My Encounter with Quantum Computing

An Introduction to Quantum Computing

All of today's computers function on certain basic principles: data is encoded as 0's and 1's, and computation is performed by using logic gates to switch 0's and 1's. Since computer memory and computational speeds have been increasing exponentially during the last 50 years, few people have questioned or altered the underlying computational model. However, it has recently become evident that this golden run cannot last forever. Increases in computing power have historically been driven by miniaturization of computer components such as transistors. However, as component size has entered the nanometer range, it has become evident that in the foreseeable future components will be so small that it will be physically impossible to shrink them further. This revelation has spurred computer scientists to explore other models and methods to improve computational performance. Quantum computing, which exploits the quantum nature of reality to perform computation, is perhaps the most promising and exciting alternative model available today. Theoretical work has shown that quantum computers will be able to perform a plethora of computational tasks unfeasible by today's standards. Using quantum computers, scientists might be able to simulate entire molecules to study their properties, communicate using unbreakable codes, or solve currently intractable equations. The applications are legion. Unfortunately, quantum computers are proving extremely difficult to create. Today's quantum computers are in a stage of infancy; it was considered a tremendous achievement when an experimental quantum computer multiplied three times five a few years ago. Yet, improvements are being made at a steady clip, and we might see a functioning quantum computer within a decade or two. There are numerous considerations that suggest that quantum computing, despite its theoretical and technological challenges, will eventually have a profound impact on our lives.

This last summer, I independently researched a subfield of quantum computing, called quantum architecture, which investigates possible organizational structures for quantum computers. The goal of quantum architectural studies is to develop designs for quantum computers that minimize technological requirements and ease the transition from small experimental quantum computers to large and functioning ones. In my research, I investigated whether current ideas on communication networks in quantum computers were realistic and scalable. My results indicated the adequacy of current models for the time being, but also showed that communications networks were limited by the error probabilities of certain processes in the quantum computers. To build large-scale quantum communication networks, researchers will constantly have to endeavor to reduce these types of error in their quantum computers.

My Personal Research Experience

I first learned about high school research at the beginning of junior year, when I came across the Intel Science Talent Search's website. I first believed it to be a type of science fair, but my assumptions were quickly swept aside when I noticed the rigor of the judging guidelines. I was especially taken aback by the fact that entrants were expected to have made meaningful discoveries in important scientific fields. The Intel competition excited me, since I had always conceived of research as a rarefied activity performed only by the highest echelons of academia. I was eager to do my own research project, so I searched the internet to find how other students had done such projects. I discovered that many had performed their work at special summer camps. These camps taught students the required background material, paired them with real researchers, and gave them the facilities and support to complete their projects. I downloaded the application materials for two such programs. I asked my teachers for letters of recommendation and I wrote my application essays. I mailed in my applications and waited eagerly for my results. I was rather disheartened to learn, near the end of junior year, that I hadn't been accepted by either one.

I found myself in a quandary: I still wanted to do research during the coming summer, but I had no idea how to go about doing it. The task of researching alone seemed daunting, and I didn't have any specified area that I wanted to research. I was still determined to write a research paper, so I decided to try my luck and move forward. As a first step, I had to choose a proper topic. Experimental work was impossible, due to lack of equipment, so I had to go with either theoretical or computer work. I was painfully aware how unprepared my one-year of calculus left me to understand work based on the higher mathematics, so I had to pick a field that didn't rely on math too heavily. I was confident in my programming abilities, so I decided to play to my strengths and research a field in computer science. I used wikipedia to investigate some interesting fields, and I eventually happened across the article on quantum computing. At the risk of sounding corny, I have to say it was love at first sight. I had long been interested in quantum mechanics, and I was surprised and excited to learn that technology had advanced to a point where we could utilize the inscrutably odd effects of the quantum world rather than just study them. The field was difficult and involved a good deal more mathematics than I liked, but I had nothing to lose so I decided to go for it.

