

# The future of molecular dynamics simulations in drug discovery

David W. Borhani · David E. Shaw

Received: 22 November 2011 / Accepted: 5 December 2011 / Published online: 20 December 2011  
© The Author(s) 2011. This article is published with open access at Springerlink.com

**Abstract** Molecular dynamics simulations can now track rapid processes—those occurring in less than about a millisecond—at atomic resolution for many biologically relevant systems. These simulations appear poised to exert a significant impact on how new drugs are found, perhaps even transforming the very process of drug discovery. We predict here future results we can expect from, and enhancements we need to make in, molecular dynamics simulations over the coming 25 years, and in so doing set out several Grand Challenges for the field. In the context of the problems now facing the pharmaceutical industry, we ask how we can best address drug discovery needs of the next quarter century using molecular dynamics simulations, and we suggest some possible approaches.

**Keywords** Drug discovery · Molecular dynamics simulations · Grand challenges · Free energy

## Introduction

Understanding the mechanisms underlying the behavior of chemical and biological systems requires scrutiny at spatial and temporal resolutions that challenge current experimental techniques. Molecular dynamics (MD) simulations

are being increasingly partnered with experiments in this quest because simulations can track system behavior across a vast spatiotemporal domain—length scales up to thousands of ångströms, with atomic precision, and timescales up to milliseconds, at femtosecond resolution. This power of simulations has been further increased by recent methodological advances. Here, we predict what the next 25 years of MD simulations may bring, especially regarding their application to the search for new drugs.

Novel computational methods, including MD simulations, have assumed an ever growing role in drug discovery over the past quarter century. Yet, despite having learned and contributed much, we face many challenges ahead. To take novel computational methods to the next level—such that they radically alter the very landscape of drug discovery—we must grapple with those challenges and rise above them.

This essay is meant to be thought provoking—we raise more questions than we answer. It is arranged such that the knowledgeable reader can easily skip to the parts of interest. We begin by reviewing the challenges, using any technique, including computational ones, in finding new drugs. Next, we compare the computational state of the art of 25 years ago with that of today, with a particular emphasis on MD simulations. Following brief comments on the nature of innovation and the art of making predictions, we make a series of straightforward predictions on the future directions of MD simulations, including particular ways simulations might be used in drug discovery; some key methodological improvements we believe will be needed for success are discussed. We then turn our attention to several “Grand Challenges” for the field, including a rather audacious goal on how simulations might be used in drug discovery. We conclude with our thoughts on how MD simulations are perceived by the larger scientific community, and on the importance of setting goals.

---

D. W. Borhani (✉) · D. E. Shaw (✉)  
D. E. Shaw Research, New York, NY 10036, USA  
e-mail: David.Borhani@DEShawResearch.com

D. E. Shaw  
e-mail: David.Shaw@DEShawResearch.com

D. E. Shaw  
Center for Computational Biology and Bioinformatics,  
Columbia University, New York, NY 10032, USA

## What makes finding a drug so difficult?

Computational chemistry methods have become deeply integrated into drug discovery over the past 25 years [1, 2], expanding significantly beyond early work on quantitative structure–activity relationships (QSAR) [3], computerized chemical structure representation [4], and computerized compound, reaction, property, and structure–activity databases [5]. Indeed, computational methods have become so much part of the very fabric from which new drugs are woven—simply ascertaining their specific impact is quite difficult—that their continued role in drug discovery has been questioned as the pharmaceutical industry faces looming scientific and economic challenges. It is thus prudent to ask whether we are focusing our efforts where they yield the most benefit. Are we addressing the key challenges facing the pharmaceutical industry?

Productivity is one of the most significant challenges facing the pharmaceutical industry, with \$50 billion spent annually to produce only ~20 new drugs [6]. Ever-shifting organizational structures hamper success [7–9], but the root problem is simply that most new drug projects fail—only ~3% of projects ever produce a marketed drug. It is crucial to understand at what stage in the pipeline projects fail, for what reasons, and then to ask how computational approaches can help avert those failures.

Drug discovery—the focus of most computational chemistry efforts—is tolerably successful but needs improvement: about 35% of discovery projects succeed, on average, in delivering experimental drugs ready for clinical testing. The project stages of target identification and screening, hit-to-lead, lead optimization, and preclinical candidate selection have individual success rates ranging from 69 to 85% [10]. When discovery projects fail, they fail for diverse reasons, notably unclear target biology, lack of appropriate leads, poor potency or selectivity, inappropriate drug-like properties, lack of efficacy, and unexpected animal toxicity.

Clinical development success rates, however, present a stark contrast, as do the more uniform reasons for clinical failure. An experimental drug entering a Phase I clinical trial has only a 10% chance of reaching the market: the success rates in each of the three trial Phases and for final regulatory approval are just 54, 34, 70, and 91% [10]. As to the reasons:

Lack of clinical efficacy has meanwhile become the most frequent cause for discontinuation of a drug development program. Consequently, attrition rates are highest in clinical Phase II, which usually includes the first evidence for pharmacodynamic action of the compound or, proof of concept. [11]

Two-thirds of recent Phase III failures are due to inadequate efficacy [12], as are more than half of Phase II [13] and ~16% of Phase I [14] failures. Toxicity and business-related failures (which often are due to inadequate efficacy, e.g. compared to competitive or existing drugs) contribute, but much less so. Poor pharmacokinetic properties, once a major issue, now account for just ~10% of clinical failures, mostly in Phase I [15], a gratifying result of the increased attention paid to these critical properties during lead optimization. The bottom line is this: All experimental drugs enter human clinical trials based on extensive pre-clinical data indicating that they *should* work; most nonetheless do not, defying our well-grounded expectations. The complexities of human biology, amplified by the limitations of the reductionist paradigm of target-based drug discovery [16], thus appear to be our industry's largest challenge.

We believe that computational chemistry methods—and in particular MD simulations—should reasonably be expected to significantly impact the trajectory of the pharmaceutical industry. Better success in clinical trials will come, in part, with an increased understanding of human biology, and simulations will increasingly make useful contributions here. Arguably, however, we should continue to focus our greatest efforts on early drug discovery: The problems there align especially well with potential computational solutions, and addressing those problems will reduce the significant resources—about 15 discovery projects on average—devoted to achieving a single product launch. Indeed, despite discovery enjoying triple the success rate of development, parametric sensitivity analysis highlights lead-optimization costs as the third-most important factor (after Phase II and III trial success rates) that dictates overall success in bringing a drug to market [10], an argument supported as well by a related analysis [17]. Improving our ability to select and design better molecules at all discovery stages, including lead optimization, is an achievable and valuable goal. We believe this discovery focus, by helping both to reduce the needed resources and to increase clinical candidate quality, may also bring significant indirect savings: the really big payoff may be a resultant improvement in clinical trial success rates.

## 25 years ago in *JCAMD*

What was the state of computational chemistry in drug design 25 years ago? Perusing early *JCAMD* issues reveals, perhaps surprisingly, that many of the pressing questions then are still of interest today, and many methods then new have become our methods of choice. What has changed—dramatically—is both our confidence in, and our

ability to execute swiftly, the algorithms underlying these computations.

