Report of the Workshop on Quantum Information Science

1 July 2009

In January 2009, the United States National Science and Technology Council issued a report entitled *A Federal Vision for Quantum Information Science*. The report proposes that "The United States ... create a scientific foundation for controlling, manipulating, and exploiting the behavior of quantum matter, and for identifying the physical, mathematical, and computational capabilities and limitations of quantum information processing systems in order to build a knowledge base for this 21st century technology."

A Workshop on Quantum Information Science, organized in response to the NSTC report, was held 23-25 April 2009 at the Tysons Corner Marriott in Vienna, Virginia. The workshop brought together leading theorists and experimentalists drawn from physical science, computer science, mathematics, and engineering, who assessed recent progress in QIS and identified major goals and challenges for future research. The workshop also included open evening sessions so that all participants could express their views concerning the priorities for a national QIS initiative. The workshop program and list of participants have been included as appendices to this report.

This report will highlight some of the recent developments that were discussed at the workshop. It is not our intent to provide a comprehensive overview of current QIS research, nor to prescribe in detail how a national QIS initiative should be structured. Instead, by emphasizing a few of the recent achievements of this field, we hope to convey that QIS research is advancing steadily across a wide front, and that a suitably broad national initiative can facilitate further high-impact advances in quantum science and technology in the near future. Much more detail can be found by perusing the presentations available at the workshop website: http://www.eas.caltech.edu/qis2009/

We do not yet have a clear picture of how QIS might influence the science and technology of the 21st century. It is likely that the most far-reaching quantum technologies have not yet been anticipated, and will emerge only as basic research in QIS continues to mature and develop. But as the NSTC report emphasizes, QIS will require long-term focused attention for a decade or more from a variety of government agencies and national laboratories if the US is to achieve and maintain a global leadership position while training a new generation of quantum scientists and engineers. We hope that the Workshop on Quantum Information Science and this report will help to nucleate a cohesive national effort that will nurture and invigorate this vitally important emerging field.

Quantum Information Science

Quantum information science (QIS) is a relatively new and rapidly developing interdisciplinary field of science and technology, drawing from physical science, computer science, mathematics, and engineering, which addresses how the fundamental laws of quantum physics can exploited to achieve dramatic improvements in how information is acquired, transmitted, and processed. Theoretical research indicates that

large-scale quantum computers, if and when they can be developed, will be capable of solving some otherwise intractable problems with far-reaching applications to, for example, cryptology, materials science, and medicine. Experimental efforts to manipulate quantum states, while still lagging far behind theoretical aspirations, are achieving steadily improving results using a variety of physical systems, such as atoms, molecules, photons, spins of electrons and nuclei, superconducting circuits, and mechanical oscillators.

While large-scale general purpose quantum computers are probably still decades away, the technological and scientific impact of QIS research is already becoming apparent. Quantum information processing is being exploited in some of the world's most accurate clocks, and in magnetic sensors that achieve an unprecedented combination of sensitivity and spatial resolution. Tools from atomic physics are being harnessed to simulate models of quantum many-body physics that are beyond the reach of today's digital computers. Perhaps most significantly, QIS is forging fruitful links among disparate fields; for example, insights from fundamental computer science are deepening our understanding of the foundations of quantum physics, and of the strongly correlated quantum systems studied by condensed matter physicists.

Increasingly, continued progress in experimental QIS research will hinge on advances in engineering of materials, devices, and systems. Some of these engineering challenges are closely related to the problems faced by today's information industries, which have pressing needs for reconfigurable, parallel, fault-tolerant architectures with low power consumption. We can expect a gradual convergence of quantum and classical systems engineering, though each will retain its own distinctive flavor.

Despite the clear relevance of QIS to national security and economic competitiveness, unsteady support from US government agencies has discouraged some young scientists from entering and remaining in the field, and has contributed to a "brain drain" in which some of the most able and successful researchers working in the US have been attracted to institutions elsewhere, for example in Canada, Europe, Asia, Australia, and Israel. This problem is further exacerbated by the profoundly interdisciplinary character of QIS research, which has the unfortunate consequence that hiring committees at US universities find it difficult to make thoughtful comparisons between QIS faculty candidates and other candidates working closer to the mainstream of their home disciplines.

A national initiative providing stable merit-based support for broad multidisciplinary curiosity-driven basic QIS research will help to reverse the brain drain, stimulate faculty hiring, and encourage talented young scientists to follow their hearts by committing to careers in QIS. Though one important function of a national program will be to provide badly needed QIS-trained people for mission-focused projects and for industry, an emphasis on broad curiosity-driven basic research is vital for several reasons. First, QIS is a young field where fundamental new discoveries could launch currently unforeseen quantum technologies. Second, beyond its technological potential, QIS is already generating penetrating insights with profound intellectual value concerning the

foundations of computation and physical science. And third, the best way to attract the brightest young people to QIS research is to give them the freedom to pursue the scientific challenges that they find most exciting.

Various government agencies, each with its own mission, goals, and culture, will have a stake in QIS research. By providing more avenues for innovative ideas to find an appropriately supportive home, the existence of diverse autonomous funding agencies can be a virtue, for US science in general and for QIS in particular. But to manage a national QIS initiative effectively, the agencies involved should maintain regular contact, and seek consensus concerning how a well balanced overall investment strategy can best serve the long-term national interest.

We anticipate that many developments in QIS will continue to be driven by individual investigators working at universities. But an important role can also be played by larger centers, which encourage excellence, promote multidisciplinary activity, and compete effectively for the best people against aggressive international competition. Industry and national laboratories, which can support big projects with infrastructure not accessible at most universities, also have a unique role to play. Whatever the research venue, graduate and postdoctoral fellowships funded as part of a national QIS initiative can help to fuel progress.

We also expect QIS to have an extensive impact on science and engineering education, at all levels. QIS researchers require training in a wide variety of disciplines, including mathematics, computer science, information theory, theoretical and experimental physics, chemistry, materials science, and systems engineering. We anticipate that students will increasingly view QIS not as an abstruse specialty to be encountered late in their educations but as a unifying principle connecting a wide variety of physical systems and tools for manipulating them. A course focusing on the core concepts of QIS can be accessible to science and engineering students at an early undergraduate level. Appropriately modified, it could even be taught in high school, exposing a new generation to a thrilling scientific and engineering adventure.

Bright young scientists are attracted to QIS because of its interdisciplinary character, intellectual boldness, and vast technological potential. We note, for example, that the Topical Group on Quantum Information (GQI) of the American Physical Society (APS), founded in 2004, now has 933 members of whom 524 are students, by far the largest student percentage for any of the APS units. This enthusiasm for QIS is a valuable national resource that should be cultivated and exploited. Investment in a national initiative supporting QIS research is bound to pay off handsomely, by providing the training and opportunity for a budding generation of scientists and engineers to conceive and develop the revolutionary new technologies that will be the foundation of our future prosperity.

Research Snapshots

These capsule summaries of 23 different themes of QIS research were distilled from the workshop presentations. We hope that they provide some useful perspective on the goals and status of current research, the pace of progress, and the daunting challenges that remain.

Quantum algorithms and complexity

The discovery of Shor's quantum algorithm for finding the prime factors of a large composite integer provides strong evidence that quantum computers are more powerful than classical computers. Speaking loosely, we say that Shor's algorithm achieves an "exponential speedup" relative to the fastest currently known classical algorithm for factoring, meaning that a quantum computer can factor using a number of steps that grows like a power of the number of digits in the number to be factored, while a classical computer requires a number of steps that grows faster than any power. The goal of quantum complexity theory is to better characterize what quantum computers can do and what they cannot do, and in particular to determine the class of problems for which exponential quantum speedups are achievable.

One particular important challenge is to advance the theory of "post-quantum cryptography." Quantum computers will be able to break public-key protocols that are widely used in our digital society; therefore today's cryptosystems will eventually need to be replaced by new cryptosystems that are efficient, convenient, and plausibly resistant to quantum attacks. Possible candidates for post-quantum cryptosystems include the "lattice-based" cryptosystems; though these are classical protocols, insights derived from quantum computing have improved their efficiency and deepened out understanding of their security. Judging which classical cryptosystems are likely to be quantum-resistant will require a deep understanding of the power and limitations of quantum computers.

