A Step Closer to Precision Oncology:

Computational, Biochemical, and Cell-Based Screening to Find Compounds that Stabilize p53

By
© 2017
Karen Khar
M.S., University of Missouri, 2008
D.C., Cleveland Chiropractic College, 1998

Submitted to the graduate degree program in the Center for Computational Biology and the Graduate Faculty of the University of Kansas in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

_____

Chair: Eric Deeds

_____

Ilya Vakser

_____

Joanna Slusky

_____

Christian Ray

_____

Mark Richter

Date Defended:
October 2, 2017

The dissertation committee for

Karen Khar

certifies that this is the approved version of the following dissertation:

A Step Closer to Precision Oncology: Computational, Biochemical, and Cell-Based Screening to Find Compounds that Stabilize p53

_____

Chair: Eric Deeds

Date Approved:
October 2, 2017

Abstract

Personalized medicine in cancer aims to tailor a treatment plan that takes into account the unique features of a patient's malignancy. One therapeutic target that has a chance to affect a large population of cancer patients is p53. p53 is a tumor suppressor that activates senescence or apoptosis in cells that have accumulated mutations that could lead to cancer. Half of all cancers have mutations in p53, which highlights the importance of its role in disease. A subset of these mutations have been shown to inhibit p53 function by destabilizing p53's core domain. This led to the hypothesis that a personalized drug for patients with this type of destabilized p53 mutation could lead to apoptosis in cancer cells.

There has been a lot of evidence supporting this hypothesis. This evidence has inspired many researchers to screen for small molecules that stabilize p53 mutants and rescue function. However, the hits discovered in these screens (with one potential exception) have not been found to be adequate drug leads for several reasons. Many have turned out to rescue function, but not by directly binding p53. Others bind p53, but either lack sufficient binding affinity or cause nonspecific cell responses. All of these are likely to induce side effects if used as part of a cancer therapeutic. This leads to the question: Is there a better way to find a small molecule stabilizer for cancer-associated mutants of p53?

Here, I present an alternative approach that focuses on finding a direct binder to p53's core domain in order to avoid off-target effects. Our initial step was a computational approach that uses the crystal structure of p53's core domain in order to virtually screen a set of small molecules for binding. I found a novel pocket on the protein structure that I predicted to be druggable, because the site readily forms pockets during simulations of the core domain. I performed a virtual screen using the DARC, a docking tool from the molecular modeling suite,

Rosetta, and selected the 28 best ranked compounds for biochemical testing with purified p53 using two different cancer-associated, destabilizing mutations. Surprisingly, I found that 11 of the 28 compounds stabilized both mutants. Further testing was done in cancer cell lines showing that 7 compounds activated p53 transcription of p21 and PUMA, which are known targets of p53. Using the fluorescent antibody pAb 1620 that binds natively folded p53, we showed that 4 of the compounds lead to a much higher concentration of folded p53 in cells.

The excitingly high hit rate was found from a modest sized initial virtual screen of only 64,000 molecules. This suggests that this novel pocket is prone to bind molecules in a manner that rescues structure and function, and should be as a starting point for a larger screen. Also, the compounds from the current screen are intriguing hits that will be further analyzed and optimized to develop new stabilizers of p53.

Acknowledgements

My deepest gratitude goes out to Dr. John Karanicolas for allowing me to be a part his lab. It was an amazing pleasure to gain from his knowledge and vast abilities as a scientist.

I would like to thank all members of my thesis committee: Dr. Ilya Vakser, Dr. Eric Deeds, Dr. Joanna Slusky, Dr. Mark Richter, and Dr. Christian Ray. I know it is a great deal of work to be on the committee and appreciate the time you put in to assure that I leave KU as an accomplished researcher.

I am thankful to my colleagues Ragul Gowthaman, David Johnson, Jittasak Khowsathit, Nan Bai, and Shipra Malhotra for their scientific input into my project and for allowing me to learn from theirs. Especially Ragul and David who gave me significant help in my early years.

Finally, I wish to thank my mother for her love and support during my return to academia and for being an inspiring role model.

Table of Contents

Introduction

p53 as a target for drug discovery in cancer treatment

It has been estimated that 40% of people will be diagnosed with cancer during their lifetime, which is a devastating number considering the fact that 171 out of 100,000 people with cancer die from it every year. In 2016, roughly 600,000 people in the US died of cancer [4]. Recent advances in our understanding of cancer have lead to an area of research known as personalized medicine. This is a method for treating patients based on their disease-causing genetic variation. This approach is ideal for cancer treatment since genetic mutations are a major factor in the development and progression of cancer. Though, the variations present in cancer are incredibly diverse.

p53 is a promising target for the personalized medicine approach because it is mutated in about half of all cancers [6]. p53 is a tumor suppressor that's normally kept at low levels in cells. When cell stressors such as cancer mutations are present in the cell, p53 levels are allowed to rise, which leads to cell senescence or apoptosis [7]. About half of the cancer-associated mutations in p53 have been found to inactivate it by destabilizing the core domain [10]. These destabilizing mutants, as a group, would be excellent targets for a personalized treatment for this subset of cancers.

Development of a drug that targets destabilizing mutants of p53 has the potential of affecting a wide range of cancers because p53 mutations are prevalent in so many cancers. It has been detected in almost all cancers, though its rate varies in individual types of cancers. Mutations in p53 have been found at high frequency in the majority of the most common cancers [11, 12]. Also, it has been found that the p53 mutations are associated with poor prognosis [13].

So finding a drug that successfully rescues these mutations could have a highly significant effect on cancer mortality.

Rescuing cancer-associated, destabilizing mutations of p53

Development of a drug lead for cancer-associated, destabilizing mutations of p53 is a highly challenging goal. Attempts to find a molecule that stabilizes p53 have been going on for two decades, but only modest advancements have been accomplished. The first suggestion that it was possible to find a successful stabilizer of p53 was the monoclonal antibody, PAb421, which binds p53. Microinjection into SW480 cancer cells restored some p53 function [15-17]. Since then we have significantly expanded our understanding of destabilizing p53 mutations and found several proof-of-concept compounds that bind, stabilize, and rescue function of p53.

Destabilizing mutations of p53 are referred to as "structural" mutants, which differentiates them from the other main type of cancer-associated mutations of p53 that are referred to as "contact" mutants [18]. Contact mutants occur directly in the DNA binding region of the core domain. These cannot be rescued by stabilizers because they lack affinity for DNA so cannot function as transcription factors. Structural mutants have been shown to decrease the stability of p53's core domain causing them to be mostly unfolded at physiological temperature.

The current hypothesis for potential rescue of p53 is that a small molecule or peptide called a reactivator could bind and increase the stability of the core domain enough for a structural mutant to remain folded at 37° C. If a cancer patient with a destabilizing mutation in p53 was treated with a drug that included a reactivator, p53 is hypothesized to maintain it's native fold long enough to trigger apoptosis in the cancer cells, unless of course, there are other disease factors that could inhibit the p53 pathway. If the reactivator had a very high affinity, it

could be effective at low concentrations and would be far less likely to have non-specific interactions, thus avoiding side effects. Not only does this proposed mechanism have the capacity to affect a broad range of cancer patients, but also it has the potential to have a focused effect on diseased tissue only.

While this hypothesis has not been proven, there is evidence that this mechanism is feasible. For example, the small molecule PhiKan083 was found during a virtual screen using the structure of the Y220C structural mutant of p53. It was found to bind and stabilize p53, though with inadequate affinity to be a drug lead. Nonetheless, it is also shown to rescue function of cancer cells containing the Y220C mutant [19]. Other compounds have been found using cell-based assays that rescue p53 function, though it has not been proven that they function by directly binding and stabilizing p53. So the hypothesis has merit, but more work needs to be done to bring it to fruition.

Better methods of finding and testing reactivators

The goal of my thesis project was to take the lessons from other researchers that have sought a reactivator and improve upon their work in the hunt for a p53 reactivator. There are 3 major areas that this project focuses on that improve on previous methods.

1) Virtual screen and direct p53 binding assays are more likely to find direct binders.

Most researchers used cell-based assays in their initial screen for reactivators, but many of the compounds discovered this way turned out not to directly stabilize p53 so are more likely to result in side effects if used as a cancer treatment. The best way to assure that a reactivator directly binds p53 as a part of its mechanism of action is to do the screening process using a

structure-based virtual screen for molecules that bind p53. Once a set of molecules has been found that are predicted to bind, the next step would be to do biochemical assays with purified p53 to show both binding and stabilization of p53. Once predicted compounds are found to truly bind and stabilize p53, then it is reasonable to move into cell-based testing to show that compounds can stabilize and rescue p53 function in vivo.

Any molecule that is successful in all these tests is more likely to be a better drug candidate because it has shown both specificity to and rescue of p53. Lead optimization can then be done to improve affinity and further reduce the likelihood of nonspecific interactions that cause side effects. Since compounds found by this method have a predicted binding mode, optimization of hits can be done in a structure-based way. So a virtual screen not only prevents off-target false positives that are common in cell-based screens, but also provides structural insights into potential optimization once hits have been verified.

2) Druggability analysis determines whether surface pockets are amenable to binding small molecules.

Success of a virtual screen is dependent on many factors. Most importantly, the target of the screen needs to be amenable to binding a small molecule. Two groups have attempted virtual screen for p53 reactivators and found binders with inadequate binding affinity, which may be because they are targeting the wrong pockets. Our lab has a method for measuring how druggable a surface pocket is [20]. It has been shown to be predictive for a set of benchmarking proteins with known druggable sites. It is based on the fact that most druggable sites go through a significant conformational change when binding a small molecule. So we computationally

predict whether pockets on a protein surface are able to perform the conformational change necessary to bind a molecule.

Druggability analysis is done by performing molecular simulations of protein conformational fluctuations in order to discern which locations readily form low energy pockets. Sites known to be druggable form pockets while non-druggable sites do not. This is done with a slight bias towards structures with larger pockets in order to overcome the energy barrier that would otherwise not be breached without the presence of a small molecule. Non-druggable sites would not open even with the bias. This analysis is necessary to discover if there is a druggable pocket and is helpful in finding energetically favorable, open-pocket conformations that can be used when docking small molecules during a virtual screen.

3) Fast, accurate virtual screening using DARC with wild type p53 increases the probability that hits will be found.

3a) DARC is more accurate when using shallow pockets.

Traditional virtual screen is designed to dock molecules in deep pockets such as those found in enzyme active sites [21]. This has made sense for most screening targets because most active sites evolved a deep pocket that is similar to its bound conformation in order to readily bind its cognate molecule. However, we are not able target a deep pocket with p53's core domain because the only deep pocket is the one that binds DNA. Our goal is to bind a stabilizing small molecule at a location that does not interfere with any interaction sites in order to allow p53 to function normally. So if there is a site on p53 that is able to bind a stabilizing molecule that does

not have a physiological binding partner, it will be in a relatively shallow pocket. Doing a virtual screen with a shallow pocket will require a different kind of docking tool.

Our lab has developed a docking tool that is designed to optimize docking in shallow pockets called DARC (Docking Approach with Ray Casting) [22]. It targets shallow pockets by scoring molecules with an algorithm that favors molecules with a maximum amount of contact with the protein surface. It does so by defining the topography on the pocket surface and compares it to the topography of the docked molecule's interaction at the surface. Then it optimizes the molecule's fit in the pocket in order to find the orientation that has the best topographical match. DARC has proven to successfully predict the binding modes of a set of molecules that bound Mcl-1, which is a similar docking challenge to p53 because the target pocket is also far more shallow, that traditional docking sites [22].

3b) Fast docking is necessary.

One of the greatest challenges of virtual screen is the computational expense. In an ideal situation, computation time would be insignificant. This would allow us to do highly complex simulations using elaborate energy functions and exhaustively sample docking orientations using every atom in the protein, the docked molecule, and includes water. However, this ideal case would require a lifetime to compute so docking tools must always find a balance between speed and comprehensive docking analysis. DARC has found that balance with a scoring algorithm that is informative, but fast.

DARC's scoring algorithm utilizes ray casting, as the name implies. Ray casting is also used in the computer industry as a method for rendering graphics rapidly. In computer graphics, rays are cast to a point that is a virtual representation of an observer of the graphics. Each ray

represents a photon or visual representation of the graphics that is observed. As all the rays reach the observer, the full image of the graphics is observed so you can think of each ray as a pixel on the screen. Ray casting is computationally faster because changes to each ray can be computed simultaneously with Graphic Processing Units (GPUs). GPUs have hundreds of computer cores that can process the same operation on all the cores, but each core performs the operation on a different ray.

DARC's ray casting is very similar. The observation/ray origin point is a spot inside the protein. Rays are cast from this point that go through the surface pocket. So DARC defines the topographies of the protein surface and the molecule docked at the surface by measuring where the ray first hits the protein surface and the molecule (**Chapter One Figure 2**). So GPUs can be used to make the scoring of each docked pose rapidly during optimization. Faster screening allows us to sample a much larger library of small molecules in a reasonable amount of time. This increases the probability of finding a molecule that truly binds p53.

3c) Docking should be done with wild type p53.

The final improvement that is necessary for an initial screening for p53 reactivators is the use of wild type p53 as the screening target. One of the only molecules found to specifically bind and stabilize p53 was found through a virtual screen with the Y220C mutation [19]. The molecule (PhiKan083) binds inside a pocket that forms due to the mutation. While this is an exciting proof that virtual screen can work for p53, it has two major drawbacks. First, the affinity of PhiKan083 is only 150uM, which is very weak for a drug lead. And secondly, while Y220C is a common cancer-associated mutation, it would be far more preferable to target the majority, if not all of the structural mutants.

The low affinity of PhiKan083 is not surprising considering the pocket used for virtual screen. The pocket is a small cleft that is formed because of the tyrosine to cysteine mutation. Such a small cleft is inadequate for binding a high affinity stabilizing drug lead. PhiKan08308 fills the cleft despite being so small (molecular weight of 275, 18 heavy atoms). Such a small molecule it is outside of commonly accepted drug-like size [23]. And since it fills the pocket, it would be difficult to optimize it by adding moieties to increase molecular contacts. So this pocket does not appear to be adequate for finding a drug lead for patients with the Y220C mutation.

Preferably, a small molecule stabilizer would not only target patients with Y220C, but also target all structural mutants of p53. This is why the wild type structure would be more useful during virtual screen. Targeting a surface pocket that is present in the wild type would more likely be present in multiple mutants. Even though structural mutants cause unfolding, there will be a percentage of the protein that is folded. And a molecule that bound the folded protein could stabilize it long enough for it to function as a tumor suppressor. So our goal is to find a location on the wild type that is distant from the most common mutation sites in order to increase the likelihood that we find a broad range stabilizer.

Chapter One

Fast Docking on Graphics Processing Units via Ray-Casting

Karen R. Khar[1], Lukasz Goldschmidt[2], and John Karanicolas[1,3*]

[1] Center for Bioinformatics and [3] Department of Molecular Biosciences,

University of Kansas, 2030 Becker Dr., Lawrence, KS 66045-7534

[2] UCLA-DOE Institute for Genomics and Proteomics, University of California, Los

Angeles, CA 90095

1.1 Abstract

Docking Approach using Ray Casting (DARC) is structure-based computational method for carrying out virtual screening by docking small-molecules into protein surface pockets. In a complementary study we find that DARC can be used to identify known inhibitors from large sets of decoy compounds, and can identify new compounds that are active in biochemical assays. Here, we describe our adaptation of DARC for use on Graphics Processing Units (GPUs), leading to a speedup of approximately 27-fold in typical-use cases over the corresponding calculations carried out using a CPU alone. This dramatic speedup of DARC will enable screening larger compound libraries, screening with more conformations of each compound, and including multiple receptor conformations when screening. We anticipate that all three of these enhanced approaches, which now become tractable, will lead to improved screening results.

1.2 Introduction

There are a number of structure-based methods for predicting small molecules that bind to specific sites on protein surfaces, most commonly active sites, intended for finding lead compounds in drug discovery efforts. High throughput docking tools for "virtual screening" aim to dock thousands of compounds and predict several that will exhibit measurable binding, as a starting point for further optimization. This computational approach can have potential advantages over complementary "wetlab" screening methods because it can be less expensive and time consuming [24]. If successful, hits from a computational structure-based screen may also provide insights that guide the subsequent medicinal chemistry optimization in directions that would not be evident from the chemical structure of the hit compound alone.

Atomistic molecular dynamics simulations and detailed docking approaches are too computationally expensive to allow their direct use for many thousands of independent ligands, as required for most virtual screening applications [25]. Accordingly, several methods have been developed to speed up docking. Some entail using a reduced representation of the receptor, thus reducing the number of calculations associated with each energy evaluation [26-29]. Most approaches fix the receptor conformation or allow only limited conformational changes during docking, to reduce the number of degrees of freedom associated with the search [30-34]. While some methods allow the ligand conformation to vary during docking [32, 35, 36], others carry out independent docking trajectories using a series of pre-built low-energy ligand conformations ("conformers") [30, 37, 38].

We have developed a docking tool called "Docking Approach using Ray Casting" (DARC), as part of the Rosetta macromolecular modeling software suite [39]. Our approach

entails casting a set of rays from the protein center of mass to a series of points mapping out a surface pocket, thus building up a description of the topography of the protein surface as viewed from the protein interior. Since a complementary small-molecule bound to this site should have a complementary topography, we then cast the same set of rays towards the candidate inhibitor. If the inhibitor is indeed complementary to the protein surface, the intersection distance of each ray with the inhibitor should closely match the distance at which the ray reaches the protein surface. In a separate study we find that DARC proves capable of identifying known inhibitors from among large sets of decoy compounds, and we use DARC to identify new compounds active in biochemical assays against the anti-apoptotic protein Mcl-1 [22].

Despite using low resolution scoring and a fast minimization method (both are described in detail below), DARC screening nonetheless remained limited by computational restrictions. Our initial deployment of DARC to screen against Mcl-1 entailed screening only 12,800 compounds (with a maximum of 100 pre-built conformers per compound), and required 152,500 CPU hours to complete this screen. We found that we could achieve a speedup of approximately 6-fold by efficiently neglecting to calculate interactions of rays guaranteed not to contribute to the total score (the "ray elimination" step described later), but DARC remained limited by the size of compounds libraries that could feasibly be screened.

Graphics processing units (GPUs) were originally designed to process parallel, multithreaded 3D graphics via ray tracing, and have since evolved hardware to enable broader types of high throughput processes. Modern GPUs can process mathematical operations, support flow control, and have floating point precision. New libraries such as Compute Unified Device Architecture [40] and Open Computing Language [41] allow development of non-graphics programs for GPUs. These enable an application running on a central processing unit (CPU) to

farm out parts of the job to a GPU. A variety of biomolecular modeling tasks have been adapted for GPU processing, from carrying out quantum calculations to calculating electrostatic surface potentials to stochastic modeling of chemical kinetics and molecular dynamics [42-47]. GPU computing has also been used to speed up certain other structure-based docking tools [1-3, 5, 8, 9, 14].

Given that the ray-casting step underlying our approach is highly analogous to the problem for which GPUs were originally developed, we reasoned that DARC would be highly amenable for porting to GPUs. Since each ray is scored separately and their scores are independent of one another, scoring is intrinsically a parallel process. Here we describe our adaptation of DARC for GPU scoring, leading to a speedup of approximately 27-fold over the corresponding calculation on a CPU alone.

