SMALL STONE IN POOL

SAMUEL R. BUSS [a] AND LESZEK ALEKSANDER KOŁODZIEJCZYK [b]

[a] Department of Mathematics, University of California, San Diego, La Jolla, CA 92093-0112, USA
e-mail address: sbuss@math.ucsd.edu

[b] Institute of Mathematics, University of Warsaw, Banacha 2, 02-097 Warszawa, Poland
e-mail address: lak@mimuw.edu.pl

ABSTRACT. The Stone tautologies are known to have polynomial size resolution refutations
and require exponential size regular refutations. We prove that the Stone tautologies also
have polynomial size proofs in both pool resolution and the proof system of regular tree-like
resolution with input lemmas (regRTI). Therefore, the Stone tautologies do not separate
resolution from DPLL with clause learning.

> *I have said it thrice; What I tell you three times is true.*
> *Fit the First — The Landing; The Hunting of the Snark*
> Lewis Carroll

1. INTRODUCTION

The Davis-Putnam-Logemann-Loveland (DPLL) proof search method [DLL62, DP60], augmented with clause learning [MSS99], has become a core method for solving the satisfiability (SAT) problem, especially for large-scale instances of SAT that arise in industrial applications. However, when restarts are not allowed, the proof strength of DPLL with clause learning relative to full resolution remains unknown. On one hand, if Γ is a set of clauses, and DPLL with clause learning can show that Γ is unsatisfiable in n steps, then Γ has a resolution refutation with size polynomially bounded by n (see [BKS04]). On the other hand, the results of [AJPU07, Urq11, BB12, BBJ14] imply that the length of DPLL with clause learning proof searches can be nearly exponentially smaller than the length of the shortest regular resolution proofs. Systems designed to correspond to DPLL with clause learning, such as pool resolution ([VG05]) and regRTI ([BHJ08]), are therefore simulated

by resolution and strictly stronger than regular resolution. Determining the exact strength of those systems is an open problem.

The two papers [AJPU07, Urq11] gave examples of three principles which have polynomial size resolution refutations, but require exponential size regular refutations. In the terminology of [BB12, BBJ14], these three principles were (1) the guarded graph tautologies, (2) the Stone tautologies, and (3) the guarded pebbling tautologies. Subsequently, [BB12, BBJ14] showed that the guarded graph tautologies and the guarded pebbling tautologies have polynomial size pool and regRTI refutations, and hence can be refuted by polynomially long DPLL search with clause learning and without restarts.

It remained open whether the same holds for the Stone tautologies. There seems to be an inherent simplicity in the irregularity introduced by "guarded" versions of combinatorial principles, such as (1) and (3). This is because the guarded principles have refutations in which all irregularities are at the initial inferences; namely, the resolution refutation can start by using resolution to remove the guard literals, and then give a (regular) refutation of the underlying principle as usual. In contrast, the prior known resolution refutations for the Stone principles of [AJPU07] use irregularity in a more essential fashion, with the irregularities distributed throughout the refutation. Because of this, it was conjectured that the Stone principles might be examples where the pool and regRTI systems, and thus DPLL with clause learning and no restarts, require exponential size refutations. That is, the Stone tautologies were viewed as candidates for separating DPLL with clause learning from resolution.

The present paper, however, refutes this conjecture and establishes that the Stone tautologies do indeed have polynomial size pool refutations and regRTI refutations. In light of this, the possibility that pool and regRTI actually simulate the full power of resolution perhaps becomes slightly more plausible. Nevertheless, even if such a general simulation result does hold, it is far from clear whether the methods we use to deal with the Stone tautologies can be of much help in proving it.

The remainder of this introduction gives a review of the basic definitions, first of the Stone principles, then of the various proof systems. It concludes by stating our main theorems about the existence of pool and regRTI refutations. The reader is encouraged to consult the introductory sections of [BB12, BBJ14] for a more extensive discussion of prior work, and to consult [AJPU07] for more on the Stone principles. A good general introduction to DPLL with clause learning is [BKS04].

Definition 1.1. A *literal* is either a propositional variable x or a negated variable $\overline{x}$. A *clause* is a set of literals, usually written as a list of literals separated by either $\vee$'s (disjunctions) or commas. A clause is interpreted as the disjunction of its members. A set Γ of clauses is interpreted as the conjunction of its members, so Γ represents a propositional formula in conjunctive normal form.

The next definition describes the Stone principle of [AJPU07] as a set of clauses. The Stone principle is a kind of induction principle. For a given directed acyclic graph (dag), it states that if each source vertex is pebbled with a red stone and if each vertex whose immediate predecessors are pebbled with red stones is also pebbled with a red stone, then the sink vertex is pebbled with a red stone.

Definition 1.2. Assume that $G = (V, E)$ is a dag with a single sink, with vertices $V = \{1, \ldots, N\}$, such that each non-source vertex of G has in-degree 2. We assume that vertices are numbered consistently with the directions of the edges of G so that if there is an edge

$(i', i) \in E$ from i' to i then $i' > i$, and so that the source nodes of G are exactly vertices $n+1, n+2, \ldots, N$ for some n. Vertex 1 is the sole sink of G. Further assume that $m \geq N$; here, m is the number of "stones". The (negation of the) Stone tautology for G and m is denoted $Stone(G, m)$ and uses variables $p_{i,j}$ to indicate that vertex $i \in V$ is marked ("pebbled") with the j-th stone and variables r_j to indicate that the j-th stone is colored red. $Stone(G, m)$ contains the following clauses:

- $\bigvee_{j=1}^{m} p_{i,j}$, for each vertex i in G. (Each vertex is pebbled by at least one stone.)
- $\overline{p}_{i,j} \vee r_j$, for each $j = 1, \ldots, m$, and each source vertex i in G. (Each stone on a source vertex is colored red.)
- $\overline{p}_{1,j} \vee \overline{r}_j$, for each $j = 1, \ldots, m$. (The sink vertex 1 is not pebbled by any red stone.)
- $\overline{p}_{i',j'} \vee \overline{r}_{j'} \vee \overline{p}_{i'',j''} \vee \overline{r}_{j''} \vee \overline{p}_{i,j} \vee r_j$, whenever i' and i'' are the two vertices such that $(i', i) \in E$ and $(i'', i) \in E$ and $j \notin \{j', j''\}$. (If the two predecessors of i in G are pebbled by red stones, then every stone pebbling vertex i is also red. These "induction clauses" are equivalent to $p_{i',j'} \wedge r_{j'} \wedge p_{i'',j''} \wedge r_{j''} \wedge p_{i,j} \to r_j$.)

It is permitted that vertices are pebbled with more than one stone; likewise, the same stone may pebble multiple vertices.

The Stone clauses are clearly inconsistent since if the source vertices are pebbled with red stones then the induction clauses imply that all other vertices are also pebbled with red stones, and this contradicts the third group of clauses asserting that the sink vertex is not pebbled with a red stone.

We next recall the definitions of various types of resolution.

Definition 1.3. Let A, B, and C be clauses, and x a literal such that $\overline{x} \notin A$ and $x \notin B$. Consider the inference

$$\frac{A \qquad B}{C}$$

The literal x is the *resolution variable*. Three kinds of inferences are defined by:

Resolution rule: We have $x \in A$, $\overline{x} \in B$, and $C = (A \setminus \{x\}) \vee (B \setminus \{\overline{x}\})$.

Degenerate resolution rule: [HBPVG08, VG05] If $x \in A$ and $\overline{x} \in B$, then C is obtained as in the resolution rule. If $x \in A$ and $\overline{x} \notin B$, then C is B. If $x \notin A$ and $\overline{x} \in B$, then C is A. Otherwise C is one of A or B.

w-resolution rule: [BHJ08] The clause C equals $(A \setminus \{x\}) \vee (B \setminus \{\overline{x}\})$.

The three different types of resolution coincide when $x \in A$ and $\overline{x} \in B$, in which case we refer to the inference as *non-degenerate*.

Definition 1.4. A *resolution derivation* $\mathcal{D}$ of a clause C from a set Γ of clauses is a sequence of clauses $C_1, \ldots, C_s = C$ and such that each C_i is either a clause from Γ or is inferred by a resolution rule from two previous clauses. If C is the empty clause, $\mathcal{D}$ is a *resolution refutation* of Γ. *Degenerate resolution* and *w-resolution* derivations and refutations are defined similarly.

The *size* of a refutation $C_1, \ldots, C_s = \perp$ is defined to be s.

Derivations are typically viewed as directed acyclic graphs. A derivation is *tree-like* provided its dag is a tree. It is well known that (tree-like) resolution is sound and complete, in that Γ has a refutation iff it is unsatisfiable.

Definition 1.5. A refutation $\mathcal{D}$ is *regular* provided that no variable is used as a resolution variable more than once along any path in the directed acyclic graph of $\mathcal{D}$. A derivation $\mathcal{D}$ of a clause C is *regular* provided that, in addition, no variable appearing in C is used as a resolution variable in $\mathcal{D}$.

We next define "regular resolution derivation trees with lemmas", or "regRTL", following [BHJ08]. The idea is that a dag-like proof can by rewritten as a tree-like proof in which clauses obtained earlier in the proof can be used freely as "learned" lemmas. This will be the key component in defining Van Gelder's notion of pool proofs.

Definition 1.6. Given a tree T, the *postorder* ordering $<_T$ of the nodes is defined as follows: if u, v, w are distinct nodes of T, v is a node in the subtree rooted at the left child of u, and w is a node in the subtree rooted at the right child of u, then $v <_T w <_T u$. The *preorder* ordering $<'_T$ is defined similarly, but stipulates that $u <'_T v <'_T w$.

Definition 1.7. A *regRTL derivation* [BHJ08] of a clause C from a set of initial clauses Γ is a tree-like resolution derivation T that fulfills the following conditions: (a) each leaf is labeled with either a clause of Γ or a clause (called a "lemma") that appears earlier in T in the $<_T$ ordering; (b) each internal node is labeled with a clause and a literal, and the clause is obtained by resolution from the clauses labeling the node's children by resolving on the given literal; (c) the proof tree is regular; (d) the root is labeled with C. If the labeling of the root is the empty clause, T is a *regRTL refutation*.

A *regWRTL derivation* [BHJ08] is defined similarly, but allowing w-resolution inferences instead of just resolution inferences.

A *pool resolution derivation* [VG05] is also defined similarly, but allowing degenerate resolution inferences.