Quantum computers can speed up exhaustive search for the solution to a problem, but not exponentially. Exponential speedups are possible only for problems with suitable structure for the quantum computer to exploit. In particular, quantum algorithms have been constructed for many "hidden subgroup problems" which have special symmetries. Shor's algorithm solves a hidden subgroup problem in which the underlying group is abelian, but growing evidence indicates that for some nonabelian groups (like the symmetric group) the hidden subgroup problem is hard even for quantum computers. Further study of these problems may lead to the formulation of useful "quantum one-way functions" which are easily computed by a classical computer but are hard for a quantum adversary to invert.

Quantum algorithms with polynomial speedups

In contrast to exponential quantum speedups, which so far have been established for only a rather narrow class of computational problems, polynomial quantum speedups are

known to be very common. Most familiar is Grover's algorithm which performs an exhaustive search through N candidate solutions in a time proportional to the square root of N. In the past five years, a variety of quantum algorithms have been constructed that achieve quantum speedups using quantum walks, which are better adapted to the structure of the search space than Grover's general purpose algorithm. For example, determining whether a set of N objects contains the same element twice requires N steps classically, but only $N^{2/3}$ steps using a quantum walk algorithm.

Discoveries in 2007 showed that quantum computers can speedup the evaluation of Boolean formulas. For example, the number of steps needed to evaluate a balanced formula with N leaves (an idealized model for the problem of determining whether there is a winning strategy in a two-player game) scales as $N^{.753}$ classically and $N^{1/2}$ quantumly. A more general and powerful characterization of the problems that admit polynomial quantum speedups, relating these to the problems that have classical "span programs", was formulated in 2009. Another advance in 2009 showed that a 9^{th} root quantum speedup is the best possible for a fully symmetric function, a substantial generalization of previously known lower bounds on quantum query complexity. Thus recent progress has brought us much closer to understanding what properties of a problem imply that a quantum computer can solve the problem faster than a classical computer. Perhaps further progress can be attained by exploring the applications of span programs.

Complexity theory and quantum many-body physics

One of the most important applications for quantum computers will be simulating quantum systems with many degrees of freedom. Such simulations can illuminate some of the most important open problems in physics and chemistry, encompassing strongly correlated electron systems, quantum antiferromagnets, exotic superconductors, complex biomolecules, nuclear matter at finite density, and perhaps even quantum gravity.

But to assess the advantage of quantum computers for simulating quantum systems, we need a better understanding of the hardness of these simulation problems for classical computers. QIS research is deepening our grasp of such issues, in particular by illuminating how classical computational hardness depends on properties like the geometry and spatial dimensionality of the quantum system, and on whether the energy spectrum has a gap between the ground state and the lowest excited state. For example, a new result proved in 2009 shows that a classical computer can efficiently simulate the adiabatic evolution of a quantum system in one dimension with a constant spectral gap, while on the other hand it was shown in 2008 that if the gap gets small as the size of the system increases then simulating adiabatic evolution in one dimension can be BQP-hard (that is, for some Hamiltonians it is as hard as any problem that can be solved efficiently by a quantum computer).

Questions regarding the computational complexity of quantum simulation in more than one dimension are still largely open and being pursued vigorously in current research. Meanwhile, insights derived from recent progress in understanding quantum entanglement have led to the formulation of new classes of many-body quantum states

that can be efficiently represented classically, such as matrix-product states and their various generalizations. These new tools reduce many quantum simulation problems of interest in condensed matter physics to classical optimization problems that can be solved efficiently, though it was also shown in 2008 that for some Hamiltonians this optimization problem is NP-hard (as hard as any problem whose solution can be verified efficiently by a classical computer).

Interactive quantum protocols and quantum games

The term "game" can be applied to any structured interaction in which a collection of players have well-defined goals, and we speak of a "quantum game" if the interaction involves quantum information in some way. Studying quantum games can provide a novel perspective on how quantum information differs from classical information.

One important class of games are interactive proof systems, in which one or more "provers" with great computational power try to convince a "verifier" with limited power that the prover(s) can solve a computational problem correctly. In the classical setting, studying interactive proofs has generated deep insights concerning the hardness of finding approximate solutions for problems that are hard to solve exactly (the Probabilistically Checkable Proof theorem); likewise, studying quantum interactive proofs may deepen our understanding of the power of quantum computing in ways that cannot be easily anticipated. In 2009 it was established that a quantum interactive proof with a single prover is no more powerful than a classical interactive proof if the quantum communication is limited to two messages exchanged between verifier and prover, but the power of quantum protocols with unlimited communications remains an open problem. Even less understood are protocols with more than one prover, where quantum entanglement shared by the provers can enhance their ability to fool the verifier.

Nonlocal games, in which two cooperating players are unable to communicate once the game begins, also provide novel insights into the properties of quantum entanglement (they are related to the Bell inequalities long studied by physicists). It was shown for the first time in 2008 that the optimal probability of winning such a game is computable, and also that finding an accurate approximation to the winning probability is NP-hard. Current research aims to illuminate how much entanglement is needed to play a nonlocal game optimally or near optimally, and to clarify how the optimal strategy is affected when many games are played in parallel.

Another interesting quantum game is quantum coin flipping, in which two players at different locations determine the outcome of a fair coin toss by making local measurements and sending qubits back and forth. A new formalism for analyzing two-player games introduced in 2007 made it possible to characterize completely the optimal strategy of a cheating player who tries to bias the coin. Protocols for quantum coin flipping can be designed such that the bias is guaranteed to be arbitrarily small, something not achievable in the classical world. These powerful new analytic tools are still cumbersome; simplifying and extending them will likely lead to further insights into quantum game theory and quantum cryptography.

How can a classical user of a quantum computer be sure that the computer's output is correct? In some cases, for example if the quantum computer solves a factoring problem, the user can check the output by doing an efficient classical computation, but the theory of quantum interactive proofs indicates that there are more clever general ways to test a quantum computer even if its output cannot be verified classically. Protocols proposed in 2008 allow a user who can generate random single-qubit quantum states, but who has no other quantum processing power, to check that a general quantum computation is correct. Using such a protocol, a nearly classical verifier can confirm that a quantum device too large to simulate classically really operates according to the laws of quantum physics. Whether the verifier can be fully classical is an intriguing open question.

Fundamental quantum physics

Quantum information science has its roots in early efforts to clarify the nature of quantum nonlocality, especially by formulating precise Bell inequalities implied by local realism, and doing experiments that aim to demonstate convincingly the violation of these inequalities. Even today, experimental Bell inequality tests have loopholes; experiments with entangled photons are limited by the imperfect efficiency of photodetectors, and experiments with trapped ions are limited because measurements are slow compared to the light-travel time between ions. An important goal, likely to be met in the near future, is to confirm Bell-inequality violation in a loophole-free experiment.

Meanwhile, experiments are observing coherent quantum behavior in physical systems of gradually increasing size, though there is still a large gap between the "Schrödinger cat" states studied in the laboratory and the scale of everyday life. In fact, quantifying the "cattiness" of a quantum superposition is an important theoretical problem, necessary for making fair comparisons among experiments done with distinct physical systems. Current data neither confirm nor exclude the "macrorealism" hypothesis --- that sufficiently large systems like real cats decohere not just in practice but as a matter of principle. Experimenters should continue to push studies of quantum coherence to larger scales, while theorists should search for plausible and testable models that accomodate intrinsic decoherence.

One of the most fascinating observations ever made about the difference between quantum and classical physics is that classical systems seem unable to simulate quantum systems efficiently. It is therefore reasonable to expect ideas from the foundations of computer science to aid the quest for a deeper understanding of the laws of Nature. We know for example that if the evolution equation for quantum states were a generic nonlinear equation rather than a linear equation, then NP-hard problems could be solved efficiently; perhaps then the Schrödinger equation is linear to eschew such unreasonable computational power. Another recent discovery is that, in the context of "probably-approximately-correct" (PAC) computational learning theory, an experimenter can learn the obervable features of an n-qubit quantum state making a number of measurements linear in n, suggesting that the resources Nature employs may not be quite so extravagant as the exponential size of Hilbert space indicates. And "information causality," a new

physical principle proposed in 2009, may point toward a long-sought deeper explanation for the nonlocal correlations realized by quantum systems.

Skeptics of the power of quantum computing are challenged to formulate credible models that are compatible with all of our currently confirmed knowlege about the quantum world, yet predict quantum computers will fail to outperform classical machines. Quantum information science may also help guide us toward a deeper grasp of the foundations of quantum gravity. Can a universal quantum computer simulate quantum-gravitational processes efficiently, and if not what computation model captures correctly the computational power that Nature allows?

