Int J STTT (2000) 3: 93–105

Regular contribution

Detecting races in Relay Ladder Logic programs

Alexander Aiken*, **Manuel Fähndrich****, **Zhendong Su***

EECS Department, University of California, Berkeley 387 Soda Hall #1776, Berkeley, CA 94720-1776, USA;
E-mail: {aiken,zhendong}@cs.berkeley.edu

Abstract. Relay Ladder Logic (RLL) [5] is a programming language widely used for complex embedded control applications such as manufacturing and amusement park rides. The cost of bugs in RLL programs is extremely high, often measured in millions of dollars (for shutting down a factory) or human safety (for rides). In this paper, we describe our experience in applying constraint-based program analysis techniques to analyze production RLL programs. Our approach is an interesting combination of probabilistic testing and program analysis, and we show that our system is able to detect bugs with high probability, up to the approximations made by the conservative program analysis. We demonstrate that our analysis is useful in detecting some flaws in production RLL programs that are difficult to find by other techniques.

Key words: Constraints – Software – Static analysis – Testing – Verification

1 Introduction

Programmable logic controllers (PLCs) are used extensively for complex embedded control applications such as factory control in manufacturing industries and for entertainment equipment in amusement parks. Relay Ladder Logic (RLL) is the most widely used PLC programming language; approximately 50% of the manufacturing capacity in the United States is programmed in RLL [6].

RLL has long been criticized for its low-level design, which makes it difficult to write correct programs [22]. Moreover, validation of RLL programs is extremely expensive, often measured in millions of dollars (for factory down-time) or human safety (for rides). One solution is to replace RLL with a higher-level, safer programming language. An alternative is to provide better programming support directly for RLL. Since there are many existing RLL applications, and many more will be written in this language, we consider the latter approach in this paper.

We have designed and implemented a tool for analyzing RLL programs. Our analyzer automatically detects some common programming mistakes that are extremely difficult to detect through inspection or testing. The information inferred by the analyzer can be used by RLL programmers to identify and correct these errors. Our most interesting result is an analysis to detect certain race conditions in RLL programs. Tested on real RLL programs, the analysis found several such races, including one known bug that originally caused four hours of factory down-time [6] (factory down-time generally costs upwards of $3,000 per minute).

In this paper, we describe the design and implementation of our RLL program analyzer for detection of *relay races*. Our analysis is *constraint-based*, meaning that the information we wish to know about a program is expressed as constraints [3, 4, 19]. The solutions of these constraints yield the desired information. Our analysis is built using a general constraint resolution engine, which allows us to implement the analysis directly in the same natural form it is specified. Constraint-based program analysis is discussed further in Sect. 4.

Our analysis is similar to ternary simulation for testing circuits. The underlying semantic model of the constraints is essentially the same as that of ternary simulation, which is Kleene's three-valued logic [21]. Ternary simulation has been suggested and applied for the detection of hazards in combinatorial and sequential circuits [9, 10, 16, 27]. A detailed discussion of the relationship be-

* Supported in part by the National Science Foundation, Grant No. CCR-9416973, by NSF Infrastructure Grant No. CDA-9401156, and a gift from Rockwell Corporation.

** *Present address:* Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399, USA; E-mail: maf@microsoft.com

tween our approach and ternary simulation is postponed to the related work section (Sect. 8).

Our system has two components: (a) a conservative data and control flow analysis captures information about a program in an initial system of constraints; and (b) additional constraints binding program inputs to actual values are added to the initial constraint system, which is then solved to obtain the desired information. Part (a) is done only once, but part (b) is done many times for randomly chosen inputs. Our underlying constraint resolution engine solves and simplifies the initial constraints generated by (a), thereby greatly improving the performance of (b).

Beyond the particular application to RLL programs, this system architecture has properties that may be of independent interest. First, the use of constraints greatly simplifies the engineering needed to factor out the information to be computed once from that which must be reevaluated repeatedly — we simply add new constraints to the initial system. Second, our system is (to the best of our knowledge) a unique blend of conservative program analysis (part (a), which approximates certain aspects of computation) and software testing (part (b), which "executes" the abstraction for concrete inputs). Third, we are able to prove that classes of program errors are detected with high probability, up to the approximations made by the conservative analysis.

We expect that the engineering advantages of using constraints will carry over to other static analysis tools. The latter two results apply directly only if the programming language has a finite domain of values (RLL has Booleans only). Thus, our approach is suitable for some other special-purpose languages (e.g., other control languages) but not necessarily for general purpose languages.

The rest of the paper is structured as follows. First, we give an overview of the language RLL (Sect. 2) and of the race analysis (Sect. 3). Then we describe the constraint language used for the analysis (Sect. 4). The rules for generating the base system of constraints come next (Sect. 5), followed by a description of the relay race analysis (Sect. 6). Finally, we present some experimental results (Sect. 7), followed by a discussion of related work (Sect. 8) and conclusions (Sect. 9).

2 Overview of RLL

By any standard RLL is a strange language, combining features of Boolean logic (combinatorial circuits), imperative programming (assignment, `goto`, procedures, and conditionals), and real-time computation (timers and counters) with an obscure syntax and complex semantics. Although widely used, RLL is not well-known in the research community. We give a brief overview of RLL together with a more detailed, but still high-level, description of our analysis system.

RLL programs are represented as *ladder diagrams*, which are a stylized form of circuits or data flow diagrams. A *ladder diagram* consists of a set of *ladder rungs* with each rung having a set of input instructions and output instructions. We explain this terminology in the context of the example RLL program in Fig. 1. In the example, there are two vertical rails. The one on the left supplies power to all crossing rungs of the ladder. The three horizontal lines are the ladder rungs of this program. This example has four kinds of RLL instructions: input (two kinds), outputs, and timer instructions. The small vertical parallel bars | | and |/| represent input instructions, which have a single bit associated with them. The bit is named in the instruction. For example, the | | instruction (an XIC for "Normally Closed Contact" instruction) in the upper-left corner of the diagram reads from the bit named A, and the |/| instruction (an XIO for "Normally Opened Contact" instruction) in the lower-left corner of the diagram reads from the bit named C. The small circles represent output instructions that update the value of their labeled bits. The bits named in input and output instructions are classified into *external* bits, which are connected to inputs or outputs external to the program, and *internal* bits, which are local to the program for temporarily storing program states. External inputs are generally connected to sensors, while external outputs are used to control actuators. The rectangular box represents a timer instruction (a TON for "Timer On-Delay" instruction), where PR (preset) is an integer representing a time interval in seconds, AR (accumulator) keeps the accumulated value, and TB (time base) is the step of each increment of the AR. The timer instructions are used to turn an output on or off after the timer has been on for a preset time interval (the PR value).

Instructions are connected by wires, the horizontal lines between instructions. We say a wire is true if power is supplied to the wire, and the wire is false otherwise.

