Parallel maximal common subgraphs with labels for molecular biology

Wilfried Agbeto[1], Camille Coti[2], and Vladimir Reinharz[1]

[1] Université du Québec à Montréal, Canada
agbeto.kossi_wilfried@courrier.uqam.ca
reinharz.vladimir@uqam.ca
[2] École de Technologie Supérieure, Canada
camille.coti@etsmtl.ca

Abstract. Advances in graph algorithmics have allowed in-depth study of many natural objects from molecular biology or chemistry to social networks. Particularly in molecular biology and cheminformatics, understanding complex structures by identifying conserved sub-structures is a key milestone towards the artificial design of novel components with specific functions. Given a dataset of structures, we are interested in identifying all maximum common connected partial subgraphs between each pair of graphs, a task notoriously NP-Hard.
In this work, we present 3 parallel algorithms over shared and distributed memory to enumerate all maximal connected common sub-graphs between pairs of arbitrary multi-directed graphs with labels on their edges. We offer an implementation of these methods and evaluate their performance on the non-redundant dataset of all known RNA 3D structures. We show that we can compute the exact results in a reasonable time for each pairwise comparison while taking into account a much more diverse set of interactions—resulting in much denser graphs—resulting in an order of magnitude more conserved modules. All code is available at https://gitlab.info.uqam.ca/cbe/pasigraph and results in the branch `results`.

1 Introduction

Graphs are increasingly being applied in various fields such as mathematics, biology, chemistry, and social network analysis. For instance, in biology, the tertiary structure of RNA can be represented as a graph, where each vertex represents a nucleotide identified by its base name and sequence number, and edges represent interactions between nucleotides, labeled by their interaction type. In this RNA graph model, searching for RNA structural motifs (which are recurring substructures appearing at non-homologous locations in one or more RNA molecules) is equivalent to searching for maximal common connected subgraphs (MCCPS).

The problem of finding the MCCPS between two graphs is a generalization that includes the connectivity aspect of the Maximum Common Partial Subgraph (MCPS) problem, which is limited to finding only the largest maximal

common partial subgraph between two graphs and is recognized in the literature as being NP-hard : for each vertex in the first graph, we find every similar vertex in the second graph. Edges of the graph are labeled with the type of interaction between the two nucleotides modeled by the graph. For each of these similar pairs, we try to extend the similar subgraph by comparing the edges connecting their neighbors, and we continue this extension as long as we find similar neighbors . Because of the computation-intensive nature of this problem, parallel computing is an attractive approach to speed-up the computation by taking advantage of the computation and memory capabilities of current high-performance computing platforms. Indeed, finding all the MCCPS requires an exhaustive exploration of the search space, and therefore has a high combinatorial complexity.

However, finding all the common subgraphs between two graphs in parallel faces multiple challenges. The MCCPS are tied with each other: extending a subgraph can have consequences on other subgraphs that might be computed on other processes. Moreover, MCCPS computation times are very unbalanced. Therefore, computing them in parallel raises non-trivial load-balancing and inter-process communication challenges. Moreover, the number of solutions of an MCCPS problem can be exponential in the number of nodes, henceforth making exhaustive search highly computation-intensive.

In this paper, we present parallel algorithms to compute the MCCPS over shared memory, distributed memory, and a hybrid approach. Our implementation works on multi-directed graphs with labeled edges; we present a performance evaluation on the dataset of all known RNA 3D structures. A limitation of previous works was the number of labels taken into account, the most efficient algorithms would fail on a computer with 8TB of RAM when all contacts between two large molecules were taken into account. The present algorithms easily overcome this limitation requiring only 186GB of RAM to process the entire dataset.

2 Related works

The problem of finding a Maximum Common Subgraph (MCS) has been studied for half a century. The problem of finding an MCS is divided into two categories of problems: MCIS (Maximum Common Induced Subgraph) [18, 22, 28] is the problem of finding a common induced subgraph with the largest number of nodes, and MCPS (Maximum Common Partial Subgraph) [3, 23] is the problem of finding a common partial subgraph with the largest number of edges. Another distinction can be made between the connected case and the disconnected case. [10, 9] have demonstrated that this problem is one of the most challenging in terms of algorithmic complexity.

There are two main types of exact approaches to solving this problem: the reduction approach to the maximum clique problem and the enumeration approach of common subgraphs (branch and bound) [16]. The reduction approach to the maximum clique problem relies on reformulating the problem into a compatibility graph. In this graph, each clique corresponds to a common subgraph. Thus, finding a maximum common subgraph is equivalent to identifying a maximum

clique in the compatibility graph. In the enumeration approach of common subgraphs, a search tree is constructed to explore all possible common subgraphs. The creation of subgraphs is halted when the algorithm determines that a subgraph cannot produce a solution larger than the best one found so far. The efficiency of these algorithms depends on the use of pruning heuristics to quickly remove unnecessary branches from the search tree [16]. The second approach is more efficient for small or sparse graphs. In all other cases, the first approach is more efficient [5].

