COMPOSE

THE DOE CSGF ANNUAL ESSAY CONTEST JOURNAL

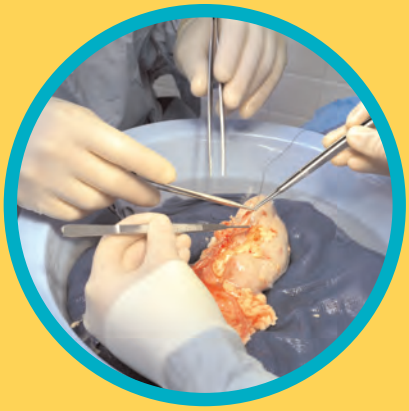

MATH MAXIMIZES ORGAN TRANSPLANTS

DOE CSGF

THE GENETIC CARRY-ON LIMIT

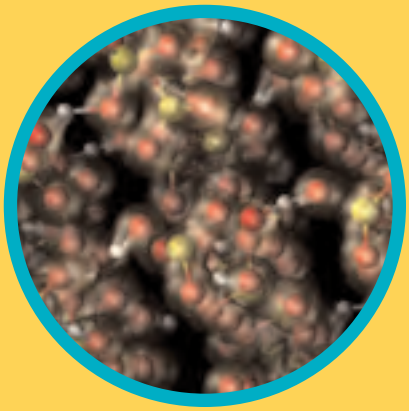

NATURE'S GREAT COMPROMISE

THE DOE CSGF ESSAY CONTEST WINNING ESSAYS

2007

The DOE CSGF Annual Essay Contest was launched in 2005 as an exciting opportunity for DOE CSGF Fellows to hone their writing skills. This contest requires Fellows to write a popular science essay on a topic of personal importance written for a non-science audience.

The DOE CSGF is proud to recognize outstanding Computational Science Graduate Fellows who have completed a non-technical writing composition on a topic in computational science. In addition to recognition and a cash prize, the winners received the opportunity to work with a professional science writer to critique and copy-edit their essays.

These copy-edited winning essays are published here, in this issue of Compose Magazine.

FOR MORE INFORMATION ON THE DOE CSGF ANNUAL ESSAY CONTEST, VISIT http://www.krellinst.org/csgf/compose/index.shtml

This year the essay submissions were judged by a three-person panel consisting of Christine Chalk, David Keyes, and Jacob Berkowitz.

Christine Chalk has been with the U.S. Department of Energy's Office of Science for more than 15 years in a variety of science policy positions. Ms. Chalk has degrees in Economics and Physics and experience on Capitol Hill. She is currently on a long-term detail to the Office of Advanced Scientific Computing Research from the Office of Budget and Planning — Division of Planning and Analysis. In addition, she has served on the screening panels for the American Association for the Advancement of Science's Science Journalism Awards the past two years. This is Ms. Chalk's second year reviewing DOE CSGF essay submissions.

David Keyes is a computational mathematician with primary interests in parallel numerical algorithms and large-scale simulations of transport phenomena – fluids, combustion, and radiation. He is the Acting Director of the Institute for Scientific Computing Research at Lawrence Livermore National Laboratory and is also the Fu Foundation Professor of Applied Mathematics in the Department of Applied Physics and Applied Mathematics at Columbia University. Dr. Keyes is active in SIAM and directs an Integrated Software Infrastructure Center in DOE's Scientific Discovery through Advanced Computing Initiative, called Terascale Optimal PDE Simulations. This is Prof. Keyes' third year as a DOE CSGF essay reviewer.

Jacob Berkowitz is a Canadian writer, journalist and playwright. He popularizes the work of leading scientists at major research-based organizations in Canada and the United States and is a long-standing contributor to DEIXIS, the DOE CSGF annual magazine. Mr. Berkowitz spoke about science writing at the 2006 DOE CSGF Annual Meeting in a talk titled, "Starting from the End: The Power of Turning Science into Story." His first book, "Jurassic Poop: What Dinosaurs and Others Left Behind," was published in 2006 and he's presently at work on a 50th anniversary follow-up to C.P. Snow's classic book on science and society "The Two Cultures." Mr. Berkowitz has been a DOE CSGF essay reviewer for three years.