An RLL program operates by first reading all the values of the external input bits and executing the rungs

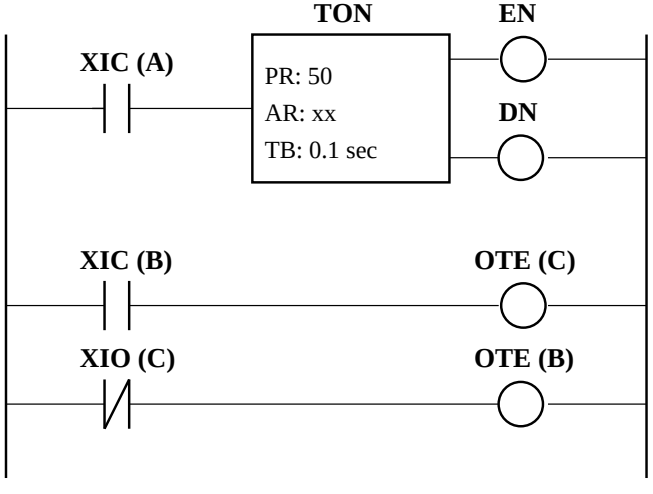

Fig. 1. An example RLL program

in sequence from top to bottom and left to right. Program control instructions may cause portions of the program to be skipped or repeatedly executed. After the last rung is evaluated, the output devices connected to the external output bits are updated. Such a three-step execution (read inputs, evaluate rungs, update outputs) of the program is called a *scan*. Programs are executed scan after scan until interrupted. Between scans, the input bit values might be changed, either because the inputs were modified by the previous scan (bits can be inputs, outputs, or both) or because of state changes in external sensors attached to the inputs. Subsequent scans use the new input values.

RLL has many types of instructions: relay instructions, timer and counter instructions, data transfer instructions, arithmetic operations, data comparison operations, and program control instructions. Examples of relay instructions are XIC, XIO, and OTE. We briefly describe how these three instructions work for the explanation of our analysis in Sect. 3. Let w_1 and w_2 be the wires before and after an instruction respectively. Further, let b be the bit referenced by an instruction.

XIC: if w_1 and b are true, w_2 is true; otherwise, w_2 is false.

XIO: if w_1 is true, and b is false, w_2 is true; otherwise, w_2 is false.

OTE: the bit b is true if and only if w_1 is true.

3 Overview of race analysis

In this section, we give a high-level overview of our RLL program analyzer for detecting relay races.

In RLL programs, it is desirable that the values of outputs depend solely on the values of inputs and the internal states of timers and counters. If under fixed inputs and timer and counter states, an output x changes from scan to scan, then there is a *relay race on* x. For example, in the program in Fig. 1, we will see later that the bit B changes value each scan regardless of its initial value. Relay races are particularly difficult to detect by traditional testing techniques, as races can depend on the timing of external events and the scan rate.

Our analysis generalizes traditional data flow analyses [1]. Instead of data flow equations, set constraints [3, 4, 19] are used. Set constraints are more expressive than data flow equations because the constraints can model not only data flow but also control flow of a program.

Our analysis consists of two steps. In the first step, we generate constraints that describe the data and control flow dependencies of an RLL program. The constraints are generated in a top-down traversal of the program's abstract syntax tree (AST). According to a set of constraint generation rules (see Sect. 5), appropriate constraints are generated for each AST node. These data and control flow constraints are solved to yield another system of simplified constraints, the *base system*. The base system models where and how a value flows in the program. The base system is a *conservative approximation* of the program. Whereas a wire can have only one value at a time during program execution, the analysis assigns a set of possible values to a wire. If during program execution, a wire or a bit can be true (false), then true (false) is in the set that denotes the values of the wire or the bit in the base system; however, false (true) may be a value in that set, too.

The second step of the relay race analysis simulates multiple scans and looks for racing outputs. We choose a random assignment of inputs and add the corresponding constraints to the base system. The resulting system is solved; its minimum solution describes the values of the outputs at the end of the scan. Since some output bits are also inputs, the input assignment of the next scan is updated using the outputs from the previous scan. Again, we add this input assignment to the base system and solve to obtain the minimum solution of the outputs after the second scan. If an output changes across scans, a relay race is detected. For example, consider the example program in Fig. 1. Since the bottom two rungs do not interfere with the first rung, consider these two rungs only. Assume that B has initial value true. Then C also is true, and so in the last rung, B becomes false. Thus, in the next scan, B is initially false. Thus, C becomes false, which makes B true at the end of this scan. Consequently, we have detected a relay race on B: after the first scan B is false, and after the second scan B is true.

The race analysis is conservative in the sense that it cannot detect all of the relay races in a program. However, any relay races the analyzer detects are indeed relay races, and we can prove that a large class of relay races is detected with high probability.

We have implemented the race analysis in Standard ML of New Jersey (SML) [24] using the BANE program analysis toolkit [2]. Our analyzer is accurate and fast enough to be practical — production RLL programs can be analyzed. The relay race analysis not only detected a known bug in a program that took an RLL programmer four hours of factory down-time to uncover, it also detected many previously unknown relay races in our benchmark programs.

4 Constraints

In this section, we describe the set constraint language we use for expressing our analysis. Our expression language consists of set variables, a least value $\bot$, a greatest value $\top$, constants $\mathbf{T}$ and $\mathbf{F}$, intersections, unions, and conditional expressions. The syntax of the expression language is

$$E ::= v \mid \bot \mid \top \mid c \mid E_1 \cup E_2 \mid E_1 \cap E_2 \mid E_1 \Rightarrow E_2,$$

where c is a constant (either $\mathbf{T}$ or $\mathbf{F}$) and $v \in V$ is a set variable.

The abstract domain consists of four elements: $\emptyset$ (represented by $\bot$), $\{\mathbf{T}\}$ (represented by $\mathbf{T}$), $\{\mathbf{F}\}$ (represented by $\mathbf{F}$), $\{\mathbf{T},\mathbf{F}\}$ (represented by $\top$) with set inclusion as the partial order on these elements. The domain is a finite lattice with $\cap$ and $\cup$ being the *meet* and *join* respectively. The semantics of the expression language is given in Fig. 2.

Conditional expressions deserve some discussion. Conditional expressions are used for accurately modeling flow-of-control (see e.g., [4]). In the context of RLL, they can be used to express Boolean relations very directly. For example, we can express the Boolean expression $v_1 \wedge v_2$ with the following conditional expression:

$$\begin{aligned}((v_1 \cap \mathbf{T}) \Rightarrow (v_2 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup \\ ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup \\ ((v_2 \cap \mathbf{F}) \Rightarrow \mathbf{F})\end{aligned}$$

To see this expression does model the $\wedge$ operator, notice that if $v_1 = \mathbf{T}$ and $v_2 = \mathbf{T}$, the above expression simplifies to

$$((\mathbf{T} \cap \mathbf{T}) \Rightarrow (\mathbf{T} \cap \mathbf{T}) \Rightarrow \mathbf{T}) \;=\; ((\mathbf{T} \Rightarrow \mathbf{T}) \Rightarrow \mathbf{T}) \;=\; \mathbf{T}.$$

One can easily check that the other three cases are also correct.

We use set constraints to model RLL programs instead of Boolean logic for two reasons. First, although the core of RLL is Boolean logic, other instructions (e.g., control flow instructions) are at best difficult to express using Boolean logic. Second, RLL programs are large and complex, so approximations are needed for performance reasons. Set constraints give us the flexibility to model certain instructions less accurately and less expensively than others, making the analysis of RLL programs more manageable.

$$\begin{aligned}\rho(\bot) &= \emptyset \\ \rho(\top) &= \{\mathbf{T},\mathbf{F}\} \\ \rho(\mathbf{T}) &= \{\mathbf{T}\} \\ \rho(\mathbf{F}) &= \{\mathbf{F}\} \\ \rho(E_1 \cap E_2) &= \rho(E_1) \cap \rho(E_2) \\ \rho(E_1 \cup E_2) &= \rho(E_1) \cup \rho(E_2) \\ \rho(E_1 \Rightarrow E_2) &= \begin{cases}\rho(E_2) & \text{if } \rho(E_1) \neq \emptyset \\ \emptyset & \text{otherwise}\end{cases}\end{aligned}$$

Fig. 2. Semantics of set expressions

5 Constraint generation

In this section, we describe how we use inclusion constraints to model RLL programs. We also describe the concrete semantics and abstract semantics of RLL. The concrete semantics is given informally, while the abstract semantics is described formally with a set of constraint generation rules (see Figs. 3 and 4). It is straightforward to present a formal concrete semantics for RLL, except for timers which require the modeling of time. As we shall see, our abstraction of timers ignores time altogether. Thus, we avoid the complications of formalizing the semantics of timers by giving an informal concrete semantics.

