Querying Distributed Digital Libraries of Mathematics

Ferruccio Guidi* and Claudio Sacerdoti Coen*

Department of Computer Science
Mura Anteo Zamboni 7, 40127 Bologna, ITALY.
{fguidi,sacerdot}@cs.unibo.it

Abstract. Several of the most effective techniques to query libraries of structured mathematics are based on pattern-matching. Among them, there are matching-based queries, unification-based queries and several kinds of queries up to isomorphisms (as associative-commutative rewriting). All of these techniques do not scale when the libraries become large or when they are distributed. In this paper we present a filtering technique that can be applied to locate a set of candidates that are likely to match the given pattern. Thus our technique can be used as a first phase that is followed by the actual matching over the set of candidates. Tests performed over the library of the Coq system (http://coq.inria.fr) (about 40.000 theorems) show that the candidates can be located rather quickly and that the number of false matches is surprisingly low.

1 Introduction

As Jorge Luis Borges pointed out in his beautiful novel "The Library of Babel" [8], a huge library is completely useless unless we have effective methods to retrieve what we are looking for. Unfortunately, the largest part of the mathematical documents available over the Web are Postscripts or PDF documents. Thus, the only reasonable way to look for a given mathematical definition or result is to rely on textual matching or to perform queries over the provided metadata (when available). In particular, the organizations of librarians have developed several useful classification schemes to index mathematical textbooks and papers [4, 1, 2]. This approach can be effective enough when we are looking for a general work on a given mathematical topic; but it can not be used to retrieve a single theorem or definition over the Web, especially when the theorem is not fundamental enough to be given a well-known name.

The forthcoming development of Web services for automated deduction and computation put high pressure for the development of alternative effective methods to retrieve a single result. Luckily, at the same time more and more mathematical knowledge is put on the web in a structured form, thanks to the introduction of standard markup languages [6, 3, 5] for the encoding of mathematical formulae and expressions. Once the statements in a library are marked up, say,

* Supported by MOWGLI: European FET Project IST-2001-33562.

in OpenMath, we can use standard pattern-matching or unification techniques to retrieve a theorem given a part of its statement. For example, the pattern $\forall x : ?_1.\ \forall y : ?_1.\ (?_2\ x\ y) = (?_2\ y\ x)$ will retrieve any theorem stating the commutative property of a binary operation. The question marks are *metavariables*, which are non-linear placeholders for any (well-typed, if enforceable) subexpression closed in the context.

To resolve a query based on pattern matching, the trivial approach consists in iterating the matching operation over the whole library. This solution does not scale well and it can not be applied in those cases where the library itself is physically distributed and no code can be run on the severs[1].

To solve this problem, we propose an approach based on four distinct phases:

- *Data-mining*: using a spider, we extract a small set of automatically computed metadata from each theorem or definition in the distributed library. These metadata are devised to capture some of the invariants of the chosen matching operation.
- *Pattern compilation:* when the user submits a pattern, we extract a set of constraints over these metadata from the pattern and we generate a low-level query[2] starting from these constraints.
- *Filtering:* we execute the generated query over the database built by the spider, obtaining our set of candidates.
- *Matching:* we iterate the matching operation only on the set of candidates.

If the filtering operation is both quick and correct (in the sense that no good candidates are dropped), we are able to achieve both accuracy and performance.

Since the metadata must capture the invariants of matching, we need to identify several kinds of metadata. In sections 2–4 we will progressively consider several examples of matching criteria and we will identify the related metadata. We will also consider the issue of correctness and a tradeoff between correctness and performance. In section 5 we will present a generic set of metadata and constraints that can subsume all the previous examples. Finally, section 6 and section 7 are devoted to comparisons with other works and future developments.

2 Use Case 1: Finding the Applicable Theorems

The first use case that we consider is knowing which theorems can be applied to prove a given goal. A goal is a sequent made of several hypotheses and a conclusion. A theorem can be applied if its conclusion, where every bound variable has been replaced with a metavariable, matches the conclusion of the sequent.

[1] Indeed, if the library is distributed, we have only two possibilities: either we can perform the matching where the theorem is located or we need to download it. In the second case the download times make the operation unfeasible already for very small libraries. In the first case we need to install software were the library is published, which is something we would rather avoid.

