

Tracing Conformational Changes in Proteins

Nurit Haspel¹, Mark Moll¹, Matthew L. Baker², Wah Chiu^{2,3} and Lydia E. Kaviraki^{1,3,4}

¹Department of Computer Science, Rice University, Houston, TX 77005 USA

²National Center for Macromolecular Imaging, Verna and Marrs McLean Department of Biochemistry and Molecular Biology, Baylor College of Medicine, Houston, TX 77030 USA

³Graduate Program in Structural and Computational Biology and Molecular Biophysics, Baylor College of Medicine, Houston, TX 77030 USA

⁴Department of Bioengineering, Rice University, Houston, TX 77005 USA

Abstract—Many proteins undergo extensive conformational changes as part of their functionality. Tracing these changes is important for understanding the way these proteins function. Traditional biophysics-based conformational search methods require a large number of calculations and are hard to apply to large-scale conformational motions. In this work we investigate the application of a robotics-inspired method, using backbone and limited side chain representation and a coarse grained energy function to trace large-scale conformational motions. We tested the algorithm on three well known medium to large proteins and we show that even with relatively little information we are able to trace low-energy conformational pathways efficiently. The conformational pathways produced by our methods can be further filtered and refined to produce more useful information on the way proteins function under physiological conditions. Contact: Lydia E. Kaviraki, kaviraki@cs.rice.edu

I. INTRODUCTION

Proteins are flexible molecules that undergo conformational changes as part of their interactions with other proteins or drug molecules [1]. Changes in torsional angles may induce localized changes or large scale domain motions. Figure 1 shows an illustration of the closed structure of the GroEL monomer taken from PDB code 1SS8 (Figure 1(a)) and the opened structure (GroEL-GroES-ADP7) taken from PDB code 1SX4 (Figure 1(b)). GroEL transitions between the closed and open conformations as part of its chaperone activity, but the structural details of the transition process are not fully understood. Tracing these changes is crucial for understanding the way these proteins perform their function. Existing physics-based computational methods that trace and simulate conformational changes in proteins include Molecular Dynamics (MD) [2], Monte Carlo (MC) [3] and their variants. These methods require large amounts of computational resources and are therefore hard to apply to conformational motions that take place over time scales larger than several hundreds of nanoseconds. In the past two decades several efficient conformational search algorithms have been developed. Some use a coarse representation of the protein molecule [4], [5], [6] and employ various efficient search methods such as Normal Mode Analysis (NMA) [7], [8], elastic network modeling [9], [10], [11], [12], [13], [14], or morphing [15], [16]. In recent years sampling based motion planning methods have been successfully applied towards an efficient exploration of protein conformational space. Motion planning is an area in robotics concerned with finding a pathway for robot-like objects in constrained environments [17], [18], [19]. When applied to

biological problems, the protein is represented as an articulated body with the degrees of freedom in all or selected torsional angles. The physical constraints are implicitly encoded in a penalty function which approximates the potential energy of the molecule. The conformational space of the protein is explored so that high energy regions are avoided and feasible conformational pathways are obtained more efficiently than with traditional simulation methods. Among the many applications of motion planning to biology are the characterization of near-native protein conformational ensembles [20], the study of conformational flexibility in proteins [21], [22], protein folding and binding simulation [23], [24], [25], modeling protein loops [26], [27], simulation of RNA folding kinetics [28] and recently the elucidation of conformational pathways in proteins, subject to pre-specified constraints [29].

The search methods described above strike a balance between accuracy and efficiency. Many of those methods are successful in sampling the conformational landscape of proteins but are often biased by the protein native conformation and some of them require additional, problem specific information. Additionally, when atomic details are skipped the conformational search process is greatly accelerated but fine details are missed.

In this work we present a prototype of a novel, efficient motion-planning based methodology to perform conformational search on proteins requiring only backbone and limited side-chain information. The molecule is mapped into a reduced representation using a small number of parameters that represent its degrees of freedom. This allows for larger motions to be explored efficiently. We aim to make the conformational search as general as possible so it can be applied with as little system specific knowledge as possible. We use a coarse-grained physics based energy function which captures low energy conformations in a realistic but efficient way [30]. We identify the flexible parts of the proteins and manipulate them to simulate the conformational changes, treating the rest of the protein as rigid. In this way we reduce the dimensionality of the search space while still capturing the essential conformational flexibility of the protein. We tested our methodology on three proteins of different lengths (214 to 525) amino acids, known to undergo extensive conformational changes. The results show that we are able to efficiently produce low energy pathways for each one of them. The method can serve as a filtering tool which can provide biologists with useful hypotheses about the way proteins transition from one

conformational state to another, and help to gain more insight about protein function.

Problem Statement

Given two conformational states of a molecule, denoted by start and goal, our objective is to find conformational pathways connecting the start and goal conformations. A pathway is a sequence of affine transformations that, when applied successively to the degrees of freedom of the start conformation, the start conformation will be brought to within a tolerance range of the goal conformation under a defined distance metric. Furthermore, the energy of each intermediate conformation along the pathway must be lower than a given threshold as measured by a potential function that approximates the protein energy. The degrees of freedom of the structures lie in the flexible parts connecting rigid structural elements. Several assumptions are made in this paper. We assume that secondary structure elements do not change significantly during domain motion and that the flexible parts are the loops connecting secondary structure elements. While this assumption is true in many cases, there are cases where secondary structure elements melt or change. In these cases, it is possible to incorporate a more detailed modeling of the flexible parts into the general framework of the algorithm without limiting the proposed procedure. It should be emphasized that the algorithm does not always produce the same conformational pathway, but rather a possible pathway. By repeating the procedure a large number of times we produce a set of feasible pathways, thus limiting the huge search space to a manageable number of possibilities. These pathways can later be clustered, refined and filtered using information about the tested systems. The size of the clusters can give us information about the likelihood of given conformations along the pathway.

