

A New Model for Pharmaceutical

*Supercomputing-
based modeling
may help validate
and accelerate
drug research.*

THE emergence in recent years of multidrug-resistant bacteria, which Livermore physicist Monte LaBute describes as “nature’s own bioterrorists,” underscores both the challenge of antibiotics development and the need for novel antibiotics. However, the pharmaceutical industry considers antibiotics in particular a poor investment because they are prescribed for only a short period, while drugs that treat a chronic condition such as high blood pressure may be prescribed for the remainder of a patient’s life. Drug companies often devote 10 to 15 years and more than a billion dollars to bring a new drug to market. This cost and time commitment can act as a powerful disincentive to developing new drug therapies.

Unless antibiotic development is revitalized, a growing portion of infections, particularly in hospitals, may become impossible to treat. In addition, bioterrorists could potentially introduce new bacterial threats, for instance, by releasing multidrug-resistant microbes in a large city. In the event of such an attack or of a sudden but natural emergence of a new bacterial strain, the current drug development process and the limited industry interest in antibiotic development would make a timely response difficult.

Pharmaceutical companies need a faster and more accurate way to identify promising drug compounds and evaluate the efficacy and safety of new drugs.

Research

Such a capability could reduce costs and risks, thus allowing companies to bring antibiotics and other drugs more quickly to market. The solution, according to Livermore researchers, may lie in supercomputer-based modeling and simulation.

While pharmaceutical companies do engage in computer-aided drug design, most only have access to personal computers or midsize clusters with a few hundred cores, which are not powerful enough to run comprehensive simulations in a reasonable time frame. The companies may also lack the expertise to use these tools effectively. Livermore computational biologist Felice Lightstone and her team in the Physical and Life Sciences Directorate are pioneering high-performance computing (HPC) modeling techniques that they hope will accelerate the development of medical countermeasures such as antibiotics.

In 2011, Lightstone's team collaborated with Trius Therapeutics, Inc., to develop what may be the first new class of antibiotics in 30 years. In the process, the team demonstrated how computer-based screening can minimize costly laboratory experiments and help shorten the chemistry phase of drug development from three or four years to just six months. Lightstone notes, "Our utility is in the speed and the number of calculations we can do. Sometimes the need is for

Computational chemistry efforts such as those by Felice Lightstone's team at Lawrence Livermore may help streamline the drug discovery process, thereby enabling researchers to bring new therapies to clinical trials and the marketplace more rapidly and with a higher rate of success.



many small calculations and other times for a massively large one. We have the computing resources and expertise at Livermore to do both.”

Through funding from Livermore’s Laboratory Directed Research and Development Program, Lightstone’s team is now in its third and final year of a project that shows promise in helping prioritize therapeutic candidates and mitigate the risk of failure in clinical trials. Computational biologist Sergio Wong explains, “Following our work with Trius involving the chemistry phase, or drug design phase, we realized that improving the success rate of drugs in clinical trials would be even more beneficial. Years of clinical trials and potential failures still lie ahead after the six-month design phase. Ninety-five percent of medicines studied in humans during clinical trials fail to be both safe and effective. Reducing that failure rate to even 70 percent would result in a huge savings in both cost and time for drug companies.” Lightstone’s project includes several thrust areas aimed at quickly zeroing in on which drug candidates are likely to be effective. Another focus is

to enumerate the side effects, known as adverse drug reactions (ADRs), of a particular drug.

Virtual Screening, In Parallel

Most drug compounds consist of molecules that work by binding with a protein receptor and either activating or inhibiting its behavior. Virtual screening is a computational technique used to identify which of the 30 million or more candidate drugs contained in publicly available chemical and pharmaceutical databases are most likely to bind to a targeted receptor. This technique can help to efficiently reduce the list of drug candidates to a manageable number for synthesis and testing.

