DIVERSE DISCIPLINES, ONE COMMUNITY

BiomedicalComputation REVIEW

Published by Simbios, a National Center for Biomedical Computing

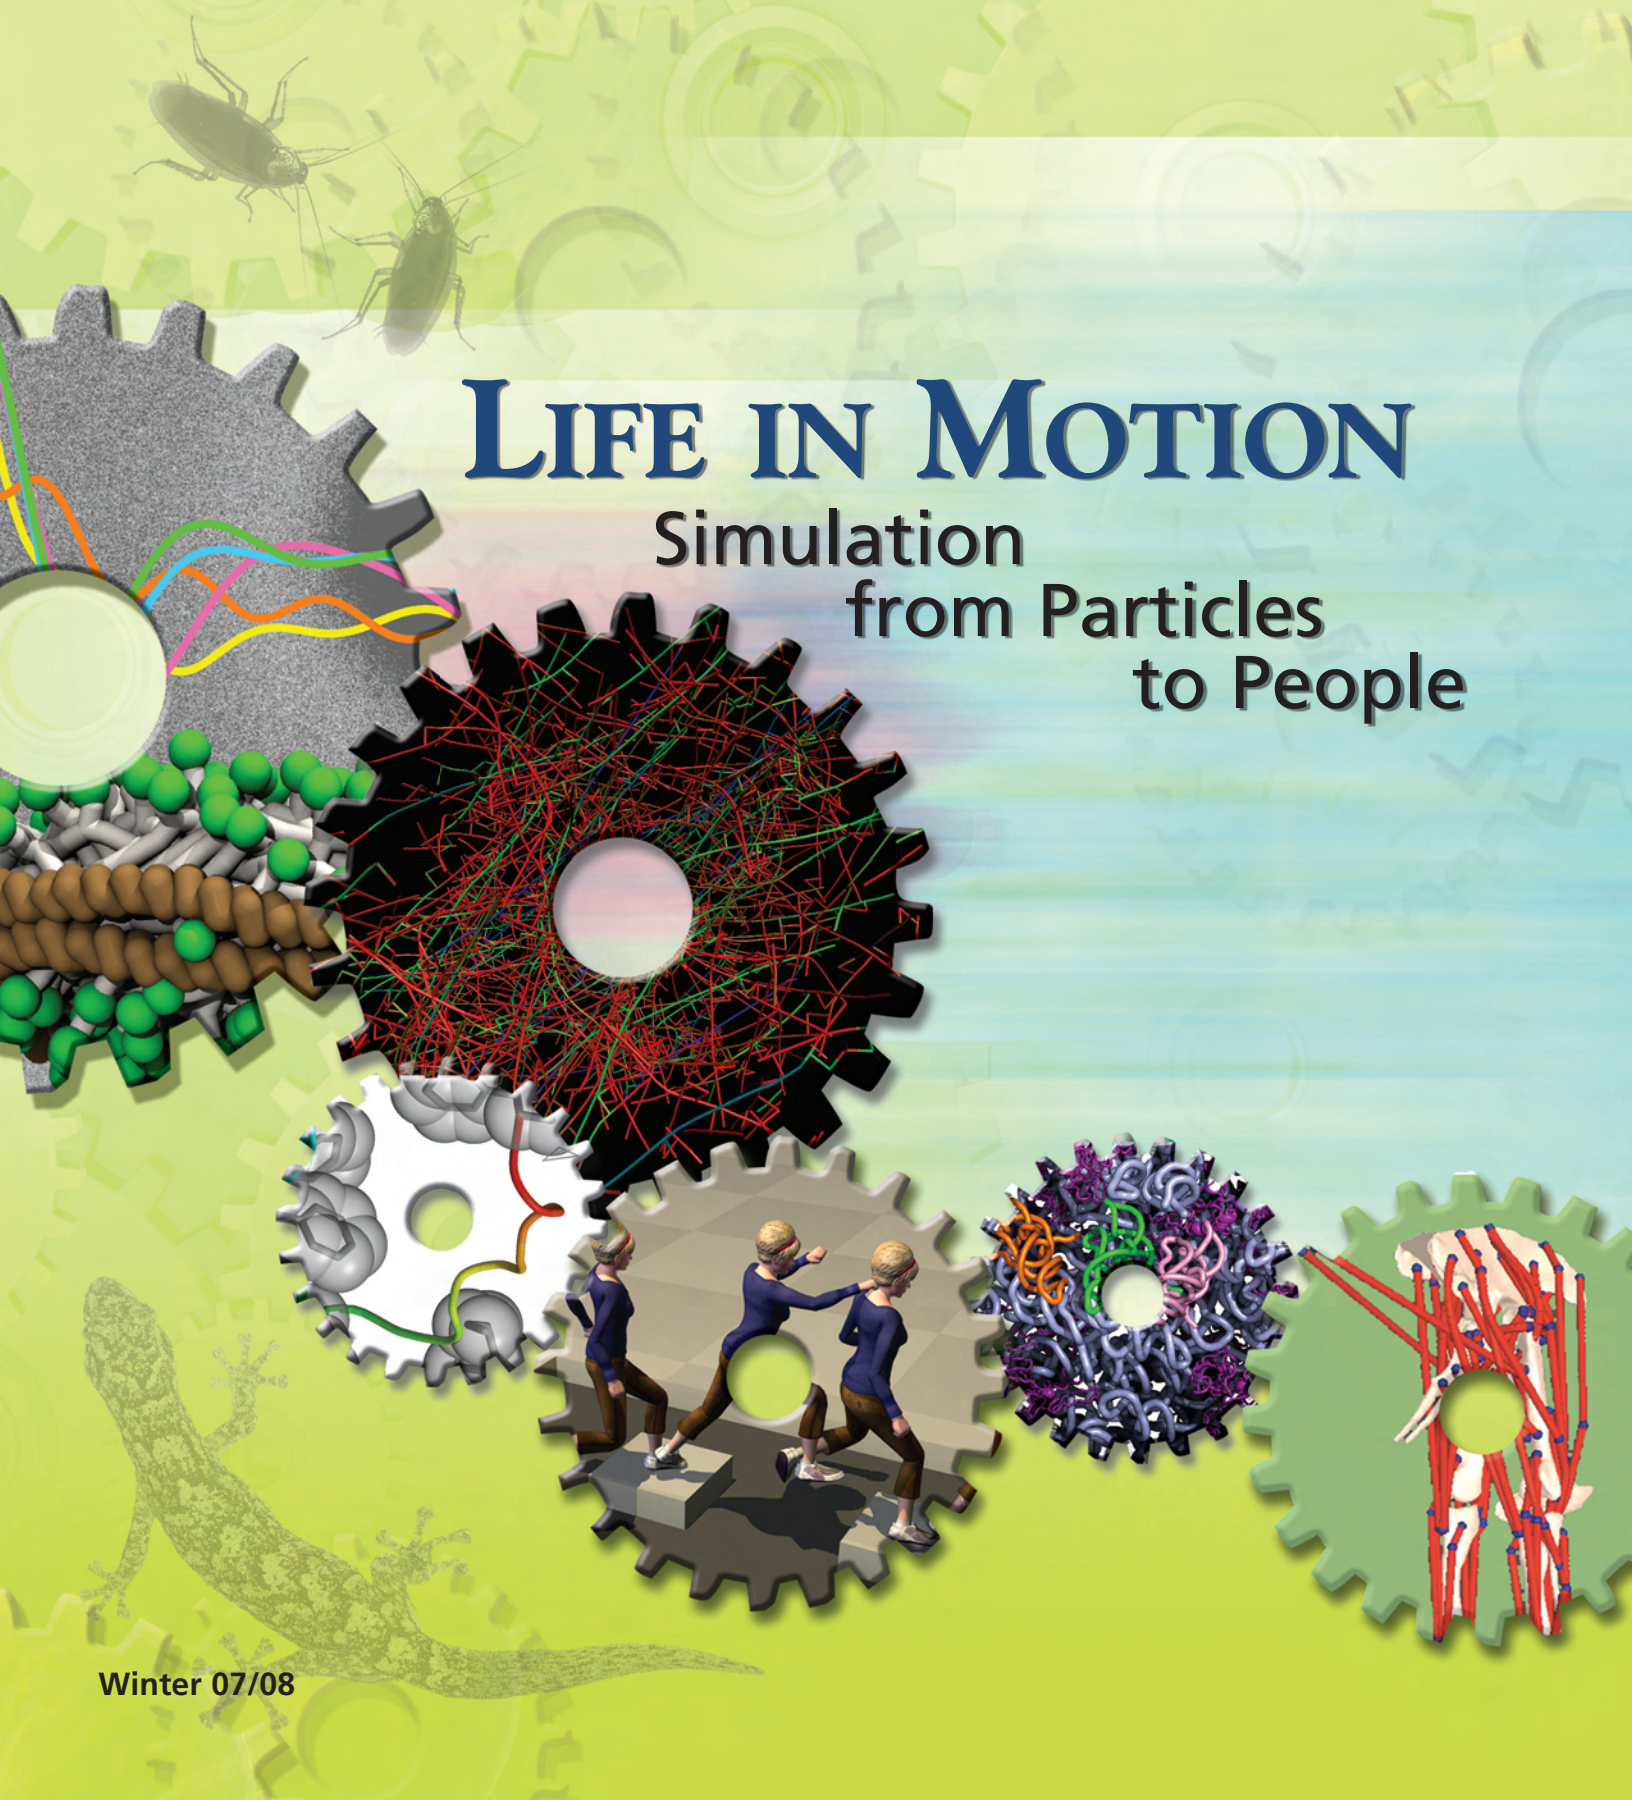

LIFE IN MOTION

Simulation from Particles to People

Winter 07/08

contents

ContentsWinter 2007/08

FEATURE

DEPARTMENTS

Cover Art by
Beth Walker of Affiliated Design

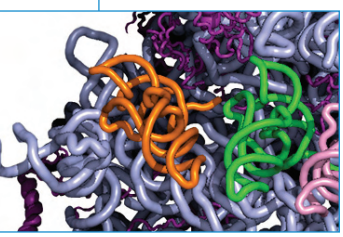

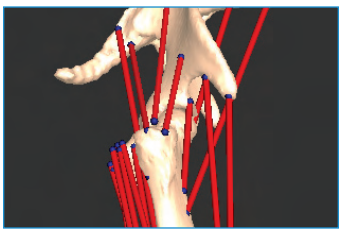

Winter 2007/08
Volume 4, Issue 1
ISSN 1557-3192

Executive Editor
David Paik, PhD

Managing Editor
Katharine Miller

Associate Editor
Joy Ku, PhD

Science Writers
Katharine Miller, John Cannon, Amber Dance, PhD, Madolyn Bowman, Alissa Poh, Rogers, PhD, Erin Digitale, PhD, Massie Santos Ballon, Hayley Rutger, Rachel Tompa, PhD, Jane Liaw, Roberta Kwok

Community Contributors
Scott Delp, PhD
Joy Ku, PhD
Alain Laederach, PhD

Layout and Design
Affiliated Design

Printing
Advanced Printing

For general inquiries, subscriptions, or letters to the editor, visit our website at www.biomedicalcomputationreview.org

Office
Biomedical Computation Review
Stanford University
318 Campus Drive
Clark Center Room S231
Stanford, CA 94305-5444

Biomedical Computation Review is published quarterly by Simbios National Center for Biomedical Computing and supported by the National Institutes of Health through the NIH Roadmap for Medical Research Grant U54 GM072970. Information on the National Centers for Biomedical Computing can be obtained from http://nihroadmap.nih.gov/bioinformatics. The NIH program and science officers for Simbios are:

Peter Lyster, PhD (NIGMS)
Jennie Larkin, PhD (NHLBI)
Jennifer Couch, PhD (NCI)
Semahat Demir, PhD (NSF)
Peter Highnam, PhD (NCRR)
Jerry Li, MD, PhD (NIGMS)
Richard Morris, PhD (NIAID)
Grace Peng, PhD (NIBIB)
David Thomassen, PhD (DOE)
Ronald J. White, PhD (NASA/USRA)
Jane Ye, PhD (NLM)
Yuan Liu, PhD (NINDS)

GuestEditorial

BY SCOTT DELP, PhD

Moon Shots in Biomedical Computation

The world changed when Neil Armstrong set foot on the moon in 1969. Humans could survive outside the earth's atmosphere! Science and engineering could achieve great things! And the nerds at the Mission Control Center in Houston were so cool. As I watched this event on TV, my brothers and I decided to order our first Heathkit, an educational electronics kit, launching us as a family of nerds.

The project of putting a man on the moon was a powerful, galvanizing force in science and engineering. The goal was challenging and clear and captivated the minds and hearts of the American public. Achieving the goal required an extremely talented and dedicated team.

Today, science needs more moon shots—projects that achieve important breakthroughs through the heroic efforts of many people. Projects that captivate the public and inspire a new generation of kids to pursue science and engineering.

What is a moon shot for biomedical computation? Fortunately, plenty of projects could fit the bill. In my own spectrum, three come to mind. The first would provide simulations that improve treatment outcomes for persons with movement disorders. Young children with cerebral palsy, for example, undergo a variety of orthopaedic and neurosurgical procedures to improve their mobility. While some experience dramatic improvements in their functional capacities, others are left with weak or dysfunctional limbs. Developing computational models that represent the neuromuscular system with sufficient accuracy to predict the outcome of these interventions and provide consistent positive results for individuals with movement disorders is a scientific and engineering challenge not unlike going to the moon. The development of these models would require collaboration among biologists, physicians, computational scientists, and bioengineers across the globe. While I was thrilled by the television images of one giant step for mankind, this could not compare to the thrill of watching a child taking his or her first steps unencumbered by disease.

A second moon shot is designing life. Almost nothing in biology is currently designed. By contrast, almost every complex product we use is designed with simulations. Dishwashers, cars, aircraft, and cell phones are all designed in software before they are implemented in physical reality. Working together, biologists, computation scientists, and design engineers could apply the same engineering capacity to design proteins, molecular machines, implantable devices, and drugs. That would be a moon shot. The ability to engineer and design biology would change the world.

Third, a collaboration among biologists, physicians, computational scientists and bioengineers could produce a digital human—a computational model of human form and function with the complexity and range of behaviors similar to a real human. A digital human would be used to study the mechanisms of disease, design biomedical devices, and predict the outcome of treatments. It could be used to teach anatomy, test drugs, and probe the basis of human behavior. A moon shot for sure.

These are just three examples that come from work in my own laboratory and from the mission of Simbios, a National Center for Biomedical Computing based at Stanford University. I encourage each of you to develop your own personal moon shots. As leaders and participants of an effort to build an infrastructure that enables biomedical computing on a broad basis, it is incumbent upon us to define clear and challenging goals that will dazzle the world. □

Photo by NASA

NewsBytes

"Sloppy" Systems Biology

Systems biologists seek to model many complex biological interactions all at once. Typically, they input tens or even hundreds of variables to produce predictions about a system—for example, how a cell might react to an environmental signal, or how an animal might respond to a drug. But, researchers have now found, many systems models are strikingly vulnerable to even small changes in the variables, according to a recent analysis of 17 such simulations.

"This pattern we see is universal," says **Ryan Gutenkunst, PhD**, who performed the research under **James Sethna, PhD**, a professor of physics at Cornell University. "It's common among all these models." The work was published in *PLoS Computational Biology* in October 2007.

Typically, modelers scan the literature or perform experiments to define the parameters of a system. But, Gutenkunst notes, such experimental data might not reflect biological reality. For example, an enzyme may function differently in a test tube than it does in a cell. And although scientists knew some models were sensitive to parameter variation, the extent of the problem was elusive.

To test how well models deal with varying parameters, Gutenkunst and his colleagues collected 17 systems biology models, including the yeast cell cycle, circadian rhythms, and others, from the literature and an online database. All 17 examples were vulnerable to producing inaccurate predictions when parameters changed only a small amount. Gutenkunst and his co-authors say this means the models are "sloppy," which doesn't necessarily mean bad. "Sloppy's a descriptive word for the fact that there's all this wiggle room," he says.