Proposition 1.8. *If Γ has a regWRTL refutation R, then Γ has a pool resolution refutation R' with the size of R' no greater than the size of R.*

The proof of Proposition 1.8 is simple. Each clause C in R corresponds to a clause C' in R' with $C' \subseteq C$. Arguing inductively, suppose that C is derived in R from the clauses C_1 and C_2 using resolution literal x. Then, it is straightforward to define C' from C'_1 and C'_2 as the unique clause that can be inferred by degenerate resolution from C'_1 and C'_2 with respect to x. □

The strategy of proving the existence of short pool refutations via constructing short regWTRL refutations is employed in the proof of Theorem 1.14 below.

Definition 1.9. ([BHJ08]). A "lemma" in clause (a) of the definition of regRTL derivations is called an *input lemma* if it is derived by an *input* subderivation, namely by a subderivation in which each inference has at least one hypothesis which is a member of Γ or a lemma. A regRTI derivation is a regRTL derivation which uses only input lemmas as lemmas.

A bit more generally, we say that a clause is "learned" provided it is available for use as a lemma by virtue of having been learned earlier in the postorder traversal of the proof, or by virtue of being an initial clause:

Definition 1.10. Suppose that R is a regRTL (respectively, a regRTI) refutation of Γ, and let C be a clause in R. The *learned clauses* of R at clause C are the clauses which are either in Γ or which have been derived in R (respectively, have been derived by an input subderivation in R) before C in the postordering of R.

Theorem 5.1 of [BHJ08] gives a polynomial equivalence between regRTI proofs and DPLL with clause learning without restarts. This equivalence, however, uses non-greedy DPLL; namely, the DPLL proof search may need to ignore contradictions during its search. Since most real-world DPLL search algorithms do not ignore contradictions, and use unit propagation whenever possible, is it natural to posit similar properties for regRTI proofs. These are formalized by the next two definitions.

Definition 1.11. Let C be a clause appearing in a regular derivation R. Following [BBJ14], we write C^{pool} to denote the clause containing the literals that appear in any clause in the path from the root of R up to and including C.

The clause C^{pool} is the same as what [BB12] calls C^+. The regularity of R ensures that C^{pool} contains no contradictory literals.

Definition 1.12. Let R be a tree-like refutation of Γ. A clause D in R is *prior-learned* for a clause C in R if either $D \in \Gamma$ or there is an occurrence of D as a learned clause which appears in R before C in both postorder and preorder.

The intuition for "prior-learned", is that, when reaching the clause C while constructing R in left-to-right, depth-first order, the prior-learned clauses are the clauses that are already available to help derive C.

Definition 1.13. (See [BB12]). A refutation R is *greedy* provided that, for each clause C of R, if C or any subclause of C is prior-learned, then C itself is a prior-learned clause and is a leaf clause of R. A refutation R is *greedy and unit-propagating* provided that, for each clause C of R, if there is an input derivation of some clause $C' \subseteq C^{\mathrm{pool}}$ from the prior-learned clauses of R at C which does not resolve on any literal in C^{pool}, then C is derived in R by such a derivation.

We can now state our main results.

Theorem 1.14. *The Stone principles $Stone(G, m)$ have regWRTL refutations, and thus pool refutations, of size $O(Nm^3)$.*

Theorem 1.15. *The Stone principles $Stone(G, m)$ have regRTI refutations of size $O(N^3 m^4)$.*

It follows from Theorem 1.15 and Theorem 5.1 of [BHJ08] that DPLL proof search with clause learning and without restarts can refute the Stone principle clauses in polynomial time. It is possible that the regRTI refutations of Theorem 1.15 can be made greedy and unit-propagating, but we have not tried to prove this.

The proofs of Theorems 1.14 and 1.15 are given in Sections 3 and 4, respectively. Section 2 first gives some preliminary resolution derivations that will be useful for both proofs.

Of course, Theorem 1.15 implies Theorem 1.14 apart from the size bounds. However, it seems useful to prove the two theorems separately, since the proof of Theorem 1.14 is substantially simpler than the proof of Theorem 1.15.

The intuition behind both proofs is similar. The reason the Stone tautologies seem highly irregular is that, in the earlier refutations given by [AJPU07], some of the variables (the r_j's) are resolved on repeatedly during the refutation. The intuition is that the regWRTL/regRTI proofs for Theorems 1.14 and 1.15 can be built in a bottom-up fashion starting from the empty clause, by first resolving on the variables $p_{i,j}$ that do not cause irregularities, and saving the problematic variables r_j to be resolved on later (higher in the

proof). This is not quite completely true, since our derivations do also resolve again on $p_{i,j}$'s at the top of the derivations; it is nonetheless a useful intuition.

2. Learning and 3-Learning

The regWRTL refutation for Theorem 1.14 and the regRTI refutation for Theorem 1.15 both work by learning the clauses $\overline{p}_{i,j}, r_j$. If i is a source vertex of G then these clauses are Stone clauses, but otherwise they must be learned.

Suppose that i is a non-leaf vertex, and i' and i'' are the two predecessors of i in G. In addition, suppose that every clause $\overline{p}_{i',j}, r_j$ and $\overline{p}_{i'',j}, r_j$ has already been learned. Fix a value of j. A derivation of $\overline{p}_{i,j}, r_j$ proceeds in the following three steps.

First, for each $j' \neq j''$, both distinct from j, derive the clause

$$\overline{p}_{i',j'}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j \tag{2.1}$$

by resolving a Stone clause against the two learned clauses $\overline{p}_{i',j'}, r_{j'}$ and $\overline{p}_{i'',j''}, r_{j''}$ using $r_{j'}$ and $r_{j''}$ as resolution variables:

$$\frac{\dfrac{\overline{p}_{i',j'}, \overline{r}_{j'}, \overline{p}_{i'',j''}, \overline{r}_{j''}, \overline{p}_{i,j}, r_j \qquad \overline{p}_{i',j'}, r_{j'}}{\overline{p}_{i',j'}, \overline{p}_{i'',j''}, \overline{r}_{j''}, \overline{p}_{i,j}, r_j} \qquad \overline{p}_{i'',j''}, r_{j''}}{\overline{p}_{i',j'}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j}$$

For $j' = j'' \neq j$, the clause (2.1) is derived in one step by resolving the Stone clause against only one of the two learned clauses $\overline{p}_{i',j'}, r_{j'}$ and $\overline{p}_{i'',j''}, r_{j''}$.

Second, for each $j'' \neq j$, resolve the Stone clause $\bigvee_{j'=1}^{m} p_{i',j'}$ against the learned clause $\overline{p}_{i',j}, r_j$ and against $m-1$ of the clauses (2.1) to obtain

$$\overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j. \tag{2.2}$$

This is shown in Figure 1.

Third, resolve the Stone clause $\bigvee_{j''=1}^{m} p_{i'',j''}$ against the learned clause $\overline{p}_{i'',j}, r_j$ and against the $m-1$ many clauses (2.2), and derive the desired clause $\overline{p}_{i,j}, r_j$. This is shown in Figure 2.

For the regWRTL proof constructed in Section 3, the clause $\overline{p}_{i,j}, r_j$ will be learned and available to use as a lemma once the above three steps have been carried out.

For the regRTI proof described in Section 4, this is not sufficient, since only clauses derived by input subderivations are learned. For regRTI proofs, the above three steps are used the first time $\overline{p}_{i,j}, r_j$ is derived. This results in the clauses (2.1) being learned as input lemmas, but not the clauses (2.2). The second time $\overline{p}_{i,j}, r_j$ is derived, only the second and third steps of the above derivation are carried out. This results in the clauses (2.2) becoming learned, but not the clause $\overline{p}_{i,j}, r_j$. The third time $\overline{p}_{i,j}, r_j$ is derived, only the third step is needed; this results in the clause $\overline{p}_{i,j}, r_j$ becoming learned as an input lemma.

This leads to the following definition, which will be useful for the regRTI derivations of Section 4:

Definition 2.1. Let i, i', i'', j be as above, so in particular all the clauses $\overline{p}_{i',j'}, r_{j'}$ and $\overline{p}_{i'',j''}, r_{j''}$ are learned. The clause $\overline{p}_{i,j}, r_j$ is called *3-learned* provided it has been learned. It is called *2-learned* if all of the clauses (2.2) for $j'' \neq j$ have been learned. It is called *1-learned* if all of the clauses (2.1) for $j' \neq j$, $j'' \neq j$ have been learned.

$$p_{i',1}, p_{i',2}, \ldots, p_{i',m} \qquad \overline{p}_{i',m}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j$$

$$p_{i',1}, p_{i',2}, \ldots, p_{i',m-1}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j \qquad \overline{p}_{i',m-1}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j$$

$$p_{i',1}, p_{i',2}, \ldots, p_{i',m-2}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j$$

$$p_{i',1}, p_{i',2}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j \qquad \overline{p}_{i',2}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j$$

$$p_{i',1}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j \qquad \overline{p}_{i',1}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j$$

$$\overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j$$

Figure 1: The derivation of a clause (2.2) follows this pattern with the one exception (not shown) that the clause $\overline{p}_{i',j}, \overline{p}_{i'',j''}, \overline{p}_{i,j}, r_j$ is not one of the $m-1$ clauses (2.1) and the learned clause $\overline{p}_{i',j}, r_j$ is used instead.

$$p_{i'',1}, p_{i'',2}, \ldots, p_{i'',m} \qquad \overline{p}_{i'',m}, \overline{p}_{i,j}, r_j$$

$$p_{i'',1}, p_{i'',2}, \ldots, p_{i'',m-1}, \overline{p}_{i,j}, r_j \qquad \overline{p}_{i'',m-1}, \overline{p}_{i,j}, r_j$$

$$p_{i'',1}, p_{i'',2}, \ldots, p_{i'',m-2}, \overline{p}_{i,j}, r_j$$

$$p_{i'',1}, p_{i'',2}, \overline{p}_{i,j}, r_j \qquad \overline{p}_{i'',2}, \overline{p}_{i,j}, r_j$$

$$p_{i'',1}, \overline{p}_{i,j}, r_j \qquad \overline{p}_{i'',1}, \overline{p}_{i,j}, r_j$$

$$\overline{p}_{i,j}, r_j$$

Figure 2: The derivation of the clause $\overline{p}_{i,j}, r_j$ follows this pattern with the one exception (not shown) that the clause $\overline{p}_{i'',j}, \overline{p}_{i,j}, r_j$ is not one of the $m-1$ clauses (2.2) and the clause $\overline{p}_{i'',j}, r_j$ is used instead.

A vertex i is defined to be K-learned, for $K = 1, 2, 3$, if and only if every $\overline{p}_{i,j}, r_j$ has been K-learned. It is also allowed that $K = 0$: every clause $\overline{p}_{i,j}, r_j$ and every vertex i is considered to be 0-learned.

