

Advanced Computer Technology for Novel Information Processing Paradigms

Clare Thiem,^{*} Steven Drager,[†] Christopher Flynn,[†] Thomas Renz,[†] and Daniel Burns,[†]
Air Force Research Laboratory/IFTC, Rome, New York 13441-4514

Advanced computer technologies for novel information processing paradigms to be applied to command and control systems of the next decades are presented in this paper. The Advanced Computing Technology Branch (IFTC) within the Air Force Research Laboratory's (AFRL) Information Directorate is exploring and developing many technological avenues to incorporate novel information processing capabilities that address future command and control (C2) systems requirements. The research and development process includes but is not limited to participation in Defense Advanced Research Projects Agency (DARPA) research programs, and utilization of United States Air Force funded research and development activity especially in-house research programs. The authors provide a snapshot of computer technologies that the military aerospace community will see in future information systems. These systems include near term technologies composed of hybrid hardware/software computing architectures incorporating Processor In Memory (PIM) into conventional computers, architectures that are capable of dynamically morphing their hardware and software when required, and high productivity computing systems. Far term systems will include biomolecular and quantum computing architectures that incorporate data storage and processing mechanisms with density, power, and speed performances far beyond state-of-the-art silicon technologies. Pursuit of the technologies presented in this paper permits researchers to explore computer architectures with greater capacity and sophistication for addressing dynamic mission objectives under constraints imposed by Command, Control, Communication and Computer (C4), Intelligence, Surveillance and Reconnaissance (ISR), and strike systems in order to establish, maintain, and exploit information superiority.

I. Introduction

GROWTH of information technology in the 21st century will be driven by advanced computing technology brought about through the development and implementation of information processing paradigms that are novel by today's standards. These advances in information technology will provide tremendous benefits for the war fighter who not only faces the enemy on the field, but struggles under an overwhelming amount of data that must be processed with split second accuracy for life and death decisions. Lessons learned from recent conflicts are providing focus for information technology research and development activity. Numerous research organizations are exploring short and long term solutions in a variety of diverse technologies to bring about new methods for handling, processing, and storing information. Within AFRL's Information Directorate this activity occurs under the auspices of the Advanced Computing Architecture Focus Area. The activity utilizes a spiral research and development process that moves technology from fundamental science to deployable technology for the war fighter.

Presented as Paper 2004-1263 at the 42nd AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, 5-8 January 2004; received 27 February 2004; revision received 23 June 2004, accepted for publication 29 June 2004. This material is a work of the U.S. Government and is not subject to copyright protection in the United States. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 1542-9423/04 \$10.00 in correspondence with the CCC.

^{*}Electronics Engineer, 26 Electronic Parkway. AIAA Senior Member

[†]Electronics Engineer, 26 Electronic Parkway.

Technologies being pursued include near-mid term silicon advances as well as far term biomolecular computing,¹⁻⁷ quantum computing,^{8,9} and hardware assisted evolutionary computing.^{10,11} The authors are participating in DARPA sponsored research including the Data Intensive Systems (DIS),^{*} Polymorphic Computing Architectures (PCA),[†] High Productivity Computing Systems (HPCS),[‡] Biofluidics on a Chip (BioFlips),[§] Simulation of Biological Systems (SIMBIOSYS),^{**} Quantum Information Science and Technology (QuIST),^{††} and Biocomputation^{**} programs. This participation has provided the authors with inroads to the various research communities resulting in a better understanding of the technologies, including potential capabilities as well as pitfalls. As a result, specific applications of the technology for near term solutions addressing Air Force and DOD information technology voids have been identified. Furthermore, the authors have been able to define technology developments that will affect the far term transformation of Air Force and DOD information systems. The information presented in this paper is basically organized according to the general timeframe from which the authors believe the technologies from a particular research area will be integrated into deployable systems starting with the near term. The exception to this is the placement of a short discussion on genetic algorithms toward the end of the paper. Short background information on each research area is presented along with potential benefits and applications of the technology. Therefore, the body of this paper will begin with the first research area Data Intensive Systems, followed by Polymorphic Computing Architectures, High Productivity Computing Systems, Quantum Computing, BioMolecular Computing, and Genetic Algorithms. The paper ends with some concluding remarks.

