

# Verified LRAT and LPR Proof Checking with cake\_lpr

Yong Kiam Tan      Marijn J. H. Heule      Magnus O. Myreen

## 1 Summary

We present the `cake_lpr` proof checker [2] which is capable of checking proofs in either Linear RAT (LRAT) or Linear PR (LPR) proof formats.<sup>1</sup> The checker is formally verified using CakeML and the HOL4 theorem prover; its formal proof is discussed in Tan et al. [2] and briefly in Section 3. The DRAT and DPR proof formats are supported using `DRAT-trim` and `DPR-trim` as preprocessing tools, respectively. We also propose to add support for binary DPR, LPR, and LRAT formats in the relevant tools, following the binary DRAT format.

**NOTE: For SAT-COMP usage, please make sure to read Section 4.**

### 1.1 Obtaining the Proof Checker

The proof checker is available at:

[https://github.com/tanyongkiam/cake\\_lpr](https://github.com/tanyongkiam/cake_lpr)

The `DRAT-trim` and `DPR-trim` tools are available at:

<https://github.com/marijnheule/drat-trim>

<https://github.com/marijnheule/dpr-trim>

### 1.2 Example

An outline of an end-to-end LRAT proof checking run is as follows:

```
# Assume the problem is input.cnf in DIMACS format
... run SAT solver on input.cnf generate input.drat ...
# Run drat-trim on the DRAT proof, generate LRAT file
drat-trim input.cnf input.drat -L input.lrat
# Run cake_lpr on the resulting LRAT proof
cake_lpr input.cnf input.lrat
```

If the proof checks successfully, `cake_lpr` will print to standard output:

```
s VERIFIED UNSAT
```

<sup>1</sup>The LPR format is a backwards-compatible extension of LRAT.

All other error messages (proof checking error, parsing error, out-of-memory error etc.) will be printed to `stderr`. Solvers capable of generating LRAT proofs directly can skip the use of `DRAT-trim`. End-to-end proof checking for LPR proofs can be done similarly, using `DPR-trim` as the preprocessor for DPR proofs. It is also possible to convert DPR proofs to DRAT, then use `DRAT-trim`, but this approach is **not recommended** as it is significantly slower than checking DPR (and LPR) proofs directly [2].

## 2 Supported Proof Formats

Formal descriptions of all proof formats are available in the cited publications [1, 2] and online. We give brief descriptions of the formats with concrete examples.

### 2.1 DRAT and LRAT

The DRAT format consists of a list of clause addition or deletion steps, one per line. All lines are terminated by 0. Each added clause must have RAT redundancy with respect to the current formula.

```
<CLAUSE> 0
d <CLAUSE> 0
```

**Concrete example:**

```
1 -2 3 0 # Add clause x_1,!x_2,x_3
d 1 2 -3 0 # Del clause x_1,x_2,!x_3
```

The `DRAT-trim` tool can be used as a preprocessor to automatically convert an input DRAT proof to LRAT format. The latter format extends DRAT with a notion of clause IDs and proof hints for each line. The input CNF is assumed to be given IDs in ascending order from 1 to  $n$  where  $n$  is the number of clauses in the file. Addition lines in LRAT have the following format, where `<ID>` is a positive integer, `<IDs>` is a list of `<ID>`, and `[...]*` denotes 0 or more repetitions of the enclosed block:

```
<ID> <CLAUSE> 0 <IDs> [-<ID> <IDs>]* 0
```

The first `<ID>` is the clause ID to be assigned to `<CLAUSE>`. If `<CLAUSE>` has RAT redundancy, then the first literal in the clause is the pivot literal. The first block of `<IDs>` lists unit propagation steps starting from the blocking assignment for `<CLAUSE>`. If `<CLAUSE>` has RAT redundancy, then this first block is followed by 0 or more `-<ID> <IDs>` blocks, where `-<ID>` refers to the `<ID>`-th clause in the RAT proof and the corresponding `<IDs>` indicate unit propagation steps for that clause.

Deletion steps are written with a list of clause IDs rather than clauses. All the clauses with IDs in `<IDs>` are deleted.

```
<ID> d <IDs> 0
```

### Concrete example:

```
# Add clause x_1,!x_2,x_3 at clause ID 15 with RAT on pivot !x_2
15 -2 1 3 0 4 13 7 10 8 -5 78 2 4 -10 41 3 5 0
# Del clause IDs 13 14 15 (the ID 16 in front of the line is ignored)
16 d 13 14 15 0
```

A complexity analysis for the LRAT proof format is given in Cruz-Filipe et al. [1, Theorem 2], where asymptotically (keeping all parameters constant except number  $n$  of steps in proofs), the complexity is reported as  $O(n^2 \log n)$ ; `cake_lpr` slightly improves the asymptotic bound to  $O(n^2)$  because it uses constant-time rather than logarithmic-time lookup data structures [2]. Empirically, we have observed that most proofs generated by solvers in the SAT competition are dominated by simple (non-RAT) steps. In that case, one may expect near-linear scaling from `cake_lpr`.

## 2.2 DPR and LPR

The DPR format extends DRAT so that added clauses are *propagation redundant* with respect to the current formula. Here, `<WITNESS>` is a list of literals which must start with the first literal in `<CLAUSE>`. Note that this is syntactically backwards compatible with DRAT (when `<WITNESS>` is empty).

```
<CLAUSE> <WITNESS> 0
```

The `DPR-trim` tool can be used as a preprocessor to automatically convert an input DPR proof to LPR format. The latter format extends DPR with clause IDs and proof hints in the same way LRAT extends DRAT. The only syntactic addition is the optional `<WITNESS>` after `<CLAUSE>`.

```
<ID> <CLAUSE> <WITNESS> 0 <IDs> [-<ID> <IDs>]* 0
```

The proof checking procedure for LPR is backwards compatible with LRAT, using `<IDs>` and `[-<ID> <IDs>]*` as unit propagation hints for propagation redundancy [2]. Deletion lines are identical for LPR and LRAT.

