Symbolic construction of GR(1) contracts for systems with full information

Ioannis Filippidis and Richard M. Murray

Abstract— This work proposes a symbolic algorithm for the construction of assume-guarantee specifications that allow multiple agents to cooperate. Each agent is assigned goals expressed in a fragment of linear temporal logic known as generalized Streett with one pair, GR(1). These goals may be unrealizable, unless each agent makes additional assumptions, about the behavior of other agents. The algorithm constructs a contract among the agents, in that only the infinite behavior of the given goals is constrained, known as liveness, not the finite one, known as safety. This defers synthesis to a later stage of refinement, modularizing the design process. We prove that there exist GR(1) games that do not admit any GR(1) contract. For this reason, we formulate contracts with nested GR(1) properties and auxiliary communication variables, and prove that they always exist. The algorithm's fixpoint structure is similar to GR(1) synthesis, enjoying time complexity polynomial in the number of states, and linear in number of recurrence goals.

I. INTRODUCTION

The design and construction of a large system relies on the ability to divide the problem into smaller ones. Each subproblem involves a subset of the system, and may itself be refined further into smaller problems. The subsystems that result from the smaller problems are considered as modules of the larger system. In many cases, the modules interact with each other, either physically, or as software, or both. For this reason, the interaction between modules needs to be constrained, to ensure that they can perform their operation as intended.

A representation with precise syntax and semantics, or *specification*, is desirable to describe each module, and its interaction with other modules. Broadly, the system designer is faced with two problems. First, to specify the modules, with detail sufficient to allow for implementation, preferably automated. Second, to look at each module's specification, construct an implementation, and assemble the results.

Ioannis Filippidis and Richard Murray are with Control and Dynamical Systems, California Institute of Technology, Pasadena, CA, 91125, USA. E-mail: {ifilippi, murray}@caltech.edu

In this work, we are interested in automating the first step, modularization of a design that has been partially specified by a human. Note that human input is necessary, in one form or another, because the algorithm cannot know what the modules are intended for. These *primitive* specifications may not yield a coherent system. Starting from these, the proposed algorithm constructs a collection of new specifications for the modules, a *contract* [1], [2]. Each new specification is implementable, and, together, they refine the original intent.

A common refinement is *synthesis*, whereby the immediate behavior of all modules is designed at once. We want to avoid doing so, because such refinement is premature. Instead, we want to restrict only the *infinite* behavior of the system. We work under the assumption that the modules can observe all variables in the problem, known as *full information*. To this, we are motivated by the undecidability of distributed synthesis [3], [4], except for restrictions of the communication architecture [5], the desired behaviors [6], or the search space [7].

The paper is organized as follows. Sec. II reviews preliminaries on logic and games, where we formulate concepts for multi-player games, in an interleaving logic representation. Contracts are defined in Sec. III, and our purpose to find realizable module specifications high in the refinement hierarchy. We want contracts as Streett(1) specifications, but prove in Sec. III-B that these do not exist in general. For a restricted case, this limitation is lifted in Sec. IV, where we introduce nested GR(1) specifications, and an algorithm to construct them. In the presence of multiple recurrence goals, we prove in Sec. IV-B that auxiliary variables must be introduced, and thus define a distributed scheme of switching between nested GR(1) games, one for each recurrence goal. Examples are given in Sec. V, relevant work and conclusions are discussed in Sec. VI.

II. PRELIMINARIES

A. Temporal logic and μ-calculus

Linear temporal logic (LTL) [8] can be used to describe sets of infinite sequences $w : \mathbb{N} \to V \triangleq 2^{\mathcal{V}}, \mathbb{N} \triangleq \{0, 1, \dots\}$ over a set $\mathcal{V} \triangleq \{p, q, \dots\}$ of

propositional variables. Let $w[i \dots]$ denote the subsequence $w[i]w[i+1]\dots$. Besides the Boolean operators "and" $\wedge$, "or" $\vee$, "not" $\neg$, LTL includes the future operators "next" $\bigcirc p$, "always" $\square p$, "eventually" $\Diamond p$, and "until" $p\,\mathcal{U}q$, and past operators "weak previous" $\ominus p$ and "historically" $\boxminus p$. The semantics of LTL [8], [9] defines whether a word w is a *model* of a formula φ at position i, denoted by $(w,i) \models \varphi$, as follows:

- $(w,i) \models \bigcirc\varphi$ iff $(w,(i+1)) \models \varphi$
- $(w,i) \models \varphi\,\mathcal{U}\psi$ iff $\exists k \in \mathbb{N}.\ \big(w,(i+k)\big) \models \psi$ and $\forall j \in \mathbb{N}_{<k}.\ (w,(i+j)) \models \varphi$,

and $\Diamond\varphi \triangleq \top\,\mathcal{U}\varphi$ and $\square\varphi \triangleq \neg\Diamond\neg\varphi$.

- $(w,i) \models \ominus\varphi$ iff $i = 0$ or $\big(w,(i-1)\big) \models \varphi$
- $(w,i) \models \boxminus\varphi$ iff $\forall j \in \mathbb{N}_{\leq i}.\ \big(w,(i-j)\big) \models \varphi$.

Given a propositional formula φ over variables in $\mathcal{V}$, define the set of models $[\![\varphi]\!]_{\mathcal{V}} \triangleq \{u \in V \mid u \models \varphi\}$. Given an LTL formula ψ, define the set of models $[\![\psi]\!]_{\mathcal{V}} \triangleq \{u \in V^{\mathbb{N}} \mid (u,0) \models \varphi\}$. For readability, occasionally we refer to a set $[\![f]\!]$ by mentioning f. We use first-order LTL, i.e., integer variables over $\mathcal{D}$ (nullary function variables [10]), with $V \triangleq \mathcal{D}^{\mathcal{V}}$. We use φ' as a shorthand for $\bigcirc\varphi$. With the Cantor topology over $V^{\mathbb{N}}$ [11], if a formula describes an open (closed) set, then it is called a liveness (safety) property [12]. For example, $\square\Diamond p$ describes liveness, and $\square(p \leftrightarrow \bigcirc\neg p)$ safety. Properties of the form $\square\Diamond p$ are known as *recurrence*, and $\Diamond\square p$ as *persistence* [13].

The modal μ-calculus [14] extends modal logic with a least (greatest) fixpoint operator μ (ν). For example, if $[\![S]\!]$ is a set of system states, and $[\![f(S)]\!]$ the set of states that can be reached in one time step from some state in $[\![S]\!]$, then $[\![\mu X.\ I \vee f(X)]\!]$ is the set of states that are inside, or reachable in a finite number of time steps, starting from within the set $[\![I]\!]$. For comparison, $[\![\mu X.\ f(I) \vee f(X)]\!]$ contains states that are reachable in a finite (nonzero) number of time steps, and $[\![\mu X.\ f(X)]\!] = \emptyset$. For a formal definition, we refer to [15], [16].