1.3 Methods

Virtual screening using DARC

An overview of the intended DARC workflow for virtual screening is diagrammed in **Figure 1**. The flow is separated into pre-DARC, DARC, and post-DARC stages.

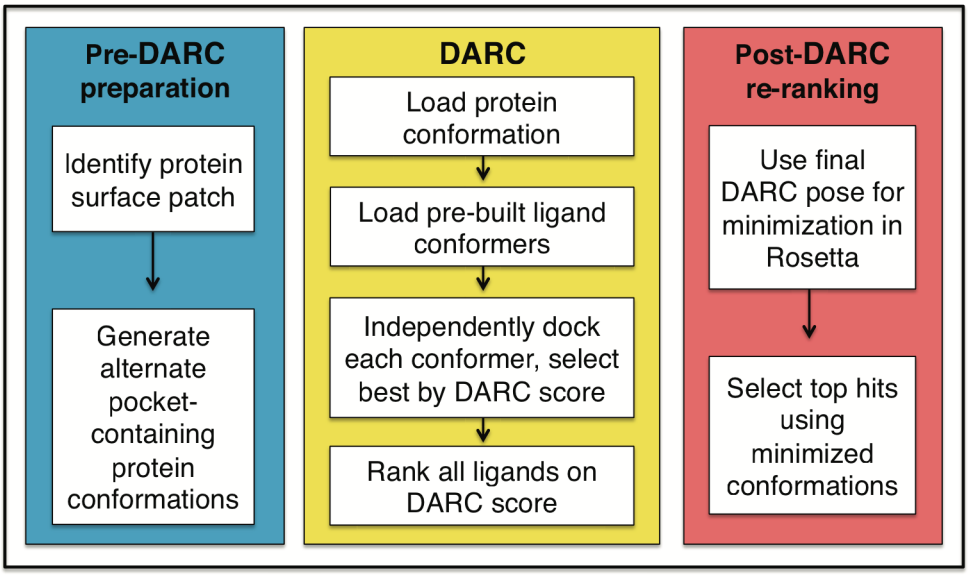

Figure 1: Docking Overview. A schematic diagram of the complete workflow split into three stages: pre-DARC preparation, DARC, and post-DARC re-ranking.

In the pre-DARC preparation stage, a target pocket on the protein is identified and protein structures are generated for use with DARC. DARC was designed for docking at shallow pockets characteristic of those used by small-molecule inhibitors of protein-protein interfaces [48, 49]. The protein conformation is not moved during docking, and can come either from an experimental derived structure or from simulations designed to generate energetically favorable structures with diverse surface pocket shapes at the target site [20].

Each of these protein conformations is then used as a starting point for docking in DARC. Briefly, DARC sequentially carries out rigid body docking for each ligand conformer using a scoring function that maximizes the complementarity of the pocket and ligand shapes when viewed from the protein interior; the following two sections will describe the DARC scoring scheme and optimization protocol in detail. DARC is used to select the optimal conformer and docked pose for every member of the compound library.

The top-scoring model complexes (typically the best 10%) serve as a starting point for further optimization using the all-atom force field in Rosetta. This final energy minimization includes all rotatable dihedral angles (in both the protein and the ligand) as degrees of freedom. Finally, these minimized complexes are re-ranked on the basis of energetic considerations (e.g. interaction energy) as well as structural considerations (e.g. number of buried unsatisfied polar groups). The top scoring compounds can then be advanced for further characterization in biochemical or cell-based assays.

Since DARC scoring considers solely shape complementarity, the intended use of DARC is *not* as a standalone tool for predicting binding free energies, or even for predicting whether any particular compound is likely to bind the target protein. Rather, DARC is intended to provide a fast, low-resolution tool for identifying the likely binding mode of a compound. Our intended workflow thus separates the extensive burden of sampling (carried out by DARC using a crude scoring scheme) from the requirement of a detailed energy function to discriminate active from inactive compounds. This approach is in contrast to complementary methods such as RosettaLigand [50-52], which carries out detailed flexible-ligand docking via Monte Carlo simulations using the all-atom Rosetta energy function but is too computationally expensive to enable routine screening of large compound libraries.

Scoring with DARC

DARC starts from a PDB file of a protein conformation, either from an experimentally derived structure or from biased "pocket optimization" simulations [20]. The shape of a surface pocket is defined using a grid-based method described in detail elsewhere [20]. Briefly, a grid is placed over the protein surface of interest. Based on the coordinates and radii of the atoms

comprising the protein, grid points are marked either “protein” (P) or “solvent” (S). Solvent points which lie on a line between two protein points are then marked as “pocket” (to denote concave regions on the protein surface); this approach was originally used in the LigASite software [53].

The pocket “shell” is identified as those pocket grid points in direct contact with the protein (**Figure 2**, *yellow squares*). Additional grid points are then added around the perimeter of the pocket shell (**Figure 2**, *red squares*), used to mark regions outside the pocket where ligand binding will not lead to favorable interactions (“forbidden” points). The direction from the pocket center of mass to the protein center of mass is defined, and a point 30 Å along this direction is defined as the origin from which rays will emanate (**Figure 2**, *white point*).

The angles and the distances expressing each of the shell points and forbidden points in spherical coordinates (relative to the origin point) are calculated and saved. The number of shell points and “forbidden” points that define the pocket – and thus the number of rays – depends both the grid spacing (typically 0.5 Å) and on the size of the surface pocket. In a typical use case, approximately 7,000 rays are used to define the protein pocket. This collection of vectors (representing points on this small region of the protein surface expressed in spherical coordinates) serves as a mapping of the protein surface topography that should be complemented by a well-docked ligand; the protein conformation and grid points are not directly used in docking beyond this point.

Given the position and orientation of a ligand to be scored, a series of rays are cast from the origin along each of the directions used to map the surface topography. For each ray, the distance at which the first intersection with the ligand occurs is calculated and subtracted from the (stored) distance at which the same ray hit the protein surface (i.e. the shell point). Each ray

contributes to the total score as follows (where c1, c2, c3, and c4 are constants set to 1.0, 1.4, 21.6, and 9.5 respectively):

Ray condition		Contribution to score
1:	Ray hits protein surface point before ligand	c1 * difference between distances
2:	Ray hits ligand before protein surface point	c2 * difference between distances
3:	Ray does not intersect with ligand	c3
4:	"Forbidden" ray intersects ligand	c4

A highly complementary ligand will fill the pocket on the protein surface exactly; each contribution to the score represents some imperfection. A ray that hits the protein surface point before the ligand (condition #1) indicates unpacking in this docked pose (**Figure 2**, *yellow rays*). Conversely, a ray that hits the ligand before the protein surface (condition #2) indicates a steric clash (**Figure 2**, *pink rays*). A ray that does not intersect the ligand (condition #3) indicates that the ligand does not fully fill the surface pocket (**Figure 2**, *orange rays*), and "forbidden" rays that intersect the ligand (condition #4) indicate that the ligand extends beyond the boundaries of the surface pocket (**Figure 2**, *red rays*). Forbidden rays that do not intersect with the ligand do not contribute to the score (**Figure 2**, *purple rays*). The score assigned to the docked pose is taken as the sum of contributions from individual rays, divided by the number of contributing rays.

This approach to scoring is notably different from commonly-used docking tools, each of which estimate energies as the sum of contributions from interacting atom-atom pairs [24].

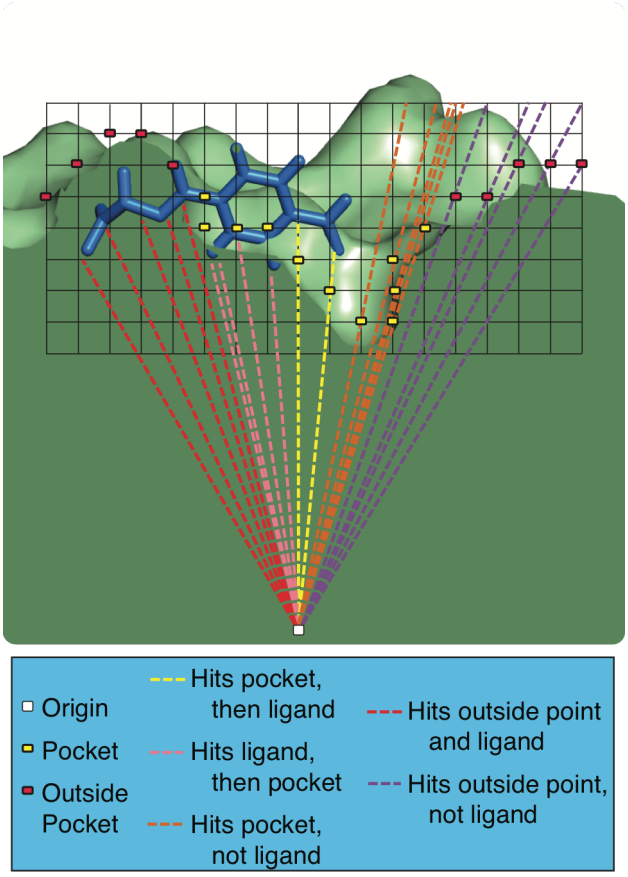

Figure 2: Docking Approach using Ray Casting. A schematic diagram of DARC scoring is shown in cross section. A grid is placed at a region of interest on a protein surface, and used to identify "deep pocket" points. Points that are not in direct contact with the protein surface are removed, leaving behind a set of points that map the topography of the protein surface pocket (*yellow squares*). An adjacent layer of points on the protein surface is then labeled "forbidden" points (*red squares*). Rays are cast from an origin point within the protein (*white square*) at each pocket point and forbidden point. To score a docked pose, the same rays are cast at the ligand (*blue*), and the first intersection (if any) is calculated. The contribution to the total score from each ray is dependent on whether the ray was defined based on a pocket point or a forbidden point, and whether the ray intersects this point before or after it intersects with the ligand. These conditions are described in detail in the main text.

Docking with DARC

Using this method for scoring poses, docking is then carried out using the particle swarm optimization (PSO) scheme [54] to optimize this objective function. Much like a genetic algorithm, this approach entails generating a set of candidate solutions (here called "particles," each of which corresponds to a different docked pose). The position and orientation of each particle is then allowed to adapt in response to the other particles, moving towards the best-

scoring local and global particles with a step size that depends on the relative scores of the particles [54]. After a number of iterations in which all particles move in response to one another, the "swarm" of particles ideally converges upon the globally optimal solution (in this case the lowest-scoring pose).

Though some docking approaches carry out sampling by greedy algorithms (such as incremental construction [55], the most common approaches involve either individual Monte Carlo trajectories that sample Cartesian space or approaches that generate optimal solutions from a population of candidate solutions [24]. The latter class of methods, which include particle swarm optimization and genetic algorithms, make use of coupling between candidate solutions that can be advantageous in guiding the search towards optimal solutions: in the case of AutoDock, for example, a genetic algorithm was found to outperform a Monte Carlo simulated annealing protocol [56]. The potential drawback of this coupling lies in the fact that the inherent need for communication may preclude running candidate solutions on multiple separate machines. In the case of DARC (and virtual screening approaches that use genetic algorithms), however, the scoring function can be evaluated sufficiently rapidly that simulation of all candidate solutions (particles) can reasonably be evaluated on a single processor. Further, in a virtual screening context, running each member of the screening library as an independent job can still allow for parallelization across multiple machines.

In a typical use case, we generate ~7,000 rays to map the protein pocket and dock ligands of ~30 (non-hydrogen) atoms, iterating 200 times over a swarm comprised of 200 particles. This requires evaluating the DARC score for 40,000 docked poses, from a total of 8.4×10^9 potential ray-atom intersections per simulation (210,000 potential ray-atom intersections per pose).

In practice, however, angular bounds can be computed from the docked pose that restrict which rays will intersect with a ligand. In other words, given a ligand atom radius and position relative to the origin, one can compute the maximum and minimum values of each angle required for intersection with this atom. Any rays that fall outside the bounds set by all atoms are guaranteed not to intersect with the ligand, and thus (in a step we call "ray elimination") can be removed from consideration before this docked pose is scored. This reduces the number of ray-atom intersections that need to be computed, and leads to a speedup of about 6-fold when running on a CPU.

DARC using GPU computing

As pointed out earlier, particles encoding the position and orientation of the ligand move collectively in response to one another, making this aspect of docking not naturally amenable to parallel computing. The scoring step, however, entails simultaneously evaluating the scores of 200 particles by summing independent contributions from a large number of rays; this represented a logical candidate for GPU computing.

DARC scoring was implemented on the GPU using the Open Computing Language (OpenCL), which allows the execution of custom programs called "kernels" on a variety of GPUs. Modern GPUs have hundreds of processing cores, thus allowing massive parallel execution of such kernels on a single GPU. Each kernel performs the same operation, but on a different data element from a large set. An important consideration for efficiently adapting DARC for GPU computing was avoiding latency associated with the cost of sending data between the CPU and the GPU.

Our GPU implementation of DARC separates score evaluation (to be carried out on the GPU) from updating particle positions (to be carried out on the CPU) (**Figure 3**). We begin by storing information pertaining to rays (i.e. angles and the distance at which these hit the protein surface) on the GPU before optimization begins: this information will persist there, since it does not change over the course of the minimization. At each iteration of the optimization, information pertaining to *all* particles (i.e. ligand position and orientation) is transferred from the CPU to the GPU in a single step. The GPU uses a first kernel to compute the score contribution for a single ray to *every* particle. In the typical use case described above, each of 7,000 processes is therefore responsible for computing the potential intersection with the 6,000 atoms comprising the swarm (200 particles with 30 atoms each). A second kernel is then applied to each of the 200 particles, to sum the 7,000 contributions from each ray to the score of this particle. Through the use of the second kernel on the GPU, only 200 scores corresponding to particles must be returned to the CPU, instead of 1,400,000 scores from individual rays. Once the scores for each of the particles have been transferred, the CPU uses these scores to update the ligand position and orientation encoded by each particle accordingly.

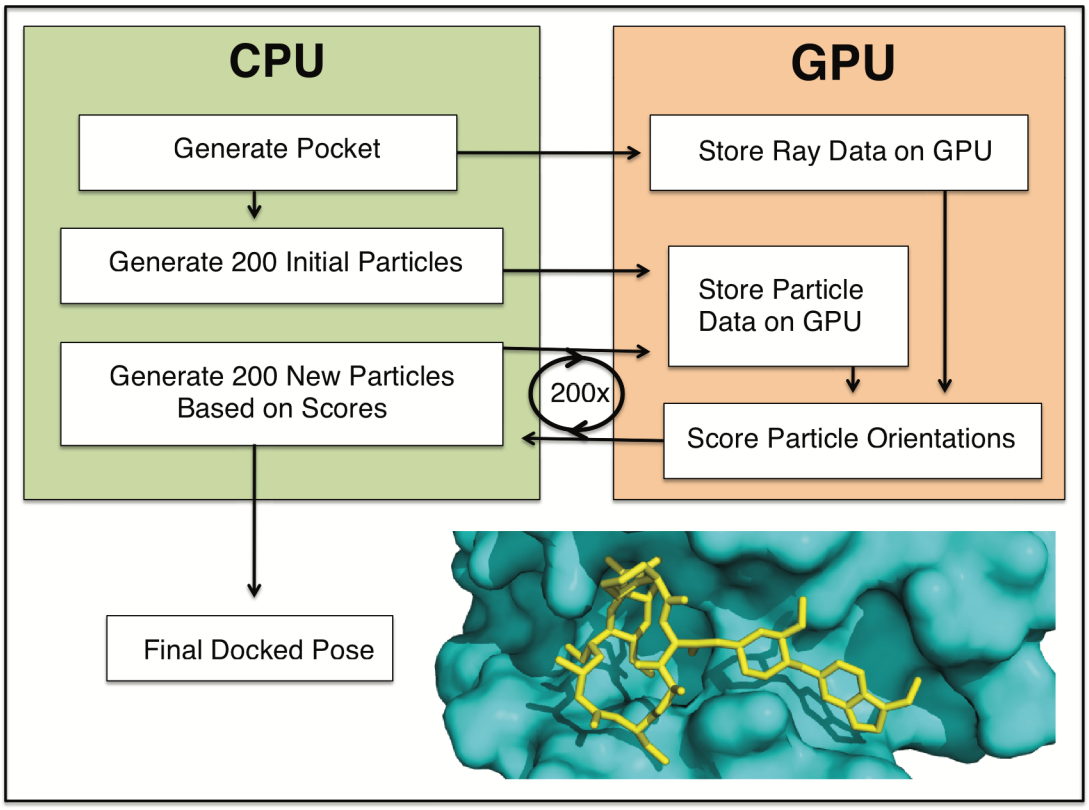

Figure 3: Control flow for GPU-enabled DARC. Control begins on the CPU. The CPU generates the pocket and casts rays at the protein surface, then stores this information on the GPU. The CPU generates 200 "particles" (independent initial ligand orientations to be used in the optimization) and passes each of these docked poses to the GPU. The GPU evaluates the DARC score of each docked pose, and passes these back to the CPU. The CPU uses these scores to update the docked poses accordingly, then sends the new poses to the GPU. This process is repeated 200 times, and the best-scoring particle is reported.

DARC PSO scoring on CPU and GPU

DARC scoring on a CPU occurs as follows:

```
Loop over Particles {
   Identify the max/min angles required for intersection with ligand
   Loop over Rays {
      Check if Ray angles may allow intersection with ligand
      If Ray may intersect with ligand {
         Loop over Atoms in current Particle {
            If Ray intersects Atom {
               Calculate distance of first intersection
               Save this distance if it is the lowest of all Atoms
            }
         }
      }
      Save the contribution of this Ray for the current Particle
   }
   Particle score = Sum of Ray scores / Number of Contributing Rays
```

```
}
```

Scoring with the GPU version occurs using two separate two kernels. The first kernel processes one ray per thread as follows:

Get rayID for this process, define current Ray

```
Loop over Particles {
   Loop over Atoms in current Particle {
      Calculate distance of first intersection with current Ray,
      if intersection occurs.
      Save this distance if it is the lowest of all Atoms in
      this Particle.
   }
   Calculate the contribution of the Ray for the current Particle,
   store it on GPU.
}
```

The second kernel processes one particle per thread as follows:

```
Get particleID for this process, define current Particle

Loop over Ray scores for this Particle {
   Add to current score
}

Particle score = Sum of Ray scores / Number of Contributing Rays
```

Running DARC in Rosetta

DARC is implemented in the Rosetta software suite [39]. Calculations described here were carried out using svn revision 52964 of the developer trunk source code. Rosetta is freely available for academic use [57], with the new features described here included in the 3.6 release.