[2] In our prototype, we generated the query in the MathQL Level 1 language [12]. Generating a query in other languages is surely possible.

We restrict ourselves to the case of *first order matching*, i.e. only one expression (the conclusion of the theorem) has metavariables in it; and the expression is rigid. For example:

- The theorem $\forall n.\ \forall m.\ \forall a.\ a \geq 0 \Rightarrow n \leq m \Rightarrow an \leq am$ can be applied to prove $2x \leq 2(x+1)$ since the pattern $?a?n \leq ?a?m$ matches $2x \leq 2(x+1)$.
- The theorem $\forall n.\ \forall m.\ \forall a. n \leq m \Rightarrow an \leq am$ can not be applied to prove $x \leq 2x$ since the pattern $?a?n \leq ?a?m$ does not match $x \leq 2x$.
- The induction principle over natural numbers $\forall P.\ (P\ 0) \Rightarrow (\forall n.\ (P\ n) \Rightarrow (P\ (n+1))) \Rightarrow \forall n.\ (Pn)$ does not match $(is_prime\ 5)$ since the pattern $(?P\ ?n)$ is not rigid.

Now we identify a set of invariants of the matching procedure that we want to capture as metadata:

- The constant in *head position* in the conclusion of the theorem. Since the pattern is rigid, any matching term must have the same head constant of the pattern.
- The list of constants in the conclusion. Since we are doing matching and not unification, any constant in the pattern must also occur in the matching term (since instantiating a metavariable can not remove an occurring constant).

Thus we identify the following set of metadata: for each theorem in the library we remember the name of the constant in head position in the conclusion and the names of all the constants occurring in the rest of the conclusion. Now, to find the theorems that are candidates for the query "give me all the theorems that can be applied to this goal", we simply compute these metadata over the conclusion of the goal and we look for every theorem in the library such that:

- the "constant in head position" metadata is the same as the computed one
- the set of "constant in conclusion" metadata is a subset of the computed one

The following theorem is a trivial consequence of the fact that our metadata reflects the matching invariants:

Theorem 1. *The filtering operation is correct, i.e. no pattern that matches the goal conclusion is filtered out.*

Let's see how it works in practice: given the goal $2x \leq 2(x+1)$

- $?a?n \leq ?a?m$ is a candidate that will match
- $?n \leq ?n^2$ is not a candidate, since the power operator does not occur in the goal
- $?n?a \leq ?m?a$ is a candidate that will not match (a *false match*)

We could also capture a much stronger invariant, that subsumes both our previous invariants: *the rigid skeleton of the term is preserved by metavariable substitutions.* Nevertheless, the stronger invariant has two drawbacks: the first one is that it produces complex metadata (contexts, i.e. terms with holes) that can not be handled efficiently within the relational database model or the RDF

model (which is often reduced to the relational model); the second drawback is that the invariant is probably too strong. In particular, it would reject the third candidate of the previous example. Even if that candidate turns out to be a false match, it is pretty obvious (at least to a human user) that the statement found is useful to proceed in the proof. Very often, indeed, the user is also interested in theorems that can "almost" be applied, since they can reveal an error in the proof (when a required assumption is missing) or they can suggest a different way to prove the goal.

Of course, the danger of not adopting the stronger invariant is that of finding too many useless false matches. To evaluate our invariants, we tried several queries over the library of the Coq proof-assistant (about 40.000 theorems). The queries were not artificial: one of the authors used the query engine in order to retrieve some useful lemmas while developing a new medium-sized Coq contribution. Even for queries involving only frequently used notions (e.g. algebraic operations), the accuracy of the filtering phase was always very high: the number of uninteresting false matches was very low and in more than one occasion the search engine found false matches that we were actually able to exploit. Thus, the productivity level of the author was definitely raised by the use of the search-engine.

