

A fundamental challenge in radiation research related to human health is to predict the biological impact of exposure to low dose (<0.1 Gy) ionizing radiation (IR). Excess cancers have been observed in the Japanese atomic-bomb survivors at doses of 0.1 to 4 Gy, which are 40 to 1600 times the average yearly background levels in the USA. The excess risks vary significantly with gender, attained age, and age at exposure for all solid cancers as a group and many individual sites as a consequence of the atomic bomb (1). It has been estimated that if radiation exposure occurs at age 30, the solid cancer rates at age 70 is increased by about 35% per Gy (90% CI 28%; 43%) for men and 58% per Gy (90% CI 43%; 69%) for women (1). However, predicting cancer risk in populations exposed to doses lower than ~0.1 Gy is limited by statistical considerations.

The most recent review of the biological effects of ionizing radiation (BEIR VII, 2007) of the National Academy of Sciences concluded that human health risks continue in a linear fashion at low doses without a threshold such that the smallest dose has the potential to increase cancer risk. The scientific rationale for linearly extrapolating radiation health effects is underpinned by biophysical theory of how energy interacts with DNA, which is thought to be the major biological target. This area of radiation biology has made significant progress in identifying the critical mechanisms, processes and pathways by which DNA is damaged, repaired or misrepaired. The efficiency and frequency by which IR induces mutations and chromosomal aberrations is thought by most to be the best surrogate of its carcinogenic potential, in part because there is a clear mechanistic understanding of these genomic modifications via energy deposition, and because these events are strongly associated with cancer. A fundamental principle of target theory is that the effect (e.g. DNA damage, cell kill, mutation) is linear or linear/linear-quadratic as a function of dose due to biophysical considerations that energy deposition (i.e. dose) is proportional to damage.

However biological responses to DNA damage quickly evolve and amplify in a non-linear manner, particularly at low doses (reviewed in (2, 3)). There are now myriad experimental reports that low dose radiation (1) alters the response of cells and tissues to subsequent challenge doses (i.e. adaptive responses, AR), (2) affects daughter cell fates such as differentiation and senescence, (3) induces long-range signals that affect non-irradiated cells, and (4) generates a state of chronic genomic instability (GIN). Although there are several definitions of non-targeted effects, we define non-targeted effects as those that are inconsistent with either direct energy deposition, such as bystander phenomenon (4-7), or those that are exhibited in the daughters of irradiated cells, but not mediated by a mutational mechanism, such as radiation-induced genomic instability (8-12) and persistent phenotypic changes (13-16). Although the extent to which these phenomena reflect different molecular mechanisms is not clear, experimental results to date suggest that significant deviation from linearity at low doses may impact the ability to predict cancer risk in humans (17-21).

Our overarching hypothesis is that cancer emerges as a result of a complex, but ultimately predictable, interplay between targeted and non-targeted radiation effects in the context of host genetics and physiology (22, 23). Just as DNA damage elicits a dramatic transition in signaling within a cell, each irradiated tissue has its own set of signals and cell types, distinct from those of un-irradiated tissue and different from other irradiated tissues. The sum of these events, occurring in different organs and highly modulated by genotype, predicates the consequence to the organism. Describing this complexity, identifying key mediators and predicting health outcomes for individuals requires new multiscale modeling of the biology in irradiated tissues.

References

1. Preston DL, Ron E, Tokuoka S, *et al.* Solid cancer incidence in atomic bomb survivors: 1958-1998. *Radiat Res* 2007;168(1):1-64.
2. Brooks AL. Paradigm Shifts in Radiation Biology: Their Impact on Intervention for Radiation-Induced Disease. *Radiation Research* 2005;454-61.
3. Wright EG, Coates PJ. Untargeted effects of ionizing radiation: Implications for radiation pathology. *Mutation Res* 2006 2006/5/11;597(1-2):119-32.
4. Kaplan HS, Carnes WH, Brown MB, Hirsch BB. Indirect Induction of Lymphomas in Irradiated Mice: I. Tumor Incidence and Morphology in Mice Bearing Nonirradiated Thymic Grafts *Cancer Res* 1956 June 1, 1956;16(5):422-5.
5. Hei TK, Wu LJ, Liu SX, Vannais D, Waldren CA, Randers-Pehrson G. Mutagenic effects of a single and an exact number of alpha particles in mammalian cells. *Proc Natl Acad Sci U S A* 1997;94(8):3765-70.
6. Barcellos-Hoff MH, Ravani SA. Irradiated mammary gland stroma promotes the expression of tumorigenic potential by unirradiated epithelial cells. *Cancer Res* 2000;60:1254-60.
7. Mothersill C, Rea D, Wright EG, *et al.* Individual variation in the production of a 'bystander signal' following irradiation of primary cultures of normal human urothelium. *Carcinogenesis* 2001;22(9):1465-71.
8. Kadhim MA, Lorimore SA, Hepburn MD, Goodhead DT, Buckle VJ, Wright EG. Alpha-particle-induced chromosomal instability in human bone marrow cells. *Lancet* 1994 1994 Oct 8;344(8928):987-8.
9. Kadhim MA, Lorimore SA, Townsend KM, Goodhead DT, Buckle VJ, Wright EG. Radiation-induced genomic instability: delayed cytogenetic aberrations and apoptosis in primary human bone marrow cells. *Int J Radiat Biol* 1995 1995 Mar;67(3):287-93.
10. Kadhim MA, Macdonald DA, Goodhead DT, Lorimore SA, Marsden SJ, Wright EG. Transmission of chromosomal instability after plutonium alpha-particle irradiation [see comments]. *Nature* 1992 1992 Feb 20;355(6362):738-40.
11. Clutton SM, Townsend KM, D.T. G, J.D A, Wright EG. Differentiation and delayed cell death in embryonal stem cells exposed to low doses of ionising radiation. *Cell Death Differ* 1996 Jan;3(1):141-8.
12. Limoli CL, Kaplan MI, Corcoran J, Meyers M, Boothman DA, Morgan WF. Chromosomal instability and its relationship to other end points of genomic instability. *Cancer Res* 1997 1997 Dec 15;57(24):5557-63.
13. Herskind C, Rodemann HP. Spontaneous and radiation-induced differentiation of fibroblasts. *Experimental Gerontology* 2000 2000/9;35(6-7):747-55.
14. Rave-Frank M, Virsik-Kopp P, Pradier O, Nitsche M, Grunefeld S, H. S. In vitro response of human dermal fibroblasts to X-irradiation: relationship between radiation-induced clonogenic cell death, chromosome aberrations and markers of proliferative senescence or differentiation. *Int J Radiat Biol* 2001;77:1163-74.
15. Park CC, Henshall-Powell R, Erickson AC, *et al.* Ionizing Radiation Induces Heritable Disruption of Epithelial Cell-Microenvironment Interactions. *Proc Natl Acad Sci* 2003;100(19):10728-33.
16. Tsai KKC, Chuang EY-Y, Little JB, Yuan Z-M. Cellular Mechanisms for Low-Dose Ionizing Radiation-Induced Perturbation of the Breast Tissue Microenvironment. *Cancer Res* 2005 August 1, 2005;65(15):6734-44.

