Automatic Analysis of Relay Ladder Logic Programs

Zhendong Su

Report No. UCB/CSD-97-969

September 1997

Computer Science Division (EECS)
University of California
Berkeley, California 94720

Automatic Analysis of Relay Ladder Logic Programs*

Zhendong Su
EECS Department
University of California, Berkeley
Berkeley, CA 94720-1776
zhendong@cs.berkeley.edu

Abstract

Relay Ladder Logic (RLL) [4] 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. We demonstrate that our analyses are useful in detecting some common programming mistakes and can be easily extended to perform other kinds of analysis for RLL programs such as some of the analyses described by Barrett [6].

1 Introduction

Programming logic controllers (PLC's) are control development systems used extensively in manufacturing industries for complex embedded control applications such as factory control and for entertainment equipment such as amusement park rides. 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 [5].

RLL has long been criticized for its low level design, which makes it difficult to write correct RLL programs [19]. Moreover, validation of RLL programs is extremely expensive, often measured in millions of dollars (for shutting down a factory) or human safety (for rides). One solution is to replace RLL with a higher-level, safer programming language. An alternative is to provide direct programming support for RLL. Since there are many existing RLL applications, and many more will be written in this language, we consider this 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 difficult, if not impossible, to detect manually. 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 cost several million dollars measured in factory down-time [5].

Our analyses are *constraint-based*, meaning that the information we wish to know about a program is expressed as constraints [17, 2, 3]. The solutions of these constraints yield the desired information. Our analyses are built using a generic constraint resolution engine, which allows our analyses to be expressed very directly. Constraint-based program analysis is discussed further in Section 2.

RLL programs are represented as *ladder diagrams*, which are a stylized form of a circuit or data flow diagram. 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 Figure 1. In the example, there are two vertical rails. The one on the left supplies power to all crossing rungs of the ladder. The five horizontal lines are the ladder rungs of this program. This example has four kinds of

*The research was funded in part by the National Science Foundation, Grant No. CCR-9416973, and by NSF Infrastructure Grant No. CDA-9401156, and supported in part by a gift from Rockwell Corporation. The information presented here does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred.

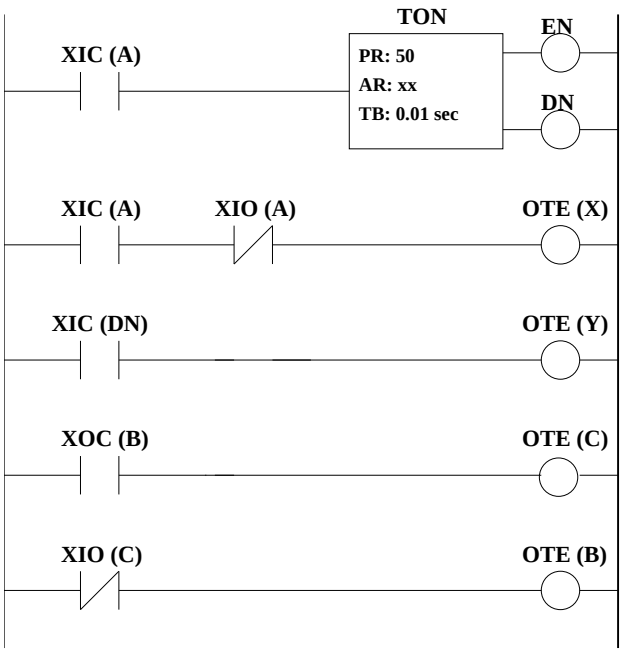

Figure 1: An example RLL program.

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 (or on) if power is supplied to the wire, and the wire is false (or off) otherwise.

An RLL program operates by first reading in all the values of the external input bits and executing the rungs in sequence from top to bottom and left to right. Program control instructions may cause portions of the program be skipped or repeatedly executed. After the last rung is evaluated, all the real 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. Thus, RLL programs are examples of reactive systems. 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. A grammar for the subset of RLL discussed in this report is in Figure 2.

Examples of relay instructions are XIC, XIO, and OTE. We briefly describe how these three instructions work for the explanation of our analyses. Let w_1 and w_2 be the wires before an instruction 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.

$$
\begin{array}{rcl}
\mathit{program} & ::= & \mathit{ladder\_files} \\
\mathit{ladder\_files} & ::= & \mathit{ladder\_files}\ \mathit{ladder\_file} \mid \mathit{ladder\_file} \\
\mathit{ladder\_file} & ::= & \mathit{rungs} \\
\mathit{rungs} & ::= & \mathit{rungs}\ \mathit{rung} \mid \epsilon \\
\mathit{rung} & ::= & \mathit{input\_list}\ \mathit{output\_list} \\
\mathit{input\_list} & ::= & \mathit{instruction}\ \mathit{input\_list} \mid \mathit{input\_branch}\ \mathit{input\_list} \mid \epsilon \\
\mathit{input\_branch} & ::= & \mathit{input\_level}\ \mathit{input\_list} \\
\mathit{input\_level} & ::= & \mathit{input\_level}\ \mathit{input\_list} \mid \mathit{input\_list} \\
\mathit{output\_list} & ::= & \mathit{instruction} \mid \mathit{output\_branch} \\
\mathit{output\_branch} & ::= & \mathit{output\_branch}\ \mathit{input\_list}\ \mathit{output\_list} \mid \epsilon \\
\mathit{instruction} & ::= & \texttt{XIC} \mid \texttt{XIO} \mid \texttt{OTE} \mid \texttt{OTL} \mid \texttt{OTU} \mid \texttt{TON} \mid \texttt{CTU} \mid \texttt{MOV} \mid \texttt{JSR}
\end{array}
$$

Figure 2: Grammar of the ladder language.

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.

In this paper, we describe the design and implementation of our RLL program analyzer. Currently the analyzer performs two different analyses. One is *constant wire analysis*, in which the analyzer detects if there is any wire in a given program that is always true or always false during the execution of a program. Constant wires indicate possible programming mistakes, because it is unlikely that a programmer would intentionally write constant-valued circuits. If a wire is always true or always false, there is no reason to put any instructions before these wires. For example, in the program in Figure 1, if we know that the wire after the XIO(A) instruction in the second rung is always false, then the two instructions XIC(A) and XIO(A) are superfluous.

Our second analysis detects *relay races*. 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. For example, in the program in Figure 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 analyses are a generalization of traditional data flow analyses [1]. Instead of data flow equations, set constraints [17, 2, 3] are used. Set constraints are more expressive than data flow equations since the constraints can model not only the data flow but also the control flow of a program.

Our analyses consist of two steps. In the first step, we generate constraints that describe the data and control flow dependences of an RLL program. The constraints are generated in a top-down traversal of the abstract syntax tree (AST) of the program. According to a set of constraint generation rules (see Section 4), appropriate constraints are generated for each AST node. These data and control flow constraints are solved to yield another set of simplified constraints. We call the set of resulting constraints the *base system*. The base system models where and how a value flows in the program. For example, the constraints S in Figure 3 are produced for the third rung of the example program in Figure 1.

