Practical Adiabatic Quantum Computing: Implications for the Simulation Community

Robert F. Lucas, Gene Wagenbreth & John J. Tran
Information Sciences Institute, Univ. of So. Calif.
Marina del Rey, California
{rflucas, genew, jtran}@isi.edu

David R. Pratt
SAIC
Kailua-Kona, Hawai'i
prattda@SAIC.com

Dan M. Davis
HPC-Education
Long Beach, California
dmdavis@acm.org

ABSTRACT

Despite the asymptotic approach to the limits of transistor-based CPUs, there remains a general expectation of improved computational performance. Quantum Computing is advanced by many as the next major breakthrough that will satisfy those expectations. The adiabatic form of quantum computing has advanced from theory to operational practice in only twelve years. The authors report early results of more than a year's experience on an Adiabatic Quantum Annealer at the University of Southern California – Lockheed Martin Quantum Computing Center, located at USC's Information Sciences Institute (ISI). This represents the emergence of a new age of Quantum Computing, which has the potential for overcoming heretofore intractable computational challenges, thereby improving simulations, enhancing decision support, and enabling innovative data management. The paper first describes quantum annealing and the theoretical orders of magnitude improvements it may deliver. It then outlines the D-Wave installation at ISI and gives examples of early results. Using these data as foundations, the potential in the realm of DoD simulation is discussed, based on the authors' extensive experience with the SAF family of battlefield simulations, including exercises such as U.S. JFCOM's Urban Resolve. The authors rely on their decades of research and operations in High Performance Computing, as well as their experience with the promise and the limits of technologies such as GPUs and Cell Processors, to give an objective over-view of what the members of the modeling and simulation community should realistically expect from this new capability. They discuss a range of the simulation-specific problems that should be amenable to this new technology and forthrightly list a few areas that they believe will not benefit from the various types of Quantum Computing. Real data will be adduced to support their conclusions and to substantiate their predictions and timelines.

ABOUT THE AUTHORS

Robert F. Lucas is the Director of the Computational Sciences Division of the University of Southern California's Information Sciences Institute and a Research Associate Professor in the Department of Computer Science at USC. There he manages research in computer architecture, VLSI, compilers and other software tools. He was the principal investigator on the JESPP project from 2002 to 2011, which first implemented GPU acceleration in high performance computing for battlefield simulations. Prior to joining ISI, he was the Head of the High Performance Computing Research for NERSC at LBNL, the Deputy Director of DARPA's ITO, and a researcher at the Institute for Defense Analyses. Dr. Lucas earned BS, MS, and PhD degrees in Electrical Engineering from Stanford University.

Gene Wagenbreth is a Systems Analyst for Parallel Processing at the Information Sciences Institute at the University of Southern California, doing research in the Computational Sciences Division. Prior positions have included Vice President and Chief Architect at Applied Parallel Research and Lead Programmer at Pacific Sierra Research, where he specialized in tools for distributed and shared memory parallelization of Fortran programs. He has also been active in benchmarking, optimization and porting of software for private industry and government labs. He has programmed on CRAY, SGI, Hitachi, Fujitsu, NEC, PC clusters, networked workstations, IBM SP2, the ILLIAC IV, and conventional machines. He received a BS in Math/Computer Science from the University of Illinois in 1971.

John J. Tran, is a Research Assistant at ISI/USC, where he is currently pursuing a doctorate in Computer Science at USC's Viterbi School of Engineering. Prior employment has included research at the Stanford Linear Accelerator Center, Safetopia, and Intel. His current research centers on Linux cluster engineering, effective control of parallel programs, and communications fabrics for large-scale computation. Major Tran is a member of the California Air National Guard and has served active duty tours at the White House Communications Center and in Iraq. He received both his BS and MS Degrees in Computer Science and Engineering from the University of Notre Dame.

David R. Pratt is a senior member of the technical staff supporting a range of new technology and business initiatives at SAIC. He was the Chief Scientist (Fellow) for SAIC's Strategies and Simulation Solutions Business unit. As a vice president for technology, his responsibilities included developing and fostering leading- edge information technology and M&S technologies. He also served as the Forces Modeling and Simulation point of contact for DoD's High Performance Computing Modernization Program (HPCMP). He received a BSEE from Duke University and MS and Ph.D. degrees in Computer Science from the Naval Postgraduate School.

Dan M. Davis is a consultant for the Information Sciences Institute, University of Southern California (USC), focusing on large-scale distributed DoD simulations, including being the Director of the JESPP project for a decade. As the Assistant Director of the Center for Advanced Computing Research at Caltech, he managed Synthetic Forces Express, bringing HPC to DoD simulations. Prior experience includes being a Director at the Maui High Performance Computing Center and as a Software Engineer at the Jet Propulsion Laboratory and Martin Marietta. He was the Chairman of the Coalition of Academic Supercomputing Centers and has taught at the undergraduate and graduate levels. He saw duty in Vietnam as a USMC Cryptologist and retired as a Commander, Cryptologic Specialty, U.S.N.R. He received B.A. and J.D. degrees from the University of Colorado in Boulder.