Since axiom clauses are considered to be learned, the source vertices $i > n$ are 3-learned by definition.

The next theorem summarizes the above construction.

Theorem 2.2. *Let i' and i'' be the two predecessors of vertex i. There is a regular tree-like derivation of the clause $\overline{p}_{i,j}, r_j$ from Stone clauses and the clauses $\overline{p}_{i',j'}, r_{j'}$ and $\overline{p}_{i'',j''}, r_{j''}$, which has size $O(m^2)$ and resolves on (only) the variables r_k for $k \neq j$ and the variables $p_{i',k}$ and $p_{i'',k}$ for all k.*

In the setting of a regRTI proof, if i' and i'' are 3-learned, and $\overline{p}_{i,j}, r_j$ was already K-learned for $K < 3$, then there is a regular tree-like derivation of $\overline{p}_{i,j}, r_j$ from learned clauses (including Stone clauses) which causes it to become $(K+1)$-learned. This derivation has size $O(m^2)$ and resolves on at most the variables r_k for $k \neq j$ and variables $p_{i',k}$ and $p_{i'',k}$ for all k.

It will also be useful to modify the derivations described above to allow side variables $\overline{r}_{j_1}, \ldots, \overline{r}_{j_\ell}$. This is summarized by the next theorem.

Theorem 2.3. *Let i, i', i'' be as above. Let $F = \{\overline{r}_{j_1}, \ldots, \overline{r}_{j_\ell}\}$ where $\overline{r}_j \notin F$. There is a regular tree-like derivation of the clause $F, \overline{p}_{i,j}, r_j$ from Stone clauses and the clauses $\overline{p}_{i',j'}, r_{j'}$ and $\overline{p}_{i'',j''}, r_{j''}$, which has size $O(m^2)$ and resolves on (only) the variables r_k for $k \notin \{j, j_1, \ldots, j_\ell\}$ and the variables $p_{i',k}$ and $p_{i'',k}$ for all k.*

The derivation for Theorem 2.3 is obtained from the derivation for Theorem 2.2 by omitting inferences that resolve on the literals r_{j_q} against the clauses $\overline{p}_{i',j_q}, r_{j_q}$ and $\overline{p}_{i'',j_q}, r_{j_q}$.

3. The pool/regWRTL refutation

This section proves Theorem 1.14 by describing regWRTL proofs of the Stone principles. Fix an instance of the Stone principle for a dag G as above with N vertices and m stones. Recall that G has $n < N$ non-source vertices.

The regWRTL refutation of $Stone(G, m)$ will be a tree with its final, empty, clause at the bottom. The main part of the regWRTL refutation above the empty clause is a "skeleton", which consists of a long branch containing n segments of length m each, as is shown in Figures 3 and 4. Each segment in the skeleton corresponds to a non-source vertex in G, and the role of the i-th segment is that clauses of the form $\overline{p}_{i,j}, r_j$ are learned on branches to the right of the segment. In keeping with the intuition discussed in the introduction, this skeleton is the bottom part of the proof which resolves on the literals $p_{i,j}$; the variables r_j (plus additional variables $p_{i',j'}$ with $i' > i$) will be resolved on above the skeleton.

The first, second, and last segments of the skeleton are pictured in Figure 3. Each of these three segments is somewhat atypical, but the i-th segment for a typical intermediate $i \in \{3, \ldots, n-1\}$ is pictured in Figure 4. In the typical situation, the idea is that for given $3 \leq i \leq n-1$ and $j < m$, the clause $\overline{p}_{i,j}, r_j$ is learned in the $\cdots \vdots \cdots$ part of the proof above the clause $\overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,j}$ on the right hand side of Figure 4. For $i = 1, 2, n$ and $j \neq m$, the idea is that the clause $\overline{p}_{i,j}, r_j$ is learned in the $\cdots \vdots \cdots$ part of the proof above the clause containing $\overline{p}_{i,j}$. There are various exceptions to this idea, and several complications, as discussed below. However, in all cases, when working in the subproof in the $\cdots \vdots \cdots$ part
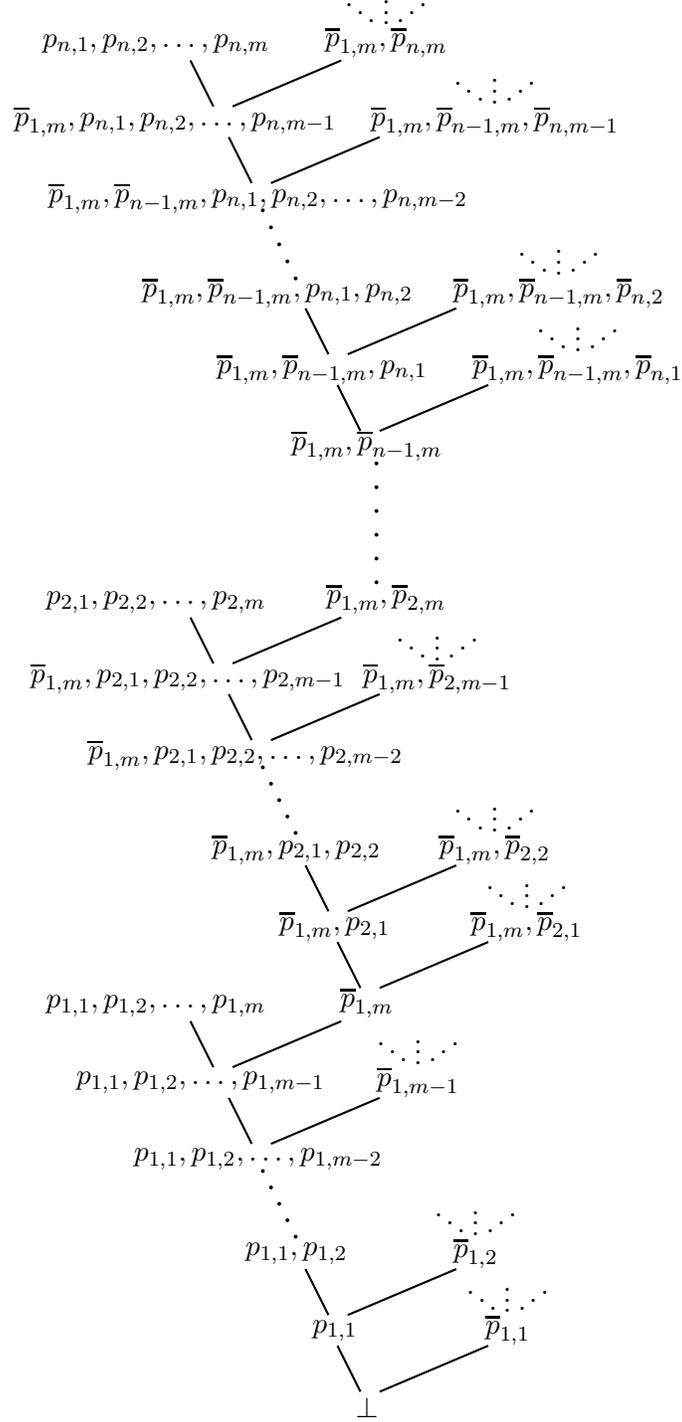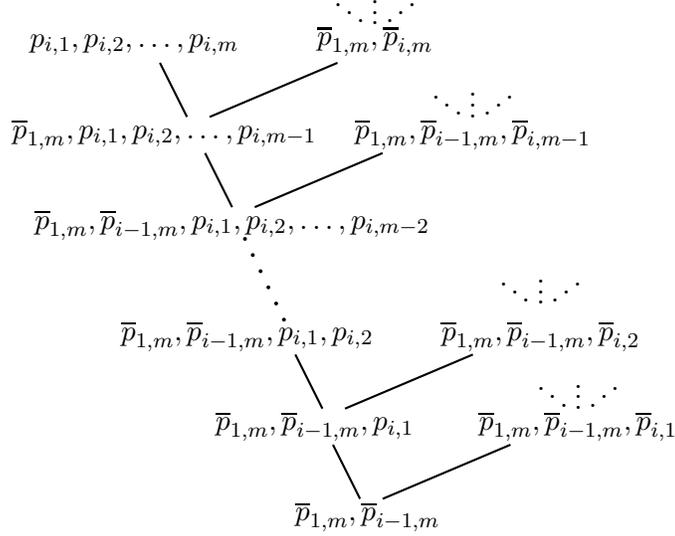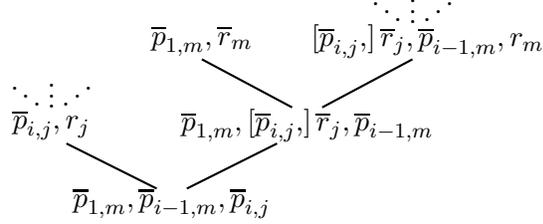
Figure 3: The "skeleton" of the regWRTL proof.

of the proof above the clause containing $\overline{p}_{i,j}$, all clauses of the form $\overline{p}_{i',j'}, r_{j'}$ with $i' > i$ have already been learned. To maintain the regularity property, this subproof must itself be regular, and will not resolve on any literals $p_{i',j'}$ with $i' \leq i$.

$$p_{i,1}, p_{i,2}, \ldots, p_{i,m} \qquad \overline{p}_{1,m}, \overline{p}_{i,m}$$

$$\overline{p}_{1,m}, p_{i,1}, p_{i,2}, \ldots, p_{i,m-1} \qquad \overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,m-1}$$

$$\overline{p}_{1,m}, \overline{p}_{i-1,m}, p_{i,1}, p_{i,2}, \ldots, p_{i,m-2}$$

$$\overline{p}_{1,m}, \overline{p}_{i-1,m}, p_{i,1}, p_{i,2} \qquad \overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,2}$$

$$\overline{p}_{1,m}, \overline{p}_{i-1,m}, p_{i,1} \qquad \overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,1}$$

$$\overline{p}_{1,m}, \overline{p}_{i-1,m}$$

Figure 4: The i-th segment of the skeleton, for $3 \le i < n$.

$$\overline{p}_{1,m}, \overline{r}_m \qquad [\overline{p}_{i,j},] \overline{r}_j, \overline{p}_{i-1,m}, r_m$$

$$\overline{p}_{i,j}, r_j \qquad \overline{p}_{1,m}, [\overline{p}_{i,j},] \overline{r}_j, \overline{p}_{i-1,m}$$

$$\overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,j}$$

Figure 5: The proof above $\overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,j}$.

We now outline how a clause $\overline{p}_{i,j}, r_j$ is learned in the most typical case, where $i = 3, \ldots, n-1$ and $j \le m-2$ or $i = n$ and $j \le m-1$. (The restriction that $j < m-2$ when $i < n-1$ is made because a more complicated construction will be needed when $j = m-1$ in order to also learn $\overline{p}_{i,m}, r_m$.) The part of the proof directly above $\overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,j}$ is presented in Figure 5.

