tain predominantly shallow inclinations. These two explanations, long-lasting octupole fields or preferentially equatorial paleolocations of the continents, could each have very important implications for paleogeography, as well as for the generation of the magnetic field and the evolution of Earth's core or alternatively for mantle dynamics and true polar wander (that is, a tumbling of the body of Earth with respect to the rotation axis) (*4*).

The most serious immediate implications arise for paleogeography and for paleomagnetists attempting to reconstruct ancient continental configurations. The figure, adapted from Kent and Smethurst (*2*), shows inclination values as a function of latitude for a purely geocentric axial dipole (GAD) model and the 0.25G3 model with the addition of a small quadrupole field (10% of the dipole, or G2 = 0.1). The magnitude of the octupole field is fairly critical for producing a good fit to the observations, but the magnitude of the quadrupole field (G2) is not, because its effects are antisymmetric about the equator and averaged by the analysis method that combines hemispheres. A paleolatitude calculated with the GAD model could differ by up to about 18° from the real paleolatitude if the field were partly nondipolar in this fashion. And if, coincidentally, two coeval paleomagnetic sites, one with an inclination I = +40° and the other with I = −50°, were to be compared, the relative paleolatitude difference estimated with the GAD model might be erroneous by more than 30°. Are there any indications that errors of such magnitude could be present? An affirmative answer would mean, of course, that there is some independent evidence for the paleolatitudes, and there's the rub: For most of Precambrian time, quantitative estimates of paleolatitude are fairly well limited to paleomagnetic data only. However, for the late Paleozoic, there is a long-standing paleogeographic problem that could well be solved by assuming a nondipole field contribution, and it could also lead to an independent test of the proposal of Kent and Smethurst.

A large continent, not internally deformed or disrupted since 250 million years ago, can provide a test of the 0.25G3 model if it (i) covered a large range of paleolatitudes and (ii) has yielded numerous coeval paleomagnetic results from widespread sites. For the Paleozoic (or Precambrian), no single continent existing today qualifies, but a supercontinental assembly such as Pangea provides some clues, granted certain assumptions about how it was configured. This mention of "assumptions" reveals, in fact, the "problem": Paleomagnetists have long been forced to modify the classical Pangea A configuration, if they wanted to keep adhering to the GAD model [for discussion, see (*5*)]. From this modification resulted proposals such as those for Pangea B or C, in which the Gondwana continents are located 3500 km (or more) to the east of their Pangea A location with respect to a fixed North America–Europe continent and some 10° to 30° to the north. However, nonpaleomagnetist paleogeographers have not adopted Pangea B or C, and if the 0.25G3 model is applicable to times in which Pangea existed (~250 million years ago), they may have been right.

Kent and Smethurst include in their paper a discussion of the implications of their 0.25G3 model for the evolution of Earth's core and the generation of the magnetic field. They speculate that growth of the solid inner core may have reached a threshold size for stabilizing the GAD field (*6*) as late as some 250 million years ago. However, their alternative explanation of the shallow bias in $|I|$, namely that continents in the Precambrian and Paleozoic were perhaps preferentially located in lower latitudes, is also of great interest. Anderson (*7*) has suggested that supercontinents induced long-lived geoid highs that could influence Earth's moment of inertia tensor in such a way as to cause true polar wander. This would imply that continental movements were "frustrated" in their attempts to reach polar latitudes, as the latter continuously receded to new locations orthogonal to the bulk of the continental positions. As interest in true polar wander in older geological times has recently been revived (*4*), this alternative will undoubtedly see considerable attention in the coming years.

References

1. M. W. McElhinny, P. L. McFadden, R. T. Merrill, *J. Geophys. Res.* **101**, 25007 (1996).
2. D. V. Kent and M. Smethurst, *Earth Planet. Sci. Lett.*, in press.
3. M. E. Evans, *Nature* **262**, 676 (1976).
4. J. L. Kirschvink, R. L. Ripperdan, D. A. Evans, *Science* **277**, 541 (1997).
5. R. Van der Voo, *Paleomagnetism of the Atlantic, Tethys and Iapetus Oceans* (Cambridge Univ. Press, Cambridge, 1993), pp. 117–121; G. Muttoni, D. V. Kent, J. E. T. Channell, *Earth Planet. Sci. Lett.* **140**, 97 (1996).
6. R. Hollerbach and C. A. Jones, *Nature* **365**, 541 (1993); *Phys. Earth Planet. Inter.* **87**, 171 (1995).
7. D. L. Anderson, *Nature* **297**, 391 (1982).