17. Baverstock K. Radiation-induced genomic instability: a paradigm-breaking phenomenon and its relevance to environmentally induced cancer. *Mutat Res* 2000;454(1-2):89-109.
18. Wright EG. Inducible genomic instability: new insights into the biological effects of ionizing radiation. *Med Confl Surviv* 2000;16:117-30.
19. Barcellos-Hoff MH, Brooks AL. Extracellular signaling via the microenvironment: A hypothesis relating carcinogenesis, bystander effects and genomic instability. *Radiat Res* 2001;156(5):618-27.
20. Huang L, Snyder AR, Morgan WF. Radiation-induced genomic instability and its implications for radiation carcinogenesis. *Oncogene* 2003;22(37):5848-54.
21. Little JB. Genomic instability and bystander effects: a historical perspective. *Oncogene* 2003 Oct 13;22(45):6978-87.
22. Barcellos-Hoff MH. Integrative radiation carcinogenesis: interactions between cell and tissue responses to DNA damage. *Semin Cancer Biol* 2005 Apr;15(2):138-48.
23. Barcellos-Hoff MH. Cancer as an Emergent Phenomenon in Systems Radiation Biology. *Radiat Env Biophys* 2007 Feb;47(1):33-8.

High Performance Computational Modeling of Neural Systems

The human brain is the most complex system that we presently know, and the only such system that stands a realistic chance of understanding its own function. In the past 50 years a combination of experimental studies, computational modeling and theory have produced remarkable progress in our understanding of the individual processing elements (neurons) and even of the function of some very simple circuits. The last two decades have produced powerful new methods for the study and manipulation of large neural populations, at the level of cellular networks or extended brain systems. Our progress is now limited by the need to deal with more realistic (and more complex) neuronal models, and to accommodate the variety and sheer numbers of neurons and the connections between them.

There are several areas of basic science, clinical medicine and engineering where high performance computational modeling is likely to have significant, even transformational impact:

- 1) Analyzing and interpreting experimental data from neural systems, typically requiring the coupling of neurophysiological and biophysical modeling techniques.
- 2) Exploring complex interactions within the system and predicting the consequences of intervention, e.g. by pharmacological treatment or electromagnetic stimulation.
- 3) Understanding information encoding and processing mechanisms that support neural computation, enabling synthetic sensory cognition, and neuromimetic and neuromorphic electronic systems.

Each of these applications requires simulation across multiple spatial and temporal scales. For example, neural simulations must incorporate geometrically realistic models of neurons, neural networks, and neural tissue to allow prediction of experimentally observable responses, including responses of individual cells, spatially resolved responses of networks of neurons provided by electrode arrays or novel optical imaging techniques, or the integrated, noninvasive responses of large neural systems such as the electroretinogram (ERG), MEG or EEG, and functional MRI. These methods are also needed to explore feasibility or to optimize proposed new methods for measurements of neural population response such as MRI techniques sensitive to neuronal currents, or optical tomography in scattering tissue. We need to develop capabilities to simulate the neural response to artificial stimulation by applied currents or magnetic fields as employed by emerging systems for electroneural prosthesis, in order to provide an engineering basis for the design of such systems.

There are existing codes for neural simulation, including Neuron GENESIS, and the Los Alamos codes SENSE/PetANNET/PetaVision. Taken together these provide many of the required basic capabilities. Some of these codes have already been applied in the HPC context; an early version of PetaVision was used to build a large network of integrate and fire neurons that was the first research code to achieve sustained petaflop performance on the Los Alamos Roadrunner computer. However new high performance codes will be specifically designed for massively parallel computing, e.g. to deal with the characteristic patterns of dense local communication and arbitrary long range connections. The system should support hierarchical models (ion channels, subcellular structures, realistic cells, networks, tissues) spanning multiple spatial and temporal scales, and allowing simplified or more complete models depending on the requirements of the computational and scientific question. Embedded links to realistic models of cellular and tissue biophysics will allow us to predict experimental observables.

