



Efficient similarity search in protein structure databases by *k*-clique hashing

Nils Weskamp^{1,2}, Daniel Kuhn², Eyke Hüllermeier^{1,*} and Gerhard Klebe²

¹Department of Mathematics and Computer Science, University of Marburg, Hans-Meerwein-Straße, 35032 Marburg, Germany and ²Institute of Pharmaceutical Chemistry, University of Marburg, Marbacher Weg 6, 35032 Marburg, Germany

Received on October 12, 2003; accepted on February 3, 2004

ABSTRACT

Motivation: Graph-based clique-detection techniques are widely used for the recognition of common substructures in proteins. They permit the detection of resemblances that are independent of sequence or fold homologies and are also able to handle conformational flexibility. Their high computational complexity is often a limiting factor and prevents a detailed and fine-grained modeling of the protein structure.

Results: We present an efficient two-step method that significantly speeds up the detection of common substructures, especially when used to screen larger databases. It combines the advantages from both clique-detection and geometric hashing. The method is applied to an established approach for the comparison of protein binding-pockets, and some empirical results are presented.

Availability: Upon request from the authors.

Contact: eyke@mathematik.uni-marburg.de

INTRODUCTION

Most ambitious among the competing techniques for the structural comparison of proteins are those that operate completely independent of sequence and fold information. These techniques are capable of revealing functional relations among proteins that are not due to phylogenetic dependency.

Unfortunately, such approaches are computationally very complex and require a reduction of the problem size, which is usually achieved by considering only the secondary structure elements of the protein (Harrison *et al.*, 2002; Koch *et al.*, 1996; Grindley *et al.*, 1993). Our approach is instead based on a recent method for the automatic extraction and description of protein binding sites by Schmitt *et al.* (2002). Following the idea that the function of a protein is determined by the shape and the physicochemical properties of the binding pocket, we assign pseudocenters—as three-dimensional (3D)-descriptors—to the cavity-flanking amino acids to represent their interaction properties. Two binding sites are regarded as

similar if they share a common spatial arrangement of assigned pseudocenters and expose similar physicochemical properties into the binding site. The binding site-composing amino acids, the attributed pseudocenters and a cavity surface are stored in the database Cavbase (Schmitt *et al.*, 2002).

The original approach by Schmitt *et al.* (2002) uses a clique-detection method that interprets the pseudocenters as nodes of a graph. Typically, this results in graphs of the size 50–150 (yet graphs with more than 1000 nodes do exist). The physicochemical properties of the centers are modeled as a coloring of the nodes. An edge is inserted between two nodes, weighted with the geometric distance between the adjacent nodes. To reduce the complexity, an edge is inserted only if its length does not exceed 12.0 Å. A standard algorithm (Bron and Kerbosch, 1973) is used to detect connected maximal common substructures in the graph representations, applying the modifications suggested in the study of (Koch *et al.*, 1996). The size of the common substructures is our measure of similarity. Additionally, the common substructures may be used to calculate a geometric superimposition of the structures and to perform more complex similarity measurements based on mutual surface–surface matches of common cavity patches. Although both these measures do not formally fulfill the properties of a metric, they were subjected to empirical validation and usually revealed good results. The implications for clustering are not in the scope of this paper and will be discussed elsewhere.

Due to the comparatively fine-grained and therefore large graph representations, the runtime complexity of the method is relatively high. To allow efficient similarity searches in large datasets and all-against-all comparisons as a basis for clustering, we developed an improved hybrid method for the detection of common substructures in binding sites. This method combines advantages from both, graph-based clique-detection and geometric hashing, as will be detailed in the following. Before that we recall the standard clique-detection approach and point out some of its drawbacks. Finally, we shall present some empirical results.

*To whom correspondence should be addressed.