B. Games

1) Representation: The behavior of a collection of n modules can be represented as a game with multiple players. Here, we consider turn-based games, where a single player moves at each change of state [17]. The game is synchronous, in that players move in a fixed order that repeats. Define the set $\mathcal{X}_j$ of variables controlled by player j. By x_j we will denote both the tuple of symbols in $\mathcal{X}_j$, as well as a tuple of values assigned to those symbols. Let $x \triangleq (x_0, x_1, \dots, x_{n-1})$ be the aggregate state. Let $\bar{x}_j$ denote (either a tuple of, or an assignment to) variables in $\bigcup_{k=0,k\neq j}^{n-1} \mathcal{X}_k$. We assume full information, meaning that each player can observe all the

variables. Let i be an auxiliary variable that signifies which player should move next, $\mathcal{V} \triangleq \{i\} \cup \bigcup_{j=0}^{n-1} \mathcal{X}_j$.

We use an interleaving representation of the game in temporal logic, with one player moving in each time step. Let $\mathcal{D} \triangleq \mathbb{N} \cup \{\text{true}, \text{false}\}$ and $V \triangleq \mathcal{D}^{\mathcal{V}}$ be the set of assignments (game "nodes") to variables in $\mathcal{V}$. Partition $V_j \subseteq V$ contains nodes with $i = j$. Player j moves only from nodes in V_j, by assigning values to variables in x_i, and incrementing i. The transition relation of player j is an action formula $\hat{\rho}_j(x, x_j')$ (Boolean-valued formula over primed and unprimed variables) [18]. In an interleaving representation of a synchronous turn-based game, the transition relations are $\rho_j(x, x_j', i) \triangleq \text{ite}\big(i \neq j,\ x_j' = x_j,\ \hat{\rho}_j \wedge (i' = i \oplus_n 1)\big)$, $\bar{\rho}_j(x, \bar{x}_j', i) \triangleq \bigwedge_{k\neq j} \rho_k(x, x_k', i)$, where $a \oplus_c b \triangleq (a + b) \mod c$, and $\text{ite}(a,b,c)$ the ternary conditional connective.

2) Winning: A *play* is an infinite sequence of nodes in V that a game visits, as players move. Given a play w, and a set W_j called *winning condition* for player j, if $w \in W_j$, then player j is a *winner* in the play w. In practice, a temporal logic formula φ_j is used to *describe* the winning condition W_j. Any LTL winning condition φ_j can be represented as a conjunction of a safety with a liveness formula [18], [12]. The complexity of solving a game depends on the liveness formula, as captured by its *Rabin index* [11], [13], [15]. Each Rabin index is associated with a syntactic fragment of LTL, with normal forms those of Streett $\bigwedge(\square\Diamond \vee \Diamond\square)$ (conjunctive), and Rabin $\bigvee(\square\Diamond \wedge \Diamond\square)$. Each formula $\square\Diamond \vee \Diamond\square$ ($\wedge$) is called a Streett (Rabin) *pair*. A *generalized* pair has the form $(\bigvee^N \Diamond\square) \vee \bigwedge^M \square\Diamond$. Games with a single generalized Streett pair (generalized reactivity(1) or GR(1)) [19] can be solved in time quadratic in the number of nodes, and linear in each of N, M.

Let the set $M_j \subseteq \mathbb{N}$ with $|M_j| < \infty$ be a finite amount of memory. A (deterministic) *strategy* for player j is a (partial) function $T : V_j \times M_j \to \bar{V}_j \triangleq V \setminus V_j$. Let $\mathcal{T}_j$ be the set of strategies for player j, and $\bar{\mathcal{T}}_j \triangleq \prod_{k\neq j} \mathcal{T}_k$. For a node $u \in V$, define the set $\text{Plays}(u, T_0, \dots, T_{n-1}) \subseteq V^{\mathbb{N}}$ of all the plays possible if, for each $j < n$, player j uses strategy T_j. For deterministic strategies, $\text{Plays}(u, T_0, \dots, T_{n-1})$ is a singleton. Let a *realization from u* be $\mathcal{L}(j, u, T_j) \triangleq \bigcup_{\bar{T}_j \in \bar{\mathcal{T}}_j} \text{Plays}(u, T_0, \dots, T_{n-1})$ [20]. Given a property φ, define the *winning set* of player j as $[\![\text{Win}(j, \varphi)]\!] \triangleq \{u \in V \mid \exists T \in \mathcal{T}_j.\ \mathcal{L}(j, u, T) \subseteq [\![\varphi]\!]\}$. For the games we consider, playing against arbitrary, but finite, strategies, and against arbitrary input sequences coincide [21, p.619]. So, the above definition is not restrictive. The notion that no player should violate their allowed transitions, before any other player can

be expressed using past LTL [16]. Here, we define this precedence in an interleaving representation, as

Definition 1: Let ρ_e, ρ_s be action, W_e, W_s liveness formulae. Define the *strict implication* operator $\overset{\text{sr}}{\Rightarrow}$ as $(\Box\rho_e \wedge W_e) \overset{\text{sr}}{\Rightarrow} (\Box\rho_s \wedge W_s) \triangleq \Box((\ominus\Box\rho_e) \to \rho_s) \wedge ((\Box\rho_e \wedge W_e) \to W_s) = (\Box\rho_e) \overset{\text{sr}}{\Rightarrow} (\Box\rho_s \wedge (W_s \to W_e))$. We are interested in games with winning conditions that have favorable complexity, in particular

Definition 2 (Generalized Streett(1) [19]): Assume that, for all $k \in \mathbb{N}_{<N}, r \in \mathbb{N}_{<M}$, $P_{jk}(x,i), R_{jr}(x,i)$ are assertions (Boolean formulae over unprimed variables). Then, the LTL formula $\varphi_j \triangleq (\Box\bar\rho_j) \overset{\text{sr}}{\Rightarrow} (\Box\rho_j \wedge (\bigwedge_k \Box\Diamond P_{jk} \to \bigwedge_r \Box\Diamond R_{jr}))$ describes a GR(1) property for player j.

III. CONTRACTS

We are interested in a method for constructing assume-guarantee specifications for a set of modules. The resulting specifications must be *realizable* [20], i.e., for some *common* initial condition(s), for each module, there should exist a strategy that implements its specification. In the context of full information, we can directly construct a strategy for each module, and thus solve this problem. Informally, a realization does not contain any quantification – it has all been removed by this point. At the cost of restricting the temporal property far beyond what is necessary. We do not want to do so. Instead, we want to refine the given specifications, in a way that fewer behaviors be removed.

