


Compact Representation of Continuous Energy Surfaces for More Efficient Protein Design

Mark A. Hallen,^{†,§} Pablo Gainza,[†] and Bruce R. Donald^{*,†,‡,§}[†]Department of Computer Science and [‡]Department of Chemistry, Duke University, Durham, North Carolina 27708, United States[§]Department of Biochemistry, Duke University Medical Center, Durham, North Carolina 27710, United States

ABSTRACT: In macromolecular design, conformational energies are sensitive to small changes in atom coordinates; thus, modeling the small, continuous motions of atoms around low-energy wells confers a substantial advantage in structural accuracy. However, modeling these motions comes at the cost of a very large number of energy function calls, which form the bottleneck in the design calculations. In this work, we remove this bottleneck by consolidating all conformational energy evaluations into the pre-computation of a local polynomial expansion of the energy about the “ideal” conformation for each low-energy, “rotameric” state of each residue pair. This expansion is called “energy as polynomials in internal coordinates” (EPIC), where the internal coordinates can be side-chain dihedrals, backbone angles, and/or any other continuous degrees of freedom of a macromolecule, and any energy function can be used without adding any asymptotic complexity to the design. We demonstrate that EPIC efficiently represents the energy surface for both molecular-mechanics and quantum-mechanical energy functions, and apply it specifically to protein design for modeling both side chain and backbone degrees of freedom.



1. INTRODUCTION

Computational design algorithms are an effective approach to engineering proteins and discovering new drugs for many biomedically relevant challenges, such as drug resistance prediction,¹ peptide-inhibitor design,² and enzyme design.³ Protein design algorithms search through large sequence and conformational spaces for sequences that will fold to a desired structure and perform a specific function. One of the key challenges in protein design is modeling and searching the many continuous conformational degrees of freedom inherent in proteins and other molecules. Protein design algorithms must estimate optimal values for all these degrees of freedom to optimize the sequence of the protein and to optimize the chemical structure of the ligand if used for drug design. Molecular dynamics simulations can be used for this purpose if the protein sequence and ligand are known, because they can move all of the molecule's degrees of freedom,⁴ but these simulations are computationally expensive and must be run separately for each sequence or ligand chemical structure. Hence, this direct simulation approach is unsuitable for searching large combinatorial design spaces. For example, many protein design problems require searching over trillions of sequences—far too many for individual molecular dynamics runs.

To address the combinatorially large sequence spaces inherent to protein design, dedicated protein design algorithms efficiently choose an amino-acid type and conformation for each residue in a protein that together minimize some energy functions.⁵ Because the side chain conformations of each amino-acid type are generally found in clusters, known as rotamers,⁶ the protein design problem has often been treated as a discrete

optimization problem. In this case, the output is a set of rotamer assignments (a rotamer, including amino-acid type, is assigned to each residue). The objective function is an energy function, which maps conformations to their energies. However, because proteins are continuously flexible and have backbone as well as side chain flexibility, some of the protein's internal coordinates will likely have functionally significant variations from the rotamer's “ideal” value (at the center of the cluster). Clashing ideal rotamers can often be converted to favorable conformations by relatively small adjustments in the side chain conformations.^{7,8} Small adjustments in the backbone conformation away from the wild-type backbone can also be functionally significant.^{9–11} As a result, modeling of continuous flexibility has been shown to dramatically improve the accuracy of structural modeling in designs,^{8,11} even using a limited set of degrees of freedom, and has led to designs that perform well experimentally.^{1–3,5,7,12} Furthermore, attempts to mimic this effect by discrete sampling at finer resolution have been shown to either poorly approximate the continuous solutions or be computationally prohibitive.⁸ Modeling additional continuous degrees of freedom, with the goal of modeling all conformational variations that significantly impact protein function, is expected to further increase the accuracy of the designs.

Modeling of continuous flexibility in protein design can still exploit our knowledge of rotamers, because rotamers provide an excellent prior estimate of where “energy wells” are likely to be in the conformational space of the protein. Residues' side

Received: November 19, 2014

Published: March 20, 2015

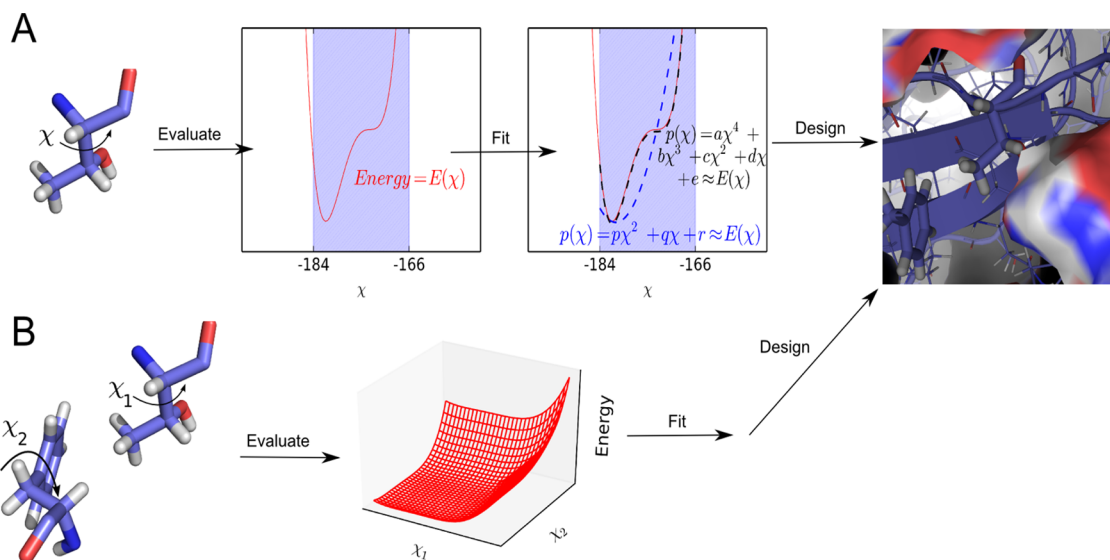


Figure 1. (A) The energy of each residue is represented by EPIC as a polynomial in the internal coordinates, such as side chain dihedrals χ . Low-degree, inexpensive polynomials (blue) are tried first, and the degree is increased as needed to achieve a good fit (black) to the actual energy function (red). These polynomials are then used for the design in place of the full energy function. (B) Interactions between pairs of residues are represented in terms of both residues' internal coordinates.

chains will usually be found in the region of conformational space fairly close (e.g., within 10–20° for side chain dihedrals) to an ideal rotamer, even with a relatively small rotamer library.¹³ As a result, if by using a “minimization-aware” search process one can find the nearest ideal rotameric conformation to the true global minimum-energy conformation (GMEC) of a protein, the GMEC itself can generally be found by local minimization initialized to that ideal rotameric conformation. Thus, protein design can fully account for continuous side chain flexibility while still functioning as a “minimization-aware” search⁷ over a discrete rotamer space. This same paradigm can be extended to continuous backbone flexibility if ideal conformations that include backbone motions, known as “residue conformations” (RCs),¹¹ are included in the search.

Such “minimization-aware” searches can take multiple forms. For example, the iMinDEE algorithm⁸ produces a gap-free, provably accurate list of rotamer assignments in order of lower bound on minimized energy. iMinDEE performs energy minimization on each of these rotamer assignments in that gap-free order until the lower bound exceeds the best minimized energy E_b enumerated to that point. At this point, any subsequently enumerated assignments would be guaranteed to have higher minimized energies than E_b ; thus, E_b is provably the global minimum energy. iMinDEE enumerates rotamer assignments efficiently using the A* search algorithm.^{14,15} A Monte Carlo search over rotamer space can also incorporate minimization^{16,17} but without any provable guarantees. A Monte Carlo search will be minimization-aware if continuous minimization is performed for sequences and conformations during the search, and the minimized energy is used in the calculation of acceptance probabilities for new rotamer assignments, as in ref 16 and the final phase of ref 17. Nevertheless, a method without provable guarantees will likely require more conformations and sequences to be minimized to obtain the same gain in accuracy as iMinDEE, because unlike in iMinDEE, the conformations being minimized are not guaranteed to be the most promising ones. Furthermore, there is no finite number N for a given protein design problem such

that enumerating N conformations by Monte Carlo is guaranteed to yield the GMEC.

Any minimization-aware method, however, will require a large number of subroutine calls to local minimization. Continuous energy minimization is computationally expensive, even with molecular mechanics-type energy functions that prioritize speed over accuracy. This causes the minimization of the energy function to be the bottleneck in protein design with continuous flexibility.

This bottleneck becomes more severe when more sophisticated energy functions are introduced. Computational protein design is typically performed with energy functions that prioritize speed over accuracy. For example, they typically use simplified implicit solvation models, such as EEF1.¹⁸ Vizcarra et al. have investigated the use of the Poisson-Boltzmann model, a much more accurate implicit solvation model, in protein design.¹⁹ They found it to be amenable to representation as a sum of residue-pair interactions, which is the form required for most protein design algorithms, but orders of magnitude more expensive than EEF1. Other methods of improving energy function accuracy are likely to face the same problem. For example, quantitatively accurate descriptions of most molecular interactions require computation of the electronic structure using quantum chemistry, but the methods that can do this are very computationally intensive.²⁰ Methods to reduce the number of calls to an energy function needed in protein design could allow more accurate energy functions to be used and thus yield more accurate results.

In protein design with only discrete flexibility, precomputation methods are typically used to reduce the number of energy function calls needed, that is, the number of conformations for which the energy must be evaluated. Before the design is started, the interaction energy of each pair of ideal, rigid rotamers at different residue positions is precomputed and stored in an energy matrix. Then, an ideal rotamer is chosen for each residue based on the energies in this matrix, and no further calls to the energy function are needed during the actual design calculation. The number of energy function calls required is

thus quadratic in the number of residues in the system (that is, it scales as the number of pairs of residues). A precomputed energy matrix is of limited use, however, if we want to model continuous flexibility. No benefit in design is gained by performing *post hoc* minimization on the best conformation found using ideal rotamers.⁸ This is true even if a high degree of flexibility is used for minimization, e.g., if molecular dynamics techniques are used, because the designed sequence is already determined before minimization is performed. In contrast, a minimization-aware search performs local continuous minimization for all rotamer assignments that might be optimal to find the true GMEC.⁷

Thus, an analogous energy matrix precomputation method for continuous flexibility would be very useful. It would ensure a polynomial number of energy function calls for minimization-aware protein design, in contrast to the exponential number of calls that may arise in minimization of all possibly optimal rotamer assignments (because the number of such assignments may be exponential with respect to the number of residues modeled). The search for rotamer assignments itself is unlikely to admit a polynomial-time algorithm, because it is NP-hard even to approximate.^{21,22} But a method to precompute pairwise energies for continuously flexible design would change the overall time cost from

$$\frac{\text{(a large rotamer search cost)} \times \text{(the energy function cost)}}{\text{to}} \\ \text{(a large rotamer search cost)} + \text{(the energy function cost)}.$$

The rotamer search cost will necessarily be exponential in the worst case if one wants to obtain the GMEC or an approximation to the GMEC within a fixed error threshold. But the energy function cost will merely be quadratic in the number of residues, indicating that the pairwise energy precomputation shifts the bottleneck away from the energy function calls. This brings the same improvement to minimization-aware designs that energy matrix precomputation brought to non-continuously flexible designs.