METHODS

Standard clique-detection techniques

Clique-detection is a well-known strategy for the structural comparison of proteins or drug-sized organic molecules. Given two labeled input graphs G_1 and G_2 , the first step is the construction of a so-called product graph P (cf. Fig. 1). Each node of P consists of a pair of nodes with identical labels from the input graphs, thus representing an ‘isomorphism’ (matching) of size 1 between the two graphs. Two nodes (u, v) and (u', v') of P are connected by an edge if the respective matchings are compatible, i.e. the edge connecting u and u' in G_1 (if any) has the same label as the edge connecting v and v' in G_2 . As shown by Levi (1972), a maximal complete subgraph (clique) of P corresponds bijectively to a maximal subgraph isomorphism between G_1 and G_2 . Thus, it is possible to detect a common substructure of G_1 and G_2 using the clique-detection algorithm of Bron and Kerbosch (1973). In the example of Figure 1 (left), this yields the mapping $A-a$, $B-b$ and so on.

Even though clique-detection methods provide a powerful tool for the structural comparison of proteins, they possess some major drawbacks: first, depending on the size of the input graphs and the distribution of the different label types, the product graph soon becomes relatively large and densely connected, which in turn leads to huge running times of the applied algorithms. This problem is particularly severe in our case, as we are indeed interested in the comparison of large graphs. Second, it is difficult to estimate the result of a comparison in advance. Hence, it is necessary to perform many superfluous calculations, as it is impossible to detect cases in which structures are obviously unrelated. In the next section, we propose a new hybrid approach that overcomes these limitations at least to some extent.

Clique-detection based on clique hashing

One reason for the large size of the product graph, P , is the fact that false-positive matches are very likely to occur for single nodes of the input graphs: P will usually contain a large number of nodes having the same label just by chance. Take for instance the node A of graph G_1 , which is isomorphic to node a of graph G_2 (Fig. 1). Apart from this isomorphism, the mapping $A-e$ has to be considered as well, as e has the same color as A . But since A and e have a completely different environment (nodes G and B may not be mapped to d and f), this is actually not necessary. The clique-detection algorithm spends most of its running time on the elimination of such false-positive matches and the assembling of real matches into larger matchings.

Intuitively, assembling individual pieces becomes simpler when starting from larger pieces. Our method therefore starts from a modified product graph P' , whose nodes represent larger local matches of size k . This substantially reduces the probability of false-positive matches and leads to smaller

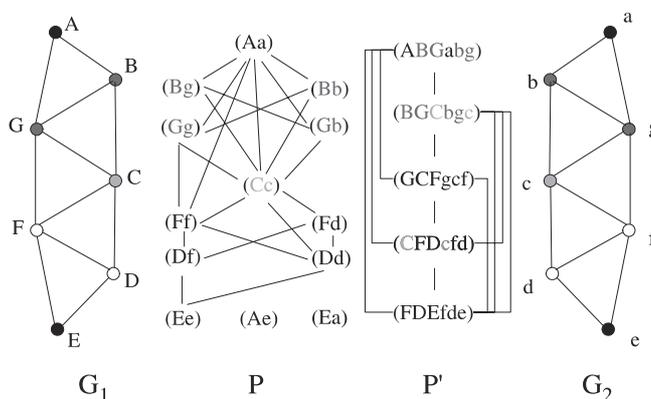


Fig. 1. Search for common substructures of G_1 and G_2 by clique-detection. P shows the standard product graph, P' the much simpler modified product graph for $k = 3$.

product graphs (Fig. 1, right). At the same time, the information content of the local matches increases. Thus, it can be hoped that the number of local matches allows an estimation of the similarity of the complete structures under consideration.

Generation of local matches

Our approach is motivated by the geometric-hashing technique (Nussinov and Wolfson, 1991). The input graphs are split into a large number of k -cliques, complete subgraphs of a fixed size k . As complete graphs of a fixed size are always isomorphic, each substructure is defined uniquely by its node and edge labels. Hence, to search for similar substructures, it is sufficient to consider only those labels. The labels are then mapped onto points in Euclidean space by using a normalization scheme, whereas the node coloring is mapped onto a discrete attribute and the edge weights are represented directly. These points are then stored in an R^* -tree (Beckmann *et al.*, 1990), a standard spatial index structure for external memory.