### Concrete example:

```
# Add clause x_1,!x_2,x_3 at clause ID 15 with PR witness !x_2,x_5
15 -2 1 3 -2 5 0 4 13 7 10 8 -5 78 2 4 -10 41 3 5 0
# Del clause IDs 13 14 15 (the ID 16 in front of the line is ignored)
16 d 13 14 15 0
```

The proof checking procedure and thus the proof checking complexity for LPR is essentially the same as LRAT, i.e., with  $O(n^2)$  complexity (assuming all other parameters are held constant).

$\vdash \text{cake\_lpr\_run } cl \ fs \ mc \ ms \Rightarrow$		(1)
$\text{machine\_sem } mc \ (\text{basis\_ffi } cl \ fs) \ ms \subseteq$ $\text{extend\_with\_resource\_limit}$		(2)
$\{ \text{Terminate Success } (\text{cake\_lpr\_io\_events } cl \ fs) \} \wedge$ $\exists \text{ out err.}$		(3)
$\text{extract\_fs } fs \ (\text{cake\_lpr\_io\_events } cl \ fs) =$ $\text{Some } (\text{add\_stdout } (\text{add\_stderr } fs \ err) \ out) \wedge$ $\text{if length } cl = 2 \text{ then}$ $\text{if inFS\_fname } fs \ (\text{el } 1 \ cl) \text{ then}$ $\text{case parse\_dimacs } (\text{all\_lines } fs \ (\text{el } 1 \ cl)) \text{ of}$ $\text{None } \Rightarrow \ out = \langle\langle \rangle\rangle$ $\mid \text{Some } fml \Rightarrow \ out = \text{concat } (\text{print\_dimacs } fml)$ $\text{else } \ out = \langle\langle \rangle\rangle$ $\text{else if length } cl = 3 \text{ then}$ $\text{if } \ out = \langle\langle \text{s VERIFIED UNSAT}\backslash\Omega \rangle\rangle \text{ then}$ $\text{inFS\_fname } fs \ (\text{el } 1 \ cl) \wedge$ $\exists fml.$ $\text{parse\_dimacs } (\text{all\_lines } fs \ (\text{el } 1 \ cl)) = \text{Some } fml \wedge$ $\text{unsatisfiable } (\text{interp } fml)$ $\text{else } \ out = \langle\langle \rangle\rangle$ $\text{else } \dots$		(4)

Figure 1: The end-to-end correctness theorem for the CakeML LPR proof checker. (Some irrelevant cases are elided with ... for brevity).

### 3 Proof Checker Verification

Our proof checker, `cake_lpr`, is formally verified down to the level of its x64 machine code implementation, which eliminates the possibility of bugs arising from, e.g., compiler errors, code extraction, or other, unverified additions to (verified) source code. This is achieved by compiling its formally verified CakeML source code implementation, with a formally verified compiler for CakeML [2].

The key correctness theorem is shown in Fig. 1. To informally summarize:

- Line (1) assumes that the `cake_lpr` binary is executed in an x64 machine environment set up according to the standard CakeML assumptions.
- Lines (2) guarantees that `cake_lpr` will terminate successfully (i.e., no out of bounds array accesses, etc.); it may run out of either heap or stack memory (resource limits).
- Lines (3) says that, according to the CakeML file system model, there will be some strings printed to standard output and standard error.
- Lines (4) says (among other things) that, IF the string “s VERIFIED UNSAT” is printed onto standard output, then the first command line

argument corresponds to a file, which parsed in DIMACS format, to a formula that is unsatisfiable. This DIMACS parser is also verified to be left inverse to our DIMACS printer.

## 4 Advanced Compilation Instructions

CakeML is an end-to-end verified programming language. Its verified runtime manages its own heap and stack space, including garbage collection. The heap and stack space for `cake_lpr` is configured at compile time, and the default is set at 4GB each in order to support easy usage on personal computers. This setting can be found in `basis_ffi.c`:

```
unsigned long cml_heap_sz = 4096 * sz;    // Default: 4 GB heap
unsigned long cml_stack_sz = 4096 * sz;   // Default: 4 GB stack
```

**If a sufficiently powerful machine is available, e.g., during the SAT competition, we recommend changing the memory limits to support the high memory requirements of certain input problems as follows:**

```
unsigned long cml_heap_sz = 65536 * sz; // Default: 64 GB heap
unsigned long cml_stack_sz = 16384 * sz; // Default: 16 GB stack
```

**Please remember to recompile `cake_lpr` after the modification.**

## References

- [1] Luís Cruz-Filipe, Marijn J. H. Heule, Warren A. Hunt Jr., Matt Kaufmann, and Peter Schneider-Kamp. Efficient certified RAT verification. In Leonardo de Moura, editor, *CADE*, volume 10395 of *LNCS*, pages 220–236. Springer, 2017. doi: 10.1007/978-3-319-63046-5\_14.
- [2] Yong Kiam Tan, Marijn J. H. Heule, and Magnus O. Myreen. `cake_lpr`: Verified propagation redundancy checking in CakeML. In Jan Friso Groote and Kim Guldstrand Larsen, editors, *TACAS*, volume 12652 of *LNCS*, pages 223–241. Springer, 2021. doi: 10.1007/978-3-030-72013-1\_12.