Because of the scan evaluation model of RLL, it is sufficient to give a model of the meaning of a single scan. We give constraint generation rules for the primitive constructs of RLL. In the rules set variables denote the values of bits and wires. Thus, a bit or wire may be assigned the abstract values $\emptyset$ (meaning no value), $\{\mathbf{T}\}$ (definitely true), $\{\mathbf{F}\}$ (definitely false) or $\{\mathbf{T},\mathbf{F}\}$ (meaning either true or false, i.e., no information). Rules have the form

$$E, I \mapsto E', S, v_1, v_2$$

where:

- E and E' are mappings of bits to their corresponding set variables. The operator $+$ extends the mapping such that $(E + \{b, v\})(b') = \begin{cases} v, & \text{if } b' = b \\ E(b'), & \text{otherwise}\end{cases}$
- I is the current instruction;
- S is the set of constraints generated for this instruction;
- v_1 and v_2 are set variables associated with the wires before and after instruction I and are used to link instructions together.

The rule can be read as "under the variable mapping E, for the instruction I, the constraint set S is generated, along with a modified variable mapping E' and two set variables v_1 and v_2 denoting the wire preceding and following I respectively." As an example, consider the following rule for the instruction XIC.

$$\frac{\begin{array}{c}v_1 \text{ and } v_2 \text{ are fresh variables} \\ v_{ct} = E(b) \\ S = \{\mathrm{AND}(v_1, v_{ct}) \subseteq v_2\}\end{array}}{E, \mathrm{XIC}(b) \mapsto E, S, v_1, v_2}$$

where $\mathrm{AND}(v_1, v_{ct})$ denotes the set expression

$$\begin{aligned}((v_1 \cap \mathbf{T}) \Rightarrow (v_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup \\ ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup \\ ((v_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{F})\end{aligned}$$

The rule says that for the instruction XIC we generate the constraint $\mathrm{AND}(v_1, v_{ct}) \subseteq v_2$, with the variable mapping unchanged. In the rule, two fresh variables v_1 and v_2 are created to denote the wires preceding and following the instruction $\mathrm{XIC}(b)$. The statement $v_{ct} = E(b)$ is used to retrieve the set variable that is associated with the bit b from the mapping E.

Figures 3 and 4 give some example inference rules for generating the constraints describing the data and control flow of RLL programs. In Figs. 3 and 4, and in the rest

$$\frac{v_1 \text{ and } v_2 \text{ are fresh variables} \quad v_{ct} = E(b) \quad S = \{((v_1 \cap \mathbf{T}) \Rightarrow (v_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup ((v_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{F}) \ \subseteq v_2\}}{E, \mathrm{XIC}(b) \mapsto E, S, v_1, v_2} \quad [\mathrm{XIC}]$$

$$\frac{v_1 \text{ and } v_2 \text{ are fresh variables} \quad v_{ct} = E(b) \quad S = \{((v_1 \cap \mathbf{T}) \Rightarrow (v_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup ((v_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{F}) \ \subseteq v_2\}}{E, \mathrm{XIO}(b) \mapsto E, S, v_1, v_2} \quad [\mathrm{XIO}]$$

$$\frac{v_1, v_2, \text{ and } v_{ct} \text{ are fresh variables} \quad E' = E + \{(b, v_{ct})\} \quad S = \{((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \ \subseteq v_{ct}\}}{E, \mathrm{OTE}(b) \mapsto E', S, v_1, v_2} \quad [\mathrm{OTE}]$$

$$\frac{v_1, v_2, \text{ and } v_{ct} \text{ are fresh variables} \quad v'_{ct} = E(b) \quad E' = E + \{(b, v_{ct})\} \quad S = \{((v'_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow (v'_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{F}) \ \subseteq v_{ct}\}}{E, \mathrm{OTL}(b) \mapsto E', S, v_1, v_2} \quad [\mathrm{OTL}]$$

$$\frac{v_1, v_2, \text{ and } v_{ct} \text{ are fresh variables} \quad v'_{ct} = E(b) \quad E' = E + \{(b, v_{ct})\} \quad S = \{((v'_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup ((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{F}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow (v'_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{T}) \ \subseteq v_{ct}\}}{E, \mathrm{OTU}(b) \mapsto E', S, v_1, v_2} \quad [\mathrm{OTU}]$$

Fig. 3. Some rules for generating constraints (part 1)

of this section, we use w_1 and w_2 to stand for the wires preceding and following an instruction respectively. Furthermore, b denotes the bit referenced by an instruction unless specified otherwise. Below, we explain these rules in more detail.

Contacts.
The instruction XIC is called "Normally Closed Contact." If w_1 is true, then b is examined. If b is true, then w_2 is true. Otherwise, w_2 is false. In the rule [XIC], two fresh set variables v_1 and v_2 represent the two wires w_1 and w_2. The set variable v_{ct} represents the referenced bit b. The constraints express that w_2 is true if and only if both w_1 and b are true. The instruction XIO, called "Normally Opened Contact," is the dual of XIC. The wire w_2 is true if and only if w_1 is true and the referenced bit b is false. The constraint generation rule for XIO is similar to the rule [XIC].

Energise coil.
The instruction OTE or "Energise Coil" is programmed to control either an output connected to the controller or an internal bit. If the wire w_1 is true, then the referenced bit b is set to true. Otherwise, b is set to false. Rule [OTE] models this instruction. The set variables v_1 and v_2 are the same as in the rule [XIC]. The set variable v_{ct} is fresh, representing a new instance[1] of the referenced bit b. The new instance is recorded in the mapping E'. Later references to b use this instance. The constraints express that b is true if and only if w_1 is true.

Latches.
The instructions OTL and OTU are similar to OTE.

[1] Due to the sequential evaluation of rungs, a particular bit can take on distinct values in different parts of a program. An instance of a bit captures the state of a bit at a particular program point.

$$\frac{\begin{array}{c}v_1, v_2, v_{dn}, v_{en}, \text{ and } v_{tt} \text{ are fresh variables} \\ E' = E + \{(\mathrm{DN}, v_{dn}), (\mathrm{EN}, v_{en}), (\mathrm{TT}, v_{tt})\} \\ S = \left\{ \begin{array}{r} ((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T} \cup \mathbf{F}) \subseteq v_{dn}, \\ ((v_1 \cap \mathbf{T}) \Rightarrow (v_{dn} \cap \mathbf{F}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup ((v_{dn} \cap \mathbf{T}) \Rightarrow \mathbf{F}) \subseteq v_{tt}, \\ ((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \subseteq v_{en} \end{array} \right\}\end{array}}{E, \mathrm{TON} \mapsto E', S, v_1, v_2} \quad [\mathrm{TON]}$$

$$\frac{\begin{array}{c}v_1, v_2, v_{dn}, \text{ and } v_{cu} \text{ are fresh variables} \\ E' = E + \{(\mathrm{DN}, v_{dn}), (\mathrm{CU}, v_{en})\} \\ S = \left\{ \begin{array}{l} ((v_1 \cap \mathbf{T}) \Rightarrow (v_1 \cap \mathbf{F}) \Rightarrow \mathbf{T}) \cup \mathbf{F} \subseteq v_{dn}, \\ ((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \subseteq v_{cu} \end{array} \right\}\end{array}}{E, \mathrm{CTU} \mapsto E', S, v_1, v_2} \quad [\mathrm{CTU}]$$

