

May 16, 2006

The prime goal of my research is the investigation of new models of computation and information processing, as manifested in the rapidly evolving area of quantum information, computation and cryptography. This field is usually called “Quantum Information Processing (QIP)”.

First, I am highly interested in finding and developing new applications in this area, and in particular “near-future” ones. Second, as the usage of the full power of quantum information processing is sometimes far beyond current technology, I focus my research on analyzing the power of more “restricted models” or “limited models” – namely, models in which various practical considerations are taken into account, and “semi-quantum” models in which the allowed states or operations are restricted in some respects. Third (and closely related to the above two), I study the complex relationship between theoretical QIP achievements, proposed models for the physical realization of quantum information units (quantum bits, quantum trits, etc.), and the actual experimental implementations of these models. My goal here is to bridge the gap between QIP achievements and experimental implementations by creating more appropriate models and by suggesting novel experiments.

My main contributions¹ to the field in recent years have been

- proving the *security* of theoretical quantum key distribution (QKD) [C4,A1], while also proving the *insecurity* of practical QKD schemes [C5,J13]
- studying the power of limited models for quantum information processing: 1.— obtaining “quantum nonlocality without entanglement” [J8,J9,J18], “quantum computing without entanglement” [J20,A4], and even some form of “classical nonlocality” [J24]. 2.— studying modeling and improving the scalability of several suggested realizations of quantum computing devices [J12,J17,Sub1]
- inventing what might become the first near-future application of quantum computing devices – the “algorithmic cooling of spins” [J17,J21,Pat1]

1 The QIP Research Field

Computer and information sciences are based on manipulating *bits*, and the implication of this fact is that these sciences are based on the principles of classical physics. However, in a world ruled by the laws of quantized systems, such classical models are not the most powerful, and other models seem to have much more power. Assuming that the rules of quantum mechanics are correct, quantum computation suggests an algorithm which factorizes large numbers exponentially faster than any known classical algorithm; quantum cryptography suggests schemes for information secure key distribution, a task which is beyond the ability of any classical scheme.

From a practical point of view, even if we were not now aware of QIP, we would have to begin studying it very soon, as the Moore laws regarding the size of a computing unit (the

¹J ≡ Journal publication; A ≡ Accepted for publication in a Journal; C ≡ publication in Conference proceedings; Sub ≡ Submitted to a journal; Pat ≡ Patent.

transistor), the energy/power consumption of a computing step, or the time of a computing step (etc.), would lead us in that direction. The Moore laws would lead us into the quantum domain of information processing and therefore into the QIP world, in about 10-15 years, at the point of having reached the 1-10 particles limit. There can be no certainty that it will ever prove possible to build a useful quantum computer that will factorize large numbers. Nevertheless, the journey towards the implementation of quantum computation is already yielding valuable spin-off technologies and applications in fields such as communication security, detection, and nano-technologies (the control of quantum systems).

While bits can only take the values 0 or 1, quantum bits (two-level systems, which we simply call *qubits*) can be in a superposition of these classical values, and unlike the classical states, an unknown quantum state *cannot* be cloned. Furthermore, while a string of bits has a particular value (composed of the values of the individual bits), a string of qubits can be in an *entangled state*, a state which shows novel non-local effects and parallelism. Entanglement and no-cloning are the pillars of non-classicality and are at the root of the extra power which quantum computer and information sciences promise.

In recent years a great deal of research has focused on various limited QIP models, sometimes due to their practicality, sometimes in order to better understand the conceptual borderline between quantum information processing and classical information processing, and always in order to bridge the world of theory and the world of experiment.

2 Research Contributions

In the following I shall discuss my major contributions in the areas of no-cloning, entanglement, the road from theory to practice, and near-future applications of quantum computing devices.

NO-CLONING: In the theoretical QKD model Alice (the sender) sends qubits to Bob (the receiver) in order to distribute a secret and random key. The security of QKD is closely related to the no-cloning principle. Yet, a formal proof of security relies on much more than just the basic no-cloning principle. Based on an “information versus disturbance” technique, symmetrization of any attack, and the law of large numbers, we [C4,A1] provided a full proof of security for the theoretical model.

As an important step towards that full proof of security we first provided a complete proof of security of QKD against all collective attacks [J16], an important class of attacks defined in [J5] that might be quite realizable to a near-future eavesdropper.

We [C5,J13] also considered a more realistic model, in which photons (rather than abstract “qubits”) are used to transmit the quantum information. We determined that a proper approximation of reality requires a model in which photons (the carriers of the information units) are not qubits but qu-hexit (six-dimensional Hilbert space replaces the two-dimensional one of the purely theoretical model). Based on this realistic model, we suggested a new attack, the photon-number-splitting attack, and we concluded that experimental schemes are actually highly or completely insecure. Our work showed that quantum information, in many realistic QKD schemes, *can be cloned*. This well-cited work created a shift in the field of practical aspects of QKD by highlighting the importance of properly describing the relationship between theory and experiment. Since then, many additional attacks along similar lines have been proposed, and relevant modifications of the QKD protocols (more resistant against such attacks) have been proposed and implemented.

We recently considered QKD in which one party is classical, asking “Can such a scenario still lead to a secure key?” We designed such a protocol and presented it at a conference. This paper is currently in preparation.