Let's face now the issue of performance, considering for instance the goal $2x \leq 2(x+1)$ which generates the constraints:

- $\leq$ `in head position in the conclusion`
- `only 2 and + in other positions in the conclusion`

The generated query (in a pseudo-language inspired by MathQL level 1 [12]) is:

```
select every t in the library such that
 t.head = '<=' and t.in_conclusion subset of {'2','+'}
```

Since the second test can not be performed in a relational DB, we are forced to extract from the DB all the theorems `t` that satisfy the first constraint together with their `in_conclusion` field and then iterate the second test. Thus, the computational cost is linear over the size of the result set of the first test. If the head position of the goal is a frequently used constant, the result set is large. For example, there is an equality in head position in the 11% of the theorems of the library of Coq, i.e. in 3660 theorems. To reduce the computational cost of the query, we propose a user-controlled tradeoff with the accuracy of the query: the user provides an additional set of constraints consisting in a list of constants that *must* appear in the conclusion of the matched theorems. To avoid confusion, we will call these constraints *must constraints* and we rename the previous set of constraints *only constraints*. In order to retrieve a non-empty subset of results, the set of constants occurring in the *must constraints* must be a subset of the constants occurring in the *only constraints*.

To better understand the *must constraints*, let us consider the goal $2x \leq 2(x+1)$. The generated *only constraints* are $\leq$ in head position and 1,2,$*$ and $+$ in other positions in the conclusion. Let's suppose that the user choice of *must constraints* is $\leq$ in head position and $*$ in other positions. Then

- $?a?n \leq ?a?m$ is a candidate that will match (a true match)
- $?n?a \leq ?m?a$ is a candidate that will not match (an interesting false match)
- $?n \leq ?n + 1$ is not a candidate since it satisfies the *only constraints* but not the *must constraints* ('$*$' does not occur). Anyway, it would have been an uninteresting false match.

Note that the *must constraints* are meaningful to the user: in the previous example, the *must constraints* are stating that the user is interested in a property of multiplication. Thus, it is not difficult to provide an user-friendly interface to select these constraints. It is also easy to provide some heuristics to automatically chose the *must constraints*[3].

A non-empty set of *must constraints* reduces the accuracy of the query, since some theorems that could have been applied are not found any longer. For example, the previous query is unable to find the transitivity principle for the $\leq$ relation, since it states nothing about multiplication[4]. The performance gain achieved can be remarkable. To understand why, let us see the query generated for the last example:

```
let S =
 select every t in the library such that
  t.head = '<=' and '*' occurs in t.in_conclusion
in
 select every t in S such that
  t.in_conclusion subset of {'1','2','+','*'}
```

The first part of the query (the first `select`) can be performed by a single call to the database. Then, as before, we need to iterate the second test over the set returned by the first select. The more *must constraints* we have, the smallest the set `S` will be and the smallest the time required to perform the query will be.

In the next two sections we will try to apply the same approach to another couple of examples, choosing the matching operation we are interested in and deriving the *must* and *only constraints* from it.

3 Use Case 2: Finding the Elimination Principles

Let us consider queries for the retrieval of all the elimination principles "à la McBride" [13] over a given datatype. Briefly, these are elimination principles that allow to prove a generic property P over an element n of a datatype T by proving that P holds in several cases, some of them under additional inductive hypotheses. For example, there are several elimination principles over lists of natural numbers. Among them, we find the structural induction principle $\forall P.\ (P\ \mathit{empty}) \Rightarrow (\forall l.\ \forall n.\ (P\ l) \Rightarrow (P\ n :: l)) \Rightarrow \forall l.\ (P\ l)$ and the

[3] As an example, we provide a predefined heuristic that chooses only the constants occurring in the first n levels of the syntactic tree of the goal thesis.

[4] In our experience, very often loosing these "very general" theorems can be considered a feature.

induction principle over ordered lists $\forall P.\ (P\ empty) \Rightarrow (\forall l.\ \forall n.\ (ordered\ l) \Rightarrow n \leq (head\ l) \Rightarrow (P\ n :: l)) \Rightarrow \forall l.\ (ordered\ l) \Rightarrow (P\ l)$.