The standard Rosetta can be built enabling GPU processing as follows (it may be necessary to alter rosetta_source/tools/build/basic.settings to add the address of individual OpenCL headers):

```
scons mode=release extras=opencl bin
```

Input files for small molecules are generated in two steps. The first involves downloading the ligand in the SMILES format from the ZINC database [58], then creating a pdb format file with multiple conformers with using the Omega software [30, 59, 60] as follows:

```
OpenEye/bin/omega2 -in molecule.smi –out molecules.pdb
–maxconfs #conformers
```

When creating multiple conformers, they can be separated by babel as follows:

```
babel –ipdb molecules.pdb –opdb molecule.pdb -m
```

In the second step, a parameter file for the ligand is created with babel and the Rosetta python app molfile_to_params, as follows:

```
babel –ipdb molecule.pdb –opdb molecule.sdf
molfile_to_params.py –c –nKHR –pmol molecule.sdf
```

The Rosetta command line used to generate a set of rays (rays.txt) that define a protein pocket topography is as follows (for target residue number 105 of protein Bcl-xL with the file 2YXJ.pdb):

```
make_rayfiles.linuxgccrelease –input_protein_file 2YXJ.pdb
–central_relax_pdb_num 105
```

The Rosetta command line used to run DARC on a GPU using these input files is as follows:

```
DARC.opencl.linuxgccrelease -input_protein_file 2YXJ.pdb
-input_ligand_file molecule.pdb -extra_res_fa molecule.params
-eggshell_triplet rays.txt -gpu 1
```

1.4 Results

Determining suitable stopping criteria

The two key parameters that determine the DARC runtime are the number of particles and the number of iterations. In order to determine the extent of sampling required for adequate convergence, we evaluated the difference in DARC score obtained from simulations of varying computational requirements against the score obtained from an intensive "gold-standard" simulation. As a model system, we randomly selected a compound from the ZINC database of commercially available compounds [58], ZINC00057615, and docked a single conformer of this compound to a pocket on the surface of the protein Bcl-xL (PDB ID 2yxj).

We initially fixed the number of particles at 200, and sequentially extended the number of iterations from 10 up to our "gold standard" value of 1000 iterations. As expected, increasing the length of our trajectories led to progressively lower final scores (**Figure 4a**), at the expense of a linear increase in (CPU) runtime (not shown). While the docked score decreased rapidly at first, much of the improvement had already been realized after 200 iterations: extending the trajectory beyond this point led only to a modest decrease in score. For this reason, we adopted 200 iterations as our "typical use" value.

We then turned to the number of particles for inclusion, and carried out an analogous experiment. Using 200 iterations in all cases, we sequentially increased the number of particles from 10 up to our "gold standard" value of 1000 particles. As expected, increasing the number of particles similarly led to better solutions (**Figure 4b**), again with a linear increase in runtime (not shown). Based on the diminishing benefit of including a large number of particles, we adopted 200 particles as our "typical use" value.

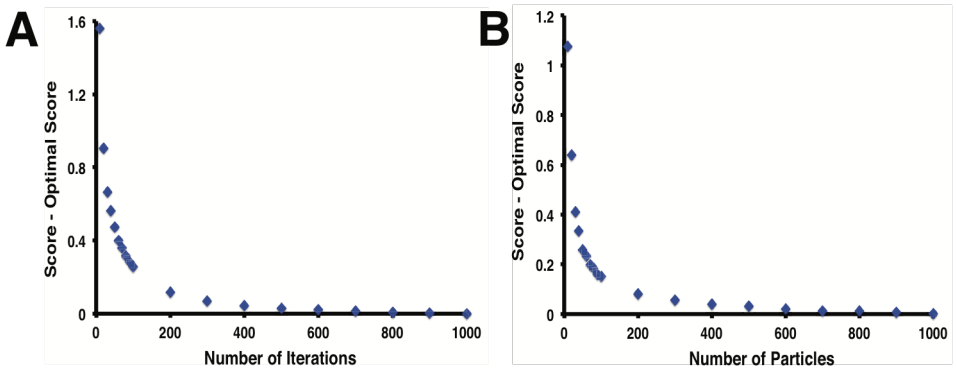

Figure 4: Effect of the number of particles and the number of iterations on DARC score. To determine the number of particles and number of iterations required for reasonable convergence of the DARC score, docking was carried out with (A) an increasing number of iterations while holding the number of particles fixed at 200, and (B) an increasing number of particles while holding the number of iterations fixed at 200. Differences in score are reported relative to the "gold standard," taken to be the most extensive simulation in the set (i.e. 1000 iterations or 1000 particles).

To put these results in the more pragmatic context of virtual screening experiment, we then compiled a set of 1000 randomly selected compound from the ZINC database [58], and evaluated how the extent of sampling would affect the ranking of these compounds against the same Bcl-xL surface pocket. We started with a "gold standard" ranking of each member of our library, by carrying out docking with DARC using 1000 particles and 1000 iterations. We marked the top-scoring 10% of the library (100 compounds) as "hits," then asked how many of these "hit" compounds would remain in the top 10% if docking was carried out using a reduced

number of iterations and particles. We found that 94 of the 100 hit compounds were recovered in the top-scoring 10% using our "typical use" parameters of 200 particles and 200 iterations (**Figure 5**), with little benefit associated with more extensive sampling. We therefore carried forward these values for the further studies described below.

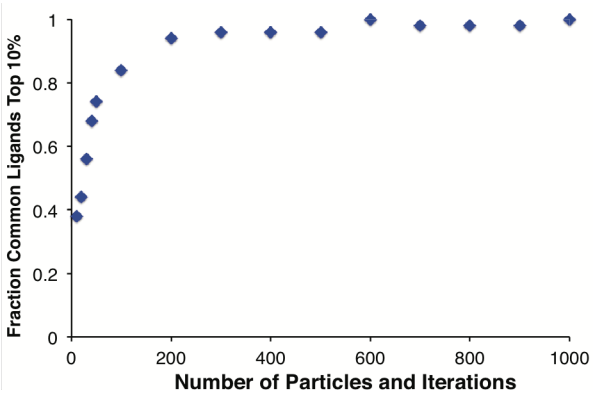

Figure 5: Effect of the number of particles and the number of iterations on the "hit" compounds selected. The most pragmatic measure of convergence is the identity of the "hits" to be advanced for further evaluation. The top scoring 10% of the compound library from the most extensive docking simulations were considered to be the "gold standard" hits. With increasing computationally intensive simulations (by together increasing the number of particles and the number of iterations), an increasing fraction of the hits are members of the "gold standard" set.

DARC speedup on Graphics Processing Units (GPUs)

All timing comparisons described below were carried out using a GeForce GTX 580 GPU, which can run 1024 threads concurrently, and a Dual Intel Xeon E5-2670 CPU using one thread.

As a first timing benchmark, we evaluated the time needed to carry out docking using the same model system described earlier: a single conformer of ZINC00057615 docked against a pocket on the surface of the protein Bcl-xL. Based on our typical grid spacing (0.5 Å) and the size of the surface pocket we would typically use about 7,000 rays to describe this pocket; for

benchmarking, we instead reduced the grid spacing to generate 93,000 initial rays then varied the number of rays used in docking by generating subsets of this large collection.

As expected, the time required to complete this calculation scales approximately linearly with the number of rays and the number of particles, whether carried out entirely on a CPU (**Figure S1a**) or with the help of a GPU (**Figure S1b**). While the scaling is similar, however, the calculations are completed much more quickly using the GPU: in a typical uses case (7,000 rays, 200 particles and 200 iterations), the CPU takes 93 seconds to carry out the calculation and the GPU takes 3.4 seconds, corresponding to a 27-fold speedup (**Figure S1c**).

Similar behavior is observed when docking a single conformer to a surface pocket at the functional site of another protein, Mdm2 (**Figure S1d-f**). Due to the different size and shape of this pocket, the same grid spacing would lead to only 3,000 rays to describe this protein surface. Under these conditions (again with the standard 200 particles and 200 iterations), the calculation would take 47 seconds using the CPU alone, or 3.2 seconds using the GPU (a 15-fold speedup).

We next tested the scaling of time with regards to the number of atoms in the ligand, docking to Mdm2 using 5,000 rays and 200 particles. We used a series of ligands containing 20 (ZINC0043625), 25 (ZINC00469420), 30 (ZINC01280234), 35 (ZINC01298436), and 40 (ZINC02091520) non-hydrogen atoms. We find that the time required for this calculation on the CPU alone is not linearly related to the number of ligand atoms (**Figure 6a**), because the geometry of the ligand dictates how much of the calculation can be avoided through the "ray elimination" step. In all cases, carrying out this calculation using the GPU results in a speedup of about 25-fold (**Figure 6b**).

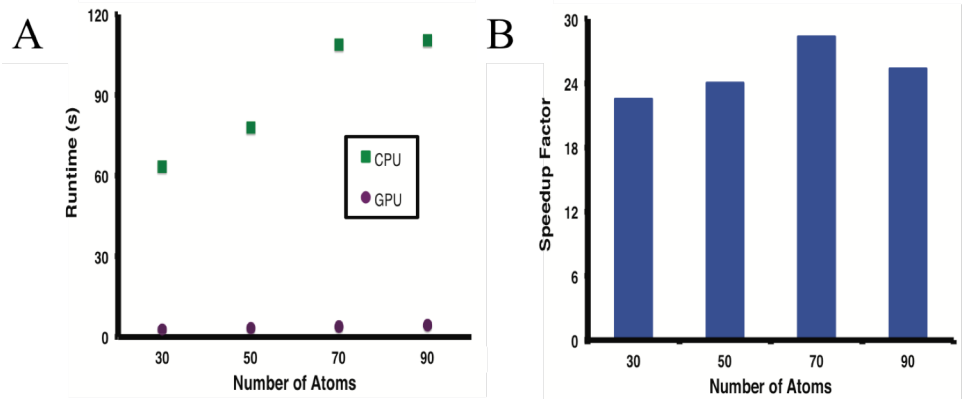

Figure 6: Dependence on number of atoms in the ligand. Ligands of varying sizes were docked using DARC. A) Time required to complete the optimization, using a CPU alone or with the GPU. B) Speedup factor, reported as the ratio of the time required using the GPU to the time required using the CPU alone.

While the typical-use speedup in the examples here is dramatic, we note that these data in fact *downplay* the true difference stemming from the use of the GPU for these calculations. In the timings we have reported above, the algorithm carried out on the CPU includes the "ray elimination" step that reduces the number of potential ray-atom intersections to be considered. The GPU calculations described above, however, do not include this step; we made a design decision not to take advantage of the potential for fewer calculations on the GPU, because the ray elimination step would cause threads to become asynchronous. This branch divergence in the kernel execution would lead to uncoalesced memory access, slowing the total time required for the calculation. For a straightforward comparison, we therefore additionally tested a variation of the CPU code that does not include the "ray elimination" step, and a variation of the GPU code that does include this step (**Figure 7**). We find that the GPU optimization requires a very similar time to reach completion regardless of whether or not the "ray elimination" step is used, justifying our design decision. As expected, the opposite holds for the CPU version: performance is significantly slower when the "ray elimination" step is not used. In a typical use case for Bcl-

xL comprising 7,000 rays, the GPU version of DARC without the "ray elimination" step is completed about 180-fold faster than the same calculation on the CPU alone.

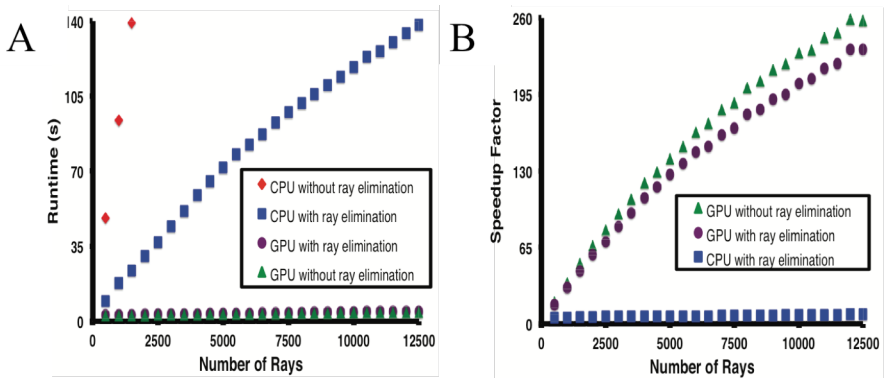

Figure 7: Comparison of DARC optimization with and without the "ray elimination" step. The "ray elimination" step is found to significantly improve performance of DARC on the CPU alone, but made little difference when the GPU is used. A) Time required to complete the optimization, using a CPU alone or with the GPU. B) Speedup factor, reported as the ratio of the time required using the GPU to the time required using the CPU alone.

Analysis and implications of DARC speedup on GPUs

As described earlier, a key motivation in adapting DARC for GPU processing stemmed from the practical limitation on the size of compound libraries that can be routinely screened: our initial deployment of DARC entailed screening only 12,800 compounds, and required vast computational resources. To test whether extending our library size would improve the quality of compounds identified – subject to the DARC objective function – we carried out an experiment to determine the effect of library size on the resulting hit compounds. Since virtual screening involves drawing those few compounds from the extreme end of the distribution of scores, we trivially anticipated that increasing library size would lead to a monotonic improvement in the score of the top-scoring compound. Accordingly, we built a library of 46,000 compounds corresponding to a drug-like subset of the ZINC database [58], then used this to build further incrementally smaller libraries (decreasing the library size 10-fold each time). We carried out a

virtual screen of each library against two protein targets, interleukin-2 (PDB ID 1m47) and Mdm2 (PDB ID 4jvr), and unsurprisingly observed a considerable decrease in the DARC score for the top-scoring compound as we increased our library size (**Figure 8**). These results serve to illustrate the fact that chemical space is not heavily covered by (random) compound libraries of this size, and that computational enhancements that enable screening of larger compound libraries are likely to enable identification of more optimal compounds for the target of interest – subject to the strong caveat that compounds with better scores may not necessarily show more activity, depending on the objective function.

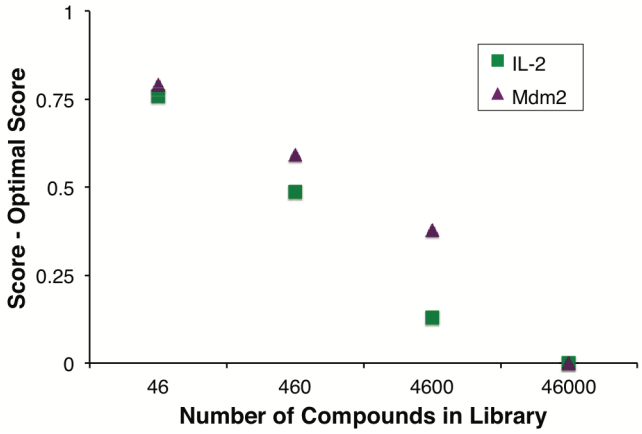

Figure 8: The GPU-enabled speedup facilitates screening of larger libraries, which in turn allows better-scoring ligands to be identified. Compound libraries of increasing size were screened against interleukin-2 and Mdm2. As expected, screening larger libraries led to identification of compounds with better scores. All scores are reported relative to the lowest scoring ligand in the largest set.

With an eye towards additional optimization of our GPU adaption of DARC in the future, we sought to better understand the rate-limiting step in our current implementation. Based on the relatively weak dependence of the GPU timing on factors that dictate the number of potential ray-atom intersections to be considered (number of rays, number of ligand atoms, and number of particles) (**Figure 9**), we surmised that GPU calculation itself was not the rate-limiting step in the overall calculation.

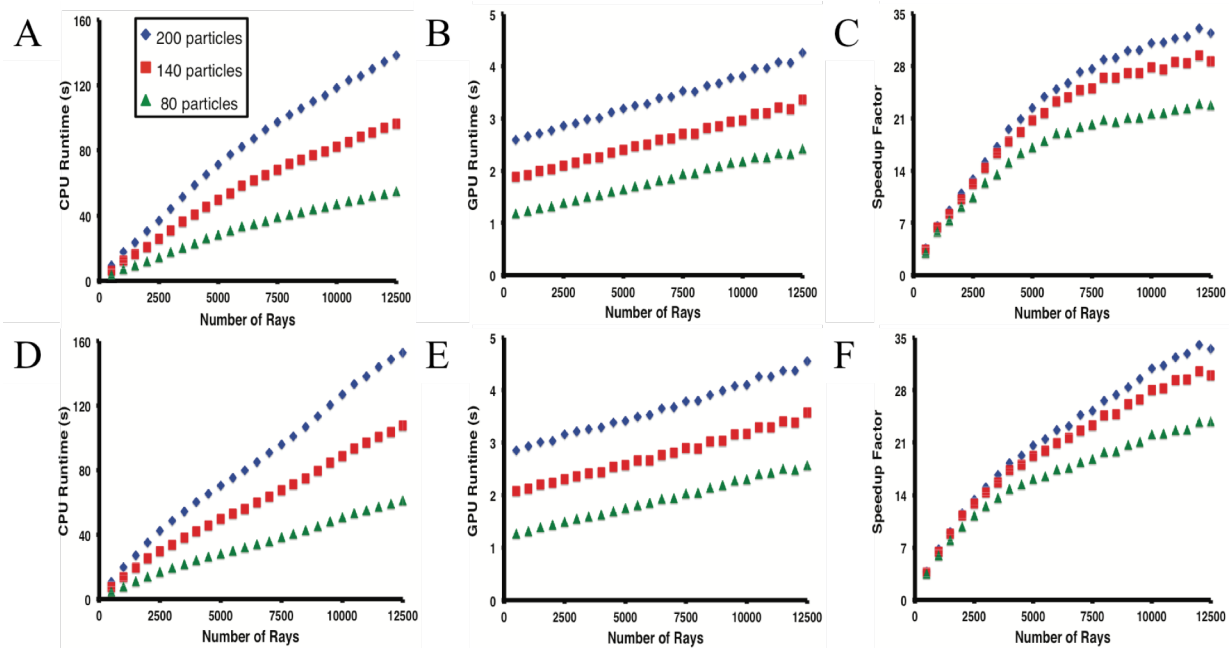

Figure 9: Dependence of simulation time on number of rays. A single ligand conformation was docked in the Bcl-xL (A-C) or Mdm2 (D-F) surface pocket, independently varying the number of rays defining the pocket and the number of particles. A,D) Time required to complete the optimization using a CPU. B,E) Time required to complete the optimization with the GPU. C,F) Speedup factor, reported as the ratio of the time required using the GPU to the time required using the CPU alone.

To test this hypothesis, we carried out minimizations of Bcl-xL (with our typical use case of 7,000 rays), but varied the number of iterations while keeping the product of the number of iterations and the number of particles was fixed. As expected from fixing the total number of potential ray-atom intersections to be computed, the CPU alone required an almost identical amount of time to complete each of these calculations, confirming that calculating ray-atom intersections was indeed rate-limiting. If the same step was rate-limiting when carried using the GPU implementation, we would expect each of these calculations to again require a fixed amount of time for completion. In contrast, the use of the GPU allowed faster calculations upon decreasing the number of iterations but using more particles: this in turn lead to a greater overall speedup with respect to the CPU implementation (**Figure 10**). We further found that up to eight

independent GPU-DARC threads running on eight (CPU) cores required the same time for completion as a single GPU-DARC thread, despite *sharing* a single GPU (not shown). Collectively these observations suggest that given a "typical use" setup in the current implementation, the portion of the calculation carried out on the GPU is *not* rate-limiting; rather the rate-limiting step lies either in the CPU-GPU communication step occurring once per iteration or, more likely given that a GPU can be effectively shared between multiple cores, lies in the few remaining calculations taking place on the CPU. The implications of these observations will be discussed further below.

