

# GRAPES-DD: Exploiting Decision Diagrams for Index-Driven Biomedical Databases Search

**Nicola Licheri**

University of Turin

**Vincenzo Bonnici** (✉ [vincenzo.bonnici@univr.it](mailto:vincenzo.bonnici@univr.it))

<https://orcid.org/0000-0002-1637-7545>

**Marco Beccuti**

University of Turin

**Rosalba Giugno**

Università degli Studi di Verona

---

## Methodology article

**Keywords:** query processing, pattern matching, subgraph isomorphism, graph indexing, decision diagrams

**Posted Date:** August 6th, 2020

**DOI:** <https://doi.org/10.21203/rs.3.rs-48943/v1>

**License:** © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

---

**Version of Record:** A version of this preprint was published at BMC Bioinformatics on April 22nd, 2021. See the published version at <https://doi.org/10.1186/s12859-021-04129-0>.

## RESEARCH

# GRAPES-DD: exploiting decision diagrams for index-driven biomedical databases search.

Nicola Licheri<sup>1</sup>, Vincenzo Bonnici<sup>2\*</sup>, Marco Beccuti<sup>1†</sup> and Rosalba Giugno<sup>2†</sup>

\*Correspondence:

[vincenzo.bonnici@univr.it](mailto:vincenzo.bonnici@univr.it)

<sup>2</sup>University of Verona, Strada le

Grazie 15, 37134 Verona, Italy

Full list of author information is available at the end of the article

<sup>†</sup>Equal contributor

## Abstract

**Background:** Graphs are mathematical structures widely used for expressing relationships among elements when representing biomedical and biological information. On top of these representations, several analyses are conducted. A common task is the search of one substructure, called query, within one graph, called target. The problem is referred to as one-to-one subgraph search, and it is shown to be NP-complete. However, heuristics and indexing techniques can be applied to facilitate the search. Such indexing techniques are also exploited in the context of searching in a collection of target graphs, referred to as one-to-many subgraph problem. Filter-and-verification methods that use indexing approaches provide a fast pruning of target graphs or parts of them that do not contain the query. The expensive verification phase is then performed only on the subset of promising targets. Indexing strategies extract graph features at a sufficient granularity level for performing a powerful filtering step. Features are memorized in data structures allowing an efficient access. Indexing size, querying time and filtering power are key points for the development of efficient subgraph searching solutions.

**Results:** An existing subgraph approach, GRAPES, has been shown to have good performance in term of speed-up for both one-to-one and one-to-many cases. However, it suffers in the size of the built index. For this reason, we propose GRAPES-DD, a modified version of GRAPES in which the indexing structure has been replaced with a Decision Diagram. Decision Diagrams are a broad class of data structures widely used to encode and manipulate functions efficiently. Our experiments on real biomedical structures and synthetic graphs have confirmed our expectation showing that GRAPES-DD has substantially reduced the memory utilization with respect to GRAPES without worsening the searching time.

**Conclusion:** The use of Decision Diagrams for searching in biomedical graphs is completely new and potentially promising thanks to their ability to encode compactly sets by exploiting their structure and regularity, and (ii) to manipulate entire sets of elements at once, instead of exploring each single element explicitly. This work shows that search strategies based on Decision Diagram makes the indexing for biochemical graphs, and not only, more scalable and affordable allowing us to potentially deal with huge and ever growing collections of biomedical structures.

**Keywords:** query processing; pattern matching; subgraph isomorphism; graph indexing; decision diagrams

## 1 Introduction

Graphs are a well-known mathematical structure used to encode relationships among elements of a set. They are employed in the representation of many biochem-

ical systems at various levels from molecular representation [1], to protein-protein or RNA-mate interaction networks [2, 3] including also disease characterization [4]. The search for substructures, also called embeddings, in biochemical structures is widely involved in many bioinformatic approaches as well as in the field of computational chemistry. It is a preliminary step in findings motifs in biological networks [5, 6, 7], or in tuning parameters in biomolecular and molecular simulations [8]. In addition, it is a base procedure in biomedical database systems [4] for drug repurposing studies [9], for prioritizing disease-associated genes [10], or for knowledge discovery and hypotheses generation in plant and crop researchers [11]. In details, the problem consists in searching a query graph within a target graph, and it is a well-studied computational problem which is known to be NP-Complete [12]. A generalization of such formulation was also proposed, in which more than one target graph is considered. This is typically referred to as *one-to-many* in contrast to the original formulation that is referred to as *one-to-one*. Techniques for solving the one-to-one problem are mainly based on heuristics to speed-up the searching of mapping function. Instead, the main efforts for solving the one-to-many problem are focused on developing a good filtering strategy for discarding target graphs belonging to the collection that do not contain the query graph. In particular, the most effective methodology for filtering strategy is the creation of an index in which features of the target graphs are stored. Then, when searching for a specific query graph, the target graphs are filtered by comparing the features of the query to those of the target graphs via the index. Thus, indexes are aimed at provide a very compact representation of the set of features and their correspondence to the target graphs. Performance in terms of construction time, size, querying time and filtering power are key concepts for their development. Such a performance is strictly related to the type of feature that is taken into account.

In details one-to-one approaches can be divided in two categories: *pure subgraph isomorphism* and *assisted solvers*. The first category is composed by algorithms that are focused on improving the performances of the combinatorial search by exploiting heuristic methods for pruning the search space, such as VF2 [13] and VF3 [14], or by changing the order in which query vertices are matched, such as RI [15]. The second category comprises algorithms able to efficiently reduce the number of target vertices that are candidate to match with query vertices. This reduction is obtained by indexing the target graph and by comparing the features assigned to target vertices with those of the query vertices. Indexing means that a predefined type of features are extracted from the graph and they are stored in a data structure in order to recognize in which part of the graph, or in which graph of a collection, a given feature occurs. Once candidates are retrieved, this information is also used for generating a quasi-optimal ordering of the query vertices. In this perspective, GraphQL [16] uses a pseudo subgraph isomorphism test, while TurboISO [17] exploits a tree-structured auxiliary index, and CSL [18] postpones cartesian products with a matching order that prioritizes the query vertices in the core structure, similar to RI.

One-to-many approaches can be differentiated by the type of features they take into account (e.g. paths, trees, cycles or subgraphs) and how they extract them.

---

[1]<https://het.io/>

GraphGrep [19], GraphGrepSX [20], GRAPES [21] and SING [22] extract paths by indexed graphs with simple enumeration procedures, but they differ in the type of data structure and additional information they use. Simple enumeration is also used by CT-Index [23] for extracting trees and cycles, and by GDIndex [24] and GCode [25] for extracting subgraphs. On the contrary, mining-based algorithms recognize *frequent* features with *ad hoc* procedures. SwiftIndex [26] and TreePi [27] extract frequent trees, as well as Tree+Delta [28] which also retrieves frequent substructures. Mining of subgraph is also performed by CP-index [29], gIndex [30], FG-Index [31] and Lindex+ [32]. Alternatively, signatures based on the pairs of vertex labels of the graphs can be exploited [33]. Mining-based approaches require high amount of time because of the mining step, however they are able to build more compact indexes w.r.t. the approaches based on the exhaustive enumeration.

In recent years, one-to-one approaches have reached a high performance. In many cases, they outperform the indexing methodologies of one-to-many approaches by simply scanning all the target graphs in a collection. However, when the number of graphs in the collection is relatively high, or when the target graphs have relatively large size, indexing techniques are still predominant, and hybrid approaches are investigated [34]. In [35], authors proposed an algorithm for the one-to-many problem which exploits a technique that it is usually embedded in one-to-one approaches, such as GraphQL, TurboIso and CFL. The technique consists in a pre-processing step for detecting the set of target vertices that are most probable to be matched with a given query vertices by looking at their connectivity. Authors have equipped the verification phase of GraphGrepSX, GRAPES and CT-Index with such a technique showing that modified one-to-many algorithm, in particular GRAPES, sensibly outperform GraphQL, TurboIso and CFL for the verification step. However, such a modification is added up to the original indexing techniques of the algorithms, thus it only helps in increasing the filtering power but it does not solve problems linked to the size and build time of the original index. Similar considerations can be done for cache-assisted frameworks [36, 37]. In this perspective, compression of the index plays a central role for both one-to-one and one-to-many approaches [38, 39].

A performance study [40] finds that index-based approaches have several issues in building indices on large graph databases in terms of number of distinct labels, number of vertices in data graphs, density of target graphs and number of target graphs due to their poor time and space efficiency of index construction. Among the tested algorithms, GRAPES showed the best performance in terms of running time, however its index requires a relatively high memory amount compared to the other approaches. GRAPES is implemented both as sequential and parallel software using symmetric multiprocessing (SMP) architectures. In addition, GRAPES has been developed for achieving good performance in collection of graphs as well as in scanning a query over a single large target graph. For this reasons, we decided to improve the performance of the sequential version of GRAPES by reducing the memory required for its index. We investigated the use of decision diagrams for reaching the goal without degrading the running time of the algorithm. Synthetic graphs were engaged for evaluating the performance of the modified version. In addition, a well established collection of biochemical graphs have been for testing.

Results show that the modified version, called GRAPES-DD, can reduce the size of the index of a factor of five order of magnitude. The reduce index size helps the algorithm in optimizing accesses to primary memory, and as a result it can speed-up the building time of GRAPES-DD w.r.t. the original algorithm in the same situations. GRAPES-DD is available at the following online repository <https://github.com/dd-user/grapes-dd>.

## 2 Background

### 2.1 Path-based graph indexing

Graph indexing strategies based on labelled paths consist in extracting all the paths in the graphs up to a given length (number of nodes which they are composed) and succinctly storing them into a data structure [41, 19, 20, 42]. These techniques show good performance in terms of filtering power and construction/querying time. However, the size of the index is still a major problem with these approaches. GRAPES [21] provides one of the most recent implementation of path-based graph indexing. It searches a query graph into one or more target graphs. For each path, GRAPES stores the identification of its starting vertices and the number of its occurrences in each graph. During searching phase, paths are extracted from the query graph and searched in the index. GRAPES by comparing the ordered sequence of labels and the count of the occurrences, effectively filters out targets graphs which are improbable to contain the query graph.

