ECAI 2020 G.D. Giacomo et al. (Eds.) © 2020 The authors and IOS Press. This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0). doi:10.3233/FAIA200178

Reuse, Reduce and Recycle: Optimizing Reiter's HS-Tree for Sequential Diagnosis

Patrick Rodler¹

Abstract. Reiter's HS-Tree is one of the most popular diagnostic search algorithms due to its desirable properties and general applicability. In sequential diagnosis, where the addressed diagnosis problem is subject to successive change through the acquisition of additional knowledge about the diagnosed system, HS-Tree is used in a stateless fashion. That is, the existing search tree is discarded when new knowledge is obtained, albeit often large parts of the tree are still relevant and have to be rebuilt in the next iteration, involving redundant operations and costly reasoner calls. As a remedy, we propose DynamicHS, a variant of HS-Tree that avoids these redundancy issues by maintaining state throughout sequential diagnosis while preserving all desirable properties of HS-Tree. Evaluations in a problem domain where HS-Tree is the state-of-the-art diagnostic method reveal stable and significant time savings achieved by DynamicHS.

1 Introduction

In model-based diagnosis, given a *diagnosis problem instance* (*DPI*)—consisting of the *system description*, the system *components*, and the *observations*—a (*minimal*) *diagnosis* is an (irreducible) set of components that, when assumed abnormal, leads to consistency between system description (predicted behavior) and observations (real behavior). In many cases, there is a substantial number of competing (minimal) diagnoses. To isolate the *actual diagnosis* (which pinpoints the *actually* faulty components), *sequential diagnosis* (minimal) methods collect additional system observations (called *measurements*) to gradually refine the set of diagnoses.

One of the most popular and widely used algorithms for the computation of diagnoses in model-based diagnosis is Reiter's HS-Tree [32]. It is adopted in various domains such as for the debugging of software [2, 53] ontologies and knowledge bases [16, 23, 27, 33], hardware [17], recommender systems [15], configuration systems [13], and circuits [32]. The reasons for its widespread adoption are that (*i*) it is broadly applicable, because all it needs is a system description in some monotonic knowledge representation language for which a sound and complete inference method exists, (*ii*) it is sound and complete, as it computes only minimal diagnoses and can, in principle, output all minimal diagnoses, and (*iii*) it computes diagnoses in best-first order according to a given preference criterion.

However, HS-Tree per se does not encompass any specific provisions for being used in an iterative way. In other words, the DPI to be solved is assumed constant throughout the execution of HS-Tree. As a consequence of that, the question we address in this work is whether HS-Tree can be optimized for adoption *in a sequential diagnosis scenario*, where the DPI to be solved is subject to successive change (information acquisition through measurements). Already Raymond Reiter, in his seminal paper [32] from 1987, asked:

When new diagnoses do arise as a result of system measurements, can we determine these new diagnoses in a reasonable way from the (...) HS-Tree already computed in determining the old diagnoses?

To the best of our knowledge, no algorithm has yet been proposed that sheds light on this very question.

As a result, state-of-the-art sequential approaches which draw on HS-Tree for diagnosis computation handle the varying DPI to be solved by re-invoking HS-Tree each time a new piece of system knowledge (measurement outcome) is obtained. This amounts to a *discard-and-rebuild* usage of HS-Tree, where the search tree produced in one iteration is dropped prior to the next iteration, where a new one is built from scratch. As the new tree obtained after incorporating the information about one measurement outcome usually quite closely resembles the existing tree, this approach generally requires substantial redundant computations, which often involve a significant number of expensive reasoner calls.

Motivated by that, we propose DynamicHS, a stateful variant of HS-Tree that pursues a reuse-and-adapt strategy and is able to manage the dynamicity of the DPI throughout sequential diagnosis while avoiding the mentioned redundancy issues. The idea is to maintain one (successively refined) tree data structure and to exploit the information it contains to enhance computations in the subsequent iteration(s), e.g., by circumventing costly reasoner invocations. The main objective of DynamicHS is to allow for a better efficiency than HS-Tree while maintaining all aforementioned advantages (generality, soundness, completeness, best-first property) of the latter. Evaluations on various knowledge-base debugging problems-a domain where HS-Tree is the prevalent means for diagnosis computation [16, 20, 23, 27, 33, 44]—prove the reasonability of DynamicHS, which exhibits significant (up to more than 70%, on average more than 50%) performance gains over HS-Tree and tops the latter in more than 98 % of the executed experiments.

2 Preliminaries

We briefly characterize technical concepts used throughout this work, based on the framework of [33] which is (slightly) more general [38] than Reiter's theory of diagnosis [32].²

Diagnosis Problem Instance (DPI). We assume that the diagnosed system, consisting of a set of components $\{c_1, \ldots, c_k\}$, is described by a finite set of logical sentences $\mathcal{K} \cup \mathcal{B}$, where \mathcal{K} (possi-

¹ University of Klagenfurt, Austria, email: patrick.rodler@aau.at

 $^{^{2}}$ The main reason for using this more general framework is its ability to handle negative measurements (things that must *not* be entailed), which are helpful, e.g., for diagnosing knowledge bases [13, 44].

bly faulty sentences) includes knowledge about the behavior of the system components, and \mathcal{B} (correct background knowledge) comprises any additional available system knowledge and observations. More precisely, there is a one-to-one relationship between sentences $ax_i \in \mathcal{K}$ and components c_i , where ax_i describes the nominal behavior of c_i (weak fault model). E.g., if c_i is an AND-gate in a circuit, then $ax_i := out(c_i) = and(in1(c_i), in2(c_i)); \mathcal{B}$ in this case might contain sentences stating, e.g., which components are connected by wires, or observed circuit outputs. The inclusion of a sentence ax_i in \mathcal{K} corresponds to the (implicit) assumption that c_i is healthy. Evidence about the system behavior is captured by sets of positive (P)and negative (N) measurements [11, 13, 32]. Each measurement is a logical sentence; positive ones $p \in P$ must be true and negative ones $n \in N$ must not be true. The former can be, depending on the context, e.g., observations about the system, probes or required system properties. The latter model properties that must not hold for the system, e.g., if \mathcal{K} is a medical knowledge base to be debugged, a negative test case might be "every tumor causes pain." We call $\langle \mathcal{K}, \mathcal{B}, P, N \rangle$ a diagnosis problem instance (DPI).

Diagnoses. Given that the system description along with the positive measurements (under the assumption \mathcal{K} that all components are healthy) is inconsistent, i.e., $\mathcal{K} \cup \mathcal{B} \cup P \models \bot$, or some negative measurement is entailed, i.e., $\mathcal{K} \cup \mathcal{B} \cup P \models n$ for some $n \in N$, some assumption(s) about the healthiness of components, i.e., some sentences in \mathcal{K} , must be retracted. We call such a set of sentences $\mathcal{D} \subseteq \mathcal{K}$ a *diagnosis* for the DPI $\langle \mathcal{K}, \mathcal{B}, P, N \rangle$ iff $(\mathcal{K} \setminus \mathcal{D}) \cup \mathcal{B} \cup P \not\models x$ for all $x \in N \cup \{\bot\}$. We say that \mathcal{D} is a *minimal diagnosis* for dpi iff there is no diagnosis $\mathcal{D}' \subset \mathcal{D}$ for dpi. The set of minimal diagnoses is representative for all diagnoses (under the weak fault model [24]), i.e., the set of all diagnoses is exactly given by the set of all supersets of all minimal diagnoses. Therefore, diagnosis approaches usually restrict their focus to only minimal diagnoses. In the following, we denote the set of all minimal diagnoses for a DPI dpi by diag(dpi). We furthermore denote by \mathcal{D}^* the *actual diagnosis* (sought solution for DPI) which pinpoints the actually faulty axioms, i.e., all elements of \mathcal{D}^* are in fact faulty and all elements of $\mathcal{K} \setminus \mathcal{D}^*$ are in fact correct. Conflicts. Useful for diagnosis computation is the notion of a conflict [11, 32]. A conflict is a set of healthiness assumptions for components c_i that cannot all hold given the current knowledge about the system. More formally, $C \subseteq \mathcal{K}$ is a *conflict* for the DPI $\langle \mathcal{K}, \mathcal{B}, P, N \rangle$ iff $\mathcal{C} \cup \mathcal{B} \cup P \models x$ for some $x \in N \cup \{\bot\}$. We call \mathcal{C} a minimal conflict for dpi iff there is no conflict $\mathcal{C}' \subset \mathcal{C}$ for dpi. In the following, we denote the set of all minimal conflicts for a DPI dpi by conf(dpi). A (minimal) diagnosis for dpi is then a (minimal) hitting set of all minimal conflicts for dpi [32]. X is a hitting set of a collection of sets **S** iff $X \subseteq \bigcup_{S_i \in \mathbf{S}} S_i$ and $X \cap S_i \neq \emptyset$ for all $S_i \in S$.

