

Synthesizing a Lego Forklift Controller in GR(1): A Case Study

Shahar Maoz Jan Oliver Ringert

School of Computer Science
Tel Aviv University

Reactive synthesis is an automated procedure to obtain a correct-by-construction reactive system from a given specification. GR(1) is a well-known fragment of linear temporal logic (LTL) where synthesis is possible using a polynomial symbolic algorithm. We conducted a case study to learn about the challenges that software engineers may face when using GR(1) synthesis for the development of a reactive robotic system. In the case study we developed two variants of a forklift controller, deployed on a Lego robot. The case study employs LTL specification patterns as an extension of the GR(1) specification language, an examination of two specification variants for execution scheduling, traceability from the synthesized controller to constraints in the specification, and generated counter strategies to support understanding reasons for unrealizability. We present the specifications we developed, our observations, and challenges faced during the case study.

1 Motivation and Context

Reactive synthesis is an automated procedure to obtain a correct-by-construction reactive system from a given specification, if one exists. The time complexity for synthesis of a reactive system from a linear temporal logic (LTL) formula is double exponential in the length of the formula [18]. However, limited fragments of LTL together with symbolic implementations exhibit more practical time complexities. One such fragment is General Reactivity of rank 1 (GR(1)), where synthesis is possible using a polynomial symbolic algorithm [3, 17].

The availability of efficient synthesis algorithms, as in the case of GR(1), and the guarantee of implementations being correct by construction, motivate applications in software engineering. GR(1) synthesis has been recently used in various application domains, including robotics [11], scenario-based specifications [14], aspect languages [13], and event-based behavior models [4], to name a few.

We conducted a case study to explore the benefits and current challenges of using existing tools and implementations of GR(1) synthesis to develop a reactive system. The research objectives of our case study were to learn about the following questions:

- What are challenges faced when using a GR(1) synthesis tool to synthesize a software controller for a robotic system?
- Is the use of LTL specification patterns helpful to formulate assumptions and guarantees?
- Is it easy to understand reasons for unrealizability?
- Do successfully synthesized controllers work as expected? If not, how can one understand why?

Our goal is to learn what reactive GR(1) synthesis offers and what it lacks to be successfully applied by software engineers in a model-based development process. On a wider scale, we want to understand what is required to make a synthesis specification language and synthesis tools available and accepted by reactive systems software engineers.

It is important to note that our case study concerns the initial development of a specification and synthesized controller. This case study was not about applying new methods and tricks for specification rewriting to make synthesis faster or to optimize the synthesized controller, e.g., as in case studies of the AMBA AHB protocol [2, 7]. When including excerpts of the developed specifications in this paper we decided to not rewrite LTL formulas in a more *clever* way but to present them as written during specification development.

The context of our case study is the synthesis of software for a robotic system. Synthesis is limited to a single controller that interacts with its environment. The specification is written and analyzed by a software engineer and the synthesized controller is used without modification for automatic, direct code generation and deployment to a robotic system. The tools used in our case study are:

- a symbolic GR(1) synthesis algorithm implementation from [3] using the JTLV framework [19] including the synthesis of counter strategies for unrealizable specifications;
- the AspectLTL [13] input language, with syntax similar to SMV, for specifying environment assumptions and system guarantees, with syntax highlighting and code completion;
- an implementation for traceability between a synthesized controller and its temporal specification based on [16]; and
- a catalog of LTL specification patterns [5] and their GR(1) templates [12].

The GR(1) synthesis problem is to find a controller that realizes a given specification over a set of environment variables and a set of system variables. A GR(1) specification consists of:

- constraints on initial assignments (constraints without temporal operators),
- safety constraints over current and next assignments (constraints of the form $\mathbf{G}(\text{exp})$ where the expression exp is limited to past time temporal operators and the **next** operator), and
- liveness constraints, more specifically justice constraints, that should hold infinitely often (constraints of the form $\mathbf{G F}(\text{exp})$ where the expression exp has no temporal operators).

All constraints are either assumptions, i.e., obligations of the environment, or guarantees, i.e., obligations of the system. Intuitively, if all assumptions are satisfied by the environment a synthesized controller satisfies all guarantees. If no such controller exists, the specification is called unrealizable. A detailed introduction to LTL, past time LTL, and GR(1) synthesis is available from [3].

In this case study we use an extension of the GR(1) synthesis problem where assumptions and guarantees may also be specified using LTL specification patterns [5]. The extension is described in [12].

Apart from the above GR(1) synthesis implementations we have used the MontiArcAutomaton language [21] for modeling the components of the robotic system and as a concrete syntax for the synthesized controllers. We have used code generators from this intermediate representation to Java for deployment to the Lego Mindstorms NXT platform¹.

2 Case Design and Variant

The task of the case study was to develop a specification for a Lego forklift robot shown on the left side of Figure 1. The forklift is an actual Lego robot² we have constructed and experimented with in our lab.