Practical Adiabatic Quantum Computing: Implications for the Simulation Community

Robert F. Lucas, Gene Wagenbreth & John J. Tran
Information Sciences Institute, Univ. of So. Calif.
Marina del Rey, California
{rflucas, genew, jtran}@isi.edu

David R. Pratt
SAIC
Kailua-Kona, Hawai'i
prattda@SAIC.com

Dan M. Davis
HPC-Education
Long Beach, California
dmdavis@acm.org

INTRODUCTION

Members of the simulation community are often seeking new capabilities to enhance their software's authenticity, increase its speed, and enrich its analytic power. The growth of capability of traditional digital computing postulated by Moore's Law (Moore, 1966) may be asymptoting at its limits. Hence the need to more seriously reevaluate new computing paradigms. One of the most discussed is quantum computing (DiVicenzo, 1995), and its first practical incarnation, adiabatic quantum annealing (Anthony, 2011). This paper is drafted in a manner that directly addresses issues of immediate consequence to the simulation practitioners and developers. It presents an overview of quantum computing concepts, a review of some early work on a D-Wave quantum annealer at the University of Southern California, and a discussion of the applicability of this technique to simulation. It is not a learned treatise on the theoretical and mathematical intricacies of quantum computing nor is it a document that advocates some specific resolution of the many issues that arise with the introduction of any new technology. The authors rely on their decades of experience in high performance computing and DoD simulations to provide assistance to the simulation community in assessing the potential utility of this advance.

BACKGROUND

According to Gordon Moore himself, the end of Moore's Law is nigh, and it is increasingly daunting to find the path forward for increasing the computing capability that can be applied to simulation and other national security challenges. The capability growth of individual processors is stagnating and the number of such cores needed is now increasing exponentially in high performance computing systems. Size and power demands now often constrain the computational power that can be brought to bear on national problems. In this environment, alternatives to commercial, off-the-shelf (COTS) technology, which would have been inconceivable for most of the last two decades, are becoming of interest. In many ways, this is the reemergence of the purpose built systems many used in earlier decades. These installations include specialized systems such as the Anton at D. E. Shaw Research (Shaw, 2008&2009) It performs certain biomolecular simulations nearly two orders-of-magnitude faster than the largest general purpose computers. They also are looking beyond CMOS to exploit other physical phenomenon, such as quantum computing.

Quantum computing has been considered an attractive extension of computational capability since the seminal paper from the Nobel Laureate Richard Feynman in 1982 (Feynman, 1982), in which he said "… with a suitable class of quantum machines you could imitate any quantum system, including the physical world.". More recently, other authors have touted its ability to produce more computing power, using terms like "magic" to stir the imagination and whet the appetites of the user community. (Gershenfeld, 1998). They point out that the capability of quantum computers arises from the way they encode information differently. Digital computers represent information with transistor-based switches having a state of 0 or 1, labeled a bit. In contrast, the basic unit of quantum computer operation, the quantum bit or qubit, can exist simultaneously as 0 and 1, with the probability for each state given by a numerical coefficient, a condition which physicists call "superposition". The theoretical quantum computer can act on all these possible states simultaneously. The authors are unaware of any such "general purpose" quantum computer that is in operation or on the horizon. However, a more manageable adiabatic quantum annealing device has been conceived, designed, produced, and delivered to the University of Southern California.

The authors have witnessed and participated in the development of high performance computing for several decades, and have developed a significant body of experience with newly introduced technologies. They were engaged in the very early introduction of parallel computing and its rivalry with sequential computing and with vector computing. In this comparison, they heard the detractors of parallel computing argue the limits of parallelism (Amdahl, 1967) and the proponents (Fox, 1994) who argued that it could be used more universally. While acknowledging there are many problems that have remained outside of the easily parallelized arena, it is evident that the majority of all large-scale computational problems are now run in parallel. This is due to the application of new techniques to decompose data and computation in effective ways (Gottschalk, 2005). Such technology has proven very useful to the simulation community (Messina, 1997; Lucas, 2006)

Further, the authors were the recipients of support from the High Performance Computing Modernization Program's (HPCMP) in the form of the first large-scale parallel computer with a general purpose graphics processing unit (GPGPU) on every computational node, installed at the Joint Experimentation Directorate of USJFCOM in Suffolk Virginia. Here again, they heard advocates assert incredible speed-ups and detractors question the utility of the technology. Taking a more pragmatic view, the authors carefully assessed the capabilities of such devices (Lucas, 2010a), measured the energy savings (Davis, D., 2009) and instructed the simulation community users (Wagenbreth, 2010). In one conference, after the presentation of a paper by one of the authors (Lucas, 2010b), a member of the audience stood and pointed out that the analysis was the only one he had heard that rigorously and definitively established both the real potential and the anticipated limits of this technology (Davis, D., 2010). The intent of this paper is to continue in that tradition.

Adiabatic Quantum Annealing

In looking at computational complexity, computer scientists often discuss problems in term of NP-hard or NP-complete. The NP stands for Non-deterministic Polynomial-time. NP-complete problems are sufficiently complex to favor solution by approximation, while NP-hard, although difficult, do have a finite solution. Many problems of concern to the Warfighter fall into this class of NP-hard problems, such as route planning, sensor assignment, and tracking, whose complexity grows too rapidly to be easily addressed using classical, digital computing algorithm ging s. Quantum annealing holds the promise of brining both power and speed to the analyst that is unheard of in digital computing, even massively parallel supercomputing.

The solution space of these types of problems is often conceptually thought of as a three-dimensional landscape. Various solutions are depicted as peaks and valleys. The challenge is to find the highest, or in our case, lowest of these, and not be misled by local minima. If the landscape is big enough, one cannot simply evaluate all of the locations to find the minimum. But one can imagine searching for it with a simple, mechanical system, *i.e.*, an analog computer.

There is a metaphor that may make this clearer. Imagine there is a physical model of this three-dimensional problem landscape. One could drop marbles on it, and watch them roll downhill. Of course, they might be stuck in local minimum, with one or more hillsides standing between them and the true, global minimum.

