

# System of Systems Complexity Addressed by Practical Adiabatic Quantum Computing

Robert F. Lucas, Dan M. Davis and Daniel P. Burns

Information Sciences Institute and Institute for Creative Technologies  
University of Southern California

*Systems of systems require more computing power than is currently available. Simulation of environments, weapons systems and individual platforms are required. Thus the Test and Evaluation community has a great need for improved computing capabilities.. Despite approaching 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 authors report early results of more than three 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). 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. Using these data as foundations, the potential in the realm of DoD Test and Evaluation is discussed. They discuss a range of the test and evaluation problems that should be amenable to this new technology and forthrightly list a few areas that they believe will not benefit from Quantum Computing.*

**Key Words:** System of Systems; Quantum Computing, Simulated Annealing; and Optimization.

According to Gordon Moore himself, the end of Moore's Law is nigh. It is increasingly daunting to continue on the current transistor-based path for increasing the traditional digital computing capability that can be applied to test and evaluation 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 defense problems. In this environment, there is a growing interest in alternatives to commercial, off-the-shelf (COTS) technology, which would have seemed inconceivable for most of the last two decades. In many ways, this may be the reemergence of the purpose built systems of earlier decades. New installations include specialized systems such as the Anton at D. E. Shaw Research<sup>1</sup> It performs certain biomolecular simulations nearly two orders-of-magnitude faster than the largest general purpose computers. Others are looking be-

yond CMOS to exploit other physical phenomenon, e.g. quantum computing.

Quantum computing has been considered a promising extension of computational capability since the semi-



Figure 1. USC-LMC  
QCC D-Wave 2

nal paper from the Nobel Laureate Richard Feynman in 1982<sup>2</sup>, 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 the University of Southern California. 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<sup>3</sup>

In recent years, 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.<sup>4</sup> They point out that the capability of quantum computers arises from the different way 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.

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 aware of its rivalry with sequential computing and with vector computing. They heard the detractors of parallel computing argue the limits of parallelism<sup>5</sup> and the proponents<sup>6</sup> 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 both data and computation in effective ways<sup>7</sup>. Such technology has proven very useful to the simulation community<sup>8, 9</sup> which has many issues identical to the test and evaluation environment.

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, advocates were heard asserting incredible speed-ups and detractors were questioning the utility of the technology. Taking a more pragmatic view, the authors carefully assessed the capabilities of such devices<sup>10</sup>, measured the energy savings<sup>11</sup> and instructed the DoD users<sup>12</sup>. In one conference, after the


presentation of a paper by one of the authors<sup>13</sup>, 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<sup>14</sup>. The intent of this paper is to continue in that tradition.

## Adiabatic Quantum Annealing

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 classical, 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 space 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, that is 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



**Figure 2. Hypothetical Simple Solution Space**

be dropped on the complex space, and watched as they rolled downhill. They might get stuck in local minimum, 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 valley, 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 known as simulated annealing. Shaking the table is equivalent to increasing the metaphorical temperature.

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<sup>15</sup>. 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 of a surface representing a solution.

## 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 or  $-273.15^\circ$  Celsius, the temperature at which entropy stops, eliminating thermal energy. 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 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 qubit D-Wave One<sup>16</sup>, 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. Small manufacturing variations and trapped flux in the superconducting circuits resulted in 503 working qubits. 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.

## 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<sup>17</sup>. 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<sup>18</sup>.

There is on-going study of the physics of the D-Wave




Figure 3. Detail of D-Wave Qubit Processor

device, among other things, trying to ascertain its performance relative to classical computers. 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 are used that will hopefully get the solution, or a close enough approximation, in a short period of time. Simulated annealing is such a heuristic in that there is no guarantee that one will not get trapped in a local minimum. Quantum annealing is too. So, a straight up comparison of the relative performance of quantum annealing to a practical alternative is to benchmark it against simulated annealing.

Figure 4 depicts such a comparison: quantum annealing on the D-Wave Two versus simulated annealing, using what we believe to be the World’s fastest such code from our colleagues at ETH. The curves plotted are the time to reach a solution as complexity increases to useful levels, for different levels of certainty, such that lower is better. In only two 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, AQC will outperform any classical computing system, of any size.