Math Maximizes Organ Transplants

My husband, surgeon Dorry Segev, still marvels at the transformation he sees in his patients after they receive a kidney transplant at the world-renowned Johns Hopkins Hospital. The three-hour operations he performs can rescue patients from slowly dying on dialysis and return them to full, healthy lives.

The essential ingredient in such a recovery is a human kidney, and those are in short supply. More than 70,000 patients in the United States are waiting for a kidney from a deceased donor. About 8 percent of those waiting die each year.

With a little help from mathematics, however, many of the waiting thousands could soon get their desperately needed transplants.

Family members of a patient with kidney failure often are willing to give one of their own kidneys, but at least one third of such offers must be abandoned because of a blood-type or tissue-type incompatibility. One procedure that helps address the problem is kidney paired donation, which connects donor-patient pairs in which the donor of one family is compatible with the patient in a second family and vice-versa. The pairs exchange kidneys in simultaneous operations. Only a few U.S. institutions track the information that would make this procedure possible.

Surgeons and patients have fought for the creation of a nationwide registry to collect information about available incompatible donor-patient pairs, but concerns that trading kidneys might be similar to illegally selling kidneys blocked the effort. It literally took an act of Congress to clarify that such a registry is legal, and in December 2007 passage of the Charlie W. Norwood Living Organ Donation Act removed that barrier.

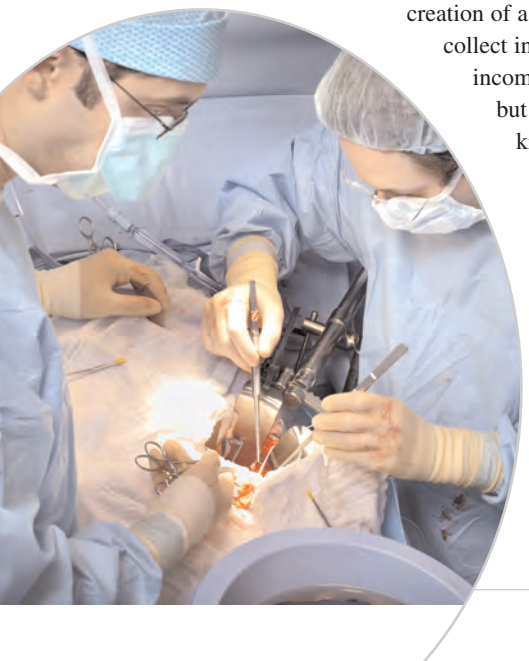

Doctors Dorry Segev, right, and Mazen Bedri of Johns Hopkins Hospital perform a kidney transplant.
Photo credit: Keith Weller

But a registry is just a start. By combining our knowledge of transplant surgery and mathematical optimization, Dorry and I have shown a branch of mathematics called graph theory can optimize the decisions surgeons make when matching donor-patient pairs. If a nationwide registry were implemented tomorrow, a graph theory tool called the Edmonds algorithm could enable surgeons to complete almost 2,000 more kidney transplants this year. We recently reported our work in the Journal of the American Medical Association.

Each patient with his incompatible donor can be visualized as a dot, or a node, in a network. A link connects any two nodes representing a compatible paired donation match. This network of nodes and links is called a graph. A paired donation graph contains links representing every possible match. Some links are more desirable than others, either because a kidney is predicted to function longer in a particular recipient or because the two families are geographically close.

> "By combining our knowledge of transplant surgery and mathematical optimization, Dorry and I have shown that a branch of mathematics called graph theory can optimize the decisions surgeons make when matching donor-patient pairs."