Unlike invertebrate neural systems, which incorporate specialized neurons with very specific patterns of connections, vertebrate systems tend to employ generic neurons and patterns of connectivity that are modified as a function of use. Because we cannot specify patterns of connectivity in detail, we must employ probabilistic, statistical descriptions of neuronal geometry, and of patterns and strengths of interconnection as a function of cellular type, geometric proximity, and patterns of activity. The substrate for these extended circuits is laid down during development, but fine-tuned by patterns of spatial-temporal covariance in activation during early experience.

Given the complexity of large neural systems, computational models offer the only path to truly understand collective function. However existing practices in neural modeling offer few tools or strategies to test and validate system models. To date, the most successful models account for detailed patterns of firing in individual cells or small networks. We have access to new forms of data, from macroscopic physiological measurements and from simultaneous measurements of many neurons in parallel (for example from electrode arrays or dynamic optical imaging) which in principle can be used both to drive and validate model development. However these require the development of coupled biophysical models beyond the scope of those presently employed.

Given the numbers of parameters and elements in even simple models, it is difficult to optimize a model, or explore the parameter space, or to characterize how robust the behavior of the system might be. We will need to incorporate and extend techniques developed for sensitivity estimation, adjoint computation, largescale optimization and probabilistic estimation of model parameters in large scale physics simulations such as climate modeling and mantle convection. The explicit prediction of experimental observables based on network simulations will allow us to optimize large scale models to account for experimental data and to conduct critical tests of these models. Characterizing the response of a large system and tracing the chain of causality that accounts for any observed response is also a substantial challenge which must be addressed.

While interpreting observational data is an important role of HPC simulations, a more exciting objective is to enable systems for synthetic cognition (especially for front end encoding, processing and segmentation of environmental data) that achieves and eventually exceeds the speed, accuracy, flexibility and reliability of human sensory systems. By identifying and emulating the computational principles and architectures of biological neural systems that enable their powerful and adaptive sensory interface to the external world, we can implement and explore *neuromimetic systems* using general-purpose parallel computing hardware. Moreover, we anticipate that the essential capabilities will be implemented in more specialized processing architectures to allow application in the demanding operational environments posed by autonomous vehicles, remote sensing platforms and distributed surveillance environments. *Neuromorphic hardware*, perhaps employing a suitable mix of analog logic and digital communications will allow us eventually to approach the remarkable specifications of size, weight and power that are achieved by biology.

While many of the essential elements (e.g. progressively more specialized feature detectors embedded in hierarchical network architectures) already appear in state of the art, biologically inspired systems for machine vision, eventual solutions will incorporate features often neglected in such models:

Spiking neurons and population encoding: Neurons spike, and to first order information is represented in the rate of firing. However, the precise timing of individual spikes, especially relative to the timing of other spikes within a cohort of neurons, can serve as an additional independent representation of relationships with a data stream, which can be encoded and subsequently extracted in patterns of synchronous firing, coincidence, and coherence within a population of cells.

Complex Architectures: Biological networks employ extensive lateral interconnections between neurons and cortical areas as well as feedback pathways that allow the system to employ specialized cortical processing areas deeper in the pathway to fully exploit all available contextual cues. These pathways also enable top down processing, to apply previously acquired world knowledge, in order to resolve ambiguities that are inevitably present in raw sensory input.

Self-organization, Learning: Cellular learning is a key element of biological networks, both in the early development of feature-based representations and the networks that extract them, and in the continuous processes of adaptation and pattern recognition. We postulate that Spike Timing Dependent Plasticity (STDP) will give rise to many of the characteristic network architectures and information representations observed in neural systems

The effort to understand the function of the brain will likely define the scientific legacy of this century, just as subatomic physics, the human genome project and the development of digital computation defined the last. *Reverse Engineering the Brain* is one of the leading Grand Challenges identified by the National Academy of Engineering. This field is likely to attract the concerted effort of neuroscientists; physicians and psychologists; physicists; material scientists; electrical, computer, and biomedical engineers; and computer and information scientists. It has already attracted substantial interest and investment both within the US and across the world, and the effort will grow.

A principal scientific objective is to understand the collective functional output and mechanisms employed by extended neural systems. With the proposed computational tools we can probe the algorithms employed by these complex systems, with the possibility of testing and validating models through experimental studies. Such computational tools will enable a number of applications of identified interest to DOE and to the nation. Tools for computational simulation of biophysical responses will enable fundamentally new paradigms for the interpretation of data provided by noninvasive techniques for functional imaging and will allow us to critically assess the prospects for enhanced and even revolutionary techniques. Models of biological networks are likely to prove critical for optimal information processing for neural prosthetic systems such as the DOE Artificial Retina. The understanding that we glean from high-resolution models has the potential to revolutionize systems for machine vision employing biomimetic algorithms and architectures, as well as systems for the precise and predictive control of actuators and mobility platforms. These building blocks will be used to assemble engineered systems capable of autonomous operation, including sensory processing and internal state monitoring, adaptation and learning, communication and self-repair. Such systems will enable important new classes of applications with transformational economic and national security impact, ranging from analysis of satellite or ground-based surveillance imagery, to autonomous vehicles and agents, to distributed sensor networks for border security, nonproliferation, and critical infrastructure monitoring.