PERSPECTIVES: QUANTUM COMPUTING

Beyond Factorization and Search

Lov K. Grover

All of the computers manufactured so far are based on the laws of classical physics. About 15 years ago, Feynman, Bennett, Deutsch, and others observed that if a computer could be built around the laws of quantum physics instead of those of classical physics, it would result in an entirely different computational structure (*1*). For a long time, it was not obvious how such a device would differ from a classical computer. Finally in 1994, Shor discovered a quantum mechanical algorithm for factorization (*2*). This aroused a lot of interest—mathematicians had been looking for efficient factoring algorithms for several decades and had not found any. Furthermore, commonly used cryptographic codes, such as RSA (Rivest-Shamir-Adleman), had been designed on the assumption that no such algorithm existed. The idea behind Shor's factorization algorithm was something every solid-state physicist knows, namely that periodic quantum systems have special properties (see figure on next page). Shor's insight was to recall the observation that the factorization problem can be converted into one of estimating the periodicity of a sequence, something that quantum systems are very good at (*2*).

The next step took 2 years. In 1996, I discovered a quantum mechanical algorithm that could rapidly search an unsorted database of N items for a single item satisfying a given condition. This took about $\sqrt{N}$ steps, which was surprising because any classical algorithm, deterministic or probabilistic, will clearly need N tries. The idea behind the quantum search algorithm was that a quantum computer can be in a superposition of states and simultaneously examine multiple items. However, for an observer outside the quantum system to be able to determine which is the satisfying item, a process known as amplitude amplification is carried out. In about $\sqrt{N}$ elementary quantum mechanical operations, the amplitude of the desired item is substantially increased at the expense of other items (*3*). The discovery of this quantum search algorithm stimulated experimentalists to explore various physical systems for which this algorithm could be implemented. Very recently, fast search algorithms have been independently implemented by two research groups (*4*).

On the theoretical side, the obvious question is can the search algorithm be further

The author is at Bell Labs, Lucent Technologies. His address is 3C-404A, 600 Mountain Avenue, Murray Hill, NJ 07974, USA. E-mail: lkgrover@bell-labs.com

improved? Bennett *et al.* showed in 1994 that it was impossible to search an unsorted list of N items in fewer than about $\sqrt{N}$ steps (*5*). Later, Boyer *et al.* proved that the number of steps required by the quantum search algorithm was within 2.62% of optimal (*6*). By further refining the mathematics, Zalka proved late last year that the quantum search algorithm is exactly optimal and cannot be improved at all (*7*). Unlike the classical case in which it is elementary to show that no algorithm can isolate the desired item with certainty in less than N steps, the proofs of the quantum mechanical lower bound are very involved—much more so than the algorithm itself! This is a rare instance in science where a result can be simply stated, yet no one has been able to find an elementary explanation.

Could similar quantum mechanical techniques be used to solve any problem? There is at least one problem—the parity problem for which quantum mechanics cannot help substantially. In this problem, a binary function, which assumes only the values 0 or 1 at each point in the domain, is specified on a domain of points; the parity problem is to determine whether the number of points at which the function is 1 is even or odd. A classical algorithm would need to look at the value of the function at each one of the points to solve this problem. Late last year, Buhrman *et al.* showed through subtle communication complexity arguments that a quantum mechanical circuit would need at least $O(N/\log N)$ steps; that is, it could be at most a factor of about $\log N$ faster than a classical algorithm (*8*). This demonstrated for the first time a problem for which quantum mechanical algorithms would not be very useful. Earlier this year, it was independently proved by two groups that a quantum mechanical algorithm would need at least $N/2$ steps; both groups also found algorithms that accomplished exactly this (*9*, *10*). Hence, for the parity problem, a quantum algorithm could only be at most twice as fast as a classical algorithm.