The clause $\overline{p}_{i,j}, r_j$ on the left hand side of Figure 5 is the clause we want to learn. This clause is derived, and learned, by the derivation given by Theorem 2.2. If i' and i'' are the two predecessors of i in G, then $i' > i$ and $i'' > i$ and thus the clauses $\overline{p}_{i',j'}, r_{j''}$ and $\overline{p}_{i'',j''}, r_{j''}$ will have all already been learned, so Theorem 2.2 is applicable.

On the right hand side of Figure 5, we do not need to learn anything. We only need to make sure that the right hand side is a well-formed proof. The notation $[\overline{p}_{i,j},]$ indicates that $\overline{p}_{i,j}$ may not be present. In fact, $\overline{p}_{i,j}$ is present exactly when i is a predecessor of $i-1$; otherwise, it is absent.

To describe the right hand side of Figure 5, first suppose that vertex i is not a predecessor of vertex $i-1$ in the graph. In this case, the literal $\overline{p}_{i,j}$ is not present, and the leaf clause $\overline{r}_j, \overline{p}_{i-1,m}, r_m$ can be proved by the proof given by Theorem 2.3. Second, suppose

that vertex i is a predecessor of $i-1$. We must give a derivation of

$$\overline{p}_{i,j}, \overline{r}_j, \overline{p}_{i-1,m}, r_m. \tag{3.1}$$

Let $i' > i$ be the other predecessor of $i-1$. The derivation proceeds as follows. First, for each $j' \notin \{j, m\}$, it resolves the learned clause $\overline{p}_{i',j'}, r_{j'}$ against the Stone clause

$$\overline{p}_{i',j'}, \overline{r}_{j'}, \overline{p}_{i,j}, \overline{r}_j, \overline{p}_{i-1,m}, r_m$$

to obtain

$$\overline{p}_{i',j'}, \overline{p}_{i,j}, \overline{r}_j, \overline{p}_{i-1,m}, r_m. \tag{3.2}$$

These steps use the resolution variables $r_{j'}$ for $j' \notin \{j, m\}$. For $j' = j$, the clause (3.2) is a Stone clause and does not need to be derived. Then, it resolves the Stone clause $\bigvee_{j'} \overline{p}_{i',j'}$ against the learned clause $\overline{p}_{i',m}, r_m$ and the $m - 1$ many clauses (3.2), resolving on the literals $p_{i',j'}$. This yields the desired clause (3.1).

That completes the description of the typical case of learning $\overline{p}_{i,j}, r_j$. In less typical cases, the changes are as follows:

- $i = 2$ and $j \leq m - 2$. As described above, except that the variables $p_{1,m}$ and $p_{i-1,m}$ coincide.
- $i = 1$, $j \leq m - 2$. There is no $p_{i-1,m}$. The clause $\overline{p}_{1,j}$ is derived from $\overline{p}_{1,j}, r_j$ and $\overline{p}_{1,j}, \overline{r}_j$. The former is learned using the derivation of Theorem 2.2, while the latter is a Stone clause.
- $i = n$ and $j = m$. The clause $\overline{p}_{1,m}, \overline{p}_{n,m}$ is derived from $\overline{p}_{n,m}, r_m$ and $\overline{p}_{1,m}, \overline{r}_m$. The former is learned via Theorem 2.2, the latter is a Stone clause.
- $2 \geq i \geq n - 1$ and $j = m - 1, m$. There is no natural place in the i-th segment of the skeleton to learn clause $\overline{p}_{i,m}, r_m$, but we must learn it somewhere. To create "room" to learn both $\overline{p}_{i,m-1}, r_{m-1}$ and $\overline{p}_{i,m}, r_m$, the clause $\overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,m-1}$ is derived by w-resolution on the resolution variable $p_{i,m}$ from the two clauses $\overline{p}_{1,m}, \overline{p}_{i-1,m}, \overline{p}_{i,m-1}$ (i.e., itself) and $\overline{p}_{1,m}, \overline{p}_{i,m}$. The derivation proceeds as in the typical case above the former clause and as over $\overline{p}_{1,m}, \overline{p}_{n,m}$ above the latter.
- $i = 1$, $j = m - 1, m$. The clause $\overline{p}_{1,m-1}$ is derived using a w-resolution inference with the resolution variable $p_{1,m}$ from itself and $\overline{p}_{1,m}$. Above those two clauses, the proof proceeds as in all other cases with $i = 1$.

We have now fully described the polynomial size refutation for the Stone principle in regWRTL. A size bound of $O(nm^3) = O(Nm^3)$ is immediate from inspection, using the bound of $O(m^2)$ for the size of the derivations of Theorems 2.2 and 2.3. This completes the proof of Theorem 1.14.

Note that our construction of the regWRTL refutation of $Stone(G, m)$ makes no use of the assumption that $m \geq N$. This is in contrast to the construction of regRTI proofs in the next section.

4. The RegRTI proof

We now give a regRTI refutation R for the $Stone(G, m)$ principles. We describe R by building it from the bottom up, constructing R in a left-to-right depth-first fashion. At each point of the construction, R is a partially formed regRTI refutation. The leaves of R are designated as either "finished" or "unfinished", and all finished leaves are to the left of all unfinished leaves. The finished leaves are learned clauses; namely, they are either valid Stone clauses or have been derived by an input subderivation earlier in the postorder

of R. Clauses $\overline{p}_{i,j}, r_j$, or vertices i are defined to be K-learned, $K = 0, 1, 2, 3$, according to whether they have been K-learned in R at the point of reaching the leftmost unfinished leaf of R.

Each unfinished leaf will contain a clause C of the form

$$\overline{p}_{i_1,j_1}, \overline{p}_{i_2,j_2}, \ldots, \overline{p}_{i_k,j_k}, \tag{4.1}$$

for $k \geq 0$. The *domain*, $\mathrm{dom}(C)$, of C is equal to $\{i_1, \ldots, i_k\}$. Clauses of the form (4.1) will be required to have distinct values for the i_ℓ's; for convenience we assume that

$$i_1 \; < \; i_2 \; < \; \cdots \; < \; i_k. \tag{4.2}$$

We let $\max \mathrm{dom}(C)$ denote the maximum member of $\mathrm{dom}(C)$, namely i_k. We let $\mathrm{dom}(C^{\mathrm{pool}})$ denote the set of i such that some $\overline{p}_{i,j} \in C^{\mathrm{pool}}$.

We define the notion of a "well-formed unfinished clause" momentarily. The intuition behind an unfinished leaf C is based on the idea that a proof is being constructed from the bottom up, by a search process which sets the values of resolution variables in such a way that the clause reached at a given point becomes false. Upon reaching C, the search process has set all of the literals in C (and C^{pool}) false, so that each vertex $i_\ell \in \mathrm{dom}(C)$ has been pebbled with stone j_ℓ. The search process will generate a derivation of C by proving that each such stone j_ℓ is red, namely that r_{j_ℓ} is true. Since well-formed initial clauses will have $i_1 = 1$, this will yield a contradiction and thereby the desired refutation.

The stone j_ℓ can be shown to be red using one of the following three scenarios: (1) i_ℓ is a source vertex in G, (2) there is $i_{\ell'} > i_\ell$ such that $j_{\ell'} = j_\ell$ and stone $j_{\ell'}$ is red, or (3) the two predecessors of i_ℓ in G are pebbled with red stones. Accordingly, for a fixed clause C of the form (4.1) satisfying (4.2), we define:

Definition 4.1. Let $1 \leq \ell \leq k$. The vertex i_ℓ is said to be *bypassed* if there is some $\ell' > \ell$ such that $j_{\ell'} = j_\ell$. For the maximum such value ℓ', the vertex $i_{\ell'}$ is called the *max-bypasser* of i_ℓ.

Now let $1 \leq \ell < \ell' \leq k$. We say that vertex $i_{\ell'}$ *directly supports* vertex i_ℓ if either (1) $i_{\ell'}$ is the max-bypasser of i_ℓ, or (2) i_ℓ is not bypassed and $i_{\ell'}$ is one of the two predecessors of i_ℓ in G. Note that i_ℓ can directly support multiple $i_{\ell'}$'s.

The "*supports*" relation is the reflexive, transitive closure of "directly supports"; namely, if $i_{\ell_1} > i_{\ell_2} > \cdots > i_{\ell_s}$, $s \geq 1$, and each i_{ℓ_q} directly supports $i_{\ell_{q-1}}$, then i_{ℓ_1} supports i_{ℓ_s}. We use $S_\ell(C)$ to denote the set of vertices in $\mathrm{dom}(C)$ which support i_ℓ. Note that $S_\ell(C) \subseteq \{i_\ell, \ldots, i_k\}$.

The construction of R starts with the empty clause as the first unfinished clause. The first step will be to generate new unfinished clauses of the form $\overline{p}_{1,j_1}$; namely of the form (4.1) with $k = 1$ and $i_1 = 1$. In subsequent steps, an unfinished clause is extended by adding literals $\overline{p}_{i_{k+1},j_{k+1}}$ where i_{k+1} is the least vertex $> i_k$ which supports the vertex $i_1 = 1$. The next definitions make this formal.

Definition 4.2. A clause C is a *well-formed unfinished clause* provided C is of the form (4.1), satisfies (4.2) with $i_1 = 1$ and $i_k \leq n$, and the following three conditions hold:

i: C^{pool} contains only literals of the form $\overline{p}_{i,j}$ for $i \leq i_k$, and for each i there is at most one j such that $\overline{p}_{i,j}$ in C^{pool}.

ii: Let $1 \leq \ell \leq k$, and consider i_ℓ. Then either

    a: The vertex i_ℓ is bypassed, or

b: The vertex i_ℓ is not bypassed, and each predecessor i' of i_ℓ in G satisfies one of the following three conditions:

(α): $i' \in \mathrm{dom}(C)$;

(β): Vertex i' is already 3-learned, and $i' \notin \mathrm{dom}(C^{\mathrm{pool}})$; or

(γ): $i' > i_k$, and i' is not 3-learned.

In case (α), i' may or may not be 3-learned.

iii: Each $i_\ell \in \mathrm{dom}(C)$ supports $i_1 = 1$. Equivalently, $S_1(C) = \mathrm{dom}(C)$.

The empty clause is also a well-formed unfinished clause.

Definition 4.3. A non-empty well-formed unfinished clause of the form (4.1) is *extendible* provided there is some i_ℓ with at least one predecessor i' in G that satisfies condition (γ). The empty clause is also extendible.