In particular, a deterministic strategy defines a property that, for each input sequence, contains at most one play. Still, for each input sequence, there remain infinitely many plays that, by this point, have been excluded from the temporal property. We want to preserve these plays as admissible behaviors. Thus, we obtain a temporal property that is higher in the partial order associated with refinement.

For the following discussion, we will need some definitions. As specification language, we choose LTL. We call specifications also properties, and define

Definition 3: Given n players, a *contract* is a partition of variables among players that defines which player controls each variable, together with an initial condition set $\llbracket I \rrbracket \neq \emptyset$, and a tuple of properties $(\varphi_0, \ldots, \varphi_{n-1})$, such that $I \to \bigwedge_j \text{Win}(j, \varphi_j)$. A GR(1) contract contains GR(1) properties.

Definition 4: Let $\varphi(\mathcal{V}), \psi(\mathcal{W})$ be two properties over variables in $\mathcal{V}, \mathcal{W}$. If $\llbracket\varphi\rrbracket_{\mathcal{V}\cup\mathcal{W}} \subseteq \llbracket\psi\rrbracket_{\mathcal{V}\cup\mathcal{W}}$, then the property φ *refines* (or implements) ψ [20]. Given φ, *synthesis from $u \in V$* is the construction of a realization ψ from u that refines φ. Synthesis is refinement. Not any refinement is synthesis.

The designer requires that the collective behavior of the players satisfy $\bigwedge_j \varphi_j(\mathcal{V})$. If, for some I, $\varphi \triangleq (\varphi_0, \ldots, \varphi_{n-1})$ is a contract, then we are done. Otherwise, we want to find a tuple $\psi \triangleq (\psi_0, \ldots, \psi_{n-1})$ that refines the given intent, defined as $\llbracket\bigwedge_j \psi_j\rrbracket_\mathcal{V} \subseteq \llbracket\bigwedge_j \varphi_j\rrbracket_\mathcal{V}$. In addition, we do not want ψ to restrict the finite behavior (safety) of the players more than φ. The limits of any finite behavior contained in a property P form its *closure* $\mathcal{C}_\mathcal{K}(P) \triangleq \{w \in (\mathcal{D}^\mathcal{K})^\mathbb{N}| \forall k \in \mathbb{N}. \exists\sigma \in P. w[0\ldots k] = \sigma[0\ldots k]\}$ [22], [23]. Therefore, the requirement to preserve in ψ any finite behavior allowed by φ (thus avoid synthesizing) can be expressed as $\mathcal{C}(\llbracket\bigwedge_j \psi_j\rrbracket_\mathcal{V}) = \mathcal{C}(\llbracket\bigwedge_j \varphi_j\rrbracket_\mathcal{V})$. This extends to team games the notion of least restrictive safety assumption [23]. A contract $\psi(\mathcal{V})$ with the previous two properties *refines* the tuple $\varphi(\mathcal{V})$ (safety-preserving refinement over $\mathcal{V}$).

Problem 5: Given a GR(1) tuple $\varphi(\mathcal{V})$, does there exist a GR(1) contract $\psi(\mathcal{V})$ that refines φ? In the following sections, we consider Prob. 5 for incrementally larger classes of behaviors.

A. Safety

Assume that, for each player j, a safety property $\varphi_j = \Box\rho_j$ is given as design intent. If a single player -1 controlled all variables in $\mathcal{V}$, then let the *cooperatively winning set* $C \triangleq \text{Win}(-1, \bigwedge_j \varphi_j)$. Note that $\llbracket C \rrbracket = \emptyset$ implies that $\llbracket\bigwedge_j \text{Win}(j, \varphi_j)\rrbracket = \emptyset$, so no contract exists in that case. A necessary condition for the existence of a contract is that $\llbracket C \rrbracket \neq \emptyset$, i.e., φ must be cooperatively winning. Define the existential predecessors of F as $\text{Pre}_j(F) \triangleq \lambda x. \lambda i. (i = j) \wedge \exists x'_j. \hat\rho_j(x, x'_j) \wedge F|_{x'_j/x_j}(\bar{x}_j, x'_j, j\oplus_n 1)$, where x'_j/x_j signifies substitution of the symbol tuple x_j by the corresponding symbols in x'_j. Denote $\text{Pre}(F) \triangleq \bigvee_j \text{Pre}_j(F)$, and the iterated predecessors $\text{Pre}^*(F) \triangleq \mu X. F \vee \text{Pre}(X)$. Given φ, the set $\llbracket C \rrbracket$ can be computed as the greatest fixpoint $\nu Z. \text{Pre}(Z)$ [19].

For $S \subseteq V$, let $\text{Safe}(S) \triangleq \{w \in V^\mathbb{N}| \forall k \in \mathbb{N}. w[k] \in S \wedge (w, 0) \models \bigwedge_j \Box\rho_j\}$. By extension of results for two-player games [23, Thm.7], it is $\text{Safe}(\llbracket C \rrbracket) = \mathcal{C}(\llbracket\bigwedge_j \varphi_j\rrbracket)$. So, the closure contains exactly those plays that remain inside C, and satisfy all ρ_j. This can be required from player j with the formula $\rho_{C,j} \triangleq C \wedge C|_{x'_j/x_j}$. Let $\bar\rho_{C,j} \triangleq \bigwedge_{k\neq j} \rho_{C,k}$. In the case of safety formulae $\varphi_j = \Box\rho_j$, the resulting specifications are $\psi_j \triangleq (\Box(\bar\rho_j \wedge \bar\rho_{C,j})) \overset{\text{sr}}{\Rightarrow} \Box(\rho_j \wedge \rho_{C,j})$. With these specifications, starting from inside C, no player can exit C before any other player, as we prove in [24, Prop.21], and, by definition of C, no player is blocked. With $I \triangleq C$, it follows that $\forall j \in \mathbb{N}_{<n}. I \to \text{Win}(j, \psi_j)$. Assumption $\llbracket C \rrbracket \neq \emptyset$ implies $\llbracket I \rrbracket \neq \emptyset$. By definition of $\text{Safe}(S)$ and ψ_j, it is $\llbracket\bigwedge_j \psi_j\rrbracket = \text{Safe}(\llbracket C \rrbracket)$. For any safety property P,

$\mathcal{C}(P) = P$ [22]. Therefore, $\mathcal{C}(\llbracket \bigwedge_j \psi_j \rrbracket) = \llbracket \bigwedge_j \psi_j \rrbracket = \text{Safe}(\llbracket C \rrbracket) = \mathcal{C}(\llbracket \bigwedge_j \varphi_j \rrbracket) = \llbracket \bigwedge_j \varphi_j \rrbracket$. So ψ is a contract, and refines tuple φ.

An analogous result ensures that the safety component of properties with recurrence goals $\Box \Diamond G_{jr}$ can be ensured by computing $\nu Z. \bigwedge_j \nu Z_j. Z \wedge \bigwedge_r \text{Pre}^*(G_{jr} \wedge \text{Pre}(Z_j))$, as we prove in [24].