II. Data Intensive Systems (DIS)

Many defense applications have large data sets that are accessed non-contiguously. This trait stresses the data access configurations in mainstream processing architectures resulting in data starved processors. In addition, the increasing gap between processor and memory speeds, as illustrated in Fig. 1 (Ref. 12), is a well-known problem in computer architectures. DRAM Speed for this discussion is DRAM “Access” Speed. Peak processor performance has been increasing at a rate of 50-60% per year while memory access times improve at merely 5-7%. The processor to memory bottleneck severely constrains the performance of these applications. The dotted lines in Fig. 1 are extrapolated out from solid lines.

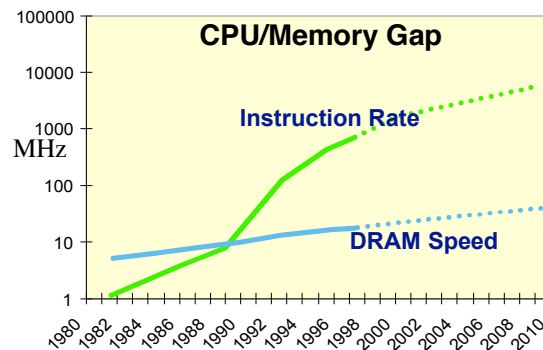


Fig. 1 Plot illustrating gap between processor and memory speed.

Computing-system designers attempt to compensate for this problem by incorporating caches or latency hiding measures such as multithreading and pre-fetching into their designs. Unfortunately, these solutions can actually increase the memory bandwidth requirements¹³ and therefore do not address the needs of data-intensive DoD applications, which ultimately operate at rates far below the peak processor capacity.

*Data available online at <http://www.darpa.mil/ipto/programs/dis/index.htm> (cited July 2004).

†Data available online at <http://www.darpa.mil/ipto/programs/pca/index.htm> (cited July 2004).

‡Data available online <http://www.darpa.mil/ipto/programs/hpcs/index.htm> (cited July 2004).

§Data available online <http://www.darpa.mil/dso/thrust/biosci/bioflips.htm> (cited July 2004).

**Data available online <http://www.darpa.mil/dso/thrust/biosci/simbios.htm> (cited July 2004).

††Data available online <http://www.darpa.mil/ipto/programs/quist/index.htm> (cited July 2004).

**Data available online at <http://www.darpa.mil/ipto/programs/biocomp/index.htm> (cited July 2004).

Recent VLSI technology trends offer one promising solution to bridging this processor-memory gap: embedded-DRAM technology integrates logic with high-density memory in a processing-in-memory chip. In-situ processing or Processor-in-Memory (PIM) architectures,¹⁴ avoid the von Neumann bottleneck in conventional machines by integrating high-density DRAM and CMOS logic on the same chip. Parallel systems based on this new technology are expected to provide higher scalability, adaptability, robustness, fault tolerance and lower power consumption than current massively parallel processors or commodity clusters. Because PIM internal processors can be directly connected to the memory banks, the memory bandwidth is dramatically increased. Latency to on-chip logic is also reduced, down to as little as one half that of a conventional memory system, for the reason that internal memory accesses avoid the delays associated with communicating off chip.

Adaptive cache management approaches attempt to achieve performance gains by making key portions of the memory hierarchy controllable by software at compile and run-time. The cache system and supporting elements are designed to enable multiple data movement structures and policies, such as multilevel caches and intelligent pre-fetching schemes. Other techniques such as tiling, where the work set size is reduced so that the problem will fit into the cache, can also be applied.

Data Intensive Systems will enable the full use of increasing processing element capabilities and reduce the under-utilization of system resources due to restricted data flow and high latency, as shown in Fig. 2, that many data rich applications encounter with current memory implementations. This will be accomplished through the development of new data-access architectures, data flow and placement concepts, and the associated chip-level technologies necessary to respond to data-starved defense applications.




Fig. 2 Typical data latency.

Data-Intensive Computing could enable a substantial increase in performance of data-starved applications. Object-oriented databases are a prime example. Searching and sorting without the need to move pointers and indices all the way up the memory hierarchy could achieve orders-of-magnitude in improvement. Object-oriented methods could execute in-place, reducing memory traffic and dramatically increasing performance. In addition, allowing applications to choreograph data movement will maximize the utilization of precious cache and memory bandwidth.

DIS architectures would benefit any application required to access or manipulate data in any manner not consistent with regular, ordered data access patterns and local working set models assumed by today's conventional architectures. Military applications such as model-based Automatic Target Recognition (ATR), Synthetic Aperture Radar (SAR) codes, large scale dynamic databases/battlefield integration, dynamic sensor-based processing, high-speed cryptanalysis, high speed distributed interactive and data intensive simulations, data-oriented problems characterized by pointer-based and other highly irregular data structures, security-sensitive applications (where protection and validation strategies are central), and real-time visualization should be leveraging this technology in six years.