The constraints S are solved and reduce to the constraints S' in Figure 3. The base system is a *conservative approximation* of the program: 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 is analysis-specific. For constant wire analysis, we use two different approaches. In the first approach, we constrain every input by both true and false and add the corresponding constraints to the base system. The resulting system is then solved, and the minimum solution is extracted. If in the minimum solution a wire w is not both true and false, we are sure that w is constant since the base system is a conservative approximation of the program. In the second approach, we use random sampling of input

$$S = \left\{ \begin{array}{rcl} & & (\mathbf{T} \in w_1) \\ (\mathbf{T} \in w_1) \Rightarrow (\mathbf{T} \in b_{DN}) & \Rightarrow & (\mathbf{T} \in w_2) \\ (\mathbf{F} \in w_1) & \Rightarrow & (\mathbf{F} \in w_2) \\ (\mathbf{F} \in b_{DN}) & \Rightarrow & (\mathbf{F} \in w_2) \\ (\mathbf{T} \in w_3) & \Rightarrow & (\mathbf{T} \in b_Y) \\ (\mathbf{F} \in w_3) & \Rightarrow & (\mathbf{F} \in b_Y) \\ (\mathbf{T} \in w_2) & \Rightarrow & (\mathbf{T} \in w_3) \\ (\mathbf{F} \in w_2) & \Rightarrow & (\mathbf{F} \in w_3) \end{array} \right\} \qquad S' = \left\{ \begin{array}{rcl} & & (\mathbf{T} \in w_1) \\ (\mathbf{T} \in b_{DN}) & \Rightarrow & (\mathbf{T} \in w_2) \\ (\mathbf{F} \in w_1) & \Rightarrow & (\mathbf{F} \in w_2) \\ (\mathbf{F} \in b_{DN}) & \Rightarrow & (\mathbf{F} \in w_2) \\ (\mathbf{T} \in w_3) & \Rightarrow & (\mathbf{T} \in b_Y) \\ (\mathbf{F} \in w_3) & \Rightarrow & (\mathbf{F} \in b_Y) \\ (\mathbf{T} \in w_2) & \Rightarrow & (\mathbf{T} \in w_3) \\ (\mathbf{F} \in w_2) & \Rightarrow & (\mathbf{F} \in w_3) \end{array} \right\}$$

where

$$\begin{array}{rcl} w_1 & : & \text{set variable denotes the wire before the instruction XIC (DN);} \\ w_2 & : & \text{set variable denotes the wire after the instruction XIC (DN);} \\ w_3 & : & \text{set variable denotes the wire before the instruction OTE (Y);} \\ b_{DN} & : & \text{set variable denotes the bit DN, a status bit of the TON instruction;} \\ b_Y & : & \text{set variable denotes the bit Y.} \end{array}$$

Figure 3: Constraint system and base system for a fragment of the example program in Figure 1.

assignments to detect constant wires. This approach gives a probabilistic guarantee that a wire is constant. The basic idea is to generate random input assignments and add corresponding constraints to the base system and solve. If a wire w takes on different values in different solutions of the respective systems, we consider that wire as "non-constant." If after some number of samples, a wire w still remains single-valued, then w is considered "constant." For example, consider again the example program of Figure 1. Since the second rung does not interfere with the other rungs, we can consider it in isolation. For this rung, whatever the value of the bit A is, the wire after the XIO (A) instruction is always false, since it requires that A to be at the same time both true and false for the wire to be true.

Relay race analysis works by simulating multiple scans and looking for racing outputs. Similar to the constant wire analysis, 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 of the output bits are also inputs, in the next scan, the input assignment is updated using the minimum solution from the previous scan. Again, we add the resulting system to the base system and solve to obtain the new minimum solution of the outputs. This process repeats. If an output changes value across scans, a relay race is detected. For example, consider the example program in Figure 1. Since the bottom two rungs do not interfere with the other three, let us consider these two rungs only. Let us 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 two analyses are conservative in the sense that they cannot detect all of the constant wires or relay races in a program. However, any constant wire detected by the first constant wire analysis are indeed constant wires, and any constant wire reported by the second constant wire analysis are constant wires with provably high probability. As to the relay race analysis, any relay races the analyzer detects are indeed relay races, and we can prove that a large class of relay races are detected with high probability.

We have implemented the two analyses described in this paper in Standard ML of New Jersey (SML) [21]. 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.

The rest of the paper is structured as follows. First, we describe the constraint language used for our analyses (Section 2). The rules for generating the base system come next (Section 3), followed by a description of our analyses (Section 4). We then discuss some techniques of using constraints to provide support for the

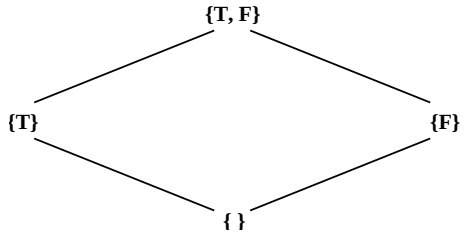

Figure 4: Our working abstract domain.

analyses (Section 5). Finally, we present some experimental results (Section 6), followed by a discussion of related work (Section 7), possible future work (Section 8) and the conclusion (Section 9).

2 Set Constraints

In this section, we describe the constraint language we use for expressing our analyses. We give its syntax and semantics.

Set constraints [17, 2, 3] are inclusion constraints between sets of terms. A set constraint is of the form $x \subseteq y$, where x and y are set expressions. Our expression language consists of set variables, a least value $\bot$, a greatest value $\top$, constant constructors $\mathbf{T}$ and $\mathbf{F}$, intersections, unions, and conditional expressions. The syntax of the expression language is given by the following grammar:

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

where c is a constant (either $\mathbf{T}$ or $\mathbf{F}$) and $\alpha \in V$ is a 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 the partial order on these elements given in Figure 4. 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 Figure 5.

Conditional expressions deserve some discussion. Conditional expressions are proposed in [3] for accurately modeling of flow-of-control. 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 **and** v_2 with the following conditional expression:

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

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

$$\begin{aligned} ((\mathbf{T} \cap \mathbf{T}) \Rightarrow (\mathbf{T} \cap \mathbf{T}) \Rightarrow \mathbf{T}) &= (\mathbf{T} \Rightarrow \mathbf{T}) \Rightarrow \mathbf{T} \\ &= \mathbf{T}. \end{aligned}$$

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 performance reasons. Set constraints give us the flexibility to model certain instructions less accurately and less expensively than others, thus, making the analysis of RLL programs more manageable.

3 Constraint Generation

In this section, we describe how we use inclusion constraints to model RLL programs. We give the constraint generation rules used to express RLL programs in our constraint language.

Because of the scan evaluation model of RLL programs, it is natural to express the meaning of a program in terms of the meaning of a single scan. The constraint generation rules we present model the meaning of

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

Figure 5: Semantics of set expressions.

a single scan of RLL programs. 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

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

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.

In this section, w_1 and w_2 denote the wires preceding and following an instruction respectively. Furthermore, b denotes the bit referenced by an instruction unless specified otherwise. Figure 6, Figure 7 and Figure 8 give the rules for generating the constraints describing the data and control flow of RLL programs. 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], we use two fresh set variables v_1 and v_2 to 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 rule [XIO] is similar to the rule [XIC].

Energise Coil

The instruction OTE is called "Energise Coil." It is programmed to control either an output connected to the controller or an internal output 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 rules [XIC] and [XIO]. The set variable v_{ct} is fresh, representing a new instance of the referenced bit b. Any later references to b use this instance. The constraints express that b is true if and only w_1 is true.