Ligand-based approaches were used, for example in the three-dimensional pharmacophore modeling of benzodiazepine receptor ligands [18], and to design nicotinic agonists using a shape matching algorithm [19]. Quantum mechanics (QM) calculations revealed an angiotensin-converting enzyme inhibitor QSAR [20]. The solution conformational energies of apomorphine analogues were correlated with their biological activities [21]. Many such conformational analyses used molecular mechanics (MM), for instance Allinger's MM2. MD simulations had progressed from the ground-breaking 9 ps *in vacuo* simulation of bovine pancreatic trypsin inhibitor [22] to encompass solvated proteins, lipids, and ion channels [23]. The determination of free energy differences using MD simulation-based thermodynamic cycle integration was reviewed [24]; within a year, free energy perturbation was used to compute the relative binding free energies of an antiviral compound to wild-type and drug-resistant human rhinovirus [25]. All of these methods, in their original or in enhanced forms, are widely used today; MD simulations in particular have advanced dramatically.

### Molecular dynamics—the current state of the art

Molecular dynamics simulations are used today to study nearly every type of macromolecule—proteins, nucleic acids, carbohydrates—of biological or medicinal interest. Simulations span wide spatial and temporal ranges and resolutions. In explicit, all-atom MD, thousands to millions of individual atoms representing, for instance, all the atoms of a protein and surrounding water molecules, move in a series of short (e.g., 2 fs), discrete time steps. At each step, the forces on each atom—determined from the “force field,” a collection of physics-based parameters that represent both bonded and non-bonded (e.g., van der Waals) inter-atomic forces—are computed and the atomic position and velocity are updated according to Newton's laws of motion [26]. This process is repeated billions of times to provide continuous atomic trajectories lasting as long as 1  $\mu$ s, or even longer. The examples below indicate some of the current capabilities of MD simulations and the insights they can provide. Additional examples may be found in a recent review on the use of MD simulations in drug discovery [27].

The biological systems studied using all-atom MD simulations can be very large, comprising millions of atoms. For instance, several such simulations of bacterial ribosomes—the pivotal RNA/protein complex that is the target of diverse antibiotics—have been carried out. In a recent example, comprising  $\sim 3.2$  million atoms, the

“accommodation” motion of the ribosome that allows aminoacyl-tRNA binding was studied [28]. Simulations of satellite tobacco mosaic virus (STMV;  $\sim 1$  million atoms) recapitulated the known stability of the complete virus and of the RNA core particle [29]. The simulations further indicated that empty STMV capsids exhibit a pronounced instability, a new finding that explained the failure of experimental efforts to prepare such empty capsids. An impressive effort to further increase the scale of MD simulations is now using 100 STMV particles—100 million atoms—as the test system [30].

MD simulations are well-suited to the study of membrane proteins, which present particular challenges for experimental methods. For instance, the control of ion channel conductance, so-called “gating,” has been studied for many channels, among them the nicotinic acetylcholine receptor. Beckstein and Sansom used MD simulations and potential of mean force calculations to determine the free-energy barrier to ion passage through the central pore of the nicotinic receptor [31]. They found an  $\sim 10$  *kT* barrier to ion passage in the constricted state, in which the hydrophobic central pore is dewetted, sufficiently high to account for effective channel closure. Their mechanistic results likely apply to the entire “Cys-loop” superfamily of ligand-gated ion channels, of which the nicotinic receptor is a much-studied prototype; these results also heralded hydrophobic gating in the structurally distinct voltage-gated ion channel superfamily [32].

Free energy calculations of ligand–receptor binding is a natural application of simulations in drug discovery. Several approaches have been used. Grand canonical Monte Carlo simulations can identify both potential ligands and their binding site(s) on the drug target [33]. In essence, the target is flooded with ligands, or more typically small fragments, which are then slowly “evaporated,” leaving behind only the most tightly bound ligands. This method has proven successful in a few cases, for instance, in the design of novel nanomolar inhibitors of p38 kinase [34].

High throughput molecular mechanics with Poisson–Boltzmann surface area (MM-PBSA) was used at Abbott Laboratories to directly estimate relative binding free energies for 308 ligands drawn from three representative drug discovery projects—the protease urokinase, the phosphatase PTP-1B, and the kinase Chk-1 [35]. The results were encouraging both for the number of ligands evaluated and the target scope, but the moderate correlations between predicted and experimental binding free energy values ( $r^2 = 0.52$ – $0.69$ ) suggest that the fast MM-PBSA method is insufficiently accurate, by itself, to guide a medicinal chemistry program.

Alchemical methods yield improved quantitative results at the cost of significantly more computation [36]. The Jorgensen group's work on non-nucleoside HIV reverse

transcriptase (HIV-RT) inhibitors (NNRTI) is a notable recent example. Compound design decisions were based, in part, on calculated estimates of binding free energy differences, determined using free energy perturbation with Monte Carlo sampling, among various Cl-substituted test compounds [37]. This optimization strategy has provided novel aminotriazines, possessing cellular  $EC_{50}$  values below 10 nM, effective against both wild-type HIV-RT and the resistant Tyr181Cys variant [38]. Using the same method, a 5  $\mu$ M virtual screening hit was transformed into a 55 pM inhibitor, apparently the most potent NNRTI reported to date [39]. The results obtained thus far on HIV-RT are quite encouraging, and the utility of this approach in other systems is an area of active investigation.

The process by which drugs bind to receptors has been studied in several systems. Benzamidine bound spontaneously to trypsin in MD simulations, achieving a good match to the crystal-structure-defined pose and revealing the binding pathway [40]. The unbiased binding of kinase inhibitors [41] and G protein-coupled receptor (GPCR) agonists [42] and antagonists [43] has also been demonstrated. For example, the endogenous cannabinoid *sn*-2-arachidonoylglycerol was found to enter the binding pocket of a CB2 receptor homology model from the lipid bilayer [42]. Notably, simulations of several  $\beta$ -blockers and a  $\beta$ -agonist binding to two  $\beta$ -adrenergic receptors revealed where along the binding pathway dehydration of the ligand and receptor—long known to a major source of ligand affinity—occurs [43]. The work further hinted that dehydration presents an unexpected kinetic barrier to binding, leading to suggestions on how ligand/receptor dehydration might be modulated to affect drug binding and unbinding kinetics.

Several related techniques leverage the power of atomistic MD simulations, extending the range of problems that can be studied. Coarse-grained simulations allow one to sacrifice spatial detail to achieve longer, more biologically relevant timescales, thereby enabling the study of processes that currently are too slow, or of systems too large, to study with atomistic simulations [44, 45]. This approach enabled, for example, the simulation of the assembly of apolipoprotein A-I and lipids into discoidal high-density lipoprotein (HDL) particles [46], and similarly the self-assembly of membrane proteins into lipid bilayers [47]. Several distinct, atomistic approaches aim to accelerate the still-insufficient sampling of protein conformational states, for example metadynamics [48], accelerated MD [49], and temperature accelerated MD [50].

Our group has recently shown, using a specialized supercomputer designed especially for MD simulations, named Anton [51], that all-atom MD simulations can now reach timescales on which much interesting biology occurs. One millisecond-long, continuous single-trajectory

simulations of small globular proteins, for instance bovine pancreatic trypsin inhibitor (58 amino acids) [52] or the fast-folding *N*-terminal fragment of ribosomal protein L9 (39 amino acids) [53], have been performed. Such simulations take a few months of elapsed time. Much larger systems (e.g., receptor tyrosine kinases, GPCRs, or voltage-gated ion channels, all embedded in lipid bilayers, totaling  $\geq 10^5$  atoms) can be simulated for hundreds of microseconds, with aggregate simulation times  $>1$  ms.