Quantum computation for chemistry

Quantum computers will be powerful tools in computational quantum chemistry. Exact full configuration interaction (FCI) computations of the energy of an n-electron molecule require time exponential in n on a classical computer, but can be done in time $O(n^5)$ on a quantum computer. Furthermore, simulations of chemical dynamics require exponential time classically but only time $O(n^2)$ quantumly. Even simulations of lattice protein folding models admit quadratic quantum speedups. Thus while classical computers are limited to *ab initio* simulations of only small molecules, quantum computers will vastly surpass these limits, with potentially profound industrial and medical applications.

Recent estimates indicate that a quantum computer with 116 logical qubits could compute the ground state energy of the water molecule, for example, more accurately than current classical supercomputers. Around 3000 logical qubits would suffice for an accurate computation of the energy of the cholesterol molecule, which is far beyond the capabilities of foreseeable classical methods. A proof-of-principle experiment in 2008 demonstrated that the ground state energy of the hydrogen molecule can be computed to better than six-digit precision using methods from quantum optics.

Much can be done to improve estimates of the resource requirements for quantum chemistry using a quantum computer, particularly when the quantum computer is subject to realistic noise. Meanwhile, the interface of quantum information with chemistry is stirring other novel insights, for example new explanations for long-lived coherent energy transfer in chlorophyll molecules, with could lead to a deeper understanding of the mechanism of photosynthesis and suggest new approaches to solar energy harvesting.

Quantum information research is also contributing powerful methods for simulating chemical systems using classical computers, based for example on matrix-product and tensor-network representations of many-particle quantum states. These methods have been used in correlated calculations of excited states for polyacenes and carotene molecules, which were beyond the reach of previous methods.

Capacities of quantum channels

In Shannon's classical information theory, a communication channel can be characterized by its capacity, a single number expressing the maximal rate at which information can be transmitted over the channel with negligible probability of error. In contrast, a quantum channel has several capacities: a classical capacity C for transmitting classical bits reliably, a quantum capacity Q for transmitting qubits reliably, and a private capacity P for transmitting classical bits privately when an eavesdropper controls the channel's environment. Furthermore, also in contrast to classical Shannon theory, auxiliary resources such as quantum entanglement shared between sender and receiver, or back communication from receiver to sender, can greatly enhance a quantum channel's capacities. In fact, the capacity for sending quantum information assisted by shared entanglement is the most easily characterized of the quantum capacities, and from that point of view seems to be the most natural quantum generalization of the classical capacity of a classical channel.

Surprising discoveries in 2008 have transformed quantum Shannon theory. We now know that in some cases two quantum channels, each with zero quantum capacity, can have a positive capacity when used together. We have also learned that for some quantum channels the classical capacity can be achieved only by sending highly entangled encoded messages.

The former development in particular reveals that quantum information is not a single commodity as formerly thought; rather the ability to deliver intact qubits can be built up from separate abilities to deliver two weaker commodities. Efforts are underway to understand what these weaker commodities could be. Another important challenge is to understand better the private capacity and the quantum capacity assisted by two-way classical communication, both of which are still very poorly characterized.

Quantum key distribution

In quantum key distribution (QKD), two parties connected by a quantum channel and an authenticated classical channel generate a string of private shared random bits that can be used as a one-time pad to encrypt and decrypt a classical message. Privacy in QKD is founded on a fundamental principle that distinguishes quantum information from classical --- gathering information about a quantum state unavoidable disturbs the state so that eavesdropping can be detected. Furthermore, unlike large-scale quantum computing, secure QKD is feasible with current technology, at least over moderate distances such as tens of kilometers of optical fiber or ground-to-satellite optical links.

While the commercial potential of QKD systems remains unclear, QKD research continues to be a rich source of theoretical insights and technological advances. A universally composable criterion for security, which ensures that the key can be safely used in other applications, was first formulated in 2005, and new analytical tools developed since then have simplified the analysis of security while at the same time

providing the first firm mathematical foundation for describing quantum sources using density operators (the Quantum De Finetti Theorem). The set of bipartite quantum states from which private shared key bits can be extracted was characterized in 2005, leading to the discovery that quantum channels with vanishing quantum capacity can be used to distribute key securely. The decoy state method, invented in 2005, established that coherent state optical sources can generate private key that is far more robust against photon loss than previously known. It was shown in 2008 that treating the infinite-dimensional Hilbert space of an optical mode that arrives at a photon detector as an idealized two-dimensional qubit can be formally justified in an analysis of security.

On the implementation side, signal sources and the ability to stabilize quantum channels have steadily improved, so that key generation rates are currently limited by the noise and inefficiency of single photon detectors. Several new ideas are spurring the development of better detectors, such as up-conversion that converts telecom wavelengths to more easily detected wavelengths, high-efficiency cryogenic superconducting detectors, and self-referencing detectors with reduced dead times.

As is the case in classical cryptography, QKD systems are potentially vulnerable to "side-channel attacks" that exploit how actual devices deviate from the idealized models used in security proofs. New security proofs appearing in 2008, based on device-independent assumptions such as the impossibility of superluminal signaling, provide a fundamental new approach to overcoming side channel attacks; however so far these device-independent proofs do not apply to realistic settings in which photon loss rates are relatively high.

Fault-tolerant quantum computing

Quantum computers are intrinsically far more vulnerable to error than classical computers. Thus our hopes that large-scale quantum computers will be built and operated someday are founded on the theory of quantum fault tolerance, which establishes that reliable quantum computation is possible when the noise afflicting the computer has suitable properties. Recent insights are broadening the class of noise models for which fault-tolerant quantum computing is provably effective, and clarifying the overhead cost of overcoming noise. Aside from its practical importance for guiding the development of quantum technologies, these theoretical developments, which address whether subtle quantum interference phenomena can be exhibited in realizable systems with many degrees of freedom, are also of fundamental interest.

Specifically, if the noise is suitably local in space and time, then quantum computing is scalable if the error rate per gate is below a critical value called the accuracy threshold. In 2008, rigorous proofs established a lower bound on the accuracy threshold of 10^{-3}, and recent numerical studies indicate that the actual value of the threshold can approach 10^{-2}, even when all quantum gates are required to be local in a two-dimensional array. Furthermore the accuracy threshold theorem was extended in 2008 to a broad class of realistic noise models that include strongly non-Markovian effects and highly asymmetric noise. This progress exploits several recently developed theoretical ideas, such as state

distillation for performing high-fidelity gates by teleportation, quantum error-correcting codes based on topological principles, and message passing algorithms for more reliable decoding of hierarchical codes.

Current work is extending the theory of quantum fault tolerance in a variety of fruitful directions, including new methods for optimizing overhead cost, incorporating fault-tolerance into the design of quantum hardware and systems, and combining fault-tolerant circuit design with methods from quantum control theory. Other intriguing problems concern the formulation of new fault-tolerant approaches applicable to non-standard models of computation, such as quantum computing by adiabatic evolution.

Topological quantum computing

Topological quantum computing is a novel approach to quantum fault tolerance that exploits the unusual physical properties of exotic states of matter with "nonabelian topological order." These systems support "anyons," particles with locally conserved charges such that quantum information can be stored in the collective quantum states of many anyons. This information can be processed by exchanging the positions of the anyons, even though the anyons never come close to one another, and it can be read out by bringing a pair of anyons together to measure the total charge of the pair. In principle, a topological quantum computation is intrinsically resistant to decoherence --- if the paths followed by the anyons execute a prescribed braid in spacetime, then the computation is guaranteed to find the right answer.

Furthermore, strong experimental evidence indicates that nonabelian states suited for topological quantum computing are among the fractional quantum Hall (FQH) liquids that can be realized by confining electrons to a two-dimensional interface between semiconductors in a high transverse magnetic field at very low temperature. The suggestion by theorists in 2005 that anyon charges can be measured in a suitably constructed double-point-contact interferometer revitalized the exploration of fraction quantum Hall physics, spurring experimenters to seek convincing experimental evidence for the existence of nonabelian anyons. Experimental confirmation of nonabelian anyons, aside from laying the foundations for a potentially viable quantum computing technology, would also be a milestone for condensed matter physics.