$$\frac{\begin{array}{c}v_1, v_2, dv_i, 0 \le i \le 15, \text{ are fresh variables} \\ E' = E + \{(\mathrm{MOV}_{sw_i}, dv_i) \mid 0 \le i \le 15\} \\ S = \{((v_1 \cap \mathbf{T}) \Rightarrow E(\mathrm{MOV}_{dw_i}) \cup (v_1 \cap \mathbf{F}) \Rightarrow E(\mathrm{MOV}_{sw_i})) \subseteq dv_i \mid 0 \le i \le 15\}\end{array}}{E, \mathrm{MOV} \mapsto E', S, v_1, v_2} \quad [\mathrm{MOV}]$$

$$\frac{\begin{array}{c}B = \text{the set of bits in the program} \\ v_1, v_2, nv_b \text{ (for all } b \in B) \text{ are fresh variables} \\ R_{fname} = \text{the rungs in the file } \mathit{fname} \\ E, R_{fname} \mapsto E', S_0 \\ E'' = \{(b, nv_b) \mid b \in B\} \\ S = ((v_1 \cap \mathbf{T}) \Rightarrow S_0) \cup \{(v_1 \cap \mathbf{T}) \Rightarrow E'(b) \cup (v_1 \cap \mathbf{F}) \Rightarrow E(b) \subseteq nv_b \mid b \in B\}\end{array}}{E, \mathrm{JSR}_{fname} \mapsto E'', S, v_1, v_2} \quad [\mathrm{JSR}]$$

$$\frac{\begin{array}{c}v \text{ is a fresh variable} \\ E, R_1 \mapsto E', S_0, v_1, v_2 \\ E', R_2 \mapsto E'', S_1, v_1', v_2' \\ S = \{(v_2 \cap \mathbf{T}) \Rightarrow \mathbf{T} \cup (v_2' \cap \mathbf{T}) \Rightarrow \mathbf{T} \cup (v_2 \cap \mathbf{F}) \Rightarrow (v_2' \cap \mathbf{F}) \Rightarrow \mathbf{F} \subseteq v\}\end{array}}{E, R_1 || R_2 \mapsto E'', S \cup S_0 \cup S_1 \cup \{v_1 = v_1'\}, v_1, v} \quad [\mathrm{PAR}]$$

Fig. 4. Some rules for generating constraints (part 2)

OTL is "Latch Coil," and OTU is "Unlatch Coil." These two instructions appear in pairs. Once an OTL instruction activates its bit b, then b remains true until it is cleared by an unlatch instruction OTU, independently of the wire w_1 which activated the latch. The unlatch coil (OTU) instruction is symmetric. In the rule [OTL], the set variable v_{ct}' represents the value of the b prior to the instruction, while the variable v_{ct} denotes the new instance of b. The constraint expresses that b is true if and only the wire w_1 is true or b is true before evaluating this instruction. The rule for OTU is similar.

Timers.

Timers (TON) are instructions that activate an output after an elapsed period of time. Three status bits are associated with a timer: the *done bit* (DN), the *timing bit* (TT), and the *on bit* (EN). The DN bit is true if the wire w_1 has remained true for a preset period of time. The bit remains true unless w_1 becomes false. The TT bit is true if the wire w_1 is true and the DN bit is false. The TT bit is false otherwise, i.e., it is false if the wire w_1 is false or the DN bit is true. The EN bit is true if and only if the wire w_1 is true. In the rule [TON], v_{dn}, v_{tt} and v_{en} are fresh

set variables representing new instances of the corresponding bits. The constraint for the DN bit is

$$((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup \mathbf{F} \subseteq v_{dn}.$$

The constraint approximates timer operation by ignoring elapsed time. The DN bit can be false (the timer has not reached its preset period), or if the wire w_1 is true, then the DN bit can be true (the timer may have reached its preset period). The constraints for the TT and EN bits are straightforward.

Remark 1. For the relay race analysis, we assume that the DN bit does not change value across scans. This assumption is reasonable since the scan time, compared with the timer increments, is infinitesimal. The DN bit essentially becomes an input bit in the race analysis, and the constraint is accordingly simplified to $E(\mathrm{DN}) \subseteq v_{dn}$.

Counters.
A counter instruction has two associated status bits: the *done bit* (DN) as in timers and the *on bit* (CU). The DN bit becomes true if the wire w_1 has made a preset number of false to true transitions across scans. The CU bit is true if and only if the wire w_1 is true. In the rule [CTU], v_{dn} and v_{cu} are fresh set variables representing new instances of the corresponding status bits. The constraint for the CU bit is the same as that for a timer's EN bit. The constraint for the DN bit is

$$((v_1 \cap \mathbf{T}) \Rightarrow (v_1 \cap \mathbf{F}) \Rightarrow \mathbf{T}) \cup \mathbf{F} \subseteq v_{dn}.$$

Notice that for the DN bit to be true, the wire w_1 must have made at least one false to true transition. The variable that models the wire w_1 is v_1. The constraint says that if v_1 has both true and false, the DN bit could be either true or false. If v_1 does not have both true and false, the DN bit is definitely false. Again, we over-estimate the value of the DN bit.

Data transfers.
The MOV instruction is used for bit transfers. If the wire w_1 is true, the source (a 16 bit word) is moved into the destination (also a 16 bit word). If w_1 is false, no action is taken. The fresh variables $dv_i, 0 \le i \le 15$ are new instances for the 16 bits of the destination. dv'_i are the variables that represent the old values of the bits in the destination. The set variables sv_i represent the 16 bits of the source. The constraints are

$$\{(v_1 \cap \mathbf{T}) \Rightarrow sv_i \cup (v_1 \cap \mathbf{F}) \Rightarrow dv'_i \subseteq dv_i \mid 0 \le i \le 15\}$$

The constraints simply say that if the wire before is true then the source is moved to the destination, otherwise there is no transfer of bits.

Subroutines.
JSR is the subroutine call instruction. If the wire w_1 evaluates to true, the subroutine (a portion of ladder rungs with label *fname* as specified in the JSR instruction) is evaluated up to a return instruction, after which execution continues with the rung after the JSR instruction. If w_1 is false, execution continues immediately with the rung after the JSR instruction. In the rule [JSR], B denotes the set of all bits in a program. If S is a system of constraints and τ a set expression, then the notation $\tau \Rightarrow S$ abbreviates the constraints

$$\{\tau \Rightarrow \tau_0 \subseteq \tau_1 \mid (\tau_0 \subseteq \tau_1) \in S\}$$

The fresh variables nv_b represent new instances of all bits $b \in B$. Constraints S_0 are generated for the ladder rungs of the subroutine together with a modified mapping E'. The constraints

$$\{(v_1 \cap \mathbf{T}) \Rightarrow E'(b) \cup (v_1 \cap \mathbf{F}) \Rightarrow E(b) \subseteq nv_b \mid b \in B\}$$

merge the two instances of every bit b from the two possible control flows. If the wire w_1 (modeled by v_1) is true, then $E'(b)$ (the instance after evaluating the subroutine) should be the value of the current instance, otherwise, $E(b)$ is the value of the current instance.