Approximate [24] and parallel [14, 7] approaches have also been proposed. Parallel approaches decompose the problem into independent sub-problems that can be solved simultaneously using domain partitioning. The MCS problem is one of the most difficult problems to parallelize due to its complex combinatorial nature and the interdependence among sub-problems. The subgraphs considered by the computation can span over multiple subdomains, leading to non-trivial communication patterns and load-balancing challenges.

Such irregular parallel applications can be dealt with using task-based programming with automatic load balancing, such as what was done for parameter space exploration. The parameter space is partitioned from points called *reference points*. From a reference point, polyhedra around it are computed. Different polyhedra can be computed in parallel as their computation is independent of each other. However, it is impossible to know in advance which points will be included in the same polyhedron, so it is important to avoid computing the same polyhedron on different computational resources. This problem can therefore be connected to the MCS problem, in which the computation time to find a common subgraph and how the domain can be decomposed cannot be known before the common subgraph has been found. Heuristics and dynamic algorithms have been designed and evaluated in [1, 2]. Their conclusion is that for small parameter spaces, a random distribution of points works well, and for larger parameter spaces, a dynamic domain decomposition with work stealing works better. [6] address a comparable parameter space partitioning problem. They compare different approaches, and the one that performs the best involves detecting and stopping redundant computations early and computing another task instead.

In this work, we will apply MCPS to the task of finding RNA motifs. An RNA molecule can be described as a sequence of nucleotides $\omega \in \{ACGU\}^n$, of a length from 10s to thousands of letters. They adopt 3D conformations through contacts (strengthened by hydrogen bonds) between the different nucleotides. The structure is known to form hierarchically [30], first canonical and wobble base pairs (A–U, G–C, and G–U contacts in the cis Watson–Crick/Watson–Crick `cWW` conformations) agregate into rigid *stems*, ladders of cWW contacts. These have experimentally determined energies [17] and are well captured by an efficient thermodynamic model in $\mathcal{O}(n^3)$, with the length of the sequence $|\omega| = n$ below 300 [12]. Yet the global configuration and binding geometries depend on finer conformations in the loops, positions outside of the *stems*, for which no energetic parameter are known. These new geometries can conveniently be described as graphs, where the nucleotides are nodes, the edges indicating

contacts between the nodes. The geometry of these contacts is classified into 12 families by Leontis–Westhof [11]. They can be of type `cWW` as others depicted in Table 2. The backbone connects the nucleotides sequentially.

Uncovering the main motifs, conserved subgraphs, is a main challenge in the field necessary to support better tools for structure prediction as the design of artificial sequences [15]. Approaches directly from 3D structure have been limited to simple loops due to the complexity of atoms-to-atoms comparisons [20]. Alternatively, graph-based algorithms have been developed and have allowed to find motifs spanning multiple loops and with hundreds of nodes and edges [8, 4, 25, 29]. Yet stackings, a core feature of RNAs, are too numerous to be analyzed with current methods. To overcome this task, we develop a memory-efficient distributed algorithm that allows to exhaustively enumerate all MCPPS when considering backbones, canonical and noncanonical contacts, as stackings.

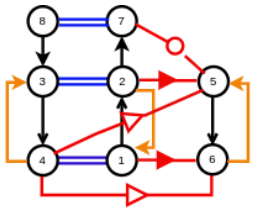

Fig. 1. RNA graph, blue double lines are `cWW` interactions forming stems, black edges connect positions sequentially, red edges are between loops, orange edges are stacks.

Edge type	Edge label	Inverse edge label
Phosphodiester bonds (Backbone)	B53	B35
Canonical and non-canonical interactions	$\alpha\beta\gamma$, $\alpha \in$\{c,t\}, βet$\gamma \in$ \{W,H,S\}	$\alpha\gamma\beta$
Stack	SAB, A et B $\in$\{5,3\}	SBA

Table 1. Labeling an RNA graph. cWW is canonical, other non-canonical

3 Algorithms

In this section, we are presenting parallel algorithms to solve the MCCPS problem. We are presenting the sequential algorithm and data structures our parallel algorithms are based on (Section 3.1), then parallel algorithms over shared memory (Section 3.2) and distributed memory (Section 3.3).

3.1 Algorithm for Maximal Common Connected Partial Subgraphs

Let g and h be two directed graphs with edge labels. The approach used to find all MCCPS between g and h involves starting from an edge match between g and h and adding neighboring edge matchings to it, considering all possibilities until reaching the maximal one. This approach is made of two steps: 1) Search for all edge matches and 2) Extension of edge matches.

