

A Novel Ensemble-Based Scoring and Search Algorithm for Protein Redesign and Its Application to Modify the Substrate Specificity of the Gramicidin Synthetase A Phenylalanine Adenylation Enzyme

RYAN H. LILIEN,^{1,2,3} BRIAN W. STEVENS,^{3,4} AMY C. ANDERSON,^{3,5}
and BRUCE R. DONALD^{1,3,5,6}

ABSTRACT

Realization of novel molecular function requires the ability to alter molecular complex formation. Enzymatic function can be altered by changing enzyme–substrate interactions via modification of an enzyme’s active site. A redesigned enzyme may either perform a novel reaction on its native substrates or its native reaction on novel substrates. A number of computational approaches have been developed to address the combinatorial nature of the protein redesign problem. These approaches typically search for the global minimum energy conformation among an exponential number of protein conformations. We present a novel algorithm for protein redesign, which combines a statistical mechanics–derived ensemble-based approach to computing the binding constant with the speed and completeness of a branch-and-bound pruning algorithm. In addition, we developed an efficient deterministic approximation algorithm, capable of approximating our scoring function to arbitrary precision. In practice, the approximation algorithm decreases the execution time of the mutation search by a factor of ten. To test our method, we examined the Phe-specific adenylation domain of the nonribosomal peptide synthetase gramicidin synthetase A (GrsA-PheA). Ensemble scoring, using a rotameric approximation to the partition functions of the bound and unbound states for GrsA-PheA, is first used to predict binding of the wildtype protein and a previously described mutant (selective for leucine), and second, to switch the enzyme specificity toward leucine, using two novel active site sequences computationally predicted by searching through the space of possible active site mutations. The top scoring *in silico* mutants were created in the wetlab and dissociation/binding constants were determined by fluorescence quenching. These tested mutations exhibit the desired change in specificity from Phe to Leu. Our ensemble-based algorithm, which flexibly models both protein and ligand

¹Dartmouth Computer Science Department, Hanover, NH 03755.

²Dartmouth Medical School, Hanover, NH 03755.

³Dartmouth Center for Structural Biology and Computational Chemistry, Hanover, NH 03755.

⁴Dartmouth Department of Biochemistry, Hanover, NH 03755.

⁵Dartmouth Department of Chemistry, Hanover, NH 03755.

⁶Dartmouth Department of Biology, Hanover, NH 03755.

using rotamer-based partition functions, has application in enzyme redesign, the prediction of protein–ligand binding, and computer-aided drug design.

Key words: protein design, enzyme design, protein flexibility, protein–ligand binding, molecular ensemble, nonribosomal peptide synthetase, fluorescence binding assay

Abbreviations used: DTT, dithiothreitol; GMEC, global minimum energy conformation; GrsA-PheA, gramicidin synthetase A–phenylalanine adenylation domain; LB, luria broth; NRPS, non-ribosomal peptide synthetase; PCR, polymerase chain reaction; PMSF, phenylmethanesulfonyl fluoride; RMSD, root mean square distance; WT, wildtype.

1. INTRODUCTION

IN A VARIETY OF FUNGI AND PARASITES, nonribosomal peptide synthetase (NRPS) enzymes complement the traditional ribosomal peptide synthesis pathway. NRPS enzymes produce peptide-like products via the incorporation of both standard and nonstandard amino acid precursors. Unlike the ribosome, many NRPS enzymes methylate or epimerize their amino acid substrates, join them with peptide or ester bonds, and sometimes cyclize their final product. NRPS products include natural antibiotics (e.g., penicillin, vancomycin), antifungals, antivirals, anticancer therapeutics, immunosuppressants, and siderophores. Enzymes of the NRPS pathway have multiple domains with individual functions acting in an assembly-line fashion (Fig. 1). It is believed that the substrate specificity of the NRPS enzymes is dictated primarily by the “gatekeeper” adenylation (A) domain that binds and acylates the incoming amino acid, forming an amino-acyl adenylate (Stachelhaus *et al.*, 1999; Challis *et al.*, 2000; Schwarzer *et al.*, 2003). Recent evidence also indicates that the condensation (C), thiolation (T), and epimerization (E) domains may carry some specificity as well, albeit to a lesser extent (Belshaw *et al.*, 1999; Weber *et al.*, 2000; Ehmann *et al.*, 2000; Linne *et al.*, 2001).

Enzyme redesign of NRPS enzymes offers the opportunity to reengineer biosynthetic pathways, greatly increasing the number and types of NRPS products (Fig. 2). Therefore, the interest in redesigning NRPS enzymes is motivated by the long-range goal of reprogramming the enzymatic pathway to achieve combinatorial biosynthesis, and the development of new libraries of antibiotics (Cane *et al.*, 1998). We explore the idea of reprogramming NRPS enzymes by introducing K^* , an ensemble-based protein redesign algorithm, to analyze and redesign the phenylalanine adenylation domain of the NRPS enzyme gramicidin synthetase A (GrsA-PheA).

1.1. Computational protein design

A variety of computational protein redesign efforts have recently been reported. Incorporation of molecular flexibility into protein design is essential; every previous structure-based protein design algorithm has included some notion of flexibility (Street and Mayo, 1999; Jin *et al.*, 2003; Jaramillo *et al.*, 2001; Hellinga and Richards, 1991; Bolon and Mayo, 2001; Offredi *et al.*, 2003; Looger *et al.*, 2003; Keating *et al.*, 2001; Shifman and Mayo, 2002). Many protein design algorithms treat the peptide backbone as rigid



FIG. 1. Gramicidin S synthetase is composed of two NRPS proteins, GrsA (3 domains) and GrsB (13 domains). Gramicidin S is produced in an assembly-line manner where two D-Phe-L-Pro-L-Val-L-Orn-L-Leu peptides are joined and cyclized. (A: adenylation, T: thiolation (peptidyl carrier protein), E: epimerization, C: condensation, TE: thioesterase)




FIG. 2. Gramicidin redesign. The GrsA-PheA domain controls incorporation of the first amino acid in the synthesis of the antibiotic gramicidin. (**Left**) The natural gramicidin construct is shown with the incorporated phenylalanine. By changing the substrate specificity of the GrsA-PheA domain to accept leucine, it may be possible to create a modified gramicidin (**right**) where the phenylalanines have been replaced by leucine.

and model amino acid side-chain flexibility with a rotamer library containing a discrete set of side-chain conformations (Lovell *et al.*, 2000; Ponder and Richards, 1987).

Design algorithms have been based on the assumption that for a given protein sequence, folding and binding can best be predicted by examining the single global minimum energy conformation (GMEC). Unfortunately, protein design by searching for the GMEC using rotamers and a pairwise energy function on a rigid peptide backbone has recently been shown to be NP-hard (Pierce and Winfree, 2002). As a result, a number of heuristic (random sampling, neural network, genetic algorithm) GMEC-based approaches for protein design have been reported (Street and Mayo, 1999; Jin *et al.*, 2003; Jaramillo *et al.*, 2001; Hellinga and Richards, 1991; Marvin and Hellinga, 2001); however, the dominant algorithm for assisting in the GMEC search has been dead-end elimination (DEE) (Desmet *et al.*, 1992; Lasters and Desmet, 1993; Pierce *et al.*, 2000). Given a protein backbone, a set of allowable mutations, and a rotamer library, DEE employs a number of sophisticated conformation pruning techniques to prune conformations that are provably not part of the GMEC. Typically, DEE will eliminate the vast majority of mutation sequences; remaining sequences can subsequently be scored and ranked. Growing evidence supports the hypothesis that protein–ligand binding can involve a number of low-energy bound states (Decanniere *et al.*, 2001; Raag and Poulos, 1991; Murthy *et al.*, 1992; Wojtczak *et al.*, 1996; Zhou *et al.*, 2000; Montfort *et al.*, 1990). Therefore, we have developed a scoring method for protein–ligand redesign based on molecular ensembles. Molecular ensembles have been successfully utilized in structure-based drug design: most commonly, molecular docking is performed against each member of an ensemble or a unified ensemble model and an average or best interaction energy between the protein and ligand may be retained (Claußen *et al.*, 2001; Österberg *et al.*, 2002; Lilien *et al.*, 2000; Bouzida *et al.*, 1999; Knegetel *et al.*, 1997; Carlson, 2002).

1.2. Previous NRPS redesign

NRPS enzyme redesign methods can be divided into two main techniques, *domain-swapping* and *active site modification through site-directed mutagenesis*. Domain-swapping techniques do not require computational analysis nor knowledge of molecular structure; NRPS enzymes are modified by swapping an adenylation domain of an existing NRPS enzyme for an adenylation domain from a second, different NRPS enzyme (carrying a different substrate specificity) (Stachelhaus *et al.*, 1995; Schneider *et al.*, 1998; Doekel and Marahiel, 2000; Mootz *et al.*, 2000). Results of domain-swapping experiments led to the hypothesis that the disruption of native domain:domain interfaces vitiates the proper transfer of synthesis intermediates thereby degrading catalytic efficiency (Linne *et al.*, 2001). Emphasis on the importance of domain:domain interactions and domain specificity has directed domain-swapping work to include simultaneous cloning of A, C, and T domains, demonstrating increased yield (Doekel and Marahiel, 2000; Mootz *et al.*, 2000).

The second method for NRPS redesign, active site modification through site-directed mutagenesis, utilizes structural information of the GrsA-PheA enzyme (IAMU [Conti *et al.*, 1997]). Sequence alignment of GrsA-PheA with 160 other known adenylation domains supports the hypothesis that NRPS adenylation domains, specific for different amino acid substrates, share a similar overall structure differing mainly in

the composition of residues lining the active site (Stachelhaus *et al.*, 1999; Eppelmann *et al.*, 2002; Challis *et al.*, 2000). A “signature sequence” can be derived for each adenylation domain by extracting those residues that align with the structurally determined substrate binding pocket of the GrsA-PheA crystal structure. By abstracting away from the GrsA-PheA crystal structure, mutations are suggested for a given amino acid substrate by sequence comparison alone. Using signature sequences, Stachelhaus *et al.* (1999) mutated two adenylation domains: PheA, successfully switching the substrate specificity from Phe to Leu; and a second adenylation domain that naturally accepts Asp to accept Asn. More recently, Eppelmann *et al.* (2002) changed the Glu adenylation domain of surfactin synthetase A to accept Gln.