During the construction of R, all unfinished leaves will be well-formed unfinished clauses. To describe the construction, we explain how to handle the leftmost unfinished leaf of the so-far constructed portion of R.

The extendible case. First suppose that the leftmost unfinished leaf is an extendible clause C of the form (4.1) with $k > 0$. Considering all vertices i_ℓ, find one with the least predecessor i' that satisfies (γ). This least i' is denoted i_{k+1}. In the special case where $k = 0$ and R contains just the empty clause (as C), let $i_{k+1} = 1$. In either case, i_{k+1} is not 3-learned and not a source vertex for G, so $i_{k+1} \leq n$.

With i_{k+1} chosen, define D_t to be the clause $C, \overline{p}_{i_{k+1},t}$. The idea is that we would like to replace C in R with a derivation of C from the Stone clause $p_{i_{k+1},1}, \ldots, p_{i_{k+1},m}$ and the clauses D_t for $t = 1, \ldots, m$ by resolving on the variables $p_{i_{k+1},t}$. The problem with this is that the D_t's may not be well-formed unfinished clauses. So, we consider the set $S_1(D_t)$, namely the set of literals that support the root vertex 1 in D_t.

Claim 4.4. $i_{k+1} \in S_1(D_t)$.

The claim is trivial for $k = 0$. To prove it when $k > 0$, first suppose that t is equal to some j_ℓ for $\ell \leq k$. Choose i_ℓ to be the least value such that $t = j_\ell$; of course i_{k+1} is a max-bypasser for i_ℓ in D_t. Using the fact that $S_1(C)$ contains every $i_{\ell'}$ for $\ell' \leq k$, a simple induction argument proves that $i_{\ell'} \in S_1(D_t)$ for every $\ell' \leq \ell$. It follows that $i_{k+1} \in S_1(D_t)$ since it is the max-bypasser of $i_\ell \in S_1(D_t)$. Second, suppose that t is distinct from all the j_ℓ values. A similar induction argument proves readily that $S_1(D_t)$ contains every i_ℓ for $\ell \leq k$. Hence $i_{k+1} \in S_1(D_t)$ since i_{k+1} is a predecessor of some non-bypassed $i_\ell \in S_1(D_t)$. $\square$

Claim 4.5. There is a t such that $S_1(D_t) = \{i_1, \ldots, i_k, i_{k+1}\}$.

To prove this, take t distinct from all j_ℓ values, and use the result from the second subcase of the previous claim. There must exist such a t since $m \geq N$, i.e., there are at least as many stones as vertices. $\square$

Define C_t to be the clause containing the literals $\overline{p}_{i_\ell, j_\ell}$ for $i_\ell \in S_1(D_t)$ with $1 \leq \ell \leq k$. Claim 4.5 shows that $C = \bigcup_t C_t$. With the aid of Claim 4.4 and the fact that C satisfies conditions ii and iii of the definition of well-formedness, it follows from the definition of C_t that the clause $C_t, \overline{p}_{i_{k+1},t}$ also satisfies the conditions ii and iii.

Now replace the clause C in R with the resolution derivation shown in Figure 6. The clauses $C_t, \overline{p}_{i_{k+1},t}$ are clearly well-formed unfinished clauses, the C_t's are subclauses of C,

$$p_{i_{k+1},1}, \ldots, p_{i_{k+1},m} \qquad C_1, \overline{p}_{i_{k+1},1}$$

$$C_1, p_{i_{k+1},2}, \ldots, p_{i_{k+1},m} \qquad C_2, \overline{p}_{i_{k+1},2}$$

$$C_2^*, p_{i_{k+1},3}, \ldots, p_{i_{k+1},m}$$

$$C_{m-2}^*, p_{i_{k+1},m-1}, p_{i_{k+1},m} \qquad C_{m-1}, \overline{p}_{i_{k+1},m-1}$$

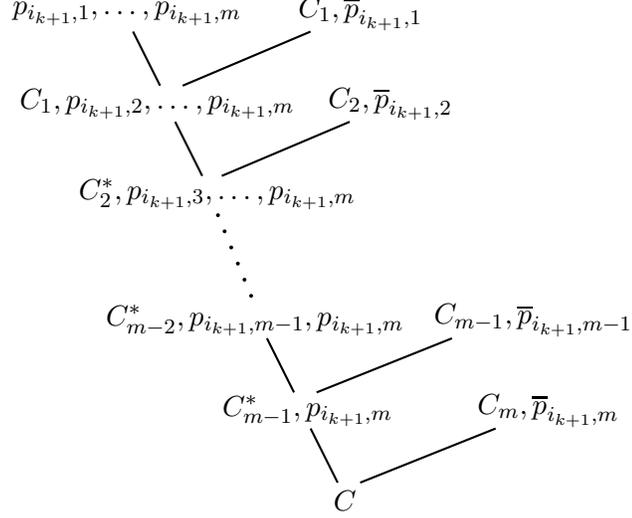$$C_{m-1}^*, p_{i_{k+1},m} \qquad C_m, \overline{p}_{i_{k+1},m}$$

$$C$$

Figure 6: The proof in the extendible case. The resolution variables are the $p_{i_{k+1},s}$'s. Each clause C_i^* is the union of the clauses $C_1, \ldots, C_i$.

and the (sub)clauses C_j^* in the figure are defined to equal $\bigcup_{t \le j} C_t$. By Claim 4.5, C_m^* is equal to C, which implies that the inference used to derive C is a valid non-degenerate resolution inference. For all the other newly added inferences, this is clear.

That completes the construction in the case where the leftmost unfinished leaf is extendible.

The non-extendible case. The construction for the case where the leftmost unfinished leaf is not extendible is summarized in the following lemma.

Lemma 4.6. *Suppose that clause C, of the form (4.1), is the leftmost unfinished leaf of R, and that C is not extendible. Then there is a regRTI derivation R_C of C from the Stone clauses and the learned clauses in R to the left of C. If the clause $\overline{p}_{i_k,j_k}, r_{j_k}$ was already K-learned for $K < 3$, then it becomes $(K+1)$-learned after the clause C is replaced in R with R_C. The only variables used as resolution variables in R_C are variables r_j, and variables $p_{i,j}$ for $i \notin \mathrm{dom}(C^{\mathrm{pool}})$. The size of R_C is $O(Nm^2)$.*

Proof. First suppose that $k = 1$. Then $i_1 = 1$, so $\overline{p}_{i_1,j_1}, \overline{r}_{i_1}$ is a Stone clause. Since C is not extendible, the two predecessors i' and i'' of i_1 are 3-learned. Theorem 2.2 gives a resolution derivation L_1 of $\overline{p}_{i_1,j_1}, r_{j_1}$ from 3-learned clauses using resolution on (at most) the variables r_j with $j \ne j_1$ and the variables $p_{i',j}$ and $p_{i'',j}$. Form the derivation R_C as

$$
\begin{array}{c}
L_1 \\
\cdots \vdots \cdots \\
\underline{\overline{p}_{i_1,j_1}, r_{j_1} \qquad \overline{p}_{i_1,j_1}, \overline{r}_{j_1}} \\
\overline{p}_{i_1,j_1}
\end{array}
\tag{4.3}
$$

Here and in the sequel, notation of the form $\overset{L}{\cdots \vdots \cdots}$ indicates that the dots should be replaced by the derivation L without its (already pictured) final clause.

By the second half of Theorem 2.2, if $\overline{p}_{i_1,j_1}, r_{j_1}$ was K-learned for $K < 3$, it becomes $(K+1)$-learned. The size of L_1 is $O(m^2)$.

We now assume that $k > 1$. This case is more involved, and takes up most of the rest of the paper.

Define B to be the set of those non-bypassed vertices $i_\ell \in \mathrm{dom}(C)$ that have at least one predecessor outside of $\mathrm{dom}(C)$. Note that since C is not extendible, a non-bypassed element of $\mathrm{dom}(C)$ belongs to B exactly if it has a predecessor satisfying condition (β) from the definition of well-formedness; in particular, the predecessor has to be 3-learned. We further split B into the sets B_1 and B_2 depending on whether one or both predecessors satisfy (β). Note that $i_1 = 1$ might be in B_1, but it cannot be in B_2, because i_2 is either a predecessor or a max-bypasser of i_1. On the other hand, i_k must be in B_2. We also define B^+ to be the set containing B and all those i_ℓ which are bypassed by a max-bypasser $i_{\ell'} \in B$. Thus, a vertex i_ℓ is in B^+ if, according to the minimal partial assignment falsifying clause C, i_ℓ is pebbled by a stone that also covers some vertex in B.

The idea behind the construction of R_C is that vertices i_ℓ that belong to $B \cup \{1\}$ determine a "partition" of $G{\restriction}_{\mathrm{dom}(C)}$ into possibly overlapping subgraphs (namely the sets $S_\ell^B(C)$ defined next). R_C will deal with these subgraphs independently. The set $S_\ell^B(C)$ is a subgraph with sink i_ℓ:

Definition 4.7. For $i_\ell \in \mathrm{dom}(C)$, the set $S_\ell^B(C)$ is the smallest set containing i_ℓ and satisfying the following whenever $i_{\ell'} \in S_\ell^B(C)$ and either $\ell' = \ell$ or $i_{\ell'} \notin B$: (1) If $i_{\ell'}$ is bypassed in C by max-bypasser $i_{\ell''} \notin B$, then $i_{\ell''} \in S_\ell^B(C)$ and (2) if $i_{\ell'}$ is not bypassed in C and $i_{\ell''}$ is a predecessor of $i_{\ell'}$ in G, then $i_{\ell''} \in S_\ell^B(C)$.

Note that the closure condition (1) does not allow $i_{\ell''} \in B$, whereas (2) does allow it.

Enumerate $B \cup \{1\}$ in decreasing order as $i_{t_1}, i_{t_2}, \ldots, i_{t_r}$ so that $t_1 = k$ and $t_r = 1$. Below, we use the convention that the index q ranges over $1, \ldots, r$, so that the vertices i_{t_q} are the members of $B \cup \{1\}$.

The overall structure of the derivation R_C is shown in Figure 7 below. R_C will be built around certain clauses related to the vertices $i_{t_q} \in B \cup \{1\}$ (the clauses C_{q-1}^*, F_q in Figure 7). To complete R_C, we will have to construct derivations of the clauses to the side of this "skeleton" (the derivations L_{t_q} of the clauses $C_{t_q}, F_q, r_{j_{t_q}}$ in Figure 7). In the case where $i_{t_q} \in B_2$, this is easy (Lemma 4.12 below), and this easy case is the one which makes $i_{t_1} = i_k$ become $(K+1)$-learned instead of K-learned. In the case where $i_{t_q} \in B_1 \cup \{1\}$, we only need to obtain a valid derivation, but this requires a relatively complex construction based on the structure of the subgraph determined by i_{t_q} (Lemma 4.13).