ENTANGLEMENT: While it is crystal-clear that entanglement is vital for “fully-powerful” QIP, it is much less clear which achievements of QIP remain if no entanglement is used or if the amount of entanglement used is severely limited. In a series of papers [J8,J9,J18], we initiated research on the topic of “quantum nonlocality without entanglement” (and a resulting phenomenon - the existence of unextendible product bases). These well-cited papers have led to a large body of research in this area.

One of the resulting directions is our more recent investigation of “quantum computing without entanglement” [J20,A4]; very recently I found a form of “classical nonlocality” — the classical analogy of the so-called “quantum remote steering” phenomenon [J24].

THE ROAD FROM THEORY TO PRACTICE: The road from theoretical QKD to practical QKD was already discussed in the previous paragraph. A closely related research direction is to look at the various realistic models for quantum computing and investigate their power, and in particular, their scalability.

Quantum physics tells us that, in principle, one can manipulate qubits by performing several operations. One can store and send them; initialize and measure their state; and transform their states, each one alone, or a few together. Unfortunately, quantum physics does not tell us *how* to manipulate qubits, and as yet, there is no practical proposal of a scalable system in which all of these basic operations can be done. Isolation of the qubits from the environment on one hand, and the ability to control the qubits on the other hand, are not easily reconciled.

A large number of interesting proposals have emerged, in which specific simple tasks (on small scale quantum computing devices composed of 3-8 qubits) can be implemented. Three of the main setups investigated in recent years (as leading candidates for quantum computation) are silicon structures, nuclear magnetic resonance (NMR), and linear-optics. I was involved in studying these three models, and in suggesting alternative models with better scalability: I am coauthor of a well-cited paper on silicon-based single-electron transistors [J12], I invented a scalable NMR quantum computer based on algorithmic cooling of spins [J17,J21,Pat1] (see the next paragraph), and we proposed scalable linked-state and cluster-state models via linear-optics [Sub1].

NEAR-FUTURE APPLICATIONS OF QUANTUM COMPUTING DEVICES:

In the process of suggesting a scalable model for NMR quantum computers we realized that our novel technique [J17,J21,Pat1] “algorithmic cooling of spins” could actually become the first near-future application of quantum computing devices. The basic idea is to use simple quantum algorithmic methods and novel data compression techniques in order to better control the entropy of quantum systems. Specifically, we invented a way to improve the signal-to-noise ratio in NMR spectroscopy (MRS) and imaging (MRI) by cooling spins via pumping entropy out of the system. Our method shows an exponential advantage [J21] over the entropy-preserving methods; we can cool a single spin (in an N -spin system) by a factor which is exponential in $N/2$, while any entropy-preserving method can at most cool that single spin by a factor of \sqrt{N} . We later found two methods for cooling by the (near-optimal) factor of $2^{(N-2)}$ [J23,A3,Sub2], and we also provided an experimental demonstration for cooling an NMR system beyond the entropy-preserving bound [Sub4].

ADDITIONAL PROJECTS: We analyzed algorithms that use imperfect gates, since such gates are important for practical quantum computation. We found a universal set of gates that can be implemented fault-tolerantly [C2,J14]. Our set of gates is extensively used in Nielsen and Chuang’s textbook. We also demonstrated [C9] how to run fault-tolerant algorithms on a limited type of quantum computing called ensemble computing, which is crucial for large-scale NMR quantum computing.

My interest in understanding the power of QKD versus protocols in modern cryptology has led me to explore what might remain of modern cryptology in a quantum environment. While the famous RSA cipher (and various other ciphers) is totally insecure in a quantum world, it is yet unclear if there is any other good candidate for “public key encryption”. People have recently begun to refer to this sub-field as “Post-Quantum Cryptography”. We looked at lattice-based ciphers (which are potentially secure even in a quantum environment), and we identified [A2] a weakness of theirs — they are insecure against chosen ciphertext attacks.

3 Research Plans

In future projects I intend to continue exploring new aspects of quantum computer and information sciences. My primary intent is to help lay the conceptual foundation which will allow us to take our field from *theory* to the world of practical implementations, where analysis of security and analysis of computing power are still in their initial stages.

I would like to prove the security of practical quantum key distribution, by combining theoretical results with analysis of practical schemes. I’m currently working on defining the most general attack, while taking into account the Hilbert spaces used by the legitimate users in realistic scenarios. My ultimate goal is either to yield a *complete* proof of security for a *running experiment*, or to provide strong evidence that such a proof cannot be reached.

Another of my goals is to reach a point at which quantum computing can be done in a system of many qubits. All current implementations suffer from severe scaling problems; In many cases of specific implementations, the variations which are better scalable are also further away from being practical (from an experimental point of view). My goal is to assist in obtaining systems in which many qubits, hopefully 50-500, can be manipulated, and potentially, several registers of qubits can be used. The ultimate goal is to reach scales where quantum computing solves problems which are beyond the ability of classical computers.

In addition, I would like to find and develop near-future applications, namely, my goal here is to replace the common question of “how can a certain technology help in advancing QIP?” by the opposite question “how can QIP help in advancing other technologies and resolving their problems?” QIP provides new methods for controlling various nano-systems (and even pico-systems). Algorithmic cooling of spins is one such example, and I would like to use algorithmic cooling for improving the MRS of bio-molecules in real-life applications such as brain research or cancer research (for which MRS would be used more often if the signal-to-noise ratio were better).

I have several additional research ideas that have not yet crystallized into concrete plans; I would like to combine my understanding in computation, information, quantum theory, thermodynamics, and spectroscopy, in order to contribute to an understanding of learning machines, artificial intelligence, and neural networks.