Search for all edge matchings (algorithm 1) Two edges in g and in h correspond together if they have the same label and direction. We use a list describing the edge types involved to treat an edge match and its reverse edge match as a single edge match. For example, for edgeTypeList=[{CWS,CSW}, {TWW,TWW}], the first element of edgeTypeList implies that if there exists an edge (a, b, CWS) in g, then there also exists an inverse edge (b, a, CSW) in g.

The function `SearchAllEdgeMatches` returns the list of all found edge matches (`edgeMatchList`), as well as two other lists: **edgeMatchIdList1** contains the identifier of the edge in g for each edge match, and **edgeMatchIdList2** contains the identifier of the edge in h for each edge match.

Extension of edge matchings Edge matchings are the elementary patterns constituting the MCCPS. In this section, we present the algorithm that extends an edge matching to find all MCCPS containing it.

Extending an edge matching involves adding neighboring edge matchings to it through a breadth-first search (BFS) traversal until maximization. Maximization is achieved when there are no more neighbors to explore.

During the extension of an edge matching, conflict events may occur. A conflict event happens when a neighboring edge matching cannot be added to the PCCS (Partial Common Connected Subgraph: subgraph under construction and not yet maximal). This occurs when a neighboring edge matching has a node ag|ah that cannot be added to the PCCS, as the latter either has a node $a'g|ah$ or a node $ag|a'h$, or both (with $a'g \neq ag$ and $a'h \neq ah$).

For each conflict, we create a new extension branch containing a new PCCS to be extended. To obtain the new PCCS, we create a copy of the MCCPS by removing the nodes and edges in conflict with the conflicting edge matching and adding the conflicting edge matching to the copy.

Extending an edge matching s, combined with conflict management, allows us to find all MCCPS containing s. Conflict management enables the extension algorithm to backtrack and explore alternatives to find additional MCCPS. However, multiple conflicts can lead to the same MCCPS, posing a challenge. To mitigate this, before creating a new extension branch, we check if the new PCCS to be extended is already included in any previously found MCCPS starting from s. If so, the branch is not created. This inclusion check involves searching for an exact subgraph isomorphism that preserves edge labels and node identifiers.

Since extending an edge matching s yields all MCCPS containing $\mathbf{s}$, then the extension of multiple edge matchings can yield identical MCCPS. To address this issue, when an edge matching $\mathbf{e}$ can be added to a PCCS resulting from extending an edge matching s_i, then we check if $\mathbf{e} \in \{s_0, s_1, s_2, ..., s_{i-1}\}$, where s_i is the edge matching at index i in **edgeMatchList**. If this is the case, the edge matching is not added to the PCCS. The subgraph obtained at the end of extending the PCCS may not be an MCCPS, meaning it may not be maximal. If at least one of the edge matchings $\mathbf{e}$, which were ignored by the PCCS, can be added to the obtained subgraph, then it is not maximal and therefore not considered as an MCCPS.

3.2 Shared Memory Algorithm

Algorithm 1 gives the parallel algorithm in the shared memory model. We based our approach on the **PCAM** method, which includes the following steps:

- **Partitioning:** The problem is decomposed into fine-grained tasks.
- **Communication:** Identification of dependencies between tasks.

- **Agglomeration:** Once the tasks and their dependencies have been determined in the previous steps, tasks can, if necessary, be combined into larger tasks to improve performance (data movements).
- **Mapping:** Distribution of tasks across processing units.

Search for all edge matchings The partitioning step is achieved using a one-dimension domain decomposition on the `edgeTypeList` to create one task for each element of `edgeTypeList`. The computation time for each of these tasks is irregular and can take advantage of a dynamic scheduling policy such as the one provided by OpenMP.

```
   Data: g, h, edgeTypeList
 1 edgeMatchList ← {} /* shared variable                                  */
 2 edgeMatchIdList1 ← {}/* shared variable                               */
 3 edgeMatchIdList2 ← {} /* shared variable                              */
 4 Parallel construction shared(g, h, edgeTypeList)
 5 edgeMatchListLocal ← {};
 6 edgeMatchIdList1Local ← {};
 7 edgeMatchIdList2Local ← {};
 8 Parallel for i ← 0 to edgeTypeList.size() do
 9 |   edgeType ← edgeTypeList[i];
10 |   label ← edgeType.edgeLabel;
11 |   edgeList1 ← getOutgoingEdge(g,label);
12 |   edgeList2 ← getOutgoingEdge(h,label);
13 |   for edge1 in edgeList1 do
14 |   |   for edge2 in edgeList2 do
15 |   |   |   matchEdge ← createMatchEdge(edge1,edge2);
16 |   |   |   edgeMatchListLocal.append(matchEdge);
17 |   |   |   edgeMatchIdList1Local.append(edge1.id);
18 |   |   |   edgeMatchIdList2Local.append(edge2.id);
19 Begin critical section
20     edgeMatchList.append(minimumSubGraphListLocal);
21     edgeMatchIdList1.append(minimumSubGraphIdList1Local);
22     edgeMatchIdList2.append(minimumSubGraphIdList2Local);
23 End critical section
24 End parallel construction;
```

Algorithm 1: SearchAllEdgeMatches Parallel

Extension of edge matchings The set of edge matches also uses a one-dimensional domain decomposition to create one task for each edge match. We obtain fine-grained tasks because extending edge matches is strongly unbalanced between the tasks. In the following, we focus on a work-stealing approach to balance the workload automatically. We are using threads as processing units. We explored the impact of the scheduling strategy being used here, and the results are presented in Section 4.1. To avoid redundant computations between threads, we used the backtracking mechanism described in section 3.1 using shared variables.