We now present a pairwise energy precomputation method that admits continuous flexibility: energy as polynomials in internal coordinates (EPIC). EPIC computes a representation of the pairwise energy for each rotamer pair, not just at the rotamers' ideal values of the internal coordinates but also for values within specified ranges around the ideal values (Figure 1). This computation is performed before the rotamer search computation is begun. This allows the rotamer search to substitute the new, quickly evaluable representation for the original energy function. EPIC is implemented in the OSPREY^{7,23,24} open-source protein design package, which has yielded many designs that performed well experimentally *in vitro*,^{1–3,25–28} *in vivo*,^{1,2,25,26} and in nonhuman primates.²⁵ EPIC provides a significant speedup when used with OSPREY's default AMBER^{29,30} and EEF1¹⁸-based energy function, but it is also shown to be suitable for representing quantum-mechanical energies.

This paper makes the following contributions:

1. A compact, closed-form representation of energy as a function of continuous internal coordinates of a protein system.
2. A modified least-squares method to compute this representation.
3. A modified implementation of the iMinDEE⁸ and DEEPer¹¹ protein design algorithms integrated into the OSPREY^{1,3,7,23,24} open-source protein design package that makes use of this representation to achieve substantial speedups. Our implementation is available for free download at <http://www.cs.duke.edu/donaldlab/osprey.php>.

4. Computational experiments showing that compact and accurate EPIC representations are possible both for the standard energy function in OSPREY and for energies obtained by quantum chemistry at the SCF and MP2 levels of theory.

5. Computational experiments showing that EPIC greatly accelerates minimization-aware protein design calculations, thus allowing designs to include not only more flexible residues but also more conformational flexibility at those residues.

2. METHODS

2.1. Preliminaries.

EPIC, like most previous protein design algorithms, is designed for pairwise energy functions. Pairwise energy functions are sums of intra+shell and pairwise terms. Intra+shell terms are functions of the amino-acid type and conformation of one residue, and pairwise terms are functions of the amino-acid types and conformations of two residues. Each pairwise term represents the interaction between a pair of flexible residues, and each intra+shell term represents the internal energy of a residue plus its interactions with nonflexible "shell" residues (those that are frozen in a single, fixed conformation throughout the entire calculation). EPIC could be easily modified to include some higher-order terms for defined combinations of more than two residues: these terms can also be represented as polynomials in their residues' degrees of freedom.

To find the GMEC, we must find an amino-acid type and conformation for each flexible residue such that the sum of the intra+shell terms for all flexible residues, plus the sum of pairwise terms for all pairs of flexible residues, is minimized. This problem is referred to as "conformational search". Conformational search can also comprise a sequence search by searching for the best conformation across many sequences' conformational spaces. Many algorithms are available for this problem, including iMinDEE⁸ and DEEPer,¹¹ which solve the problem with provable accuracy. EPIC can be used in any conformational search algorithm that models continuous flexibility because it provides polynomials that can be directly substituted for intra+shell and pairwise terms of the energy function.

The essence of EPIC is to exploit the fact that for each pairwise or intra+shell energy, the energy in the vicinity of the minimum can be described well by a relatively low-degree polynomial (usually quadratic total degree, sometimes higher; see Figure 5B). This description is computed using a modified least-squares method. We will refer to the states to which a residue may be assigned as residue conformations (RCs, cf. DEEPer¹¹). In the absence of backbone flexibility, each RC will correspond to a side chain rotamer. Within a residue conformation, the residue's continuous energy variations can be described by a set of internal coordinates, which are subject to box constraints (i.e., bounds on each internal coordinate).

Herein, the word "polynomial" will be used in two very different senses in the description of EPIC below. First, EPIC is a polynomial representation of the energy, namely, a polynomial function with respect to the internal coordinates that is explicitly constructed by the EPIC algorithm. Second, a measure m of the computational complexity of an algorithm can be described as "polynomial"^{5,31} if, for input size n , m grows no faster than n^d for a fixed exponent d . In this case, one can construct a polynomial with respect to the size of the input that will be an upper bound on the computational cost, no matter how large the input is (n). For this purpose, we consider either the time or the number of energy function calls as m —these

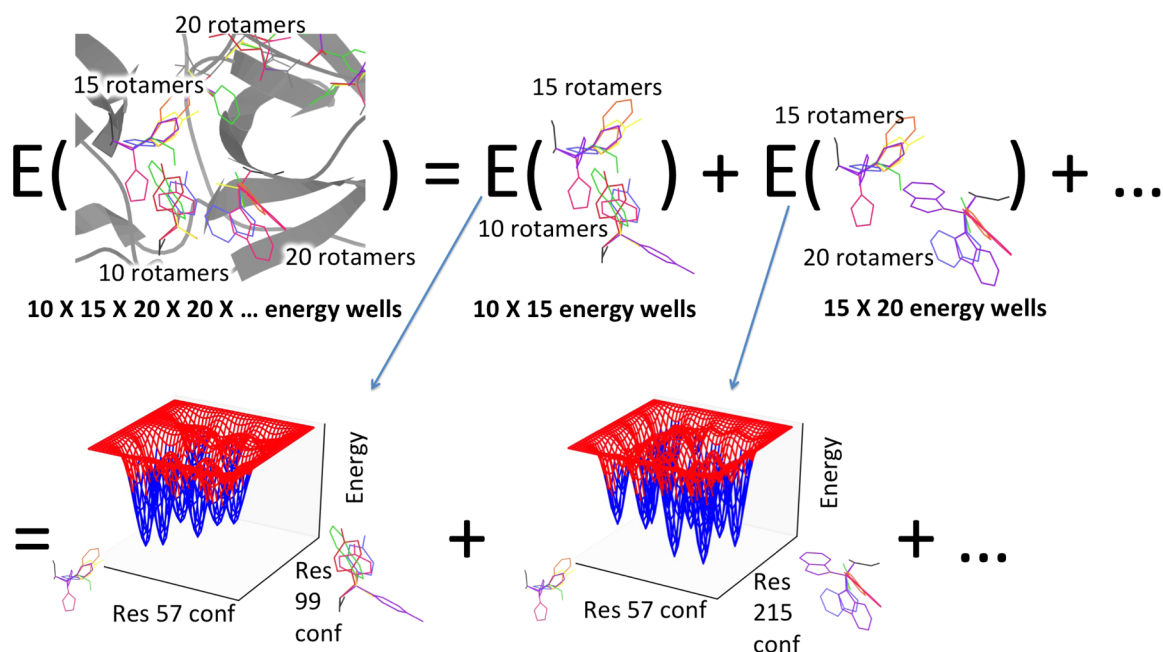


Figure 2. The number of energy wells in a protein system scales exponentially with the number of flexible residues, leading to an exponential number of energy function calls. EPIC can replace most of these calls with quick evaluations of low-degree polynomials. (Top) A protein may have an energy well for every combination of rotamers (rainbow) at different residues. The global minimum-energy conformation (GMEC) of a protein may be in any of these wells. We model the energy as a sum of pairwise energy terms. Each pairwise term will have wells for pairs of rotamers, but there are far fewer wells of this kind—a number quadratic in the number of residues. We can easily afford the energy function calls needed to characterize each pairwise well. (Bottom) By precomputing a polynomial representation (blue) of the energy within each well of each pairwise term (red), we enable the computation of any pairwise term in any pairwise well and thus of the full protein energy in any energy well of the protein solely by a quick evaluation of polynomials.

two measures of computational complexity are related to each other by a constant factor in current protein design algorithms with continuous flexibility because energy function calls generally dominate the cost of these algorithms. For example, precomputation of an EPIC representation must be performed only once for every pair of RCs at different residues, and thus the total number of polynomial fits (and thus the total time for precomputation of EPIC fits) does not grow faster than the square of the number of residues being modeled. On the other hand, protein design itself has been shown to be NP-hard,^{21,22} which means no polynomial-time algorithm is likely to exist for it. In other words, for every polynomial $p(n_r)$ in the number n_r of residues, there are protein design problems of n_r residues that are not expected to be solvable in time $p(n_r)$. A problem only solvable by exponential-time algorithms (i.e., those that take time scaling as b^n , where b is a constant and n is the size of the input) would typically be considered NP-hard.

Energy function calls are typically the bottleneck in protein design algorithms that model continuous flexibility. EPIC, however, ensures that the number of energy function calls in a protein design calculation is linear in the number of RC pairs and thus polynomial in the size of the input (Figure 2).

2.2. Basic Least-Squares Method. Consider two interacting residues i and j . Let us start with the “well-behaved” case where there exists a low-degree polynomial representation of a pairwise energy throughout the allowed ranges of both residues’ internal coordinates (of the form in eq 1).

We employ the notation introduced in the DEEPer algorithm.¹¹ Suppose we have RCs i_r and j_s with pairwise energy $E(i_r, j_s, \mathbf{x})$, where \mathbf{x} is the vector of internal coordinates (for example, dihedrals) affecting residues i and j when they

have amino-acid types corresponding to i_r and j_s . Let $E_{\ominus}(i_r, j_s) = E(i_r, j_s, \mathbf{x}_0(i_r, j_s)) = \min_{\mathbf{x}} E(i_r, j_s, \mathbf{x})$, where the minimum is taken with respect to the internal coordinates over their allowed ranges for the current RCs, and $\mathbf{x}_0(i_r, j_s)$ is the set of internal coordinates that minimizes the pairwise energy. This definition of E_{\ominus} is consistent with iMinDEE⁸ and DEEPer.¹¹

We seek a multivariate polynomial $p_{i_r, j_s}(\mathbf{x})$ such that

$$E_{\ominus}(i_r, j_s) + p_{i_r, j_s}(\mathbf{x} - \mathbf{x}_0(i_r, j_s)) \quad (1)$$

is a good approximation to $E(i_r, j_s, \mathbf{x})$. This multivariate polynomial is approximately a finite, low-degree Taylor expansion about the minimum. However, we use least-squares fits because we have found that they perform much better than Taylor expansions that are based on numerical derivatives. The fits are performed using a training set with ten times as many samples as there are parameters (polynomial coefficients) in the fit; the sampling procedure is described in section 2.7. The fits are cross-validated with an independent set of samples (section 2.6). The constraint $p(\mathbf{0}) = 0$ is applied, so the real energy and the polynomial will agree exactly at the minimum-energy point. This constraint is easily implemented by not including a constant term in the polynomial and reflects the need for the highest accuracy to be attained for the lowest-energy and thus the most biophysically reasonable conformations. As a result of this constraint, all values of the polynomial on its domain will be non-negative. Fitting begins with a multivariate quadratic fit and then moves to higher degrees as needed (see section 2.6). Because polynomials are linear with respect to their coefficients, the fitting is a linear least-squares problem.