Note that, since the pattern $(?P\ ?x)$ is not rigid, every elimination principle can always be used to progress in any proof, independently of the goal. Since there exist literally thousands of elimination principles, it is not a good idea to include them as results of the query of the first use case we considered[5]. Nevertheless, we are interested in giving the user the possibility to retrieve all elimination principles over a given datatype. In particular, given a datatype T, we want to retrieve all the theorems whose statement matches[6] the schema:

$$\forall \overline{x} : \overline{S}.\ \forall P : T \to Prop.\ \forall \overline{y} : \overline{S'}.\ \forall x : T. \forall \overline{z} : \overline{S''}.\ (P\ x)$$

We identify the following set of constraints:

- The sort `Prop` (i.e. the type of every proposition) must occur in the head position of an hypothesis which is a product[7] of length 1.
- The datatype must occur in the head position of an hypothesis which is a product of length 0 and and also in another hypothesis (not in main position).
- The head of the conclusion must be an occurrence of a bound variable.

Every constraint captures a matching invariant, since context metavariables substitution can not change the shape of the conclusion nor the shape of an hypothesis (which are all rigid parts). As in the previous use-case, there exist stricter invariants: the set of constraints identified here has been fine-tuned by hand to be able to retrieve also a few useful false matches. Even in this case, though, the final set of invariants was identified quite soon and without any major effort.

Note that, unlike the previous use case, we derive the matching pattern from the user input and not from the statements of the theorems in the library. In section 2 we had just one term to match against several different patterns; here we have just one pattern to match against the theorems in the library. This duality is reflected in the way the constraints are used to generate the query: in the previous case the identified constraints were *only constraints* (only the constants occurring in the goal can occur in the statement of the theorem); in this case the constraints are *must constraints* (all the constants in the pattern must occur in the statement of the theorem)[8].

[5] That is the reason why we constrained ourselves to first order unification in section 2.

[6] We do not specify formally here the kind of match we are interested in. Anyway, note that the "vector parts" of the schema play the role of *context metavariables*, in the sense that they are contexts of formulae that must be instantiated with a context made of repeated quantifications or implications, concluded by a single contextual hole.

[7] The terminology comes from type-theory and is derived by the Curry-Howard isomorphism, where universal quantifications are seen as (dependent) products. A product of length n is a term of the form $\forall x_1.\ \ldots\ \forall x_n.\ t$ where t does not start with an universal quantification. An implication is just a special case of universal quantification: it is a non-dependent quantification where the bound variable does not occur in the sub-expression.

[8] Both sentences are instances of the following invariant that is true for any kind of matching: "all the constants in the pattern must occur in the matched term".

The query generated from the constraints is the following:

```
select every t in the library such that
 Prop occurs in head position at depth 1 in an hypothesis of t and
 T occurs in head position at depth 0 in an hypothesis of t and
 T occurs not in head position in an hypothesis of t and
 a bound variable occurs in head position in the conclusion of t
```

Unlike the case of section 2, the generated query is already very efficient: this time we do not need any trade-off to improve the performance.

4 Use Case 3: Finding the Proofs of a Statement

The last use case we present is locating a proof (or a definition) whose statement (i.e. its type) is known. The query seems quite uninteresting, but the inability to perform it may have very bad implications in terms of waste of efforts. Indeed, every time we need to prove a lemma, we would like to issue a query that checks if somebody else already proved it; the Coq system provides no way to perform this operation and, as a result, it is not at all unusual to find simple arithmetical properties proved five, eight or even more times by different persons working in remotely located teams[9].

The kind of matching we are interested in is extremely simple: no metavariables are present in both the statement of the theorem and the statements in the library. Close enough false matches, though, play an important role: for example, when looking for a proof of $\forall n.\ n = n * 1$, any of the following statements would be satisfactory: $\forall n.\ n = 1 * n$; $\forall n.\ n * 1 = n$; $\forall n. 1 * n = n$

We identify the following set of metadata, which constitute a close approximation of the first two levels of the abstract syntax tree of the pattern:

- For each occurrence of a constant, we record its *position*, as an element of the following set: {`MainConclusion, InConclusion, MainHypothesis, InHypothesis`}. The prefix `Main` means that the constant occurs in head position in the `Conclusion/Hypothesis` of the theorem.
- For each occurrence of a bound variable, we record its *position* in the set {`MainConclusion, MainHypothesis`}. Since bound variables occur almost in every statement, they are not very effective for searching. Thus we are interested only in second order or higher order variables, that are much less frequent. We syntactically recognize a subset of these variables by the fact that they occur in head position: these are the only occurrences we record in the metadata.