In summary, previous NRPS redesign methods include domain-swapping and site-directed mutagenesis from active site signature sequences. Our method, active site manipulation by site-directed mutagenesis from a computational mutation search utilizing ensemble docking, adds to the armamentarium of techniques available for protein redesign and confers some significant advantages over existing NRPS redesign methods. Signature sequence methods project active site information into a consensus sequence, thus losing structural information. Because molecular structure is not explicitly considered during redesign, successful redesign is more difficult if there are significant structural differences affecting the overall active site shape between the A domains accepting the natural and target substrates. Our method builds mutations into the high-resolution structure of the wildtype enzyme, thus mitigating potential problems arising from these structural differences. In contrast to domain-swapping, our method is more likely to preserve the NRPS enzyme’s native modular structure, thus maintaining crucial specific domain:domain interface regions. Finally, unlike either signature sequence or domain-swapping techniques, our method can propose mutations for substrates for which no existing adenylation domain sequences are known.

When considering the use of molecular ensembles for protein design, a major challenge has been the development of ensemble-based redesign algorithms that efficiently prune mutations and conformations. In this paper, we introduce the K^* method which generalizes Boltzmann-based scoring to ensembles and applies the result to protein design. The following contributions are made in this work:

1. Introduction of the first ensemble-based protein redesign algorithm;
2. Development of ϵ -approximation algorithms for K^* capable of pruning the vast majority of conformations from more computationally expensive consideration thereby reducing execution time and making the mutation search computationally feasible;
3. The use of K^* to reproduce known adenylation domain binding experiments;
4. The use of K^* to predict novel mutation sequences capable of switching substrate specificity of GrsA-PheA;
5. Confirmation of the K^* method by the creation of predicted protein mutants in the wetlab and testing of their specificity by fluorescence quenching binding assays.

2. METHODS

2.1. Ensemble scoring method (K^*)

Protein–ligand binding in a single mutation is modeled using the following K^* equation:

$$K^* = \frac{q_{PL}}{q_P q_L}. \quad (1)$$

K^* is derived to be an approximation to the true association (binding) constant K_A by expressing each species’ chemical potential as a function of the species partition function q (Hill, 1956; McQuarrie, 1976) and solving for the equilibrium condition (full details are provided in Appendix A). Unfortunately, it is not currently possible to compute exact partition functions for complex molecular species. This would require integrating an exact energy function over a molecule’s entire conformational space. We therefore approximate these partition functions with the use of rotamerically-based conformational ensembles:

$$q_{PL} = \sum_{b \in B} \exp(-E_b/RT), \quad q_P = \sum_{f \in F} \exp(-E_f/RT), \quad q_L = \sum_{l \in L} \exp(-E_l/RT), \quad (2)$$

where B , F , and L represent rotamer-based ensembles for the bound protein–ligand complex (PL), the free protein (P), and the free ligand (L), respectively, E_s is the energy of conformation s , R is the gas constant, and T is the temperature in Kelvin. The accuracy with which K^* approximates K_A is proportional to the accuracy of the partition function approximation used.

2.2. Mutation search

When applying K^* to a protein–ligand system, a number of choices must be made with respect to ensemble generation and single-structure scoring. *Single-structure scoring* is the method by which each individual member of an ensemble is scored. These individual scores are then combined using Equations (1) and (2) to compute an ensemble score. The choices made in ensemble scoring should strike a balance between fidelity to the underlying physical biochemistry, and computational feasibility even in the inner loop of a combinatorially expensive search. A more detailed single structure scoring model takes longer to evaluate and typically cannot be used with large or complex molecules. We present one implementation of the K^* scoring function; however, alternate schemes for both ensemble generation and single-structure scoring can be explored and utilized within the K^* framework. First, in this section, we present a brute-force algorithm that does not utilize filters on mutation space nor pruning of conformation space. The algorithm is then extended to utilize both mutation space filters and conformation space pruning in Section 2.3. The specific application of K^* to redesign GrsA-PheA is described in Section 3.

In the general case, for each allowable active site sequence, K^* is computed using the following three steps.

Step 1: Ensemble generation. Molecular ensembles are generated by fixing the protein backbone and using the Lovell *et al.* (2000) rotamer library to vary side-chain conformation. Each flexibly modeled residue is allowed to sample all allowable rotamers. Steric clash is determined by computing the distance that two atoms' AMBER van der Waals' (vdW) spheres penetrate each other. Conformations containing any pair of atoms with more than 1.5 Å of steric clash are discarded. Finally, if the volume of the active site is large relative to the size of the ligand, multiple translations of the ligand in the active site should be generated and included in the bound ensemble. Because the ligand binding site in GrsA-PheA tightly fits the amino acid ligand, when modeling GrsA-PheA a maximum translation of 1.2 Å (and a maximum translation step size of 0.1 Å) was used. To model other larger molecular systems, we utilized a larger three-dimensional grid placed over the active site to generate additional ligand translations.

Step 2: Ensemble scoring (single-structure scoring). In the brute-force algorithm, all conformations that pass step 1 are energy-minimized using our implementation of the AMBER energy function (containing electrostatic, vdW, and dihedral terms) (Weiner *et al.*, 1984; Cornell *et al.*, 1995). In our model, hydrogen atoms are added to all residues by the LEAP module of the AMBER distribution and are used in computing the electrostatic but not vdW energies. We perform a constrained minimization (analogous to voxel minimization [Tucker-Kellogg, 2002; Rienstra *et al.*, 2002]) on each rotameric conformation where side chain dihedrals on flexible residues (including the ligand when present) may move by up to $\pm 9^\circ$ and, for the bound states, the entire ligand may rotate and translate in the active site. This allows our algorithm to sample a larger region of conformation space while not allowing one rotamer to minimize into another.

Step 3: K^ scoring.* The three partition functions q_{PL} , q_P , and q_L are computed separately. The energy minimized scores from Step 2 are used to compute each partition function which is then combined to compute K^* (Equation 1).

After computing K^* for each mutation, the top mutations (those with the largest K^*) are examined graphically and selected for testing in the wetlab. It is worth noting that although the AMBER scoring function we use for single structure scoring contains only enthalpic terms, the K^* method of ensemble scoring encompasses conformational entropy through use of the partition function over ligand and side chain conformations.

2.3. Efficient algorithms for mutation search

Because protein redesign is NP-Hard (Pierce and Winfree, 2002), there is most likely no way, in the worst case, to avoid having to potentially examine an exponential number of conformations. That is to say, the run time of a protein design algorithm that returns the optimal mutation sequence (under a given metric) is likely to be inherently exponential. In contrast, a random sampling mutation search algorithm

can run in subexponential time, yet the mutation sequence returned is not guaranteed to be optimal. The combinatorial nature of protein redesign is exacerbated when utilizing an ensemble-based scoring function because multiple low-energy states must be considered for each mutation sequence. Therefore, in designing an ensemble-based mutation search algorithm it is necessary to prune unlikely conformations and mutations and also to reduce the runtime constant. Consequently, our mutation search algorithm utilizes the following methods. First, volume and steric filters prune a combinatorial number of conformations from consideration. Thus, analogous to DEE, our pruning techniques eliminate the majority of conformations early in the mutation search. Second, our K^* approximation algorithm quickly discards the majority of remaining conformations by placing provable bounds on each conformation's contribution to the partition function; only those conformations that significantly contribute to the partition function are further considered. This algorithm reduces the average amount of time spent examining each conformation and is essential in practice; without such optimizations the ensemble-based mutation search would not be possible.

Sequence-space filters. A residue type filter restricts the mutation search to include only a subset of amino acids based on compatibility with the target substrate. A volume filter removes mutations that significantly over- or underpack the substrate-bound active site relative to the wildtype.

Deterministic approximation algorithm. In a Boltzmann distribution, conformations with large energies are not likely to be assumed and contribute only a vanishingly small amount to the partition function. We therefore prune conformations from consideration (and hence minimization) when we know that they will contribute only a small percentage to the total partition function. In Section 2.3.2, we show that the true partition function can be provably approximated to arbitrary precision.

2.3.1. Conformation generation and the steric filter. Active site rotameric conformations are generated by traversing a conformation tree in a depth-first search order. In a conformation tree (Fig. 3), the rotamers of flexible residue i are represented by the branches at depth i . For example, in Fig. 3, residue A has four rotamers, residue B has two rotamers, and residue C has three rotamers. Internal nodes of a conformation tree represent partially assigned conformations. For example, in Fig. 3, node **e** represents the partially assigned conformation where residue A has assumed rotamer 2 and residue B has assumed rotamer 1; no rotamer has yet been assigned for residue C. Nodes of the conformation tree are visited in a depth-first search order. As each node is visited, conformations with more than 1.5 \AA of preminimization steric overlap are not further considered, thereby pruning a branch of the conformation tree. For example, when generating the children of node **b**, rotamer 2 is not assigned to residue B since it causes steric clash. That branch of the search tree is pruned and not considered further. Steric clash identified at higher levels of the conformation tree prunes more conformations than the identification of steric clash at lower nodes of the conformation tree. Pruning at depth i eliminates $O(c^{n-i})$ conformations, where c is the average number of rotamers per amino acid type and n is the total number of flexible residues.

2.3.2. Intramutation pruning. We now derive the energetic-based pruning method and quantify the total error accrued by ignoring pruned conformations when computing a single partition function. Because




FIG. 3. An example conformation tree. The rotamers of flexible residue i are represented by the branches at depth i . Internal nodes of a conformation tree represent partially-assigned conformations. Red X's represent nodes of the conformation tree where steric clash has been identified among a partially assigned conformation. All children of X nodes are pruned and not considered.

this technique is applied during the computation of a single partition function for a single mutation, it is called *intra-mutation pruning*. We show that any desired approximation accuracy to the true partition function can be guaranteed.