In 2008 theorists proposed measurement-only topological quantum computation, in which interferometric charge measurements alone, without any additional anyon braiding steps, are adequate for quantum information processing. The theory of how quantum point contacts work in FQH liquids, which is critically important for interpreting the experimental results, also advanced in 2008 based on the insight that tunneling across the liquid can be formulated in terms of renormalization-group flow from one boundary conformal field theory to another. On the experimental front, measurements in 2008 of how quasiparticle tunneling depends on temperature strengthened the case for the interpretation of the filling-factor 5/2 FQH liquid as a nonabelian state, while the observation that interference fringes are reproducible after a one week delay suggests that error rates due to thermal activation of trapped anyons are extremely low.

These encouraging experimental developments are still tentative and will need to be confirmed by other labs and using different experimental set ups. Faster methods for data collection will help to clarify the story and will be needed in any case as a step toward larger-scale measurement-only quantum computing. Theorists should refine proposals for achieving computational universality using the 5/2 FQH state based on gate teleportation and state distillation. Meanwhile, other physical systems that might realize nonabelian topological phases, using for example trapped ultracold atoms or molecules, should be pursued by theorists and experimenters.

Quantum control

Methods from quantum control theory will be an essential tool in the quest for quantum hardware that surpasses the threshold of accuracy. Furthermore, apart from the potential applications to quantum computing, quantum effects need to be properly taken into account to control a broad range of mesoscopic devices. While quantum control theory draws heavily on the corresponding classical theory, there are also important conceptual differences; indeed, viewing quantum mechanics as a basis for the design of devices and systems provides a fresh perspective on foundational issues like the interpretation of quantum measurement.

Control theory divides loosely into open-loop control in which the system is steered but not monitored, and closed-loop control which includes continous real-time feedback. Open-loop quantum control, with important antecedents in the study of chemical dynamics and nuclear magnetic resonance, provides general procedures for dynamically decoupling a system from its environment. While earlier open-loop theory assumes hard control pulses, more recent work has formulated effective analytic methods based on pulses of bounded strength and non-negligible width. It has also been shown recently that suppression of decoherence can be improved by varying the time interval between pulses, and analytic methods have been developed for optimizing the pulse sequence. These methods do not require detailed knowledge of the properties of the bath that couples to the system, but when that knowledge is available much shorter pulse sequences can be found using numerical methods, such as the Gradient Ascent Pulse Engineering (GRAPE) algorithm developed in 2005.

Closed-loop quantum control differs from classical theory because quantum measurement back-action must be included. Adaptive measurement schemes have been developed that can attain the Heisenberg uncertainty limit in optical phase measurement, for example. A current frontier is coherent-feedback control, in which quantum information collected during continuous monitoring is subjected to unitary quantum processing rather than amplified via quantum measurement. Coherent-feedback control was demonstrated experimentally in 2008 using a quantum-optical system, and experiments are already pushing beyond the reach of current theory.

Quantum architecture

Current experimental work on quantum hardware focuses on the quest for qubits with long-lived coherence and devices that execute high-fidelity gates. Eventually, though, we will face the challenge of assembling quantum devices into reliable systems. The path from devices to systems has been well charted by computer scientists, but quantum computation poses some novel challenges. Perhaps most fundamentally, quantum systems will rely on quantum error correction and fault tolerance, at a fine-grained level of architecture, and new research will be needed to understand the optimal balance of space, time, and energy for implementing a quantum computation given underlying device capabilities.

Recent work on quantum architectures includes detailed case studies of large-scale quantum computers based on realistic projections of trapped ion, quantum dot, and photonic quantum gate technologies. These studies predict the resources needed for lasers, classical control, and cooling, and have introduced many new concepts, such as using entangled states as power supplies for quantum chips, and providing inter-chip and intra-chip communication via teleportation. Among the key architectural building blocks are "software factories," in which special resource states needed for computational universality are prepared and verified.

Predictive tools for designing and verifying large-scale quantum information processing systems are badly needed to steer the development of key technologies. Research on quantum architecture design may also benefit the design of classical systems, which will increasingly rely on fault-tolerant operations and methods for minimizing power consumption.

Trapped-ion quantum information science

Experiments with trapped ions continue to yield impressive achievements in the coherent manipulation of quantum systems. Among the tasks demonstrated are qubit teleportation, error correction, the Deutsch-Josza algorithm, the Grover search algorithm, a Toffoli gate, the quantum Fourier transform, entangled state purification, dense coding, entanglement-assisted detection, preparation of the eight-qubit W state and the six-qubit GHZ state, and generation of arbitrary motional state superpositions. So far, however, these gates and algorithms have not attained the stringent accuracy requirements for fault-tolerant quantum computing, and current research aims to improve operation fidelities and set the stage for scaling up to large numbers of qubits. These improvements are desired not just to enhance the power of trapped-ion quantum computing, but also for other applications to quantum information science, for example in metrology.

The pursuit of quantum computation using ion traps is opening new avenues of research. Entanglement of distantly separated ions has been achieved using two-photon interference; the demonstrated data rate is low but could be substantially enhanced with optical cavities. Such optical links between ions could be used to couple qubits in a large

quantum computer, or for long-distance quantum communication in a quantum network. Lithographically prepared surface-electrode traps are now available, but may need to run at low temperature to suppress anomalous heating of ion motion. In these small traps, strong radio-frequency magnetic field gradients rather than lasers can couple an ion's motion to its internal state, dramatically reducing the laser power requirements for quantum information processing.

As in other areas of high-precision physics, ion trap experiments push the limits of currently available classical control systems. Specifically, control of laser beam intensity, waveform control, and predictable switching are all big problems, and increasingly complex classical computer control will be required as quantum algorithms grow more complicated. These engineering issues must be addressed for trapped-ion quantum computing to continue to progress.

Quantum information science with cold neutral atoms

Like trapped ions, cold neutral atoms are ideal quantum systems with excellent coherence and accurately known Hamiltonians. Furthermore, by exploiting Bose condensates and optical lattices, many qubits can be initialized and massively parallel operations can be performed. For example, entangling gates applied in parallel to neighboring atoms in a lattice can prepare a many-qubit cluster state, a potentially powerful resource for measurement-based quantum computing.

Experiments using optical lattices of double-well potentials have demonstrated entangling operations between pairs of atoms, such as the "square-root-of-SWAP" operation, with reasonable fidelity. In the initial experiments of this type, many qubits or pairs of qubits were detected simultaneously, but more recently developed tools should allow atoms in an optical lattice to be addressed individually with focused laser beams. Focused laser beams used as optical tweezers can also hold and manipulate individual atoms and entangle them with other similarly confined atoms. While neutral atoms in their ground states have only short-range interactions, so that entangling operations can be applied only between atoms in close contact, atoms in highly excited Rydberg states have strong long-range interactions, so that atoms microns apart can be entangled in much less than a microsecond.

It remains challenging to load atoms into a lattice without defects, and to keep the atoms cold as a many-qubit state is processed. Recent atom sorting experiments using conveyor belts have demonstrated new tools for repairing defects in an atomic storage register, and other promising repair protocols have been studied theoretically. One potentially promising alternative to optical lattices is provided by magnetic microtraps integrated on atom chips.

Spin qubits

Mesoscopic semiconductor devices occupy the middle ground between the quantum world of individual atoms and the classical world of everyday objects. Unlike atoms, no

two such devices are identical, but semiconductor devices have other potential advantages, such as high clock rate, scalability, controllability, and compatibility with current electronics. There are a variety of options for encoding qubits in semiconductors, for example using electron spins controlled either optically or electrically, using nuclear spins, using excitons, or using the electron's charge.

Most experiments with spin qubits have been done with gallium arsenide (GaAs) devices, in which electrons can be accurately controlled and individual electron spins can be well isolated in quantum dots. In these systems it is important to control decoherence driven by the hyperfine coupling of the electron spin to nuclear spins in the material. One useful trick is to encode the qubits using the spin-singlet and spin-triplet states of a double quantum dot, for which quantum gates can be performed in a few nanoseconds, single-shot measurements are fast, and coherence times greater than a microsecond have been achieved. Current experiments with two double quantum dots aim to demonstrate high-fidelity two-qubit quantum gates. Eventually, decoherence due to the nuclear spin bath might be suppressed further using quantum dots in materials with zero nuclear spin, such as isotopically pure silicon-28. So far, single electron spins have not yet been isolated in silicon and silicon-germanium quantum dots, but the gap between GaAs and SiGe technology may narrow.

