

Specification and Synthesis of Reactive Protocols for Aircraft Electric Power Distribution

Huan Xu, Ufuk Topcu, and Richard M. Murray, *Fellow, IEEE*

Abstract—The increasing complexity of electric power systems leads to challenges in integration and verification. We consider the problem of designing a control protocol for an aircraft electric power system that meets a set of requirements describing the correct behaviors of the system and reacts dynamically to changes in internal system states. We formalize the requirements by translating them into a temporal logic specification language and apply game-based, temporal logic formal methods to automatically synthesize a controller protocol that satisfies these overall properties and requirements. Through a case study, we perform a design exploration to show the benefits and tradeoffs between centralized and distributed control architectures.

Index Terms—Aerospace and electronic systems, control system synthesis, control systems.

I. INTRODUCTION

THE MOVE from conventional to more-electric aircraft architectures within the aerospace industry has been motivated by advancements in electronics technology that can result in aircraft with improved reliability, optimized system-level performance, and decreased life-cycle costs [1]. As more subsystems become reliant on electric power, the ability of the electric power system to function properly is critical in flight operations. However, this increased reliance also increases the complexity of the electric power system, requiring more components, actuators, and more complex architecture. Presently, standard methodologies for designing such systems are costly both in terms of time and money. Previous work has focused on the analysis of aircraft performance and power optimization by using modeling libraries and simulations on proposed designs [2], [3]. Analysis of all faults or errant behaviors in models, however, is difficult due to the high complexity of systems and subsystem interactions.

Manuscript received May 12, 2014; revised August 5, 2014; accepted October 28, 2014. Date of publication February 6, 2015; date of current version June 16, 2015. This work was supported in part by the Multiscale Systems Center (MuSyC), in part by the Boeing Corporation, and in part by the AFOSR AwardFA9550-12-1-0302. Recommended by Associate Editor S. Azuma.

H. Xu is with the Institute for Systems Research and Aerospace Engineering, University of Maryland, College Park, MD 20740 USA (e-mail: mumu@umd.edu).

U. Topcu is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: utopcu@seas.upenn.edu).

R. M. Murray is with Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: murray@cds.caltech.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCNS.2015.2401174

Since the process of verifying the correctness of a system with respect to specifications is expensive, this difficulty has led to a greater emphasis on the use of formal methods within a model-based systems engineering framework to aid in safety and performance certification. The work described in [4] provides a methodology for the overall design of an electric power system, beginning from requirements to real-time simulation. In this paper, we examine the control logic portion of the methodology and “specify and synthesize” a solution to the design problem instead of “design then verify.” In other words, instead of piecing together legacy designs that are then simulated and verified, we can synthesize control logic that is provably correct with respect to system requirements.

Building on previous work [5], we apply formal synthesis of control protocols that enable dynamic reconfiguration of power in more-electric aircraft that react to uncontrolled moves from an environment (or adversary) [6]. We begin by writing English-based specifications in linear temporal logic (LTL), and then use a combination of tools from computer science formal method domains for the automatic synthesis of control protocols. The use of LTL in designing controllers has been of great interest, and [7] provides a survey of recent contributions. The use of synthesis methods follows from their successful integration in verification of hardware and software systems in computer science, engineering, and robotics domains [8]–[12]. Applications of synthesis tools are limited to small problems due to the state space explosion issue. To address this challenge, we utilize previous work on the compositional design of correct-by-construction, distributed protocols for an electric power system [13], [14]. Distribution of the design and implementation of the electric power system will reduce the computational complexity and allow for the design of flexible control architectures in terms of modularity, fault-tolerance, and integrability [15].

The main contributions of this paper are twofold: 1) We investigate how specifications typically imposed in the design of controllers for electric power distribution can be translated into temporal logic. This formalization allows automation of synthesis of reactive control protocols that ensure the correctness of the underlying specifications. We provide an end-to-end demonstration of such a “specify and synthesize” design flow on an example. 2) As a step toward addressing industrial-scale problems, we examine the application of the compositional synthesis approaches, which we developed in our prior work, for the design of distributed control protocols. We show the utility of this compositional method on a distributed control architecture. The long-term potential for incorporating the “specify and synthesize” will be 1) faster design times, as problems can be

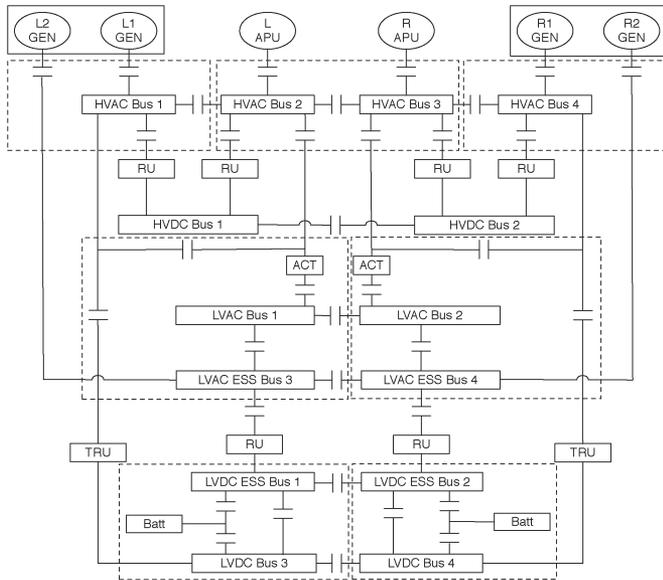


Fig. 1. Single-line diagram of an electric power system adapted from a Honeywell, Inc. patent [17]. Two high-voltage generators, two APU-mounted generators, and two low-voltage generators serve as power sources for the aircraft. Power can be routed from sources to buses through contactors, rectifier units, and transformers. Buses are connected to subsystem loads. Batteries can be used to provide emergency power.

found earlier in the design cycle, and 2) less costly systems, due to the ability to save time and perform design space exploration in software rather than in hardware.

The remainder of this paper is structured as follows: We describe a standard electric power system, including components, connectivity, and typical design considerations in Section II. Sections III and IV present a technical description of the linear temporal logic and the specifications for an electric power system. Section V details the synthesis procedure. A case study is presented in Section VI with results for a centralized and distributed control architecture.

II. ELECTRIC POWER DISTRIBUTION SYSTEM

A. System Components

A standard electric power system comprises generators, batteries, rectifier units, and batteries that supply power to a set of loads through buses. Fig. 1 is a single-line diagram of an electric power system. The following text is a brief description of components referenced throughout the book [16].

High- and low-voltage ac (HVAC, LVAC) buses and high and low-voltage dc (HVDC, LVDC) buses deliver power to loads and power conversion equipment. These buses can be classified as essential (denoted by ESS) and nonessential. Essential buses supply loads that should always remain powered, such as the flight actuation subsystem, while nonessential buses have loads that may be shed, such as cabin lighting, in the case of a failure. Generators connected to the left and right engine (L1 GEN, L2 GEN, R1 GEN, R2 GEN), as well as generators connected to the left and right auxiliary power unit (L APU, R APU) can operate at either high voltages, which can connect to the high-voltage ac buses, or low voltages, which feed directly to

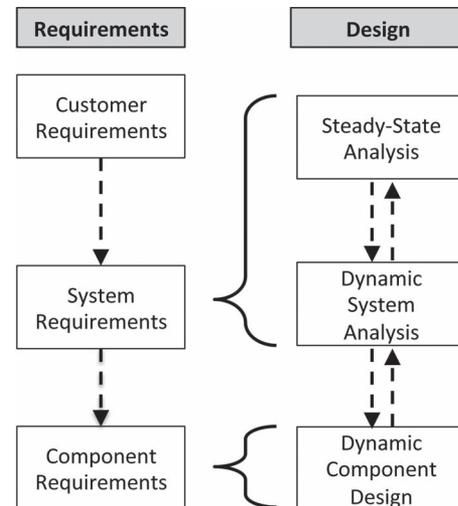


Fig. 2. Typical aerospace system design flow depicting requirements and design streams.

the low-voltage buses. Contactors are high-power electronic switches that connect the flow of power from sources to buses and loads. Contactors provide the actuation for reconfiguration of the topology of the electric power system, hence, changing the paths through which power is delivered. Rectifier units (RUs) convert three-phase ac power to dc power. Transformers (ACT) convert ac power to dc power, and transformer rectifier units (TRUs) both convert and lower voltages.