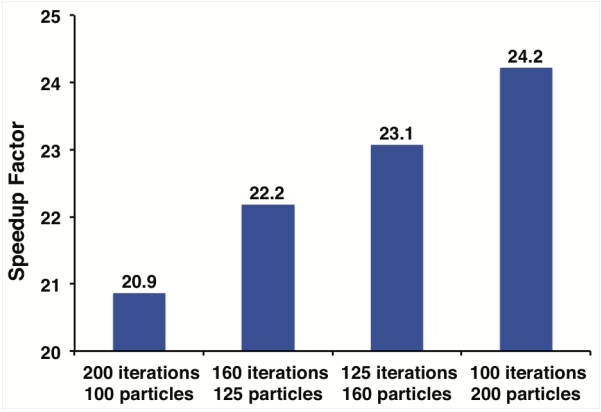

Figure 10: Runtime dependence on the number of particles and the number of iterations. A series of optimizations are compared in which the number of calculations (and thus the total time required) on the CPU is constant, and the speedup factor is reported as the ratio of the time required using the GPU to the time required using the CPU alone. The benefit of using the GPU is enhanced when individual GPU tasks are larger (more particles), allowing fewer CPU-GPU communication steps.

1.5 Discussion

Here we describe a faster implementation of the DARC ligand-docking program enabled by GPU computing. By carrying out the scoring step on GPUs, we achieve a speedup a 180-fold speedup over the same calculation carried out on a CPU alone. This calculation could be carried out 6-fold faster on the CPU by eliminating certain interactions from consideration before

scoring, but this algorithmic difference did not affect timing on the GPU. Accordingly, the GPU-enabled code is therefore 27-fold faster than our fastest CPU-only code. This speedup was achieved using a modern GPU that is relatively inexpensive (less than $500).

Several other docking tools have recently been adapted to make use of GPU computing, leading to reported speedups in ranging from 2-fold to 100-fold (**Table 1**). Methods that require long serial trajectories, such as those built upon molecular dynamics [1, 2], require frequent CPU-GPU communication. This in turn leads to latency that limits the speedup achievable through GPU computing. A feature common to tools that achieve dramatic speedup is the ability to break up tasks into parallel subtasks that are either very numerous (i.e. DARC, PLANTS, GPUperTrAmber) or else individually computationally intensive (i.e. AutoDock Vina): either approach leads to long stretches of computing carried out exclusively on the GPU without the need for communication with the CPU. By extension, for applications such as DARC in which the objective function can be easily ported for calculation on the GPU, optimization schemes that simultaneously consider multiple candidate solutions (such as genetic algorithms and particle swarm optimization) are exceptionally well-suited to achieve dramatic speedups through relatively minor code changes.

Table 1: Comparison of GPU-enabled docking tools. Docking methods have been adapted for GPU computing using a variety of strategies. These require different degrees of CPU-GPU communication, and accordingly enable varying speedups relative to the analogous CPU-only protocol.

Docking tool	GPU enabled functionality	Speedup
Molecular dynamics combined with docking	Molecular dynamics	2-3x [1]
DOCK6	Amber scoring (molecular dynamics)	6.5x [2]
ZDOCK / PIPER / Hex	Fast Fourier Transforms	15x [3]
MolDock	Initially only scoring, then also differential evolution	27x [5]
DARC	*Simultaneously scoring multiple particles*	*27x*
PLANTS	Concurrent grid-based search	60x [8]
AutoDock Vina	Runs docking concurrently from different starting orientations	62x [9]
GPUperTrAmber	Scoring very large systems by decomposition	100x [14]

In the case of our GPU-enabled DARC implementation, these insights provide inspiration by which further speedups may be possible. As noted earlier, the fact that all particles move collectively in response to one another does not make porting the entire PSO calculation to the GPU an attractive approach for achieving further speedup. However, the fact that eight CPU cores can share a single GPU without noticeable slowing implies that the GPU remains under-utilized in our current implementation; this in turn suggests that the current framework could be adapted by increasing the size of the problem allocated to the GPU at each iteration. Through further careful examination of the relationship between the number of particles and the number of iterations (**Figure 10**), it may prove possible to achieve equivalent convergence more quickly more particles and fewer iterations. Alternatively, further parallelization may be realized by bundling particles corresponding to different ligand conformers for simultaneous scoring on the GPU, rather than carry out separate (serial) optimization of each conformer. The fact that

additional calculations can be likely carried out on the GPU with little additional cost also offers the opportunity to fundamentally change the DARC scoring paradigm: either by simultaneously using multiple sets of rays originating from distinct origins within the protein, and/or by adding new components to capture effects of electrostatics. In short, any enhancement that increases the computational burden per iteration that is carried by the GPU is likely to yield further speedup relative to the CPU alone.

Given fixed computational resources allocated for completion of a project, the ability to carry out docking more rapidly will have profound implications for applications of DARC. In the most obvious case, this speedup will allow screening against very large libraries that previously may not have been tractable, for example the complete ZINC database [58] or a library of hypothetical compounds likely amenable to straightforward synthesis [61]. Even in cases in which a relatively small library of interest is to be screened (for example, computational screening of a library of compounds currently available in-house), this speedup will allow an increase in the number of conformers screened per compound; this in turn is expected to reduce the number of false negatives in the screen, by increasing the likelihood of including an active conformer. This speedup may further allow the use of multiple pre-built receptor conformations for docking [62-68], providing a means to implicitly represent receptor flexibility and thus allow further diversity in collection of hits identified.

1.6 Acknowledgements

We thank Ragul Gowthaman for valuable discussions. We are grateful to OpenEye Scientific Software (Santa Fe, NM) for providing an academic license for the use of Omega.

Chapter Two

A Universal Binding Site that Allows Rescue of Multiple Destabilized p53 Mutants

Karen R. Khar[1,2], Atul Ranjan[3], Tomoo Iwakuma[3], and John Karanicolas[1,2,4*]

[1] Center for Computational Biology, University of Kansas, Lawrence KS

[2] Program in Molecular Therapeutics, Fox Chase Cancer Center, Philadelphia, PA

[3]Department of Cancer Biology, University of Kansas Medical Center, Kansas City KS

[4]Department of Molecular Biosciences, University of Kansas, Lawrence KS

2.1 Abstract

Mutations in the tumor suppressor p53 are found in about half of all human cancers. Many of these mutations deactivate p53 by reducing the stability of its "core domain", which is the DNA binding domain. Certain small molecules that stabilize the native fold, known as "reactivators", have been shown to rescue activity of these p53 mutants. This has spurred efforts in search of potent stabilizing compounds, since these could serve as a starting point for development of new drug leads. Here we describe our identification of a novel druggable pocket on the surface of p53's core domain. We used structure-based virtual screening to select 28 candidate ligands to complement this pocket, and found that 11 of these stabilize recombinant p53 core domain in biochemical assays. When applied to cancer cell lines that harbor destabilized mutant forms of p53, 4 of these compounds restore p53's folded structure and lead to upregulation of its target genes. Because compounds designed against a single site can rescue activity of many different destabilizing mutations, this surface pocket may represent a starting point for development of a new class of "universal" drugs that re-activate a broad spectrum of p53 mutants.

2.2 Introduction

The p53 cell pathway plays a key role in human cancer [69, 70]. p53 is a tumor suppressor, and in healthy cells is continuously produced but maintained at a low level by MDM2/MDMX ubiquitination [7, 71]. In the presence of cell stressors such as DNA damage, however, p53 levels are allowed to increase [72]. This increase leads to a variety of responses; the most well-characterized of these is through p53's activity as a transcription factor, which directly upregulates genes that lead to cell cycle arrest, apoptosis, DNA repair, senescence and inhibition of metastasis [73-76]. Thus, the p53 pathway is responsible for preventing cells from becoming transformed into tumor or cancer cells.

Disruption of the p53 pathway in human cancers can occur by upregulation of MDM2/MDMX, or by direct genetic alterations to the p53 gene. p53 is the single most-frequently mutated gene in many different cancers [11]. Further, cancers with p53 mutations are often associated with resistance to therapy and poor prognosis [77]. Most genetic alterations to p53 are missense mutations that occur in p53's DNA-binding (or "core") domain [78, 79]. Analysis of these mutations in the core domain has led the field to classify p53 mutations into two broad categories. The first type are "DNA contact" mutations, which occur at residues making direct contact with DNA. These lead to significant loss of affinity for the p53 response element, so that it can no longer function as a transcription factor. This in turn allows the unchecked replication that leads to cancer [18].

The second category of p53 mutations are known as "structural" or "conformational" mutations [18], and it is estimated that 30-40% of cancer-associated mutations fall into this class [10]. p53's core domain is only moderately thermodynamically stable at physiological

temperature [80]; mutations in this class often decrease p53's stability by 3-4 kcal/mol, corresponding to a decrease in melting temperature of 5-7 °C [80, 81]. Because of this the protein is mostly unfolded at 37 °C, but at lower temperatures it can recover its native fold and thus bind DNA [82, 83].

This mechanistic understanding of p53 mutations naturally led to new approach for drug discovery. It has been shown that exogenously restoring expression of p53 can lead to tumor regression *in vivo*, in a variety of tumor models [84]. In principle, a compound that selectively binds and stabilizes natively-folded p53 could restore wild-type activity to these structural mutants: this would be expected to have an analogous effect as exogenous addition of the wild-type protein, and thus could stimulate apoptosis / growth arrest in cancer cells.

An early proof of concept study in support of this idea came from a peptide, CDB3, that was derived from p53's binding partner 53BP2. Addition of this peptide was found to restore DNA-binding to the p53's R249S structural mutant, and thus represents one of the first examples of what has become known as "pharmacological chaperones" for p53 [85]. Other groups have since identified small molecules that bind and stabilize other structural mutants which rescued p53 function and lead to tumor regression [86-88]. These stabilizing compounds are now typically called "reactivators": they rescue by binding the native conformation of p53's structural mutants, and shifting the equilibrium towards the active conformation [89].

To date there have been 11 important reactivators of p53 that have been described in the literature [90], though the term reactivator in this context has since come to encompass not only compounds that directly bind and rescue p53, but also compounds that rescue p53 activity function *indirectly* without binding to p53. Of these 11 reactivators, eight of them have been shown to be effective against one or two structural mutants (Chetomin, NSC319726, p53R3,

Prima-1, PK7088, Stictic Acid, STIMA-1, WR-1065) and very few have shown activity against three or more mutants (Mira-1, RITA, SCH529074). RITA was shown to reactivate the largest number of p53 mutants (seven), but it also induces apoptosis in a cell line expressing wild type p53. At this stage the precise mechanism of action for RITA and several of the other reactivators (p53R3, SCH529074, WR-1065, STIMA-1) still remains unclear, including even whether or not they interact directly with p53 itself [90]. Stictic acid was identified from a computational screen against a surface pocket near Cys124 on p53, but it has not yet been confirmed that the compound indeed interacts with this surface of p53 (let alone whether any potential interaction is covalent or non-covalent).

That said, the mechanism of action for some of these compounds has been established, and these vary widely. For example, NSC319726 reactivates p53 by serving as a zinc metallochaperone: it increases the zinc concentration in the cell, which can compensate for p53 mutations that reduce binding affinity for zinc (needed for p53's structural integrity) [91]. Chetomin reactivates p53 by increasing levels of heat shock protein 40 (HSP40), which in turn binds p53 and facilitates folding to the active conformation [90].

PhiKan083 and PK7088 represent the best-characterized direct stabilizers of p53's core domain. PhiKan083 was identified through structure-based computational screening [19], whereas PK7088 derived from biophysical screening [92]. Nonetheless, both interact with p53 in a very similar manner: the destabilizing mutation Y220C introduces a new cleft on the surface of p53, and both of these compounds target this feature. Because of this, both classes of compound are selective for this Y220C mutant: this is advantageous because it avoids potential on-target toxicity with wild-type p53, but is also a limitation because these compounds cannot rescue other destabilized mutant. The potency of these compounds are also potentially limited by this small

and exposed binding mode: while medicinal chemistry optimization led to a derivative that improved upon PhiKan083's activity 5-fold [93], both compounds nonetheless started with binding affinity weaker than 100 μM.

In light of the limitations inherent to these previously described reactivators, we set out to identify a new druggable pocket on the surface of p53's core domain. If successful, we anticipated that a new site could enable identification of compounds that bind more potently and selectively than those described to date, and also reactivate a broader spectrum of destabilized p53 mutants. Accordingly, we therefore used structure-based computational tools to search for candidate druggable sites that would not require covalently binding compounds, at locations not overlapping with the most frequently occurring cancer-associated mutations.

2.3 Results

Selecting a "Druggable" Pocket on p53's DNA Binding Domain

Our goal is to find a molecule that binds to p53 with high specificity, so we are using rational design with a high-resolution crystal structure to find small molecules that dock in a surface pocket with high steric and chemical complementarity. Structure-based virtual screens with p53's DNA Binding Domain have been attempted, yet they have found few hits, which had low affinity. This may be because these screens are not targeting regions of the protein that are amenable to binding a small molecule. Choosing a location on a protein that is druggable has proven to be a challenging problem. However, our lab has recently developed a method for predicting druggability of surface pockets.

It has been hypothesized that druggability correlates with the propensity of a surface location to form larger and more varied pockets. This is because a protein will often have to

undergo some conformational changes in order to be able to bind a molecule. Therefore, it would have to be more flexible than a site that is not druggable. Our lab's method computationally samples a protein's flexibility using Rosetta Relax, which involves backbone perturbation and full atom minimizations in a Monte Carlo simulation. However, our variation has a bias towards larger pockets by including an additional energy term that improves the scores for conformations with larger pockets. We refer to this method as a Pocket Optimization. We found a correlation between sites known to be druggable and the site's capacity to form large pockets using Pocket Optimization. And just as importantly, the remaining sites on the protein surface would not form pockets [20].

We measured pocket volumes globally across p53's DNA Binding Domain crystal structure (excluding the DNA binding site) and found three pocket sites (**Figure 1a-c**). The measured pocket volumes only include deeper regions in order to prevent bias towards very shallow pockets so the available volume that a molecule could bind is actually much larger. Using these three sites, we ran 3-6,000 simulations per pocket with Pocket Optimization. Pocket One was clearly able to open pockets of significant size, which were much higher than 150 $Å^3$ in over 70% of it's decoys while the others were only able to reach 150 $Å^3$ in 10-20% of their decoys (**Figure 1d**). We also find it to be an appropriately druggable pocket because it is a significant distance from the DNA binding region, the zinc binding region, and the most common structural mutation sites (**Figure 1e-f**). Small molecule binding at or near the DNA binding region could inhibit DNA binding which would prevent p53 from being able to function as a transcription factor. It would also be detrimental to inhibit zinc binding because loss of zinc is known to be destabilizing [94]. And it is preferable for the pocket to be distant from common cancer-associated mutations because they may alter the pocket shape and loose affinity for a

stabilizer that bound there. We aim to find a stabilizer with affinity for a majority of destabilizing mutations in order to target more cancer patients.

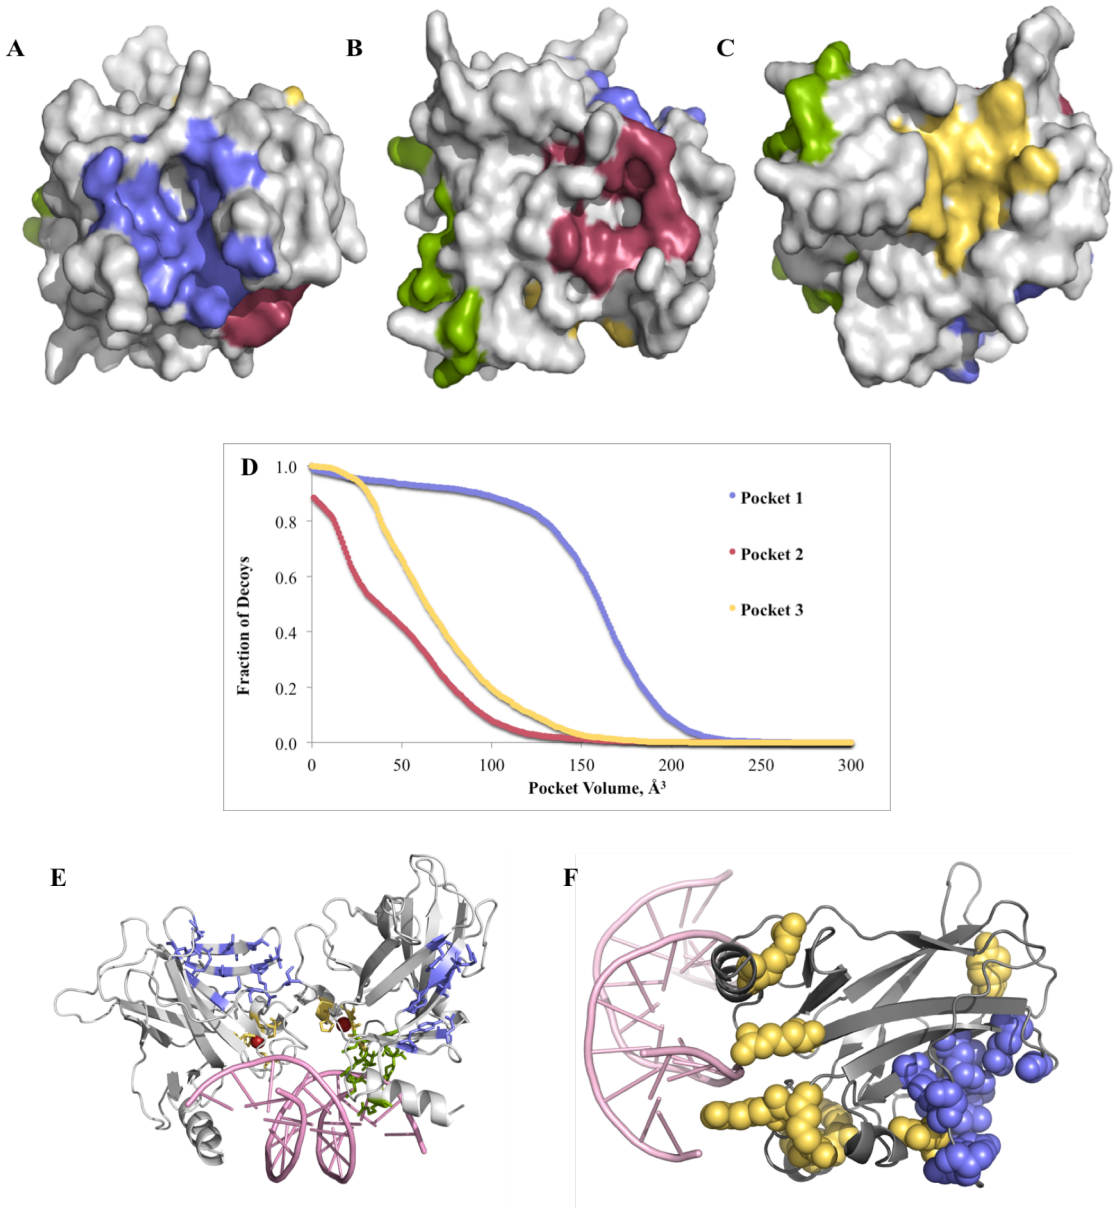

Figure 1: Analyzing Pockets in the Crystal Structure. A-C) Protein is shown in surface view in grey with 3 rotations of the structure to highlight pockets. Residues that contact DNA are colored green while residues that make up pocket one (A), two (B), and three (C) are blue, red, and yellow respectively. D) Fraction of decoys from Pocket Optimization simulations that had pocket volumes from 0 to 300 $Å^3$ for each pocket. Volumes exclude the shallowest regions by top shaving 2 Å. E-F) Structure show p53 in cartoon bound to DNA. Pocket residues in blue, residues that contact DNA in green, and zinc contacting residues in yellow. DNA is in pink and zinc sphere is red. F) Top 10 most common cancer-associated mutations in the TP53 database are shown as yellow spheres. Pocket residues are shown in blue spheres.