[9] E.g. the proof of the fact that 1 is the neutral element of multiplication was given five times. Of course everybody suspects that a proof should already exist. Nevertheless, proving it again is faster than guessing who gave the proof and in what library. As a consequence, the theorem is proved again and the new proof is stored in the library under development by the user, which is likely to be unrelated to arithmetical properties. Thus, without an automatic searching procedure, it is virtually impossible to know that a proof of the theorem was put in that library.

- For each occurrence of a sort *Prop* or *Set*[10], we record its *position* in the set {`MainConclusion, MainHypothesis`}
- For each occurrence in {`MainConclusion, MainHypothesis`} we record its *depth*, i.e. the number of universal quantifiers in the statement or hypothesis.

On the previous metadata we automatically impose both *must constraints* and *only constraints*[11] to achieve both performance and accuracy at the same time[12]. As in the previous cases, our experiments over the Coq library show that the chosen set of constraints is a good compromise between the strictness of the query (in terms of false matches) and the need to get results close to the expected ones (interesting false matches).

Expert users of the system can also edit the set of generated constraints, relaxing both the *must constraints* or the *only constraints*. As a result, they are able to answer similar queries such as: *Retrieve every binary relation* or *Retrieve any n-ary relation over natural numbers.*

Instead of describing here these and other use cases with the introduction of new ad-hoc constraints, we will present in the next section a generic set of constraints that have been identified so far in the projects HELM (Hypertextual Electronic Library of Mathematics, http://helm.cs.unibo.it) and MoWGLI (Mathematics on the Web: Get it by Logic and Interfaces, European Project IST-33562, http://mowgli.cs.unibo.it). The set of constraints is general enough to subsume the constraints identified in the previous use-cases. Moreover it is relatively stable, since no new constraints have been added in the last few months. The current work is focused on the automatic generation of larger and larger classes of queries based on these constraints. For example, we are now able to write by hand queries to retrieve every theorem in the library of Coq stating the associative or commutative property of any binary operator on any datatype. The number of false matches for these kind of queries is surprisingly low[13]. It is surely possible to generate these queries automatically starting from the formal definition of the wanted property.

5 A Generic Class of Metadata and Constraints

There is an obvious common pattern for the metadata identified in the use cases above: they provide an approximate description of the abstract syntax tree of

[10] *Prop* is the type of every proposition and *Set* is the type of any data-type. They are used to capture second order and higher-order quantifications over properties and types. They can also be used to capture axiom schemas in first-order theories.

[11] With *only constraints* we mean that every occurrence of a constant, bound variable or sort in the matched term must occur in the *only constraints*. In section 2 this test was restricted to the occurrences in the conclusion.

[12] There is no tradeoff here: the *must constraints* can not filter out any good candidate since there are no metavariables in the pattern.

[13] For example, looking for every associative property in the Coq library gives 82 good matches and just 6 false matches.

the statements[14] in the library in terms of the constants and bound variables occurring in them, together with the occurrence locations. This observation leads us to describe our metadata model using *objects* having a list of *references*: the objects represent the statements in the library at the metadata level, i.e. they are the entities the metadata are about, while the references are the metadata describing the occurrences of constants and bound variables in the statements. As we said, an object may have three kinds of references. Namely:

- `refObj` describes the occurrence of a primitive or defined constant in terms of an `occurrence`, a `position` and a `depth` (see below).
- `refSort` describes the occurrence of a sort in terms of a `sort`, a `position` and a `depth` (see below). A sort is the type of propositions or data-types in higher-order logics (i.e. the second order quantification "*for every property* P" can be thought as the typed universal quantification $\forall P : Prop$). Sorts can also be used to capture axiom schemas in first-order theories.
- `refRel` describes the occurrence of a bound variable in terms of a `position` and a `depth` (see below).