Now that I had a field, I had to learn the actual material. I bought myself the leading textbook on quantum computing and decided to devote the first month of summer to learning the field and choosing an appropriately narrow research topic. I optimistically expected that a month of study would give me a good handle on the subject, but the pages of highly mathematical derivations, proofs, and analysis had quite an intimidating effect on me. I had to give myself a crash course in the notation and basic math principles used before I could proceed any further (for those interested, quantum computing relies heavily on quantum computing and some topics of discrete math). I struggled through the material, and sometimes spent hours staring at single pages of especially difficult theorems. After two or three weeks of wrestling with the book, I decided that I had learned enough, and I started to read research papers. I had no delusions that I understood quantum computing as well as professional researchers did, but I knew enough to keep on moving. I had access to the IEEE and ACM computing journals, so I was able to read multiple research papers free of charge. I skimmed through several papers, trying to get a feel for the different types of quantum computing research that were going on. Eventually, I came across Quantum Architecture, which deals with possible layouts for quantum computers. This particular subfield of quantum computing didn't use too many mathematical equations, but instead dealt with issues of organization and structure. Even better, the researchers in this field seemed to rely heavily on computer simulation. I was confident that I could handle the creation of a simulation program. I started reading all the papers I could find in the area. After another couple of weeks, I started to feel like I had a general idea of what I was going to do. I noticed that many papers didn't extensively discuss communications networks. Instead, they often seemed to brush the topic aside after superficial analyses. I thought this a mistake, since the communications networks were crucial and complex components of quantum computing architectures. I saw an opportunity there for doing research, so I decided to test previous ideas on communications networks to see if their claims and assumptions held up under more rigorous analysis. From that point, moving forward was much easier as I gained momentum. It took me a month of long hours (often ten hours a day, seven days a week) to flesh out my ideas and create the computer programs required to investigate them. After finally finishing my work, I started typing up my results into a research paper format, and by the time the submission date rolled around for Intel, I had everything prepared. I submitted my paper and sat back to wait for the results. To my happy surprise, I was selected as a semifinalist in the competition. It was very satisfying to know that my work had been deemed important enough for me to be awarded a thousand dollar prize for it. My experiences have motivated me to seek out more research opportunities and I'm looking forward to being able to the possibility of working with real researchers when I go to UC Berkeley this fall.

Analysis of Quantum Computing Interconnection Networks

In Quantum computing, there is a schism between theoretical and experimental research. Theoretical researchers have uncovered numerous powerful quantum algorithms, but experimental researchers have built only the smallest of quantum computers. Therefore, research efforts have recently focused on quantum architectural studies, which seek to discover how efficient designs for quantum computers can speed the transition from very small experimental systems to full-size functioning ones. My particular research project focused on possible communication networks in quantum computers. I investigated whether ideas about communications networks introduced in previous architectural studies of quantum computers could scale properly as the size of the quantum computer increased. My results revealed that they would scale properly, but with certain caveats. Moreover, I also discovered a new relationship between error and communication networks.

For me to discuss my work in depth, I'll first have to introduce some basic concepts about quantum computers. The basic unit of the quantum computer is called the qubit. The qubit is the quantum equivalent of a bit, but unlike a normal bit, which is limited to being a zero or a one, a qubit is capable of being both 0 and 1 at the same time. This bizarre property is due to a phenomenon called superposition. Another important fact about qubits is that they are only abstractions; a real qubit is actually implemented as any one of a variety of atomic or subatomic particles. The most promising current implementation uses ions as qubits.

An important concern in creating a design for quantum computers is the creation of an efficient communications system. To perform quantum algorithms, it is necessary for qubits to interact directly. This direct interaction requires that the interacting qubits be physically adjacent. Since qubits are (in the implementation we will discuss here) electrically charged ions, the obvious solution would be to move them around using electric fields. Unfortunately, when qubits are moved in this fashion, the data they hold is quickly corrupted. Therefore, prior work advocates using an alternative transportation method. This method, called teleportation, transfers the data that a qubit holds to another qubit already at the destination location without physically transporting the original qubit. Unlike the "Beam me up, Scotty" type of teleportation, there is no instantaneous transport involved. Real teleportation requires a series of steps, which include physical movement of auxiliary qubits to source and destination. To perform a quantum teleportation, it is necessary to create a pair of qubits, called an EPR pair, which is in a special quantum state. One member of the EPR pair is moved to the original qubit's location and the other member is moved to the desired location. Then, a series of quantum procedures is performed on the qubits at the original location and then on the qubit at the destination. Once these processes are finished, the data in the original qubit has been erased and transferred to the EPR pair qubit at the destination. For all intents and purposes, it's as though the original qubit had been physically transported to the destination.