Virtual Screen for Reactivators

We performed a structure-based screen using a modest-sized small-molecule library using our lab's docking tool, DARC (Docking Approach with Ray Casting). Most docking tools are designed to dock molecules in deep pockets such as those found in enzyme active sites [21]. But compared to enzymes, this pocket is shallow. However, DARC is ideal for a surface pocket because it was designed to optimize contacts of a small molecule to a shallow surface [22].

DARC uses a rigid body structure of a pocket for docking a library of molecules. We used the pocket that was found to be druggable on the crystal structure for docking. However, a crystal structure doesn't account for dynamic changes that are likely to occur in vivo and fails to show specific conformational changes that occur when a protein binds a molecule. For many ligands, the protein would only form the bound conformation in the presence of that ligand because the ligand forms contacts that induce the necessary conformational change that would otherwise be energetically unfavorable. However, the Pocket Optimization simulations done when selecting a druggable site essentially act as an implicit ligand by biasing for open pockets. So it can find conformations similar to what a bound protein would form. Therefore, the Pocket Optimization simulations were needed not only to find druggable pockets, but were employed to find open conformations to use as rigid body structures for virtual screening.

Pocket Optimized structures were analyzed in order to select representative structures that had diverse, large pockets. We were looking for larger pockets in order to complement to larger molecules in order to maximize stabilizing contacts and affinity. Also, once binders are found, further optimization of the hits can be done so a larger pocket allows room for additional moieties to be added in order to increase affinity of hits during the medicinal chemistry optimization phase. Visual inspection of pockets formed during the simulations showed that

there was significant heterogeneity in pocket shapes. This allows us to find more diverse binders in chemical space in order to find a wider range of molecules to improve the probability of finding a stabilizer.

We hypothesized that there were a few general pocket conformations present among the simulations. So we used hierarchical clustering with a natural number stopping criteria to group Pocket Optimized structures by pocket shape similarity. For each group, the clustroid was selected to be the representative structure of the group. We selected simulated structures from the two major clusters for virtual screen in addition to the crystal structure (**Figure 2**). The crystal structure pocket was deeper, while the two pockets from simulation were more spread out.

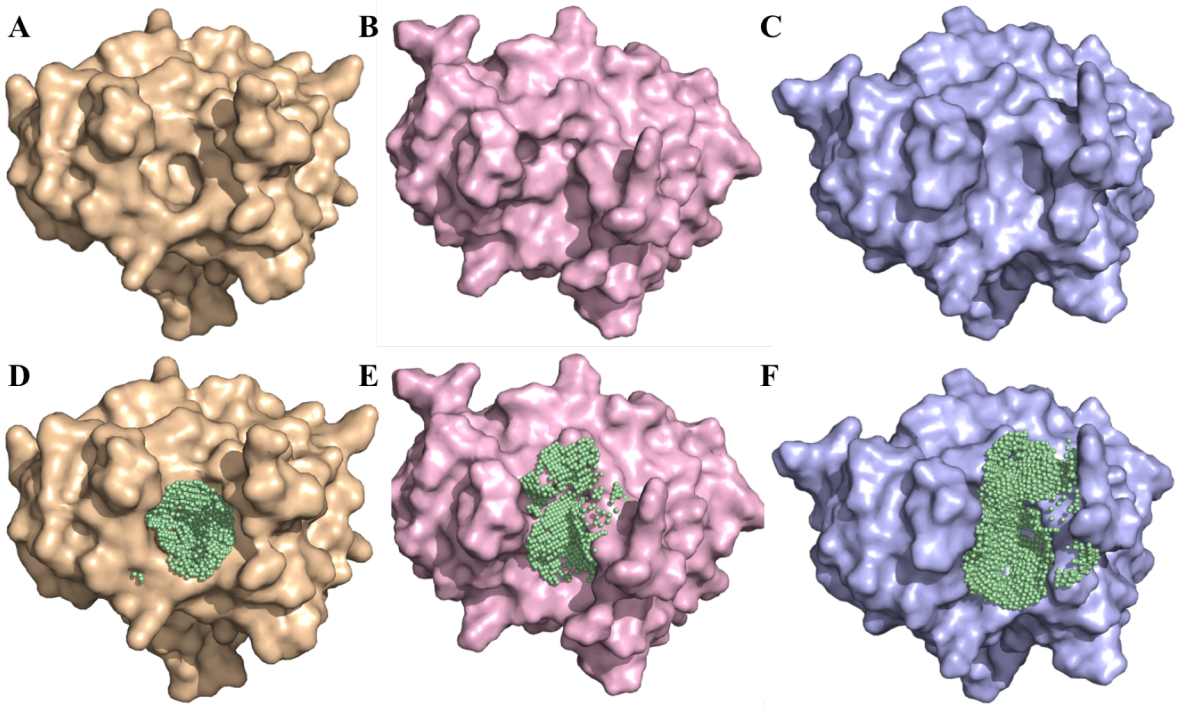

Figure 2: Selecting Multiple Conformations of a Pocket for Virtual Screen. Pockets at a single location from Pocket Optimization simulations were clustered by shape into 2 bins. Representative structures of each were chosen for virtual screen. The crystal structure is shown in a and d. Representative structure of cluster one is in b and e, structure from cluster two is shown in c and e. Docking region is shown in d-f.

We used the crystal structure and two representative pockets from the biased simulation for screening with a virtual library of 64,000 small molecules. We performed the screen with an early version of DARC that used steric minimization. We took the final docked pose from DARC for each pocket/molecule pair and performed additional optimization with a Rosetta minimization. For each pocket, scores were ranked in terms of total Rosetta Energy and Interaction Energy. We visually inspected the top hits and took into consideration the number of hydrogen bonds, unsatisfied polar bonds, and solvent accessible surface area reported by Rosetta. One simulated pocket (**Figure 2c**) lacked favorable hits in comparison to the crystal structure and the other simulated structure so was excluded. Following visual inspection of the top 300 hits, twenty-eight compounds were selected (**Table 1**).

Table 1: Summary of Docking Results. List of small molecules predicted to bind p53 DNA Binding Domain. Final docked pose is scored in term of total and interface energy in terms of Rosetta Energy Units (REUs). Ranked by Interaction Energy.

Zinc ID	Code Name	Source	Total REU	Interaction REU
ZINC12439266	Spiderman	crystal	-475.8	-19.8
ZINC32972442	Thor	simulation	-499.2	-19.3
ZINC22024197	Han Solo	crystal	-483.0	-18.9
ZINC02759838	Ripley	simulation	-498.8	-18.4
ZINC00811037	Sara Connor	crystal	-476.1	-18.2
ZINC04554233	Huxley	crystal	-483.3	-18.1
ZINC05471605	Oprah	crystal	-484.2	-17.8
ZINC02433954	Daenerys	crystal	-480.9	-17.7
ZINC00246066	Katniss	crystal	-484.6	-17.6
ZINC04100029	The Doctor	crystal	-478.4	-17.2
ZINC00883497	Spartacus	simulation	-499.3	-17.1
ZINC14115481	Capt. Hammer	crystal	-478.6	-17.1
ZINC22018013	Buffy	crystal	-480.7	-16.9
ZINC12736934	Shaft	crystal	-480.8	-16.4
ZINC13682491	Bruce Lee	simulation	-496.4	-16.4
ZINC08762382	Xena	simulation	-496.3	-16.3
ZINC07467929	Starbuck	simulation	-497.6	-15.8
ZINC13497494	Evita	crystal	-480.1	-15.4
ZINC12608763	Buck Rogers	simulation	-496.6	-15.4
ZINC77438884	Athena	simulation	-496.2	-14.3
ZINC65505827	Che Guevara	crystal	-472.5	-13.5
ZINC17730245	John	simulation	-494.7	-13.3
ZINC04613946	Jeeves	simulation	-496.4	-12.9
ZINC00641309	Jayhawk	simulation	-493.7	-12.8
ZINC02860908	George	simulation	-495.6	-12.1
ZINC13755423	Jambunathon	simulation	-495.4	-11.9
ZINC12858876	Dumbledore	simulation	-494.2	-11.8
ZINC76062101	The Bride	simulation	-494.2	-11.5

Biochemical Testing with Recombinant p53

In order to verify that compounds specifically bind p53 and stabilize structure, we used the purified core domain of p53. Many researchers looking for reactivators test their compounds in cells because purified p53 is difficult to work with due to its instability and because working with small molecules often causes noisy results in biochemical assays. However, a biochemical assay is the optimal approach for determining if the mechanism of action involves direct binding.

We created recombinant p53 with two separate cancer-associated, destabilizing mutations. These were used for stability experiments and further testing was done in cell lines.

The initial biochemical screen was done with the core domain of p53 with the H179Y destabilizing mutation. Stability in the presence of compound was measured with the protease Trypsin. Trypsin cleaves at arginine and lysine residues, which are prevalent throughout p53. However, trypsin would only be able to cleave exposed residues so most residues would not be cleaved when p53 is in its folded conformation, so increased p53 cleavage is a measure of unfolding. At 37° C, even wild type p53 is moderately stable so trypsin would be able to cause cleavage to some percent of the p53 molecules. With the H179Y mutant, a noticeable amount of p53 is cleaved by 10 minutes at 37° C when visualized by western blot. In the presence of a stabilizing compound, we expect fewer p53 molecules will have exposed cleavage sights so more intact p53 would be present.

Before testing began, we hypothesized that we would measure the amount of intact p53 as a measure of stability, however it was difficult to calculate the levels of intact protein remaining. In order to differentiate degrees of stability at lower compound concentration, we found that measuring relative amounts of the cleavage fragments was a better indicator of whether a compound bound p53. We used Maltose Binding Protein (MBP) fused p53 in order to keep p53 in solution after partial cleavage has taken place in order to capture partial cleavage states.

For the pulse proteolysis assay with trypsin, we measured the band intensity of the largest p53 cleavage fragment relative to all cleavage fragments (**Figure 3a)**. In the presence of a compound that binds p53, there is relatively less of this fragment. So the ratio of the largest fragment compared to all fragments will be smaller with a stabilizing compound. MBP is highly

resistant to trypsin at the concentrations that was used so all fragments that have the N terminus are kept in solution and these are seen on the western blot for analysis.

The initial screen with H179Y was done with all compounds at 40uM, except one compound that was found to be a trypsin inhibitor. Five additional compounds (LM10-15) were tested as negative controls, which were similar to our compounds because they come from the same virtual screen compounds library, but are compounds predicted to bind Mcl-1 instead of p53. All 5 of these had similar ratios as DMSO showing that the assay does not readily have false positives. Of the 27 compounds predicted to bind p53 from virtual screen, a surprisingly large percent of them showed signs of stabilizing H179Y-p53. We found that 26% of them had a relative decrease in the largest fragment by a third. The three best hits (Han Solo, Ripley, and Spiderman) decreased the ratio nearly in half (**Figure 3b**).

Given the high percent of success, we looked to another destabilizing mutant to select hits that were stabilizing to multiple mutants. This is because we intend our stabilizer to function as a broad range reactivator of structural p53 mutants. We selected eleven compounds that had the strongest stabilizing effect for H179Y and tested them with the structural mutant R175H, which is the most prevalent of the cancer-associated, structural mutants in the TP53 database [6]. We had hypothesized that this assay would narrow our focus to the best hits, but all eleven turned out to be stabilizing on R175H as well (**Figure 3c**).

Of course, we also feared that the apparent stabilization of p53 was actually inhibition of trypsin so we measured inhibition by compounds with the BaAMC assay. This is a highly sensitive assay that can capture even slight inhibition. We also biased the assay towards inhibition by using 5x the amount of trypsin that is used in pulse proteolysis and 5x the compound concentration that we test. So even minor inhibition would show significant decrease

in BaAMC results. Of the 28 compounds analyzed, only one is found to be a strong trypsin inhibitor (Jeeves) so was not included in pulse proteolysis **(Figure 3d**). As an indicator of the test's sensitivity, we found that Dumbledore's BaAMC result was only 24% that of DMSO despite being in the range of the negative control in the pulse proteolysis assay with p53.

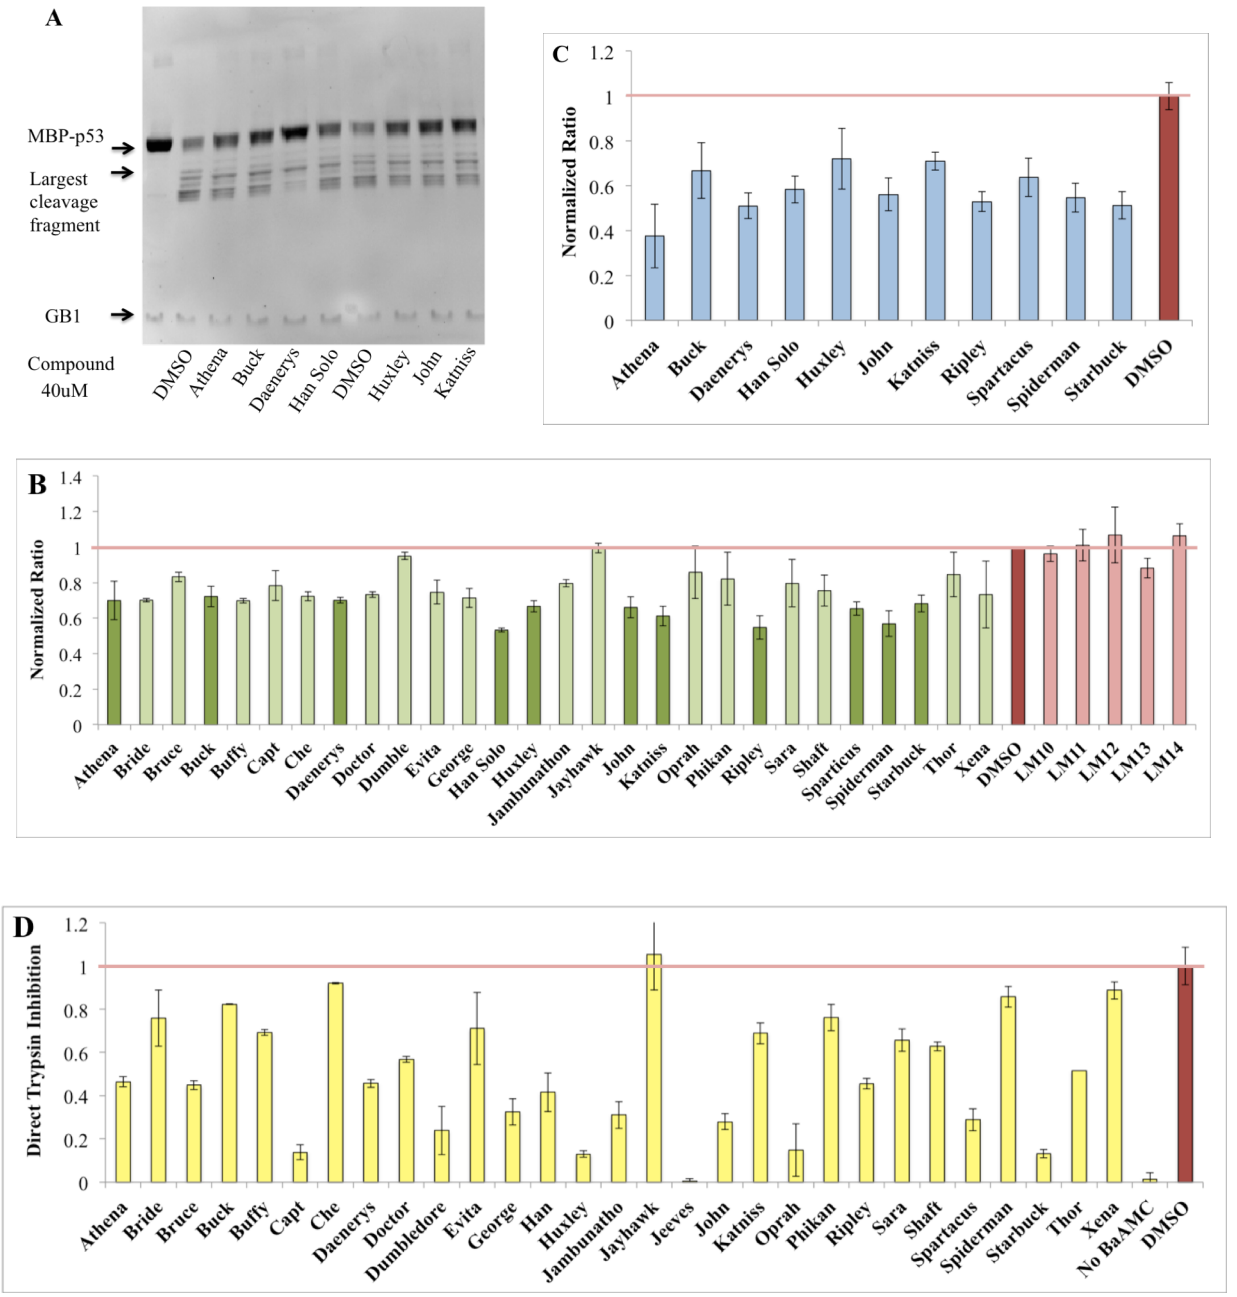

Figure 3: Trypsin cleavage used as a measure of protein stability. A) AntiHis$_6$ Western Blot of gels following pulse proteolysis shows rescue with seven compounds with MBP-H179Y-p53. First lane has no trypsin. Largest cleavage fragment is regularly less prevalent when rescue is present. So ratio of fragment density compared to other fragments is measured in B-C. B) Ratio is normalized to 1 for DMSO (red) for destabilizing mutant H179Y using compounds at 40uM. Negative control compounds shown in light red are similar to our compounds, but not predicted to bind p53. Compounds shown in green. Best compounds are shown in darker green, C) which are the compounds used to test in an additional destabilizing mutant R175H. D) The highly sensitive trypsin inhibition assay, BaAMC, was used with 5x the amount of compound and trypsin.

Cell experiments

In order to test whether molecules restored wild type function of mutant p53, we used multiple cancer cell lines that had wild type, mutant, and knockout p53 to measure downstream mRNA levels. We also used a luciferase assay with the response element targeted by p53. We visualized wild type p53 activation by fluorescent microscopy using antibodies whose epitope only occurs in either native or unfolded p53 in order to validate a stabilizing effect of compounds. This work focused on seven compounds that performed well in biochemical assays and PhiKan083, which is known to stabilize and rescue p53 mutant Y220C.

Two well-characterized p53 target genes are p21 (cdkn1a) and PUMA (BBC3) [95]. We used quantitative Real time PCR (QT-RT-PCR) to measure their mRNA levels in the presence of compounds. p21 is a cyclin dependent kinase inhibitor, which will cause the cell to go into G0/senescence when it reaches high levels. It is directly upregulated by the activation of wild type p53, which acts as a transcription factor [96]. PUMA is also directly upregulated by the transcription factor, p53. PUMA is a proapoptotic protein in the Bcl-2 protein family [97]. We measured p21 and PUMA mRNA levels in cell lines with two destabilizing mutations of p53 (Y220C, R175H) and one DNA binding mutation (R273H) in order to differentiate between stability specific rescue of p53 and nonspecific rescue. Most compounds showed significant upregulation of p21 and/or PUMA in comparison to DMSO (**Figure 4a**). The largest increase for the destabilizing mutant lines was in the presence of Thor.