Narrative for the Blue Brain Project presentation

Sean Hill and Felix Schürmann

The Blue Brain Project aims to simulate an entire human brain of approximately 100 billion neurons and over 10 trillion synapses at the cellular level for basic and medical research. Ultimately, this simulation should be adapted to an individual patient to provide personalized treatment. Currently, there is no drug for which the brain-scale effect is understood. The effectiveness of many pharmacological agents depends on the unique configuration of an individual brain and which receptors and ion channels are coded for in each individual's DNA. Similarly, individual variations in development and experience can influence the organization and functioning of the brain. By integrating genetic, functional and structural data from an individual patient, the model brain can provide personalized medicine and therapeutics, increasing the likelihood of successful treatments and reducing negative secondary effects.

The simulation of an entire human brain at the cellular level of detail will require an estimated 1 exaflop of computation. The simulation of this whole brain model would occur at the electrical level of detail – capturing the dynamics of ion channels and synapses distributed across 100 billion unique neurons. The molecular-scale activity of gene expression, biochemical signaling, protein-protein interactions, the vasculature, glial cells, ion channels, receptors and synapses would be linked to the electrical and cellular scale, capturing the effect of pharmacological agents. Thus, simulation of large-scale molecular-level models as well as large-scale cellular brain-scale models is required.

To iteratively approach simulations at the human brain scale, the Blue Brain toolchain will be extended to build models of rat, mouse, cat and primate brains. Recently, the project has completed a functioning prototype of the neocortical column – the template circuit of the neocortex, which consists of 10,000 physiologically detailed three-dimensional neurons and 30,000,000 synapses based on data from the somatosensory cortex of the young rat and simulated on a 8,192 core 4-rack BlueGene/L system. Currently, the Blue Brain Project is working toward the goal of simulating an entire rat brain on a petascale supercomputer. An entire toolchain for databasing, building, simulating, visualizing and analyzing the neocortical column has been developed and is being used to continually validate the model while integrating additional experimental measurements.

Biological experimental data provides the key validation criteria for the simulations and providing an interactive environment for the neuroscience community will provide an important means for ongoing validation and refinement of the model. The present process includes validating all models by replicating experimental protocols and data including: ion channels, neuron firing behavior, synapses, dendritic integration, morphological parameters, connectivity, polysynaptic loops and emergent network activity.

Whole brain simulation is a new field and only recently have claims been made of simulating brain-scale systems. In contrast to modern physics simulations, the precise benchmarks for what constitutes a valid whole brain simulation and accepted measures to characterize computational efficiency have not yet been established. Furthermore, the canonical algorithms and simulation architectures for establishing whole brain simulations are still under development. The Blue Brain Project is working to establish the canonical algorithms and benchmarks for detailed physiological simulations at the whole brain level, in close collaboration with the author of NEURON, Michael Hines. However, there remains much work to establish a community for verification, validation and performance of whole brain simulation architectures.

In the first phase of the Blue Brain Project we saw that in addition to the simulator a tremendous attention needs to be paid to integrate databasing of the source data with construction of the model, simulation of the model, and analysis of the model in a way that a domain-scientist (not an HPC expert) can operate and understand it. While each of the technical domains in the peta and exascale range represent solid challenges by themselves, only the combination with the application domain, in our case neurobiology, will allow consistency, validation and eventually the much needed impact for personalized medicine.

In addition to the advancements in computing, such a project will necessarily involve international scientific collaborations and industrial scale data acquisition.

Michael Hines

Large scale spiking neural network simulations

Computational Neuroscience aims to understand learning and behavior through simulation of the cells and massive cell connectivity that comprise the nervous system. Models are necessarily simplified since the number of cells in the human brain exceeds 100 billion and number of connections exceeds tens of trillions. Each cell has a complicated tree shape and the cell membrane is also very complex, incorporating hundreds of distinct ion channels and activity dependent biochemical kinetics.

From the viewpoint of communication, large scale spiking neural networks consist of computational units called neurons connected by one-way delay lines to many other neurons. Neurons generate logical events, called spikes, at various moments in time, to be delivered to many other neurons with some constant propagation delay which can be different for different connections. Neurons generally send their spikes to thousands of neurons and receive spikes from thousands of neurons. During time intervals between input events, the neuron is typically defined by a system of continuous ordinary differential equations along with a threshold detector which watches one of the states and determines when the output event is generated.

Spiking neural network models vary greatly in the computational complexity of their neurons. The Blue Brain project for example uses neuron models derived from 3-d reconstructions of large dendritic trees with complex membrane properties due to the presence of several dozen types of nonlinear voltage gated channels selective to the passage of sodium, potassium, and calcium, and whose permeability is also sensitive to the calcium concentration adjacent to the internal surface of the membrane. The largest cell in the Blue Brain project's 10,000 cell neocortical column model is represented using approximately 40,000 coupled ODE's. More typical are simplified nonlinear conductance based single compartment models described by a few to a dozen or so equations. Also commonly used, and the least computationally expensive, are abstract "integrate and fire" models which have simple closed form solutions for their states between events. Such models do not require numerical integration of equations but involve updating their state variables only at input events based on their state at the previous input event. Of course, the simpler the neurons, the larger the network that can be simulated with a given resource, and it is not unusual for models to contain millions of neurons with a total of billions of connections. Nevertheless, the neurons are sparsely connected. Also, such networks remain small when compared to the number of neurons in mammalian brains.