## POTENTIAL IMPACT ON THE DOD TEST AND EVALUATION COMMUNITY

Scientists at USC's QCC are examining practical applications of this technology. The D-Wave personnel have suggested many possible candidates, set forth in Table 1. They argue that even the concept of scientific

<b>Data Mgt.</b>	<b>Behaviors</b>	<b>Analysis</b>
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

discovery itself is an optimization problem, in which the researcher is trying to ascertain the optimal configuration<sup>1</sup> of parameters that would contribute to a scientific mathematical expression which comports with real world observations<sup>19</sup>. Their investigations have so far included software verification and validation (V&V), model checking, sensor assignment, and tracking.

As discussed above, quantum computers (QC) may well provide significant potential advantages over classical computational systems. In this section we will discuss 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 these discussions to fruition. Some of the more pessimistic estimate that it may be four more years before practical systems are seen. It may be even longer, if at all, before production AQC-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.

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.

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, keeping 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, 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<sup>20</sup>, discussed more generally on line<sup>21</sup>. For analysis, each ship could be considered as a node wherein 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.* 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<sup>22</sup>.

Having conceived the problem using this graph theory approach, it is now possible to optimize the plan using AQC 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 configuration that produces 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, by 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 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 test and evaluation 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.<sup>23</sup>

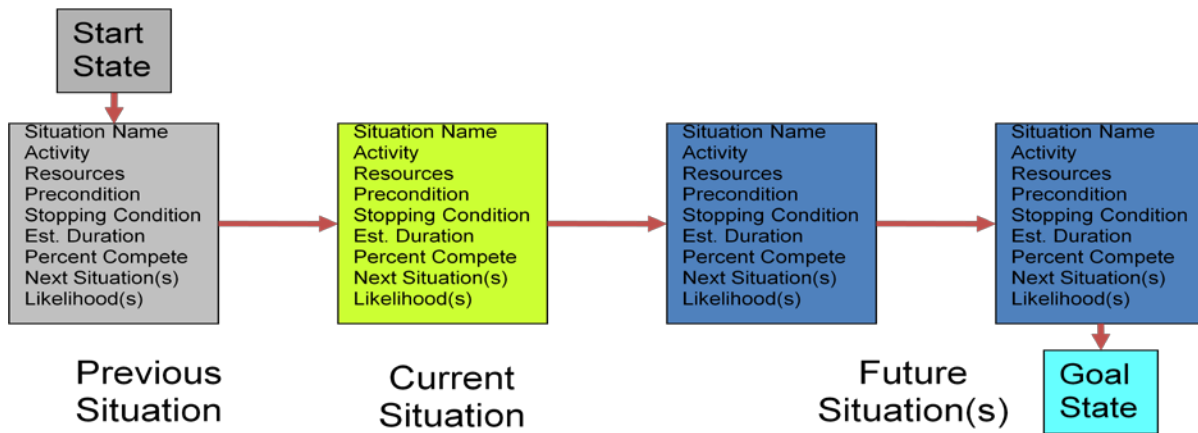
## A Hypothetical Quantum Annealer Implementation

Given that the future universal acceptance of quantum computing is beyond the scope of this paper, the next focus will be on the mapping of CGF system issues against the class of problems that are best suited to AQC. Problems in which there are large amounts of uncertainty have been identified as likely candidates<sup>24</sup>. In essence, the use of AQC allows the user to rapidly explore and refine the possible solution spaces using a combination of brute force and iterative techniques. Thus, any problem in the CGF space will best be expressed in such a way to make it a search-space-problem in order to make it amenable to AQC. 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.

## The structure of the Potential Future Graph

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<sup>25</sup>. 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, therefore, could make the CGF domain manageable by AQC. In the PFG representation, each future could be represented as series of states, or situations, as are notionally shown in Figure 5 below.