These simulations have demonstrated the *de novo* folding of proteins, long recognized as an important problem in biophysics [54]. Initial work has focused on fast-folding proteins, for instance the WW domain protein Fip35; WW domains, which comprise a three-stranded  $\beta$ -sheet arranged as two  $\beta$ -hairpins, bind proline-rich sequences. Fip35 folds with an experimental time constant of 14  $\mu$ s [55], making it the fastest-folding WW domain known when our work began; this rapidity has made it an attractive simulation target [56, 57]. In our simulations of Fip35—initiated from the extended state—the protein achieved the folded state, with a backbone root-mean-squared deviation (RMSD) of  $\sim 1$  Å from the crystal structure [52]. The simulations were carried out under conditions where the folded and unfolded states exist in reversible equilibrium; repeated folding/unfolding barrier crossings followed a single well-defined pathway, with kinetics that closely match experiment. Elucidation of the folding transition state allowed an even faster-folding Fip35 variant to be designed, which was subsequently confirmed experimentally [58]. These folding results have been extended to encompass 12 small proteins of diverse structure— $\alpha$ -helical,  $\beta$ -sheet, and mixed  $\alpha/\beta$ —with 8 of the 12 proteins reaching RMSD values less than 2 Å from the respective crystal structure [53]. It is noteworthy that all 12 of these folding simulations used a single, physics-based molecular mechanics force field—a modified version of the CHARMM force field [59]—indicating an increased level of accuracy that enables simulation of large conformational changes.

### The art of making predictions

Making predictions is hard, and, as Niels Bohr famously said, it becomes especially difficult when we seek to predict the future. The varied forms that technological advances take suggest why this is so: Innovations range from obvious extrapolations beyond current practice, to unexpected new technologies or experimental “We’re not sure what this is good for...” ideas, to active, purposeful invention of the future. Innovations also vary dramatically in impact. Most are incremental advances, with predictable effects. The effects of “disruptive” innovations, which

upset the way things are done, are harder to predict. And a very few innovations are truly “revolutionary”—they alter life in manifold, unforeseeable ways.

A brief history of the computer age illustrates these points. Computers permeate modern life in ways that would rightly be viewed as revolutionary a century ago. The underlying technology, the transistor, was itself the result of experimentation and purposeful invention. The transistor is both a disruptive innovation—a new approach to electronic switching that displaced the relays and vacuum tubes of early computers—and a revolutionary innovation: It enabled creation of the integrated circuit, thus launching the dramatic “Moore’s law” increases in computational power. This trend, which relies upon continuous innovation, often incremental, and an intentional drive to achieve a now self-fulfilling prophesy, has enabled the tremendous computational advances upon which all areas of human endeavor—including computer-aided molecular design—have become so reliant.

This history exhibits both the determination of Edison—he said the light bulb was “one percent inspiration, ninety-nine percent perspiration”—and Alan Kay’s belief that the best way to predict the future is to invent it. For our task here, the import of this history is that it seems largely obvious in hindsight, and yet it would have been very difficult, if not impossible—say, for Shockley, Brattain, and Bardeen, in 1947—to predict.

### Molecular dynamics—predictions for the next 25 years

Computer power has tremendously increased in the past 25 years. Many studies from 1987 were performed on Digital Equipment Corporation VAX 11/780 computers (1978,  $\sim 0.0001$  GFLOPs); the rhinovirus study mentioned above [25] used a Cray X-MP (1982, 0.4 GFLOPs). Today, personal computers far more powerful are commonplace (2011 Intel Core i7,  $\sim 110$  GFLOPs), and most MD simulations are run on commodity clusters (*tera*FLOPs). Supercomputers such as IBM’s Blue Gene/L (2005, 0.3 *peta*FLOPs), deliver greater performance on diverse tasks. Tailoring the hardware to a specific task—MD simulations, for instance—enables even higher performance: the Anton supercomputer we created [51] increases the speed of individual MD simulations by nearly two orders of magnitude.

Computational power in the year 2037 may be as much as one million-fold greater than it is today. We make several assumptions: Continuation of Moore’s law—not unreasonable in light of current industry 15-year projections and a 50-year history of surmounting technological hurdles—suggests that processor computational power may increase as much as one thousand-fold. Enhanced

processor integration, and architectural and software advances (including MD-specific algorithmic improvements), will yield further increases. Using ever more processors in parallel will compensate for limits in individual processor performance; such computers, like Anton, will increasingly favor very much larger, rather than very much longer, simulations. Given these considerations, what are some implications of such large increases in computational power, both for MD simulations in general, and for the use of simulations in drug discovery?

Simulations will become much larger and will reach longer timescales. Today’s not atypical simulation size, a box  $100 \text{ \AA}$  on each edge—a volume of  $10^{-6} \mu\text{m}^3$ —might scale up to  $1 \mu\text{m}^3$ , the volume of *Haemophilus influenzae*, a small bacterium. Of course, simulating an entire bacterium for a millisecond or so probably wouldn’t teach us much, in part because of the limited diffusion of individual macromolecular assemblies on this timescale. Smaller simulations comprising a substantial fraction of a cell (e.g., RNA polymerase with associated transcriptional factors), however, on biologically significant timeframes such as one second (i.e., long enough to transcribe a small gene), may be feasible. This reach for ever larger and longer simulations will also be increasingly aided by improved algorithmic methods to increase the sampling of conformational space.

We predict that we will determine, computationally, the three-dimensional folded structure of proteins from their amino acid sequence. Put another way, we will be able to observe the Central Dogma of molecular biology—DNA  $\rightarrow$  RNA  $\rightarrow$  protein—using simulations. Atomistic MD simulations have already demonstrated today the *de novo* folding of small (up to 80 residues) protein domains. The effect of force field quality on folding was mentioned [59]; it seems likely that continuous force field improvement (see below) will enable further progress. With those caveats, extrapolation suggests that folding of more typical ( $\sim 300$  residue) single-domain proteins ( $>10$  ms simulations of  $\sim 10^5$  atoms) will be feasible within 10 years. Folding of large, multi-domain proteins—for instance,  $\beta$ -galactosidase, comprised of four 1,024-residue, five-domain subunits—will likely be accessible within 25 years. Beyond delivering folded structures, useful in their own right, especially in drug discovery, simulations will help us to understand basic folding mechanisms. Our hope is that by folding many proteins using a physics-based approach, these simulations will “provide the data for developing abstract models at a conceptual level that describe general and unambiguous features of the protein-folding mechanisms” [60]; brute force folding simulations might then no longer be needed, thereby increasing our capabilities even further.

A benefit of our ability to fold proteins will be that structure-guided drug design can be extended to new

targets. Some will simply be proteins for which no crystal or NMR structure is available, or for which the available structures are perhaps not in a biologically relevant form. More significant, perhaps, will be those mis-folded or aggregated proteins—especially prions, amyloidogenic proteins, and intrinsically disordered proteins (which often partially fold on binding a partner)—that are thought to underlie major diseases, for instance Alzheimer's disease. The tractability of targeting these proteins will be increased if we can, for example, observe computationally both the atomic details of their (mis)folding and the modulating impact of candidate drug molecules.

Of great interest will be the computational assembly of interacting macromolecules. It is increasingly clear that macromolecule–macromolecule interactions drive much of biology: the sheer number of pairs, or larger assemblies, of macromolecules—and hence their regulatory capacity—far exceeds the limited number of individual human gene products (<25,000). In accord with this notion, the unbiased assessment of protein–protein and protein–nucleic acid interactions, in particular, will present new drug discovery opportunities. Nature seems to have already leveraged this key aspect of biology—many natural product drugs and signal mediators bind at macromolecule–macromolecule interfaces [61]—and we are beginning to catch on with some synthetic and natural product-derived antibiotics and anti-cancer agents that bind at, and stabilize, interfaces among and between proteins and nucleic acids. MD simulations can address two aspects of this problem. First, how and where do the partners interact? And as they interact, are novel ligand binding sites created, at the interface or at allosteric sites? MD simulations seem to be good at identifying low-energy protein conformations that harbor cryptic drug binding sites [27]. Second, how can small molecule ligands modulate those interactions? How can we optimize their binding and drug-like properties?