In principle, many spin qubits could be integrated on a single chip; efforts are underway to map out the control electronics that would be needed to operate a large-scale system. A quantum computer using semiconductor quantum dots will necessarily operate at sub-Kelvin temperatures, but recent advances in cryogen-free refrigeration (without liquid helium) should reduce the difficulty of low-temperature experiments with semiconductors, superconductors and perhaps even ion traps.

Nitrogen-vacancy (NV) centers in diamond provide another promising type of spin qubit, which has a long coherence time even at room temperature. An NV center is a lattice impurity accompanied by an electron spin that can be manipulated using optical techniques borrowed from molecular spectroscopy, and also coupled to nearby nuclear spins. Already a useful quantum algorithm has been implemented using NV centers --- a ten-fold improvement in single-electron-spin detection was achieved by repeatedly reading out nuclear spins. In contrast to a quantum computer using semiconductor quantum dots, with foreseeable advances in materials engineering a solid-state quantum computer based on NV centers might work someday at room temperature.

Qubits in superconductors

Like spin qubits, superconducting qubits have notable strengths and weaknesses. On the one hand, they can be mass produced, are electronically controllable, have strong tunable interactions, and couple readily to traveling photonic qubits. On the other hand, they require careful calibration, and are subject to 1/f noise of mysterious origin. Quantum information can be encoded in three ways: the two states of a phase qubit are the ground or first excited state in an anharmonic potential, the two states of a flux qubit have clockwise or counterclockwise current circulating in a superconducting loop, and the two

states of a charge qubit have n or n+1 Cooper pairs on a small superconducting island. Qubits can be coupled with a capacitor, an inductor, or a microwave transmission line.

Coherence times exceeding a microsecond have now been demonstrated, thanks to a combination of better qubit designs that suppress sensitivity to charge fluctuations and improved materials and fabrication methods that reduce dielectric losses. In 2009, single-qubit gates with a 1.2% error rate were demonstrated, and a simple quantum circuit of one-qubit and two-qubit gates, implementing a rudimentary form of Grover's search algorithm, was executed with reasonable fidelity. Furthermore, by coupling a phase qubit to a microwave resonator, a variety of multiphoton states with up to six photons were prepared and their Wigner functions measured, a procedure requiring dozens of highly accurate pulses with sub-nanosecond timing. This research illustrates the promise of a new arena for strongly-coupled cavity quantum electrodynamics and continuous-variable quantum information processing.

Further progress with superconducting qubits would be facilitated by a variety of engineering improvements, such as reliable dilution refrigerators, cheap waveform generators, electronically controlled couplings with high on/off ratio, quantum-limited detectors in the gigahertz range for qubit readout, and low-electrical-loss device fabrication to suppress decoherence.

Nuclear magnetic resonance quantum information processing

In nuclear magnetic resonance (NMR) quantum computing, qubits are stored in nuclear spins, which have long coherence times, and gates are executed using radio frequency pulses that modify the spin evolution. In contrast to other types of quantum hardware, qubits in room-temperature liquid-state NMR are very weakly polarized, and quantum information is actually stored in an ensemble containing many molecules. For these and other reasons, the NMR quantum computer is not expected to be scalable beyond at best a few tens of qubits. Nevertheless, NMR remains a powerful tool for investigating a variety of ideas concerning quantum information processing, such as noise characterization, quantum control, and quantum error correction. Thinking about NMR computing has also inspired new theoretical ideas that may be broadly applicable, such as quantum algorithms for cooling and new computational models which, while weaker than full-blown quantum computing, might still be able to solve some problems that are beyond the reach of classical computers.

Experiments in 2009 established that using liquid state NMR single-qubit quantum gates with error rate per gate approaching 10^{-4} and multi-qubit gates with error rate below 1% can be achieved, an unprecedented level of control in quantum systems. Furthermore, highly entangled states of up to 12 qubits have been prepared, using pulse sequence design methods that are likely to be applicable to other settings. NMR experiments have also provided a proof of principle for algorithmic cooling methods, in which a warm thermal state of several qubits is mapped to a colder state of fewer qubits.

New tools are emerging from hybrid approaches that combine the virtues of NMR and electron spin resonance. Electron spins are easier to polarize and detect than nuclear spins, but they also decohere faster. By coherently transferring quantum information between the two, new types of quantum sensors and actuators can be realized, which might prove useful for probing biological systems at the single molecule level.

Optical quantum information processing

The experimental foundations of quantum information science have their roots in quantum optics. Entangled photons were the first systems used for demonstrations of quantum nonlocality, quantum erasure, quantum teleportation, decoherence-free subspaces, and entanglement distillation. And of course photons are used in all realizations of quantum key distribution; some of the challenges facing photonic quantum communication are discussed in other sections of this report concerning quantum key distribution and quantum networks.

.

Though some small-scale quantum algorithms have been demonstrated using photons, these are all based on post-selection methods which are not truly scalable to large systems. Remarkably, however, scalable quantum computing is possible in principle if linear optics is augmented by single-photon sources and projective measurements. While early estimates of the resource requirements for this scheme were discouraging, these can be substantially reduced using a cluster-state architecture, which is also quite robust against photon loss.

Even with these improved protocols, the technical requirements for optical quantum computing are demanding. Efforts are underway to develop the on-demand single-photon sources, number-resolving photon detectors, and reliable quantum memories that would enhance scalability. Meanwhile, theorists are challenged to conceive more efficient architectures and error correction schemes that are better tailored to the noise in photonic devices. One potentially more efficient scheme uses linear optical elements to process qubits encoded in squeezed states of light; the catch is that the resource states needed for this scheme are complicated quantum superpositions that are hard to prepare.

Quantum networks

The distance range of secure high-rate quantum key distribution is currently limited by photon losses in optical fiber to a few tens of kilometers. For a truly global quantum Internet, quantum repeaters will be needed that use quantum error correction or entanglement purification at intermediate nodes of the network to extend reliable quantum communication to longer distances. Indeed, quantum repeaters will be an important application for modest scale quantum information processing, potentially attainable well before the advent of large-scale quantum computers that solve hard computational problems.

The key challenge in realizing a quantum network is developing quantum interconnects that coherently and reversibly map quantum information from a memory comprised of

matter qubits at one node to a photon that carries the information to an adjacent node. It was shown in 2008 that an entangled state of light can be deterministically mapped to a pair of atomic ensembles several meters apart, and then mapped back to photonic modes with reasonable efficiency after a programmable delay. Ions a meter apart have also been entangled recently using two-photon interference in a probabilistic protocol. Other emerging hybrid technologies may prove suitable for realizing quantum repeaters, based for example on solid-state quantum memories using NV centers in diamond or semiconductor quantum dots.

A practical quantum network would combine efficient light-matter connections, long-term few-qubit memories, small-scale quantum logic, and telecom fibers all in one integrated system. With known protocols, resource requirements scale polynomially with the communication distance, but better theoretical ideas might improve speed and efficiency substantially.

Aside from their potential relevance to practical quantum communication, quantum networks can be viewed as novel quantum many-body systems, with interactions among nodes mediated by quantum channels. These systems can exhibit subtle collective effects, including quantum phase transitions, that invite exploration by both theorists and experimentalists.

Quantum metrology

Quantum entanglement, aside from enabling computational speedups, can also be exploited to enhance precision in spectroscopy and atomic clocks. For example a “cat” state of n qubits evolves n times faster than a single qubit, which can improve resolution in measurements of frequencies or time intervals. Quantum information processing can also improve detection and sensing, for instance by mapping quantum information from the system we wish to measure to another system that can be read out more easily. These developments have great potential implications for fundamental physics, since improvements in measurement precision often lead to new discoveries, and also beyond, since better sensing and imaging could have widespread applications to nanotechnology and biomolecular systems.

An optical-frequency “quantum logic clock,” which exploits the frequency stability of an aluminum ion by transferring its quantum state to a beryllium ion that can be more easily detected with lasers, was demonstrated in 2008 and is now among the world’s most accurate. Also in the past year, “spin-squeezed” states of an ensemble of atoms have been prepared and used to achieve an entanglement-induced suppression of quantum noise. Furthermore, an adaptive version of the quantum phase estimation algorithm has been applied to surpass the shot-noise limit on the measurement precision of optical phase.

Some of these simple tricks based on quantum information insights are on the verge of becoming routine tools in optical measurement science. Quantum information ideas can also be fruitfully applied to atom interferometers, which already far outperform optical methods in measurements of acceleration and gravitational field gradients. One

possibility is that interactions among atoms in a Bose-Einstein condensate can be exploited to exceed the naïve limits on phase measurement that arise from the Heisenberg uncertainty principle. At any rate, interpreting a measurement procedure as a quantum circuit clarifies the conceptual basis of limitations on precision while also suggesting novel strategies for improving precision.