A technique to improve this method is to shake the table whenever a marble comes to a stop. If the marble is in a valley with a shallow pass, the shaking may cause the marble to roll uphill out of the valley, and then go downhill until it reaches another, lower minimum.

The combination of dropping thousands of marbles and shaking the table in a controlled fashion is akin to the process know as simulated annealing. Shaking the table is equivalent to increasing the temperature in simulated annealing.

Quantum annealing represents an even more powerful heuristic, in which a mechanism is provided to allow one to "tunnel through" the walls which separate local minor minima from the global minimum. No longer does one have to climb the walls and traverse the surface of an optimization function, as required by classical annealing algorithms.

Of course, real problems usually contain a surface with many more than three dimensions. But, given N coordinates of a point it is easy to calculate a value for that point. An N dimensional surface where N is much larger than three is difficult for most to visualize, but the annealing described above, typically realized today as simulated annealing, can be used to find the minimum value of a defined surface via digital simulation.

D-Wave

D-Wave is a small company that makes an adiabatic quantum annealing device which operates at a temperature of below 20 milliKelvin. This is barely above absolute zero, the temperature at which entropy stops, eliminating thermal energy. Absolute zero is defined as 0° Kelvin, or -273.15° Celsius. Published papers are available to detail the technical issues faced and overcome to produce an operating quantum annealer. This paper will not dwell on that here. A good compendium of detailed technical papers is to be found on line at http://www.dwavesys.com/en/publications.html . For the purposes of this paper, it is sufficient to say that the company is a Canadian firm that was founded in 1999 for the express purpose of developing a quantum computing capability.

As early as 2007, D-Wave was demonstrating an operating 28 qubit machine. In 2011, D-Wave announced the 128 cubit D-Wave One (Johnson, 2011), and Lockheed Martin acquired one for the USC – Lockheed Martin Quantum Computing Center (QCC), at USC's Information Sciences Institute (ISI). This has since been upgraded to a D-Wave Two, 512 qubit system.

Figure 1 – D-Wave Two at the USC-Lockheed Martin Quantum Computing Center

Small manufacturing variations and trapped flux in the superconducting circuits resulted in a yield of 503 working qubits. Figure 1 shows the D-Wave Two, as installed at ISI in Marina del Rey. Figure 2 depicts the 128 qubit chip used in the D-Wave One.

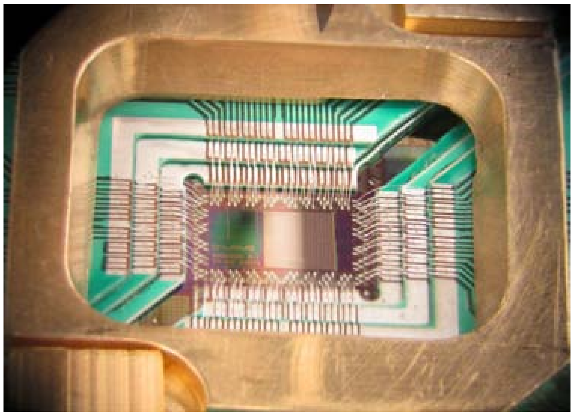

Figure 2 – Detail of Qubit Processor by D-Wave. (D-Wave photo)

Early Results and Analyses

One of the interesting issues raised by skeptics has been whether the D-Wave is actually doing anything "quantum" in its operation. USC and Lockheed Martin are now performing enough calculations at the USC- Lockheed Martin Quantum Computing Center to answer that question and explore the potential of adiabatic quantum annealing. The scientists there, together with their colleagues, have independently verified that the D-Wave is in fact an adiabatic quantum annealer (Boixo, 2012).

There is continued study of the physics of the D-Wave device, for example, trying to ascertain its performance relative to classical computers and also looking for practical applications of AQC. The D-Wave personnel suggest many possible candidates, set forth in Table 1 below. They argue that even the concept of scientific discovery itself is an optimization problem, in which the researcher is trying to ascertain the optimal 'configuration' of parameters that would contribute to a scientific mathematical expression which comports with real world observations. (D-Wave, 2013a). The investigations have so far included software verification and validation (V&V), model checking, sensor assignment, and tracking.

Table 1. Possible Applications for the AQC

Proposed Uses of Quantum Annealing		
Data Mgt.	**Behaviors**	**Analysis**
Labeling Images	Extracting News Stories	Creating/ Testing Hypotheses
Scanning Data for Correlations or Anomalies	Natural Language Performance	Object Detecting in Imagery
Correlating Bio-Informatics	Factor Analysis of Intelligence	Verifying Computer Codes
…	**…**	**…**

POTENTIAL IMPACT ON THE DOD SIMULATION COMMUNITY

As discussed above, quantum computers (QC) may well provide significant potential advantages over classical computational systems. In this section we will discussion some of the areas within the domain of computer generated forces (CGF) where we see the greatest potential for significant advances in the overall performance of the CGF system. These discussions are caveated with observation that there still is significant work to be done in the development of the AQC hardware and even more work, some of it fundamental new algorithms and programming paradigms, to bring some of the discussions to fruition. Some of the more pessimistic estimate that it may be five more years before prototype systems are seen. It may be even longer, if at all, more before production QC based CGF systems are developed. The authors experience is that only time will tell how quickly this new capability will be adopted. If history is a guide, those needing the most power may adopt the technology very soon, even at the cost of having to resolve problems and invent new approaches.