Simulations of macromolecular assembly will also extend to include very large complexes, such as the nuclear pore complex or even entire organelles. It should be possible, for example, to observe the passage of cargo proteins as they transit through the nuclear pore. Intentional mis-assembly of complexes, for instance of bacterial flagella or viral capsids, again presents opportunities for the targeting of new drugs. We also think that simulations of enzyme catalytic cycles, complete with bond making and breaking, will become much more common, enabled by methodological improvements that accelerate QM/MM simulations.

Putting these predictions into practice

These predictions—really just straightforward extrapolations—rest on a few key assumptions. We mentioned raw processor computational power above. The semiconductor

industry, having hit some fundamental barrier, may at some point cease to deliver. Gordon Moore has bet against his own law several times; he has, however, each time been wrong. We think it likely that the innovative spirit of device architects, coupled with one or more cutting edge technologies that are published in science journals today but will be discussed in engineering journals tomorrow, will prevail. More within the control of the computational chemistry community are some clear improvements we need to make, or in some cases adopt now, in how we carry out MD simulations, especially with regard to force fields.

Lou Allinger recently made a compelling case for carefully crafted small-molecule molecular mechanics force fields, such as MM4, in the prediction of molecular structure [62]. We believe that MD simulations would benefit from use of analogous, necessarily more-complex force fields. MM4 is a so-called “Class 3” force field—it includes all significant off-diagonal force matrix (“cross”) terms—making it more complicated, and more realistic, than the typical, “Class 1” diagonal force fields used in MD simulations. The MMFF94 force field, widely used for drug-like molecules because of its broad parameterization [63], is similar in spirit to MM4, although it uses point charges (like most widely used MD force fields) rather than induced dipoles. Although we have been able to simulate—on millisecond time scales—processes as complex as small protein folding and ion permeation through channels using (tweaked [59]) Class 1 force fields, we will reach a point soon where we must be more realistic in our modeling. Two salient examples are how force fields handle nucleic acids and (especially divalent) cations. Another example is modeling of cation- $\pi$  interactions [64]. These interactions—driven by electric quadrupole moments and polarization effects—are now recognized to be quite important to both protein structure (e.g., arginine–tryptophan ladders) and protein function, for instance, in a wide variety of protein–ligand interactions [65]. Similarly, London dispersion forces between hydrogen atoms—which appear to contribute significantly to the energetics of branched and strained alkanes [66]—may impact ligand–receptor binding energetics as well.

Polarizable models should enable us to more accurately describe (inter)molecular interactions, and indeed several polarizable force fields have begun to demonstrate their value. The AMOEBA force field, for instance, includes some cross terms and, significantly, polarizable atomic multipoles (up to quadrupoles) replace fixed partial charges [67]; the QMPFF3 force field uses similar functional forms [68]. The two differ in parameterization strategy (chemically sensible groups versus individual atom types, respectively). AMOEBA-based simulations performed well in the recent SAMPL2 hydration free energy challenge sponsored by OpenEye [69]. Simulations using QMPFF3

have performed with impressive accuracy in modeling aromatic–aromatic interactions in the gas, liquid, and solid phases [68], and most notably in ligand–receptor relative free energy binding calculations (average  $r^2 = 0.9$ ) [70]. Force fields such as these will see widespread use.

Here are a few other possible areas for improvement, which in principle should improve MD simulation accuracy:

- Simple fixes—for example, improved van der Waals combining rules, such as those of Waldman and Hagler [71]—should be evaluated, as should replacement of point charges by smeared charges of some form (e.g., multipoles, exponentials, Gaussians). The use of more accurate van der Waals and charge models may remove the need for complicated, and effectively arbitrary, torsional potentials.
- Constant pH simulations should become standard. The current approach—multiple simulations launched from distinct, unchanging protonation states—models reality poorly. Both implicit solvent (e.g., [72]) and explicit solvent (e.g., [73]) approaches currently afford comparable results for test proteins such as hen egg white lysozyme. The explicit approach may prove more robust in the end. Nature takes advantage of the diffusion of protons through water [74], and so should we.

### Molecular dynamics—grand challenges

Grand Challenges are goals that, if achieved, will have revolutionary impact. We present here several such goals we believe to be worthy of significant effort.

Free energy calculations must become reliable and rapid, for both macromolecule–ligand and macromolecule–macromolecule interactions. The importance of free energy to every aspect of drug discovery cannot be overemphasized [75]: free energy dictates, for instance, the strength of interactions, accessible macromolecule and ligand conformations, drug binding both to targets and to anti-targets, and passive and active drug transport properties. For free energy calculations to be of consistent use in a drug discovery environment—particularly for quantitative binding predictions during lead optimization—we will need to achieve an accuracy equal to or better than typical experimental binding or activity assays, that is, correct to within a factor of two ( $\sim 0.4$  kcal/mol). These calculations also need to be turn-key; automated methods must provide ligand parameters of quality equivalent to those of the target (macromolecule) force field.

Improving free energy calculations has been hard because the two key issues—insufficient sampling of configurational space, and inadequate force fields—are

impossible to test independently of each other. If the calculation has not converged, how can one say the force field is at fault, and vice versa? Sampling and convergence will naturally increase with longer simulation times, but other, more clever approaches may prove useful or even necessary. The more advanced force fields of today appear to work well in limited cases [67, 70], but their generality, especially in protein–ligand binding, remains unproven. We believe that the force field improvements mentioned above, or others of a similar or completely novel nature, should help significantly. Converged calculations will enable rigorous determination of both force field accuracy and the need for specific force field improvements. Chodera and co-authors recently suggested that although the simulation “field has been extraordinarily productive in generating new algorithmic ideas and advancing technologies to facilitate the development of more accurate force fields, it has failed to produce an effective set of tools for the design of small molecules. To do so, it is necessary for the field to begin a shift from a research focus to an engineering focus” [76]. We agree. Blind tests such as the OpenEye SAMPL challenges [69] will continue to be especially useful, because they allow us to gauge our successes and failures in an unbiased manner.

A consequence of accurate free energy calculations: General protein or ligand design will be achievable. “What if I change this atom” types of questions, for both small molecule ligands and macromolecules (e.g., antibody design) must become completely tractable.

We need much more efficient ways to sample conformational space. Must we track Newtonian dynamics at *femtosecond* resolution when the events of interest occur over, say, *milliseconds*? Twelve orders of magnitude: that’s equivalent to tracking the advance and retreat of the glaciers of the last Ice Age—tens of thousands of years—by noting their locations each and every *second*. Perhaps we don’t need to resolve all the fast motions (cf. the widely used SHAKE algorithm [77], and variants). Through some kind of clever dynamics or averaging, perhaps there is a way to “sample,” without bias, conformational space well enough to gain an full understanding of the biological phenomena being studied.

Two related examples: We will conduct much larger simulations—on biologically relevant timescales—in the coming quarter century. Say we are simulating a mitochondrion. Do we really need to compute electrostatic interactions from one end of it to the other? Likewise, enormous computational power will be “wasted” on water molecules. How should we gradate water, for instance—from near-scale atomistic, QM-like particles to a far-scale bulk-like phase—correctly and effectively [44]?