Assume that a partition function is to be computed over n sterically allowed conformations. Let $C_k = \{c_1, c_2, \dots, c_k\}$ be the subset containing the first k conformations such that C_n contains all the conformations. Let q_k be the partial partition function computed by evaluating the first k conformations,

$$q_k = \sum_{c \in C_k} \exp(-E_c/RT).$$

Let q_k^* be an approximation to q_k determined by examining a subset of states S_k , such that $S_k \subseteq C_k$ and $q_k^* \leq q_k$,

$$q_k^* = \sum_{s \in S_k} \exp(-E_s/RT).$$

If S_k contains most of the low energy conformations of C_k , then q_k^* will represent a good approximation to q_k . Let p_k be the partition function of the pruned residues ($C_k - S_k$) equal to the difference $q_k - q_k^*$ such that

$$p_k = \sum_{s \in (C_k - S_k)} \exp(-E_s/RT).$$

One method of ensuring that q_k^* is a good approximation to q_k is to maintain an invariant throughout the computation requiring that q_k^* be a good approximation to q_k . Therefore, we maintain that at any point k , q_k^* must be a good approximation to q_k , that is,

$$q_k^* \geq (1 - \epsilon)q_k, \quad \forall k \leq n. \quad (3)$$

Here, ϵ is the desired approximation constant ($\epsilon < 1$). When Equation (3) holds, we say that q_k^* is an ϵ -approximation to q_k . We can maintain Equation (3) by ensuring that $p_k \leq q_k \epsilon$. Since we know that $q_k^* \leq q_k$, we can also maintain Equation (3) and therefore that q_k^* remains an ϵ -approximation to q_k by ensuring that

$$p_k \leq q_k^* \epsilon. \quad (4)$$

To determine a pruning criterion, we assume that we have already considered k conformations and that S_k contains the subset of fully evaluated conformations. To prune conformation c_{k+1} , we must first guarantee that after pruning c_{k+1} , the value q_{k+1}^* is an ϵ -approximation to q_{k+1} . We know that

$$q_{k+1} = q_k^* + p_k + \exp(-E_{c_{k+1}}/RT).$$

If we prune c_{k+1} , then

$$p_{k+1} = p_k + \exp(-E_{c_{k+1}}/RT), \quad (5)$$

$$q_{k+1}^* = q_k^*. \quad (6)$$

To maintain the invariant, we need to ensure $p_{k+1} \leq q_{k+1}^* \epsilon$, which can be rewritten using Equations (5) and (6) as

$$p_k + \exp(-E_{c_{k+1}}/RT) \leq q_k^* \epsilon. \quad (7)$$

Solving Equation (7) for $E_{c_{k+1}}$, we get that

$$E_{c_{k+1}} \geq -RT \ln(q_k^* \epsilon - p_k). \quad (8)$$

Therefore, if the energy of conformation c_{k+1} satisfies Equation (8), then we can prune conformation c_{k+1} while maintaining $q_{k+1}^* \geq (1 - \epsilon)q_{k+1}$ (that is, q_{k+1}^* is an ϵ -approximation to q_{k+1}). Of course, the purpose of pruning is to avoid the computationally expensive energy minimization required to obtain $E_{c_{k+1}}$. In practice, we use a *pairwise energy matrix* to compute a lower bound on $E_{c_{k+1}}$ (full details are in Appendix D) and compare this bound to Equation (8). Because we don't compute $E_{c_{k+1}}$, we cannot maintain an exact value for p_k . Therefore, the lower bounds on $E_{c_{k+1}}$ are used to compute p_k^* , an upper bound on p_k , namely

$$p_k \leq p_k^* = \sum_{s \in (C_k - S_k)} \exp(-B(s)/RT),$$

where $B(s)$ returns a lower bound on the energy of conformation s . Because $-RT \ln(q_k^* \epsilon - p_k^*) \geq -RT \ln(q_k^* \epsilon - p_k)$, the approximation p_k^* can be used to determine the practical pruning criteria; that is, conformation c_{k+1} can be pruned if

$$B(c_{k+1}) \geq -RT \ln(q_k^* \epsilon - p_k^*). \quad (9)$$

This leads to the following lemma:

Lemma 1. *If Equation (9) is satisfied by conformation c_{k+1} , then conformation c_{k+1} can be pruned and q_{k+1}^* is guaranteed to be an ϵ -approximation to q_{k+1} .*

Maintaining the invariant that after considering k conformations q_k^* is an ϵ -approximation to q_k leads to the following lemma:

Lemma 2. *If q_k^* is an ϵ -approximation to q_k for all k ($0 < k \leq n$), then by induction, at the end of the computation, q_n^* will be an ϵ -approximation to q_n .*

The intra-mutation pruning algorithm is shown in Fig. 4.

When using K^* to perform a mutation search, we can bootstrap the pruning condition for improved efficiency (by caching partition functions, we can exploit K^* bounds from other mutations in the same search). Our search algorithm has the desirable property that provably accurate ϵ -approximations are computed for top-ranking mutations, while the bounds we can prove on the quickly-computed K^* values for lower-ranked mutations do not enjoy the same degree of accuracy. Our algorithm requires an ϵ -approximation only for those mutations with $K^* \geq \gamma \max_{i \leq m} K_i^*$, where K_i^* is the K^* value of the i^{th} mutation, m is the total number of mutations, and γ ($0 < \gamma \leq 1.0$) is a user-specified constant defining a range of K^* values that must be computed accurately. Since the value of the largest (best) K^* is not known during

```

Let  $n \leftarrow$  Number of Rotameric Conformations
Let  $c \leftarrow$  Rotameric Conformations
Initialize:  $q^* \leftarrow 0$ ,  $p^* \leftarrow 0$ 
for  $k = 1$  to  $n$ 
    if  $B(c_k) \leq -RT \ln(q^* \epsilon - p^*)$ 
         $q^* \leftarrow q^* + \exp(-\text{ComputeMinEnergy}(c_k)/RT)$ 
    else
         $p^* \leftarrow p^* + \exp(-B(c_k)/RT)$ 
Return  $q^*$ 

```

FIG. 4. Intramutation pruning. Here q^* is the running approximation to the partition function, and p^* is an upper bound on the partition function of the pruned conformations. The function $B(\cdot)$ computes a lower energy bound for the given conformation. The function $\text{ComputeMinEnergy}(\cdot)$ returns the energy of the energy-minimized conformation as computed using steepest-descent minimization and our implementation of the AMBER energy function (as described in Step 2 of Section 2.2). At the end, q^* represents an ϵ -approximation to the true partition function q such that $q^* \geq (1 - \epsilon)q$.

the mutation search, we use K_o^* , the largest K^* value encountered so far. Similarly to Equation (9), after examining k of n total conformations of a partition function q_{PL} , evaluation of conformation c_{k+1} can be skipped if

$$B(c_{k+1}) \geq -RT \ln(\psi_k - p_k^*), \quad (10)$$

where $\psi_k = \max(\epsilon K_o^* q_L \gamma q_P, q_k^* \epsilon)$. The full derivation of Equation (10) is in Appendix C.

In summary, four levels of approximation are used in computing K^* scores. First, a rotamer library (Lovell *et al.*, 2000) is used to model side chain conformations. Second, prior to computing partition function values, sterically disallowed rotamer-based conformations are combinatorially pruned from consideration (Section 2.3.1). Third, intra- and inter-mutation pruning methods (Section 2.3.2 and Appendix C) are used to skip evaluation of conformations not required to compute ϵ -approximations to the partition function and K^* values. Fourth, in computing these partition functions and K^* values, our algorithm starts with discrete rotamers and then performs bounded minimization (Step 2, Section 2.2). The third and fourth of these approximation techniques provide provable bounds on the accuracy of the computed partition function and K^* values. All four of these methods are used in performing the mutation search of Section 3. Our implementation includes a multithreaded distributed parallel application for K^* -based mutation search which allows the mutation search to utilize any Unix-based machine with access to our filesystem. Approximately 18 processors were used in the Leu redesign of Section 3.4.

3. RESULTS AND DISCUSSION

3.1. Structural model

Our structural model employs the previously solved structure of GrsA-PheA (1AMU) (Conti *et al.*, 1997) and consists of the 9 active site residues (D235, A236, W239, T278, I299, A301, A322, I330, C331) (Fig. 5), the 30 residues with at least one atom within 8 Å of the active site (termed the *steric shell*), the amino acid substrate, and the AMP cofactor. The steric shell allows us to compute the interaction energy of the active site residues with neighboring regions of the protein and constrains the active site residues from assuming conformations that would sterically clash with the body of the PheA protein.

3.2. Comparison to wildtype PheA

We performed a series of experiments to confirm the proper implementation of the steric filter, the rotamer library, the AMBER energy function, and the minimization algorithm. The first test was designed to confirm that we could find an accepted (crystallographically confirmed) conformation for Phe in the PheA active site. The bound partition function q_{PL} was computed for Phe in the GrsA-PheA wildtype (WT) protein. The energy calculated for the best minimized rotamer structure (i.e., the lowest computed energy) and that calculated for the crystal structure are within 5% of each other and have a non-hydrogen atom RMSD of 0.66 Å (Fig. 5A). In the bound molecular ensemble, approximately 39 conformations of the PheA:Phe complex have energies within 5% of the minimum (Fig. 5B). This observation supports the hypothesis that multiple structures have energies contributing to the weighted ensemble. This test confirmed that we were able to generate structures compatible with the X-ray structure and therefore demonstrated the feasibility of both the rotamer search strategy with minimization and the scoring scheme.