III. Polymorphic Computing Architectures (PCA)

Current DoD embedded information computing systems can be characterized as fixed in nature, relying on hardware driven heterogeneous point-solutions that represent fixed architectures and software optimizations. Today's embedded computing systems were developed for fixed mission scenarios and cannot provide the robust embedded processing capability necessary to fully support retargetable and multi-mission systems. This results in a

lack of versatility to address dynamic mission requirements and poorly matched processing performance. A unique processing design for each specific mission's sensor configuration can not be afforded due to the cost that such an approach requires in the multiplicity of platforms and the inability to accurately define and predict mission variations prior to deployment.

Polymorphic Computing Architectures will institute a paradigm shift from static open loop to reactive closed loop mission algorithms, application software, and hardware implementation. The processing capability is mission and technology invariant yet highly optimizable for each specific in-mission and multi-mission and/or technology instantiation providing for tactical and strategic tempo opportunities as well as technical upgradeability.

Measurement metrics are being developed in support of processing system design and optimizations for size, weight, energy, performance, and time (SWEPT). Figure 3 illustrates the concept that PCAs will perform adequately across the spectrum of applications. Polymorphic architectures will provide optimal SWEPT efficiencies across the different types of embedded processing possibilities. FPGAs, for example, might exhibit high SWEPT efficiencies for structured bit operations of some search routines, while General Purpose (GP) processors might exhibit high SWEPT efficiencies for symbolic operations found in AI applications, but neither is very efficient for some types of signal processing applications. FPGAs are a little more flexible in terms of their SWEPT efficiencies; you can tailor them at the start for speed, power dissipation, etc. They cannot, however, morph on the fly as the scenario changes. For example, if it is determined that the battery is getting low and power conservation is necessary for survival, they are not capable of on-demand reconfiguration. Polymorphic architectures are expected to be able to reconfigure themselves and exhibit high SWEPT efficiencies across the spectrum. They may not be able to achieve the SWEPT efficiency of an ASIC designed for a specific task, but an ASIC does not provide any flexibility, are costly to design and fabricate and probably won't be able to take advantage of future advances in technology. ASICs are developed to have limited SWEPT range, they may be fast, but consume lots of power. The green ASIC line in Fig. 3 can be moved left and right on the x-axis depending upon the specific ASIC. Due to their programmability, GP's are more flexible than ASICs, but less than FPGAs and are more optimal for a different class of problems than FPGAs. The idea is that a PCA can perform adequately for all elements of SWEPT.

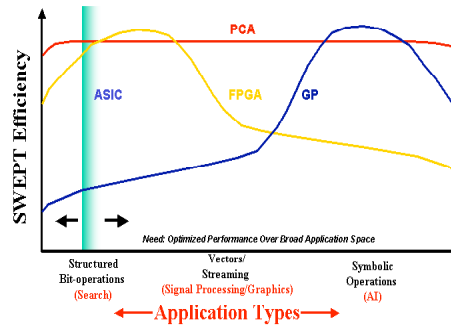


Fig. 3 PCA's high SWEPT efficiencies vs ASIC's, FPGA's, and general purpose processors.

The mission payoff is the ability to react to collaborative information centric strategies using a common highly optimized (based upon SWEPT) processing architecture regardless of the mission dynamics or sensor suites. The result will be the capability to provide optimized mission processing implementations (pre-mission, in-mission, and post-mission) without the need of custom development effort including associated time and cost.

Polymorphic architectures are intended to satisfy the processing requirements for many embedded applications. As the mission lifetimes of DoD weapon systems increases, PCAs will provide the flexibility for rapid upgradeability as well as being used in existing DoD platforms as replacements for aging avionics. Their flexibility makes them ideal candidates for multi-mission multi-sensor weapon systems. The polymorphic processing capability will enable SAR, Moving Target Indicator (MTI), and Signal Intelligence (SIGINT) to be processed with the same hardware, reducing the need for multiple hardware types for the various processing types. Future combat systems will benefit from PCAs in the many ways already described.

Polymorphous Computing Architectures provide a revolutionary approach to implementing embedded computing systems supporting reactive multi-mission, multi-sensor, and in-flight retargetable missions. Payload adaptation, optimization, and verification can be reduced from years to days to minutes. Polymorphic architectures break the current failure prone development approach of hardware first and software last by moving beyond

conventional silicon to flexible polymorphous computing systems. Demonstrations of this technology area should be seen within two years and systems leveraging the technology being fielded about four years later.