In other words, are we simulating biological molecules in the best way? How much of our current edifice is really

needed, and how much of it needs to change? As Phillip Windley wrote [78]: “Cathedrals have one millionth the mass of pyramids. The difference was the arch. Architecture demands arches.” Windley’s math may be off, but his point is apt. Are we missing the MD simulation “Arch”? The issues raised above are just a few obvious ones, without current solutions. The “Arch”—and with it, that glorious, open space underneath—awaits someone with better vision to see it.

The ultimate challenge—put here for completeness, though it is beyond our 25 year timeframe—is Al Gilman’s vision:

The premise: Someday there will be a computer labeled “A Cell,” and it will accurately predict all details of the behavior of a normal cell, as well as that perturbed by exogenous regulatory influences, drugs, mutations, and so on. I think I still believe the premise, but my time line for the prediction has expanded considerably. [79]

In the meantime, let us focus our attention on this key question: “Which molecule should I make next?” This is *the* question most important to a medicinal chemist, and it is the question we can, if we frame our studies carefully, answer in a useful and timely manner.

### The future of molecular dynamics simulations in drug discovery

Antoine de Saint-Exupery wrote “As for the future, your task is not to foresee it, but to enable it.” Here is our final Grand Challenge—our Goal:

*Create the computational methods to enable in silico drug design*

Drug design is fiendishly complex, and the universe of potential drugs is uncharted. Nearly one billion drug-like compounds comprising just 13 heavy atoms (C, N, O, F, S) exist [80], yet less than 70 million compounds, of any size, have been made. We have, in essence, explored only the very center of this multidimensional chemical universe—a universe in which essentially all the volume lies in the dark, unexplored corners [81]. Our computational chemistry tools, impressive and helpful as they are (e.g., shape-based methods [82]), do not yet provide comprehensive drug design solutions. If we can advance computation in drug design to the engineering level it enjoys in the aerospace, architectural, automotive, and electronics industries—industries for which simulations are now critical to success—then we will have unleashed the full power of computers to complement and enhance our own insights and intuitions.

All the key computational tools needed to reach this goal are already used—in nascent form—in drug discovery programs today. Each of these tools needs sharpening, by means of algorithmic innovation coupled with thorough experimental validation. We believe that computational tools and experimental methods should be used in concert, each according to its particular strengths. These key computational tools are sketched out below:

- The selection of a drug target or interaction partners in a particular metabolic or signaling pathway—a critical step in any project—would be guided by extensive genetics and bioinformatics input. This approach is being used today; as genomic information and associated (non-simulation) computational methods mature, it will become even more powerful.
- The structure of the target (or target complex), if not available, would be obtained using folding simulations. Alternatively, simulations may be used to prepare homology models as accurate as an experimental structure.
- Simulations would be used, along with complementary computational chemistry tools, to identify novel drug binding sites, including allosteric sites. These simulations may be carried out on the target alone or in the presence of suitable fragments (cf., current ligand binding or grand canonical Monte Carlo simulations).
- Fragment libraries would be allowed to bind to the site(s) of interest. As in current practice, library construction should be guided by sound medicinal chemistry principles (e.g., the so-called “Rule of 3” [83]).
- Candidate fragments would be scored using various methods. A triage scoring function, with little gradation, would be followed by binding free energy calculations (including of bioisosteres); as noted above, our ability to rapidly and accurately calculate binding free energies is critical [75].
- Fragments would be grown [84], and possibly linked [85], within the binding site. This growing process must be designed from the outset to produce drug-like molecules: ligands would be penalized (but not eliminated) for falling outside of known molecular descriptor bounds [86, 87]; transformations that increase molecular diversity could be favored; and the growing process would be completely synthetically “aware” [88–90]—synthetic tractability must be a given.
- Selectivity would be assessed, in awareness of the project’s particular pharmacological goals (single or multiple targets).
- It is crucial that we produce drugs, not simply inhibitors. Thus, ligands would be scored for ADME (absorption, distribution, metabolism, excretion) and



toxicological properties. Knowledge-based methods may play a dominant role here, but simulations could also be used to test directly ligand binding to both metabolic activators (cytochrome P450 oxidases, as recently demonstrated [91]; glucuronidases; etc.) and selected anti-targets (e.g., hERG, P-glycoprotein and other MDR transporters).

Each of these tools will likely mature at different times in the coming years. Eventually, all may be linked into an iterative process, with concomitant ligand–target dynamics and continual rescoring. Such a process would proceed in parallel on multiple chemotypes, to compensate for later attrition. Devising robust, multi-dimensional optimization processes capable of handling enormous numbers of candidate ligands—search tree pruning must be apt, not too soon, not too late—will present significant challenges to extending existing computational synthetic methods [88].

Achieving this audacious goal will require the combined efforts of computational scientists and engineers working hand-in-hand with experimentalists to ensure that these computational tools, as they become more powerful, truly address the key issues in drug discovery. The idea is not to displace experimental methods—after all, aircraft are not designed by just “telling” a computer to “do it”—but rather to advance computation such that it and experiment become fully complementary partners in the search for new drugs.

### Achieving wider acceptance of molecular dynamics simulations

Despite their power, molecular dynamics simulations of biological systems struggle with two issues of perception among the broader scientific community. First, many experimental biologists and (medicinal) chemists do not trust that MD simulation results are necessarily correct—they simply don't find the results compelling in the absence of thorough experimental validation. Simulations are received very differently elsewhere, for example in most engineering fields, astrophysics, condensed matter physics, weather and climate prediction, and fluid dynamics. We rely on and trust computational results in those fields. New aircraft, for instance, enter a wind tunnel, if at all [92], only after extensive, integrated simulations have provided flight-ready (or nearly so) designs. Do MD simulations really differ so significantly from computational fluid dynamics (CFD) simulations, and if so, how?

These methodological and epistemological aspects of computer simulations have been deeply pondered by Eric Winsberg [93]. The methods of CFD simulations—in level of detail, in modeling viscosity, in performing numerical

integration—are, in truth, as arbitrary and approximate as those underlying MD simulations. Quantitatively different, yes, but impossible to *prove* qualitatively superior. One key empirical distinction helped aircraft designers become convinced decades ago of the correctness and value, and occasional limitations, of CFD simulations: simulations could capture the essential phenomena on timeframes overlapping with wind-tunnel tests. This overlap enabled CFD simulations to be validated by experiment, and experimental observations to be explained by, and then predicted by, CFD simulations. MD simulations of biological systems are now entering just such an overlapping timeframe regime—orders-of-magnitude extrapolations are rapidly becoming a thing of the past. How can we strengthen, then, our partnerships with experimental colleagues to accelerate the improvement of MD simulations through iterative cycles of predictive simulation and experimental validation, as was done with CFD simulations? And, how can we better communicate the positive and reliable aspects of current simulations to the broader biological and pharmaceutical community?

The second, related, issue has to do with how we tend to conduct our research. Fifty years ago, John Platt made a strong case for the practice of what he called “strong inference”—a systematic method of scientific thinking—in enabling rapid research progress [94]. Were Platt to look at MD simulations today, he would likely perceive an over-emphasis on ever more precise quantitative (but ultimately unrevealing) measurements rather than qualitative results that actively disprove alternative hypotheses. Platt put it this way:

Organic chemistry has been the spiritual home of strong inference from the beginning. Do the bonds alternate in benzene or are they equivalent? If the first, there should be five disubstituted derivatives; if the second, three. And three it is<sup>19</sup>. This is a *strong-inference* test—not a matter of measurement, of whether there are grams or milligrams of the products, but a matter of logical alternatives.