3.3. T278M/A301G double mutant

To further test our model, we simulated the biochemical activity assays of L-Phe and L-Leu against wildtype PheA and the T278M/A301G double mutant (Stachelhaus *et al.*, 1999). The T278M/A301G double mutant was designed by signature sequence homology modeling by Stachelhaus *et al.* (1999) to be similar to a known Leu adenylation domain. The K^* scores were computed for each substrate in each active site and compared with activity assays performed by Stachelhaus *et al.* (1999). Because the experimentally measured solvation energy of Leu and Phe are similar (2.3 vs. 2.4 kcal/mol) (Eisenberg and McLachlan, 1984; Fauchère and Pliška, 1983), in these experiments we chose not to determine q_L computationally but rather treat the q_L values of Leu and Phe as equivalent. Stachelhaus *et al.* normalized the activity

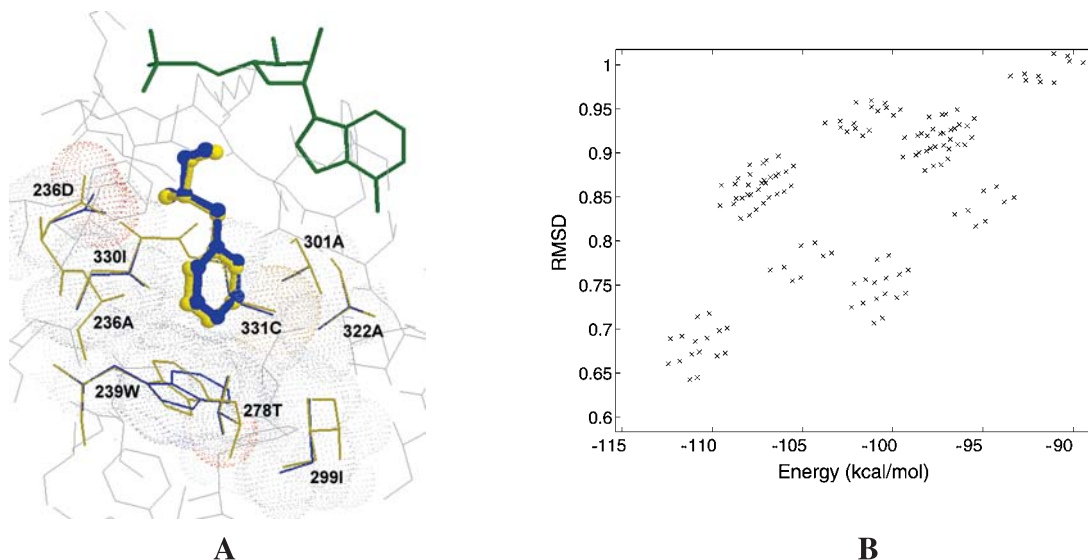


FIG. 5. Comparison of K^* -predicted and crystal structures of GrsA-PheA. (A) The lowest-energy ensemble member of Phe in WT GrsA-PheA is shown with the crystal structure (RMSD: 0.66 Å). Active site residues are yellow (predicted) and blue (crystal), Phe substrate is ball and stick, AMP is green, and residues immediately surrounding the active site (the steric shell) are grey wireframe. The energy-minimized rotamer-based prediction accurately reproduces the crystallographic conformation. The main structural difference occurs in 239W where a small difference in the (C_α - C_β - C_γ) bond angle (118.3 [resp., 115.1] degrees in the crystal structure [resp., prediction]) prevents the rotamer-based conformation from more closely matching the crystallographic conformation. (B) Conformational energy (kcal/mol) vs. RMSD for conformations in the K^* generated ensemble for Phe in wildtype GrsA-PheA. RMSDs are computed between each ensemble conformation, and the crystal structure 1AMU (Conti *et al.*, 1997) using all non-hydrogen atoms. Low-energy conformations have a lower RMSD than high energy conformations, and conformations with low RMSDs have lower energies than conformations with larger RMSDs.




FIG. 6. The top two K^* -predicted mutations and the mutation frequency in the 40 top-ranking mutations. Shown are the lowest energy ensemble members of the bound partition functions for (A) A301G/I330W and (B) A301G/I330F. Residue 301G sterically allows for the Leu C_δ atoms while residues 330W and 330F both stack on residue 239W and serve to fill a void at the bottom of the active site created by the difference in size between Phe and Leu. Mutated residues are shown in orange. (C) The fraction of the top 40 K^* -ranked sequences involving the specified residues. If mutations were randomly distributed, one would expect that each residue would mutate 2/9 (22.2%) of the time (indicated by the vertical black line). Therefore, residues 235, 236, 330, and 331 tend to assume the wildtype amino acid in the K^* -predicted distribution for the Leu-binding mutation search.

TABLE 1. CONFORMATIONAL PRUNING^a

	<i>Conformations remaining</i>	<i>Pruning factor (%)</i>	<i>Pruning type</i>
Initial	6.8×10^8	—	—
Volume filter	1.98×10^8	3.43 (70.9)	C
Steric filter	8.22×10^5	240.86 (99.6)	C
Energy filter	7.99×10^4	10.28 (90.3)	CF

^aThe initial number of conformations for the GrsA-PheA 2-residue Leu mutation search is shown with the number of conformations remaining after the application of volume, steric, and energy (Equation 10) pruning. The pruning factor represents the ratio of the number of conformations present before and after the given pruning stage. The pruning-% (in parentheses) represents the percentage of remaining conformations eliminated by the given pruning stage. The combined pruning factor of all filters is 8,510. The pruning type column indicates if the pruning represents a combinatorial (C) or a constant factor (CF) speedup.

of each protein such that the substrate with the most activity was assigned a specificity of 100%. The wildtype enzyme has a normalized specificity of 100% for Phe and approximately 10% for Leu while the T278M/A301G double mutant has a normalized specificity of approximately 40% for Phe and 100% for Leu (Stachelhaus *et al.*, 1999). Our normalized K^* results closely agree with these specificity scores. For the wildtype enzyme, PheA has a normalized K^* for Phe of 100% and for Leu of 6%. The double mutant enzyme has a K^* of 10% for Phe and 100% for Leu. Although K^* is a binding constant approximation, its results qualitatively agree with the activity assays of Stachelhaus *et al.*

3.4. Redesign for Leu

A K^* mutation search was performed to redesign GrsA-PheA to bind and adenylate Leu instead of Phe. The 9 active site residues, 30 active site neighboring residues, substrate, and AMP were modeled as described above. We performed a 2-residue mutation search, where any two of the nine active site residues were allowed to mutate to any of the hydrophobic residues (GAVLIFYWM) (2,916 possible mutations, approximately 6.8×10^8 conformations). Each mutation was checked against the volume filter. Active sites that were over- or underpacked relative to Phe in the PheA wildtype by more than 30 \AA^3 were eliminated. Altogether 1,011 mutations (35% of the total), containing 1.98×10^8 conformations, passed the volume filter and were fully evaluated. Of the 822,061 conformations that passed the steric pruning step, 742,116 (90.3%) were pruned based on minimum energy bounds (Equation 10) leaving only 79,945 conformations of the original 6.8×10^8 that were then energy minimized and scored (Table 1). The 2-residue mutation search took less than one day on a cluster of 18 1.6 GHz Athlon processors. Without energetic pruning, approximately 10 times as many conformations would require minimization taking approximately 10 times longer to execute. Computationally, for comparison, a three-point mutation search has been run in 10 days; consequently, due to the computation required in minimizing each rotameric conformation, execution without pruning is impractical. In practice, the accuracy of the computed solution is significantly higher than that guaranteed by the approximation. When a 3% approximation (accuracy = 97%) is requested, the accuracy achieved is over 99% which suggests that we may relax the pruning criteria and still maintain an excellent approximation.

The two mutation sequences with the best K^* scores are A301G/I330W and A301G/I330F: these novel mutations are unknown in nature and have never been tested before. The lowest energy predicted conformations of the bound ensemble for the best two mutation sequences are shown in Fig. 6. The first mutation in both sequences, A301G, sterically allows for the difference in position of the C_δ atoms between Phe and Leu. Residue 301G appears in 69% of the top 40 K^* -ranked mutations¹ (Fig. 6C) and is also present in

¹The histogram of amino acid type usage in Fig. 6C is first-order and does not currently examine correlations among mutation pairs. It would be interesting to extend our analysis to identify residues which tend to mutate together and amino acid type pairs which are more common for a particular mutation search. This information could be helpful in deciding which mutation sequences to create and clone in the wetlab.

TABLE 2. DISSOCIATION CONSTANTS^a AND STANDARD ERRORS (IN PARENTHESES) MEASURED BY FLUOROSPECTROPHOTOMETRY FOR AMINO ACID BINDING TO GRSA-PHEA AND THE NOVEL MUTANTS 301G/330W AND 301G/330F

Substrate	Dissociation constant (μM) K_D		
	WT	301G/330W	301G/330F
L-Phe	26 (2.6)	94 (17.2)	53 (4.9)
L-Leu	50 (11.6)	45 (9.5)	30 (2.4)

^aA lower dissociation constant is associated with tighter binding.




FIG. 7. Change in fluorescence vs. Leu concentration for 301G/330F. (Inset) Reciprocal plot of the fluorescence quenching data and the linear regression (correlation coefficient of 0.977) showing a dissociation constant of 30 μM . Dissociation constants and standard errors were computed using SPSS Rel. 11.0.2 (SPSS Inc., Chicago, IL).

all 19 known native Leu adenylation domains (Challis *et al.*, 2000). The second mutation (I330W, I330F) fills the bottom of the substrate binding pocket accounting for the difference in size between Phe and Leu. Both mutations I330W and I330F form a staggered stacked ring structure (Samanta *et al.*, 1999) with the existing residue 239W (Fig. 6). The previously-reported T278M/A301G mutant (Stachelhaus *et al.*, 1999) is ranked 12th out of 2,916 by K^* , thus demonstrating that a known Leu binding mutation is ranked highly in our mutation search.

The GrsA-PheA gene was cloned into the QE60 vector using PCR. The *E. coli* M15 cell lines were transformed with the constructed plasmid. These genes were then modified to incorporate the two desired mutations, 301G/330W and 301G/330F. Complete expression and purification details are in Appendix E. For both mutations, the presence of Trp239 in the active site allowed us to determine the dissociation constants for substrate binding (K_D) by measuring the change in fluorescence of Trp at 340 nm after titrating substrate and exciting at 280 nm. The dissociation constant K_D is inversely proportional to the binding constant $K_A = 1/K_D$. Hence, a smaller K_D is associated with tighter binding. Both mutations clearly exhibit stronger binding for Leu than for Phe (Table 2 and Fig. 7), and the K_D measured for Leu in both redesigned proteins is approximately half that for Phe, strongly demonstrating the success of the protein redesign.

4. COMPARISON TO GMEC SEARCH

For comparison, we explored whether the mutation sequences suggested by the K^* mutation search were indeed different from those that would have been found using a GMEC type search. We therefore compared the K^* -based mutation sequence ranking to two non-ensemble-based scoring techniques. The first alternative mutation scoring method scores each mutation sequence by the lowest minimum energy bound conformation (among all allowable rotameric conformations for the given mutation sequence). We refer to this brute-force approach as the *minimum energy (ME)* technique. The second alternative scoring method simulates the rankings that would be returned by a DEE-type search. This method consists of two stages. The first stage is an initial pruning step based on the lowest-energy bound rotameric conformation

(no energy minimization is performed). In the second stage, a ranking is computed for each remaining mutation based on an energy-minimized binding energy ($E_{bound} - E_{unbound}$). We therefore refer to this method as the *minimized binding energy (MBE)* technique. In computing the MBE ranking, the top 10% scoring mutations from the first stage are scored in stage two. The rankings produced by this method therefore approximate those of a DEE-type search in that the vast majority of mutation sequences are pruned in the DEE stage based on the lowest energy bound rotameric conformation, and unpruned mutations are subsequently scored using a more sophisticated method. To compare the mutation sequence scoring technique rather than the energy function, both the ME and MBE techniques use the same implementation of the AMBER empirical energy function and steepest descent energy minimization used by K^* .