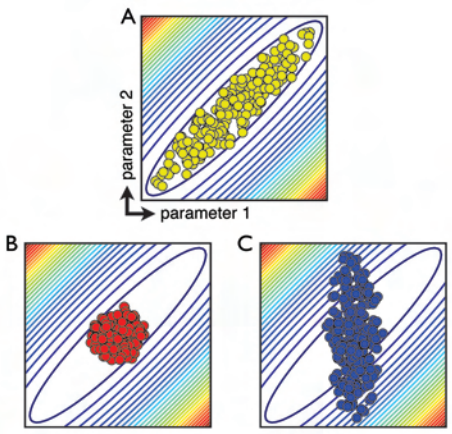

In these diagrams of model input, the best predictions result when the input parameters (colored points) are within the central ellipse. The red points in (B) represent what happens when all parameters are measured accurately. In (C), the blue points show how the predictions can become inaccurate when even one parameter is off. In (A), working backward from real data to constrain parameters results in all of the yellow points staying within the ellipse, where predictions are most trustworthy. Courtesy of Ryan Gutenkunst.

The traditional approach to modeling is akin to basic arithmetic: If every number on the left-hand side of an equation (i.e., the parameters) is known, then the answer (the prediction) is calculable. Gutenkunst and his co-authors support an alternative more like algebra: There are unknown variables on the left-hand side, but using a known answer on the right, it's possible to work backwards to define them.

"You can still get good useful predictions out of these models," Gutenkunst says. He suggests plugging in real-life information—the right-hand side of the equation—and searching for parameters that give the correct result. For example, modelers could use experimental data on how yeast grow to determine what parameters will work on the left-hand side of their cell-cycle equation. The researchers found that even if they can't define the parameters precisely, they still get useful predictions.

While the algebra approach to modeling is not new, the notion that "sloppiness" pervades biological modeling will apply to many researchers, says **Nathan Price, PhD**, a systems biologist at the University of Illinois at Urbana-Champaign. "What they argue is that it's not even very worthwhile to try to know all these parameters in advance," Price says. "It's a very broad message."

—By Amber Dance, PhD

"Sloppy is a descriptive word for the fact that there's all this wiggle room," says Ryan Gutenkunst.

Turning Therapeutic Antibodies into Better Drugs

The word "antibody" conjures images of our bodies fighting off bacteria and viruses. But because they can latch onto their targets with great precision, antibodies are also used to treat non-infectious diseases such as cancer. Researchers at Massachusetts Institute of Technology have now designed a computer algorithm that manipulates antibodies to predict which forms will bind their targets more tightly. These predictions are then confirmed in

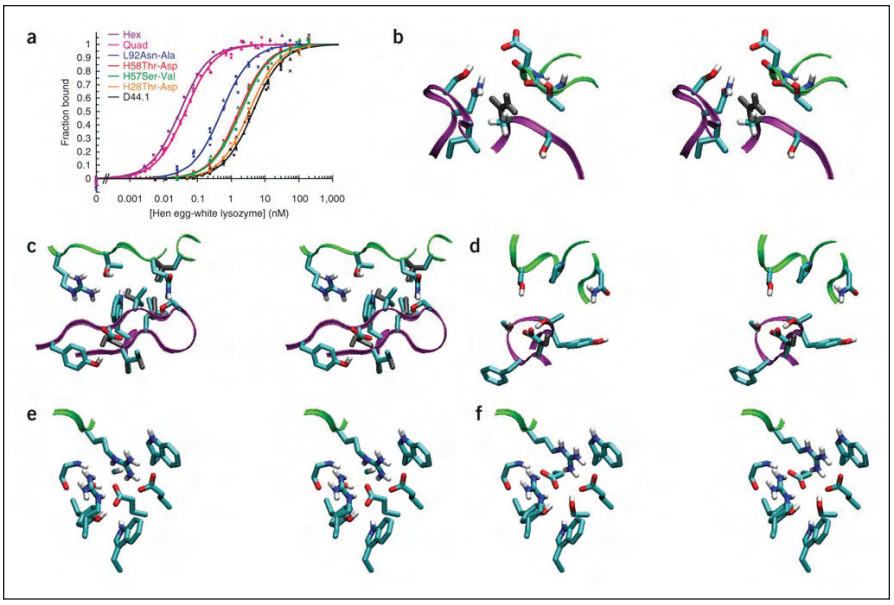

The high specificity of antibodies makes them valuable as drugs, but the conventional process of developing antibody-based drugs is tedious. MIT's new computational approach identifies all possible amino-acid changes in a particular antibody, such as D44.1 (depicted above), predicts the binding strengths of any introduced mutations (a), and models the structures of these mutations (b – f). With this approach, researchers can design a customized antibody that binds more tightly to its target. Courtesy of Bruce Tidor. Reprinted by permission from Macmillan Publishers Ltd: Nature Biotechnology 25, 1171-1176 (2007).

the laboratory. The group's work could lead to significantly improved antibody-based drug design.

"Part of the effectiveness of an antibody-based drug is related to how tightly it binds its target," explains **Bruce Tidor, PhD**, professor of biological engineering and computer science at MIT. He co-authored the work with **K. Dane Wittrup, PhD**, professor of chemical engineering and bio-engineering at MIT, and **Shaun Lippow, PhD**, the paper's lead author and a joint graduate student of both Tidor and Wittrup at the time the work was done. The research was published in the October issue of *Nature Biotechnology*.

Like all proteins, antibodies aren't rigid; they are more like Play-Doh than wooden building blocks. It doesn't take much to affect an antibody's shape. For example, substituting any of the amino acids strung together in a protein chain may alter its final folded shape markedly. Such changes, in turn, impact how strongly the antibodies bind to other molecules.

The biggest challenge in antibody-based drug design has been tweaking amino acid sequences to obtain that 'just-so' fit with the target. Traditional methods miss many possible amino-acid changes that might make the altered antibody bind more tightly. MIT's approach, combining computational structure analysis and experimental lab chemistry, may provide the missing link.

The computer algorithm works by first modeling the physical interactions that make a particular antibody latch onto its target. It then rapidly identifies all possible amino-acid substitutions for that antibody and calculates which of those changes will tighten binding. Researchers can introduce mutations that improve antibody function but might never arise naturally or with conventional techniques, and they can predict the effectiveness of these mutations.

The researchers experimentally verified their model on a drug called cetuximab (trade name Erbitux®, used to treat colorectal cancer). With guidance from the computer program, they synthesized a new version that binds 10 times more strongly to its target, a molecule called epidermal growth factor receptor. They also created a revised version of an antibody (D44.1) that is useful in laboratory experiments. It has a 140-fold improvement in binding affinity.

"This represents an interactive collaboration between calculation and experiment," says Tidor. "The ability to have tight feedback cycles between predictions and testing was essential to our success in this work."

Janna Wehrle, PhD, program director of the biophysics branch at the National Institute of General Medical Sciences, is enthusiastic about the new model. "Dr. Tidor and his team have developed a method that will allow much of the design work to be done on the computer, saving months or years in the lab," she says.

—By Alissa Poh

Protein Structure Prediction: Getting it Right

When nature folds an amino acid sequence into a protein, it usually knows that just one conformation is the right one. But when a computer tries to do the same thing, it often predicts multiple possible shapes. Now, a team of scientists at the University of Washington, led by biochemistry professor **David Baker, PhD**, have made a significant advance toward predicting which of the multiple structures is correct. They also accurately predicted a small protein's structure without relying on X-ray crystallography.

To predict a protein's structure, researchers must find the arrangement of the individual amino acids that represents the lowest energy form. It's kind of like gravity, notes Baker. "If you drop a ball on a hill, it rolls to the bottom of the hill." For proteins, that spot represents the most settled overall shape, a compact blob of amino acids linked into helices and sheets. In the past, it was hard to figure out when a predicted structure truly reached its lowest possible energy, not just an intermediate step. "If you

"I crunch for Rosetta because I believe in this project wholeheartedly," says blogger Antony Magnus.

drop a ball on a bumpy landscape, it may get stuck in a [higher] valley," Baker says.

For a number of years, the Baker team's primary tool for predicting protein structures has been Rosetta@Home, a program that relies on a staggering amount of computing power. "We employ the computers of about 150,000 volunteers," Baker says. Volunteers install Rosetta@Home on their computers. It runs like a screen saver while the computer is otherwise idle. The program calculates many possible structures for an amino-acid chain and sends promising structures to the researchers. A central computer then searches for the lowest-energy structure, in which the chain curls up most comfortably.

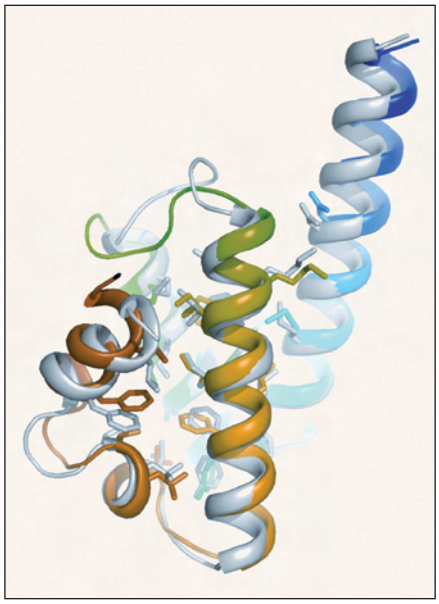

Using Rosetta@Home, a program that runs on the personal computers of 150,000 volunteers worldwide, David Baker's team predicted the structure of a 112-amino-acid protein from scratch. The predicted structure (gray) closely mimics the true protein structure (in color).

His team's new research, published online in *Nature* on October 14, 2007, describes a major refinement to Rosetta@Home that searches for a way out of "energy valleys." The team fine-tuned how Rosetta analyzes the toughest protein sections. If the program consistently predicts the same folded shape, the answer is probably correct. But when Rosetta churns out many different solutions, the program now recalculates those error-prone regions in search of the lowest possible energies—and more robust final shapes.

The refined method makes it easier to get useful data from traditional protein-structure experiments, in which researchers blast X-rays at protein crystals. Baker's lab also used the method to predict an accurate structure for a small protein (112 amino acids) with no X-ray data, an achievement noted in a *Nature* commentary as "a real breakthrough."

The new research is "a significant milestone in the development of methods to model protein structure from amino-acid sequence," comments **John Moult, D. Phil.**, a professor of computational biology and biophysics at the Center for Advanced Research in Biotechnology in Maryland.

Rosetta@Home's clan around the world savors the success. As volunteer **Antony Magnus** wrote in an online message board: "I crunch for Rosetta because I believe in this project wholeheartedly." Volunteers interested in participating in Rosetta@home can sign up at boinc.bakerlab.org/rosetta.

—By Erin Digitale, PhD

Extinct Sabercat Brought to Life

Wildlife biologists can watch a lion stalk its prey, but paleontologists must examine fossils to understand how the extinct saber-toothed cat hunted. Researchers now have modeled an American sabercat's skull with software designed for stress testing in engineering, building the highest resolution vertebrate animal model to date. They found that the sabercat's massive teeth belied a surprisingly weak bite.

On a computer, "you can crash test a biological design," says **Colin McHenry**, a doctoral candidate at the University of Newcastle, and lead author of the work. His team built a virtual sabercat skull that could display the effects of stress down to cubic millimeter resolution. Stress resistance indicates how hard the cat could bite and which muscles contributed the most force. The study appeared in the October 9, 2007 issue of the *Proceedings of the National Academy of Sciences*.

Despite more than 150 years spent studying sabercats, scientists have yet to agree on the animal's biting power and the relative importance of head and neck muscles. In recent years, researchers have turned to computer simulations to reconstruct the musculature of extinct animals. They use the finite element method (FEM), a system originally designed to test aeronautical designs under stress. Until now, FEM studies featured animal skulls modeled as though bone has the same strength and density throughout—which it doesn't. And they didn't account for moving jaws.

To create a more lifelike simulation, McHenry's team used a standard medical imaging technology, computed tomography, to build high-resolution FEM models of sabercat and (for comparison) lion skulls. The individual elements that make up this 3D model mimicked realistic bits of bone with different strengths. The team then added musculature, estimating the sabercat's muscle sizes and strengths from the skull's geometry. After subjecting both

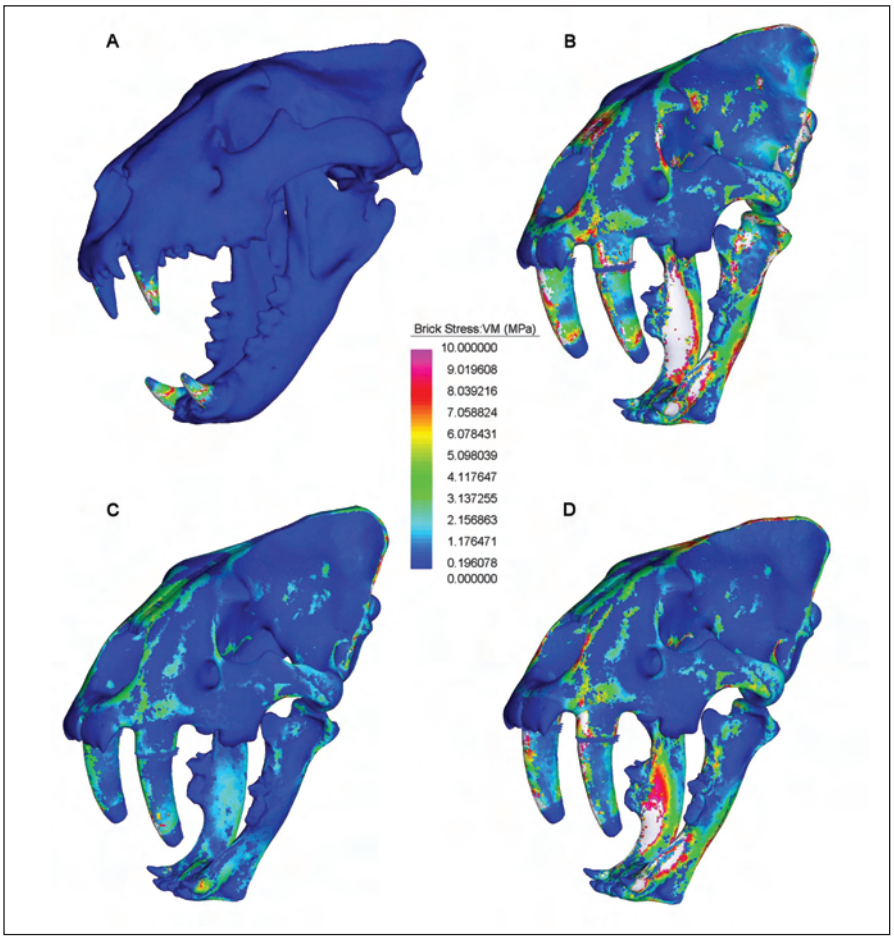

A high-resolution model of a lion's skull (A) shows little stress compared to a model of an American sabercat skull (B) when researchers apply lateral forces to simulate thrashing prey. Twisting forces (C) and forces pulling forward on the canine teeth (D) also illustrate the stresses a sabercat might have encountered while killing animals. Courtesy of Colin McHenry.

models (sabercat and lion) to the forces of struggling prey and the pull of the animal's own muscles, they mapped the resulting stresses.

The sabercat skull generally handled forces poorly, while the lion skull took them like a tank. The researchers concluded that the sabercat didn't land powerful bites, and that the jaw muscles may have required help from the neck muscles to puncture prey. These results support existing arguments that sabercats killed with piercing canine tooth bites, but there was still debate about the bite force, says McHenry. The sabercat probably bit one-third as hard as a comparably sized lion, the team concluded.

The next step is to account for the way bone responds to pressure from different directions, a method called anisotropic modeling, says **Lior Horesh, PhD**, a post-doctoral research fellow in Emory University's department of mathematics and computer science. Horesh calls the team's research "one good step forward."

McHenry and his colleagues soon will apply FEM modeling to biomedical questions, including mechanical evaluation of surgical planning procedures and stress-testing of prosthetic devices. "I think the medical community can learn a lot from paleontologists and biologists," he says.

—By Hayley Rutger

Center of Mass Controls Balance

Bumped from behind, a person may step forward to avoid falling. Perhaps her arms fly out as well. To the untrained eye, these movements seem like the result of the brain controlling individual nerve and muscle reflexes. Yet an elegant new model of balance control suggests the brain only cares about one thing: the body's center of mass. This possibility, modeled for the first time, could help rehabilitation experts design better treatments to suit the specific needs of each balance-impaired patient.

"People had theories about the center of mass being important, but they hadn't actually demonstrated in a causal sense that it was critical," says **Lena Ting, PhD**, assistant professor of biomedical engineering at Georgia Institute of Technology and Emory University and co-author of the work published in the October 2007 issue of *Nature Neuroscience*. "We've shown that the nervous system controls the arms and legs to regulate center of mass motion."