There are an enormous number of ways to match a group of, say, 250 incompatible pairs. If they were all printed out, the resulting stack of paper would be more than 500 feet high. So how can doctors decide which matches among all those pairs would give the greatest benefit in terms of match quality and the total number of transplanted kidneys? One simple answer would be to first choose the most desirable link, then choose

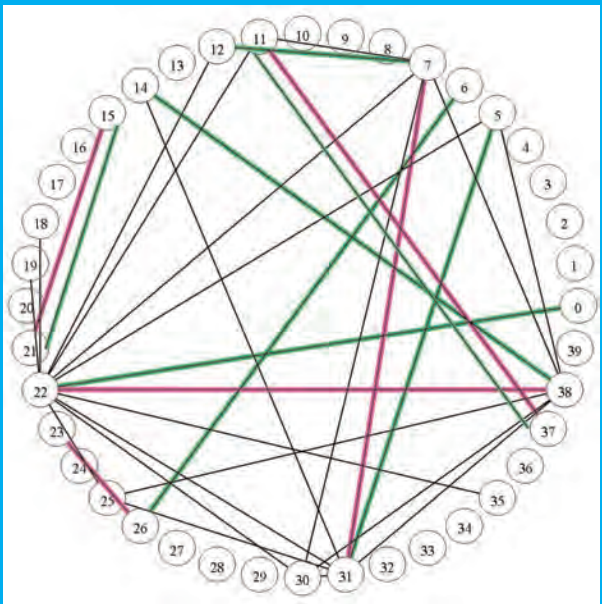

In this 40-node graph, black lines represent possible kidney paired donation matches. Pink lines show one possible set of matches, but only 10 pairs get transplants. Green lines show the best possible set of matches, with 14 pairs receiving transplants.

the second most desirable, et cetera, until no more paired donations are possible.

This naïve approach would ignore the other connections in the graph. But the graph's structure actually holds the key to increasing the number of transplants. For instance, some patients and their donors will be linked to a large number of potential matches, while some patients might be linked to only a handful of other pairs in the entire nation. Those with very few links are unlikely to receive a paired donation if the system fails to take the number of links per patient into account. The decision to match two pairs together affects every other patient in the graph who had a link to either of those two pairs. Deciding which patients and their donors should be matched, while considering all of the connections in the graph, is a job for the Edmonds algorithm.

Mathematicians know that the Edmonds algorithm is the best way to decide which links in a graph should be chosen so that the greatest number of nodes (patients and donors) is matched. However, the gold standard in medical research has been the clinical trial: a multi-year study involving a large number of patients randomly assigned to receive different proposed treatments. We propose that some medical research should instead be conducted by computational trial: an experiment that uses computers, along with known probabilities of patient and donor blood types and tissue characteristics, to predict the results of medical decision systems. Computer simulations use random numbers to create databases of virtual patients and donors that resemble real patients and donors.

Our computational trial simulated patients across the U.S. who need kidneys and compared paired donation algorithms for matching them. We predict that a national registry using the Edmonds algorithm could facilitate nearly 1,900 additional kidney transplants – an increase of about 12 percent over current transplant rates. A registry that did not use graph theory might miss about 300 of those transplant opportunities. By our calculations, paired donation also will save half a billion dollars in medical costs, because transplantation is less expensive than dialysis.

Before this computational trial, transplant physicians suspected that paired donation had promise, but underestimated how many patients it would affect. Now they know exactly how many thousands of patients could end the slow death of dialysis through paired donation. That knowledge has influenced lawmakers and doctors to make it happen.

Dorry and I are thrilled to see how modeling the impact of kidney paired donation has motivated the medical community to act quickly on an unparalleled opportunity to increase donation. When experts in math and medicine think collaboratively about organ transplants, patients win.

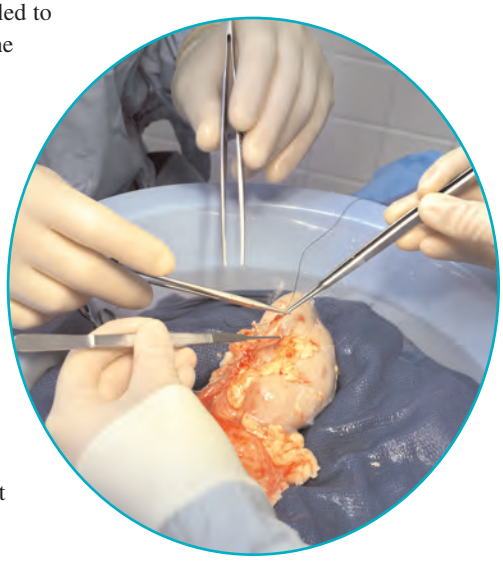

Surgeons prepare a kidney for transplant.