Parallel wires.
The rule [PAR] describes the generation of constraints for parallel wires. Parallel wires behave the same as the disjunction of two Boolean variables, i.e., the wire after the parallel wires is true if any one of the two input wires is true. In the rule $v_1 = v'_1$ is an abbreviation for the two constraints $v_1 \subseteq v'_1$ and $v'_1 \subseteq v_1$. The fresh variable v is used to model the wire after the parallel wires. The constraint

$$\begin{aligned} (((v_2 \cap \mathbf{T}) &\Rightarrow \mathbf{T}) \cup \\ ((v'_2 \cap \mathbf{T}) &\Rightarrow \mathbf{T}) \cup \\ ((v_2 \cap \mathbf{F}) &\Rightarrow (v'_2 \cap \mathbf{F}) \Rightarrow \mathbf{F})) \subseteq v \end{aligned}$$

says that the wire after the parallel wires is true if one of the parallel wires is true. There are other rules for linking instructions together. These rules are similar to [PAR] and are also straightforward.

All solutions of the generated constraints conservatively approximate the evaluation of RLL programs. However, the best approximation is the least solution (in terms of set sizes). We now present a theorem which states that the constraints generated from an RLL program together with constraints for restricting the inputs have a least solution.

Theorem 1 (Existence of least solution). *For any RLL program $\mathcal{P}$, let S be the constraint system generated by the rules given in Figs. 3 and 4. Further let c be an input configuration for $\mathcal{P}$. The constraint system S together with the corresponding constraints of c has a least solution, Sol_{least}.*

Next, we state a soundness theorem of our model of RLL programs, namely that our model is a safe approximation of RLL.

Theorem 2 (Soundness). *Let $\mathcal{P}$ be an RLL program and S be the constraint system generated by the rules given in Figs. 3 and 4. Further let c be an input configuration for $\mathcal{P}$. The least solution Sol_{least} to the constraint system S together with the constraints restricting the inputs* safely approximates *the values of the wires and bits in one scan, meaning that if an instance of a bit or a wire is true (false) in an actual scan, then true (false) is a value in the set representing this instance.*

Theorem 1 and Theorem 2 are proven in [26].

6 Relay race analysis

In this section, we describe our analysis for detecting relay races in RLL programs. In RLL programs, it is desirable if the values of outputs depend solely on the values of inputs and the internal states of timers and counters. If under fixed inputs and timer and counter states, an output b changes from scan to scan, then there is a relay race on b.

Before describing our analysis, we give a more formal definition of the problem. Consider an RLL program P. Let **IN** denote the set of inputs, and let **OUT** denote the set of outputs[2]. Formally, an RLL program P is a function mapping $\mathbf{IN} \to \{\mathbf{T}, \mathbf{F}\}$ to $\mathbf{OUT} \to \{\mathbf{T}, \mathbf{F}\}$. Here we follow the convention that for any two sets S and T, $S \to T$ denotes the set of total functions from S to T. Let $C = \mathbf{IN} \to \{\mathbf{T}, \mathbf{F}\}$ denote the set of all possible input configurations. Further, let

$$\Psi_i : \mathbf{OUT} \to \{\mathbf{T}, \mathbf{F}\}$$

be the mapping from the set of outputs to their corresponding values at the end of the i-th scan.

Definition 1. *An RLL program P is* race free *if for any input configurations $c \in C$, by fixing c, it holds that for all $i \geq 1, \Psi_i = \Psi_1$. Otherwise, we say the program has a race.*

Definition 1 states under what conditions a program exhibits a race. Note that this definition assumes that outputs should stabilize after a single scan.

For any set S, we denote its powerset by $\wp(S)$.

Definition 2. *Let*

$$P : (\mathbf{IN} \to \{\mathbf{T}, \mathbf{F}\}) \to (\mathbf{OUT} \to \{\mathbf{T}, \mathbf{F}\})$$

be an RLL program. An approximation

$$A : (\mathbf{IN} \to \wp(\{\mathbf{T}, \mathbf{F}\})) \to (\mathbf{OUT} \to \wp(\{\mathbf{T}, \mathbf{F}\}))$$

is an abstraction of P such that, for any configuration $c : \mathbf{IN} \to \{\mathbf{T}, \mathbf{F}\}$ and bit $b \in \mathbf{OUT}$ of P, at the end of any scan, the following condition holds: $P_c(b)$ (the value of b in the program P) is contained in $A_c(b)$ (the value of b in the abstraction A), i.e., $P_c(b) \in A_c(b)$.

Let A be an approximation of P. Let

$$\Phi_i : \mathbf{OUT} \to \wp(\{\mathbf{T}, \mathbf{F}\})$$

be the mapping from the set of outputs to their corresponding values at the end of the i-th scan in A.

Definition 3. *An approximation A of an RLL program P is* race free *if for any fixed initial input configuration $c \in C$, and the resulting infinite sequence of abstract scans $S_1, S_2, S_3, \ldots$, there exists*

$$\Psi^* : \mathbf{OUT} \to \{\mathbf{T}, \mathbf{F}\}$$

such that $\Psi^*(b) \in \Phi_i(b)$, for all $b \in \mathbf{OUT}$ and $i \geq 1$.

Lemma 1. *Let P be an RLL program and A an approximation of P. If P is race free, then so is A. In other words, if A exhibits a race, so does P.*

Lemma 1 states that if our analysis detects a race under some input c, then the program will exhibit a race under input c. We now deal with the problem of detecting races in our approximation of RLL programs.

Theorem 3. *For any approximation A of an RLL program P and input $c \in C$, the approximation A races under c if and only if there exists $b \in \mathbf{OUT}$ such that $\bigcap_{i \geq 1} \Phi_i(b) = \emptyset$.*

Since two scans are necessary and sufficient for detecting any races in an RLL program, one may suspect that the same holds for any abstract model of the program as well.

Conjecture 1. Let A be an approximation of a program P. If A has a race under the input configuration c, then there exists an input configuration c', under which

$$\Phi_1(b) \cap \Phi_2(b) = \emptyset$$

for some bit b.

Surprisingly, Conjecture 1 does not hold, and we give a counter example. Consider the example given in Fig. 6. The truth table representation of a program is given in Fig. 6a, and that for its approximation is given in Fig. 6b. For the approximation, only under the input configuration

$$\{x = \mathbf{F}, y = \mathbf{T}\}$$

does the approximation have a race, exhibited with three scans:

$$\{x = \mathbf{F}, y = \mathbf{T}\} \xrightarrow{1} \{x = \mathbf{T}, y = \top\} \xrightarrow{2} \{x = \top, y = \mathbf{F}\} \xrightarrow{3} \{x = \mathbf{F}, y = \mathbf{F}\}$$

[2] Note that **IN** = set of external inputs + internal bits, and **OUT** = set of external outputs + internal bits.

where $\xrightarrow{i}$ denotes the transition of the i-th scan. Notice that the race on x is detected after the third scan. For the other three input configurations, the approximation does not exhibit a race. Thus, the conjecture does not hold.

In principle, for any given input assignment, it may be necessary to simulate scans until a repeating sequence of output configurations is detected, which may require a number of scans exponential in the number of inputs. However, the following lemma shows that two scans are sufficient to uncover the common case.

Lemma 2. *Let A be an approximation of a program P. If A has a race of bit b under input configuration c, such that $\Phi_i(b) \cap \Phi_{i+1}(b) = \emptyset$ for some scan i, then there exists another input configuration c' such that $\Phi_1(b) \cap \Phi_2(b) = \emptyset$ under c', i.e., it is sufficient to use two scans on every input configuration to uncover the race on b.*

We have verified experimentally that performing only two scans works well; an experiment in which we performed ten scans per initial input configuration detected no additional races. Theorem 3 and Lemma 2 thus lead naturally to the algorithm in Fig. 5 for detecting relay races. The general strategy for the analysis is:

1. Generate the base system of constraints using the constraint generation rules presented in Sect. 5.
2. Add constraints that assign random bits to the inputs.
3. Check whether the program races under this input assignment.
4. Repeat 2–3.