The fastest spike communication method on supercomputer clusters of order ten thousand processors is also the simplest and is based on the MPI_Allgather collective. The allgather method, used by the NEURON simulator as well as other simulators, is based on the fact that network connection delay intervals, typically in the neighborhood of 1 ms, which generally includes axonal and synaptic delay, are generally quite large compared to integration time steps, typically 0.1 ms or smaller, and small compared to interval between spikes of a single neuron, at least 1ms and typically 10--1000 ms. The computations are segregated into integration intervals which are less than or equal to the minimum interprocessor network connection delay. Therefore any spike generated in an interval does not have to be delivered to the target cells until after the end of that interval. All processors work on the same interval, synchronizing only at the end of the current integration interval. Spikes generated by cells on a given processor are stored in a buffer list of (cell identifier, spiketime) pairs and, at the end of an integration interval, the spike count in the buffer along with a fixed size portion of each buffer is exchanged with every other processor using MPI_Allgather. If the number of spikes is larger than the fixed size buffer, the overflow is sent using MPI_Allgatherv.

For less than ten thousand processors, most processors need most spikes, and so allgather performance is usually better than when using point to point exchange methods. However, above that range, i.e. above the ceiling on number of connections per cell, four considerations suggest that the allgather will exhibit poor scaling for very large neural network simulations which, nevertheless, have sparse cell to target processor connectivity. First, MPI_Allgather itself requires twice the time when the number of processors double. Second, all incoming processor buffers must be examined for spikes, even if the spike count for a given source processor is 0. Third, every incoming spike requires a search in a table for whether or not the spike is needed by at least one cell on the processor. Fourth, it is not possible to overlap computation and communication. None of these issues apply to point to point exchange methods using non-blocking sends. As an aside, it is worth noting that because of the biological ceiling on number of connections to a neuron for very large network models, event queue size also begins to scale with the number of cells per processor and so exhibits strong scaling behavior.

Performance tests comparing the allgather and multisend method on a Blue Gene/P for a 256K cell random artificial net show that, for 1k connections per cell, runtime is the same on 8K processors but the multisend method continues to scale linearly up to the largest number (32K processors) we used whereas runtime begins to increase with the

allgather method between 16K and 32K processors. For 10k connections per cell, MPI_Allgather continues to have better performance, even with 32K processors, than a multiseed method that utilizes the persistent DCMF_Multicast (Deep Computing Messaging Framework) implemented using remote direct memory access. Trends show an expected turnover at 64K processors.

Thus, the computational neuroscience community is well placed to take advantage of the coming availability of very large parallel computers and it seems likely that incremental improvements in existing load balance and spike exchange methods will allow efficient use of machines consisting of millions of processors. At such scales, it will probably be necessary to roughly distribute the neurons in a fashion topographically similar to the biological arrangement so as to keep connectivity as local as possible to avoid unnecessary sharing of interprocessor bandwidth. The multiseed method will undoubtedly benefit from a machine specific implementation that avoids sending 10K messages from the source processor to the target processors for each spike in favor of a more distributed transmission method that takes advantage of the fact that biological delay is greater when cells are farther away.

The dynamics of blood flow in the human brain depend upon a complex network of vessels under a variety of temporal and spatial constraints. Abnormalities in the delivery of blood to the brain clearly underlie the pathophysiology of stroke, vasospasm, traumatic brain injury, vascular dementias, and probably conditions such as migraine and hydrocephalus. Clinical decisions are often made on the basis of steady state conditions (e.g., mean intracranial pressures, mean cerebral blood flow, etc), but there is clearly a risk that ignoring the range of spatial and temporal scales present may limit understanding, and hence clinical effectiveness. Exascale computing can facilitate the development of computational multiscale models of the cerebral vasculature that includes all blood vessels from the circle of Willis to the arteriolar tree, to the capillary bed, and even detailed spectrin-level models of the red blood cell (RBC). An even more exciting prospect is the new possibility of a full brain model coupling of the neuronal and vascular trees that will lead to understanding and treatment of devastating diseases such as Alzheimer’s, meningitis and multiple sclerosis. A full-scale model with patient-specific geometry and conditions at all scales, including spectrin-level RBC models, will facilitate *physiologically correct* simulations of brain perfusion and associated pathologies (e.g., malaria, sickle cells anemia) for *realistic* future studies.

A Multiscale Brain Vascular Model – The cardiovascular system of the human body is the envy of every engineer. In just one minute, the average heart beats about 70 times, pumping the entire blood supply of 5 liters through 62,000 miles of vessels, that is one-fourth of the distance between the Moon and the Earth! The human brain, in particular, although less than 2% of the body weight, receives about 20% of the resting cardiac output of blood and 25% of the body’s oxygen supply [1]. Interactions of blood flow in the human brain occur between *different scales*, determined by flow features in the large arteries (diameter of 0.5 mm or larger), the smaller arteries and arterioles (500 μm to 10 μm), and the capillaries (mean diameter of 5 μm) – all being coupled to cellular and sub-cellular biological processes. While many biological aspects have been studied systematically, surprisingly less effort has been put into studying blood flow patterns and oxygen transport within the brain, i.e., the fundamental biomechanical processes of the *integrated* vascular network. However, recent pioneering 3D imaging of the *human* brain by Cassot et al. in [2] and of the *mouse* brain by Choe et al. [3] provides statistical information for constructing realistic topological models on which future brain simulations will be based. The main observation is that arterioles down to 10 μm follow a *tree-like* structure (governed by a fractal law) whereas the capillary bed (below 10 μm) follows a *net-like* structure, i.e., a mesh, see figure 1.