Neuroscientists have tested a variety of hypotheses such as whether balance originates from motions of the head or the ankle. But these hypotheses have not consistently predicted which muscles would spring into action when a person loses balance. Ting's previous experiments found that the only way to foretell muscle reaction accurately was to monitor the direction of the fall, not individual joint angles. This suggested that the body's reflexes during a fall involve a higher level of control: If the center of mass is off-kilter, the nervous system will act to bring it back to balance.

To explore this idea in action, the researchers placed cats on a moving platform that made them lose their balance. The team also induced sensory damage in the cats that triggered balance-control problems. A computer simulation created by the researchers accurately predicted the reactions of the cats' muscles based on the motions of their centers of mass.

In addition, the cats with sensory damage regained balance within a few days. They used different sensory pathways to do the same balancing tasks, resulting in unique patterns of muscle activity. While these muscular adjustments were clinically "abnormal," they were close to optimum for the balancing task at hand. This result should earn notice from balance rehabilitation professionals, Ting says; they now have an accurate, unique goal toward which they can aim each patient's rehabilitation efforts.

Biomedical engineer Lena Ting (right) prepares to measure how a volunteer's muscles react when her balance is disturbed by a moving platform. Courtesy of Georgia Institute of Technology.

Ting's group has provided an attractive, simple model of posture control, says **Fay Horak, PhD**, a senior scientist at the Neurological Sciences Institute of the Oregon Health & Science University. "The big implication is that something this complicated, that involves many, many joints and muscles, could be controlled by the nervous system regulating a single parameter," she says.

Horak believes researchers need more data before applying Ting's results to humans with balance disorders. Ting concurs, noting that her team has started using the moving platform to test human balance reactions.

—By Jane Liaw

A Model of Epstein—Barr Virus

During our lives, most of us will come in contact with the Epstein-Barr virus, commonly known as that bane of teenagers, infectious mononucleosis. Now, a new simulation mimics the virus's infection cycle on the tonsils, shedding some light on how the infection spreads.

"The actual biology is so complicated," says **David Thorley-Lawson, PhD**, professor of pathology at Tufts University who co-leads the project with **Karen Duca, PhD**, a biophysicist formerly at the Virginia Bioinformatics Institute (now at Kwame Nkrumah University, Ghana). "But what we got out of the simulation looks remarkably like a real infection." They and others developed the Pathogen Simulation (PathSim) model published in the October 2007 issue of *PLoS Pathogens*.

The potential benefits of modeling infection seem endless. Scientists can raise the viral load in ways that would be unethical in humans; and insight into the dynamics of infection could lead to novel therapies. But the question remains: Do the models truly replicate how infection spreads in the body?

Until now, most computer modeling only reproduced general properties of the immune system, or involved the use of differential equations to provide more specific insights, such as with some HIV models. Using a well-studied virus like Epstein-Barr as a guide, Thorley-Lawson and colleagues believed they could create a model that rivaled the sophistication of HIV models while remaining comprehensible to nonspecialists.

"One of the main goals was to have models that biologists could look at and say, 'Oh, I get that,'" says Thorley-Lawson. To accomplish this goal, the researchers created a ring of tonsils—the point of attack for Epstein-Barr virus—on a virtual grid.

During the simulation, the virus infected cells at about the same rate it does in a person. "This suggests we're not missing huge parts of the biology," says **Michael Shapiro, PhD**, co-lead author and lecturer in pathology at Tufts.

Already, the model has helped explain a clinical puzzle. Several years ago, Thorley-Lawson and his team found that when Epstein-Barr causes infection *in vivo*, only 0.5 to 1.0 percent of the host's B cells—key sentinels of the immune system—replicate the virus. At the time, the researchers didn't know why that replication rate was so low.

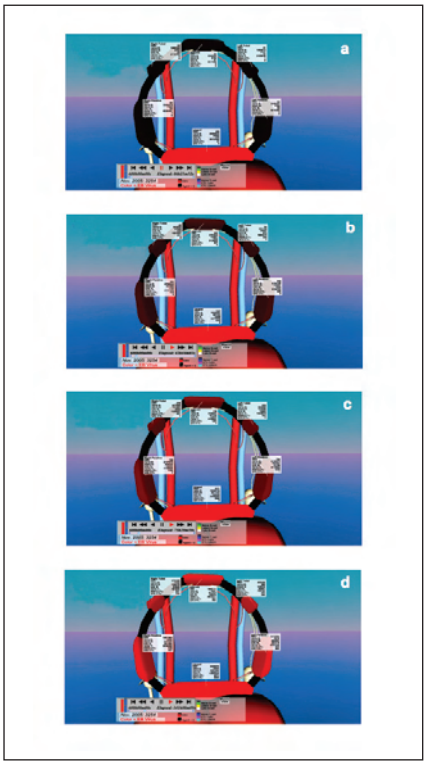

The PathSim model predicts that infection begins on the lingual tonsil (at the base of the ring) and spreads evenly (increasing red color) to the other tonsils through the bloodstream, instead of spreading through the ring directly from one tonsil to the next. Courtesy of David Thorley-Lawson.

When an almost identical proportion of the model's B cells were active at the same stage in PathSim, Thorley-Lawson and his colleagues increased the number of B cells replicating the virtual virus to see what would happen. This tinkering "killed" the virus's host by overwhelming it with infected B cells. According to PathSim, the virus has honed in on the speediest possible replication rate while still keeping its host alive—thus ensuring its further spread.

"This is a very nice first step," says **Alan Perelson, PhD**, a biophysicist at Los Alamos National Laboratory. He acknowledges that as PathSim becomes more complex, it will rely on more biological assumptions. Still, he says, "The model looks like it's driving some new experimentation."

—By John Cannon

A Digital Human Could Advance Medicine

Science and medicine have fractured the human body into pieces: the cardiovascular system, the immune system, the endocrine system. Now a European initiative seeks to put the jigsaw puzzle back together by developing a computer model of a complete human being. The Virtual Physiological Human (VPH) would encompass all the knowledge we've gathered, from genetic interactions to systems biology, into one integrated digital package.

"If you thought the genome project was big work, this is probably a million times more complicated," says Marco Viceconti of the Virtual Physiological Human.

"If you thought the genome project was big work, this is probably a million times more complicated," says **Marco Viceconti, PhD**, scientific officer of the Strategy for The EuroPhysiome (STEP), a coalition of leaders from research, industry, and clinical practice who hope to create the virtual human. "This is not something you will ever finish."

In recent years other groups have begun digital modeling of various human processes, including two other worldwide projects to assemble our physiome—a complete description of human physiology. To prevent fragmentation, all projects are communicating under an umbrella organization called the World Integrative Research Initiative. They will share technology and computer infrastructure, and will agree on common terminology.

Researchers in the VPH initiative, funded by the European Commission, plan to build the virtual human piecemeal by linking each model as it's created. A brain aneurysm model is already under way. If clinicians could predict which aneurysms are unlikely to rupture, they might avoid unnecessary brain surgeries. Scientists working on the project (called @neurIST) are gathering genetic and metabolic data from patients with aneurysms, which they will feed into a computer model to develop a predictive algorithm.

Candidates for other initial projects include disease models for diabetes and osteoporosis. The European Commission is running an evaluation process to select projects for funding starting in early 2008.

It may take at least a decade before a complete virtual human exists, but Viceconti hopes some pioneering applications such as the aneurysm model will soon deliver benefits. An early goal is reducing the costs and risks of drug development by first testing drugs on a virtual patient to gauge harmful side effects. Eventually, physicians could use the virtual human for better diagnosis and treatment by programming it with a patient's specific data, yielding a unique assessment of how certain drugs might affect him or her.

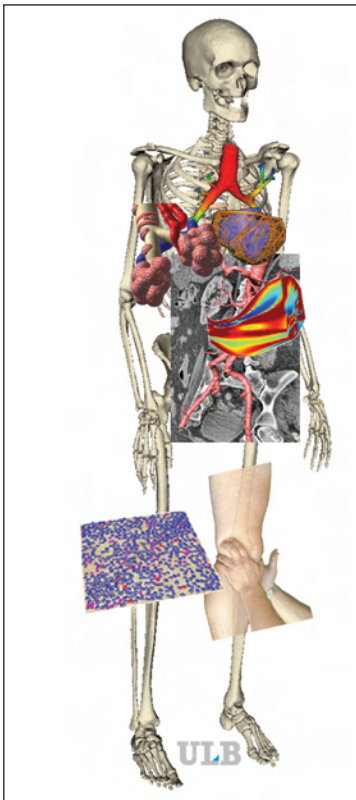

The Virtual Physiological Human will integrate digital modeling at all levels—genetic, cellular, tissues, organs, and systems—into one complete package that will be useful in medicine and research. Courtesy of Serge Van Sint Jan, Université Libre de Bruxelles.

"There's a very strong focus in the EuroPhysiome on modeling for clinical applications," says **Peter Hunter, PhD**, director of the Bioengineering Institute at the University of Auckland and representative of the IUPS Physiome Project, one of the other international physiome initiatives. He sees the EuroPhysiome as complementary to his project.

A strong spirit of international collaboration will help the EuroPhysiome succeed, says Viceconti. "This is definitely a team science exercise."

—By Madolyn Bowman Rogers, PhD

Virtual Genomic Scans with Real Data

Trying to find the genetic causes of a human disease requires lots of data. These days, researchers scan the genomes of people who do and don't have a particular disease and look for genome-wide associations between a particular disease and a gene or genes. But they'd like to know if their findings are statistically valid. Moreover, the variety of disease models currently in use have led to debates over which work best. Now, researchers have developed a new tool that they hope will help resolve these issues and will also work with any genotyping platform in use. Their software generates large simulated populations using present-day genetic information from specific populations.

"The main challenge is working out how you draw from real data to mimic what you expect to happen in a disease model situation," says **Fred Wright, PhD**, a biostatistician at the University of North Carolina and senior author of a study published online in September 2007 in the journal *Bioinformatics*. "Because of that, we developed a method that's simple, almost dumb, in the way it approaches it."

Current statistical simulations either work backward to generate genetic "histories" that might give rise to present-day forms, or else they go forward, simulating genetic data from the distant past until the present day.

To present a more accurate simulation grounded in real data, Wright's method—called HAP-SAMPLE—now offers a third option: using data from a real population to generate a large sample set against which genes of interest can be checked. Because the data already contain realistic historic mutations, there's no need to let the population evolve (developing new mutations) over time. Instead, HAP-SAMPLE generates simulated populations solely by meiosis and its associated crossovers—it's that simple.

The real genetic data is supplied by HapMap, an international database that catalogs 10 million common genetic variations (single nucleotide polymorphisms or "SNPs") within three populations—Caucasian European, Chinese /Japanese, and Nigerian.

HAP-SAMPLE is potentially valuable to researchers who have identified a possible gene-disease association and want to see how it would play out in a larger population. For example, would the same SNP still be a significant contributor to the disease of interest in a larger population? By comparing the resulting simulated data against known SNPs, they can figure out how good their statistical methods are.

HAP-SAMPLE is potentially valuable to researchers who have identified a possible gene-disease association and want to see how it would play out in a larger population.

"HAP-SAMPLE is great because it takes real data as the template for the simulation," says **Marylyn Ritchie, PhD**, a computational geneticist at Vanderbilt University, whose lab developed a complementary simulation tool. Still, she adds, HAP-SAMPLE's usefulness is limited by HapMap's small chromosome pool: Fewer than 400 people represent the three populations. For some researchers, having a real data template might not overcome the problem of limited population size, Ritchie cautions.

"What they're asking for is just a broader population base," Wright responds. His team does plan to augment HAP-SAMPLE soon with updates from other genetic databases.

—By Massie Santos Ballon

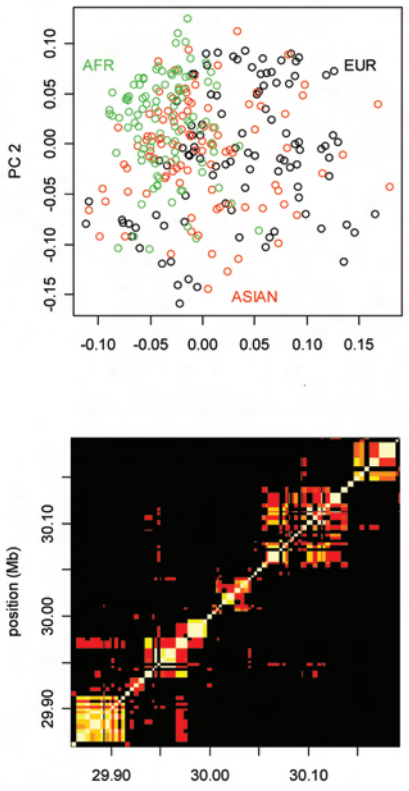

HAP-SAMPLE simulations accurately reflect population ancestry and how often different SNPs are inherited together. At the top, we see that for the disease gene in question, three distinct populations (AFR= Nigerian samples; EUR= Caucasian European samples, ASIAN= Chinese/Japanese samples) are only mildly different from one another. The heatmap at bottom plots SNPs against one another based on their chromosomal positions. Areas of brightness (white is strongest) indicate SNPs that are likely to be co-inherited in the Caucasian European data. Courtesy of Fred Wright.

Homing in on the Minimum Genome

Scientists have long wondered how many genes are necessary to support life. This knowledge could be used to construct new forms of artificial life to efficiently produce better biofuels or drugs.

Now, computer scientists are using hypothetical synthetic bacteria to screen out inessential genes as a way to home in on the "minimum genome." The remnants, they hope, should make good candidates for synthesizing artificial organisms and reduce the number of costly experiments required to achieve that goal.

"If your hypothetical organism does not survive the simulation, the chances are high that it would not survive in reality," says **Roberto Marangoni, PhD**, professor of bioinformatics at the University of Pisa and senior author of the study in the September issue of *PLoS Computational Biology*.

Scientists plan to build the minimum genome by culling unnecessary genes from *Mycoplasma genitalium*, a bacterium with one of the smallest known genomes. At just 521 genes, *Mycoplasma's* genome is about one fiftieth the size of the human genome. In an earlier attempt to find an "essential" set of genetic instructions (published in 2006), **J. Craig Venter, PhD**, of the J. Craig Venter Institute, and his colleagues shaved that number from 521 to 382 by disrupting each gene one at a time. They excluded from their hypothetical minimum genome all the genes whose disruption did not kill the bacteria.

Eventually, if lab scientists try to build artificial life from scratch, testing potential minimum genomes would be a time-consuming and expensive trial-and-error process. Researchers will need a way to increase the chances of hitting the right set of genes on an early try, Marangoni says.

To address that problem, Marangoni and his team created a computer simulation to test the viability of theoretical bacteria with a variety of possible minimum genomes. They gathered all chemical reactions known to take place inside *Mycoplasma* and ran recurring simulations of all of these reactions, assuming different sets of genes were present. The team looked for genomes that, over the course of many repeated reactions, produced a life-friendly balance of the reaction products in the cell. The simulations of some virtual creatures resulted in wildly fluctuating chemical levels, or levels that bottomed out almost immediately—conditions that would not sustain actual life.

