

TBox Classification in Parallel: Design and First Evaluation

Mina Aslani and Volker Haarslev

Concordia University, Montreal, Canada
{m.aslani,haarslev}@cse.concordia.ca

Abstract. One of the most frequently used inference services of description logic reasoners classifies all named classes of OWL ontologies into a subsumption hierarchy. Due to emerging OWL ontologies from the web community consisting of up to hundreds of thousand of named classes and the increasing availability of multi-processor and multi- or many-core computers, we extend the work on parallel TBox classification and propose a new algorithm that is sound and complete and demonstrates in a first experimental evaluation a low overhead in the number of subsumption tests due to parallel execution.

1 Motivation

Due to the recent popularity of OWL ontologies in the web one can observe a trend toward the development of very large or huge OWL-DL ontologies. For instance, well known examples from the bioinformatics or medical community are UMLS, GALEN or even ontologies with more concepts. Some (versions) of the ontologies consist of more than hundreds of thousands of named concepts/classes and have become challenging even for the most advanced and optimized description logic (DL) reasoners. Although specialized DL reasoners for certain sublogics (e.g., CEL for EL++) and OWL-DL reasoners such as FaCT++, Pellet, HermiT, or RacerPro could demonstrate impressive speed enhancements due to newly designed optimization techniques, one can expect the need for parallelizing description logic inference services in the near future in order to achieve a web-like scalability. Our research is also strongly motivated by recent trends in computer hardware where processors feature multi-cores (2 to 8 cores) or many-cores (tens or even hundreds of cores). These processors promise significant speed-ups for algorithms exploiting so-called thread-level parallelism. This type of parallelism is very promising for DL reasoning algorithms that can be executed in parallel but might share common data structures (e.g., and/or parallelism in proofs, classification of TBoxes, ABox realization or query answering).

First approaches on more scalable reasoning algorithms for ABoxes (sets of declarations about individuals) were investigated with the Racer architecture [11] where novel instance retrieval algorithms were developed and analyzed, which exploit a variety of techniques such as index maintenance, dependency analysis, precompletion generation, etc. Other research focused on scalable ABox reasoning with optimization techniques to partition ABoxes into independent parts

and/or creating condensed (summary) ABoxes [8, 9, 6]. These approaches rely on the observation that the structure of particular ABoxes is often redundant and these ABoxes contain assertions not needed for ABox consistency checking or query answering.

Parallel algorithms for description logic reasoning were first explored in the FLEX system [3] where various distributed message-passing schemes for rule execution were evaluated. The only other approach on parallelizing core description logic reasoning [13] reported promising results using multi-core/processor hardware, where the parallel treatment of disjunctions and individual merging (due to number restrictions) is explored. While there exists some work on parallel DL algorithms, on parallel reasoning for first-order theorem proving (with completely different proof techniques based on resolution), and on parallel distributed RDF inferencing (e.g., [16]), parallel TBox classification has only been addressed in [1]. There has also been substantial work on reasoning through modularity and partitioning knowledge bases (e.g., [7, 5, 4]) that might be applicable to our work.

In the following we extend the work on parallel TBox classification [1] and propose a new algorithm that is sound and complete although it runs in parallel. The implemented prototype system simulates parallel TBox classification with various parameters such as number of threads, size of partitions assigned to threads, etc. First results from a preliminary evaluation look very promising and indicate a very low overhead in the number of subsumption tests due to parallel execution.

2 The New Parallel TBox Classifier

This section describes the architecture of the implemented system and its underlying *sound and complete* algorithm for parallel classification of DL ontologies. To compute the hierarchy in parallel, we developed a simulator using a multi-threaded architecture providing control parameters such as number of threads, number of concepts (also called partition size) to be inserted per thread, and strategies used to partition a given set of concepts. The simulator reads an input file containing a list of concept names to be classified and information about them. The per-concept information available in the file includes its name, parents (in the complete taxonomy), told subsumers, told disjoints, and pseudo model information. The information about parents is used to compute the set of ancestors and descendants of a concept. Told information consists of subsumers and disjoints that can be easily extracted from axioms without requiring proof procedures, e.g. the axiom $A \sqsubseteq B \sqcap \neg C$ would result in information asserting B as told subsumer of A and C as told disjoint of A . With the exception of told subsumers this information is only used for (i) emulating a tableau subsumption test, i.e., by checking whether a possible subsumer (subsumee) is in the list of ancestors (descendants) of given concept, and (ii) in order to verify the completeness of the taxonomy computed by the parallel classifier. The input information substitutes for an implemented tableaux reasoning procedure, hence

makes the parallel classifier independent of a particular DL logic or reasoner. Currently, Racer is used to generate this file for a given OWL-DL ontology after performing TBox classification.