B. Current Design Practice

Fig. 2 depicts the design flow for a typical aerospace system. In current practice, there are two concurrent work streams. In the requirements workstream (left column in Fig. 2), a set of English-based customer requirements is manually decomposed into system-level and component-level specifications. In the second workstream, steady-state analysis is first performed on the system. Modeling tools are used to select relevant properties or set points. The methods typically rely on reuse where possible and domain expertise to ensure the entire envelope is covered. Once this is done, a dynamic system analysis is performed on low-fidelity models to assess operational scenarios and validate control performance. Finally, detailed models are used to analyze at the component level. System and component requirements are mapped or linked to the models (from the left column to the right column). All three levels of abstraction are constructed by different groups using design reviews to ensure consistency. If an inconsistency is found, then the design gets iterated until the problem is corrected. Dashed arrows thus represent manual checks between groups to guarantee consistency.

The disadvantage to current practice is that exchanges between these design layers are manual and text based, thus requiring design reviews, multiple iterations, and verification of results. Problems that are found late in the design process are difficult to diagnose, and costly to change. An alternative to designing a system and then verifying its correctness is to “specify and synthesize.” Requirements are specified using a

mathematical language, and then a control law is automatically synthesized to be correct with respect to those requirements. The potential benefits of such an approach are twofold. First, synthesis is useful to automatically construct complex control logic that is guaranteed to satisfy requirements. Second, the use of formal specification languages to synthesize controllers allows for easier and systematic ways to identify improper designs, and allows designers to capture system requirements (including temporal requirements) in an unambiguous manner.

One of the major barriers to implementing synthesis and formal methods techniques has been the aversion to techniques and tools that are new or have a high learning curve. Past methods for design of systems have worked. Thus, the need for a more model-based approach was seen as unnecessary or risky. This mentality has been slowly changing as newer systems have proven to be much more complex to design and difficult to verify [18], [19], where delays in production have lasted years and have cost millions of dollars. The second barrier to synthesis is the issue of scalability. As systems increase in size and complexity, synthesis tools become time and memory intensive for industrial scale problems. Current synthesis tools are difficult to apply to industrial-scale problems as the full reactive synthesis problem is double exponential in size of the state space and specifications. Distributed synthesis is one way to alleviate this limitation as the size of the state space is reduced by solving smaller problems. A second way is to restrict the type of specifications used. Efficient algorithms and corresponding tools could be developed specifically tailored for this problem domain. Moreover, the electric power system topology is suitable for relatively straightforward decomposition of specifications. As seen in Fig. 1, the topology is divided into physical panels (as depicted by the dotted boxes). Synthesizing controllers for individual panels, while simultaneously leveraging the symmetry of the topology, could decrease the overall computational time. The approach described in this paper serves as a foundation on which to develop ways to fully address industrial-scale problems, which is the subject of ongoing work.

C. Design Considerations

The control protocol design problem considers how the system shall reconfigure as a function of the changes in flight conditions and faults in the components. Typically such reconfiguration takes place in multiple layers. Generation and primary distribution involves the startup or shut down of high-voltage (HV) generators or APUs in addition to the reconfiguration of contactors in order to route power to HV and LV buses and their respective loads.

The networked structure of the electric power system requires a protocol to account for all possible interleavings of faults or events throughout a flight. These combinations and sequences of failures may be rare but can be catastrophic. This structure likewise introduces and allows for redundancy in the electric power system in hardware and software. In the remainder of this paper, we focus on the dynamic reconfiguration of generation and primary distribution systems by designing a control protocol for contactors. Based on the status of gener-

ators and buses, the protocol ensures the proper switching of contactors to guarantee buses will remain powered.

III. FORMAL SPECIFICATION USING LINEAR TEMPORAL LOGIC

Given a topology of an electric power system similar to that in Fig. 1, the main design problem becomes determining all correct configurations of contactors for all flight conditions and faults that can occur in the system. For a configuration to be “correct” means that it satisfies system requirements, also referred to as specifications. We now discuss a formal specification language that will be utilized for the synthesis of control protocols later in this section.

In reactive systems (i.e., systems that react to a dynamic, *a priori* unknown environment), correctness will depend not only on inputs and outputs of a computation, but also on execution of the system. Temporal logic is a branch of logic that incorporates notions of temporal ordering to reason about the correctness of propositions over a sequence of states, and is well suited for problems in which the system must react to an adversary or environment. First used as a specification language by Pnueli [20] in the 1970s, it has been utilized to specify properties in a number of applications, including embedded systems, robotics, and controls [8]. In this paper, we consider a version called linear temporal logic (LTL).

LTL’s main building block is the atomic proposition, which is a statement on a valuation of variables that has a unique truth value (*True* or *False*). Consider, for example, the health status of generators g_1 and g_2 , where $\{g_1 = \text{healthy}\}$, and $\{g_2 = \text{unhealthy}\}$ are atomic propositions. The truth values of each proposition can be determined for a given configuration of the electric power system. LTL combines logical connectives such as negation (\neg), disjunction (\vee), conjunction (\wedge), and material implication (\rightarrow) alongside temporal operators such as *always* (\square), *eventually* (\diamond), *next* (\bigcirc), and *until* (\mathcal{U}) to create complex specifications for a system. These logical connectives and temporal operators can be combined to specify a number of complex requirements.

Safety formulas assert that a state or sequence of states will not be reached. In particular, we use a subclass of safety formulas referred to as invariants throughout this paper. Invariant formula assert that a property will remain true throughout the entire execution, and ensure that nothing bad will happen. A safety specification for the electric power system could take the form $\square(\neg \text{bus}_i \text{unpowered})$ where i is the bus index.

Progress formulas guarantee that a property holds infinitely often in an execution. This property ensures that the system will make progress. For example, always eventually ensure that Bus 1 is powered and can be written as $\square \diamond \text{bus}_1 \text{powered}$.

Response formulas state that at some point in the execution following a state where a property is true, there exists a point where a second property is true. Response properties describe how systems need to react to changes in operating conditions. For example, if a generator fails, then at some point a corresponding contactor should open: $\square((\text{gen}_j \text{not_healthy}) \rightarrow \diamond(\text{contactor}_k \text{open}))$ where j, k represent indices for generators and contactors, respectively.