And:

It consists of asking in your own mind, on hearing any scientific explanation or theory put forward, “But sir, what experiment could *disprove* your hypothesis?”

Glance at any current issue of *Cell* to see exactly what Platt is talking about—in nearly every paper, multiple, competing hypotheses are generated, put to the test, and then discarded on the way toward a higher truth. Are MD simulations not compelling to some because we, too often, shy away from this rigorous and, admittedly, grueling way of conducting science? Should we be doing more predicting, and (computational or experimental) testing, and less

observing and explaining? Without doubt, a new instrument that enables things previously unseen to be seen demands surveys—MD simulation has been rightly called a “computational microscope.” But, surveys cannot continue for long—rapid progress in science comes from formulating and rigorously testing hypotheses.

### On goals

Two closing thoughts regarding goals. First, as we set future goals for MD simulations, others will not be standing still—alternative approaches will also advance, raising the bar for our success higher than we might now think. The Rosetta molecular modeling approach, for instance, has proven to be very powerful (see, e.g., [95, 96]). For now, this approach is complementary to physics-based MD simulations, but this may not always be true—one approach may prove to be simply more effective than the other. And, in the pharmaceutical arena, the next quarter century may bring significant advances in alternative therapeutic approaches for which MD simulations have less to offer. It is not so far-fetched to think that intracellular antibody delivery or gene therapy—or some unanticipated, revolutionary innovation—may become commonplace, decreasing dramatically our dependence on conventional drugs. We should be mindful of these considerations as we apply MD simulations in the search for new drugs.

Second, we must set audacious goals. With the dogged pursuit of goals will come success, even if from unexpected directions. IBM set out to build a computer that could beat a grandmaster at chess—and succeeded. They then set their sights on winning *Jeopardy!*—and succeeded. We set out to build a computer that could speed up MD simulations by orders of magnitude—and succeeded. We have presented a few audacious goals for the next 25 years; not the best, perhaps, but a start, because without clear goals to serve as our guiding star, we are unlikely to succeed. By setting audacious goals, and by making a plan to achieve them, we lay a solid foundation for future success.

**Acknowledgments** We would like to thank our DESRES colleagues Adam Butts, Albert Pan, Alex Donchev, Andrew Taube, Dan Arlow, Huafeng Xu, John Salmon, Justin Gullingsrud, Kim Palmo, Kresten Lindorff-Larsen, Mike Eastwood, Ron Dror, and Stefano Piana for their many insightful discussions and abundance of helpful suggestions. We also thank Anthony Klon for useful discussions. We are especially grateful to Phil Miller and John Arrowsmith of Thomson Reuters for sharing their unpublished data on the reasons underlying Phase I clinical trial failures.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

### References

1. Marshall GR (1987) Computer-aided drug design. *Annu Rev Pharmacol Toxicol* 27:193–213
2. Van Drie JH (2007) Computer-aided drug design: the next 20 years. *J Comput Aided Mol Des* 21:591–601
3. Martin YC (1981) A practitioner’s perspective of the role of quantitative structure-activity analysis in medicinal chemistry. *J Med Chem* 24:229–237
4. Corey EJ, Wipke WT (1969) Computer-assisted design of complex organic syntheses. *Science* 166:178–192
5. Martin YC, Danaher EB, May CS, Weininger D (1988) MENTHORE, a database system for the storage and retrieval of three-dimensional molecular structures and associated data searchable by substructural, biologic, physical, or geometric properties. *J Comput Aided Mol Des* 2:15–29
6. Munos B (2009) Lessons from 60 years of pharmaceutical innovation. *Nat Rev Drug Discov* 8:959–968
7. Cuatrecasas P (2006) Drug discovery in jeopardy. *J Clin Invest* 116:2837–2842
8. Cuatrecasas P (2009) Pedro Cuatrecasas. *Nat Rev Drug Discov* 8:446
9. Firestone RA (2011) Lessons from 54 years of pharmaceutical research. *Nat Rev Drug Discov* 10:963
10. Paul SM, Mytelka DS, Dunwiddie CT, Persinger CC, Munos BH, Lindborg SR, Schacht AL (2010) How to improve R&D productivity: the pharmaceutical industry’s grand challenge. *Nat Rev Drug Discov* 9:203–214
11. Schäfer S, Kolkhof P (2008) Failure is an option: learning from unsuccessful proof-of-concept trials. *Drug Discov Today* 13:913–916
12. Arrowsmith J (2011) Trial watch: phase III and submission failures: 2007–2010. *Nat Rev Drug Discov* 10:87
13. Arrowsmith J (2011) Trial watch: phase II failures: 2008–2010. *Nat Rev Drug Discov* 10:328–329
14. Phase I data from the 2010 Global R&D performance metrics programme, CMR International, a Thomson Reuters business; courtesy of Phil Miller, Director, Product Strategy
15. Kola I, Landis J (2004) Can the pharmaceutical industry reduce attrition rates? *Nat Rev Drug Discov* 3:711–715
16. Maggiora GM (2011) The reductionist paradox: are the laws of chemistry and physics sufficient for the discovery of new drugs? *J Comput Aided Mol Des* 25:699–708
17. Dimitri N (2011) An assessment of R&D productivity in the pharmaceutical industry. *Trends Pharmacol Sci* 32:683–685
18. Tebib S, Bourguignon J-J, Wermuth C-G (1987) The active analog approach applied to the pharmacophore identification of benzodiazepine receptor ligands. *J Comput Aided Mol Des* 1:153–170
19. Sheridan RP, Venkataraghavan R (1987) Designing novel nicotinic agonists by searching a database of molecular shapes. *J Comput Aided Mol Des* 1:243–256
20. Saunders MR, Tute MS, Webb GA (1987) A theoretical study of angiotensin-converting enzyme inhibitors. *J Comput Aided Mol Des* 1:133–142
21. Pettersson I, Liljefors T (1987) Structure-activity relationships for apomorphine congeners. Conformational energies vs. biological activities. *J Comput Aided Mol Des* 1:143–152
22. McCammon JA, Gelin BR, Karplus M (1977) Dynamics of folded proteins. *Nature* 267:585–590
23. Wendoloski JJ, Wasserman ZR, Salemme FR (1987) Computer simulation of biological interactions and reactivity. *J Comput Aided Mol Des* 1:313–322
24. van Gunsteren WF, Berendsen HJC (1987) Thermodynamic cycle integration by computer simulation as a tool for obtaining free

