

Towards Reactive Synthesis as a Programming Paradigm

Leyi Cui * ^{†1}, Raven Rothkopf * ^{‡1} and Mark Santolucito ^{§1}

¹Barnard College, Columbia University, New York, USA

Abstract

Reactive program synthesis from logical specifications has yet to match the user-friendly approach of example-based programming for spreadsheets, despite its success in specific domains. A main challenge hindering the broader adoption of reactive synthesis is in the complexity of specification engineering in temporal logics. We map out challenges and tools that arise as users write temporal logic specifications in Temporal Stream Logic. Our goal is to provide a roadmap for future usability work that can elevate temporal specification engineering for synthesis to match the usability support available for software engineering. By generalizing these concepts, we can gain a deeper insight into the challenges people face when reasoning about the temporal behavior of their systems.

1 Introduction

FlashFill [1] is a notable success in program synthesis, allowing users to generate Excel data manipulation programs from simple examples. For example, given a column of names *ex. AngelCui*, FlashFill can generate a program to extract the initials *ex. AC* from each entry. This approach, known as programming-by-example, is just one approach within program synthesis. For cases where examples are impractical, formal logic-based synthesis offers a viable alternative. This is particularly true for reactive systems operating on infinite input and output streams, where examples can be limiting. These systems are typically instead defined using temporal logic, with reactive synthesis creating a controller that ensures correct interaction between the system and environment, responding aptly to all input scenarios. A simple example of this is the synthesis of a reactive 3D animation. The small specification in Table 1 synthesizes a concrete implementation of a controller for a reactive, animated cube. The cube reacts to the left and right arrow key presses, spinning about the y-axis depending on the key pressed. Detailed walkthroughs of this example specification can be found in Sec. 2.4. The synthesis of the AMBA bus protocol from Linear Temporal Logic exemplifies the success of formal logic-based synthesis [2]. Recent advancements extend reactive synthesis to educational programs [3], FPGA game development [4], musical interfaces [5], and interactive animation creation [6].

Despite temporal logic specifications' wide array of application domains and substantial research progress, they remain unfamiliar and challenging to software developers to write. Recent work has explored these challenges [3], [7], yet solutions remain unclear. Unlike software engineering, which has a wealth of debugging techniques, specification engineering for temporal logics lacks similar resources. This work aims to define reactive synthesis challenges in Temporal Stream Logic specification writing, suggests tools that would aid with the development and debugging of Temporal Stream Logic, and highlights areas needing further research.

We focus on Temporal Stream Logic (TSL) [8]—a high-level, temporal logic specification language used in reactive synthesis. It extends Linear Temporal Logic (LTL) [9] with updates and predicates over arbitrary function terms. The synthesis of a TSL specification yields concrete program code corresponding to the reactive system by capturing both reactive properties and data manipulations. TSL has shown promising initial results that could introduce reactive synthesis as a path towards more expressive low/no-code platforms, which currently suffer from limitations in language expressivity [10]. The use of TSL has extended reactive synthesis to new application domains including music [5], video games [11], mobile apps [8], animation [6], and autonomous vehicle controllers [8]. Whether with TSL, LTL, or any other reactive logic [12], [13], there is a critical need for a structured framework for the

PLATEAU

13th Annual Workshop at the
Intersection of PL and HCI

Organizers:

Sarah Chasins, Elena
Glassman, and Joshua
Sunshine

This work is licensed under a
Creative Commons Attribution
4.0 International License.

* Authors contributed equally

[†]Email: lc3542@barnard.edu

[‡]Email: rgr2124@barnard.edu

[§]Email: msantolu@barnard.edu

development of temporal logic specification engineering support tools.

In this work, we present the following key contributions:

1. Based on anecdotal first-hand experience, we outline common challenges users have about reactive synthesis when writing specifications in Temporal Stream Logic (TSL).
2. We outline a list of tools that assist in the development and debugging of reactive synthesis.
3. We propose a mapping between the challenges and tools to define the space for future work in the usability of reactive synthesis.
4. To provide motivation, we explore the real-world repercussions of TSL challenges using a benchmark specification from a recently published paper. We then highlight the tools that would have prevented those challenges.

2 Preliminaries

To make this work self-contained, we formally define some technical terms that are referenced later in our descriptions of the challenges and tools. Note that understanding the formal definitions of Temporal Stream Logic is not critical to understanding the main contributions of this work. However, it provides extra context for the origin of the challenges and the technical challenges of the tools. For a full exposition of TSL's formal background, we refer the reader to prior work [8].

2.1 TSL

Temporal Stream Logic (TSL) [8] is a high-level, logical specification language that describes the behavior of a reactive system over discrete time. TSL extends Linear Temporal Logic (LTL) [9] with updates and predicates over arbitrary function terms. With TSL, one can specify a reactive system that reacts to an infinite stream of inputs to produce an infinite stream of outputs. TSL specification can synthesize implementable reactive programs written in JavaScript or Python.