The told subsumer information is passed to a preprocessing algorithm which creates a taxonomy skeleton based on the already known (i.e. told) subsumptions and generates a topological-order list (e.g. depth-first traversal). Using a topological sorting algorithm, the partial order can be serialized such that a total order between concept names (or sets of concept names) is defined. During classification, the concept names are processed in the order of the topological order. In our topological order list, from left to right, parent concepts precede child concepts. To manage concurrency in our system, at least two shared-memory approaches could be taken into account by using either (i) sets of local trees (so-called ParTree approach) or (ii) one global tree. In the ParTree algorithm [14] a local tree would be assigned to each thread, and after all the threads have finished the construction of their local hierarchy, the local trees need to be merged into one global tree. TBox classification through a local tree algorithm would not need any communication or synchronization between the threads. ParTree is well suited for distributed systems which do not have shared memory. The global tree approach was chosen because it implements a shared space which is accessible to different threads running in parallel and avoids the large scale overhead of ParTree on synchronizing local trees. To ensure data integrity a lock mechanism for single nodes is used. This allows a proper lock granularity and helps to increase the number of simultaneous write accesses to the subsumption hierarchy under construction.

Most TBox classification algorithms are based on two (symmetric) tasks (e.g., see [2]). The first phase (top search) determines the parents of a given concept to be inserted into the subsumption tree. It starts with the top concept (\top) and tries to push the given concept below the children of the current concept and repeats this process with the goal to push the given concept as much to the bottom of the subsumption tree as possible. Whenever a concept in the tree subsumes the given concept, it is pushed below this subsumer. The second phase (bottom search) determines the children of a given concept. It starts with the bottom concept (\perp) and tries to move the given concept above the parents of the current concept and repeats this process with the goal to move the current concept up in the tree as much as possible. Whenever a concept in the tree is subsumed by the given concept, it is moved above of this subsumee. Eventually, the given concept is correctly positioned in the current subsumption hierarchy. Both phases tag nodes of the tree ('visited', 'positive', 'negative') to prune the search and avoid visiting already processed nodes. For instance, 'positive' is used to tag nodes already known as (told) subsumers and 'negative' for already known as (told) disjoints.

The work in [2] is an example for algorithms that incrementally construct a subsumption tree and are highly optimized for sequential execution. In [10] some of these techniques were extended to better deal with huge TBox hierarchies but these algorithms are still based on a sequential execution. A recent approach

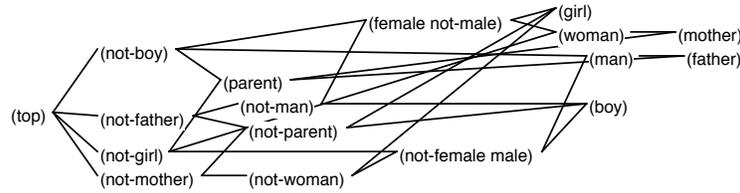


Fig. 1. Complete subsumption hierarchy for yaya-1

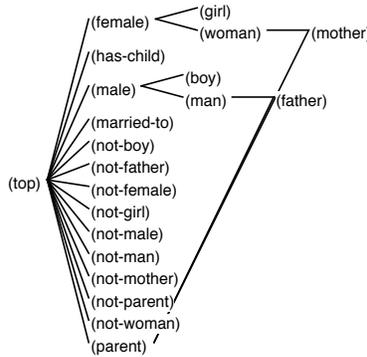


Fig. 2. Told subsumer hierarchy for yaya-1

[15] on TBox classification exploits partial information about OWL subclass relationships to reduce the number of subsumption tests and, thus, improves the algorithms presented in [2].

2.1 Example Scenario

In [1] the degree of incompleteness caused by classifying partitions of concepts in parallel was tested. For a variety of ontologies it turned out that a surprisingly few number of subsumptions were missed. This motivated the work in this paper. In the following we illustrate two scenarios which may cause that a concept is misplaced in the taxonomy due to parallel classification. We use a very small ontology named yaya-1 with 16 concepts (see Fig. 1 and 2).

