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 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 (USC) – Lockheed Martin Quantum Computing Center, located at USC's Information Sciences Institute (ISI). This device heralds 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' substantial experience with the SAF family of battlefield simulations, including experiments such as U.S. JFCOM's Urban Resolve. The decades of research and operations in High Performance Computing, as well as experience with the promise and the limits of technologies such as GPUs and Cell Processors, are used to give an objective over-view of what the members of the modeling and simulation community should realistically expect from this new capability. A range of the simulation-specific problems that should be amenable to this new technology are listed, along with a few areas that the authors 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 a Deputy Director of the Information Sciences Institute where he leads the Computational Sciences Division and he is a Research Associate Professor in the Department of Computer Science at USC. At ISI he manages research in computer architectures, 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 directed High Performance Computing Research for NERSC at LBNL, was 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 USC's ISI, 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, producing tools for distributed and shared memory parallelization of FORTRAN code. Gene has been programming for high performance computers since 1972 and has taught courses in GPGPU acceleration. He received a BS in Math/Computer Science from the University of Illinois.

John J. Tran is a member of the technical staff at the Aerospace Corporation and a Research Assistant at USC, where he is currently pursuing a doctorate in Computer Science. Prior employment has included research at the Stanford Linear Accelerator Center, Safetopia, and Intel. 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 at SAIC, supporting a range of new technology and business initiatives. 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 was commissioned a 2nd Lieutenant in the United State Marine Corps. After his tour of duty, Dr. Pratt later earned 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, 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 serving as a Director at the Maui High Performance Computing Center and as a Software Engineer at the Jet Propulsion Laboratory and Martin Marietta. He has served as the Chairman of the Coalition of Academic Supercomputing Centers and has taught at the undergraduate and graduate levels. As early as 1971, Dan was writing programs in FORTRAN on one of Seymour Cray's CDC 6500's. 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 in need of new capabilities to enhance their software's authenticity, to increase its speed, and to enrich its analytic power. The growth of transistor density on computer chips that was postulated in Moore's Law (Moore, 1966) and the related capability of traditional digital computing are approaching their limits asymptotically. Hence the need to more seriously reevaluate new computing paradigms. One of the most often discussed computational advances is quantum computing (DiVicenzo, 1995) and its first practical incarnation, adiabatic quantum annealing (Anthony, 2011). This paper focuses on issues of concern to simulation application developers. It presents an overview of quantum computing, a review of some early work on a D-Wave quantum annealer at the University of Southern California (USC), and a discussion of the applicability of this technology to simulation. It is not a learned treatise on the theoretical and mathematical intricacies of quantum computing nor does it advocate some specific resolution of the many issues that arise with the introduction of any new technology. Instead, the authors rely on decades of experience in both high performance computing and DoD simulations to offer aid to the simulation community in assessing the potential utility of this advance. The few technical discussions that are included will be of interest to the more mechanically inclined, but will not be vital to those who wish simply to review the analysis at the end of the paper.

BACKGROUND

According to Gordon Moore himself, the end of Moore's Law is nigh. It is increasingly challenging to remain on the current transistor-based path. That path has been enhancing the capability of traditional digital computing and that power has been applied to simulation and other national security challenges. The processing speed of individual processors is now stagnating and the number of such cores needed is increasingly expensive in high performance computing systems. Size and power constraints now often limit the computational power that can be brought to bear on defense problems. In this environment, there is a growing interest in alternatives to commercial, off-the-shelf (COTS) technology. This would have seemed inconceivable for most of the last two decades. Now, even purpose built systems are becoming more attractive. New installations include specialized systems such as the Anton at D. E. Shaw Research (Shaw, 2008 & 2009), which performs certain biomolecular simulations two orders-of-magnitude faster than the largest general purpose computers.

