

IN PROGRESS //



Painkiller Snails


University of Utah scientist Tom Cheatham is gaining new insight into snail toxins that can treat human pain

Although they come in pretty shells shaped like ice-cream cones and are popular for jewelry, cone snails are carnivorous hunters, and some of the larger species can be deadly to humans. They are equipped with a barbed, hollow tooth that launches like a harpoon. A neurotoxic venom paralyzes small fish almost instantly.

The good news is that these cone-snail venoms, known as conotoxins, show promise in humans as potent painkillers. One conotoxin-derived drug, Ziconotide, received FDA approval in 2004, and others are being tested as possible treatments for Alzheimer's disease, Parkinson's disease and epilepsy.

University of Utah pharmaceutical chemist Tom Cheatham and visiting student Pawel Gruszczynski from the University of Gdansk in Poland are using Pople, PSC's shared-memory SGI Altix, to investigate a family of conotoxins called Mu-conotoxins. "These toxins," says Cheatham, "are tissue specific and act by binding to voltage-sensitive sodium channels." Using long-timescale molecular dynamics (MD), with software called AMBER, Cheatham's group is working to predict and refine the structure of these conotoxins. With Pople, says Cheatham, the aggregate throughput was excellent. "We benefit from the fastest and most balanced machines."

Their results, analyzed in collaboration with University of Utah experimentalists Greg Bulaj and Toto Olivera, make it possible to map the



These pictures of a Mu-conotoxin show representative structures from clustering that occur with 200 nanoseconds of MD simulation. "Flexibility in the side-chain is evident," notes Cheatham, "although the core backbone structure is largely intact."

pharmacophore — the essential molecular features responsible for the drug's biological activity. This work aims toward identifying the precise atom-by-atom details of how conotoxin molecules — short amino-acid chains called peptides — have their potential pain-killing effect.

Shared Memory Poker

Carnegie Mellon computer scientists are using PSC's Pople to scale up their champion poker program and solve other problems in game theory

You won't find them at the Vegas casinos and what they do is hard to call gambling, but it's fair to say that Tuomas Sandholm and grad student Andrew Gilpin of Carnegie Mellon's School of Computer Science are professional poker players. This July in Chicago — at the AAAI (Association for the Advancement of Artificial Intelligence) Computer Poker Competition, involving 19 programs from six countries — they walked away with no pile of cash but, nevertheless, were the biggest winners.

Their field, game theory, in which Sandholm's work is internationally recognized, describes conflict in which the payoff is affected by actions and counter-actions of intelligent opponents. Head-to-head poker between two players is a prominent example of what's called two-person zero-sum games: One player wins what the other loses.

In recent years, poker has emerged as an AI challenge much as chess was for many years, but poker is far more demanding, says Sandholm: "In poker, a player doesn't know which cards the other player holds or what cards will be dealt in the future."

Like many games, poker can be formulated mathematically, but the formulations are huge. Two-player poker has a game-tree of a billion-billion (10^{18}) nodes, notes Gilpin. "To solve that requires massive computational resources. Our research is on scaling up game-theory solution techniques to those large games, and new algorithmic design." The most computationally intensive portion of their algorithm is a matrix-vector product, where the matrix is the payoff matrix and the vector is a strategy for one of the players. This operation accounts for more than 99-percent of the computation, and is a bottleneck to applying game theory to many problems of practical importance.

To drastically increase the size of problems the algorithm can handle, Gilpin and Sandholm devised an approach that can potentially exploit massively parallel systems of non-uniform memory-access architecture, such as Pople, PSC's SGI Altix. By making all data addressable from a single process, shared memory simplifies a central, non-parallelizable operation performed in

conjunction with the matrix-vector product. Sandholm and Gilpin have revised their code to run on Pople, and are doing experiments to learn how much parallelism can help, and possibly point to areas for further algorithmic improvement.



To Stop the Pandemic

Modeling with PSC shared-memory resources will improve decision-making in deploying vaccines to stop disease outbreaks

With support of a \$10-million grant from the Bill & Melinda Gates Foundation, a team of researchers, including PSC scientist Shawn Brown, is using PSC resources to explore computational models that simulate the spread of infectious disease. The grant funds computer simulations of epidemics to show worst-case and best-case outbreak scenarios. Findings will be used to evaluate new vaccine technologies and modes of vaccine delivery. The goal is more informed decision-making in deploying vaccine technologies in pandemic outbreaks and, in general, to support World Health Organization Communicable Disease Control programs.