¹LEGO Mindstorms website: <http://www.lego.com/en-us/mindstorms>

²Robot based on these building instructions: <http://www.nxtprograms.com/NXT2/forklift/steps.html>

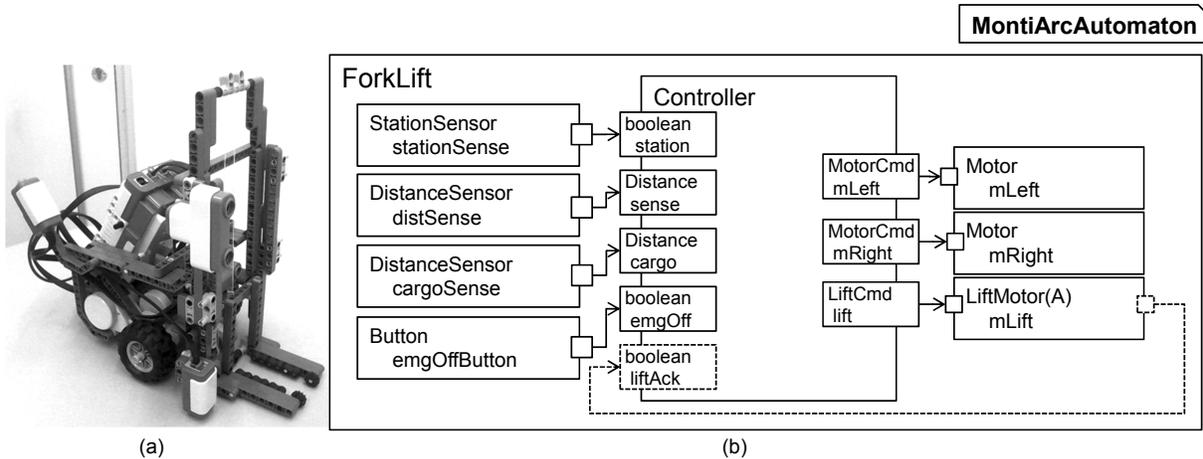


Figure 1: (a) The Lego forklift robot with four sensors and three actuators. (b) The logical software architecture of the robot with wrappers for sensors and actuators and the main component `Controller` to be synthesized (data types of input and output ports defined in Figure 2). The dashed elements describe a second variant with feedback of the lift motor to acknowledge completion of lifting or dropping the fork.

<i>enum</i> Distance	<i>enum</i> MotorCmd	<i>enum</i> LiftCmd
CLOSE	FWD	LIFT
FAR	STP	DROP
	BWD	NIL

Figure 2: Data type definitions for ports of component `Controller` shown in Figure 1

A criterion for success of synthesis was not only to obtain and inspect a synthesized controller but also to deploy it to the real robot and see that its behavior makes sense.

The forklift shown in Figure 1 has a sensor to determine whether it is at a station, two distance sensors to detect obstacles and cargo, and an emergency button to stop it. It has two motors to turn the left and right wheels and one motor to lift the fork. The case definition consists of an initial set of informal requirements for the behavior of the forklift:

1. Do not run into obstacles.
2. Only pick up or drop cargo at stations.
3. Do not attempt to lift cargo if cargo is lifted.
4. Always keep on delivering cargo.
5. Never drop cargo at the station where it was picked up.
6. Stop moving if emergency off switch is pressed.

Formalization and refinement of these requirements into guarantees and the elicitation of suitable assumptions was part of the case study execution.

Together with the above list of informal requirements the logical software architecture of the forklift was defined before case study execution. It is depicted as a component and connector model in Figure 1 (b). The components on the left side are hardware wrappers that read sensor values and publish

them as messages on their output ports. The output ports of the sensor components are connected to input ports of component `Controller`. The output ports of component `Controller` are connected to three components on the right that receive commands and encapsulate access to the motors of the forklift. The datatypes of input and output ports as well as their names are written on the ports of component `Controller`. Datatypes other than `boolean` are defined as enumerations in Figure 2.

The execution of the robot is performed in a cycle: read sensor data, execute controller, perform actions. When deploying the synthesized controller on the forklift robot the output produced by transitions between states of the controller manipulates the robots environment through its actuators. We decided on two scheduling variants to integrate the synthesized controller and the real world.

V1.Delayed In the first variant the execution of transitions of the synthesized controller is delayed to give the physical robot enough time to execute tasks of the actuators. The delay has to be large enough to, e.g., completely lift or drop the fork, complete a 90 degrees turn, or back up from an obstacle. We set the delay to 2000ms.

V2.Continuous In the second variant the controller is executed continuously without any delay. This variant uses a technique inspired by Raman et al. [20] to synthesize a reactive controller for continuous control. The setting requires the controller to be aware whether actions have completed or not. In our case driving actions get feedback from distance sensors but an additional feedback signal had to be added from the lift motor to acknowledge completion of its actions (shown as a dashed line in Figure 1).

Depending on the scheduler and actuator implementations more variants are possible, e.g., a variant where the scheduler pauses execution of the synthesized controller while an actuator performs a task (e.g., similarly by Kress-Gazit et al. [10]).

3 Resulting Specifications

Development of the specifications started with the first variant **V1.Delay** in 7 versions with incremental addition of features. When the first variant was almost complete the development of variant **V2.Continuous** started based on the existing specification. The final specifications for both variants are available from [22].

The first and common part to both specifications is a schematic translation of the input and output ports of component `Controller` shown in Figure 1 to environment variables (`VARENV`) and system variables (`VAR`) declared in Listing 1. The following excerpts of the two specifications refer to variable names and short names defined in the `DEFINE` block of Listing 1. The only difference between the variable declarations for **V1.Delay** and **V2.Continuous** is the environment variable `liftAck` in line 6 added only in variant **V2.Continuous**.