The value of a reference is a record made of several fields, whose semantics is:

- `occurrence` specifies the referred defined constant. Its value is an object.
- `sort` specifies the referred sort. Its value belongs to a predefined collection of entities representing the possible sorts at the metadata level.
- `position` specifies the position of the described occurrence in the statement. Its value belongs to a predefined collection of entities representing the possible positions at the metadata level. According to the previous use cases, we identified the following positions:
  - `MainHypothesis`: in head position of a statement premise.
  - `InHypothesis`: in another position of a statement premise.
  - `MainConclusion`: in head position of the statement conclusion.
  - `InConclusion`: in another position of the statement conclusion.
  - `InBody`: not in the statement (for instance in its proof).
- `depth` specifies a depth index associated to the position of an occurrence in the statement. Its value is a natural number. As we saw in the use cases, the `depth` of a reference is defined only when its `position` is set to `MainHypothesis` or to `MainConclusion` and it represents the number of premises of that hypothesis or conclusion.

The above description shows that every reference has some fields for locating the corresponding occurrence in the statement (`position` and `depth`) and may have one field specifying the referred entity (`occurrence` and `sort`)[15].

The queries we generate on these metadata are once again inspired by the use cases we saw in the previous sections. Our general approach concerning query

[14] With statements here we mean the statements of axioms and theorems, and also the types of definitions.

[15] `refRel` does not have this field because, up to now, we did not meet any use case in which we need to discriminate between the occurrences of different bound variables.

generation is that complex queries should be obtained composing basic queries generated from the imposed atomic constraints on the wanted objects.

We identify the following two classes of basic queries:

A. **The wanted objects must have a reference to a given object R (or to a given primitive constant S or to a bound variable) in a given position P with a given depth index D.**
Here we search for the objects having the assigned values in corresponding fields of at least one of their references. Using the pseudo-language exploited in the previous sections, a query of this class looks like:

```
select every t in the library such that
 the set {
  select every r in t.refObj such that
   r.occurrence = R and r.position = P and r.depth = D
 } is not empty
```

This class of queries includes what we called the basic *must constraints*.

B. **The wanted objects may have a reference to an object (or to a primitive constant or to a bound variable) only if its position is not included in a given set U of positions, or if it concerns a given object R (or a primitive constant S, or a bound variable) in a given position P with a given depth index D.**
Here we query the objects such that there are no references whose fields do not have the assigned values. A query of this class may look like:

```
select every t in the library such that
 the set {
  select every r in t.refObj such that
   not (r.occurrence = R and r.position = P and r.depth = D)
   and r.position belongs to U
 } is empty
```

This formulation exploiting the logical double negation is more convenient, since it can be expressed more easily in SQL. This class of queries includes what we called the basic *only constraints*.

Note that the parameters R, S, P, D, U are always optional.

Coming to the issue of how the above queries should be composed, let us suppose that we are interested in finding the objects satisfying a given set (call it K) of "basic constraints". If the set K contains p constraints in the class A, the wanted objects are meant to satisfy all of them. Since a query of the class A (a *must constraint*) is of the general form

```
M(i) =
 select every t in the library such that
  the set m(i) is not empty
```

the composition of M(1) ... M(p) is:

```
M = M(1) + ... + M(p) =
 select every t in the library such that
  the set m(1) is not empty and ... and the set m(p) is not empty
```

Dually if the set K contains some constraints in the class B, the wanted objects are meant to satisfy at least one of them. Since a query of the classes B (an *only constraint*) is of the general form:

```
O(j) =
 select every t in the library such that
  the set o(j) is empty
```

the composition of O(1) ... O(q), provided that they concern the same kind of reference, say `refObj`, is:

```
O = O(1) + ... + O(q) =
 select every t in the library such that
  the set o(1) intersected ... intersected o(q) is empty
```

and the same holds for the queries O'(1) ... O'(q') concerning `refSort` and for the queries O"(1) ... O"(q") concerning `refRel`. Now the final query is obtained composing M, O, O' and O" conjunctively:

```
M + O + O' + O'' =
 select every t in the library such that
  the set m(1) is not empty
    and ... and the set m(p) is not empty  and
  the set o(1) intersect ... intersect o(q) is empty  and
  the set o'(1) intersect ... intersect o'(q') is empty  and
  the set o''(1) intersect ... intersect o''(q'') is empty
```

6 Comparison with Other Works

The MBase [14] system resolves a query by matching, by iterating the matching operation over each thesis in its data-base of (which now holds about 6,600 theorems). Thanks to the clever implementation of the unification procedure, the time spent to perform the query is still very low, in the order of $0.8 - 0.9$ seconds. Nevertheless, a library of just 6,600 theorems is to be considered still very small; in fact the libraries of the Coq proof-assistant and of the Mizar proof-assistant already hold 40,000 theorems each and they cover just a small subset of the overall mathematical knowledge. This approach will definitely not scale to much larger libraries. To solve the problem, the MBase team is now working on a solution that is extremely close to the one proposed in this paper: the theorems will be indexed in order to avoid unifications that will surely fail. To the authors knowledge, no details on the indexing technology has been published so far.