LTL is naturally used for specifying properties with temporal ordering. For specifications involving hard-timing constraints, LTL can be used for synchronous systems where all processes (i.e., components) proceed in a lock-step manner. The next operator has a “time” measure so that, for a given property φ , $\bigcirc\varphi$ signifies at the next time instant, φ is true. To specify a property occurring at some point in the future, multiple next operators can be used, such that $\bigcirc^k\varphi \triangleq \bigcirc\bigcirc\dots\bigcirc\varphi$ asserts that property φ holds after k time instants. Alternatively, the “timed” specifications in the electric power system uses a clock variable to define an equivalent property. A more detailed discussion of LTL and other temporal logics can be found in [21] and [22].

IV. FORMAL SPECIFICATIONS FOR AIRCRAFT ELECTRIC POWER SYSTEMS

The following list details the temporal logic specifications for synthesizing control protocols for electric power systems.

Environment Assumption: Let \mathcal{G} and \mathcal{R} represent the set of all generators and rectifier units in a given topology. Let Boolean variables g and r denote the health status of generator $G \in \mathcal{G}$ and rectifier $R \in \mathcal{R}$, respectively. The lowercase symbol represents the health status of the component denoted by the corresponding uppercase symbol. We use a similar convention between upper and lowercase symbols in the remainder of this paper. In order for the system to be nontrivial, we assume that at least one generator and one rectifier unit must always be healthy, that is, have a status of 1, at any given time. This environment assumption is written as

$$\square \left\{ \bigvee_{G \in \mathcal{G}} (g = 1) \wedge \bigvee_{R \in \mathcal{R}} (r = 1) \right\}. \quad (1)$$

Contactor Delays: To capture the physical delay in contactor actuation times, we introduce the Boolean variable \tilde{c} , denoting controller intent (command) for contactor status c . The delay between intent and contactor status is handled by additional clock variables x_C^{close} (closing time) and x_C^{open} (opening time) for each contactor C , where each “tick” of the clock represents δ time. If the intent is to open and a contactor status is closed, then the contactor opens within time bound $[T_{o_{\min}}, T_{o_{\max}}]$. Similarly, if the intent is to close and a contactor status is open, the contactor closes within time bound $[T_{c_{\min}}, T_{c_{\max}}]$. If the contactor status does not match the intent, at the next step clock, x_C increments by δ . The controlled variable is the contactor intent, while the status is modeled as an environment variable. The following specifications are thus included in the environment assumption.

If the status and intent match, then in the next step, the clock resets to zero, which is written as

$$\square \{ (\bigcirc c = \tilde{c}) \rightarrow (\bigcirc x_C = 0) \}. \quad (2)$$

If the intent is the same as the contactor status, then contactor status remains the same in the next step, i.e.,

$$\square \{ (\tilde{c} = c) \rightarrow (\bigcirc c = c) \}. \quad (3)$$

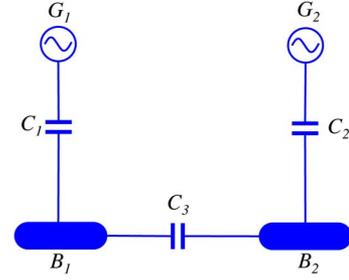


Fig. 3. Single-line diagram with two generators, two buses, and three contactors. Paralleling of ac sources can occur if the statuses of all three contactors C_1 , C_2 , and C_3 are all closed.

The assumption to capture the delay in closing time between a contactor intent and status is given by

$$\square \{ (\tilde{c} = 1 \wedge c = 0 \wedge (x_C^{\text{close}} < T_{c_{\min}})) \rightarrow (\bigcirc c = 0 \wedge \bigcirc x_C^{\text{close}} = x_C^{\text{close}} + \delta) \}, \quad (4)$$

$$\square \{ (\tilde{c} = 1 \wedge c = 0 \wedge (x_C^{\text{close}} \geq T_{c_{\min}})) \rightarrow (\bigcirc c = 1 \vee \bigcirc x_C^{\text{close}} = x_C^{\text{close}} + \delta) \}. \quad (5)$$

Specifications to capture the delay in opening between a contactor intent and status are given by

$$\square \{ (\tilde{c} = 0 \wedge c = 1 \wedge (x_C^{\text{open}} < T_{o_{\min}})) \rightarrow (\bigcirc c = 1 \wedge \bigcirc x_C^{\text{open}} = x_C^{\text{open}} + \delta) \}, \quad (6)$$

$$\square \{ (\tilde{c} = 0 \wedge c = 1 \wedge (x_C^{\text{open}} \geq T_{o_{\min}})) \rightarrow (\bigcirc c = 0 \vee \bigcirc x_C^{\text{open}} = x_C^{\text{open}} + \delta) \}. \quad (7)$$

The clock values for x_C^{close} and x_C^{open} are known *a priori* for each contactor. Finally, we guarantee that clocks x_C^{close} and x_C^{open} never exceed the opening and closing time bounds, written as

$$\square (x_C^{\text{close}} \leq T_{c_{\max}}), \quad (8)$$

$$\square (x_C^{\text{open}} \leq T_{o_{\max}}). \quad (9)$$

Unhealthy Sources: We require any contactor adjoining a generator or rectifier unit to open when that component becomes unhealthy. Let \mathcal{C} represent the set of all contactors in the electric power system. Sets $\mathcal{N}(G_i) \subseteq \mathcal{C}$ and $\mathcal{N}(R_i) \subseteq \mathcal{C}$ represent the contactors directly connected to, or neighboring, generator G_i and rectifier R_i , respectively. In Fig. 3, for example, sets $\mathcal{N}(G_1)$ and $\mathcal{N}(G_2)$ consist of contactors C_1 and C_2 , respectively. For a contactor C , let c be its status (for example, 0 represents an open contactor and 1 is a closed contactor). If a generator’s status is unhealthy, then contactors connecting to it should be commanded open, that is, the contactor intent variable should take the value of 0. This can be written as

$$\bigwedge_{G \in \mathcal{G}} \square \left\{ (g = 0) \rightarrow \bigwedge_{C \in \mathcal{N}(G)} (\tilde{c} = 0) \right\}. \quad (10)$$

Similar specifications hold for disconnecting rectifier units.

No Paralleling of AC Sources: A mismatch in ac generator frequencies and voltages can lead to a loss of availability and even damage to the distribution system. To avoid difficulties of generator synchronization, we disallow any paralleling of ac sources, that is, no bus should be powered by multiple ac generators at the same time. In the example shown in Fig. 3, paralleling could occur if the status of contactors C_1 , C_2 , and C_3 were all closed at the same time. A specification would then be to never allow all contactors along a path to close at the same time if that path could connect two ac sources.

Let X_{G_i, G_j} represent the set of components along a path between generators $G_i, G_j \in \mathcal{G}$ and $i \neq j$. We disallow configurations where all contactors $C \in X_{G_i, G_j}$ create a *live path* (i.e., all contactors are closed along a path). These specifications are written as

$$\square \bigwedge_{G_i, G_j \in \mathcal{G}} \left\{ \neg \bigwedge_{C \in X_{G_i, G_j}} (c = 1) \right\}. \quad (11)$$

Power Status of Buses: An ac bus is powered if there exists a live path that connects the bus to a healthy generator. A dc bus is powered if there exists a live path between the dc bus and a healthy rectifier unit, which itself is connected to a powered ac bus. Let $X_{G, B}$ denote the set of all components (contactors, buses, and rectifier units) along a path between bus B and generator G , excluding B and G . Bus B is powered if there exists a live path between B and G , written as

$$\square \left\{ \bigvee_{G \in \mathcal{G}} \left((g = 1) \wedge \bigwedge_{X \in X_{G, B}} (x = 1) \right) \rightarrow (b = 1) \right\}. \quad (12)$$

If no live path exists, then B will be unpowered, written as

$$\square \left\{ \neg \bigvee_{G \in \mathcal{G}} \left((g = 1) \wedge \bigwedge_{X \in X_{G, B}} (x = 1) \right) \rightarrow (b = 0) \right\}. \quad (13)$$

Essential Buses: Let \mathcal{B}_s be the set of all safety-critical buses that must remain powered. Denote the allowable length of time a bus can remain unpowered as T . Typical values for T fall in the 50 ms range [16]. Time in this formulation is implemented through an additional clock θ_B associated with bus B , where each increment represents δ time. For each safety-critical bus in $B \in \mathcal{B}_s$, these specifications can be written as follows.