Dynamic Distribution Each thread receives a task corresponding to the extension of an edge match. When it is done, it moves on to another task. In this approach, tasks are independent, so it is not necessary to communicate or synchronize

between tasks. However, this approach suffers from significant load imbalance when the distribution of the number of MCCPS per edge matching is highly unbalanced.

Nested Threads are organized in groups. The master thread of each group receives a task corresponding to the extension of an edge match, and the branch extension tasks resulting from conflict resolution are distributed among the threads of the same group. Threads within the same group must synchronize to avoid extending a branch that will yield an already-found MCCPS. When it is done, the group moves on to another task.

Work Stealing This approach is similar to the dynamic distribution one, except that when there is no more tasks that extend an edge match are exhausted, idle threads steal branch extension tasks from active threads. This approach allows for better workload balancing. To reduce task synchronization overhead, threads do not directly share their branch extension tasks; they only do so when at least one other thread is available. A thread knows if there is at least one other available thread through the shared variable *numThreadFree*. When there are no more tasks to extend edge matches, idle threads set the value at index *rankThread* in *threadStateList* to 1, indicating they are available, and increment *numThreadFree* by 1. When at least one thread is available, only one of the busy threads shares its branch extension tasks: the thread with the smallest thread id using the function *isYourRound* that takes *threadStateList* and *rankThread* as input, traverses *threadStateList*, and returns 1 if it finds only 1, indicating that the calling thread should share its tasks. Otherwise, the function returns 0.

3.3 Distributed Memory Algorithm

Search for All Edge Matches The starting point of the algorithm is edge matching. This can be distributed easily like on shared memory: every process starts from a set of labels and communication is necessary at this step. Each process executes the same algorithm for shared memory (Algorithm 1).

Extension of Edge Matches This step follows a master-worker scheme. The master sends a batch of starting points to the workers. Each worker implements a work-stealing approach to balance the load between threads. Since every process processes a set of starting points, statistically, the load unbalance between processes is reduced compared to a finer-grain approach. Moreover, the master-worker scheme provides an automatic load-balancing mechanism.

If two processes start from two different points that are part of the same common subgraph, they will perform the same computation. To avoid that, at the end of the first step, every process has an ordered list of all the edge matches. Each edge match is associated with an index i. The master sends a range $[i, i']$ to each worker. It is important to notice that we are not discovering a subgraph label by label, but instead, each process finds the list of PCCSs each edge is part of. The master is in charge of merging these PCCSs and creating the subgraphs during the post-processing step.

A work-stealing approach can also be implemented. However, it requires frequent inter-process interactions, so synchronization between branch extension tasks would result in a significant amount of inter-process communication.

3.4 From MCCPS to biological RNA motifs

A few constraints are necessary to transition from the mathematical object to a relevant biological motif. It is known that interesting RNA motifs have a cyclic structure, and each node is involved in a non-canonical interaction (interaction labeled not `CWW` nor `B53`) [19, 13].

As a postprocessing step, for each of the MCPPS, we iteratively remove nodes that do not fulfill the following conditions: (1) nodes must be connected by an interaction other than B53, (2) if two nodes n_1, n_2 are connected by a `cWW`, there must be a node n_3 neighbor to n_1 or n_2 involved in a non-canonical or stacking interaction with a fourth node (i.e. we do not extend stacks of only canonical base pairs), and (3) each node belongs to a cycle.

It is clear that all maximal recurrent RNA motifs must belong to an MCPPS and be extractable from it, else another bigger RNA motif must contain it, a contradiction.

4 Performance evaluation

We implemented these algorithms using OpenMP (4.5) provided with gcc 11.4 and OpenMPI (4.1.2). The code was optimized by g++ using -O3. The performance data presented in this section was measured on Calcul Québec's machine Narval[3]. The set of 1722 non-redundant RNA structures was taken from RNA3DHub version 3.269 [21]. The graphs for each structure were built with annotations of the contact geometries from FR3D [27]. The smallest graph has 2 nodes and 2 edges, the biggest 16 770 nodes and 37 582 edges. Over 200 graphs have a node count greater than or equal to 100.