The HCT116 cell line with knockout p53 was used for isogenic analysis with p53 null, wild type, or knockin of the DNA-binding mutant of p53 (R273H). Again we measured rescue with QT-RT-PCR of p21 and PUMA with seven compounds and PhiKan083. Since PhiKan083

binds the Y220C mutant, it is a negative control in this assay. It had some nonspecific increases in both genes and all isogenic variations despite its specificity for Y220C. The seven compounds had a similar response, showing that these compounds do not rescue in cells without destabilized p53 more than the negative control, except for the case of John in cells with R273H (**Figure 4b**).

We performed the luciferase assay with BxPC-3 cells transfected with either the PG13 plasmid, which contain wild type p53 response element upstream of luciferase gene or the MG15 plasmid, which contains a mutated variation. These transfected cells were treated with Athena, John, and Thor. All three compounds showed a significant increase in luciferase activity of PG13 while showing no effect in MG13 when compared with DMSO control. (**Figure 4c**). Again, Thor had the strongest response (**Figure 4d**).

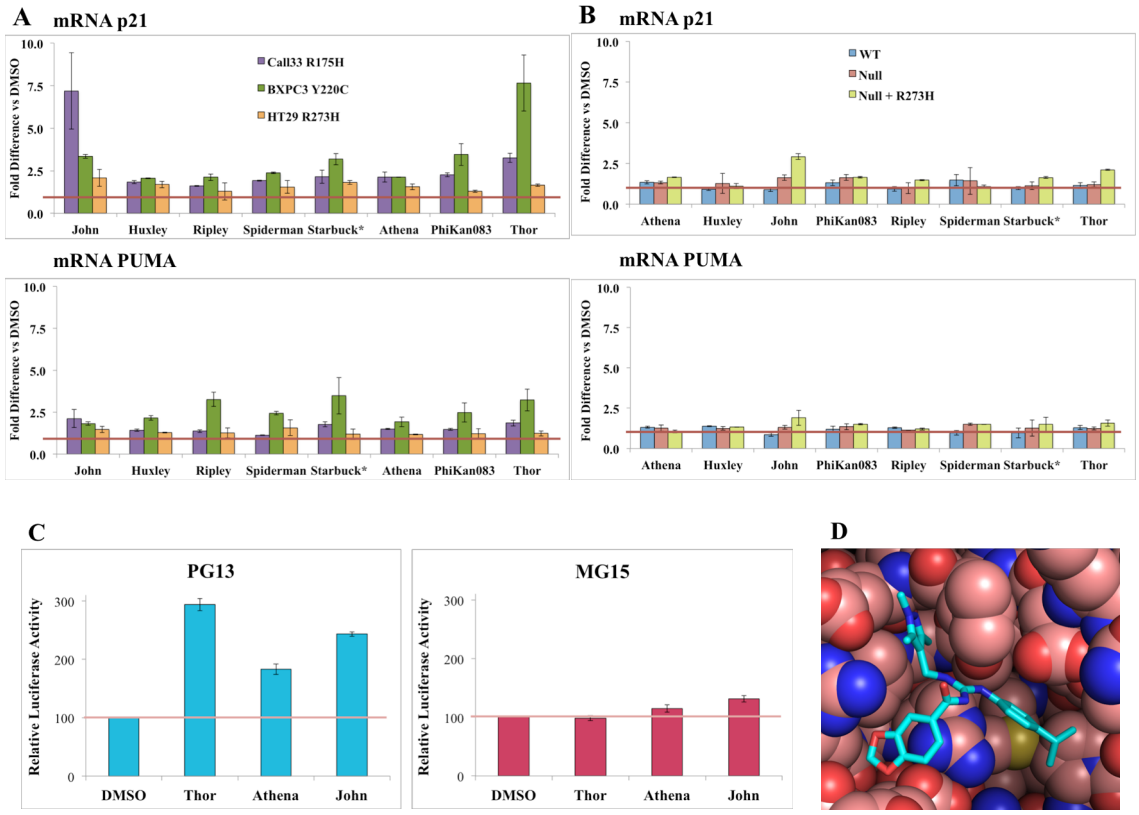

Figure 4: Compound rescue in cancer cell lines. Tops hits in biochemical tests used in cells.
A) Compounds tested at 40uM with two destabilizing mutant (R175H, Y220C) and one DNA binding mutant (R273H). B) Isogeneic p53 knockout cell line used compared to same line with WT or R273H knockin. C) Luciferase assay with p53 response element (PG13) or a mutated response element (MG15). D) Final docked pose of Thor (blue) with p53 (pink) in spheres.
*Starbuck tested at 10uM

Final evidence that four of the compounds improved stability of p53 mutant Y220C in cells came when using the antibody p1620 whose epitope was only present in folded p53. With DMSO, there is clearly minimal folded p53 present in cells, but all four compounds had a significant shift of equilibrium towards folded p53 (**Figure 5a**). In order to verify that there is a shift in equilibrium from unfolded to folded p53, we treated cells with Thor and added the antibody p240, which only binds an epitope of the unfolded conformation of p53. Thor clearly caused a large shift towards the folded population (**Figure 5b**).

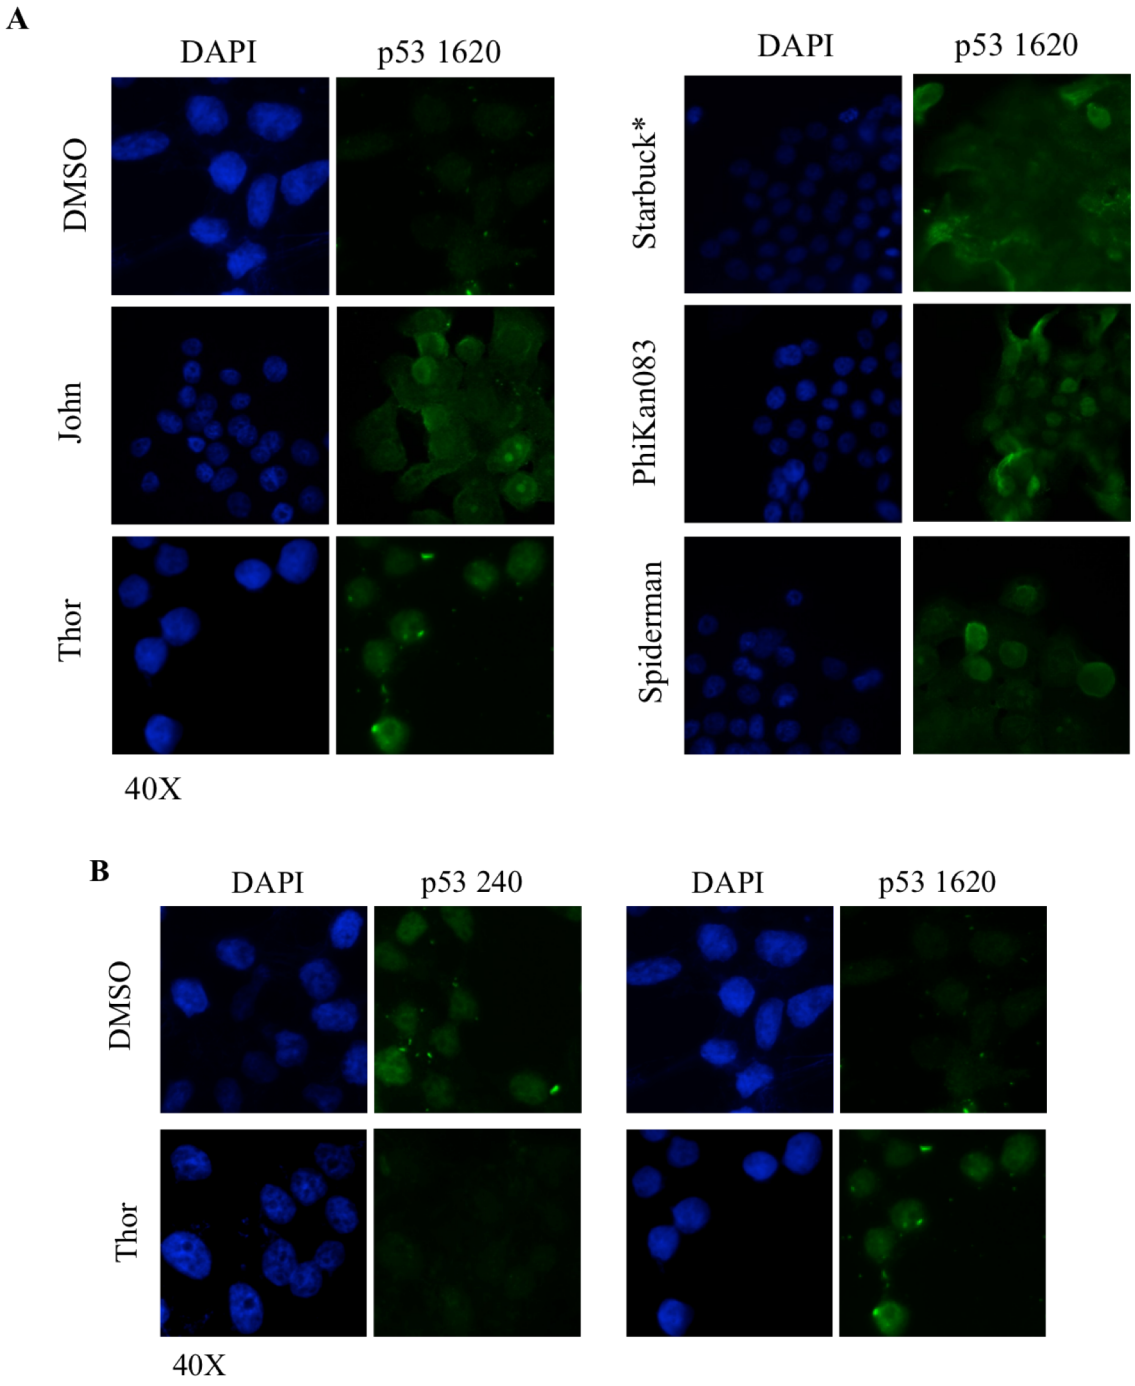

Figure 5: Immunofluorescence of p53 cells. Immunofluorescence in BxPC3 (Y220C p53) cells after treatment for 24 hours of compound at 40uM. **A)** pAb 1620 in green binds folded p53. Cells treated with DMSO, compound, or PhiKan083, a known stabilizer of p53. B) Antibodies pAb 1620 and pAb 240 bind correctly folded and unfolded p53 respectively with cells treated with DMSO or Thor.
*Starbuck tested at 10uM

2.4 Discussion

Compounds that reactivate destabilized p53 mutants have the potential to serve as unique chemical tools in the treatment of cancer patients. In order to unlock that potential, however, such molecules should act in a highly selective manner, and via a clearly established mechanism of action. Here, we have identified a novel druggable site on the surface of p53, and demonstrate the potential utility of this site through compounds designed to bind here. Through biochemical assays we have demonstrated that these compounds interact with p53, and in subsequent cell-based assays we have begun to probe their biological effects.

Exploring druggable sites on p53

Two previous studies have utilized structure-based virtual screening in search of compounds that reactivate p53 mutants, but these have yielded low affinity hits (Stictic Acid and PhiKan083). In retrospect, the sites targeted in these prior studies may explain this. Stictic Acid was discovered using the surface pocket found during molecular dynamics centered at a cysteine suspected to be a potential PRIMA-1 adduct location. In the case of PhiKan083, the pocket was discovered on the crystallographically observed cleft formed by the Y220C mutation [19, 98].

In contrast to earlier efforts, here we began by exhaustively searching the surface of p53 for its optimally druggable site. To do so, we considered not only the conformation observed in crystal structures of p53, but also low-energy excited states generated through biased simulations. Interestingly, analysis of (unbiased) molecular dynamics simulations in search of druggable pockets on the surface of p53 did not reveal this site; this may have simply resulted from the different energy functions used in these studies, or alternatively may reflect the

challenges in sampling these pocket-containing conformations without explicitly biasing the simulations to do so [98, 99].

It is encouraging to note that the site is distant from the 5 most common cancer-associated, destabilizing mutation sites in p53 (R175, G245, R249, R282, Y220), which make up 15% of somatic mutations in the TP53 database [6]. Because this druggable site is not in proximity to these frequently-mutated residues, it is likely that this pocket will not be distorted by these mutations – further implying that compounds acting at this site are expected to restore activity to each of them.

We further note that the pocket we have described here is centered around Ser215; interestingly, increased phosphorylation at this position by mitotic kinase Aurora A and by PAK4 (p21 Activated Kinase 4) have also been shown to deactivate p53 in certain cancers [100-102]. Based on sterics, we anticipate that compounds binding at this site may also prevent phosphorylation at this site: if so, the same compounds may additionally prove effective for rescuing p53 activity in cancers that retain wild-type p53 but are instead driven by aberrant upregulation of oncogenes Aurora A and/or PAK4. In future, we look forward to testing this hypothesis by evaluating the effect of the tool compounds presented here on additional carefully selected cancer cell lines.

Inhibition of phosphorylation at Ser215 may account for the apparent non-specific rescue of John in the DNA-binding mutant cell line (**Figure 4b**). Treatment with John led to rescue of p53 in the R273H knockin compared to DMSO. Even though this mutation has decreased affinity to DNA, it is likely that it can still bind DNA if its concentration was allowed to increase. So inhibition of phosphorylation could increase the levels of functional p53. Further testing is required to verify this hypothesis.

Therapeutic opportunities for p53 reactivators

Finding a reactivator that rescues a wide range of structural mutants of p53 has the potential to help a large number of cancer patients. Since about half of all cancers have a mutation in p53 and 30-40% of those are structural mutants, a broad spectrum stabilizer could rescue the p53 pathway in up to 20% of cancers [6, 10]. While that is an optimistic figure, the reactivator approach for cancer treatment is promising. Reactivators have been found to rescue p53 activity in cancer cell lines and in animal models, and one of these compounds, PRIMA-1, has recently entered clinical trials [89]. In fact, PRIMA-1 is currently being tested in three ongoing Phase 1b/2 clinical trials [103], to evaluate its promise for targeting oesophageal carcinomas, platinum sensitive recurrent high-grade serous ovarian cancers, and myeloid neoplasms [6].

Despite the promising outlook of existing compounds, there remains a strong need to find additional p53 reactivators for structural mutants. PRIMA-1's progress is exciting, but clinical trials have just begun and so their outcome is not yet certain. Further, structure activity relationship analysis around PRIMA-1 confirms that it almost certainly involves covalent attachment to some target, but it remains controversial whether the relevant target is indeed p53 [104]. Beyond this, PRIMA-1 also causes an increase in reactive oxygen species (ROS) due to its interactions with thioredoxin reductase 1 and glutathione [105-107], implying that unanticipated side effects may emerge as studies continue.

While it is impossible to rule out adverse effects before a drug candidate is thoroughly tested through clinical trials, side effects are far less likely for compounds that unambiguously engage the intended target. Also, side effects are less likely if a candidate lacks chemically

reactive functional groups that form covalent adducts with proteins. While a number of p53 reactivators have been described in the literature, many have been identified through phenotypic (cell-based) screens, and did not turn out to interact with p53 directly. Outside of the compounds presented here only a limited subset have been demonstrated through biochemical assays to bind to p53: these include PRIMA-1, MIRA-1, PhiKan083, PK7088, and stictic acid. Among these though, the former two rely on reactive functional groups and the latter three have very low binding affinity [92, 98, 108, 109].

By using structure-based design, and by following up with biochemical assays before moving into cells, we have ensured that the hit compounds from our study indeed interact with p53 directly. Furthermore, this approach has allowed us to rescue many different p53 mutants by targeting a specific site on the protein surface: this would not have been possible without structure-based design. Using only a modest library, we have already found several hit compounds: this suggests that the site is highly amenable to reactivator binding, and should be the focus of a larger screen. Rescue of p53 activity has long been perceived to be an effective avenue for development of new cancer therapeutics, and now this druggable site may lead to development of potent compounds to realize this promise [84].

Even beyond therapeutic applications, we also foresee that a potent stabilizer with favorable *in vivo* properties could also play a role in cancer prevention. As many as 10% of all cancers are thought to arise from unfavorable genetic variations that mark certain people with a predisposition to cancer [6].

These members of the population are then more susceptible to additional somatic mutations, such as destabilizing mutations in p53. If a potent p53 reactivator can be identified that can be administered to healthy patients without side effects, it could be provided to those

people most susceptible to future cancers. This preventative approach could trigger apoptosis in precancerous somatic cells that have accumulated DNA damage, yet lack the normal p53 response due to mutations in p53. Such an approach would be particularly effective in patients with Li-Fraumeni syndrome, in whom p53 is already systemically deactivated.

2.5 Methods

The computational methods described here are implemented in the Rosetta software suite [39]. Rosetta is freely available for academic use.

Selecting druggable pocket

The structure was downloaded for wild type p53 DNA Binding Domain without DNA (2OCJ.pdb) from The Protein Database. The volume of surface pockets was calculated using the method described in 2013 *Johnson et al*. Briefly, a grid is created around a residue on the protein surface. Grid points are considered protein, surface, or solvent. Pocket grid points are when a line of the grid goes from surface-solvent-surface so that the volume of grid squares encompassed by pocket grid points can be calculated. Residues that had overlapping pockets were grouped into one pocket.

Each pocket was analyzed separately. Rosetta relax was run with an energy term that biases towards larger pockets, also described in *Johnson et al*. [20]. The proportionality constant was modified to be -0.25 Rosetta energy units per $Å^3$. We ran 3-6,000 simulations per pocket depending on shape and 1,000 simulations without the bias were run for reference energy. Biased structures with energy higher than reference were excluded.

Creating Rigid Body Pockets for Docking

We generated structures when running simulations with Rosetta Relax for pocket selection. These were used to create a pocket shell for each structure using the grid defined in Selecting Druggable Pockets. Each shell was aligned in order to create a pairwise distance score. The score = the number of overlapping grid points divided by the number of grid points in the smaller pocket. Scores were put into a distance matrix, which was used with hierarchical clustering. We found two clusters of general pocket shapes present within the simulations and selected the clustroid from each as representative structures for docking. These two clustroid structures and the crystal structure were used to create three pockets for docking, as described in *Khar et al* [110].

Virtual Screen

We used the Rosetta tool, DARC, using the method defined in *Gowthaman et al*. Briefly; we created a file that defined the topology of the pocket that DARC uses for scoring. A small molecule library was created using the ZINC database to download smiles that were used by OMEGA to create pdb and params files with up to 200 conformers as described in *Khar et al* [110]. These were used by DARC, which generated a final docked pose of the most favorable conformer of each small molecule in its optimal conformation in each pocket in terms of DARC scoring. The final poses were minimized and scored with Rosetta. For each pocket, Z scores were generated for Rosetta total energy, Rosetta interaction energy, number of hydrogen bonds, number of unsatisfied polar bonds, and solvent exposed surface area. We used the sum of these Z scores to rank the final poses, examined the top few hundred structures manually in pymol, and selected 28 commercially available compounds for testing.