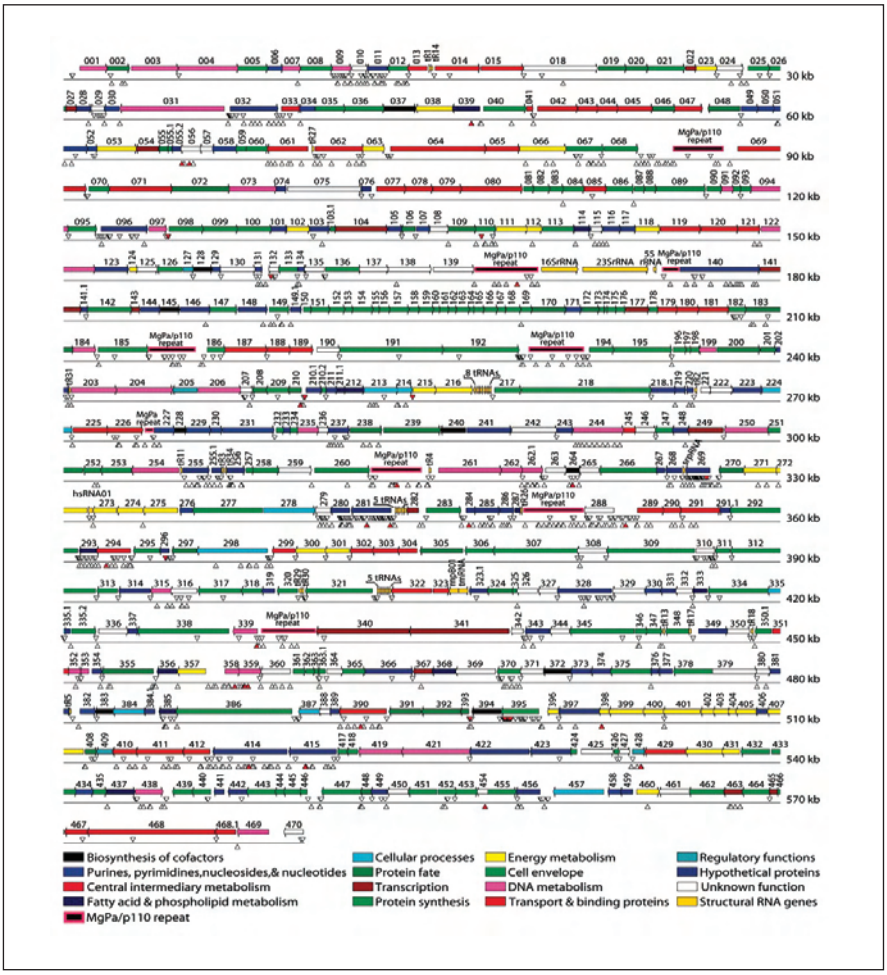

Scientists hope to create synthetic organisms using the minimal gene set possible to sustain life. They plan to base this minimum genome on that of the bacteria* Mycoplasma genitalium*, which has a paltry set of 521 genes. In this graphic, each of* Mycoplasma*'s genes is a colored bar. Courtesy of Hamilton Smith and John Glass, J. Craig Venter Institute.

Some previously proposed minimum genomes failed this test. These creatures' energy supplies went to zero very quickly, Marangoni says, which is probably what "killed" them.

Marangoni's work is promising, but not the final word, says **Arcady Mushegian, PhD**, director of bioinformatics at the Stowers Institute for Medical Research, "There will have to be more computer estimates of genes and how they fit together in the metabolic puzzle," he says. "And of course ultimately the actual organism should be engineered."

—By Rachel Tompa, PhD

Simulating a Scaffold for Bone Growth

Designing a scaffold, the internal structure that helps patients regenerate bone, is a delicate balancing act. The scaffold must be strong enough to protect the injury, porous enough to allow nutrients to pass through, and fast-dissolving enough to make room for new tissue. Now, using a 3-D computer model, scientists have simulated stem cells growing within a scaffold to predict which combination of these properties will produce the most bone.

"It's the first 3-D computational work that takes account of stem cells" in scaffold design, says senior author **Patrick J. Prendergast, PhD**, a professor of bioengineering at the University of Dublin, Trinity College.

Patients rely on scaffolds to support bone regeneration after surgeries such as bone grafts, cartilage repair, or tumor removal that requires bone to be cut away. Most scaffolds are either made of gels, which tend to be weak, or stiffer materials such as coral, which may not completely dissolve. In the past, scientists evaluated scaffold materials by testing them on animals. At the same time, computational biologists devised algorithms predicting how a patient's stem cells might differentiate into new types of tissue during healing. But until now, no one had simulated the scaffold alongside the stem cells as a way of improving scaffold design.

Prendergast's group created a 3-D computer lattice model of a scaffold, then planted "seeds" inside the lattice to represent the patient's stem cells. In their simulations, the cells multiplied, spread, and eventually transformed into bone, cartilage, or connective tissue depending on the strain and fluid pressure affecting each cell.

"It's the first 3-D computational work that takes account of stem cells" in scaffold design, says Patrick J. Prendergast.

Meanwhile, the program tracked the progress of the scaffold as it slowly dissolved and became more porous, clearing room for new tissue. The scientists tried various combinations of scaffold properties and tested the system under high and low load-bearing conditions to simulate injuries in different parts of the body. A leg bone, for instance, bears more load than an arm bone and might heal differently.

The researchers found that the scaffold only works if it has the right balance of pore size and disintegration rate. If both are too high, "It won't be long before the whole thing dissolves away," says Prendergast. They also found that the load on the area changes how scaffolds perform, suggesting that scaffold designers should tailor their materials for specific patients and body parts. The work appears in the December 2007 issue of *Biomaterials*.

Studying the interplay between cells and synthetic materials is promising because most people focus on only one, says **Christopher Jacobs, PhD**, director of the Cell and Molecular Biomechanics Laboratory at Stanford University. "I think that's a very creative concept," he says. The work needs to be verified with further experiments, says Jacobs, but could potentially direct the design of better scaffolds that both offer enough support and dissolve completely into the body.

Toward that end, Prendergast and his colleagues plan to simulate blood vessels growing in scaffolds, which can affect bone regeneration.

—By Roberta Kwok □

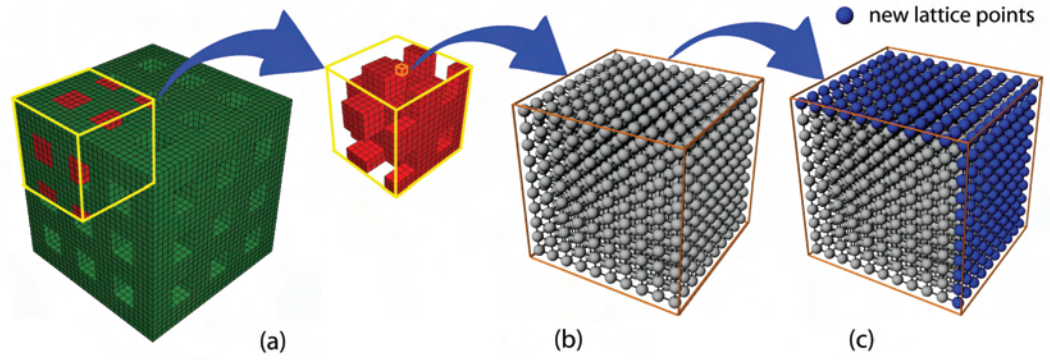

Granulation tissue (red), a mixture of tissue matrix and cells that develops early in wound healing, fills a bone regeneration scaffold (green) that's 50 percent porous. Zooming into small portions of the tissue, we see a lattice (grey) containing stem cells, which multiply and develop into new tissue (blue) that replaces the dissolving scaffold. Courtesy of Damien P. Byrne. This article was published in Biomaterials, Volume 28, Damien P. Byrne, Damien Lacroix, Josep A. Plannel, Daniel J. Kelly, Patrick J. Prendergast, "Simulation of tissue differentiation in a scaffold as a function of porosity, Young's modulus and dissolution rate: Application of mechanobiological models in tissue engineering," p. 5544-5554, Copyright Elsevier (2007).

BY JOY KU, PhD

BCATS: Not Your Usual Biomedical Computation Conference

Outwardly, the Biomedical Computation at Stanford (BCATS) conference resembles other academic conferences: Researchers converge to hear about the latest developments in their field and to exchange ideas with colleagues. But behind the scenes at BCATS, you find an unusual situation—students, and not faculty, are the ones in charge.

For the last eight years, students have organized this annual conference where Stanford University students and post-doctoral fellows share their latest research in the field of biomedical computation. The quality and breadth of the research represented at BCATS draw hundreds of individuals from across the campus and the community. And the latest BCATS, held on October 27, 2007, at Stanford University was no exception.

Leighton Read, MD, a partner in the life sciences group at Alloy Ventures, says, "BCATS is one of the highest quality one-day conferences I can think of and it's because it's student-run." Alloy Ventures, a venture capitalist firm, has supported BCATS every year since its inception in 2000. It's not just the quality that attracts sponsors, though. BCATS "touches on everything we do," emphasizes Read.

Scott Delp, PhD, a professor of bioengineering and mechanical engineering at Stanford University and one of the principal investigators of Simbios, another organization that sponsors BCATS, agrees. "Simbios is one slice of biocomputation at Stanford," Delp points out, "but BCATS is the whole pie. I think it's really important not to lose that breadth."

This year BCATS had ten student speakers and 51 poster presenters. And the research topics spanned the field: prediction of cancer genotypes from imaging data; automatic generation of machine-readable summaries of biomedical literature; blood velocity detection with a new ultrasound transducer; simulation of bone growth in tennis players.

Though working on seemingly unrelated problems, the students share a general interest in biomedical computation. And BCATS brings them together to discover their commonalities. For example, **Karen Sachs, PhD**, a post-doctoral student at Stanford and this year's winner of the BCATS Best Talk Award, is a computational biologist who works in immunology. "The types of interactions I have with computational biologists are very different from those I have with immunologists and that's very valuable to me," she says.

For the student organizers—**Annie Chiang, PhD, Yael Garten, Jen Hicks, Marc Schaub**, and **Marina Sirota**—the conference was valuable in ways they hadn't anticipated. They learned firsthand about the peer review process and managed a large budget. They also became much more familiar with the biomedical computational research going on at Stanford.

And then there's that euphoric feeling that comes from creating something that has impact. Chiang says there is a "sense of pride to have brought forth all these interactions and collaborations." □

WANT TO FIND OUT MORE ABOUT BCATS?

The BCATS website (http://bcats.stanford.edu) lists abstracts and information from all previous BCATS conferences, and provides information about next year's event. If you are interested in helping next year, e-mail the organizers at bcats-2007-organizers@lists.stanford.edu.

BCATS WINNERS

Five individuals received the Outstanding Poster Award this year: **Gilwoo Choi** (abdominal aortic 3D deformations); **Rebecca Taylor** (bone growth modeling in tennis players); **Gennadiy Chuyeshov** (stereo imagery for guidewire localization during endovascular interventions); **Melinda J. Cromie** (effects of posterior cruciate ligament removal in total knee arthroplasty); **Aaron S. Wang** (image-based models of blood flow in the human upper extremity arteries).

You can check out their posters, along with Karen Sachs' award-winning presentation on "Learning Signaling Pathway Structures from Single Cell Measurements of Network Subsets," at http://biomedicalcomputationreview.org/4/1/posters.html.