Latches

The instructions OTL and OTU are similar to OTE. OTL is "Latch Coil," and OTU is "Unlatch Coil." These two instructions appear in pairs. In latch coil, the bit b remains true even though the

$$\frac{\begin{array}{c}v_1 \text{ and } v_2 \text{ are fresh variables} \\ v_{ct} = E(XIC_{ct}) \\ S = \{((v_1 \cap \mathbf{T}) \Rightarrow (v_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup ((v_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{F}) \ \subseteq v_2\}\end{array}}{E, XIC \rightarrow E, S, v_1, v_2} \quad \text{[XIC]}$$

$$\frac{\begin{array}{c}v_1 \text{ and } v_2 \text{ are fresh variables} \\ v_{ct} = E(XIO_{ct}) \\ S = \{((v_1 \cap \mathbf{T}) \Rightarrow (v_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup ((v_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{F}) \ \subseteq v_2\}\end{array}}{E, XIO \rightarrow E, S, v_1, v_2} \quad \text{[XIO]}$$

$$\frac{\begin{array}{c}v_1,\ v_2, \text{ and } v_{ct} \text{ are fresh variables} \\ E' = E + \{(OTE_{ct}, v_{ct})\} \\ S = \{((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \ \subseteq v_{ct}\}\end{array}}{E, OTE \rightarrow E', S, v_1, v_2} \quad \text{[OTE]}$$

$$\frac{\begin{array}{c}v_1,\ v_2, \text{ and } v_{ct} \text{ are fresh variables} \\ v'_{ct} = E(OTL_{ct}) \\ E' = E + \{(OTL_{ct}, v_{ct})\} \\ S = \{((v'_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow (v'_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{F}) \ \subseteq v_{ct}\}\end{array}}{E, OTL \rightarrow E', S, v_1, v_2} \quad \text{[OTL]}$$

$$\frac{\begin{array}{c}v_1,\ v_2, \text{ and } v_{ct} \text{ are fresh variables} \\ v'_{ct} = E(OTU_{ct}) \\ E' = E + \{(OTU_{ct}, v_{ct})\} \\ S = \{((v'_{ct} \cap \mathbf{F}) \Rightarrow \mathbf{F}) \cup ((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{F}) \cup ((v_1 \cap \mathbf{F}) \Rightarrow (v'_{ct} \cap \mathbf{T}) \Rightarrow \mathbf{T}) \ \subseteq v_{ct}\}\end{array}}{E, OTU \rightarrow E', S, v_1, v_2} \quad \text{[OTU]}$$

Figure 6: Part one of rules for generating constraints.

bits that caused this output to be true have changed (i.e., it is a latch). The bit b is true if w_1 is true or it is true before the instruction executes. Otherwise, b is false. The unlatch coil (OTU) instruction is symmetric. In both the rules [OTL] and [OTU], 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. In the rule [OTL], the constraints express that b is true if and only the wire w_1 is true or b is true before evaluating this instruction. The rule [OTU] is similar.

Timers

Timers (TON) and counters (CTU) are output instructions that control a device after an elapsed period of time or an expired count. They are normally internal output instructions with some associated status bits that may cause other outputs to be on (true) or off (false).

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. It 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 says that if the wire w_1 is true, then the DN bit is either true or false, i.e., we do not have any information of whether it is true or of whether it is false. If the wire w_1 is false, then the DN bit is definitely false. Notice that in this constraint, we over-estimate the value of the DN bit, meaning that additional values may be assumed for the bit besides its actual value. The constraints for the TT and EN bits are straightforward.

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

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

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 set of constraints and τ a set expression, then the notation $\tau \Rightarrow S$ abbreviates the set of 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.

The rules in Figure 8 specify the order of evaluation of RLL programs. Constraints are generated in this same order. The order of generating constraints is important because the correct instances of wires and bits should be used.

The rule [RUNG] specifies that the constraints are generated rung by rung in order. The rule [NORUNG] is straightforward, simply saying that no constraints need to be generated.

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

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

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

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

Figure 7: Part two of rules for generating constraints.

The rule [IO] describes the generation of constraints for a single rung. The constraints for the input instructions are generated and then the constraints for the output instructions are generated. Notice that, in the rule, the constraint

$$\mathbf{T} \subseteq v_1.$$

The constraint says that, in a rung, the wire before the first instruction is always true. The constraint

$$v_2 \subseteq v_1'$$

is for connecting inputs and outputs.

Rules [INO] and [IBRANCH] are similar to the rule [IO], except that v_1 is not always true. The rule [NOINPUT] is straight forward. Similar to the rule [NORUNG], it says that no constraint is generated.

The rule [ILEVEL] describes the generation of constraints for parallel inputs — inputs of the form:

In the rule

$$v_1 = v_1'$$

is an abbreviation for the two constraints

$$v_1 \subseteq v_1' \text{ and } v_1' \subseteq v_1.$$

The fresh variable v is used to model the wire after the parallel wires. The constraint

$$(v_2 \cap \mathbf{T}) \Rightarrow \mathbf{T} \cup (v_2' \cap \mathbf{T}) \Rightarrow \mathbf{T} \cup (v_2 \cap \mathbf{F}) \Rightarrow (v_2' \cap \mathbf{F}) \Rightarrow \mathbf{F} \subseteq v$$

says that the wire after the parallel wires is true if one of the parallel wires is true.

The rule [OBRANCH] describes the generation of constraints for parallel outputs — outputs of the form:

The rule says that the parallel levels of outputs are evaluated from top to bottom. Note that Figure 6 and Figure 7 only give a partial list of all the instructions in RLL. The rules for most other instructions are straightforward. We now present a theorem which states that the constraints generated from an RLL program together with constraints for restricting the inputs has a least solution.

Theorem 3.1 (Existence of Least Solution) *For any RLL program $\mathcal{P}$, let S be the constraint system generated by the rules given in Figure 6, Figure 7 and Figure 8. 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 3.2 (Soundness) *Let $\mathcal{P}$ be an RLL program and S be the constraint system generated by the rules given in Figure 6, Figure 7 and Figure 8. 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), then true (false) is a value in the set representing this instance.*

Theorem 3.1 and Theorem 3.2 are proven in Appendix A and Appendix B respectively.

rungs

$$\frac{E, rungs \rightarrow E', S_0 \quad E', rung \rightarrow E'', S_1}{E, rungs\ rung \rightarrow E'', S_0\ \cup S_1} \quad \text{[RUNG]}$$

$$\frac{}{E, \epsilon \rightarrow E, \emptyset} \quad \text{[NORUNG]}$$

rung

$$\frac{E, input\_list \rightarrow E', S_0, v_1, v_2 \quad E', output\_list \rightarrow E'', S_1, v_1', v_2'}{E, input\_list\ output\_list \rightarrow E'', S_0\ \cup S_1\ \cup \{v_2 \subseteq v_1', \mathbf{T}\ \subseteq\ v_1\}, v_1, v_2'} \quad \text{[IO]}$$

input_list

$$\frac{E, instruction \rightarrow E', S_0, v_1, v_2 \quad E', output\_list \rightarrow E'', S_1, v_1', v_2'}{E, instruction\ output\_list \rightarrow E'', S_0\ \cup S_1\ \cup \{v_2 \subseteq v_1'\}, v_1, v_2'} \quad \text{[INO]}$$

$$\frac{E, input\_branch \rightarrow E', S_0, v_1, v_2 \quad E', input\_list \rightarrow E'', S_1, v_1', v_2'}{E, input\_branch\ input\_list \rightarrow E'', S_0\ \cup S_1\ \cup \{v_2 \subseteq v_1'\}, v_1, v_2'} \quad \text{[IBRANCH]}$$

$$\frac{v \text{ fresh}}{E, \epsilon \rightarrow E, \emptyset, v, v} \quad \text{[NOINPUT]}$$