Sequential Diagnosis. Given multiple minimal diagnoses for a DPI, a sequential diagnosis process can be initiated. It involves a recurring execution of four phases, (*i*) computation of a set of leading (e.g., most probable) minimal diagnoses, (*ii*) selection of the best measurement based on these, (*iii*) conduction of measurement actions, and (*iv*) exploitation of the measurement outcome to refine the system knowledge. The goal in sequential diagnosis is to achieve sufficient diagnostic certainty (e.g., a single or highly probable remaining diagnosis) with highest efficiency. At this, the overall efficiency is determined by the costs for *measurement conduction* and for *computations of the diagnosis engine*. Whereas the former—which is not the topic of this work—can be ruled by proposing appropriate (low-cost, informative) measurements [11, 30, 35, 39, 40, 41, 44], the latter is composed of the time required for *diagnosis computation*, for *measurement selection*, as well as for the *system knowledge update*. We

Algorithm 1 Sequential Diagnosis

- **Input:** DPI $dpi_0 := \langle \mathcal{K}, \mathcal{B}, P, N \rangle$, probability measure pr, number $ld (\geq 2)$ of minimal diagnoses to be computed per iteration, heuristic heur for measurement selection, boolean dynamic (use DynamicHS if true, HS-Tree otherwise)
- **Output:** $\{\mathcal{D}\}$ where $\hat{\mathcal{D}}$ is the only remaining diagnosis for the extended DPI $\langle \mathcal{K}, \mathcal{B}, \mathcal{P} \cup \mathcal{P}', \mathcal{N} \cup \mathcal{N}' \rangle$

1: .	$P' \leftarrow \emptyset, N' \leftarrow \emptyset$	> performed measurements
2:]	$\mathbf{D}_{\checkmark} \leftarrow \emptyset, \mathbf{D}_{\times} \leftarrow \emptyset, \text{state} \leftarrow \langle [[]], [], \emptyset, \emptyset \rangle \triangleright$	initial state of DynamicHS
3: 1	while true do	
4:	if dynamic then	
5:	$\langle \mathbf{D}, state \rangle \leftarrow DYNAMICHS(dpi_0, P', N',$	$pr, ld, \mathbf{D}_{\checkmark}, \mathbf{D}_{\times}, state)$
6:	else	
7:	$\mathbf{D} \leftarrow \text{HS-TREE}(dpi_0, P', N', pr, ld)$	
8:	if $ \mathbf{D} = 1$ then return \mathbf{D}	
9:	$mp \leftarrow \text{computeBestMeasPoint}(\mathbf{D}, dpi_0, I)$	P', N', pr, heur)
10:	$outcome \leftarrow PERFORMMEAS(mp) > oracle$	e inquiry (user interaction)
11:	$\langle P', N' \rangle \leftarrow \text{ADDMEAS}(mp, outcome, P', N)$	T')
12:	if dynamic then	
13:	$\langle \mathbf{D}_{\checkmark}, \mathbf{D}_{\times} \rangle \leftarrow \text{ASSIGNDIAGSOKNOK}(\mathbf{D}, d)$	$lpi_0, P', N')$

address the efficiency optimization problem in sequential diagnosis by suggesting new methods (DynamicHS algorithm) for the diagnosis computation and knowledge update processes.

3 DynamicHS Algorithm

Inputs and Outputs. DynamicHS (Alg. 2) accepts the following arguments: (1) an initial DPI $dpi_0 = \langle \mathcal{K}, \mathcal{B}, P, N \rangle$, (2) the already accumulated positive and negative measurements P' and N', (3) a probability measure pr (allowing to compute the probability of diagnoses), (4) a stipulated number $ld \geq 2$ of diagnoses to be returned, (5) the set of those diagnoses returned by the previous DynamicHS run that are consistent (\mathbf{D}_{\checkmark}) and those that are inconsistent (\mathbf{D}_{\times}) with the latest added measurement, and (6) a tuple of variables state (cf. Alg. 2), which altogether describe DynamicHS's current state. It outputs the ld (if existent) minimal diagnoses of maximal probability wrt. pr for the DPI $\langle \mathcal{K}, \mathcal{B}, P \cup P', N \cup N' \rangle$.

Embedding in Sequential Diagnosis Process. Alg. 1 sketches a generic sequential diagnosis algorithm and shows how it accommodates DynamicHS (line 5) or, alternatively, Reiter's HS-Tree (line 7), as methods for iterative diagnosis computation. The algorithm iterates in a while-loop (line 3) until the solution space of minimal diagnoses includes only a single element. This is the case iff a diagnosis set **D** with $|\mathbf{D}| = 1$ is output (line 8) since both DynamicHS and HS-Tree are complete and always attempt to compute at least two diagnoses ($ld \ge 2$). On the other hand, as long as $|\mathbf{D}| > 1$, the algorithm seeks to acquire additional information to rule out further elements in **D**. To this end, the best next measurement point mp is computed (line 9), using the current system information— dpi_0 , **D**, and acquired measurements P', N'—as well as the given probabilistic information pr and some measurement selection heuristic heur (which defines what "best" means, cf. [35]). The conduction of the measurement at mp (line 10) is usually accomplished by a qualified user (oracle) that interacts with the sequential diagnosis system, e.g., an electrical engineer for a defective circuit, or a domain expert in case of a faulty ontology. The measurement point mp along with its result *outcome* are used to formulate a logical sentence m that is either added to P' if m constitutes a positive measurement, and to N' otherwise (line 11). Finally, if DynamicHS is adopted, the set of diagnoses **D** is partitioned into those consistent $(\mathbf{D}_{\checkmark})$ and those inconsistent (\mathbf{D}_{\times}) with the newly added measurement m (line 13). Reiter's HS-Tree. DynamicHS inherits many of its aspects from Reiter's HS-Tree. Hence, we first recapitulate HS-Tree and then focus on the differences to and idiosyncrasies of DynamicHS.

HS-Tree computes minimal diagnoses for $dpi = \langle \mathcal{K}, \mathcal{B}, P \cup \rangle$

875

 $P', N \cup N'\rangle$ in a sound, complete³ and best-first way. Starting from a priority queue of unlabeled nodes \mathbf{Q} , initially comprising only an unlabeled root node, the algorithm continues to remove and label the first ranked node from \mathbf{Q} (GETANDDELETEFIRST) until all nodes are labeled ($\mathbf{Q} = []$) or ld minimal diagnoses have been computed. The possible node labels are minimal conflicts (for internal tree nodes) and *valid* as well as *closed* (for leaf nodes). All minimal conflicts that have already been computed and used as node labels are stored in the (initially empty) set \mathbf{C}_{calc} . Each edge in the constructed tree has a label. For ease of notation, each tree node nd is conceived of as the set of edge labels along the branch from the root node to itself. E.g., the node at location (\mathbf{D}) in iteration 1 of Fig. 2 is referred to as $\{1, 2, 5\}$. Once the tree has been completed ($\mathbf{Q} = []$), i.e., there are no more unlabeled nodes, it holds that $\mathbf{diag}(dpi) = \{\mathbf{nd} \mid \mathbf{nd} \text{ is labeled by valid}\}.$

To label a node nd, the algorithm calls a labeling function (LABEL) which executes the following tests in the given order and returns as soon as a label for nd has been determined:

- (L1) (*non-minimality*): Check if nd is non-minimal (i.e. whether there is a node n with label *valid* where $nd \supseteq n$). If so, nd is labeled by *closed*.
- (L2) (*duplicate*): Check if nd is duplicate (i.e. whether nd = n for some other n in Q). If so, nd is labeled by *closed*.
- (L3) *(reuse label):* Scan \mathbf{C}_{calc} for some \mathcal{C} such that $\mathsf{nd} \cap \mathcal{C} = \emptyset$. If so, nd is labeled by \mathcal{C} .
- (L4) (compute label): Invoke FINDMINCONFLICT, a (sound and complete) minimal conflict searcher, e.g., QuickXPlain [22], to get a minimal conflict for $\langle \mathcal{K} \setminus \mathsf{nd}, \mathcal{B}, P \cup P', N \cup N' \rangle$. If a minimal conflict \mathcal{C} is output, nd is labeled by \mathcal{C} . Otherwise, if 'no conflict' is returned, then nd is labeled by valid.

All nodes labeled by *closed* or *valid* have no successors and are leaf nodes. For each node nd labeled by a minimal conflict L, one outgoing edge is constructed for each element $e \in L$, where this edge is labeled by e and pointing to a newly created unlabeled node nd $\cup \{e\}$. Each new node is added to \mathbf{Q} such that \mathbf{Q} 's sorting is preserved (INSERTSORTED). \mathbf{Q} might be, e.g., (*i*) a FIFO queue, entailing that HS-Tree computes diagnoses in minimum-cardinality-first order (*breadth-first search*), or (*ii*) sorted in descending order by pr, where most probable diagnoses are generated first (*uniform-cost search*; for details see [33, Sec. 4.6]).