II. METHODS

In the following section we describe in depth the conformational search method including the data structures, distance metric, energy function and search algorithm used to perform the conformational search and produce low-energy pathways.

A. Data Representation

We use a coarse grained representation used successfully by our research group in the past [30]. The proteins are stripped of their side-chain and hydrogen atoms and represented at the backbone + C- β level (Glycine is represented by its backbone only). The amino acids are grouped into secondary structure elements. The secondary structures can be assigned by the PDB header or using a secondary structure assignment algorithm such as DSSP [31]. Loops are assigned to the nearest secondary structure element. To save computational time, it is possible to cluster several secondary structure elements into one rigid element if their positions are known not to change with respect to one another during the conformational transition. Alternatively, to gain accuracy, in highly movable regions of the protein such as flexible loops or if some secondary structure elements are known to break or change,

their structural representation can be refined and broken down to smaller sub-structures. This refinement is not considered in the context of this paper but it can be applied without loss of generality. The high-level data structure that represents a conformation is a graph $G = (V, E)$ such that each secondary structure element is a node $v \in V$ in the graph. Two secondary structure elements are connected by an edge $e \in E$ if there is at least one pair of consecutive amino acids u, v , such that $u \in S_1$ and $v \in S_2$. Based on the graph we construct a spanning tree $T = (V, X)$ where X is a subset of E using a greedy approach. The root of the tree is specified as the structure that moves the least during the search as determined by aligning the start and goal structures and measuring the least RMSD between corresponding secondary structure elements. Each one of the root's neighbors forms a child node in the tree, and at each stage the selected node and its adjacent edges are removed from the graph. The process repeats iteratively until all the secondary structure elements are represented in the tree.

B. Distance Between Structures

Motion planning methods need a distance measure to estimate the progress of the search. In the case of proteins a distance measure is hard to define due to the complexity of protein structures and the high dimensionality of the problem equal to the number of the rotational degrees of freedom of the protein. To overcome this challenge and because the structure is manipulated at the flexible loops connecting secondary structure level, we devised a distance measure that estimates the similarity between two protein structures based on the positioning of their secondary structure elements with respect to one another.

Given a conformation C , we first define a score for each secondary structure element i in C :

$$score(C^i) = \sum_{j \in K} (|\alpha_{ij} - \alpha'_{ij}| \times w_i + |d_{ij} - d'_{ij}| \times w'_i). \quad (1)$$

The summation is over the set K of secondary structures in C excluding i , α_{ij} is the angle and d_{ij} is the distance between secondary structure element i and secondary structure element j in C , α'_{ij} is the angle and d'_{ij} is the distance between the corresponding secondary structure elements in the goal structure, and w_i and w'_i are weight factors proportional to the size of secondary structure element i , such that the angle and distance components will be brought to the same order of magnitude. In the current implementation we use the values of 1 for w_i and 5 for w'_i , which seem to give the best results. An angle between two secondary structure elements is defined as the angle between the two vectors representing them. A vector representing a helix is the least square straight line that passes through the helix atoms, and a vector representing a sheet is the normal to the surface best representing the sheet. The distance between two secondary structure elements is defined as the distance between their centers of masses.

We then compute for a conformation C a feature vector:

$$v_C = \langle score(C^1), score(C^2), \dots, score(C^k) \rangle \quad (2)$$

where the components of the vector are the scores of the K secondary structure elements of the conformation.

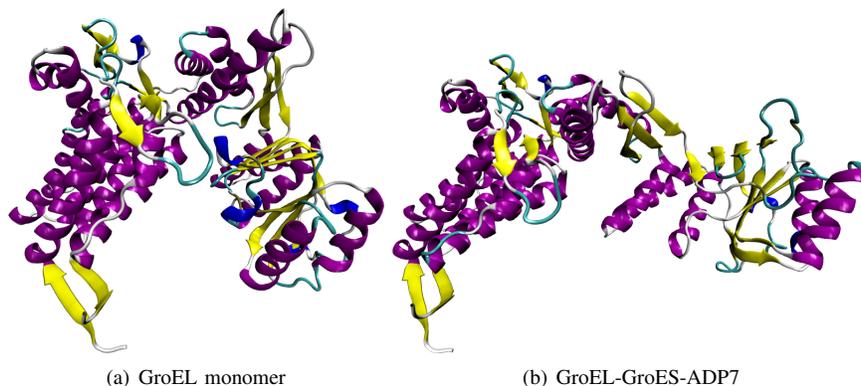


Fig. 1. (a) The GroEL monomer (PDB structure 1ss8). (b) The GroEL-GroES-ADP7 monomer (PDB structure 1sx4).

The distance between two conformations, C_1 and C_2 is defined as the Euclidean distance between their feature vectors, i.e., $\|v_{C_1} - v_{C_2}\|^2$. By definition, when C_2 is the goal structure, the *score* of C_1 is the magnitude of its vector representation. Therefore, the lower the score for a given conformation, the more similar it is to the goal structure. This vector is used as a projection of the conformation to a lower dimension subspace which is used to measure coverage of the search space by the search method described below. It should be noted that other distance measures exist [32], but after extensive experimentation the measure described above produced the best results.