This method can be generalized without modification to intra+shell energies as well as to any continuous degrees of freedom, such as newly modeled backbone perturbations¹¹ or rigid-body motions of ligands. In every case, the number of variables for the polynomial will be the number of continuous degrees of freedom that define the conformation of the residue or residue pair of interest. For example, in a pairwise energy computation for a rotamer of lysine and a rotamer of valine with only side chain flexibility, the polynomial will have five variables (the four dihedrals of lysine and one dihedral of valine). The polynomial coefficients are real numbers.

Let \mathbf{r} be an RC assignment, represented as a tuple of RCs with one RC for each residue. Let i_r be the rotamer in \mathbf{r} at residue i . To approximate the minimized energy of an enumerated conformation \mathbf{r} , instead of minimizing the full energy

$$E_{\mathbf{r}}(\mathbf{x}) = \sum_i E(i_r, \mathbf{x}) + \sum_{j<i} E(i_r, j_r, \mathbf{x}) \quad (2)$$

with respect to the system's continuous degrees of freedom \mathbf{x} , we simply minimize the polynomial approximation

$$q_{\mathbf{r}}(\mathbf{x}) = \sum_i E_{\ominus}(i_r) + p_{i_r}(\mathbf{x} - \mathbf{x}_0(i_r)) + \sum_{j<i} E_{\ominus}(i_r, j_r) + p_{i_r, j_r}(\mathbf{x} - \mathbf{x}_0(i_r, j_r)) \quad (3)$$

with respect to \mathbf{x} . These least-squares approximations achieve high accuracy for the low-energy wells of rotamers and local backbone motions (i.e., the portions of conformational space where both the continuous degrees of freedom and the energy are relatively close to the local minimum of the pairwise energy). Higher energies may also be found close in conformational space to the local minimum, but these energies indicate strained conformations unlikely to be seen in nature. Thus, for a “well-behaved” energy term whose energy is unstrained throughout the bounds on continuous degrees of freedom that define our current RCs, EPIC simply performs a least-squares fit of the energy to represent it as a multivariate polynomial with respect to the continuous degrees of freedom.

Many RCs do however contain both regions with feasible energies and regions with higher energies that represent biophysically inaccessible conformations, such as steric clashes. These RCs present difficulties for the basic least-squares fit, but the following algorithmic modification avoids this problem.

2.3. Modified Least-Squares Method. To handle RCs with high-energy regions, we note that we do not necessarily need the polynomial to be a good approximation for the energy throughout the entire region allowed by the box constraints. We merely require that eq 2 be a good approximation for eq 3 when used with biophysically feasible, minimized values of \mathbf{x} . In particular, we can expect that the optimal, minimized structure has no clashes or other particularly large local strains. We have the advantage that even though interaction energies in proteins can rise steeply toward infinity in the case of steric clashes, there are no physical phenomena that will cause interaction energies to decrease steeply toward negative infinity. Thus, energies are relatively well-behaved in low-energy regions. We can thus effectively partition the conformational space into relatively smooth, low-energy regions that we approximate accurately and high-energy regions that can be ruled out.

Let us denote the energy relative to the minimum as $E'(i_r, j_s, \mathbf{x} - \mathbf{x}_0(i_r, j_s)) = E(i_r, j_s, \mathbf{x}) - E_{\ominus}(i_r, j_s)$ in the pairwise case or $E'(i_r, \mathbf{x} - \mathbf{x}_0(i_r)) = E(i_r, \mathbf{x}) - E_{\ominus}(i_r)$ in the intra

+shell case. Our requirements for a “good approximation” of the energy can be defined rigorously in terms of two upper bounds b_1 and b_2 that we place on E' . For each intra+shell or pairwise energy term, we estimate an upper bound b_1 on E' that we expect to hold for all minimized conformations that we want to output (the GMEC, or the lowest-energy c conformations if we are computing a c -conformation ensemble). The algorithm will be able to determine if b_1 is valid or not, so we can try again with a higher b_1 if necessary. Additionally, we need a second, possibly looser, upper bound b_2 on E' that we are confident will be valid for all minimized conformations that we compute during our search, whether they turn out to be the GMEC or not. The value of b_2 must be the same for all intra+shell and pairwise terms (b_1 can be term-specific, though in practice a single value for b_1 is convenient).

If EPIC is being used with the iMinDEE algorithm for the conformational search,⁸ we can provably obtain the GMEC without considering any conformations whose energies E' exceed the pruning interval,⁸ an upper bound computed by iMinDEE for the difference between the lowest conformational energy lower bound (based on pairwise minimum energies) and the GMEC. Thus, when running iMinDEE, we can set b_2 equal to the iMinDEE pruning interval. When b_2 is set equal to the pruning interval, we know it is a valid upper bound on E' for all minimized conformations computed during the search, and thus our GMEC calculation is provable. We can also do this when running DEEPPer,¹¹ which is essentially a backbone-flexible version of iMinDEE. For other algorithms, we may want to set b_2 based on knowledge of the system being designed. Setting $b_2 = 2b_1$ is likely to be an acceptable heuristic.

Our polynomial only needs to be a good fit to E' for values of the internal coordinates where $E' \leq b_1$. For $E' > b_1$, we will require that the polynomial lie above b_1 . This will ensure that when we enumerate conformations in order of minimized energy computed using polynomials, as long as the thresholds b_1 are chosen correctly, we will obtain nonclashing conformations before conformations with clashes, and these nonclashing conformations' energies will be accurately represented by the polynomial fits. Furthermore, we will require that for $b_1 < E' < b_2$ the polynomial should be a lower bound on E' (Figure 3). This will ensure that regardless of what thresholds b_1 were used, we never overestimate a conformational energy that is below the threshold b_2 and thus never exclude it from the enumerated list of conformations. The requirement to be a lower bound is easy to satisfy because clashing van der Waals interactions are very steep and thus tend to rise much more quickly than the polynomial fits. Thus, when we perform polynomial fits using thresholds, we know we will be getting a gap-free list of conformations in order of energy. If the thresholds b_1 were chosen to be too low, some higher-energy conformations with underestimated energy might be included as well, but these will be limited to minimized conformations containing energy terms with $E' > b_1$. This condition can easily be checked. If desired, the run can be redone with increased b_1 thresholds to eliminate this error. Thus, the choice of b_1 affects the ultimate speed of the algorithm but not its correctness.

For our experiments in this work (section 3), we have set b_1 to 10 kcal/mol. This threshold was found to be sufficient for all experiments described in this work as well as most other EPIC designs that we have tried. Physically, any pair of residues whose interaction energy is 10 kcal/mol worse than the optimal interaction for its current RC pair is likely in a highly strained conformation, such as a steric clash. Thus, a design requiring b_1




Figure 3. (A) For each energy value E' , there is a range of “ideal” values for the EPIC fit (green). For energies below cutoff b_1 , which may be found in favorable conformations, this range is just the energy (the range has zero width). For higher energies, the range is defined using cutoffs b_1 and b_2 . For fitting purposes, EPIC fit values are penalized by the amount they lie outside the ideal range (the purple point represents a sample conformation for a given EPIC fit incurring the penalty indicated in red). (B) Example of curves satisfying these conditions. The EPIC fit closely matches the energy up to cutoff b_1 , after which it deviates from the energy but stays in the target region shown in A by staying below the energy. Once the energy is over b_2 , the EPIC fit can be either above or below the true energy without leaving the target region.

greater than 10 kcal/mol is likely to be biologically infeasible. For example, the protein is likely to unfold or undergo a large and unexpected structural change rather than suffer this local strain.

Let us use \mathbf{z} to denote a vector in the domain of our polynomial fit p , which is considered a good representation of the energy if, for some small $\varepsilon > 0$, the following conditions are satisfied:

1. For \mathbf{z} such that $E'(i_r, j_s, \mathbf{z}) \leq b_1$, $|p_{i_r, j_s}(\mathbf{z}) - E'(i_r, j_s, \mathbf{z})| < \varepsilon$
2. For \mathbf{z} such that $b_1 < E'(i_r, j_s, \mathbf{z}) < b_2$, $b_1 - \varepsilon < p_{i_r, j_s}(\mathbf{z}) < E'(i_r, j_s, \mathbf{z}) + \varepsilon$
3. For \mathbf{z} such that $b_2 \leq E'(i_r, j_s, \mathbf{z})$, $b_1 - \varepsilon < p_{i_r, j_s}(\mathbf{z})$

These conditions are illustrated in Figure 3. They can be achieved using a modified least-squares fit, using special “one-sided” penalties to enforce the inequalities in conditions 2 and 3 along with usual (two-sided) least-squares penalties to enforce condition 1. The objective function is the sum of terms from each sample in the training set. For a sample \mathbf{z} such that $E'(i_r, j_s, \mathbf{z}) \leq b_1$, the objective function term is $(p_{i_r, j_s}(\mathbf{z}) - E'(i_r, j_s, \mathbf{z}))^2$ (as is typical for least-squares). A term of this form is also used if the lower-bounding condition is violated (i.e., if $E' < b_2$ but $p > E'$). Otherwise, the objective function term for \mathbf{z} is $(p_{i_r, j_s}(\mathbf{z}) - b_1)^2$ for $p_{i_r, j_s}(\mathbf{z}) < b_1$ and 0 for $p_{i_r, j_s}(\mathbf{z}) \geq b_1$.

If the modified least-squares method is applied to a set of samples that mostly have $E' > b_1$, then overfitting to the few points with $E' < b_1$ may occur regardless of the number of samples. As an extreme case, if all samples have $E' > b_2$, then almost any polynomial with very large values throughout its domain will give a 0 value for the objective function, but this may still provide a very poor description of the energy landscape. For this situation to be avoided, when a test set of n samples is being drawn and $n/2$ samples with $E' > b_1$ have already been drawn, then if more samples come up with $E' > b_1$, they are redrawn to ensure that a sufficient number of samples

with $E' \leq b_1$ are available (section 2.7). Minimization-aware dead-end elimination pruning⁸ (both singles and pairs pruning) is performed before computation of the polynomial fits because the pruned rotamers and pairs will not be needed during enumeration. This pruning usually eliminates the clashing rotamers and pairs, leaving rotamers and pairs that are well suited for simple polynomial representations.

This objective function can be optimized efficiently because it is convex with respect to the polynomial coefficients (see section 2.4). However, we found general-purpose convex minimizers to be rather time-consuming for the higher-order fits. To address this, we developed the algorithm described in section 2.4. It exploits specific properties of the objective function to obtain a more efficient and reliable fit than a general-purpose convex minimizer would be likely to obtain.

2.4. A Fast Algorithm for Modified Least-Squares Fitting. The following algorithm performs a modified least-squares fit, providing a useful polynomial for energy terms that include both low-energy regions, where an accurate polynomial representation of the energy surface is required, and high-energy regions that must be excluded from our search.