B. Recurrence

Advancing from Sec. III-A, we consider tuples φ with $\varphi_0 = \Box \rho_0 \wedge \Box \Diamond G$ and, for $j > 0$, $\varphi_j = \Box \rho_j$. We prove that, unlike safety, for a single recurrence $\Box \Diamond G$ there may not exist any GR(1) contract ψ over $\mathcal{V}$. For two players, Prob. 5 means that we want to restrict the infinite behavior of player 1, without removing any transitions from any player, thus

Proposition 6: ASSUME: Define the transition relations ρ_0, ρ_1 by the game graph of Fig. 1, the set of nodes $V \triangleq \{s_0, \ldots, s_7\}$, and the goal $\llbracket G \rrbracket \triangleq \{s_6\}$ of player 0. PROVE: For all sets $\llbracket P \rrbracket \subseteq V$, with $\psi_1 \triangleq (\Box \rho_0) \overset{\text{sr}}{\Rightarrow} (\Box \rho_1 \wedge \Box \Diamond P)$ and $\psi_0 \triangleq (\Box \rho_1 \wedge \Box \Diamond P) \overset{\text{sr}}{\Rightarrow} (\Box \rho_0 \wedge \Box \Diamond G)$, it is $\llbracket \text{Win}(0, \psi_0) \rrbracket \cap \llbracket \text{Win}(1, \psi_1) \rrbracket = \emptyset$.

PROOF SKETCH: For any P that does not include any node from $\{s_0, \ldots, s_3\}$, player 1 cannot win, because player 0 can force, and keep, the play outside of P. For similar reasons, P should intersect each of $\{s_0, s_1\}$ and $\{s_2, s_3\}$. If P intersects $\{s_0, s_1\}$ and $\{s_2, s_3\}$, then player 1 can win by always moving from s_4 to s_1, when the play comes to s_4. This forces a visit to either both s_0 and s_1, or both s_2 and s_3. Therefore, there does not exist a P, such that both players have a winning strategy. A formal [25] proof can be found in [24].

Prop. 6 proves that, even in the case of one recurrence goal for only one player, there may not exist any GR(1) contract. Also, this implies that, even if a problem is cooperatively realizable by the players, existing approaches for constructing GR(1) assumptions [26], [27], [28] may not find any GR(1) assumptions, because none exist.

One may wonder why we didn't consider $\Box \Diamond s_4$ as antecedent in ψ_1, since with $\llbracket P \rrbracket = \{s_5\}$, this may seem to work. The problematic aspect in Fig. 1 is that, in GR(1), $\Box \Diamond s_4$ would also have to be conjoined into the consequent of ψ_0, to tell player 0 to do so. Player 0 can violate its liveness goals, thus $\Box \Diamond s_4$, as long as that falsifies the antecedent $\Box \Diamond P$ of ψ_0, in accord with Def. 2. This is known as *trivial realizability* [23], [16], [29]. In other words, there is *circularity* of dependencies for the liveness properties [30], [31], [32], [22] that each player needs to assume about the other. In Sec. IV, we propose an algorithm that constructs nested GR(1) specifications, which enjoy complexity properties similar to GR(1) games.


Fig. 1: Game where a liveness assumption realizable by player 1, and sufficient for player 0, does not exist. Player 0 (player 1) moves from disks (boxes).

IV. NESTED GR(1) GAMES

A. Single recurrence

In Sec. III-B we proved that Prob. 5 can fail to have a solution. To overcome this limitation, we extend the class of properties to

Definition 7: If, in syntax, $\varphi_j = \Box \bar{\rho}_j \overset{\text{sr}}{\Rightarrow} \Big(\Box \rho_j \wedge \bigwedge_m \Box ((P_m \wedge \bigwedge_l \Box \Diamond \neg \eta_{ml}) \to \Diamond Q_m) \wedge \Box (Q_m \to \bigwedge_l \Box \Diamond \neg \xi_{ml}) \Big)$, and, in semantics, $\llbracket P_{m-1} \rrbracket \subseteq \llbracket Q_m \rrbracket \subseteq \llbracket P_m \rrbracket$, $\llbracket \eta_{ml} \rrbracket \subseteq \llbracket P_m \wedge \neg Q_m \rrbracket$ and $\llbracket \xi_{ml} \rrbracket \subseteq \llbracket Q_m \wedge \neg P_{m-1} \rrbracket$, then we call φ_j a *nested GR(1)* property for player j.

If $a \subseteq b \subseteq \cdots \subseteq c$, then the sets $a, b, \ldots, c$ form a *chain*. A nested GR(1) property can be regarded as a request-response chain, in analogy to Rabin and Streett chains (parity) [33, p.13] [34]. Note though that synthesis for a nested GR(1) property is similar to GR(1), whereas for parity games it is unknown whether a polynomial time algorithm exists.

A nested GR(1) property enables to isolate conditional assumptions, by partitioning the game into smaller ones. Each conjunct game becomes active in P_m, it has its own assumptions $\Box \Diamond \neg \eta_{ml}$, and the reachability goal Q_m, tighter than the games with larger m, on which it depends. This prevents circularity of liveness dependencies. The assumptions $\Box \Diamond \neg \eta_{ml}$ are unconditional. Assumptions that themselves depend on other liveness assumptions become objectives in their own game. The approach of nested games is reminiscent of McNaughton's recursive algorithm for solving parity games [35].

We propose Alg. 1, which computes a stack of nested games that cover the cooperatively winning set $\llbracket C \rrbracket$, and define a contract. So, a later visit to G is always possible, from any node in C. Define the *controllable predecessors* of F for player j as $\text{CPre}_j(F) \triangleq \lambda x. \lambda i. \bigvee_{k=0}^{n-1} ((k = i) \wedge \neg^{k \neq j} \exists x'_k. \hat{\rho}_k(x, x'_k) \wedge \neg^{k \neq j} F|_{x'_k/x_k}(\bar{x}_k, x'_k, k \oplus_n 1))$. The *attractor* $\text{Attr}_j(F) \triangleq \mu X. F \vee \text{CPre}_j(X)$ [36] contains those nodes from where player j is in, or can force a future visit, to F. Define the *controlled-escape* subset of S as $\text{Trap}_j(S, E) \triangleq \nu X. E \vee (\text{CPre}_j(X) \wedge S)$ that contains those nodes, from where player j can force either to remain inside $\text{Trap}_j(S, E)$, or move to E, or is already in E.

Fig. 2: Predicates computed by UNCONDASM, Alg. 1.