It may be useful here to consider a contrived example, but one that has an appropriate historical foundation. One could posit the existence of a hypothetical unit of the armed forces, say a naval squadron, and further posit the need to split up the armada into two operational units, with the knowledge that the improper allocation of resources would inevitably result in diminished capabilities. A CSG (carrier strike group) typically has a number of smaller ships, AOs, AOEs, *etc*. Each of these provides service to the other ships in the armada, *e.g*. fuel, food, communications, *etc*.

One way to look at this problem is to use a graph theory approach (Lucas, A., 2012). Readers wishing to pursue this theory further could read an introductory article on graph theory that can be found on line (Wikipedia, 2013). For analysis, each ship could be considered as a node and the relationship between the nodes, *e.g.* food, fuel, *etc.*, are called edges by graph theorists. When it is exigent to split up the armada into two CSGs of comparable size, it is prudent to make sure that the assets (entourages) provide sufficient services between the two groups; *i.e*. we want to minimize the number of cross-group dependencies or degradations of operational capability in either group. It would be counterproductive to have a food ship travel additional and unnecessary thousands of miles in the Pacific Ocean to service the two separated CSGs.

This resource partitioning problem is considered to be NP-Complete, the formal mathematical description of which is described in work by Karp (Karp, 1972). In other words if a naval logistics officer were to sit down and map out all the possible combinations to optimize the partitioning, then his task may well be hopelessly convoluted, hence the common identification of this quandary as a combinatorial problem.

Having conceived the problem using this graph theory approach, it is now possible to optimize the solution using AQC. The various parameters are configured as nodes and edges of a graph and the data is submitted to the annealer which can calculate the optimal ground state or states. The D-Wave will return a histogram of the configuration that produces the necessary values for the given configuration. This will lead to a solution in a fraction of the time necessary for a digital computer. More than one run may be required in some situations and not all problems will have a solution.

Framing problems this way is not familiar to many programmers. To make AQC more approachable, D-Wave is developing an advanced interface, which they call a Black Box. While the programmer is always capable of programming the D-Wave device directly, it would be a daunting task for most simulation professionals. Therefore, its developers state that they felt that abstracting away from the overwhelming underlying complexity is critical for making the system broadly accessible. (D-Wave, 2013b)

A Hypothetical Quantum Annealer Implementation

Given that the widespread acceptance of quantum computing is beyond the scope of this paper, part of the focus will be on the issues of mapping the CGF system problem space to the domain of problems that are suited to AQC. Given the newness of the whole quantum computing domain writ large, let alone the only recent access to operations on its subset of quantum annealing, there still remain significant areas of research as to what the ideal problems are. Problems where there are large amounts of uncertainty, *e.g.* the areas of traditional and deferred solution search space problems, have been identified as likely candidates (Grover, 2013). The abilities of AQC to compute multiple solutions concurrently and quickly allows for rapid exploration of both known and possible alternatives. In essence, the use of AQC allows us to rapidly explore and refine the possible solution space using a combination of brute force and iterative techniques. Thus, any problem in the CGF space will have to be expressed in such a way to make it a search space problem to make it amenable to AQC.

As discussed earlier, the AQC *forte* is the NP hard and NP complete problem set. Additionally, the use of quantum annealing as a fundamental technique leads us to problems that can be expressed as graph traversals. One such problem is that of the military planning process where the selection of a course of action that would lead to desired end state; it is NP-hard, similar to the naval forces portioning problem discussed above. This problem can be thought of as a variation of the classic Turing Halting Problem (Davis, M., 1980) were the current state and enemy actions are the inputs and reaching an end state (where complete annihilation or victory has been achieved) is the halting condition.

The structure of the Potential Future Graph

As part of the Deep Green effort, the SAIC Team developed an abstract plan representation that characterizes the plan, and other possible excursions, as a tree data structure that is rooted in current reality with the goal state(s) being a stopping state that represents successful execution of the plan (Pratt, 2008). The plan is the initial thread and the excursions are represented as branches in what becomes potential futures graph (PFG). The traversal through this graph could be thought of as a modified path planning problem, where there are multiple end states, of varying benefit, creating a multiple set of possible stopping conditions. The advantage of this representation is that it allows the formulation of the CGF as a graph search space problem and, therefore, makes the CGF domain amicable to AQC computational architecture. In the PFG representation, each future is represented as series of states, or situations, as are notionally shown in Figure 4 below.

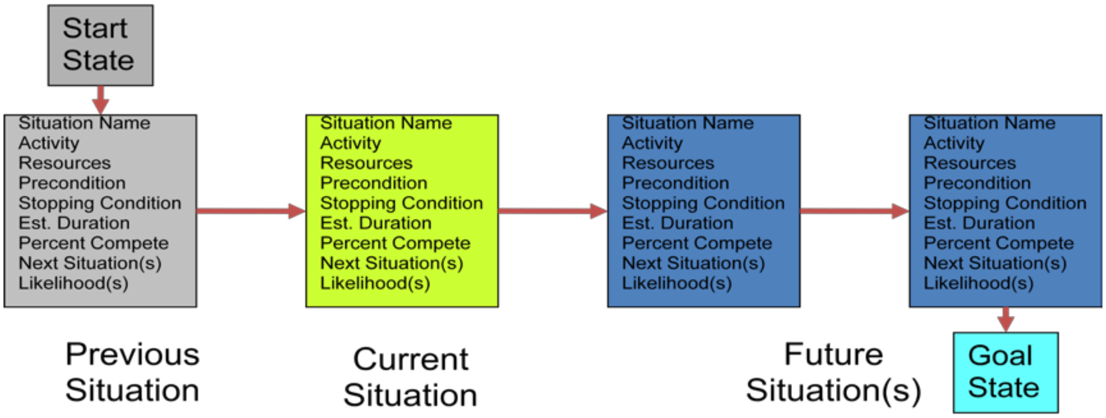

Figure 4. An Alternative Future can be represented as a series of situations and the events that cause transitions between them

The resulting structure in Exhibit 5 (below) allows the user to pinpoint the critical uncertainties by examining where the branches occur. As time unfolds and more data is gathered, more insights are gathered reducing the critical uncertainties and probabilities of many of the branches and effectively trimming the futures graph.