The two system variants are however very different. The main difference between both systems is that the first variant of the forklift is scheduled in steps of 2000ms (it waits for 2 seconds after each read input, compute, write output iteration), while the second variant does not use any delay. To illustrate this difference: when executing the controller synthesized from **V1.Delay** without the delay the forklift runs over cargo it was expected to have lifted up; it may initiate dropping cargo at a station but only completes dropping it after leaving the station; during this process it detects the dropping cargo as an obstacle and tries to drive clear from it.

```

1  VARENV
2  cargo : { CLEAR, BLOCKED };
3  sense : { CLEAR, BLOCKED };
4  station : boolean;
5  emgOff : boolean;
6  liftAck : boolean; -- only in V2.Continuous
7
8  VAR
9  mLeft : { FWD, STOP, BWD };
10 mRight : { FWD, STOP, BWD };
11 lift : { LIFT, DROP, NIL };
12
13 DEFINE -- defines to ease reading specs
14 backing := mLeft = BWD & mRight = BWD;
15 stopping := mLeft = STOP & mRight = STOP;
16 turning := mLeft = BWD & mRight = FWD | mLeft = FWD & mRight = BWD;
17 forwarding := mLeft = FWD & mRight = FWD;
18 dropping := lift = DROP;
19 lifting := lift = LIFT;

```

Listing 1: Environment variables (VARENV), system variables (VAR), and definitions of short names (DEFINE) of the controller to synthesize as an implementation of component `Controller` from Figure 1 (b)

3.1 Specifications

We give an overview of both specifications developed as part of the case study. We start with **V1.Delay** and later discuss the changes in **V2.Continuous** compared to **V1.Delay**.

V1.Delay: Assumptions and Guarantees

The specification document for **V1.Delay** contains 7 assumptions, 12 guarantees, and 1 auxiliary variable.

Of the 7 assumptions, 1 is a safety constraint, and 6 are LTL specification patterns: 5 instances of the response pattern P26 and 1 instance of a bounded existence pattern P15 (patterns numbered as appearing on the website of [5] and in our catalog of GR(1) templates [12]). An example for a typical response pattern is shown in Listing 2, ll. 1-3. The pattern expresses the assumption that if the robot is backing or turning it will reach a state where both distance sensors are clear unless it decides to go forward or stop, i.e., the robot may *escape* obstacles. A safety assumption about the environment is that the reading of the sensor `station` will not change when the forklift is stopping (see Listing 2, ll. 5-6). Another assumption is formulated using pattern P15 is shown in Listing 2, l. 8-9. It expresses that the robot will encounter at most two low obstacles between stations.

Of the 12 guarantees of **V1.Delay**, 1 constrains the initial state, 8 are a safety constraints, 1 is a justice constraint, and 2 are LTL specification patterns (P09 and P20). One auxiliary variable is defined (we treat the variable definition and assignments as a special case although the assignments are currently implemented as four safety guarantees). Our current implementation distinguishes manually added auxiliary variables from other variables by the prefix `spec_` preceding the name of the variable.

The specification uses pattern P09 as a guarantee that the forklift has to leave the station between lifting cargo and delivering it (Listing 3, ll. 1-2). Another guarantee uses pattern P20 to express that after leaving a station cargo cannot be dropped until the forklift reaches a station (Listing 3, ll. 4-5).

```

1 ASSUMPTION -- backing or turning clears sensors
2   Globally (backing | turning) leads to
3     ((sense=CLEAR & cargo=CLEAR) | forwarding | stopping)); --P26
4
5 ASSUMPTION -- station does not change when stopping
6   G (stopping -> station = next(station));
7
8 ASSUMPTION -- at most blocked twice between stations
9   After (!atStation) have at most two (lowObstacle) until (atStation); --P15

```

Listing 2: Three assumptions of variant **V1.Delay** using LTL specification patterns [5] with equivalent representations in GR(1) [12]

```

1 GUARANTEE -- have to leave station to deliver
2   (!atStation) occurs between (lift=LIFT) and (lift=DROP); --P09
3
4 GUARANTEE -- don't drop cargo after leaving station until arriving
5   Globally (lift!=DROP) after (!atStation) until (atStation) --P20

```

Listing 3: Two guarantees of of variant **V1.Delay** using LTL specification patterns [5] with equivalent representations in GR(1) [12]

```

1 VAR -- new auxiliary variable to "remember" when cargo is loaded
2   spec_loaded : boolean;
3 GUARANTEE
4   !spec_loaded; -- initial value false
5   G (lifting -> next (spec_loaded)); -- set loaded when lifting
6   G (dropping -> ! next (spec_loaded)); -- unset loaded when dropping
7   G (lift = NIL -> next (spec_loaded) = spec_loaded); -- preserve value