The next question was whether most problems were of the search and factorization kind where quantum computing gives a substantial speed-up or of the parity kind where it does not substantially help. By an ingenious counting argument that consists of enumerating all possible functions versus all possible quantum circuits, Beals *et al.* found that almost all problems are of the parity kind; search and factorization are rare exceptions (*10*). A similar result has also been independently derived by Ozhigov using an information theory type of averaging over all possible problems (*11*). However, to mathematically carry out these calculations, the class of problems that has to be included is so large that it does not tell us anything about how useful quantum mechanical algorithms will be for problems of practical interest to computer scientists.

What other problems apart from search and factorization are quantum mechanical algorithms good for? The quantum search algorithm was recently shown to be a particular case of a large class of algorithms based on amplitude amplification. In fact, any algorithm that gives a probabilistic solution can be used as part of a quantum mechanical algorithm (*12*). The number of steps required by this quantum algorithm is the square root of the number of steps the classical probabilistic algorithm needs—just like quantum search. One example of such an application is search in the presence of partial information, for example, retrieving an image from a noisy data transmission.

Another natural application of quantum mechanical algorithms is in the field of statistics. Because a quantum mechanical system can simultaneously be in multiple states and thus simultaneously examine multiple pieces of data, it seems plausible that statistical applications might constitute a natural niche for quantum algorithms. Indeed, the amplitude amplification technique of quantum search adapts well to these applications as well. Efficient algorithms for estimating the mean and median of a population have recently been invented (*13*). The number of steps required by these quantum algorithms is the square root of the best possible classical algorithm. It was recently proved that these were within a logarithmic factor of the fastest possible quantum algorithm (*14*).

A major achievement of classical computer science has been the theory of NP-complete problems. These are problems that occur in a variety of different contexts ranging from optimization (traveling salesman problem) to theoretical physics (Ising model). No efficient solution to any of these has been found. A solution to any one of these will immediately yield efficient solutions to the whole class of problems (*15*). Furthermore, considering the number of researchers from different fields who have invested considerable time and effort on this problem, the consensus among computer scientists is that no efficient classical solutions to these problems exist (see figure, right panel). What about possible quantum mechanical solutions?

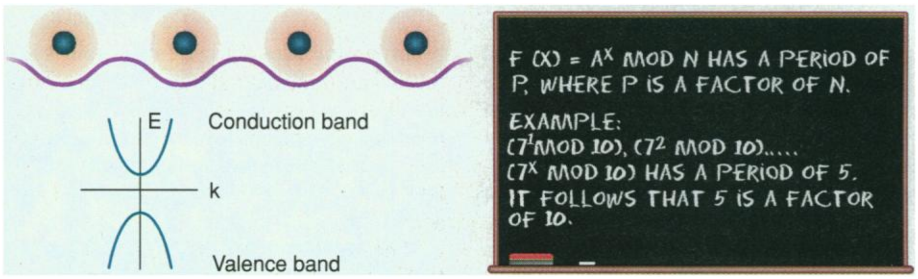

Quantum periodicity. The propagation of electrons in periodic structures **(left)** is one of the wonders of quantum mechanics. Shor's factorization algorithm uses quantum mechanics to estimate the periodicity of a sequence of numbers **(right)**. From this periodicity, the factors of a given number are detected.

The theory of quantum mechanical algorithms is very recent, and so there have been relatively few attempts to solve NP-complete problems with quantum computing. A variant of quantum search has been applied to graph coloring (an NP-complete problem). By making use of the structure of the problem, the quantum algorithm is able to solve the problem in a number of steps that is less than the square root of the total number of possibilities (*16*). However, the number of steps required is still exponential in the problem size, and by the yardstick of complexity theory, it is still not an efficient algorithm.