Subcloning, Site-Directed Mutation, and Protein Expression

Wild type p53's DNA Binding Domain (residues 94-312) was subcloned into pTBSG-MAL (which was provided courtesy od Dr. Philip Gao) by Ligation Independent Cloning [111] as described in Lea et al [112]. This vector forms a construct with an N-terminal His tag followed by Maltose Binding Protein, a TEV recognition site, and C-terminal p53. It was transformed into E. coli NEB 5-alpha cells and plated on LB agar with 50 mg/mL ampicillin. A single colony was used to inoculate 10ml of LB Broth with 50 mg/ml ampicillin and grown overnight at 37°C. Plasmids were purified with Quiagen miniprep kit. Upon sequence verification, plasmid was transformed into E. coli Rosetta2 (2DE3) pLysS cells and plated on LB agar with 50 mg/mL ampicillin. A single colony was used to inoculate 10ml of LB Broth with 50 mg/ml ampicillin and grown overnight at 37°C. In the morning, 0.5ml of growth was added to 0.5ml glycerol and moved to -80°C storage. Site Directed Mutation was done in order to create structural mutant constructs including R175H and H179Y.

For expression, an LB agar plate with 50 mg/mL ampicillin was inoculated with glycerol stock. A single colony was used to inoculate 100ml LB Broth with 50 mg/mL ampicillin for 37°C overnight growth. This was used to inoculate four liters of LB Broth with 50 mg/mL ampicillin and allowed to grown until 600nm OD reaches ~0.6 and induced with 1mM IPTG. Temperature was shifted to 21°C and it was allowed to grow overnight.

Cells were centrifuged at 5,000g for 10 minutes, then lysed in 50mM Tris pH 7.2, 50mM Imidazole, 1M NaCl, 1M Urea, and 1mM DTT. This was centrifuged for 1 hour at 27,000g. Lysate was run on Ni-NTA and eluted with 50mM Tris pH 7.2, 500mM Imidazole, 0.3M NaCl, and 1mM DTT. This was dialyzed into 50mM Tris pH 7.2, 150mM NaCl, and 10mM DTT.

Finally, the protein was concentrated with an Amicon-15 Filter with 10,000MW cutoff until concentration is over 5mg/ml. Glycerol is added to bring it up to 50%, then protein is aliquoted and stored at -20°C.

Pulse Proteolysis

Stock solution of trypsin was prepared fresh daily at 1mg/ml. Each sample had a final volume of 45ul. Buffer was 50mM Tris pH 7.2, 150mM NaCl, 10mM DTT. DMSO was kept at 4%. Compounds were dissolved in DMSO and used at 40uM. p53 concentration was 5mM. Final concentration of trypsin was 3ug/ml. Samples were mixed on ice in 1.5ml tubes. Trypsin was added last. Samples were transferred to 37°C water bath for 10 minutes, then transferred to ice. 15ul of 4x SDS loading buffer with BME was added to stop reaction and to denature for SDS PAGE, which was run immediately. Gel was washed with DI water, then protein was transferred to nitrocellulose membrane for 7 minutes. Nitrocellulose was washed three timed for 10 minutes with TBS. Then it was blocked for one hour using 5% alkali soluble casein. The TBS wash was repeated, then IR800 Anti-His antibody was diluted 50,000x into 5% alkali soluble casein and sat overnight on rocker. TBS wash was repeated, then it was visualized with infrared 800 light.

Images were analyzed with ImageJ. Protein bands were selected in order to calculate the intensity. A diagnostic band (largest cleavage fragment shown in figure 3a) is shown to be weaker when stabilizing compounds are present, which is found to be a 60kDa fragment composed of MBP and the C-terminal 16kDa portion of p53. The total intensity of all p53 bands is summed in order to calculate the ratio of diagnostic band intensity/total fragment intensity of cleavage bands. This ratio is normalized to DMSO samples so values less than 1 indicate rescue.

Trypsin inhibition assay was done with N-Alpha-Benzoyl-L-Arginine 7-Amido-4-Methylcoumatin HCl (BaAMC), which is a fluorogenic substrate for trypsin. Excitation was 340nm with 40nm slit, emission measured at 460nm with 40nm slit. Trypsin was 15ug/ml, which is 5x higher than what is used in pulse proteolysis assay. Compounds were tested at 200uM. The slope of emission is measured during first 20 minutes for 2 replicates for all compounds, but 5 replicates of positive and negative controls.

Cell Assays

Human cell lines used in this study are BxPC3 $p53^{Y220C}$(pancreatic adenocarcinoma provided by D. R. Welch at University of Kansas Medical Center, USA), HCT116 $p53^{WT}$, HCT116 p^{53null}, HCT116 $p^{53null+273}$(R273H exogenous, isogenic knockout colon carcinoma provided by B. Vogelstein at Johns Hopkins Medicine, USA.), H2087 $p53^{V157F}$ (lung adenocarcinoma provided by T. Komiya at University of Kansas Medical Center, USA), HT29 $p53^{R273H}$ (colorectal adenocarcinoma provided by D. A. Dixon and S. Anant, respectively, at University of Kansas Medical Center, USA), MiaPaCa $p53^{R248W}$ (pancreatic carcinoma provided by D. R. Welch at University of Kansas Medical Center, USA) and CAL33 $p53^{R175H}$ (tongue squamous cell carcinoma provided by S. Thomas at University of Kansas Medical Center, USA). Cell lines were maintained in Dulbecco's modified Eagle's medium (DMEM) or Roswell Park Memorial Institute (RPMI) 1640 with 10% fetal bovine serum (FBS) and 1% penicillin–streptomycin. The HCT116 subcell lines, p53null+R175H and p53null+R273H, were generated by infecting HCT116 p53null cells with retroviral vectors encoding p53R175H and p53R273H cDNAs, respectively. All cell lines were authenticated by autosomal STR profiles provided by the University of Arizona Genetics Core. All cell lines were tested negative for mycoplasma.

None of the cell lines used was found in the database of commonly misidentified cell lines that are maintained by ICLAC and NCBI Biosample.

RNA was isolated using the RNA-Quick MiniPrep (Zymo Research). Total RNA (1 μg) was reversed transcribed to cDNA using M-MLV reverse transcriptase (Amresco), according to the manufacturer's instructions, and TaqMan assays were performed with ViiA7 (Life Technologies). TaqMan assay primers and probes were purchased from Life Technologies using the following assay numbers: human p53, Hs00153349_m1; mouse p21 Mm00432448_m1; Bax Mm00432050_m1, Puma Mm00519268_m1; Gapdh, Mm99999915_g1. Taqman assay for human GAPDH was purchased from Integrated DNA Technologies (Hs.PT.58.40035104). The mRNA levels were normalized to those of GAPDH.

For Luciferase assay as BxPC3 $p53^{Y220C}$ were transfected with PG13 plasmid with response element for p53 and MG15 plasmid that has a mutated variation, along with control plasmid for transfection efficiency. After 36 hour of transfection media was changed and different drug treatment was done was 24 hour and then after Luciferase activity was measured using The Dual-Glo® Luciferase Assay System (Promega, Madison, WI) with BioTek Synergy H4 multifunctional plate reader (BioTek).

Cells grown onto poly-D-lysine/laminin-coated glass cover-slips (BD Biosciences) were fixed with 4% paraformaldehyde for 20 min and then permeabilized with 0.3% Triton X-100 for 5 min, followed by blocking in 1% BSA in PBS-T for 1 h and incubation with PAb1620 p53 antibody (1:25 dilution) overnight at 4 °C. Goat anti-mouse IgG was used as a secondary antibody. Samples were mounted in the ProLong Gold Antifade Reagent with DAPI (Invitrogen), followed by analysis with a Nikon epifluorescence microscope.

2.6 Acknowledgements

We thank Wendy Lea for assistance and useful discussions. We are grateful to OpenEye Scientific Software (Santa Fe, NM) for providing an academic license for the use of OMEGA. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) allocation MCB130049, which is supported by National Science Foundation grant number ACI-1053575. This work was supported by grants from the National Institute of General Medical Sciences of the National Institutes of Health (R01GM099959 and R01GM112736).

Conclusions

Developing a cure for cancer is one of the greatest scientific challenges of our time. Each different type of cancer has its own set of disease variations, which is further complicated by variations within each individual patient. This makes cancer treatment an indescribably complex problem. Researchers have recently developed the hypothesis that clinical outcomes would improve by utilizing Precision Oncology. This is a personalized medicine approach that tailors a cancer remedy based on an individual's genetic profile [113]. This approach would require genotyping an individual's malignant cells and searching that genome for an established set of genetic mutations associated with cancer, then prescribing a set of anticancer drugs that target those specific mutations.

One set of mutations that may soon become a component of this approach is the destabilizing mutations of p53 that are known to be strongly associated with cancer. These are the targets of my thesis research. The long term goal of this project is to find a drug that rescues the structure and function of destabilizing mutations of p53 in cancer patients. The immediate ambition is that the methods developed during this project are an improvement over previous methods of finding a p53 stabilizing molecule, so will be a better example for future work. Additionally, we suspect that the stabilizing compounds found here are potential drug leads.

The first stage of my thesis project was to optimize the docking tool DARC by enabling GPU processing of the scoring function. Enabling GPU computation resulted in an 180-fold increase in the time it takes to score individual poses during docking. This allowed us to perform the entire docking process 27 times faster than when calculating on CPUs alone. This speedup in important because the amount of time it takes to run a virtual screen is a limiting factor. A faster screening process allows us to sample a significantly larger library of molecules during virtual

screen, which drastically increases the probability that we will be able to find molecules that bind to their targets.

Once our virtual screening method was optimized, my goal was to use that method on p53. I performed druggability analysis on the core domain and found a novel pocket that I used for my screen. I found it to be druggable by doing Rosetta Relax simulations that sampled conformational variations in the protein in order to find low-energy structures that had larger pockets. Pockets opened easily at the novel pocket site, but not on the rest of the protein surface (excluding the DNA binding region). This pocket appeared to be an ideal screening spot because it is distant from the DNA binding site, the zinc binding site, and the most common cancer-associated mutation sites. So it should not interfere with normal function and would be present (at last transiently) in multiple cancer-associated, destabilizing mutants of p53.

After selecting a druggable pocket on p53, I performed a virtual screen using DARC on the novel pocket using a modest-sized library of compounds. DARC takes a rigid-body pocket and samples up to 200 conformers of a small molecule in the pocket in thousands of orientations. This is sufficient to find the optimal scoring orientation of the molecule in the pocket, but it doesn't account for variation in pocket shape. We use a static pocket in order to decrease the degrees of freedom, thus decreasing the number of computations. This makes the process quicker, which (as I just mentioned) is a rate-limiting factor when doing virtual screen. We made up for the lack of pocket degrees of freedom by running the screen multiple times using several energetically-favorable pocket shapes.

During the Rosetta simulations for druggability analysis, a variety of pocket shapes were found at the novel site. Since these shapes represent the dynamic changes that the protein is likely to sample in vivo, I analyzed the pocket shapes in order to capture several representative

pockets that I could use as rigid-body structures during docking in the virtual screen. I clustered all structures created in the simulation based on their pocket shapes. I found two general shapes were present among the majority of the structures and selected a representative structure from each cluster to use in the virtual screen, in addition to the crystal structure. I found that all 3 structures had significantly varied pocket shape that would make them unlikely to bind the same molecules.

I performed a virtual screen using DARC on the novel pocket on p53 using a library of 64,000 compounds and selected 28 compounds that ranked well. My next goal was to do biochemical testing to discover which of these compounds bind and stabilize p53. I used the purified core domain of p53 that had either the H179Y or R175H mutations, which are both cancer-associated, destabilizing mutations. Initial testing was done with H179Y using a pulse proteolysis assay. Surprisingly, I found that the majority of the compounds appeared to stabilize p53. I selected the 11 most stabilizing of the 28 compounds and repeated the assay on R175H, which is the most common cancer-associated, destabilizing mutation of p53. I found that all 11 of the compounds stabilized R175H.

We further tested of the compounds in cell based assays. We used mRNA levels of p21 and PUMA as indicators of p53 rescue because they are directly upregulated by wild type p53. Using two cell lines with destabilizing mutations of p53 and one line with a DNA binding mutation, we found that 7 compounds showed p53 rescue in the cell lines with destabilizing mutations. Most notably, several compounds had stronger rescue than PhiKan083 in the Y220C cell line. PhiKan083 is a known stabilizer of this p53 mutant. The best compound in this assay was the molecule nicknamed Thor, which upregulated p21 and PUMA by 7.6 and 3.2-fold respectively. Thor and 2 other compounds were also found to rescue p53 with the luciferase

assay that used a p53 response element. Again Thor was the best, and had almost a 3-fold increase in the luciferase response. Thor also showed signs of rescue in the final cell test, which used pAb 1620, a immunofluorescent antibody that binds p53. A cell line with Y220C mutant of p53 had a highly noticeable increase in the amount of pAb1620 in the presence of Thor and 3 other compounds.

While the evidence presented here strongly suggests that we have identified multiple p53 stabilizers, there is still a lot of work to be done before we can be confident that these are good hits. First, we will to use NMR and/or XRay crystallography to verify that these compounds directly bind p53 and to confirm their binding modes. Next, we will establish binding affinity with additional biochemical assays such as surface plasmon resonance or differential scanning fluorimetry. This will allow us to determine which of these hits bind most potently. Lastly, we will perform medicinal chemistry optimization in order to improve affinity and obtain the physicochemical properties necessary for our hits to be drug-like compounds.

During my thesis work, I optimized the docking tool DARC, discovered a novel, druggable pocket on p53, performed a virtual screen with that pocket, and demonstrated that many of the predicted p53 binders not only bind, but rescue the structure and function of multiple cancer-associated, destabilizing mutants of p53. These hits may turn out to be strong drug leads, and just as importantly, the method used to find these hits appears to be superior to previous methods and should be repeated with a larger virtual screen. This brings us a step closer to Precision Oncology.

References

1. Zhu T, Lee H, Lei H, Jones C, Patel K, Johnson ME, Hevener KE: **Fragment-Based Drug Discovery Using a Multidomain, Parallel MD-MM/PBSA Screening Protocol**. *J Chem Inf Model* 2013, **53**(3):560-72.
2. Yang H, Zhou Q, Li B, Wang Y, Luan Z, Qian D, Li H: **GPU acceleration of Dock6's Amber scoring computation**. *Adv Exp Med Biol* 2010, **680**:497-511.
3. Ritchie DW, Venkatraman V: **Ultra-fast FFT protein docking on graphics processors**. *Bioinformatics* 2010, **26**(19):2398-2405.
4. http://www.cancer.gov/about-cancer/understanding/statistics.
5. Morris RJ, Najmanovich RJ, Kahraman A, Thornton JM: **Real spherical harmonic expansion coefficients as 3D shape descriptors for protein binding pocket and ligand comparisons**. *Bioinformatics* 2005, **21**(10):2347-2355.
6. http://www.cancer.gov/about-cancer/causes-prevention/genetics/genetic-testing-fact-sheet
7. Kubbutat MH, Jones SN, Vousden KH: **Regulation of p53 stability by Mdm2**. *Nature* 1997, **387**(6630):299-303.
8. Korb E, Finkbeiner S: **Arc in synaptic plasticity: from gene to behavior**. *Trends Neurosci* 2011, **34**(11):591-598.
9. Moitessier N, Englebienne P, Lee D, Lawandi J, Corbeil CR: **Towards the development of universal, fast and highly accurate docking/scoring methods: a long way to go**. *Br J Pharmacol* 2008, **153 Suppl 1**:S7-26.
10. Joerger AC, Fersht AR: **The tumor suppressor p53: from structures to drug discovery**. *Cold Spring Harb Perspect Biol* 2010, **2**(6):a000919.
11. Kandoth C, McLellan MD, Vandin F, Ye K, Niu B, Lu C, Xie M, Zhang Q, McMichael JF, Wyczalkowski MA *et al*: **Mutational landscape and significance across 12 major cancer types**. *Nature* 2013, **502**(7471):333-339.
12. Duffy MJ, Synnott NC, Crown J: **Mutant p53 as a target for cancer treatment**. *Eur J Cancer* 2017, **83**:258-265.
13. Robles AI, Harris CC: **Clinical outcomes and correlates of TP53 mutations and cancer**. *Cold Spring Harb Perspect Biol* 2010, **2**(3):a001016.
14. Heinzerling L, Klein R, Rarey M: **Fast force field-based optimization of protein-ligand complexes with graphics processor**. *J Comput Chem* 2012, **33**(32):2554-2565.
15. Caron de Fromentel C, Gruel N, Venot C, Debussche L, Conseiller E, Dureuil C, Teillaud JL, Tocque B, Bracco L: **Restoration of transcriptional activity of p53 mutants in human tumour cells by intracellular expression of anti-p53 single chain Fv fragments**. *Oncogene* 1999, **18**(2):551-557.
16. Abarzua P, LoSardo JE, Gubler ML, Spathis R, Lu YA, Felix A, Neri A: **Restoration of the transcription activation function to mutant p53 in human cancer cells**. *Oncogene* 1996, **13**(11):2477-2482.
17. Abarzua P, LoSardo JE, Gubler ML, Neri A: **Microinjection of monoclonal antibody PAb421 into human SW480 colorectal carcinoma cells restores the transcription activation function to mutant p53**. *Cancer Res* 1995, **55**(16):3490-3494.
18. Bullock AN, Fersht AR: **Rescuing the function of mutant p53**. *Nat Rev Cancer* 2001, **1**(1):68-76.

19. Boeckler FM, Joerger AC, Jaggi G, Rutherford TJ, Veprintsev DB, Fersht AR: **Targeted rescue of a destabilized mutant of p53 by an in silico screened drug**. *Proc Natl Acad Sci U S A* 2008, **105**(30):10360-10365.
20. Johnson DK, Karanicolas J: **Druggable protein interaction sites are more predisposed to surface pocket formation than the rest of the protein surface**. *PLoS Comput Biol* 2013, **9**(3):e1002951.
21. Lionta E, Spyrou G, Vassilatis DK, Cournia Z: **Structure-based virtual screening for drug discovery: principles, applications and recent advances**. *Cur Top Med Chem* 2014, **14**(16):1923-1938.
22. Gowthaman R, Miller SA, Rogers S, Khowsathit J, Lan L, Bai N, Johnson DK, Liu C, Xu L, Anbanandam A *et al*: **DARC: Mapping Surface Topography by Ray-Casting for Effective Virtual Screening at Protein Interaction Sites**. *Journal Med Chem* 2016, **59**(9):4152-4170.
23. Ghose AK, Viswanadhan VN, Wendoloski JJ: **A knowledge-based approach in designing combinatorial or medicinal chemistry libraries for drug discovery. 1. A qualitative and quantitative characterization of known drug databases**. *Journal Comb Chem* 1999, **1**(1):55-68.
24. Cheng T, Li Q, Zhou Z, Wang Y, Bryant SH: **Structure-based virtual screening for drug discovery: a problem-centric review**. *The AAPS journal* 2012, **14**(1):133-141.
25. Perez-Sanchez H, Wenzel W: **Optimization methods for virtual screening on novel computational architectures**. *Curr Comput Aided Drug Des* 2011, **7**(1):44-52.
26. Zhou H, Skolnick J: **FINDSITE(comb): a threading/structure-based, proteomic-scale virtual ligand screening approach**. *J Chem Inf Model* 2013, **53**(1):230-240.
27. Lee HS, Zhang Y: **BSP-SLIM: a blind low-resolution ligand-protein docking approach using predicted protein structures**. *Proteins* 2012, **80**(1):93-110.
28. Vorobjev YN: **Blind docking method combining search of low-resolution binding sites with ligand pose refinement by molecular dynamics-based global optimization**. *J Comput Chem* 2010, **31**(5):1080-1092.
29. Oost TK, Sun C, Armstrong RC, Al-Assaad AS, Betz SF, Deckwerth TL, Ding H, Elmore SW, Meadows RP, Olejniczak ET *et al*: **Discovery of potent antagonists of the antiapoptotic protein XIAP for the treatment of cancer**. *Journal Med Chem* 2004, **47**(18):4417-4426.
30. McGann M: **FRED Pose Prediction and Virtual Screening Accuracy**. *J Chem Inf Model* 2011, **51**(3):578-596.
31. Trott O, Olson AJ: **AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading**. *J Comput Chem* 2010, **31**(2):455-461.
32. Lang PT, Brozell SR, Mukherjee S, Pettersen EF, Meng EC, Thomas V, Rizzo RC, Case DA, James TL, Kuntz ID: **DOCK 6: combining techniques to model RNA-small molecule complexes**. *Rna* 2009, **15**(6):1219-1230.
33. Friesner RA, Banks JL, Murphy RB, Halgren TA, Klicic JJ, Mainz DT, Repasky MP, Knoll EH, Shelley M, Perry JK *et al*: **Glide: a new approach for rapid, accurate docking and scoring. 1. Method and assessment of docking accuracy**. *Journal Med Chem* 2004, **47**(7):1739-1749.
34. Neves MA, Totrov M, Abagyan R: **Docking and scoring with ICM: the benchmarking results and strategies for improvement**. *J Comput Aided Mol Des* 2012, **26**(6):675-686.
35. Zhou Z, Felts AK, Friesner RA, Levy RM: **Comparative performance of several flexible docking programs and scoring functions: enrichment studies for a diverse set of pharmaceutically relevant targets**. *J Chem Inf Model* 2007, **47**(4):1599-1608.
36. Verdonk ML, Cole JC, Hartshorn MJ, Murray CW, Taylor RD: **Improved protein-ligand docking using GOLD**. *Proteins* 2003, **52**(4):609-623.