Let us represent our polynomial fit $p(\mathbf{z})$ as $\mathbf{p} \cdot \mathbf{y}(\mathbf{z})$, where \mathbf{p} is the polynomial’s vector of coefficients, $\mathbf{y}(\mathbf{z})$ is the corresponding vector of monomials built from the degree-of-freedom values \mathbf{z} , and \cdot is the standard inner product. For example, if \mathbf{z} consists of the two dihedrals z_1 and z_2 , and we are performing a quadratic fit, then $\mathbf{y}(\mathbf{z})$ will have the elements 1, z_1 , z_2 , z_1^2 , z_2^2 , and $z_1 z_2$. For each sample s in our training set of samples (see section 2.7), let \mathbf{z}_s be the vector of degree-of-freedom values and $\mathbf{y}_s = \mathbf{y}(\mathbf{z}_s)$ be the corresponding vector of monomials. Let E'_s be the energy for the sample, where the minimum-energy point is defined to have 0 energy. Then, a modified least-squares fit consists of minimizing the objective function f to obtain the best-fit polynomial coefficients \mathbf{p}_b

$$\mathbf{p}_b = \arg \min_{\mathbf{p}} \sum_{s \mid E'_s \leq b_1} (E'_s - \mathbf{p} \cdot \mathbf{y}_s)^2 + \sum_{\substack{s \mid \mathbf{p} \cdot \mathbf{y}_s \geq E'_s, \\ b_1 < E'_s < b_2}} (E'_s - \mathbf{p} \cdot \mathbf{y}_s)^2 + \sum_{\substack{s \mid E'_s > b_1, \\ \mathbf{p} \cdot \mathbf{y}_s < b_1}} (b_1 - \mathbf{p} \cdot \mathbf{y}_s)^2 \quad (4)$$

where $\{s \mid E'_s \leq b_1\}$ denotes the set of samples whose energies are less than or equal to b_1 . If we define P_1 to be the set of sample points such that either

$$E'_s \leq b_1 \quad (5)$$

or

$$\mathbf{p} \cdot \mathbf{y}_s \geq E'_s, b_1 < E'_s < b_2 \quad (6)$$

and we define P_2 to be the set of sample points such that

$$E'_s > b_1, \mathbf{p} \cdot \mathbf{y}_s < b_1 \quad (7)$$

then our objective function f becomes

$$\sum_{s \in P_1} (E'_s - \mathbf{p} \cdot \mathbf{y}_s)^2 + \sum_{s \in P_2} (b_1 - \mathbf{p} \cdot \mathbf{y}_s)^2 \quad (8)$$

Thus, if we know P_1 and P_2 , minimizing the objective function is a basic least-squares problem and can be solved analytically. Like basic least-squares, this algorithm operates on a single “training” set of samples and provably minimizes the objective function (i.e., the error) for that training set.

We can show the objective function is convex with respect to \mathbf{p} by noting that the contribution from each sample s is a function of the single linear combination $u = \mathbf{p} \cdot \mathbf{y}_s$ of the elements of \mathbf{p} . This contribution depends on E'_s , but it is always convex (and piecewise quadratic). If $E'_s \leq b_1$, the contribution is just the parabola $(E'_s - u)^2$. If $b_1 < E'_s < b_2$, it is the “truncated” or “flat-bottomed” parabola given by $(b_1 - u)^2$ for $u \leq b_1$, 0 for $b_1 \leq u \leq E'_s$, and $(E'_s - u)^2$ for $u \geq E'_s$. Otherwise (if $b_2 \leq E'_s$), the contribution is the “one-sided” parabola given by $(b_1 - u)^2$ for $u < b_1$ and is 0 otherwise. Hence, the objective function is a sum of convex functions, making it convex itself. Thus, minimizing the objective function to find \mathbf{p} is tractable, with any local minimum being the global minimum. As a result, we know that if, for any sets of samples P_1 and P_2 , we have coefficients \mathbf{p} that minimize eq 8 and satisfy the conditions of eqs 5–7, then the coefficients \mathbf{p} are globally optimal.

The algorithm finds P_1 and P_2 iteratively. As an initial guess, P_1 can be initialized to s such that $E'_s \leq b_1$, and P_2 to be empty. This corresponds to assuming that the one-sided restraints can all be satisfied perfectly. This is followed by performing the basic least-squares computation of minimizing eq 8, which returns coefficients \mathbf{p} , and recalculating P_1 and P_2 from \mathbf{p} using the conditions from eqs 5–7. This procedure is then repeated using the new P_1 and P_2 until a self-consistent solution is found. Generally, only a small minority of the samples will be moved in and out of the least-squares problem at each iteration; thus, the least-squares matrix can be updated quickly at each step, which is useful because forming this matrix is the bottleneck. Typically, only a few iterations are needed.

This algorithm is actually a special case of Newton’s method because its estimate for the objective-function minimum at each iteration is the minimum of the local quadratic Taylor expansion of the objective function. This minimum can be found analytically because the local expansion is convex.

In our implementation of this algorithm, by far the bulk of its time cost is spent in forming the matrix for the first basic least-squares fit (with the initial P_1 and P_2). The subsequent fits are much faster because they are only sparse updates. Thus, the modified least-squares fitting is negligibly more expensive than the first basic least-squares fitting.

2.5. Sparse Atom-Pair Energies (SAPE). SAPE is a method to reduce the degrees of polynomials needed by EPIC by including some nonpolynomial terms in the representation of the energy.

The need for higher-order polynomial fits is driven by large values of higher derivatives. These values are contributed primarily by a small number of van der Waals (vdW) terms between pairs of atoms that are very near each other. It is possible to obtain substantial time and memory savings by evaluating these terms explicitly and fitting the rest of the energy function to a polynomial. To select atom pairs whose vdW terms are to be evaluated explicitly, we chose a cutoff distance of 3 or 4 Å (see section 2.6). Then, an atom pair’s vdW terms are evaluated explicitly if and only if the atoms can be found within that distance of each other within the bounds on internal coordinates for the given residue conformations. These terms are not polynomials in the degrees of freedom because they are inverse powers of distances between atoms, and the atom coordinates themselves are in general not polynomial functions of the degrees of freedom. For example, the expressions for atom coordinates in a side chain in terms of the side chain dihedral angles will include sines and cosines of those angles.

Once we decide to evaluate vdW terms for a given pair of atoms, it costs negligible extra time and memory to also

calculate the electrostatic interaction between these atoms (because we already have the distance between the atoms).

2.6. Attaining the Required Accuracy. We will now describe the methods used to choose polynomial degrees for EPIC fits and ensure that fits of sufficient accuracy are obtained.

Fit accuracy is checked and controlled using cross-validation. For cross-validation purposes, a mean-square error is computed with absolute error used below $E' = 1$ kcal/mol and relative error above. This can be seen as a weighting of the error terms: the weight is 1 for $E' \leq 1$ and $1/\min(E', b_1)$ for $E' \geq 1$ (this levels off at b_1 to avoid excessive underweighting of the one-sided constraints). These weights, which are continuous with respect to E' , are also used during the least-squares fitting.

Cross-validation is used to select the degree of the polynomial that is fit. Low-degree polynomials save time and memory both during the A*/enumeration step and during the precomputation step but may not provide a sufficiently good representation. Hence, we proceed through a sequence of increasingly expensive fits (Figure 1A), and each time a fit is completed, it is cross-validated with an independently drawn set of samples. Like the training set, this cross-validation sample set has ten times as many samples as fit parameters. If the mean-square error is below a specified threshold, the fit is stored, and if it is above, we proceed to the next method. The default threshold value is set to 10^{-4} . However, limited investigation suggests that larger thresholds still tend to keep the errors in the conformations’ minimized energies small compared to thermal energy, and thus are likely acceptable as well. It is also useful to avoid doing fits with a very large number of parameters, as these have enormous time and memory costs in both the enumeration and precomputation steps. Thus, OSPREY is currently set to refuse to do fits with over 2000 parameters—this way, computations that would have prohibitive time costs may still be satisfactorily completed with a slightly higher error threshold than usual.

Some of the fits use lower-degree terms for all degrees of freedom and higher-order terms for selected degrees of freedom. These selected degrees of freedom are eigenvectors \mathbf{v}_k of the Hessian from a modified least-squares quadratic fit (step 1 in the list of steps below). Letting λ_k be the eigenvalue corresponding to \mathbf{v}_k , we define

$$D_q = \left\{ \mathbf{v}_k \mid |\lambda_k| \geq \frac{\max_i |\lambda_i|}{q} \right\} \quad (9)$$

for $q > 0$. Let us define f_n to be a polynomial fit of total degree d (e.g., f_2 is a quadratic fit), $f_d(D_q)$ to be a fit to a polynomial of total degree d in all degrees of freedom plus terms of total degree $d+1$ and $d+2$ in the degrees of freedom in D_q , and $s(n,c)$ to be a polynomial fit of total degree d plus SAPE with a cutoff of $c \lambda$. Fits were tried in the following order: $f_2, s(2,3), f_2(D_{10}), f_2(D_{100}), f_4, s(4,4), f_4(D_{10}), f_4(D_{100}), f_6$, and $s(6,4)$.

The Stone–Weierstrass theorem³² guarantees that a sufficiently high-degree polynomial can approximate any continuous function on any closed and bounded portion of Cartesian space to any desired accuracy. In other words, it guarantees that any energy function can be represented by EPIC to arbitrary accuracy if we allow sufficiently high-degree polynomials. The basis of Bernstein polynomials can be used to construct such approximations with guaranteed convergence to any function.³³ However, for the purpose of energy representation for protein design, modified least-squares is likely to provide good approximations

using much lower-degree polynomials than we would obtain using the Bernstein basis, because we do not need close approximations of the high energy in clashing regions. In these regions, we only need a reasonable lower bound that is much higher than the rotameric wells. This strategy keeps the polynomial degrees low enough to be practical.

2.7. Sampling To Train and Validate Least-Squares Fits. Training and validation sets for EPIC fits consist of sample conformations of the residue(s) involved, specified as vectors of internal coordinates, drawn from throughout the allowed region of conformational space.

By default, samples for both the training and validation sets were sampled uniformly (i.e., each degree of freedom was sampled uniformly and independently from the interval corresponding to the current rotamer or RC). Ten samples were always used in each of these training and validation sets for each parameter in a fit. However, if most of the samples corresponded to energies above the threshold b_1 (see section 2.3), then overfitting could result because, for such samples, there are infinitely many polynomial values that yield zero error. To avoid this, we need sufficient samples from the set B of conformations with energies below b_1 ; B is the set of conformations where the polynomial needs to be quantitatively accurate. We ensure sufficient samples from B by rejecting samples outside of B whenever we desire n samples in total and we already have $n/2$ samples outside B , and thus drawing the rest of our samples uniformly from B by rejection sampling. If 10000 samples are rejected consecutively, indicating that B is too small for efficient rejection sampling, then the Metropolis algorithm³⁴ is used to sample from B .

We have confidence in the parameters obtained by fitting to the training samples for three reasons. First, a useful measure of the accuracy of a polynomial approximation to the energy surface is that there is a low probability that any region of the energy surface deviates significantly from the polynomial approximation (except for high-energy regions approximated by similarly high values of the polynomial). Because our cross-validation of each polynomial fit uses a large number of independent samples (ten times the number of parameters), we are left with a very low chance that our cross-validation samples will miss any such regions. Thus, an insufficiently accurate polynomial surface will be detected upon cross-validation and remedied by an increase in polynomial degree. Second, errors in the minimized energies obtained using polynomial approximations are consistently low, as shown in our computational experiments (Table 1). Third, we expect the energy function to be relatively smooth in the vicinity of a minimum because the gradient must be zero at the minimum, and thus, we expect a polynomial of relatively low order (e.g., the Taylor series of the energy) to yield a good approximation in the vicinity of a minimum.