Theorem 8: ASSUME: A cooperatively realizable tuple $\varphi = (\varphi_0, \ldots, \varphi_{n-1})$ with $\varphi_0 = \Box\rho_0 \wedge \Box\Diamond G$ and $\forall k > 0.\ \varphi_k = \Box\rho_k$. PROVE: After $\mathcal{O}(n\,|V|^2)$ calls to CPre_j, the call GAMESTACK$(0, G, C, s)$ in Alg. 1 returns in list s a contract of nested GR(1) properties that refines the tuple φ, with $I = C$.

PROOF SKETCH: Function UNCONDASM finds the largest set r from where player j can keep the play inside the k-attractor of the j-attractor of goal g, as shown in Fig. 2. This yields an *unconditional* assumption that player j makes about player k. Iteration l of L3 produces an assumption $\Box\Diamond\neg\eta_{ml}$ of player j, with $\eta = \text{trap} = r$, which is also a guarantee $\Box\Diamond\neg\xi_{m'l}$ by player k, with $\xi_{m'l} = \text{trap}$. The fixpoint $\nu.\mu.\mu$ for computing r in UNCONDASM has alternation depth 1 [37], so it invokes $\text{CPre}_?$ $\mathcal{O}(|V|)$ times. Due to determinacy [11], [21], and the definition of the cooperatively winning set C, for every n nested calls at L15, at least one call to UNCONDASM at L5 removes a node from *uncovered*. So, GAMESTACK makes $\mathcal{O}(n\,|V|^2)$ calls to $\text{CPre}_?$. A formal proof can be found in [24].

Note that searching for fewer assumptions, inducing a smaller winning set, can be exponentially expensive, as proved for syntactic recurrence formulae in [27]. Conceptually, the nesting of games has common elements with modular game graphs [38] and open temporal logic [39].

Example 1: Let us revisit the example of Fig. 1, to observe the algorithm's execution. Player 0 wants $\Box\Diamond G$. The first call to GAMESTACK will call UNCONDASM. Player 0 can force a visit to s_6 from the attractor $A = \text{Attr}_0(s_6) = s_5 \vee s_6$. Player 1 can force A from $B = \text{Attr}_1(A) = s_4 \vee s_5 \vee s_6$. But $r = \bot$, because player 1 can escape to s_1. So, a nested game is constructed over $s_0 \vee s_1 \vee s_2 \vee s_3 \vee s_4$, with player 1 wanting $\Diamond(s_5 \vee s_6)$. In the nested game, $A = \text{Attr}_1(s_5 \vee s_6) = s_4 \vee s_5 \vee s_6$. The attractor $B = \text{Attr}_0(s_4 \vee s_5 \vee s_6) = \top$, and player 0 can keep player 1 in there, until player 0 visits $s_4 \vee s_5 \vee s_6$. So, in the nested game, player 1 makes the assumption $\Box\Diamond\neg(s_0 \vee s_1 \vee s_2 \vee s_3)$. This covers the cooperative winning set, in this example the entire game graph.

B. Conjoined recurrence

Let us consider Prob. 5 again, this time for any GR(1) tuple φ. The difference with Sec. IV-A is that

---

Algorithm 1 Construction of nested GR(1) specification, for a single recurrence goal G.

```
1: def GAMESTACK(j, G, uncovered, stack)
2:     trap ← ⊤, goal ← G, stack ← set()
          ▷ Create unconditional assumptions η_ml
3:     while [[trap]] ≠ ∅ :            ▷ until stuck
4:         for k ≠ j :
5:             attr, trap ← UNCONDASM(j, k, goal)
6:             goal ← attr ∨ trap
7:             assumptions.add((k, □◇¬trap))
8:             if [[trap]] ≠ ∅ :
9:                 break
10:        game ← (j, goal, G, assumptions)
11:        stack.append(game)
12:        uncovered ← uncovered ∧ ¬goal
          ▷ Covered cooperatively winning set?
13:        if [[uncovered]] = ∅ :
14:            return
          ▷ Construct a nested game
15:        GAMESTACK(j ⊕_n 1, goal, uncovered, stack)
16: def UNCONDASM(j, k, g)
17:     A ← Attr_j(g)
18:     B ← Attr_k(A)
19:     r ← ¬A ∧ B ∧ Trap_j(B, A)
20:     return A, r
```

there may be more than one recurrence goals in φ. In this case, we prove with the counterexample of Fig. 3 that a GR(1) contract may not exist. This includes *any* GR(1) contract, not only those that refine φ. In other words, no transitions can be removed, and no additional recurrence goals can produce two GR(1) properties that both players can realize from the same set.

Proposition 9: ASSUME: Define the transition relations ρ_0, ρ_1 as in Fig. 3, and goals $G_1 \triangleq \{s_6\}, G_2 \triangleq \{s_0\}$ of player 0. PROVE: For all properties P such that $\psi_1 \triangleq (\Box\rho_0) \overset{\text{sr}}{\Rightarrow} (\Box\rho_1 \wedge P)$ and $\psi_0 \triangleq (\Box\rho_1 \wedge P) \overset{\text{sr}}{\Rightarrow} (\Box\rho_0 \wedge \Box\Diamond G_1 \wedge \Box\Diamond G_2)$ be GR(1) properties, it is $[[\text{Win}(0, \psi_0)]] \cap [[\text{Win}(1, \psi_1)]] = \emptyset$.

A proof can be found in [24]. The reason for this limitation is that auxiliary variables are needed for coordination between players, to ensure they play the same nested GR(1) game. This relates to the requirement for memory in strategies for GR(1) games [36], [16]. Each player can have their own memory, but if these local memories are not coordinated, the situation is similar to one player with two independent memories over a partition. In other words, additional *common* memory is required.

However, Prob. 5 does not apply any more, because ψ_j will be defined over $\mathcal{W}$, which is $\mathcal{V}$ aug-

Fig. 3: No GR(1) contract exists for this example.

mented with auxiliary variables. Given a property P over variables $\mathcal{W} \supseteq \mathcal{K}$, define projection on $\mathcal{K}$ as $\pi_\mathcal{K}(P) \triangleq \{w \in (\mathcal{D}^\mathcal{K})^\mathbb{N} | \exists \sigma \in P. \forall k \in \mathbb{N}. w[k] = \sigma[k] \cap \mathcal{D}^\mathcal{K}\}$. Define closure preservation over the original variables $\mathcal{V}$ as $\pi_\mathcal{V}\big(\mathcal{C}_\mathcal{W}([\![\bigwedge_j \psi_j]\!]_\mathcal{W})\big) = \mathcal{C}_\mathcal{V}([\![\bigwedge_j \varphi_j]\!]_\mathcal{V})$. If $\mathcal{V} \subseteq \mathcal{W}$, $\psi(\mathcal{W})$ is a contract, $[\![\bigwedge_j \psi_j]\!]_\mathcal{W} \subseteq [\![\bigwedge_j \varphi_j]\!]_\mathcal{W}$ and closure is preserved over $\mathcal{V}$, then contract $\psi(\mathcal{W})$ *refines tuple φ with memory $\mathcal{W} \setminus \mathcal{V}$.*

Problem 10: Does there exist a class of contracts that admit synthesis in polynomial time, such that, given a GR(1) tuple $\varphi(\mathcal{V})$, there always exists a contract $\psi(\mathcal{W})$ that refines φ with memory $\mathcal{W} \setminus \mathcal{V}$?