Figure 1.USC- LMC QCC D-Wave 2

Others are looking beyond today’s CMOS technology to exploit other physical phenomenon, *e.g.* quantum computing. Quantum computing has been considered a promising advance in computational capability since the seminal paper from 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.”. The authors are unaware of any such “general purpose” quantum computer that is even nearing operation. However, a more manageable adiabatic quantum annealing device has been conceived, designed, produced, and delivered to USC. Figure 1 shows the D-Wave Two, as installed in the USC – Lockheed Martin Quantum Computing Center (QCC) at the Information Sciences Institute (ISI) in Marina del Rey.

Other authors have touted quantum computing's ability to produce more 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 different way in which they encode information. 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 of each being given by a numerical coefficient, a condition physicists call "superposition". The quantum computer can act on all these possible states simultaneously, a major advantage presented by the introduction of this technology.

The development of high performance computing over several decades has generated a significant body of experience with several other freshly introduced technologies. The early introduction of parallel computing created a rivalry with sequential computing and with vector computing. The detractors of parallel computing argued the limits of parallelism (Amdahl, 1967) and the proponents (Fox, 1994) 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 both decompose data and distribute computation in effective ways (Gottschalk, 2005). Such technology has proven very useful to the simulation community (Messina, 1997; Lucas, 2006)

Further, there was support for simulation from the High Performance Computing Modernization Program (HPCMP) in the form of providing the first large-scale parallel computer with a general purpose graphics processing unit (GPGPU) on every computational node to the Joint Experimentation Directorate of USJFCOM in Suffolk Virginia (Davis, D., 2010a). Here again, advocates were heard asserting incredible speed-ups and detractors were questioning the utility of the GPGPU 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., 2010b). The intent of this paper on quantum computing is to continue in that tradition.

Adiabatic Quantum Annealing (AQA)

Computer scientists often discuss computational complexity in terms of NP-hard or NP-complete. The NP stands for Non-deterministic Polynomial-time. Many problems of concern to the Warfighter fall into the class of NP problems, *e.g*. route planning, sensor assignment, and tracking. Their complexity grows too rapidly to be easily and efficiently addressed using classic digital computing algorithms. Quantum annealing holds the promise of bringing both power and speed to the analyst that is unheard of in digital computing, even massively parallel supercomputing.

The solution spaces of these types of problems can conceptually be thought of as a three-dimensional landscape. Various solutions are depicted as peaks and valleys. In the classic minimization problem, the challenge is to find the highest, or in this 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.

There is a metaphor that may make this clearer. Imagine there is a table-top model of this three-dimensional problem landscape, with countless peaks and depressions representing the extrema, *i.e*. the maxima and minima of the solution. If there were numerous such peaks and valleys, similar to the simple peak and valley shown in Figure 2, marbles could be dropped on the complex space, and watched as they roll downhill. They might get stuck in local minima, with one or more hillsides standing between them and the true, global minimum. A technique to improve this method would be to shake the table whenever a marble comes to a stop. If the marble is in a shallow enough valley, the shaking may cause the marble to roll upslope 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 metaphorical temperature of the system.

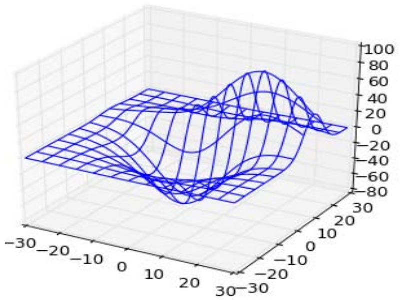

Figure 2. Graph of a Hypothetical Simple Solution Space

Quantum annealing represents an even more powerful heuristic, in which a mechanism is provided that is capable of "tunneling through" the walls which separate local minor minima from the global minimum. No longer is it necessary 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. An N-dimensional surface where N is much larger than three is difficult for most to visualize, but the annealing described above, can be used to find the minimum value on a surface representing a solution.

D-Wave

D-Wave is a small company that has developed an adiabatic quantum annealer which operates at a temperature of less than 20 MilliKelvin. This is barely above absolute zero or -273.15° Celsius, the temperature at which entropy stops, eliminating thermal energy. Published papers are available to detail the technical issues that have been faced and overcome to produce an operating quantum annealer. This paper will not dwell on that process here. A good compendium of detailed technical papers is to be found at http://www.dwavesys.com/en/publications.html .