The three scoring methods were applied to all 1,011 2-residue mutation sequences that passed the volume filter during Leu redesign. Of the top 40 K^* -ranked mutations, only two (5%) appear among the top 40 ME-ranked mutations, and only seven (17.5%) appear among the top 100 ME-ranked mutations. Conversely, of the top 40 ME-ranked mutations, only two (5%) appear among the top 40 K^* -ranked mutations, and only eight (20%) appear among the top 100 K^* -ranked mutations. When compared to the MBE method, of the top 40 K^* -ranked mutations only 10 (25%) appeared in either the top 40 or the top 100 MBE-ranked mutations. Conversely, of the top 40 MBE-ranked mutations, 10 (25%) appear among the top 40 K^* -ranked mutations and 23 (57.5%) appear among the top 100 K^* -ranked mutations. Perhaps most interesting is the result that neither of the top two K^* -ranked mutations (A301G/I330W and A301G/I330F) were found among the top 100 ME- or MBE-ranked mutations. Furthermore, the previously reported (T278M/A301G) Leu binding mutation (Stachelhaus *et al.*, 1999) is ranked 80th by the ME method, 3rd by the MBE method, and 12th by the K^* method.

Most of the top K^* -, ME-, and MBE-ranked mutations remain biologically untested, thus precluding an exhaustive comparison of the mutations predicted by the three scoring techniques; however, we can conclude that the top mutation sequences returned by K^* are different from those returned by either ME or MBE. The tested top K^* mutations, shown in this paper to have Leu binding specificity (by wetlab experiments), provide evidence that the K^* rankings provide an additional and effective method for ranking mutation sequences.

5. CONCLUSIONS

The K^* ensemble scoring method presented here was successfully applied to redesign the Phe adenylation domain of gramicidin synthetase A. This represents the first use of an ensemble-based scoring function for enzyme redesign. Despite the inherently exponential nature of ensemble-based scoring, a deterministic approximation algorithm for computing each partition function enables sufficient pruning to make the search feasible. The redesigned enzymes demonstrate a specificity switch from Phe to Leu in binding affinity, and we are now pursuing enzyme activity assays to determine the rate of amino-acyl adenylate formation for the designed proteins. Our ensemble-based mutation search algorithm represents a novel and effective alternative to both domain swapping and the use of signature sequences for NRPS adenylation domain modification.

Many previous modeling algorithms have used biophysically motivated scoring functions to rank results (i.e., LUDI score [Böhm, 1994], DOCK score [Kuntz *et al.*, 1982; Lorber and Shoichet, 1998]). Although based on biophysical phenomenon, these scores often do not provide accurate absolute binding information but rather are useful in predicting relative binding. The K^* scoring method represents a similar type of scoring function. At present, K^* can best provide relative binding to rank mutations for a given ligand. In the future, we hope to enhance K^* to provide more information on absolute binding.

While GMEC-based approaches remain the dominant algorithm for protein design because of their ability to handle the design of large proteins, we propose that, when active site flexibility and multiple binding modes must be considered for the redesign of a moderate-sized system, K^* represents an accurate and feasible approach to ensemble-based redesign. The ease with which the K^* -based mutation search algorithm can be parallelized allows for the utilization of additional hardware resources for larger enzyme redesigns. We have recently scaled the K^* -based mutation search to perform a 3-residue mutation search using over one hundred processors (data not shown) to switch specificity towards Val and Tyr. In this paper, we demonstrate K^* as a stand-alone algorithm; however, in larger systems, DEE can be used to reduce

the exponential number of mutation sequences and conformations prior to K^* scoring. In this manner, K^* could be used to efficiently rank those mutation sequences that survive DEE pruning, leading to a DEE- K^* hybrid search. Enhancements to our pruning methods should increase both the fraction of sequence space searched during protein design and the size of active site for which K^* is feasible.

It would be interesting to extend our algorithm to create “sloppy” adenylation domains capable of adenylation several types of amino acids thereby facilitating combinatorial biosynthesis. Such enzymes could, potentially, play a role in drug synthesis analogous to “generic operations” in computer science. Modified synthesis pathways will create multiple final synthesis products, each demonstrating slight variations on the designed product (Cane *et al.*, 1998). Such biosynthetic combinatorial diversity should prove useful during the lead-discovery phase of pharmaceutical development.

APPENDIX

In Appendix A, K^* is derived from first principles by examining the sum of chemical potentials at chemical equilibrium. Details of our structural model are provided in Appendix B. Appendix C presents the derivation of intermutation pruning (Equation 10). A description of how a bound is computed on a conformation’s minimum energy is presented in Appendix D. Finally, in Appendix E we provide details on the cloning, mutation, expression, and purification of our novel mutant proteins as well as details of the fluorescence quenching experiments.

A. Detailed derivation of K^*

Function K^* represents a biophysically motivated scoring function over molecular ensembles. By using the Boltzmann probability distribution, K^* satisfies the Ergodic hypothesis and can be proved to approximate the true association (binding) constant K_A . If K^* is computed using exact partition functions, then K^* will equal K_A . In practice, we sample conformation space, replacing the continuous integral with a discrete summation, and use a molecular mechanics scoring function to compute the energy of each conformation. Hence, our algorithm represents an approximation to the true association (binding) constant.

We describe K^* for the protein–ligand binding reaction $P + L \rightleftharpoons PL$, where P represents the protein and L can represent either a small molecule in protein–ligand binding or a complete protein in the case of protein–protein binding. Our scoring method represents an approximation to the association constant, K_A , by

$$K^* = \frac{\sum_{b \in B} \exp(-E_b/RT)}{\sum_{l \in L} \exp(-E_l/RT) \sum_{f \in F} \exp(-E_f/RT)}, \quad (11)$$

where B is the set of bound protein states, F is the set of unbound (free) protein states, L is the set of unbound ligand states, E_s is the energy of conformation s , R is the gas constant, and T is the temperature in Kelvin. We will now motivate this equation and describe our physically derived approximation. We first note that for the enzyme/ligand system the true association (binding) constant is defined as

$$K_A = \frac{[PL]}{[P][L]}.$$

It is known in statistical mechanics (Hill, 1956; McQuarrie, 1976) that at chemical equilibrium the sum of the chemical potentials, μ , is equal to zero. In our ligand binding example,

$$\mu_p + \mu_L - \mu_{pL} = 0 \quad (12)$$

where μ_p , μ_L , and μ_{pL} are the chemical potentials for the free protein, free ligand, and protein–ligand complex respectively. The chemical potential, μ_J , for a species J of indistinguishable particles is

$$\mu_J = -kT \ln \left(\frac{q_J(V, T)}{N_J} \right), \quad (13)$$

where k is Boltzmann's constant and $q_J(V, T)$ is the partition function for the N_J molecules of species J at constant volume V and temperature T . The partition function includes all allowable states of a system. By substituting the chemical potentials Equation (13) into the equilibrium condition Equation (12), we obtain the result that at equilibrium

$$\frac{q_{PL}(V, T)}{q_P(V, T)q_L(V, T)} = \frac{N_{PL}}{N_P N_L} = K_A. \quad (14)$$

Thus, the association constant K_A is the quotient of the individual species partition functions. Unfortunately, it is not currently possible to compute exact partition functions for a complex molecular species. This would require integrating an exact energy function over a molecule's entire conformational space. We therefore approximate these partition functions with the use of rotamerically based conformational ensembles. The partition function is approximated by our rotamerically based conformations as:

$$q_{PL} = \sum_{b \in B} \exp(-E_b/RT), \quad q_P = \sum_{f \in F} \exp(-E_f/RT), \quad q_L = \sum_{l \in L} \exp(-E_l/RT), \quad (15)$$

where B , F , and L represent our rotamer-based ensembles for the bound protein–ligand complex, the free protein, and the free ligand conformations, respectively. When combined, this gives the binding constant approximation K^* in Equation (11).

The accuracy with which K^* approximates K_A is proportional to the accuracy of the partition function approximation used. There are two components to an ensemble-based approximation to the partition function: the set of conformations used in the ensemble and the method used to score each conformation. While a molecular ensemble can be generated by multiple techniques (rotamers, multiple NMR structures, multiple crystal structures, molecular dynamics) (Philippopoulos and Lim, 1999; Lovell *et al.*, 2000; Knegt *et al.*, 1997), it is important that the ensemble sample all appropriate regions of conformation space. For example, an ensemble of structures generated from NMR experiments on an apo protein may not sample regions of protein conformation space that are compatible with ligand binding. Rotameric-based ensembles have the potential to sample this space more evenly. When ensembles containing a large number of conformations are used, it is important to choose an energy function that can be computed efficiently.

B. Details of the structural model

Our model consists of a portion of the GrsA-PheA protein (pdb: 1AMU [Conti *et al.*, 1997]) including the active site and a shell of surrounding residues (termed the *steric shell*). The residues of the active site modeled as flexible using rotamers and subject to energy minimization include 235D, 236A, 239W, 278T, 299I, 301A, 322A, 330I, and 331C. The steric shell was selected to include all residues not modeled as flexible and that contain at least one atom within 8 Å of the active site. The steric shell residues include 186Y, 188I, 190T, 210L, 213F, 214F, 230A, 234F, 237S, 238V, 240E, 243M, 279L, 300T, 302G, 303S, 320I, 321N, 323Y, 324G, 325P, 326T, 327E, 328T, 329T, 332A, 333T, 334T, 515N, and 517K. In addition to the active site flexible and steric shell residues, the model also includes the substrate and AMP.

Flexible residues are represented by rotamers from the Lovell *et al.* rotamer library (Lovell *et al.*, 2000). Each rotameric based conformation in B , F , and L is minimized by steepest descent minimization using the AMBER energy function (electrostatic, vdW, and dihedral energy terms) (Weiner *et al.*, 1984; Cornell *et al.*, 1995) and is then combined using Equations (11) and (15) above.