TSL uses the usual LTL operators: next $\bigcirc$ and until $\mathcal{U}$. Additionally, the syntax of TSL contains *predicate terms* τ_P , *function terms* τ_F , and *update terms* τ_U , as defined in the following grammar:

$$\varphi := \tau \in \mathcal{T}_P \cup \mathcal{T}_U \mid \neg\varphi \mid \varphi \wedge \varphi \mid \bigcirc\varphi \mid \varphi \mathcal{U} \varphi$$

$$\tau_F := \mathbf{s} \mid \mathbf{f}(\tau_F^0, \tau_F^1, \dots, \tau_F^{n-1})$$

$$\tau_P := \mathbf{p}(\tau_F^0, \tau_F^1, \dots, \tau_F^{n-1})$$

$$\tau_U := [\mathbf{s} \leftarrow \tau_F]$$

TSL also uses the standard derived operators, such as *always* $\Box\varphi \equiv \perp \mathcal{R} \varphi$, *eventually* $\Diamond\varphi \equiv \text{true} \mathcal{U} \varphi$, *weak until* $\varphi \mathcal{W} \psi \equiv (\varphi \mathcal{U} \psi) \vee (\Box\varphi)$, and *release* $\varphi \mathcal{R} \psi \equiv \neg(\neg\varphi \mathcal{U} \neg\psi)$.

TSL specifies the behavior of a reactive system by utilizing signals, indicated as $\mathbf{s}$, which contain data values of arbitrary type. Through a TSL specification, it is possible to describe how functions are applied to these signals over time. Signals can either be pure outputs or *cells* that memorize data values so that the outputs of a specific time t are provided as inputs for time $t + 1$. These characteristics establish the semantics of TSL, which follow the conventional LTL semantics while integrating predicate evaluations, function evaluations, and update terms. A formal definition of TSL semantics is given in [8].

TSL synthesis can be modelled as a two-player game between the system (choosing moves in the form of outputs $\mathcal{O}$) and the environment (choosing moves in the form of inputs $\mathcal{I}$). A winning strategy of the system in this game will be a finite automaton that can produce output values that satisfy φ in reaction to all possible inputs from the environment. Likewise, a winning strategy of the environment is a finite automaton which provides a sequence of inputs that no system can satisfy.

With TSL, the user may provide constraints on the system and the environment players. Constraints on the environment restrict the possible inputs given to the system and are called assumptions. Assumptions are specified in TSL's *assume* block. Guarantees are specified in TSL's *guarantee* block, which are specified by the constraints on the system. These two blocks are desugared into a single

TSL formula of the form *assume* $\Rightarrow$ *guarantee*. Notice that if an assumption is violated, the formula becomes $\perp \Rightarrow$ *guarantee*, and so the guarantees are trivially satisfied. Thus, the environment cannot violate an assumption and still win the game.

2.2 TSL Synthesis

The realizability problem of TSL is stated as: given a TSL formula φ , is there a strategy $\sigma \in \mathcal{I}^+ \rightarrow \mathcal{O}$ mapping a finite input stream (since the beginning of time) to an output (at each particular timestep), such that for any input stream $\iota \in \mathcal{I}^\omega$, and every possible function interpretation (some concrete implementation) $\langle \cdot \rangle : \mathbb{F} \rightarrow \mathcal{F}$, the execution of that strategy over the input $\sigma \wr \iota$ satisfies φ , i.e.,

$$\exists \sigma \in \mathcal{I}^+ \rightarrow \mathcal{O}. \forall \iota \in \mathcal{I}^\omega. \forall \langle \cdot \rangle : \mathbb{F} \rightarrow \mathcal{F}. \sigma \wr \iota, \iota \models_{\langle \cdot \rangle} \varphi$$

If such a strategy σ exists, we say that σ realizes φ .

The synthesis problem of TSL asks for a concrete implementation of σ . In TSL synthesis, this model σ can be turned into a Control Flow Model (CFM), an abstract representation of the system that covers all possible behaviors, which can then be represented as program code. A formal definition of the TSL realizability and synthesis is available in a previous work [8].

Note here that our synthesis procedures asks for one such strategy σ - depending on the specification, there may be many such strategies that will satisfy φ . It is the task of the synthesis engine to choose one among these solutions (typically, the minimal solution is chosen, especially when using bounded synthesis [14]).

2.3 TSL Counter Strategies

A *counter strategy* for reactive synthesis is a strategy that the environment can use to falsify the system no matter the moves the system makes. This strategy can only be generated from an unrealizable TSL formula. For the reactive synthesis problem, a counter strategy can be formally defined as the strategy σ_C that satisfies the TSL formula $\neg\varphi$, where $\mathcal{I}_C = \mathcal{O}$ and $\mathcal{O}_C = \mathcal{I}$ are the inputs and outputs for σ_C , generated by the system and the environment.

2.4 TSL Example

We demonstrate the advantages of TSL as a specification language for real-world development using the example application of 3D animation. The small specification depicted in Table 1 synthesizes a concrete implementation of a controller for a reactive, animated cube. In this specification, the cube reacts to the left and right arrow key presses, spinning about the y-axis depending on the key pressed. The arrow key values, `pressLeft(e)` and `pressRight(e)`, are delivered as Boolean input streams to the system: `true` while the arrow key is pressed and `false` otherwise. The cube's rotation, `cube.rotation`, is emitted from the system as an integer output stream. Finally, we have one additional signal, an internal signal, or *cell*, `x`. `x` is an integer output value that appears as an input in the next time step.