- If the status of bus B is unpowered, in the next step, the clock variable θ_B will increment by one unit

$$\square \{(b = 0) \rightarrow (\bigcirc \theta_B = \theta_B + \delta)\}. \quad (14)$$

- If the status of bus B is powered, in the next step, the clock variable θ_B is reset to 0

$$\square \{(b = 1) \rightarrow (\bigcirc \theta_B = 0)\}. \quad (15)$$

- The clock variable θ_B will never be greater than the maximum-allowable unpowered time T . This is implemented by

$$\square \{\theta_B \leq T\}. \quad (16)$$

Remark 1: Specifications for dc components are the same as ac specifications except for two simplifications: 1) The nonparalleling of ac sources specification may be ignored and 2) no dc bus may ever be unpowered (i.e., unpowered time $T = 0$). In addition, rectifier units are unidirectional and, thus, power cannot flow from dc buses to ac buses. For simplicity of the synthesis problem, contactors and buses are assumed to behave ideally (i.e., they cannot fail). In addition, we assume that sensors (not depicted in the topology) are able to perfectly measure all variables.

V. SYNTHESIS OF REACTIVE CONTROL PROTOCOLS

The correctness of the system is not merely a function of the states of the controlled variables. It needs to be interpreted in conjunction with the statuses of the externalities that interact with the system that cannot be controlled. Furthermore, it is necessary to incorporate information on potential environment conditions under which the system is expected to operate. If the environment variables are not properly constrained, then the resulting control protocol may be overly conservative, and it may not be possible to construct a protocol that ensures the satisfaction of the system requirements. If all generators simultaneously stay unhealthy for a long enough time, then it is not possible to satisfy the condition that the essential buses shall not be unpowered longer than some prespecified period. Consequently, the overall goal is to design a protocol that determines how controlled variables shall move at each point of the execution as a function of the behaviors of the controlled and environment variables so far in the execution as long as the environment assumptions are satisfied.

A. Reactive Synthesis

Equipped with LTL as a specification language, we now formally state the reactive synthesis problem. Let the system S be comprised of E and P , which are sets of environment and controlled variables, respectively. Let state $s = (e, p) \in \text{dom}(E) \times \text{dom}(P)$ be a valuation of the system. Consider an LTL specification φ of assume-guarantee form

$$\varphi = \varphi_e \rightarrow \varphi_s \quad (17)$$

where φ_e is the conjunction of LTL specifications that characterize the assumptions on the environment and φ_s is the conjunction of LTL specifications that characterize the system requirements. The synthesis problem is then concerned with constructing a strategy, that is, a partial function $f : (s_0 s_1 \dots s_{t-1}, e_t) \mapsto p_t$ that chooses the move of the controlled variables based on the state sequence so far and the behavior of the environment so that the system satisfies φ_s as long as the environment satisfies φ_e . The synthesis problem can be viewed as a two-player game between the environment and controlled plant: the environment attempts to falsify the specification in (17) and the controlled plant tries to satisfy it.

For general LTL, the synthesis problem has a doubly exponential complexity in the size of the specification [8]. For a subset of LTL called generalized reactivity (1) (GR(1)),

Piterman *et al.*, have shown that synthesis can be solved in polynomial time (polynomial in the number of valuations of the variables in E and P) [6]. GR(1) specifications restrict φ_e and φ_s to take the following form, for $\alpha \in \{e, s\}$

$$\varphi_\alpha := \varphi_{\text{init}}^\alpha \wedge \bigwedge_{i \in I_1^\alpha} \square \varphi_{1,i}^\alpha \wedge \bigwedge_{i \in I_2^\alpha} \square \diamond \varphi_{2,i}^\alpha$$

where $\varphi_{\text{init}}^\alpha$ is a propositional formula characterizing the initial conditions; $\varphi_{1,i}^\alpha$ are transition relations characterizing safe, allowable moves and propositional formulas characterizing invariants; and $\varphi_{2,i}^\alpha$ are propositional formulas characterizing states that should be attained infinitely often. Many interesting temporal logic specifications can be expressed or easily transformed into GR(1) specifications. See [6] and [23] for a more precise treatment of GR(1) synthesis, and [5], [6], [23], and [24] for case studies in hardware synthesis, motion planning for autonomous vehicles, and vehicle-management systems.

Given a GR(1) specification, the digital design synthesis tool implemented in JTLV (a framework for developing temporal verification algorithm) [25] generates finite automaton that represents a switching strategy for the system. The temporal logic planning (TuLiP) toolbox, a collection of python-based code for automatic synthesis of correct-by-construction embedded control software provides an interface to JTLV [26]. For examples discussed in this paper, we use TuLiP.

B. Distributed Synthesis

The control protocols discussed in Section V-A are centralized in that the controller has access to measurements of all controlled and environment variables, and is able to determine the evolution of all controlled variables in order to satisfy a set of specifications. We now detail a few reasons for migrating to distributed control architectures.

Hardware challenges: A centralized controller onboard an aircraft requires wiring from a central processing unit to all components. Local controllers allow for shorter wires and increased efficiency due to this reduction in weight. *Increased resilience to failure:* By distributing the implementation of the controller, the electric power system can be more robust to failures, that is, if one portion of the electric power system malfunctions, the other sections are unaffected and can still be fully operational. *Reduction of computational complexity:* With an increased number of electric components, the combination of configurations the controller must account for quickly becomes intractable for current verification and synthesis tools as well as testing. A distributed controller design correctly decomposes the design task into smaller subproblems each of which may be easier to cope with.

Advantages from the distribution of the control design come with increased importance of reasoning about the interfaces between the controlled subsystems. There is relatively extensive literature on compositional reasoning [27]–[29]. Here, we follow the exposition in recent work from [13]. Fig. 4 illustrates the decomposition of global specifications into local specifications. For ease of presentation, consider the case where system SYS is decomposed into two subsystems SYS_1 and

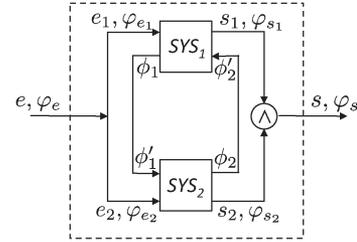


Fig. 4. Schematic for the decomposition of global specifications into distributed controllers for two subsystems. The overall environment assumptions φ_e and system guarantees φ_s are distributed into the two subsystems SYS_1 and SYS_2 . Each subsystem has its own local environment assumptions and system guarantees. In addition, SYS_1 has an extra set of local guarantees ϕ_1 that interact with SYS_2 as environment assumptions ϕ_1' , while SYS_2 guarantees contained in ϕ_2 act as environment assumptions ϕ_2' for SYS_1 .