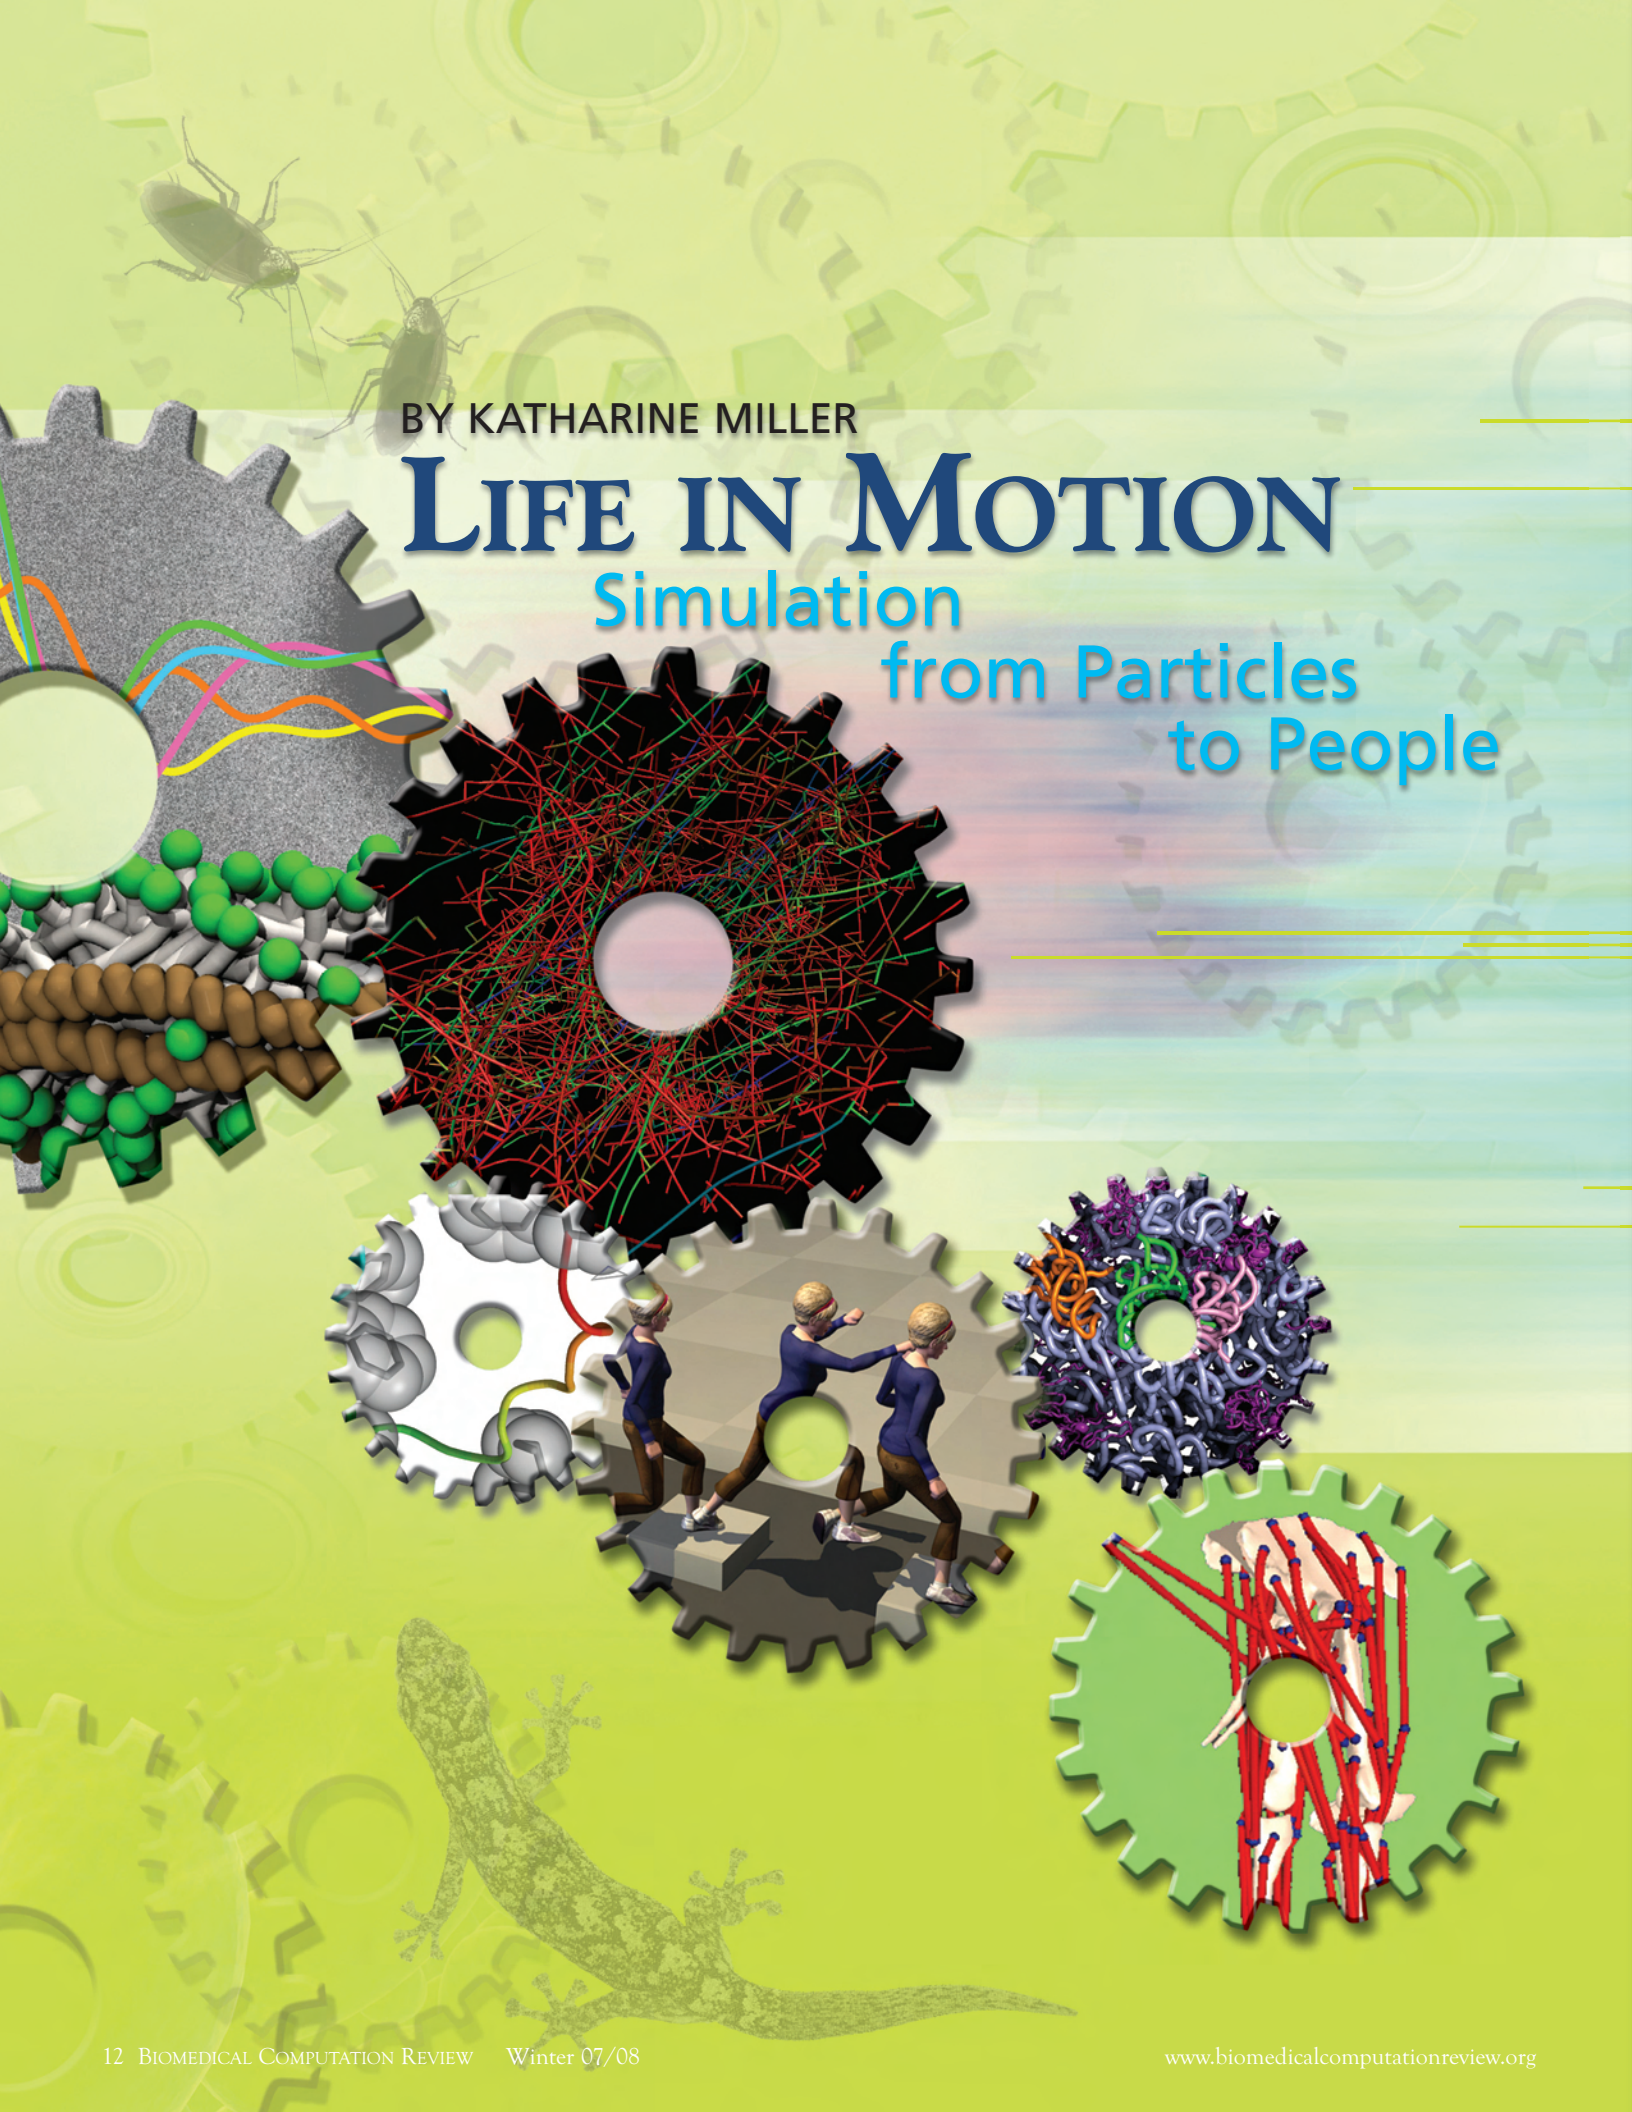

BY KATHARINE MILLER

Life in Motion

Simulation from Particles to People

From atoms and molecules to insects, dinosaurs, and humans, computational researchers are finding that much of life can be understood in mechanical terms. Indeed, the machines of life are well-tuned.

"Because nature has evolved forms that naturally have the desired functions, you don't have to bend over backwards to steer, control or coerce the structures to do their jobs," says **Russ Altman, MD, PhD**, chair of the bioengineering department at Stanford University. "They do them naturally." And that's true at all scales, as computational researchers are discovering.

At the most basic level, charged atoms push and pull on one another to control the inner workings of every living thing. Cellular machines called ribosomes use ratchets and springs to translate coded messages into the workhorses of the cell: proteins. And small movements made by these proteins act as cellular signals that give directionality and function to developing tissues. Combinations of tissues then produce appendages designed to carry entire organisms, including humans, through the natural environment.

In a feedback loop with laboratory experiments, computational simulations of life in motion at every scale—molecular, cellular, tissue-level, and whole organism—are boosting our understanding of the role mechanics plays in controlling life. Such simulations were the focus of the "Life in Motion" symposium sponsored by Stanford's Simbios Center and BioX Program this past October. That symposium served as the foundation for the ten stories presented here.

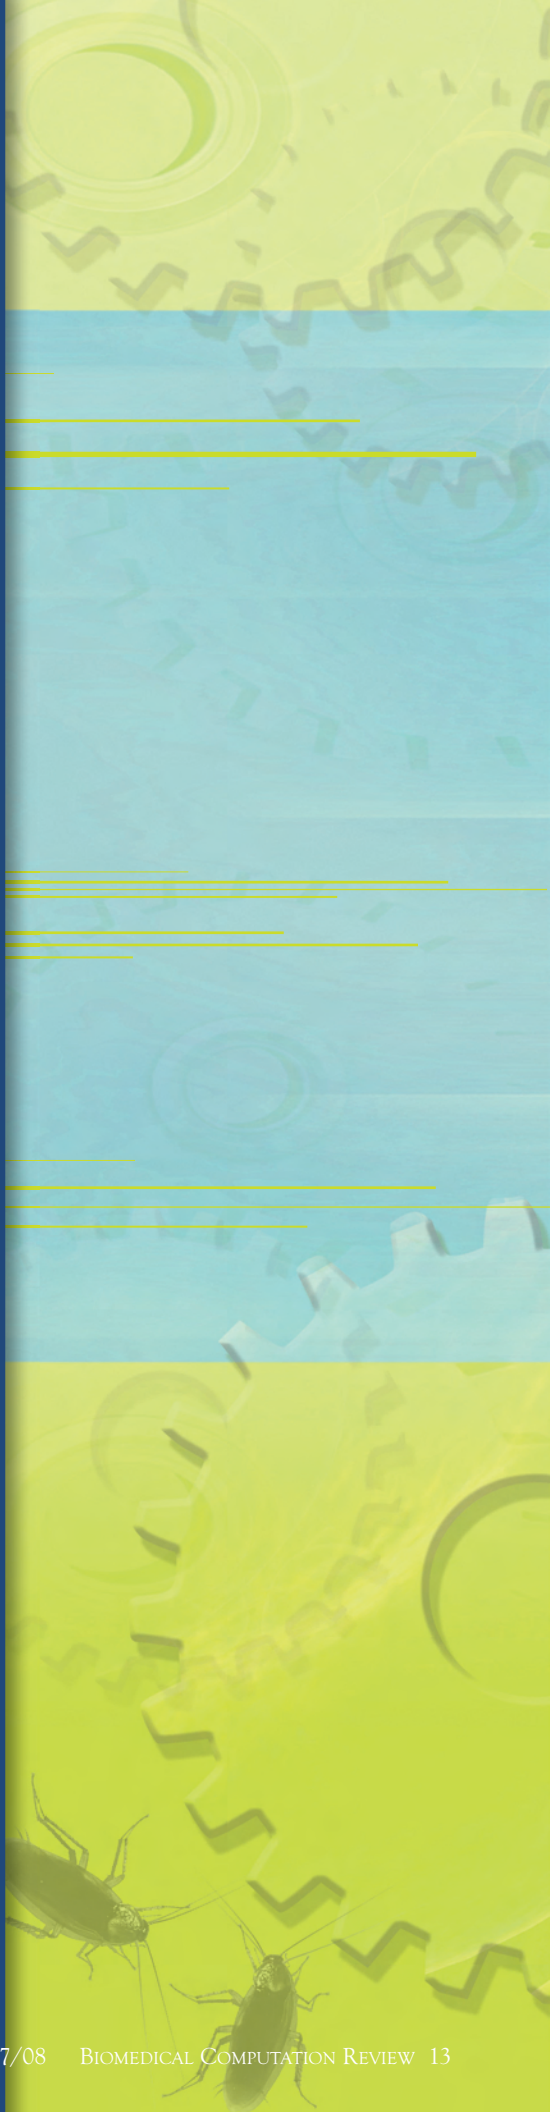

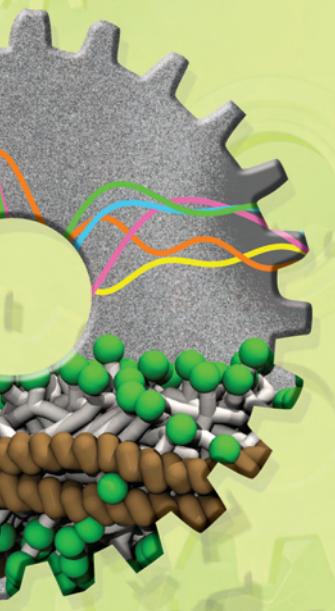

Jiggling Molecules

"Everything that living things do can be understood in terms of the jiggling and wiggling of atoms," said **Richard Feynman** in his seminal 1963 *Lectures on Physics*. Guided by the laws of physics, the atoms that make up the molecules of life assemble themselves into essential forms to do a wide variety of tasks.

Today, researchers can observe those molecules interacting by simulating them on a computer. Using this approach, known as molecular dynamics, computational researchers are going beyond what experimentalists can do to understand the way life works at the nanoscale, says **Klaus Schulten, PhD**, professor of physics at the University of Illinois at Urbana-Champaign.

For example, when Schulten's group simulated the structure of a water channel, they learned that water molecules pass through the channel in a very specific orientation: oxygen first. X-ray crystallography—a standard experimental method for studying atom arrangements—could not determine the orientation of the molecules, Schulten says. "Computation gives additional insight into the system."

Sometimes Schulten and his colleagues beat experimentalists to the punch. For example, in 1999, their molecular dynamics simulations of the largest known protein, titin, explained how the protein gives muscles stretchability—under force, nine hydrogen bonds are disrupted in a reversible way.

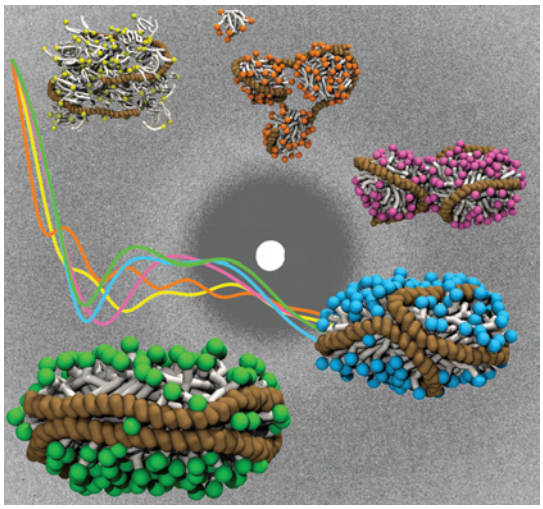

This graphic shows several steps in a coarse-grained molecular dynamics simulation of a lipoprotein nanodisc assembly. At the start, two semi-circular membrane scaffold proteins (brown) are surrounded by randomly scatttered lipids (small tailed objects shown here in a different color at each step in the simulation). During a 10 microsecond simulation, the lipids quickly glom together. The fusion of these lipid micelles draws the two protein strands (brown) together eventually forming a single lipoprotein particle. The aggregation, driven by hydrophobic effects, is followed by a much slower protein tertiary structure rearrangement. Eventually the protein strands rearrange themselves to form a double-belted nanodisc. Courtesy of Klaus Schulten.

Schulten and his colleagues have simulated interactions among all of the atoms in a variety of biomolecular systems ranging in size from water channels and lipoproteins (on the order of 10^5 atoms) to an entire virus particle (10^6 atoms) and, most recently, a bacterial flagellum (10^9 atoms).

"With computation, we can take experimental data with limited meaning and, using what I call a "computational microscope," turn it into information about the chemical structure of the system under investigation," Schulten says.

Three years later that finding was confirmed in a lab. Likewise with ankyrin, a molecule that's important for hearing. Simulations showed that the protein was a very soft spring. "It was very stretchable in the computer," Schulten says. "Put a feather on it and it stretches to the ground." The computational results were published before the experimental results came in. "The lab researchers confirmed the computational work," he says.

Recently, Schulten's group has been taking steps toward simulating larger systems. "We've moved from single protein sports in the cell to describing team sports," Schulten says.

Lipoproteins, which contain both proteins and lipids, cannot be crystallized because there is so much disorder in the lipids. "To get a picture of the molecule, you need the computer," Schulten says. But simulating the assembly of the lipoprotein molecule at atomic resolution would have required simulations in excess of 100 milliseconds—more than their computer could do. So the team simplified the system—a process known as coarse-graining—to effectively cut down the number of elements. A series of pictures showing the self-assembling lipoprotein appears on the previous page.

Schulten's team is currently using molecular dynamics simulations to generate movement of a bacterial flagellum. They've already created an all-atom model and simplified it using coarse graining. "Now we hope to rotate the entire flagellum around the base to see what kind of motions we get," Schulten says. "It's a very challenging 10 microsecond rotational period, and we're not done yet."

"With computation, we can take experimental data with limited meaning and, using a "computational microscope," turn it into information about the chemical structure of the system under investigation," says Klaus Schulten.

Molecular Movement Takes Time

One of the challenges of simulating moving particles is that movement takes time, and simulating over time requires significant computational resources. For example, with 20,000 processors, Schulten says he can simulate 100 nanoseconds of molecular movement a day; and he needs 100 days to get to the microsecond level.

But **Vijay Pande, PhD**, associate professor of chemistry and of structural biology at Stanford, wants to cover longer time scales than the typical nano- and micro-second simulation. He's interested in processes that take milliseconds or seconds—or even protein activation, which can take minutes or hours. "Our interest has been to push as hard as possible into this area," he says. "If you have something that takes a millisecond but you're simulating it for 1,000 times less time, you'll probably be missing things."

Pande points to simulations of the headpiece for the protein villin. Initially, simulations of all possible future positions of each particle (known as trajectories) lasted only a brief time-step with gaps between. It was impossible to make sense of what was happening. "Now that we have thousands of trajectories, each on these long time scales, we can see what it looks like," Pande says. And the details matter: In a movie, the headpiece folds, unfolds a bit, tries again, gets some things right but maybe not all, unfolds a bit again, and so on until eventually it makes its final shape. Yet, says Pande, "Every step of the way it creates more and more native-like structure."

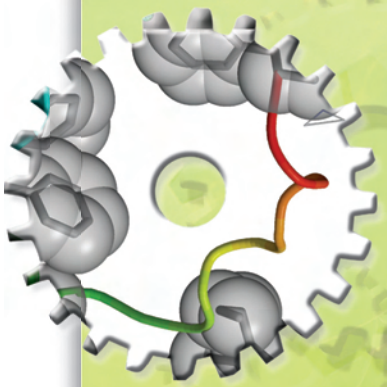

"Typical simulations cover events happening in nano- and micro-seconds," says Vijay Pande. "But when I look at the events I want to go after, they occur more slowly—at milliseconds or seconds."