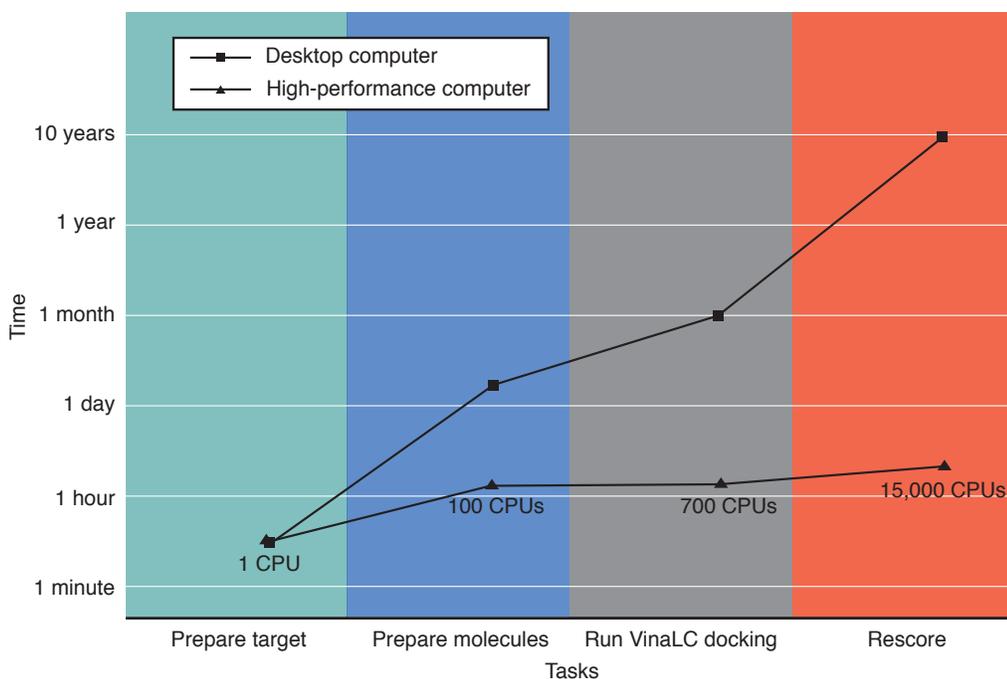
Virtual screening includes two main steps: molecular docking (see *S&TR*, April 2001, pp. 4–11; June 2002, pp. 4–10) and rescoring based on docking results. Postdoctoral researcher Xiaohua Zhang explains, “Molecular docking predicts which molecules are most likely to interact favorably with a particular protein receptor. Then a more accurate and computationally demanding

molecular mechanics method reevaluates the top-ranking combinations.”

Most molecular docking programs run on personal computers or small workstations and scale poorly to high-performance systems, thus limiting the number of molecules that can be screened within a reasonable time frame. To expand the use of this potentially powerful research tool, Zhang and colleagues have optimized a popular docking program for parallel computing. The resulting application—called VinaLC, where LC stands for Livermore Computing—implements a hybrid programming scheme to coordinate work and keep available processors busy. It also collects results from thousands of processes without overwhelming interprocess communication channels.

In tests on several Livermore supercomputers, including the 20-petaflop (quadrillion floating-point operations) Sequoia system, VinaLC demonstrated excellent scaling results. It easily outperformed the few other parallel docking programs that can run on HPC systems without sacrificing accuracy.

High-performance computing (HPC) systems can drastically reduce the current virtual screening time frame and increase the feasibility of more accurately screening extremely large compound databases. In this example, researchers used Livermore’s parallel docking program VinaLC to test about 40,000 molecules against a single drug target. They then used a more computationally intensive method to rescore (reevaluate) the top 20 docking poses. These high-fidelity simulations calculated the free energy in the system. (CPU is central processing unit.)