The Mizar team is now developing a user-friendly query language for the Mizar library [7]. The query language is not based on unification, but it nevertheless quite useful in retrieving interesting results from the library. To retrieve

a theorem, the user combines a set of filters over the whole library, until the wanted result is found. The filters are not applied to the library itself, but to a set of indexes (metadata) automatically extracted from the library. Thus the Mizar approach is extremely close to our approach: the main difference is in the set of metadata that, in the case of Mizar, are often very peculiar to the system. For instance, in Mizar it is possible that several symbols share the same notation and there is a query to retrieve the symbols represented by a given notation. Other metadata, instead, are less refined versions of the ones we just presented.

The problem of speeding up matches and unification has already been extensively studied in the literature. In particular, several term indexing techniques have been proposed to speed up matching and unification by storing the whole library in an ad-hoc data-structure particularly optimized for matching. Among the most successful structures, there are discrimination trees, substitution trees and, more recently, coded context trees [11]. The common idea behind these data-structures is to share the term structure as much as possible: two shared substructures are syntactically equal and thus they match one another. Matching and substitutions implemented over these structures can be several orders of magnitude more efficient that iterating unification over the whole library.

Fast term indexing techniques are not incompatible with our approach. In particular, they can be seen as the limit case where the collected metadata are so abundant to describe all the statements in the library and the information about their sharing. Of course, the choice of an efficient data-structure to collect the metadata is a necessary requirement to achieve the gain of fast term indexing. Even if fast term indexing techniques are surely much more performant of our implementation, our general approach has some merits. First of all, since our metadata are RDF triples and not term contexts, they can be easily stored in a common relational data base, while term indexing techniques require specialized libraries implementing complex data-structures. As a consequence, we can easily benefit from the advancements of the data base community, comprising transparent distribution techniques. Moreover, writing spiders to collect new metadata and implementing new kinds of queries is a minor effort in our approach and leaves complete freedom to the implementor. On the contrary, even if the great majority of queries can be reduced to unification, unification is the only operation that is natively provided by term indexing data structures. Moreover, the experience of the Mizar group, also supported by our implementative evidence, shows that unification is often overkilling: the results obtained by querying a very small set of computed metadata are often precise enough to skip the unification step. Finally, the interesting false matches provided by our implementation do provide useful information "for free". To retrieve the same information from term indexing data structures, additional effort must be spent.

7 Conclusions and Future Work

In this paper we presented a general methodology for speeding up queries by matching over large and distributed libraries of mathematical knowledge.

Our approach consists in capturing some of the matching invariants by means of metadata, which can be automatically computed by a batch process (a Web spider). Every time the user submits a pattern to be matched, we use these metadata to filter a small subset of the library items which is made of matching candidates: we compute a query over the metadata from the matching pattern; the results of the query are the matching candidates. The correctness of the filtering operation, which is the property of not loosing matching terms, is a trivial consequence of the fact that our metadata capture the matching invariants.

Depending on the kind of matching we are interested in, we need to identify a different set of metadata and query over them. A second outcome of this work has been to identify a rather reasonable and accurate set of metadata that cover several common matching procedures. They are described in the Section 5.

To judge the efficacy of our approach, we implemented a prototype working over the HELM distributed library, made of about 40.000 theorems coming from the Coq proof-assistant library. The prototype has been implemented using the MathQL Level 1 query language [12] and has been turned into a Web service. A Web interface to the prototype is available following the appropriate link at the address: http://helm.cs.unibo.it/library.html. The prototype has been extensively used by one of the authors while developing a medium-sized Coq contribution. The experiment outcome was extremely positive: the author was able to locate and reuse several theorems that were given in parts of the library where the author did not expect them to be. As a consequence, the library reusage factor was improved and the development effort was reduced.