A clause appearing in R_C will contain some literals of the form $\overline{p}_{i,j}$, some literals of the form $\overline{r}_j$, and at most one literal of the form r_j. These literals have to be selected so as to avoid irregularities and degenerate inferences. R_C is defined using four special types of clauses: C_ℓ, C_q^*, E_ℓ, and F_q. (The C_ℓ's are different from the C_t's used for the extendible case.) The C_ℓ's and C_q^*'s consist of $\overline{p}_{i,j}$'s, while the E_ℓ's and F_q's consist of $\overline{r}_j$'s. As suggested by the notation, the C_q^*'s and F_q's are parametrized by vertices $i_{t_q} \in B \cup \{1\}$, whereas the C_ℓ's and E_ℓ's are parametrized by vertices $i_\ell \in \mathrm{dom}(C)$.

Definition 4.8. If i_ℓ is not bypassed in the clause C, then C_ℓ is the clause containing the literals $\overline{p}_{i_{\ell'},j_{\ell'}}$ for $i_{\ell'} \in S_\ell^B(C)$. If i_ℓ is bypassed in C, then C_ℓ is the same set except the literal $\overline{p}_{i_\ell,j_\ell}$ is omitted.

Definition 4.9. C_q^* equals $\{\overline{p}_{1,j_1}\} \cup \bigcup_{q'>q} C_{t_{q'}}$.

The reason we have to explicitly include $\overline{p}_{1,j_1}$ in C_q^* is that $\overline{p}_{1,j_1} \notin C_1$ if vertex 1 is bypassed.

Definition 4.10. If $i_\ell \in B$, then the clause E_ℓ is the set of literals $\overline{r}_{j_{\ell'}}$ such that $\ell' \neq \ell$ and $i_{\ell'} \in S_\ell^B(C) \cap B^+$. For $i_\ell \notin B$, the clause E_ℓ contains the literals $\overline{r}_{j_{\ell'}}$ such that $i_{\ell'} \in S_\ell^B(C) \cap B^+$.

Informally, E_ℓ contains the literals $\overline{r}_{j_{\ell'}}$ for $i_{\ell'} \in B^+$ a source (leaf) vertex above i_ℓ relative to the subgraph $S_\ell^B(C)$.

The cases where i_ℓ is bypassed by a max-bypasser $i_{\ell'}$ deserve special mention. Of course, $j_\ell = j_{\ell'}$. If $i_{\ell'} \in B$, then $i_\ell \in B^+ \setminus B$, and we have $S_\ell^B(C) = \{i_\ell\}$, $C_\ell = \emptyset$, and $E_\ell = \{\overline{r}_{j_\ell}\}$. On the other hand, if $i_{\ell'} \notin B$, then $i_\ell \notin B^+$, and we have $S_\ell^B(C) = \{i_\ell\} \cup S_{\ell'}^B(C)$, $C_\ell = C_{\ell'}$, and $E_\ell = E_{\ell'}$.

Before defining the F_q's, we prove a lemma listing some basic properties of the C's, C^*'s, and E's:

Lemma 4.11.

(a): *Each $i_\ell \neq 1$ is a member of some $S_{t_q}^B(C)$. Thus, C_0^* is equal to C.*
(b): *If 1 is bypassed by max-bypasser i_ℓ, then $\ell = 2$.*
(c): $\overline{r}_{j_{t_{r-1}}} \in E_1$.
(d): $E_{t_q} \subseteq \{\overline{r}_{j_{t_1}}, \ldots, \overline{r}_{j_{t_{q-1}}}\}$, *for all q.*

Proof. (a) Given $\ell \neq 1$, consider a chain $i_\ell = i_{\ell_1} > i_{\ell_2} > \cdots > i_{\ell_s} = 1$ of directly supporting vertices from i_ℓ to 1. Let i_{ℓ_a} be the first member of $B \cup \{1\}$ in this sequence. Then $\ell_a = t_{q'}$ for some q', and we have $i_\ell \in S_{t_{q'}}^B(C)$. Thus, $\overline{p}_{i_\ell, j_\ell}$ is in $C_{t_{q'}}$ and hence C_0^*.

(b) Suppose that 1 has max-bypasser i_ℓ with $\ell > 2$. Then $i_\ell > i_2$. There must exist a chain of directly supporting vertices from i_2 to 1, but this contradicts $i_\ell > i_2$.

(c) Consider a chain $i_{t_{r-1}} = i_{\ell_1} > i_{\ell_2} > \cdots > i_{\ell_s} = 1$ of directly supporting vertices from $i_{t_{r-1}}$ to 1. By the descending order of the t_q's, none of $i_{\ell_2}, \ldots, i_{\ell_{s-1}}$ is in B. If $i_{t_{r-1}}$ is not the max-bypasser of i_{ℓ_2}, then $i_{t_{r-1}} \in S_1^B(C)$ and thus $\overline{r}_{j_{t_{r-1}}} \in E_1$. Suppose instead that $i_{t_{r-1}}$ is the max-bypasser of i_{ℓ_2}, so $i_{\ell_2} \in B^+ \setminus B$. Then $i_{\ell_2} \in S_1^B(C)$, so $\overline{r}_{j_{\ell_2}} \in E_1$. By the definition of bypasser, $j_{\ell_2} = j_{t_{r-1}}$, so $\overline{r}_{j_{t_{r-1}}} \in E_1$.

(d) Let $\overline{r}_{j_\ell} \in E_{t_q}$, so $\ell \geq t_q$ and $i_\ell \in B^+$. If $i_\ell \in B$, then $\ell > t_q$, and thus $\ell = t_{q'}$ for some $q' < q$. If $i_\ell \in B^+ \setminus B$, then i_ℓ has max-bypasser $i_{\ell'} \in B$. Then $\ell' > \ell \geq t_q$ and thus ℓ' equals $t_{q'}$ for some $q' < q$, and by the definition of bypasser $j_{\ell'} = j_\ell$. Therefore $\overline{r}_{j_\ell}$ is the same as $\overline{r}_{j_{t_{q'}}}$. $\qquad\square$

Define F_q to be the clause

$$F_q := \begin{cases} E_{t_r} = E_1 & \text{if } q = r \text{ and } 1 \notin B \\ \overline{r}_{j_{t_1}}, \ldots, \overline{r}_{j_{t_{q-1}}} & \text{otherwise} \end{cases}$$

By Lemma 4.11(d), $E_{t_q} \subseteq F_q$. Note that $F_1 = \emptyset$, because $r \neq 1$ and so in evaluating F_1 we use the second case of the definition of F_q.

We can now describe the derivation R_C in detail. As mentioned, the general structure of R_C is shown in Figure 7. The parts of the derivation displayed in brackets are omitted if $1 \in B^+ \setminus B$.

Note that C_0^*, F_1, the final clause of R_C, is the same as C. For $q < r$, the inference

$$\frac{C_{t_q}, F_q, r_{j_{t_q}} \qquad C_q^*, F_{q+1}}{C_{q-1}^*, F_q}$$

resolves on $r_{j_{t_q}}$, and it is non-degenerate by the definitions of the F_q's and C_q^*'s. (By Lemma 4.11(c,d), this is true even in the case where $q = r - 1$ and $1 \notin B$.) It follows that the resolution variables on the path from $C_{t_q}, F_q, r_{j_{t_q}}$ to C are exactly $r_{j_{t_1}}, \ldots, r_{j_{t_q}}$. The derivation L_{t_q} is described below (cf. Lemmas 4.12 and 4.13).

For $q = r$, we have t_q equal to 1. If $1 \notin B^+ \setminus B$, then, as shown in Figure 7, the clause C_{r-1}^*, F_r is derived by:

$$\frac{\overset{L_1}{\underset{\ddots \; \vdots \; \cdot}{}} \atop \dfrac{C_1, F_r, r_{j_1} \qquad \overline{p}_{1,j_1}, \overline{r}_{j_1}}{C_{r-1}^*, F_r}}{}$$

Again, the final inference deriving C_{r-1}^*, F_r is non-degenerate. The upper right clause $\overline{p}_{1,j_1}, \overline{r}_{j_1}$ is a Stone clause. The derivation L_1 is described in Lemma 4.13.

On the other hand, if $1 \in B^+ \setminus B$ and thus 1 has max-bypasser $i_2 \in B$, then C_{r-1}^* is $\overline{p}_{1,j_1}$ and $F_r = E_1$ is $\overline{r}_{j_1}$. Therefore clause C_{r-1}^*, F_r is equal to the Stone clause $\overline{p}_{1,j_1}, \overline{r}_{j_1}$, and there is no need to add to R_C anything above it.

To finish the description of R_C, we must describe the derivations L_{t_q}. This is done by the next two lemmas.

Lemma 4.12. *Suppose that $i_{t_q} \in B_2$. Then there is a regRTI proof L_{t_q} of $C_{t_q}, F_q, r_{j_{t_q}}$ of size $O(m^2)$. The variables used as resolution variables in L_{t_q} are (at most) the variables r_j where $\overline{r}_j \notin F_q \cup \{\overline{r}_{j_{t_q}}\}$ and the variables $\overline{p}_{i',j}$ and $\overline{p}_{i'',j}$ for i' and i'' the predecessors of i_{t_q} in G.*

In the special case of $q = 1$, so $t_1 = k$, if $\overline{p}_{i_k,j_k}, r_{j_k}$ was K-learned for $K < 3$, then $\overline{p}_{i_k,j_k}, r_{j_k}$ becomes $(K+1)$-learned in L_k.

Note that the values i' and i'' are not in $\mathrm{dom}(C^{\mathrm{pool}})$ by the definition of B_2; therefore, the resolution variables $\overline{p}_{i',j}$ and $\overline{p}_{i'',j}$ do not violate the regularity condition. Also note that, since $1 \notin B_2$, we have $t_q \neq 1$ and the condition that $\overline{r}_j \notin F_q \cup \{\overline{r}_{j_{t_q}}\}$ is equivalent to $j \notin \{j_{t_1}, \ldots, j_{t_q}\}$. This is precisely what is needed to ensure regularity.