IV. High Productivity Computing Systems (HPCS)

High performance computing is at a critical juncture. Over the past three decades, this important technology area has provided crucial superior computational capability for many important national security applications. Government research, including substantial DoD investments, has enabled major advances in computing, contributing to the U.S. dominance of the world computer market. Unfortunately, current trends in commercial high performance computing, future complementary metal oxide semiconductor (CMOS) technology challenges, and emerging threats are creating technology gaps that threaten continued U.S. superiority in important national security applications.

Recent DoD studies¹⁵⁻¹⁷ indicate that there is a national security requirement for high productivity computing systems. Without government research and development (R&D) and participation, high-end computing will be available only through commodity manufacturers primarily focused on mass-market consumer and business needs. This solution would be ineffective for important DoD and national security applications.

For the DoD, revolutionary change in High Productivity Computing Systems means making supercomputing resources easier to use and easier to program. It means accessing large data repositories located around the world and merging scientific computing with pervasive computing on the battlefield. It requires ubiquitous multilevel security, autonomous systems management and applying real-world requirements to the challenges of HPCS to make everything work.

Key to the research and development of new HPCS systems is the ability to measure and understand critical performance characteristics for the entire system – both hardware and software. The ability to characterize and predict performance will provide a clearer picture of future hardware and software requirements and serves as a critical basis for evaluation and development of high-end systems. Therefore, it is vital that a broad spectrum of potential HPCS applications be analyzed to extract the key HPCS system design characteristics, parameters, constraints and programming environments. Applications currently being studied include: operational weather and ocean forecasting; planning exercises related to analysis of the dispersion of airborne contaminants; cryptanalysis; military platform analysis; survivability/stealth design; intelligence/surveillance/reconnaissance systems; virtual manufacturing/failure analysis of large aircraft, ships, and structures; emerging biotechnologies and C3 applications, see Fig. 4.




Fig. 4 Representative DoD challenge applications for HPCS.

Focused research and development will create new generations of high end programming environments, software tools, architectures, and hardware components to realize a new vision of high end, highly productive computing systems. Currently there are three HPCS phase 2 development teams and a productivity team. The HPCS development teams are developing proprietary new hardware including CPUs, memory structures, interconnect, and cooling technologies. In addition they are developing programming languages, program development environments, etc. that are in most cases planned to be open source. Their efforts will address the issues of low efficiency, scalability, software tools and environments, and growing physical constraints, such as size, heat and power. New system software and programming environments are being developed to increase programmer productivity and facilitate performance modeling, measurement and prediction. The productivity team is collaborating with the development teams in developing benchmarks, graduate courses in high performance programming, and high

performance SW developments lessons learned. All this will produce viable high productivity computing system solutions to fill the DoD high-end computing gap between today's late 80's based technology High Performance Computers and the promise of quantum computing. Simulations of this technology should occur in four years with prototypes following two years later. HPCS systems should be fully leveraged in 7-10 years.

V. Quantum Computing

Quantum computers are being researched as a possible heir to today's silicon computers 20 plus years out. Quantum computers offer theoretical computational power not available in classical computers, as the principle of superposition allows an extraordinarily large number of computations to be performed simultaneously, in true parallel fashion. Classical parallel processor computers, only truly do one thing at a time: there are just two or more of them doing it. The computational power of quantum computing brings the promise of the ability to efficiently solve some of the most difficult problems in computational science. Included in this problem set are integer factorization, discrete logarithms and quantum modeling.

As opposed to the transistors used in a classical computer, a quantum computer uses the characteristics of particles at the sub-atomic level to perform computations. Characteristics such as the polarization states of a photon, energy levels of an electron, or the spin directions of an electron are used to perform computations.

Binary information is described in the form of two-state systems for a quantum computer, with a single bit of information being known as a qubit. As mentioned earlier, the theory of superposition allows a large number of computations to be performed simultaneously. Superposition is the ability for the particle to exist in more than one state at a time. When a particle is in a superposition of states, it behaves as if it were in both states simultaneously. Each qubit is both 0 and 1 at the same time. Thus, the number of computations that a quantum computer could undertake is 2^n , where n is the number of qubits used. A classical 2-bit register stores only one of four binary configurations (00, 01, 10, or 11) at any given time. However, a 2-qubit register in a quantum computer can store all four numbers simultaneously. As more qubits are added the capacity of a quantum computer increases exponentially.