input_level

$$\frac{\begin{array}{c} v \text{ is a fresh variable} \\ E, input\_level \rightarrow E', S_0, v_1, v_2 \quad E', input\_list \rightarrow E'', S_1, v_1', v_2' \\ S = \begin{array}{c}\{\ (v_2 \cap \mathbf{T}) \Rightarrow \mathbf{T}\ \cup\ (v_2' \cap \mathbf{T}) \Rightarrow \mathbf{T} \cup \\ (v_2 \cap \mathbf{F}) \Rightarrow (v_2' \cap \mathbf{F}) \Rightarrow \mathbf{F}\ \subseteq\ v\}\end{array} \end{array}}{E, input\_level\ input\_list \rightarrow E'', S \cup\ S_0\ \cup\ S_1\ \cup\ \{v_1 = v_1'\}, v_1, v} \quad \text{[ILEVEL]}$$

output_branch

$$\frac{\begin{array}{c} E, output\_branch \rightarrow E', S_0, v_1, v_2 \quad E', input\_list \rightarrow E'', S_1, v_1', v_2' \\ E'', output\_list \rightarrow E'', S_2, v_1'', v_2'' \\ S = S_0\ \cup S_1\ \cup S_2 \cup \{v_1 = v_1', v_2' \subseteq v_1''\} \end{array}}{E, output\_branch\ input\_list\ output\_list \rightarrow E''', S, v_1, v_2} \quad \text{[OBRANCH]}$$

Figure 8: Part three of rules for generating constraints.

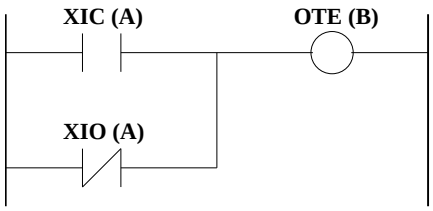

Figure 9: An example RLL program.

4 Analyses

In this section, we describe our analyses for detecting constant wires and relay races in RLL programs. The general strategy for each analysis is

1. generate the base system using the constraint generation rules presented in the previous section.
2. add constraints that restrict the inputs to the base system to express the desired information.

In both analyses, we make the assumption that all input assignments are possible. Our analyses can be made more accurate if additional information about the possible input values are available.

4.1 Constant Wire Analysis

We first describe the analysis for detecting constant wires in an RLL program. Recall that the problem is detecting wires that are constant over all possible program executions. Since such a wire contributes nothing to any run of the program, the existence of such a wire usually indicates a programming mistake.

Our approach is to compute both an upper and a lower bound on the set of constant wires. For the lower bound, we constrain every input variable v by

$$\top \subseteq v.$$

These constraints are added to the base system. The least solution for the resulting constraint system is then computed. If a variable v is *not* $\top$ in the least solution, then we know that the variable *must* only have one value: either $\mathbf{T}, \mathbf{F}$ or $\bot$ (undefined). We call this analysis **LB**.

The drawback of **LB** is that it is very inaccurate in the sense that most wires are considered non-constant; in practice, it is a very coarse approximation. Consider the example in Figure 9. It is clear that the wire before the instruction OTE(B) is always true. However, this simple analysis cannot detect this fact. The inaccuracy of **LB** results from its inability to capture interdependencies between quantities, for example between a variable and its negation. The base system for this program is given in the top of Figure 10.

Since bit A is the only input bit, we add the constraint

$$\top \subseteq b_A$$

to the base system. The minimum solution of the resulting system is presented in the table of Figure 10 (column 3). We see that **LB** does not detect the constant wire before OTE(B).

Any constant wires that are computed by **LB** are guaranteed to be constant. Thus, it gives a *lower bound* on the number of constant wires in an RLL program. To get more accurate information, we must model concrete inputs as closely as possible. One possibility is to exhaustively test each possible input configuration, which is just a **T** or **F** assignment for each input variable. Since the number of input variables are usually large, and there are 2^n input configurations of n inputs, exhaustive testing is impractical. However, exhaustive testing is not necessary because we are interested not in what the system computes but whether there are any constant wires. Thus, we can choose input configurations uniformly at random, compute the value for each wire under this input configuration, and union the values of the same wire over all configurations. If the union for a wire turns to be $\top$, the wire is not constant.

The intuition behind this analysis is that after a relatively small number of samples, there are few single-valued wires remaining, and they are likely to be constant wires. Since there are only a small number of

$$\left\{\begin{array}{rcl} \mathbf{T} & \subseteq & w_0 \\ \mathbf{T} & \subseteq & w_1 \\ ((\mathbf{T} \cap b_A) \Rightarrow \mathbf{T}) \cup ((\mathbf{F} \cap b_A) \Rightarrow \mathbf{F}) & \subseteq & w_2 \\ ((\mathbf{T} \cap b_A) \Rightarrow \mathbf{F}) \cup ((\mathbf{F} \cap b_A) \Rightarrow \mathbf{T}) & \subseteq & w_3 \\ ((\mathbf{T} \cap w_2) \Rightarrow \mathbf{T}) \cup ((\mathbf{T} \cap w_3) \Rightarrow \mathbf{T}) \cup (((\mathbf{F} \cap w_2) \Rightarrow (\mathbf{F} \cap w_3)) \Rightarrow \mathbf{F}) & \subseteq & w_4 \\ w_4 & \subseteq & w_5 \\ ((\mathbf{T} \cap w_5) \Rightarrow \mathbf{T}) \cup ((\mathbf{F} \cap w_5) \Rightarrow \mathbf{F}) & \subseteq & b_B \end{array}\right\}$$

bit or wire	variable	**LB** results	**UB** results
wire preceding XIC(A)	w_0	$\mathbf{T}$	$\mathbf{T}$
wire preceding XIO(A)	w_1	$\mathbf{T}$	$\mathbf{T}$
wire following XIC(A)	w_2	$\top$	$\top$
wire following XIO(A)	w_3	$\top$	$\top$
wire following the joint	w_4	$\top$	$\mathbf{T}$
wire preceding OTE(B)	w_5	$\top$	$\mathbf{T}$
the bit A	b_A	$\top$	$\top$
the bit B	b_B	$\top$	$\mathbf{T}$

Figure 10: Base system for the example in Figure 9.

them, a programmer should be able to check each individual wire. The constraint solver can compute a backward slice [22] for a wire to tell what inputs affect it, along with a boolean function of the wire in terms of these inputs. This information can help the programmer to determine whether a wire is constant and, if it is, the reason it is constant. We call this analysis **UB**.

For the example in Figure 9, the analysis **UB** will include the wire before the instruction OTE(B) as possibly constant, since whatever value (either $\mathbf{T}$ or $\mathbf{F}$) the bit A assumes, the wire before OTE(B) is always $\mathbf{T}$. The base system is the same as that for **LB**. The bit A is the only input bit. There are two input configurations: A is true, or A is false. For the input configuration that A is true, we add the following constraint to the base system:

$$\mathbf{T} \subseteq b_A$$

In the minimum solution of these constraints, we know that w_5 is $\mathbf{T}$. For the input configuration that A is false, the following constraint is added to the base system:

$$\mathbf{F} \subseteq b_A$$

We now see that w_5 is $\mathbf{T}$ in the new minimum solution, too. Therefore, the wire before OTE(B) is considered constant by **UB**. The result for **UB** is presented in the table of Figure 10 (column 4). The number of wires that are considered possibly constant by **UB** gives an *upper bound* on the number of constant wires under our model of RLL programs.

4.2 The Effectiveness of Random Sampling

In RLL programs, a bit or a wire usually only depends on a small number of inputs, typically around 10 [1]. This fact makes random sampling in **UB** more effective than one might expect. After a relatively small number of samples of input assignments, we are confident that almost all possible input assignments affecting each input are covered.