C. Energy Function

In order to approximate the potential energy of the produced conformations we suggest a simplified energy function which includes the following components:

$$E_{total} = E_{soft-vdW} + E_{HB} + E_{burial} + E_{water} + E_{bond} + E_{angle} \quad (3)$$

The first four terms in this coarse-grained energy function are a part of a function successfully used in our group in the past [30]. The compaction term mentioned in [30], which biases the energy towards folded, compact structures, was removed from our implementation since we are not simulating protein folding. The bond and angle terms are taken from the AMBER ff03 force field [33]. If the structural manipulation causes the energy to be at least 100 kcal/mol higher than the energy of the starting structure 20 minimization steps are performed over the bond, angle and van der Waals degrees of freedom of the manipulated secondary structure elements using a steepest descent scheme [34].

D. Search Methodology

The search is performed using a sampling based motion planning algorithm. Motion planning algorithms have been applied extensively in the past to solve biological problems due to the analogy between protein chains and robotic articulated mechanisms [23], [24], [25]. The search methodology applied in this paper is based on the Path-Directed Subdivision Tree (PDST) planner [35], [36]. We chose this algorithm because of its good performance with articulated systems with complex

dynamics moving in physically constrained environments. We adapted the algorithm to model protein motions. In our adaptation, the planner iteratively constructs a tree of conformational pathways as the search progresses. The input to the algorithm consists of the start and end conformations of a molecule, represented as sets of articulated secondary structures as discussed in Data Representation above. The root of the search tree is a “pathway” of length 0 consisting only of the starting structure. At every iteration a previously generated pathway is selected for propagation using a deterministic scoring scheme described below. From a random conformation along that pathway, a new pathway is propagated by applying a small random rotation to the ϕ or ψ backbone dihedral angle of a residue that resides on a loop connecting two randomly chosen secondary structure elements. A molecular motion is sampled by applying the rotation until a high energy conformation is reached. The coarse grained energy function described above is used to determine when a high energy conformation is encountered. A high energy conformation is defined as being more than 50 kcal/mol above the starting energy. The algorithm maintains a subdivision of the low-dimensional projection of the conformational space (described in Distance Between Structures above) into cells, such that no sample spans more than one cell in the subdivision. The goal of the subdivision is to guarantee coverage of the search space [35]. After a sample is selected for propagation, the cell containing that sample is subdivided into two cells. The algorithm keeps track of how many samples are contained in each cell to estimate how dense the sampling is in different areas of the space. It maintains a scoring scheme that gives selection preference to samples residing in large, empty cells, thus pushing the exploration towards unvisited areas in the conformational space. Probabilistic completeness is obtained via a scoring scheme that favors the selection of samples contained in larger cells and leads to unexplored areas of the search space. The sample scores are updated in a way that guarantees that every sample in the tree will eventually be selected for propagation and avoids over-sampling of parts of the space. A previous study in path-directed motion planning algorithms [37] showed that employing a biasing scheme in a small percentage of the iterations greatly improves the performance of the planner. Motivated by these results [37], we employed biasing at 10%

of the iterations. During these iterations the scoring scheme described above is ignored and a sample is chosen out of a pool of conformations closest to the goal conformation, which gives the planner a better chance to successfully terminate the search. We found that the biasing improves the performance of the algorithm. Our top level algorithm runs PDST iteratively. Each iteration runs until a generated conformation is closer to the goal conformation than a pre-specified intermediate distance threshold, where the distance threshold is determined by the distance measure described above. We found that a threshold of 0.8-0.9 of the distance between the start and goal conformations is usually sufficient to achieve good results. The iterative runs of the PDST planner help reduce memory use and improve performance, as also shown in [38]. To produce the results shown in this paper, three PDST cycles, each of 20000 iterations, were allowed per run of the algorithm for each example.

III. RESULTS AND DISCUSSION

We ran the PDST-based search algorithm mentioned above on three test cases, Adenylate Kinase (AdK), Ribose binding protein (RBP) and the GroEL monomer. These proteins have been chosen for the following reasons: all undergo extensive conformational transitions, they are well studied and have an abundance of data for testing and comparison.

For comparison purposes, we produced conformational pathways using a random walk Monte Carlo like algorithm [3]. In order to make the two methods as comparable as possible, we used the same representation, similarity score and potential function described in our algorithm. The random walk algorithm differs from the common use of Monte Carlo in protein conformational search. Rather than optimizing the energy, it optimizes the similarity score (see Distance between Structures subsection under Methods for definition) in order to simulate a conformational pathway from the start to the goal conformation. The energy, while not optimized, is used to filter out non-feasible conformations. The random walk implementation uses the Metropolis criterion for the selection of steps. At each iteration a random conformational pathway is generated from the current conformation by applying a small random transformation to either the ϕ or ψ dihedral angle of one of the degrees of freedom connecting secondary structure elements, in a similar way to the one used to generate new conformations described in the Search Methodology subsection above. If a step brings the similarity score of the generated conformation closer to the goal it will be accepted. Otherwise it is accepted with a probability proportional to $e^{\Delta S}$ where ΔS is the difference in the similarity score of the current step and the previous step. In practice, this criterion accepts all “good” steps while allowing a very small fraction of “bad” steps.

In order to compare the performance of the two methods by an objective standard, each algorithm was run a 100 times per example and the least RMSD (IRMSD) of the closest conformation to the goal at that given time step was measured. IRMSD is the root mean square deviation between two conformations after alignment. In our implementation, only C- α atoms were considered for the IRMSD measurement.

IRMSD was measured after 60 and 120 minutes and in the end of the run. All runs were allowed to continue for a maximum of 8 hours or until a generated conformation is closer to the goal structure than a specified threshold, varying according to the tested protein. on the Rice Cray XD1 Cluster, where each node runs at 2.2 Ghz and has 8 GB RAM. Table I summarizes the IRMSD statistics over 80 of the 100 test runs for each algorithm and protein test case, having the top and bottom 10% outliers excluded from the calculation.