```

Listing 4: Definition of the auxiliary variable spec_loaded

```

1 ASSUMPTION -- find station to drop when forwarding (unless backing or stopping)
2   Globally (forwarding & spec_loaded) leads to
3             ((station & cargo=CLEAR) | backing | stopping));
4
5 ASSUMPTION -- find station and cargo when forwarding (unless backing or stopping)
6   Globally (forwarding & !spec_loaded) leads to
7             ((station & cargo=BLOCKED) | backing | stopping));
8
9 GUARANTEE -- do not lift again when cargo is loaded
10  G (spec_loaded -> !lifting);

```

Listing 5: An assumption and a guarantee of variant **V1.Delay** using the auxiliary variable `spec_loaded` introduced in Listing 3

```

1 GUARANTEE -- emergency stop
2   G (emgOff -> (stopping & lift=NIL));
3
4 GUARANTEE -- main goal to always eventually deliver cargo
5   G F ((lift = DROP) | emgOff | lowObstacle);

```

Listing 6: More guarantees of variant **V1.Delay**: the emergency stop feature and the main justice constraint of the forklift

In Listing 4 a new auxiliary variable `spec_loaded` is introduced for the purpose of simplifying the specification. This variable is used by the specifier to keep track on whether the forklift has loaded cargo or not. The value of the variable is determined by safety constraints. The variable is used in the specification in both assumptions and guarantees as shown in Listing 5. The assumption in lines 1-3 expresses that if the forklift goes forward and has cargo loaded it will eventually find a station where it can deliver cargo unless it stops or goes backward.

Two additional guarantees in Listing 6 describe the emergency stop feature of the forklift and the main justice constraint to always eventually deliver cargo. The main justice constraint of the forklift is shown in Listing 6, ll. 4-5. It does not only contain the expected `lift = DROP`, i.e., delivering cargo, but also the alternatives `emgOff` and `lowObstacle`. These alternatives represent environment actions that if occurring infinitely often liberate the forklift from its obligation.

V2.Continuous: Assumptions and Guarantees

We now describe the differences between the specification of the first variant **V1.Delay** and the second variant **V2.Continuous** of the forklift specifications. The second variant applies a method inspired by Raman et al. [20] for continuous control. The main idea is to add a new variable for every action which signals completion of the action. In our case study this method is only necessary for the completion of lifting and dropping cargo. The environment variable `liftAck` is added (see Listing 1) to signal completion of the actions of component `LiftMotor` (see Figure 1).

The specification document for variant **V2.Continuous** contains 9 assumptions, 14 guarantees, and 2 auxiliary variable definitions. None of the assumptions or guarantees of **V1.Delay** were removed, three assumptions were added, and two guarantees were added.

Of the 2 added assumptions, 1 is a safety constraint, and 1 is an instance of pattern P26. The auxiliary variable `spec_waitingForLifting` is declared and completely defined in Listing 7, ll. 9-15 in the same way as `spec_loaded`. The variable is `true` iff a lifting command was issued and completion has not been acknowledged yet. The assumption in Listing 8, ll. 1-2 expresses that an acknowledgment

```

1  VAR
2    spec_loaded : boolean;
3  GUARANTEE -- updated assignments of var spec_loaded
4    !spec_loaded;
5    G (liftAck & !spec_loaded -> next(spec_loaded));
6    G (liftAck & spec_loaded -> next(!spec_loaded));
7    G (!liftAck -> spec_loaded = next(spec_loaded));
8
9  VAR -- variable to keep track of waiting for lifting
10   spec_waitingForLifting : boolean;
11  GUARANTEE
12    !spec_waitingForLifting;
13    G ((lift=LIFT | lift=DROP) -> next(spec_waitingForLifting));
14    G (liftAck -> ((lift=LIFT | lift=DROP) | next(!spec_waitingForLifting)));
15    G (!(lift=LIFT | lift=DROP) | liftAck) -> spec_waitingForLifting = next(
    spec_waitingForLifting);

```

Listing 7: Auxiliary variable definitions in **V2.Continuous**: the assignments of `spec_loaded` were modified and variable `spec_waitingForLifting` was added

```

1  ASSUMPTION -- always eventually complete lifting
2    Globally (spec_waitingForLifting) leads to (liftAck);
3
4  ASSUMPTION -- only acknowledge if waiting for it
5    G (next(liftAck) -> (spec_waitingForLifting & !liftAck));

```

Listing 8: New assumptions added in variant **V2.Continuous** about acknowledgments of executing the LIFT and DROP commands

`liftAck` will eventually follow every waiting. A second assumption ensures that acknowledgments are only sent if `spec_waitingForLifting` is `true`. The assignment of variable `spec_loaded` has been modified. The assignment now depends on the previous value of variable and the new input `liftAck` (see ll. 1-6).

Both added guarantees are safety constraints, as shown in Listing 9. The first expresses that the forklift stops when it waits for the completion of a lifting action (see Listing 9, ll. 1-2), and the second expresses that it does not issue new lifting commands when waiting for completion (see Listing 9, ll. 4-5).