We make the assumption that all input bits are independently assigned **T** or **F** uniformly at random. Under this assumption, all input assignments are possible. In practice, because of the nature of the external devices and sensors the program interacts with, there may be dependencies between inputs that make some input configurations unrealizable. Our analysis can be made more accurate by excluding unrealizable configurations, if information about these dependencies is available.

We use the example in Fig. 1 to demonstrate how the race detection algorithm works. Consider the last two rungs in the example RLL program in isolation. The base system for these two rungs is given in the top of Fig. 7.

```
1   for every output b
2       B_sum(b) := {T, F};
3   S_input := random assignment;
4   for Scan := 1 to 2
5       B_current := Sol_least(S_base ∪ S_input);
6       S_input := GetInput(B_current);
7       B_sum := B_sum ∩ B_current;
8       if B_sum(b) = ∅ for some output b
9           then output b is racing;
```

Fig. 5. Algorithm for detecting races

x	y	x'	y'
F	**F**	**F**	**F**
F	**T**	**T**	**F**
T	**F**	**F**	**F**
T	**T**	**F**	**F**

(a) Concrete Program

x	y	x'	y'
F	**F**	**F**	**F**
F	$\top$	$\top$	$\top$
F	**T**	**T**	$\top$
$\top$	**F**	**F**	**F**
$\top$	$\top$	$\top$	$\top$
$\top$	**T**	$\top$	$\top$
T	**F**	**F**	**F**
T	$\top$	$\top$	**F**
T	**T**	$\top$	**F**

(b) Approximation

Fig. 6. Truth tables for an example and an approximation

Assume the bit B is initially true. Adding the constraint $\mathbf{T} \subseteq b_{B_0}$ to the base system and solving the resulting system, we obtain its least solution at the end of the first scan (column 3 in Fig. 7). We see that at the end of the first scan, the bit B is false. In the second scan, we add the constraint $\mathbf{F} \subseteq b_{B_0}$ to the base system. The resulting system is solved, and its least solution is shown in column 4 of Fig. 7. We intersect the values of the output bits, i.e., bits B (the last instance) and C, in the least solutions from the first two scans. Since the intersections are empty, we have detected a race.

If our analysis finds a race, then the program does indeed exhibit a race. The absence of races cannot be proven by our analysis due to approximations and due to the finite subspace of input assignments we sample. However, we can analyze the coverage of our random sampling approach using the well-known Coupon Collector's Problem. Consider a hat containing n distinct coupons. In a trial a coupon is drawn at random from the hat, examined, and then placed back in the hat. We are interested in the expected number of trials needed to select all n coupons at least once. One can show that the expected number of trials is $n \ln n + \mathcal{O}(n)$, and that the actual number of trials is sharply concentrated around this expected value (for any constant $c > 0$, the probability that after $n(\ln n + c)$ trials there are still coupons not selected is approximately $1 - e^{-e^{-c}}$). Notice that $1 - e^{-e^{-c}} \approx 0.05$ when $c = 3$, and this probability is independent of n. See Appendix A for more details on the Coupon Collector's Problem.

Recall that we assume that the inputs bits are independently assigned **T** or **F** uniformly at random. Therefore, any assignment of n input bits restricted to $k \leq n$ bits corresponds to an input assignment, selected uniformly at random, of these k bits. Thus, we have the following theorem, which states that without many trials, any race depending on a small number of inputs is detected with high probability.

Theorem 4. *Using the analysis of the Coupon Collector's Problem, after approximately $2^k \ln(2^k + 3)$ random*

$$\begin{aligned}
\mathbf{T} &\subseteq w_0 \\
((\mathbf{T} \cap b_{B_0}) \Rightarrow \mathbf{T}) \cup ((\mathbf{F} \cap b_{B_0}) \Rightarrow \mathbf{F}) &\subseteq w_1 \\
((\mathbf{T} \cap w_1) \Rightarrow \mathbf{T}) \cup ((\mathbf{F} \cap w_1) \Rightarrow \mathbf{F}) &\subseteq w_2 \\
((\mathbf{T} \cap w_2) \Rightarrow \mathbf{T}) \cup ((\mathbf{F} \cap w_2) \Rightarrow \mathbf{F}) &\subseteq b_C \\
\mathbf{T} &\subseteq w_3 \\
((\mathbf{T} \cap b_{B_0}) \Rightarrow \mathbf{F}) \cup ((\mathbf{F} \cap b_{B_0}) \Rightarrow \mathbf{T}) &\subseteq w_4 \\
((\mathbf{T} \cap w_4) \Rightarrow \mathbf{T}) \cup ((\mathbf{F} \cap w_4) \Rightarrow \mathbf{F}) &\subseteq w_5 \\
((\mathbf{T} \cap w_5) \Rightarrow \mathbf{T}) \cup ((\mathbf{F} \cap w_5) \Rightarrow \mathbf{F}) &\subseteq b_{B_1}
\end{aligned}$$

bit or wire	*variable*	*value after the first scan*	*value after the second scan*
wire preceding XIC(B)	w_0	**T**	**T**
wire following XIC(B)	w_1	**T**	**F**
wire preceding OTE(C)	w_2	**T**	**F**
wire preceding XIO(C)	w_3	**T**	**T**
wire following XIO(C)	w_4	**F**	**T**
wire preceding OTE(B)	w_5	**F**	**T**
first instance of B	b_{B_0}	**T**	**F**
last instance of B	b_{B_1}	**F**	**T**
the bit C	b_C	**T**	**F**

Fig. 7. Base system for the last two rungs of the example program in Fig. 1 with the least solutions at the end of the first and the second scans given in the table

samples, any race depending on a fixed set of k or fewer inputs has been detected with high probability (95%), up to the approximations due to conservative analysis and performing only two scans.

Note that the expected number of trials depends only on the number of inputs participating in the race, not on the total number of inputs. For example, the number of trials required to find races involving 5 inputs with 95% probability is 200 whether there are 100, 1000, or 10000 inputs to the program.

One alternative to random trials is the approach taking by, for example, logic programming. A program is described as a set of logic formulae. One can query for what values, if any, the formulae is satisfied. This approach can be easily adapted to find races. Our base system corresponds to a set of ternary logic formulae, which can be represented as a ternary function f. In principle, we can compute the function $f \circ f$, which is f composed with f itself. We then construct a goal formula g describing that the program has a race. Finally we ask whether there is a satisfying assignment for the formulae f, $f \circ f$, and g, i.e., whether $f \wedge (f \circ f) \wedge g$ is satisfiable. Although possible in principle, we suspect that this approach is too expensive. Since we need to compute the composition $f \circ f$, the resulting formula is potentially very large. Thus, it may be quite expensive to apply BDD-based (Binary Decision Diagram) or SAT-based decision procedures to find satisfying assignments for $f \wedge (f \circ f) \wedge g$.

7 Experimental results

We have implemented our analysis using a general constraint solver [2]. Inputs to our analysis are abstract syntax tree (AST) representations of RLL programs. The ASTs are parsed into internal representations, and constraints are generated using the rules in Figs. 3 and 4. The resulting constraints are solved and simplified to obtain the base system.

7.1 Benchmarks

Four large RLL programs were made available to us in AST form for evaluating our analysis.

- **Mini Factory**
  This is an example program written and used by RLL programmers and researchers working on tools for RLL programming.
- **Big Bak**
  This is a production RLL program.
- **Wdsdflt(1)**
  Another production application, this program has a known race.
- **Wdsdflt(2)**
  This program is a modified version of Wdsdflt(1) with the known race eliminated. The program is included for comparing its results with the results from the original program.