We have several requirements for our system to control the cube movements depending on the left and right arrow keys. First, the system can never receive two input signals at the same time. The mutual exclusion of inputs from the environment, $\neg(\text{pressLeft}(e) \ \&\& \ \text{pressRight}(e))$, is specified in the *assume* block. Second, once an arrow key is pressed we want the cube to continuously spin in the direction specified by the input until the complementary key is pressed. To satisfy this behavior, we use the weak until operator, $\varphi \mathcal{W} \psi$, which states that either φ is true until ψ is true, or φ is true forever. The line of the plaintext TSL formula, $\text{pressLeft}(e) \rightarrow ([x \leftarrow x + 1] \mathcal{W} \text{pressRight}(e))$, specifies that when the system receives the input `pressLeft(e)`, the cell `x` will be continuously incremented until the system receives the input `pressRight(e)`. In the same time step, `cube.rotation` is updated with the value of `x` that depends on the input, consequently spinning the cube to either the left or the right. By the semantics of TSL, it is already ensured that assignments to the same cell are mutually exclusive, i.e., `x` can never be incremented and decremented at the same time, further enforced by the mutual exclusion of the system inputs.

TSL	JavaScript
<pre> always assume { !(pressR(e) && pressL(e)); } always guarantee { pressR(e) -> ([x <- x + 1] W pressL(e)); pressL(e) -> ([x <- x - 1] W pressR(e)); [cube.rotation <- x]; } </pre>	<pre> if (pressL(e) && pressR(e)) { currentState = 0 } else if (!pressL(e)) { x = x + 1 cube.rotation = x currentState = 0 } else if (pressL(e) && !pressR(e)) { x = x - 1 cube.rotation = x currentState = 1 } } else if (currentState === 1) { if (pressL(e) && pressR(e)) { currentState = 0 } else if (!pressL(e) && pressR(e)) { x = x + 1 cube.rotation = x currentState = 0 } } } </pre>

Table 1. The following example shows a TSL spec and a section of its synthesized JavaScript code.

When the specification is synthesized using our TSL synthesis tool [15], the result is an automatically generated CFM that satisfies the specified control behavior or the reactive rotation of the 3D cube. This CFM is then used to generate the system’s implementation in program code shown in the second column of Table 1.

3 TSL Challenges

We summarize key challenges and challenges faced by developers in writing TSL. We derived these categories from our personal experience of writing TSL specifications. Identifying what makes TSL hard to write further defines the problem space for future work on improving the learnability and usability of TSL.

3.1 Non-well Separation

The basic form of a reactive specification is separated into two kinds of specification: assumptions in the *assume* block and guarantees in the *guarantee* block. Conceptually, assumptions impose restrictions on the inputs to the system from its environment and guarantees provide specifications for how the system should react to those inputs and behave in the environment.

Reactive synthesis studies [3] have shown that the task of knowing what to specify in each block is error-prone. Users often struggle with differentiating between variables managed by the environment and variables managed by the system. Knowing which properties to specify in the *assume* block is particularly challenging. These conceptual hurdles often lead to *non-well-separation* [16]. A non-well-separated specification is one in which the system can avoid satisfying its guarantees by preventing the environment from being able to satisfy its assumptions. Although a non-well-separated specification is realizable in the sense that there exists a system that satisfies the specification, the system does not behave as the user intended.

A common mistake in TSL that leads to a non-well-separated specification is putting system updates on the right side of an implication in the *assume* block. Such a specification would allow the system to choose a set of updates that violates the assumptions of the specification, and since these blocks are desugared into the form $assume \Rightarrow guarantee$, the predicate evaluates to false, meaning that any consequent will yield true for the entire formula (false implies anything). Take the specification $\Box \neg[x \leftarrow y] \Rightarrow \Box(z \wedge \neg z)$ as a concrete example. The guarantee is impossible to satisfy on its own, but if the system choose to always update $[x \leftarrow y]$, the specification is $\perp \Rightarrow \Box(z \wedge \neg z)$, which is trivially satisfiable.

3.2 Temporal Operator Semantics

In order to use TSL to generate reactive systems, developers must first understand the logic they are writing. This process of understanding can be challenging because writing temporal logic requires a shift in mindset from the process of writing code [3]. If a user holds a challenge about the semantics of temporal logic operators, it can quickly lead to unrealizable specifications or realizable specifications that do not have the intended behavior. Furthermore, there are no precautions that address these challenges. The synthesis tool will blindly apply properties and check or generate the requested behavior, whether it is the desired one or not. Therefore, it is critical to help users accurately understand TSL.

We outline three concrete challenges about the semantics of TSL’s temporal operators, built upon our own experiences and the results of recent work.