First, the transition relations should be restricted to ensure closure (safety component), as in Sec. III-A. As shown in Sec. IV-A, each recurrence goal can be reached from anywhere in the cooperatively winning set C, by using nested GR(1) properties for the players. Therefore, playing consecutively games that cycle through all the tuples of nested GR(1) properties, one for each recurrence goal, ensures that the play repeatedly visits each recurrence goal.

For this purpose, we introduce the *auxiliary counter variables* $g_j \in \mathbb{N}_{\leq N_j}$, where N_j the number of recurrence goals G_{jr} of player j. Each counter is used to indicate a recurrence goal of player j that is actively pursued by the players, each one using the strategy for the nested GR(1) game that corresponds to that goal. The behavior of the counters is described by the formulae $\bigwedge_r \square\big((g_j = r) \rightarrow \text{ite}(G_{jr}, \ g'_j = g_j + 1, \ g'_j = g_j)\big), \square\big((g_j = N_j) \rightarrow \text{ite}(\ominus(g_{j \oplus_n 1} = N_{j \oplus_n 1} - 1) \wedge (g_{j \oplus_n 1} = N_{j \oplus_n 1}), \ g'_j = 0, \ g'_j = N_j)\big)$ with initial conditions $(g_0 = 0) \wedge (\forall k > 0. \ g_k = N_k)$. The nested GR(1) game with index r is activated by conjoining $(g_j = r)$ with the antecedents P_m and Q_m in φ_j of Def. 7.

V. EXAMPLES

The proposed algorithms have been implemented [40] in PYTHON, using our packages `omega` [41], [42], for symbolic solution of games, and `dd`, for binary decision diagrams (BDD) [43], [42]. `dd` offers both a pure PYTHON implementation, and an identical interface via CYTHON to the C library CUDD [44], for demanding applications. They are available under a BSD license. The predicates in the constructed temporal properties are represented by BDDs.

1) Grid world: The first example involves two robots A, B that are moving in a narrow passage,

and need to cross each other. Their discrete world is described by cells, as in Fig. 4a, and moves are possible between adjacent cells. Define variable a (b) as the cell that robot A (B) currently is in. Each robot can remain stationary, and shall not move from, or to, a cell where another robot is, e.g., $\square(a \neq b \neq a')$ for A. Robot A (B) can move from cell $0 \mapsto 1 \mapsto 2 \mapsto 3 \mapsto 4 \mapsto 5 \mapsto 0$ ($0 \mapsto 5 \mapsto 4 \mapsto 3 \mapsto 2 \mapsto 1 \mapsto 0$, and $2 \mapsto 6 \mapsto 2$, and $0 \mapsto 7 \mapsto 0$). In other words, robot A (B) moves (counter)clockwise, and A cannot go to cells 7 and 6.

The construction of a contract for the property tuple $\varphi \triangleq (\varphi_a, \varphi_b)$ with $\varphi_a \triangleq \square\rho_a \wedge \square\diamond(a = 4)$ and $\varphi_b \triangleq \square\rho_b$. starts with the closure $\mathcal{C}([\![\varphi_a \wedge \varphi_b]\!])$, as discussed in Sec. III-A. This removes nodes of the game graph that are unsafe for $\square\diamond(a = 4)$. For example, $(a = 3) \wedge (b = 4)$ is such a node, because there is no way for A to go to $a = 4$ from there, even if both A, B were controlled by one player. The relation $\rho_{C,j}$, computed to ensure that both robots remain in the cooperatively winning set C, removes only transitions that "cross" the boundary from C to $\neg C$, as in [23]. For example, node $u \triangleq (a = 3) \wedge (b = 5) \wedge (i = 1)$ is in C, but $v \triangleq (a = 3) \wedge (b = 4) \wedge (i = 0)$ is not (recall that i signifies who should move next, with 0 for a and 1 for b). The transition from $u \mapsto v$ satisfies $\square\rho_a \wedge \square\rho_b$, but violates $\square\rho_a \wedge \square\rho_b \wedge \square\diamond(a = 4)$. In the following, ρ_a, ρ_b are constrained with $\rho_{C,a}, \rho_{C,b}$.

Secondly, Alg. 1 runs. It produces the specification ψ_a with $P_0 = C$, $Q_0 = (a = 4) \vee (b \in \{0, \ldots 3, 6, 7\})$ and one assumption $\eta_{0,0} = \big((a = 1) \wedge \big((b \in \{2, \ldots 5\}) \vee (b = 0 \wedge i = 1) \vee (b = 6 \wedge i = 1)\big)\big) \vee \big((a = 5) \wedge \big((b \in \{0, 1\}) \vee (b = 7 \wedge i = 1)\big)\big)$. Recall from Def. 7 that ψ_a includes the conjunct $\square\big((P_0 \wedge \square\diamond\neg\eta_{0,0}) \rightarrow \diamond Q_0\big)$. So, robot A has to eventually reach $[\![Q_0]\!] \subseteq [\![a = 4]\!]$ from any node in $[\![P_0]\!]$. Suppose that the system is at $(a = 0) \wedge (b = 3)$. In order for A to reach 4, B must come to 6, then A cross it, and continue to reach 4. Without any liveness assumptions, A cannot do so, because B can forever stay in 3. If B continues to sit at 3, it satisfies $\rho_{C,b}$. In order to force B to move with a safety constraint, we would have to remove transitions (namely, the self-loop of B at cell 3). This would happen when applying distributed synthesis, and it would restrict the movement of B, every time it visits 3. We do not want to do so.

The generated assumption $\square\diamond\eta_{0,0}$ makes it possible to force B to leave cell 3, and eventually come to cell 6. Robot A "enters" $[\![\eta_{0,0}]\!]$ by going to $a = 1$. It waits there, and the only way for B out of $[\![\eta_{0,0}]\!]$ is to eventually move to 6. Note that, at 6, B exits $[\![\eta_{0,0}]\!]$ as soon as A's turn comes, i.e., $(a = 1) \wedge (b = 6) \wedge (i = 1)$ is in $\eta_{0,0}$, but $(a = 1) \wedge (b = 6) \wedge (i = 0)$ is not. So,
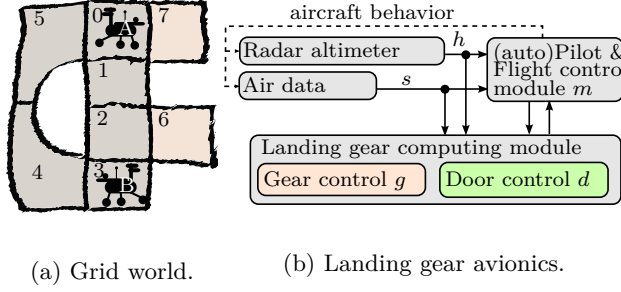
(a) Grid world.　　　(b) Landing gear avionics.