Finally, note the *statelessness* of Reiter's HS-Tree, reflected by \mathbf{Q} initially including *only an unlabeled root node*, and \mathbf{C}_{calc} being initially *empty*. That is, a HS-Tree is built from scratch in each iteration, every time for different measurement sets P', N'.

Dynamicity of DPI in Sequential Diagnosis. In the course of sequential diagnosis (Alg. 1), where additional system knowledge is gathered in terms of measurements, the DPI is subject to gradual change—it is dynamic. At this, each addition of a new (informative) measurement to the DPI also effectuates a transition of the sets of (minimal) diagnoses and (minimal) conflicts. Whereas this fact is of no concern to a stateless diagnosis computation strategy, it has to be carefully taken into account when engineering a stateful approach.

Towards Stateful Hitting Set Computation. To understand the necessary design decisions to devise a sound and complete stateful hitting set search, we look at more specifics of the conflicts and diagnoses evolution in sequential diagnosis:⁴

Property 1. Let $dpi_i = \langle \mathcal{K}, \mathcal{B}, P, N \rangle$ be a DPI and let T be Reiter's

HS-Tree for dpi_j (executed until) producing the diagnoses **D** where $|\mathbf{D}| \geq 2$. Further, let dpi_{j+1} be the DPI resulting from dpi_j by adding an informative⁵ measurement m to either P or N. Then:

- T is not a correct HS-Tree for dpi_{j+1}, i.e., (at least) some node labeled by valid in T is incorrectly labeled. That is, to reuse T for dpi_{j+1}, T must be updated.
- 2. For all $\mathcal{D}' \in \operatorname{diag}(dpi_{j+1})$ there is some $\mathcal{D} \in \operatorname{diag}(dpi_j)$ such that $\mathcal{D}' \supseteq \mathcal{D}$.

That is, minimal diagnoses can grow or remain unchanged, but cannot shrink. Consequently, to reuse T for sound and complete minimal diagnosis computation for dpi_{j+1} , existing nodes must never be reduced—either a node is left as is, deleted as a whole, or (prepared to be) extended.

3. For all $C \in \operatorname{conf}(dpi_j)$ there is some $C' \in \operatorname{conf}(dpi_{j+1})$ such that $C' \subseteq C$.

That is, existing minimal conflicts can only shrink or remain unaffected, but cannot grow. Hence, priorly computed minimal conflicts (for an old DPI) might not be minimal for the current DPI. In other words, node-labeling conflicts in T can, but do not need to be, correct for dpi_{i+1} .

4. (a) There is some $\mathcal{C} \in \mathbf{conf}(dpi_j)$ for which there is a $\mathcal{C}' \in \mathbf{conf}(dpi_{j+1})$ with $\mathcal{C}' \subset \mathcal{C}$, and/or

(b) there is some $C_{new} \in \mathbf{conf}(dpi_{j+1})$ where $C_{new} \not\subseteq C$ and $C_{new} \not\supseteq C$ for all $C \in \mathbf{conf}(dpi_{j})$.

That is, (a) some minimal conflict is reduced in size, and/or (b) some entirely new minimal conflict (not in any subsetrelationship with existing ones) arises. Some existing node in T which represents a minimal diagnosis for $dpi_j(a)$ can be deleted since it would not be present when using C' as node label in Twherever C is used, or (b) must be extended to constitute a diagnosis for dpi_{i+1} , since it does not hit C_{new} .

Major Modifications to Reiter's HS-Tree. Based on Property 1, the following principal amendments to Reiter's HS-Tree are necessary to make it a properly-working stateful diagnosis computation method:

(Mod1) Non-minimal diagnoses (test (L1) in HS-Tree) and duplicate nodes (test (L2)) are stored in collections D_{\supset} and Q_{dup} , respectively, instead of being closed and discarded.

Justification: Property 1.2 suggests to store also non-minimal diagnoses, as they might be (sub-branches of) minimal diagnoses in the subsequent iteration. Further, Property 1.4(a) suggests to record all duplicates for completeness of the diagnosis search. Because, some node nd representing this duplicate in the current tree could become obsolete due to the shrinkage of some conflict, and the duplicate might be non-obsolete and eligible to replace nd in the tree.

(Mod2) Each node nd is no longer identified with the *set* of edge labels along its branch, but as an *ordered list* of these edge labels. In addition, an ordered list of the conflicts used to label internal nodes along this branch is stored in terms of nd.cs. E.g., for node nd at location O in iteration 1 of Fig. 1, nd = [2, 5] and nd.cs = [$\langle 1, 2 \rangle$, $\langle 1, 3, 5 \rangle$].

Justification: (Property 1.4) The reason for storing both the edge labels and the internal node labels as lists lies in the replacement of obsolete tree branches by stored duplicates. In fact, any duplicate used to replace a node must correspond to the same *set* of edge labels as the replaced node. However, in the branch of the obsolete node, some node-labeling conflict has been reduced to make the node redundant, whereas for a suitable duplicate replacing the node, no redundancy-

³Unlike Reiter, we assume that only *minimal* conflicts are used as node labels. Thus, the issue pointed out by [18] does not arise.

⁴See [33, Sec. 12.4] for a more formal treatment and proofs.

⁵That is, adding m to the (positive or negative) measurements of the DPI effectuates an invalidation of some diagnosis in **D**.

Algorithm 2 DynamicHS

Input: tuple $\langle dpi, P', N', pr, ld, \mathbf{D}_{\checkmark}, \mathbf{D}_{\times}, \mathsf{state} \rangle$ comprising

- a DPI $dpi = \langle \mathcal{K}, \mathcal{B}, P, N \rangle$
- the acquired sets of positive (P') and negative (N') measurements so far
- a function pr assigning a fault probability to each element in \mathcal{K}
- the number *ld* of leading minimal diagnoses to be computed
- the set \mathbf{D}_{\checkmark} of all elements of the set \mathbf{D}_{calc} (returned by the previous DYNA-MICHS run) which are minimal diagnoses wrt. $\langle \mathcal{K}, \mathcal{B}, P \cup P', N \cup N' \rangle$
- the set \mathbf{D}_{\times} of all elements of the set \mathbf{D}_{calc} (returned by the previous DYNA-MICHS run) which are no diagnoses wrt. $\langle \mathcal{K}, \mathcal{B}, P \cup P', N \cup N' \rangle$
- state = $\langle \mathbf{Q}, \mathbf{Q}_{dup}, \mathbf{D}_{\supset}, \mathbf{C}_{calc} \rangle$ where
- **Q** is the current queue of unlabeled nodes,
- \mathbf{Q}_{dup} is the current queue of duplicate nodes,
- \mathbf{D}_{\supset} is the current set of computed non-minimal diagnoses,
- \mathbf{C}_{calc} is the current set of computed minimal conflict sets.

Output: tuple $\langle \mathbf{D}, \mathsf{state} \rangle$ where

- D is the set of the *ld* (if existent) most probable (as per *pr*) minimal diagnoses wrt. (K, B, P ∪ P', N ∪ N')
- state is as described above

1: procedure DYNAMICHS($dpi, P', N', pr, ld, \mathbf{D}_{\checkmark}, \mathbf{D}_{\times}, \mathsf{state}$)