Photo credit: Keith Weller

The Genetic Carry-On Limit

Flying through Philadelphia is practically the same as *asking* to lose your luggage. As I wait at the baggage claim for a suitcase I know won't come, I wish I had listened to my mother and squeezed everything into a carry-on.

Packing for a flight reminds me of a game I played in elementary school, in which we had to list five items we would want if we were stranded on a desert island. I always listed television, my best friend and ice cream, but, for some reason, never a boat.

It's the same question we face every time we travel. What should we bring, and what should we leave behind? Deciding what to pack is almost equivalent to the philosophical question of what you need to survive and be comfortable.

I think about problems like this a lot. As a biochemical engineer in the Shuler research group at Cornell University, I'm working on finding the shortest possible list of genes necessary to support bacterial life. We do this by asking what a bacterial cell would choose to "carry on" if it were subject to size and weight restrictions like an airline passenger.

Knowing the bare minimum packing list for bacterial genes could help scientists connect those genes to the cell's behavior, answering questions like why certain bacteria make people sick and others do not, or which bacteria would be best for producing a medically important compound like human insulin.

Escherichia coli cells imaged using Differential Interference Contrast (DIC) microscopy. The bacteria's DNA is stained red.

Photo credit: Tricia Echtenkamp

Bacteria must constantly figure out how to pack everything they need to survive into a small space. In their case, the bag is a chromosome and the items packed into it are genes. The solution was determined through millions of years of trial and error, a process we understand as evolution. When a gene no longer serves a purpose, the bacteria eventually may lose it because carrying around that unnecessary burden slows the organism's growth. Bacteria also can acquire new genes as their environmental circumstances change.

Studying what bacteria *do* pack into their genomes can provide insights into which biological functions are required for life. Our research group's approach is to design a computational model of a cell with the minimum number of genes necessary to grow and divide. We call this the *minimal cell model.*

One strategy to decide what to pack in your carry-on might be to ask everyone at the airport what they packed and then copy the most frequent choices. Some items vary by location, but some are universal – you don't need your winter boots in San Diego, but everyone packs a toothbrush. Similarly, we can search online databases that list currently sequenced bacterial genomes and find common features. A researcher can reasonably assume that the genes or genetic features that appear most frequently are required parts of a bacterium's life-support machinery. For example, researchers have compared the genomes of the distantly related bacterial species *Mycoplasma genitalium*, the smallest living bacterium, and *Haemophilus influenzae*, the first organism to have its genome sequenced, and found they have approximately 250 genes in common. We can interpret these genes as the 250 necessary to support life. This represents less than 10 percent of the 1,000 to 5,000 genes found in an average bacterium. Other researchers have tried to compare larger collections of bacteria to obtain a more representative picture.

Our group uses the opposite approach to design a minimal cell. Rather than taking existing bacteria and removing genes until only necessary ones remain, we start with zero genes on our list and add those that will perform required functions. We can pick genes from any bacterial species to meet these requirements. Similarly, when you pack your carry-on, you know you need toothpaste, but the brand usually doesn't matter as long as you can brush your teeth.

When designing a minimal gene set, we use computational searches to find genes that can perform multiple functions in the cell. Including these genes in our list is like packing a toothbrush that can also comb your hair. Packing that multi-function item means you can take one less article in your carry-on.

The minimal cells we build exist only on a computer, but we may see a minimal gene set in a living bacterium in the near future. Gene cartographer J. Craig Venter of Synthetic Genomics, Inc., believes he is just a few experiments away from creating the world's first free-living artificial cell. While such an organism may not be minimal because it could house genes that are not necessary to support life, the project will benefit from the understanding of what *is* on the list of essential genes. A minimal cell model with a completely defined genome provides both a platform to test questions about how cells control themselves and a base recipe for modeling existing cells.