Another useful theory is that of quantum entanglement. Particles that have interacted at some point retain a type of connection and can be entangled with each other in pairs. Quantum entanglement allows qubits that are separated by incredible distances to interact with each other instantaneously. The speed of interaction is not limited to the speed of light. No matter how great the distance between the correlated particles, they will remain entangled as long as they are isolated. With this, instantaneous communications become possible.¹⁸

To date, quantum computers exist as small laboratory demonstrations. Figure 5 shows dicarbonylcyclopentadienyl (perfluorobutadien-2-yl) iron ($C_{11}H_5F_5O_2Fe$ - also known as pentafluorobutadienyl cyclopentadienyldicarbonyl-iron complex). It formed the 7-qubit quantum computer used by IBM researchers who were the first to demonstrate Shor's factoring algorithm.* Significant research is occurring in developing scalable quantum computer architectures and new demonstrations and discoveries are occurring daily.

One potential application of quantum computers is image processing. One such application would be physics model-based automatic target recognition, which is a classically computationally intensive task. Figure 6 is a simple hexagon with 12 vertices. Classically, to verify this object with a high degree of probability one might check orientations by 1 degree over the x, y and z planes. This leads to $360 \times 360 \times 180$ or 23,328,000 orientations which must be verified. Application of Grover's algorithm tells us that in the quantum domain, one would need to verify only 4829 orientations, or $O(\sqrt{N})$. This is a significant reduction in the processing which must be performed.




Fig. 5 $C_{11}H_5F_5O_2Fe$ used in quantum computing demonstration.

*Data available online at IBM Research News, http://www.research.ibm.com/resources/news/20011219_quantum.shtml, 19 Dec. 2001 (cited July 2004).




Fig. 6 Illustration of a simple hexagon.

VI. Biomolecular Computing

BioMolecular Computing (BMC) is the field of computing which uses biological materials in the fabrication of the computer. The biological material may constitute the active device which performs the fundamental compute operation. It may also serve as passive infrastructure material or a combination of both.

We are reaching fundamental limitations to continue Moore's law progress in increasing computing power per dollar spent, creating the need to find methods for progress beyond scaling.* The inclusion of copper metallization and low-k dielectrics recently accelerated the pace of change in materials while maintaining Moore's law progress. As features approach molecular size revolutionary changes in materials and fabrication techniques will be required to continue progress. A growing community is looking at biomolecular materials as the solution to continue Moore's law progress towards faster, cheaper, lower power consuming computers after the current paradigm runs into difficulties around 2010-2012. While most researchers agree significant change is necessary, many are not sure how it will be accomplished on a system scale.

BMC takes advantage of considerable knowledge developed to understand and predictably modify the properties of biomaterials, such as DNA. There have been numerous paradigms proposed for BMC as defined above, using recombinant DNA technology,^{3,6,19} bioengineered proteins²⁰ or cell-like entities including neural networks.²¹

The application of mathematical Formal Language Theory to describe DNA splicing systems introduced the use of DNA recombination as a formal information system.^{22,23} This led to DNA recombination as computation, by synthetic biological processes. Tremendous strides in the development of laboratory techniques for manipulation of DNA, developed for biotechnology and molecular biology, have cleared or dramatically reduced several of the roadblocks to setting up and operating engineered biomolecular computing systems. Alternates to DNA include RNA, enzymes and engineered proteins.

Many accomplishments related to BMC have occurred in the last decade. The first DNA computing system was demonstrated in 1994 and performed a difficult type of calculation important to logistics planning.¹⁹ The light sensitive rhodopsin protein was proposed in 1994 as a memory substrate for a molecular computer.²⁰ A functioning memory was demonstrated using bioengineered protein. The system could be adapted to associative storage and retrieval of images as well as volumetric storage of data.²⁴




Fig. 7 DNA Barcode Lattice. Courtesy of John Reif, Duke University.

The use of DNA as a structural material has opened many new possibilities for engineered nano-scale fabrication of computing systems. Beginning with efforts to improve crystallization of DNA for determination of natural structure,²⁵ the development of synthetic DNA as a structural material for nano-scale fabrication has proceeded to the demonstration of a nano-scale barcode and various two-dimensional and three-dimensional structures,²⁶ see Fig.

*Data available online at International Technology Roadmap for Semiconductors public report, <http://public.itrs.net/> (cited July 2004).