4.1 Shared memory performance

First, we compared the performance of the three shared memory algorithms presented in section 3. We picked two graphs randomly; one graph has 390 vertices and 1003 edges, and the other one has 1800 vertices and 4282 edges. The performance comparison is presented in Figure 2. We can see that all three algorithms benefit from the parallelization when a few threads are used, and beyond 4 threads, the performance of the nested and dynamic algorithms stalls whereas the work-stealing algorithm keeps scaling.

We picked randomly 7 pairs of graphs from the RNA3DHub dataset and we compared the performance obtained with the three algorithms on these graphs using 10 threads. As presented in Figure 3, the work-stealing approach was faster on every graph we used.

4.2 Hybrid algorithm

Current parallel architectures are hierarchical: they are made of multi-core nodes connected by a high-speed network. As a consequence, parallel programs can follow a hybrid approach and use shared-memory algorithms on multiple cores

[3] https://docs.alliancecan.ca/wiki/Narval/en

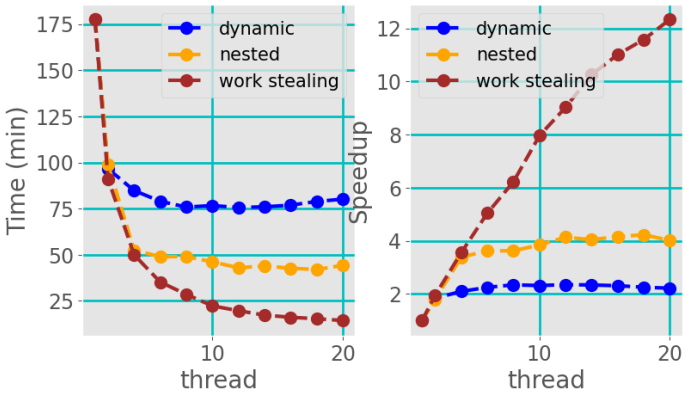

Fig. 2. Scalability of the shared memory algorithms.

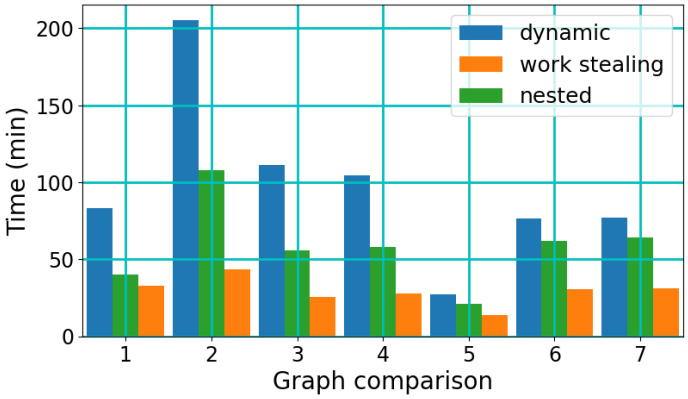

Fig. 3. Performance comparison between the shared memory algorithms.

of the same node, and distributed-memory algorithms between nodes. Since distributed-memory algorithms can be used on several cores of a given node, the distribution of the processing units (processes and threads) between shared memory and distributed memory is a tuning parameter of the program.

In our case, we used a multi-thread algorithm over shared memory and multiple processes over distributed memory. We evaluated the impact of the threads *times* processes combination on the performance. Since Narval's nodes feature 48 cores, we compared the performance obtained using one process and 48 threads per process, and all integer combinations up to 48 single-threaded processes.

Based on the observations made in section 4.1, we used the work-stealing algorithm on the multi-thread part of the computation. The performance comparison on a single node is presented in Figure 4. The speed-up is indicated over every bar. We can see that the approach using 48 threads gives the best performance.

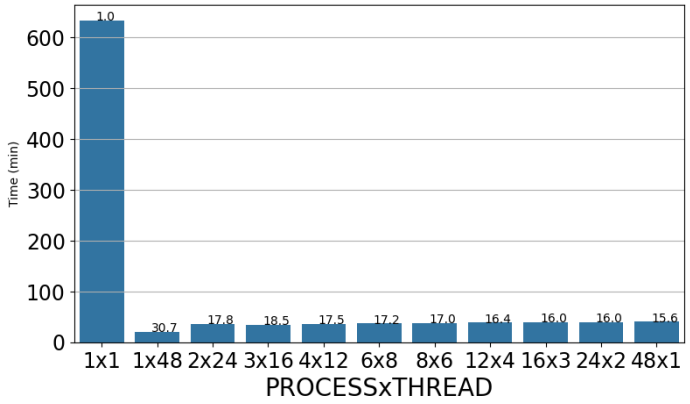

Fig. 4. Hybrid algorithm: number of threads per process for a given number of cores.