Exascale computing can facilitate the development of an integrated model of the vascular network in the human brain (cerebrovasculature) characterized by three distinct spatial length scales (see figs. 1 and 2):

- (1) *The macrovascular network (MaN) consisting of large arteries*, down to diameter of 0.5 mm, which are patient-specific and can be reconstructed from CT/MR imaging. Typically, about 100^1 such arteries start from the circle of Willis, which is formed downstream of the four main arterial inlets at the neck (two carotids and two vertebral arteries).
- (2) *The mesovascular network (MeN) consisting of small arteries and arterioles*, from 500 μm down to 10 μm , which follow a tree-like structure governed by specific fractal laws [2, 5]. The human brain contains about 10 million² small arteries and arterioles.
- (3) *The microvascular network (MiN) consisting of the capillary bed*, which follows a net-like struc-

¹Due to arterial variations in a human this number is patient-specific [4].

²The number of vessels stated here is computed based on Murray’s law [6, 7] with a modified index ($q = 2.5$) and asymmetric structure.

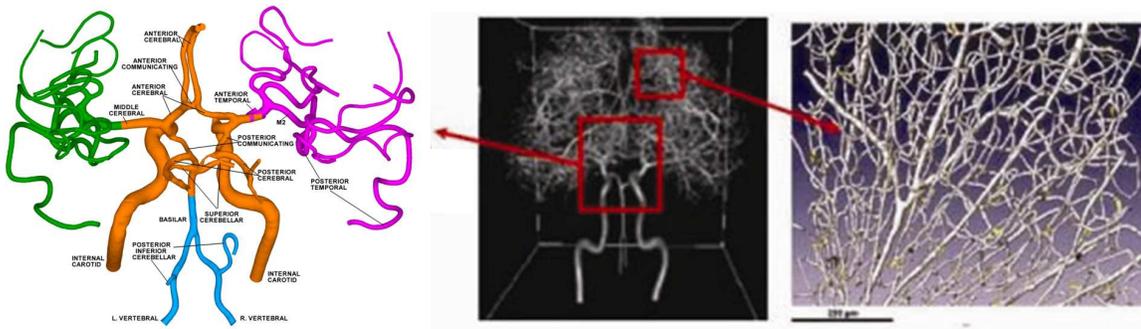


Figure 1: Schematic of multiscale modeling of the cerebrovasculature: Left - MaN: Large arteries and circle of Willis (our data). Middle - MeN: Arteriolar tree model ([5]). Right - MiN: Capillary bed ([2]). The different colors in MaN represent the different “patches” in the two-level domain decomposition method (see “Technical Approach”). The geometry is described statistically in MeN/MaN.

Macrovascular Network (MaN): Large Arteries (100)	Mesovascular Network (MeN): Arteriolar Tree (10M)	Microvascular Network (MiN): Capillary Bed (1B)
3D Navier–Stokes	1D Stochastic PDE Model	* Stochastic Darcy’s Law * DPD Modeling $\rightarrow K_{eff}$
6mm	0.5mm	10 μ m < 10 μ m

Figure 2: Coupling of MaN-MeN-MiN and corresponding mathematical models. The numbers in parenthesis show the approximate number of vessels; K_{eff} denotes an “effective” permeability of the capillary bed, which will be extracted via *upscaling* of DPD simulation results obtained on stochastic replicas.

ture; its topological statistics have been recently quantified for the human brain in [2]. The typical number of capillary segments in the brain is more than 1 billion.

More comprehensively, the simulations can be divided into two regimes: The first one involves all arteries that can be accurately imaged clinically at the present time (see (1) above and preliminary results in figure 3), whereas the second regime involves the “subpixel” dynamics as described by (2) and (3) above.

There are three main reasons for coupling all three networks and not simply model MaN as in figure 3: (1) to provide a closure for MaN modeling, and (2) to model brain perfusion, and (3) to form the foundation for neurovascular coupling and modeling of the blood-brain-barrier (BBB).

A Multiscale Neuro-Vascular Model – While for many years the study of ischemic brain injury and repair focused on the neural tissue, translating these laboratory results into clinically effective stroke treatments or diagnosing the early onset of Alzheimer’s disease, still remain major challenges. It is now recognized that even relatively small changes in blood flow propagate downstream in veins and can give rise to spurious activation of sites remote from neuronal activity.

A rational approach to meet these challenges starts by considering the “neurovascular unit” (NVU) concept, see figure 4, which encompasses the cellular and functional interactions among the capillaries, glia, and neurons of the brain [8]. Targeting the NVU based on biological considerations provides an integrative view to ischemic brain damage which may be closer to the clinical reality. By employing data from Magnetic Resonance Imaging (MRI), computational models can be developed that simulate the changes in NVU in vivo in order to understand the integration of cerebrovascular and neurobiological mechanisms in patients with severe ischemic stroke or degenerative disease. Exaccalle