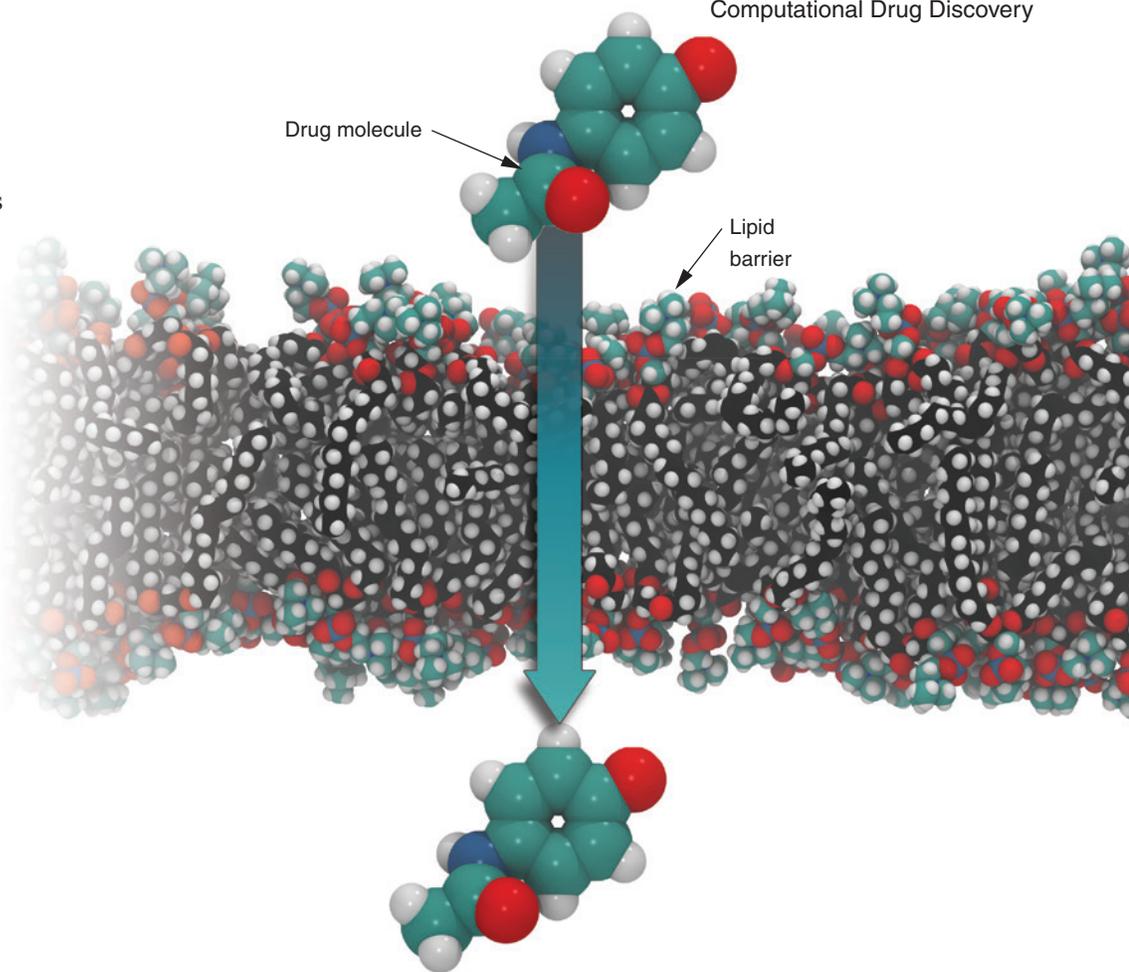
The code completed 1 million flexible compound molecular docking calculations in just 1.4 hours using 15,000 central processing units (CPUs) on Livermore's Sierra supercomputer.

Lightstone's team has also integrated VinaLC into its in-house, high-volume virtual screening system. This system automates and expedites molecule preparation, binding-site selection, docking calculations, and rescoring of the top pairings based on docking results. All applications in the system have been optimized for HPC. With Sierra, the researchers rescored a total of 700,000 compounds against 38 protein targets, in what may be the largest rescoring calculation to date. With such a large number of calculations, the team was able to statistically determine the optimal number of docking poses—the orientation of the compound with respect to the target—that should be kept for rescoring. The team found that 5 to 10 poses per compound provided a good compromise between accuracy and computational expense.

These efforts demonstrated that running calculations on a supercomputer dramatically reduces the time it takes to perform virtual screening, enabling the comprehensive evaluation of larger pools of candidate drugs. The team is now working to enhance the virtual screening system's speed, accuracy, and ease of use.