Once this index structure is built for all structures of the dataset in a preprocessing stage, it is possible to perform an efficient similarity search. Given a query structure, for all k -cliques in its graph representation, Q , a window query in the index is performed, resulting in a number of k -cliques with the same node coloring and a similar (up to a parameter ϵ) edge weighting. Each such hit indicates a local match (‘isomorphism’ of size k) between query structure Q and a hit structure H from the indexed dataset. The number of local matches found for each hit structure allows a rough estimation of the similarity between Q and H . Moreover, the local matches may be used to build the modified product graph mentioned earlier. Note that the decomposition of the graph structures is highly redundant, and hence it is neither necessary to use all possible k -cliques for the index generation nor is it required to launch all possible queries associated with the k -cliques of the query pocket. Instead, random sampling strategies may be

used to reduce the index size and to accelerate the querying process.

A critical point concerns the selection of the parameter k , which determines the size of the substructures under consideration. On the one hand, it would be preferable to choose a high value of k to reduce the probability of false-positive matches. On the other hand, the number of k -cliques in a graph depends exponentially on k , and thus large values of k lead to huge index structures and immense numbers of queries that have to be performed. To overcome this problem, we developed a filtering step for the identification of obviously false-positive chance matches. Using this technique, which is detailed below, it is possible to use relatively small values for k without being overwhelmed by a huge amount of insignificant hits. For our experiments, we used $k = 3$.

Assembly of local matches

The clique-detection technique for the modified product graph intuitively tries to assemble the local matches into larger matchings and is thus somewhat related to the problem of DNA fragment assembly. As in fragment assembly, the assembling is guided by overlaps. We define a neighborhood relation for the substructures of a structure as follows: two k -cliques of a graph descriptor are overlapping (neighbored) if and only if they share at least one common node. If two substructures of the query structure are overlapping, only those hits for these structures that are also overlapping are of interest. Other hits can be discarded: they cannot be merged into larger matchings since this would lead to an ambiguous mapping of graph nodes. Figure 2a shows an example of two matches that do not maintain the neighborhood relation: the three-clique BFG and CEF are neighbors as they share the common node F . The mapping of BFG to abg and of CEF to cde as indicated by the arrows is discarded as abg and cde are not overlapping and thus a merging of those local matchings would lead to an ambiguous mapping of F .

This filtering step is carried out using a novel construct called a hit list (Fig. 2b). Each node of the query graph is endowed with such a list. The latter collects all possible matching partners for the respective query node during the index querying process. Each time a hit is found for a k -clique containing the respective node, this hit ‘votes’ for the mapping of the respective node to the associated node of an indexed structure. This ‘voting’ is stored in the hit list. At the end of the whole querying process, each hit list contains a number of possible mappings of the respective node to different nodes of structures from the indexed dataset. Mappings with a high number of ‘votes’ are supported by a large number of overlapping hit substructures because these hits all share a common node (i.e. the one involved in the respective mapping). Such high-scoring hit list entries are therefore likely to participate in a larger local matching. In other words, there is a low probability that these entries are false-positive matches. In the example of Figure 2b, there are five votes for the mapping of

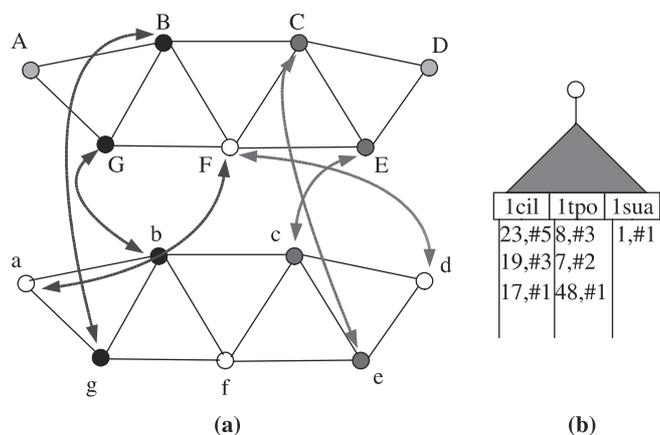


Fig. 2. (a) A combination of two false-positive matches, indicating an ambiguous mapping of F to a and d . (b) The hit list for the comparison of one pseudocenter against a dataset of three binding pockets.

