Available online at www.sciencedirect.com
ScienceDirect
Current Opinion in Chemical Biology

Large-scale integrated super-computing platform for next generation virtual drug discovery

Wayne Mitchell[1,2] and Shunji Matsumoto[3]

Traditional drug discovery starts by experimentally screening chemical libraries to find hit compounds that bind to protein targets, modulating their activity. Subsequent rounds of iterative chemical derivitization and rescreening are conducted to enhance the potency, selectivity, and pharmacological properties of hit compounds. Although computational docking of ligands to targets has been used to augment the empirical discovery process, its historical effectiveness has been limited because of the poor correlation of ligand dock scores and experimentally determined binding constants. Recent progress in super-computing, coupled to theoretical insights, allows the calculation of the Gibbs free energy, and therefore accurate binding constants, for usually large ligand–receptor systems. This advance extends the potential of virtual drug discovery. A specific embodiment of the technology, integrating *de novo*, abstract fragment based drug design, sophisticated molecular simulation, and the ability to calculate thermodynamic binding constants with unprecedented accuracy, are discussed.

Addresses
[1] Experimental Therapeutics Centre, Agency for Science and Technology Research (A*STAR), 31 Biopolis Way, #03-01 Nanos, Biopolis, Singapore 138669, Singapore
[2] Nanyang Technological University, School of Computer Engineering, Division of Information Sciences, 50 Nanyang Avenue, Singapore 639798, Singapore
[3] *In-silico* Drug Discovery Research Division, Bio-IT Business Development Unit, Fujitsu Limited, 1-9-8 Nakase, Mihama-ku, Chiba-city, Chiba 261-8588, Japan

Corresponding author: Matsumoto, Shunji (shun@jp.fujitsu.com)

Current Opinion in Chemical Biology 2011, **15**:553–559

This review comes from a themed issue on
Next Generation Therapeutics
Edited by Alex Matter and Thomas H. Keller

Available online 30th June 2011

1367-5931/$ – see front matter

DOI 10.1016/j.cbpa.2011.06.005

Introduction

Drug discovery is distressed. The number of approved new molecular entities has declined steadily for 15 years [1], the cost per new approved compound has breached the one billion United States Dollar (USD) benchmark [2] and, by one informed estimate, the financial return provided by all therapeutic product categories does not even recover the capital costs of their development. Moreover, when biologics are removed from this model, the net present value of pharmaceutical Research and Development (R&D) investment is actually negative [3]. Disturbingly, poor research performance occurs in the context of increasing social need. The rising incidence of drug resistant bacteria is well documented [4], viral pathogens such as Severe Acute Respiratory Syndrome (SARS), Dengue and H1N1 pose pandemic threats [5], and an aging and more affluent global population drives up the prevalence of chronic diseases that are if anything more difficult than infectious diseases to drug. In the United States, where by 2023 the Census Bureau projects ~15% of males and ~19% of females will be 65 years old or older, rates of cancer are expected to rise from 33% to 62%, of diabetes from 33% to 53%, of cardiovascular complaints from 6% to 39%, of mental disorders from 35% to 54%, for a total increase in chronic morbidity in the U.S. population from 17% to 42% [6]. An increase in successful drug discovery, especially against difficult chronic targets, is clearly desirable. However, progress must materialize in the framework of reduced per compound cost and enhanced efficiency, since capital will not continue to flow into a sector that offers net negative returns. Computational or 'virtual' drug discovery strategies, potentially cheaper and faster, offer attractive alternative, or at least complimentary, routes to improved R&D performance in the therapeutics sector [7].

Computational drug discovery

The physiological effect of a drug is mediated by electrostatic and geometrical interactions of the atoms of the ligand with the atoms of its corresponding receptor, interactions which conform to the laws of physics and quantum chemistry, and which can therefore be described by predictive mathematical models [8•]. Although these models are complex (and in the quantum case inherently non-exact), researchers active in the computer intensive field of molecular graphics realized thirty years ago that *in silico* assessment of drug-receptor binding could be deployed to accelerate drug discovery [9,10]. They also realized that continuing operation of Moore's Law, with 18 month doublings in computing efficiency and economy, implied an on-going improvement and refinement in computational techniques. The crux of the computational drug discovery paradigm is this coupling of the fundamental laws of biophysics with the accelerating technological performance of the semiconductor industry. The former inspires faith that the approach can work in theory, the latter that it will work in practice.