7. The potential is for computing systems where DNA serves as a self-assembling scaffold that positions active components, which may or may not be DNA (Refs. 27-28) see Fig. 8. The use of biomaterials will provide self-assembly of 10^{15} molecular scale components in parallel and could be the most cost effective engineering solution to fabrication of nanotube based ballistic electron and quantum computers. Synthetic DNA has been proposed to assemble quantum dot and nanotube based computing systems. Biomaterials will also provide the ability to fabricate unique architectures and information processing schemes. One example is intrinsic contextual addressing where information is stored and operated on, based on the information's context or semantic meaning. Understanding context is important for increasing the intelligence of computer decisions.




Fig. 8 Open 2 dimensional lattice of DNA. Hooks could be attached at vertices to anchor devices, e.g. RAM array. Courtesy of John Reif, Duke University.

Work is now in progress to metallize bio-molecules, producing conductors and semiconductors, for ballistic electron operations. Two promising paradigms involve plating DNA to create a nanowire²⁹ or replacing hydrogen ions in the DNA base pairs with metal, creating a conductive DNA strand.³⁰ Metallization of biomolecules will enable efficient interfaces between chemical operations and ballistic electronic operations in hybrid systems and facilitate self-assembly of multi-material / multi-physics molecular scale systems.

In the Advanced Computing Architectures Focus Area, the emphasis for BMC is on development of beyond device level architectures, development of metrics and evaluation of technologies, development of modeling, simulation and design tools for hybrid systems, and advanced information architectures for increased cognizance. BMC, as envisioned by the authors, is not expected to be a fully functional technology for another 20 years. The current strategy is to concentrate on algorithm development, information assurance, hybrid information and hardware architecture development, and metrics and evaluation of new computation paradigms.

VII. Genetic Algorithms

The new paradigms of biomolecular and quantum computing offer the long-term promise of an exponential speedup in computations through their inherent parallel processing mechanisms that imply extreme scalability. Consequently, new models of computation need to be developed and implemented. To bridge the gap between near term silicon solutions and future biomolecular and quantum information systems, novel architectures applying Evolutionary Computing methods modeled on nature are appealing. Two methods, Genetic Algorithms (GA) and Genetic Programming (GProg) have an inherent potential for parallelization as they use simple heuristics that may be more easily implemented in hardware than other solution methods using complex heuristics. In addition, recent research indicates that non-traditional algorithms used by GAs and GProg may be implemented utilizing emerging biomolecular computing technologies. Such use would be applicable to solving various hard, NP-complete optimization problems.^{31,32} For example, genetic algorithm methods may be applied to parameterizing sets of differential equations,^{33,34} discovering Hamiltonian paths, Boolean Satisfiability, and Traveling Salesman Problems as well as network and distributed data base design. GProg methods may be applied to discovering the form of the models themselves,³⁵ for tuning silicon circuit designs,³⁶ and for the composition of reduced order models that mirror the predictions of complex (but slow) simulation methods. Predicting toxic plume behavior as a function of weather and topology is one example where the need for reduced order models is critical. Specific problem domains of interest to Information Directorate include complex mathematical model parameterization, communications network design with constrained resources, scheduling and assignment of routes for planes, and target sequencing.

Two hypotheses are currently being examined in an in-house effort. The first hypothesis is that the classical optimization tools being used today will not scale well to solve the complex modeling and systems simulation tasks being envisioned for the future (e.g. designing bio-systems for data input/output, data storage, and computational mechanisms using protocols involving DNA, RNA, and even living bacteria). The second hypothesis is that classical algorithms based on complex, deep heuristics will not translate well to hardware and subsequently run as fast as the

relatively simple GAs and GProg algorithms. The in-house effort is aimed at identifying both in-house and externally developed codes for solving such problems, and evaluating their performances compared to classical methods. This research will utilize several types of computing platforms ranging from a single PC running general purpose software, to a PC with time intensive portions of the code embedded in an FPGA, to a cluster computer, and to a cluster architecture incorporating FPGAs at each node. This project will help frame longer-term questions regarding the potential for scaling and miniaturization of such computing architectures.

VIII. Conclusion

The authors have provided a glimpse at computer technology that is being explored as a means to obtain novel information processing paradigms for future information systems. The technologies presented in this paper represent only a small portion research and development being conducted in the Advanced Computing Architectures Focus Area within AFRL's Information Directorate. The computer technology is being pursued through a spiral research and development process to address current needs, prepare the next generation of technology for upcoming requirements, and foster basic research to plant the seeds of future innovation. The authors believe future information systems are likely to be hybrid systems of the technologies described above. Not only will they be able to process information faster, they will acquire new attributes. Creating hybrid systems by combining General, Electronic, Silicon computing Systems, (GESS), with non-GESS systems, such as biomolecular and quantum computing systems, will require new architectures. The non-GESS system will have an intrinsic architecture requiring interface hardware and information transducing architectures and may be unique GESS architectures enabled by the input from the non-GESS addition. This complicates the problem of integration and presents an opportunity for architecture level modeling and simulation. There is a definite need for additional research and development.