Quantum simulation

Understanding and predicting the exotic collective behavior of quantum many-body systems is of great intrinsic scientific interest, and may also guide the discovery of new materials with important technological implications. Large-scale quantum computers will be powerful tools for simulating highly-correlated quantum many-body systems, which are hard to simulate classically in many cases of interest. Remarkably, though general purpose quantum computers may still be decades away, quantum many-body systems that are beyond the reach of currently available computers can already be simulated today using experimental tools developed by atomic physicists.

For example, an optical lattice formed by a web of interfering laser beams can be uniformly filled with ultracold atoms, and the interactions among the atoms can be tuned by adjusting the strength, frequency, and polarization of the laser light. Using these tools, a reversible quantum phase transition from a conducting state (with long-range phase coherence) to an insulating state (with atoms localized at lattice sites) was first achieved in 2002. An experiment in 2008 precisely locating the phase transition was consistent with a recent computer simulation, using the quantum Monte Carlo method, of the two-dimensional Bose-Hubbard model. Other recent quantum simulation demonstrations include: the mapping of phase diagrams for quantum phase transitions in antiferromagnetic spinor condensates, determination of the Berezinskii-Kosterlitz-Thouless phase boundary between superfluid and non-superfluid states in a two-dimensional Bose gas at nonzero temperature, and the observation of ferromagnetic ordering in a Fermi gas. Experiments in the near future will investigate the phases of the Fermi-Hubbard model, which may provide insights into the still mysterious mechanism of high-temperature superconductivity, and probe the behavior of charged bosons and fermions in simulated magnetic fields.

Though in some ways quantum simulation with ultracold atoms is in the midst of merging with mainstream condensed matter physics, the subject also maintains strong connections with quantum information science. For one thing, quantum information research provides helpful guidance concerning which simulation tasks are likely to be hard for classical digital computers; also, coherent processing tricks borrowed from quantum computing can expand a quantum simulator's capabilities. For example, some exotic phases of matter, like spin liquids and other topological phases, are most easily realized as ground states of Hamiltonians that include complicated terms acting collectively on three or more particles, and these many-body terms can be induced using quantum logic. From another perspective, a quantum simulator might be handy for preparing highly entangled "resource states" that are suitable for measurement-based quantum computation or other coherent processing tasks.

Other physical systems besides ultracold atoms might also be useful tools for quantum simulation. For example, trapped polar molecules have strong long-range interactions that can be flexibly tailored to realize a variety of many-body Hamiltonians. Josephson junction arrays have already been used to explore many-body phenomena, though not yet in a fully coherent regime. And ion traps may be well suited for simulations of spin systems on graphs with Ising-like interactions, even when the simulated graph has a geometry much different than the actual physical layout of the ions.

Quantum aspects of mechanical motion

A central goal of quantum information science is to exhibit and control quantum effects in physical systems that are much larger than atomic size, both to address fundamental scientific questions and to realize new technological capabilities. A particularly ambitious goal is to generate, detect, and exploit quantum states of motion in mesoscopic mechanical systems. Mechanical systems, coupled to microwave resonators or to optical fields, can be used as ultra-sensitive force detectors, as buses in quantum information processors, or as probes of quantum behavior in ordinary matter. Mechanical oscillators have already been used to detect the magnetic force (of order 10^{-18} newtons) due to a single electron spin; mechanical detection of a single nuclear spin, for which the magnetic force is weaker by yet another three orders of magnitude, is a tempting but more distant goal.

Cooling a mesoscopic mechanical oscillator to its quantum (zero-phonon) ground state is a great challenge, but there has been rapid progress within the past year. For example, a mean phonon number of 12 has been achieved by back-action cooling of a mechanical oscillator parametrically coupled to a microwave resonator, and a mean phonon number of about 60 has been reached using resolved sideband laser cooling in a cavity optomechanical system. Furthermore, for the first time nanomechanical motion has been measured with precision surpassing the standard quantum limit, using a microwave analog of a Mach-Zehnder interferometer.

In the relatively near future, it should be possible to cool a variety of mesoscopic mechanical oscillators to their motional ground states, to entangle these systems with qubits and with optical cavities, and to measure decoherence rates for squeezed motional states and other superpositions of motional Fock states. These investigations will test quantum mechanics and decoherence theory in a previously inaccessible regime, and provide powerful new tools for ultra-sensitive measurement of very weak forces.

Hybrid quantum systems

Qubits based on different physical systems are each best suited for different tasks. Photons are useful for transmitting quantum information, but difficult to store and process. Nuclear spins are excellent quantum memories but hard to access. Qubits encoded using electric charges can be fabricated en masse and are easy to manipulate and measure, but they decohere quickly. To realize practical quantum technologies, we

should learn to combine the complementary advantages of different qubits in integrated systems, for example by using solid state devices for quantum processing, and atomic-molecular-optical (AMO) devices for quantum memory and communication. To take full advantage of the complementary strengths of solid-state and AMO tools, however, we will need reliable coherent quantum interfaces between these different systems.

Realizing good interfaces is a challenge, but recently there have been a variety of promising proposals. Part of the difficulty is that, in order to couple strongly to a (cold) solid-state quantum device, an atom may need to be trapped close to a surface inside a cryostat. One idea is to achieve cavity-mediated strong coupling of a single atom trapped in an optical lattice with an oscillating mechanical membrane. Coupling distantly separated trapped ions or neutral atoms using photons is another potentially important hybrid technology. Various possibilities are being pursued for coupling superconducting circuits to spins, ions, molecules, and mechanical resonators. And spin qubits in diamond, manipulated by optical techniques, have been used to demonstrate a nanoscale magnetic sensor with an unprecedented combination of sensitivity and spatial resolution.

The quest for optimal qubits draws heavily from theory, experiment, engineering, and materials science, and the emerging hybrid approaches highlight the importance of developing a wide variety of quantum information platforms, some of which may find niches that are hard to anticipate. Hybrid quantum systems might provide a path toward powerful general purpose quantum computers, but even if not the new technological capabilities that emerge are bound to have widespread implications for science and engineering.

Some open questions in quantum information science

Here we list some open questions, mostly drawn from the workshop presentations, that are being addressed by current research. The questions listed are merely representative examples; they are not necessarily more interesting or more important than questions that are omitted. For further context, see the Research Snapshots and the online workshop presentations.

The questions range from more theoretical questions toward the beginning of the list to questions relating more to experiment and technology toward the end. We have divided the list into a few broad categories, but the boundaries between categories are fuzzy, and some of the questions might easily have been classified differently.

Algorithms

Can we find new quantum algorithms for solving problems that are believed to be hard for classical computers but not NP-hard, such as graph isomorphism, lattice problems, the unknotting problem, computing Nash equilibria, and computing partition functions of statistical-mechanical models?

Can we find new quantum algorithms that achieve exponential speedups relative to classical computers based on representation theory, on the Fourier transform, and on other transforms?

Can we find new quantum algorithmic speedups using physics-inspired methods such as quantum walks on graphs and adiabatic quantum computation?

Can we find entirely new techniques for constructing quantum algorithms that achieve significant speedups?

Complexity

How does BQP, the class of problems solvable in polynomial time on a quantum computer, relate to classical complexity classes, and in particular what is the smallest classical complexity class that contains BQP?

How powerful is the class QMA, the quantum analog of NP, for which a complete problem is computing approximately the ground-state energy of a local Hamiltonian? How is the hardness of the local Hamiltonian problem affected when, for example, the Hamiltonian is "frustration free" or a sum of commuting terms?

The classical Probabilistically-Checkable-Proof (PCP) theorem indicates that it is hard for classical computers to find approximate solutions to classical constraint satisfaction problems. Is there a quantum version of the PCP theorem, and if so what are its implications?

What properties of a computational problem imply that the speedup achieved by a quantum computer over a classical computer is at best polynomial?

How powerful are multi-prover quantum interactive proof systems?

Are there quantum correlations other than entanglement that can be exploited for quantum information processing, for example in mixed-state quantum computers?

Cryptography / Communication

Are there efficient classical public key cryptosystems that are plausibly resistant to attacks by quantum computers, and more concretely, can quantum computers break lattice-based cryptosystems?