A. Adenylate kinase (AdK)

AdK is a monomeric phosphotransferase enzyme that catalyzes reversible transfer of a phosphoryl group from ATP to AMP. The structure of AdK, which contains 214 amino acids, is composed of the three main domains, the CORE (residues 1–29, 68–117, and 161–214), the ATP binding domain called the LID (residues 118–167), and the NMP binding domain (residues 30–67). AdK assumes an “open” conformation in the unligated structure and a “closed” conformation. The IRMSD between the two structures is 6.95Å. Supposedly, during the transition from the “open” to “closed” form, the largest conformational change occurs in the LID and NMP domain with the CORE domain being relatively rigid. Our model contains 10 rigid elements where most of the CORE domain was modeled as one large segment and was considered fixed, since it does not undergo a large-scale motion. The distance measure threshold for successful termination of the algorithm was 0.84 of the distance between the start and the goal conformation. Figure 2(a) shows an example of a pathway from the start to the end conformation. The C α RMSD from the goal structure is 2.07Å. As seen in table I the resulting average IRMSD was 2.53Å. Random walk performed significantly worse compared to our planner with an average IRMSD of 3.65Å. The average running time was 3 hours, 58 minutes.

B. Ribose binding protein (RBP)

RBP is a sugar-binding bacterial periplasmic protein whose function is associated with large conformational changes upon binding to ribose. It is a 271 residue protein made of two domains, the first containing residues 1–99 and 238–260 and the second containing residues 104–233. The domains are linked by a three stranded hinge spanning residues 100–103, 234–237 and 261–271. The IRMSD between the two conformations is 4.06Å. We modeled the closed state to open state motion using PDB codes 2DRI and 1URP for the closed and open states respectively. Our model contains 3 rigid elements where most of the N- and C-terminal domains were modeled as rigid segments and the hinge was modeled as a separate domain. The distance measure threshold for successful termination of the algorithm was 0.92 of the distance between the start and the goal conformation. As seen in table I, the resulting average RMSD was approximately 1.38Å. Random walk performed poorly comparing to our planner and the average RMSD in the end of the run was 2.59Å. In this example, as well as the AdK example above, the vast majority of the progress was achieved during the first 60 minutes of the run. Figure 2(b) shows an

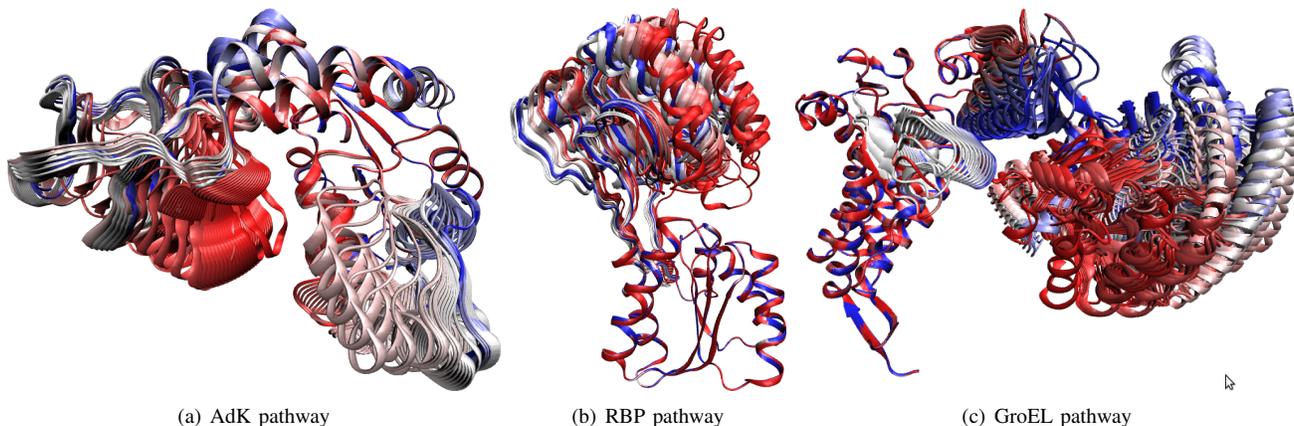


Fig. 2. Illustration of the results for AdK (a), RBP (b) and GroEL (c): The conformational pathways are obtained after side chain completion and basic energy minimization. The conformation colors are interpolated on the red (start) to blue (goal) scale.

example of a pathway from the start to the end conformation. The $C\alpha$ RMSD from the goal structure is 0.76\AA . The average run time was approximately 1 hour and 40 minutes.

C. GroEL monomer

The GroEL protein belongs to the chaperonin family and is found in a large number of bacteria [39]. It is required for the correct folding of many proteins. GroEL requires the lid-like cochaperonin protein complex GroES. Binding of substrate protein, in addition to binding of ATP, induces an extensive conformational change that allows association of the binary complex with GroES. We modeled the epical domain movement from the GroEL monomer (modeled from chain A of PDB code 1SS8) to the GroEL-GroES-ADP7 monomer (modeled from chain A of PDB code 1SX4). The monomer contains 525 amino acids, and our model contains 13 rigid elements where most of the equatorial domain, whose structure does not change significantly, was modeled as one large segment and was considered fixed. The distance measure threshold for successful termination of the algorithm was 0.82 of the distance between the start and the goal conformation. The initial IRMSD between the $C\alpha$ atoms of the two monomers is 12.21\AA . Table I shows that our method significantly outperforms random walk both in runtime and average IRMSD. The average IRMSD between the resulting structures and the goal structure was 4.04\AA after two hours compared to 6.17\AA for MC. Many runs produced low IRMSD results in the order of magnitude of $3\text{--}4\text{\AA}$ RMSD or less from the goal structure. The average run time was approximately 5 hours. It should be noted that the conformational pathway for the monomer may not be indicative of the pathway for the entire 7 membered GroEL ring, but the experiment in this study was done for testing purposes. Producing a pathway for the entire GroEL complex is the subject of on-going work.