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, in Marina del Rey. This has since been upgraded to a D-Wave Two, which is a 512 qubit system. It yields only 503 consistently working qubits, due to small manufacturing variations and trapped flux in the superconducting circuits. While this size is capable of generating interesting results, it is not yet big enough to set world records against gargantuan clusters. Figure 3 depicts the 128 qubit chip used in the D-Wave One. It is based on Josephson Junction technology.

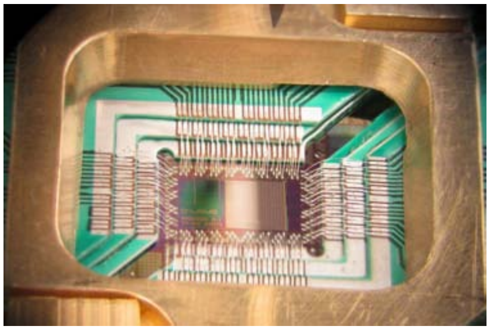

Figure 3. Detail of D-Wave Qubit Processor

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 have now performed enough calculations at the Quantum Computing Center to answer that question in the affirmative (Boixo, 2013).

There is on-going study of the D-Wave device to ascertain its performance relative to classical computers. As mentioned above, the NP-hard problems that the D-Wave machine is designed to solve are in general, too hard to solve exactly in any reasonable amount of time on a normal computer. Instead, heuristics has been used that will regularly get the solution, or a close enough approximation, in a short period of time. Simulated annealing on a digital machine is such a heuristic, as is quantum annealing. In both, there is no guarantee that one will not get trapped in some local minimum. In any case, one effective comparison of the relative performance of quantum annealing to a digital alternative would be to benchmark the performance of AQA against simulated annealing.

Figure 4 depicts such a comparison: quantum annealing on the D-Wave Two versus simulated annealing. The adiabatic quantum annealing timing data were collected by researchers at USC. The simulated annealing data are from Swiss colleagues at Eidgenössische Technische Hochschule (ETH) in Zurich. They have what is believed to be the World's fastest such annealing code. The curves plotted are the time to reach a solution as complexity increases to useful levels. Each line represents different levels of certainty, and lower time is better. In just two short generations, quantum annealing has matched the performance of an eight-core microprocessor at modest complexity levels. It is quite possible that in the next generation, AQA will outperform any classical computing system, of any size.

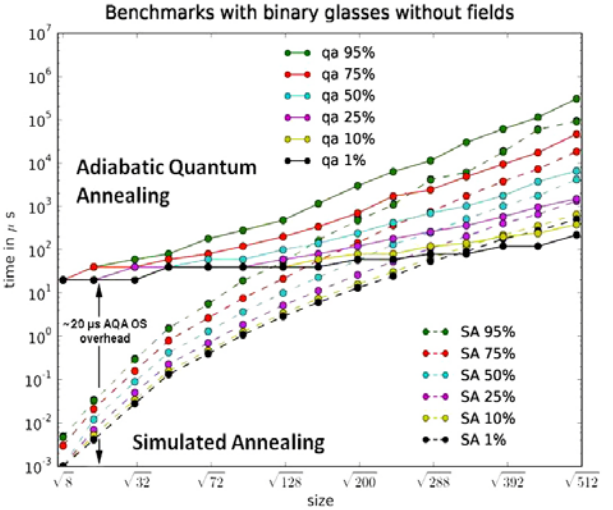

Figure 4. Relative performance of D-Wave Two Quantum Annealing vs. Standard Digital Simulated Annealing

POTENTIAL IMPACT ON THE DOD SIMULATION COMMUNITY

Scientists at USC's QCC are examining practical applications of this quantum technology. D-Wave personnel have suggested many possible candidates, set forth in Table 1. 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). Their investigations have so far included testing operational software for verification and validation (V&V), model checking, sensor assignment, and tracking.