For this example, we configured our system so that it runs with 4 threads and 3 number-of-tasks-per-thread. As explained previously, in parallel classification the topological sort order divides concept partitions between the threads (e.g. round-robin). For instance, in Fig. 3 a list of concepts allocated to each thread is

- thread#1 → (female not-male), girl, parent
- thread#2 → woman, mother, (male not-female)
- thread#3 → man, boy, father
- thread#4 → not-boy, not-father, not-girl
- thread#1 → not-man, not-mother, not-parent, not-woman

Fig. 3. Concept assignments to each thread for classifying yaya-1

shown. The two possible scenarios that may lead to a situation where the correct place of a concept in the hierarchy is overlooked are described as follows.

Scenario I: In top search, as the new concept is pushed downward, right after the children of the current concept have been processed, at least one new child is added by another thread. In this scenario, the top search for the new concept is not aware of the recent change and this might cause missing subsumptions if there is any interaction between the new concept and the added children. The same might happen in bottom search if the bottom search for the new concept is not informed of the recent change to the list of parents of the current node.

Scenario II: Between the time that top search has been started to find the location of a new concept in the taxonomy and the time that its location has been decided, another thread has placed at least one concept into the hierarchy which the new concept has an interaction with. Again, this might cause missing subsumptions and is analogously also applicable to bottom search.

In our example (yaya-1), due to the small size of the taxonomy, scenario I was not encountered, however, scenario II occurred in our experiments because thread#1 inserted (female not-male)¹ and thread#2 added woman independently into the taxonomy and due to the parallelism each thread did not have any information regarding the latest concept insertion by other threads (see also Fig. 3). Hence, both (female not-male) and woman were initially placed under the top concept although woman should be a child of (female not-male) (see Fig. 1). This was discovered and corrected by executing lines 6-7, 16-17, and 25-36 in Algorithm 2 as shown below.

2.2 Algorithms for Parallel Classification

The procedure `parallel_tbox_classification` is sketched in Algorithm 1. It is called with a list of named concepts and sorts them in topological order w.r.t. to the initial taxonomy created from the already known told ancestors and descendants of each concept (using the told subsumer information). The classifier assigns in a round-robin manner partitions with a fixed size from the concept list to idle threads and activates these threads with their assigned partition using the procedure `insert_partition` outlined in Algorithm 2. All threads work in parallel with the goal to construct a global subsumption tree (taxonomy). They also share a global array `inserted_concepts` indexed by thread identifications. Nodes in the global tree as well as entries in the array will be locked for modification.

The procedure `insert_partition` inserts all concepts of a given partition into the global taxonomy. For updating a concept or its parents or children, it locks the corresponding nodes. It first performs for each concept *new* the top-search phase (starting from the top concept) and possibly repeats the top-search phase for *new* if other threads updated the list of children of its parents. Then, it sets the parents of *new* and adds *new* for each parent to its list of children. Afterwards the bottom-search phase (starting from the bottom concept) is performed.

¹ This notation indicates that the concepts `female` and `not-male` are synonyms for each other.

Algorithm 1: parallel_tbox_classification(*concept_list*)

```

topological-order-list  $\leftarrow$  topological_order(concept_list)
repeat
  wait until an idle thread  $t_i$  becomes available
  select a partition  $p_i$  from topological-order-list
  run thread  $t_i$  with insert_partition( $p_i, t_i$ )
until all concepts in topological-order-list are inserted
compute ratio and overhead
print statistics