37. Sauton N, Lagorce D, Villoutreix BO, Miteva MA: **MS-DOCK: accurate multiple conformation generator and rigid docking protocol for multi-step virtual ligand screening**. *BMC bioinformatics* 2008, **9**:184.
38. Lorber DM, Shoichet BK: **Flexible ligand docking using conformational ensembles**. *Protein Sci* 1998, **7**(4):938-950.
39. Leaver-Fay A, Tyka M, Lewis SM, Lange OF, Thompson J, Jacak R, Kaufman K, Renfrew PD, Smith CA, Sheffler W *et al*: **ROSETTA3: an object-oriented software suite for the simulation and design of macromolecules**. *Methods Enzymol* 2011, **487**:545-574.
40. Nicholls A: **What do we know and when do we know it?** *J Comput Aided Mol Des* 2008, **22**(3-4):239-255.
41. **OpenCL,** http://www.khronos.org/opencl.
42. Daga M, Feng WC: **Multi-dimensional characterization of electrostatic surface potential computation on graphics processors**. *BMC bioinformatics* 2012, **13 Suppl 5**:S4.
43. Hofinger S, Acocella A, Pop SC, Narumi T, Yasuoka K, Beu T, Zerbetto F: **GPU-accelerated computation of electron transfer**. *J Comput Chem* 2012, **33**(29):2351-2356.
44. Komarov I, D'Souza RM, Tapia JJ: **Accelerating the Gillespie tau-Leaping Method using graphics processing units**. *PloS one* 2012, **7**(6):e37370.
45. Friedrichs MS, Eastman P, Vaidyanathan V, Houston M, Legrand S, Beberg AL, Ensign DL, Bruns CM, Pande VS: **Accelerating molecular dynamic simulation on graphics processing units**. *J Comput Chem* 2009, **30**(6):864-872.
46. Stone JE, Phillips JC, Freddolino PL, Hardy DJ, Trabuco LG, Schulten K: **Accelerating molecular modeling applications with graphics processors**. *J Comput Chem* 2007, **28**(16):2618-2640.
47. Sanchez-Linares I, Perez-Sanchez H, Cecilia JM, Garcia JM: **High-Throughput parallel blind Virtual Screening using BINDSURF**. *BMC bioinformatics* 2012, **13 Suppl 14**:S13.
48. Fuller JC, Burgoyne NJ, Jackson RM: **Predicting druggable binding sites at the protein-protein interface**. *Drug Discov Today* 2009, **14**(3-4):155-161.
49. Wells JA, McClendon CL: **Reaching for high-hanging fruit in drug discovery at protein-protein interfaces**. *Nature* 2007, **450**(7172):1001-1009.
50. Meiler J, Baker D: **ROSETTALIGAND: protein-small molecule docking with full side-chain flexibility**. *Proteins* 2006, **65**(3):538-548.
51. Davis IW, Baker D: **RosettaLigand docking with full ligand and receptor flexibility**. *J Mol Biol* 2009, **385**(2):381-392.
52. Kaufmann KW, Meiler J: **Using RosettaLigand for small molecule docking into comparative models**. *PloS one* 2012, **7**(12):e50769.
53. Huang B, Schroeder M: **LIGSITEcsc: predicting ligand binding sites using the Connolly surface and degree of conservation**. *BMC structural biology* 2006, **6**:19.
54. Call ST, Zubarev DY, Boldyrev AI: **Global minimum structure searches via particle swarm optimization**. *J Comput Chem* 2007, **28**(7):1177-1186.
55. Rarey M, Kramer B, Lengauer T, Klebe G: **A fast flexible docking method using an incremental construction algorithm**. *J Mol Biol* 1996, **261**(3):470-489.
56. Morris GM GD, Halliday RS, Huey R, Hart WE, Below RK, Olson AJ,: **Automated Docking Using a Lamarckian Genetic Algorithm and an Empirical Binding Free Energy Function** *J Comput Chem* 1998, **19**:1639-1662.
57. McGovern SL, Shoichet BK: **Information decay in molecular docking screens against holo, apo, and modeled conformations of enzymes**. *Journal Med Chem* 2003, **46**(14):2895-2907.
58. Irwin JJ, Sterling T, Mysinger MM, Bolstad ES, Coleman RG: **ZINC: A Free Tool to Discover Chemistry for Biology**. *J Chem Inf Model* 2012 **52**(7):1757-68.

59. Hawkins PC, Skillman AG, Warren GL, Ellingson BA, Stahl MT: **Conformer generation with OMEGA: algorithm and validation using high quality structures from the Protein Databank and Cambridge Structural Database**. *J Chem Inf Model* 2010, **50**(4):572-584.
60. Hawkins PC, Nicholls A: **Conformer generation with OMEGA: learning from the data set and the analysis of failures**. *J Chem Inf Model* 2012, **52**(11):2919-2936.
61. Koes D, Khoury K, Huang Y, W. W, Bista M, Popowicz GM, Wolf S, Holak TA, Domling A, Camacho CJ: **Enabling Large-Scale Design, Synthesis and Validation of Small Molecule Protein-Protein Antagonists**. *PloS one* 2012, **7**(3):e32839.
62. Cavasotto CN, Kovacs JA, Abagyan RA: **Representing receptor flexibility in ligand docking through relevant normal modes**. *J Am Chem Soc* 2005, **127**(26):9632-9640.
63. Bottegoni G, Kufareva I, Totrov M, Abagyan R: **Four-dimensional docking: a fast and accurate account of discrete receptor flexibility in ligand docking**. *Journal Med Chem* 2009, **52**(2):397-406.
64. Eyrisch S, Helms V: **Transient pockets on protein surfaces involved in protein-protein interaction**. *Journal Med Chem* 2007, **50**(15):3457-3464.
65. Brown SP, Hajduk PJ: **Effects of conformational dynamics on predicted protein druggability**. *ChemMedChem* 2006, **1**(1):70-72.
66. Leis S, Zacharias M: **Efficient inclusion of receptor flexibility in grid-based protein-ligand docking**. *J Comput Chem* 2011, **32**(16):3433-3439.
67. Osguthorpe DJ, Sherman W, Hagler AT: **Generation of receptor structural ensembles for virtual screening using binding site shape analysis and clustering**. *Chem Biol Drug Des* 2012, **80**(2):182-193.
68. Fan H, Irwin JJ, Webb BM, Klebe G, Shoichet BK, Sali A: **Molecular docking screens using comparative models of proteins**. *J Chem Inf Model* 2009, **49**(11):2512-2527.
69. Lane D, Levine A: **p53 Research: the past thirty years and the next thirty years**. *Cold Spring Harb Perspect Biol* 2010, **2**(12):a000893.
70. Yue X, Zhao Y, Xu Y, Zheng M, Feng Z, Hu W: **Mutant p53 in Cancer: Accumulation, Gain-of-Function, and Therapy**. *J Mol Biol* 2017, **429**(11):1595-1606.
71. Haupt Y, Maya R, Kazaz A, Oren M: **Mdm2 promotes the rapid degradation of p53**. *Nature* 1997, **387**(6630):296-299.
72. Baptiste N, Friedlander P, Chen X, Prives C: **The proline-rich domain of p53 is required for cooperation with anti-neoplastic agents to promote apoptosis of tumor cells**. *Oncogene* 2002, **21**(1):9-21.
73. Bieging KT, Attardi LD: **Deconstructing p53 transcriptional networks in tumor suppression**. *Trends Cell Biol* 2012, **22**(2):97-106.
74. Meek DW: **The p53 response to DNA damage**. *DNA repair* 2004, **3**(8-9):1049-1056.
75. Su X, Chakravarti D, Cho MS, Liu L, Gi YJ, Lin YL, Leung ML, El-Naggar A, Creighton CJ, Suraokar MB *et al*: **TAp63 suppresses metastasis through coordinate regulation of Dicer and miRNAs**. *Nature* 2010, **467**(7318):986-990.
76. Adorno M, Cordenonsi M, Montagner M, Dupont S, Wong C, Hann B, Solari A, Bobisse S, Rondina MB, Guzzardo V *et al*: **A Mutant-p53/Smad complex opposes p63 to empower TGFbeta-induced metastasis**. *Cell* 2009, **137**(1):87-98.
77. Brachmann RK: **p53 mutants: the achilles' heel of human cancers?** *Cell cycle* 2004, **3**(8):1030-1034.
78. Gorlov IP, Gorlova OY, Amos CI: **Predicting the oncogenicity of missense mutations reported in the International Agency for Cancer Research (IARC) mutation database on p53**. *Human Mut* 2005, **26**(5):446-454.
79. Hainaut P, Hollstein M: **p53 and human cancer: the first ten thousand mutations**. *Adv Cancer Res* 2000, **77**:81-137.

80. Bullock AN, Henckel J, Fersht AR: **Quantitative analysis of residual folding and DNA binding in mutant p53 core domain: definition of mutant states for rescue in cancer therapy**. *Oncogene* 2000, **19**(10):1245-1256.
81. Ang HC, Joerger AC, Mayer S, Fersht AR: **Effects of common cancer mutations on stability and DNA binding of full-length p53 compared with isolated core domains**. *J Biol Chem* 2006, **281**(31):21934-21941.
82. Dearth LR, Qian H, Wang T, Baroni TE, Zeng J, Chen SW, Yi SY, Brachmann RK: **Inactive full-length p53 mutants lacking dominant wild-type p53 inhibition highlight loss of heterozygosity as an important aspect of p53 status in human cancers**. *Carcinogenesis* 2007, **28**(2):289-298.
83. Shiraishi K, Kato S, Han SY, Liu W, Otsuka K, Sakayori M, Ishida T, Takeda M, Kanamaru R, Ohuchi N *et al*: **Isolation of temperature-sensitive p53 mutations from a comprehensive missense mutation library**. *J Biol Chem* 2004, **279**(1):348-355.
84. Ventura A, Kirsch DG, McLaughlin ME, Tuveson DA, Grimm J, Lintault L, Newman J, Reczek EE, Weissleder R, Jacks T: **Restoration of p53 function leads to tumour regression in vivo**. *Nature* 2007, **445**(7128):661-665.
85. Friedler A, Hansson LO, Veprintsev DB, Freund SM, Rippin TM, Nikolova PV, Proctor MR, Rudiger S, Fersht AR: **A peptide that binds and stabilizes p53 core domain: chaperone strategy for rescue of oncogenic mutants**. *Proc Natl Acad Sci U S A* 2002, **99**(2):937-942.
86. Zandi R, Selivanova G, Christensen CL, Gerds TA, Willumsen BM, Poulsen HS: **PRIMA-1Met/APR-246 induces apoptosis and tumor growth delay in small cell lung cancer expressing mutant p53**. *Clin Cancer Res* 2011, **17**(9):2830-2841.
87. Foster BA, Coffey HA, Morin MJ, Rastinejad F: **Pharmacological rescue of mutant p53 conformation and function**. *Science* 1999, **286**(5449):2507-2510.
88. Synnott NC, Murray A, McGowan PM, Kiely M, Kiely PA, O'Donovan N, O'Connor DP, Gallagher WM, Crown J, Duffy MJ: **Mutant p53: a novel target for the treatment of patients with triple-negative breast cancer?** *Int J Cancer* 2017, **140**(1):234-246.
89. Nguyen D, Liao W, Zeng SX, Lu H: **Reviving the guardian of the genome: Small molecule activators of p53**. *Pharmacol Ther* 2017.
90. Parrales A, Iwakuma T: **Targeting Oncogenic Mutant p53 for Cancer Therapy**. *Front Oncol* 2015, **5**:288.
91. Yu X, Blanden AR, Narayanan S, Jayakumar L, Lubin D, Augeri D, Kimball SD, Loh SN, Carpizo DR: **Small molecule restoration of wildtype structure and function of mutant p53 using a novel zinc-metallochaperone based mechanism**. *Oncotarget* 2014, **5**(19):8879-8892.
92. Liu X, Wilcken R, Joerger AC, Chuckowree IS, Amin J, Spencer J, Fersht AR: **Small molecule induced reactivation of mutant p53 in cancer cells**. *Nucleic Acids Res* 2013, **41**(12):6034-6044.
93. Bauer MR, Jones RN, Baud MG, Wilcken R, Boeckler FM, Fersht AR, Joerger AC, Spencer J: **Harnessing Fluorine-Sulfur Contacts and Multipolar Interactions for the Design of p53 Mutant Y220C Rescue Drugs**. *ACS Chem Biol* 2016, **11**(8):2265-2274.
94. Butler JS, Loh SN: **Structure, function, and aggregation of the zinc-free form of the p53 DNA binding domain**. *Biochemistry* 2003, **42**(8):2396-2403.
95. Benchimol S: **p53-dependent pathways of apoptosis**. *Cell Death Differ* 2001, **8**(11):1049-1051.
96. el-Deiry WS, Harper JW, O'Connor PM, Velculescu VE, Canman CE, Jackman J, Pietenpol JA, Burrell M, Hill DE, Wang Y *et al*: **WAF1/CIP1 is induced in p53-mediated G1 arrest and apoptosis**. *Cancer Res* 1994, **54**(5):1169-1174.
97. Nakano K, Vousden KH: **PUMA, a novel proapoptotic gene, is induced by p53**. *Mol Cell* 2001, **7**(3):683-694.
98. Wassman CD, Baronio R, Demir O, Wallentine BD, Chen CK, Hall LV, Salehi F, Lin DW, Chung BP, Hatfield GW *et al*: **Computational identification of a transiently open L1/S3 pocket for reactivation of mutant p53**. *Nat Commun* 2013, **4**:1407.

99. Bromley D, Bauer MR, Fersht AR, Daggett V: **An in silico algorithm for identifying stabilizing pockets in proteins: test case, the Y220C mutant of the p53 tumor suppressor protein**. *Protein Eng Des Sel : PEDS* 2016, **29**(9):377-390.
100. Liu Q, Kaneko S, Yang L, Feldman RI, Nicosia SV, Chen J, Cheng JQ: **Aurora-A abrogation of p53 DNA binding and transactivation activity by phosphorylation of serine 215**. *J Biol Chem* 2004, **279**(50):52175-52182.
101. Sen S, Zhou H, White RA: **A putative serine/threonine kinase encoding gene BTAK on chromosome 20q13 is amplified and overexpressed in human breast cancer cell lines**. *Oncogene* 1997, **14**(18):2195-2200.
102. Xu HT, Lai WL, Liu HF, Wong LL, Ng IO, Ching YP: **PAK4 Phosphorylates p53 at Serine 215 to Promote Liver Cancer Metastasis**. *Cancer Res* 2016, **76**(19):5732-5742.
103. Deneberg S, Cherif H, Lazarevic V, Andersson PO, von Euler M, Juliusson G, Lehmann S: **An open-label phase I dose-finding study of APR-246 in hematological malignancies**. *Blood cancer J* 2016, **6**(7):e447.
104. Lambert JM, Gorzov P, Veprintsev DB, Soderqvist M, Segerback D, Bergman J, Fersht AR, Hainaut P, Wiman KG, Bykov VJ: **PRIMA-1 reactivates mutant p53 by covalent binding to the core domain**. *Cancer Cell* 2009, **15**(5):376-388.
105. Arner ES: **Focus on mammalian thioredoxin reductases--important selenoproteins with versatile functions**. *Biochim Biophys Acta* 2009, **1790**(6):495-526.
106. Peng X, Zhang MQ, Conserva F, Hosny G, Selivanova G, Bykov VJ, Arner ES, Wiman KG: **APR-246/PRIMA-1MET inhibits thioredoxin reductase 1 and converts the enzyme to a dedicated NADPH oxidase**. *Cell Death Dis* 2013, **4**:e881.
107. Tessoulin B, Descamps G, Moreau P, Maiga S, Lode L, Godon C, Marionneau-Lambot S, Oullier T, Le Gouill S, Amiot M *et al*: **PRIMA-1Met induces myeloma cell death independent of p53 by impairing the GSH/ROS balance**. *Blood* 2014, **124**(10):1626-1636.
108. Bykov VJ, Issaeva N, Zache N, Shilov A, Hultcrantz M, Bergman J, Selivanova G, Wiman KG: **Reactivation of mutant p53 and induction of apoptosis in human tumor cells by maleimide analogs**. *J Biol Chem* 2005, **280**(34):30384-30391.
109. Bykov VJ, Issaeva N, Shilov A, Hultcrantz M, Pugacheva E, Chumakov P, Bergman J, Wiman KG, Selivanova G: **Restoration of the tumor suppressor function to mutant p53 by a low-molecular-weight compound**. *Nature medicine* 2002, **8**(3):282-288.
110. Khar KR, Goldschmidt L, Karanicolas J: **Fast docking on graphics processing units via Ray-Casting**. *PloS one* 2013, **8**(8):e70661.
111. Qin H, Hu J, Hua Y, Challa SV, Cross TA, Gao FP: **Construction of a series of vectors for high throughput cloning and expression screening of membrane proteins from Mycobacterium tuberculosis**. *BMC biotechnology* 2008, **8**:51.
112. Lea WA, O'Neil PT, Machen AJ, Naik S, Chaudhri T, McGinn-Straub W, Tischer A, Auton MT, Burns JR, Baldwin MR *et al*: **Chaperonin-Based Biolayer Interferometry To Assess the Kinetic Stability of Metastable, Aggregation-Prone Proteins**. *Biochemistry* 2016, **55**(35):4885-4908.
113. Garraway LA: **Genomics-driven oncology: framework for an emerging paradigm**. *J Clin Oncol* 2013, **31**(15):1806-1814.