Selective But Predictable Barrier

The easiest way for drug therapy to travel from the bloodstream to tissues is through gaps between the cells in blood vessel walls. Blood vessels in the brain, however, have especially tight cell junctions, creating a highly selective barrier. Predicting whether a given compound will successfully travel from the bloodstream through these cells to reach the brain is a crucial part of drug design. Some drugs, such as those targeting Alzheimer's disease or depression, need to enter the brain to be



This simulation shows a drug compound crossing the lipid barrier between the bloodstream and the brain. Predicting such behavior could help drug researchers better gauge drug efficacy and catch potentially serious side effects at an early stage of drug development.

effective, but drugs targeting other parts of the body must be kept out, lest they cause unexpected and potentially serious side effects.

Testing whether a drug will cross the blood-brain barrier typically occurs five to six years into a drug development project, during in vitro or animal studies. Livermore researchers have been evaluating whether computer modeling can predict this behavior at a much earlier stage. In their modeling protocol, the barrier is represented by two leaflets of closely packed phospholipid molecules, which are surrounded by water. This thin membrane has a hydrophobic center and hydrophilic outer edges.

The researchers investigated the model's predictive capability by simulating

a chemically diverse set of 12 compounds, all of which have been well studied experimentally. The technique they used, called umbrella sampling, involves subdividing each drug's interaction with the phospholipid bilayer into 100 separate molecular dynamics simulations. Each 45-nanosecond simulation features a single drug compound at a slightly different position within the system. Together, the simulations follow a compound's progression through the bilayer and out the other side. The model applies a force to the drug to keep it close to the center of the simulation, although it is free to move laterally and rotate to its optimal orientation. Based on a measure of how much force is needed to keep the compound in position, the scientists can

estimate the relative energy values across the system for that specific compound and thus how likely the compound is to diffuse through the barrier.

Most of the drugs required more force to keep them in the middle of the bilayer than in the water. Computational biochemist Tim Carpenter, who is leading the simulation effort, says, “In general, we found that if the compound had a positive free energy in the middle of the bilayer, it would not cross the barrier. However, if the compound had a negative value, it would cross the bilayer and enter the brain.” Results correlated well with experimental data and also compared favorably with existing computational techniques, most

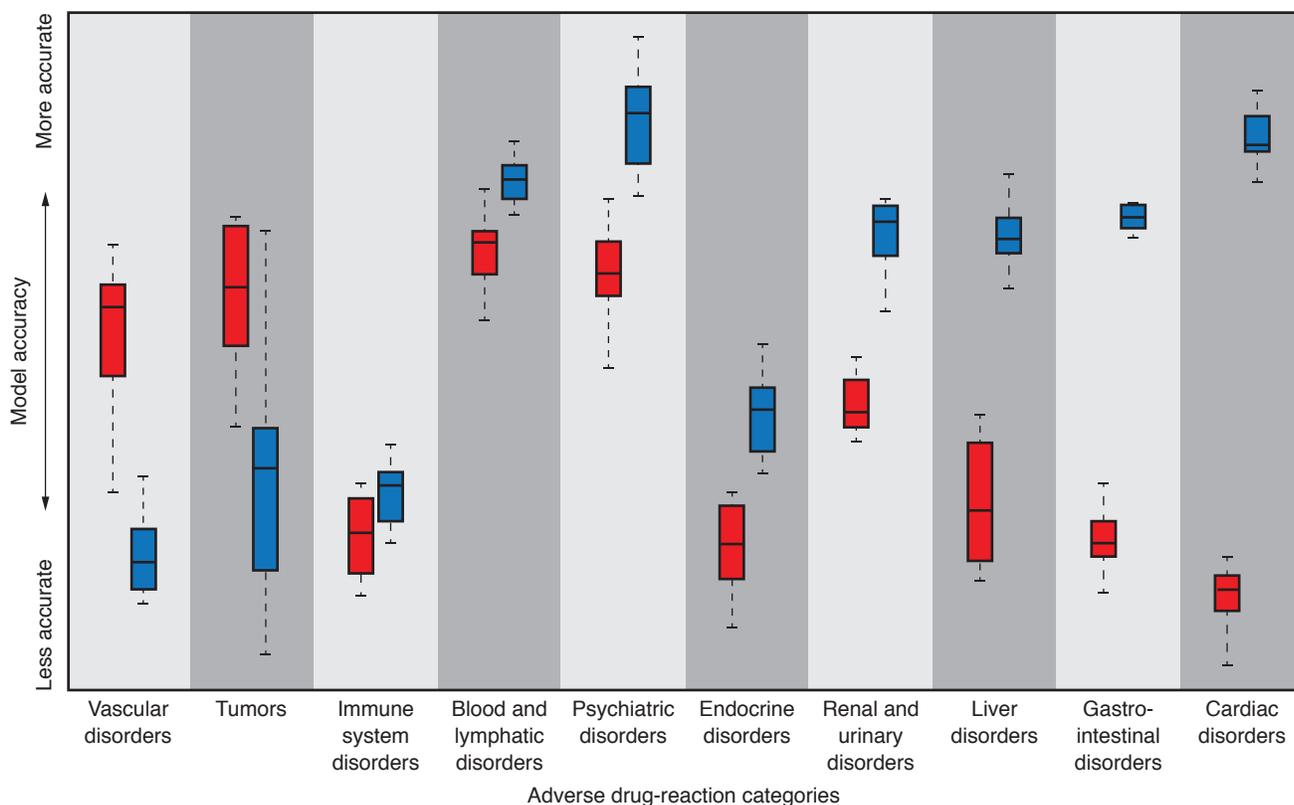
of which use empirical rather than first-principles methods.