```

Algorithm 2: insert_partition(*partition, id*)

```

1: lock(inserted_concepts(id))
2: inserted_concepts(id)  $\leftarrow$   $\emptyset$ 
3: unlock(inserted_concepts(id))
4: for all new  $\in$  partition do
5:   parents  $\leftarrow$  top_search(new,  $\top$ )
6:   while  $\neg$  consistent_in_top_search(parents, new) do
7:     parents  $\leftarrow$  top_search(new,  $\top$ )
8:   lock(new)
9:   predecessors(new)  $\leftarrow$  parents
10:  unlock(new)
11:  for all pred  $\in$  parents do
12:    lock(pred)
13:    successors(pred)  $\leftarrow$  successors(pred)  $\cup$  {new}
14:    unlock(pred)
15:  children  $\leftarrow$  bottom_search(new,  $\perp$ )
16:  while  $\neg$  consistent_in_bottom_search(children, new) do
17:    children  $\leftarrow$  bottom_search(new,  $\perp$ )
18:  lock(new)
19:  successors(new)  $\leftarrow$  children
20:  unlock(new)
21:  for all succ  $\in$  children do
22:    lock(succ)
23:    predecessors(succ)  $\leftarrow$  predecessors(succ)  $\cup$  {new}
24:    unlock(succ)
25:  check  $\leftarrow$  check_if_concept_inserted(new, inserted_concepts(id))
26:  if check  $\neq$  0 then
27:    if check = 1  $\vee$  check = 3 then
28:      new_predecessors  $\leftarrow$  top_search(new,  $\top$ )
29:      lock(new)
30:      predecessors(new)  $\leftarrow$  new_predecessors
31:      unlock(new)
32:    if check = 2  $\vee$  check = 3 then
33:      new_successors  $\leftarrow$  bottom_search(new,  $\perp$ )
34:      lock(new)
35:      successors(new)  $\leftarrow$  new_successors
36:      unlock(new)
37:  for all busy threads  $t_i \neq id$  do
38:    lock(inserted_concepts( $t_i$ ))
39:    inserted_concepts( $t_i$ )  $\leftarrow$  inserted_concepts( $t_i$ )  $\cup$  {new}
40:    unlock(inserted_concepts( $t_i$ ))

```

Analogously to the top-search phase the bottom search is possibly repeated and sets the children of *new* and updates the parents of the children of *new*. After finishing the top and bottom search for *new* it is checked again whether other threads updated its entry in *inserted_concepts* and the top and/or bottom search needs to be repeated. This step needs to be done only once. Finally, *new* is added to the entries in *inserted_concepts* of all other busy threads.

In order to avoid unnecessary tree traversals and tableau subsumption tests when computing the subsumption hierarchy, the parallel classifier adapted the enhanced traversal method [2], which is an algorithm that was designed for sequential execution. Algorithm 3 and 4 outline the traversal procedures for the top-search phase.

Algorithm 3: top_search(*new*,*current*)

```

mark(current, 'visited')
pos_succ ← ∅
captured_successors(new)(current) ← successors(current)
for all y ∈ successors(current) do
  if enhanced_top_subs(y,new) then
    pos_succ ← pos_succ ∪ {y}
if pos_succ = ∅ then
  return {current}
else
  result ← ∅
  for all y ∈ pos_succ do
    if y not marked as 'visited' then
      result ← result ∪ top_search(new,y)
  return result

```

The procedure top_search outlined in Algorithm 3 recursively traverses the taxonomy top-down from a current concept and tries to push the new concept down the taxonomy as far as possible by traversing the children of the current concept. It uses an auxiliary procedure enhanced_top_subs (outlined in Algorithm 4) which itself uses an auxiliary procedure subsumes (not specified here) that implements a subsumption test.

In a symmetric manner the procedure bottom_search traverses the taxonomy bottom-up from a current concept and tries to push the new concept up the taxonomy as far as possible. It uses an auxiliary procedure enhanced_bottom_subs. Both procedures are omitted for ease of presentation.

To resolve the possible incompleteness caused by parallel classification, we utilize Algorithms 5, 6 and 7. The procedure consistent_in_bottom_search is not shown here because it mirrors consistent_in_top_search.

Algorithms 5 and 6 illustrate the solution for scenario I described in Section 2.1. As already described, in top search we start traversing from the top concept to locate the concept *new* in the taxonomy. At time *t1*, when top_search is called, we capture the children information “captured_successors” of the concept *current*; the children information is stored relative² to the concept *new* being

² Otherwise a different thread could overwrite captured_successors for node *current*. This is now prevented because each concept (*new*) is inserted by only one thread.