3.2 Synthesis Times and Controller Sizes

For both variants a controller that implements the specification can be synthesized in a few seconds using the Java-based BDD engine that comes with JTLV [19]. The sizes of both specification variants are summarized in the upper half of Table 1. We report the assumption and guarantee constraints, the total number of variables including auxiliary variables (the state space referred to as N in the time complexity $O(nmN^2)$ of GR(1) synthesis is 2 to the power of the number of Boolean variables required to repre-

```

1  GUARANTEE -- while lift performs action stop
2    G (spec_waitingForLifting -> stopping);
3
4  GUARANTEE -- do not send another request when waiting for completion
5    G (spec_waitingForLifting -> lift=NIL);

```

Listing 9: Added guarantees of variant **V2.Continuous**

	V1.Delay	V2.Continuous
Assumptions (patterns)	1 safety, 5 times P26, P15	2 safety, 6 times P26, P15
Guarantees (patterns)	1 initial, 8 safety, 1 justice, P09, P20	1 initial, 8 safety, 1 justice, P09, P20
Boolean Variables (auxiliary)	4 environment, 6 system, 1 manual, 12 pattern	5 environment, 6 system, 2 manual, 13 pattern
Checking Realizability	0.2 sec	0.7 sec
Controller Construction	1.8 sec	1.3 sec
States of Controller	3412	2888

Table 1: For both variants we report the size of the specification, times for checking realizability and controller construction in seconds, and the size of a synthesized controller.

sent environment, system, and auxiliary variables). Environment and system variables amount to 2^{10} in **V1.Delay** and to 2^{11} in **V2.Continuous**. For the specifications of variant **V1.Delay** and **V2.Continuous** developed in this case study 13 and 15 auxiliary variables were automatically (and implicitly) added to support specification patterns or explicitly added as auxiliary variables in the specification. The GR(1) synthesis problem after the translation of patterns to GR(1), via the templates described in [12], had 6 environment and 3 system justice goals in variant **V1.Delay** and 7 environment and 3 system justice goals in variant **V2.Continuous**.

In the lower half of Table 1 we report the time it took the synthesis algorithm to decide realizability and the additional time consumed by the controller construction phase, in seconds. Times are shown as reported by JTLV running on an ordinary laptop with Java 7, Windows 7 64bit, 8GB RAM, and an Intel i7 CPU with 3.0 GHz. Synthesis times for all intermediate versions during specification development confirmed that synthesis is conveniently fast.

3.3 Running Synthesized Controllers on the Forklift

We used code generators implemented in our group for the NXJ LeJOS platform³ to generate code and directly deploy the software components shown in Figure 1 (b) to the LEGO NXT forklift shown in Figure 1 (a). The components `StationSensor` (light value measured on the ground), `DistanceSensor` (ultrasonic distance sensor), and `Button` (touch sensor) have generic implementations in Java that wrap the LeJOS sensor API. The components `Motor` and `LiftMotor` have generic implementations that wrap the LeJOS motor API.

To execute the synthesized controller of variant **V1.Delay** on the forklift we had to adapt a set of platform specific parameters: number of degrees motors `mLeft` and `mRight` rotate backward and forward, number of degrees the lifting motor rotates, distance to obstacles detected by `distSense`, and distance to cargo detected by `cargoSense`. In the second variant **V2.Continuous** the motors `mLeft` and `mRight` do not rotate a fixed amount of degrees but move continuously.

The components on the robot are executed in execution cycles. Every execution cycle starts with reading all sensor values, executes the controller, and ends with executing all actuators. In variant **V1.Delay** we added a fixed delay of 2000ms after the execution of the actuators to allow the forklift finishing all actions. In variant **V2.Continuous** a delay was not added and the execution time of one

³Website of LeJOS NXJ: <http://www.lejos.org/nxj.php>

cycle on the robot was around 50ms.

We observed that variant **V1.Delay** provided more reliable sensor readings than **V2.Continuous** because it measured values after completing its actions in a resting position. For variant **V2.Continuous** erroneous sensor readings had serious effects. As one example, when the station sensor falsely detected a station the forklift stopped to drop cargo. A second reading when stopped did not report the station and thus caused a violation of an environment assumption (Listing 2, ll. 5-6), and so the forklift went into an infinite loop of stopping. As another example, unstable readings of the ultrasonic sensor led to detecting and not detecting obstacles at the sensor’s threshold level. This happened when the forklift was not at a station and led to a violation of the assumption that at most two obstacles occur (Listing 2, ll. 8-9). In this case the forklift kept going forward in an infinite loop not stopping at obstacles anymore.

Some works have addressed the challenge of synthesizing controllers which are more robust to assumption violations, e.g., Ehlers and Topcu [6] suggested an approach where the synthesized controller allows violations of safety assumptions up to some constant number of times.

4 Observations and Challenges

We now report some of our observations and challenges faced during specification development.

O1: Differences between V1.Delay and V2.Continuous

The way that the synthesized controllers of variants **V1.Delay** and **V2.Continuous** are executed and interact with their environment is fundamentally different. Thus, we found it surprising that their GR(1) specifications are still very similar. Only three assumptions and two guarantees were added. One reason for the similarity is that **V2.Continuous** is based on **V1.Delay**. It is still interesting that most existing assumptions and guarantees also remain valid for the second variant, although it is based on a continuous execution scheduling, without delays. The main reason we did not have to adapt many assumptions seems to be the use of the response patterns already in **V1.Delay** (five instances) instead of explicitly referring to immediate successor states.