What are the implications of quantum cryptography for cryptographic tasks beyond key distribution, such as anonymous voting and counterfeit-proof money?

Can we formulate a rigorous, comprehensive theory of quantum games with multiple-round interactions, and if so, what are the implications for cryptography, computational complexity, communication complexity, and distributed computation?

Are there general protocols that allow a classical verifier to check that a quantum computer is operating correctly?

Can we prove the security of practical quantum key distribution against side-channel attacks based on device-independent assumptions?

Can symmetric ciphers be used to expand quantum-distributed keys while remaining resistant to quantum and classical cryptanalysis?

Can we extend the range of practical, secure quantum communication using quantum repeaters?

How can quantum protocols using currently available technology improve the security of practical cryptographic tasks under the assumption, presumably valid in the near term, that large-scale quantum computers for cryptanalysis are not available?

Can we find useful formulas for the quantum capacities, classical capacities, and private capacities of quantum channels?

Simulation

What quantum many-body systems of interest in physics and chemistry can be simulated using classical computers? Using quantum simulators? Using large-scale quantum computers?

For ground states of gapped local Hamiltonians in one spatial dimension, there is a succinct classical description in terms of matrix-product states, and the entanglement entropy of a subsystem scales like the size of its boundary (the "area law"). How do these results generalize to other physically relevant quantum many-body systems, particularly in higher dimensions?

Can exotic nonabelian topological phases of quantum matter be realized using ultracold atoms or molecules, and can the nonabelian statistics of the quasiparticles be confirmed?

Are there efficient mechanisms for cooling ultracold atoms to the low temperatures needed for studying strongly correlated physics, and what are the fundamental limits on heating rates in such systems?

Can quantum-many-body systems be simulated reliably by analog quantum simulators without using error correction?

What quantum phase transitions and other exotic collective phenomena can occur in quantum networks?

Can quantum computers efficiently simulate models of quantum field theory and quantum gravity?

Physics foundations

Will as yet undiscovered principles of fundamental physics prevent large-scale quantum computers from ever working?

Can the strength of nonlocal correlations allowed in quantum theory be derived from deeper physical or mathematical principles?

Can insights about quantum information deepen our understanding of the foundations of quantum statistical physics and thermodynamics?

Can violations of local realism be confirmed in loophole-free experiments?

Systems

What useful quantum information processing tasks can be performed with a small number of logical qubits?

What are the overriding systems-control challenges for large-scale quantum computing, and can these be met?

Will fault-tolerant quantum computing work against all noise mechanisms in realizable quantum processors?

Can we build self-correcting quantum memories that protect quantum states physically, without active error correction?

Can we find systems in two or three dimensions, either "naturally" occurring or purposefully engineered, that are suited for topologically protected universal quantum computation?

How well can we control quantum systems using coherent feedback?

Implementation / Hardware

Can any of the proposed quantum processors reach the accuracy requirements for scalable fault-tolerant quantum computing?

What are pros and cons of gallium arsenide, silicon, carbon, and other materials as substrates for spin qubits, and how can various decoherence effects be minimized in these materials?

What methods for encoding logical spin qubits are best suited for minimizing decoherence, improving control, and realizing error correction?

What are the pros and cons of phase qubits, flux qubits, and charge qubits in superconducting circuits, and how can we improve the coherence times and controllability of these qubits?

Can we do useful continuous-variable quantum information processing with superconducting circuits and microwave resonators?

Can we perform quantum information processing in a solid-state system at room temperature, for example using nitrogen-vacancy centers in diamond?

Can we confirm nonabelian statistics in fractional quantum Hall systems, and demonstrate measurement-only topological quantum computing?

Can the optical control required for large-scale trapped-ion or trapped-neutral-atom quantum computing be achieved, and can the laser power requirements be substantially reduced, for example by executing quantum gates using radio-frequency magnetic fields?

Can optical links be used to distribute quantum entanglement efficiently between separated ion trap processors, between trapped-neutral-atom processors, or between such processors and other quantum processors?

Can accurate long-distance entangling gates be performed on pairs of neutral atoms using, for example, highly excited Rydberg states?

Can we realize a practical, scalable quantum computer that processes quantum information encoded in photons?

Can we cool mesoscopic mechanical oscillators to their motional ground states, squeeze them, and entangle them?

Implications

What are the advantages in principle of entangled and other non-classical states for metrology, and can these produce practical gains? For example, can we exploit quantum entanglement to improve the sensitivity of atomic clocks and gravity gradiometers by orders of magnitude?

Can we improve the sensitivity of quantum sensors using novel methods for error and noise suppression?

Can we exploit the enhanced communication capabilities of quantum channels to improve communication rates over presently installed telecommunication systems?

Can we find powerful practical applications of quantum computers to quantum chemistry, for example by speeding up simulations of chemical dynamics and computations of molecular ground-state and excited-state energies?

Can quantum simulations using ultracold atoms in optical lattices or arrays of trapped ions deepen our understanding of high-temperature superconductivity, and suggest new materials that are superconducting at higher temperatures?

Can concepts from quantum information science help us understand natural processes like energy transfer in photosynthetic complexes?

Can spin-offs from research on fault-tolerant quantum systems facilitate the realization of low-power fault-tolerant classical computers built from unreliable nanoscale components.

Can algorithmic cooling enhance magnetic resonance imaging?

Can quantum information methods help us to image biomolecular systems at the atomic scale?

What totally novel and unexpected applications can we find for the essentially quantum phenomena of superposition and entanglement?

Appendix A

Workshop on Quantum Information Science
23-25 April 2009, Vienna, VA
Program

Thursday Morning
8:30 -- 8:45 Carl Williams, The NSTC, OSTP, and the SQIS
8:45 – 9:20 Charles Bennett, Quantum information theory
9:20 – 9:55 Charles Marcus, Semiconductor qubits
9:55 – 10:30 Umesh Vazirani, Quantum algorithms and complexity
10:50 – 11:25 Robert Schoelkopf, Quantum computing with superconducting circuits
11:25 – 12:00 Michael Freedman, Topological quantum computing

Thursday Afternoon
1:45 – 2:20 John Preskill, Fault-tolerant quantum computation
2:20 – 2:55 Mikhail Lukin, Hybrid approaches to quantum information science
2:55 – 3:30 Leonard Schulman, Quantum algorithms with exponential speedups
3:50 – 4:25 Mark Kasevich, Atom interferometry
4:25 – 5:00 Scott Aaronson, Quantum complexity and fundamental physics
5:00 – 5:35 Isaac Chuang, Quantum architecture: from devices to systems

Thursday Evening
7:30 – 9:00 Open session

Friday Morning
8:45 – 9:20 Peter Zoller, Quantum information science with AMO
9:20 – 9:55 David Wineland, Quantum information processing and metrology with ions
9:55 – 10:30 Andris Ambainis, Quantum algorithms with polynomial speedups
10:50 – 11:25 William Phillips, Quantum information, computing, and simulation with cold atoms
11:25 – 12:00 Anthony Leggett, Testing quantum mechanics towards the level of everyday life: recent results and current prospects

Friday Afternoon
1:45 – 2:20 Dorit Aharonov, Quantum Hamiltonian complexity
2:20 – 2:55 Raymond Laflamme, NMR quantum information processing: successes and challenges
2:55 – 3:30 Barbara Terhal, Complexity of simulating quantum systems on classical computers
3:50 – 4:25 Paul Kwiat, Optical quantum information processing
4:25 – 5:00 Alán Aspuru-Guzik, Quantum computation for chemistry
5:00 – 5:35 Anne Broadbent, Universal blind quantum computation

Friday Evening
7:30 – 9:00 Open session

Saturday Morning
8:45 – 9:20 Norbert Lütkenhaus, Quantum key distribution
9:20 – 9:55 Jeff Kimble, Quantum networks: the interface of light and matter
9:55 – 10:30 John Watrous, Modeling quantum interactions as games
10:50 – 11:25 Keith Schwab, Experimental pursuit of the quantum aspects of mechanical motion
11:25 – 12:00 Birgitta Whaley, Quantum control of qubits and quantum systems

Appendix B

Workshop on Quantum Information Science
23-25 April 2009, Vienna, VA
Participants