Algorithm 4: `enhanced_top_sub(current,new)`

```

if current marked as ‘positive’ then
  return true
else if current marked as ‘negative’ then
  return false
else if for all  $z \in \text{predecessors}(\textit{current})$ 
  enhanced_top_sub(z,new)
  and subsumes(current,new) then
  mark(current,‘positive’)
  return true
else
  mark(current,‘negative’)
  return false

```

Algorithm 5: `consistent_in_top_search(parents,new)`

```

for all  $pred \in \textit{parents}$  do
  if successors(pred)  $\neq$  captured_successors(new)(pred) then
     $\textit{diff\_children} \leftarrow \text{successors}(\textit{pred}) \setminus \text{captured\_successors}(\textit{new})(\textit{pred})$ 
    for all  $\textit{diff} \in \textit{diff\_children}$  do
      if check_interactions(diff,new) then
        return false
  return true

```

inserted (we use an array of arrays) and captures the successors of the concept *current* (see Algorithm 3). As soon as `top_search` is finished at time t_2 , and the parents of the concept *new* have been determined, we check if there has been any update on the children list of the computed parents for *new* between t_1 and t_2 (e.g., see Algorithm 5 on how this is discovered). If there is any inconsistency and also if there is a subsumption possible³ between *new* and any concept newly added to the children list, we rerun `top_search` until there is no inconsistency (see line 6 in Algorithm 2).

The same process as illustrated in Algorithm 5 happens in bottom search. The only difference is that parents information is captured when bottom search starts; and when bottom search finishes, the inconsistency and interaction is checked between the parents list of the computed children for *new* and the “`captured_predecessors`”.

Algorithms 6 and 7 describe the solution for scenario II; every time a thread inserts a concept in the taxonomy, it notifies the other threads by adding the concept name to their “`inserted_concepts`” list. Therefore, as soon as a thread finds the parents and children of the *new* concept by running `top_search` and `bottom_search`; it checks if there is any interaction between *new* concept and

³ This is checked by `subsumption_possible` using pseudo model merging [12], where a sound but incomplete test for non-subsumption on cached pseudo models of named concepts and their negation is utilized.

Algorithm 6: `check_interactions(diff,new)`

return `subsumption_possible(diff,new)`

Algorithm 7: `check_if_concept_inserted(new,inserted_concepts)`

```

if inserted_concepts = ∅ then
  return 0
else
  for all concept ∈ inserted_concepts do
    if check_interactions(concept,new) then
      if check_interactions(new,concept) then
        return 3
      else
        return 1
    else if check_interactions(new,concept) then
      if check_interactions(concept,new) then
        return 3
      else
        return 2
  return 0

```

the concepts located in the “inserted_concepts” list. Based on the interaction, `top_search` and/or `bottom_search` need to be repeated accordingly.

Proposition 1 (Completeness of Parallel TBox Classifier) *The proposed algorithms are complete for TBox classification.*

TBox classification based on top search and bottom search is complete in the sequential case. This means that the subsumption algorithms will find all subsumption relationships between concepts of a partition assigned to a single thread. The threads lock and unlock nodes whenever they are updating the information about a node in the global subsumption tree. Thus, we need to consider only the scenarios where two concepts C and D are inserted in parallel by different threads (e.g., `thread#1` inserts concept C while `thread#2` inserts concept D). In principle, if top (bottom) search pushed a new concept down (up), the information about children (parents) of a traversed node E could be incomplete because another thread might later add more nodes to the parents or children of E that were not considered when determining whether the concept being inserted subsumes or is subsumed by any of these newly added nodes. This leads to two scenarios that need to be examined for incompleteness.

W.l.o.g. we restrict our analysis to the case where a concept C is a parent of a concept D in the complete subsumption tree (CT). Let us assume that our algorithms would not determine this subsumption, i.e., in the computed (incomplete) tree (IT) the concept C is not a parent of D .

Case I: top_search incomplete for D : After D has been pushed down the tree IT as far as possible by top search (executed by `thread#2`) and top search has traversed the children of a concept E and E has become the parent of D , C

Table 1. Characteristics of the used test ontologies.