Having proven the feasibility of their approach, the researchers now hope to develop a more realistic model, incorporating a greater variety of lipids. The current model simulates only passive diffusion of molecules. Future models may also incorporate active uptake, because some drugs mimic compounds the brain needs to fool certain proteins into pumping the compounds across the blood–brain barrier.

Searching for Unexpected Targets

“A drug introduced into the body interacts with many things in addition

to those intended,” says Wong. “It is the sum of these effects that gives us the clinical results we see.” Interactions may include off-target drug bindings, when the drug binds with more protein types than just the target protein, or more complex disruptions to a biological pathway, either of which could produce unanticipated ADRs. Enabling drug researchers to identify and evaluate the severity of these reactions early in the drug development process, when redesign is less costly and time consuming, is the goal of an ongoing collaboration between Lightstone’s team and LaBute. Together, they are exploring the feasibility of predicting ADRs with statistical modeling.



A Livermore-developed virtual adverse-drug-reaction (ADR) screening program integrates in-house molecular docking data with publicly available data sets. Models tested with VinaLC docking calculation data (red) and with experimental binding data (blue) are roughly comparable at predicting ADRs. This graph displays model quality averaged across 10 rounds of cross validation for 10 categories of serious or lethal ADRs, for a total of about 500 drugs. The VinaLC docking calculations evaluated about 500 protein receptors per drug.

The team's primary approach has been to use VinaLC to calculate binding possibilities between known drugs and all available human protein structures. Then using data on known ADRs for these drugs, the researchers are attempting to statistically associate reactions with binding events identified through molecular docking. Molecular docking should allow the team to spot more off-target effects than in vitro testing. "Pharmaceutical companies screen between a dozen and 50 receptors for ADRs in vitro, but that's a minuscule fraction of human proteins," notes LaBute.

For comparison purposes, the team is also analyzing correlations found between existing publicly available data sets, including those for drug structures, known ADRs, biological pathways, and proteins that cause side effects. Unfortunately, the available data have limitations. Public information on known ADRs consists mostly of on-target reactions and thus is not always a complete listing for a given drug. Also, Livermore researchers are focusing on serious or lethal ADRs. Those reactions have a clearer cause and effect and are more likely to cause a drug to fail. However, serious ADRs are underrepresented in the data because documented drug failures are often proprietary information—not available to the public.

LaBute has prepared and tested his models through cross validation, an iterative technique often used when data are sparse or incomplete. Modeling results so far are encouraging. The on-target data set was better at predicting certain types of ADRs, particularly more complex reactions such as endocrine and gastrointestinal disorders. The off-target molecular docking data were better at predicting other ADR categories, such as vascular disorders and cancers. In time, the off-target predictive approach may prove a good complement to existing ADR screening methods. The team is also investigating how to enhance predictions by using machine learning to identify new