O2: Manually adding auxiliary variables helpful

During the development of the two specifications we found it very helpful to add auxiliary variables, as they assisted us in making relevant states explicit. As an example, the auxiliary variable `spec_loaded` (Listing 4) is defined to be `true` iff the forklift has loaded cargo. This information is derived and not provided as a sensor input. The variable `spec_loaded` appears in two assumptions and two guarantees. The past LTL formula

$$\text{PREV } (\text{lift} \neq \text{DROP} \text{ SINCE } \text{lift} = \text{LIFT})$$

could replace the auxiliary variable but we believe that it helps readability of the formulas to make the property explicit with a new name.

Technically, adding an auxiliary variable to the specification requires that its value is determined by complete and deterministic safety constraints. This mechanism is discussed by Bloem et al. [3] for supporting past LTL in GR(1) synthesis.

O3: Support for patterns helpful

LTL specification patterns [5] allow one to express high-level temporal patterns in a convenient way that are otherwise complicated and error-prone to express in LTL. In the limited fragment of GR(1) correctly expressing these behaviors becomes even more complicated due to the limitation of available operators and no nesting. We provide GR(1) templates for 52 of the 55 LTL specification patterns [12].

During the case study we found that using patterns is very helpful for expressing more complicated temporal properties. Moreover, using patterns gave us better confidence that the specification matches our intention.

A specifically useful pattern in assumptions was the response pattern P26. Instances of this pattern appeared as 5 out of 7 assumptions in **V1.Delay** (6 out of 9 in **V2.Continuous**). This pattern seems to be well suited for describing robotic systems where one should assume that some actions of actuators eventually have an impact on sensor values.

It is important to note that the addition of patterns comes at a price. Most patterns cannot be expressed without the addition of auxiliary variables. In both variants of our case study more than half of the Boolean variables encoding the statespace were auxiliary variables, implicitly added in the translation of the patterns to the GR(1) form based on our templates. The resulting synthesis problems of the case study are however still solved in a few seconds.

C1: Environment vs. real environment

During the case study it turned out that it is difficult to use assumptions to describe realistic environments. A very simple example that appeared early during development of **V1.Delay** is the assumption that turning of the robot will make the distance sensor signal clear:

```
G (turning -> next (sense=CLEAR));
```

Our inspection of the synthesized controller concluded positively. When deploying the controller to the forklift it turned and stopped in a corner of the room. The assumption was too strong for the forklift's real environment. In a corner, turning 90 degrees leaves the robot facing another wall. The response pattern **Globally** (turning) **leads to** (sense=CLEAR) would solve this particular problem but the expressed assumption is still too strong and we observed another undesired behavior: after being blocked the forklift turned once but then stopped and waited for the obstacle to disappear. A corrected version (also handling driving backwards) is shown in Listing 2, ll. 1-3.

A different challenge appeared when adding a guarantee that the robot should pick up cargo and not drop it between stations:

```
Globally (lift!=DROP SINCE lift=LIFT)
           after (!atStation) until (atStation);
```

The addition of this guarantee made the specification unrealizable. The past time formula `lift!=DROP SINCE lift=LIFT` embedded in the specification pattern requires that the forklift leaves a station only when cargo has been picked up. We decided that it might be reasonable to assume that at every station the forklift may pick up cargo. It turned out too complicated for us to formulate this assumption in GR(1). A possible assumption might require encoding the area of each station and valid navigation of the forklift to allow it to explore the station for cargo. Finally, we weakened the above guarantee to the one in Listing 3, ll. 4-5. The modified version allows the forklift to leave stations without cargo.

To summarize, in the case study we faced both the challenge of formulating assumptions that are too strong for the real environment of the forklift and the challenge of not being able to formulate reasonable assumptions due to difficulties in expressing them. Both cases are not easy to address. Specifically the first can lead to successful synthesis of a controller that fails in a real environment.

C2: Undesired realizable case: system forces environment to violate assumptions

When the system can force the environment to violate its assumptions the controller often does not act as the engineer would have expected it to act. Consider the following two assumptions from an early version of **V1.Delay**. The first assumption is that if the forklift does not move the station sensor value will not change. The second assumption is that the forklift will eventually leave a station if it moves forward.

```
G (stopping -> station = next(station));
Globally (forwarding) leads to (!station);
```

We synthesized a controller and in some cases the forklift running the controller stopped and did not continue to move. To understand this behavior in the example above we enabled our synthesis tool to annotate every transition with one of three reasons for it to be included in the controller. The possible reasons are [3]: satisfying a justice constraint, working towards satisfying a justice constraint, or preventing a justice constraint of the environment. The annotation of each transition of the controller includes the justice constraint from the specification (see [16] for more details about traceability). Using this traceability information, which links transitions in the controller to elements of the specification, we learned the reasons for the synthesized strategy of the controller. We found out that in the example above, the forklift goes forward on a station. If it is still on the station in the next step it stops forever and thus prevents the environment from satisfying the justice constraint of the response pattern.

In this simple example with a controller of 80 states we were able to find an explanation we could trace to the specification and modify the response pattern as shown in Listing 5, ll. 5-7. It would have been helpful to automatically find cases where the system forces the environment to violate assumptions for an engineer to decide whether this behavior is desired or not. Towards the end of our specification development process synthesized controllers simply became too large for manual inspection.

Klein and Pnueli [8] defined environments where the system cannot force a violation as *well-separated*. They suggest a modified GR(1) game to check whether an environment is well-separated. This may be a direction towards addressing this challenge.