Ontology	DL language	No. of named concepts
Embassi-2	\mathcal{ALCHN}	657
Embassi-3	\mathcal{ALCHN}	1,121
Galen	\mathcal{SHN}	2,730
Galen1	\mathcal{ALCH}	2,730
Galen2	\mathcal{ELH}	3,928
FungalWeb	$\mathcal{ALCHIN}(\mathcal{D})$	3,603
Umls-2	$\mathcal{ALCHIN}(\mathcal{D})$	9,479
Tambis-2a	\mathcal{ELH}	10,116

is inserted by thread#1 as a new child of E . In line 6 of Algorithm 2 `top_search` is iteratively repeated for the concept new as long as `consistent_in_top_search` finds a discrepancy between the captured and current successors of the parents of the newly inserted concept new . After finishing top and bottom search, Algorithm 2 checks again in lines 27-28 whether top search needs to be repeated due to newly added nodes. If any of the newly added children of D would subsume C and become a parent of C , the repeated execution of `top_search` would find this subsumption. This contradicts our assumption.

Case II: bottom_search incomplete for C : After C has been pushed up the tree IT as far as possible by bottom search (executed by thread#1) and bottom search has traversed the parents of a concept E and E has become a child of C , D is inserted by thread#2 as a new parent of E . In line 16 of Algorithm 2 `bottom_search` is iteratively repeated for the concept new as long as `consistent_in_bottom_search` finds a discrepancy between the captured and current predecessors of the children of the newly inserted concept new . After finishing top and bottom search, Algorithm 2 checks again in lines 32-33 whether bottom search needs to be repeated due to newly added nodes. If C would subsume any of the newly added parents of D and it would become a child of C , the repeated execution of `bottom_search` would find this subsumption. This contradicts our assumption.

3 Evaluation

The Parallel TBox Classifier has been developed to speed up the classification time especially for large ontologies by utilizing parallel threads sharing the same memory. The benchmarking can be configured so that it runs various experiments over ontologies. We evaluated it with a collection of 8 mostly publicly available ontologies. Their name, size in number of named concepts, and used DL is shown in Table 1. As mentioned in the previous section, two parameters influence the parallel TBox classification, namely number of tasks/concepts per thread and number of threads; the number of tasks/concepts per thread was set to 5 and number of threads to 2 in our empirical experiments.

To better compare the performance between the sequential and parallel case, we assume that every subsumption test runs in time $t1$ and in the sequential and parallel case the same amount of time is used for an executed subsumption test.

Table 2. Subsumptions tests and their ratio for the test ontologies.

	Embassi-2	Embassi-3	Galen	Galen1
Subs. Tests in sequent.	154,034	420,912	2,706,412	2,688,107
Subs. Tests in thread#1	76,267	217,324	1,363,321	1,367,302
Subs. Tests in thread#2	77,767	214,633	1,354,297	1,348,281
Worst Case Ratio	50.48%	51.63%	50.37%	50.86%
Overhead	1.64%	2.62%	0.41%	1.02%
	Galen2	FungalWeb	Umls-2	Tambis-2a
Subs. Tests in sequent.	5,734,976	4,996,932	87,423,341	36,555,225
Subs. Tests in thread#1	2,929,276	2,518,676	44,042,203	18,342,944
Subs. Tests in thread#2	2,893,716	2,490,329	44,025,988	18,261,532
Worst Case Ratio	51.07%	50.40%	50.37%	50.17%
Overhead	1.53%	0.24%	0.73%	0.13%

Subsumption tests can be expensive and, hence, are preferred to be avoided by optimization techniques such as pseudo model merging [12].

$$Ratio = \frac{MaxOfSubsTests}{TST_s} \quad (1) \quad Overhead = \frac{TST_p - TST_s}{TST_s} \quad (2)$$

The ratio illustrated in Equation 1 uses TST_s , the number of times a subsumption test was computed in the sequential case, and $MaxOfSubsTests$, the maximum number of subsumption tests performed in all threads. Similarly, Equation 2 defines the overhead (where the index p refers to the parallel case).

Table 2 shows an excellent performance increase and a surprisingly small overhead when using the Parallel TBox Classifier. Using two threads the maximum number of subsumption test for all ontologies could be reduced to roughly one half compared to the sequential case. The overhead as defined in Equation 2 varies between 0.13% and 2.62%. The overhead is mostly determined by the quality of the told subsumers and disjoints information, the imposed order of traversal within a partitioning, and the division of the ordered concept list into partitions. In general, one should try to insert nodes as close as possible to their final order in the tree using a top to bottom strategy.