 $\mathbf{D}_{\mathit{calc}} \gets \emptyset$ state \leftarrow UPDATETREE $(dpi, P', N', pr, \mathbf{D}_{\checkmark}, \mathbf{D}_{\times}, \text{state})$ 3 while $\mathbf{Q} \neq [] \land (|\mathbf{D}_{calc}| < ld)$ do 4 $\mathsf{node} \leftarrow \mathsf{GETANDDELETEFIRST}(\mathbf{Q})$ ▷ node is processed 5 if node $\in \mathbf{D}_\checkmark$ then $\triangleright \mathbf{D}_{\checkmark}$ includes only min... \triangleright ...diags wrt. current DPI 6 7: $L \leftarrow valid$ 8: 9: else $\langle L, \mathsf{state} \rangle \leftarrow \mathsf{DLABEL}(dpi, P', N', pr, \mathsf{node}, \mathbf{D}_{calc}, \mathsf{state})$ if L = valid then $\mathbf{D}_{calc} \leftarrow \mathbf{D}_{calc} \cup \{\mathsf{node}\} \triangleright \mathsf{node}$ is a min diag wrt. current DPI 10: 11: else if L = nonmin then $\mathbf{D}_{\supset} \leftarrow \mathbf{D}_{\supset} \cup \{\text{node}\}$ 12: 13: ▷ node is a non-min diag wrt. current DPI else for $e \in L$ do $\triangleright L$ is a min conflict wrt. current DPI node_e \leftarrow APPEND(node, e) \triangleright node_e is generated 14: 15: 16: 17 $\mathsf{node}_e.\mathsf{cs} \leftarrow \mathsf{APPEND}(\mathsf{node.cs}, L)$ 18: if $\mathsf{node}_e \in \mathbf{Q} \lor \mathsf{node}_e \in \mathbf{D}_{\supset}$ then 19: \triangleright node_e is (*set-equal*) duplicate of some node in **Q** or **D**_{\supset} 20: $\mathbf{Q}_{dup} \leftarrow \text{INSERTSORTED}(\mathsf{node}_e, \mathbf{Q}_{dup}, card, <)$ $\begin{array}{l} \textbf{else} \\ \mathbf{Q} \leftarrow \texttt{INSERTSORTED}(\mathsf{node}_e, \mathbf{Q}, pr, >) \end{array}$ 21: 22 23: return $\langle \mathbf{D}_{calc}, \mathsf{state} \rangle$ 24: procedure DLABEL($\langle \mathcal{K}, \mathcal{B}, P, N \rangle, P', N', pr, \mathsf{node}, \mathbf{D}_{calc}, \mathsf{state}$) $\begin{array}{l} \text{for nd} \in \mathbf{D}_{calc} \text{ do} \\ \text{ if node } \supset \text{ nd then} \end{array}$ 25 26: ▷ node is a non-min diag 27 **return** (*nonmin*, state) $\triangleright \mathbf{C}_{calc}$ includes conflicts wrt. current DPI 28: for $\mathcal{C} \in \mathbf{C}_{calc}$ do $\text{ if } \mathcal{C} \cap \mathsf{node} = \emptyset \text{ then }$ $\bigcap \text{ node } = \emptyset \text{ then } \qquad \triangleright \text{ reuse (a subset of } \mathcal{C} \text{ to label node } \\ X \leftarrow \text{FINDMINCONFLICT}(\langle \mathcal{C}, \mathcal{B}, P \cup P', N \cup N' \rangle)$ 29 30: if X = C then return $\langle C, \text{state} \rangle$ 31: 32: 33: else $\triangleright X \subset \mathcal{C}$ state $\leftarrow \text{PRUNE}(X, \text{state}, pr)$ 34: 35: **return** $\langle X, state \rangle$ $L \leftarrow \text{FINDMINCONFLICT}(\langle \mathcal{K} \setminus \text{node}, \mathcal{B}, P \cup P', N \cup N' \rangle)$ 36: 37: if L = 'no conflict' then ⊳ node is a diag return (valid, state) 38: $\triangleright L$ is a *new* min conflict ($\notin \mathbf{C}_{calc}$) 39: else $\mathbf{C}_{calc} \leftarrow \mathbf{C}_{calc} \cup \{L\}$ 40: **return** $\langle L, \text{state} \rangle$ 41: 42: procedure UPDATETREE($dpi, P', N', pr, \mathbf{D}_{\checkmark}, \mathbf{D}_{\times}, \mathsf{state}$) for $nd \in D_{\times}$ do ▷ search for redundant nodes among invalidated diags 43: if REDUNDANT(nd, dpi) then 44: 45. $\mathsf{state} \leftarrow \mathsf{PRUNE}(X, \mathsf{state}, pr)$ \triangleright add all (non-pruned) nodes in \mathbf{D}_{\times} to \mathbf{Q} for $\mathsf{nd} \in \mathbf{D}_{\times}$ do 46 $\mathbf{Q} \leftarrow \texttt{INSERTSORTED}(\mathsf{nd}, \mathbf{Q}, pr, >)$ 47. $\mathbf{D}_{\times} \leftarrow \mathbf{D}_{\times} \setminus \{\mathsf{nd}\}$ 48 49 for $\mathsf{nd} \in \mathbf{D}_{\supset}$ do \triangleright add all (non-pruned) nodes in \mathbf{D}_{\supset} to \mathbf{Q} , which... 50: $nonmin \leftarrow false$ \triangleright ... are no longer supersets of any diag in \mathbf{D}_{\checkmark} for $\mathsf{nd}' \in \mathbf{D}_\checkmark$ do 51: if $nd \supset nd'$ then $nonmin \leftarrow true$ break 52: 53: 54: 55: if nonmin = false then 56 $\mathbf{Q} \leftarrow \text{INSERTSORTED}(\mathsf{nd}, \mathbf{Q}, pr, >)$ $\mathbf{D}_{\supset} \leftarrow \mathbf{D}_{\supset} \setminus \{\mathsf{nd}\}$ 57 \triangleright add known min diags in \mathbf{D}_{\checkmark} to \mathbf{Q} to find diags... 58 for $\mathcal{D} \in \mathbf{D}_{\checkmark}$ do $\mathbf{Q} \leftarrow \text{INSERTSORTED}(\mathcal{D}, \mathbf{Q}, pr, >) \mathrel{\triangleright} \dots \text{in best-first order (as per pr)}$ 59 return state 60

causing changes to conflicts along its branch have occurred. By storing only sets of edge labels, we could not differentiate between the redundant and the non-redundant nodes.

(Mod3) Before reusing a conflict C (test (L3) in HS-Tree), a minimality check for C is performed. If it leads to the identification of a conflict $X \subset C$ for the current DPI, X is used to prune obsolete tree branches, to replace node-labeling conflicts that are supersets of X, and to update C_{calc} in that X is added and its supersets are deleted. *Justification:* (Property 1.3) Conflicts in C_{calc} and those appearing as labels in the existing tree (elements of lists nd.cs for tree nodes nd) might not be minimal for the current DPI (as they might have been computed for a prior DPI). This minimality check helps both to prune the tree (reduction of number of nodes) and to ensure that extensions to the tree use only minimal conflicts wrt. the current DPI as node labels (avoidance of the introduction of unnecessary new edges).

(**Mod4**) Execution of a tree update at start of each DynamicHS execution, where the tree produced for a previous DPI is adapted to a tree that allows to compute minimal diagnoses for the current DPI in a sound, complete and best-first way.

Justification: Property 1.1.

(**Mod5**) State of DynamicHS (in terms of the so-far produced tree) is stored over all its iterations executed throughout sequential diagnosis (Alg. 1) by means of the tuple state.

Justification: Statefulness of DynamicHS.

DynamicHS: Algorithm Walkthrough. Like HS-Tree, DynamicHS (Alg. 2) is processing a priority queue \mathbf{Q} of nodes (while-loop). In each iteration, the top-ranked node node is removed from \mathbf{Q} to be labeled (GETANDDELETEFIRST). Before calling the labeling function (DLABEL), however, the algorithm checks if node is among the known minimal diagnoses \mathbf{D}_{\checkmark} from the previous iteration (line 6). If so, the node is directly labeled *valid* (line 7). Otherwise DLABEL is invoked to compute a label for node (line 9).

DLABEL: First, the non-minimality check is performed (lines 25-27), just as in (L1) in HS-Tree. If negative, a conflict-reuse check is carried out (lines 28-35). Note, the duplicate check ((L2) in HS-Tree) is obsolete since no duplicate nodes can ever be elements of Q in DynamicHS (duplicates are identified and added to \mathbf{Q}_{dup} at node generation time, lines 18–20). The conflict-reuse check starts equally as in HS-Tree. However, if a conflict C for reuse is found in C_{calc} (line 29), then the minimality of C wrt. the *current* DPI is checked using FINDMINCONFLICT (line 30). If a conflict $X \subset C$ is detected (line 33), then X is used to prune the current hitting set tree (line 34; PRUNE function, see below). Finally, the found minimal conflict (C or X, depending on minimality check) is used to label node (lines 32, 35). The case where there is no conflict for reuse is handled just as in HS-Tree (lines 36-41, cf. (L4)). Finally, note that DLABEL gets and returns the tuple state (current tree state) as an argument, since the potentially performed pruning actions (line 34) might modify state.

The output of DLABEL is then processed by DynamicHS (lines 10–23) Specifically, node is assigned to \mathbf{D}_{calc} if the returned label is *valid* (line 11), and to \mathbf{D}_{\supset} if the label is *nonmin* (line 13). If the label is a minimal conflict L, then a child node node_e is created for each element $e \in L$ and assigned to either \mathbf{Q}_{dup} (line 20) if there is a node in \mathbf{Q} that is *set*-equal to node_e, or to \mathbf{Q} otherwise (line 22). At this, node_e is constructed from node via the APPEND function (lines 16 and 17), which appends the element e to the list node, and the conflict L to the list node.cs (cf. (Mod2) above).