4.3 Distributed memory scalability

We evaluated the scalability of the distributed memory part of the algorithm. We set a given number of threads per process (4, 6, 8, and 10) with one process per node, and we increased the number of nodes. We can see the performance Figure 5. We can see that the performance improves when we increase the number of computation resources used. For a given total number of resources, the performance is very similar for each number of threads.

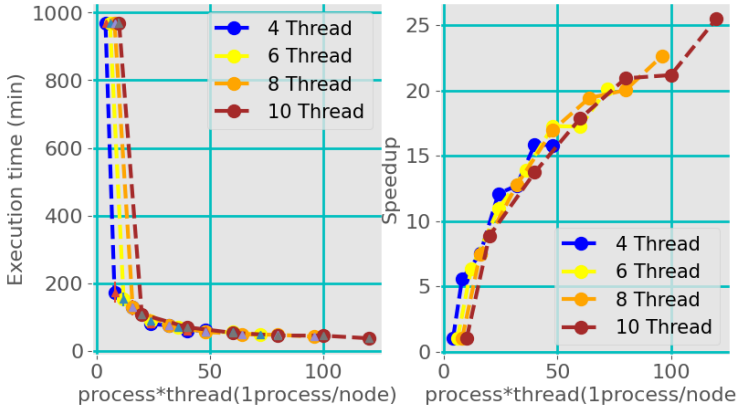

Fig. 5. Scalability of the hybrid approach over multiple nodes.

4.4 Performance of the hybrid algorithm

Figure 6 presents the scalability of the hybrid approach. We used all the cores of each node, with different combinations of the number of processes per node

and the number of threads per process. As explained in the previous sections, we used the work-stealing approach. We can see that using one process per node and 48 threads per process gives the best performance, confirming the expectation presented in Section 4.2. The performance improves with the number of nodes up to 6 nodes (288 cores). This size was sufficient for us to obtain the results we needed: the aggregated memory of these 6 nodes was enough to perform the computation, and the execution time was reasonable.

We tried to speed up the computation more by using a larger number of cores, and the performance worsened. Thanks to the performance we reached with the distributed version, it was possible to complete the computation and get results for the bioinformatics community. However, getting results faster is always desirable, for instance to complete more computations in a given makespan. Performance analysis suggests when the number of processes is large enough, the workload is distributed between them over a few steps of the master-worker scheme, and the automatic load balancing mechanism cannot compensate for load unbalances between nodes. Future works will investigate other distributed approaches, for example with more dynamic load-balancing approaches.

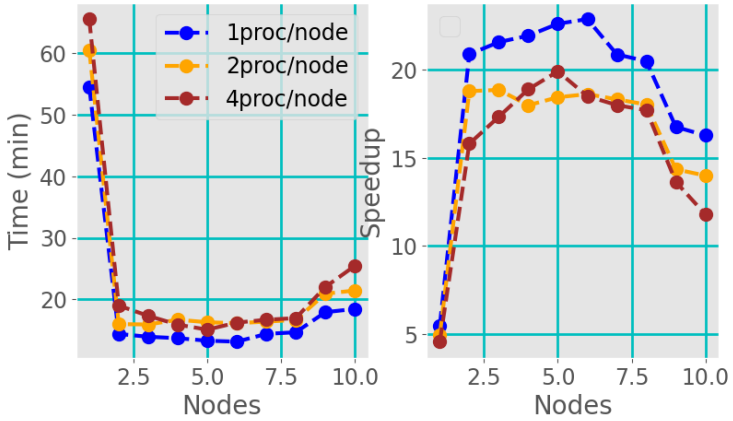

Fig. 6. Scalability of the hybrid approach.

4.5 RNAs

We used our parallel program to perform a pairwise comparison of all the non-redundant existing RNA structures to identify all the structural motifs they contain (see Sec. 3.4. It took 4 weeks on Narval, which is significantly faster than previous works [26, 29]. It required a total of 186 GB of memory. It is important to note that the results obtained were not achievable with the methods proposed in previous work, even with 8TB of RAM. Additionally, we found a larger number of structural motifs. We found 157 344 structural motifs for a total of 209 750 474 occurrences. The smallest motif has 3 nodes and 6 edges, and the largest 1135 nodes and 5566 edges.

The A-minor motif is well known for its structural importance as a general tool to join loops, yet this versatility makes it hard to predict in practice 7.

We show in Fig. 7 the first map of stacking variants of the A-minor motif. The numbers on the figure indicate motif coverage.