It's important to note that the above description is greatly simplified. It turns out that simply creating one EPR pair and moving its members to source and destination would corrupt the data. Instead, a more complex protocol, that can use dozens or even hundreds of qubits, is necessary to ensure that the members of one pure EPR pair are present at source and destination. Since EPR pairs are themselves qubits, they can be corrupted if they are physically moved for long distances. Therefore, the protocols require that EPR pairs themselves be teleported in a set of small teleportation steps. The figure below (drawn from my research paper) schematically represents one possible teleportation protocol, along with some of the required resources (note: the dotted arrows represent physical movement of EPR pairs with electric fields, and black boxes represent teleportations).

Figure2: The Teleportation Protocol used in[8]

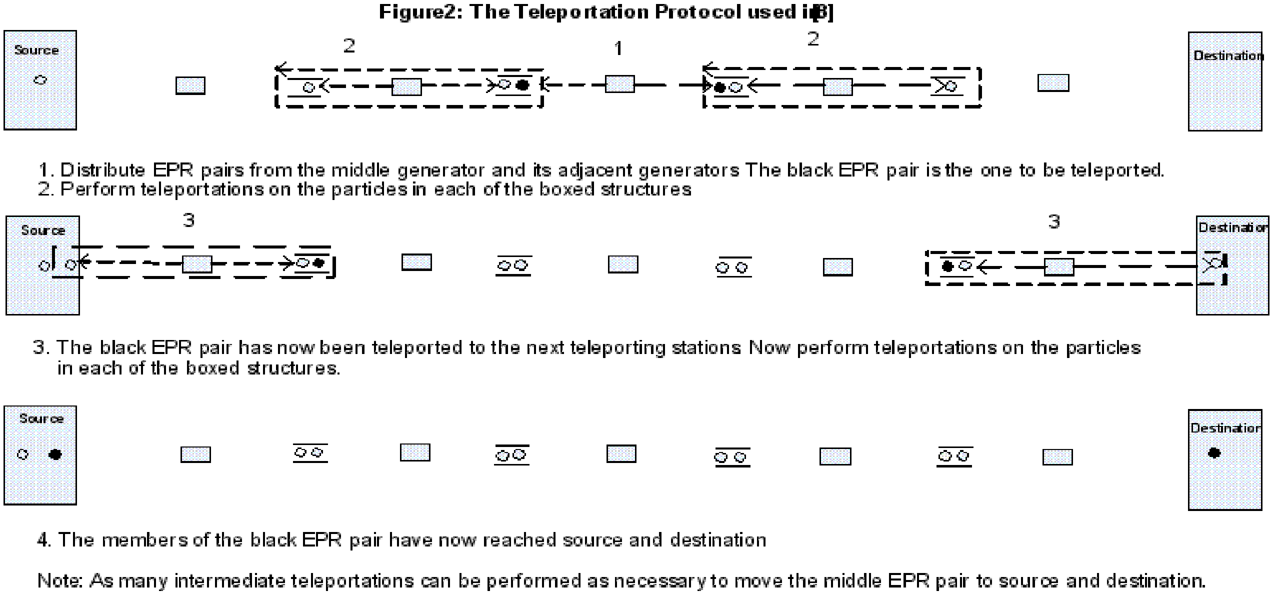

It turns out that moving EPR pairs even small distances can corrupt them. Therefore, it is necessary to perform a purification procedure on them before they are used in teleportation procedure (an important note: this procedure only works on EPR qubits and not on data-holding qubits). The purification procedure works by destroying the states of two EPR pairs to create one purified EPR pair. This procedure will work when the members of an EPR pair are separated, but requires that corresponding members of the two pairs are in the same physical location. Performing the procedure one time doesn't usually grant adequate purity, so multiple rounds of purification are often required. Unfortunately, the hierarchical structure of the purification rounds (see following figure) results in an exponential growth in required EPR pairs with increasing rounds of purification. In the figure below, the structures marked *Pur* are purifiers, which take in two impure EPR pairs, and output a purer one:

Fig. 4: Representation of2-Round Purification Process

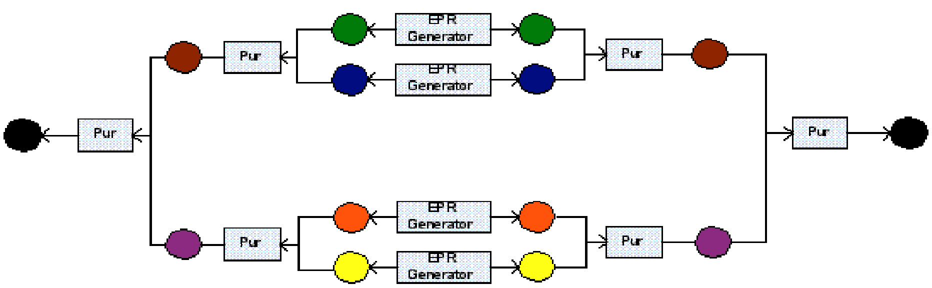

Purification Description
Each box labeled Pur represents a purifier. Each purifier takes in two EPR particles from different EPR pairs and outputs one pair of higher fidelity (we have represented the members of an EPR pair as having the same color). Here, we represent the purification processes required for two purification rounds

A detail that further complicates communications networks is the fact that the "qubits" used in quantum algorithms must actually be made up of dozens or hundreds of ions due to error concern. Therefore, to teleport one "qubit" for an algorithm, dozens of ions must be simultaneously teleported. Thus, in a fairly large quantum computer, thousands or tens of thousands of individual teleportations could be simultaneously occurring. I believed that such a level of complexity merited a proper study, but I found that previous papers on quantum architecture did not do the problem justice. Such papers did provide ideas on how communications networks could be implemented, but they didn't investigate how their schemes scaled as communications networks grew larger and more complex. Upon noticing this point, I decided to investigate the issue for my research paper. I approached the problem from two different angles. First, I decided to analyze the teleportation protocol itself, to see whether its resource requirements scaled reasonably. Second, I decided to analyze the performance of an entire system running the given protocols, to ensure that competition for limited resources didn't become inhibitive.

To accomplish my first objective, I analyzed certain equations that governed resource requirements for the teleportation process. By slightly reworking equations provided in earlier work, I obtained these equations:

$$F_{ORIG} = (1 - P_{1Q}) \times (1-P_{2Q}) \times (1 - P_{MV})^{(D/2)} \quad (1)$$

$$F_{REQ} = [4/3(F_{TEL}) -1/3]^{1/N} \times \{.75 / [(1-P_{1Q}) \times (1-P_{2Q}) \times (4/3\ (1-P_{MS})^{2} -1/3)]\} + ¼ \quad (2)$$

F_{ORIG} is the fidelity of the EPR pairs before they are purified and F_{REQ} is the fidelity the intermediate EPR pairs must be at before they are teleported (fidelity is a quantitative measure of how pure a qubit is, with 0 being complete corrupted and 1 completely pure). P_{1q}, P_{2q}, and P_{MS} are the error inherent in the experimental techniques researchers use to perform various quantum processes necessary for quantum computation. N and D are measures of the total distance between source and destination. Using F_{ORIG} and F_{REQ}, it is possible to calculate the number of purification rounds required to ensure purity. I used a computer program to analyze how this number scaled as the total distance teleported increased. For the most part, I discovered that resource requirements increased linearly with increasing teleportation distance, which vindicated the assumptions of the earlier papers. However, and more importantly, I found that in certain instances, teleportation became fundamentally impossible beyond a given distance. In these cases F_{REQ} became larger than 1, a physical impossibility. By studying equations (1) and (2), I discovered that the only way to increase the teleportation distance is to reduce the values of P_{1q},

P_{2q}, and P_{MS} by developing more powerful and less error-prone techniques to manipulate qubits. This was a most major result of my work.