Table 1. 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

As discussed above, AQA may well provide significant potential advantages over classic computational systems. This section will discuss some of the areas within the domain of computer generated forces (CGF) where some of the greatest potential is seen for significant advances in the overall performance of those systems. These discussions are caveated with observation that there still is significant work to be done in the development of the AQA hardware and even more work, some of it fundamental new algorithms and programming paradigms, to bring some of these developments to fruition. Some of the more pessimistic commentators estimate that it may be four more years before practical systems are seen; others see utility today. It may be even longer, if at all, before production AQA-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 this new power the most may adopt the technology very soon, even at the cost of the users having to be the ones to resolve problems and invent new approaches.

It may be useful here to consider a contrived example, but one that has appropriate historical analogs. 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. There is a need to keep in mind the fact that the improper allocation of resources would inevitably result in diminished capabilities. A CSG (carrier strike group) typically has a number of smaller ships. Each of these provides service to the other ships in the armada, *e.g.* fuel, food and communications.

One way to look at this problem (Lucas, A., 2012) is to use a graph theory approach, discussed more generally on line (Wikipedia, 2013). For analysis, each ship could be considered as a node wherein the relationship between the nodes, *e.g.* food and fuel, 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 for the two groups. The number of cross-group dependencies should be minimized and degradations of operational capability in either group avoided. It would be counterproductive to have a food ship travel unnecessary thousands of miles to service the two separated CSGs. 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. This resource partitioning problem is considered to be NP-Complete, the formal mathematical description of which is described in work by Karp (Karp, 1972).

Having conceived the problem using this graph theory approach, it is now possible to optimize the plan using AQA to produce a significantly more efficient solution. 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 configurations that produced the desired 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, but the total of these runs should be a fraction of the time required for standard computing systems. Framing problems this way is not familiar to many programmers. To make AQA 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, the D-Wave developers state that they felt that abstracting away from the overwhelming underlying complexity is critical for making the system broadly accessible. (D-Wave, 2013b) The developer would then be able to use standard programming methods and still access AQA's power.

The annealing process typically identifies a set of local minima, the smallest of which is likely the global minimum. The D-Wave returns a histogram of solutions it finds, and most of them will be at, or near, the global minimum. If the global minimum is the only answer of interest, these other data may be discarded. In the case of decision support for the battlefield commander, the location of the local minima across the solution sets may be of significant interest. The quantum annealer can produce output establishing the location of these minima in the n-dimensional solution space. The analyst would then be able to equate varying outcomes with varying input parameters, *e.g*. strength of forces engaged, plans of attack, terrain, weather, *etc*. After all, given ambiguity in the inputs provided, the global minimum may not in fact be the desired outcome. The authors' experience in combat zones would suggest most commanders would prefer knowing a list of probable outcomes and possible surprises for a proposed course of action instead of a single oracle-like pronouncement.

A Hypothetical Quantum Annealer Implementation

The next focus will be on the mapping of CGF system issues against the class of problems that are well suited to AQA. Problems in which there are large amounts of uncertainty have been identified as likely candidates (Grover, 2013). In essence, the use of AQA allows the user to rapidly explore and refine the possible solution spaces using a combination of brute force and iterative techniques. So, any problem in the CGF space will best be expressed in such a way to make it a search-space-problem, thus making it most amenable to AQA. One such problem is that of the military planning process where the selection of a course of action that would lead to one or more desired end states, similar to the naval forces portioning problem discussed above.

As part of the Deep Green effort, the SAIC Team developed an abstract plan representation that characterized the intentions and possible excursions as a tree data structure. It initially reflected current reality and its end state(s) contained a stopping state that represented successful execution of the plan (Pratt, 2008). The plan was the initial thread and the excursions were represented as branches in what became a potential futures graph (PFG). The traversal through this graph could be thought of as a modified path planning problem, where there were multiple end states, of varying benefit, creating a multiple set of possible stopping conditions. The advantage of this representation was that it allowed the formulation of the CGF as a graph search space problem and manageable by AQA. In the PFG representation, each future could be represented as series of states, or situations, as are notionally shown in Figure 5 below.