C. Detailed derivation of intermutation pruning

In Section 2.3.2, we described conditions under which a conformation could be pruned when computing a single partition function for a single mutation. When performing a mutation search, we can bootstrap the pruning condition for improved efficiency. As in Section 2.3.2, pruned conformations are not energy-minimized, thereby saving time in the overall mutation search. We will show how, in a mutation search, the ϵ -approximation pruning conditions derived below make use of the partition functions previously computed for other, different mutation sequences evaluated earlier in that search. Therefore, we call this pruning *intermutation pruning*. The intuition is that we assume a lower bound on the partition function q_n^* that allows us to prune more conformations earlier in the search. Intermutation pruning can only be

applied during the computation of a bound partition function, q_{PL} ; the unbound partition function q_p must be computed using the intramutation pruning method of Section 2.3.2.

During a mutation search, our primary goal is to compute a provably accurate ϵ -approximation for the top-ranking mutations while quickly computing K^* values for lower-ranked mutations that do not require the same degree of accuracy. As each mutation is examined in the mutation search, it suffices to compute an ϵ -approximation for only those mutation sequences with $K^* \geq \gamma \max_{i \leq m} K_i^*$, where K_i^* is the K^* value of the i^{th} mutation, m is the total number of mutations, and γ ($0 < \gamma \leq 1.0$) is a user-specified constant defining a range of K^* values that must be computed accurately. Setting $\gamma = 1.0$ will compute an ϵ -approximation for only the best scoring K^* value. Setting $\gamma = 0.0$ will compute an ϵ -approximation for all K^* values. We typically set $\gamma = 0.01$ which causes the approximation algorithm to compute ϵ -approximations for all mutation sequences with K^* scores within two orders of magnitude of the best. Since the value of $\max_{i \leq m} K_i^*$ is not known during the mutation search, we compute ϵ -approximations for all $K^* \geq \gamma K_o^*$, where K_o^* is the largest (best) K^* value seen thus far in the mutation search. By definition, all values of K_o^* satisfy the inequality $K_o^* \leq \max_{i \leq m} K_i^*$ (in other words, all local maxima must be less than or equal to the global maximum). As a result, by computing an ϵ -approximation for all mutations with $K^* \geq \gamma K_o^*$, we will have computed ϵ -approximations for all mutations with $K^* \geq \gamma \max_{i \leq m} K_i^*$. When

$$K^*(1 - \epsilon) \leq \tilde{K}^* \leq K^* \frac{1}{1 - \epsilon}, \quad (16)$$

we say \tilde{K}^* is an ϵ -approximation to K^* . To prove that the computed \tilde{K}^* is an ϵ -approximation to K^* , we first show that intermutation pruning can compute an ϵ -approximation for q_{PL} and then combine the result with the intramutation pruning of Section 2.3.2. The following proof builds from the ideas used in the intramutation pruning of Section 2.3.2. Assume that we've computed q_p (Equation 2) using intramutation pruning and now want to efficiently compute q_{PL} (Equation 2). As stated in the previous paragraph, it is only necessary to compute q_{PL} accurately for mutation sequences with corresponding K^* values that are larger than our minimum accepting score (γK_o^*). That is, we require an ϵ -approximation to q_{PL} when

$$\frac{q_{PL}}{q_p q_L} > \frac{q'_{PL}}{q'_p q'_L} \gamma, \quad (17)$$

where q_{PL} , q_p , and q_L are the partition functions used to compute K^* and q'_{PL} , q'_p , and q'_L are the partition functions used to compute K_o^* . Since we are performing a mutation search to find good mutation sequences for a single ligand, we know $q'_L = q_L$. Therefore, it is only necessary to compute q_{PL} accurately when

$$q_{PL} > \frac{q'_{PL}}{q'_p} \gamma q_p. \quad (18)$$

For notational convenience, we define $K^\dagger = \frac{q_{PL}}{q_p}$ and $K_o^\dagger = \frac{q'_{PL}}{q'_p}$.

Proposition 1. *The algorithm in Fig. 8 computes an ϵ -approximation q_{PL}^* for a bound partition function, q_{PL} , when $q_{PL} > K_o^\dagger \gamma q_p$. If $q_{PL} \leq K_o^\dagger \gamma q_p$, then an $(\epsilon + \delta)$ -approximation ($\delta \geq 0$) is computed.*

Proof. To prove Proposition 1, we consider two cases. Case 1, Equation (18) holds thus requiring an ϵ -approximation. Case 2, Equation (18) does not hold thus requiring only an $(\epsilon + \delta)$ -approximation. We derive pruning criteria for Case 1. The pruning criteria will compute a correct ϵ -approximation for Case 1 and will compute an $(\epsilon + \delta)$ -approximation ($\delta \geq 0$) for Case 2. We will show that the $(\epsilon + \delta)$ -approximation holds only for partition functions falling into Case 2 and that these are situations for which we do not require an ϵ -approximation (see Proposition 2).

After computing an ϵ -approximation to the partition function, it must be the case that $q_n^* \geq (1 - \epsilon)q_n$, which implies that $p_n \leq \epsilon q_n$. If we assume Equation (18) holds (Case 1), then $q_n \geq K_o^\dagger \gamma q_p$, and we can


```

Let  $n \leftarrow$  Number of Rotameric Conformations
Let  $c \leftarrow$  Rotameric Conformations
Initialize:  $\psi \leftarrow \epsilon K_o^\dagger \gamma q_p$ ,  $q^* \leftarrow 0$ ,  $p^* \leftarrow 0$ 
for  $k = 1$  to  $n$ 
   $\psi \leftarrow \text{Max}(\psi, q^* \epsilon)$ 
  if  $B(c_k) \leq -RT \ln(\psi - p^*)$ 
     $q^* \leftarrow q^* + \exp(-\text{ComputeMinEnergy}(c_k)/RT)$ 
  else
     $p^* \leftarrow p^* + \exp(-B(c_k)/RT)$ 
Return  $q^*$ 

```

FIG. 8. Intermutation pruning used in computing the bound partition function q_{PL} . The values q^* and p^* and the functions $B(\cdot)$ and $\text{ComputeMinEnergy}(\cdot)$ are as described in Fig. 4. For all mutations with $q_{PL} \geq K_o^\dagger \gamma q_p$, the computed q^* will represent an ϵ -approximation to the true partition function q such that $q^* \geq (1 - \epsilon)q$.

conservatively conclude that $p_n \leq \epsilon q_n$ if

$$p_n \leq \epsilon K_o^\dagger \gamma q_p. \quad (19)$$

In reality, p_n can be as large as ϵq_n , but during the conformation search we don't yet know the value of q_n . Therefore, given p_k , we can prune conformation, c_{k+1} if p_{k+1} remains less than $\epsilon K_o^\dagger \gamma q_p$ thereby satisfying Equation (19), i.e.,

$$p_{k+1} = p_k + \exp(-B(c_{k+1})/RT) \leq \epsilon K_o^\dagger \gamma q_p. \quad (20)$$

If we solve for $B(c_{k+1})$, then the pruning criterion becomes

$$B(c_{k+1}) \geq -RT \ln(\epsilon K_o^\dagger \gamma q_p - p_k). \quad (21)$$

Because $p_k \leq p_k^*$, Equation (21) can be rewritten as

$$B(c_{k+1}) \geq -RT \ln(\epsilon K_o^\dagger \gamma q_p - p_k^*). \quad (22)$$

Equation (22) is the same as Equation (9) if $(\epsilon K_o^\dagger \gamma q_p)$ is substituted for $(q_k^* \epsilon)$. Therefore, when computing q_{PL} during a mutation search, we use the pruning criterion

$$B(c_{k+1}) \geq -RT \ln(\psi_k - p_k^*), \quad (23)$$

where $\psi_k = \max(\epsilon K_o^\dagger \gamma q_p, q_k^* \epsilon)$ and p_k^* is an upper bound on the partition function of the pruned conformations ($C_k - S_k$) as described in Section 2.3.2.

Lemma 3. *When Equation (18) holds, any conformation c_{k+1} that satisfies Equation (23) can be pruned during computation of the bound partition function while maintaining the invariant that q_k^* is an ϵ -approximation to q_k (where q_k is q_{PL} computed through conformation k and q_k^* is q_{PL}^* computed through conformation k).*

By maintaining the invariant throughout computation of the bound partition function, the following lemma holds:

Lemma 4. *When a bound partition function is computed for a mutation sequence satisfying Equation (18) while maintaining the invariant that q_k^* is an ϵ -approximation to q_k for all k ($0 < k \leq n$), then, by induction, at the end of the computation, q_n^* will be an ϵ -approximation to q_n .*

The intermutation pruning algorithm is shown in Fig. 8.

We have shown that when $q_{PL} > K_o^\dagger \gamma q_P$ (Case 1) pruning using Equation (23) will produce an ϵ -approximation q_{PL}^* . When $q_{PL} \leq K_o^\dagger \gamma q_P$, it is possible that Equation (23) will prune the wrong conformations resulting in $q_{PL}^* < (1 - \epsilon)q_{PL}$. Thus, the computed q_{PL}^* will be too small, and we will not have computed an ϵ -approximation. However, by definition q_{PL} will have been from a mutation whose corresponding $K^\dagger \leq K_o^\dagger \gamma$ ($K^* \leq K_o^* \gamma$) (Case 2), and thus (by Proposition 1) it was not necessary to compute an ϵ -approximation. These partition functions are easily identified by their magnitude. This completes the proof of Proposition 1. ■

We now have the necessary tools to prove Proposition 2.