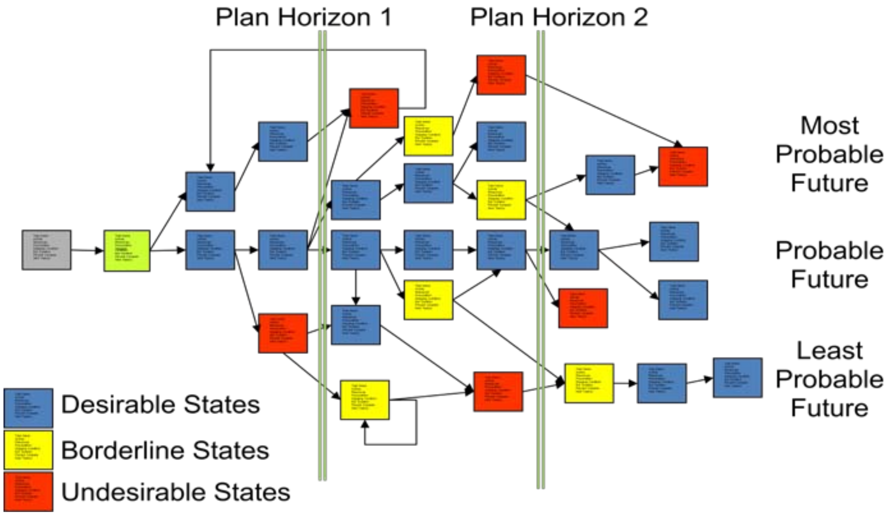

Figure 5. The Potential Future Graph (PFG) representation of the alternative futures shows the events, branches, and desired states.

Likewise, as time unfolds, the root of the graph, the current state, moves; this also changes the graph. Thus, the graph constantly changes and requires constant reevaluation for this approach to be practical. This creates a demand for an incredible amount of computational power within the digital computational paradigm.

Applying QC to Alternative futures

The mapping of the CGF problem to a search space problem would require significant computational resources to allow the user to do a direct mapping to the AQC processing models. However, there still remain significant research issues. Foremost of these is the mapping the nodes to a qubit representation. Since the qubits can be in a 0, 1 or superposition state, it is possible to map sections of the next state to the value of the one or more the qubits. Thus, any path through the PFG can be represented by a string of bits. Advantage can be taken of the uncertainty of the superposition state to represent the inherent uncertainty of the outcomes of processing at the PFG nodes. As discussed above, the two primary types of processing are interactions and decisions. The remainder of this section will discuss the possible techniques for processing of the nodes.

A traditional method for determining the result of combat interactions is the use of some variation of the Lanchester equations. The original equations, and many of their derivations, are based upon quantifying the ratio of the forces to each other in order to determine the attrition and the result of the interactions. While the use of those equations has largely been recently superseded in favor of entity-based models, the equations have been used for years as the basis for aggregate-level models. Non-determinism is introduced into the system via the use of random number generators and output threshold values. Per the current literature, this is the type of computational processes that are ideally suited quantum computing. However, at this time, the results of such implementations are still being analyzed by the authors using emerging data.

The other nodes in the PFG represent decision points. At the aggregate levels; these are staff actions/orders that result from either a preplanned action or in response to environmental stimuli or opposition action. Given that a wide range of artificial intelligence (AI) techniques are based either search or constraint satisfaction, it stands to reason that that some types of AI may also benefit from AQC implementation.

Thus, the planning process might be cast as an NP Hard Turing Halting Problem, that is a problem wherein the question is whether the program will run to completion or run forever. That would possibly make use of the AQC's simulated annealing processing ability to provide graph traversals in significantly less time than traditional digital computation methods.

Given all this, it seems likely that, when constructed as search space system with QC amenable processing elements, the user can map the proposed CGF system architecture to the QC computational architecture. As discussed elsewhere in the paper, there are the programming paradigms that have yet to be developed and job control system that are as yet not implemented. If the existing cryptology work is any indication of the level of effort, we can expect this effort to require new processing algorithms and programming paradigms to implement the versions Lanchester equations and AI elements. Initially, these would probably be toy implementations to prove the principles and the user should expect significant co-evolution of the CGF software elements with the QC hardware elements.

In doing this optimization, a system can be created that could provide the computational power and the resulting insights needed to make the Deep Green initiative a reality (Kenyon, 2007). The effort will require the advent of a new CGF architecture that rephrases the current emphasis on emulation of reality into one that trades precision accuracy for processes speed. By using traditional multiple executions and analyzing the results, the resulting system is expected to more rapidly provide enhanced insights into the solution space within the needed operational timelines.

Adoption of New Technology

Assuming that quantum annealing does in fact prove to be a capable new tool for solving problems arising in simulation and CGF, that is not enough to ensure its successful deployment and adoption. There will have to be significant changes in the current codes and adaptations of simulation paradigms; it is unlikely that simply adding an AQC capability to a codebase will suffice.