Formally, a path in a graph is an ordered sequence of vertices such that each vertex is connected with the next vertex by an edge. Thus, given a graph  $G = (V, E)$ , where  $E : V \times V$ , a path  $p$  of length  $l$  is a vector  $(v_1^p, v_2^p, \dots, v_l^p)$  such that  $v_i^p \in V$ , for  $1 \leq i \leq l$ , and  $(v_i^p, v_{i+1}^p) \in E$ , for  $1 \leq i < l$ . Given a set of labels  $\Sigma$ , a graph is labelled via a function  $f_\sigma : \Sigma \mapsto V$  which maps each vertex to a label in  $\Sigma$ . The same label can be associated with different vertices. A labelled path  $\hat{p}$  is obtained by mapping the vertices of a path to their corresponding labels via the  $f_\sigma$  function, thus  $\hat{p} = (f_\sigma(v_1^p), f_\sigma(v_2^p), \dots, f_\sigma(v_l^p)) = (\sigma_1^p, \sigma_2^p, \dots, \sigma_l^p)$ . In GRAPES labelled paths are stored in a trie, a tree structure which compacts paths by their longest common prefixes. Given two labelled paths,  $\hat{p} = (\sigma_1^p, \sigma_2^p, \dots, \sigma_l^p)$  and  $\hat{q} = (\sigma_1^q, \sigma_2^q, \dots, \sigma_l^q)$ , that share the first  $i$  labels,  $(\sigma_1^p, \sigma_2^p, \dots, \sigma_i^p) = (\sigma_1^q, \sigma_2^q, \dots, \sigma_i^q)$ , a branch, starting from the root of the tree, is built in order to represent the shared part of the paths. Then, the branch is split into two different branches that represent the non shared suffixes of the paths,  $(\sigma_{i+1}^p, \dots, \sigma_l^p)$  and  $(\sigma_{i+1}^q, \dots, \sigma_l^q)$ . Information regarding the starting vertices,  $v_1^p$  and  $v_1^q$ , is stored on the corresponding leaves of the branches, as well as the number of time each path occurs in each target graph. If only paths of the same length were extracted, the information would reside only on the leaves of the trie. By considering paths of variable length up to a maximum length  $l_p$ , the information also resides on intermediate nodes of the trie.

### 2.2 Filter and verification

During querying phase, labelled paths are extracted from the query. Similarly to the extraction of paths from target graphs, for each path the number of times it occurs in the query graph is retrieved. Initially, all the target graphs are candidates to contain the query graph. Query paths are searched in the index in order to recognize the

target graphs that contain the same paths of the query. For each path, the number of occurrences within the target graph must be equal or exceed the number of its occurrences in the query graph. By using the starting nodes of the paths stored in the index, the initial structures of target graphs are skimmed in order to extract only the vertices that are the starting point of paths that are present in the query graph. Thus, the filtering procedure produces two different results, a list of graphs that may contain the query (since each selected graph contains the same labelled path of the query with the same amount), and for each selected graph the list of vertices that are candidate to match with the query vertices. The verification phase is performed with the VF2 algorithm[13] which solves the subgraph isomorphism problem. The problem of searching a query graph within a target graph consists in finding a mapping between the vertices of the query and target graphs such that constraints are satisfied. Constraints regard the compatibility of labels assigned to the vertices and the existence of the query edges between the corresponding query-target mapped vertices.

### 2.3 Decision diagrams

Decision diagrams (DDs) are a broad class of data structures widely used to encode and manipulate functions efficiently [43]. Initially, they were proposed for industrial hardware verification due to their ability of encoding complex Boolean functions on very large domains. Then, they were successfully applied in different research fields ranging from the network reliability analysis [44] to the performance evaluation of stochastic systems [45]. In these contexts DDs have proven to be an effective tool (i) to encode compactly structured sets by exploiting their structure and regularity; (ii) to manipulate entire sets of elements at once, instead of exploring each single element explicitly.

Formally a DD is a directed acyclic graph that represents functions of type  $f : \mathcal{X}_1 \times \dots \times \mathcal{X}_k \rightarrow \mathcal{R}$  where each variable  $\mathcal{X}_k$  can assume a set of values over a finite domain  $\mathcal{D}_k$ . Hence, a DD node can be either *non-terminal*, in which case it is labeled by a variable  $\mathcal{X}_k$ , or *terminal*, which has no outgoing edges and is labeled with a constant  $c \in \mathcal{R}$  representing the return value of the function. Obviously, different variable types, variable domain, and function co-domain lead to various types of DDs.

For instance, a *Binary Decision Diagram* (BDD) represents a boolean function on boolean variables, of the form  $f : \{0, 1\}^k \rightarrow \{0, 1\}$ . Over BDDs the set operations (e.g. union, intersection, complement, ...) are defined to manipulate them. A *Multi-Way DD* (MDD) is an extension of a BDD representing a boolean function on multivalued variables, of the form  $f : \mathcal{X}_1 \times \dots \times \mathcal{X}_k \rightarrow \{0, 1\}$ , where  $\mathcal{D}_i$  is the domain of variable  $\mathcal{X}_i$ . Similar to BDDs, MDDs can be manipulated through set operators. A *Multi-Terminal MDDs* (MTMDD) is instead an extension of an MDD representing a multivalued discrete functions on multivalued variables, of the form  $f : \mathcal{X}_1 \times \dots \times \mathcal{X}_k \rightarrow \mathcal{R}$ , with either  $\mathcal{R} \subset \mathbb{N}$  or  $\mathcal{R} \subset \mathbb{R}$ . Over MTMDDs are defined both logical and arithmetic operations.

A DD is denoted *ordered DD* (ODD) when an order is defined on its variables (i.e.,  $\mathcal{X}_l > \mathcal{X}_k \Rightarrow l > k$ ) such that every path through the DD visits nodes according to this ordering. An ODD without duplicates nodes (i.e., non-terminal nodes labeled

with the same variable and with identical children) is a *canonical representation*, which means that given an ordering of the variables, there is a one-to-one correspondence between the DD and the function it represents. This is a very desirable property for determining formal equivalence.

The evaluation of the function represented by such an ODD, for a given assignment of its variables, can be determined by following a path from its root node to one of the terminal nodes. It is important to notice that the choice of an ordering for the variables of a DD can strongly affect the size of the DD, i.e. its number of nodes. Unluckily, finding the optimal variable ordering is known to be a co NP problem [46].

Different reduction rules have been proposed to reduce significantly the number of nodes without affecting the represented function. An ordered decision diagram is called *reduced* (RODD) if it is an ODD and it does not contains any *redundant* nodes. Redundant nodes are non-terminal nodes having all identical children; in this case, the value of the function does not depend on the value of the variable represented by the node. A common variation of the above reduction rule is to allow redundant nodes but no duplicates. An ODD is called *quasi-reduced* (QRODD) if contains no duplicate nodes and if all paths from the root node to a terminal node visit exactly one node for each variable.

In the literature different software libraries implementing decision diagrams were proposed, such as CUDD [47], LibDDD [48] and Meddly [49]. In this work, we chose to use Meddly because of its efficiency and its ease of use. In fact, it provides a simple interface in which the complex aspects of using DDs (e.g. caching and garbage collection, ...) are automatically handled. Meddly, short for Multi-way and Edge-valued Decision Diagram LibrarY, is an open-source software library supporting natively all the types of DDs introduced above, as well as others such as Matrix Diagrams, Edge-Valued MDDs, etc. All DDs represented in Meddly are ordered and canonical.

A named collection of decision diagrams associated with the same domain is called a *forest*. Within a forest, Meddly automatically removes duplicate nodes by means of a *unique table*[50], imposes reduction rules and handles memory management of the nodes.

Meddly provides two different user interfaces: a *basic interface* which provides the basic operators to easily create and manipulate DDs, and an *expert interface* which allows user to extend the existent operators and/or to integrate new ones. In this work, we implemented our tool by taking advantage of the basic interface of Meddly; in particular the following operators were exploited:

- `createEdge()` creates a new DD in the given forest by explicitly stating the return values for a set of variable assignments. Unspecified assignments are assumed to return a default value, which depends on the forest type (usually it is 0). For example, given the forest  $F$  and some variable assignments  $Y = (y_1, \dots, y_k)$ ,  $W = (w_1, \dots, w_k)$  and  $Z = (z_1, \dots, z_k)$ , a call to `F.createEdge(Y, W, Z, a, b, c)` creates a new DD within  $F$  representing the

function

$$f(x_1, \dots, x_k) = \begin{cases} a & \text{if } x_1 = y_1 \wedge \dots \wedge x_k = y_k \\ b & \text{if } x_1 = w_1 \wedge \dots \wedge x_k = w_k \\ c & \text{if } x_1 = z_1 \wedge \dots \wedge x_k = z_k \\ 0 & \text{otherwise} \end{cases}$$

- **evaluate()** determines the value of the function represented by the DD for a given assignment of its variables. Then, the call to **evaluate**(*dd*,  $x_1, \dots, x_k$ ) returns the terminal value linked to the path  $x_1, \dots, x_k$  of the decision diagram *dd*.
- **apply()** is used to manipulate DD applying on it a specific DD operator. Meddly supports both unary and binary operators and imposes that operands of binary operators must have the same domain, but they can live in different forests. Given a binary operator  $\diamond$  and two DDs *dd*<sub>1</sub> and *dd*<sub>2</sub> representing the functions  $f_1$  and  $f_2$ , respectively, a call to **apply**( $\diamond$ , *dd*<sub>1</sub>, *dd*<sub>2</sub>) creates a new DD encoding the function  $f_1(x_1, \dots, x_k) \diamond f_2(x_1, \dots, x_k)$ . For instance, the MTMDD multiplication between *dd*<sub>1</sub> and *dd*<sub>2</sub> builds an MTMDD where the terminal node of a path *p* is equal to the product of the evaluations of *p* in *dd*<sub>1</sub> and *dd*<sub>2</sub>.  
To speed-up the apply operation a cache, namely *computed table*, is exploited to avoid recomputing the recent performed computations [50].

### 3 Methods

GRAPES uses a trie, also known as prefix tree, to store the indexed graphs, since it provides a compact representation of a set of strings by taking advantage of their common prefixes, considerably reducing the data redundancy. Nevertheless, the tree structure of a trie (i.e. only a single edge can point to a node) makes hard to exploit other types of symmetries present in the indexed graphs, as for instance the sharing of the same starting vertices and/or same relative occurrence number, as well as the sharing of common substrings which are not prefixes.

To deal with these aspects, in this work we propose to encode the indexed graphs into a DD, specifically an MTMDD: a trie generalization in which the requirement to have a tree structure is relaxed allowing multiple arcs to point to the same node. The main advantage of this is a potentially more compact representation due to the MTMDD ability to better exploit the regular structure of the data, such as common substrings present in the indexed graph paths. This allows the proposed methodology not only to reduce the memory utilization required to build and store the index, but also to reduce the time required for the pruning phase.

The GRAPES-DD workflow is reported in Figure 1. The workflow is composed by two main phases: (i) the *indexing building phase* in which an MTMDD indexing the collection of target graphs is created, and (ii) the *filtering phase* in which, given a query graph, the set of target graphs is potentially restricted to those sub-graphs probably containing the query. The GRAPES-DD verification phase remains as in the original version of the software (see Section 2.2 and [21] for details).



### 3.1 Indexing building phase

The indexing building phase takes as input a collection of target graphs  $\mathbb{G}^t = \{G_1^t, G_2^t, \dots, G_n^t\}$  and the maximum path length  $l_p$  and returns as output an MTMDD that maps each path to the total number of times it appears in the graph for a specific input vertex. We will refer to this MTMDD as *index MTMDD*.