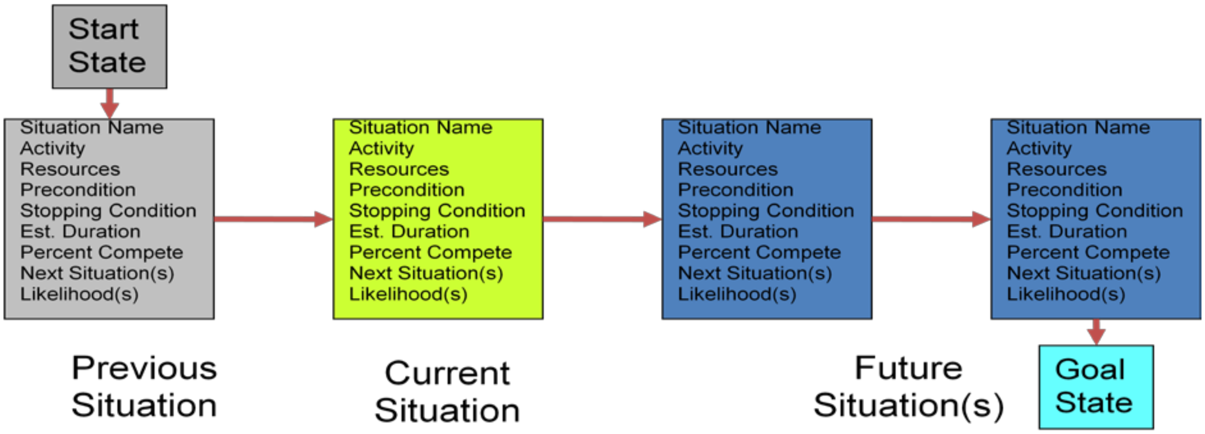

Figure 5. An example of battlefield futures that can be represented as a series of situations and the events that cause transitions between them

The resulting structure in Figure 6 (below) allows the user to pinpoint the critical uncertainties by examining where the branches occur. 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 that challenges the limits of the digital computational paradigm.

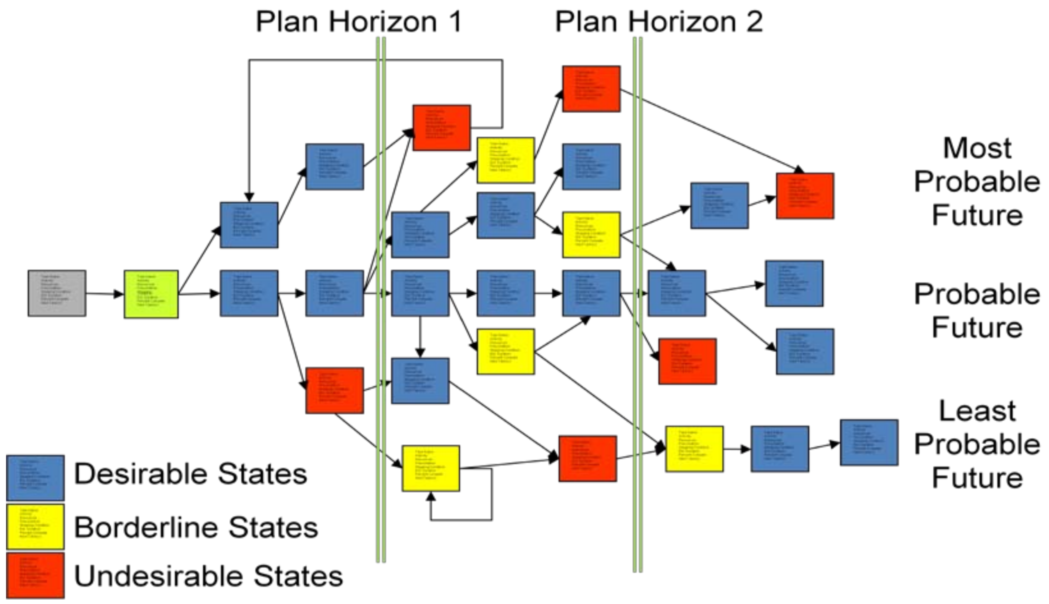

Figure 6. The Potential Future Graph (PFG) representation of the alternative battlefield futures shows the events, branches, and desired states.