There are other questions we can answer using the minimal gene set. For example, bacteria depend on some genes more than others, just as some items in our carry-on suitcases are more critical than others. How upset are you if you get to the hotel and discover your shampoo is missing? How about your prescription medicine? If the minimal gene set contains *only* those necessary for a bacterial cell to survive, then we can measure the benefits of non-essential genes by adding them to the minimal cell model. Insights gained from such simulations could help revolutionize processes for producing life-saving medications or alternative energy sources by finding the best possible combination of genes to make what we need.

The computer model that will let us skip the baggage claim may be long in coming, but learning which functions bacteria need to pack in their genetic carry-on luggage will bring us closer to understanding what drives all life.

Brandon Wood– 2007 DOE CSGF Essay Contest Honorable Mention

Nature's Great Compromise

I hate making decisions. Aisle or window? Soup or salad? Top or bottom bunk? The world as we know it is fraught with choices, and I usually find myself wishing I could compromise and have a little of both. I've recently found comfort in discovering that nature, in its infinite wisdom, apparently agrees with me.

The choices I must make are similar to those humans make as they seek to understand the physical world by classifying matter into discrete categories. Mendeleev, for instance, organized the known naturally occurring elements into his "periodic table" based on similarities in a variety of properties, including whether the elements conduct electricity or occur naturally as solids, liquids or gases. Such taxonomy certainly bears merit, but it turns out that nature is not always so discerning. In fact, certain materials can successfully bridge multiple categories, foregoing classification to have a little of everything instead.

One especially curious type of matter that defies conventional description is a superionic. Superionic materials are generally composed of two or more elements, certain of which behave as conventional solids and others as conventional liquids. Imagine a wet sponge, with a solid matrix (the sponge) filled with a free-flowing liquid (water). In superionics, however, the interface between solid and liquid occurs at a scale of only one billionth of a meter. Such materials are vital components in just about every portable electronic device from laptops to cell phones. They also play important roles in automobile catalytic converters and next-generation fuel cells.

Despite their technological importance, however, most research into new superionics is by trial-and-error and generally involves creating permutations of existing technologies in hopes of creating new materials. (In fact, the superionic phenomenon was discovered by accident.) This is because superionic activity itself is poorly understood. The underlying atomistic mechanisms are just too complex and too tiny to probe using traditional means.

Enter computer simulations, which are not burdened by the same limitations as ordinary experiments.

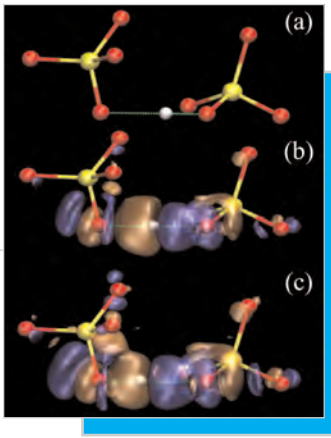

Regions of addition (orange) and depletion (blue) of electronic charge as a proton migrates, or hops, freely across sulfate (SO4) groups in cesium hydrogen sulfate (CsHSO4), a superionic fuel-cell electrolyte.

A carbon monoxide molecule binds to a superionic ceria surface in a catalytic converter similar to those used in cars and trucks.

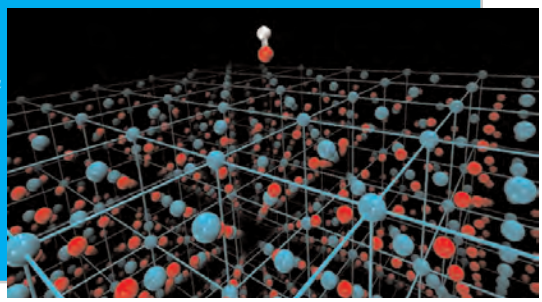

In order to truly predict what goes on inside a material, we need to understand how the atoms that comprise it behave. In matter, all atomic movements are governed by complex interactions between nuclei and the electrons that orbit them, both within each atom and between nearby atoms. If we could decipher these interactions and use them to compute all of the forces that push or pull each atom at each instant, we could reliably predict where and how each will move. This principle underlies an extremely powerful type of predictive computer simulation known as quantum molecular dynamics (QMD).