In details, the first level of the index MTMDD stores the identification of the vertices of the indexed graphs. Then, the labelled paths are stored from the second to the last level of the MTMDD, one label per level starting from the first label of the path. Finally, the total occurrence number of the labelled paths in each indexed graph resides on the terminal nodes of the MTMDD.

An example of such a data structure is reported in Figure 1(b). This MTMDD is created considering the three target graphs,  $G_1$ ,  $G_2$  and  $G_3$  in Figure 1(a) and  $l_p = 3$ . Since in the first level of the MTMDD all the vertices of the target graphs are enumerated then its domain is  $[1, 15]$ . Then, for each vertex in  $G_i^t$  all the labelled paths up to length 3 starting from it are added in the MTMDD. Special nodes, namely *unlabelled* nodes and colored gray in Figure 1(b), are introduced for deal with labelled paths having length smaller than  $l_p$ . This is needed because MTMDDs cannot directly encode paths with different lengths.

Practically, the index MTMDD is created in an incremental manner, processing one graph at the time. Thus, the vertices of each graph are initially grouped based on their labels. Hence, for each of these groups, all the labelled paths containing up to  $l_p$  vertices are retrieved by depth-limited search on the graph. These paths are stored temporally in a trie to efficiently count the occurrence of each labelled path in  $G_i^t$ . Before considering a new group, the trie is explored to create the corresponding MTMDD using the `creatEdge` operator of Meddly and then is discarded. The created MTMDD is merged to the index MTMDD using the `apply` function with operator `addition`. Once the building phase is over, the MTMDD index is persisted on a text file reporting also the labels and vertices of graphs mapping the paths.

Figure 2 shows the building steps for the target graph  $G_1$  with  $l_p = 3$ . The vertices of  $G_1$  are indeed divided into four groups (i.e.  $\{1, 5\}$ ,  $\{2, 3\}$ ,  $\{4\}$  and  $\{6\}$ ) depending on their label. For each group the corresponding temporary trie and MTMDD are reported. The original GRAPES index is the union of the tries shown in Figure 2. Thus, it contains the root node and 25 ( $9 + 8 + 4 + 4$ ) nodes representing labelled paths, plus 13 nodes for storing the occurrences count, 13 nodes for storing the ids of the starting vertices in correspondence of each path, for a total of 51 nodes. On the contrary, the final MTMDD structure contains a single node storing all the six vertices ids, 13 nodes for the labelled paths and one node for the occurrence count, for a total of 15 nodes. In addition, the trie stores 39 links between its nodes, while the MTMDD stores 23 links.

### 3.2 Filtering phase

The *filtering phase* takes as input the index MTMDD and a query graph  $G^q$  that has to be searched within the graph collection. As output, the filtering phase provides for each graph the list of vertices candidate to match the query. Therefore, only these vertices will be subsequently tested using the subgraph isomorphism algorithm.

The algorithm initially builds the *query MTMDD* to represent the query graph through the use of its features (i.e. paths), which is shown in Figure 1(d). The

vertices of the query graph are not represented in the corresponding MTMDD. The first level of such MTMDD contains all the vertices ids of  $G_i^t$ , meaning that initially each vertex of  $G_i^t$  is candidate to match any path of  $G^q$ .

The multiplication operator (using the `apply` function) is applied between the index and the query MTMDDs, in order to extract from the index the information about the vertices really involved in the current query. We called *pruned MTMDD* the result of such multiplication, which is depicted in Figure 1(e). We see that only the subgraph composed by the vertices  $\{1,2,6\}$  is kept from  $G_1$ ,  $G_2$  is entirely kept and  $G_3$  is totally discarded because it does not contain any feature in  $G^q$ .

The set of candidate vertices obtained is then filtered in order to keep only those graphs whose vertices satisfy the constraints imposed by the query. For each vertex  $v_q$  of the query and a potentially matching vertex  $v$  of a target graph  $G_i$ , the algorithm verifies that (i) any path starting from  $v_q$  also starts from  $v$  and that (ii) the occurrence number of each path in the target graph is not less than the occurrence number of the same path in the query graph. Figure 1 shows that the graph  $G_1$  is filtered out because the occurrence number of its features are not sufficient to satisfy the constraint imposed by  $G^q$ , while all the vertices of  $G_2$  passed the filtering phase.

Finally, for each vertex of the query, the algorithm reports the list of the matchable vertices of the indexed graphs passing the pruning phase. The overall effect is that the algorithm extracts from the graph collection all the maximally connected components composed only by the vertices involved in the query graph. Over these components, the GRAPES subgraph isomorphism algorithm can be run to find all the occurrences of the query graph [21].

## 4 Results

### 4.1 Datasets description

For this study, we considered six different types of graphs. Four of them are real graphs widely used as a benchmarks in the fields of bioinformatics and computational chemistry, the others are synthetically generated by means of the Barabasi-Albert's and the Forest-Fire models. the choice of such two synthetic models has been taken according to their properties of the topologies to be similar the graphs used in biomedical databases. Differently from collections of real graphs, synthetic topologies allow us to investigate the performance of compared methods in relation to the parameters of such models, and thus to the properties of the produced topologies.

#### *Biochemical structures*

The collection of biochemical graphs was initially used for evaluating the performance of one-to-one subgraph isomorphism algorithms [51], and, nowadays, it is a well-established benchmark for graph theory problems linked to the subgraph isomorphism [52]. These four datasets that compose the collection are described in what follows.

*AIDS* is the standard database for Antiviral Screen [53], and it consists of 40k chemical structures representing small molecules. Vertices are atoms and edge are the chemical bounds linking them. Vertices are labelled by the atomic element they

represent, and there are a total of 62 distinct elements. The average number of vertices per graph is 44.98, and the average degree is 4.17.

*PDBS* is a benchmark composed of 600 target graphs representing the topological structure of proteins [54, 55]. Vertices are the atoms and edges are chemio-physical bounds between them. These graphs have up to 16,431 vertices and 33,562 edges, with an average degree over the whole dataset equal to 4.27. There are a total of 10 unique labels, corresponding to the atomic types.

*PCM* is composed of three-dimensional structures of protein, called protein contact maps [56]. Vertices represent the amino acids of a protein and edges inform about the spatial proximity of amino acids. The dataset contains 200 target graphs having up to 883 vertices and 18,832 edges, with an average of 376 vertices per graph and 44.78 edges per vertex. There are a total of 21 labels of which 18 appears on average in each graph.

*PPI* is a dataset of 20 protein-protein interaction target graphs of 5 different species: *Caenorhabditis elegans*, *Drosophila melanogaster*, *Mus musculus*, *Saccharomyces cerevisiae* and *Homo sapiens* [57]. Vertices are proteins and edges are predicted physical interactions between them. For each species, different thresholds on the accurateness of the prediction were applied, ranging from 0.4, 0.5, 0.6 to 0.7. Vertices are labelled according to their functional category, for a total of 45 distinct categories. The dataset contains graphs up to 10,186 vertices and 179,348 edges, an average degree of 18.46 and an average number of distinct labels per graph equal to 28.45.

For all of the biochemical datasets, queries were extracted from the target graphs by fixing the desired number of edges, from 4, 8, 18 to 32, and such that the topological structure of the extracted graph reflects the properties of the graph of origin.

### *Synthetic graphs*

The Barabasi-Albert's model is able to reproduce graph with an observed stationary scale-free distribution, which reflects many of the structures that can be encountered in nature [58]. Starting from an initial set of vertices,  $m_0$ , the model inserts one vertex at time to the graph. At each insertion, new edges are added in order to connect the new vertex with existing ones. The probability of an edge with vertex  $i$  is  $p_i = k_i^\alpha$ , where  $k$  is the vertex degree and  $\alpha$  is a user defined parameter. The benchmark contains 384 target graphs which were generated by fixing a desired number of vertices and average degree. Generated graphs have 200, 500, 1k, 5k, 10k and 20k vertices. In addition, three copies of each generated network are made in order to provide a labelled version of the initial structure with three different percentages of distinct labels, 0.1%, 1% and 10%. Labels are assigned randomly to vertices.

The second type of synthetic graphs were generated according to the Forest-Fire model [59], that is inspired by forest growing behaviours, and which imposes a geometric distribution with mean  $p/(1-p)$  for the probability of linking two distinct vertices. This benchmark contains 160 target graphs having the same number of vertices of the Barabasi-Albert benchmark, and they were labelled in the same way of the previous model. Moreover, the graphs were generated by varying the value of the model parameter  $p$  as 0.1, 0.3, 0.5, 0.7 and 0.9.

For both synthetic benchmarks, query graphs were extracted from the generated target graphs. The extraction was performed by fixing the number of desired vertices, ranging from 4, 8, 24, 32 to 64, and by extracting all edges among the selected vertices.

For a more details regarding the two sets of synthetic benchmarks, the reader can also refer to [60].

#### 4.2 Experimental setup and output

We evaluated the performance of GRAPES-DD, w.r.t. its predecessor GRAPES, taking into account both space and time requirements. The analysis was mainly focused on the index construction phase because it is the main difference between the two approaches. They share the same methodology for what concerns the matching phase. In addition, the filtering time can be considered negligible w.r.t. the total querying time.

Both GRAPES-DD and GRAPES have been containerized in a Docker[61] image in order to ensure both functional and computational reproducibility of the experiments. The dockerfile to build the image is provided together with the source code. Both the tools were implemented in C++ and compiled with gcc 6.3.0. Then, the experiments have been carried out on a server equipped with four processors AMD Opteron 6167 2.20 Ghz and 502 GB of RAM. Since GRAPES is a natively parallel software while GRAPES-DD is sequential, the experiments were executed using GRAPES with a single-thread.

For both types of benchmark, biochemical and synthetic, we evaluated the capacity of GRAPES-DD to construct a more compact index w.r.t. GRAPES, and the difference in the running time of the two approaches. We evaluated the amount of primary memory that the two approaches require during the execution, that is reported as *memory peak*, as well as the space needed to store the built index in the hard disk, that is reported as *index size*. In addition, we compared the running time required by the two approaches for building the index.

Figures 3 and 4 show memory peak and index size on the synthetic datasets obtained by indexing one target graph at time. Values are calculated taking into account three different grouping strategies that reflect the way in which the datasets are generated. Plots were generated via the *Pandas* framework available for Python [2]. In details, datasets were grouped by (i) percentage of distinct labels w.r.t. the total number of vertices of the graph, (ii) number of vertices and (iii) value of the Barabasi-Albert model parameter  $\alpha$  or Forest-Fire parameter  $p$ .

Results show that, independently from the label percentage and model parameters, the performance of GRAPES-DD improves as the number of vertices of the indexed graph increases. In fact, for graphs having less than 5k vertices, the memory peak required by GRAPES-DD is higher than the peak of GRAPES, resulting in a ratio between the two values less than 1. The out-performance of GRAPES-DD can reach two/three orders of magnitude w.r.t. GRAPES, for graphs with 20k vertices, which means that the memory requirement of GRAPES-DD is one hundredth that of GRAPES.