Perhaps the only general statement one can make about protein folding, Pande says, is that it's a stochastic process, meaning it involves chance or probability. "Trying to understand this means we might want to rethink how we simulate dynamics," Pande says. "The question is: if a handful of trajectories don't really describe the system, how are we going to capture all the complexity and interest that even a small protein molecule might have?"

So Pande suggests a paradigm shift. Instead of running simulations evenly, giving each trajectory equal attention, he proposes using Bayesian statistical methods to figure out which areas really need to be simulated longer in order to gain insight. This can yield huge speed increases, making protein folding simulations 10 to 1000 times faster.

To magnify that speed increase, Pande relies on large amounts of computing power—specifically, grid computing. He has 250,000 processors participating in his distributed computing program, including graphics processors like Sony's Playstation® 3—which he says give a 20 to 50 times speed increase. By combining this with Bayesian methods, "we hope to get millions times the speed of what you can do with one computer," he says.

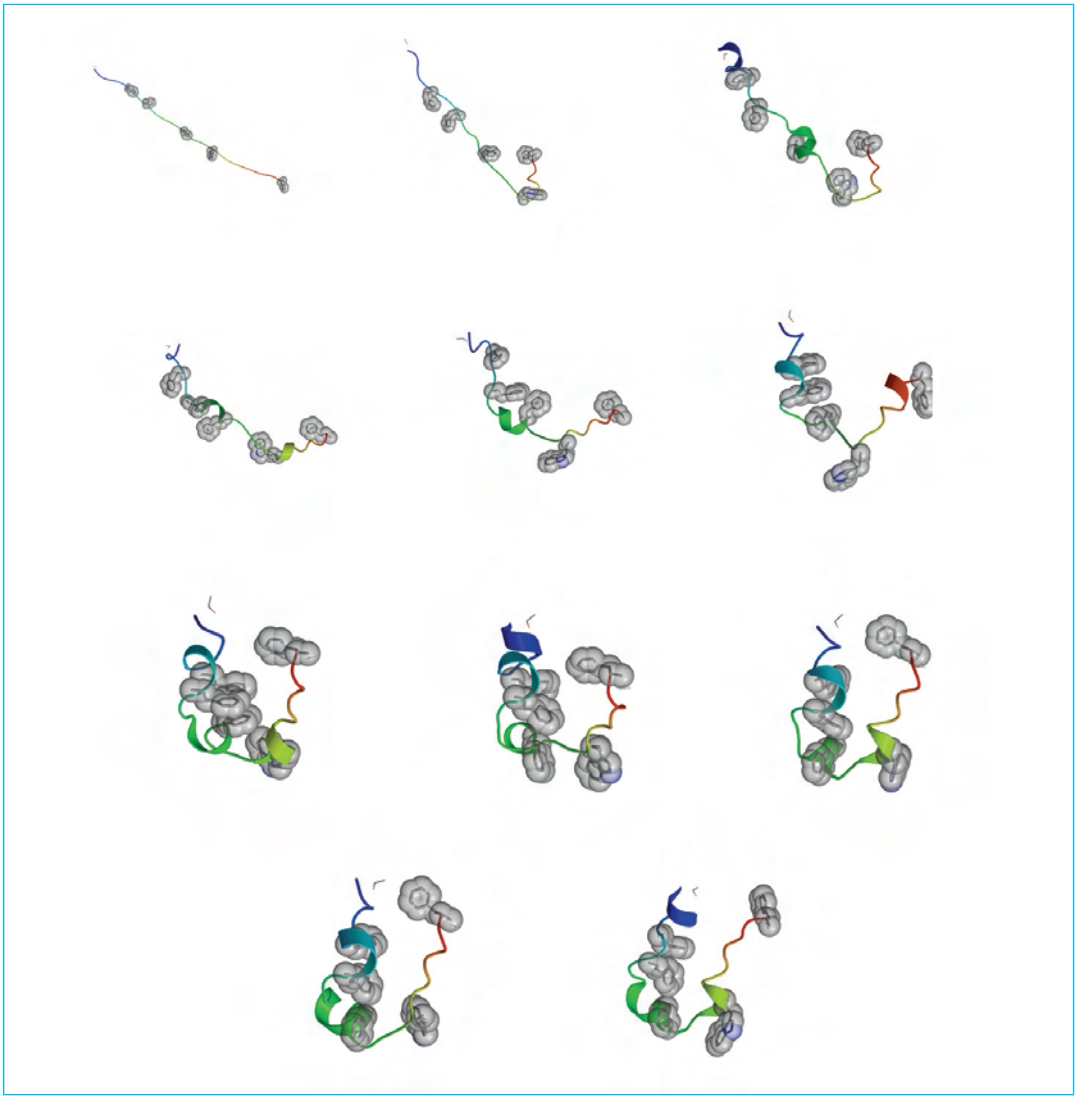

The headpiece of villin, an actin-binding protein, makes for interesting molecular dynamics simulations because it is quite small and folds quickly. That speed allowed Pande's group to obtain thousands of trajectories showing the headpiece folding and unfolding while continuing to make progress toward its final state. Courtesy of Vijay Pande.

The Ratcheting, Springing, Ejecting Ribosome

The ribosome is a magnificent piece of cellular machinery. It's where the DNA code (carried by messenger RNA) gets translated into useful proteins. Three different kinds of molecules move through the ribosome during this process: messenger RNA (mRNA); numerous transfer RNAs (tRNA) (each one carrying an amino acid); and the growing protein that's being assembled from those amino acids. For each pathway through the ribosome, unique mechanical features help ensure an efficient and accurate process.

Studying the dynamics of such a molecular machine is a complex business. But by combining cryo-electron microscopy, or cryo-EM, with an array of interdisciplinary methods, researchers have made tremendous headway. **Joachim Frank, PhD**, professor in the School of Public Health and Biomedical Sciences at State University of New York, Albany, is one important contributor.

Frank's team starts with a soup of ribosomes and other essential ingredients for translation including mRNA, amino-acid-carrying tRNAs, various elongation factors, and amino acids. Using antibiotics or other chemical means, the researchers stop the translation process at a particular step. In this manner, all the ribosomes in the sample are trapped in the same conformation—e.g., with tRNA snapped into place or not; or with mRNA in a particular stage of movement. The electron microscope then makes tens of thousands of projections of these ribosomes that must be assembled into a single 3-D map. "The job of the computer is to make sense of all these projections," Frank says.

Although the resolution of the 3-D maps created with this cryo-EM approach is getting better and better, it's still shy of atomic resolution. So Frank's team then docks existing crystal structures into the cryo-EM map. At this point, Frank says, "We have only discrete stops along the way. So we have to figure out what happens in between the snapshots." Multiple interdisciplinary techniques then come into play: Kinetics data, molecular dynamics simulations, normal-mode analysis, single-molecule FRET, and other approaches create a more complete picture of the ribosomal molecular machine in motion.

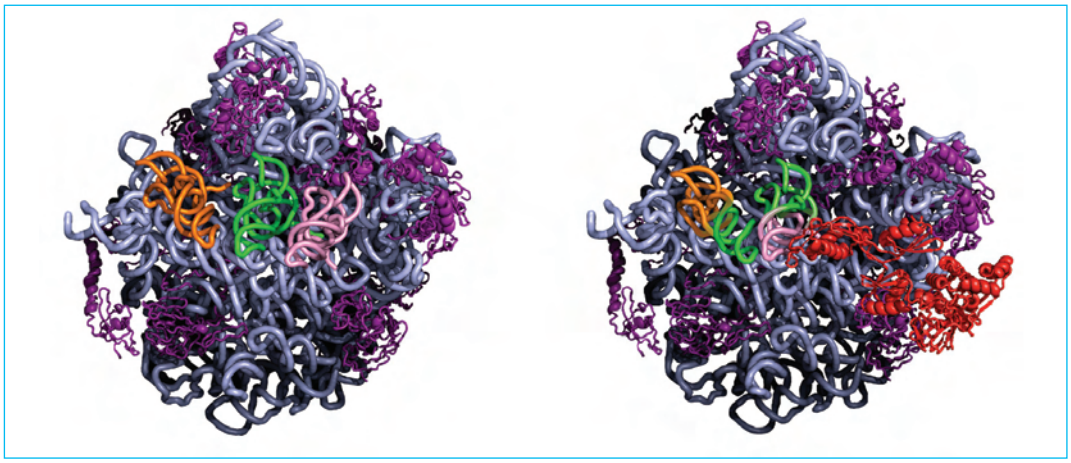

Here, the large ribosomal subunit is shown alone (the small subunit has been removed) in the process of translocation. On the left, three tRNAs sit in the three slots inside the ribosome: orange prepared to exit; green in the middle; and pink having just arrived and in a position to link its amino acid to the growing peptide chain. When (at right) elongation factor G (red) binds to the ribosome, it induces a ratcheting motion—the small subunit (not shown) twists down and away, which causes the mRNA to shift over by one codon. In addition, the three tRNAs shift toward their next position, as shown. When the subunits a spot for a new tRNA to enter. Image courtesy of Joachim Frank and Haixiao Gao.

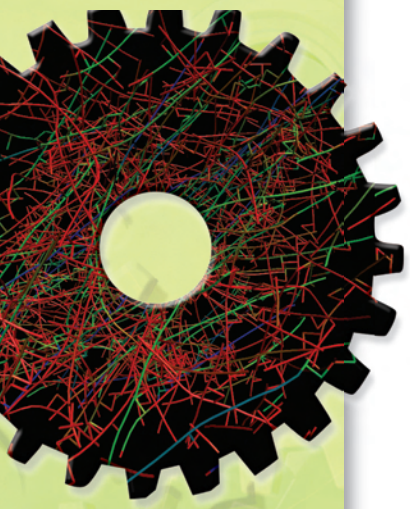

The results have answered important questions about how the ribosome works, including the processes known as tRNA selection and translocation. In tRNA selection, the ribosome must choose among twenty different tRNAs each attempting to deliver a different amino acid package to the growing peptide. How does the ribosome choose the right one accurately? Frank's cryo-EM analysis shows that the tRNA undergoes a conformational change when it enters the ribosome to try out its match: it gets bent into a molecular spring with high energy. Only with a match between the tRNA anticodon and the mRNA codon does the spring snap into place. "tRNA makes an enormous move going into its place in the ribosome," Frank says.

The work also elucidated the process of how mRNA and tRNA move through the ribosome (translocation). When elongation factor G (EF-G) binds to a site between the ribosome's large and small subunits, the small subunit moves in a ratcheting motion to one side and (after EF-G departs) back again. "There is an enormous movement of the bridges that connect the two subunits," Frank says. As a result, mRNA shifts over by one codon and, at the same time, tRNA moves stepwise from one of three positions to the next.

"The ribosome's dynamic properties follow from the molecular architecture," Frank says. And that's something Vijay Pande has thought about in his investigations of the polypeptide's exit tunnel through the ribosome. Pande wonders why the growing protein departs through the center of the ribosome and what interactions occur along the way.

To investigate, Pande began by considering tunnels generally. His team simulated a helical peptide inside different sized nanotubes with a small number of water layers. The result: "Most people would expect that you'd have a more stable helix in the smaller tube because you're removing the unfolded states," he says. "But there were fewer helical residues and more protein-hydrogen bonds." His hypothesis—in such a small space, water denatured the protein!

"Thinking about water is really important," Pande says. "If you didn't think about it explicitly here, you'd have gotten the opposite result."

Taking that idea to the ribosome, Pande's simulations show that protein helices mostly remain coiled inside the ribosomal tunnel. Currently, his group is simulating whether the polypeptide interacts with the ribosomal tunnel on its exit journey. "Perhaps the ribosome itself goes through some changes during the process," he says. "Even residues down at the end of the tunnel might affect what's going on higher up." Results are expected soon.

The Stretching Cytoskeleton

Inside a cell, the cytoskeleton creates a scaffold for essential activities such as cell migration, division and signaling. To function well, it must be flexible and stretchable. But we have a poor understanding of how it works mechanically.

To study the cytoskeleton's mechanical properties, **Roger Kamm, PhD**, professor of mechanical engineering and biological engineering at Massachusetts Institute of Technology, uses both experimental and computational approaches.

As you increase strain on a cell, Kamm says, cytoskeletal filaments get stiffer, a process known as strain hardening. Eventually, these filaments reach a critical point and there's a sudden drop in stiffness; the material gets much more fluid-like. But what, Kamm wondered, causes both the hardening and the drop?

To explore that question, Kamm and his colleagues developed an actin cytoskeleton *in silico* consisting of a network of actin filaments in a 500

"We developed an actin cytoskeleton *in silico*," says Roger Kamm, "so that we can study mechanical processes in a more systematic way where we can probe things in detail."

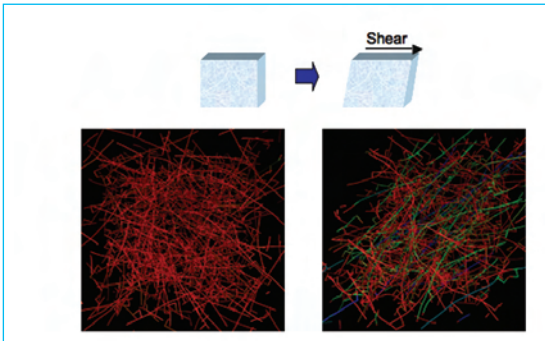

A simulated cytoskeleton shown without (left) and with (right) shear stress applied. The green filaments (at right) are the ones that "feel" the sheer stress. Courtesy of Roger Kamm.

nanometer cube. And because strain hardening only occurs in experimental networks with cross-linking proteins, they threw two cross-linkers into the simulation—one that connects the filaments in parallel, and another that joins them at right angles. "The nice thing is that we can start simple and add in different types of cross-linkers as we go," he says. Taking a slice through the cube, the network appears reasonably similar to that seen in slices through real cells.

Kamm's group then modeled the *in silico* matrix with and without shear stress. The key result: "At higher concentrations of cross-linking protein, you get this dramatic strain stiffening behavior," Kamm says. On the other hand, he didn't see the catastrophic drop in stiffness seen experimentally. For this, the cross-links need to rupture under force, an effect that is now being incorporated into the model.

Eventually, Kamm hopes this kind of iterative research using experimental and computational approaches will lead to a better understanding of how forces are transmitted across the cell membrane and within the cell.

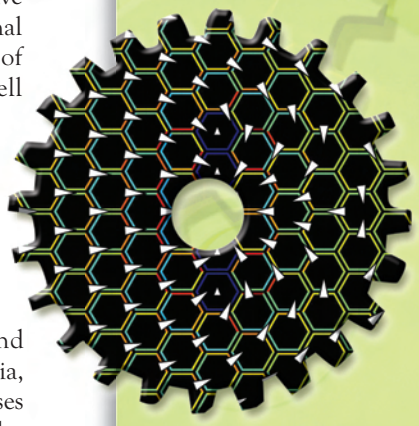

One-Way Tissue

Many types of cells and tissues develop a kind of directionality called cell polarity: certain events happen toward one end of the cell or tissue. When disrupted in humans, a variety of disorders may result: congenital deafness, respiratory diseases involving cilia, neural tube defects, and even cancer.

To study cell polarity, some researchers turn to the little hairs that grow from the distal (far) side of each cell on flies' wings. "How does a single cell in the midst of thousands of cells, identify the distal side?" asks **Claire Tomlin, PhD**, professor of aeronautics and astronautics at Stanford University and professor of electrical engineering and computer sciences at the University of California, Berkeley. To address that question, Tomlin uses mathematical models to bridge from molecular level understanding to tissue-level effects.

A fly's wing hairs form between 18 and 34 hours after the pupa forms. Before 18 hours, key proteins in the cell are homogeneously distributed. After that, they localize—two (Dsh and Fz) to the distal side and two (Pk and Vang) to the proximal side. The hairs form where Dsh concentrations are highest. But various mutants exhibit unusual

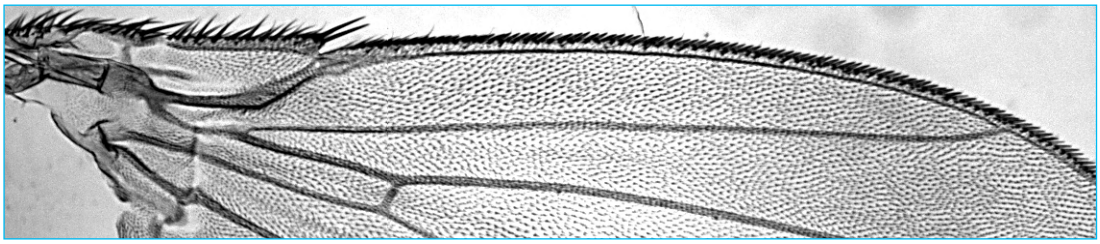

The hairs on a fly's wing grow from the distal side of each cell (the side away from the fly's body), displaying what's called planar cell polarity. Courtesy of Claire Tomlin.

characteristics. For example, mutant cells with no Fz grow more than one hair from the cell's center. In addition, wild-type cells surrounding those mutants grow hairs pointing toward the mutants, suggesting some sort of signaling between cells.