To be more precise, assume N is the number of inputs and

$$M = \max_{v \in \mathbf{VAR}} |\mathbf{DEP}(v)|,$$

where **VAR** is the set of variables and $\mathbf{DEP}(v)$ of a variable v is the set of inputs that v depends on. In other words, for all variable v, it depends on no more than M variables. Let $k = 2^M$.

[1] This information is obtained from experiments with a few production size RLL programs.

Theorem 4.1 *For any variable v, the expected number of samples to draw to get all the possible truth assignments of the inputs in* $\mathbf{DEP}(v)$ *is no more than* $k \ln k + \mathcal{O}(k)$.

Proof. Notice this problem is just a variation of the Coupon Collector's Problem (See Appendix C). □

We know from the analysis of the Coupon Collector's Problem that the actual value is sharply concentrated around this expected value.

Theorem 4.2 *For any variable v and $c > 0$, the probability that after $k(\ln k + c)$ random samples that there are truth assignments missing from the samples is approximately* $1 - e^{-e^{-c}}$.

We also present some empirical measurements of the effectiveness of random sampling in Section 6.2.

4.3 Relay Race Analysis

Our second analysis detects relay races. 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]. Let C be the set of all possible input configurations. Further, let $\Psi_i : \mathbf{OUT} \rightarrow \{\mathbf{T}, \mathbf{F}\}$ be the mapping from the set of outputs to their corresponding values at the end of the ith scan.

Definition 4.3 *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 4.3 states under what conditions a program exhibits a race. Note that this definition assumes that outputs should stabilize after a single scan.

Definition 4.4 *Let P be an RLL program. An* approximation *A of P is an abstraction of the RLL program satisfying for any input configuration c and bit b of P, $P_c(b)$ (the value of b in the program P with input c) at the end of one scan is contained in $A_c(b)$ (the value of b in the abstraction A with input c), i.e., $P_c(b) \in A_c(b)$.*

Let A be an approximation of P. Let $\Phi_i : \mathbf{OUT} \rightarrow \wp(\{\mathbf{T}, \mathbf{F}\})$ be the mapping from the set of outputs to their corresponding values at the end of the ith scan in A.

Definition 4.5 *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 $\Psi^* : \mathbf{OUT} \rightarrow \{\mathbf{T}, \mathbf{F}\}$ such that $\Psi^*(b) \in \Phi_i(b)$, for all $b \in \mathbf{OUT}$ and $i \geq 1$.*

Lemma 4.6 *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.*

Proof. Since P is race free, by Definition 4.3, we have $\Psi_i = \Psi_1$ for all $i \geq 1$. Since A is an approximation of P, by Definition 4.4, $\Psi_i(b) \in \Phi_i(b)$ for all $i \geq 1$. Thus, $\Psi_1(b) \in \Phi_i(b)$ for all $i \geq 1$, and by Definition 4.5, the approximation A is also race free. □

Lemma 4.6 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 4.7 *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 $\bigcap_{i \geq 1} \Phi_i(b) = \emptyset$.*

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

```
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;
```

Figure 11: Algorithm for detecting races.

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 $b \in \mathbf{OUT}$, we have $\bigcap_{i\geq 1} \Phi_i(b) \neq \emptyset$. 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. □

In principle, for any given input assignment, it is 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 4.8 *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.*

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 input configuration c' prior to scan i. Now chose any configuration c'', s.t. $c''(b) \subseteq c'(b)$ for all b. Since our analysis is monotone in the input (Theorem 3.1), 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''. □

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 4.7 and Lemma 4.8 thus lead naturally to the algorithm in Figure 11 for detecting relay races. The general strategy for the analysis is:

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

We make the assumption that all input assignments are possible. In practice, there may be dependencies between inputs that make some input configurations unrealizable. Our analysis can be made more accurate if information about these dependencies is available.

We use the example in Figure 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 Figure 12. 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 Figure 12). 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 Figure 12.

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

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**

Figure 12: Base system for the last two rungs of the example program in Figure 1 with the least solutions at the end of the first and the second scans given in the table.

```
1   I_current = a racing configuration;
2   while there exists input v that is not checked
3       I = I_current ∪ {⊤ ⊆ v};
4       run relay race analysis with I as the input;
5       if the same races are observed
6           I_current = I;
7       else
8           The input v contributes to the races;
```

Figure 13: Algorithm for computing the set of inputs causing a race.

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.

The algorithm in Figure 11 detects whether an output races or not under a given input. To help the RLL programmers to find the cause of a race, it is important also to report the relevant inputs. For each input v, we add the constraint

$$\top \subseteq v$$

to the base constraint system and leave the other inputs unchanged. We run the algorithm in Figure 11 with this modified input configuration. If the same race is observed, we know that v is not one of the inputs causing the race. Otherwise, the input v does contribute to the race. This process repeats until all inputs have been checked. The algorithm is given in Figure 13.

While simple, the algorithm in Figure 13 is an expensive way to compute the inputs that cause a race. Another way of getting the information is presented in Figure 14. The input to the algorithm is the base constraint system and a set of bits that are racing. The algorithm outputs a set of inputs that affect the set of racing bits. The algorithm first computes the inputs that affect B_{racing} in one scan using the facility provided by the constraint solver. Since some of the inputs might be internal, these bits may be affected by other inputs from previous scans. We need to compute what inputs affect these bits by another backward slice. This process repeats until the set I does not grow.

For the relay race analysis, we need to modify the rules [TON] since the status bits of the timers are assumed to be the same for all scans under a given input. This assumption is reasonable since the scan time,

```
1       B_racing:       the set of racing bits
2       C_base:         the base system
3       I:              the set of inputs that affect the bits in B_racing transitively
4
5       (* compute the set of inputs that affect the bits in B_racing in a scan *)
6       I = SLICE_backward(B_racing, C_base)
7
8       repeat
9               (* C is the set of the last instances of the bits in B *)
10              C = LAST(I);
11              I = I ∪ SLICE_backward(C, C_base);
12      until I does not change
```

Figure 14: A more efficient algorithm for computing the inputs that cause a race.

compared with the timer increments, is infinitesimal. The modified rule is given by

$$\frac{\begin{array}{c} v_1,\ v_2,\ v_{en}, \text{ and } v_{tt} \text{ are fresh variables} \\ E' = E + \{(TON_{en}, v_{en}), (TON_{tt}, v_{tt})\} \\ S = \left\{ \begin{array}{r} ((v_1 \cap \mathbf{T}) \Rightarrow (E(TON_{dn}) \cap \mathbf{F}) \Rightarrow \mathbf{T}) \ \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \ \cup ((E(TON_{dn}) \cap \mathbf{T}) \Rightarrow \mathbf{F}) \ \subseteq \ v_{tt}, \\ ((v_1 \cap \mathbf{T}) \Rightarrow \mathbf{T}) \ \cup ((v_1 \cap \mathbf{F}) \Rightarrow \mathbf{F}) \ \subseteq \ v_{en} \end{array} \right\} \end{array}}{E, TON \rightarrow E', S, v_1, v_2}$$

By the analysis of the Coupon Collector's Problem, after approximately $2^k \ln(2^k) = 2^k \cdot k \ln 2 < k \cdot 2^k$ scans, we have detected, in our approximation, all races of k inputs with high probability. These are actual races in the original RLL program.