SYS_2 . For $i = 1, 2$, let E_i and P_i be the environment variables and controlled variables for SYS_i such that $P_1 \cup P_2 = P$ and $P_1 \cap P_2 = \emptyset$. Let φ_{e1} and φ_{e2} be LTL formulas containing variables in E_1 and E_2 , respectively. Similarly, let φ_{s1} and φ_{s2} be LTL formulas in terms of $E_1 \cup P_1$ and $E_2 \cup P_2$, respectively. The following conditions are assumed to be true for a decomposition:

- 1) any execution (i.e., sequence of actions) from the environment that satisfies φ_e also satisfies $(\varphi_{e1} \wedge \varphi_{e2})$;
- 2) any execution of the system that satisfies $(\varphi_{s1} \wedge \varphi_{s2})$ also satisfies φ_s ;
- 3) there exist two control protocols that realize the local specifications $(\varphi_{e1} \rightarrow \varphi_{s1})$ and $(\varphi_{e2} \rightarrow \varphi_{s2})$.

Then, by a result in [13], implementing these control protocols leads to a system where the global specification is met.

It is possible that even if the centralized problem is realizable, the local distributed synthesis may be unrealizable. Subsystems may need to interact with each other through shared variables (either information or physical values) in order to become realizable. As seen in Fig. 4, subsystem SYS_1 provides additional guarantees ϕ_1 to subsystem SYS_2 , evaluated as an environment assumption and denoted as ϕ_1' . The same interaction applies to the interface between SYS_2 , which sends its own local guarantees ϕ_2 to SYS_1 . Additional interface refinements may be added to local specifications in order to ensure realizability of local controllers. Thus, if the following local specifications (and interface refinements) hold:

$$\phi_2' \wedge \varphi_{e1} \rightarrow \varphi_{s1} \wedge \phi_1, \quad (18)$$

$$\phi_1' \wedge \varphi_{e2} \rightarrow \varphi_{s2} \wedge \phi_2 \quad (19)$$

then the global specification $\varphi_e \rightarrow \varphi_s$ is realizable.

While the initial decomposition of specifications is performed in an ad-hoc manner, more systematic methods can be used to refine specifications in order to synthesize controllers for each individual subsystem. Recent work from [30] uses a counter-strategy guided refinements method. The algorithm synthesizes a set of candidate assumptions from a counter-strategy that, if added to the original decomposition, would restrict the environment's behaviors from acting according to the counter strategy. More specifically, given a specification

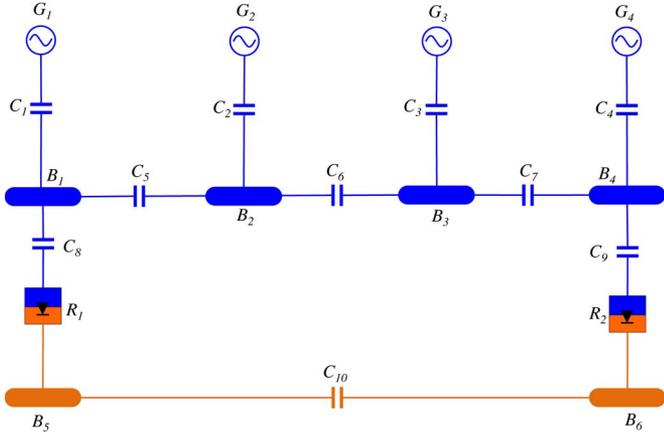


Fig. 5. Simplified single-line diagram. Four generators connect to four ac buses. Two rectifier units convert power from ac to dc, and connect to two dc buses.

$\varphi = \varphi_e \rightarrow \varphi_s$ that is unrealizable, the problem is to find a refinement $\psi = \bigwedge_i \psi_i$ consisting of environment assumptions ψ_i to which $\varphi_e \wedge \psi \rightarrow \varphi_s$ is realizable.

VI. CASE STUDY

We examine a simplified single-line diagram as shown in Fig. 5. For all allowable failures, the synthesis problem is to design a control logic that will reconfigure contactors so that power will be delivered to buses.

A. Variables

Environment Variables: The health statuses for generators G_1, G_2, G_3, G_4 , and rectifier units R_1 and R_2 can each take values of healthy (1) or unhealthy (0). Again, we distinguish component variables and status variables by upper and lower cases, for example, the first generator is represented by G_1 , while its health status is denoted by g_1 . To implement delays between a contactor intent and status, we define contactor statuses c_1, c_2, c_3 , and c_4 with values of open (0) and closed (1). For simplicity, all other contactors are considered to have no delays (i.e., at each time step, the physical state of the contactor will always match the intent of the contactor).

Controlled Variables: Contactor intent $\tilde{c}_1, \tilde{c}_2, \tilde{c}_3$, and \tilde{c}_4 for contactors connecting generators to buses can each take values of intend to open (0) or intend to closed (1). All other contactors are directly controlled and can either be 0, denoting an open contactor, or 1, signifying a closed contactor.

Dependent Variables: The statuses of buses B_1, B_2, B_3, B_4, B_5 , and B_6 can be either powered (1) or unpowered (0) depending on the status of neighboring contactors, rectifier units, and generators.

B. Specifications

Given the topology in Fig. 5, the specifications described in Section IV reduce to the following specifications used in the synthesis problem for the simplified single-line diagram.

Environment Assumption: From (1), the assumption that at least one generator and rectifier unit is always healthy becomes

$$\Box\{(g_1 = 1) \vee (g_2 = 1) \vee (g_3 = 1) \vee (g_4 = 1)\}, \quad (20)$$

$$\Box\{(r_1 = 1) \vee (r_2 = 1)\}. \quad (21)$$

Because of delays between contactor intents and statuses, we prevent generators and rectifier statuses from rapidly switching between healthy and unhealthy. Without such a restriction, the intent of the contactor could constantly change before the status would have time to react, thus rendering the synthesis problem unrealizable. A simple assumption is that once a generator becomes unhealthy, it will remain unhealthy. This is equivalent to assuming that even if the generator were to become healthy again, the contactor would not allow the generator to provide power to the rest of the system, which is true in practice as well, because in this case study, the contactors connected to rectifier units have no delay, and there is no need to similarly restrict their environment behavior. Therefore, the additional specification is

$$\Box \bigwedge_{i=1}^4 \{(g_i = 0) \rightarrow (\bigcirc g_i = 0)\}. \quad (22)$$

Contactor Delays: We set each clock tick such that $\delta = 1$, and consider the delay times to open and close all contactors to be the same, with minimum and maximum delay times to be 0 and 2. Thus, $T_{c_{\min}} = T_{o_{\min}} = 0$ and $T_{c_{\max}} = T_{o_{\max}} = 2$, and (2)–(9) are used to specify contactor C_i delays, with $i \in [1, 4]$.

Unhealthy Sources: In Fig. 5, the set of neighboring contactors to generators are $\mathcal{N}(G_1) = C_1$, $\mathcal{N}(G_2) = C_2$, $\mathcal{N}(G_3) = C_3$, and $\mathcal{N}(G_4) = C_3$. Neighboring contactors to rectifier units are $\mathcal{N}(R_1) = C_8$ and $\mathcal{N}(R_2) = C_9$. If, for example, the status of rectifier R_1 becomes unhealthy, its neighboring contactor status intent should be set to open (0). This is written as

$$\Box\{(r_1 = 0) \rightarrow (\tilde{c}_8 = 0)\}. \quad (23)$$

Similar specifications hold for all neighboring contactors to generators and rectifier units.