C3: Unrealizable case

In many cases adding a new feature expressed as a set of assumptions and guarantees led to an unrealizable specification. In some cases we were able to find mistakes quickly in the added assumptions and guarantees. Many times we initially forgot to add alternatives in safety constraints leading to their unsatisfiability by the system. In more complicated cases we asked our tool to synthesize a counter strategy that represents an environment forcing all system strategies to lose. We (interactively, as described in [16, 15]) executed moves of our intended system and learned where our strategy loses against the environment. This often led to better understanding of the reasons for unrealizability.

In some cases it was however very difficult to understand the reason for unrealizability, trace it back to the specification, and fix it. As an example, consider the introduction of the new environment variable

`liftAck` and the auxiliary variable `spec_waitingForLifting` (see Listing 7). We added the assumption that lifting can only be acknowledged if the controller is expecting it:

```
G (next(liftAck) -> spec_waitingForLifting);
```

The specification of variant **V2.Continuous** with the above assumption is unrealizable. A synthesized counter strategy has 3735 states. The Java code generated by our synthesis tool for interactive exploration of the counter strategy failed to compile due to its size. Printing the counter strategy including information on successors disabled by safety properties ran out of memory. It was not easily possible to reduce the synthesis problem by removing irrelevant parts from the specification. Lifting of cargo is related to movement of the motors and the cargo sensor. The movement of motors is related to the distance sensor.

We manually executed the counter strategy by inspecting the generated text output. Many steps were repetitive (long chains of apparently similar states) as we were working towards forcing the environment to present a station with cargo, and after lifting forcing it to acknowledge lifting. Right after acknowledging lifting the environment acknowledged lifting again. The double acknowledgment set and immediately afterwards disabled the variable `spec_loaded` (see Listing 7, ll. 5-6). The double acknowledgment was possible because `liftAck` disables `spec_waitingForLifting` only in the next step (see Listing 7, ll. 14). We adapted the above assumption as shown in Listing 8, l. 5.

While counter strategies help understanding reasons of unrealizability their handling by our tools turned out to be insufficient for larger specifications. In the future we plan to examine how recent work by others, e.g., [1, 9], may help in addressing the unrealizability challenge.

5 Conclusion

We have presented a case study of the development of a software controller for a forklift robot using GR(1) synthesis tools. Rather than examining how to write most elegant and efficient specifications we focused on challenges for software engineers in the process of specification development. We showed the specifications of two variants of the controller. On the one hand, our observations are that extensions of the specification language with auxiliary variables and higher-level specification patterns support writing specifications with better confidence. On the other hand, with growing specification size, understanding reasons for synthesized behavior and for unrealizability turned out to be major challenges.

This case study is part of our larger project on bridging the gap between the theory and algorithms of reactive synthesis on the one hand and software engineering practice on the other. In many aspects it demonstrates the different challenges awaiting us.

Acknowledgments The authors thank the anonymous reviewers of the SYNT 2015 workshop for their helpful comments. Jan O. Ringert acknowledges support from a postdoctoral Minerva Fellowship, funded by the German Federal Ministry for Education and Research. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 638049, SYNTECH).

References

- [1] Rajeev Alur, Salar Moarref & Ufuk Topcu (2013): *Counter-strategy guided refinement of GR(1) temporal logic specifications*. In: *Formal Methods in Computer-Aided Design, FMCAD 2013, Portland, OR*,