References

- ¹Kampfner, R. R., "Integrating Molecular and Digital Computing: An Information Systems Design Perspective," *BioSystems* 35, Elsevier, 1995, pp. 229-232.
- ²Rambidi, N. G., "Biomolecular Computer: Roots and Promises," *BioSystems* 44, Elsevier, 1997, pp. 1-15.
- ³Garzon, M. H., and Deaton, R. J., "Biomolecular Computing and Programming," *IEEE Transactions on Evolutionary Computation*, Vol. 3, No. 3. Sept. 1999, pp. 236-250.
- ⁴Brenneman, A., and Condon, A., "Strand Design for Biomolecular Computation," *Theoretical Computer Science*, Elsevier, Vol. 287, 2002, pp. 39-58.
- ⁵Amos, M., Paun, G., Rozenberg, G., and Salomaa, A., "Topics in the Theory of DNA Computing," *Theoretical Computer Science*, Elsevier, Vol. 287, 2002, pp. 3-38.
- ⁶Reif, J. H., "The Emerging Discipline of Biomolecular Computation in the US," *New Generation Computing*, Ohmsha Ltd. and Springer-Verlag, Vol. 20, 2002, pp. 217-236.
- ⁷Reif, J. H., "Alternative Computational Models: A Comparison of BioMolecular and Quantum Computation," *Proceedings of the 18th International Conference on Foundations of Software Technology and Theoretical Computer Science*, 1998.
- ⁸Vandersypen, L., and Chuang, I., "NMR Quantum Computing – Lessons for the Future," *Quantum Information and Computation*, Rinton Press, Vol. 2, Special 2001, pp. 134-142.
- ⁹Nielsen, M., and Chuang, I., *Quantum Computation and Quantum Information*, Cambridge Univ. Press, 2000.
- ¹⁰Graham, P., and Nelson, B., "Genetic Algorithms in Software and in Hardware—A performance analysis of workstation and custom computing machine implementations," *1996 Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines*, April 1996, pp. 216 – 225.
- ¹¹Higuchi, T., "Evolvable Hardware and its Industrial Applications," *2002 Genetic and Evolutionary Computing Conference*, (Tutorial track), International Society for Genetic and Evolutionary Computation, New York, July 2002, pp. 437-456.
- ¹²Musmanno, J., "Data Intensive System (DIS) Benchmark Performance Summary," Final Technical Report, Report Number AFRL-IF RS-TR-2003-198, Contract F30602-99-C-0153, Aug. 2003.
- ¹³Burger, D., Goodman, J., and Kagi, A., "Memory Bandwidth Limitations of Future Microprocessors," *Proceedings of the 23rd International Symposium on Computer Architecture (ISCA)*, Association for Computing Machinery Special Interest Group on Computer Architectures and the IEEE Computer Society, May 1996.
- ¹⁴Hall, M., Kogge, P., Koller, J., Diniz, P., Chame, J., Draper, J., LaCoss, J., Granacki, J., Brockman, J., Srivastava, A., Athas, W., Freeh, V., Shin, J., and Park, J., "Mapping Irregular Applications to DIVA, a PIM-based Data-Intensive Architecture," *Proceedings of Super Computing 1999*, Nov. 1999.
- ¹⁵Report of the Defense Science Board, Task Force on DoD Supercomputing Needs, 11 Oct. 2000.
- ¹⁶Defense Research and Engineering, "Survey and Analysis of the National Security High Performance Computing Architectural Requirements," 4 June 2001.
- ¹⁷IHEC 2002 Study, "Report on High Performance Computing for the National Security Community," July 2002.
- ¹⁸Phister, P., Drager, S., and Plonisch, I., "EP Communications: The Next Revolutionary Step in Global Communications," *2003 IEEE Aerospace Conference*, 2003, pp. 6-3021–6-3026.