Jeff Axelrod, PhD, associate professor of pathology at Stanford University, proposed a feedback model among the various players to explain this and other mutant outcomes. But the model was controversial. Some felt it couldn't explain certain phenomena such as why cells that over-express Pk produce increased Dsh at the boundary.

Enter Tomlin who developed mathematical models to determine if Axelrod's model was plausible. She used continuous partial differential equations to represent the observed diffusion. Simulations of various knockout scenarios produced results just like those seen in the lab. Even the controversial result was captured: Overexpressing Pk brought more Dsh to the boundary. The lesson, Tomlin says: "Feedback can be nonintuitive." Here, it turned out that overexpressed Pk led to overcompensation by the rest of the feedback loop.

Tomlin's model also performed remarkably well in a blind challenge. A researcher in England asked her to attempt a variety of "funky" knockouts. After the simulations were complete, he showed her his lab results for the same experiments. They all matched. "It supports the plausibility of the model," Tomlin says. "And as we get more and more information from experiments to guide our model, our model can also help guide experiments."

"We're using a mathematical model to bridge the gap between a molecular-level hypothesis and its tissue-level effects," says Claire Tomlin.

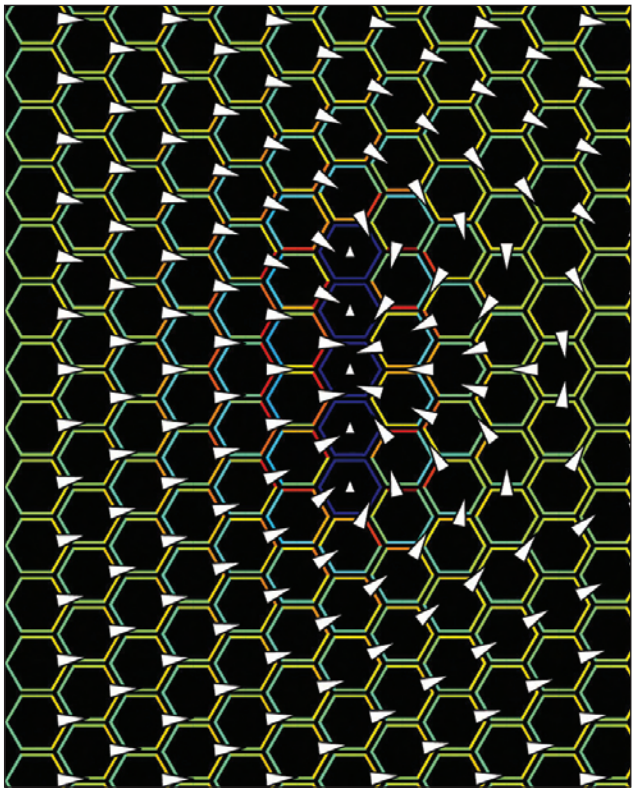

In this model of a fly's wing, some cells (in the center of this image) contain a non-functioning clone of the gene "frizzled." As a result, hairs (white triangles) grown from neighboring cells point toward the Fz-deficient cells. Courtesy of Claire Tomlin.

Swimming Larvae in the Natural World

On a larger scale, some researchers study entire organisms moving in their natural environments. "They evolve in the messy natural world," says **Mimi Koehl, PhD**, professor of integrative biology at the University of California, Berkeley. "So we should look at how organisms interact mechanically with the world around them."

Koehl has studied how crabs move in and out of water (they crouch lower and wider in water) and how lobsters gather odors from the environment (they flick their antennules to "sniff" the surrounding water). Recently, she teamed with some engineers to study whether sea slug larvae have any control over where they land on the sea floor. Ocean currents carry these microscopic larvae, but these creatures need to land on a coral reef in order to metamorphose into their adult form.

"What if you're really tiny and not a great swimmer," Koehl asks, "How can you recruit a suitable habitat?"

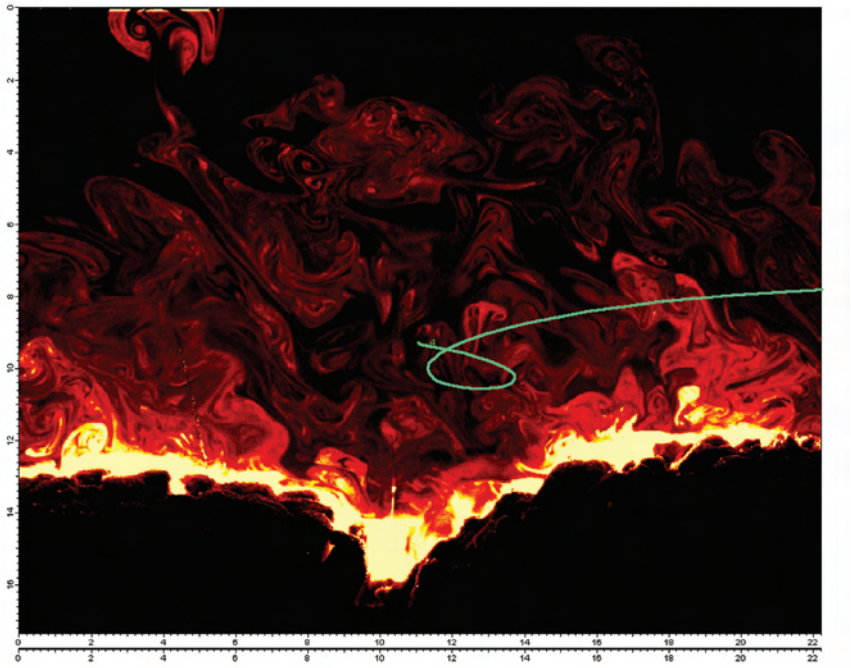

A computational larval sea slug (blue) finds its way to a coral reef by swimming when it senses no coral aroma (in black filament) and sinking when it senses such cues (in red water). As a result, the larva follows a spiralling path that enhances the chance of landing on the reef. This model was built by simulating realistic turbulence and waves in a laboratory wave tank. The black branching structures at the bottom of the image are corals exposed to turbulent flow. The yellow and orange filaments swirling around in the water above the corals are fluorescent dye leaching off the surfaces of the corals (just as the aroma of corals leaks out of corals in nature). A thin optical slice of the dye (aroma) plume is illuminated by a thin sheet of laser light. The brighter and lighter the dye in this image, the higher its concentration. The aqua line is the simulated trajectory of a microscopic larva of a sea slug carried in turbulent water flow. The trajectory was calculated using a computer simulation of the larval behavior (in odor-free water it swims; in water with coral-aroma it sinks) as well as the water movement (waves and turbulence), and the changing field of filaments of aroma swirling around in the water. Photo by M. Reidenbach; trajectory calculated by J. Strother. Courtesy of Mimi Koehl.

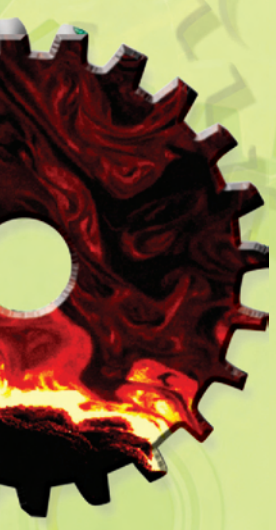

In a laboratory dish, Koehl's collaborator, **Michael Hadfield, PhD**, at the University of Hawaii, had seen that chemical cues can trigger sea slug larvae to settle on the bottom. But in nature where they drift over coral reefs, they have to contend with turbulence and waves. Wouldn't the coral odors be washed away and wouldn't larval behavior be overwhelmed by the ambient flow, Koehl wondered?

To mimic the conditions in nature, Koehl and her colleagues measured water flow over coral reefs and then worked with **Jeffrey Koseff, PhD**, and **Matthew Reidenbach, PhD**, at Stanford University to recreate a coral reef in a wave tank, complete with realistic wave movements and turbulence. They painted a fluorescent dye on the model corals to represent chemical cues released from coral surfaces. As the fluorescent dye dissolved into the water column, the researchers shined a skinny sheet of laser light vertically through the dye plume so that they could look at how the dye was distributed in the water on the fine scale that would be encountered by a tiny larva. The fine slice revealed that the fluorescent dye wasn't merely a diffuse cloud. It was made up of fine filaments swirling around with odor-free water. "A tiny dot the size of a larva is going to be in no odor, then odor, on-off, on-off as he swims around through this plume," Koehl says.

What do larvae do in response to this situation? In filaments of odor-free water they swim; in filaments containing coral aromas (above a threshold concentration) they turn off their cilia, pull in their swimming gear, and sink. When larvae exit the cue, they resume swimming.

One question remained: Does this behavior help larvae to land on the reef? To study this, James Strother, an undergraduate physics student at University of California, Berkeley, worked with Koehl to create a computational model of larvae over a reef. Mathematical larvae were placed in the video records of dye swirling over corals in waves. The larvae were programmed to sink if surrounded by a cue concentration greater than a pre-determined threshold, and swim if immersed in a lower cue concentration. The larva's velocity depended on its swimming or sinking speed plus the velocity of the waves and the effect of turbulence. "It's carried by the fluid but also sinking or swimming by its own volition," Koehl says.

In the simulation, the researchers saw larvae following a spiraling trajectory, eventually landing on the reef. When they calculated the trajectories of thousands of larvae they found the settlement rate into the reef increased about 30 percent because of the larvae's sink/swim behavior. Larval behavior made a statistical difference to larval survival. "Even if you're a tiny, weak swimmer, you can bias how the environment moves you," says Koehl.

Crawling Creatures Unplugged

Running creatures with two, four, and six legs veer toward stability. Indeed, they seem mechanically designed to cope with varied and unpredictable terrain and to recover from trips and jolts that disrupt them along the way.

To understand how runners achieve such stability, one might assemble a model of multiple skeletal supports, hundreds of muscles and millions of neurons. But that would be an extremely complex model with many variables. Moreover, it would be hard to make any general statements about what's going on.

So the system should be simplified down to its essence, says **Robert Full, PhD**, professor of integrative biology at the University of California, Berkeley. But that essence should be anchored in a realistic physical representation of an animal.

The essence of running—distilled down to a simple, dynamical system in one plane—can be represented by a pogo stick. "It's a mass sitting on top of a spring," Full says. "And it's the same for two, four, six or eight-legged animals."

In addition to a vertical springing motion, running involves movement in the horizontal plane, Full says. He collaborated with a math colleague at Princeton to produce a spring-mass model that bounces side to side as well as up and down. Then they perturbed the model. The result: heading, velocity, orientation and rotational velocity all remained stable. "This is a passive mechanical self-stabilizing system with almost no neural feedback," says Full. "The stabilization is built into the tuning of the mechanical system."

"Morphology makes a big difference with respect to control," says Robert Full.

His team then tested the spring by gradually adding physical characteristics of animals—first legs, then a simple muscle model and some damping, and finally some programmable leg and hip positions that could control joint torque. Each addition revealed a new characteristic—stabilization with respect to inertia, then speed, then stride. The lesson: animals opt for a combination of speed and stride frequency based on stability. "The most stable region is where the animal actually functioned. It didn't venture into the unstable regions," Full says.

A postdoc in Full's lab went on to produce a variety of hopping models and learned that two legs add lateral stability; three increase stability in all directions. "Morphology makes a big difference with respect to control," Full says.

The next step was to add a model of neural control. Full's team had experimented with real animals—cockroaches wearing jet packs or tripping over a step—and found that little neural control was needed to respond to perturbations. So they expanded their model to inclu de only a very simple neural control model: an oscillator (one for the whole system or one for each leg) coupled to a mass supported by legs.

This modeling and animal experimentation inspired the design of a physical model, a robot. The insert-like robot consisted of six springy legs coupled to a body without any external neural sensing. Yet, surprisingly, it is remarkably stable and can negotiate varied terrain, including climbing steps. Most of the control resides in the body and not the brain, Full says.

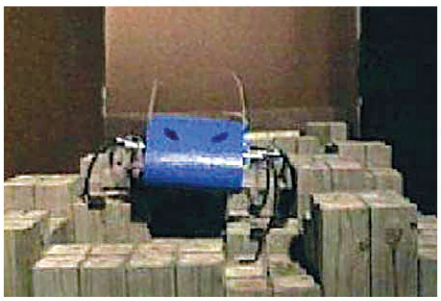

At top, a cockroach moves across irregular terrain containing obstacles three times its leg height. At bottom, a biologically-inspired hexapedal robot successfully traverses a scaled-up version of the same landscape without any sensory control. From Koditschek, et al., Arthropod Structure & Development 33 (2004) 251–272 with permission from Elsevier. Courtesy of Robert Full.

Running Dinosaurs

Like Full, **John Hutchinson, PhD**, a lecturer at the Royal Veterinary College, University of London, uses simplified models to understand animal movement. But that's in part because he has to: The animals he studies are extinct.

Specifically, he's interested in theropod dinosaurs, a group that walked on only two legs and includes the largest bipeds that have ever lived. Living animals appear to have a speed limit—at a certain size, getting bigger no longer means getting faster. Hutchinson wants to know: Was this true for *Tryrannosaurus rex*? Could this massive dinosaur run?

Empirical data give limited information about dinosaur locomotion. Skeletons can tell researchers what positions the animals couldn't take—for example, poses that would disarticulate the knee. Footprints allow some estimates of dinosaur speeds. And comparisons to living animals also give some clues. However, in the case of *T. rex*, "There are no six ton bipedal land animals alive today," Hutchinson says.

So he uses computer modeling and simulation to go beyond what he can see in fossils, footprints, and analogous live animals.

Initially, he created a simple model to determine how much muscle mass it would take for *T. rex* to sustain the ground reaction force of normal running (2.5 times body weight). "No matter what posture we put into the model, *T. rex* could not have carried enough muscle to run quickly—even if you gave it incredibly big muscles," Hutchinson says. According to the model, *T. rex* would max out at no more than 15 to 25 miles per hour.

Next, he created more complex 3-D models. He started with 6 million possible poses and then eliminated the unlikely ones based on principles of how living animals move. For example, he imposed reasonable limits on such things as limb bending, muscle mass, and the size of the moment arm about each joint. After carving down the possibilities, 3000 poses (.05 percent) remained that could sustain 1.5 times body weight (at the boundary between walking and running). These were fairly straight-legged poses, like those seen in living large mammals.

John Hutchinson created these muscle-activated models (from left to right above) of an ostrich, elephant and Tyrannosaurus rex in order to compare their running ability. When a typical living animal runs fast, the ground reaction force (at midpoint stance) peaks at about 2.5 times body weight. Hutchinson looked at how much muscle the living animals would need in order to sustain that force. He found that birds—small and large—all had enough muscle to run. In addition, large birds such as the ostrich could run quickly because they have big muscles, good muscle leverage and use straight legs, which give them a mechanical advantage. Elephants could also run. But T. rex doesn't appear to have been able to sustain a ground reaction force of 2.5 times body weight. Courtesy of John R. Hutchinson.

He then used this model to simulate *T. rex* at the midpoint in a walking stride and asked the computer: How much ground reaction force could have been produced at the foot per one unit of muscle force? It's a measure of mechanical advantage.

What he found: As the *T. rex* posture gets more erect, the mechanical advantage goes up until it reaches a plateau. "It's possible that perfectly straight legs aren't that much better than a little bit of bend in the legs," he says.

For dinosaurs, so much is unknown—the posture, the moment arms, and the dimensions of the animal's muscles and bodies. But, Hutchinson says, "The unknown values have to be within some range....So perhaps it doesn't matter what we assume, but what we do with an assumption and how much we vary it. The modeling tools make all this careful inquiry possible; otherwise we'd just be left guessing."

"The modeling tools make all this careful inquiry possible; otherwise we'd just be left guessing," says John Hutchinson of his *T. rex* models.