---

[2] <https://pandas.pydata.org/pandas-docs/stable/reference/api/pandas.DataFrame.boxplot.html>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

Similar trends are observed for the size of the index when it is stored into the hard disk. In this case, the ratio can achieve five orders of magnitude as it is shown for the Forest-Fire graphs with 20k vertices. In general, best ratios are obtained for the Forest-Fire graphs with a high number of vertices, however, this behaviour is counterbalanced by the fact that on average Forest-Fire graphs with less than 5k vertices are also those with the lowest ratios.

For what concerns the memory peak, we can observe that the label percentage is a more crucial factor for the Barabasi-Albert model rather than for the Forest-fire model. More in general, a low label percentage is to the advantage to the trie structure of GRAPES because the extracted paths share and relatively high number of labels. Opposite trends are observed for what concerns the storing of the index.

As for the label percentage, model parameters produce less variation compared to the number of vertices. The Barabasi-Albert model produce scale-free networks where the distribution of the degrees of the vertices follows a power law. A value greater than 1 increases the skewness of the resultant distribution, while a value less than 1 makes the distribution more flatten. Thus higher values trend to produce a more sparse graph. Results in Figure 3 show that GRAPES-DD performs better for dense graphs, namely for low values of the  $\alpha$  parameter. The trend is confirmed by the results regarding the Forest-fire models (Figure 4), where higher values of the  $p$  parameters produce more dense graphs.

Subsequently, we evaluated the performance of exploiting the MTMDD structure for indexing 14 collections of synthetic graphs (see Table 1). The first three collections are obtained by grouping Barabasi-Albert graphs by the label percentage, such that graphs having the same percentage are contained in the same collection. Similarly, Forest-Fire graphs were grouped into other three collections. The grouping procedure was also performed by taking into account the  $\alpha$  and  $p$  model parameters. As for the previous analysis, the ratio is computed by dividing the values measured for the trie structure of GRAPES with those registered for the MTMDD of GRAPES-DD. As it has been shown for the single-graph analysis, the percentage of distinct labels (w.r.t. the total number of vertices in the graphs) is a discriminant factor for the compression gain obtained by the MTMDD structure. Also the trends relative to the parameters of the models are confirmed. In general, the MTMDD structure is on average more convenient on the Forest-Fire graphs for what concerns the memory peak. Barabasi-Albert graphs with  $\alpha = 0.5$  are an exception to this trend, since they reach the maximum registered ratio equal to 45. In contrast to the single-graph analysis, the space required for storing the index into the hard disk does not provide the same advantage to the MTMDD structure. In fact, in the single-graph analysis the ratio reaches a value of  $10^5$  that is two order of magnitude higher of the ratios obtained for the memory peak. On the contrary, these experiments show an inversion of the ratio such that the MTMDD structure reaches best results for the memory peak. It is notable to report that, while the trie structure requires a maximum of 94Gb of memory, the process for building the MTMDD-based index does not reaches the 9Gb of requirement, making it suitable for common personal computers.

Table 1 also shows the running time of the two approaches for building the index and for storing it. The MTMDD structure requires more time for its construction,

the compression capability of the MTMDD must come with an unavoidable additional cost. However, the growth time is only a few minutes and the construction of the index is performed in a preprocessing phase, only once and reused for each query search.

Table 2 reports the complete set of experiments that were performed on the biochemical graphs. The experiments regards the indexing of the four different collections of real graphs. For this benchmark, ratios are less prominent w.r.t. synthetic graphs, however many of them are higher than 1, confirming a gain in using the MTMDD structure rather than the trie. The trend for which paths extracted from more dense and more uniform graphs are better compacted by the MTMDD structure is confirmed. In fact, the best ratio is obtained for the *PCM* collection that contains the most dense graphs. However, the *PCM* collection is also the one with the lowest number of labels and a relatively small number of vertices. Thus, it seems that the density of the graphs is the key factor for the good performance of GRAPES-DD in biochemical graphs. In addition, in contrast with the results on the synthetic graphs, the running time of GRAPES-DD for the construction of index is generally faster than the time required by GRAPES. In these cases, the compression capability of the MTMDD comes without additional cost.

Lastly, Table 3 reports the results regarding the PPI networks obtained by indexing one PPI at time, since PPI networks are often analysed stand-alone. Similarly to the synthetic networks, the increase of the graph size, i.e. number of vertices  $|V|$  and number of edges  $|E|$ , results in a better performance of GRAPES-DD w.r.t. GRAPES. However, comparing the ratios obtained for *M. musculus* and *H. sapiens* we can deduce that as expected there is not a fixed correlation between the graph size and the performance. Therefore, the intrinsic nature of the graph is also responsible for these results. PPI networks are also the targets for which running times of GRAPES-DD are comparable to those of GRAPES, and some times they are even better. The GRAPES-DD building approach includes the construction of partial tries but without merging them. The cost for traversing a single whole trie may limit GRAPES.

## 5 Discussion

In this study, we afford the problem of reducing the indexing size of biomedical graph searching systems to make them effective with the increasing size of the structures. We have shown that the indexing of labelled graphs can take the advantages of newly adapted data structures based on decision diagrams. These techniques allow already existing methodologies to increase their compression power, in terms of memory consumption, without significantly increasing the searching time requirement.

We examined synthetic graphs because they offer a more systematic way of investigating performance of indexing using decision diagrams. Since the type of the generated graphs reflects the structures that are found in nature, their analysis can be exploited for inferring performance behaviour of real biomedical structures. The results have shown that relevant indexing compression ratio can be obtained in relation with the size and the topological structure of the graphs and the distribution of labels within them. Moreover, the larger are the indexed graphs, the higher is the gain that the proposed methodology can reach than to the exploiting of the Decision Diagram data structure.

A well-established benchmark was also used for evaluating the performance on real graphs. The size of the considered graphs are relatively small, compared with the synthetically generated ones, however trends of gain ratio are confirmed. This must be considered in the perspective of future applications of the proposed indexing technique, because the continuous development of new technologies for extraction biological information leads to the construction of biomedical relational systems that constantly increase in size.

## 6 Conclusions

Nowadays, graphs are fundamental structures for representing and for investigating the current biological and biomedical knowledge. In this work we investigated the possibility to improve the performance of the cutting-edge algorithms for searching substructures in graphs based on indexing, by addressing one of their disadvantages which is the size of the index. To this aim, we developed GRAPES-DD, a new version of GRAPES tool, whose strength is the use of decision diagrams to substantially reduce the size of the index. Experimental results performed on a set of synthetic and real benchmarks reported clearly that the use of this data structure allows us to substantially reduce the memory footprint of the index (i.e up to 5 orders of magnitude smaller) w.r.t. the original version of GRAPES without impacting the running time of the algorithm.

Further enhancement of GRAPES-DD will be to re-implement the building phase to allow thread-based parallelization as in the original GRAPES implementation. Moreover, since the efficacy of decision diagram techniques is strictly dependent on the variable order, we will investigate how different algorithms for variable orderings behave and we will evaluate the possibility of developing meta-heuristics to identify a-priori the best variable ordering depending by the features of each target graph.

## 7 Declarations

### Ethics approval and consent to participate

Not applicable.

### Consent to publish

Not applicable.

### Availability of data and materials

Project name: GRAPES-DD

Project home page: <https://github.com/dd-user/grapes-dd>

Operating system(s): Linux

Programming language: C++, Docker

License: MIT

Any restrictions to use by non-academics: licence needed

### Competing interests

The authors declare that they have no competing interests.

### Funding

NL and MB were partially supported by "Creation of a computational framework to model and study West Nile Disease" project supported by CRT foundation (PI Marco Beccuti). VB and RG were partially supported by GNCS-INDAM and JPND 2019-466-037.

### Authors' Contributions

All the authors equally contributed to the development of the methodology. NL and VB contribute do the developed of the software. All the authors equally contributed to the writing of the manuscript. MB and RG supervised the work. All the authors have read and approved the final manuscript.

### Acknowledgements

Not applicable.

## Author details

<sup>1</sup>University of Turin, Turin, Italy. <sup>2</sup>University of Verona, Strada le Grazie 15, 37134 Verona, Italy.