- energy differences in molecular chemistry. *J Comput Aided Mol Des* 1:171–176
25. Lybrand TP, McCammon JA (1989) Computer simulation study of the binding of an antiviral agent to a sensitive and a resistant human rhinovirus. *J Comput Aided Mol Des* 2:259–266
  26. Karplus M, Kuriyan J (2005) Molecular dynamics and protein function. *Proc Natl Acad Sci USA* 102:6679–6685
  27. Durrant JD, McCammon JA (2011) Molecular dynamics simulations and drug discovery. *BMC Biol* 9:71
  28. Whitford PC, Geggier P, Altman RB, Blanchard SC, Onuchic JN, Sanbonmatsu KY (2010) Accommodation of aminoacyl-tRNA into the ribosome involves reversible excursions along multiple pathways. *RNA* 16:1196–1204
  29. Freddolino PL, Arkipov A, Larson SB, McPherson A, Schulten K (2006) Molecular dynamics simulations of the complete satellite tobacco mosaic virus. *Structure* 14:437–449
  30. Mei C, Sun Y, Zheng G, Bohm EJ, Kalé LV, Phillips JC, Harrison C (2011) Enabling and scaling biomolecular simulations of 100 million atoms on petascale machines with a multicore-optimized message-driven runtime. In: Proceedings of the 2011 ACM/IEEE conference on supercomputing (SC11). IEEE Computer Society Press, Washington
  31. Beckstein O, Sansom MS (2006) A hydrophobic gate in an ion channel: the closed state of the nicotinic acetylcholine receptor. *Phys Biol* 3:147–159
  32. Jensen MØ, Borhani DW, Lindorff-Larsen K, Maragakis P, Jogini V, Eastwood MP, Dror RO, Shaw DE (2010) Principles of conduction and hydrophobic gating in K<sup>+</sup> channels. *Proc Natl Acad Sci USA* 107:5833–5838
  33. Clark M, Guarnieri F, Shkurko I, Wiseman J (2006) Grand canonical Monte Carlo simulation of ligand–protein binding. *J Chem Inf Model* 46:231–242
  34. Moffett K, Konteatis Z, Nguyen D, Shetty R, Ludington J, Fujimoto T, Lee KJ, Chai X, Nambodiri H, Karpus M, Dorsey B, Guarnieri F, Bukhtiyarova M, Springman E, Michelotti E (2011) Discovery of a novel class of non-ATP site DFG-out state p38 inhibitors utilizing computationally assisted virtual fragment-based drug design (vFBDD). *Bioorg Med Chem Lett* 21:7155–7165
  35. Brown SP, Muchmore SW (2009) Large-scale application of high-throughput molecular mechanics with Poisson–Boltzmann surface area for routine physics-based scoring of protein–ligand complexes. *J Med Chem* 52:3159–3165
  36. Michel J, Essex JW (2010) Prediction of protein–ligand binding affinity by free energy simulations: assumptions, pitfalls and expectations. *J Comput Aided Mol Des* 24:639–658
  37. Zeevaert JG, Wang L, Thakur VV, Leung CS, Tirado-Rives J, Bailey CM, Domaal RA, Anderson KS, Jorgensen WL (2008) Optimization of azoles as anti-human immunodeficiency virus agents guided by free-energy calculations. *J Am Chem Soc* 130:9492–9499
  38. Jorgensen WL, Bollini M, Thakur VV, Domaal RA, Spasov KA, Anderson KS (2011) Efficient discovery of potent anti-HIV agents targeting the Tyr181Cys variant of HIV reverse transcriptase. *J Am Chem Soc* 133:15686–15696
  39. Bollini M, Domaal RA, Thakur VV, Gallardo-Macias R, Spasov KA, Anderson KS, Jorgensen WL (2011) Computationally-guided optimization of a docking hit to yield catechol diethers as potent anti-HIV agents. *J Med Chem*. Nov 14. 2011. doi:10.1021/jm201134m
  40. Buch I, Giorgino T, De Fabritiis G (2011) Complete reconstruction of an enzyme–inhibitor binding process by molecular dynamics simulations. *Proc Natl Acad Sci USA* 108:10184–10189
  41. Shan Y, Kim ET, Eastwood MP, Dror RO, Seeliger MA, Shaw DE (2011) How does a drug molecule find its target binding site? *J Am Chem Soc* 133:9181–9183
  42. Hurst DP, Grossfield A, Lynch DL, Feller S, Romo TD, Gawrisch K, Pitman MC, Reggio PH (2010) A lipid pathway for ligand binding is necessary for a cannabinoid G protein–coupled receptor. *J Biol Chem* 285:17954–17964
  43. Dror RO, Pan AC, Arlow DH, Borhani DW, Maragakis P, Shan Y, Xu H, Shaw DE (2011) Pathway and mechanism of drug binding to G-protein-coupled receptors. *Proc Natl Acad Sci USA* 108:13118–13123
  44. Kamerlin SC, Vicatos S, Dryga A, Warshel A (2011) Coarse-grained (multiscale) simulations in studies of biophysical and chemical systems. *Annu Rev Phys Chem* 62:41–64
  45. Tozzini V (2010) Minimalist models for proteins: a comparative analysis. *Q Rev Biophys* 43:333–371
  46. Shih AY, Arkipov A, Freddolino PL, Sligar SG, Schulten K (2007) Assembly of lipids and proteins into lipoprotein particles. *J Phys Chem B* 111:11095–11104
  47. Scott KA, Bond PJ, Ivetac A, Chetwynd AP, Khalid S, Sansom MS (2008) Coarse-grained MD simulations of membrane protein–bilayer self-assembly. *Structure* 16:621–630
  48. Laio A, Parrinello M (2002) Escaping free-energy minima. *Proc Natl Acad Sci USA* 99:12562–12566
  49. Hamelberg D, Mongan J, McCammon JA (2004) Accelerated molecular dynamics: a promising and efficient simulation method for biomolecules. *J Chem Phys* 120:11919–11929
  50. Maragliano L, Vanden-Eijnden E (2006) A temperature accelerated method for sampling free energy and determining reaction pathways in rare events simulations. *Chem Phys Lett* 426:168–175
  51. Shaw DE, Dror RO, Salmon JK, Grossman JP, Mackenzie KM, Bank JA, Young C, Deneroff MM, Batson B, Bowers KJ, Chow E, Eastwood MP, Ierardi DJ, Klepeis JL, Kuskin JS, Larson RH, Lindorff-Larsen K, Maragakis P, Moraes MA, Piana S, Shan Y, Towles B (2009) Millisecond-scale molecular dynamics simulations on Anton. In: Proceedings of the ACM/IEEE conference on supercomputing (SC09). IEEE Computer Society Press, Washington
  52. Shaw DE, Maragakis P, Lindorff-Larsen K, Piana S, Dror RO, Eastwood MP, Bank JA, Jumper JM, Salmon JK, Shan Y, Wrighers W (2010) Atomic-level characterization of the structural dynamics of proteins. *Science* 330:341–346
  53. Lindorff-Larsen K, Piana S, Dror RO, Shaw DE (2011) How fast-folding proteins fold. *Science* 334:517–520
  54. Wimmer E (1987) Future in biomolecular computation. *J Comput Aided Mol Des* 1:283–290
  55. Liu F, Du D, Fuller AA, Davoren JE, Wipf P, Kelly JW, Gruebele M (2008) An experimental survey of the transition between two-state and downhill protein folding scenarios. *Proc Natl Acad Sci USA* 105:2369–2374
  56. Freddolino PL, Liu F, Gruebele M, Schulten K (2008) Ten-microsecond molecular dynamics simulation of a fast-folding WW domain. *Biophys J* 94:L75–L77
  57. Ensign DL, Pande VS (2009) The Fip35 WW domain folds with structural and mechanistic heterogeneity in molecular dynamics simulations. *Biophys J* 96:L53–L55
  58. Piana S, Sarkar K, Lindorff-Larsen K, Guo M, Gruebele M, Shaw DE (2011) Computational design and experimental testing of the fastest-folding  $\beta$ -sheet protein. *J Mol Biol* 405:43–48
  59. Piana S, Lindorff-Larsen K, Shaw DE (2011) How robust are protein folding simulations with respect to force field parameterization? *Biophys J* 100:L47–L49
  60. Duan Y, Harvey SC, Kollman PA (2001) Protein folding and beyond. In: Keinan E, Schechter I (eds) *Chemistry for the 21st century*. Wiley, Weinheim
  61. Pommier Y, Cherfils J (2005) Interfacial inhibition of macromolecular interactions: nature’s paradigm for drug discovery. *Trends Pharmacol Sci* 26:138–145