Fig. 4: Examples where liveness assumptions are constructed for cooperation of the modules.

A cannot force violation by B of the ψ_b guarantee $\Box\Diamond\neg\xi_{0,0}$ with $\xi_{0,0} = \eta_{0,0}$. B can keep staying at 6, until A decides to move to 2.

An observation of primary importance is that no finite behavior allowed by the given tuple φ was precluded in ψ. In other words, ψ is a contract over the cooperatively winning set C, and refines φ, in the sense of Sec. III. Thus, we have preserved infinitely many behaviors for the system. For example, B may be a solar-powered robot, with slow motors. At the time of contract creation, it doesn't matter what motors and batteries the designer will choose, nor what weather will be prevalent at the deployment site. As long as a robot B capable of realizing the property ψ_b is available, contract ψ can be realized by the system. This approach avoids to constrain the irrelevant details prematurely. Also, we didn't have to come up with assumptions manually, a task that can be challenging, error prone, and over-constraining. We specified in φ only the *primitive* goals that we are interested in.

2) Landing gear: We consider now a simplified modular avionics system for landing gear control [45], [46]. There are two modules, as shown in Fig. 4b, one representing the (auto)pilot and flight control module, another for landing gear control, which runs two independent control applications, one for control of the gear g (retracted: 0, moving: 1, fully extended: 2), another for control of the gear doors d (closed: 0, moving: 1, open: 2). The (auto)pilot controls airspeed $s \in \{0, \ldots, 1000\}$ km/h and absolute altitude $h \in \{0, \ldots, 100\}$ 100m, and the mode m (landing: 0, cruise: 1, takeoff: 2). The safety specification for all components requires that gear be fully extended below 300m $\Box((g \neq 2) \to (h > 3))$, doors be closed above 300 km/h $\Box((s > 300) \to (d = 0))$, doors be open if gear not retracted, $\Box((g \neq 0) \to (d = 2))$, gear be extended on ground, $\Box((h = 0) \to (g = 2))$, the flight control module cannot enter landing mode, unless the gear is fully extended $\Box((m = 0) \to (g = 2))$, and the gear must be retracted during cruise $\Box((m = 1) \to (g = 0))$. The (auto)pilot's goal is to

be able to enter all modes $\bigwedge_{k=0}^{2} \Box\Diamond(m = k)$.

For the goal $\Box\Diamond(m = 0)$ three levels of nesting are constructed. At the bottom one, the doors assume that the pilot will eventually reduce the speed below 300 km/h. If so, the doors open, leading to $Q_{d,0}$ for the doors, which is $P_{g,0}$ for the gear. When the doors are open, the landing gear can start moving, and eventually extend fully. This leads to $Q_{g,0}$, which is $P_{a,0}$, from where the flight control module can switch to landing mode. To enter cruise mode, $\Box\Diamond(m = 1)$, the nesting is different. At the bottom, the gear control assumes that eventually the height is above 300m, to retract. Then, the doors close, leading to the autopilot's attractor of cruise mode.

VI. RELATED WORK AND CONCLUSIONS

Synthesis has been applied widely to robotic and multiagent systems [47], [48], [49], [50]. The approach we proposed investigates construction of non-constraining properties, instead of synthesis. Most relevant to the work presented here are approaches for constructing GR(1) and LTL assumptions [26], [27], [28], [23], and for analyzing unrealizable specifications [51], [52], [53]. The syntactic methods in [26], [27] are enumerative in nature, in that a counter-strategy is iteratively synthesized, and consider only non-circular dependencies between modules. The semantic approach proposed here is symbolic (using BDDs), automated, and addresses circularity. The construction of safety assumptions before liveness is based on [23]. The assume-guarantee framework originates from [54], [55] that evolved to [56], [57], [58], [2]. Ideas about *temporal proof lattices* [31] and *temporal well-founded induction* [30] led to Alg. 1. The nonexistence of memoryless GR(1) contracts is related to the need for auxiliary variables in constructing refinement maps [59].

REFERENCES

[1] B. Meyer, "Applying "design by contract"," *Computer*, vol. 25, no. 10, pp. 40–51, 1992.
[2] A. Benveniste, B. Caillaud, D. Nickovic, R. Passerone, J. Raclet, P. Reinkemeier, A. Sangiovanni-Vincentelli, W. Damm, T. Henzinger, and K. Larsen, "Contracts for systems design," INRIA, Tech. Rep. 8147, 2012.
[3] A. Pnueli and R. Rosner, "Distributed reactive systems are hard to synthesize," in *FOCS*, vol. 2, 1990, pp. 746–757.
[4] S. Schewe and B. Finkbeiner, "Synthesis of asynchronous systems," in *LOPSTR*, 2007, pp. 127–142.
[5] B. Finkbeiner and S. Schewe, "Uniform distributed synthesis," in *LICS*, 2005, pp. 321–330.
[6] K. Chatterjee, T. Henzinger, J. Otop, and A. Pavlogiannis, "Distributed synthesis for LTL fragments," in *FMCAD*, 2013, pp. 18–25.

[7] B. Finkbeiner and S. Schewe, "Bounded synthesis," *STTT*, vol. 15, no. 5–6, pp. 519–539, 2013.

[8] A. Pnueli, "The temporal logic of programs," in *FOCS*, 1977, pp. 46–57.

[9] Z. Manna and A. Pnueli, "The anchored version of the temporal framework," in *Linear time, branching time and partial order in Logics and models for concurrency*, ser. LNCS. Springer, 1989, vol. 354, pp. 201–284.

[10] Z. Manna, *Mathematical theory of computation.* McGraw-Hill, 1974.

[11] W. Thomas and H. Lescow, "Logical specifications of infinite computations," in *A decade of concurrency reflections and perspectives*, 1993, pp. 583–621.

[12] B. Alpern and F. B. Schneider, "Defining liveness," *IPL*, vol. 21, no. 4, pp. 181–185, 1985.

[13] Z. Manna and A. Pnueli, "A hierarchy of temporal properties," in *PODC*, 1990, pp. 377–410.

[14] D. Kozen, "Results on the propositional μ-calculus," *TCS*, vol. 27, no. 3, pp. 333–354, 1983.

[15] K. Schneider, *Verification of reactive systems: formal methods and algorithms.* Springer, 2004.