- ¹⁹Adleman, L., "Molecular Computation of Solutions to Combinatorial Problems," *Science*, Vol. 266, 1994, pp. 1021-1024.
- ²⁰Birge, R., "Protein-Based Optical Computing and Memories," *IEEE Computer*, Nov. 1992, pp. 56-67.
- ²¹DeMarse, T., Wagenaar, D., Blau, A., and Potter, S., "The Neurally Controlled Animat: Biological Brains Acting With Simulated Bodie," *Autonomous Robots*, Vol. 11, 2001, pp. 305-310.
- ²²Head, T., "Formal Language Theory and DNA: An Analysis of the Generative Capacity of Specific Recombinant Behaviors," *Bulletin of Mathematical Biology*, Vol. 49, 1987, pp. 737-759.
- ²³Kari, L., "DNA Computing – The Arrival of Biological Mathematics," *The Mathematical Intelligencer*, Vol. 19, No. 2, 1997, pp. 9-22.
- ²⁴Birge, R., Gillespie, N., Izaguirre, E., Kusnetzow, A., Lawrence, A., Singh, D. Song, W., Schmidt, E., Stuart, J., Seetharaman, S., and Wise, K., "Biomolecular Electronics: Protein-based Associative Processors and Volumetric Memories," *Journal of Physical Chemistry B*, Vol. 103, No. 49, 1999, pp. 10746-10766.
- ²⁵Seeman, N., "Nucleic Acid Junctions: Building Blocks for Genetic Engineering in Three Dimensions," *Biomolecular Stereodynamics*, ed. by R Sarma, Adenine Press, New York, 1981, pp. 269-277.
- ²⁶Yan, H., LaBean, T., Feng, L., and Reif, J., "Directed Nucleation Assembly of Barcode Patterned DNA Lattices," *PNAS*, Vol. 100, No. 14, 8 July 2003, pp. 8103-8108.
- ²⁷Winfree, E., "Algorithmic Self-Assembly of DNA: Theoretical Motivations and 2D Assembly Experiments," *Journal of Biomolecular Structure and Dynamics*, Conversation Vol. 11, No. 2, 2000.
- ²⁸Carbone, A., and Seeman, N., "Circuits and Programmable Self-Assembling DNA Structures," *Proceedings of National Academy of Science*, Vol. 99, 2002, pp. 12577-12582.
- ²⁹Aeren, K., Krueger, M., Gilad, R., Ben-Yoseph, G., Sivan, U., and Braun, E., "Sequence Specific Molecular Lithography on Single DNA Molecules," *Science*, Vol. 297, 2002, pp. 72-75.
- ³⁰Rakitin, A., Aich, P., Papadopoulos, C., Kobzar, Y., Vedeneev, A., Lee, J., and Xu, J., "Metallic Conduction Through Engineered DNA: DNA Nanoelectronic Building Blocks," *Physics Review Letters*, Vol. 86, No. 16, 2001, pp. 3670-3673.
- ³¹Aporntewan, C., and Chongstitvatana, P., "A Hardware Implementation of the Compact Genetic Algorithm," *Proceedings of the 2001 Congress on Evolutionary Computation*, Vol. 1, 2001, pp. 624 -629.
- ³²Ferguson, M. I., Zebulum, R., Keymeulen, D., and Stoica, A., "An Evolvable Hardware Platform based on DSP and FPGA," *2002 Genetic and Evolutionary Computing Conference (Late Breaking Papers volume)*, International Society for Genetic and Evolutionary Computation, New York, July 2002, pp. 145-152.
- ³³Coello, C., Van Veldhuizen, D., and Lamont, G., *Evolutionary Algorithms for Solving Multiobjective Problems*, Kluwer, 2002.
- ³⁴Iba, H., and Sakamoto, E., "Inference of Differential Equation Models by Genetic Programming," *2002 Genetic and Evolutionary Computing Conference (Proceedings)*, International Society for Genetic and Evolutionary Computation, New York, July 2002, pp. 788-795.
- ³⁵Koza, J. A., Mydlowec, W., Lanza, L., Yu, J., and Keane, M. A., "Reverse Engineering and Automated Synthesis of Metabolic Pathways From Observed Data Using Genetic Programming," *Stanford Medical Informatics Technical Report SMI-2000-0851*, 2000.
- ³⁶Zebulon, R. S., Guo, R., Stoica, A., Keymeulen, D., and Ferguson, M., "Case Studies on the Evolutionary Synthesis of Computational Circuits, Filters, and Complete Fuzzy Systems," *2002 Genetic and Evolutionary Computing Conference (Late Breaking Papers volume)*, International Society for Genetic and Evolutionary Computation, New York, July 2002, pp 520-527.