2.8. Application in Protein Design Algorithms. Once the polynomials are computed, they can be used in protein design algorithms wherever an energy function would ordinarily be called. The GMEC will simply be the set of rotamers for which the minimized value of eq 3 with respect to \mathbf{x} has the lowest possible value.

The simplest method to provably find the GMEC using EPIC is to use a protein design algorithm that enumerates conformations in order of a lower bound, and then instead of minimizing the full energy (eq 2), merely minimizing the polynomial-based energy (eq 3) to compute the energy for each enumerated conformation (Figure 2). For example,

iMinDEE/A*⁸ can be used for this enumeration process, and we use this algorithm in our computational experiments (section 2.10).

EPIC can also be applied in free energy calculations using the K^* algorithm,^{7,12} which approximates binding constants as ratios of partition functions computed from low-energy conformations enumerated by A^* . During these calculations, one can simply use the polynomials instead of the energy function to compute the partition function, given the enumerated RC assignments. This method gives a constant-time speedup determined by the ratio of time to evaluate the energy function versus the EPIC energy.

An additional speedup is possible for branch-and-bound protein design algorithms (e.g., $A^{*14,15}$) that use a tree structure for the conformational search. These algorithms build nodes that each represent a subset of conformational space and are scored using a lower bound on the conformational energies in that space. In each node's conformational space, some residues are restricted to a single RC; these RCs are referred to as "assigned" to their respective residues. At each level of the tree, an RC is assigned to one more residue. One can use the EPIC polynomials to improve the lower-bound energy for each of these nodes. At each node, we need to compute a lower bound L for the conformational energy $q_{\mathbf{r}}$, which is defined in eq 3. That is, we compute L such that

$$L \leq q_{\mathbf{r}}(\mathbf{x}) = \sum_i E_{\Theta}(i_{\mathbf{r}}) + p_{i_{\mathbf{r}}}(\mathbf{x} - \mathbf{x}_0(i_{\mathbf{r}})) + \sum_{j < i} E_{\Theta}(i_{\mathbf{r}}, j_{\mathbf{r}}) + p_{i_{\mathbf{r}}, j_{\mathbf{r}}}(\mathbf{x} - \mathbf{x}_0(i_{\mathbf{r}}, j_{\mathbf{r}})) \quad (10)$$

for all RC assignments \mathbf{r} and all degree-of-freedom values \mathbf{x} that are part of the node's conformational space. If \mathbf{r} is known (i.e., if RCs are fully assigned at all residue positions), then a tight lower bound can be computed trivially by local minimization with respect to \mathbf{x} . Otherwise, we let $q_{\mathbf{r}}(\mathbf{x}) = E_{\Theta}(\mathbf{r}) + E_p(\mathbf{r}, \mathbf{x})$, where E_p consists only of EPIC polynomials

$$E_{\Theta}(\mathbf{r}) = \sum_i E_{\Theta}(i_{\mathbf{r}}) + \sum_{j < i} E_{\Theta}(i_{\mathbf{r}}, j_{\mathbf{r}}) \quad (11)$$

$$E_p(\mathbf{r}, \mathbf{x}) = \sum_i p_{i_{\mathbf{r}}}(\mathbf{x} - \mathbf{x}_0(i_{\mathbf{r}})) + \sum_{j < i} p_{i_{\mathbf{r}}, j_{\mathbf{r}}}(\mathbf{x} - \mathbf{x}_0(i_{\mathbf{r}}, j_{\mathbf{r}})) \quad (12)$$

Now, if we compute lower bounds L_{Θ} and L_p such that $L_{\Theta} \leq E_{\Theta}(\mathbf{r})$ and $L_p \leq E_p$ for all \mathbf{r}, \mathbf{x} in our conformational space, then $L = L_{\Theta} + L_p$ will satisfy eq 10, giving us a valid lower bound. Computation of L_{Θ} has been described previously, because lower bounds of this form are computed in iMinDEE⁸ and DEEPer.¹¹ To compute L_p , we use the fact that EPIC polynomials are always non-negative; thus, for any \mathbf{r} and \mathbf{x} and any subset S of the residues we are modeling,

$$\sum_{i \in S} p_{i_{\mathbf{r}}}(\mathbf{x} - \mathbf{x}_0(i_{\mathbf{r}})) + \sum_{\substack{j \in S, \\ j < i}} p_{i_{\mathbf{r}}, j_{\mathbf{r}}}(\mathbf{x} - \mathbf{x}_0(i_{\mathbf{r}}, j_{\mathbf{r}})) \leq E_p(\mathbf{r}, \mathbf{x}) \quad (13)$$

If we let S be the set of residues with fully assigned RCs, then there is only one possible RC $i_{\mathbf{r}}$ for each residue $i \in S$, and so we can find the minimum of eq 13

$$\min_{\mathbf{r}, \mathbf{x}} \left(\sum_{i \in S} p_{i_{\mathbf{r}}}(\mathbf{x} - \mathbf{x}_0(i_{\mathbf{r}})) + \sum_{\substack{j \in S, \\ j < i}} p_{i_{\mathbf{r}}, j_{\mathbf{r}}}(\mathbf{x} - \mathbf{x}_0(i_{\mathbf{r}}, j_{\mathbf{r}})) \right) \quad (14)$$

Table 1. GMEC Calculation Test Cases^a

protein name	PDB code	mutable residue count	min. speedup with SAPE	min. speedup without SAPE	A* speedup with SAPE	A* speedup without SAPE	fit time (min)	tot. time (min)	GMEC energy error (kcal/mol)
scorpion toxin	1aho	7	26.58	13.60	26.88	48.84	42	49	0.01
scorpion toxin	1aho	9	23.59	14.53	105.70	143.52	435	456	0.03
scorpion toxin	1aho	12	29.66	fit DNF	>142.9 ^b	n/a	566	1119	0.04
scorpion toxin	1aho	14	26.49	fit DNF	>4.30 ^b	n/a	8079	19446	0.05
cytochrome c553	1c75	6	41.03	24.48	102.23	73.78	57	63	0.01
Atx1 metallochaperone	1cc8	7	36.49	bad fit	75.97	bad fit	332	630	0.05
bucandin	1f94	7	15.89	5.03	250.13	179.70	243	427	0.03
nonspecific lipid-transfer protein	1fk5	6	35.25	144.14	24.17	24.87	4	6	0.01
transcription factor IIF	1i27	7	27.66	fit DNF	320.41	fit DNF	2205	3793	0.03
ferredoxin	1iqz	9	59.36	132.22	96.37	127.82	36	41	0.02
Trp repressor	1jhg	7	41.16	bad fit	139.39	bad fit	162	367	0.04
fructose-6-phosphate aldolase	1l6w	6	194.30	178.60	152.26	77.73	197	264	0.03
cephalosporin C deacetylase	1l7a	8	157.71	fit DNF	>3170 ^b	n/a	3341	4784	0.04
PA-I lectin	1l7l	6	58.21	141.90	30.77	32.60	41	55	0.01
phosphoserine phosphatase	1l7m	7	69.84	340.73	810.68	998.47	246	624	0.03
α -D-glucuronidase	1l8n	5	355.82	443.00	0.99	0.93	691	1410	0.01
granulysin	1l9l	7	50.66	34.25	2.54	13.60	94	119	0.01
γ -glutamyl hydrolase	1l9x	5	111.57	468.00	1.86	1.87	103	264	0.00
ferritin	1lb3	5	103.05	33.31	51.68	49.33	148	174	0.01
cytochrome c	1m1q	8	80.79	61.94	320.23	486.77	39	97	0.05
hypothetical protein Ycil	1mwq	8	60.05	48.30	56.10	16.41	60	133	0.00
ponsin	2o9s	14	68.32	bad fit	>45.81 ^b	n/a	218	1732	0.41
scytovirin	2qsk	10	64.94	35.45	296.06	329.99	95	188	0.12
dihydrofolate reductase	2rh2	14	54.92	30.28	771.86	521.44	9	37	0.01
putative monooxygenase	2ril	8	159.98	580.00	28.36	66.72	7	11	0.01
dpy-30-like protein	3g36	4	36.44	17.96	5.77	10.00	23	30	0.05
HIV gp120	3u7y	16	141.17	bad fit	>14.33 ^b	n/a	1773	7762	0.06

^aEPIC with SAPE solved each of these cases. In most cases, A* without EPIC was slower or did not finish within the time limit^b. Minimization (min.) speedup (ratio of single-conformation minimization time without EPIC relative to time with EPIC) is reported for EPIC with and without SAPE. Similarly, A* speedup denotes the ratio of total A* time without EPIC relative to time with EPIC. Fit times (times to calculate the EPIC polynomials) and total calculation times (tot. time) are reported for EPIC calculations with SAPE, calculated without parallelization of fitting. GMEC energy error is the absolute difference between GMEC energies calculated with and without EPIC. Fit DNF means the time limit was exceeded during the polynomial fit precomputation, and bad fit means that fits for EPIC without SAPE, even at the maximum allowed polynomial degree, did not accurately represent the energy; this issue is rescued by SAPE in these cases. When neither A* without EPIC nor EPIC without SAPE were successful, A* speedup without SAPE cannot be calculated and is listed as n/a. EPIC was performed with minimization of partial conformations in all cases. ^bA* without EPIC did not finish within the time limit, so a lower bound on the speedup is reported: the ratio of the time limit (17 days) to the A* time with EPIC and SAPE.

exactly by local minimization with respect to \mathbf{x} . Equation 14 is a lower bound on $E_p(\mathbf{r}, \mathbf{x})$, and thus we set L_p equal to it, giving us a score for our node. We note that L_p is strictly non-negative because eq 13 and thus eq 14 are always non-negative.

Because this continuous minimization with respect to \mathbf{x} is more expensive and has to be performed separately at each node, it is evaluated in a lazy³¹ fashion in our A* implementation. Nodes are assigned the traditional, discrete lower bound L_Θ when they are generated; this bound is fast to compute. The A* priority queue contains nodes both with and without the polynomial contribution L_p included. When a new node is popped from the queue, we check whether or not L_p is present. If it is, we expand the node; if it is not, we compute L_p and insert the node back into the priority queue. This ensures that nodes come off the priority queue in order of their complete lower bound $L_\Theta + L_p$, as is necessary for A* to function correctly. However, it also ensures that we do not waste time computing L_p for nodes whose L_Θ is high enough to preclude expansion. This method gives a combinatorial speedup,

because a high polynomial contribution for a partial conformation can effectively prune an entire branch of the A* tree. In practice, though, the large constant speedup from EPIC minimization of fully assigned conformations tends to be more significant (section 3.1, Table 2).