5 Implementation Techniques

In this section, we discuss some ways in which we use constraints either to limit the size of the information one needs to examine or to obtain useful information from the constraint system. This illustrates that constraints are useful for providing programming support not directly related to the analyses, such as freeing programmers from examining irrelevant information and providing explanation for the causes of certain behaviors of the programs.

5.1 Filter Values

Recall in the constant wire analysis, after the least solution is computed, we need to determine which wires or bits have only values either $\perp$, $\mathbf{T}$, or $\mathbf{F}$. In order to obtain this information, we test whether

$$\top \subseteq \bigcup Sol_{least}(v),$$

where v ranges over the instances of a wire w or a bit b. If the subset relation holds, we know that w or b can be both true and false. On the other hand, if the relation does not hold, w or b has one of the other three possible values. With the simple test above, some irrelevant wires or bits may be left for inspection by the programmer. These wires or bits consist of two kinds: the inputs and the left-most wire of each rung.

With random sampling, each input bit is either true or false. To avoid examining these bits, we add a special set constructor **input** to our expression language with semantic value $\{\mathbf{input}\}$. Each input bit has the value $\mathbf{input} \cup \mathbf{T}$ or $\mathbf{input} \cup \mathbf{F}$. Similarly for the beginning wires, we add another special set constructor **initial** to our expression language with semantic value $\{\mathbf{initial}\}$. Each start wire has the value $\mathbf{initial} \cup \mathbf{T}$.

Again to determine the wires and bits to inspect, we perform the following test:

$$\begin{array}{rcll} \mathbf{T} \cup \mathbf{F} & \subseteq & V & \textbf{or} \\ \textbf{input} & \subseteq & V & \textbf{or} \\ \textbf{initial} & \subseteq & V, & \end{array}$$

where V denotes $\bigcup Sol_{least}(v)$. If the test fails, we need to inspect the corresponding bit or wire. Since in the constraint generation rules **input** or **initial** are not propagated from the inputs or the beginning wires, only the inputs have the value **input** and the beginning wires have the value **initial**. Thus, if an input has the value **input**, we know it must be an input, and if a wire has the value **initial**, it must be a beginning wire.

5.2 Counter Wires

In this section, we describe another method to reduce the number of irrelevant wires to be inspected by a programmer.

Recall that a counter (CTU) counts how many times the wire preceding the instruction makes false to true transitions. The done bit (DN) associated with a counter becomes true if the preceding wire has made a preset number of false to true transitions across scans. The constraint for the done bit is given by

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

where v_1 and v_{dn} are the set variables for the wire preceding the counter instruction and the done bit respectively.

Notice that for v_{dn} to have the value $\mathbf{T}$, v_1 must be both true and false in some samples. Suppose in the program, the wire corresponding to v_1 can be true and false. Then the done bit can be true in some execution sequence of the program. Assume, however, in our approximation of the program, for all samples, v_1 is always true or false, but not both. Then v_{dn} only has the value false. Thus, the done bit is considered constant. In addition, many wires and bits affected by this done bit may be considered constant as well because the done bit is always false. To remove these irrelevant wires and bits, we keep a record of *counter wires*, wires that immediately precede counter instructions. We add not only the constraints corresponding to a sample configuration to the base system, but also the constraints

$$\{V_{union}(w) \ \subseteq \ w \mid w \text{ is a counter wire}\}$$

where $V_{union}(w)$ gives the union of the values of w up to the current sample. With the addition of these constraints, the problem with the done bit is readily solved.

5.3 Backward Slicing

Let v be a given variable. It is desirable to know the set of inputs that affect v. This set of inputs is called a *backward slice* for v [22]. The constraint solver we are using can provide us with this information by computing a backward slice. The solver not only provides us with the set of inputs that affect v, but also a boolean formula that describes how v depends on these inputs. This information can help an RLL programmer to determine whether a wire is indeed constant, and, if the wire is constant, possible causes of the problem. The slice of a variable v can be simply computed by recursively replacing the intermediate variables by their lower bounds until all the variables in the lower bound of v are inputs. This lower bound can be simplified, and the inputs left are the slice of v and the simplified lower bound is effectively a boolean formula describing how these inputs affect v.

6 Experimental Results

We have implemented our analyses using a general constraint solver [14]. The analyses are implemented in SML. Inputs to our analyses are abstract syntax tree (AST) representations of RLL programs. The ASTs are parsed into internal representations, and constraints are generated using the rules in Figure 6, Figure 7, and Figure 8. The resulting constraints are solved to obtain the base system.

Program	Size	Num. of Vars.	Secs / Scan
Mini Factory	9,267	4,227	0.5
Big Bak	32,005	21,596	4
Wdsdflt(1)	58,561	22,860	3
Wdsdflt(2)	58,561	22,860	3

Figure 15: Benchmark programs for evaluating our analyses.

Program	Lower Bound	Upper Bound	Number of samples
Mini Factory	0	0	500
Big Bak	0	0	30
Wdsdflt(1)	32	868	1000
Wdsdflt(2)	32	868	1000

Figure 16: Results from the constant wire analysis.

6.1 Benchmarks

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

- **Mini Factory**
  This program is an example program that has been studied and tested by RLL programmers and researchers working on program analysis for RLL programs.
- **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.

Figure 15 gives a table showing the size of each program in terms of number of lines in abstract syntax tree form, number of variables that are in its base system, and the time it takes our analyses to analyze one scan. All measurements reported here were done on a Sun Enterprise-5000 with 512MB of main memory (using only one of the eight processors).

6.2 Constant Wire Analysis

We performed the two kinds of constant wire analyses on the four benchmark programs. The results from the analyses are given in Figure 16. In the table, we give, for each program, the number of constant wires from **LB** the number of constant wires from **UB** and the number of samples that **UB** used.

For Mini Factory and Big Bak, both **LB** and **UB** do not detect any constant wires. In one run of Mini Factory, after around 500 samples, there were no "constant" wires left. In one run of Big Bak, after 30 random samples, there were no "constant" wires left. The reason Big Bak requires so few samples is that there are many arithmetic instructions in Big Bak, which are not easily modeled accurately without drastically increasing the number of constraints. As a result, the inaccurate modeling of arithmetic operations resulted in most wires being inferred to be both true and false rather quickly. Thus, **UB** terminated much earlier on Big Bak than on Mini Factory. For the two Wdsdflt programs, **LB** detected some constant wires. However, these were not bugs, but rather an artifact of some debugging code in the program that is normally turned off. Because of this debugging code, **UB** reported many wires as possibly constant, as shown in the table.

Figure 17 shows the effectiveness of the idea of random sampling in reducing the number of wires to examine in Mini Factory. The x-axis is the number of random samples. The y-axis shows the number of

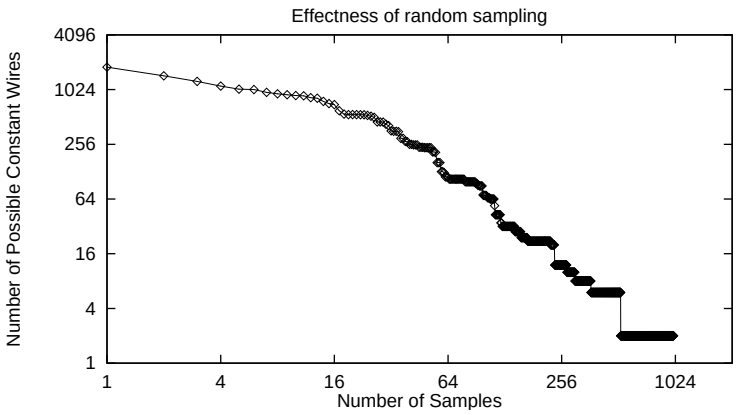

Figure 17: Effectiveness of random sampling.