62. Allinger NL (2011) Understanding molecular structure from molecular mechanics. *J Comput Aided Mol Des* 25:295–316
63. Halgren TA (1996) Merck molecular force field. I. Basis, form, scope, parameterization, and performance of MMFF94. *J Comput Chem* 17:490–519
64. Ma JC, Dougherty DA (1997) The cation- $\pi$  interaction. *Chem Rev* 97:1303–1324
65. Zacharias N, Dougherty DA (2002) Cation- $\pi$  interactions in ligand recognition and catalysis. *Trends Pharmacol Sci* 23: 281–287
66. Schreiner PR, Chernish LV, Gunchenko PA, Tikhonchuk EY, Hausmann H, Serafin M, Schlecht S, Dahl JE, Carlson RM, Fokin AA (2011) Overcoming lability of extremely long alkane carbon-carbon bonds through dispersion forces. *Nature* 477:308–311
67. Ponder JW, Wu C, Ren P, Pande VS, Chodera JD, Schnieders MJ, Haque I, Mobley DL, Lambrecht DS, DiStasio RA Jr, Head-Gordon M, Clark GN, Johnson ME, Head-Gordon T (2010) Current status of the AMOEBA polarizable force field. *J Phys Chem B* 114:2549–2564
68. Donchev AG, Galkin NG, Pereyaslavets LB, Tarasov VI (2006) Quantum mechanical polarizable force field (QMPFF3): refinement and validation of the dispersion interaction for aromatic carbon. *J Chem Phys* 125:244107
69. Skillman AG, Geballe MT, Nicholls A (2010) SAMPL2 challenge: prediction of solvation energies and tautomer ratios. *J Comput Aided Mol Des* 24:257–258
70. Khoruzhii O, Donchev AG, Galkin N, Illarionov A, Olevanov M, Ozrin V, Queen C, Tarasov V (2008) Application of a polarizable force field to calculations of relative protein-ligand binding affinities. *Proc Natl Acad Sci USA* 105:10378–10383
71. Waldman M, Hagler AT (1993) New combining rules for rare gas van der Waals parameters. *J Comp Chem* 14:1077–1084
72. Williams SL, de Oliveira CA, McCammon JA (2010) Coupling constant pH molecular dynamics with accelerated molecular dynamics. *J Chem Theory Comput* 6:560–568
73. Machuqueiro M, Baptista AM (2008) Acidic range titration of HEWL using a constant-pH molecular dynamics method. *Proteins* 72:289–298
74. Eigen M (1964) Proton transfer, acid-base catalysis, and enzymatic hydrolysis. *Angew Chem Int Ed Engl* 3:1–19
75. Connelly PR, Vuong TM, Murcko MA (2011) Getting physical to fix pharma. *Nat Chem* 3:692–695
76. Chodera JD, Mobley DL, Shirts MR, Dixon RW, Branson K, Pande VS (2011) Alchemical free energy methods for drug discovery: progress and challenges. *Curr Opin Struct Biol* 21:150–160
77. Ryckaert J-P, Ciccotti G, Berendsen HJC (1977) Numerical integration of the cartesian equations of motion of a system with constraints: molecular dynamics of *n*-alkanes. *J Comp Phys* 23:327–341
78. Windley, PJ (2006) Alan Kay: is computer science an oxymoron? [http://www.windley.com/archives/2006/02/alan\\_kay\\_is\\_com.shtml](http://www.windley.com/archives/2006/02/alan_kay_is_com.shtml). Accessed 18 Nov 2011
79. Gilman AG (2011) Silver spoons and other personal reflections. *Annu Rev Pharmacol Toxicol*. doi:10.1146/annurev-pharmtox-010611-134652
80. Blum LC, Reymond JL (2009) 970 million druglike small molecules for virtual screening in the chemical universe database GDB-13. *J Am Chem Soc* 131:8732–8733
81. Hayes B (2011) An adventure in the *N*th dimension. *Am Sci* 99:442–446
82. Nicholls A, McGaughey GB, Sheridan RP, Good AC, Warren G, Mathieu M, Muchmore SW, Brown SP, Grant JA, Haigh JA, Nevins N, Jain AN, Kelley B (2010) Molecular shape and medicinal chemistry: a perspective. *J Med Chem* 53:3862–3886
83. Congreve M, Carr R, Murray C, Jhoti H (2003) A ‘rule of three’ for fragment-based lead discovery? *Drug Discov Today* 8: 876–877
84. Moon JB, Howe WJ (1991) Computer design of bioactive molecules: a method for receptor-based de novo ligand design. *Proteins* 11:314–328
85. Ichihara O, Barker J, Law RJ, Whittaker M (2011) Compound design by fragment-linking. *Mol Inf* 30:298–306
86. Walters WP, Green J, Weiss JR, Murcko MA (2011) What do medicinal chemists actually make? A 50-year retrospective. *J Med Chem* 54:6405–6416
87. Lipinski CA, Lombardo F, Dominy BW, Feeney PJ (2001) Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv Drug Deliv Rev* 46:3–26
88. Todd MH (2005) Computer-aided organic synthesis. *Chem Soc Rev* 34:247–266
89. Judson P (2009) Knowledge-based expert systems in chemistry. Not counting on computers. Royal Society of Chemistry Publishing, Cambridge
90. Cook A, Johnson AP, Law J, Mirzazadeh M, Ravitz O, Simon A (2011) Computer-aided synthesis design: 40 years on. *WIREs Comput Mol Sci* 0:1–29. doi:10.1002/wcms.61
91. Li J, Schneebeli ST, Bylund J, Farid R, Friesner RA (2011) ID-Site: an accurate approach to predict P450-mediated drug metabolism. *J Chem Theory Comput* 7:3829–3845
92. Arguably the most innovative aircraft of the 21st century, SpaceShipOne was designed by perhaps the most creative aircraft designer now living, Burt Rutan, at Scaled Composites LLC. This spacecraft won the \$10 million Ansari X Prize. The prize, set up to encourage the development of space tourism, went to the team that safely sent a piloted aircraft to a height of 100 kilometers (62 miles) twice in a single week. On October 4, 2004, SpaceShipOne accomplished that feat. It was flown 14 times; maximum speed was Mach 3.09. The craft has many unique elements, including a fuel system that supplies explosive rocket power by combining nitrous oxide and solid rubber, and wings that fold up to provide high drag for reentry. In a nod to modernity, Rutan’s team used no wind-tunnel testing; all design refinements and performance predictions were derived from computational fluid dynamic tools. See: <http://www.scaled.com/projects/tierone/faq/>, <http://www.astronautix.com/craft/spaipone.htm>, and <http://www.pbs.org/wgbh/nova/concorde/inno-nf.html>. Accessed 18 Nov 2011
93. Winsberg EB (2010) Science in the age of computer simulation. University of Chicago Press, Chicago
94. Platt JR (1964) Strong inference: certain systematic methods of scientific thinking may produce much more rapid progress than others. *Science* 146:347–353
95. Siegel JB, Zanghellini A, Lovick HM, Kiss G, Lambert AR, St Clair JL, Gallaher JL, Hilvert D, Gelb MH, Stoddard BL, Houk KN, Michael FE, Baker D (2010) Computational design of an enzyme catalyst for a stereoselective bimolecular Diels-Alder reaction. *Science* 329:309–313
96. Fleishman SJ, Whitehead TA, Ekiert DC, Dreyfus C, Corn JE, Strauch EM, Wilson IA, Baker D (2011) Computational design of proteins targeting the conserved stem region of influenza hemagglutinin. *Science* 332:816–821