1	Aaronson	Scott	MIT
2	Abo-Shaeer	Jamil	Booz Allen
3	Aharonov	Dorit	Hebrew University, Jerusalem
4	Ahmed	Ergin	Temple University
5	Aizenman	Morris	National Science Foundation
6	Alsing	Paul	Air Force Research Laboratory
7	Altepeter	Joseph	Northwestern University
8	Ambainis	Andris	University of Latvia
9	Arrowood	Drew	Microsoft
10	Aspuru-Guzik	Alán	Harvard University
11	Aubrey	Joysree	Los Alamos National Lab
12	Barnum	Howard	Los Alamos National Laboratory
13	Barton	Daniel	Sandia National Laboratories
14	Beavan	Sarah	NIST / ANU
15	Behrman	Elizabeth	Wichita State University
16	Behunin	Ryan	University of Maryland
17	Bennett	Charles H.	IBM Research Yorktown
18	Berberian	John	Berberian & Company
19	Bienfang	Joshua	NIST
20	Boisvert	Ronald	NIST
21	Boshier	Malcolm	Los Alamos National Laboratory
22	Brandt	Howard	Army Research Laboratory
23	Braun	Daniel	University Toulouse and JQI
24	Broadbent	Anne	IQC, University of Waterloo
25	Brown	Ben	DOE Office of Science
26	Brown	Roger	JQI
27	Buice	Michael	NIH
28	Calder	Austin	National Security Agency
29	Caldwell	Denise	NSF
30	Campbell	Wes	U. Maryland / JQI
31	Carroll	Malcolm	Sandia National Labs

32	Chapuran	Thomas	Telcordia
33	Chen	Jun	NIST
34	Chijioke	Akobuije	NIST
35	Chuang	Ike	MIT
36	Clark	Charles	NIST/ONR
37	Côté	Robin	UConn
38	Cross	Andrew	SAIC
39	Curcic	Tatjana	AFOSR
40	Deng	Lu	NIST
41	Deutsch	Ivan	University of New Mexico
42	Dowling	Jonathan	Louisiana State University
43	Draper	Thomas	National Security Agency
44	Dutt	Gurudev	University of Pittsburgh
45	Dutton	Zachary	BBN Technologies
46	Eastin	Bryan	NIST
47	Economou	Sophia	Naval Research Lab
48	Edwards	Mark	Georgia Southern University
49	Flammia	Steve	Perimeter Institute
50	Fleming	Chris	UMD:CP
51	Franson	James	UMBC
52	Freedman	Michael	Microsoft
53	Fuller-Mora	Wendy	NSF
54	Galang	Jemellie	NIST
55	Goldberg	Lawrence	National Science Foundation
56	Goldfield	Evelyn	National Science Foundation
57	Goldschmidt	Elizabeth	Joint Quantum Institute
58	Graber	James	Library of Congress
59	Habif	Jonathan	BBN Technologies
60	Hall	Tracy	Brigham Young University
61	Han	Siyuan	University of Kansas
62	Harrington	Jim	Los Alamos National Laboratory
63	Hearne	Sean	Sandia National Labs
64	Heiligman	Mark	IARPA
65	Ho	Kwan-yuet	University of Maryland
66	Houck	Andrew	Princeton University
67	Hu	Anzi	Joint Quantum Institute,UMD/NIST
68	Hu	Xuedong	JQI, University of Maryland
69	Hughes	Richard	Los Alamos National Lab
70	Jacobs	Bryan	Johns Hopkins
71	Jingyun	Fan	NIST
72	Joynt	Robert	Univ. of Wisconsin-Madison
73	Kaminsky	William	MIT
74	Kannan	Sampath	NSF
75	Kasevich	Mark	Stanford
76	Kim	Kihwan	JQI
77	Kimble	Jeff	Caltech

78	Kosloski	Jon	LTS
79	Kotochigova	Svetlana	Temple University
80	Krause	Jeff	Department of Energy
81	Kruger	Marvin	Lab for Physical Sciences
82	Kuo	Paulina	NIST
83	Kwiat	Paul	Univ. Illinois
84	Lababidi	Mahmoud	GMU
85	Lackey	Brad	National Information Assurance Research Laboratory
86	Laflamme	Raymond	University of Waterloo
87	Landahl	Andrew	University of New Mexico
88	Lanzagorta	Marco	ITT Corporation
89	Lee	Jason	George Mason University
90	Leggett	Tony	U. Illinois at Urbana-Champaign
91	Leibholz	Stephen	TechLabs
92	Lin	Yu-Ju	NIST
93	Ling	Alexander	JQI and NIST
94	Lippel	Philip	WTEC
95	Liu	Yingmei	JQI / NIST
96	Logan	George	L. P. S.
97	Lukin	Mikhail	Harvard University
98	Lundblad	Nathan	NIST/JQI/UMD
99	Luo	Le	University of Maryland
100	Luong	Bao	DoD
101	Lütkenhaus	Norbert	Institute for Quantum Computing, Waterloo
102	Luvaul	Michael	McMurry University
103	Lyyra	Marjatta	Temple University
104	Malkova	Natalia	NIST/JQI
105	Manferdelli	John	Microsoft Corporation
106	Marcus	Charles	Harvard University
107	Maska	Maciej	University of Silesia
108	Maslov	Dmitry	National Science Foundation
109	Mathey	Ludwig	NIST
110	Matsukevich	Dzmitry	Univ. of Maryland
111	Maunz	Peter	JQI / U of Maryland
112	McCracken	James	SAIC
113	Meisner	Bob	NNSA
114	Metcalfe	Michael	NIST
115	Migdall	Alan	NIST, JQI
116	Miller	Keith	Laboratory for Physical Sciences
117	Miller	Warner	Florida Atlantic University
118	Mitra	Kaushik	NIST, Gaithersburg
119	Mizel	Ari	Science Applications International Corporation
120	Monroe	Christopher	JQI and Univ. Maryland
121	Mote	Safa	UMD
122	Muniz	Sergio	JQI - UMD/NIST
123	Nam	Sae Woo	NIST

124	Obenland	Kevin	SAIC
125	Orozco	Luis	Joint Quantum Institute
126	Peters	Nick	Telcordia Technologies
127	Pfister	Olivier	U. of Virginia
128	Phillips	Bill	NIST
129	Phillips	Nate	College of William & Mary
130	Plunk	Gabriel	University of Maryland
131	Polyakov	Sergey	NIST/JQI
132	Porto	Trey	NIST
133	Preskill	John	Caltech
134	Quraishi	Qudsia	University of Maryland/JQI
135	Restelli	Alessandro	NIST
136	Reynolds	Peter	ARO
137	Roenigk	Karl	ODNI
138	Rohlfing	Celeste	NSF
139	Rokhinson	Leonid	Purdue University
140	Rolston	Steven	JQI - Univ. of MD
141	Rudolph	Terry	Imperial College
142	Rutter	Natalia	Georgetown University
143	Santamore	Deborah	Temple University
144	Schlosser	Malte	NIST
145	Schneeberger	Will	NIARL
146	Schoelkopf	Robert	Yale University
147	Schulman	Leonard	Caltech
148	Schwab	David	UCLA
149	Schwab	Keith	Caltech
150	Shah	Jag	DARPA
151	Simons	Matt	Coll. of William & Mary
152	Sinha	Kanupriya	University of Maryland
153	Sofge	Don	Naval Research Laboratory
154	Srinivasan	Kartik	NIST
155	Stwalley	William	Univ. of CT
156	Subasi	Yigit	University of Maryland
157	Sulcoski	Mark	DoD
158	Svore	Krysta	Microsoft Research
159	Tahan	Charles	DARPA/Booz Allen
160	Tarman	Tom	Sandia National Labs
161	Terhal	Barbara	IBM Research
162	Thorbeck	Ted	JQI: NIST and U. of Maryland
163	Tian	Ming	George Mason University
164	Tiesinga	Eite	JQI/NIST
165	Troupe	James	U.S. Navy
166	Tsai	Chin-Chun	Cheng Kung University, Taiwan
167	van Dam	Wim	UC Santa Barbara
168	Vazirani	Umesh	u.c. berkeley
169	Walck	Scott	Lebanon Valley College

170	Warchall	Henry	National Science Foundation
171	Watrous	John	University of Waterloo
172	Whaley	Birgitta	UC Berkeley
173	Wilde	Mark	Science Applications International Corporation
174	Williams	Carl	Executive Office of the President
175	Wineland	David	NIST
176	Wu	Saijun	University of Maryland
177	Yard	Jon	Los Alamos National Laboratory
178	Zoller	Peter	University of Innsbruck