = 1$.

More generally, a string of n qubits can exist in any state of the form

$$\psi = \sum_{x=00\ldots0}^{11\ldots1} \psi_x |x\rangle$$

where the ψ_x are complex numbers such that $\sum_x |\psi_x|^2 = 1$. In other words, a quantum state of n qubits is represented by a unit vector in a 2^n-dimensional complex Hilbert space, defined as the tensor product of the n copies of the 2-dimensional Hilbert space representing the state of a single qubit. It is the exponentially large dimensionality of this space which distinguishes quantum computers from classical computers. Whereas the state of a classical system can be completely described by separately specifying the state of each part, the overwhelming majority of states in a quantum computer are "entangled," i.e., not representable as a direct product of the states of its individual qubits. As stated in [3], "the ability to preserve and manipulate entangled states is the distinguishing feature of quantum computers, responsible both for their power and for the difficulty in building them."

The crux of Shor's factoring algorithm (after reducing the problem of factoring N to that of determining for a random X coprime to N, the order of $X\,(modulo\ N)$, is a brilliant application of the discrete Fourier transform in such a way as to have all the incorrect candidate orders (quantum mechanically) cancel out, leaving only (multiples) of the correct order of X appearing (with high probability) when the output is finally measured. I heartily recommend that the reader consult the paper of Shor in this Volume, or [26, 31] for more details.

Of course, complicated quantum systems are delicate creatures and any substantial interaction with the external environment can cause rapid "decoherence," which then can result in the system collapsing to some classical state, thereby prematurely terminating the ongoing computation. This was the basis for the strong initial skepticism that any serious quantum computer could actually ever be built. However, Shor's subsequent contributions changed this situation substantially. His paper [27] in 1995 announced the discovery of quantum error-correcting codes, cutting through some widely held misconceptions about quantum information, and showing that suitable measurements of a quantum system can acquire sufficient information for detecting and correcting errors *without* disturbing any of the encoded information. These ideas were further developed in [6, 7] to produce a new theory of quantum error-correcting codes for protection against multiple errors, using clever ideas from orthogonal geometry and properties of the recently discovered ordinary (as opposed to quantum) codes over $GF(4)$.

Finally, any quantum computer which is actually built will be composed of components which are not completely reliable. Thus, it will be essential to create algorithms which are "fault-tolerant" on such computers. In yet another path-breaking paper [28], Shor in 1996 showed how this indeed could be done.

Not only does Peter Shor's work on quantum computation during the past four years represent scientific achievements of the first rank, but in my mind it holds out the first real promise that non-trivial quantum computers may actually exist in our lifetimes.

REFERENCES

[1] P. K. Agarwal, M. Sharir and P. Shor, Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences, *J. Comb. Theory (A)* 52 (1989), 228-274.

[2] C. H. Bennett, Logical reversibility of computation, *IBM J. Res. Develop.* 17, (1973), 525-532.

[3] C. H. Bennett and P. W. Shor, Quantum information theory, *(to appear)*

[4] P. Benioff, Quantum mechanical models of Turing machines that dissipate no energy, *Phys. Rev. Letters* 48 (1982), 1581-1585.

[5] J. Blum, On convergence of empirical distribution functions, *Ann. Math. Stat* 26 (1955), 527-529.

[6] A. R. Calderbank and P. W. Shor, Good quantum error-correcting codes exist, *Phys. Review A* 54 (1995), 1098-1106.

[7] A. R. Calderbank, E. M. Rains, P. W. Shor and N. J. A. Sloane, Quantum error correction via codes over $GF(4)$, *IEEE Trans. Inf. Theory* 44 (1998), 1369-1387.

[8] K. L. Clarkson and P. W. Shor, Applications of random sampling in computational geometry, II, *Discrete and Comput. Geom.* 4 (1989), 387-421.

[9] E. G. Coffman, Jr. and F. T. Leighton, A provably efficient algorithm for dynamic storage allocation, *J. Comput. System Sci.* 38 (1989), 2-35.

[10] E. G. Coffman, Jr., M. R. Garey, and D. S. Johnson, Approximation algorithms for bin-packing – an updated survey, *Algorithm design for computer system design*, CISM Courses and Lectures 284, Springer-Verlag, Vienna (1984), 49-106.