Proposition 2. *An ϵ -approximation, \tilde{K}^* , is computed for K^* when $K^* > \gamma K_o^*$.*

Proof. We start by determining bounds on the computed K^* for Case 1 and Case 2. In both cases, the unbound partition function q_p^* was computed according to the intramutation pruning of Section 2.3.2, and the bound partition function q_{PL}^* was computed according to the intermutation pruning described above in this section. Note that because each term in the summation of each partition function (Equation 2) is positive, all approximations q_n^* computed by omitting terms must be less than q_n . Thus for both intra- and intermutation pruning, when q_n^* is an ϵ -approximation to q_n , we know that $q_n \geq q_n^* > (1 - \epsilon)q_n$. Therefore,

$$\begin{aligned} q_p &\geq q_p^* > (1 - \epsilon)q_p && \text{and} \\ \text{Case 1) } q_{PL} &\geq q_{PL}^* > (1 - \epsilon)q_{PL} \\ \text{Case 2) } q_{PL} &\geq q_{PL}^* > 0. \end{aligned}$$

The resulting \tilde{K}^\dagger (the approximation to K^\dagger) for high-scoring mutations (Case 1) will fall within the range

$$\left[K^\dagger(1 - \epsilon), K^\dagger \frac{1}{1 - \epsilon} \right],$$

which implies that \tilde{K}^* (the approximation to K^*) lies in the range

$$\left[K^*(1 - \epsilon), K^* \frac{1}{1 - \epsilon} \right],$$

as desired. For lower-scoring mutations (i.e., mutations with $K^* \leq \gamma K_o^*$) (Case 2), the resulting \tilde{K}^\dagger will fall within the range

$$\left[0, K^\dagger \frac{1}{1 - \epsilon} \right],$$

which implies that \tilde{K}^* lies in the range

$$\left[0, K^* \frac{1}{1 - \epsilon} \right].$$

Therefore, an ϵ -approximation \tilde{K}^* is computed for K^* when $K^* > \gamma K_o^*$ and an $(\epsilon + \delta)$ -approximation ($\delta \geq 0$) is computed when $K^* \leq \gamma K_o^*$. This completes the proof of Proposition 2. ■

D. Computing a bound on a conformation's minimum energy

In our structural model, we treat some residues as rigid while others have a rigid backbone but flexible side-chains. If we let R represent the set of all rigid atoms (all atoms of the steric shell and backbone atoms of the flexible residues), then for a system with k flexible residues the energy of this system can be computed as

$$E_0 = A_{00} + \sum_{i \leq k} A_{0i} + \sum_{i \leq k} A_{i0} + \sum_{i \leq k} \sum_{i < j \leq k} A_{ij}, \quad (24)$$

where A is a precomputed residue-indexed pairwise *energy matrix*. A residue-indexed pairwise energy matrix is a $(k + 1) \times (k + 1)$ matrix composed of energy terms describing each residue's interaction energy with the backbone, itself, and all other residues. A detailed description of the matrix elements is as follows: A_{00} is the sum of all energy terms exclusively involving atoms in R , A_{0i} is the sum of all energy terms involving at least one atom of residue i and at least one atom from set R , A_{i0} is the sum of all energy terms involving only atoms of residue i (intraresidue energy), and A_{ij} is the sum of all energy terms involving at least one atom from each of residues i and j .

Similarly, one can define the energy E_m of a minimized conformation, using a different matrix M of pairwise energy terms evaluated on the energy minimized conformation:

$$E_m = M_{00} + \sum_{i \leq k} M_{0i} + \sum_{i \leq k} M_{i0} + \sum_{i \leq k} \sum_{i < j \leq k} M_{ij}, \quad (25)$$

where M_{00} , M_{0i} , M_{i0} , and M_{ij} are the analogues of A_{00} , A_{0i} , A_{i0} , and A_{ij} except that they are computed based on the positions of the atoms in the energy-minimized structure.

During both intra- and intermutation pruning (see Section 2.3.2 and Appendix C), a lower bound on the energy of the energy-minimized conformation is required. If we replace M_{00} , M_{0i} , M_{i0} , and M_{ij} in Equation (25) with lower bounds D_{00} , D_{0i} , D_{i0} , and D_{ij} such that $D_{00} \leq M_{00}$, $D_{i0} \leq M_{i0}$, $D_{0i} \leq M_{0i}$, and $D_{ij} \leq M_{ij}$, then we can compute a bound

$$E_b = D_{00} + \sum_{i \leq k} D_{0i} + \sum_{i \leq k} D_{i0} + \sum_{i \leq k} \sum_{i < j \leq k} D_{ij}, \quad (26)$$

such that $E_b \leq E_m$.

Because E_b is simply the sum of $O(k^2)$ pairwise energy terms, if a precomputed residue-indexed lower-bound pairwise energy matrix is available, then E_b can be computed in time $O(k^2)$. The use of a precomputed residue-based pairwise energy matrix thus avoids the costly computation of $O(a^2)$ energy terms, where a is the total number of atoms in the system and $k \ll a$.

In our implementation of the K^* algorithm, we precompute a residue-indexed lower-bound pairwise energy matrix V over all rotamers for each residue. This matrix contains a lower bound (V_{ij}) on the energy for all allowed pairs of rotamers as well as lower bounds V_{00} , V_{0i} , and V_{i0} on the shell-shell, shell-residue, and residue-self energies (respectively, as described above). The matrix we use in computation is thus slightly different than those described in Equations (24), (25), and (26). The matrix V contains lower-bound energy terms for all rotamers for all residues. Therefore, the matrix has size $(km + 1) \times (km + 1)$ where k is the number of flexibly-modeled residues and m is the number of allowed rotamers (spanning all residue types). When computing an energy bound, terms corresponding to the currently assigned rotamers are used in a manner similar to those described for Equations (24), (25), and (26).

To prevent one rotamer from minimizing into another, the maximum dihedral movement allowed during energy minimization is bounded (as described in *Step 2* of Section 2.2). As a result, one can easily compute the terms of matrix V by examining all pairs of residues in their active site specified relative orientations. The lower-bound energy matrices are precomputed before K^* evaluation or a mutation search is performed.

E. Cloning, mutation, expression, purification, and fluorescence

Cloning. GrsA-PheA was cloned from GrsA by PCR as described previously (Stachelhaus and Marahiel, 1995). PCR reactions were performed using PfuTurbo DNA Polymerase (Stratagene) per manufacturer's

directions. PCR products and pQE-60 were digested with 10U NcoI and 10U BamHI for 2 hrs at 37°C (50 mM Tris-HCl pH 8.0, 10 mM MgCl₂, 100 mM NaCl). PCR products and linearized vector were gel purified and recovered using the QIAquick gel extraction kit (Qiagen), and ligated (2U T4 ligase, 50 mM Tris-HCl pH 7.6, 10 mM MgCl₂, 1 mM ATP, 1 mM DTT, 5% (w/v) PEG-8000, 1 hr at 24°C). *Escherichia coli* M15(pRep4) were transformed with the PheA pQE-60 construct and selected by growth on Luria Broth (LB) supplemented with 50 μg/ml ampicillin and 30 μg/ml kanamycin.

Mutation. Mutations were introduced by site directed mutagenesis using the QuikChange Site-Directed Mutagenesis Kit (Qiagen). Protocols were carried out per manufacturers instructions with the following primers: 301G: CGTTAATTACAGGAGGCTCAGCTACC ($T_m = 72^\circ\text{C}$) 330F: CCTACGGAAACAACCTTTTGTGCGACTACATGG ($T_m = 76^\circ\text{C}$) 330W: GGCCCTACGGAAACAACCTGGTGTGCGACTACATGG ($T_m = 77^\circ\text{C}$).

Expression. 1 Liter of LB was inoculated and grown at 37°C until OD₆₀₀ = 1.5–1.8. IPTG was added (1.5 mM) and cultures grown for an additional 3 hr. Cells were harvested by centrifugation at 2,500 g for 40 min, resuspended in 20 mL LB, pelleted by centrifugation at 2,000 g for 30 min, and frozen at –80°C.

Purification. Cells were thawed, resuspended in buffer (20 mM Tris-HCl pH 7.4, 50 mM NaCl, 50 μM PMSF) and lysed by sonication. Cell extract was clarified by centrifugation at 40,000 g for 40 min and 0.45 μm syringe-driven filtration. PheA-His₆ was purified by Ni²⁺ affinity chromatography using a gradient of 0–100 mM imidazole. Pure PheA-His₆ was dialyzed in 50 mM Hepes pH 7.6.

Fluorescence. Each protein solution (50 mM Hepes pH 7.6) was supplemented with 2 mM dithiothreitol (DTT), 100 mM NaCl, and 10 mM MgCl₂. The excitation wavelength was 280 nm. Substrate was titrated into the protein solution and fluorescence quenching was measured at 340 nm.

ACKNOWLEDGMENTS

We thank Drs. C. Bailey-Kellogg, C. Langmead, L. Wang, T. Lozano-Pérez, R. Mettu and B. Tidor, Mr. A. Yan, Ms. E. Werner-Reiss, and all members of the Donald and Anderson Labs for helpful discussions and comments on drafts. We are grateful to Dr. H. Higgs for use of his ISS PC1 spectrofluorometer.

This work is supported by grants to B.R.D. from the National Institutes of Health (R01 GM-65982), and the National Science Foundation (IIS-9906790, EIA-0102710, EIA-0102712, EIA-9818299, EIA-9802068, and EIA-0305444), and to A.C.A. from the NSF (0133469) and the Research Corporation (RI0857).

REFERENCES

- Belshaw, P., Walsh, C., and Stachelhaus, T. 1999. Aminoacyl-CoAs as probes of condensation domain selectivity in nonribosomal peptide synthesis. *Science* 284, 486–489.
- Böhm, H. 1994. On the use of LUDI to search the fine chemicals directory for ligands of proteins of known three-dimensional structure. *J. Comput. Aided Mol. Des.* 8, 623–632.
- Bolon, D., and Mayo, S. 2001. Enzyme-like proteins by computational design. *Proc. Natl. Acad. Sci. USA* 98, 14274–14279.
- Bouzida, D., Rejto, P., Arthurs, S., Colson, A., Freer, S., Gehlhaar, D., Larson, V., Luty, B., Rose, P., and Verkhivker, G. 1999. Computer simulations of ligand–protein binding with ensembles of protein conformations: A Monte Carlo study of HIV-1 protease binding energy landscapes. *Int. J. Quantum Chem.* 72, 73–84.
- Cane, D., Walsh, C., and Khosla, C. 1998. Harnessing the biosynthetic code: Combinations, permutations, and mutations. *Science* 282, 63–68.
- Carlson, H. 2002. Protein flexibility and drug design: How to hit a moving target. *Curr. Opin. Chem. Biol.* 6, 447–452.
- Challis, G., Ravel, J., and Townsend, C. 2000. Predictive, structure-based model of amino acid recognition by nonribosomal peptide synthetase adenylation domains. *Chem. Biol.* 7, 211–224.
- Claußen, H., Buning, C., Rarey, M., and Lengauer, T. 2001. FlexE: Efficient molecular docking considering protein structure variations. *J. Mol. Biol.* 308, 377–395.
- Conti, E., Stachelhaus, T., Marahiel, M., and Brick, P. 1997. Structural basis for the activation of phenylalanine in the non-ribosomal biosynthesis of Gramicidin S. *EMBO J.* 16, 4174–4183.
- Cornell, W., Cieplak, P., Bayly, C., Gould, I., Merz, K., Ferguson, D., Spellmeyer, D., Fox, T., Caldwell, J., and Kollman, P. 1995. A second generation force field for the simulation of proteins, nucleic acids and organic molecules. *J. Am. Chem. Soc.* 117, 5179–5197.