For the second half of my analysis, I needed to analyze how the resource requirements of the entire communications network scaled when using teleportation protocols. I specifically had to investigate whether competition for limited resources dragged down performance. To do so, I wrote a simulation program in Java that accurately simulated the runtime of a quantum computer. This simulator accepted as input the communications network resources available along with the quantum algorithm to be run on the computer, and outputted the time that the quantum computer would require to perform the given algorithm with the given resources. Using this simulator, I simulated various inputs and discovered the optimal resources required for a variety of quantum computer sizes (optimal resources is defined as the amount of resources beyond which extra resources make no difference in the performance of the quantum computer). I found that optimal resources increased quickly with increasing computer size under my requirements. Upon investigating, I found that this increase was caused by the definition of optimal resources. If I was willing to settle for slightly suboptimal resources, I could achieve nearly constant resource requirements (See *Figure 11* below, which depicts required resources for increasing quantum computer sizes).

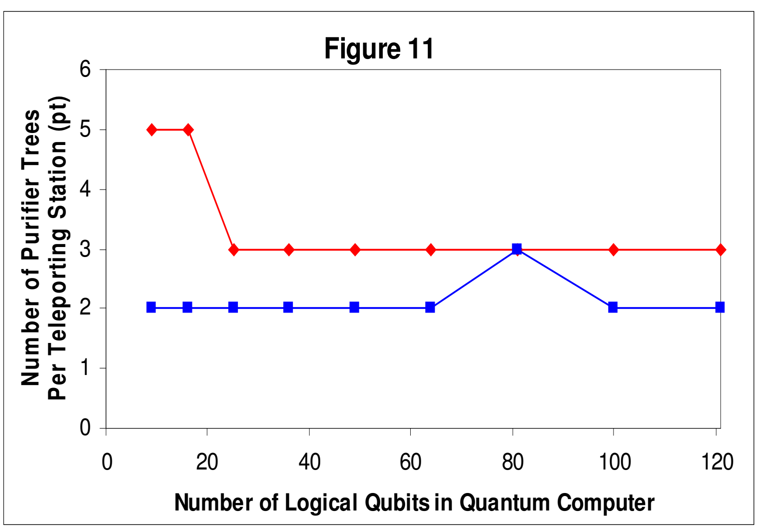

You might note that *Figure 11* has two separate lines. These represent the respective performances of “Mobile Qubit” (red line) and “Home Base” (blue line) architectures. The latter architect demands that qubits teleport back to their original location after the initial teleportation while the former lets them stay at the destination. My tests using the simulator revealed that "Home Base" architecture offered better performance, despite the extra number of teleportations required.

Advice to Prospective High School Researchers

I would strongly recommend any high school student seeking to do scientific or mathematical research to develop a solid foundation in mathematics. My travails with the math of quantum physics have left me with a profound respect for the elegance and utility of the higher mathematics. Therefore, I encourage any prospective high-school researcher to take the core undergraduate classes in mathematics (multivariable calculus, differential equations, linear algebra, and abstract algebra). There are numerous accredited online courses which will teach you the material. One such program (which I sadly learned about after finishing my research) is EPGY, an accredited online program run by Stanford that offers all the classes I've mentioned, along with several even more advanced courses. Gaining a strong foundation in mathematics will open up numerous scientific doors and research opportunities. I would also recommend developing good programming skills. Computer analysis is ubiquitous nowadays in the sciences, and the ability to program simplifies many tedious and complex tasks. I also suggest finding a research topic early, perhaps in ninth or tenth grade, and following it during multiple summers. As you work for longer periods on a topic, you'll be able to come up with more elegant and useful ideas. As a final note, I encourage you to enjoy yourself as you work. At times, it can be stressful trying to think up new ideas within a time limit, but it is important to remember that researching is more rewarding when you do it for the joy of discovery rather than for rewards from a competition.