the respective node to node 23 of structure *1cil* but only one vote for the mapping to node 17. Thus, a mapping to node 23 is more likely. Note that the whole filtering process depends only on the query structure and is independent of the hit structure. Thus, queries may be processed using solely information from the query and the index structures. This is a major advantage when working with large datasets as it is not necessary to keep large portions of the dataset in main memory. At the same time, the filtering step is linear in the (usually large) number of hits and thus more efficient than an explicit check of the overlap, which would require an enumeration of all pairs of hits from the same dataset entry.

Whether an entry of a hit list qualifies for further consideration depends on a user-defined threshold parameter, M . The choice of this parameter has a considerable influence not only on the runtime of the method but also on the quality of the query results as it is responsible for the separation of significant and insignificant matches. We are currently building a probabilistic model for the method in order to derive statistical distributions for the number of votes. These distributions will provide the basis for an optimal choice of the M parameter. If a sampling strategy is applied, the parameter has to be adjusted as fewer hits will lead to a smaller number of ‘votes’.

The list of qualified entries for each pair of binding pockets is finally used to build the product graph for the clique-detection process. As each entry of this list is part of a larger local matching, the product graph will be relatively small, containing only few false-positive matches. The size of the matching generated by clique-detection is then used as a measure of similarity. As noted earlier, it is also possible to calculate a superimposition from the mapping and to perform a more complex grid-based analysis of the resemblance of the surfaces.

RESULTS

The implementation of our approach has been carried out in C++ and is based on the Cavbase extension of the Relibase system (Schmitt *et al.*, 2002). We use Cavbase for the data handling and exploit some of its infrastructure. A dataset of 2138 binding pockets originating from 1560 Protein Data Bank (PDB) entries has been derived and used to assess the performance and to validate our approach. The dataset is intended to be a representative subset of all enzyme binding sites in the PDB. The main selection criteria have been (i) the availability of an EC-classification for the respective protein structure from the ENZYME-classification database (Bairoch, 2000) and (ii) whether the binding pocket contains a complete ligand structure. The dataset is highly diverse and covers more than 200 different EC numbers. Typically, not only the catalytic pockets of a structure are included but also binding sites of co-factors, etc.

The preprocessing step—which has to be performed only once for each dataset—required a few hours on five standard Linux computers. The index was built using a sampling rate of 25% (i.e. only one of four k -cliques was included in the index).

After the index is built, similarity searches may be performed efficiently in the dataset. Depending on the configuration of the applied hardware and on the size of the query structure, a single (one against all) search may be performed in a few minutes. Again, it is possible to reduce this runtime by making use of a random sampling strategy: Figure 3 summarizes some query results for different types of binding pockets and compares them with the SCOP-classification (Murzin *et al.*, 1995) of the respective protein structures. The runtimes are given for the original approach as well as for the optimized approach combined with different (query) sampling rates. Remember that a fixed sample rate of 25% has been used while building the index. To allow a fair comparison of the results, both implementations used the grid-based surface-overlap measure of the original approach for the scoring. All entries of the dataset with a similarity score above 5.0 were considered as hits. Typically, only 2–3% of all entries score above this threshold. For each query, the number of entries from the same SCOP-family among the hits is shown.

Obviously, it is possible to achieve a significant speed-up (~6-fold, in some cases even one order of magnitude) relative to the original approach without missing too many hits. Note that even a sampling rate of 20%—which means that only 5% of the implicit comparisons of the k -cliques of two binding pockets are actually performed—usually yields sufficiently good results.