Applying AQA to Alternative Futures

The mapping of the CGF problem to a search space problem may require significant computational and intellectual resources to allow the user to receive the maximum benefit from AQA processing models. Foremost among the remaining research issues is the mapping of 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. Using this capability 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 interaction is some variation of the Lanchester equations. The original equations were based upon establishing the ratio of the forces to each other in order to determine the attrition and the result of the interactions. While recently the use of those equations has been largely superseded by entity-based models, the equations have been used for years as the basis for aggregate-level models. Non-determinism is often 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 may be ideally suited to quantum computing. 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 a response to environmental stimuli or opposition action. Given that a wide range of artificial intelligence (AI) techniques are based on either search or constraint satisfaction, it stands to reason that that some types of AI may also benefit from AQA adoption. Thus, the planning process might be cast as questioning if the analytic program will run to completion or cycle forever; *i.e*. an NP Hard Turing Halting Problem. That would possibly make excellent use of quantum annealing processing to provide graph traversals in significantly less time.

Given all this, it seems likely that, when constructed as a search space system with AQA-amenable processing elements, the user can map the proposed CGF system architecture to the AQA 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. This effort is expected to require new processing algorithms and programming paradigms to implement the versions Lanchester equations and AI 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 may 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 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 AQA Technology

Assuming that quantum annealing does in fact prove to be a capable new tool for solving problems arising in simulation, even 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. With some 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 (GPGPU) in the mid-2000s. The promises of performance breakthroughs using the newfound 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 ostensibly offered a potential source for significant speed up of two of the most computational intensive elements associated with CGF systems: geometric inter-visibility (often referred to line of sight (LOS)) and route planning. Under test conditions, this approach worked remarkably well achieving up to a 20x speed up in execution times of the target subfunction. Under more normal operating conditions, 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, *e.g.* JESPP at JFCOM (Lucas, 1993 and Davis, D, 2010a).

One of the authors proposed the use of the CELL processor board to host the environmental runtime component (ERC) and database interfaces (Pratt, 2007). That paper proposed to port the ERC to the CELL processor board, in order to bring the software and computational architectures in line. The anticipated impact was not realized for two primary reasons that are reflective of why many of the promising new technologies are not adopted: fickle market forces and poor lifecycle support. The base market was changing too fast to yield a "standard" GPU implementation and the CELL processor failed to establish a mass market impact outside PlayStation. Perhaps, for any CGF system to adopt quantum computing, AQA itself would have to become a mainstream technology with a body of standards.

This history supports the belief that, unless there are significant markets and lifecycle supports, AQA may be most effective in purpose built systems. As mentioned, a combination of energy costs and time to solution requirements 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 start from behind, but eventually entirely subsume some major segment of the computational spectrum (Christensen, 1997). The growth of the number of qubits will play a dominant role in establishing AQA's utility. That is envisioned by D-Wave in Figure 7.

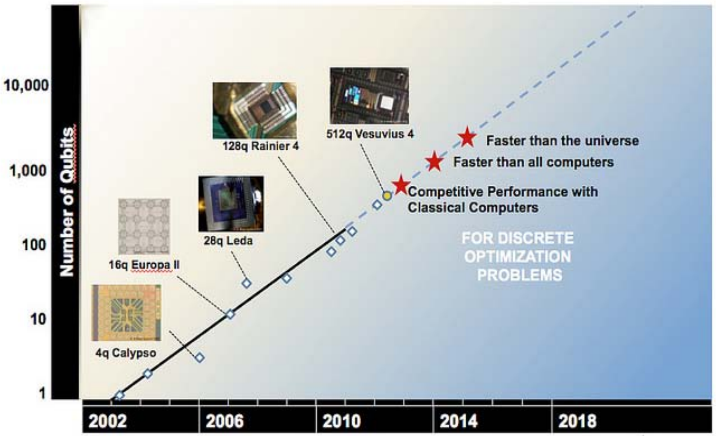

Figure 7. D-Wave's Projection of their Development Path

ANALYSIS

One of the major responsibilities of computational scientists is the requirement 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 part of the DoD simulation community itself, the authors feel 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 insightful comments of Dr. Andy Ceranowicz, who expanded the initial list of simulation grand challenges significantly (Ceranowicz, 2013).