## References

- Bonchev, D.: *Chemical Graph Theory: Introduction and Fundamentals* vol. 1. CRC Press, ??? (1991)
- Mering, C.v., Huynen, M., Jaeggi, D., Schmidt, S., Bork, P., Snel, B.: String: a database of predicted functional associations between proteins. *Nucleic acids research* **31**(1), 258–261 (2003)
- Bonnici, V., De Caro, G., Constantino, G., Liuni, S., D'Elia, D., Bombieri, N., Licciulli, F., Giugno, R.: Arena-idb: a platform to build human non-coding rna interaction networks. *BMC bioinformatics* **19**(10), 350 (2018)
- Goh, K.-I., Cusick, M.E., Valle, D., Childs, B., Vidal, M., Barabási, A.-L.: The human disease network. *Proceedings of the National Academy of Sciences* **104**(21), 8685–8690 (2007)
- Grochow, J.A., Kellis, M.: Network motif discovery using subgraph enumeration and symmetry-breaking. In: *Annual International Conference on Research in Computational Molecular Biology*, pp. 92–106 (2007). Springer
- Rinnone, F., Micale, G., Bonnici, V., Bader, G.D., Shasha, D., Ferro, A., Pulvirenti, A., Giugno, R.: Netmatchstar: an enhanced cytoscape network querying app. *F1000Research* **4** (2015)
- Sadeghi, M., Ordway, B., Rafiei, I., Borad, P., Fang, B., Koomen, J.L., Zhang, C., Yoder, S., Johnson, J., Damaghi, M.: Integrative analysis of breast cancer cells reveals an epithelial-mesenchymal transition role in adaptation to acidic microenvironment. *Frontiers in Oncology* **10**, 304 (2020)
- Allison, J.R., Welsh, I.D.: Cherrypicker: An algorithm for the automated parameterisation of large biomolecules for molecular simulation. *Frontiers in chemistry* **7**, 400 (2019)
- Himmelstein, D.S., Lizee, A., Hessler, C., Brueggeman, L., Chen, S.L., Hadley, D., Green, A., Khankhanian, P., Baranzini, S.E.: Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *Elife* **6**, 26726 (2017)
- Himmelstein, D.S., Baranzini, S.E.: Heterogeneous network edge prediction: a data integration approach to prioritize disease-associated genes. *PLoS Comput Biol* **11**(7), 1004259 (2015)
- Hassani-Pak, K., Singh, A., Brandizi, M., Hearnshaw, J., Phillips, A.L., Doonan, J.H., Rawlings, C.: Knetminer-intelligent search and visualisation of connected data to explain complex traits and diseases. *BioRxiv* (2020)
- Cook, S.A.: The complexity of theorem-proving procedures. In: *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, pp. 151–158 (1971)
- Cordella, L.P., Foggia, P., Sansone, C., Vento, M.: An improved algorithm for matching large graphs. In: *3rd IAPR-TC15 Workshop on Graph-based Representations in Pattern Recognition*, pp. 149–159 (2001)
- Carletti, V., Foggia, P., Saggese, A., Vento, M.: Introducing vf3: A new algorithm for subgraph isomorphism. In: *International Workshop on Graph-Based Representations in Pattern Recognition*, pp. 128–139 (2017). Springer
- Bonnici, V., Giugno, R.: On the variable ordering in subgraph isomorphism algorithms. *IEEE/ACM transactions on computational biology and bioinformatics* **14**(1), 193–203 (2016)
- He, H., Singh, A.K.: Graphs-at-a-time: query language and access methods for graph databases. In: *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, pp. 405–418 (2008)
- Han, W.-S., Lee, J., Lee, J.-H.: Turboiso: towards ultrafast and robust subgraph isomorphism search in large graph databases. In: *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pp. 337–348 (2013)
- Bi, F., Chang, L., Lin, X., Qin, L., Zhang, W.: Efficient subgraph matching by postponing cartesian products. In: *Proceedings of the 2016 International Conference on Management of Data*, pp. 1199–1214 (2016)
- Giugno, R., Shasha, D.: Graphgrep: A fast and universal method for querying graphs. In: *Object Recognition Supported by User Interaction for Service Robots*, vol. 2, pp. 112–115 (2002). IEEE
- Bonnici, V., Ferro, A., Giugno, R., Pulvirenti, A., Shasha, D.: Enhancing graph database indexing by suffix tree structure. In: *IAPR International Conference on Pattern Recognition in Bioinformatics*, pp. 195–203 (2010). Springer
- Giugno, R., Bonnici, V., Bombieri, N., Pulvirenti, A., Ferro, A., Shasha, D.: Grapes: A software for parallel searching on biological graphs targeting multi-core architectures. *PLoS one* **8**(10) (2013)
- Di Natale, R., Ferro, A., Giugno, R., Mongiovi, M., Pulvirenti, A., Shasha, D.: Sing: Subgraph search in non-homogeneous graphs. *BMC bioinformatics* **11**(1), 96 (2010)
- Klein, K., Kriege, N., Mutzel, P.: Ct-index: Fingerprint-based graph indexing combining cycles and trees. In: *2011 IEEE 27th International Conference on Data Engineering*, pp. 1115–1126 (2011). IEEE
- Williams, D.W., Huan, J., Wang, W.: Graph database indexing using structured graph decomposition. In: *2007 IEEE 23rd International Conference on Data Engineering*, pp. 976–985 (2007). IEEE
- Zou, L., Chen, L., Yu, J.X., Lu, Y.: A novel spectral coding in a large graph database. In: *Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology*, pp. 181–192 (2008)
- Shang, H., Zhang, Y., Lin, X., Yu, J.X.: Taming verification hardness: an efficient algorithm for testing subgraph isomorphism. *Proceedings of the VLDB Endowment* **1**(1), 364–375 (2008)
- Zhang, S., Hu, M., Yang, J.: Treepi: A novel graph indexing method. In: *2007 IEEE 23rd International Conference on Data Engineering*, pp. 966–975 (2007). IEEE
- Zhao, P., Yu, J.X., Philip, S.Y.: Graph indexing: Tree+  $\delta$  = graph. In: *VLDB*, vol. 7, pp. 938–949 (2007)
- Xie, Y., Yu, P.S.: Cp-index: on the efficient indexing of large graphs. In: *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, pp. 1795–1804 (2011)
- Yan, X., Yu, P.S., Han, J.: Graph indexing: a frequent structure-based approach. In: *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, pp. 335–346 (2004)
- Cheng, J., Ke, Y., Ng, W., Lu, A.: Fg-index: towards verification-free query processing on graph databases. In: *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, pp. 857–872 (2007)



32. Yuan, D., Mitra, P.: Lindex: a lattice-based index for graph databases. *The VLDB Journal* **22**(2), 229–252 (2013)
33. Pal, D., Rao, P., Slavov, V., Katib, A.: Fast processing of graph queries on a large database of small and medium-sized data graphs. *Journal of Computer and System Sciences* **82**(6), 1112–1143 (2016)
34. Katsarou, F., Ntarmos, N., Triantafillou, P.: Hybrid algorithms for subgraph pattern queries in graph databases. In: 2017 IEEE International Conference on Big Data (Big Data), pp. 656–665 (2017). IEEE
35. Sun, S., Luo, Q.: Scaling up subgraph query processing with efficient subgraph matching. In: 2019 IEEE 35th International Conference on Data Engineering (ICDE), pp. 220–231 (2019). IEEE
36. Wang, J., Ntarmos, N., Triantafillou, P.: Indexing query graphs to speed up graph query processing (2016)
37. Wang, J., Ntarmos, N., Triantafillou, P.: Graphcache: a caching system for graph queries (2017)
38. Gouda, K., Hassaan, M.: Compressed feature-based filtering and verification approach for subgraph search. In: Proceedings of the 16th International Conference on Extending Database Technology, pp. 287–298 (2013)
39. Luaces, D., Viqueira, J.R., Pena, T.F., Cotos, J.M.: Leveraging bitmap indexing for subgraph searching. In: EDBT, pp. 49–60 (2019)
40. Katsarou, F., Ntarmos, N., Triantafillou, P.: Performance and scalability of indexed subgraph query processing methods. *Proceedings of the VLDB Endowment* **8**(12), 1566–1577 (2015)
41. James, C.A., Weininger, D., Delany, J.: Daylight theory manual. daylight chemical information systems. Inc., Irvine, CA (1995)
42. Sæther, A.A.: Path-based graph indexing for keyword search on rdf data-discovering concepts through community detection. Master's thesis, NTNU (2017)
43. Burch, J.R., Clarke, E.M., McMillan, K.L., Dill, D.L., Hwang, L.-J.: Symbolic model checking: 1020 states and beyond. *Information and computation* **98**(2), 142–170 (1992)
44. Veeraraghavan, M., Trivedi, K.S.: An improved algorithm for the symbolic reliability analysis of networks. In: Proceedings Ninth Symposium on Reliable Distributed Systems, pp. 34–43 (1990). IEEE
45. Ciardo, G., Zhao, Y., Jin, X.: Ten years of saturation: A petri net perspective. In: Transactions on Petri Nets and Other Models of Concurrency V, pp. 51–95. Springer, ??? (2012)
46. Bollig, B., Wegener, I.: Improving the variable ordering of obdds is np-complete. *IEEE Transactions on computers* **45**(9), 993–1002 (1996)
47. Somenzi, F.: Cudd: Cu decision diagram package release 2.3. 0. University of Colorado at Boulder (1998)
48. Thierry-Mieg, Y., Poitrenaud, D., Hamez, A., Kordon, F.: Hierarchical set decision diagrams and regular models. In: Kowalewski, S., Philippou, A. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, pp. 1–15. Springer, Berlin, Heidelberg (2009)
49. Babar, J., Miner, A.: Meddly: Multi-terminal and edge-valued decision diagram library. In: 2010 Seventh International Conference on the Quantitative Evaluation of Systems, pp. 195–196 (2010). IEEE
50. Brace, K.S., Rudell, R.L., Bryant, R.E.: Efficient implementation of a bdd package. In: 27th ACM/IEEE Design Automation Conference, pp. 40–45 (1990). IEEE
51. Bonnici, V., Giugno, R., Pulvirenti, A., Shasha, D., Ferro, A.: A subgraph isomorphism algorithm and its application to biochemical data. *BMC bioinformatics* **14**(S7), 13 (2013)
52. Carletti, V., Foggia, P., Vento, M.: Performance comparison of five exact graph matching algorithms on biological databases. In: International Conference on Image Analysis and Processing, pp. 409–417 (2013). Springer
53. National cancer institute. <http://www.nci.nih.gov/>. Accessed: 2013 march 10
54. Berman, H.M., Battistuz, T., Bhat, T.N., Bluhm, W.F., Bourne, P.E., Burkhardt, K., Feng, Z., Gilliland, G.L., Lyte, L., Jain, S., et al.: The protein data bank. *Acta Crystallographica Section D: Biological Crystallography* **58**(6), 899–907 (2002)
55. Huehne, R., Suehnel, J.: The jena library of biological macromolecules-jenalib. *Nature Precedings*, 1–1 (2009)
56. Vehlow, C., Stehr, H., Winkelmann, M., Duarte, J.M., Petzold, L., Dinse, J., Lappe, M.: Cmviz: interactive contact map visualization and analysis. *Bioinformatics* **27**(11), 1573–1574 (2011)
57. Szklarczyk, D., Franceschini, A., Kuhn, M., Simonovic, M., Roth, A., Minguéz, P., Doerks, T., Stark, M., Müller, J., Bork, P., et al.: The string database in 2011: functional interaction networks of proteins, globally integrated and scored. *Nucleic acids research* **39**(suppl.1), 561–568 (2010)
58. Barabási, A.-L., Albert, R.: Emergence of scaling in random networks. *science* **286**(5439), 509–512 (1999)
59. Leskovec, J., Kleinberg, J., Faloutsos, C.: Graphs over time: densification laws, shrinking diameters and possible explanations. In: Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining, pp. 177–187 (2005)
60. Aparo, A., Bonnici, V., Micale, G., Ferro, A., Shasha, D., Pulvirenti, A., Giugno, R.: Fast subgraph matching strategies based on pattern-only heuristics. *Interdisciplinary Sciences: Computational Life Sciences* **11**(1), 21–32 (2019)
61. Merkel, D.: Docker: Lightweight linux containers for consistent development and deployment. *Linux J.* **2014**(239) (2014)

## Figures

**Figure 1** GRAPES-DD workflow with path length  $l_p = 3$ .

**Figure 2** GRAPES-DD indexing of a target graph using a MTMDD built from partial tries.

**Table 1** Indexing comparison of GRAPES and GRAPES-DD of synthetic graphs in terms of RAM requirement, Storage requirement, and Building time.