- Decanniere, K., Transue, T., Desmyter, A., Maes, D., Muyldermans, S., and Wyns, L. 2001. Degenerate interfaces in antigen-antibody complexes. *J. Mol. Biol.* 313, 473–478.
- Desmet, J., Maeyer, M., Hazes, B., and Lasters, I. 1992. The dead-end elimination theorem and its use in protein side-chain positioning. *Nature* 356, 539–542.
- Doekel, S., and Marahiel, M. 2000. Dipeptide formation on engineered hybrid peptide synthetases. *Chem. Biol.* 7, 373–384.
- Ehmann, D., Trauger, J., Stachelhaus, T., and Walsh, C. 2000. Aminoacyl-SNACs as small-molecule substrates for the condensation domains of nonribosomal peptide synthetases. *Chem. Biol.* 7, 765–772.
- Eisenberg, D., and McLachlan, A. 1984. Solvation energy in protein folding and binding. *Nature* 319, 199–203.
- Eppelmann, K., Stachelhaus, T., and Marahiel, M. 2002. Exploitation of the selectivity-conferring code of nonribosomal peptide synthetases for the rational design of novel peptide antibiotics. *Biochemistry* 41, 9718–9726.
- Fauchère, J., and Pliška, V. 1983. Hydrophobic parameters Π of amino-acid side chains from the partitioning of n-acetyl-amino-acid amides. *Eur. J. Med. Chem. Chim. Ther.* 18, 369–375.
- Hellinga, H., and Richards, F. 1991. Construction of new ligand binding sites in proteins of known structure: I. Computer-aided modeling of sites with pre-defined geometry. *J. Mol. Biol.* 222, 763–785.
- Hill, T. 1956. *Statistical Mechanics: Principles and Selected Applications*, 1–17, McGraw-Hill, New York.
- Jaramillo, A., Wernisch, L., Héry, S., and Wodak, S. 2001. Automatic procedures for protein design. *Comb. Chem. High Throughput Screen.* 4, 643–659.
- Jin, W., Kambara, O., Sasakawa, H., Tamura, A., and Takada, S. 2003. De novo design of foldable proteins with smooth folding funnel: Automated negative design and experimental verification. *Structure* 11, 581–591.
- Keating, A., Malashkevich, V., Tidor, B., and Kim, P. 2001. Side-chain repacking calculations for predicting structures and stabilities of heterodimeric coiled coils. *Proc. Natl. Acad. Sci. USA* 98, 14825–14830.
- Knetgel, R., Kuntz, I., and Oshiro, C. 1997. Molecular docking to ensembles of protein structures. *J. Mol. Biol.* 266, 424–440.
- Kuntz, I., Blaney, J., Oatley, S., Langridge, R., and Ferrin, T. 1982. A geometric approach to macromolecular-ligand interactions. *J. Mol. Biol.* 161, 269–288.
- Lasters, I., and Desmet, J. 1993. The fuzzy-end elimination theorem: Correctly implementing the side chain placement algorithm based on the dead-end elimination theorem. *Protein Eng.* 6, 717–722.
- Lilien, R., Sridharan, M., Huang, J., Bushweller, J., and Donald, B. 2000. Computational screening studies for core binding factor beta: Use of multiple conformations to model receptor flexibility. *Poster—8th Int. Conf. on Intelligent Systems for Molecular Biology*, ISMB 2000.
- Linne, U., Doekel, S., and Marahiel, M. 2001. Portability of epimerization domain and role of peptidyl carrier protein on epimerization activity in nonribosomal peptide synthetases. *Biochemistry* 40, 15824–15834.
- Looger, L., Dwyer, M., Smith, J., and Hellinga, H. 2003. Computational design of receptor and sensor proteins with novel functions. *Nature* 423, 185–190.
- Lorber, D., and Shoichet, B. 1998. Flexible ligand docking using conformational ensembles. *Protein Sci.* 7, 938–950.
- Lovell, S., Word, J., Richardson, J., and Richardson, D. 2000. The penultimate rotamer library. *Proteins* 40, 389–408.
- Marvin, J., and Hellinga, H. 2001. Conversion of a maltose receptor into a zinc biosensor by computational design. *PNAS* 98, 4955–4960.
- McQuarrie, D. 1976. *Statistical Mechanics*, Harper and Row, New York.
- Montfort, W., Perry, K., Fauman, E., Finer-Moore, J., Maley, G., Hardy, L., Maley, F., and Stroud, R. 1990. Structure, multiple site binding, and segmental accommodation in thymidylate synthase on binding dUMP and an anti-folate. *Biochemistry* 29, 6964–6977.
- Mootz, H., Schwarzer, D., and Marahiel, M. 2000. Construction of hybrid peptide synthetases by module and domain fusions. *Proc. Natl. Acad. Sci. USA* 97, 5848–5853.
- Murthy, K., Winborne, E., Minnich, M., Culp, J., and Debouck, C. 1992. The crystal structures at 2.2-Å resolution of hydroxyethylene-based inhibitors bound to human immunodeficiency virus type 1 protease show the inhibitors are present in two distinct orientations. *J. Biol. Chem.* 267, 22770–22778.
- Offredi, F., Dubail, F., Kischel, P., Sarinski, K., Stern, A., Van de Weerd, C., Hoch, J., Prospero, C., François, J., Mayo, S., and Martial, J. 2003. De novo backbone and sequence design of an idealized α/β -barrel protein: Evidence of stable tertiary structure. *J. Mol. Biol.* 325, 163–174.
- Österberg, F., Morris, G., Sanner, M., Olson, A., and Goodsell, D. 2002. Automated docking to multiple target structures: Incorporation of protein mobility and structural water heterogeneity in autodock. *Proteins* 46, 34–40.
- Philippopoulos, M., and Lim, C. 1999. Exploring the dynamic information content of a protein NMR structure: Comparison of a molecular dynamics simulation with the NMR and X-Ray structures of *Escherichia coli* ribonuclease HI. *Proteins* 36, 87–110.
- Pierce, N., Spriet, J., Desmet, J., and Mayo, S. 2000. Conformational splitting: A more powerful criterion for dead-end elimination. *J. Comput. Chem.* 21, 999–1009.
- Pierce, N., and Winfree, E. 2002. Protein design is NP-hard. *Protein Eng.* 15, 779–782.

- Ponder, J., and Richards, F. 1987. Tertiary templates for proteins: Use of packing criteria in the enumeration of allowed sequences for different structural classes. *J. Mol. Biol.* 193, 775–791.
- Raag, R., and Poulos, L. 1991. Crystal structures of cytochrome P-450CAM complexed with camphane, thiocamphor, and adamantane: Factors controlling P-450 substrate hydroxylation. *Biochemistry* 30, 2674–2684.
- Rienstra, C., Tucker-Kellogg, L., Jaroniec, C., Hohwy, M., Reif, B., McMahon, M., Tidor, B., Lozano-Pérez, T., and Griffin, R. 2002. De novo determination of peptide structure with solid-state magic-angle spinning NMR spectroscopy. *Proc. Natl. Acad. Sci. USA* 99, 10260–10265.
- Samanta, U., Pal, D., and Chakrabarti, P. 1999. Packing of aromatic rings against tryptophan residues in proteins. *Acta Cryst. D* 55, 1421–1427.
- Schneider, A., Stachelhaus, T., and Marahiel, M. 1998. Targeted alteration of the substrate specificity of peptide synthetases by rational module swapping. *Mol. Gen. Genet.* 257, 308–318.
- Schwarzer, D., Finking, R., and Marahiel, M. 2003. Nonribosomal peptides: From genes to products. *Nat. Prod. Rep.* 20, 275–287.
- Shifman, J., and Mayo, S. 2002. Modulating calmodulin binding specificity through computational protein design. *J. Mol. Biol.* 323, 417–423.
- Stachelhaus, T., and Marahiel, M. 1995. Modular structure of peptide synthetases revealed by dissection of the multifunctional enzyme GrsA. *J. Biol. Chem.* 270, 6163–6169.
- Stachelhaus, T., Mootz, H., and Marahiel, M. 1999. The specificity-conferring code of adenylation domains in nonribosomal peptide synthetases. *Chem. Biol.* 6, 493–505.
- Stachelhaus, T., Schneider, A., and Marahiel, M. 1995. Rational design of peptide antibiotics by targeted replacement of bacterial and fungal domains. *Science* 269, 69–72.
- Street, A., and Mayo, S. 1999. Computational protein design. *Structure* 7, R105–R109.
- Tucker-Kellogg, L. 2002. *Systematic Conformational Search with Constraint Satisfaction*. Ph.D. thesis, Massachusetts Institute of Technology.
- Weber, T., Baumgartner, R., Renner, C., Marahiel, M., and Holak, T. 2000. Solution structure of PCP, a prototype for the peptidyl carrier domains of modular peptide synthetases. *Structure* 8, 407–418.
- Weiner, S., Kollman, P., Case, D., Singh, U., Ghio, C., Alagona, G., Profeta, S., and Weiner, P. 1984. A new force field for molecular mechanical simulation of nucleic acids and proteins. *J. Am. Chem. Soc.* 106, 765–784.
- Wojtczak, J., Cody, V., Luft, J., and Pangborn, W. 1996. Structures of human transthyretin complexed with thyroxine at 2.0 Å resolution and 3', 5'-Dinitro-*N*-acetyl-L-thyronine at 2.2 Å resolution. *Acta Cryst. D* 52, 758–765.
- Zhou, G., Ferrer, M., Chopra, R., Kapoor, T., Strassmaier, T., Weissenhorn, W., Skehel, J., Oprian, D., Schreiber, S., Harrison, S., and Wiley, D. 2000. The structure of an HIV-1 specific cell entry inhibitor in complex with the HIV-1 gp41 trimeric core. *Bioorg. Med. Chem.* 8, 2219–2228.

Address correspondence to:

Bruce R. Donald

or Amy C. Anderson

Department of Computer Science

Department of Chemistry

Dartmouth College

6211 Sudikoff Laboratory

Hanover, NH 03755-3510

E-mail: Bruce.R.Donald@dartmouth.edu, Amy.C.Anderson@dartmouth.edu