What if the laws of quantum mechanics were slightly different? In an intriguing paper, Abrams *et al.* have shown that if there was even the slightest amount of nonlinearity in quantum mechanics, no matter what, it would be possible to modify the amplitude amplification scheme of quantum search to obtain an efficient algorithm for NP-complete problems (*17*). Unfortunately, it is generally believed that such a nonlinearity probably does not exist because it would lead to faster-than-light communication, noncausality, and other violations of fundamental physical principles.

There are two main ideas that have resulted in efficient quantum mechanical algorithms: the first is periodicity estimation that leads to a factorization algorithm, and the second is amplitude amplification that leads to quantum search and a host of related algorithms. As physicists have observed in the last 70 years, there are several quantum mechanical phenomena that lead to puzzling and even paradoxical results. Most of these still remain to be investigated from a quantum computing perspective. It is possible that one of these might result in an algorithm that could solve NP-complete problems. Such a solution would greatly increase the interest in building a quantum computer.

References

1. D. Deutsch, in *The Fabric of Reality* (Penguin, New York, 1997), pp. 194–222.
2. P. W. Shor, in *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science* (IEEE Computer Society Press, Los Alamitos, CA, 1994), pp. 124–134.
3. L. Grover, *Phys. Rev. Lett.* **79**, 325 (1997); G. Brassard, *Science* **275**, 627 (1997).
4. I. Chuang *et al.*, *Phys. Rev. Lett.* **80**, 3408 (1998); J. Jones *et al.*, *Nature* **393**, 344 (1998); L. Grover, *Science* **280**, 228 (1998); J. Jones, *ibid.*, p. 229.
5. C. Bennett *et al.*, *SIAM J. Comput.* **26**, 1510 (1997).
6. M. Boyer *et al.*, *Fortschr. Phys.* **46**, 493 (1998), e-print available at http://xxx.lanl.gov/abs/quant-ph/9605034. Other e-prints referenced below can be obtained at xxx.lanl.gov with the equivalent URL.
7. C. Zalka, quant-ph/9711070.
8. H. Buhrman *et al.*, in *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing* (ACM Press, Dallas, TX, 1998), pp. 63–68, quant-ph/9802040.
9. E. Farhi *et al.*, quant-ph/9802045.
10. R. Beals *et al.*, in *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science, 1998*, (IEEE Computer Society Press, Los Alamitos, CA, in press) quant-ph/9802049.
11. Y. Ozhigov, quant-ph/9803064.
12. L. Grover, *Phys. Rev. Lett.* **80**, 4329 (1998), quant-ph/9712011; G. Brassard *et al.*, quant-ph/9805082.
13. L. Grover, in *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing* (ACM Press, Dallas, TX, 1998), pp. 53–62, quant-ph/9711043.
14. A. Nayak *et al.*, quant-ph/9804066.
15. M. Garey and D. Johnson, in *Computers and Intractability–A Guide to the Theory of NP-Completeness* (Freeman, New York, 1979), pp. 1–4.
16. N. Cerf *et al.*, quant-ph/9806078.
17. D. Abrams *et al.*, quant-ph/9801041.

PERSPECTIVES: DEVELOPMENTAL NEUROSCIENCE

Choline, a Vital Amine

Jan Krzysztof Blusztajn

In April of this year choline was classified as an essential nutrient for humans by the Food and Nutrition Board of the Institute of Medicine of the National Academy of Sciences. For the first time, recommendations were issued for the adequate intake of this small molecule (*1*). This official recognition of the importance of choline in the human diet likely foreshadows new public health initiatives on choline nutrition.

Enhanced online at www.sciencemag.org

In 1912 Casimir Funk coined the term "vital amine" to describe organic compounds that are required in small amounts in the diet for the maintenance of normal health. The term later evolved into the more familiar word, vitamin. Choline fits the original definition: It is an amine and for normal health must be consumed in the diet, even though humans can biosynthesize small amounts. Choline is present in most foods (*2*) but is found in particularly high amounts in eggs, liver, peanuts, and a variety of meats and vegetables.