No Paralleling of AC Sources: For the single-line diagram in Fig. 5, there are six ac generator pairs $\{G_1, G_2\}$, $\{G_1, G_3\}$, $\{G_1, G_4\}$, $\{G_2, G_3\}$, $\{G_2, G_4\}$, and $\{G_3, G_4\}$. For example, a live path for the pair $\{G_1, G_3\}$ exists if contactors C_1, C_5, C_6 , and C_3 are all closed. To avoid instances of paralleling ac sources, the specification to disallow the live path would be

$$\Box \neg\{c_1 = 1 \wedge c_5 = 1 \wedge c_6 = 1 \wedge c_3 = 1\}. \quad (24)$$

Similar specifications to avoid nonparalleling are implemented for all ac generator pairs.

Power Status of Buses: Consider ac bus B_1 in Fig. 5. Bus B_1 will be powered if a live path exists between itself and

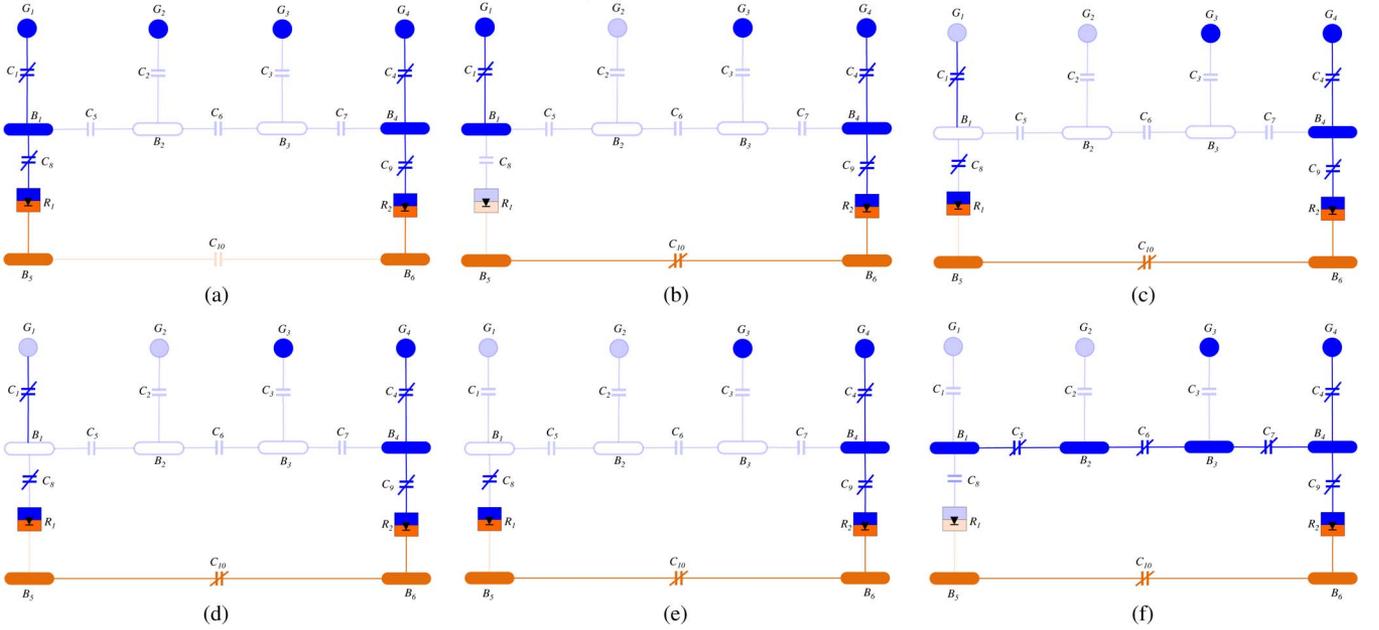


Fig. 6. Single simulation trace for a centralized controller for the electric power system single-line diagram case study. Each subfigure represents a single time step. Shaded generators and rectifier units are healthy, shaded contactors are closed, and shaded buses are powered. (a) $t = 0$. (b) $t = 1$. (c) $t = 2$. (d) $t = 3$. (e) $t = 4$. (f) $t = 5$.

a generator. This requirement is captured by the following equation:

$$\square\{(c_1 = 1 \wedge g_1 = 1) \rightarrow b_1 = 1\}, \quad (25)$$

$$\square\{(c_5 = 1 \wedge b_2 = 1 \wedge c_2 = 1 \wedge g_2 = 1) \rightarrow (b_1 = 1)\}, \quad (26)$$

$$\square\{(c_5 = 1 \wedge b_2 = 1 \wedge c_6 = 1 \wedge b_3 = 1 \wedge c_3 = 1 \wedge g_3 = 1) \rightarrow (b_1 = 1)\}, \quad (27)$$

$$\square\{(c_5 = 1 \wedge b_2 = 1 \wedge c_6 = 1 \wedge b_3 = 1 \wedge c_7 = 1 \wedge b_4 = 1 \wedge c_4 = 1 \wedge g_4 = 1) \rightarrow (b_1 = 1)\}. \quad (28)$$

If none of the four conditions to the left side of the implication are true, then the status of bus B_1 is unpowered. Similar specifications hold for ac buses B_2 , B_3 , and B_4 .

DC bus B_5 will be powered if it is connected to a healthy rectifier unit and a live path exists between itself and a powered ac bus. This requirement is captured by the following equation:

$$\square\{(r_1 = 1 \wedge c_8 = 1 \wedge b_1 = 1) \rightarrow (b_5 = 1)\}, \quad (29)$$

$$\square\{(c_{10} = 1 \wedge b_6 = 1 \wedge r_2 = 1 \wedge c_9 = 1 \wedge b_4 = 1) \rightarrow (b_5 = 1)\}. \quad (30)$$

If neither of the above two conditions are true, then the status of bus B_5 is unpowered. Similar specifications hold for B_6 .

Essential Buses: In this problem, we consider buses B_1 and B_4 to be connected to safety-critical loads, and can be unpowered for no longer than three time steps. Each increment of the clock variable θ_{B_1} and θ_{B_4} represents one time step $\delta = 1$. A safety specification for B_1 is of the following form:

- If bus status b_1 is unpowered, then at the next time step, clock θ_{B_1} increments by one, such that

$$\square\{(b_1 = 0) \rightarrow (\bigcirc\theta_{B_1} = \theta_{B_1} + 1)\}. \quad (31)$$

- If bus status b_1 is powered, then at the next time step, reset clock θ_{B_1} to zero. This is written as

$$\square\{(b_1 = 1) \rightarrow (\bigcirc\theta_{B_1} = 0)\}. \quad (32)$$

- To ensure that the status of B_1 is never unpowered for more than three steps, (16) becomes

$$\square\{\theta_{B_1} \leq 3\}. \quad (33)$$

The requirement that all dc buses must always remain powered can be implemented by setting $T = 0$ or, more compactly, written as

$$\square\{b_5 = 1 \wedge b_6 = 1\}. \quad (34)$$

C. Results

Consider a system model with environment variables E that includes generators $G_1 - G_4$ and rectifier units R_1, R_2 , and control variables P consisting of contactors $C_1 - C_{10}$ and buses $B_1 - B_6$. Given environment assumption φ_e from (2)–(9), (20)–(22), and φ_s as the conjunction of all specifications from (23)–(33), the synthesis problem is finding a control protocol such that (17) holds. The output of the synthesis procedure is a discrete planner represented as a finite-state automation where states are valuations of environment and control variables.