Animating Realistic People on the Go

Jessica Hodgins, PhD, professor of computer science and robotics at Carnegie Mellon University, has no problem gathering ample data about her subjects: humans. But, like Hutchinson, she relies on a process of elimination to rule out unlikely human movements in favor of realistic ones. Her goal: to create better computer animations of people and make it easier for casual computer users to create such animations. Ultimately, she'd like to see her work have some practical impact. For example, it might help clinicians plan and implement physical therapy programs.

In 1995, Hodgins made her first efforts to simulate human motion. She relied on models of simple springs that were proposed in the biomechanics literature at the time; her own observations of people; and a healthy dose of intuition. The resulting animations weren't terrible, but it was easy to see they weren't realistic. "Our standards for human motion are really high," she says. "We know when something's wrong."

But this early effort taught Hodgins something: human control laws are hard to design. It's not enough to get the physics right. Accurately determined forces impacting a rigid body do not automatically produce an appealing human character. Moreover, they make everyone look the same. "We don't have a language for stylistic subtleties," Hodgins says.

She decided she needed to know more about how people move. So, starting in 2000, Hodgins created a motion capture lab. By placing numerous reflective markers on an individual person in motion and capturing the 3-D locations of those markers, she created a database of possible movements for a number of ordinary people as well as some professionals such as gymnasts and clowns.

Hodgins and her PhD student **Alla Savanova** could then organize the data in a different sequence to create an animated figure that moves in ways different from the original subjects. "A few examples of a given behavior can be generalized to quite different examples of that behavior and still look realistic," Hodgins says.

For example, using limited motion capture data, she can realistically animate a person following a randomly chosen circuitous path while avoiding various unexpected obstacles. To do this, she takes each pose from the motion capture data and puts it into a graph data structure: Similar poses and velocities land in similar locations on the graph. She can then search the graph for a path that approximates where the person is walking along a specified route. "Any path through this graph should produce natural motion," Hodgins says.

This procedure was quite efficient, but Hodgins made it even more efficient by reducing the number of possible poses to only those that look natural, and then interpolating between them. The physics are still correct for the full behavior, but the only motions the character can make are selected from a smaller set of possible motions.

Comparing animations based on the full motion capture data set with animations based on the low dimensional set, there's almost no difference. "So we're getting a lot of generality out of limited data," Hodgins says.

The team also worked to optimize the range of possible movements. For example, a person approaching an object to pick it up might bend over too soon or too late, looking unnatural. To find the most realistic options, Hodgins' group did a significant amount of culling of non-optimal trajectories on the graph. "Optimal solutions look so much more natural," she says.

But the movement of soft tissue is still missing from these rigid body animations: Muscles don't jiggle and feet don't compress with each step. Hodgins is now gathering data on such flesh and muscle movements by putting 400 sensors all over the bodies of weight lifters, belly dancers, and ordinary people with a variety of body types. Eventually she hopes to add models of these data to her animations.

Hodgins' work is top down: grabbing data and trying to mine it for the principles of human motion. The other option is to find out how the system works from the bottom up. "It will be nice when we meet in the middle," she says.

"A few examples of a given behavior can be generalized to quite different examples of that behavior and still look realistic," says Jessica Hodgins.

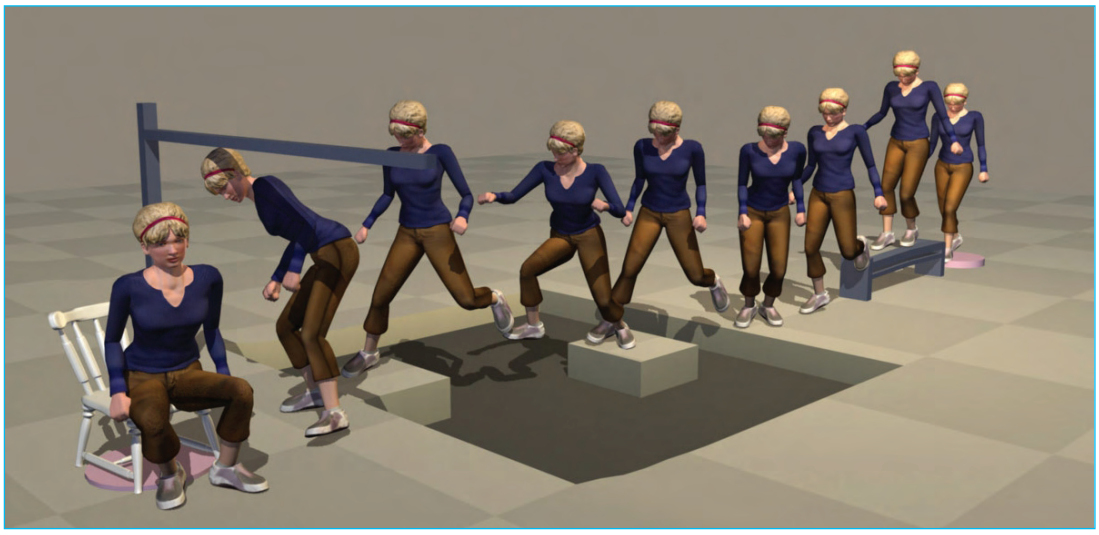

Jessica Hodgins captures the motions of living subjects and then manipulates them to produce realistic-looking movements in new situations. These snapshots, taken from an animation, show a woman walking on a balance beam, leaping between stepping stones, ducking under a bar, and then seating herself in a chair. Courtesy of Alla Safonova and Jessica Hodgins, Carnegie Mellon University.

Standing Up, Learning, and Taking Trains

Demetri Terzopoulos, PhD, the Chancellor's Professor of Computer Science at the University of California, Los Angeles, takes a comprehensive "artificial life" approach to animating humans and lower animals in a realistic manner. His characters are built from a basic biomechanical framework in which physics-based concepts such as joint torques and gravity—and not motion capture data—control movements. On top of that, he uses machine learning techniques to build in additional capabilities—learning, perception, behavior, and cognition—so that his characters can act autonomously.

A simple example is that of an articulated skeleton that tries to remain standing while being pulled by a virtual rope tied to its middle. Sometimes the skeleton falls; sometimes it stays up. The yank is repeated until the character "learns" which responses enable it to remain standing. After this type of training, the character can react autonomously, making a protective step to avoid a fall, for example, or getting back up after tumbling to the ground.

Terzopoulos' simulations go beyond just the physics and locomotion to also mimic animal perception and behavior. The models used in these simulations have a biomechanical component, but also include a set of sensors and a brain with motor control, perception, behavior, and learning centers. With this more complex model, Terzopoulos and his students created a biomechanical model of fish that can learn how to swim. The fish also avoid collisions with other fish, forage for food, and engage in more complex behaviors such as mating.

From this, Terzopoulos and his team developed a formulation of learning as an optimization problem. "It's trivially simple," he says. "Even a dumb animal can do it through trial and error." His program trained an artificial shark to swim by finding the most energetically efficient way to move given the physics of the environment (gravity in water) and the limits of its physique (e.g., muscle arrangement and strength). The shark begins by twitching, but soon, says Terzopoulos, "It discovers the proper gait given its body structure, and then refines it until it becomes an efficient swimmer."

Most recently, Terzopoulos and his team have applied their "artificial life" approach to human characters in a virtual train station (a model of the original Penn Station in New York City). Using the same layers (locomotion, sensors and a brain, this time including a cognition center), his team created a realistic simulation of several thousand autonomous pedestrians commuting through the

Virtual pedestrians move through this virtual train station by making autonomous decisions. They automatically negotiate tight spaces, observe social conventions (staying to the right), and queue up at the ticket line. They know when their train will leave, and decide what to do while waiting—sit and read; watch some performers; or grab a bite to eat. Each individual is constructed of several layers: biomechanics, sensors, and a "brain" able to perceive and learn. Courtesy of Demetri Terzopoulos.

station. "Each pedestrian is a highly capable individual with things he or she must do, such as purchase a ticket and proceed to the correct train platform at the appropriate time," Terzopoulos says. It's then possible to sit back and watch the train station dynamics on a large scale. "It's order and disorder at the same time," he says. "It's highly complex."

Eventually, Terzopoulos would like to create a whole city of autonomous virtual humans. The UCLA Urban Simulation Lab has created a detailed 3-D model of Los Angeles. "Wouldn't it be wonderful to populate it with as many people as possible?" he asks.

The UCLA Urban Simulation Lab has created a detailed 3-D model of Los Angeles... "Wouldn't it be great to populate it with as many people as possible?" Demetri Terzopoulos asks.

The Versatility of F=ma

When researchers simulate life in motion, they rely on a powerful rule of nature: Newton's Second Law of Motion. At all scales, the basic rule that force equals mass multiplied by acceleration (F=ma) helps predict how life moves.

Simbios, a National Center for Biomedical Computing at Stanford, was founded on the premise that this commonality—F=ma—should allow development of a common software toolkit for creating simulations at all scales of life.

And that's why the Simbios-sponsored "Life in Motion" symposium brought together the ten researchers described here. "Its pretty impressive that such a vast array of problems can be attacked with essentially the same tool," says Altman.

While highlighting the versatility of physics-based computer simulation, the symposium also fostered cross-fertilization among different disciplines that all use motion simulation. "This will help to build a meta-community of scientists with a common interest in understanding how biological matter moves, and how that motion can be simulated in computers," Altman says. □

Under TheHood

BY ALAIN LAEDERACH, PhD, STANFORD UNIVERSITY, DEPARTMENT OF GENETICS

Elastic Computing, Virtualization, and the Demise of the Common Cluster

Alain Laederach, PhD

After a hard day in "Lab," it is rare to come home with the feeling of actually having accomplished anything tangible. Sure, the software I am developing has some new feature (with the accompanying 10 new bugs), and perhaps an algorithm is slightly more likely to yield biological insight, but all in all, I've mostly pushed and pulled electrons around. I often envy my experimental collaborators, who at the end of the day have some DNA in their Eppendorf, or a new plasmid that will express a protein. Those are tangible, day-to-day results.

The other day, however, I came home having the feeling of true accomplishment. We had spent the entire day in the server room reorganizing our clusters and servers: We pulled power and Ethernet cords, tightened screws, and inserted RAM chips. But this might be the last time I ever upgrade RAM on a rack mounted computer because Amazon (yes, the online book store where you bought the latest Harry Potter) now offers another option.

Until now, computing has been about machines—a fixed cost. Labs buy a certain number of computers with a new grant, and that is it. The number of jobs queued on the hardware is highly correlated with conference deadlines, but most of the time, the hardware keeps itself busy running daily cron jobs.

Now there is a new player in the game, it is called EC2 (for Elastic Cloud 2) and it is available at http://www.amazon.com/ec2 along with more Harry Potter paraphernalia than I ever imagined. Amazon is the first company to sell computing as a true commodity independent of hardware. Large companies like Google, Amazon, Oracle and Microsoft are building data centers all over the country to meet their own huge CPU needs. But Amazon is the first to realize that their in-house technology (cheap commodity computing) can make them money—probably a lot of money.

EC2 offers cheap and simple pricing (10 cents per CPU hour on a 3 GHz equivalent processor with 1.7 GB of RAM, and 160 GB of storage). But perhaps more important for computational research, it will mean no more queues before conference deadlines. Amazon will worry about load balancing the world's computing resources; we'll just pay for computation as we go.

To those of you who are skeptical about scales and complexity, consider this: Overnight on July 21, 2007, Amazon shipped 2.2 million pre-orders for the latest (and final) Harry Potter novel. This is a company that has some experience with load balancing. Some will still say that the loss in performance with the added layer of virtualization is unacceptable, and that interconnects between virtual machines will never be fast enough for their highly parallel problems. But I'm betting that, eventually, they'll be buying CPU time too. The EC2 cloud already allows purchases of "large" instances (15 GB of memory, 8 CPUs, 1690 GB of instance storage, 64-bit platform) for 80 cents an hour. The price may sound steep, but consider the fact that you can create 1000 such instances almost instantaneously for $800. That's some cheap supercomputing, IMHO. □

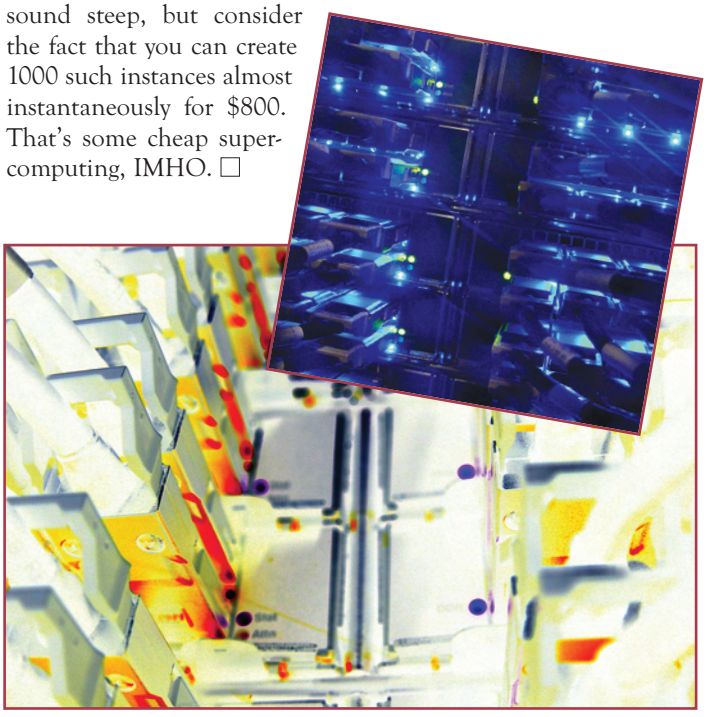

Photo credit: Shirley Wu

DETAILS

Alain Laederach, PhD, is a post-doc in Russ Altman's lab at Stanford University. He recently accepted a faculty position at the Wadsworth Center in Albany, NY, and he is not now and has never been associated with Amazon in any way!

You can reach him at alain@helix.stanford.edu

DIVERSE DISCIPLINES, ONE COMMUNITY
BiomedicalComputation REVIEW
Published by **Simbios**, a National Center for Biomedical Computing

Nonprofit Org.
U.S. Postage Paid
Permit No. 28
Palo Alto, CA

Biomedical Computation Review
Simbios A National Center for Biomedical Computing
Stanford University
318 Campus Drive
Clark Center Room S231
Stanford, CA 94305-5444

seeing science

SeeingScience

BY JOY KU, PhD

Simulation Simplified

Typically, researchers who simulate life in motion—from particles to people—start by describing the motion of each part of an object independently of the other parts. Additional equations then limit the motion of one part relative to another. The result: Computers must then solve a large number of awkwardly coupled equations. That often means long simulation times or, worse, equations that won't converge to a solution.

To sidestep these problems, **Michael Sherman** and his team at Simbios, a National Center for Biomedical Computing based at Stanford University, developed a multi-body dynamics software toolset called Simbody. Simbody introduces the concept of a "mobilizer," which directly expresses a part's movement, however complex, purely in relation to another part. In this way, fewer and simpler equations are needed to simulate the part's motion. □

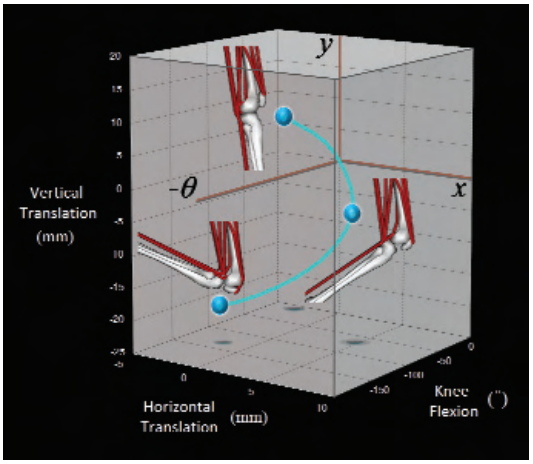

Traditionally, researchers have used five equations to describe knee motion–one equation for each of the ways a knee can move (forward and back; side-to-side; and rotational) and two equations to relate the translation of the knee to its rotation. With Simbody, a single equation, represented by the blue line, can simulate the complex motion of the knee. Courtesy of Ajay Seth, PhD, of Stanford University.

DETAILS

Simbody is part of the SimTK Core toolkit, an open-source C++ application programming interface (API) to computational tools and algorithms for biological simulations. A workshop on using Simbody and the SimTK Core will be held at Stanford University on March 20-21, 2008. For more information, contact Blanca Pineda, bpineda@stanford.edu.

Pre-release source code for the SimTK Core toolkit, including Simbody, can be freely accessed at http://simtk.org. A full release is planned for March 2008.