D. Analysis of the Results

1) *Potential Energy Measurement*: In order to provide initial validation for our results, we tested whether our algorithm produces biologically reasonable, low energy pathways when using an all-atom force field. Such an analysis was done in

an earlier work [29], where the authors used a similar method to show that their conformational search was reasonable. Side chain information was completed for the resulting pathways using the algorithm described in [40]. The resulting full-atomic structures were minimized for 1000 Steepest Descent steps using the AMBER energy minimization package [2] and subject to a harmonic restraining force of 10 kcal/mol/\AA^2 . The minimization was done for a relatively small number of steps and was restrained in order to resolve initial clashes but not cause large conformational changes to the structures. The purpose of this test is not to provide a fully minimized pathway, but to show that the algorithm produces pathways with reasonable conformations whose clashes can be resolved within a small number of minimization steps. Figure 2 shows an example of a pathway for AdK, RBP and GroEL and Figure 3(a-c) shows the potential energy plots of the corresponding pathways. In each case the pathway chosen for figures 2 and 3 (a-c) was the result of the run with the lowest final IRMSD from the goal structure. For clarity, the conformations shown in the figures were sampled at approximately 1 distance measure unit from one another (see Distance Between Structures section for definition). As seen, even with a small number of energy minimization steps all the intermediate structures exhibit low potential energies, below -6000 kcal/mol for AdK, below -7000 kcal/mol for RBP and around -15000 kcal/mol for GroEL, as measured by AMBER.

2) *Free Energy Profile for AdK*: To provide further evidence that the produced paths are reasonable, we refer to a study [41] which provided an extensive analysis of the conformational pathway of AdK. The authors generated a conformational pathway using a Nudged Elastic Band (NEB) simulation [15]. Their large-scale analysis of the pathway included a free energy profile using umbrella sampling over a number of reaction coordinates. One of the reaction coordinates used for the free energy calculation was ΔD_{RMSD} which is defined, given conformation C , as:

$$\Delta D_{\text{RMSD}}(C) = \text{RMSD}(C, C_{\text{open}}) - \text{RMSD}(C, C_{\text{closed}}) \quad (4)$$

We characterized the free energy profile along this reaction coordinate using our results. The data points were obtained

TABLE I

PERFORMANCE STATISTICS FOR THE AdK, RBP AND GroEL MONOMER EXAMPLES. THE AVERAGE \pm (STANDARD DEVIATION) LRMSD DATA WERE TAKEN OVER 80 RUNS WHERE THE TOP AND BOTTOM 10% OUTLIERS WERE REMOVED FROM THE ORIGINAL SET OF 100 RUNS.

| | AdK | AdK RW [†] | RBP | RBP RW [†] | GroEL | GroEL RW [†] |
|-------------------------|-----------------|---------------------|-----------------|---------------------|------------------|-----------------------|
| Initial IRMSD (Å) | 6.95 | 6.95 | 4.06 | 4.06 | 12.21 | 12.21 |
| Residues | 214 | 214 | 271 | 271 | 525 | 525 |
| IRMSD after 1 hour.(Å) | 2.69 \pm 0.21 | 3.81 \pm 0.49 | 1.48 \pm 0.25 | 2.35 \pm 0.52 | 4.68 \pm 0.467 | 6.39 \pm 1.7 |
| IRMSD after 2 hours.(Å) | 2.55 \pm 0.2 | 3.68 \pm 0.46 | 1.34 \pm 0.21 | 2.23 \pm 0.45 | 4.26 \pm 0.73 | 6.33 \pm 1.67 |
| Final IRMSD (Å) | 2.53 \pm 0.2 | 3.65 \pm 0.47 | 1.26 \pm 0.15 | 2.22 \pm 0.49 | 4.04 \pm 0.9 | 6.17 \pm 1.61 |

[†] Random walk. See Results section for a discussion.