When the hitting set tree has been completed ($\mathbf{Q} = []$), or ld diagnoses have been found ($|\mathbf{D}_{calc}| = ld$), DynamicHS returns \mathbf{D}_{calc} along with the current tree state state.

UPDATETREE: The goal is to adapt the existing tree in a way

it constitutes a basis for finding all and only minimal diagnoses in highest-probability-first order for the *current* DPI. The strategy is to search for non-minimal conflicts to be updated, and tree branches to be pruned, among the minimal diagnoses (\mathbf{D}_{\times}) for the previous DPI that have been invalidated by the latest added measurement.

Regarding tree pruning, we call a node nd *redundant* (wrt. a DPI dpi) iff there is some index j and a minimal conflict X wrt. dpi such that the conflict $nd.cs[j] \supset X$ and the element $nd[j] \in nd.cs[j] \setminus X$. Moreover, we call X a *witness of redundancy for* nd (wrt. dpi). Simply put, nd is redundant iff it would not exist given that the (current) minimal conflict X had been used as a label instead of the (old, formerly minimal, but by now) non-minimal conflict nd.cs[j].

If a redundant node is detected among the elements of \mathbf{D}_{\times} (function REDUNDANT), then PRUNE (see below) is called given the witness of redundancy of the redundant node as an argument (lines 43–45). After each node in \mathbf{D}_{\times} has been processed, the remaining nodes in \mathbf{D}_{\times} (those that are non-redundant and thus have not been pruned) are re-added to \mathbf{Q} in prioritized order (INSERTSORTED) according to pr (lines 46–48). Likewise, all non-pruned nodes in \mathbf{D}_{\supset} (note that pruning always considers all node collections \mathbf{Q}_{dup} , \mathbf{Q} , \mathbf{D}_{\checkmark} , \mathbf{D}_{\times} and \mathbf{D}_{\supset}) which are no longer supersets of any known minimal diagnosis, are added to \mathbf{Q} again (lines 49–57). Finally, those minimal diagnoses returned in the previous iteration and consistent with the latest added measurement (the elements of \mathbf{D}_{\checkmark}), are put back to the ordered queue \mathbf{Q} (lines 58–59). This is necessary to preserve the best-first property, as there might be "new" minimal diagnoses.

PRUNE: Using its given argument X, the tree pruning runs through all (active and duplicate) nodes of the current tree (node collections \mathbf{Q}_{dup} , \mathbf{Q} , \mathbf{D}_{\supset} and \mathbf{D}_{calc} for call in line 34, and \mathbf{Q}_{dup} , \mathbf{Q} , \mathbf{D}_{\supset} , \mathbf{D}_{\times} and \mathbf{D}_{\checkmark} for call in line 45), and

- (relabeling of old conflicts) replaces by X all labels nd.cs[i] which are proper supersets of X for all nodes nd and for all i = 1,..., |nd.cs|, and
- (*deletion of redundant nodes*) deletes each redundant node nd for which X is a witness of redundancy, and
- (potential replacement of deleted nodes) for each of the deleted nodes nd, if available, uses a suitable (non-redundant) node nd' (constructed) from the elements of Q_{dup} to replace nd by nd'.

A node nd' qualifies as a *replacement node for* nd iff nd' is nonredundant and nd' is *set*-equal to nd (i.e., the *sets*, not lists, of edge labels are equal). This node replacement is necessary from the point of view of completeness (cf. [18]). Importantly, duplicates (\mathbf{Q}_{dup}) must be pruned prior to all other nodes ($\mathbf{Q}, \mathbf{D}_{\supset}, \mathbf{D}_{\times}, \mathbf{D}_{\checkmark}, \mathbf{D}_{calc}$), to guarantee that all surviving nodes in \mathbf{Q}_{dup} represent possible *nonredundant* replacement nodes at the time other nodes are pruned.

Additionally, the argument X is used to update the conflicts stored for reuse (set C_{calc}), i.e., all supersets of X are removed from C_{calc} and X is added to C_{calc} .

Table 1: Example DPI stated in propositional logic.

\mathcal{K} =	$ \begin{cases} ax_1 : A \to \neg B & ax_2 \\ ax_4 : B \to C & ax_5 \end{cases} $	$ \begin{array}{l} : A \to B ax_3 : A \to \neg C \\ : A \to B \lor C \end{array} \} $
$\mathcal{B}=\emptyset$	P = 0	$N = \{\neg A\}$

Example 1 Consider the propositional DPI dpi_0 in Tab. 1. The goal is to locate the faulty axioms in the KB \mathcal{K} that prevent the satisfaction of given measurements P and N (here, only one negative measurement $\neg A$ is given, i.e., $\neg A$ must not be entailed by the correct KB).

Table 2: Evolution of minimal diagnoses and minimal conflicts after successive extension of the example DPI dpi_0 (Tab. 1) by positive (P') or negative (N') measurements m_i shown in Figures 1 and 2. Newly arisen conflicts (cf. C_{new} in Property 1.4) are framed.

iteration j	P'	N'		$\mathbf{diag}(dpi_{j-1})$	$conf(dpi_{j-1})$
1	-	-		[1, 3], [1, 4], [2, 3], [2, 5]	$\left \begin{array}{c} \langle 1,2 \rangle , \langle 2,3,4 \rangle , \\ \langle 1,3,5 \rangle , \langle 3,4,5 \rangle \end{array} \right $
2	-	$A \rightarrow C$		[1, 4], [2, 5]	$\begin{array}{c} \langle 1,2\rangle, \langle 2,4\rangle, \\ \langle 1,5\rangle, \langle 4,5\rangle \end{array}$
3	-	$A \rightarrow C, A -$	$\rightarrow \neg B$ [[1, 4], [1, 2, 3, 5]	$\langle 1 \rangle, \langle 2, 4 \rangle, \\ \langle 3, 4 \rangle, \langle 4, 5 \rangle$
4	$A \to \neg C$	$A \to C, A -$	$\rightarrow \neg B$	[1, 4]	$\langle 1 \rangle, \langle 4 \rangle$
		_		C	
	$\odot/2$	$1 \downarrow C \stackrel{1}{\leftarrow} (1)$	$\bigcup \langle 1, 2 \end{vmatrix}$	$\sim \frac{2}{2}$	$o \in C$
	$(2)\langle 2, 3\rangle$	$\begin{pmatrix} 4 \\ 4 \end{pmatrix}$			5, 5
(5)(3, 4)	$\downarrow 3$ $5\rangle^C$ (6)	າ √ " _{(D+1}) (7)√	(\mathcal{D}_{2})	(4) <i>dup</i> (8)√	າ ຊ (*) (໑√*)
3,/		$(\mathcal{D}_1) \bigcirc 5$	(D_2)	\bigcirc	$(D_3) \bigcirc (D_4)$
$(10) \times (\neg D)$	4 (1)	$\langle (\neg \mathcal{D}_2) \rangle$ (12	$\mathbb{D} \times (\neg \tau)$	24)	
0 (52	I) Uterat	(JU_2)	$\Rightarrow m_1$	$A \to C$ to	N′
		1 1	$\langle 1, 2 \rangle$	2	
	2 $\langle 2, \cdot \rangle$	$4\rangle$		\searrow $(1,$	$5\rangle_{5}$
<u> </u>	<i>1</i> /	<i>Y</i>			<i>Y</i>
(5)(3	$,4,5\rangle$	$(\mathcal{I}) \checkmark (\mathcal{D}_2)$		(4)dup	$(\underline{9}) \checkmark (\mathcal{D}_4)$
34			.		
(13(4,5))		$\langle (\supset \mathcal{D}_2) \rangle$ (12))×(⊃1	(2_4)	
4¥ ○ ×		,			
$(14) \times (\supset \mathcal{D})$	2) (15)	$(\supset \mathcal{D}_4)$			
	Iteratio	$ on 2 (dpi_1) = $	$\Rightarrow m_2$	$A \rightarrow \neg B$ to	> N'
'		1 (1	$\hat{\Gamma}(1)$		
	$(2)\langle 2, \cdot$	$_{4\rangle}$			
	2 /	4			
(5)(3)	$,4,5\rangle$	$\bigcirc \checkmark_{(\mathcal{D}_2)}$			
3	4	5			
(3)(4,5)	\rangle $(1) \rangle$	$(\supset \mathcal{D}_2)$	$\langle 3, 4 \rangle$	\rangle^{C}	
4	5		3	4	
$\mathrm{A}_{(\mathcal{D}\mathcal{D})}$	2) 18	$\checkmark^*_{(\mathcal{D}_5)}$	(1) dup	0 19×(⊃1	(D_2)
	Iterati	on $3(dpi_2) =$	$\Rightarrow m_3$	$: A \to \neg C$ to	$\circ P'$
	A	(1) = 1	$(\underline{I})\langle 1\rangle$		
	Ð	4			
		× (7)√ (7))			
		Iterat	tion 4 (a	lpi ₃)	

Figure 1: DynamicHS executed on example DPI given in Tab. 1.