With some minor exceptions for research and one-off systems, the authors feel comfortable generalizing that the basic software architecture of the CGF systems has not significantly changed since the advent of Janus / Joint Theater Level Simulation (JTLS) and Modular Semi-Automated Forces (ModSAF) in the early 1990's. The reason for this is that the prevalent computational architecture has not changed since that time. Granted, there have some enhancements, notably the migration from mainframes to work stations to PCs and the advent of object and component based systems, but the basic structure of commodity processing has remained the same. Thus, the software architecture that encapsulated the problem definition remained the same to enable its mapping on the "standard" computational platforms.

That is not to say there haven't been some architectural excursions that have made use of new technologies. A notable instance of a change in the available computational architecture was the advent of user programmable general purpose graphics processing units in the mid-2000s. The promises of using the new found computational engine were quite prevalent. Based upon the GPGPU performance data contained in (Manocha *et al.*, 2004) and as discussed in (Verdesca, 2005), the GPGPU provided a potential source for significant speed up of two of the most computational intensive elements associated with CGF systems: geometric intervisibility (often referred to line of sight (LOS)) and route planning. While the detailed algorithms are contained in the paper, the basic concept was to offload the processing of the LOS and path planning to the GPGPUs as co-processors. Under test conditions, this approach worked remarkably well achieving up to a 20x speed up in execution times of the target subfunction.

So, why aren't these systems in greater use? It turned out that under normal operating conditions, the implementation of the acceleration of the route planning function was often interrupted when the GPGPU was required to multitask between the LOS processing and rendering the image on the screen. In large scale, high performance computers, this is not the case, as rendering is left to the workstations consoles in front of the users. The use of the GPU in the workstations for screen image processing, typically limited the benefit of such system to those environments were the systems were used in the headless mode, *e.g.* the Joint Experimentation on Scalable Parallel Processors work at JFCOM (Lucas, 1993). Furthermore, when used operationally in machines processing both visualization and simulation, the CGF system would be requesting LOS computation for radically different areas of the terrain database in rapid succession. This significantly reduced the ability of the GPGPU to load the processing pipeline and dramatically increased the demands on the memory bus. The net effect of this programming difficulty was, in many cases, to actually reduce the performance of the system due to the miss match between the problem architecture and the underlying computational architecture.

In one attempt to circumvent this mismatch, the authors proposed the use of the CELL processor board to host the environmental runtime component (ERC) and database interfaces (Pratt, 2007). In that paper one of the authors proposed to port the ERC to the CELL processor board, in order to bring the software and computational architectures in line. This option was not pursued for the other reason that many of the excursions that show promise fail to become operational systems: market forces and lifecycle support. From a lifecycle support point of view, the GPU and CELL excursions were not great successes. The underlying market was changing too fast to have a "standard" GPU implementation and the CELL processor failed to make a mass market (other than the PlayStation) impact. Thus, for any CGF system on QC to become mainstream, QC themselves will likely have to become a mainstream technology and have an associated body of standards.

With this history in mind, some believe that unless there are significant market and lifecycle reasons, AQC may wind up being limited to purpose built systems. However, a combination of energy costs and time to solution constraints are creating an environment where even such exotic systems are gaining acceptance again. Then again, quantum annealing may turn out to be one of Clayton Christianson's disruptive technologies that will entirely subsume some major segment of the computational spectrum (Christensen, 1997). The growth of the number of qubits will play a significant role in AQC's increasing utility.

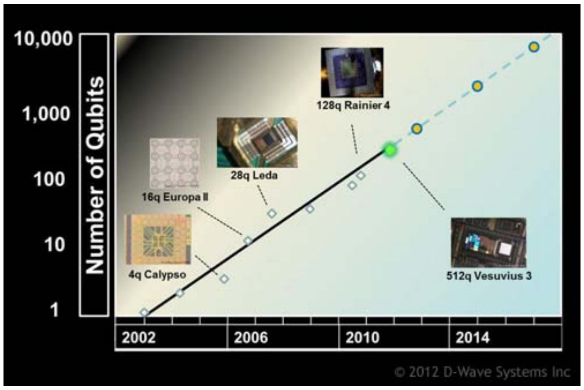

Figure 6. Development Path
(As per D-Wave)

ANALYSIS

One of the major objectives for computational scientists is the assurance that they fully grasp the most pressing needs of their user community. The authors have witnessed and participated in many high performance computing initiatives that have produced amazing results, only to find that those results were not central to the potential users' immediate concerns. As the authors are part of the DoD simulation community itself, there is a sense that such a misdirected effort is less likely in this case, but members of the community have been carefully surveyed to insure a representative view of current concerns. Most illuminating were the comments of Dr. Andy Ceranowicz, who expanded an initial list of simulation grand challenges significantly (Ceranowicz, 2013).

- Augmentation of multi-agent path planning
- Facilitation of entity inter-visibility
- Generation of environmental weather models
- Acceleration of weather radar modeling
- Implementation of better sonar weather modeling
- Injection of Master Scenario Entity List (MSEL)
- Scenario production for preordained MSELs
- Data reduction of multiple-source data
- Production of coherent parametric data
- Data abstraction to evaluate goal achievement
- Production of realistic intelligence reports
- Instantiation of human networks and populations
- Generation of realistic intelligence backgrounds
- Elimination of role players
- Stimulation of players with natural interaction mechanisms such as natural language speech and messages

Experience on the Quantum Annealer would indicate that if the core of each of the challenges listed can be abstracted to an appropriate mathematical representation, the annealer might provide significant breakthroughs. From their own experiences with entity-level simulations such as the SAF family, the authors feel that many of these challenges have major components that could be enhanced and accelerated by the use of quantum annealing, but that is not the end of the cost/benefit analysis.