2.9. Complexity of Energy Evaluations. The speedup due to EPIC can be explained in terms of the asymptotic costs of EPIC polynomial evaluations compared to direct energy function calls. The cost of evaluating an EPIC polynomial scales as the number of terms in the polynomial. This cost is itself a polynomial (usually quadratic) in the number of internal coordinates of the residue pair (or single residue, in the intra+shell case) of interest. In contrast, the cost of evaluating a molecular mechanics-based energy function is generally quadratic in the number of atoms involved because distances between all pairs of atoms need to be considered. EPIC achieves a marked speedup because most residues have far more atoms than significantly flexible internal coordinates. For example, most protein residues have two or fewer side chain

Table 2. Partition Function Calculation Results^a

protein name	PDB ID	mutable residue count	A* speedup due to EPIC ^b	A* speedup due to partial ^c
histidine triad protein	2cs7	14	>48.48 ^d	1.52
ponsin	2o9s	14	>1131 ^d	0.96
transcriptional regulator AhrC	2p5k	11	>256.0 ^d	1.34
scytovirin	2qsk	10	5698.04	1.01
hemolysin	2r2z	12	>217.7 ^d	1.24
dihydrofolate reductase	2rh2	14	>50.45 ^d	2.16
putative monooxygenase	2ril	8	1121.34	1.01
α -crystallin	2wj5	15	>29.06 ^d	0.80
cytochrome c555	2zxy	14	>8.87 ^d	3.32
high-potential iron–sulfur protein	3a38	13	>1916 ^d	1.40
ClpS protease adaptor	3dnj	12	>50.25 ^d	0.93
putative monooxygenase	3fgv	10	>9639 ^d	1.57
protein G	3fl	14	>1475 ^d	1.31
viral capsid	3g21	15	>7.84 ^d	1.21
dpy-30-like protein	3g36	4	211.87	0.91
Hfq protein	3hfo	10	>194.8 ^d	1.45
cold shock protein	3i2z	14	>872.8 ^d	1.24
trypsin	3pwc	10	>624.2 ^d	1.04
trypsin	3pwc	11	>131.8 ^d	2.07

^aEPIC was performed with SAPE in all cases. As in Table 1, EPIC solved all cases, and cases for which A* without EPIC did not finish are denoted.^d ^bRatio of total A* time without EPIC relative to total A* time with EPIC and minimization of partial conformations. ^cRatio of total A* time with EPIC but no minimization of partial conformations relative to total A* time with EPIC and minimization of partial conformations. ^dA* without EPIC did not finish within the time limit, so a lower bound on the speedup is reported: the ratio of the time limit (17 days) to the A* time with EPIC.

dihedrals but over ten atoms. The remaining internal coordinates (e.g., bond length, angles, etc.) are relatively inflexible. Polynomial evaluations are also performed entirely by addition and multiplication, which are much faster than the more complicated elementary operations (e.g., trigonometric functions, square roots, etc.) needed to evaluate molecular mechanics energy terms.

When quantum-mechanical energy functions are introduced, all of the electrons must be accounted for explicitly, and even fairly approximate quantum-chemical methods have time costs that are higher-order polynomials with respect to the number of electrons. For example, any method that accounts for repulsions between all atomic orbitals (e.g., Hartree-Fock and all post-Hartree-Fock methods) must calculate repulsion integrals for all quadruples of atomic orbitals, and there are at least as many atomic orbitals as electrons. Furthermore, there are far more electrons than there are atoms and far more atoms than internal coordinates, giving EPIC an extreme performance advantage. Yet EPIC can represent the same energy surface to a high degree of accuracy once the EPIC polynomials have been precomputed.

Whether or not EPIC is used, these types of pairwise energy evaluations must be performed for every pair of residues in a design system. In general, this means the number of pairwise energy evaluations needed is quadratic with respect to the number of residues in the system. This number can be reduced if a cutoff is applied to remove interactions between distant residues. However, this speedup applies equally for EPIC and non-EPIC calculations.

2.10. Computational Experiments. Protein design calculations were performed in OSPREY^{1,3,7,23,24} with and without EPIC to investigate (a) what previously intractable systems become newly tractable with EPIC, (b) what speedups EPIC brings to conformational enumeration for previously tractable systems, and (c) what types of polynomial representations are needed for these purposes. EPIC runs were performed with SAPE and with conformational minimization for partially assigned conformations during A* search. For comparison, runs with each of these features omitted were also performed.

Times were compared for the A* search, including conformation enumeration and minimization, because this is the portion of the design that is not guaranteed to be completed in polynomial time and is thus the bottleneck. As part of the EPIC runs, GMEC energies were also computed using the regular energy function for comparison to the EPIC results, and the ratio of minimization times with and without EPIC was computed. For runs with multiple conformations very close in energy to each other (within the error range of EPIC, typically <0.1 kcal/mol), the time ratios were averaged.

All minimizations were performed using a cyclic coordinate descent minimizer, which is now included in OSPREY. Default OSPREY energy function settings were used where applicable: AMBER with EEF1 solvation and a distance-dependent dielectric constant of 6. Rotamers were determined using the Penultimate rotamer library.¹³

Test systems were chosen to evaluate both partition function and GMEC calculations and to include all three types of continuous degrees of freedom used in OSPREY: side chain dihedrals, backbone perturbation (shear and backrub) parameters,¹¹ and rigid-body rotations and translations of strands. Some of the tests are intended to be within the scope of previous methods, allowing a quantitative comparison of running times, whereas others are intended to show that EPIC can compute previously intractable GMECs and partition functions with provable accuracy.

For GMEC calculations (Table 1), the first set of systems used was taken from Gainza, Roberts, and Donald,⁸ and featured only side chain dihedral flexibility. The structures for these correspond to PDB IDs 2o9s, 2qsk, 2rh2, 2ril, and 3g36. The second set of systems was taken from Hallen, Keedy, and Donald,¹¹ and included both side chain and backbone flexibility. The structures' PDB codes were 1aho, 1c75, 1cc8, 1f94, 1fk5, 1i27, 1iqz, 1jhg, 1l6w, 1l7a, 1l7l, 1l7m, 1l8n, 1l9l, 1l9x, 1lb3, 1m1q, and 1mwq. Three variants of the 1aho system with more residues were tried as well. Finally, a GMEC calculation was performed for the complex of the HIV surface protein gp120 with the broadly neutralizing antibody NIH45-46 (PDB code 3u7y³⁵).

To investigate the application of EPIC to partition function calculations (Table 2), we first chose systems with only side chain dihedral flexibility from Gainza, Roberts, and Donald⁸ and calculated a partition function for the unliganded protein, with wild-type amino acids at all residue positions, to within 97% guaranteed accuracy. Partition function calculations such as these are the key operation in K^* ^{7,12} calculations. The structures for these correspond to PDB IDs 2cs7, 2o9s, 2p5k, 2qsk, 2r2z, 2rh2, 2ril, 2wj5, 2zxy, 3a38, 3dnj, 3fgv, 3fl, 3g21, 3g36, 3hfo, and 3i2z. Furthermore, a K^* run is presented for trypsin with a small-molecule inhibitor (PDB ID: 3pwc); the run is tractable with EPIC but fails to finish without it. Unlike the calculations for the other, monomeric structures, the K^*

run for trypsin involves the calculation of three partition functions: one each for the protein, ligand, and complex.

Each design was allowed 17 days of total runtime, after which those that had not finished were deemed to have exceeded the time limit and were terminated. A* times with EPIC ranged from 0.7 s to 4 days, and the speedups due to EPIC are shown in Tables 1 and 2.

In these experiments, fitting was performed without parallelization. However, the computation of the EPIC polynomial for each pair of RCs is an independent operation, so each can be done in parallel, meaning that parallelization to p processors will give a p -fold speedup as long as p does not approach the number of RC pairs. OSPREY currently supports computation of each residue pair in parallel, so the speedup holds as long as p does not approach the number of residue pairs. In practice, however, pruning and A* take longer than the polynomial fitting for large systems, which means that this parallelization may not be necessary. Additionally, once the EPIC fits have been computed for a system, there may be a large number of computations that can be performed using it, including calculation of partition functions for many sequences, computation of GMECs for various subsets of the sequence space, etc. These extensive reuses of the fits may be especially desirable when designing a library of sequences for experimental testing—if one performs various optimizations with different assumptions and tests top sequences from each optimization, the results will be more robust to errors in the assumptions.

To investigate the ability of EPIC to represent quantum-mechanical energy functions, we performed EPIC calculations on the aspartame dipeptide (extracted from PDB ID: 1a8j³⁶) with the energies for EPIC samples evaluated using NWChem³⁷ instead of OSPREY's usual energy function. Calculations were performed at the SCF level of theory with STO-3G and with 6-31G** basis sets and also at the MP2 level of theory with a STO-3G basis set.²⁰ For each rotamer of each residue, dihedrals were sampled within the allowed range for the rotamer, and the total energy of the dipeptide was fit to a polynomial.

2.11. Applications of EPIC to Other Algorithms. For this study, EPIC was implemented in the context of the protein design package OSPREY^{1,3,7,23,24} to run along with the algorithms (iMinDEE, DEEPer, and K^*) and pairwise energy functions already implemented in OSPREY. However, EPIC would enable some other capabilities in different implementations.

First, EPIC can be applied in the context of other protein design algorithms. For example, it can be applied to an iterative algorithm, such as FASTER³⁸ or Monte Carlo,³⁴ that tries to find a suitably low-energy conformation by accepting or rejecting rotamer changes based on the energies of conformations with these changes. Whenever the energy is needed for a rotamer assignment, the EPIC energy for the protein can be locally minimized starting at the ideal internal coordinate values for that rotamer assignment. In this case, the matrix of EPIC polynomials substitutes directly for the matrix of pairwise rotamer energies commonly used to calculate conformational energies for these algorithms in the absence of continuous flexibility. EPIC could even be used for molecular dynamics because most residue pairs in a molecular dynamics trajectory⁴ will spend most of their time in fairly relaxed conformations, which is the region of conformational space modeled by EPIC. In all these cases, EPIC energy evaluations would be markedly faster than regular energy function calls, particularly for expensive energy functions. Thus, EPIC would provide




Figure 4. Mutable residues in the redesign of the surface of the HIV surface protein gp120 in complex with the broadly neutralizing antibody NIH45-46 (PDB ID: 3u7y³⁵). This design finished only when EPIC was used. Mutable residues, blue backbone and pink side chains; gp120, black backbone; NIH45-46 heavy chain, green backbone; NIH45-46 light chain, brown backbone.

a substantial speedup for any algorithm whose bottleneck is energy function calls.

For docking algorithms³⁹ that require side chain optimization or local backbone optimization, EPIC can be used for both conformation scoring and optimization using any of the above algorithms.

2.12. EPIC Can Accommodate Higher-than-Pairwise Energies. EPIC was implemented in this work to handle pairwise energy functions (in the sense of a sum of 1-body and 2-body energies with no terms dependent on three or more residues' degrees of freedom) because these are currently typical for protein design and include the AMBER, CHARMM, and EEF1 energy functions we use in OSPREY. However, the true energy of proteins is not exactly pairwise decomposable, and EPIC could easily accommodate higher-order terms. EPIC simply requires that each energy term correspond to a set D of degrees of freedom, constrained to a region in which they are relatively well-behaved (e.g., dihedrals at each residue constrained to a single rotamer); we can sample D subject to the constraints and then fit the energies as a polynomial function with domain D . Thus, for any set R of more than two interacting residues, for each RC assignment to those residues, we can fit an EPIC polynomial and thus describe the energy terms for R .