|             |          |      | RAM req. (MB) |        |       | Storage req. (MB) |        |       | Build time (s) |        |
|-------------|----------|------|---------------|--------|-------|-------------------|--------|-------|----------------|--------|
|             |          |      | GRAPES-DD     | GRAPES | ratio | GRAPES-DD         | GRAPES | ratio | GRAPES-DD      | GRAPES |
| Barabasi-A. | <i>l</i> | 0.1% | 3,649         | 7,935  | 2.2   | 305               | 3,493  | 11.5  | 470            | 109    |
|             |          | 1%   | 8,229         | 66,838 | 8.1   | 1,543             | 28,772 | 18.6  | 646            | 214    |
|             |          | 10%  | 7,876         | 81,552 | 10.4  | 10,071            | 34,368 | 3.4   | 668            | 265    |
|             | <i>α</i> | 0.5  | 2,103         | 94,654 | 45.0  | 24,915            | 40,281 | 1.6   | 516            | 330    |
|             |          | 1    | 8,351         | 58,519 | 7.0   | 16,602            | 25,145 | 1.5   | 766            | 213    |
|             |          | 1.5  | 1,447         | 3,068  | 2.1   | 1,246             | 1,219  | 1.0   | 144            | 26     |
| Forest-Fire | <i>l</i> | 0.1% | 834           | 3,929  | 4.7   | 121               | 1,689  | 13.9  | 63             | 29     |
|             |          | 1%   | 1,308         | 21,255 | 16.2  | 738               | 9,229  | 12.5  | 77             | 66     |
|             |          | 10%  | 1,167         | 24,936 | 21.4  | 5,351             | 10,650 | 2.0   | 62             | 70     |
|             | <i>p</i> | 0.1  | 147           | 1,426  | 9.7   | 2,882             | 585    | 0.2   | 10             | 7      |
|             |          | 0.3  | 188           | 2,451  | 13.1  | 3,358             | 1,024  | 0.3   | 14             | 9      |
|             |          | 0.5  | 281           | 4,966  | 17.7  | 3,922             | 2,109  | 0.5   | 26             | 16     |
|             |          | 0.7  | 487           | 11,694 | 24.0  | 4,535             | 5,020  | 1.1   | 57             | 37     |
|             |          | 0.9  | 988           | 29,565 | 29.9  | 5,386             | 12,840 | 2.4   | 139            | 88     |

**Figure 3** GRAPES/GRAPES-DD ratios of memory peak (as a RAM requirement) and index size (as a storage requirement), obtained by indexing Barabasi-Albert graphs

**Figure 4** GRAPES/GRAPES-DD ratios of memory peak (as a RAM requirement) and index size (as a storage requirement), obtained by indexing Forest-Fire graphs

**Tables**

**Table 2** Indexing comparison of GRAPES and GRAPES-DD of biochemical datasets in terms of RAM requirement, Storage requirement, and Building time

|             | RAM req. (MB) |        |       | Storage req. (MB) |        |       | Build time (s) |          |
|-------------|---------------|--------|-------|-------------------|--------|-------|----------------|----------|
|             | GRAPES-DD     | GRAPES | ratio | GRAPES-DD         | GRAPES | ratio | GRAPES-DD      | GRAPES   |
| <i>AIDS</i> | 5,304         | 1,064  | 0.20  | 164               | 39     | 0.24  | 170.12         | 16       |
| <i>PDBS</i> | 532           | 556    | 1.04  | 22                | 17     | 0.78  | 176.00         | 10.07    |
| <i>PCM</i>  | 512           | 7,057  | 13.77 | 253               | 1,392  | 5.51  | 617.24         | 754.56   |
| <i>PPI</i>  | 629           | 1,698  | 2.70  | 166               | 665    | 4.00  | 2,514.18       | 2,906.65 |

**Table 3** Indexing comparison of GRAPES and GRAPES-DD of single PPI network in terms of RAM requirement, Storage requirement, and Building time

| Species                | V      | E      | RAM req. (MB) |        |       | Storage req. (MB) |        |       | Build time (s) |        |
|------------------------|--------|--------|---------------|--------|-------|-------------------|--------|-------|----------------|--------|
|                        |        |        | GRAPES-DD     | GRAPES | ratio | GRAPES-DD         | GRAPES | ratio | GRAPES-DD      | GRAPES |
| <i>S. cerevisiae</i>   | 4,709  | 40,284 | 38.1          | 91.1   | 2.39  | 10.5              | 26.2   | 2.49  | 60.93          | 63.49  |
|                        | 5,230  | 53,699 | 56.3          | 136.3  | 2.42  | 11.7              | 42.4   | 3.63  | 604.47         | 630.92 |
|                        | 5,762  | 76,482 | 61.1          | 150.9  | 2.47  | 12.7              | 47.5   | 3.73  | 846.24         | 879.90 |
|                        | 5,936  | 89,674 | 46.5          | 121.2  | 2.61  | 12.9              | 36.8   | 2.86  | 128.52         | 135.01 |
| <i>C. elegans</i>      | 1,557  | 2,472  | 7.1           | 5.7    | 0.80  | 0.2               | 0.3    | 1.45  | 0.24           | 0.20   |
|                        | 2,421  | 3,981  | 7.8           | 7.9    | 1.02  | 0.5               | 1.0    | 2.06  | 0.52           | 0.44   |
|                        | 3,664  | 7,005  | 10.4          | 14.6   | 1.40  | 1.0               | 2.8    | 2.65  | 1.33           | 1.13   |
|                        | 6,173  | 26,184 | 25.1          | 58.5   | 2.33  | 4.3               | 16.0   | 3.75  | 34.16          | 34.19  |
| <i>D. melanogaster</i> | 1,185  | 2,008  | 8.0           | 12.1   | 1.51  | 0.6               | 1.5    | 2.30  | 0.31           | 0.26   |
|                        | 2,488  | 6,151  | 12.1          | 32.1   | 2.65  | 1.9               | 6.6    | 3.47  | 3.28           | 3.21   |
|                        | 2,729  | 7,235  | 13.3          | 36.2   | 2.73  | 2.3               | 7.9    | 3.37  | 4.30           | 4.22   |
|                        | 7,928  | 37,542 | 52.3          | 198.7  | 3.80  | 13.1              | 64.3   | 4.90  | 144.05         | 156.07 |
| <i>M. musculus</i>     | 1,810  | 2,413  | 8.0           | 13.1   | 1.64  | 0.7               | 2.2    | 2.94  | 0.42           | 0.36   |
|                        | 3,255  | 5,424  | 11.0          | 31.0   | 2.81  | 1.9               | 7.1    | 3.68  | 2.52           | 2.50   |
|                        | 3,758  | 6,853  | 13.1          | 43.6   | 3.33  | 2.6               | 11.2   | 4.30  | 4.47           | 4.61   |
|                        | 6,875  | 23,779 | 41.1          | 193.6  | 4.71  | 12.4              | 62.1   | 5.02  | 76.64          | 81.56  |
| <i>H. sapiens</i>      | 4,638  | 10,665 | 17.3          | 55.0   | 3.18  | 3.6               | 14.6   | 4.07  | 5.60           | 5.36   |
|                        | 8,728  | 31,164 | 53.5          | 215.4  | 4.02  | 13.1              | 70.8   | 5.42  | 65.14          | 68.26  |
|                        | 9,826  | 48,835 | 87.5          | 351.4  | 4.02  | 21.5              | 120.1  | 5.59  | 213.95         | 230.16 |
|                        | 10,186 | 51,484 | 89.2          | 391.8  | 4.39  | 22.0              | 134.1  | 6.10  | 191.63         | 209.63 |

Figure 1


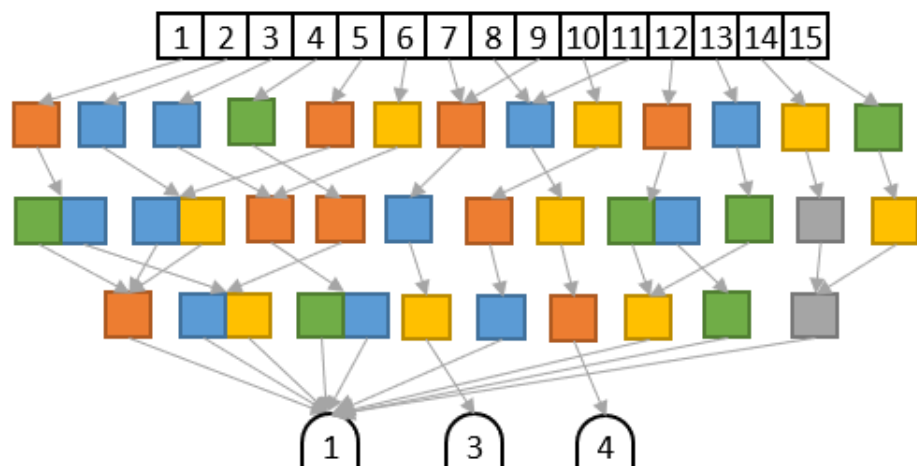


**(a) Target graphs collection****(b) Index MTMDD****(c) Query graph****(e) Pruned MTMDD****Filtering phase****(d) Query MTMDD****(f) Candidate vertices**Indexing  
building  
phaseFiltering  
phase