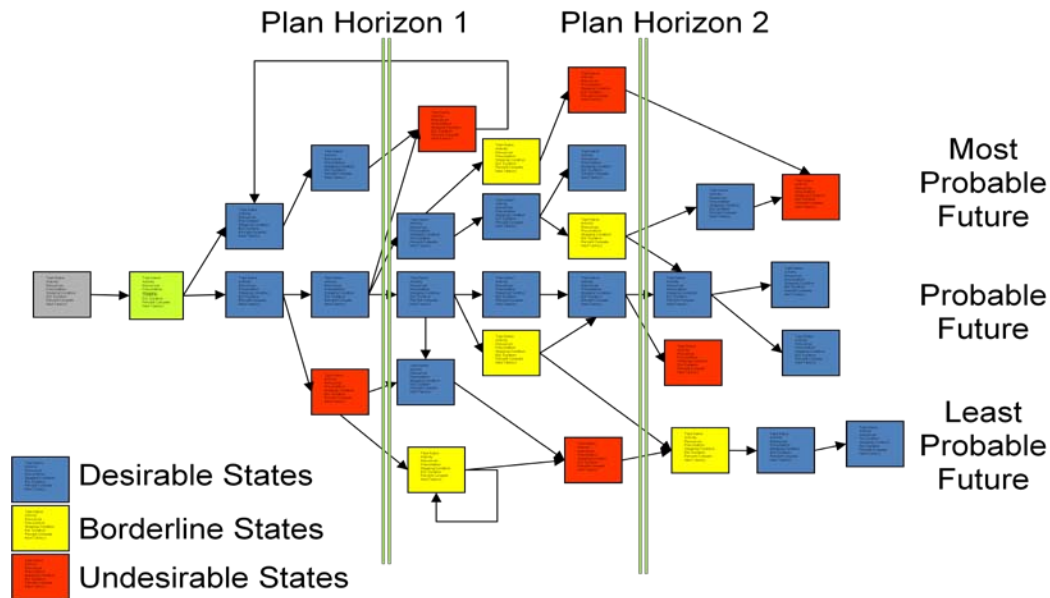


**Figure 5. An Alternative Future can be represented as a series of situations and the events that cause transitions between them**

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 futures shows the events, branches, and desired states.**

### Applying AQC to Alternative Futures

The mapping of the CGF problem to a search space problem would require significant computational and intellectual resources to allow the user to do a direct mapping to the AQC processing models. Foremost among the remaining research issue 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. 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 are based upon establishing 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 by 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 is 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 AQC adoption. Thus, the planning process might be cast as questioning if the 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 search space system with AQC amenable processing elements, the user can map the proposed CGF system architecture to the AQC 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<sup>26</sup>. 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 AQC 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 test and evaluation 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 been some enhancements, notably the migration from mainframes to workstations 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<sup>27</sup> and as discussed in<sup>28</sup>, the GPGPU provided a potential source for significant speed up of two of the most computationally 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 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<sup>29</sup>.

One of the authors proposed the use of the CELL processor board to host the environmental runtime component (ERC) and database interfaces<sup>30</sup>. That paper proposed to port the ERC to the CELL processor board, in order to bring the software and computational architectures in line. That option was not pursued for two of the primary reasons that many of the promising new technologies fail to be adopted: market forces and lifecycle support. The underlying market was changing too fast to yield a "standard" GPU implementation and the CELL processor failed to establish a mass market (other than the PlayStation) impact. Some argue that for any CGF system to adopt quantum computing, AQC itself would best become a mainstream technology with a body of standards.

With this history in mind, some believe that unless there are significant market and lifecycle reasons, AQC may be most effective as 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 start from behind, but eventually entirely subsume some major segment of the computational spectrum<sup>31</sup>. The growth of the number of qubits, as it is envisioned by D-Wave in Figure 7<sup>32</sup>, will play a dominant role in establishing AQC's utility.

## 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 test and evaluation 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 insightful comments of Dr. J. Michael Barton, who significantly expanded the initial list of test and evaluation grand challenges<sup>33</sup>.