- USA, October 20-23, 2013, IEEE, pp. 26–33. Available at http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=6679387.
- [2] Roderick Bloem, Swen Jacobs & Ayrat Khalimov (2014): *Parameterized Synthesis Case Study: AMBA AHB*. In Krishnendu Chatterjee, Rüdiger Ehlers & Susmit Jha, editors: *Proceedings 3rd Workshop on Synthesis, SYNT 2014, Vienna, Austria, July 23-24, 2014.*, EPTCS 157, pp. 68–83. Available at <http://dx.doi.org/10.4204/EPTCS.157.9>.
- [3] Roderick Bloem, Barbara Jobstmann, Nir Piterman, Amir Pnueli & Yaniv Sa’ar (2012): *Synthesis of Reactive(1) Designs*. *J. Comput. Syst. Sci.* 78(3), pp. 911–938. Available at <http://dx.doi.org/10.1016/j.jcss.2011.08.007>.
- [4] Nicolás D’Ippolito, Víctor A. Braberman, Nir Piterman & Sebastián Uchitel (2013): *Synthesizing nonanomalous event-based controllers for liveness goals*. *ACM Trans. Softw. Eng. Methodol.* 22(1), p. 9. Available at <http://doi.acm.org/10.1145/2430536.2430543>.
- [5] Matthew B. Dwyer, George S. Avrunin & James C. Corbett (1999): *Patterns in Property Specifications for Finite-State Verification*. In: *ICSE*, ACM, pp. 411–420.
- [6] Rüdiger Ehlers & Ufuk Topcu (2014): *Resilience to intermittent assumption violations in reactive synthesis*. In Martin Fränzle & John Lygeros, editors: *17th International Conference on Hybrid Systems: Computation and Control (part of CPS Week), HSCC’14, Berlin, Germany, April 15-17, 2014*, ACM, pp. 203–212, doi:10.1145/2562059.2562128. Available at <http://doi.acm.org/10.1145/2562059.2562128>.
- [7] Yashdeep Godhal, Krishnendu Chatterjee & Thomas A. Henzinger (2013): *Synthesis of AMBA AHB from formal specification: a case study*. *STTT* 15(5-6), pp. 585–601, doi:10.1007/s10009-011-0207-9. Available at <http://dx.doi.org/10.1007/s10009-011-0207-9>.
- [8] Uri Klein & Amir Pnueli (2010): *Revisiting Synthesis of GR(1) Specifications*. In Sharon Barner, Ian G. Harris, Daniel Kroening & Orna Raz, editors: *Hardware and Software: Verification and Testing - 6th International Haifa Verification Conference, HVC 2010, Haifa, Israel, October 4-7, 2010. Revised Selected Papers, Lecture Notes in Computer Science 6504*, Springer, pp. 161–181. Available at http://dx.doi.org/10.1007/978-3-642-19583-9_16.
- [9] Robert Könighofer, Georg Hofferek & Roderick Bloem (2013): *Debugging formal specifications: a practical approach using model-based diagnosis and counterstrategies*. *STTT* 15(5-6), pp. 563–583, doi:10.1007/s10009-011-0221-y. Available at <http://dx.doi.org/10.1007/s10009-011-0221-y>.
- [10] Hadas Kress-Gazit, Georgios E. Fainekos & George J. Pappas (2007): *Where’s Waldo? Sensor-Based Temporal Logic Motion Planning*. In: *2007 IEEE International Conference on Robotics and Automation, ICRA 2007, 10-14 April 2007, Roma, Italy*, IEEE, pp. 3116–3121, doi:10.1109/ROBOT.2007.363946. Available at <http://dx.doi.org/10.1109/ROBOT.2007.363946>.
- [11] Hadas Kress-Gazit, Georgios E. Fainekos & George J. Pappas (2009): *Temporal-Logic-Based Reactive Mission and Motion Planning*. *IEEE Trans. Robotics* 25(6), pp. 1370–1381. Available at <http://dx.doi.org/10.1109/TRO.2009.2030225>.
- [12] Shahar Maoz & Jan Oliver Ringert (2015): *GR(1) Synthesis for LTL Specification Patterns*. In: *ESEC/FSE*, ACM. <http://smlab.cs.tau.ac.il/syntech/patterns/>.
- [13] Shahar Maoz & Yaniv Sa’ar (2011): *AspectLTL: an aspect language for LTL specifications*. In Paulo Borba & Shigeru Chiba, editors: *AOSD*, ACM, pp. 19–30. Available at <http://doi.acm.org/10.1145/1960275.1960280>.
- [14] Shahar Maoz & Yaniv Sa’ar (2012): *Assume-Guarantee Scenarios: Semantics and Synthesis*. In: *MODELS, LNCS 7590*, Springer, pp. 335–351. Available at http://dx.doi.org/10.1007/978-3-642-33666-9_22.
- [15] Shahar Maoz & Yaniv Sa’ar (2013): *Counter play-out: executing unrealizable scenario-based specifications*. In: *ICSE, IEEE / ACM*, pp. 242–251. Available at <http://dl.acm.org/citation.cfm?id=2486821>.

- [16] Shahar Maoz & Yaniv Sa'ar (2013): *Two-Way Traceability and Conflict Debugging for AspectLTL Programs*. *T. Aspect-Oriented Software Development* 10, pp. 39–72. Available at http://dx.doi.org/10.1007/978-3-642-36964-3_2.
- [17] Nir Piterman, Amir Pnueli & Yaniv Sa'ar (2006): *Synthesis of Reactive(1) Designs*. In: *VMCAI, LNCS 3855*, Springer, pp. 364–380. Available at http://dx.doi.org/10.1007/11609773_24.
- [18] Amir Pnueli & Roni Rosner (1989): *On the Synthesis of a Reactive Module*. In: *POPL*, ACM Press, pp. 179–190.
- [19] Amir Pnueli, Yaniv Sa'ar & Lenore D. Zuck (2010): *JTLV: A Framework for Developing Verification Algorithms*. In: *CAV, LNCS 6174*, Springer, pp. 171–174. Available at http://dx.doi.org/10.1007/978-3-642-14295-6_18.
- [20] Vasumathi Raman, Nir Piterman & Hadas Kress-Gazit (2013): *Provably correct continuous control for high-level robot behaviors with actions of arbitrary execution durations*. In: *2013 IEEE International Conference on Robotics and Automation, Karlsruhe, Germany, May 6-10, 2013*, IEEE, pp. 4075–4081, doi:10.1109/ICRA.2013.6631152. Available at <http://dx.doi.org/10.1109/ICRA.2013.6631152>.
- [21] Jan Oliver Ringert, Bernhard Rumpe & Andreas Wortmann (2014): *Architecture and Behavior Modeling of Cyber-Physical Systems with MontiArcAutomaton*. *Aachener Informatik-Berichte, Software Engineering 20*, Shaker Verlag.
- [22] *SYNTECH forklift website*. <http://smlab.cs.tau.ac.il/syntech/forklift/>.