Figure 2

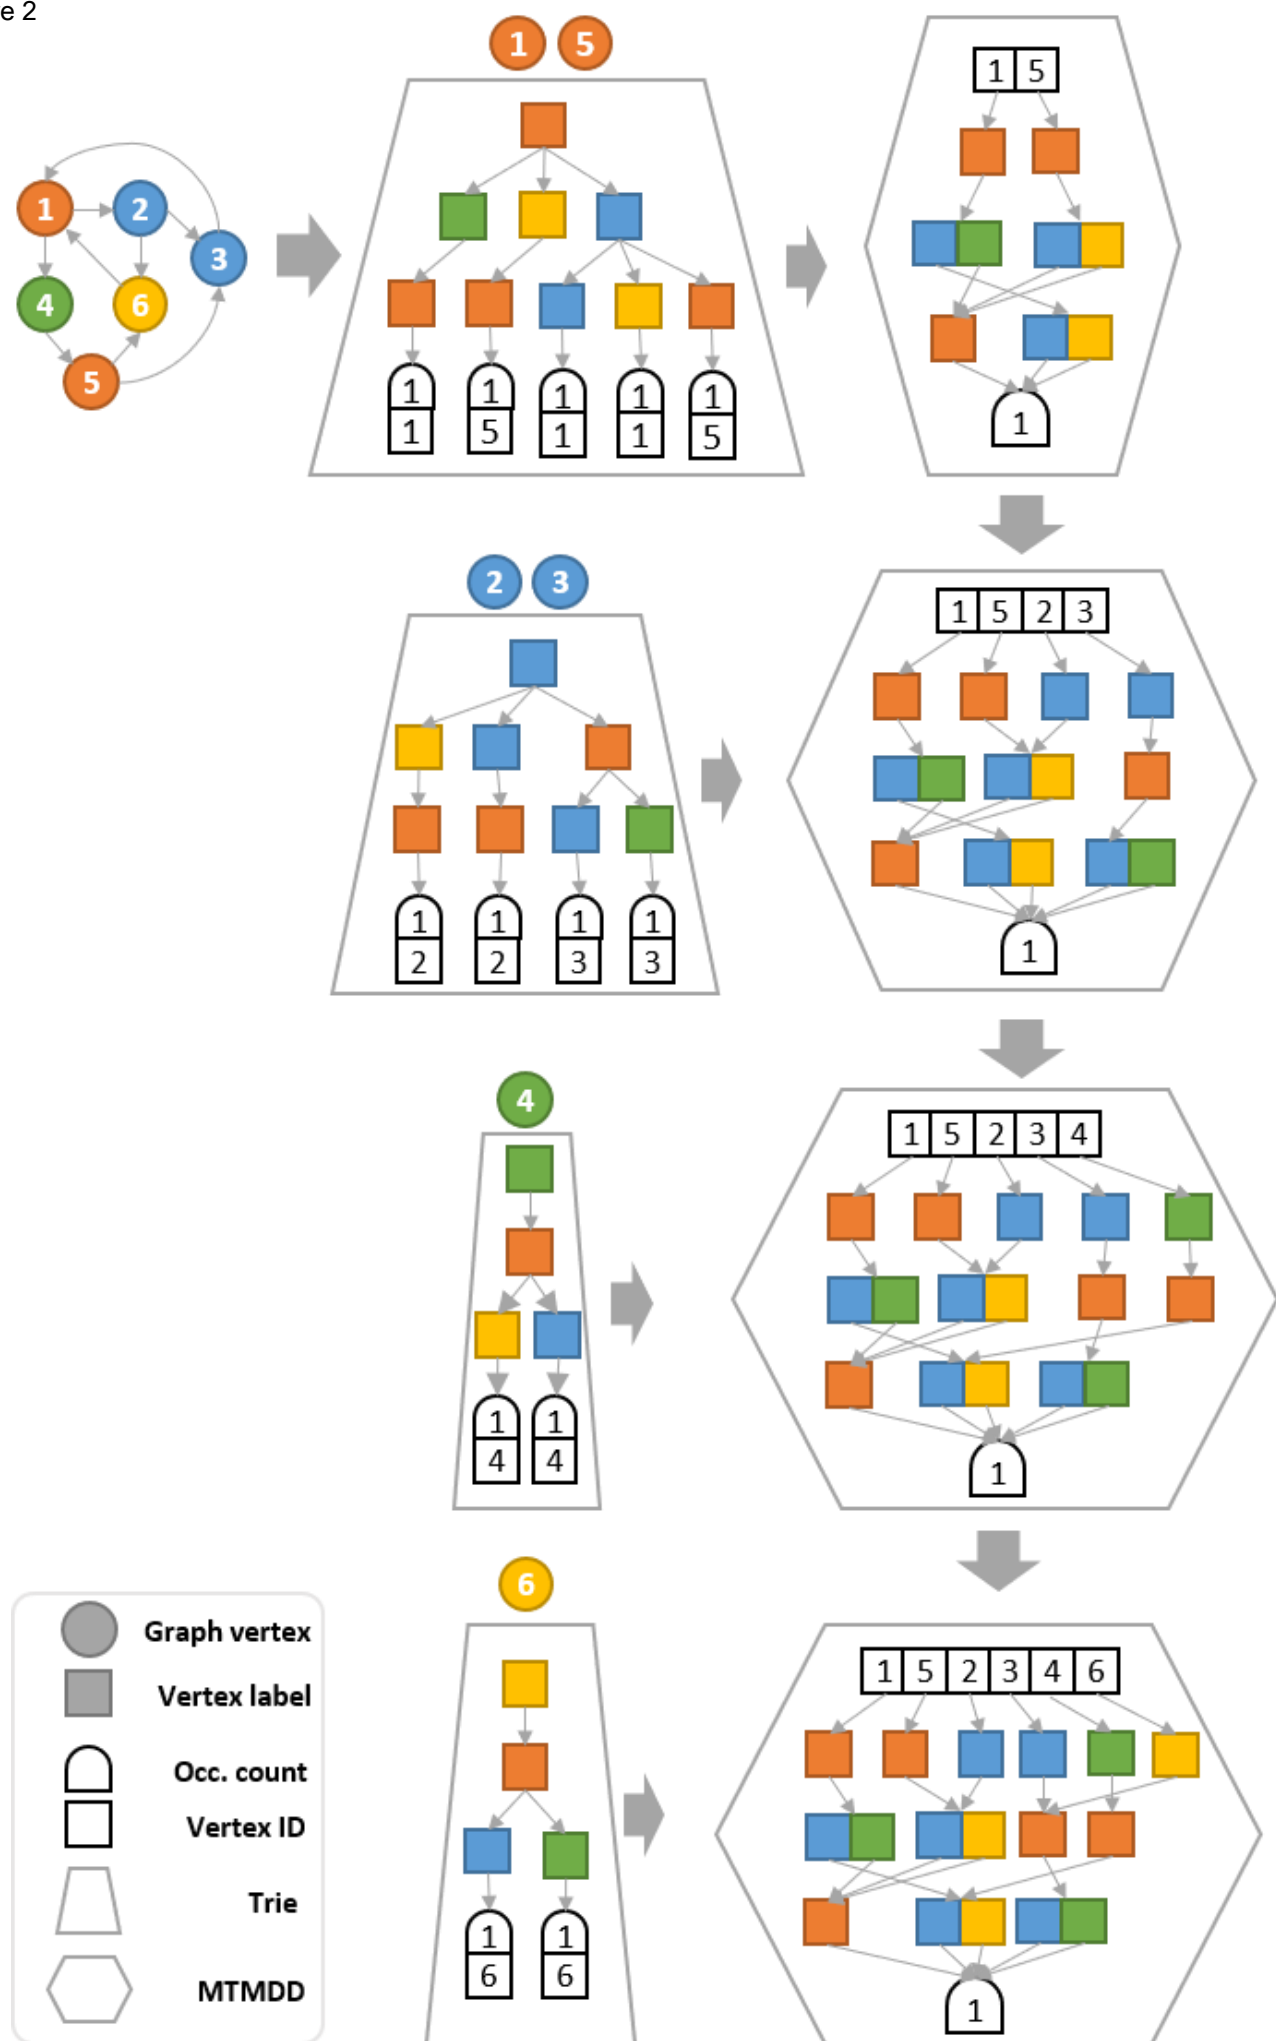





Figure 3

RAM requirement ratio



Storage requirement ratio


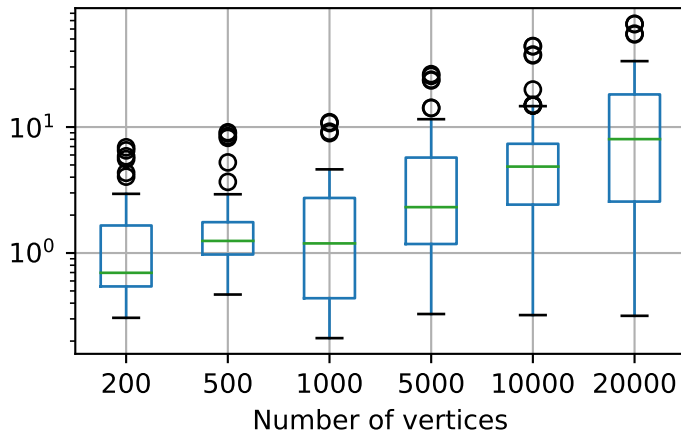





Figure 4

RAM requirement ratio



Storage requirement ratio



# Figures

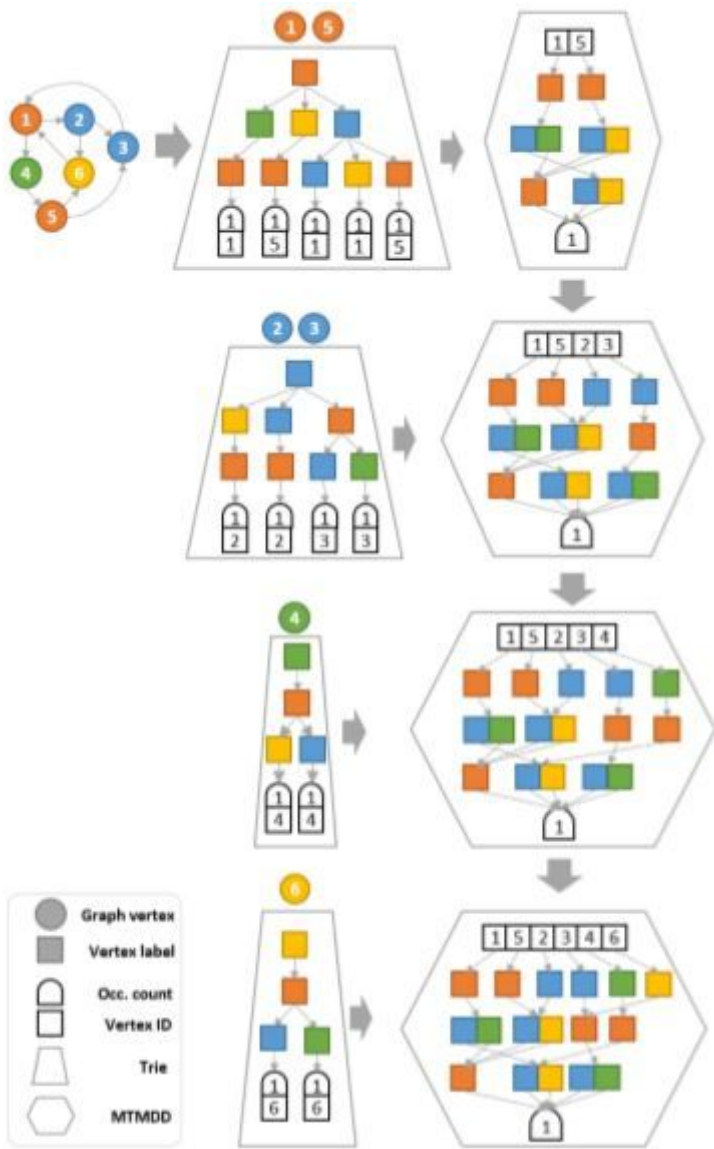


Figure 1

GRAPES-DD workflow with path length  $l_p = 3$ .