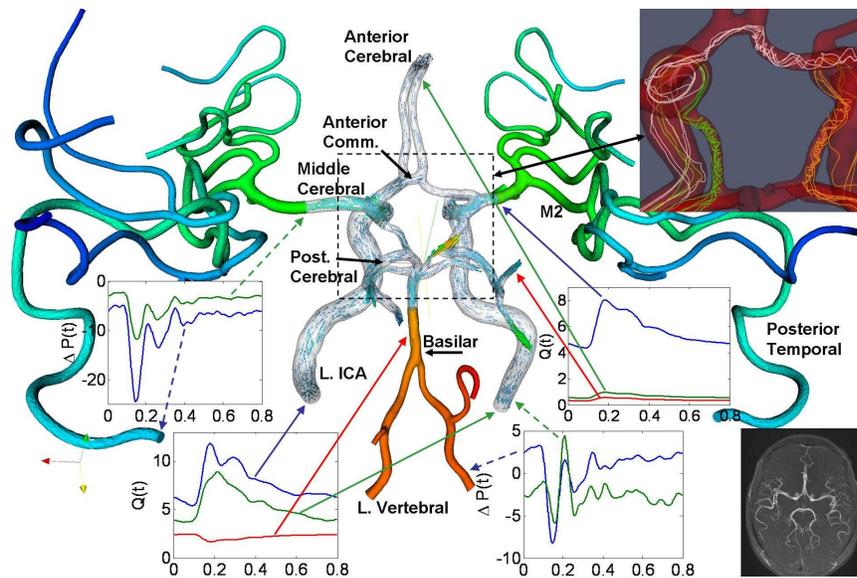


Figure 3: MaN simulations involving a “complete” Circle of Willis consisting of 65 arteries. Colors represent pressure (inner part is transparent to show velocity vectors and not pressure levels) arrows velocity, XY plots depict flowrate in ml/s and pressure drop ΔP in $mmHg$. Top right: instantaneous streamlines shown swirling flow in communicating arteries. Bottom right MRA image of the cranial arterial system provided by our collaborator Dr. J.R. Madsen (Harvard Medical School).

computing can facilitate the development and implementation of such models and in particular the scaling up of a single unit to a representative region of a human brain.

An Exascale Application – Based on detailed estimates and ongoing simulations in our group, we can project that a sustained 10 Petaflops performance can lead to MaN-MeN coupled simulations of the full brain within 24 hours. On the other hand, simulating MiN requires atomistic (mesoscopic) approaches to capture the RBC dynamic accurately in the capillary bed and for 1 billion capillaries we estimate that we will require simulations with 10^{13} particles, which can be performed efficiently on a 10-100 Petaflops platform. Taken together these estimates along with estimates for running NEURON to account for neurovascular coupling, we can project that an exaflop (1000 petaflops) sustained performance will lead to the first ever simulation of a full brain.

References

- [1] R.D. Bevan and J.A. Bevan. *The Human Brain Circulation*. Humana Press, Inc, 1994.
- [2] Cassot F., Lauwers F., Fuard C., Prohaska S., and Lauwers-Cances V. A novel three-dimensional computer-assisted method for a quantitative study of microvascular networks of the human cerebral cortex. *Microcirculation*, 13:1–18, 2006.
- [3] D. Mayerich, J. Kwon, Y. Choe, L. Abbott, and J. Keyser. Constructing high-resolution microvascular models. In *Proceedings of the 3rd International Workshop on Microscopic Image Analysis with Applications in Biology (MIAAB 2008)*, 2008.
- [4] H. Lippert and R. Pabst. *Arterial Variations in Man: Classification and Frequency*. J.F. Bergmann, 1985.

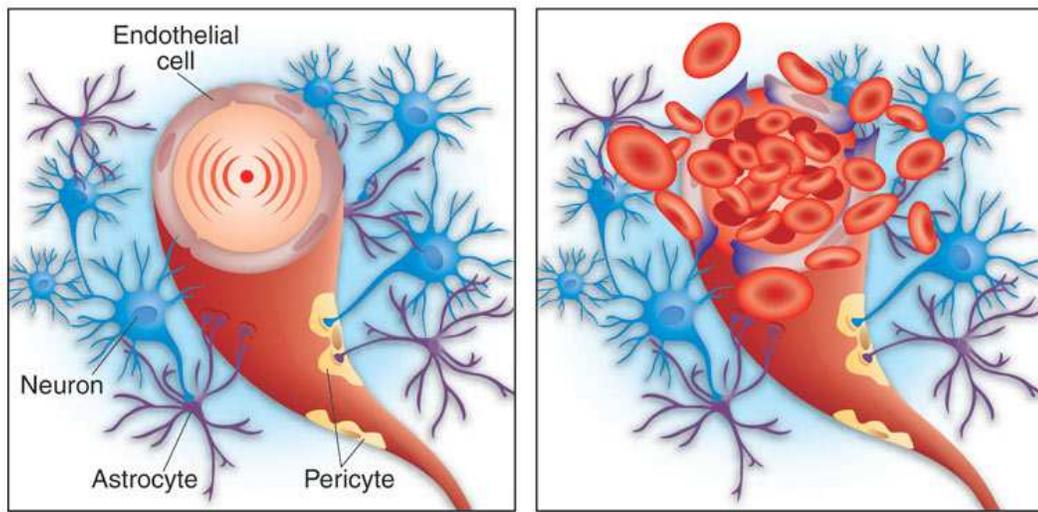


Figure 4: Neurovascular coupling: Flow in the capillaries and associated brain perfusion are strongly coupled to the neuronal activity.

- [5] M. Zamir. On fractal properties of arterial trees. *J. Theor. Biol.*, 197:517–526, 1999.
- [6] C.D. Murray. The physiological principle of minimum work. I The vascular system and the cost of blood volume. *Proc. Natl. Acad. Sci. USA*, 12:207–214, 1926.
- [7] T.F. Sherman. The meaning of Murray’s law. *J. Gen. Physiol.*, 78:431–453, 1981.
- [8] B.T. Hawkins and T.P. Davis. The blood-brain barrier/neurovascular unit in health and disease. *Pharmacological Reviews*, 57(2):173–185, 2005.