Program	Size	#Vars.	Secs/Scan	Ext. Races	Int. Races	#Samples	Time (s)
Mini Factory	9,267	4,227	0.4	55	186	1,000	844
Big Bak	32,005	21,596	4.0	4	6	1,000	7,466
Wdsdflt(1)	58,561	22,860	3.0	8	163	1,000	7,285
Wdsdflt(2)	58,561	22,860	3.0	7	156	1,000	7,075

Fig. 8. Benchmark programs for evaluating our analysis

Figure 8 gives a table showing the size of each program as number of lines in abstract syntax tree form, number of set variables in the base system, and the time to analyze one scan. All measurements reported here were done on a Sun Enterprise 5000 with 512 MB of main memory (using only one of the eight processors).

Our analysis discovered many relay races in these programs. The results are presented in Fig. 8. For each program, we show the number of external racing bits (bits connected to external outputs), the number of internal racing bits (bits internal to the program), the number of samples, and the total analysis time in seconds. By Theorem 4, 1000 trials are sufficient to uncover races involving 7 or fewer inputs.

No relay races were known for the Mini Factory program. Our analysis detected 55 external races, some of which were subsequently verified by running a model factory under the corresponding inputs. Fewer races were found in Big Bak, even though it is a much larger program. Two likely reasons for this situation are that Big Bak uses arithmetic operations heavily (which our analysis approximates rather coarsely) and that Big Bak is a production program and has been more thoroughly debugged than Mini Factory. Our analysis discovered the known relay race in Wdsdflt(1) (fixed in Wdsdflt(2)) among 8 external and 163 internal races. Note that some of the reported races may be unrealizable if they depend on input configurations that cannot occur in practice.

8 Related work

In this section, we discuss the relationship of our work to work in ternary simulation of combinatorial and sequential circuits, data flow analysis, model checking, and testing.

Ternary simulation. Ternary simulation was introduced by Yoeli and Rinon [27] to analyze static hazards using three-valued logic [21] in combinatorial circuits. The method was extended by Eichelberger [16] to handle general hazards in combinatorial circuits, and races and oscillations in sequential circuits. The method was further developed by Brzozowski and Yoeli [10] at the gate level and by Bryant [9] at the transistor level.

Besides 0 (false) and 1 (true), three-valued logic has an additional value $\frac{1}{2}$ having the informal meaning "unknown", "don't care", or "transient" depending on the context. The value $\frac{1}{2}$ corresponds to $\{\mathbf{T}, \mathbf{F}\}$ in our semantic domain. One slight difference is the use of $\emptyset$ in our semantic domain. One might argue that since for the relay race analysis, each bit or wire can only be assigned an abstract value $\{\mathbf{T}\}$, $\{\mathbf{F}\}$, or $\{\mathbf{T}, \mathbf{F}\}$, the value $\emptyset$ is never used. Thus, there is no fundamental difference. However, one can imagine that $\emptyset$ is useful for finding uninitialized values.

It appears that ternary simulation has been used exclusively for circuit analysis. While RLL has a circuit programming metaphor, it is really a fairly complete programming language with `goto`, procedures, and pointers. Although it is possible to model these constructs with ternary logic formulae, constraints give a much more natural model.

Data flow analysis. Data flow analysis is used primarily in optimizing compilers to collect variable usage information for optimizations such as dead code elimination and register allocation [1]. It has also been applied for ensuring software reliability [17, 18]. Our approach differs from classical data flow analysis in two points. First, we use conditional constraints [4], which are essential for modeling both the Boolean instructions and control flow instructions. Second, the use of constraints gives us the flexibility to analyze many input configurations by adding constraints to a base system, instead of performing a global dataflow analysis repeatedly. Our approach is more efficient because the base system can be solved and simplified once and then used repeatedly on different input configurations.

Model checking. Model checking [12, 13] is a branch of formal verification that can be fully automated. Model checking has been used successfully for verifying finite state systems such as hardware and communication protocols [7, 8, 14, 15, 20]. Model checkers exploit the finite nature of these systems by performing exhaustive state space searches. Because even these finite state spaces may be huge, model checking is usually applied to some abstract models of the actual system. These abstract systems are symbolically executed to obtain information about the actual systems. Our analysis for RLL programs is similar to model checking in that our abstract models are finite, whereas RLL programs are in general infinite state systems. Similar to model checking, we make tradeoffs between modeling accuracy and efficiency, our abstraction approximates timers, counters, and arithmetic. It is through these approximations that we obtain a simpler analysis that is practical for production codes. On the other hand, due to these approximations our analysis cannot guarantee the absence of errors. Our approach differs from model checking in the way abstract models

are obtained. In model checking, abstract models are often obtained manually, while our analysis automatically generates the model.

Testing. Testing is one of the most commonly used methods for assuring hardware and software quality. The I/O behaviors of the system on input instances are used to deduce whether the given system is faulty or not [23]. Testing is non-exhaustive in most cases due to a large or infinite number of test cases. One distinction of our approach from testing is that we work with an abstract model of the actual system. There are advantages and disadvantages to using an abstract model. A disadvantage is that there is loss of information due to abstraction. As a result, the detection of an error may be impossible, whereas testing the actual system would show the incorrect I/O behavior. Abstract models have the advantage that a much larger space of possible inputs can be covered, which is important if the set of inputs exhibiting a problem is a tiny fraction of all possible inputs. An abstract model is also advantageous when it is very difficult or very expensive to test the actual system. Both of these advantages of abstract modeling apply in the case of detecting relay races in RLL programs. [11] discusses some other tradeoffs of using the actual system and abstract models of the system for testing.

9 Conclusion

In this paper, we have described a relay race analysis for RLL programs to help RLL programmers detect some common programming mistakes. We have demonstrated that the analysis is useful in statically catching such programming errors. Our implementation of the analysis is accurate and fast enough to be practical –

production RLL programs can be analyzed. The relay race analysis not only detected a known bug in a program that took an RLL programmer four hours of factory down-time to uncover, it also detected many previously unknown relay races in our benchmark programs.

Acknowledgements. We would like to thank Jim Martin for bringing RLL to our attention and for making this work possible. We would also like to thank Anthony Barrett for information on RLL, providing us with abstract syntax trees of RLL programs, and running some experiments to validate our results. Finally, we thank the anonymous referees for the helpful comments.