Negation refers to a challenge in the understanding of the $!$ operator. Observed from our previous experience with TSL, a user confused $! F(x)$ with $F(! x)$. The first specification guarantees that x will never be true, while the second specification guarantees that x will not be true for at least one time-step. The syntactic similarity makes it difficult for new users to understand the significant semantic difference.

WeakU [7] refers to a challenge that confuses the $\mathcal{U}$ operator with its weak variant, $\mathcal{W}$. $\mathcal{U}$ guarantees that its second subterm will eventually hold, while $\mathcal{W}$ does not. As an example, take the TSL specification from Table 1, but with one minor change. In this specification, shown in Figure 1,

```
always assume {  
  ! ( pressR(e) && pressL(e) );  
}  
always guarantee {  
  pressR(e)  $\rightarrow$  ( [x <- x + 1]  $\mathcal{U}$  pressL(e) );  
  pressL(e)  $\rightarrow$  ( [x <- x - 1]  $\mathcal{U}$  pressR(e) );  
  [cube.rotation <- x];  
}
```

Figure 1. Modified specification from Table 1, demonstrating a **WeakU** challenge.

we have replaced $\mathcal{W}$ with $\mathcal{U}$ on lines 5 and 6. This specification is now unrealizable—or has no solution to the synthesis problem—because $\mathcal{U}$ guarantees that once the first subterm is true, the second subterm must eventually be true. Our specification has no assumption on the environment’s inputs that fulfill the guarantee of the $\mathcal{U}$ operator. We explore how to debug this class of challenge in Sec. 4.9.

Prior work [7] explores other semantic challenges such as **BadStateQuantification**. As with any logic, complex combinations of operators can always lead to confusion, especially for novice users.

3.3 (Non) Reactive Systems

The power of TSL is best utilized when synthesizing reactive, multi-state systems. When looking only at a TSL formula, it can be difficult to determine if a system is indeed reactive. A non-reactive system written in TSL is one without inputs from the environment - the system cannot “react” without a stream of inputs from the environment. While the specification may be realizable, likely, the specification does not capture the user’s intentions. Both TSL formulas written for other publications [17] (cf. our Case Study in Sec. 6) and formulas written by members of our lab have unintentionally specified non-reactive systems.

3.4 Precedence and Syntactic Challenges

Although the semantic complexity of temporal logic is the most significant hurdle to overcome in the development process, TSL’s syntactic complexity can also be non-trivial, especially for novices. Users have to deal with the new language paradigm of TSL with its unfamiliar operators and structural form. Syntactic challenges and even simple typos can inhibit developers from focusing on the semantic complexity of their specifications.

Prior work has shown that one of the key syntactic challenges developers have when working with reactive synthesis is precedence [7]. An operator precedence challenge is defined as a specification that is correct up until missing parentheses. We have found in our experience that other syntactic issues

that arise in software engineering also arise in temporal logic specification engineering. As another example, variable shadowing can easily turn input signals into cell signals, as illustrated in our Case Study in Sec. 6.

3.5 Unrealizability

Initial attempts of users to write a temporal logic specification are often unrealizable, meaning no system implements the specification. Debugging unrealizable specifications is challenging because they cannot be executed or simulated (similar to code that does not compile or has a runtime error). A common reason for unrealizability is that assumptions about the environment are incomplete. While unrealizability is not a challenge, many challenges lead to unrealizable specifications. Unrealizability is a well-studied phenomenon [18]–[20], and recent work has begun to investigate how unrealizability impacts users in their specification development process [3]. Just as software engineers have tools to debug code that does not fully compile or run (type errors, unit test, break points, etc.), so too is there a need for techniques for fixing unrealizable specifications.

3.6 Underspecification

Even in the case where the user has provided a realizable specification, and every part of that specification correctly matches the users intentions, there is still room for error. In particular, the goal of synthesis is that the user may leave some part of the problem unspecified and the synthesis engine will complete that part of the solution automatically. The goal is that the user specifies all the parts of the problem they care about and leaves other parts unspecified. However, if a user had in mind some constraint on the system and did not include that constraint as part of the specification, one of two things might happen. Either the synthesis engine will find a solution that matches the users intentions (either by luck or some intention inference), or the synthesis engine will find a solution that does not match the users missing part of the specification.

As users develop expertise in reactive synthesis, we have noticed they are often able to correctly guess the behavior of the synthesis engine and strategically omit parts of the specification that can be handled well by the tool. However, novice user do not have this expertise and need ways to inspect the output of the tool to confirm that they did not miss some part of the specification.

4 TSL Tools

We outline a list of tools that aid in the reactive synthesis development and debugging process. We then explain the motivation behind the mapping from challenges to the proposed tools.

4.1 Syntax Highlighting/Warnings

Similar to syntax highlighting in an IDE for software development, we can also use syntax highlighting for TSL. This can be used to help users recognize particular syntactic constructions in TSL, as well as to alert users to potential semantic issues. As a prototype of such a tool, we built a VSCode extension for TSL. This extension does basic syntax highlighting, and is also capable of alerting users to potential semantic errors. For example, issues with the construction of assumptions, as described in Sec. 3.1, would appear as shown in the screenshot of our tool shown in Fig. 2.