Virtual compound screening

Virtual screening, whether of compounds or molecular fragments, has two stages. First, the algorithms attempt to find the correct conformation and position the ligand in the active site of the receptor, and then they try to quantify the quality of particular atomic arrangements by assigning a score. Several technically different approaches to predicting ligand–receptor interactions have been developed, but all are known as 'docking' algorithms after the suggestively named primogenitor program, 'DOCK' [9]. The modeling of ligand–receptor atomic interactions presupposes an accurate three-dimensional molecular structure of the receptor so that inter-atomic forces can be calculated. Since protein folding cannot yet be modeled, this means having an X-ray crystal or NMR structure of the receptor, or a homology model which maps a related protein sequence onto a known structure. *A priori* one might suppose the experimentally determined crystal structure to be inevitably superior. Surprisingly, a meta-analysis of DOCKing studies concluded that in some cases virtual screening was more successful on homology models compared to experimental structures [11••]. This seems counter-intuitive, but may indicate that the relaxed precision of the homology models indirectly capture conformational flexibility that is lost in 'frozen', possibly subtly distorted crystal structures. In any case, an homology model must start from a closely related experimental structure, so an important contributing factor in the increased utility of computational drug discovery is the rapid growth in the number of available protein structures (currently approaching 75,000 structures in the Protein Data Base http://www.pdb.org/pdb/statistics/contentGrowthChart.do?content=total&seqid=100), a number which in turn reflects improvements in protein production, robotic crystallization regimens, and the wide availability of sophisticated advanced light sources [12]. Another positive development in virtual screening infrastructure is the creation of curated virtual compound databases that provide large prebuilt sets of virtual representations of commercially available molecules suitable for input to virtual screens. ZINC at the University of California San Francisco [13], and EDULISS at Edinburgh University [14], are two examples. The much larger Chemical Universe Database GDB-13 takes a different approach, attempting to construct the universe of 'synthetically plausible', rather than 'available' compounds [15].

Molecular docking: successes and limitations

At a high level the performance of dock programs can be measured by two criteria: 'DOCKing power' (the ability to identify the correct experimental ligand binding pose in a collection of incorrect, computer generated 'decoy' poses, i.e. the ability to correctly position ligands in the active site, or to 'pose' them); and 'scoring power' (the ability to produce dock binding scores that correlate with experimentally determined binding affinities). In the past decade a large number of comparative studies of the performance of various dock programs have been undertaken, in both academic and pharmaceutical settings [11,16,17••,18•,19–22]. Despite the diverse backgrounds of the investigators, and although these studies differ in methodology and are not directly comparable, they nevertheless unanimously agree on two points. One, dock algorithms fairly accurately *pose* ligands in the active site, and two, the same dock algorithms poorly *score* those ligands' affinity. In other words, dock programs correctly identify the geometry of ligand–receptor systems, but, do not in general accurately predict the binding energy, and therefore cannot predict ligand potency. To make this concrete, a typical dock screen might produce 1000 'hit' compounds, but, the most potent compounds are as likely to be ranked at the bottom of that list by the scoring function as they are to appear near the top. This is a significant deficiency since the expected potency of a compound often will be the operational feature of interest, for example in prioritizing compounds for medicinal chemistry. In sum, Dock algorithms can 'pose' ligands well but they 'score' them poorly.

Beyond dock scores: accurate binding affinity from thermodynamic calculation with MAPLE CAFEE

$\Delta G = RT \ln K_d$ exactly relates the computed Gibbs free energy difference ΔG and experimentally measured dissociative constant K_d under temperature T (where, R is the gas constant). Free energy differences between bound and unbound equilibrium states of a protein–ligand–water system (ΔG) gives the binding affinity of the ligand, which in general translates into drug efficacy. In other words, *correct* computation of ΔG values for a series of ligands leads immediately to a ligand list accurately ranked by potency.

Computational methods to perform the ΔG calculation have been studied enthusiastically since the late 1990s when it was proved that a nonequilibrium process in finite-time can derive the binding free energy exactly [23–25]. This theoretical insight was followed up in 2005 when ΔG was shown to be approachable by massively parallel computation. Scaling up to thousands of concurrent CPUs reduced the computational requirement from years to days, and allowed binding free energy of real molecular systems to be computed. Access to this high performance computational resource made possible an important series of proof of concept experiments that in turn produced calculated binding affinities in excellent agreement with corresponding experimental values [26,27]. Subsequently, the computational methodology has been improved, a better force field refining method has been implemented, and the platform, christened Massively Parallel Computation of Absolute binding Free Energy with well Equilibrated system (MAPLE