References

1. Aho, A.V., Sethi, R., Ullman, J.D.: Compilers, Principles, Techniques and Tools. Reading, MA: Addison-Wesley, 1986
2. Aiken, A., Fähndrich, M., Foster, J., Su, Z.: A toolkit for constructing type- and constraint-based program analyses. In: Proc. 2nd Int. Workshop on Types in Compilation (TIC '98), pp. 78–96, March 1998
3. Aiken, A., Wimmers, E.: Type inclusion constraints and type inference. In: Proc. 1993 Conference on Functional Programming Languages and Computer Architecture, pp. 31–41, Copenhagen, Denmark, June 1993
4. Aiken, A., Wimmers, E., Lakshman, T.K.: Soft typing with conditional types. In: 21st Annual ACM Symposium on Principles of Programming Languages, pp. 163–173, Portland, OR, January 1994
5. Allen–Bradley, Rockwell Automation.: SLC 500 and MicroLogix 1000 Instruction Set
6. Barrett, A.: Private communication
7. Browne, M., Clarke, E.M., Dill, D.: Checking the correctness of sequential circuits. In: Proc. IEEE Int. Conf. on Computer Design, pp. 545–548, 1985
8. Browne, M., Clarke, E.M., Dill, D., Mishra, B.: Automatic verification of sequential circuits using temporal logic. IEEE Trans. Comput. 35(12): 1035–1044, 1986
9. Bryant, R.E.: Boolean analysis of mos circuits. IEEE Transactions on Computer-aided Design 6(4): 634–649, July 1987
10. Brzozowski, J.A., Yoeli, M.: On a ternary model of gate networks. IEEE Trans. Comput. C-28: 178–184, 1979
11. Carver, R.H., Durham, R.: Integrating formal methods and testing for concurrent programs. In: Proc. 10th Annual Conference on Computer Assurance, pp. 25–33, New York, June 1995
12. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching time temporal logic. In: Proc. Workshop on Logics of Programs 131. Berlin, Heidelberg, New York: Springer-Verlag, 1981, pp. 52–71
13. Clarke, E.M., Emerson, E.A., Sistla, A.P.: Automatic verification of finite-state concurrent systems using temporal logic specifications. ACM Transactions on Programming Languages and Systems 8(2): 244–263, 1986
14. Clarke, E.M., Grumberg, O., Hiraishi, H., Jha, S., Long, D.E., McMillan, K.L., Ness, L.A.: Verification of the futurebus+ cache coherence protocol. In: Claesen, L. (ed.): Proc. 11th Int. Symposium on Computer Hardware Description Languages and their Applications, North-Holland, April 1993
15. Dill, D., Clarke, E.M.: Automatic verification of asynchronous circuits using temporal logic. In: Proc. IEEE 133: 276–282, 1986
16. Eichelberger, E.B.: Hazard detection in combinatorial and sequential switching circuits. IBM J. Res. Div. 9: 90–99, 1965
17. Fosdick, L.D., Osterweil, L.J.: Data flow analysis in software reliability. ACM Computing Surveys 8(3): 305–330, September 1976
18. Harrold, M.J.: Using data flow analysis for testing. Technical Report 93-112, Department of Computer Science, Clemson University, 1993
19. Heintze, N.: Set Based Program Analysis. PhD. thesis, Carnegie Mellon University, 1992
20. Holzmann, G.: Design and Validation of Computer Protocols. Englewood Cliffs, NJ: Prentice-Hall, 1991
21. Kleene, S.C.: On a notation for ordinal numbers. J. Symbolic Ligic 3: 150–155, 1938
22. Krigman, A.: Relay ladder diagrams: we love them, we love them not. In: Tech, pp. 39–47, October 1985
23. Lee, D., Yannakakis, M.: Principles and methods of testing finite state machines-a survey. In: Proc. IEEE, pp. 1090–1123, August 1996
24. Milner, R., Tofte, M., Harper, R.: The Definition of Standard M. MIT Press, 1990
25. Motwani, R., Raghavan, P.: Randomized Algorithms. Cambridge University Press, 1995
26. Su, Z.: Automatic analysis of relay ladder logic programs. Technical Report UCB/CSD-97-969, University of California at Berkeley, 1997
27. Yoeli, M., Rinon, S.: Application of ternary algebra to the study of static hazards. J. ACM 11: 84–97, 1964

Appendix A: The Coupon Collector's Problem

In the Coupon Collector's Problem, there are n different coupons. At each trial a coupon is drawn uniformly at random. The selected coupon is put back with the rest of the coupons after it has been examined. We are interested

in the expected number of trials needed to select all of the n coupons.

Theorem 5 (Expected Value). *The expected number trials to select all the n coupons is $n \ln n + \mathcal{O}(n)$.*

Proof. Let X be a random variable defined to be the number of trials needed to collect all of the n coupons. Define a *success* to be a trial in which a new coupon is collected. Define the random variables X_i, for $0 \leq i \leq n-1$, to be the number of trials that follows the i-th success and ends on the trial that collects the $(i+1)$-th coupon. Thus, we have

$$X = \sum_{i=0}^{n-1} X_i.$$

Let p_i be the probability of success on any trial after the i-th coupon has been collected. This is the probability of drawing one of $n-i$ coupons from a pool of n coupons, so that

$$p_i = \frac{n-i}{n}.$$

The random variable X_i is geometrically distributed with parameter p_i. Thus, its expectation

$$\mathbf{E}[X_i] = \frac{1}{p_i} = \frac{n}{n-i}.$$

By linearity of expectation, we have that

$$\begin{aligned}\mathbf{E}[X] &= \mathbf{E}[\sum_{i=0}^{n-1} X_i] \\ &= \sum_{i=0}^{n-1} \mathbf{E}[X_i] \\ &= \sum_{i=0}^{n-1} \frac{n}{n-i} \\ &= n \sum_{i=1}^{n} \frac{1}{i} \\ &= nH_n\end{aligned}$$

where H_n is the n-th Harmonic number. Since $H_n = \ln n + \Theta(1)$, we have

$$\mathbf{E}[X] = n \ln n + \mathcal{O}(n).$$

The next theorem states that the actual value is sharply concentrated around this expected value.

Theorem 6 (Sharp Threshhold). *Let the random variable X denote the number of trials for collecting each of the n types of coupons. We have, for any real constant c, and $m = n \ln n + cn$,*

$$\lim_{n\to\infty} \mathbf{Pr}[X > m] = 1 - e^{-e^{-c}}.$$

A proof for the above theorem can be found in [25].

Appendix B: Proofs of Lemmas and Theorems

B.1 Proof of Lemma 1

Proof. Since P is race free, by Definition 1, we have $\Psi_i = \Psi_1$ for all $i \geq 1$. Since A is an approximation of P, by Definition 2, $\Psi_i(b) \in \Phi_i(b)$ for all $i \geq 1$. Thus, $\Psi_1(b) \in \Phi_i(b)$ for all $i \geq 1$, and by Definition 3, the approximation A is also race free.

B.2 Proof of Theorem 3

Proof. Let $b \in \mathbf{OUT}$ be an output such that

$$\bigcap_{i\geq 1} \Phi_i(b) = \emptyset.$$

Since A is an approximation of the program P, we have $\Phi_i(b) \neq \emptyset$. Thus, there exist positive integers $i \neq j$ such that $\Phi_i(b) = \{\mathbf{T}\}$ and $\Phi_j(b) = \{\mathbf{F}\}$. Therefore, there does not exist a $\Psi^* : \mathbf{OUT} \rightarrow \{\mathbf{T}, \mathbf{F}\}$ such that $\Psi^*(b) \in \Phi_i(b)$ for all $b \in \mathbf{OUT}$ and for all $i \geq 1$. Hence, A has a race under c.
Conversely, suppose for all $b \in \mathbf{OUT}$, $\bigcap_{i\geq 1} \Phi_i(b) \neq \emptyset$ holds. Then, let $\Phi(b) = \bigcap_{i\geq 1} \Phi_i(b)$ for all $b \in \mathbf{OUT}$. Clearly there exists a $\Psi^* : \mathbf{OUT} \rightarrow \{\mathbf{T}, \mathbf{F}\}$ such that $\Psi^*(b) \in \Phi(b)$ for all $b \in \mathbf{OUT}$. Therefore, A does not race under input c.

B.3 Proof of Lemma 2

Proof. Let $\Phi_i^c(b)$ denote the value of b at the end of the ith scan starting with input configuration c. Without loss of generality, assume $\Phi_i^c(b) = \{\mathbf{T}\}$ and $\Phi_{i+1}^c(b) = \{\mathbf{F}\}$. Consider the values of the inputs c_i prior to scan i. Now choose any configuration c', s.t. $c'(b) \subseteq c_i(b)$ for all b. Since our analysis is monotone in the input (Theorem 1), we have $\Phi_1^{c'}(b) = \{\mathbf{T}\}$ and $\Phi_2^{c'}(b) = \{\mathbf{F}\}$. Hence, the race on bit b can be detected within two scans, starting from a configuration c'.