In practice, the number of sets of residues that can interact significantly is quite limited, because residues typically must be physically near each other to have significant higher-than-pairwise interactions. For example, if we have a Ramachandran-based potential, its terms each depend on the ψ and ϕ backbone dihedrals of a certain residue r and thus depend on the conformations of the three residues $r - 1$, r , and $r + 1$. Likewise, the conformation of a residue i can induce polarization effects in a nearby residue j that will affect the interactions of j with another residue k , and this effect can be quantified using quantum chemistry, but i and j have to be physically very close to each other ($\ll 1$ nm) for this effect to be significant (and j and k have to be fairly close too, probably subnanometer as well, because

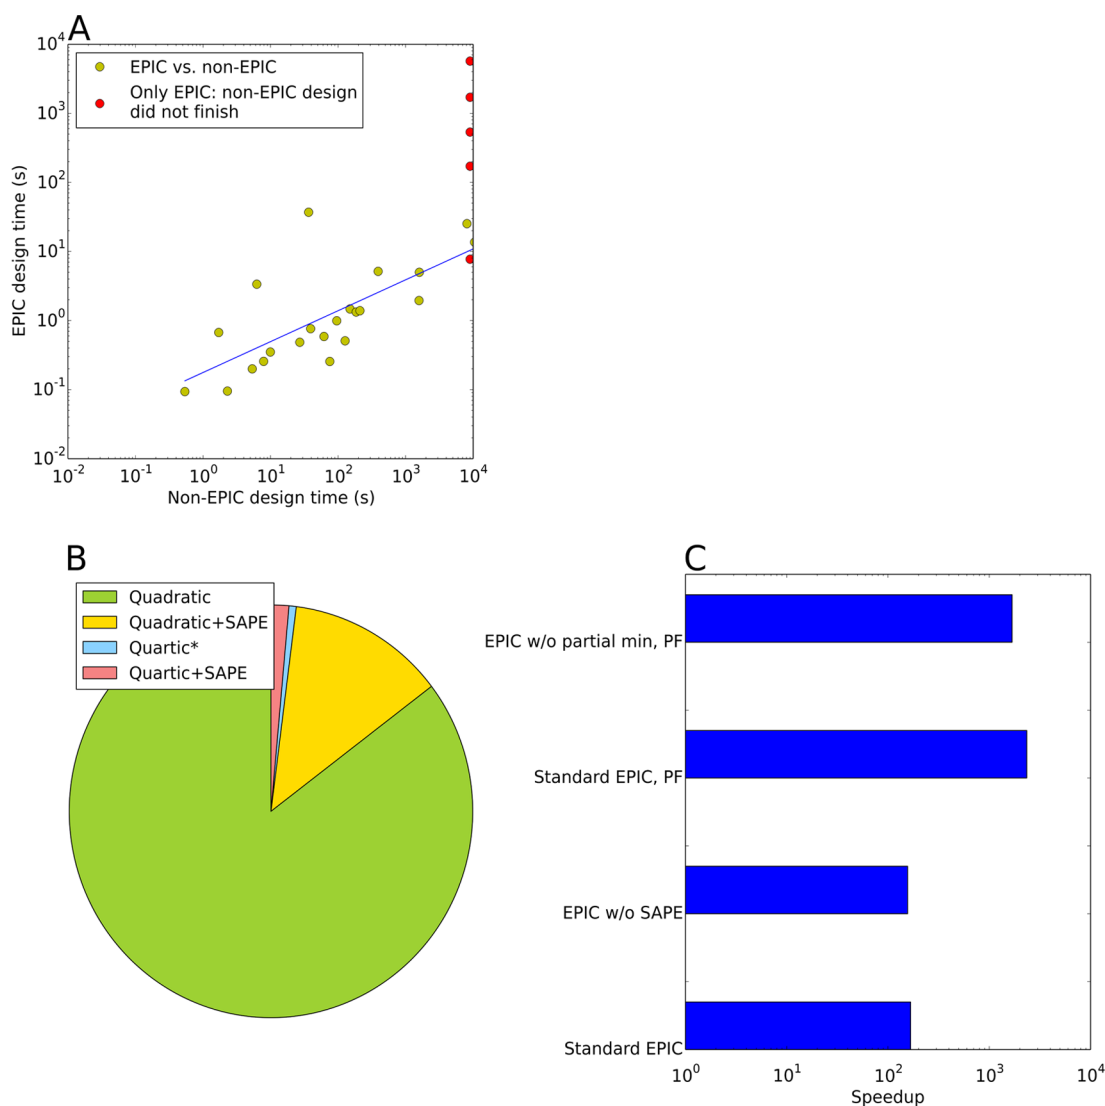


Figure 5. (A) A* times with and without EPIC. Five designs that did not finish without EPIC are shown on the right in red. (B) Proportions of each type of fit (see section 2.6) required in EPIC calculations. The “quartic*” category includes both full quartic fits and quadratic fits with quartic terms added for D_{10} or for D_{100} . Fits were all made as high-degree as needed to obtain a residual below 0.0001, as described in section 2.6. Some fits have substantially lower residuals, especially quadratic fits without SAPE, because no lower fit degrees were allowed. (C) Speedups due to different EPIC methods compared to A* based on pairwise lower-bound energies; standard EPIC includes both SAPE and minimization of partial conformations. PF denotes partition function calculations; the others are GMEC calculations.

the potential of an induced dipole falls off faster than $1/d^2$ with distance d). Hence, EPIC can be used to model any realistic energy function by accounting for all sets of residues with significant energetic interactions.

3. RESULTS

Computational experiments were performed to measure what kinds of polynomial fits are necessary to accurately model different proteins with different degrees of freedom and energy functions, and what speedups EPIC brings to DEE/A* and K* calculations. The results demonstrate that EPIC provides a substantial speedup to design calculations when proteins are modeled in the same manner as previous OSPREY designs.^{1,3,8,11,23} They also show that EPIC efficiently represents energies calculated by quantum chemistry and is a potentially decisive tool for using both realistic, continuous flexibility and quantum-mechanical energy functions in protein design.

3.1. Application to Protein Designs. First, computational experiments were performed to compare GMEC search with and without EPIC, as described in section 2.10. Key portions of the design calculation were timed with and without EPIC to determine the speedup for these portions in particular (Table 1). On average, minimization of fully enumerated conformations was 79-fold faster using EPIC than with traditional energy function calls. Overall, A* speedups due to EPIC averaged 167-fold (Figure 5A). The overall A* speedup is likely greater than the minimization speedup because of the way OSPREY’s standard energy function is implemented. Each time the energy function is run on a new sequence, setup time (e.g., initialization of the energy function) is required to identify electrostatic, van der Waals, and solvation terms that will be necessary for that sequence. This setup time is eliminated by EPIC and is not counted as part of the minimization time here, but it may be performed an exponential number of times without EPIC because minimizations may be required for an

exponential number of sequences. Runs that did not finish without EPIC are not included in these averages. Eighty-five percent of the fits in these experiments were quadratic with no SAPE needed (Figure 5B). GMECs from EPIC runs showed good agreement between energies from minimization of EPIC energies and energies from the actual energy function. The average energy difference was 0.04 kcal/mol, which is less than one-tenth of thermal energy at room temperature (0.592 kcal/mol, calculated as the universal gas constant times a room temperature of 298 °K) and thus functionally insignificant.

Five of the 27 systems finished only with EPIC, demonstrating that EPIC allows for the design of larger and more diverse systems than were previously possible. For example, a redesign of the complex of HIV surface protein gp120 with the antibody NIH45-46 did not finish when run without EPIC but finished with EPIC using about one day of A* time (Figure 4). This redesign allowed the mutation of 16 residues all over the gp120 surface in the interface—five in the D-loop of gp120,⁴⁰ which is central to the interaction with NIH45-46, and the other 11 scattered through other parts of the interface in various types of secondary structures. Redesigns of the gp120 surface to achieve specific binding to particular antibodies has been instrumental in the development of probes to isolate these antibodies from sera.²⁸ Redesigning the antibody surface of a gp120–antibody complex has also been effective in optimizing antibody affinity,²⁵ which is useful for passive immunization and immunogen design. Interestingly, the redesign of the NIH45-46 complex yielded 12 top conformations within 0.06 kcal/mol of each other—two from the top sequence and ten from a double mutant. This high density of favorable conformations suggests that the complex is entropically favored, a result consistent with the observed high affinity of NIH45-46 for gp120, attained through extensive affinity maturation of the antibody.

Two variations of EPIC were also tried for these systems. It was found that minimization of partial conformations during A* (section 2.8) provides a speedup of 2.3-fold on average, though it is not nearly as great as the speedup from faster minimization of fully assigned conformations (Figure 5). Furthermore, EPIC without SAPE was often effective; however, under some circumstances, it was unable to provide accurate fits (as usual, trying only the polynomial degrees described in section 2.6). In systems where EPIC without SAPE was effective, it averaged insignificantly (1.1-fold) slower than EPIC with SAPE for A*. However, there were four systems that required SAPE to give accurate results (out of 27; Table 1), including the HIV gp120 complex with antibody NIH45-46 (Figure 4). There were also four systems that exceeded the time limit during fitting. These limitations on EPIC without SAPE do not indicate a fundamental theoretical barrier, because in principle the Stone-Weierstrass theorem guarantees an accurate fit if the polynomial degree is increased sufficiently (see section 2.6). However, they do indicate that EPIC without SAPE may sometimes require polynomial degrees that are prohibitively time-consuming for typical protein designs and/or a higher numerical precision than the double precision efficiently supported in Java and thus used in OSPREY.

Experiments were also performed to compare partition function calculations with and without EPIC (Table 2). To obtain a provably good approximation to the partition function,^{7,12} many more conformations must be enumerated and

minimized than for GMEC calculations. As a result, marked speedups were achieved by EPIC (Figure 5C). Out of the 19 systems for which EPIC finished, only three finished without EPIC (2000-fold average speedup). The speedup in EPIC designs from minimization of partial conformations was modest at 1.4-fold.

With an A* speedup of 2–3 orders of magnitude, designs that would previously take years can be performed with EPIC in days. This will allow many designs that would otherwise be considered intractable to be completed using EPIC.

3.2. Quantum-Mechanical Energies. In addition to classical mechanics-based energy functions, we also used EPIC to fit conformational energies of the aspartame dipeptide calculated using quantum-mechanical models of electronic structure. EPIC fits for aspartame showed that quantum-mechanical energies and AMBER and EEF1 energies can be represented by polynomials of very similar degree (i.e., energy surfaces from quantum chemistry are just as polynomial-like as energy surfaces from molecular mechanics). SAPE was not found to significantly increase the accuracy of the fits, and thus were not included, though it is likely that reparameterized van der Waals and/or electrostatic terms (or other specially fit functions of the atom-pair distance) would be able to improve the fit quality. This discrepancy indicates that the atom-pair energies used in SAPE are a poor approximation to the interactions between the same atom pairs predicted by quantum mechanics and, thus, that energies returned by quantum-mechanical and molecular-mechanics methods are substantively different.