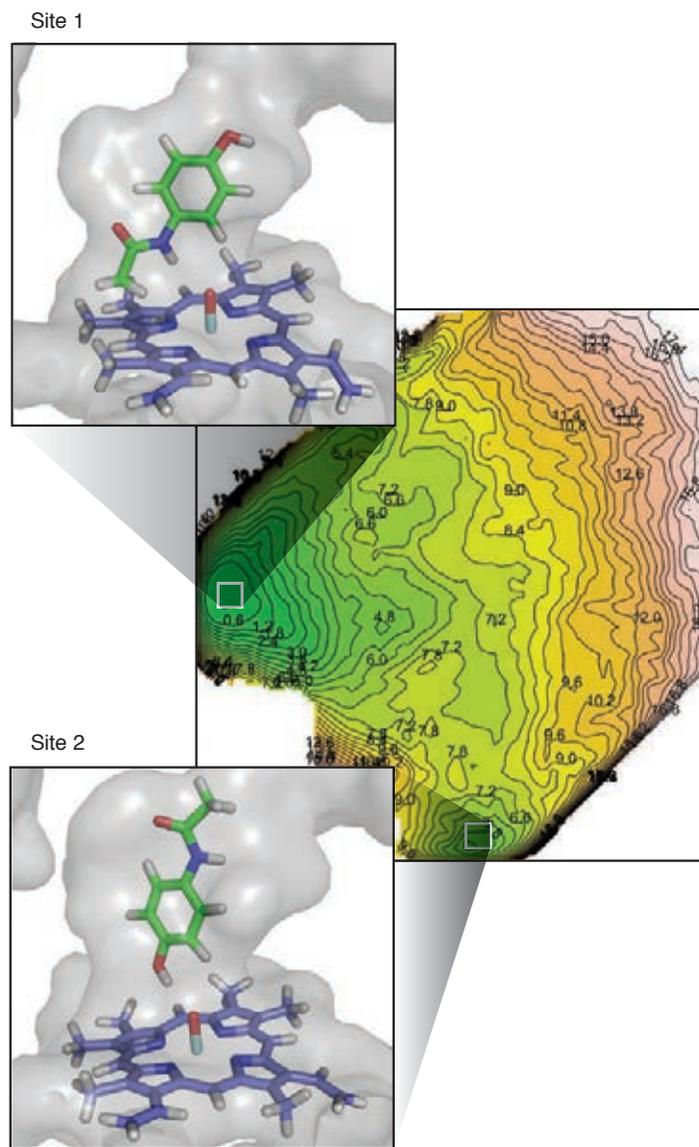
patterns in existing data sets or by mining the scientific literature for relevant information to expand the data pool.

Mapping Free Energy

When breaking down a drug, the body may produce one or more by-products, called metabolites. Some of these metabolites may be toxic, even when the drug itself is not. Postdoctoral researcher Yue Yang is leading an effort to use HPC for predicting at the early stage of drug discovery which metabolites will be produced and in what relative quantities. Yang's team has extensively modeled acetaminophen, a popular and well-studied pain reliever.

The cytochromes P450 (CYPs), a family of proteins, play a crucial role in breaking down drugs in the body. During acetaminophen metabolism, several CYPs produce mainly N-acetyl-p-benzoquinone imine (NAPQI), a toxic by-product associated with liver damage, while other CYPs produce mainly nontoxic metabolites. Livermore researchers have hypothesized that the reason different P450 members produce different metabolites must have to do with how and where the CYPs bind with acetaminophen during drug metabolism.

To explore this idea, Yang's team performed molecular docking calculations to find the top-ranked binding poses for five CYP protein types. With Sierra, the researchers then ran large-scale molecular dynamics simulations and two-dimensional umbrella sampling to determine the leading binding arrangements. Calculating binding



Livermore researchers calculated the most favorable binding arrangements and locations for several key proteins during acetaminophen metabolism to help determine which reactions produce which by-products. Free-energy mapping (right) and binding calculations (insets) show that the protein CYP2E1 prefers site 1, a configuration and binding location leading to N-acetyl-p-benzoquinone imine (NAPQI), a by-product of acetaminophen that can cause liver damage. Site 2, another possibility, would produce a nontoxic by-product; however, it has a higher free-energy value, so a reaction is much less likely.

probability across large binding sites such as those on the CYPs is computationally intensive. Each CYP required about 280 10-nanosecond simulations. The whole project took 800,000 CPU hours to complete.