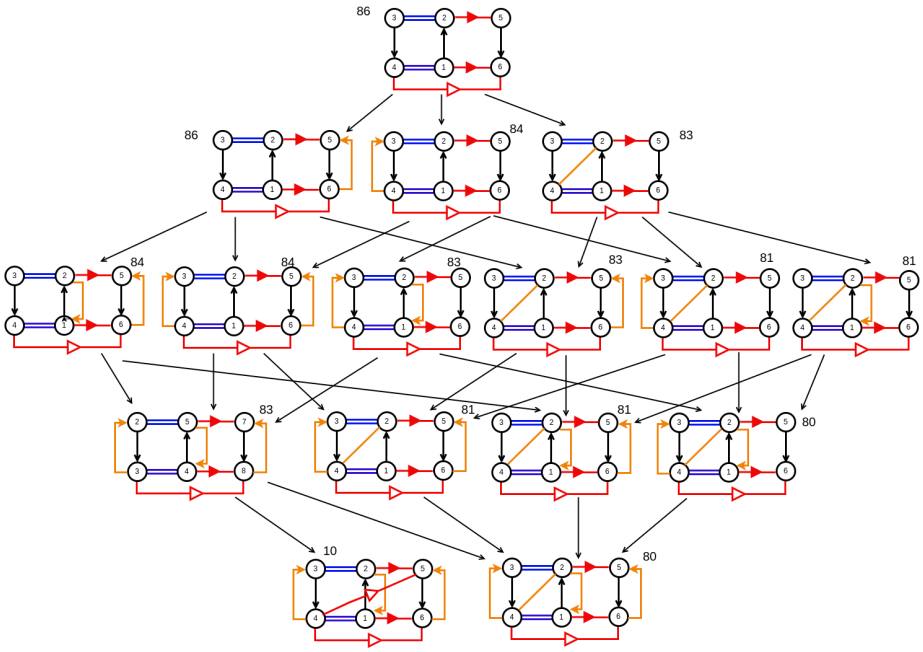

Fig. 7. Variation of A-Minor motifs with stacking. The orange lines, oriented and unoriented, represent stacking interactions S53 and S55, respectively.

5 Conclusion

In this paper, we present novel parallel algorithms for the exhaustive enumeration of Maximum Common Partial Subgraphs between two arbitrary multi-directed graphs, admitting labels over edges. This problem has a high computation complexity, but its irregular nature makes it non-trivial to parallelize. We offer an open-source implementation in C and evaluate it using complex RNA structure graphs as benchmarks. We show that our algorithms scale over shared and distributed memory using multiple threads and multiple processes, and with efficient memory usage.

In molecular biology, MCPS can be indicative of important functional motifs. In the case of RNA structures that have an efficient graph representation, motifs finding algorithms always ignored different key features such as stackings, known to be important but too numerous to be fully taken into account, until now. With the algorithms presented here, we can now exhaustively list all recurrent motifs creating novel maps of natural variants with stackings. Such information will be leveraged in structure prediction, molecular design, and binding target prediction.

The abundance and size of motifs raise new questions that will require efficient parallel algorithms to solve. Studying RNA structures is not a static endeavor, every week new structures emerge and have to be analyzed and merged with our dataset. With over 150 000 motifs now, their organization and clustering are another challenge. Annotating all their positions in the few thousands of known RNAs requires several terabytes of data. While the number of known RNA structures remains low, the amount of information to analyze has been increasing over the last decades and will continue to do so.

RNA structures are one limited example of what can be achieved now using the computation power offered by these parallel algorithms. Labels on nodes are currently ignored, but with new experimental techniques allowing annotation of chemical modifications in RNAs using nanopore sequencing, we can expect such a feature to be highly desirable.

References

1. Étienne André, Coti, C., Nguyen, H.G.: Enhanced distributed behavioral cartography of parametric timed automata. In: Proceedings of The 17th International Conference on Formal Engineering Methods (ICFEM 2015). pp. 319–335 (November 2015)
2. André, É., Coti, C., Nguyen, H.G.: Enhanced Distributed Behavioral Cartography of Parametric Timed Automata (Informal Presentation). In: André, É., Frehse, G. (eds.) 2nd International Workshop on Synthesis of Complex Parameters (SynCoP'15). OpenAccess Series in Informatics (OASIcs), vol. 44, pp. 104–105. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2015)
3. Bahiense, L., Manić, G., Piva, B., De Souza, C.C.: The maximum common edge subgraph problem: A polyhedral investigation. Discrete Appl. Math. **160**(18), 2523–2541 (dec 2012)
4. Chojnowski, G., Waleń, T., Bujnicki, J.M.: RNA Bricks—a database of RNA 3D motifs and their interactions. Nucleic acids research **42**(D1), D123–D131 (2014)
5. Conte, D., Foggia, P., Vento, M.: Challenging complexity of maximum common subgraph detection algorithms: A performance analysis of three algorithms on a wide database of graphs. J. Graph Algorithms Appl. **11**, 99–143 (01 2007)
6. Coti, C., Monniaux, D., Yu, H.: Parallel parametric linear programming solving, and application to polyhedral computations. In: Rodrigues, J.M.F., Cardoso, P.J.S., Monteiro, J., Lam, R., Krzhizhanovskaya, V.V., Lees, M.H., Dongarra, J.J., Sloot, P.M. (eds.) Computational Science – ICCS 2019. pp. 566–572. Springer International Publishing, Cham (2019)
7. Depolli, M., Konc, J., Rozman, K., Trobec, R., Janezic, D.: Exact parallel maximum clique algorithm for general and protein graphs. Journal of chemical information and modeling **53** (08 2013)
8. Djelloul, M., Denise, A.: Automated motif extraction and classification in RNA tertiary structures. RNA **14**(12), 2489–2497 (2008)
9. Huang, X., Lai, J., Jennings, S.: Maximum common subgraph: Some upper bound and lower bound results. BMC bioinformatics **7 Suppl 4**, S6 (02 2006)
10. Kann, V.: On the approximability of the maximum common subgraph problem. pp. 377–388 (02 1992)
11. Leontis, N.B., Westhof, E.: Geometric nomenclature and classification of RNA base pairs. Rna **7**(4), 499–512 (2001)