1) *Centralized Controller Design:* We now present the results for the centralized case of the electric power system design problem with variables and specifications discussed in the previous section. Fig. 6 shows the simplified single-line diagram used in the problem formulation over a sequence of time. Each subfigure represents a single step of the simulation, starting at step 0 and ending with step 5. For clarity of exposition, the

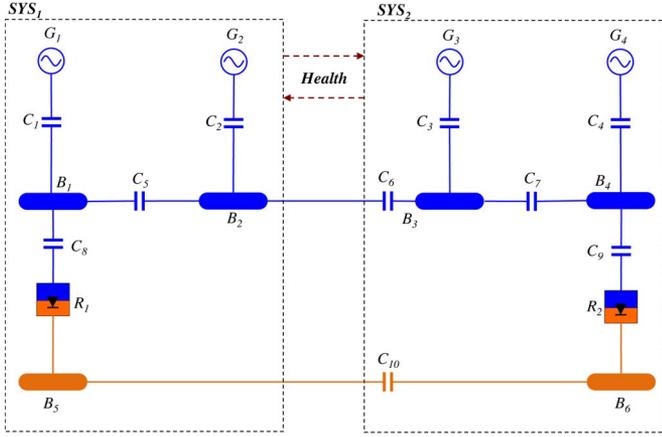


Fig. 7. Distributed controller decomposition. Components enclosed within the dashed rectangles are controlled by their own respective controllers. The dashed arrow represents information flow, in the form of a health status variable, directed from SYS₁ to SYS₂. The solid arrow represents the physical transfer of power from SYS₂ to SYS₁.

results presented only implement a contactor delay for C_1 , and all other contactors have instant actuation.

At each step, depicted in Fig. 6(a)-(f), the statuses of environment variables can switch between healthy and unhealthy, subject to the assumption that at least one generator and one rectifier always remains healthy. Furthermore, because of the contactor delay implemented for C_1 , once the status for generator G_1 becomes unhealthy, that status will remain unhealthy. A healthy generator or rectifier unit is denoted by a shaded component. In addition, a shaded bus is powered, and a shaded contactor is closed.

Fig. 6(a) shows the initial configuration at time $t = 0$. All four generator statuses are healthy, as are both rectifier units. Contactor statuses c_1, c_4, c_8 , and c_9 are closed, and bus statuses b_1, b_4, b_5 , and b_6 are powered. At time $t = 1$ [Fig. 6(b)], the status for generator G_2 and rectifier R_1 becomes unhealthy. As a result, contactor status c_8 switches to open, c_{10} closes, and bus B_5 is powered through B_6 . At time $t = 2$, generator status g_1 becomes unhealthy. Because of the contactor delay time for C_1 , status c_1 does not open until $t = 4$. For safety-critical buses B_1 and B_4 , their statuses are never unpowered for more than three time steps throughout the entire simulation sequence. This specification is not imposed on the middle two buses, however, and, thus, b_3 can remain unpowered for five steps without violating any system requirements. Bus status b_1 remains unpowered for three time steps, and then becomes powered again at time $t = 5$.

The synthesis process produces a control protocol in the form of a finite state automation. The resulting automaton for the electric power system centralized controller takes 258 s to solve on a 2.9-GHz Intel Core Duo processor with 8-GB memory, and has 2049 states.

2) *Distributed Control Architecture*: The physical decomposition of the electric power system for a distributed control architecture is shown in Fig. 7. Let SYS₁ represent subsystem on the left, and SYS₂ represent the subsystem on the right.

For the decomposition in Fig. 7, SYS₂ can control the supply of power via contactors C_6 and C_{10} . SYS₁ can only receive

power when SYS₂ provides it. We decompose the global environment assumption, where at least one power source must remain healthy at each step, such that

$$\varphi_{e_2} = \square((g_3 = 1 \vee g_4 = 1) \wedge r_2 = 1); \varphi_{e_1} = \square(\text{true}).$$

The specification for φ_{e_1} states that there are no restrictions on the behavior of φ_{e_1} . The assumption placed on φ_{e_2} ensures that for any allowable sequence of environment actions, the controller for SYS₂ is able to supply power to SYS₁ at any step. Health status information for G_1 and G_2 is sent to SYS₂ via a health status variable h_1 . The variable is set to 0 if neither source is healthy, and 1 if either g_1 or g_2 is healthy. In addition, the status of rectifier R_1 from SYS₁ is also sent to SYS₂ as an information variable.

The specifications for this distributed system lead to a synthesis problem that is unrealizable. Using the counter-strategy guided refinement method discussed in Section V-B, we find that the original decomposition was insufficient for SYS₁ to guarantee power to its buses since there were no restrictions on generators G_1 and G_2 . Furthermore, while generators G_3 and G_4 are able to generate power at all steps, the controller for SYS₂ must also be able to *guarantee* that power can be delivered to SYS₁. One candidate solution is to impose that at least one of the generator statuses will be healthy, $\square g_1 = 1 \vee g_2 = 1$. Focusing on interface refinements, however, a better candidate solution is to impose an environment restriction on contactors C_6 and C_{10} to guarantee that power from buses B_3 and B_6 can be supplied to SYS₁ if health status h_1 is unhealthy.

We introduce ϕ_2 as a guarantee for controller SYS₂, and denote ϕ'_2 as an assumption for controller SYS₁. Since one subsystem controls the flow of power, a single-sided refinement is sufficient for the design problem to be realizable, and we can set $\phi_1 = \text{true}$. The refinement specifications are of the following form: bus status b_3 is never unpowered for a prespecified period of time T . In other words, B_3 becomes an essential bus, and we introduce a variable θ_{B_3} that is used as a counter such that

$$\square\{(b_3 = 0) \rightarrow (\bigcirc\theta_{B_3} = \theta_{B_3} + 1)\} \wedge \square\{(b_3 = 1) \rightarrow (\bigcirc\theta_{B_3} = 0)\} \wedge \square\{\theta_{B_3} \leq T\}.$$

Moreover, because of the requirement that dc buses must always be powered, then maximum unpowered time $T = 0$ in order to guarantee that power can always be provided to SYS₁.

If health status $h_1 = 0$, that is, g_1 and g_2 are unhealthy, then, whenever b_3 is powered, \tilde{c}_6 will be set to close

$$\square\{((h_1 = 0) \wedge (b_3 = 1)) \rightarrow (\tilde{c}_6 = 1)\}.$$

Similarly, if health status $r_1 = 0$, then \tilde{c}_{10} will be set to close. Counter-strategy guided methods are not fully automated to decompose specifications. It does, however, provide a more systematic way to “debug” and refine interfaces. Methods to fully automate decomposition are currently a work in progress.

3) *Scalability*: For $T = 0$, that is, when essential buses must always be powered, the resulting distributed controllers for each subsystem take 5 s to synthesize and result in automata of approximately 120 states each. For the full topology depicted in Fig. 5, the computation time increases to 10 s. Adding an

additional row of components (i.e., generator, rectifier, ac bus, ac contactor, dc bus, and dc contactor) to the topology increases the computation time to 1000 s. Even for the simple example shown in the case study, decreasing the number of components in the synthesis problem decreases the computation time by an order of magnitude. Given existing synthesis tools, distributed synthesis techniques, such as these demonstrated, are one way to address industrial-scale problems.