The total computation productivity effort needs to be taken into account when assessing the value of adopting a new technology (Kepner, 2006). The authors have been pleasantly surprised with other new technologies, when seemingly intractable programming problems have been quickly overcome to facilitate the use of the new technology. However, they have been disappointed to an equal degree when users' legacy codes, thought to be naturals for porting to the new technology, turned out to be more troublesome than the adoption was beneficial.

The best advice based on those previous experiences would be for the user to carefully consider how disruptive the current road blocks are, how much the overall system performance would improve from the relief of those road blocks, what the total costs of such attempts are, and what other resolutions to the problems are imminent. Paying close heed to future developments in the quantum computing discipline will also give the user increasingly useful approaches these cost/benefit analyses, as data on methods, successes, and failures become available from others' efforts.

Ever mindful of the perils of predictions, the authors are nevertheless willing to give their current impressions of the amenability of the listed simulation grand challenges to AQC resolution, all the while expecting to be surprised in both directions. Beginning with the extremes, there is reasonable consensus that AQC holds promise in the approach to multi-agent path planning and that there is little hope it will be useful in inter-entity visibility, that issue being left to the introduction of Feynman's more general purpose quantum computer. With decreasing consensus, some see useful roles for AQC in weather modeling: environmental, radar and sonar, as well as the data management challenges. Then there is the middle ground where more experience with MSELs and scenario production may identify a role for AQC that the authors do not yet see. Less hope is held out for the challenges involving human behavior emulation, *i.e.* intelligence background and reports, role player elimination, and interaction mechanisms. The users will, no doubt, more clearly see opportunities for those difficult challenges, as well as hurdles for the challenges for which promise has otherwise been seen.

SUMMARY

AQC is a powerful new capability, realized in the D-Wave open system adiabatic quantum annealer, which should be available to the CGF community in the near future. It will likely initially be used to solve hard, combinatorial optimization problems, which can include sensor assignment, tracking, and constructing strong classifiers for recognizing specific features in large data sets. It will not replace the ensembles of PCs that serve the simulation community, but rather augment them. In an increasingly distributed, heterogeneous computing environment, it will just be another tool capable of providing solutions in near real time to problems that might otherwise be seen as intractable.

ACKNOWLEDGEMENTS

The authors would like to thank the Lockheed Martin Corporation for having the extraordinary vision to be the first to acquire a quantum computing device and apply it to practical engineering challenges. Those of us at the University of Southern California are proud to be their partners in this endeavor. Special thanks are extended to Dr. Andy Ceranowicz for his contributions to the simulation grand challenges list. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the University of Southern California, Science Applications International Corporation, or the U.S. Government.

REFERENCES

Anthony, S., (2011), “First Ever Commercial Quantum Computer Now Available for $10 Million,” Extreme Tech Website, News announcement retrieved from http://www.extremetech.com/computing/84228-first-ever-commercial-quantum-computer-now-available-for-10-million, on 05 May 2013.

Amdahl, G.M., (1967), “Validity of the single-processor approach to achieving large scale computing capabilities,” In AFIPS Conference Proceedings, vol. 30. AFIPS Press, Reston, Va.; pp. 483-485.

Boixo, S.; Ronnow, T.F.; Isakov, S.V.; Wang, Z.; Wecker, D.; Lidar, D.A.; Martinis, J.M.; & Troyer, M. , (2013), “Quantum annealing with more than one hundred qubits,” arXiv:1304.4595 [quant-ph], and subsequent Blog posts, April, 2013.

Ceranowicz, A., (2013), “Simulation Grand Challenges,” private correspondence between Dr. Ceranowicz and the authors in the spring of 2013

Christensen, Clayton (1997). “The Innovator's Dilemma: The Revolutionary Book That Will Change the Way You Do Business,” Harvard Business Review Press, Cambridge, Massachusetts

D-Wave, (2013a), “Quantum Computer Software Tutorial - Introduction,” User tutorial from D-Wave staff, retrieved from http://www.dwavesys.com/en/dev-tutorial-intro.html on 30 Apr 2013.

D-Wave, (2013b), “Quantum Computer Software Tutorial - Software,” User tutorial from D-Wave staff, retrieved from http://www.dwavesys.com/en/dev-tutorial-software.html on 30 Apr 2013.

Davis, D. M.; Lucas, R. F.; Gottschalk, T. D.; Wagenbreth, G.; & Agalsoff, J., (2009), “FLOPS per Watt: Heterogeneous-Computing’s Approach to DoD Imperatives,” in the Proceedings of the Interservice/Industry Simulation, Training and Education Conference, Orlando, Florida, 2009

Davis, D.M, (2010), personal notes taken while attending conference: the comment from the audience came from Jeremy Kepner of MIT who at that time led the DARPA project investigating High Productivity Computing.