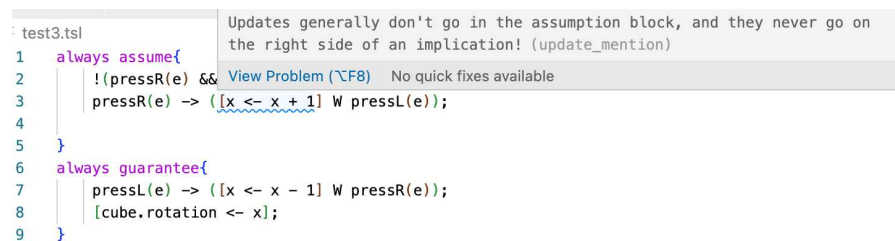


Figure 2. The TSL specification from Table 1 demonstrates a non-well-separation challenge, and can be aided by syntax highlighting and error checking.

4.2 Block-based Editors

Block-based structure editors, such as Scratch [21], have become widely accepted as tools for learning programming. These tools lighten the cognitive load experienced when programming by guarding against syntactic mistakes. Anecdotally, we noticed that when writing TSL, novice users experience a similar distracting cognitive load of worrying about syntax. When writing temporal logic, users must recall the syntactic and grammatical structure of the new language paradigm of TSL. This reliance on recall inhibits users from focusing on the semantics of their specifications. Displaying all operators and terms in a syntactically correct format enables users to recognize what they need to use in their specifications instead of remembering it, shifting the working mental model from recall to recognition.




Figure 3. The TSL specification from Table 1 built with the tslBlocks structure editor.

Recent work has proposed a block-based structure editor for TSL, tslBlocks [6]. In the specification editor, users can drag and drop temporal logic operators, update terms and predicates into locations that are restricted by the grammar of TSL. Fig. 3 demonstrates an example of a specification built in the tslBlocks editor. In Fig. 3, temporal operators, update terms, input/output signals, domain-specific functions, and mathematical operators are all distinguished by color and shape. This tool is well suited to tackle the syntactic challenges from Sec. 3.4, as well as any syntactic mistakes.

4.3 State Machine Visualization

We propose the ability to interactively test a synthesized solution while visualizing the generated Mealy machine as another tool. This tool is well suited to tackle the Temporal Operator Semantics (3.2) and Underspecification (3.6) challenges. This is particularly helpful to identify specifications that result in fewer states than expected. We built a prototype of a tool that automatically generates an interactive state machine during the synthesis process, as shown in Fig. 4.




Figure 4. A visualization of the synthesized automaton from the specification in Table 1.

4.4 Reactive System Visualizer

Understanding the inputs, outputs, and cells of a reactive system is a critical part of building a TSL specification. We built a reactive system visualizer to illustrate the behavior of reactive systems that visualizes input, output, and cell values from the TSL specification. While the state machine tool

described in Sec. 4.3 aims to provide an understanding of the fine-grained behavior of a synthesized system, the reactive system visualization specifically aims to act as a sanity check for ensuring the overall structure of the reactive system matches the user’s intention.

4.5 Minimal Assumption Cores

In reactive synthesis, assumptions are constraints put on the environment of the reactive system. Assumptions restrict the possible inputs that the environment can provide to the system. These assumptions are provided by the user, but certain assumptions may not be necessary to realize a specification. Similar to a warning about unused library imports, identifying a minimal assumption core can help prune extraneous assumptions. However, the use case is slightly more subtle—if an assumption is not necessary for realizability there is no guarantee that it will have an impact on the synthesis result. As such, extraneous assumptions may indicate a challenge the user has with another part of the specification.

4.6 Code Pruning

Optimizing the synthesized code from a TSL specification itself would aid in the development and debugging of TSL. With readable code, users can map changes in their specifications to changes in the synthesized code and debug their specifications. This mental mapping could allow users better understand the semantics of temporal operators as they observe how changing operations impact the synthesized code.

Pruning invalid transitions generated from violated assumptions is one area of potential optimization of the synthesized code. TSL synthesis generates code to handle violations of assumptions because the synthesized controller must be a system that can handle any stream of input, no matter if that input violates an assumption of the environment or not. This behavior is due to the hardware synthesis roots of reactive synthesis, where the system (a circuit) will continue to operate no matter the given inputs. In the software setting, we now face the question of how to handle invalid transitions.

The generated Mealy machine contains transitions telling the system what to do with inputs that violate the assumptions of the environment (from the guarantee block of the TSL specification). These transitions are invalid in the sense that they violate user-specified assumptions. Consequently, this leads to substantial portions of non-executable and perplexing code, particularly in complex systems with numerous states and assumptions. One solution to achieve readable and more optimized synthesized code is to prune the code results so that repetitive states and transitions or those that violate the assumptions of the specification are removed from the output [6].

This tool is well suited to tackle the Temporal Operator Semantics (3.2) challenge.