With the simulation data, the team could analyze the binding free-energy landscape for each acetaminophen–CYP reactive complex and create a map of the spatial positions of the interacting molecules in the system along with the corresponding energy levels. A system usually seeks to achieve a minimum of free energy, so the lower the binding free energy, the more probable the binding mode. Using the maps, the researchers identified several of the most stable binding sites on each CYP and the metabolite each binding reaction would produce.

CYP2E1 was of particular interest. Although this protein is not expressed in great quantities in the human body, it is responsible for a large portion of acetaminophen metabolism and is the primary producer of NAPQI. Interestingly, CYP2E1's top-scoring pose from the docking calculations is one that should lead to a nontoxic metabolite, but the molecular dynamics simulations and the free-energy profile indicated that a different site and configuration were more energetically favorable. This alternate mode generates NAPQI consistent with experimental results.

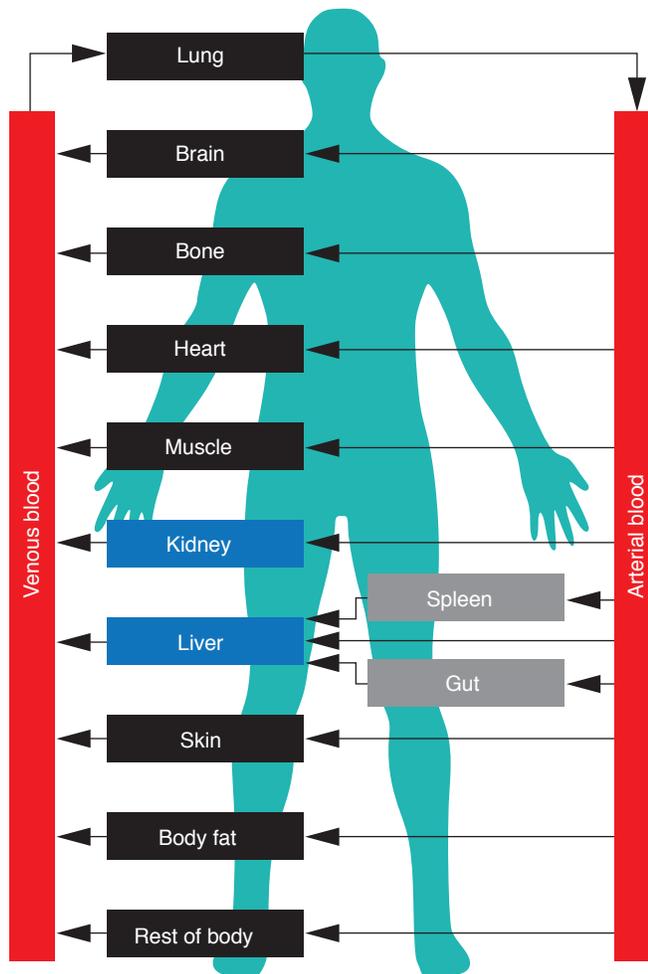
The study successfully reproduced the metabolites reported experimentally for each CYP, suggesting that the technique will be able to identify possible unwanted reactions in new drug entities. "This study shows that the relative binding free energy for the drug and different binding modes play an important role in determining the distribution of toxic and nontoxic metabolites," says Yang. "Previously, researchers looked only at the chemical reaction

barrier to determine the most probable metabolites. Now, we can more rationally predict toxic metabolites *in silico*, before drug testing, which is the most expensive step in the drug discovery and development process."

Existing Data, New Conclusions

Acetaminophen also served to validate a study that examined ADRs and drug metabolism on a larger scale. Lightstone and her team are developing

a first-principles kinetics model to better understand how drugs behave and change in the body, based on the drug's structure. This physics-based model will aid them in assessing efficacy and toxicity early in a drug's development. The model represents the body as a series of compartments (organs) linked by blood. It can be used to track the flow of a drug into and out of the compartments and determine how the drug interacts with specific proteins in each organ.



In Livermore's 14-compartment physics-based kinetic model of the human body, various tissues are connected by the circulatory system, and key processes in the human response to drugs are included. This model will help researchers assess toxicity risks early in the drug development process.