For Phe 2 of aspartame, the same types of polynomial fits were needed for Hartree-Fock with a STO-3G basis set, Hartree-Fock with a 6-31G** basis set, and the usual AMBER/EEF1 energy function (Figure 6). These were quadratic fits for three rotamers and a quadratic fit plus quartic fits on D_{10} for the fourth. MP2 with an STO-3G basis set also required quadratic fits for the first three rotamers and required a quadratic fit plus quartic fits on D_{100} for the fourth.




Figure 6. Intra-residue energy calculated for Phe 2 of aspartame using Hartree-Fock theory with an STO-3G basis set, and quadratic EPIC fit, as a function of the two side chain dihedrals. The fit is very close to the energy surface, though a slight discrepancy is visible in the upper right-hand corner ($\chi_1 \approx 70^\circ$, $\chi_2 \approx 100^\circ$).

For Asp 1, AMBER/EEF1 required quadratic terms plus quartic terms on D_{10} for three rotamers and on D_{100} for two. Both Hartree-Fock and MP2 with an STO-3G basis set required slightly simpler fits: only quadratic for one rotamer, quadratic plus quartic terms on D_{10} for three, and quadratic plus quartic terms on D_{100} for one.

These results show that quantum and polarization-type effects can be represented effectively by polynomial fits, and the polynomial degrees required are essentially the same as for AMBER/EEF1.

4. CONCLUSIONS

EPIC eliminates the current bottleneck in minimization-aware protein design by performing energy function calls only in a precomputation step. It thus opens several avenues for more accurate and efficient design calculations. More residues can now be mutated, and more ligands can be tested. Additional continuous conformational degrees of freedom (e.g., in the backbone) can be modeled, and minimization can be performed over a greater range for each degree of freedom when appropriate. In this sense, EPIC helps protein design algorithms emulate the extensive continuous flexibility of molecular dynamics algorithms, while searching an exponentially large sequence space that would be intractable for molecular dynamics-based design. Furthermore, more accurate but (previously) slower energy functions can be incorporated without any asymptotic increase in computation time.

The polynomial representation of energy provided by EPIC could also allow dedicated algorithms for polynomials to be used in energy calculations. For example, because exact derivatives of polynomials are trivial to compute, EPIC is very amenable to rapid calculation of energy derivatives with respect to internal coordinates, which are used in many minimization algorithms.^{41,42} Note that the gradient of the EPIC polynomials may not approximate the gradient of the energy (i.e., forces) to the same degree of accuracy as the polynomials approximate the energy because the polynomials are fit to the energies rather than the forces. Thus, differentiation of the fit polynomial may amplify noise. However, when we are numerically minimizing the polynomial approximation to the energy, we require derivatives of the polynomials themselves.

EPIC fits are tractable, accurate approximations that provide a new understanding of the energy landscape of proteins in the vicinity of ideal rotamers and more generally in the vicinity of energy minima. This is useful because the high dimensionality of conformational space makes direct visualization difficult.

By enabling better modeling both of conformational space and of conformational energies, EPIC moves us closer to the goal of algorithms that can produce reliable predictions for our biomedically and biologically important protein and drug design problems. EPIC is thus offered to the protein design community both as an immediate speedup in designs and as an enabling technology for future improvements.

AUTHOR INFORMATION

Corresponding Author

*E-mail: brd+jctc14@cs.duke.edu.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

We thank all members of the Donald lab for helpful comments and the PhRMA foundation (M.A.H.) and NIH (Grant R01-GM-78031 to B.R.D.) for funding.

REFERENCES

- (1) Frey, K. M.; Georgiev, I.; Donald, B. R.; Anderson, A. C. *Proc. Natl. Acad. Sci. U.S.A.* **2010**, *107*, 13707–13712.
- (2) Roberts, K. E.; Cushing, P. R.; Boisguerin, P.; Madden, D. R.; Donald, B. R. *PLoS Comput. Biol.* **2012**, *8*, e1002477.
- (3) Chen, C.-Y.; Georgiev, I.; Anderson, A. C.; Donald, B. R. *Proc. Natl. Acad. Sci. U.S.A.* **2009**, *106*, 3764–3769.
- (4) Rapaport, D. C. *The Art of Molecular Dynamics Simulation*, 2nd ed.; Cambridge University Press: Cambridge, England, 2004.
- (5) Donald, B. R. *Algorithms in Structural Molecular Biology*; MIT Press: Cambridge, MA, 2011.
- (6) Janin, J.; Wodak, S.; Levitt, M.; Maigret, B. *J. Mol. Biol.* **1978**, *125*, 357–386.
- (7) Georgiev, I.; Lilien, R. H.; Donald, B. R. *J. Comput. Chem.* **2008**, *29*, 1527–1542.
- (8) Gainza, P.; Roberts, K.; Donald, B. R. *PLoS Comput. Biol.* **2012**, *8*, e1002335.
- (9) Georgiev, I.; Donald, B. R. *Bioinformatics* **2007**, *23*, i185–i194.
- (10) Georgiev, I.; Keedy, D.; Richardson, J. S.; Richardson, D. C.; Donald, B. R. *Bioinformatics* **2008**, *24*, i196–i204.
- (11) Hallen, M. A.; Keedy, D. A.; Donald, B. R. *Proteins: Struct., Funct., Bioinf.* **2013**, *81*, 18–39.
- (12) Lilien, R. H.; Stevens, B. W.; Anderson, A. C.; Donald, B. R. *J. Comput. Biol.* **2005**, *12*, 740–761.
- (13) Lovell, S. C.; Word, M. J.; Richardson, J. S.; Richardson, D. C. *Proteins: Struct., Funct., Genet.* **2000**, *40*, 389–408.
- (14) Hart, P. E.; Nilsson, N. J.; Raphael, B. *IEEE Transactions on Systems Science and Cybernetics* **1968**, *4*, 100–107.
- (15) Leach, A. R.; Lemon, A. P. *Proteins: Struct., Funct., Bioinf.* **1998**, *33*, 227–239.
- (16) Li, Z.; Scheraga, H. A. *Proc. Natl. Acad. Sci. U.S.A.* **1987**, *84*, 6611–6615.
- (17) Wang, C.; Schueler-Furman, O.; Baker, D. *Protein Sci.* **2005**, *14*, 1328–1339.
- (18) Lazaridis, T.; Karplus, M. *Proteins: Struct., Funct., Bioinf.* **1999**, *35*, 133–152.
- (19) Vizcarra, C. L.; Zhang, N.; Marshall, S. A.; Wingreen, N. S.; Zeng, C.; Mayo, S. L. *J. Comput. Chem.* **2008**, *29*, 1153–1162.
- (20) Cook, D. B. *Handbook of Computational Quantum Chemistry*; Dover Publications, Inc: Mineola, NY, 2005.
- (21) Pierce, N. A.; Winfree, E. *Protein Eng.* **2002**, *15*, 779–782.
- (22) Chazelle, B.; Kingsford, C.; Singh, M. *INFORMS Journal on Computing* **2004**, *16*, 380–392.
- (23) Gainza, P.; Roberts, K. E.; Georgiev, I.; Lilien, R. H.; Keedy, D. A.; Chen, C.-Y.; Reza, F.; Anderson, A. C.; Richardson, D. C.; Richardson, J. S.; Donald, B. R. *Methods Enzymol.* **2013**, *523*, 87–107.
- (24) Georgiev, I.; Roberts, K. E.; Gainza, P.; Hallen, M. A.; Donald, B. R. *OSPREY (Open Source Protein Redesign for You) User Manual*. Available online: www.cs.duke.edu/donaldlab/software.php (2009). Updated **2012**, 94 pages.
- (25) Rudicell, R. S.; et al. *J. Virol.* **2014**, *88*, 12669–12682.
- (26) Gorczynski, M. J.; et al. *Chemistry and Biology* **2007**, *14*, 1186–1197.
- (27) Stevens, B. W.; Lilien, R. H.; Georgiev, I.; Donald, B. R.; Anderson, A. C. *Biochemistry* **2006**, *45*, 15495–15504.
- (28) Georgiev, I.; Acharya, P.; Schmidt, S.; Li, Y.; Wycuff, D.; Ofek, G.; Doria-Rose, N.; Luongo, T.; Yang, Y.; Zhou, T.; Donald, B. R.; Mascola, J.; Kwong, P. *Retrovirology* **2012**, *9*, P50.
- (29) Weiner, P. K.; Kollman, P. A. *J. Comput. Chem.* **1981**, *2*, 287–303.
- (30) Cornell, W. D.; Cieplak, P.; Bayly, C. I.; Gould, I. R.; Merz, K. M.; Ferguson, D. M.; Spellmeyer, D. C.; Fox, T.; Caldwell, J. W.; Kollman, P. A. *J. Am. Chem. Soc.* **1995**, *117*, 5179–5197.

- (31) Cormen, T. H.; Leiserson, C. E.; Rivest, R. L.; Stein, C. *Introduction to Algorithms*, 3rd ed.; MIT Press: Cambridge, MA, 2009.
- (32) Stone, M. H. *Mathematics Magazine* **1948**, *21*, 167–184.
- (33) Bernstein, S. N. *Communications de la Société Mathématique de Kharkoff* **1912**, *13*, 1–2.
- (34) Metropolis, N.; Rosenbluth, A. W.; Rosenbluth, M. N.; Teller, A. H.; Teller, E. J. *Chem. Phys.* **1953**, *21*, 1087–1092.
- (35) Diskin, R.; Scheid, J. F.; Marcovecchio, P. M.; West, J.; Anthony, P.; Klein, F.; Gao, H.; Gnanapragasam, P. N. P.; Abadir, A.; Seaman, M. S.; Nussenzweig, M. C.; Bjorkman, P. J. *Science* **2011**, *334*, 1289–1293.
- (36) Edmunson, A. B.; Manion, C. V. *Clin. Pharmacol. Ther.* **1998**, *63*, 580–593.
- (37) Valiev, M.; Bylaska, E. J.; Govind, N.; Kowalski, K.; Straatsma, T. P.; van Dam, H. J. J.; Wang, D.; Nieplocha, J.; Apra, E.; Windus, T. L.; de Jong, W. A. *Comput. Phys. Commun.* **2010**, *181*, 1477–1489.
- (38) Desmet, J.; Spriet, J.; Lasters, I. *Proteins: Struct., Funct., Bioinf.* **2002**, *48*, 31–43.
- (39) Cerqueira, N. M. F. S. A.; Bras, N. F.; Fernandes, P. A.; Ramos, M. J. *Proteins: Struct., Funct., Bioinf.* **2009**, *74*, 192–206.
- (40) Kwong, P. D.; Wyatt, R.; Robinson, J.; Sweet, R. W.; Sodroski, J.; Hendrickson, W. A. *Nature* **1998**, *393*, 648–659.
- (41) Hestenes, M. R.; Stiefel, E. J. *Res. Natl. Bur. Stand.* **1952**, *49*, 409–436.
- (42) Nocedal, J.; Wright, S. J. *Numerical Optimization*, 2nd ed.; Springer-Verlag: Berlin, Germany, 2006.