The work is coordinated through the Vaccine Modeling Initiative (VMI), headquartered at Pitt's Graduate School of Public Health (GSPH). VMI is a research partnership among infectious-disease modeling teams at the University of Pittsburgh, The Pennsylvania State University and Imperial College London. VMI also involves collaborations with leading infectious disease experts and public health officials at Johns Hopkins University, Médecins Sans Frontières Epicentre, University of Georgia, Resources for the Future and the World

Health Organization. The project will exploit the shared-memory environment of Pople, PSC's SGI Altix, to build extensive agent-based models of disease spread in a number of third-world populations. The group expects that a large shared-memory system will improve their ability to complete full studies in a timely fashion.

"Infectious diseases create an enormous burden on the world's population, from both a human suffering and an economic development perspective," says Donald S. Burke, M.D., principal investigator of the grant and dean of GSPH. "One of the major challenges we face in stopping infectious disease outbreaks is predicting how control strategies, such as vaccines, will work. By using computer models to conduct 'epidemiology in silicon,' we will be able to test the impact of new candidate vaccine technologies and select the most effective strategies."

Initially, the project will focus on evaluation of new vaccine technologies for influenza, measles and dengue,

a mosquito-borne infection, diseases that affect millions of people globally. Later, the project will develop vaccine models of epidemic pertussis, rotavirus, polio, pneumococcus, malaria and tuberculosis.

Day 35 of an influenza outbreak in the United States as simulated with an agent-based model (written by Neil Ferguson, Imperial College, London) on Pople, PSC's SGI Altix. The image shows the infected population (red) and population segments that have recovered (green). PSC scientist Shawn Brown created this visualization with a PSC-developed interface to Google Earth.

Hitting the Fight-or-Flight Target


With help from BigBen, Thomas Huber of Rockefeller University makes progress toward understanding the molecular mechanism of therapeutic drugs

Drugs have targets. They're called receptors — usually large proteins lodged in cell membranes. When another molecule of the right shape, size and electrical charge arrives — called a ligand, it binds to and activates the receptor, igniting a cascade of biochemical processes. Most ligands are from within the body (endogenous), such as neurotransmitters. Drug research involves finding other ligands, either to activate receptors or to block them from being activated by endogenous ligands.

One "superfamily" of receptors — called G protein-coupled receptors, GPCRs for short — is estimated to be the target for 27-percent of approved therapeutic drugs. Within the GPCR family, an important sub-group is adrenergic receptors (ARs), which interact with the "fight-or-flight" hormones epinephrine (adrenaline) and norepinephrine. Many cells have ARs, and binding of these hormones increases heart rate, diverts blood to muscle tissue, and mobilizes energy.

Thomas Huber of Rockefeller University has been using PSC's BigBen to better understand the biophysical principles of ligand binding to GPCRs. In recent work he has focused on one of the ARs — beta(2)-adrenergic receptor, BAR for short — the main target for drugs that treat asthma, by relaxing the pulmonary muscles. The crystal structure of BAR in 2007 was a long-awaited breakthrough in the field. Huber used molecular dynamics (MD) software to simulate BAR structures as soon as they became available for study. Working with PSC scientist Phil Blood to take advantage of BigBen's parallelism, he was able to run four independent

This representation of Huber's simulations shows a BAR structure (seven vertical helical coils) embedded within a bilayer membrane (cyan spheres represent the membrane-water interface), with a ligand called carazolol (amid the vertical coils) that blocks BAR activity. The overall model is about 50,000 atoms. The image shows superimposed structures of protein backbone snapshots every nanosecond over 60 nanoseconds of simulation.



64-processor jobs at one time. "For the work I'm doing, this is the fastest machine available," says Huber. "The queues are relatively predictable and allow good job throughput."

On models of BAR embedded in a membrane with receptor that encompasses 50,000 atoms, he completed 600 nanoseconds of MD simulation. One set of simulations used the beta-blocker drug carazolol, which binds to and blocks activation of the receptor in pharmacological experiments. The second set used the activating drug epinephrine. Findings so far show details of how BAR's ligand-binding pocket changes when bound alternatively to these two drugs. While the structural changes are too small to account for receptor activation, the long simulations on BigBen demonstrate that these changes are statistically significant. "Information on this timescale," says Huber, "shows how receptor activating drugs pull on the receptor in an attempt to toggle its switch." This information can help point the way to new GPCR-targeted drugs.