Proof. The clause C_{t_q} consists of the single literal $\overline{p}_{i_{t_q},j_{t_q}}$. Thus the derivation L_{t_q} is the derivation given by Theorem 2.3. When $q = 1$ and $t_q = k$, we have $F_1 = \emptyset$ and C_k is the clause $\overline{p}_{i_k,j_k}$. In this case, L_k is the derivation given by the second part of Theorem 2.2. $\square$

Lemma 4.13. *Suppose that $1 \leq q \leq r$, and $i_{t_q} \in B_1 \cup \{1\}$ and $i_{t_q} \notin B^+ \setminus B$. Let $N_q = |S_{t_q}^B(C)|$.*

(a): *There is a regular dag-like derivation L'_{t_q} of size $O(N_q)$ which contains each clause $C_\ell, E_\ell, r_{j_\ell}$ for $i_\ell \notin B^+$ such that $i_\ell \in S_{t_q}^B(C)$. The subderivation L'_ℓ of L'_{t_q} that ends with $C_\ell, E_\ell, r_{j_\ell}$ uses as resolution variables precisely the variables $r_{j_{\ell'}}$ such that $j_\ell \neq j_{\ell'}$, $i_{\ell'} \in S_\ell^B(C)$, and $i_{\ell'} \notin B^+$.*

If $i_{t_q} \notin B_1$, then $t_q = 1$ and $q = r$ and the final clause of L'_{t_q} is C_1, E_1, r_{j_1} (that is, with $\ell = t_q = 1$). Suppose instead that $i_{t_q} \in B_1$ and i_ℓ is the (only) predecessor of i_{t_q}

$$\begin{bmatrix} L_1 & \cdots \vdots \cdots \\ & C_1, F_r, r_{j_1} \end{bmatrix} \quad [\overline{p}_{1,j_1}, \overline{r}_{j_1}]$$

$$L_{t_{r-1}} \cdots \vdots \cdots$$

$$C_{t_{r-1}}, F_{r-1}, r_{j_{t_{r-1}}} \qquad C_{r-1}^*, F_r$$

$$C_{r-2}^*, F_{r-1}$$

$$L_{t_2} \cdots \vdots \cdots$$

$$C_{t_2}, F_2, r_{j_{t_2}}$$

$$L_{t_1} \cdots \vdots \cdots$$

$$C_{t_1}, F_1, r_{j_{t_1}} \qquad C_1^*, F_2$$

$$C_0^*, F_1$$

Figure 7: The structure of the derivation R_C.

such that $i_\ell \in \mathrm{dom}(C^{\mathrm{pool}})$. If $i_\ell \in B^+$, then $S_{t_q}^B(C) \subseteq B^+$ and L'_{t_q} is empty. If $i_\ell \notin B^+$, then the final clause of L'_{t_q} is $C_\ell, E_\ell, r_{j_\ell}$.

(b): There is a regRTI derivation L_{t_q} of $C_{t_q}, F_q, r_{j_{t_q}}$ of size $O(N_q^2 + m)$. The variables used as resolution variables in L_{t_q} are at most the variables r_j where $\overline{r}_j \notin F_q \cup \{\overline{r}_{t_q}\}$ and, if $i_{t_q} \in B_1$, the variables $p_{i'',j}$ where i'' is the predecessor of i_{t_q} such that $i'' \notin \mathrm{dom}(C^{\mathrm{pool}})$.

Note that $i_{t_q} \in B^+ \setminus B$ only if $t_q = 1$, and this is the case where L_1 is not needed.

Proof. The regular dag-like refutation L'_{t_q} for part (a) is constructed by induction on ℓ such that $i_\ell \in S_{t_q}^B(C) \setminus B^+$. We add the clauses $C_\ell, E_\ell, r_{j_\ell}$ to L'_{t_q} for larger values of ℓ first, making sure that the condition on resolution variables remains satisfied. The inductive argument splits into four cases.

Case 1: i_ℓ is bypassed by max-bypasser $i_{\ell'}$. Since $i_\ell \notin B^+$, we have $i_{\ell'} \notin B^+$. The induction hypothesis tells us that $C_{\ell'}, E_{\ell'}, r_{j'_\ell}$ appears in the already constructed portion of L'_{t_q}. ¿From the remarks after the definitions of C_ℓ and E_ℓ, we have $j_\ell = j_{\ell'}$ and $C_\ell = C_{\ell'}$ and $E_\ell = E_{\ell'}$. Thus $C_{\ell'}, E_{\ell'}, r_{j'_\ell}$ is exactly the same clause as $C_\ell, E_\ell, r_{j_\ell}$. So no further resolution inferences need to be added to L'_{t_q} to handle i_ℓ, and $L'_\ell = L'_{\ell'}$. Since $S_\ell^B(C) = \{i_\ell\} \cup S_{\ell'}^B(C)$ and $i_\ell \notin B^+$, the subderivation L'_ℓ satisfies the condition about which resolution variables are used.

The remaining cases all assume that i_ℓ is not bypassed.

Case 2 (the base case): both of i_ℓ's predecessors $i_{\ell'}$ and $i_{\ell''}$ are in B^+. Then C_ℓ is $\overline{p}_{i_{\ell'},j_{\ell'}}, \overline{p}_{i_{\ell''},j_{\ell''}}, \overline{p}_{i_\ell,j_\ell}$. Also, E_ℓ is $\overline{r}_{j_{\ell'}}, \overline{r}_{j_{\ell''}}$. Therefore the Stone clause (4.4) is the same as $C_\ell, E_\ell, r_{j_\ell}$, so L'_ℓ consists of just this clause.

Case 3: neither of i_ℓ's predecessors $i_{\ell'}$ and $i_{\ell''}$ is in B^+. The already constructed part of L'_ℓ contains the clauses $C_{\ell'}, E_{\ell'}, r_{j_{\ell'}}$ and $C_{\ell''}, E_{\ell''}, r_{j_{\ell''}}$. Assume for the moment that $j_{\ell'} \neq j_{\ell''}$. W.l.o.g., the highest index s'' such that $j_{s''} = j_{\ell''}$ is strictly greater than highest index s' such that $j_{s'} = j_{\ell'}$; otherwise interchange ℓ' and ℓ''. It follows from the induction step for Case 1 that $L'_{\ell''}$ equals $L'_{s''}$. Therefore, by the assumption that $s' < s''$ and the inductive condition on resolution variables, $r_{j_{\ell'}}$ is not resolved on in $L'_{\ell''}$.

Using the Stone clause

$$\overline{p}_{i_{\ell'},j_{\ell'}}, \overline{r}_{j_{\ell'}}, \overline{p}_{i_{\ell''},j_{\ell''}}, \overline{r}_{j_{\ell''}}, \overline{p}_{i_\ell,j_\ell}, r_{j_\ell} \tag{4.4}$$

form L'_ℓ as

$$
\cfrac{
\cfrac{\overline{p}_{i_{\ell'},j_{\ell'}}, \overline{r}_{j_{\ell'}}, \overline{p}_{i_{\ell''},j_{\ell''}}, \overline{r}_{j_{\ell''}}, \overline{p}_{i_\ell,j_\ell}, r_{j_\ell} \qquad \cfrac{\begin{smallmatrix}\ddots\ \vdots\ \cdot^{\cdot^{\cdot}} \end{smallmatrix}\, L'_{\ell''}}{C_{\ell''}, E_{\ell''}, r_{j_{\ell''}}}}
{\overline{p}_{i_{\ell'},j_{\ell'}}, \overline{r}_{j_{\ell'}}, \overline{p}_{i_{\ell''},j_{\ell''}}, C_{\ell''}, E_{\ell''}, \overline{p}_{i_\ell,j_\ell}, r_{j_\ell}} \qquad \cfrac{\begin{smallmatrix}\ddots\ \vdots\ \cdot^{\cdot^{\cdot}} \end{smallmatrix}\, L'_{\ell'}}{C_{\ell'}, E_{\ell'}, r_{j_{\ell'}}}
}
{C_\ell, E_\ell, r_{j_\ell}}
\tag{4.5}
$$

We have $E_\ell = E_{\ell'} \cup E_{\ell''}$. The literals $\overline{p}_{i_{\ell'},j_{\ell'}}$ and $\overline{p}_{i_{\ell''},j_{\ell''}}$ may or may not appear in $C_{\ell'}$ and $C_{\ell''}$ (respectively), but in any case C_ℓ is the same as $\overline{p}_{i_{\ell'},j_{\ell'}}, C_{\ell'}, \overline{p}_{i_{\ell''},j_{\ell''}}, C_{\ell''}, \overline{p}_{i_\ell,j_\ell}$. Note that the proof as pictured above might be slightly misleading: we are constructing a dag-like derivation, not a tree-like derivation and the subderivations $L'_{\ell'}$ and $L'_{\ell''}$ need not be disjoint. The regularity of L'_ℓ follows from the induction hypotheses and the fact that $r_{j_{\ell'}}$ is not a resolution variable of $L'_{\ell''}$. The condition on which resolution variables are used in L'_ℓ follows from $S^B_\ell(C) = \{i_\ell\} \cup S^B_{\ell'}(C) \cup S^B_{\ell''}(C)$.

For the remaining part of case 3, suppose that $j_{\ell'} = j_{\ell''}$. Let s be the highest index such that $j_s = j_{\ell'} = j_{\ell''}$. By the remarks after the definitions of C_ℓ and E_ℓ, we have $C_{\ell'} = C_{\ell''} = C_s$ and $E_{\ell'} = E_{\ell''} = E_s$. Thus also $L'_{\ell'} = L'_{\ell''} = L'_s$. Now, argue as in the previous case, but replace the inferences (4.5) with the inference

$$
\cfrac{\overline{p}_{i_{\ell'},j_s}, \overline{p}_{i_{\ell''},j_s}, \overline{r}_{j_s}, \overline{p}_{i_\ell,j_\ell}, r_{j_\ell} \qquad \cfrac{\begin{smallmatrix}\ddots\ \vdots\ \cdot^{\cdot^{\cdot}} \end{smallmatrix}\, L'_s}{C_s, E_s, r_{j_s}}}
{C_\ell, E_\ell, r_{j_\ell}}
$$

The left hypothesis is a Stone clause. The rest of the argument for this subcase is as in the previous case.