4.7 Source Mapping

To better understand the semantics of a temporal logic formula in relation to the synthesized output, we propose source mapping between the specification and the generated code. Similar to the way source mapping can be used to explore code generation through compiler passes [22], this tool might be used to better understand the synthesis process. For a compiler, source mapping is a one-to-many mapping—a single line of code may generate multiple lines in the compiled target language. However, a key challenge of source mapping in reactive synthesis is that, unlike code compilation, the mapping from specification to target is many-to-many. That is, not only does a single line of the specification have an impact on many lines of code, a single line of generated code can be attributed to many (non-contiguous) lines of the specification. An effective source mapping algorithm for Temporal Stream Logic synthesis is still an open problem.

4.8 Unrealizable Cores

Another tool used in reactive synthesis to debug unrealizable specifications is the generation of unrealizable cores [23]. Unrealizable cores can be viewed as a fault-localization problem. The core of an unrealizable spec is the minimal set of guarantees—or constraints on the system—that render the synthesis problem unrealizable. After obtaining this minimal set, users have more information

about what is causing their specifications to be unrealizable. This tool is well suited to tackle the Unrealizability (3.5) challenge.

4.9 Counter Strategies

One of the main challenges of reactive synthesis is dealing with unrealizable specifications or specifications of systems that have no correct controller implementation. After writing a specification, attempting synthesis, and getting that the specification is unrealizable, users may not know where to begin to debug their spec. As a result, substantial tooling has been developed to combat the problem of unrealizability [24], [19], including the generation of counter-strategies specifically for reactive systems [25]. A counter-strategy in program synthesis is a strategy the environment can use to falsify the specification no matter how the system moves. For developers debugging an unrealizable specification, a counter-strategy can provide information about the inputs from the environment that can cause the system to fail and what areas of the specification must be refined.

If a challenge about a temporal operator leads to an unrealizable specification, localizing the fault may lead to a correction in the understanding of temporal operator semantics. This tool is well suited to tackle the Temporal Operator Semantics (3.2) and Unrealizability (3.5) challenges. The debugging process outlined above may not have been possible without that use of a counter strategy. However, how to best present counter strategies to users is still an open question.




Table 2. This table proposes a mapping between common challenges that developers have in their understanding of TSL and current tools that can be utilized to combat those challenges.

5 Mapping TSL Challenges to Tools

To summarize, we propose a map, Table 2, between common challenges with Temporal Stream Logic and tools to address those challenges. This map can serve as a space for future work in the field of advancing the learnability and usability of reactive synthesis and temporal logic specification languages in practice.

6 Case Study

To further demonstrate the repercussions of TSL challenges in reactive synthesis, we provide a practical example taken from a recently published paper [17]. As commented in their benchmark GitHub repository [26], the TSL specification, shown in Fig. 5, aims to specify “a ‘bidirectional’ (i.e. ‘full duplex’) escalator capable of moving up or down. Its direction is controlled by the number of people waiting to take the escalator; when there are none, the escalator stops.” As we will show below, this natural language specification does not match the TSL specification. Using our framework of challenges and tools, we identify the challenges that cause the mismatch between the natural language specification and the TSL specification, as well as the tools that could have helped to correct this specification.

```

// #LIA#
always assume {
  [steps <- bottom()] <-> ![bottom <- add bottom c1()];
  [steps <- up()] <-> ![top <- add top c1()];
}

always guarantee {
  // Bottom movements
  [bottom <- add bottom c1()] && [top <- top] <-> [ users <- add users c1() ];
  [bottom <- bottom] && [top <- sub top c1()] <-> [ users <- sub users c1() ];

  // Top movements
  [bottom <- bottom] && [top <- add top c1()] <-> [ users <- add users c1() ];
  [bottom <- sub bottom c1()] && [top <- top] <-> [ users <- sub users c1() ];

  // Directional components
  eq users c0() && eq bottom c0() && eq top c0() -> [steps <- stop()];
  eq users c0() && gt bottom top -> [steps <- up()];
  eq users c0() && !(gt bottom top) -> [steps <- down];
}

```

Figure 5. The faulty TSL specification for an escalator from a recent paper [17].

Recall that the `always assume` block specifies constraints on the input the system gets from the environment. Here, the spec restricts people from entering the bottom of the escalator when the escalator is moving down, and vice versa. Recall that the `always guarantee` block describes the system's response to input values. Lines 9-12 aim to describe the movement of the escalator from the bottom to the top. Lines 9 and 10 aim to describe that if one person enters the bottom of the escalator, the number of people on the bottom will increase by one and the number of people on the escalator will also increase by one. At the same time, the number of people on the top of the escalator remains the same. Lines 11 and 12 aim to describe that if one person exits the top of the escalator, the number of people on the top will decrease by one, the number of people on the escalator will decrease by one, and the number of people on the bottom of the escalator remains the same. Lines 15-18 follow the same pattern for when the escalator moves up. Lines 21-24 aim to describe the directional components of the escalator. If the number of users on the escalator, the number of users at the bottom, and the number of users on top are all equal to 0, the escalator will stop. Otherwise, if there are more users waiting at the bottom than the top, the escalator will go up, and vice versa.

While this specification is realizable, three conceptual challenges lead to a mismatch the natural language specification and the TSL specification.