We now illustrate the workings of DynamicHS (Fig. 1) and HS-Tree (Fig. 2) in terms of a sequential diagnosis session for dpi_0 while assuming that [1, 4] is the actual diagnosis. The evolution of minimal conflicts and diagnoses during the session can be read from Tab. 2.

Inputs (Sequential Diagnosis): We set ld := 5 (if existent, compute five diagnoses per iteration), heur to be the "split-in-half" heuristic [44] (prefers measurements the more, the more diagnoses they eliminate in the worst case), and pr in a way the hitting set trees are constructed breadth-first.

Notation in Figures: Axioms ax_i are simply referred to by i (in node



Figure 2: Part of the HS-Tree execution on example DPI in Tab. 1.

and edge labels). Numbers (a) indicate the chronological node labeling order. We tag conflicts $\langle 1, \ldots, k \rangle$ by C if they are freshly computed (FINDMINCONFLICT call, line 36, DLABEL), and leave them untagged if they result from a redundancy check and subsequent relabeling (lines 44–45, UPDATETREE). Leaf nodes we label as follows: $\sqrt{(D_i)}$ for a minimal diagnosis, named D_i , stored in \mathbf{D}_{calc} ; × for a duplicate in HS-Tree (see (L2) criterion); $\times_{(\supset D_i)}$ for a non-minimal diagnosis (stored in \mathbf{D}_{\supset} by DynamicHS), where D_i is one diagnosis that proves the non-minimality; and *dup* for a duplicate in DynamicHS (stored in \mathbf{Q}_{dup}). Branches representing minimal diagnoses are additionally tagged by a * if logical reasoning (FINDMINCONFLICT call, line 36, DLABEL function) is necessary to prove it is a diagnosis, and untagged otherwise (i.e., branch is diagnosis from previous iteration, stored in \mathbf{D}_{\checkmark} ; only applicable to DynamicHS).

Iteration 1: In the first iteration, HS-Tree and DynamicHS essentially build the same tree (see "Iteration 1" in Fig. 1; HS-tree not shown). The only difference is that DynamicHS stores the duplicates and non-minimal diagnoses (labels dup and $\times_{(\supset D_i)}$), whereas HS-Tree discards them. The diagnoses computed by both algorithms are $\{D_1, D_2, D_3, D_4\} = \{[1, 3], [1, 4], [2, 3], [2, 5]\}$ (cf. Tab. 2). Notably, the returned diagnoses *must* be equal for both algorithms in any iteration (given the same parameters ld and pr) since both are sound, complete and best-first minimal diagnosis searches. Thus, when using the same measurement selection (and heuristic heur), both methods *must* also give rise to the same proposed next measurement point in each iteration.

First Measurement: Accordingly, both algorithms lead to the first measurement point $mp_1 : A \to C$, which corresponds to the question about the correct KB "Does A imply C?". Say mp_1 turns out negative, e.g., by consulting a domain expert for \mathcal{K} , and is therefore added to N'. This effectuates a transition from dpi_0 to a new DPI dpi_1 (with the additional element $A \to C$ in N'), and thus a change of the relevant minimal diagnoses and conflicts (see Tab. 2).

Tree Update: Starting from the second iteration (dpi_1) , HS-Tree and DynamicHS behave differently. Whereas the former constructs a new tree from scratch for dpi_1 , DynamicHS runs UPDATETREE to make the existing tree (built for dpi_0) reusable for dpi_1 . In the course of this tree update, two witnesses of redundancy (minimal conflicts $\langle 2, 4 \rangle$, $\langle 1, 5 \rangle$) are found while analyzing the (conflicts along the) branches of the two invalidated diagnoses [1, 3] and [2, 3] (6 and (8)). E.g., nd = [1, 3] is redundant since the conflict nd.cs[2] = $\langle 2, 3, 4 \rangle$ is a proper superset of the *now minimal* conflict $X = \langle 2, 4 \rangle$ and nd's outgoing edge of nd.cs[2] is nd[2] = 3 which is an element of nd.cs[2] \ $X = \{3\}$. Since stored duplicates (here: only [2, 1]) do not allow the construction of a replacement node for any of the redundant branches [1, 3] and [2, 3], both are removed from the tree. Further, the witnesses of redundancy replace the non-minimal conflicts at (2) and (3), signified by the new numbers (2) and (3).

Other than that, only a single further change is induced by UPDA-TETREE. Namely, the branch [1, 2, 3], a non-minimal diagnosis for dpi_0 , is returned to **Q** (unlabeled nodes) because there is no longer a diagnosis in the tree witnessing its non-minimality (both such witnesses [1, 3] and [2, 3] have been discarded). Note that, first, [1, 2, 3] is in fact no longer a hitting set of all minimal conflicts for dpi_1 (cf. Tab. 2) and, second, there is still a non-minimality witness for all other branches ((\underline{n}) and (\underline{n})) representing non-minimal diagnoses for dpi_0 , which is why they remain labeled by $\times_{(\supset D_i)}$.

Iteration 2: Observe the crucial differences between HS-Tree and DynamicHS in the second iteration (cf. Figs. 1 and 2).

First, while HS-Tree has to compute all conflicts that label nodes by (potentially expensive) FINDMINCONFLICT calls (C tags), DynamicHS has (cheaply) reduced existing conflicts during pruning (see above). Note, however, not all conflicts are necessarily always kept up-to-date after a DPI-transition (lazy updating policy). E.g., node (5) is still labeled by the non-minimal conflict (3, 4, 5) after UPDA-TETREE terminates. Hence, the subtree comprising nodes (13), (14), (15) is not present in HS-Tree. Importantly, this lazy updating does not counteract sound- or completeness of DynamicHS, and the overhead incurred by such additional nodes is generally minor (all these nodes must be non-minimal diagnoses and are thus not further expanded). Second, the verification of the minimal diagnoses $(\mathcal{D}_2, \mathcal{D}_4)$ found in iteration 2 requires logical reasoning in HS-Tree (* tags of \checkmark nodes) whereas it comes for free in DynamicHS (storage and reuse of D_{\checkmark}). *Remaining Execution:* After the second measurement m_2 is added to N', causing a DPI-transition once again, DynamicHS reduces the conflict that labels the root node. This leads to the pruning of the complete right subtree. The left subtree is then further expanded in iteration 3 (see generated nodes (6), (17), (18) and (19) until the two leading diagnoses $\mathcal{D}_2 = [1, 4]$ and $\mathcal{D}_5 = [1, 2, 3, 5]$ are located and the queue of unlabeled nodes becomes empty (which proves that no further minimal diagnoses exist). Eventually, the addition of the third measurement m_3 to P' brings sufficient information to isolate the actual diagnosis. This is reflected by a pruning of all branches except for the one representing the actual diagnosis [1, 4].

Comparison (expensive operations): Generally, calling FINDMIN-CONFLICT (FC) is (clearly) more costly than REDUNDANT (RD), which in turn is more costly than a single consistency check (CC). In this example, HS-Tree requires 14/0/9, and DynamicHS only 6/4/5 FC/RD/CC calls. This reduction of costly reasoning is one crucial advantage of DynamicHS over HS-Tree (cf. Sec. 5).

DynamicHS: Properties. A proof of the following correctness theorem for DynamicHS can be found in [33, Sec. 12.4]:

Theorem 1. DynamicHS is sound, complete and best-first, i.e., it computes all and only minimal diagnoses in descending order by probability as per the probability measure pr.

4 Related Work

Literature offers a wide variety of diagnosis computation algorithms, motivated by different diagnosis problems, domains and challenges. These algorithms can be compared along multiple dimensions, e.g., *best-first* [8, 11, 18, 32, 33, 36] vs. *any-first* (no particular order on output diagnoses) [26, 31, 45], *complete* [11, 18, 19, 26, 32, 36, 45, 49, 54] vs. *incomplete* [1, 9, 12, 25, 47, 51], *conflict-based* (i.e., building diagnoses as hitting sets of conflicts) [10, 11, 18, 26, 32, 33, 36, 49, 52, 54] vs. *direct* (i.e., building diagnoses without reliance on conflicts, e.g., through divide-and-conquer or compilation techniques) [5, 14, 28, 45, 50], *stateful* [11, 33, 36, 45, 46] vs. *stateless* [10, 18, 32, 52, 54], *black-box* (reasoner used as oracle) [18, 32, 33, 36, 45, 52] vs. *glass-box* (reasoner

optimized or modified for diagnosis) [4, 20, 23, 43], and *on-the-fly* [18, 32, 36, 52] vs. *preliminary* [1, 10, 25, 26, 31, 54] (conflict computation). Our study of these existing works suggests two different things. First, the best choice of algorithm, in general, depends largely on the particular tackled problem (domain and requirements). Consequently, there is little hope to find an algorithm that comes anywhere near improving *all* of the existing ones. Second, performance improvements for algorithms are often achieved at the cost of losing desirable properties (e.g., completeness or the best-first guarantee). Hence, it is particularly noteworthy that DynamicHS aims to improve HS-Tree *while preserving all its favorable properties*.