In the body, choline subserves several biological functions (*2*) (see figure, right). It is the precursor of phosphatidylcholine and sphingomyelin, two phospholipids that serve as components of biological membranes and as precursors for intracellular messengers such as diacylglycerol or ceramide. Choline is also the precursor of two signaling lipids, platelet-activating factor and sphingosylphosphorylcholine, and of a neurotransmitter, acetylcholine (*3*). Furthermore, choline can be enzymatically oxidized to betaine and the methyl groups of betaine then used to resynthesize methionine from homocysteine, thereby providing methionine for protein synthesis and transmethylation reactions. This last pathway is also an alternative to one that uses the cofactor methyltetrahydrofolate, and thus spares methyltetrahydrofolate for its role in the synthesis of nucleic acids.

The importance of choline for maintaining health in adults has been recognized for some time (*4*), but recent work points to its critical role in brain development. Meck, Williams, and their colleagues have tested the relation between choline availability during fetal development and brain function, initially using behavioral measurements (*5–7*). Various amounts of choline were included in the diets of pregnant rats during embryonic days 11 (E11) to E17 (a period of high cell division and programmed cell death in fetal brain). Pregnant rats ingested either no choline, control amounts of choline, or approximately three times control amounts of choline. For the remainder of their lives both mothers and offspring ate a normal diet. These alterations in choline availability during the second half of gestation resulted in life-long behavioral changes. As adults, offspring of mothers that had received choline supplements were more adept at tasks that measured spatial and temporal memory and attention (*5–7*). If the mother had received no choline, the

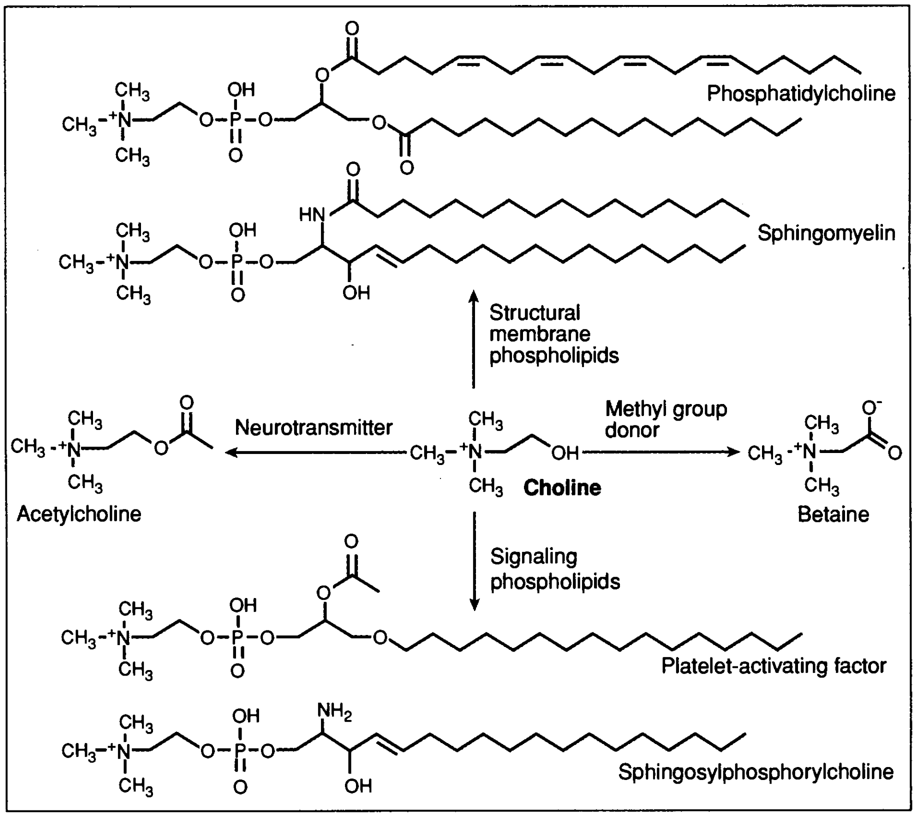

Multiple fates of choline.

The author is in the Department of Pathology and Laboratory Medicine, Boston University School of Medicine, Boston, MA 02118, USA. E-mail: jbluszta@bu.edu