Case 4: i_ℓ has predecessors $i_{\ell'} \notin B^+$ and $i_{\ell''} \in B^+$. The induction hypothesis says that the clause $C_{\ell'}, E_{\ell'}, r_{j_{\ell'}}$ has already been derived. Then C_ℓ is $\overline{p}_{i_{\ell'},j_{\ell'}}, \overline{p}_{i_{\ell''},j_{\ell''}}, \overline{p}_{i_\ell,j_\ell}, C_{\ell'}$, and E_ℓ is $\overline{r}_{j_{\ell''}}, E_{\ell'}$, so we can form L'_ℓ as

$$
\cfrac{\overline{p}_{i_{\ell'},j_{\ell'}}, \overline{r}_{j_{\ell'}}, \overline{p}_{i_{\ell''},j_{\ell''}}, \overline{r}_{j_{\ell''}}, \overline{p}_{i_\ell,j_\ell}, r_{j_\ell} \qquad \cfrac{\begin{smallmatrix}\ddots\ \vdots\ \cdot^{\cdot^{\cdot}} \end{smallmatrix}\, L'_{\ell'}}{C_{\ell'}, E_{\ell'}, r_{j_{\ell'}}}}
{C_\ell, E_\ell, r_{j_\ell}}
$$

using resolution on $r_{\ell'}$. The regularity of L'_ℓ and the conditions on which resolution variables are used in L'_ℓ follow from the induction hypothesis for $L'_{\ell'}$ and the fact that $S^B_\ell(C) = \{i_\ell, i_{\ell''}\} \cup S^B_{\ell'}$.

That completes the proof of part (a). The size bound $O(N_q)$ on L'_ℓ follows from the fact that each of the four cases in the construction of L'_ℓ added $O(1)$ clauses.

To apply part (a) in the proof of (b), we need regRTI derivations L''_{t_q}, instead of the regular dag-like refutations L'_{t_q}. For this, Theorem 3.3 of [BHJ08] states that the desired derivation L''_{t_q}, containing exactly the same clauses as L'_{t_q}, can be constructed from L''_{t_q} with the size of L''_{t_q} bounded by the product of the size and the height of L'_{t_q}. Thus, the size of L''_{t_q} is $O(N_q^2)$. Furthermore, L''_{t_q} uses the same resolution variables as L'_{t_q}.

We now prove part (b). First suppose that $q = r$ and $1 \notin B$. Then $t_q = 1$ and $F_q = E_1$. In this case, the regRTI derivation L''_1 obtained from part (a) is already the desired derivation. This only uses resolution variables r_{j_ℓ} such that $i_\ell \notin B^+$ and hence $j_\ell \notin \{j_{t_1}, \ldots, j_{t_q}\}$.

Otherwise, $i_{t_q} \in B_1$, and hence $F_q = \{\overline{r}_{j_{t_1}}, \ldots, \overline{r}_{j_{t_{q-1}}}\}$. Let $i_{\ell'}$ and i'' be the predecessors of i_{t_q}; so, i'' is 3-learned and $i'' \notin \mathrm{dom}(C^{\mathrm{pool}})$. Note that $j_{\ell'}$ may or may not be in $\{j_{t_1}, \ldots, j_{t_q}\}$. Consider the $m - 2$ many Stone clauses for $j'' \notin \{j_{t_q}, j_{\ell'}\}$:

$$\overline{p}_{i_{\ell'}, j_{\ell'}}, \overline{r}_{j_{\ell'}}, \overline{p}_{i'', j''}, \overline{r}_{j''}, \overline{p}_{i_{t_q}, j_{t_q}}, r_{j_{t_q}}.$$

Resolving these with the 3-learned clauses $\overline{p}_{i'', j''}, r_{j''}$ for $j'' \notin \{j_{t_1}, \ldots, j_{t_q}, j_{\ell'}\}$ and the Stone clause $\bigvee_{j''} p_{i'', j''}$ gives the clause

$$\overline{p}_{i_{\ell'}, j_{\ell'}}, \overline{r}_{j_{\ell'}}, F_q, \overline{p}_{i_{t_q}, j_{t_q}}, r_{j_{t_q}} \tag{4.6}$$

by resolution on the variables $r_{j''}$ for $j'' \notin \{j_{t_1}, \ldots, j_{t_q}, j_{\ell'}\}$ and the variables $p_{i'', j''}$ for all j''.

Suppose that $i_{\ell'} \in B^+$ and thus $\overline{r}_{j_{\ell'}} \in F_q$ and C_{t_q} is the clause $\overline{p}_{i_{\ell'}, j_{\ell'}}, \overline{p}_{i_{t_q}, j_{t_q}}$. Then (4.6) is the same as $C_{t_q}, F_q, r_{j_{t_q}}$ and the construction of L_{t_q} for part (b) is complete. In this case, L_{t_q} has size $O(m)$.

Alternately, suppose that $i_{\ell'} \notin B^+$, so $\overline{r}_{j_{\ell'}} \notin F_q$. By the last assertion of part (a), L''_{t_q} is a regRTI derivation of $C_{\ell'}, E_{\ell'}, r_{j_{\ell'}}$ of size $O(N_q^2)$. Form L_{t_q} by resolving this against (4.6) on the variable $r_{j_{\ell'}}$ to obtain $C_{t_q}, F_q, r_{j_{t_q}}$. This is a valid resolution inference since C_{t_q} is $\overline{p}_{i_{t_q}, j_{t_q}}, \overline{p}_{i_{\ell'}, j_{\ell'}}, C_{\ell'}$, and $E_{\ell'} \subseteq F_q$. The size of L_{t_q} is $O(N_q^2 + m)$.

This completes the description of L_{t_q}, and the proof of Lemma 4.13. □

This also completes the construction of the derivation R_C of Lemma 4.6. Since the construction of R_C invokes Lemmas 4.12 and 4.13 at most N times, the size of R_C is bounded by $O(Nm^2 + N(N^2 + m)) = O(Nm^2)$.

Finishing the proof. All that remains to finish the proof of Theorem 1.15 is to bound the size of the refutation R. As described above, R is built up from a derivation containing only the empty clause by applying the constructions of Figure 6 and Lemma 4.6, always to the leftmost currently unfinished leaf.

If the clause at that leaf is non-extendible, it is dealt with using Lemma 4.6, and no new unfinished leaves appear.

Otherwise, the leftmost unfinished leaf contains an extendible clause, and it is dealt with using the construction of Figure 6, leading to m new unfinished leaves, at least one of them labeled with a clause $C_t, \overline{p}_{i_{k+1}, t}$ such that $\overline{p}_{i_{k+1}, t}, r_t$ has not been 3-learned.

A clause of the type allowed to appear in an unfinished leaf can be iteratively extended at most n times, so at some point we have to reach a situation in which the construction of Figure 6 produces only non-extendible clauses. When Lemma 4.6 is applied to those clauses,

at least one K-learned clause of the form $\overline{p}_{i,j}, r_j$ becomes $(K+1)$-learned. There are nm clauses of the form $\overline{p}_{i,j}, r_j$ to be learned. Once all of them have been 3-learned, all remaining unfinished leaves become non-extendible and can be dealt with using the construction of Lemma 4.6.

Consider the set of clauses in R which at some point were unfinished leaves during the construction of R. Call such a clause *green* if it was handled with Lemma 4.6 and thereby a clause of the form $\overline{p}_{i,j}, r_j$ became $(K+1)$-learned instead of K-learned. The remaining clauses are called *non-green*: these clauses were either handled by Lemma 4.6 but without any clause $\overline{p}_{i,j}, r_j$ becoming $(K+1)$-learned, or were handled with the construction of Figure 6. These green and non-green clauses inherit a tree structure from R. It follows from the discussion above that this tree is m-branching, has depth at most n, and contains at most $3nm$ green leaves. Moreover, each node in the tree is a green leaf, is an ancestor or sibling of a green leaf, or is the sibling of an ancestor of a green leaf. It is straightforward to prove that such a tree has at most $O(n^2m^2)$ leaves. Since each of these leaves corresponds to a single application of Lemma 4.6, the size of R is at most $O(N^3m^4)$. This completes the proof of Theorem 1.15. $\qquad\square$

Since the Stone tautologies contain $O(Nm^3)$ many symbols, and since $N \le m$, the size upper bound $O(N^3m^4)$ is quadratic in the size of the clauses being refuted.

The refutation R described above may not be greedy. Although we lack a proof, it is possible that R can be made greedy by omitting subderivations of already learned clauses. It is also possible that R is, or could be made to be, unit-propagating. In particular, note that the only unit clauses that appear in R are the literals $\overline{p}_{1,j}$ that appear as unfinished clauses in the very first step of the construction of R. However, we have not tried to formally analyze the greedy or unit-propagating properties of R.

Acknowledgement

> *The proof is complete, If only I've stated it thrice.*
> *Fit the Fifth – The Beaver's Lesson, The Hunting of the Snark*
> Lewis Carroll

References

[AJPU07] Michael Alekhnovich, Jan Johannsen, Toniann Pitassi, and Alasdair Urquhart. An exponential separation between regular and general resolution. *Theory of Computing*, 3(5):81–102, 2007.

[BB12] Maria Luisa Bonet and Samuel R. Buss. An improved separation of regular resolution from pool resolution and clause learning. In *Proc. 15th International Conference on Theory and Applications of Satisfiability Testing – SAT 2012*, Lecture Notes in Computer Science #7317, pages 45–57, 2012.

[BBJ14] Maria Luisa Bonet, Samuel R. Buss, and Jan Johannsen. Improved separations of regular resolution from pool resolution and clause learning. *Journal of Artificial Intelligence Research*, 49:669–703, 2014.

[BHJ08] Samuel R. Buss, Jan Hoffmann, and Jan Johannsen. Resolution trees with lemmas: Resolution refinements that characterize DLL-algorithms with clause learning. *Logical Methods in Computer Science*, 4, 4:13(4:13):1–18, 2008.

[BKS04]   Paul Beame, Henry A. Kautz, and Ashish Sabharwal. Towards understanding and harnessing
             the potential of clause learning. *Journal of Artificial Intelligence Research*, 22:319–351, 2004.
[DLL62]   Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem prov-
             ing. *Communications of the ACM*, 5(7):394–397, 1962.
[DP60]   Martin Davis and Hillary Putnam. A computing procedure for quantification theory. *Journal
             of the Association for Computing Machinery*, 7(3):201–215, 1960.
[HBPVG08]   Philipp Hertel, Fahim Bacchus, Toniann Pitassi, and Allen Van Gelder. Clause learning can
             effectively p-simulate general propositional resolution. In *Proc. 23rd AAAI Conf. on Artificial
             Intelligence (AAAI 2008)*, pages 283–290. AAAI Press, 2008.
[MSS99]   João P. Marques-Silva and Karem A. Sakallah. GRASP — A new search algorithm for satisfi-
             ability. *IEEE Transactions on Computers*, 48(5):506–521, 1999.
[Urq11]   Alasdair Urquhart. A near-optimal separation of regular and general resolution. *SIAM Journal
             on Computing*, 40(1):107–121, 2011.
[VG05]   Allen Van Gelder. Pool resolution and its relation to regular resolution and DPLL with clause
             learning. In *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2005)*, Lec-
             ture Notes in Computer Science 3835, pages 580–594. Springer-Verlag, 2005.