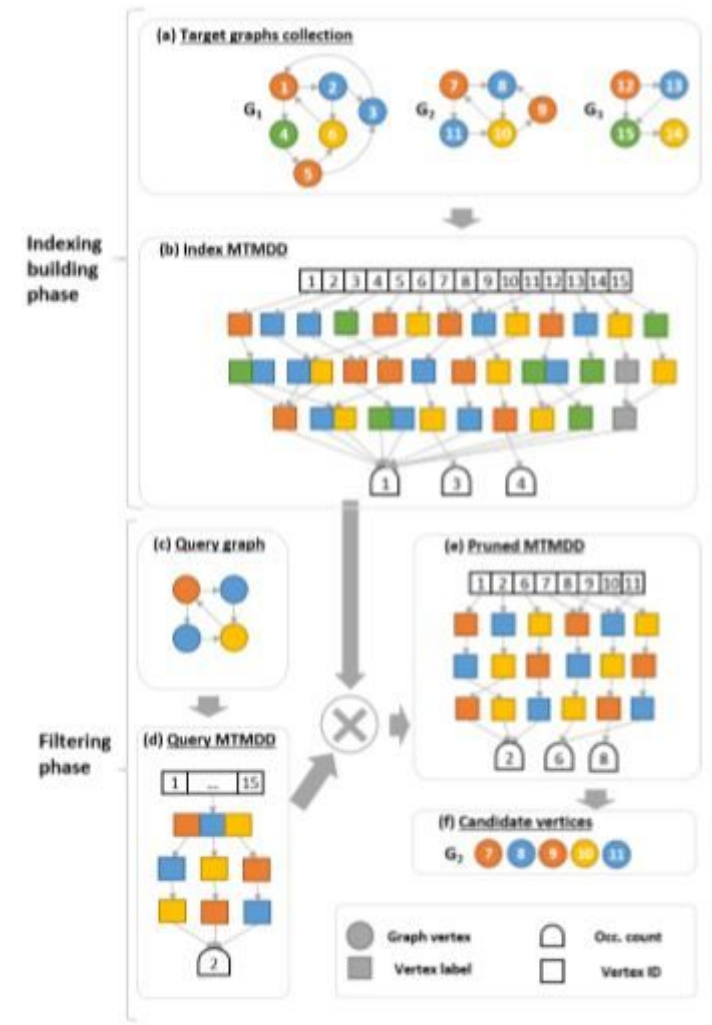
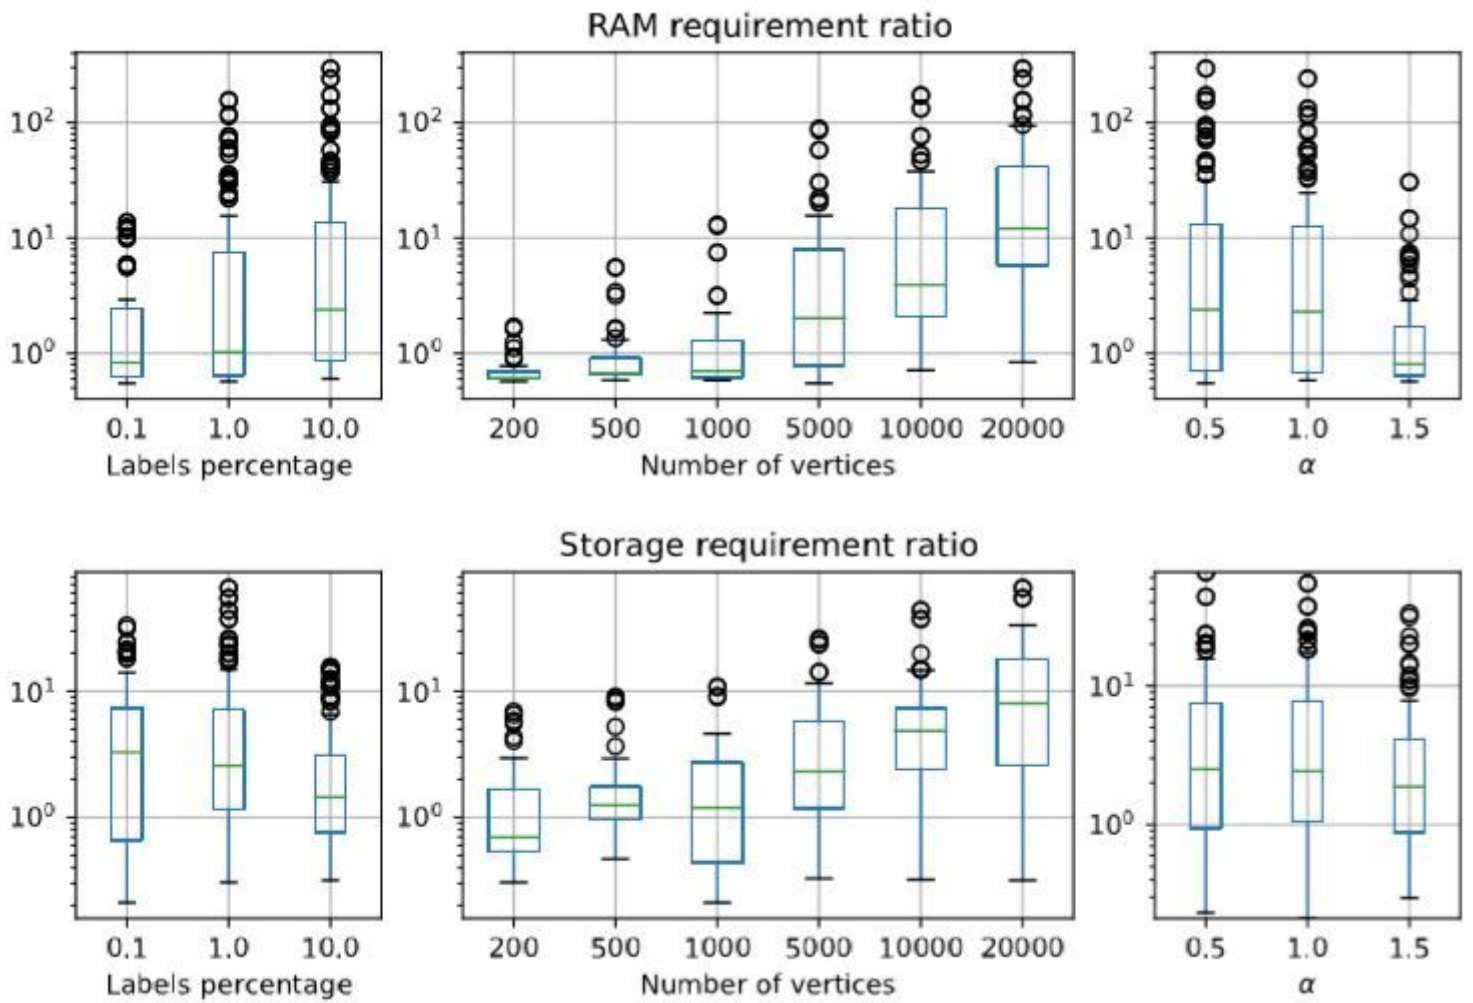


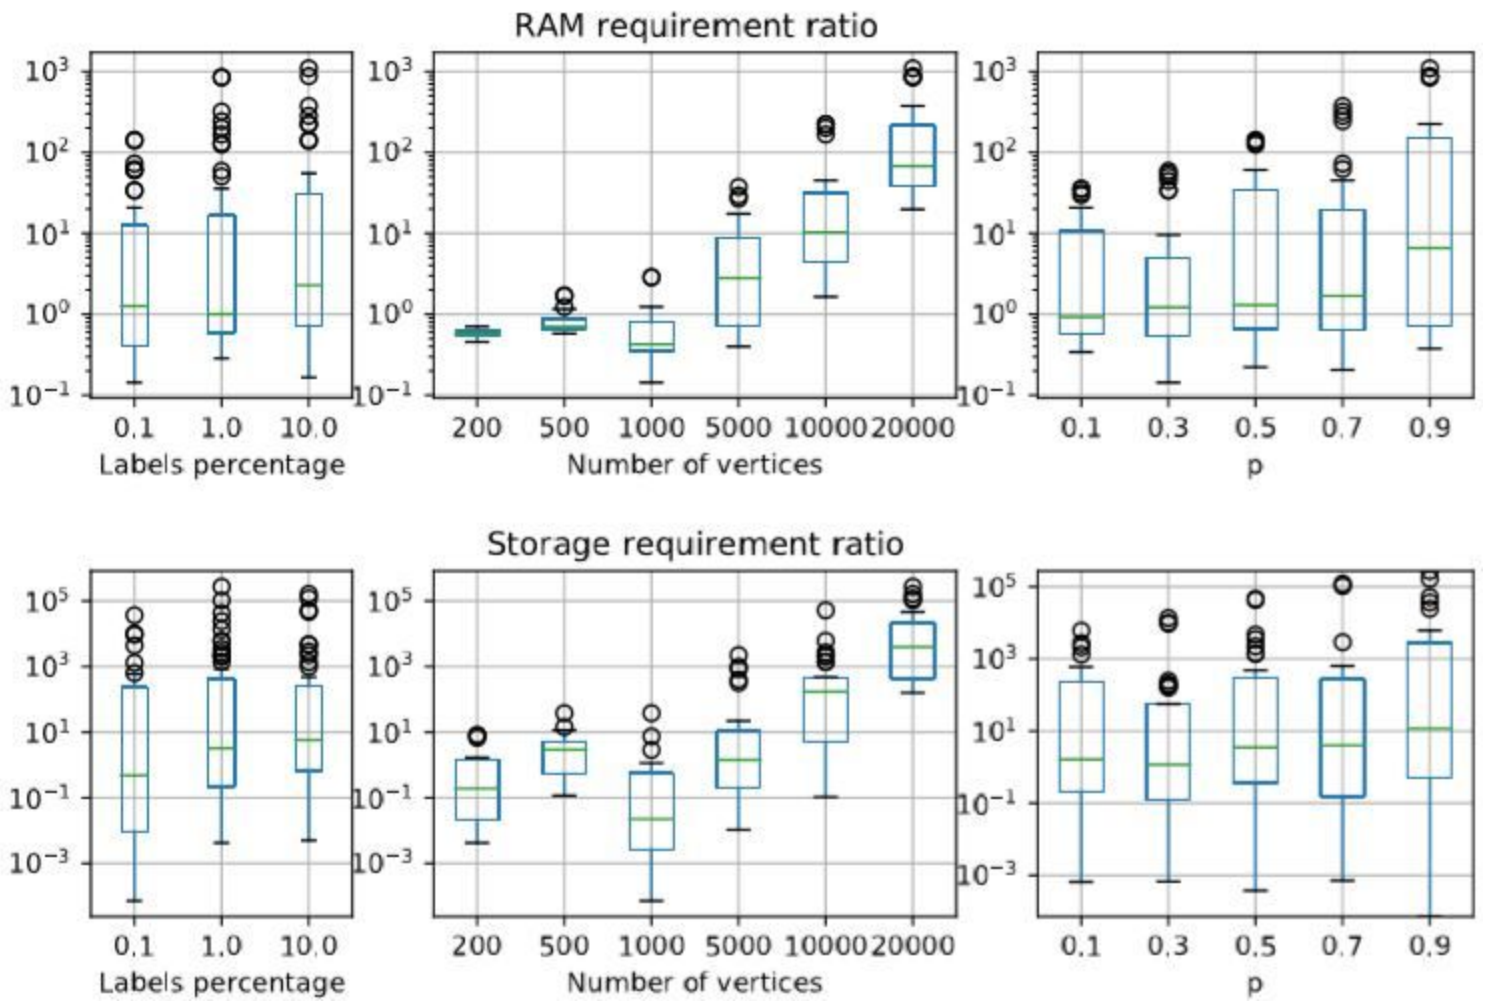
Figure 2

GRAPES-DD indexing of a target graph using a MTMDD built from partial tries.



**Figure 3**

GRAPES/GRAPES-DD ratios of memory peak (as a RAM requirement) and index size (as a storage requirement), obtained by indexing Barabasi-Albert graphs



**Figure 4**

GRAPES/GRAPES-DD ratios of memory peak (as a RAM requirement) and index size (as a storage requirement), obtained by indexing Forest-Fire graphs