Davis, M., (1980) "What is a computation", in Mathematics Today, Lynn Arthur Steen, ed., Vintage Books, Random House, New York, NY

David P. DiVincenzo (1995). "Quantum Computation". Science 270 (5234): 255–261

Feynman, R. P. (1981), “Simulating Physics with Computers,” International Journal of Theoretical Physics, Vol 21, Nos. 6/7

Fox, G.C, Williams, R.D.; & Messina, P.C., (1994), “Parallel Computing Works!”, Morgan Kaufman, New York, New York

Gershenfeld, N. & Chuang, I.L., (1998), “Quantum Computing with Molecules,” Scientific American, (278), pp. 66-71 (June 1998)

Gottschalk, T.; Amburn, P. &Davis, D., (2005), "Advanced Message Routing for Scalable Distributed Simulations," The Journal of Defense Modeling and Simulation, San Diego, California

Grover, L. K., (2013),” Quantum computers can search rapidly by using almost any transformation,” arxiv.org/pdf/quant-ph/9712011, Retrieved 14 April, 2013 http://arxiv.org/pdf/quant-ph/9712011.pdf

Johnson, M.W., (2011), M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson & G. Rose, (2011) “Quantum annealing with manufactured spins,” Nature 473, 194–198 (12 May 2011)

Karp, R.M., (1972), “Reducibility among combinatorial problems,” Complexity of Computer Computations, Miller and Thatcher, eds.; Plenum Press, New York, NY, pp.85-104

Kenyon, H.S., (2007), “Deep Green Helps Warriors Plan Ahead,” Signals, downloaded 02 May 2013 from http://www.afcea.org/content/?q=node/1418

Kepner, J., ed., (2006), “High Productivity Computing Systems and the Path Towards Usable Petascale Computing: User Productivity Challenges,”, CT Watch, Vol 2, Number 4A, November 2006

Lucas, A., (2012), “Graph Approach to Combinatorial Problems,” independent student research presented informally at USC in the Fall of 2012, slides at: http://www.hpc-educ.org/ALucasGraphSlides.pdf

Lucas, R.; & Davis, D., (2003),"Joint Experimentation on Scalable Parallel Processors," in the Proceedings of the Interservice/Industry Simulation, Training and Education Conference, Orlando, Florida, 2003

Lucas, R. F.; Wagenbreth, G.; Davis, D. M. & Grimes, R. G., (2010a), “Multifrontal Computations on GPUs and Their Multi-core Hosts”, In the *Proceedings of VECPAR’10*, Berkeley, California

Lucas, R. F.; Wagenbreth, G. & Davis, D. M., (2010b), “System Analyses and Algorithmic Considerations in CUDA Implementations for Complex Simulations”, in the Proceedings of the ITEA Annual Technology Review, Charleston, South Carolina

Manocha, D.; Salomon, B.; Gayle, R.; Yoon, S-E.; Sud, A.; Bauer, M.; Verdesca, M.; & Macedonia, M., (2004), “Accelerating LOS Computations using GPUs.” (Brochure), Department of Computer Science: UNC.

Messina, P.; Brunett, S.; Davis, D.; Gottschalk, T.; Curkendall, D.; & Seigel, H., (1997) "Distributed Interactive Simulation for Synthetic Forces," in the Proceedings of the 11th International Parallel Processing Symposium, Geneva, Switzerland, April

Pratt, D.R. , Franceschini, R.W.; Burch, R.B.; & Alexander, R.S. (2008), “A Multi Threaded and Resolution Approach to Simulated Futures Evaluation”, Winter Simulation Conference, Miami, Florida, December, 2008.

Pratt, D. R., (2007), “White Paper on the Use of IBM’s Cell Broadband Processor for Military Simulation,” SAIC White Paper, January, 2007 (available from author).

Shaw, D.E.; Deneroff, M.M.; Dror, R.O.; Kuskin, J.S.; Larson, R.H.; Salmon, J.K.; Young, D.; Batson, B.; Bowers, K.J.; Chao, J.C.; Eastwood, M.P.; Gagliardo, J.; Grossman, J.P.; Ho, C.R.; Ierardi, D.J.; Kolossváry, I.; Klepeis, J.L.; Layman, T.; McLeavey, C.; Moraes, M.A.; Mueller, R.; Priest, E.C.; Shan, Y.; Spengler, J.; Theobald, M.; Towles, B.; & Wang, S.C., (2008). "Anton, A Special-Purpose Machine for Molecular Dynamics Simulation". Communications of the ACM (ACM) 51 (7): 91–97

Shaw, D.E.; Dror, R.O.; Salmon, J.K.; Grossman, J.P.; Mackenzie, K.M.; Bank, J.A.; Young, C.; Deneroff, M, M.; Batson, B.; Bowers, K. J.; Chow, E.; Eastwood, M.P.; Ierardi, D. J.; Klepeis, J. L.; Kuskin, J.S.; Larson, R.H.; Lindorff-Larsen, K.; Maragakis, P.; Moraes, M.A.; Piana, S.; Shan, Y. and Towles, B., (2009), "Millisecond-Scale Molecular Dynamics Simulations on Anton," Proceedings of the Conference on High Performance Computing, Networking, Storage and Analysis (SC09), New York, NY: ACM

Verdesca, M.; Munro, J.; Hoffman, M.; Bauer, M.; & Manocha, D., (2005), “Using Graphics Processor Units to Accelerate OneSAF: A Case Study in Technology Transition,” Interservice/ Industry Training, Simulation, and Education Conference (I/ITSEC) December, 2005

Wagenbreth, G.; Davis, D. M. & Lucas, R. F., (2010), “GPGPU Programming Courses: Getting the Word Out to the Test and Evaluation Community”, in the *Proceedings of the ITEA Annual Technology Review*, Charleston, South Carolina

Wikipedia, (2013), “Graph Theory,” retrieved on 07 May 2013 from Wikipedia, the free encyclopedia: http://en.wikipedia.org/wiki/Graph_theory