Extreme Scale Biology Workshop

Samuel Lang

Technical Lead, Parallel Virtual File System (PVFS)

Research Interests and Challenges

The work of the systems software group in MCS includes providing high performance I/O and storage software for parallel applications, where concurrent, large I/O and many file accesses are common. Leadership class storage systems today provide aggregate I/O rates in the range of 50-100 gigabytes per second to files often as large as a few terabytes. [2] Storage systems of the future will likely be an order of magnitude larger, but only a few times faster due to the physical limits of the hardware. We are investigating new approaches to storage software that allow applications to reach the I/O limits of next generation systems.

Research in parallel I/O software has traditionally focused on parallel applications in application domains other than computational biology, such as climate modeling or astrophysics simulations. But recent research in computational biology has begun to utilize high performance I/O techniques for improving the performance of sequence search. [3] As the I/O and data requirements of computational biology applications increase, new approaches to high performance I/O will be needed. More efficient access to many small files will become important. [1] Active storage models will become increasingly important as well, allowing applications to reduce the I/O bandwidth requirements of the system by performing much of the computation where the storage is located, improving data locality. [5, 4] Further research on active storage techniques integrated within parallel file systems is needed to understand the role that it can play at extreme scales on improving the overall performance of computational biology applications.

Many of the challenges to improved I/O performance for computational biology at extreme scale are in the gap between the I/O patterns that applications prefer, and the larger I/O patterns that get the best performance on leadership-class storage systems. Many open questions still remain. Will leadership-class systems of tomorrow provide the bandwidth and capacity necessary to tackle the hardest problems in computational biology? Can evolving models for improved locality of computation and storage be utilized within the domain of computational biology? Can interfaces to storage be built that give computational biology applications more efficient data access? These are some of the challenges we focus on and hope to tackle for biology applications running at extreme scale.

References

- [1] P. Carns, S. Lang, R. Ross, M. Vilayannur, J. Kunkel, and T. Ludwig. Small file accesses in parallel file systems. In *Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, March 2009.
- [2] S. Lang, P. Carns, R. Latham, R. Ross, K. Harms, and W. Allcock. I/O performance challenges at leadership scale. In *Proceedings of Supercomputing*, November 2009. **(To Appear)**.
- [3] H. Lin, X. Ma, P. Chandramohan, A. Geist, and N. Samatova. Efficient data access for parallel blast. In *IEEE International Parallel & Distributed Processing Symposium*, April 2005.
- [4] S. Patil, G. A. Gibson, G. R. Ganger, J. Lopez, M. Polte, W. Tantisiroj, and L. Xiao. In search of an API for scalable file systems: Under the table or above it? In *USENIX HotCloud Workshop 2009*, June 2009.
- [5] J. Piernas, J. Nieplocha, and E. J. Felix. Evaluation of active storage strategies for the lustre parallel file system. In *Proceedings of Supercomputing*, November 2007.

Gayle E. Woloschak
Extreme Scale Biology Workshop

My laboratory is focused on two major areas of research: The first involves studies of irradiation on the development of late tissue toxicities including cancer. For this project, we have available to us both the dataset, paraffin embedded tissues, and pathology records from over 49,000 mice and 7,000 dogs that were exposed to varying doses of radiation of different qualities and given at different dose rates. We are currently analyzing data from the datasets for a better understanding of late tissue toxicities and tissues from the archive to understand basic mechanisms of radiation effects. The second project involves the development of nanoparticles that can be used for cancer imaging and therapy. These nanoparticles are of different sizes, shapes and are modified with different biological materials to permit targeting in biological systems.

Based on my experience there are several challenge areas that arise for modeling. With regard to radiation studies, there are a few groups that are trying to model the interaction of radiation with DNA and other subcellular components. Few have been able to build up to the full cell level much less the tissue and organism level. If effects of radiation can be modeled on tissues, this would have dramatic impact for radiation oncologists who are trying to avoid radiation toxicity in normal tissues while they are treating tumors. Modeling of the effects of different radiation qualities on tissue toxicities and other secondary endpoints would also be of value; low LET radiations such as gamma-rays and x-rays would be of value to most clinicians. The recent development of p+ beams for therapy has increased demand for understanding tissue toxicities following such treatments. NASA is concerned about space radiation effects including radiation from solar particle events (mostly p+) and galactic cosmic rays (mostly high Z radiation). Finally, being able to model the interaction of radiation with different toxicities would be of benefit to a broad community. Radiation is associated with many environmental toxins in waste sites including heavy metals and others. Radiation is usually administered with chemotherapeutic agents in the clinic, and the interactions are often not predictable.

From the nanotechnology perspective, much funding and work is going into understanding effects of toxicities of nanoparticles. Right now, each nanoparticle is being tested one by one in investigators' labs and at the NanoCharacterization Lab in Frederick, MD. Some ability to model how size, shape, biological modification, etc. will affect the toxicities of each nanoparticle would be of tremendous benefit to the field and would enhance the ability to move this work from the research lab into the clinic. Prediction of trafficking in animals and humans would be of value particularly since so many nanoparticles are being developed for imaging purposes. Some nanoparticles enter tumors by the enhanced permeability and retention mechanism which occurs because of leaky vasculature in tumor cells; others are specifically targeted to tumor cells. Being able to model this would be of great use to the imaging and therapy communities. In addition, there are current concerns about nanoparticle contamination from synthesis and from environmental concerns which would also be impacted by modeling capabilities.