[16] R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Sa'ar, "Synthesis of reacive(1) designs," *JCSS*, vol. 78, no. 3, pp. 911–938, 2012.

[17] R. Alur, T. Henzinger, and O. Kupferman, "Alternating-time temporal logic," in *JACM*, 2002, pp. 672–713.

[18] L. Lamport, "The temporal logic of actions," *TOPLAS*, vol. 16, no. 3, pp. 872–923, 1994.

[19] N. Piterman, A. Pnueli, and Y. Sa'ar, "Synthesis of reactive(1) designs," in *VMCAI*, 2006, pp. 364–380.

[20] M. Abadi, L. Lamport, and P. Wolper, "Realizable and unrealizable specifications of reactive systems," in *ICALP*, 1989, pp. 1–17.

[21] O. Kupferman, Y. Lustig, M. Vardi, and M. Yannakakis, "Temporal synthesis for bounded systems and environments," in *STACS*, 2011, pp. 615–626.

[22] M. Abadi and L. Lamport, "Composing specifications," *TOPLAS*, vol. 15, no. 1, pp. 73–132, 1993.

[23] K. Chatterjee, T. Henzinger, and B. Jobstmann, "Environment assumptions for synthesis," in *CONCUR*, 2008, pp. 147–161.

[24] I. Filippidis and R. M. Murray, "Symbolic construction of GR(1) contracts for synchronous systems with full information," Caltech, Tech. Rep. arXiv:1508.02705, 2015.

[25] L. Lamport, "How to write a 21st century proof," *Journal of fixed point theory and applications*, vol. 11, no. 1, pp. 43–63, 2012.

[26] R. Alur, S. Moarref, and U. Topcu, "Counter-strategy guided refinement of GR(1) temporal logic specifications," in *FMCAD*, 2013, pp. 26–33.

[27] ——, "Pattern-based refinement of assume-guarantee specifications in reactive synthesis," in *TACAS*, 2015, pp. 501–516.

[28] W. Li, L. Dworkin, and S. A. Seshia, "Mining assumptions for synthesis," in *MEMOCODE*, 2011, pp. 43–50.

[29] N. D'Ippolito, V. Braberman, N. Piterman, and S. Uchitel, "Synthesizing nonanomalous event-based controllers for liveness goals," *ACM TSEM*, vol. 22, no. 1, pp. 9:1–9:36, 2013.

[30] A. Pnueli, "In transition from global to modular temporal reasoning about programs," in *Logics and models of concurrent systems*, 1985, pp. 123–144.

[31] S. Owicki and L. Lamport, "Proving liveness properties of concurrent programs," *TOPLAS*, vol. 4, no. 3, pp. 455–495, 1982.

[32] K. L. McMillan, "Circular compositional reasoning about liveness," in *CHARME*, 1999, pp. 342–346.

[33] P. T. M. Varghese, "Parity and generalized Büchi automata," Ph.D. dissertation, Univ. of Liverpool, 2014.

[34] W. Thomas, "On the synthesis of strategies in infinite games," in *STACS*, 1995, pp. 1–13.

[35] R. McNaughton, "Infinite games played on finite graphs," *Ann. Pure & Appl. Logic*, vol. 65, no. 2, pp. 149–184, 1993.

[36] W. Thomas, "Solution of Church's problem: A tutorial," *New Perspectives on Games and interaction*, vol. 5, 2008.

[37] E. A. Emerson and C.-L. Lei, "Model checking in the propositional μ-calculus," University of Texas at Austin, Tech. Rep. 6, 1986.

[38] R. Alur, S. L. Torre, and P. Madhusudan, "Modular strategies for recursive game graphs," *TCS*, vol. 354, no. 2, pp. 230–249, 2006.

[39] A. Banerjee and P. Dasgupta, "The open family of temporal logics: Annotating temporal operators with input constraints," *TODAES*, vol. 10, no. 3, pp. 492–522, 2005.

[40] "Contract construction implementation," 2015. [Online]. Available: https://github.com/johnyf/contract_maker

[41] "omega: Symbolic algorithms for solving games of infinite duration (PYTHON package)." [Online]. Available: https://github.com/johnyf/omega

[42] I. Filippidis, R. M. Murray, and G. J. Holzmann, "A multi-paradigm language for reactive synthesis," in *SYNT*, 2015.

[43] "dd: Decision diagrams (PYTHON package)." [Online]. Available: https://github.com/johnyf/dd

[44] F. Somenzi, "CUDD: CU Decision Diagram package - v2.5.0," *Colorado University at Boulder*, 2012.

[45] I. Moir, A. Seabridge, and M. Jukes, *Civil avionics systems*, 2nd ed. Wiley, 2013.

[46] F. Boniol, V. Wiels, Y. Ait Ameur, and K.-D. Schewe, Eds., *ABZ 2014: The landing gear case study*, 2014.

[47] M. Kloetzer, S. Itani, S. Birch, and C. Belta, "On the need for communication in distributed implementations of LTL motion specifications," in *ICRA*, 2010, pp. 4451–4456.

[48] S. Karaman and E. Frazzoli, "Linear temporal logic vehicle routing with applications to multi-UAV mission planning," *IJRNC*, vol. 21, no. 12, pp. 1372–1395, 2011.

[49] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon temporal logic planning," *TAC*, vol. 57, no. 11, pp. 2817–2830, 2012.

[50] N. Özay, U. Topcu, and R. M. Murray, "Distributed power allocation for vehicle management systems," in *CDC*, 2011, pp. 4841–4848.

[51] R. Könighofer, G. Hofferek, and R. Bloem, "Debugging formal specifications: A practical approach using model-based diagnosis and counterstrategies," *STTT*, vol. 15, no. 5–6, pp. 563–583, 2013.

[52] V. Raman and H. Kress-Gazit, "Explaining impossible high-level robot behaviors," *TRO*, vol. 29, no. 1, pp. 94–104, 2013.

[53] G. E. Fainekos, "Revising temporal logic specifications for motion planning," in *ICRA*, 2011.

[54] R. W. Floyd, "Assigning meanings to programs," in *Symposia in Applied Mathematics*, vol. 19, 1967, pp. 19–32.

[55] C. Hoare, "An axiomatic basis for computer programming," *CACM*, vol. 12, no. 10, pp. 576–580, 1969.

[56] L. Lamport, "The "Hoare logic" of concurrent programs," *Acta Informatica*, vol. 14, pp. 21–37, 1980.

[57] N. Francez and A. Pnueli, "A proof method for cyclic programs," *Acta Informatica*, vol. 9, pp. 133–157, 1978.

[58] J. Misra and K. Chandy, "Proofs of networks of processes," *TSE*, vol. 7, no. 4, pp. 417–426, 1981.

[59] M. Abadi and L. Lamport, "The existence of refinement mappings," *TCS*, vol. 82, no. 2, pp. 253–284, 1991.