### Grand Challenges in Test and Evaluation

- Creation of range instrumentation plans
- Exploration of very large datasets
- Augmentation of comprehensive V-Model
- Implementation of better sensor models
- Facilitation of simulated inter-visibility
- Production of scenarios for tests
- Reduction of multiple source data
- Production of coherent parametric test data
- Abstraction of data to evaluate spec achievement
- Expansion of forms of data discovery & fusion
- Acceleration of V-Model evolutions
- Elimination of role players in test simulations

Experience on the Quantum Annealer would indicate that if the core computational segment of each of the challenges listed can be abstracted to an appropriate mathematical representation, the quantum annealer might provide significant breakthroughs. From their own experiences with both computer test and evaluation, as well as a passing acquaintance with physical, range testing, the authors feel that many of these challenges have major components the solutions for which 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 the listed test and evaluation 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 for the enhancement of multi-agent path planning, but that there is less hope it will be useful in enhancing data conversion calculations, that issue being left to the introduction of Feynman's more general purpose quantum computer. With decreasing consensus among the authors, useful roles for AQC are seen in data abstraction, data reduction, and should be considered anywhere else that machine learning is currently used, which may include the V-Modeling challenges. Then there is the middle ground where more experience with data reduction and simulation inter-visibility may identify a future role for AQC that the authors do not yet see. Less hope is held out for the challenges involving specific transaction computation tasks, e.g. archiving data, conversions, and logging.

The users 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<sup>34</sup>. With other new technologies, the authors have been pleasantly surprised when seemingly intractable programming problems have been quickly overcome to permit their use. However, they 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 road blocks are, how much their overall system performance would improve from the relief of those road blocks, 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 these cost/benefit analyses, as new data on methods, successes, and failures become available from others' efforts.

## SUMMARY

AQC is a powerful new capability, realized in the D-Wave open system adiabatic quantum annealer, which should be available to the test and evaluation 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 not expected to replace the ensembles of computers that serve the test and evaluation 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 problems that might otherwise be seen as intractable.

*Robert F. Lucas is a Deputy Director of the Information Sciences Institute at the University of Southern California and leads the Computational Sciences Division. He is a Research Associate Professor in the USC Department of Computer Science. At ISI he manages research in computer architectures, VLSI, compilers, and other software tools. Prior to joining ISI, he did tours as the Director of High Performance Computing Research for NERSC, the Deputy Director at DARPA's, and a researcher at the Institute for Defense Analyses. Dr. Lucas earned BS, MS, and PhD degrees in Electrical Engineering from Stanford University.*

*Dan Davis is a consultant for the Information Sciences Institute, University of Southern California, focusing on large-scale distributed simulations. There, he led the JESPP project for a decade. As Assistant Director of CACR at Caltech, he managed Synthetic Forces Express, introducing HPC to DoD simulations. He was the Chairman of the Coalition of Academic*



*Supercomputing Centers and has taught at the collegiate level. Dan started writing FORTRAN programs in 1971 on Seymour Cray's CDC 6500's. He served in Vietnam as a USMC Cryptologist and retired as a Commander, U.S.N.R. He received B.A. and J.D. degrees from the University of Colorado.*

*Daniel P. Burns is a lifelong Systems Engineer, first with the Active Duty Navy, then SAIC, and small business. He served as Naval Chair and Professor of Practice in Systems Engineering at the Naval Postgraduate School (NPS). Captain Burns served as the as the Military Associate Dean and as acting Dean of the Graduate School of Engineering and Applied Sciences at NPS. His research interests center on analyses of both human and resource utilization in defense efforts. Captain Burns received a BS degree from the U.S. Naval Academy and an MS from the Naval Postgraduate School. He is currently finishing his dissertation for a PhD from Southern Methodist University.*

---

<sup>1</sup> 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

<sup>2</sup> Feynman, R. P. (1981), "Simulating Physics with Computers," *International Journal of Theoretical Physics*, Vol 21, Nos. 6/7

---

<sup>3</sup> ISI, (2015). *D-Wave computer at Lockheed-Martin-USC Quantum Computing Center*, Marina del Rey, California, photo from Information Sciences Institute, University of Southern California.