Specifically, in terms of the above-mentioned properties, DynamicHS is best-first, complete, stateful, conflict-based, black-box, and on-the-fly. We now briefly discuss the algorithms most similar to DynamicHS in terms of these features and point out crucial differences:⁶ (i) StaticHS [36]: same features, but different (orthogonal) focus, which is on reducing measurement conduction time, instead of diagnosis computation time as in the case of DynamicHS (cf. Sec. 1). Technically, StaticHS extends one and the same "static" hitting set tree-the one for the original DPI-and uses measurement outcomes (at first) only to prune this static tree, while a new such static tree-for the new DPI including all so far collected measurement outcomes-is built only if a single diagnosis is left or the static tree exceeds a critical size. So, the intuitive idea is a lazy DPI update, i.e., to withhold the DPI update and exploit one static tree for multiple iterations. (ii) Inv-HS-Tree [45]: same features, but stateless, direct and any-first; the latter property prevents reasonable early termination⁷ in sequential diagnosis, since the best diagnosis is generally unknown; but, using a linear-space search, it can handle problems with high-cardinality diagnoses that HS-Tree (and thus DynamicHS) cannot. Algorithmically, Inv-HS-Tree leverages the duality between diagnoses and conflicts and thus builds a tree in depth-first fashion where nodes are minimal diagnoses and branches are conflicts. (iii) GDE [11]: same features, but not black-box (in that it uses bookkeeping [6] of entailment-justifications computed by the reasoner); while highly optimized and particularly well-performing for physical devices. GDE has not found widespread adoption in some other diagnosis domains such as knowledge-base debugging, most likely⁸ due to the typically richer and often highly non-local justification-structure⁹ of knowledge-based problems that might be challenging for bookkeeping approaches.

Table 3: Dataset used in the experiments.

j	$\operatorname{KB} \mathcal{K}_j$	$ \mathcal{K}_j $	expressivity 1)	#D/min/max ²⁾
1	Chemical (C)	144	$\mathcal{ALCHF}^{(D)}$	6/1/3
2	Koala (K)	42	$\mathcal{ALCON}^{(D)}$	10/1/3
3	MiniTambis (M)	173	\mathcal{ALCN}	48/3/3
4	University (U)	50	$SOIN^{(D)}$	90/3/4
5	Economy (E)	1781	$\mathcal{ALCH}^{(D)}$	864/4/8
6	Transportation (T)	1300	$\mathcal{ALCH}^{(D)}$	1782/6/9
7	IT	140	SROIQ	1045/3/7
8	UNI	142	SROIQ	1296/5/6

Description Logic expressivity, cf. [3]; the higher the expressivity of a logic, the higher is the complexity of reasoning for this logic.
 #D, win, max denote the number, the min, and max, size of minimal diag-

2): #D, min, max denote the number, the min. and max. size of minimal diagnoses for the initial DPI dpi_0 resulting from each input KB \mathcal{K}_j .

⁶We do not discuss HS-Tree again here, see Sec. 3.

⁷I.e., the stopping of a diagnosis session before a *single* diagnosis is left. ⁸Based on a plenary discussion at the Int'l Workshop on Principles of Diagnosis (DX'17) in Brescia, September 2017.

⁹Locality refers to the property that components interact (only) with physically adjacent components [7].



Figure 3: Experiment Results: x-axis shows KBs and parameter ld.

5 Evaluation

In comprehensive experiments on real-world diagnosis problems we have evaluated the practical usefulness of DynamicHS.

Dataset. As discussed, different domains typically require different approaches as regards the computation of diagnoses. Since DynamicHS aims to improve HS-Tree, we chose the application domain of knowledge-base debugging for our experiments, where HS-Tree is state-of-the-art [20, 23, 27, 33, 44]. As a dataset (Tab. 3), we used inconsistent Description Logic KBs (2nd column) that were also investigated in the studies [44] ($\mathcal{K}_1, \ldots, \mathcal{K}_6$) and [37] ($\mathcal{K}_7, \mathcal{K}_8$), and which cover a spectrum of different problem sizes (3rd col.), logical languages (4th col.), as well as diagnostic structures (last col.).

Experimental Setting. We compared DynamicHS against HS-Tree in a multitude of diagnosis scenarios, by varying the inputs to the sequential diagnosis algorithm (Alg. 1). More precisely, for each $dpi_0 := \langle \mathcal{K}_i, \emptyset, \emptyset, \emptyset \rangle$, where \mathcal{K}_i is a KB from Tab. 3, we (i) assigned uniform random failure probabilities to all components (KB axioms), and then (ii) for each number $ld \in \{2, 4, 6, 10, 20, 30\}$ of leading minimal diagnoses to be computed per sequential iteration (while-loop, Alg. 1), (iii) for each of three popular measure*ment selection heuristics* heur \in {SPL, ENT, MPS} in the field (cf. [34, 35, 39, 44]),¹⁰ (iv) we performed 20 full sequential sessions (until a single minimal diagnosis remained, cf. Alg. 1), (v) each with a randomly selected actual diagnosis \mathcal{D}^* specified as the target solution to be found, (vi) executed for both DynamicHS and HS-Tree. The oracle for measurement conduction (or question answering; cf. Ex. 1) was simulated so as to always answer in a way \mathcal{D}^* remains valid as a diagnosis. As a reasoner, we adopted Pellet [48].

Experiment Results.¹¹ The results for $\bar{\mathcal{K}}_1, \ldots, \mathcal{K}_6$ are shown in Fig. 3. We see that, in all (KB,ld,heur)-scenarios, DynamicHS leads to significant average time savings (of up to 70%) compared to HS-Tree. For the (non-depicted) KBs $\mathcal{K}_7, \mathcal{K}_8$, results were even more in favor of DynamicHS (efficiency gains between 37 and 54 % for \mathcal{K}_7 , and between 56 to 70 % for \mathcal{K}_8). Closer analyses of further measured parameters revealed that the main cause of these substantial and constant improvements is that the number of (and thus the time expended for) reasoner calls can be remarkably decreased by DynamicHS. As to the absolute runtimes, KBs ranged from easy (max. diagnosis computation times below 1s, KB \mathcal{K}_2) to fairly hard (max. times of 510s for HS-Tree vs. 221s for DynamicHS; KB \mathcal{K}_5). Finally, we stress two things: DynamicHS (1) outperformed HS-Tree in 98.75 % of all single sequential sessions we executed (cf. blue line in Fig. 3), and (2) has exactly the same desirable properties (cf. Theorem 1) as HS-Tree (and exhibited similar space requirements in our tests).

¹⁰In brief, SPL (*split-in-half*) / ENT (*entropy-based*) / MPS (*most probable singleton*) select a measurement with best worst-case diagnosis elimination rate / highest information gain / highest probability of eliminating all but one (known) diagnoses.

¹¹Please find the raw data at http://isbi.aau.at/ontodebug/evaluation.

6 Conclusion and Future Work

We propose a sound, complete and best-first diagnosis computation method for sequential diagnosis based on Reiter's HS-Tree. It aims at reducing expensive reasoning by the maintenance of a search data structure throughout a diagnostic session. Thereby, we answer a long-standing research question posed by Ray Reiter in his seminal 1987 paper [32]. Experimental results are very promising and manifest that the new method constantly and significantly outperforms HS-Tree in a domain where the latter is state-of-the-art.

Future work topics include the integration of DynamicHS into our ontology debugging tool *OntoDebug*¹² [42] for the popular ontology editor *Protégé*¹³ [29], and the evaluation of DynamicHS in other domains, such as software [53] and spreadsheet debugging [21].

Acknowledgments: Thanks to Manuel Herold for implementing DynamicHS and the experiments. This work was partly supported by the Austrian Science Fund (FWF), contract P-32445-N38.