So how can one reliably compute these atomic forces? In doing so we must remember that systems at the atomic scale don't always play by the usual rules. Instead, particles that are very tiny and move extremely fast – notably, the electrons inside an atom – are governed by a very odd set of laws we call quantum mechanics. Like superionics, quantum objects embody nature's best-of-both-worlds mentality. In certain instances, a quantum object might behave much like any particle from our everyday experience, reacting to contact with a material by simply bouncing off it. But oftentimes a quantum object is more like a broad collection of ripples in a pond, lacking any concreteness and simultaneously occupying large sections of space. These quantum wave-objects can exhibit interference much like radio waves. They can even sometimes pass through seemingly solid barriers. In physics, this dual nature of quantum particles is referred to as the wave-particle duality.

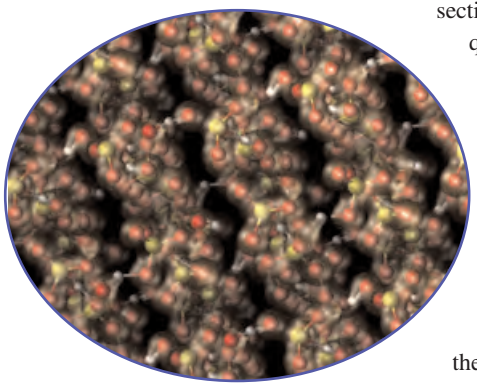

This visualization shows proton conduction channels in cesium hydrogen sulfate (CsHSO4), a superionic fuel-cell electrolyte.

Its conclusions may seem unimaginably counterintuitive, but quantum mechanics nevertheless offers a conclusive recipe for straightforwardly deciphering the complexities of the tiny electrons inside an atom. It also gives us a convenient formula for predicting the forces that act on an atom at any instant, making QMD possible. And since quantum-mechanical calculations are extremely accurate, running a simulation is every bit as good as running a real experiment. Moreover, simulations aren't limited by humanly accessible length- or time-scales. This means we can actually "observe" every atomic-level phenomenon, no matter how tiny or how fast. And since many of nature's puzzles have remained unsolved precisely because they involve mechanisms that are too small or too fast to be observed experimentally, the potential of QMD in such instances cannot be overestimated.

Despite innumerable recent advances in hardware and software, QMD is still computationally intensive: It is generally feasible to simulate atomic motions for only a fraction of a billionth of a second of real time. This may sound unimaginably short, but atoms in a material move so fast that many key events can be sampled within this duration, giving us valuable insight into material behavior. And given the astounding pace at which computing technology advances, simulation's role in the future of scientific research is assured.

So what of nature's fence-sitting superionics? Quantum molecular dynamics already has proven instrumental in deciphering why several known materials exhibit this unusual behavior. And although research continues, it seems certain that, armed with a proper understanding of the motivations for superionicity, we can at last abandon blind research in favor of directed, systematic improvement in device technologies. This is a truly remarkable feat, when one considers it's all happening inside a computer, without any real-world experiments.

Perhaps there is a lesson to learn in realizing that we, like nature, need not be bound by traditional dichotomies. After all, computer simulation successfully bridges the age-old gap between theory and experiment by employing the robust, systematic approach of experiment to draw conclusions that are grounded in the elegant simplicity of theory. I like to think of it as the theory-experiment duality. And in a quantum world where particles and waves, as well as solids and liquids, coexist, perhaps it's time to embrace a little indecision.

For more information, contact
the Krell Institute

WEB:
www.krellinst.org

PHONE:
515-956-3696

FAX:
515-956-3699

MAIL:
1609 Golden Aspen Drive
Suite 101
Ames, IA 50010

EMAIL:
csgf@krellinst.org

KRELL
institute