Since the prototype is implemented as a Web service, it can be easily exploited by theorem provers or proof-planners to retrieve lemmas or theorems and thus our search engine may become an important component in the Calculemus perspective. For instance, when a CAS needs to know whether a denominator is never zero, it can require a proof to a proof-planner that can use the search engine to retrieve some existent lemmas used to complete the proof.

An important feature of our filtering technology is that the filtering operation is not too strict, in the sense that wrong candidates that are "close enough" to matching terms are not dropped. These wrong candidates are very interesting, since the human user can often exploit them reducing his thesis to the one found; other times, instead, they can reveal missing assumptions in the user statement.

The interest in the wrong candidates, especially the ones that can be reduced to the user statements, is not a surprise. Indeed, several authors [10, 9] already pointed out that queries by matching are often too strict. For example, we would like the pattern $\forall x.x \geq 0$ to be able to match also the statement $\forall x.0 \leq x$. Other typical examples are patterns involving binary operators: often we would like to match up to associative or commutative rewriting. The most general framework for this kind of query is provided by the so called "query up to isomorphisms" [10]: both the pattern and the statements in the library are reduced to a decidable normal form before the matching; the reduction is such that a proof of the pattern can be automatically produced from the results of the query.

As a future work, we plan to apply our approach in a "query up to isomor-

phisms" setting. The required modifications are really minimal: Firstly the spider should reduce every term to its normal form *before* extracting the metadata, and secondly the query over the metadata should be generated from the pattern *after* reducing it to normal form. Even in the "query up to isomorphisms" setting, though, we still believe that "almost right" wrong candidates are important to spot false user conjectures and suggest new hypothesis or conditions to be required to make the theorem true.

Other future works involve identifying larger and larger sets of queries that can be automatically generated, possibly enlarging the set of metadata. Applying our technology to other libraries is also interesting to evaluate the scalability of our approach and the level of precision of our metadata. Indeed, the metadata identified in the previous sections are not peculiar of the Coq system, and should be retrieved from any library of mathematical content; however, it may be the case that some libraries could be better indexed with more ad-hoc set of metadatas, as the ones used for Mizar in [7].

References

1. Dewey Decimal Classification
   http://www.oclc.org/dewey
2. Library of Congress Classification Scheme,
   http://www.loc.gov
3. Mathematical Markup Language (MathML) Version 2.0, W3C Reccomendation 21 February 2001,
   http://www.w3.org/TR/MathML2
4. Mathematical Subject Classification, American Mathematical Society,
   http://www.ams.org/msc
5. OMDOC: A Standard for Open Mathematical Documents,
   http://www.mathweb.org/omdoc/omdoc.ps
6. The OpenMath Standard,
   http://www.openmath.org/cocoon/openmath/standard/index.html
7. Grzegorz Bancerek, Piotr Rudnicki, Information Retrieval in MML. In Proceedings of MKM2003. Springler-Verlag LNCS 2594, 2003.
8. Jorge Luis Borges, *The Library of Babel*, in Ficciones, Grove Press 1942.
9. David Delahaye, Roberto di Cosmo, Information Retrieval in a Coq Proof Library using Type Isomorphisms. In Proceedings of TYPES 99, Lökeberg. Springler-Verlag LNCS, 1999.
10. Roberto di Cosmo, *Isomorphisms of Types: from Lambda Calculus to Information Retrieval and Language Design*, Birkhauser, 1995, IBSN-0-8176-3763-X.
11. Harald Ganzinger, Robert Nieuwehuis, Pilar Nivela. Fast Term Indexing with Coded Context Trees. Journal of Automated Reasoning. To appear.
12. Ferruccio Guidi, *Searching and Retrieving in Content-Based Repositories of Formal Mathematical Knowledge*, Ph.D. Thesis in Computer Science, Technical Report UBLCS 2003-06, University of Bologna, March 2003.
13. Conor McBride, *Dependently Typed Functional Programs and their Proofs*, LFCS Ph.D. Thesis, 2000.
14. Michael Kohlhase, Andreas Franke. MBase: Representing Knowledge and Context for the Integration of Mathematical Software Systems, Journal of Symbolic Computation 23:4 (2001), pp. 365–402.