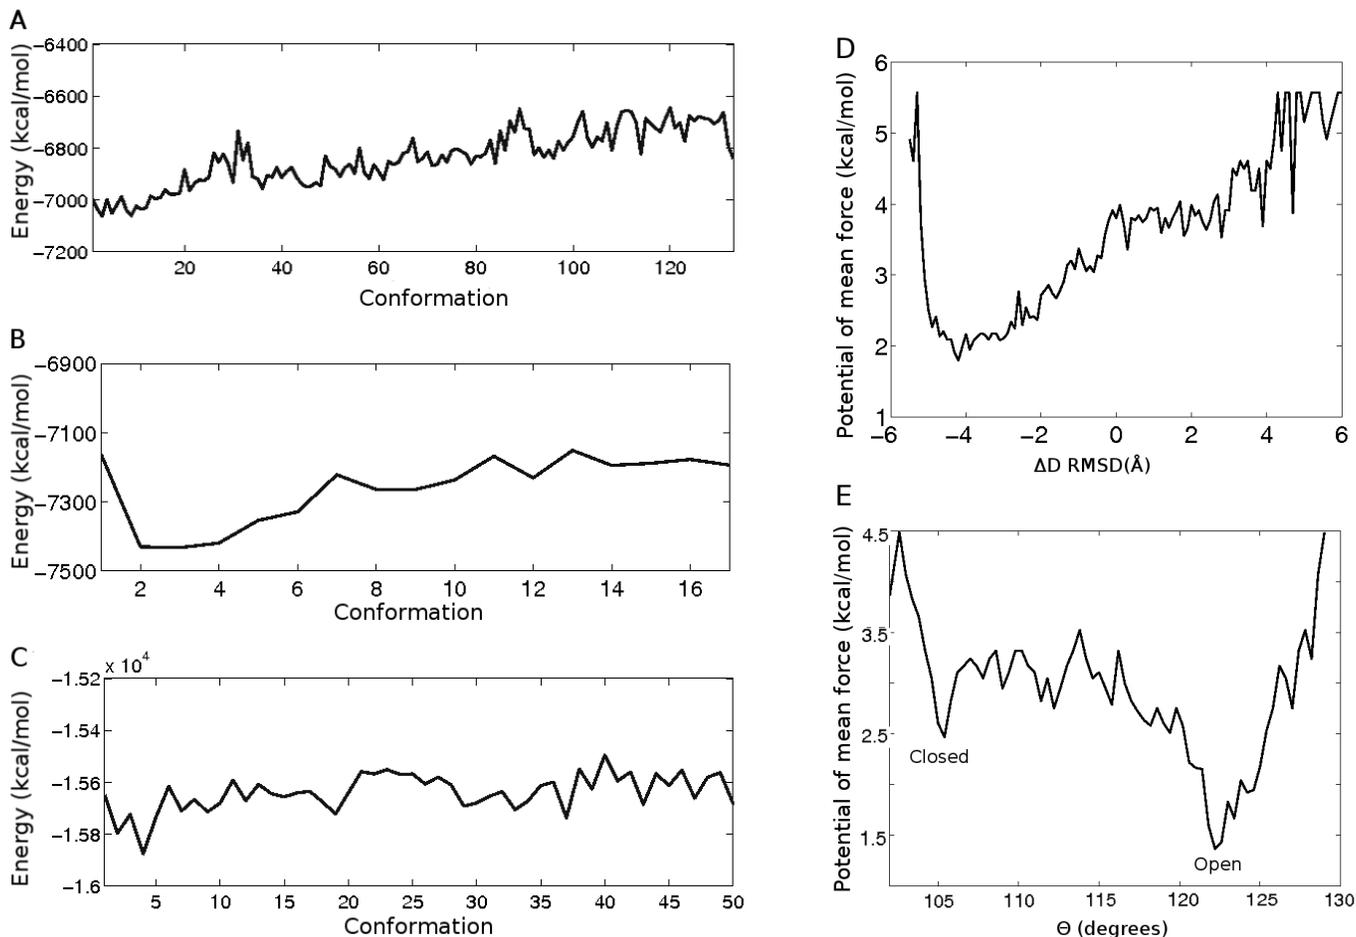


Fig. 3. Energetic profiles of the resulting pathways: A potential energy plot for AdK (A), RBP (B) and GroEL (C) slightly minimized conformational pathways. Notice the different potential energy scale and different path lengths. Free energy along the ΔD_{RMSD} reaction coordinate of the AdK pathway (D) and along the θ reaction coordinate for the RBP (E) pathway. See Results section for the definition of the reaction coordinates.

by running the algorithm on AdK for 200 times. For each resulting pathway we recorded the ΔD_{RMSD} value for the conformations along the pathway. To generate sets of uncorrelated conformations as required for free energy calculations, we sampled each pathway in spaces of 1 distance unit (see definition of the distance measure in the Methods section). Overall approximately 7500 conformations were included in the calculation. The free energy was calculated along the ΔD_{RMSD} reaction coordinate using the Weighted Histogram Analysis Method (WHAM) [42]. It should be noted that the calculation was carried out under a number of assumptions: we used only backbone and C- β and a relatively small number of samples. Therefore, our “pseudo free energy”

results should be interpreted with caution. Also, our sampling method and potential of mean force calculation parameters differ significantly from the ones used in [41]. For these reasons, we can expect only qualitative similarity to the free energy profile obtained by that work and the absolute free energy values do not have the same meaning. The free energy profile shown in Figure 3(d) exhibits a qualitatively similar pattern to that shown in Figure 2(a) in [41] for the free conformational pathway: high free energy around a ΔD_{RMSD} of 3 to 6 (closed conformation), and a low energy basin around the open conformation, at ΔD_{RMSD} of -5 to -4 . The spikes shown in the profile are the result of a relatively small number of samples and non-uniform sampling at some

areas in the search space, whereas NEB provides an initially uniform interpolation. These results show that the sampling the algorithm provides along the conformational pathway is qualitatively similar to the one provided by NEB.

3) *Free Energy Profile for RBP*: To provide further validation of our results we compare with another study which analyzed RBP [43]. The authors simulated the opening motion of the RBP protein and characterized the free energy profile using the reaction coordinate θ , which is the angle between the two domain, defined as the angle formed by the following three points: the center-of-mass (CM) of the N-terminal domain, the CM of the C-terminal domain and the CM of the hinge. The values of θ are 109 and 130 in the closed and open conformation, respectively. Our free energy calculations as a function of θ were conducted in a similar manner to the calculations described above for ADK. The result is shown in the bottom of Figure 3(e). Two minima are shown: one local minimum around 106 degrees and one global minimum at 123 degrees, very similar to the pattern shown in Figure 3(a) of [43]. It should be noted that we did not simulate the RBP mutant pathway discussed in [43], and therefore our plot ends at approximately 130 degrees.

In general, knowledge about intermediate states is needed in order to provide a case-specific validation, but this knowledge does not always exist. With the advances in structural detection and simulation methods, one can expect to have more information about intermediate states in the future. It should be noted that several intermediate structures already exist for AdK and a recent study makes use of those structures to validate their low energy profile calculations [44]. This is an important way to validate computational results and is the subject of present and future work. In cases where such information is not available, this algorithm can be viewed as an efficient initial filtering tool that reduces the tremendously high-dimensional space of possible conformations into a relatively small number of possible pathways. Refinement can then be made by other tools or indirect experimental knowledge to select biologically feasible pathways out of these possibilities. In the future we plan to apply clustering methods on the resulting pathways to extract more knowledge about feasible conformations and gain insight about the likelihood of each conformation along the resulting pathways.