REFERENCES

- R. Abreu and A.J. Van Gemund, 'A low-cost approximate minimal hitting set algorithm and its application to model-based diagnosis', in 8th Symposium on Abstraction, Reformulation, and Approximation, (2009).
- [2] R. Abreu, P. Zoeteweij, and A.J. Van Gemund, 'Simultaneous debugging of software faults', J. Syst. Softw., 84(4), 573–586, (2011).
- [3] The Description Logic Handbook, eds., F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, CUP, 2007.
- [4] F. Baader and R. Penaloza, 'Axiom pinpointing in general tableaux', J. Log. Comput., 20(1), 5–34, (2008).
- [5] A. Darwiche, 'Decomposable negation normal form', *J. ACM*, **48**(4), 608–647, (2001).
- [6] J. de Kleer, 'Problem solving with the ATMS', Artif. Intell., 28(2), 197– 224, (1986).
- [7] J. de Kleer, 'Compiling devices and processes', in 4th Int'l Workshop on Qualitative Physics, (1990).
- [8] J. de Kleer, 'Focusing on probable diagnoses', in AAAI'91, (1991).
- [9] J. de Kleer, 'Mininimum cardinality candidate generation', in *DX'09*, (2009).
- [10] J. de Kleer, 'Hitting set algorithms for model-based diagnosis', in DX'11, (2011).
- [11] J. de Kleer and B.C. Williams, 'Diagnosing multiple faults', Artif. Intell., 32(1), 97–130, (1987).
- [12] A. Feldman, G.M. Provan, and A.J. Van Gemund, 'Computing minimal diagnoses by greedy stochastic search', in AAAI'08, (2008).
- [13] A. Felfernig, G. Friedrich, D. Jannach, and M. Stumptner, 'Consistency-based diagnosis of configuration knowledge bases', *Artif. Intell.*, **152**(2), 213–234, (2004).
- [14] A. Felfernig, M. Schubert, and C. Zehentner, 'An efficient diagnosis algorithm for inconsistent constraint sets', AI EDAM, 26(1), 53–62, (2012).
- [15] A. Felfernig, E. Teppan, G. Friedrich, and K. Isak, 'Intelligent debugging and repair of utility constraint sets in knowledge-based recommender applications', in *IUI'08*, pp. 217–226, (2008).
- [16] G. Friedrich and K. Shchekotykhin, 'A General Diagnosis Method for Ontologies', in *ISWC'05*, pp. 232–246, (2005).
- [17] G. Friedrich, M. Stumptner, and F. Wotawa, 'Model-based diagnosis of hardware designs', *Artif. Intell.*, **111**(1-2), 3–39, (1999).
- [18] R. Greiner, B. Smith, and R. Wilkerson, 'A correction to the algorithm in Reiter's theory of diagnosis', *Artif. Intell.*, 41(1), 79–88, (1989).
- [19] R. Haenni, 'Generating diagnoses from conflict sets.', in *FLAIRS'98*, pp. 420–424, (1998).
- [20] M. Horridge, Justification based Explanation in Ontologies, Ph.D. dissertation, Univ. of Manchester, 2011.
- [21] D. Jannach, T. Schmitz, B. Hofer, and F. Wotawa, 'Avoiding, finding and fixing spreadsheet errors–a survey of automated approaches for spreadsheet QA', J. of Syst. Softw., 94, 129–150, (2014).

- [22] U. Junker, 'QuickXPlain: Preferred Explanations and Relaxations for Over-Constrained Problems', in AAAI'04, pp. 167–172, (2004).
- [23] A. Kalyanpur, Debugging and Repair of OWL Ontologies, Ph.D. dissertation, Univ. of Maryland, College Park, 2006.
- [24] J. de Kleer, A.K. Mackworth, and R. Reiter, 'Characterizing diagnoses and systems', *Artif. Intell.*, 56, (1992).
 [25] L. Li and J. Yunfei, 'Computing minimal hitting sets with genetic algo-
- [25] L. Li and J. Yunfei, 'Computing minimal hitting sets with genetic algorithm', in DX'02, (2002).
- [26] L. Lin and Y. Jiang, 'The computation of hitting sets: Review and new algorithms', *Inf. Process. Lett.*, 86(4), 177–184, (2003).
- [27] C. Meilicke, Alignment Incoherence in Ontology Matching, Ph.D. dissertation, Univ. Mannheim, 2011.
- [28] A. Metodi, R. Stern, M. Kalech, and M. Codish, 'A novel SAT-based approach to model based diagnosis', *JAIR*, **51**, 377–411, (2014).
- [29] N.F. Noy, M. Sintek, S. Decker, M. Crubézy, R.W. Fergerson, and M.A. Musen, 'Creating Semantic Web Contents with Protégé-2000', *IEEE Intell. Syst.*, 16(2), 60–71, (2000).
- [30] K.R. Pattipati and M.G. Alexandridis, 'Application of heuristic search and information theory to sequential fault diagnosis', *IEEE Trans. Syst. Man. Cybern.*, 20(4), 872–887, (1990).
- [31] I. Pill and T. Quaritsch, 'Optimizations for the Boolean approach to computing minimal hitting sets', in ECAI'12, pp. 648–653, (2012).
- [32] R. Reiter, 'A Theory of Diagnosis from First Principles', *Artif. Intell.*, 32(1), 57–95, (1987).
- [33] P. Rodler, Interactive Debugging of Knowledge Bases, Ph.D. dissertation, Univ. Klagenfurt, 2015. http://arxiv.org/pdf/1605.05950v1.pdf.
- [34] P. Rodler, 'Towards better response times and higher-quality queries in interactive knowledge base debugging', Tech. rep., Univ. Klagenfurt, (2016). http://arxiv.org/pdf/1609.02584v2.pdf.
- [35] P. Rodler, 'On active learning strategies for sequential diagnosis', in DX'17, (2018).
- [36] P. Rodler and M. Herold, 'StaticHS: A variant of Reiter's hitting set tree for efficient sequential diagnosis', in SoCS'18, (2018).
- [37] P. Rodler, D. Jannach, K. Schekotihin, and P. Fleiss, 'Are query-based ontology debuggers really helping knowledge engineers?', *Knowl.-Based Syst.*, **179**, 92–107, (2019).
- [38] P. Rodler and K. Schekotihin, 'Reducing model-based diagnosis to knowledge base debugging', in DX'17, pp. 284–296, (2018).
- [39] P. Rodler and W. Schmid, 'On the impact and proper use of heuristics in test-driven ontology debugging', in *RuleML+RR'18*, (2018).
- [40] P. Rodler, W. Schmid, and K. Schekotihin, 'A generally applicable, highly scalable measurement computation and optimization approach to sequential model-based diagnosis', *CoRR*, abs/1711.05508, (2017).
- [41] P. Rodler, K. Shchekotykhin, P. Fleiss, and G. Friedrich, 'RIO: Minimizing User Interaction in Ontology Debugging', in *RR'13*, (2013).
- [42] K. Schekotihin, P. Rodler, and W. Schmid, 'OntoDebug: Interactive ontology debugging plug-in for Protégé', in *FoIKS'18*, (2018).
- [43] S. Schlobach, Z. Huang, R. Cornet, and F. Van Harmelen, 'Debugging incoherent terminologies', J. Autom. Reason., 39(3), 317–349, (2007).
- [44] K. Shchekotykhin, G. Friedrich, P. Fleiss, and P. Rodler, 'Interactive Ontology Debugging: Two Query Strategies for Efficient Fault Localization', J. Web Semant., 12-13, 88–103, (2012).
- [45] K. Shchekotykhin, G. Friedrich, P. Rodler, and P. Fleiss, 'Sequential diagnosis of high cardinality faults in knowledge-bases by direct diagnosis generation', in *ECAI'14*, pp. 813–818, (2014).
- [46] S. Siddiqi and J. Huang, 'Sequential diagnosis by abstraction', JAIR, 41, 329–365, (2011).
- [47] S. Siddiqi, and J. Huang, 'Hierarchical diagnosis of multiple faults', in *IJCAI'07*, pp. 581–586, (2007).
- [48] E. Sirin, B. Parsia, B.C. Grau, A. Kalyanpur, and Y. Katz, 'Pellet: A practical OWL-DL reasoner', J. Web Semant., 5(2), 51–53, (2007).
- [49] R. Stern, M. Kalech, A. Feldman, and G. Provan, 'Exploring the duality in conflict-directed model-based diagnosis', in AAAI'12, (2012).
- [50] P. Torasso and G. Torta, 'Model-based diagnosis through OBDD compilation: A complexity analysis', in *Reasoning, Action and Interaction* in AI Theories and Systems, 287–305, (2006).
- [51] B. Williams and R. Ragno, 'Conflict-directed A* and its role in modelbased embedded systems', *Discrete Appl. Math.*, 155(12), (2007).
- [52] F. Wotawa, 'A variant of Reiter's hitting-set algorithm', Inf. Process. Lett., 79(1), 45–51, (2001).
- [53] F. Wotawa, 'Fault localization based on dynamic slicing and hitting-set computation', in *QSIC'10*, pp. 161–170, (2010).
- [54] Z. Xiangfu and O. Dantong, 'A method of combining SE-tree to compute all minimal hitting sets', *Prog. Nat. Sci.*, 16(2), 169–174, (2006).

¹²See http://isbi.aau.at/ontodebug.

¹³See https://protege.stanford.edu/.