Challenge: (Non) Reactive Systems. A (Non) Reactive Systems challenge refers to confusion about the optimal use of TSL (Sec. 3.3). Synthesizing non-reactive systems may not make a specification unrealizable, but poses risks of misalignment with user intentions. For instance, the escalator specification results in a non-reactive system as it does not incorporate environmental inputs. This challenge, mainly due to modeling issues, arises because `top` and `bottom` are treated as cells instead of inputs, as well as the fact that there are no temporal operators in the specification.

Challenge: Non-well-separation. In TSL, non-well-separation can occur when there is confusion about content allocation between the `assume` and `guarantee` blocks (Sec. 3.1). Here, lines 3 and 4 with system updates specified on the right side of an implication are in fact controlled by the system.

Challenge: Syntactic Problem. In line 24, there is a missing set of parenthesis after `down`, which makes it an input signal rather than a constant function. In this specification, `top` and `bottom` should be inputs, and `up()` and `down()` should be function terms that control the movement of the escalator. Similarly, in line 3 there is a syntactic mistake—`[steps <- bottom()]` should instead be `[steps <- down()]` in order to update the steps with the `down()` command.

7 Related Work

There have been other efforts to develop tools for logical specification engineering, including a block-based visual programming language for Linear Temporal Logic (LTL) [27] and IDE support for Property Specification Language (PSL) [28]. The Spectra Language has also integrated support for specification engineering, including tools for debugging unrealizable specifications and non-well-separation [29]. Recent work [7] focused on identifying and categorizing common challenges about Linear Temporal

Logic for reactive synthesis. Building upon this, we focus on Temporal Stream Logic, and construct more generalized challenges that can be addressed through concrete tools in the form of developer support tools. There is an increasing awareness of the need to address the user experience of temporal logic specification engineering [3], [6], including using natural language and large language models to generate temporal logic specifications [30], [31]. This will potentially have an outsize impact on the viability of teaching temporal logic in classroom settings.

8 Conclusion and Future Work

Our goal with this work is to give an initial accounting of the various challenges and opportunities facing the synthesis community as we seek to give specification engineering the same first-class treatment as software engineering. We leave for future work a full user study to properly evaluate each of these tools and verify the claims made in Table 2. Additionally, we have focused our work on specifications in TSL, in the belief that many of these concepts will generalize to other temporal logics. There are however temporal logics that are more expressive than TSL. For example, TSL is a fragment of Temporal Stream Logic Modulo Theories (TSL-MT) [17]. We hypothesize that when working with richer temporal logic languages, there may be further challenges and tools—one example is the issue of type safety that arises in TSL-MT.

Acknowledgments

This work was funded in part by NSF Grant CCF-2105208. This work also benefitted from discussions at Dagstuhl Seminar 23391: The Futures of Reactive Synthesis.

References

- [1] S. Gulwani, “Automating string processing in spreadsheets using input-output examples,” *ACM Sigplan Notices*, vol. 46, no. 1, pp. 317–330, 2011.
- [2] R. Bloem, S. Jacobs, and A. Khalimov, “Parameterized synthesis case study: Amba ahb (extended version),” *arXiv preprint arXiv:1406.7608*, 2014.
- [3] D. Ma’ayan and S. Maoz, “Using reactive synthesis: An end-to-end exploratory case study,”
- [4] G. Geier, P. Heim, F. Klein, and B. Finkbeiner, “Syntroids: Synthesizing a game for fpgas using temporal logic specifications,” in *2019 Formal Methods in Computer Aided Design (FMCAD)*, 2019, pp. 138–146. DOI: 10.23919/FMCAD.2019.8894261.
- [5] W. Choi, M. Vazirani, and M. Santolucito, “Program synthesis for musicians: A usability testbed for temporal logic specifications,” in *Asian Symposium on Programming Languages and Systems*, Springer, 2021, pp. 47–61.
- [6] R. Rothkopf, A. L. Cui, H. T. Zeng, A. Sinha, and M. Santolucito, “Towards the usability of reactive synthesis: Building blocks of temporal logic,” *Plateau Workshop*, 2023.
- [7] B. Greenman, S. Saarinen, T. Nelson, and S. Krishnamurthi, “Little tricky logic: Misconceptions in the understanding of ltl,” *arXiv preprint arXiv:2211.01677*, 2022.
- [8] B. Finkbeiner, F. Klein, R. Piskac, and M. Santolucito, “Temporal stream logic: Synthesis beyond the booleans,” in *International Conference on Computer Aided Verification*, Springer, 2019.
- [9] A. Pnueli, “The temporal logic of programs,” in *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, iee, 1977, pp. 46–57.
- [10] Y. Luo, P. Liang, C. Wang, M. Shahin, and J. Zhan, “Characteristics and challenges of low-code development: The practitioners’ perspective,” in *Proceedings of the 15th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, ser. ESEM ’21, Bari, Italy: Association for Computing Machinery, 2021, ISBN: 9781450386654. DOI: 10.1145/3475716.3475782. [Online]. Available: <https://doi.org/10.1145/3475716.3475782>.
- [11] G. Geier, P. Heim, F. Klein, and B. Finkbeiner, “Syntroids: Synthesizing a game for fpgas using temporal logic specifications,” in *2019 Formal Methods in Computer Aided Design (FMCAD)*, IEEE, 2019, pp. 138–146.