IV. CONCLUSIONS

We present a prototype for a novel method for exploring large scale conformational changes in proteins represented at the backbone level, using relatively little information. The search methodology is based on robot motion planning, and it strikes a balance between an efficient coverage of the conformational space and fast exploration towards the goal structure. A relatively simple potential function is used to guide the search. This representation and potential function make the computation tractable and especially useful in cases where side chain information is missing or if a detailed search is computationally infeasible. The goal of this paper is to provide an initial proof of concept for our method. Therefore, we tested our algorithm on the following three well

studied proteins: the GroEL monomer, Ribose binding protein and Adenylate Kinase. We show that our method performs significantly better than random walk by producing low energy pathways with resulting structures closer to the goal structure. We believe this is an important first step towards a larger scale modeling of more complex biological systems such as the entire GroEL complex.

Acknowledgements

We thank Dr. C. Clementi, Dr. T. Ju, Dr. E. Plaku, Dr. A. Shehu, A. Heath and I Şucan. This work was supported in part by NIH grant No. GM078988. Computational experiments were conducted on the Rice Computational Research Cluster funded by the NSF under grant No. CNS-0421109 and grant No. CNS-0454333, and a partnership between Rice University, AMD and Cray.

REFERENCES

- [1] M. Gerstein, A. M. Lesk, and C. Chothia, "Structural mechanisms for domain movements in proteins," *Biochemistry*, vol. 33, pp. 6739–6749, 1994.
- [2] D. A. Case, T. Cheatham, T. Darden, H. Gohlke, R. Luo, K. M. Merz Jr., A. Onufriev, C. Simmerling, B. Wang, and R. Woods, "The Amber biomolecular simulation programs," *J. Computat. Chem.*, vol. 26, pp. 1668–1688, 2005.
- [3] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, pp. 671–680, 1983.
- [4] T. Head-Gordon and S. Brown, "Minimalist models for protein folding and design," *Curr. Opin. Struct. Biol.*, vol. 13, no. 2, pp. 160–167, 2003.
- [5] P. C. Whitford, O. Miyashita, Y. Levy, and J. N. Onucic, "Conformational transitions of adenylate kinase: Switching by cracking," *Journal of Molecular Biology*, vol. 366, no. 5, pp. 1661–1671, 2007.
- [6] O. Miyashita, P. G. Wolynes, and J. N. Onucic, "Simple energy landscape model for the kinetics of functional transitions in proteins," *Journal of Physical Chemistry B*, vol. 109, no. 5, pp. 1959–1969, 2005.
- [7] G. Schroeder, A. T. Brunger, and M. Levitt, "Combining efficient conformational sampling with a deformable elastic network model facilitates structure refinement at low resolution," *Structure*, vol. 15, pp. 1630–1641, 2007.
- [8] A. Schuyler, R. Jernigan, P. Qasba, B. Ramakrishnan, and G. Chirikjian, "Iterative cluster-nma: A tool for generating conformational transitions in proteins," *Proteins*, vol. 74, pp. 760–776, 2009.
- [9] W. Zheng and B. Brooks, "Identification of dynamical correlations within the myosin motor domain by the normal mode analysis of an elastic network model," *J. Mol. Biol.*, vol. 346, no. 3, pp. 745–759, 2005.
- [10] O. Keskin, I. Bahar, D. Flatow, D. Covell, and R. Jernigan, "Molecular mechanisms of chaperonin groel-groes function," *Biochemistry*, vol. 414, pp. 491–501, 2002.
- [11] Y. Wang, A. Rader, I. Bahar, and R. Jernigan, "Global ribosome motions revealed with elastic network model," *J. Struct. Biol.*, vol. 147, no. 3, pp. 302–314, 2004.
- [12] N. Temiz, E. Meirovitch, and I. Bahar, "Escherichia coli adenylate kinase dynamics: comparison of elastic network model modes with mode-coupling (15)n-nmr relaxation data," *Proteins*, vol. 57, pp. 468–480, 2004.
- [13] H. Gohlke and M. Thorpe, "A natural coarse graining for simulating large biomolecular motion," *Biophysical Journal*, vol. 9, pp. 2115–2120, 2006.
- [14] M. Thorpe, "Comment on elastic network models and proteins," *Phys. Biol.*, vol. 4, pp. 60–63, 2007.
- [15] H. Jónsson, G. Mills, and K. W. Jacobsen, "Nudged elastic band method for finding minimum energy paths of transitions," in *Classical and Quantum Dynamics in Condensed Phase Simulations*, B. J. Berne, G. Ciccoti, and D. F. Coker, Eds. Singapore: World Scientific, 1998, pp. 385–404.
- [16] D. Weiss and M. Levitt, "Can morphing methods predict intermediate structures?" *J. Mol. Biol.*, vol. 385, pp. 665–674, 2009.
- [17] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, 2005.
- [18] S. M. LaValle and J. J. Kuffner, "Randomized kinodynamic planning," *Intl. J. Robotics Research*, vol. 20, no. 5, pp. 378–400, May 2001.