<sup>4</sup> Gershenfeld, N. & Chuang, I, (1998), "Quantum Computing with Molecules," *Scientific American*, (278), pp. 66-71

<sup>5</sup> 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

<sup>6</sup> Fox, G., Williams, R.; & Messina, P., (1994), "Parallel Computing Works!," Morgan Kaufman, New York, NY

<sup>7</sup> Gottschalk, T.; Amburn, P. & Davis, D., (2005), "Advanced Message Routing for Scalable Distributed Simulations," *The Journal of Defense Modeling and Simulation*, San Diego, California

<sup>8</sup> 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

<sup>9</sup> Lucas, R. F., Davis, D. M. & Wagenbreth, G., (2007), "Implementing a GPU-Enhanced Cluster for Large-Scale Simulations," in the *Proceedings of the Interservice/Industry Simulation, Training and Education Conference*, Orlando, Florida, 2007

<sup>10</sup> 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

<sup>11</sup> Davis, D. M.; Lucas, R. F.; Gottschalk, T. D.; Wagenbreth, G.; & Agaloff, 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

<sup>12</sup> 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

<sup>13</sup> 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

<sup>14</sup> 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.

<sup>15</sup> Boixo, S., Smelyanskiy, V. N., Shabani, A., Isakov, S. V., Dykman, M., Denchev, V. S., ... & Neven, H. (2015). Computational role of multiqubit tunneling in a quantum annealer. arXiv preprint arXiv:1502.05754

<sup>16</sup> 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)

<sup>17</sup> 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

<sup>18</sup> Albash, T., Vinci, W., Mishra, A., Warburton, P. A., & Lidar, D. A. (2015). "Consistency tests of classical and quantum models for a quantum annealer". *Physical Review A*, 91(4), 042314.

<sup>19</sup> D-Wave, (2015a), "D-Wave Software Architecture," User tutorial from D-Wave staff, retrieved from <http://www.dwavesys.com/software.html> on 30 Jun 2015

<sup>20</sup> 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>

<sup>21</sup> Wikipedia, (2015), "Graph Theory," retrieved on 07 May 2015 from Wikipedia, the free encyclopedia: [http://en.wikipedia.org/wiki/Graph\\_theory](http://en.wikipedia.org/wiki/Graph_theory)

<sup>22</sup> Karp, R.M., (1972), "Reducibility among combinatorial problems," *Complexity of Computer Computations*, Miller and Thatcher, eds.; Plenum Press, New York, NY, pp.85-104

<sup>23</sup> D-Wave, (2015b), "Programming with D-Wave: Map Coloring Problem," User tutorial from D-Wave staff, retrieved from <http://www.dwavesys.com/sites/default/files/Map%20Coloring%20WP2.pdf> on 30 Jun 2015.

<sup>24</sup> Grover, L. K., (2015), "Quantum computers can search rapidly by using almost any transformation," [arxiv.org/pdf/quant-ph/9712011](http://arxiv.org/pdf/quant-ph/9712011), Retrieved 14 April, 2015 <http://arxiv.org/pdf/quant-ph/9712011.pdf>

<sup>25</sup> 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

<sup>26</sup> Kenyon, H.S., (2007), "Deep Green Helps Warriors Plan Ahead," *Signals*, downloaded 02 May 2015 from <http://www.afcea.org/content/?q=node/1418>

<sup>27</sup> 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: University of North Carolina

<sup>28</sup> 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

<sup>29</sup> 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

<sup>30</sup> 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).

<sup>31</sup> Christensen, Clayton (1997). "The Innovator's Dilemma: The Revolutionary Book That Will Change the Way You Do Business," Harvard Business Review Press, Cambridge, Massachusetts

<sup>32</sup> D-Wave, (2015), "D-Wave Development Path and Performance Projection," *Graph from Introduction to the D-Wave Quantum Hardware*. Retrieved from <http://www.dwavesys.com/tutorials/background-reading-series/introduction-d-wave-quantum-hardware.html> on 30 Jun 2015.

<sup>33</sup> Barton, J. M., (2015), "Test and Evaluation Grand Challenges," private correspondence between Dr. Barton and authors

<sup>34</sup> 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