More broadly, the model enables scientists to see how processes such as absorption, desorption, distribution, metabolism, and elimination change the concentration and chemistry of a drug in the body over time. The researchers used the model to simulate how acetaminophen is metabolized in the body and whether certain behaviors—fasting, binge drinking, and chronic overuse—increase the likelihood of liver damage resulting from acetaminophen consumption.

How much damage a drug causes depends on its concentration and the length of time any toxic by-products of the drug stay in the system. A normal dose of acetaminophen, for instance, results in a small amount of NAPQI that can quickly be detoxified by an antioxidant called glutathione in the liver. In fact, the simulation showed that even as much as five times the recommended acetaminophen dose should be safe for healthy and well-fed individuals. However, both fasting and high levels of alcohol consumption can prevent glutathione from regenerating at a normal rate. For people whose glutathione supply is depleted, the simulation indicated that even small overdoses could cause NAPQI to accumulate in the liver, bind to various liver proteins, and cause damage.

Livermore's kinetic model was validated by comparing predictions with clinical outcomes. The model accurately replicated some existing clinical results and provided model validation, but it also shed light on metabolism of this common drug. Ali Navid, who led the study, notes that computational tools enable drug research that might be difficult or even unethical to conduct on patients. For instance, a researcher would never ask a patient to binge drink or overdose on acetaminophen simply for study purposes, but understanding the physiological and pharmacological

changes associated with these fairly common activities is crucial for toxicity prevention and treatment. Furthermore, without modeling insights, doctors may struggle to weigh the effects of patients' daily habits and to pinpoint why, for instance, a mild acetaminophen overdose causes liver damage in a particular patient.

Notes Navid, "That's the beauty of developing computational tools for pharmaceutical research—we can use information gathered by physicians and clinicians over the past 25 years and draw new conclusions from it. For example, we found that taking acetaminophen after an extended period of fasting results in more damaging side effects than taking it while consuming a bottle of vodka. These results aren't necessarily intuitive." Such findings may help clinicians settle a long-standing debate on the role of alcohol consumption in acetaminophen-induced liver damage. Next, Navid will use the same framework to simulate how drug–drug interactions such as ciprofloxacin and caffeine affect the body.

Multiscale and Multidisciplinary

The kinetic model, the cornerstone for Lightstone's project, will allow her team to integrate many different levels and types of simulations. Brain-barrier permeability estimates, metabolism site predictions, and mappings of off-target effects, for instance, will eventually feed into the kinetic model to gain a more dynamic understanding of how a drug compound will interact with the human body.

Enhancing the model is an enduring goal for the team. "We are a long way from simply uploading a person's genome into a computer program that tailors a therapy, such as was depicted in *Star Trek*," says Navid. "However, we may be as close as 25 years from using the structure of a drug to predict a patient-specific outcome. We have made large

strides in determining the blueprint of human beings and other organisms. We still need more knowledge of specific dynamic characteristics such as how environmental conditions or personal history affect biological interactions. These interactions form the foundations of our models and are crucial for predicting the outcome of therapeutic treatments."

As this project nears conclusion, other ideas that would marry biomedical research with HPC are gaining momentum at the Laboratory. One is an HPC for biology incubator, modeled on the successful hpc4energy incubator. (See *S&TR*, June 2013, pp. 4–12.) A biology incubator would partner computational biologists and biomedical companies with HPC resources to solve a problem or advance a company's research. Also, in a follow-on project to the successful Cardiod heart modeling (see *S&TR*, September 2012, pp. 22–25), Livermore is helping launch a national initiative to model the human body at multiple scales, which would greatly advance digital drug and biological countermeasure development.

"Throughout our project, we've drawn on such areas as machine learning, chemistry, biology, and physics," says LaBute. "This multiscale, multidisciplinary project was well suited for a national laboratory."

—Rose Hansen

Key Words: adverse drug reaction (ADR), blood–brain barrier, computational chemistry, cytochrome P450 (CYP), drug development, high-performance computing (HPC), kinetic model, molecular docking, molecular dynamics simulation, N-acetyl-p-benzoquinone imine (NAPQI), Sequoia, Sierra, supercomputer, umbrella sampling, virtual screening.

For further information contact Felice Lightstone (925) 423-8657 (lightstone1@llnl.gov).