4 Conclusion

In this paper, we described an architecture for parallelizing well-known algorithms for TBox classification. Our work is targeted for ontologies where independent partitions cannot be easily constructed; therefore we did not use the previously mentioned modularity approaches in our system. The first experimental evaluation of our techniques shows very promising results because the overhead for ensuring completeness is surprisingly small. In our next steps we plan to extend our tests with different configurations of threads and partition sizes and a larger variety of test ontologies. We intend to feed recorded runtimes for performing single subsumption tests into our simulator in order to make the computation of the overhead more accurate. We also plan to implement and test our approach in a multi-core and multi-processor environment.

References

1. Aslani, M., Haarslev, V.: Towards parallel classification of TBoxes. In: Proc. of the 2008 Int. Workshop on Description Logics, Dresden, Germany, May 13-16 (2008)
2. Baader, F., Franconi, E., Hollunder, B., Nebel, B., Profitlich, H.: An empirical analysis of optimization techniques for terminological representation systems or: Making KRIS get a move on. *Applied Artificial Intelligence. Special Issue on Knowledge Base Management* 4(2), 109–132 (1994)
3. Bergmann, F., Quantz, J.: Parallelizing description logics. In: Proc. of 19th Ann. German Conf. on Artificial Intelligence. pp. 137–148. LNCS, Springer-Verlag (1995)
4. Cuenca Grau, B., Halaschek-Wiener, C., Kazakov, Y.: History matters: Incremental ontology reasoning using modules. In: Proc. of the 6th Int. Semantic Web Conf. (ISWC 2007), Busan, South Korea, Nov. 11-15 (2007)
5. Cuenca Grau, B., Horrocks, I., Kazakov, Y., Sattler, U.: A logical framework for modularity of ontologies. In: In Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007), Busan, South Korea, Nov. 11-15. pp. 298–303 (2007)
6. Dolby, J., Fokoue, A., Kalyanpur, A., Kershenbaum, A., Ma, L., Schonberg, E., Srinivas, K.: Scalable semantic retrieval through summarization and refinement. In: 21st Conf. on Artificial Intelligence (AAAI). pp. 299–304. AAAI Press (2007)
7. Eyal, A., McIlraith, S.: Partition-based logical reasoning for first-order and propositional theories. *Artificial Intelligence* 162(1-2), 49–88 (2005)
8. Fokoue, A., Kershenbaum, A., Ma, L., Schonberg, E., Srinivas, K.: The summary ABox: Cutting ontologies down to size. In: Proc. of Int. Semantic Web Conf. (ISWC). LNCS, vol. 4273, pp. 343–356. Springer-Verlag (2006)
9. Guo, Y., Heflin, J.: A scalable approach for partitioning OWL knowledge bases. In: Proc. 2nd Int. Workshop on Scalable Semantic Web Knowledge Base Systems, Athens, USA. pp. 47–60 (2006)
10. Haarslev, V., Möller, R.: High performance reasoning with very large knowledge bases: A practical case study. In: Proc. of the 17th Int. Joint Conf. on Artificial Intelligence, IJCAI-01, Aug. 4-10, Seattle, USA. pp. 161–166 (2001)
11. Haarslev, V., Möller, R.: On the scalability of description logic instance retrieval. *Journal of Automated Reasoning* 41(2), 99–142 (2008)
12. Haarslev, V., Möller, R., Turhan, A.Y.: Exploiting pseudo models for TBox and ABox reasoning in expressive description logics. In: Proc. of the Int. Joint Conf. on Automated Reasoning, June 18-23, 2001, Siena, Italy. pp. 61–75 (2001)
13. Liebig, T., Müller, F.: Parallelizing tableaux-based description logic reasoning. In: Proc. of 3rd Int. Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS '07), Vilamoura, Portugal, Nov 27. LNCS, vol. 4806, pp. 1135–1144. Springer-Verlag (2007)
14. Shan, H., Singh, J.P.: Parallel tree building on a range of shared address space multiprocessors: Algorithms and application performance. In: 12th Int. Parallel Processing Symposium (IPPS '98), March 30 - April 3, 1998, Orlando, Florida, USA. pp. 475–484 (1998)
15. Shearer, R., Horrocks, I.: Exploiting partial information in taxonomy construction. In: Proc. of the 8th International Semantic Web Conference (ISWC 2009) (2009)
16. Urbani, J., Kotoulas, S., Oren, E., van Harmelen, F.: Scalable distributed reasoning using MapReduce. In: International Semantic Web Conference. pp. 634–649 (2009)