Program	External Races	Internal Races	Number of samples
Mini Factory	55	186	1000
Big Bak	4	6	1000
Wdsdflt(1)	7	156	1000
Wdsdflt(2)	8	163	1000

Figure 18: Results from the relay race analysis.

wires that are still possibly constant. After about 200 samples, the number of possibly constant wires drops to 20. Initially there are approximately 3500 wires.

6.3 Relay Race Analysis

We also performed our relay race analysis on the four benchmarks. This analysis produced more interesting results than the constant wire analysis. It discovered many relay races in our benchmark programs. The results from the analysis are presented in Figure 18. In the table, for each program, we show the number of external racing bits — bits that are connected to external outputs, and the number of internal racing bits — bits that are internal to the program, and the number of total samples run. The analysis were run for 1000 samples for all the programs. By the analysis of the Coupon Collector's Problem, 1000 trials are sufficient to uncover all races involving 7 or fewer inputs.

For the Mini Factory program, there were no known relay races in the program, but our analysis detected many such races. Some of the races were subsequently verified by running the program. From the 1000 samples, 55 external races and 186 internal races were reported. For Big Bak, 4 external races and 6 internal races were reported. Although Big Bak is a much bigger program than Mini Factory, the inaccuracy in modeling of arithmetic operations may be one reason why fewer races were found. For Wdsdflt(2), there were 7 external races and 156 internal races reported. The Wdsdflt(1) program has a known relay race, which took the programmer who developed this program four hours to find [5]. Our analysis discovered this bug among 8 external and 163 internal races. Notice that some reported races may be unrealizable if the corresponding input configuration cannot be realized. There is no way without additional information about the possible inputs to characterize which relay races may actually happen.

7 Related Work

In this section, we discuss the similarity and differences of our analyses from work in data flow analysis, model checking, and testing.

Data Flow Analysis Data flow analysis has been traditionally used in optimizing compilers to collect variable usage information for optimizations such as dead code elimination and efficient register allocation

[1]. It has also been applied for ensuring software reliability [15, 16]. There are two main distinctions of our approach from data flow analysis. One is the use of conditional constraints [3], which are essential for modeling both the boolean instructions and control flow instructions. The other one is the flexibility of our analyses to add additional constraints to the base system to get desired information, instead of solving the whole constraint system repeatedly. Our approach is more efficient because we work with an initially simplified constraint system.

Model Checking Model checking [10, 11] 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, 13, 18, 12]. 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. Our analyses for RLL programs use similar techniques. Although RLL programs in general are infinite state systems, our abstract models of RLL programs are finite-state. These abstract systems are symbolically executed to obtain information about the actual systems. In this sense, our analyses are similar to model checking. However, there are some differences. The main difference of our analyses from model checking lies in the way abstract models are obtained and how accurately these systems correspond to the actual system. In model checking, an abstract model of a concrete system is often obtained manually, while our analyses automatically generate the model. With respect to the modeling accuracy, model checking strives to produce an model which has no observable difference from the concrete system from the point of the properties to be checked, i.e., the model is a complete characterization of the actual system. However set constraints (because the use of sets) give us the flexibility to model certain parts of the system more accurately than others for analyzing large scale systems.

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 [20]. 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. Carver and Durham [9] discuss some other tradeoffs of using the actual system and abstract models of the system for testing.

8 Discussion and Future Work

In this section, we discuss our experience of analyzing RLL programs using constraints, and point out some possible directions for future work.

In designing and implementing the analyzer, we came up with the correct rules for most of the instructions and constructs fairly quickly. It is just a matter of understanding the semantics of each instruction. The language manual provides a rather good description of the instructions. However, for some instructions such as timers and counters, we spent some time considering different approaches to model them more accurately. To implement the analyzer, we spent a fair amount of time writing and correcting the parser of the AST files provided to us since there is no grammar available to us. The other parts of the implementation did not take very much time. This experience confirmed what I have believed, that it is a good idea to use a generic constraint solver versus some ad hoc analyses. It reduces the amount of effort to write an analysis tool considerably, and the analysis should be comparably efficient if the constraint solver is designed and implemented with care.

There are a few possible ways to extend this work.

- In our analyzer, we model timers and counters very conservatively. It would be interesting to model them more accurately to see how much improvement it makes to the constant wire analysis. One

possible approach is to use timed automata. We have not looked this possibility very much. It is not clear how to fit the two formalisms nicely.

- Barrett [6] discusses some other analysis problems that RLL programmers consider useful. We believe most of them can be designed and implemented quickly without much change to the existing system. It will be interesting to see how these analyses perform on real programs.
- In the analyses, we assume conservatively that all input configurations are possible. In practice, it is conceivable that some inputs cannot be turned on or off at the same time. We can imagine to have some constraints that restrict certain inputs. These constraints allow us not only to model RLL programs more realistically, but also to improve the accuracy of the analyses.

9 Conclusions

In this paper, we have described two analyses — the constant wire and relay race analyses — for RLL programs using set constraints to help RLL programmers to detect some common programming mistakes. We have demonstrated that these analyses are useful in statically catching some kinds of programming errors. Our implementation of the analyses 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.

Acknowledgments

First and foremost, I would like to thank my advisor Alex Aiken for carefully reading and commenting on drafts of this paper. Discussions with Alex Aiken and Manuel Fahndrich resulted in several ideas in the paper. We thank them. Thanks also go to Tony Barrett with Allen-Bradley, Rockwell Automation for information on RLL, providing us with abstract syntax trees of RLL programs, and running some experiments to validate our results. Finally, we thank Professor Susan Graham for the helpful comments.

References

[1] A.V. Aho, R. Sethi, and J.D. Ullman. *Compilers, Principles, Techniques and Tools.* Addison-Wesley, 1986.

[2] A. Aiken and E. Wimmers. Type inclusion constraints and type inference. In *Proceedings of the 1993 Conference on Functional Programming Languages and Computer Architecture*, pages 31–41, Copenhagen, Denmark, June 1993.

[3] A. Aiken, E. Wimmers, and T.K. Lakshman. Soft typing with conditional types. In *Twenty-First Annual ACM Symposium on Principles of Programming Languages*, pages 163–173, Portland, Oregon, January 1994.

[4] Allen–Bradley, Rockwell Automation. *SLC 500 and MicroLogix 1000 Instruction Set.*

[5] T. Barrett. Private communication.

[6] T. Barrett. Ladder logic analysis survey. Unpublished manuscript.

[7] M. Browne, E.M. Clarke, and D. Dill. Checking the correctness of sequential circuits. In *Proc. IEEE Internat. Conf. on Computer Design*, pages 545–548, 1985.

[8] M. Browne, E.M. Clarke, D. Dill, and B. Mishra. Automatic verification of sequential circuits using temporal logic. *IEEE Trans. Comput.*, 35(12):1035–1044, 1986.

[9] R.H. Carver and R. Durham. Integrating formal methods and testing for concurrent programs. In *Proceedings of the Tenth Annual Conference on Computer Assurance*, pages 25–33, New York, NY, USA, June 1995.

[10] E.M. Clarke and E.A. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proc. Workshop on Logics of Programs*, volume 131, pages 52–71, Berlin, 1981. Springer.

[11] E.M. Clarke, E.A. Emerson, and A.P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.