- [12] V. Raman, A. Donzé, D. Sadigh, R. M. Murray, and S. A. Seshia, "Reactive synthesis from signal temporal logic specifications," in *Proceedings of the 18th international conference on hybrid systems: Computation and control*, 2015, pp. 239–248.
- [13] D. J. Fremont and S. A. Seshia, "Reactive control improvisation," in *Computer Aided Verification: 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14–17, 2018, Proceedings, Part I*, Springer, 2018, pp. 307–326.
- [14] C. Gerstacker, F. Klein, and B. Finkbeiner, "Bounded synthesis of reactive programs," in *Automated Technology for Verification and Analysis: 16th International Symposium, ATVA 2018, Los Angeles, CA, USA, October 7–10, 2018, Proceedings 16*, Springer, 2018, pp. 441–457.
- [15] B. C. P. L. Lab. "Tsltools: Library and tools for the tsl specification format." accessed on 8-18-2022, Barnard College. (2021), [Online]. Available: <https://barnard-pl-labs.github.io/tsltools/>.
- [16] U. Klein and A. Pnueli, "Revisiting synthesis of gr (1) specifications," in *Haifa Verification Conference*, Springer, 2010, pp. 161–181.
- [17] W. Choi, B. Finkbeiner, R. Piskac, and M. Santolucito, "Can reactive synthesis and syntax-guided synthesis be friends?" In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 2022, pp. 229–243.
- [18] A. Cimatti, M. Roveri, V. Schuppan, and A. Tchaltev, "Diagnostic information for realizability," in *Verification, Model Checking, and Abstract Interpretation: 9th International Conference, VMCAI 2008, San Francisco, USA, January 7–9, 2008. Proceedings 9*, Springer, 2008, pp. 52–67.
- [19] R. Könighofer, G. Hofferek, and R. Bloem, "Debugging formal specifications: A practical approach using model-based diagnosis and counterstrategies," *International journal on software tools for technology transfer*, vol. 15, no. 5–6, pp. 563–583, 2013.
- [20] V. Schuppan, "Towards a notion of unsatisfiable and unrealizable cores for ltl," *Science of Computer Programming*, vol. 77, no. 7–8, pp. 908–939, 2012.
- [21] M. Resnick, J. Maloney, A. Monroy-Hernández, et al., "Scratch: Programming for all," *Communications of the ACM*, vol. 52, no. 11, pp. 60–67, 2009.
- [22] M. Godbolt, *Compiler explorer*, <https://godbolt.org/>, 2023.
- [23] S. Maoz and R. Shalom, "Unrealizable cores for reactive systems specifications," in *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, IEEE, 2021, pp. 25–36.
- [24] W. Li, L. Dworkin, and S. A. Seshia, "Mining assumptions for synthesis," in *Ninth ACM/IEEE International Conference on Formal Methods and Models for Codesign (MEMPCODE2011)*, IEEE, 2011.
- [25] R. Alur, S. Moarref, and U. Topcu, "Pattern-based refinement of assume-guarantee specifications in reactive synthesis," in *Tools and Algorithms for the Construction and Analysis of Systems: 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11–18, 2015, Proceedings 21*, Springer, 2015.
- [26] B. C. P. L. Lab. "Tslmt: Benchmarks." archived by Jan 28, 2023, Barnard College. (2022), [Online]. Available: <https://github.com/Barnard-PL-Labs/temos/blob/art-eval-pldi22/benchmarks/escalator/bidirectional/bidirectional.tslmt>.
- [27] S. Yamashita, M. Tsunoda, and T. Yokogawa, "Visual programming language for model checkers based on google blockly," in *Product-Focused Software Process Improvement*, M. Felderer, D. Méndez Fernández, B. Turhan, M. Kalinowski, F. Sarro, and D. Winkler, Eds., Cham: Springer International Publishing, 2017, pp. 597–601, ISBN: 978-3-319-69926-4.
- [28] R. Armoni, L. Fix, A. Flaisher, et al., "The forspec temporal logic: A new temporal property-specification language," in *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, Springer, 2002, pp. 296–311.
- [29] S. Maoz and J. O. Ringert, "Spectra: A specification language for reactive systems," *Software and Systems Modeling*, vol. 20, no. 5, pp. 1553–1586, 2021.
- [30] J. X. Liu, Z. Yang, B. Schornstein, et al., "Lang2LTL: Translating natural language commands to temporal specification with large language models," in *Workshop on Language and Robotics at CoRL 2022*, 2022. [Online]. Available: <https://openreview.net/forum?id=VxfjGZzrdn>.

- [31] M. Cosler, C. Hahn, D. Mendoza, F. Schmitt, and C. Trippel, *Nl2spec: Interactively translating unstructured natural language to temporal logics with large language models*, 2023. arXiv: 2303.04864 [cs.LG].