Attracting Future Scientists

An innovative science-learning program called CMIST is helping teachers bring science to life for high-school students



From the CMIST module "Enzyme Structure and Function": A virtual laboratory designed and rendered with Blender includes common lab equipment related to MCell simulations, which in this case show red blood cells in a heart vessel delivering oxygen to muscle cells.

"The visualizations allow them to understand the concepts as they actually occur," says Rebecca Day, who teaches high-school biology near Pittsburgh. "It's one thing to read about it, another thing entirely to see it happening," says Marian Opest, who teaches science in the Pittsburgh suburb of Penn Hills. They're talking about CMIST (Computational Modules In Science Teaching), an innovative science-learning program for high-school students developed by a team at PSC.

"CMIST addresses the challenge of science learning for the video-gamed, TV-nation, multi-tasking population of students," says Pallavi Ishwad, education outreach specialist for PSC's National Resource for Biomedical Supercomputing (NRBSC). Ishwad, NRBSC director Joel Stiles, scientific e-learning specialist Jenda Domaracki and visualization specialist Jacob Czech created and developed CMIST as a learning tool that could reach the majority of students, not just the self-motivated few. The hook is appealing content in an easily usable form.

In contrast to many other teaching tools, CMIST modules are produced with realistic modeling and simulation software, such as developed and used in research at PSC. The pilot module, "Molecular Transport in Cells," produced with software called MCell and DreAMM (co-authored by Stiles and used in research centers

around the world) presents important principles of osmosis and diffusion with 3D examples. Recent new modules include "Brownian Motion" and "Enzyme Structure and Function." CMIST is distributed as ready-to-use DVDs that include complete lecture slides, animations, a lesson plan aligned to national and state standards, and worksheets with answer keys.

Funded in large part by NIH, the CMIST team has presented the program to regional high-school science teachers in half-day workshops at which teachers have given enthusiastically positive feedback. PSC has also presented CMIST to teachers and administrators at state and national meetings, and is introducing web-based distribution of the prototype module, while also planning additional modules on metabolism, synaptic physiology, meteorology, geophysics and astronomy, informed by some of the forefront research at PSC.

PSC scientist Shawn Brown demonstrates WiiMD at a National Science Foundation Open House in Arlington, Virginia. "The main goal" says Brown, "is providing a means of education and outreach that connects people to supercomputing and simulations in an engaging and familiar way."

Kids and Adults Say "Wow"

The intuitive technology of Nintendo's video console applied to molecular dynamics makes a game of science learning



With an innovative hybrid of game technology and supercomputing, a team at PSC is winning new friends, young and old, to the excitement of discovery and learning. Their interface, called WiiMD, integrates the controller of the popular Nintendo Wii (pronounced "we") game console, the WiiMote, with molecular dynamics simulation software. The result for a wide range of audiences is an experience both entertaining and informative.

The package of simulations includes, for instance, buckyball bowling, in which participants use the WiiMote to impel a "buckyball" carbon molecule to smash a 10-pin formation of other buckyballs. The team has demonstrated WiiMD at various venues, including SCo7 in Reno, the National Science Foundation Open House in Arlington, VA, and at Pittsburgh's Carnegie Science Center SciTech Festival. Children are captivated by the colorful display of molecular shapes that move and change in response to the WiiMote. Along with students from K-12, WiiMD has drawn enthusiastic attention from teachers and the general public.

At the heart of WiiMD are sophisticated biomolecular simulation technologies (NAMD and VMD) developed by Klaus Schulten and colleagues at the University of Illinois Urbana-Champaign and used by researchers world-wide on the largest supercomputing resources available. Jordan Soyke, a University of Pittsburgh student, adapted the WiiMote controller for use with VMD to enable fully 3-D control. Further work by PSC scientists Shawn Brown and Philip Blood — relying on PSC-developed software called PDIO — connected the programs with BigBen, PSC's CrayXT3.

With initial implementation of WiiMD complete, the team is developing a curriculum to go with it. Marylou Kunkle, like Soyke an undergrad at the University of Pittsburgh, is adapting the Linux-based program to run on desktop platforms, such as Windows and Mac OS X, in order to allow easier use by the general public.