[12] E.M. Clarke, O. Grumberg, H. Hiraishi, S. Jha, D.E. Long, K.L. McMillan, and L.A. Ness. Verification of the futurebus+ cache coherence protocol. In L. Claesen, editor, *Proceedings of the Eleventh International Symposium on Computer Hardware Description Languages and their Applications*, North-Holland, April 1993.

[13] D. Dill and E.M. Clarke. Automatic verification of asynchronous circuits using temporal logic. In *Proceedings of the IEEE*, volume 133, pages 276–282, 1986.

[14] M. Fahndrich and A. Aiken. Making set-constraint based program analyses scale. Technical Report UCB/CSD-96-917, University of California at Berkeley, 1996.

[15] L.D. Fosdick and L.J. Osterweil. Data flow analysis in software reliability. *ACM Computing Surveys*, 8(3):305–330, September 1976.

[16] M.J. Harrold. Using data flow analysis for testing. Technical Report 93-112, Department of Computer Science, Clemson University, 1993.

[17] N. Heintze. *Set Based Program Analysis*. PhD thesis, Carnegie Mellon University, 1992.

[18] G. Holzmann. *Design and Validation of Computer Protocols*. Prentice-Hall International Editions, 1991.

[19] A. Krigman. Relay ladder diagrams: we love them, we love them not. In *Tech*, pages 39–47, October 1985.

[20] D. Lee and M. Yannakakis. Principles and methods of testing finite state machines-a survey. In *Proceedings of the IEEE*, pages 1090–1123, August 1996.

[21] R. Milner, M. Tofte, and R. Harper. *The Definition of Standard ML*. MIT Press, 1990.

[22] M. Weiser. Program slicing. *IEEE Transaction on Software Engineering*, SE-10(4):352–357, July 1984.

A Existence of the Least Solution

In this section, we prove Theorem 3.1.

Proof. Notice that every constraint is of the form $e \subseteq v$, where e is a set expression and v is a variable. Thus we can obtain a solution of the constraint system by assigning each variable $\{\mathbf{T}, \mathbf{F}\}$. To see that there is a least solution, we show that if S_1 and S_2 are two solutions, then $S_1 \cap S_2$ is also a solution, where $S_1 \cap S_2 = \{S_1(v) \cap S_2(v) \mid \text{for all variable } v\}$. First we show by induction that for any set expression e and any two variable assignments S_1 and S_2 the following holds:

$$(S_1 \cap S_2)(e) \subseteq (S_1(e) \cap S_2(e)).$$

- Base cases:
  - $e = \bot$: straight forward.
  - $e = \top$: straight forward.
  - $e = \mathbf{T}$: straight forward.
  - $e = \mathbf{F}$: straight forward.
  - $e = v'$, where v' is a variable:
    $(S_1 \cap S_2)(v') = S_1(v') \cap S_2(v')$ by the definition of $S_1 \cap S_2$.
- Inductive cases:
  - $e = (e_1 \cap e_2)$:
    We have
    $$(S_1 \cap S_2)(e_1) \subseteq (S_1(e_1) \cap S_2(e_1))$$
    and
    $$(S_1 \cap S_2)(e_2) \subseteq (S_1(e_2) \cap S_2(e_2)).$$
    Thus, we have
    $$\begin{aligned}(S_1 \cap S_2)(e_1 \cap e_2) &= (S_1 \cap S_2)(e_1) \cap (S_1 \cap S_2)(e_2) \\ &\subseteq (S_1(e_1) \cap S_2(e_1)) \cap (S_1(e_2) \cap S_2(e_2)) \\ &= (S_1(e_1 \cap e_2) \cap S_2(e_1 \cap e_2)).\end{aligned}$$
  - $e = (e_1 \cup e_2)$:
    We have
    $$(S_1 \cap S_2)(e_1) \subseteq (S_1(e_1) \cap S_2(e_1))$$
    and
    $$(S_1 \cap S_2)(e_2) \subseteq (S_1(e_2) \cap S_2(e_2)).$$
    Thus, we have
    $$\begin{aligned}(S_1 \cap S_2)(e_1 \cup e_2) &= (S_1 \cap S_2)(e_1) \cup (S_1 \cap S_2)(e_2) \\ &\subseteq (S_1(e_1) \cap S_2(e_1)) \cup (S_1(e_2) \cap S_2(e_2)) \\ &\subseteq (S_1(e_1 \cup e_2) \cap S_2(e_1 \cup e_2)).\end{aligned}$$
  - $e = (e_1 \Rightarrow e_2)$:
    We have
    $$(S_1 \cap S_2)(e_1) \subseteq (S_1(e_1) \cap S_2(e_1))$$
    and
    $$(S_1 \cap S_2)(e_2) \subseteq (S_1(e_2) \cap S_2(e_2)).$$

Thus, we have

$$\begin{aligned}(S_1 \cap S_2)(e_1 \Rightarrow e_2) &= (S_1 \cap S_2)(e_1) \Rightarrow (S_1 \cap S_2)(e_2) \\ &\subseteq (S_1(e_1) \cap S_2(e_1)) \Rightarrow (S_1(e_2) \cap S_2(e_2)) \\ &\subseteq (S_1(e_1) \Rightarrow S_1(e_2)) \cap (S_2(e_1) \Rightarrow S_2(e_2)) \\ &= (S_1(e_1 \Rightarrow e_2) \cap S_2(e_1 \Rightarrow e_2)).\end{aligned}$$

Now, let S_1 and S_2 be two solutions to the constraint system $S \cup c$. For each constraint $e \subseteq v$, we have

$$(S_1 \cap S_2)(e) \subseteq (S_1(e) \cap S_2(e)) \subseteq (S_1(v) \cap S_2(v)) = (S_1 \cap S_2)(v).$$

Thus, $S_1 \cap S_2$ is also a solution to the constraint system $S \cup c$. Therefore there exists a least solution, namely the intersection of all solutions. □

B Soundness

In this section, we prove Theorem 3.2.

Proof. Notice that the constraint system can be represented as a directed, acyclic constraint graph [3]. Thus we can prove the theorem with an induction on this graph from its sources to its sinks.

- Base case:
  The input variables have the same values as the wires or the bits that they model.
- Inductive case:
  Consider the constraint $e \subseteq v$, assuming all the variables in e approximate their corresponding instances of bits or wires. Suppose the constraint $e \subseteq v$ is generated by an application of the rule [XIC]. The proof for the other rules is similar. We thus have v_1 and v_{ct} approximate the values of XIC_{wb} and XIC_{ct}. There are four cases:
  - If $XIC_{wb} = true$ and $XIC_{ct} = true$, then $true \in v_1$ and $true \in v_{ct}$. Thus, simplifying the set expression that restricts v_2, we have $true \in v_2$.
  - If $XIC_{wb} = true$ and $XIC_{ct} = false$, then $true \in v_1$ and $false \in v_{ct}$. Thus, simplifying the set expression that restricts v_2, we have $false \in v_2$.
  - If $XIC_{wb} = false$ and $XIC_{ct} = true$, then $false \in v_1$ and $true \in v_{ct}$. Thus, simplifying the set expression that restricts v_2, we have $false \in v_2$.
  - If $XIC_{wb} = false$ and $XIC_{ct} = false$, then $false \in v_1$ and $false \in v_{ct}$. Thus, simplifying the set expression that restricts v_2, we have $false \in v_2$.

□

C 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 C.1 *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

[3] This is not true if there are backward jump instructions in an RLL program. In that case, we can do a similar induction on the strongly connected component graph of the constraint graph representing the constraint system.

$0 \leq i \leq n-1$, 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

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

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

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

□