We also expected to find structural similarities not automatically covered by SCOP or comparable classification methods as these are based on comparisons in sequence and fold space. We found the well-known examples of convergent evolution as shown by serine proteases, NAD(P)-dependent proteins, zinc-binding proteins, etc. One typical example is the

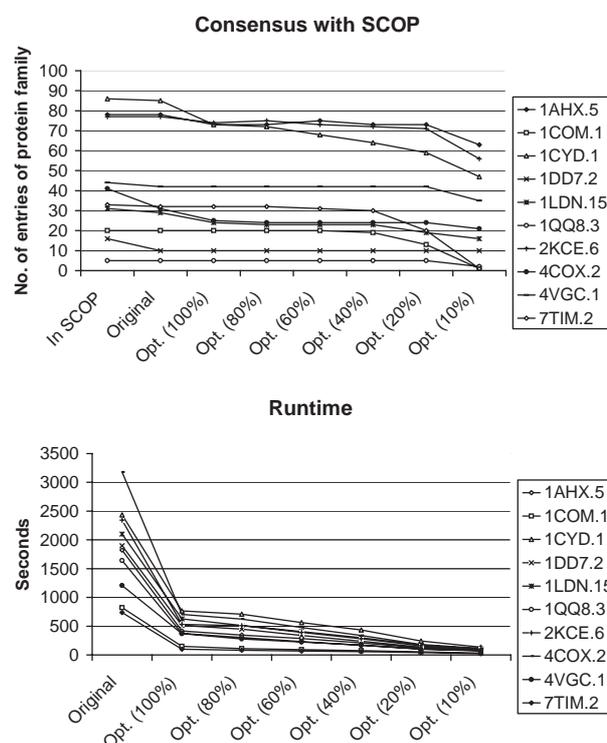


Fig. 3. Results of some similarity searches in a dataset of 2138 binding pockets, compared with the SCOP-classification. For the optimized approach, we show the results for different sampling rates.

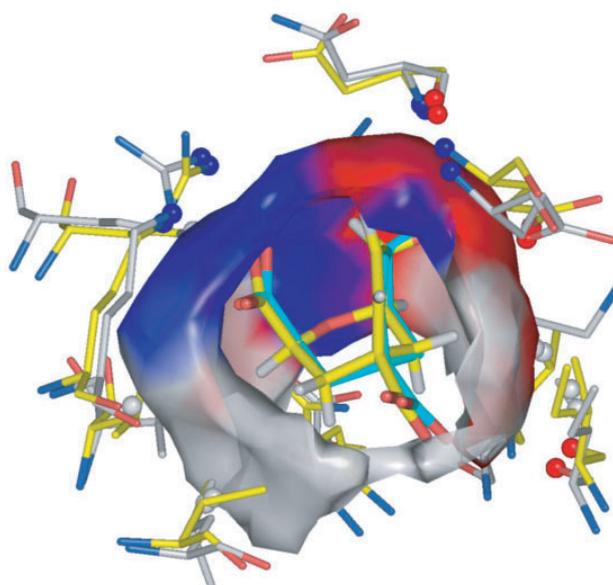


Fig. 4. Superposition of the binding pockets from the chorismate mutases *1ecm* and *4csm* (carbons are colored in yellow). The ligands superimpose convincingly well, although no ligand information was used in the similarity calculations. The figure was prepared using PyMol (DeLano, 2002).

chorismate mutases from *Escherichia coli* and *Saccharomyces cerevisiae*, which show no significant similarity in their sequence and fold (Rosen et al., 1998). We used the pocket of 4csm as query and found 1ecm as the highest ranked non-homologous hit (cf. Fig. 4).