- [19] L. E. Kavradi, P. Švestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, pp. 566–580, 1996.
- [20] A. Shehu, L. Kavradi, and C. Clementi, "On the characterization of protein native state ensembles," *Biophysical Journal*, vol. 92, no. 5, pp. 1503–1511, 2007.
- [21] J. Cortes, T. Siméon, V. Ruiz de Angulo, D. Guieysse, M. Remaud-Siméon, and V. Tran, "A path planning approach for computing large-amplitude motions of flexible molecules," *3th Annual International conference on Intelligent Systems for Molecular Biology (ISMB)*, 2005.
- [22] S. Kirillova, J. Cortes, A. Stefaniu, and T. Siméon, "An NMA-guided path planning approach for computing large-amplitude conformational changes in proteins," *Proteins: Structure, Function and Bioinformatics*, vol. 70, pp. 131–143, 2008.
- [23] S. Thomas, X. Tang, L. Tapia, and N. M. Amato, "Simulating protein motions with rigidity analysis," *J. Comp. Biol.*, vol. 14, no. 6, pp. 839–855, 2007.
- [24] S. Thomas, G. Song, and N. M. Amato, "Protein folding by motion planning," *Phys. Biol.*, vol. 2, pp. S148–S155, 2005.
- [25] T. H. Chiang, M. S. Apaydin, D. L. Brutlag, D. Hsu, , and J.-C. Latombe, "Using stochastic roadmap simulation to predict experimental quantities in protein folding kinetics," *J. Comp. Biol.*, vol. 14, no. 5, pp. 578–593, 2007.
- [26] P. Yao, A. Dhanik, N. Marz, R. Propper, C. Kou, G. Liu, H. van den Bedem, J.-C. Latombe, I. Halperin-Landsberg, and R. B. Altman, "Efficient algorithms to explore conformation spaces of flexible protein loops," *IEEE/ACM Trans Comput Biol Bioinform.*, vol. 5, no. 4, pp. 534–545, 2008.
- [27] J. Cortés, T. Siméon, V. R. de Angulo, D. Guieysse, M. Remaud-Siméon, and V. Tran, "A path planning approach for computing large-amplitude motions of flexible molecules," *Bioinformatics*, vol. 21 Suppl. 1, pp. i116–i125, 2005.
- [28] X. Tang, S. Thomas, L. Tapia, and N. M. Amato, "Tools for simulating and analyzing rna folding kinetics," in *Proc. Int. Conf. Comput. Molecular Biology (RECOMB)*, San Francisco, CA, USA, April 2007, pp. 268–282.
- [29] B. Raveh, A. Enosh, O. Furman-Schueler, and D. Halperin, "Rapid sampling of molecular motions with prior information constraints," *Plos Comp. Biol.*, vol. in press, 2009.
- [30] A. Shehu, L. Kavradi, and C. Clementi, "Multiscale characterization of protein conformational ensembles," *Proteins: Structure, Function and Bioinformatics*, 2009.
- [31] W. Kabsch and C. Sander, "Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features," *Biopolymers*, vol. 22, pp. 2577–2637, 1983.
- [32] P. J. Ballester and W. G. Richards, "Ultrafast shape recognition to search compound databases for similar molecular shapes," *J. Comput. Chem.*, vol. 28(10), pp. 1711–1723, 2007.
- [33] Y. Duan, C. Wu, S. Chowdhury, M. Lee, G. Xiong, W. Zhang, R. Yang, P. Cieplak, R. Luo, and T. Lee, "A point-charge force field for molecular mechanics simulations of proteins based on condensed-phase quantum mechanical calculations," *J. Comput. Chem.*, vol. 24, pp. 1999–2012, 2003.
- [34] G. Arfken, *The Method of Steepest Descents*, 3rd ed. Academic Press, 1985, pp. 428–436.
- [35] A. M. Ladd, "Motion planning for physical simulation," Ph.D. dissertation, Dept. of Computer Science, Rice University, Houston, TX, Dec. 2006.
- [36] K. I. Tsianos, I. A. Şucan, and L. E. Kavradi, "Sampling-based robot motion planning: Towards realistic applications," *Computer Science Review*, vol. 1, no. 1, pp. 2–11, August 2007.
- [37] I. A. Şucan, J. F. Kruse, M. Yim, and L. E. Kavradi, "Reconfiguration for modular robots using kinodynamic motion planning," in *ASME Dynamic Systems and Control Conference*, Ann Arbor, Michigan, October 2008.
- [38] K. Tsianos and L. E. Kavradi, "Replanning: A powerful planning strategy for hard kinodynamic problems," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nice, France, September 2008, pp. 1667–1672.
- [39] J. Zeilstra-Ryalls, O. Fayet, and C. Georgopoulos, "The universally conserved GroE (Hsp60) chaperonins," *Annu. Rev. Microbiol.*, vol. 45, pp. 301–325, 1991.
- [40] A. P. Heath, L. E. Kavradi, and C. Clementi, "From coarse-grain to all-atom: Toward multiscale analysis of protein landscapes," *Proteins: Structure, Function and Bioinformatics*, vol. 68, no. 3, pp. 646–661, August 2007.
- [41] K. Arora and C. L. Brooks III, "Large scale allosteric conformational transitions of adenylate kinase appear to involve a population-shift mechanism," *Proc. Nat. Acad. Sci.*, vol. 104, no. 47, pp. 18 496–18 501, 2007.
- [42] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, "The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method," *J. Comput. Chem.*, vol. 13, no. 8, pp. 1011–1021, 1992.
- [43] K. P. Ravindranathan, E. Gallicchio, and R. Levy, "Conformational equilibria and free energy profiles for the allosteric transition of the ribose-binding protein," *J. Mol. Biol.*, vol. 353, pp. 196–210, 2005.
- [44] Y. Feng, L. Yang, A. Kloczkowski, and R. Jernigan, "The energy profiles of atomic conformational transition intermediates of adenylate kinase," *Proteins*, vol. in press, 2009.