Grand Challenges in Simulation

- 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)
- Production of scenarios for preordained MSELs
- Data-reduction of multiple-source data
- Production of coherent parametric data
- Abstraction of data to evaluate goal achievement
- Promulgation of realistic intelligence reports
- Instantiation of human networks and populations
- Creation of realistic intelligence backgrounds
- Elimination of role players
- Stimulation of players with interaction mechanisms such as natural language speech and messages

Experience on the Quantum Annealer would indicate that if the core computational segment of any of the challenges listed above can be abstracted to an appropriate mathematical representation, the quantum annealer might provide significant breakthroughs. From experiences with entity-level simulations such as the SAF family, as well as aggregate level models, it is felt that many of these challenges have major components that could be enhanced or accelerated by the use of quantum annealing.

Ever mindful of the perils of predictions, the authors are nevertheless willing to give their current impressions of the amenability of these grand challenges to AQA resolution, all the while expecting to be surprised in both directions. Beginning with both extremes, there is reasonable consensus that AQA holds promise for the enhancement of multi-agent path planning, but that there is little hope it will be useful in calculating inter-entity visibility. That issue should be left to the introduction of Feynman's more general purpose quantum computer. With decreasing consensus among the authors, useful roles for AQA are seen in data abstraction, data reduction, and should be considered anywhere machine learning is currently used. This may include the weather, radar and sonar modeling. Then there is the middle ground where more experience with MSELs and scenario production may identify a role for AQA that the authors do not yet see. Less hope is held out for the challenges involving specific human behavior emulation, *e.g*. intelligence background, reports, role-player elimination, and interaction mechanisms.

The user community will, no doubt, more clearly see opportunities to resolve those difficult challenges, as well as seeing hurdles for the challenges for which promise might otherwise be seen. However, that is not the end of the analysis. The total computation productivity effort needs to be taken into account when assessing the value of adopting a new technology (Kepner, 2006). With other new technologies, there have been pleasant surprises when seemingly impossible programming problems have been quickly overcome to permit their use. However, others have been disappointed to an equal degree when some 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 their current roadblocks are, how much relief would improve their overall system performance, what the total costs of such attempts would be, and what other resolutions to their problems may be imminent. Paying close heed to future developments in the quantum computing discipline will also give the user increasingly useful approaches to these cost/benefit analyses, as new data on methods, successes, and failures become available from others' efforts.

SUMMARY

AQA is a powerful new capability, realized in the D-Wave open system adiabatic quantum annealer, which should become even more available to the simulation community in the near future, as some porting and benchmarking has already taken place. Initially, it will likely be used to solve hard, combinatorial optimization problems, which can include the assignment of sensors, the tracking of targets, and the construction of strong classifiers for recognizing specific features in large data sets. It not expected to replace the ensembles of computers that serve the simulation community, but rather augment them. In an increasingly distributed, heterogeneous computing environment, it will be yet another tool capable of providing solutions in near real time to otherwise intractable problems.

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 USC are proud to be their partners in this research. 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 USC, SAIC, 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 Blog posts, April

Ceranowicz, A., (2013), “Simulation Grand Challenges,” private correspondence from Dr. Ceranowicz to authors

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., Lucas, R. F., Wagenbreth, G., Roberts, D. W. & Brewton, C., (2010a), “The Future Uses for the GPGPU-Enhanced Cluster at JFCOM”, in the Proceedings of the HPCMP Users Group Conference

Davis, D.M, (2010b), 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., Williams, R.; & Messina, P., (1994), “Parallel Computing Works!”, Morgan Kaufman, New York, NY

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

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

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

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", *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