DISCUSSION AND CONCLUSIONS

We presented an efficient two-step approach for the search of functionally related proteins in large datasets. The two steps are connected by a novel filtering step for the elimination of false-positive matches. The filtering step does not rely on structural superimposition. It is therefore independent of rigid-body assumptions and able to handle conformational flexibility. Additionally, it depends only on information that is available from the query structure and from the index. This is a major advantage when dealing with large datasets that may not be kept in main memory. We applied our method to a medium-sized dataset and presented some of the obtained results. In comparison with the original approach, we achieved a significant speed-up and retained most of the desired hits.

Clearly, the approach presented is intended mainly for larger datasets, that will be accessed by a high number of queries—e.g. for the development of a public Web server, that allows the comparison of a query pocket against a representative subset of the PDB. For searching smaller datasets, the index generation might be too costly and a direct clique-detection implementation is more appropriate.

Existing approaches for the detection of common side chain patterns based on efficient indexing of triplets or triangles (Russell, 1998; Hamelryck, 2003) have to be very selective to avoid a large number of false-positive matches. Therefore, they usually detect only highly conserved resemblances. Using our method for post-processing, their criteria could eventually be relaxed and thus more distant relationships among structures detected.

As clique-detection techniques are widely used for the structural comparison of proteins and small organic molecules and also for docking applications, our approach might also be applicable to related problems.

REFERENCES

- Bairoch, A. (2000) The ENZYME database in 2000. *Nucleic Acids Res.*, **28**, 304–305.
- Beckmann, N., Kriegel, H.-P., Schneider, R. and Seeger, B. (1990) The R*-tree: an efficient and robust access method for points and rectangles. *Proceedings of the 1990 ACM SIGMOD International Conference on Management of Data*, May 23–26, Atlantic City, NJ, USA. ACM Press, pp. 322–331.
- Bron, C. and Kerbosch, J. (1973) Algorithm 457: finding all cliques of an undirected graph. *Commun. ACM*, **16**, 575–577.
- DeLano, W.L. (2002) *The PyMOL User's Manual*. DeLano Scientific, San Carlos, CA, USA.
- Grindley, H.M., Artymiuk, P.J., Rice, D.W. and Willet, P. (1993) Identification of tertiary structure resemblance in proteins using a maximal common subgraph isomorphism Algorithm. *J. Mol. Biol.*, **229**, 707–721.
- Hamelryck, T. (2003) Efficient identification of side-chain patterns using a multidimensional index tree. *Proteins*, **51**, 96–108.
- Harrison, A., Pearl, F., Mott, R., Thornton, J. and Orengo, C. (2002) Quantifying the similarities within fold space. *J. Mol. Biol.*, **323**, 909–926.
- Koch, I., Lengauer, T. and Wanke, E. (1996) An algorithm for finding maximal common subtopologies in a set of protein structures. *J. Comp. Biol.*, **3**, 289–206.
- Levi, G. (1972) A note on the derivation of maximal common subgraphs of two directed or undirected graphs. *Calcolo*, **9**, 341–352.
- Murzin, A.G., Brenner, S.E., Hubbard, T. and Chothia, C. (1995) SCOP: a structural classification of proteins database for the investigation of sequences and structures. *J. Mol. Biol.*, **247**, 536–540.
- Nussinov, R. and Wolfson, H.J. (1991) Efficient detection of three-dimensional structural motifs in biological macromolecules by computer vision techniques. *Proc. Natl Acad. Sci., USA*, **88**, 10495–10499.
- Rosen, M., Lin, S.L., Wolfson, H. and Nussinov, R. (1998) Molecular shape comparisons in searches for active sites and functional similarity. *Prot. Eng.*, **11**, 263–277.
- Russell, R.B. (1998) Detection of protein three-dimensional side-chain patterns: new examples of convergent evolution. *J. Mol. Biol.*, **279**, 1211–1227.
- Schmitt, S., Kuhn, D. and Klebe, G. (2002) A new method to detect related function among proteins independent of sequence and fold homology. *J. Mol. Biol.*, **323**, 387–406.