[11] K. Corrádi and S. Szabó, A combinatorial approach for Keller's conjecture, *Period. Math. Hungar.* 21 (1990), 91-100.

[12] D. Deutsch, Quantum theory, the Church-Turing principle, and the universal quantum computer, *Proc. Royal Soc. London* A400 (1985), 97-117.

[13] R. Feynman, Simulating physics with computers, *Internat. J. Theoret. Phys.* 21 (1982), 467-488.

[14] G. Hajós, Über einfache und mehrfache Bedeckung des n-dimensionalen Raumes mit einen Würfelgitter, *Math. Z.* 47 (1942), 427-467.

[15] S. Hart and M. Sharir, Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes, *Combinatorica* 6 (1986), 151-177.

[16] J. Kahn and G. Kalai, A counterexample to Borsuk's conjecture, *Bull. Amer. Math. Soc. (New Series)* 29 (1993), 60-62.

[17] R. M. Karp, M. Luby and A. Marchetti-Spaccamela, Probabilistic analysis of multi-dimensional bin packing problems, *Proceedings 25th ACM Symp. on Theory of Computing* (1984), 289-298.

[18] O. H. Keller, Über einfache und mehrfache Bedeckung des n-dimensionalen Raumes mit einen Würfelgitter, *J. Reine Angew. Math.* 163 (1930), 231-248.

[19] J. C. Lagarias and P. W. Shor, Keller's cube-tiling conjecture is false in high dimensions, *Bull. Amer. Math. Soc.* (New Series)27 (1992), 279-283.

[20] F. T. Leighton and C. E. Leiserson, Wafer-scale integration of systolic arrays, *IEEE Trans. on Computers* C-34 (1985), 448-461.

[21] T. Leighton and P. Shor, Tight bounds for minimax grid matching with applications to the average case analysis of algorithms, *Combinatorica* 9 (1989), 161-187.

[22] O. Perron, Über lückenlose Ausfüllung des n-dimensionalen Raumes durch kongruente Würfel, I, II, *Math. Z.* 46 (1940), 1-26, 161-180.

[23] P. W. Shor, *Random Planar Matching and Bin Packing*, Ph. D. Thesis, MIT Math. Dept., 1985.

[24] P. W. Shor, The average-case analysis of some on-line algorithms for bin packing, *Combinatorica* 6 (1986), 179-200.

[25] P. W. Shor, Algorithms for quantum computation: discrete logarithms and factoring, *Proc. 35th Symp. Foundations of Comp. Sci.*, IEEE Computer Society Press, (1994), 124-134.

[26] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. Computing* 26 (1997), 1484-1509.

[27] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Phys. Review A* 52 (1995), R2493-R2496.

[28] P. W. Shor, Fault-tolerant quantum computation, *Proceedings 37th Symp. Foundations Comp. Science* IEEE Computer Society Press, Los Alamitos, CA (1996), 56-65.

[29] S. Stein and S. Szabó, Algebra and Tiling, Carus Math. Monograph no. 25, Math. Assoc. America, Washington, 1994.

[30] M. Talagrand, Matching theorems and empirical discrepancy computations using majorizing measures, *J. Amer. Math. Soc.* 7 (1994), 455-537.

[31] C. Williams and S. Clearwater, *Explorations in Quantum Computing*, Springer-Verlag, Santa Clara, CA, 1998.

Ronald Graham
AT&T Labs
Florham Park, NJ
and
UCSD
La Jolla, CA
USA

PETER W. SHOR

Information Sciences Center, AT&T Labs–Research, Florham Park, NJ, USA

Born:	August 14, 1959, New York City, USA
Nationality:	US Citizen
Marital Status:	married, one daughter
1977–1981	Undergraduate at California Institute of Technology, Pasadena
1981–1985	Ph.D. in Mathematics, Massachusetts Institute of Technology
1985–1986	Postdoctoral Fellow at Mathematical Sciences Research Institute, Berkeley
1986–1996	Member of Technical Staff, AT&T Bell Labs, Murray Hill
1996	Principal Research Staff Member, AT&T Labs, Florham Park
Fields of Interest:	Theoretical Computer Science, Combinatorics

Ronald Graham and Peter W. Shor