12. Lorenz, R., Hofacker, I.L., Stadler, P.F.: RNA folding with hard and soft constraints. Algorithms for Molecular Biology **11**, 1–13 (2016)
13. Loyer, G., Reinharz, V.: Concurrent prediction of RNA secondary structures with pseudoknots and local 3D motifs in an integer programming framework. Bioinformatics **40**(2), btae022 (2024)
14. McCreesh, C., Prosser, P.: Multi-threading a state-of-the-art maximum clique algorithm. Algorithms **6**, 618–635 (10 2013)
15. Miao, Z., Adamiak, R.W., Antczak, M., et al.: RNA-Puzzles Round IV: 3D structure predictions of four ribozymes and two aptamers. Rna **26**(8), 982–995 (2020)
16. Minot, M., Ndiaye, S.N., Solnon, C.: Recherche d'un plus grand sous-graphe commun par décomposition du graphe de compatibilité. In: Onzièmes Journées Francophones de Programmation par Contraintes (JFPC). pp. 1–11. Bordeaux, France (Jun 2015)
17. Mittal, A., Turner, D.H., Mathews, D.H.: NNDB: An Expanded Database of Nearest Neighbor Parameters for Predicting Stability of Nucleic Acid Secondary Structures. Journal of Molecular Biology p. 168549 (2024)
18. Ndiaye, S.N., Solnon, C.: Cp models for maximum common subgraph problems. In: Proceedings of the 17th International Conference on Principles and Practice of Constraint Programming. p. 637–644. CP'11, Springer-Verlag, Berlin, Heidelberg (2011)
19. Parisien, M., Major, F.: The MC-Fold and MC-Sym pipeline infers RNA structure from sequence data. Nature **452**(7183), 51–55 (2008)
20. Petrov, A., Zirbel, C., Leontis, N.: Automated classification of rna 3d motifs and the rna 3d motif atlas. RNA (New York, N.Y.) **19** (08 2013)
21. Petrov, A.I.: Rna 3d motifs: Identification, clustering, and analysis (2012)
22. Quer, S., Marcelli, A., Squillero, G.: The maximum common subgraph problem: A parallel and multi-engine approach. Computation **8**, 48 (05 2020)
23. Raymond, J., Gardiner, E., Willett, P.: Rascal: Calculation of graph similarity using maximum common edge subgraphs. Comput. J. **45**, 631–644 (04 2002)
24. Raymond, J., Willett, P.: Maximum common subgraph isomorphism algorithms for the matching of chemical structures. Journal of computer-aided molecular design **16**, 521–33 (08 2002)
25. Reinharz, V., Soulé, A., Westhof, E., Waldispühl, J., Denise, A.: Mining for recurrent long-range interactions in RNA structures reveals embedded hierarchies in network families. Nucleic acids research **46**(8), 3841–3851 (2018)
26. Reinharz, V., Soulé, A., Westhof, E., Waldispühl, J., Denise, A.: Mining for recurrent long-range interactions in RNA structures reveals embedded hierarchies in network families. Nucleic Acids Research **46**(8), 3841–3851 (03 2018)
27. Sarver, M., Zirbel, C.L., Stombaugh, J., Mokdad, A., Leontis, N.B.: FR3D: finding local and composite recurrent structural motifs in RNA 3D structures. Journal of Mathematical Biology **56**, 215–252 (2007)
28. Schmidt, R., Klein, R., Rarey, M.: Maximum common substructure searching in combinatorial make-on-demand compound spaces. Journal of Chemical Information and Modeling **62**(9), 2133–2150 (2022)
29. Soulé, A., Reinharz, V., Sarrazin-Gendron, R., Denise, A., Waldispühl, J.: Finding recurrent RNA structural networks with fast maximal common subgraphs of edge-colored graphs. PLOS Computational Biology **17**(5), 1–28 (05 2021)
30. Tinoco Jr, I., Bustamante, C.: How RNA folds. Journal of molecular biology **293**(2), 271–281 (1999)