VII. CONCLUSION AND FUTURE WORK

The use of a specification language to capture requirements and a synthesis procedure to construct control protocols partially automates the design process, allowing for faster design times and more efficient ways to identify errors. This paper demonstrates how text-based specifications can be captured using linear temporal logic with timing constraints in a representative case study. Given a topology for an electric power system and a set of system requirements, we automatically synthesize a control protocol for an electric power system on an aircraft. The controller reacts to changes in the environment and is guaranteed, by construction, to satisfy the desired properties even in the presence of generator failures.

The computational complexity of synthesis makes solving industrial-scale problems difficult for current tools. We thus examined the use of distributed control protocols, which take less computational time to synthesize due to fewer components within each subsystems and, thus, smaller state spaces. They are, however, more conservative than a centralized controller in terms of length of time when nonessential buses are powered.

From the basis of this, there are a number of directions for practical and theoretical future work that apply to electric power systems as well as other application areas that span other networked control systems. While counter-strategy guided refinements partially automate decomposition of specifications, future work will focus on fully automating the decomposition of overall system specifications into subsystem specifications. Timing specifications, (e.g., safety and contactor open/closing times) in the electric power system problem are addressed with the use of clocks by way of an additional counter variable. This discretization of time further adds to the difficulties arising from the state space explosion. We are currently examining the use of timed verification and synthesis tools, such as UPPAAL and UPPAAL-TIGA [31]. One open issue not addressed is the level of abstraction needed for modeling, design, and specifications of an electric power system. Control of the power quality from generators is considered at a continuous level of abstraction. Load management and load shedding are considered at a discrete low-level of abstraction. Both of these problems, although at different levels of abstraction, should be interfaced with the primary distribution problem discussed in this paper.

ACKNOWLEDGMENT

The authors would like to thank N. Ozay from the University of Michigan, and R. Poisson and Eelco Scholte from United Technologies Aerospace Systems for their insight and helpful discussions.

REFERENCES

- [1] J. Rosero, J. Ortega, E. Aldabas, and L. Romeral, "Moving towards a more electric aircraft," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 22, no. 3, pp. 3–9, Mar. 2007.
- [2] J. Zumbege *et al.*, "Integrated aircraft electrical power system modeling and simulation analysis," *SAE Tech. Paper*, pp. 01–1804, 2010.
- [3] T. Wu, S. Bozhko, G. Asher, and D. Thomas, "Fast functional modelling of the aircraft power system including line fault scenarios," in *Proc. 5th IET Int. Conf. Power Electron., Mach. Drives*, 2010, pp. 1–7.
- [4] P. Nuzzo *et al.*, "A contract-based methodology for aircraft electric power system design," *IEEE Access*, vol. 2, pp. 1–25.
- [5] T. Wongpiromsarn, U. Topcu, and R. Murray, "Formal synthesis of embedded control software: Application to vehicle management systems," in *Proc. AIAA Infotech Aerospace Conf.*, St. Louis, MO, USA, 2011.
- [6] N. Piterman, A. Pnueli, and Y. Sa'ar, "Synthesis of reactive (1) designs," in *Verification, Model Checking, Abstract Interpretation*. New York, USA: Springer, 2006, pp. 364–380.
- [7] C. Belta *et al.*, "Symbolic planning and control of robot motion [grand challenges of roboticspeyl1a]," *IEEE Robot. Autom. Mag.*, vol. 14, no. 1, pp. 61–70, Mar. 2007.
- [8] A. Pnueli, "Applications of temporal logic to the specification and verification of reactive systems: A survey of current trends," *Current Trends Concurrency*, pp. 510–584, 1986.
- [9] A. Galton, *Temporal Logics and their Applications*. London, U.K.: Academic Press, 1987.
- [10] G. J. Holzmann, "The model checker spin," *IEEE Trans. Softw. Eng.*, vol. 23, no. 5, pp. 279–295, May 1997.
- [11] M. Kloetzer and C. Belta, "A fully automated framework for control of linear systems from temporal logic specifications," *IEEE Trans. Autom. Control*, vol. 53, no. 1, pp. 287–297, Feb. 2008.
- [12] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Trans. Robot.*, vol. 25, no. 6, pp. 1370–1381, Dec. 2009.
- [13] N. Ozay, U. Topcu, and R. M. Murray, "Distributed power allocation for vehicle management systems," in *Proc. IEEE Conf. Dec. Control Eur. Control Conf.*, 2011, pp. 4841–4848.
- [14] N. Ozay, U. Topcu, R. M. Murray, and T. Wongpiromsarn, "Distributed synthesis of control protocols for smart camera networks," in *Proc. IEEE/ACM Int. Conf. Cyber-Phys. Syst.*, 2011, pp. 45–54, IEEE.
- [15] T. Laengle, T. C. Lueth, and U. Rembold, "A distributed control architecture for autonomous robot systems," *Ser. Mach. Perception Artif. Intell.*, vol. 21, pp. 384–402, 1995.
- [16] I. Moir and A. Seabridge, *Aircraft Systems: Mechanical, Electrical, Avionics Subsystems Integration*. West Sussex, U.K.: AIAA Education Series, 2001.
- [17] R. G. Michalko, "Electrical starting, generation, conversion and distribution system architecture for a more electric vehicle," U.S., Patent 7 439 634, 10, 2008.
- [18] A. Shalal-Esa, "Pentagon says F-35 fighter delayed, costs rise 4.3 percent," *Chicago Tribune*, Mar. 2012. [Online]. Available: http://articles.chicagotribune.com/2012-03-29/news/sns-rt-us-lockheed-fighterbre82t03r-20120329_1_f-35-costs-rise-pentagon-report
- [19] S. Creedy, "Dreamliner's slow flight," *The Australian*, Oct. 2013. [Online]. Available: <http://www.theaustralian.com.au/news/features/dreamliners-slow-flight/story-e6frg6z6-1226733805717>
- [20] A. Pnueli, "The temporal logic of programs," in *Proc. IEEE 18th Annu. Symp. Foundations Comput. Sci.*, 1977, pp. 46–57.
- [21] Z. Manna and A. Pnueli, *The Temporal Logic of Reactive and Concurrent Systems: Specifications*, vol. 1. New York, USA: Springer, 1992.
- [22] C. Baier and J. Katoen, *Principles of Model Checking*. Cambridge, MA, USA: MIT Press, 2008.
- [23] R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Saar, "Synthesis of reactive (1) designs," *J. Comput. Syst. Sci.*, vol. 78, no. 3, pp. 911–938, 2012.
- [24] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon temporal logic planning," *IEEE Trans. Autom. Control*, vol. 57, no. 11, pp. 2817–2830, Nov. 2012.
- [25] A. Pnueli, Y. Sa'ar, and L. Zuck, "Jtlv: A framework for developing verification algorithms," in *Computer Aided Verification*. New York, USA: Springer, 2010, pp. 171–174.
- [26] T. Wongpiromsarn, U. Topcu, N. Ozay, H. Xu, and R. M. Murray, "Tulip: A software toolbox for receding horizon temporal logic planning," in *Proc. 14th Int. Conf. Hybrid Syst.: Comput. Control*, 2011, pp. 313–314.
- [27] M. Mukund, "From global specifications to distributed implementations," *Synthesis Control Discrete Event Syst.*, pp. 19–34, 2002.
- [28] E. Filiot, N. Jin, and J.-F. Raskin, "Antichains and compositional algorithms for ltl synthesis," *Formal Meth. Syst. Design*, vol. 39, no. 3, pp. 261–296, 2011.

- [29] P. Madhusudan and P. Thiagarajan, "Distributed controller synthesis for local specifications," *Automata Languages Program.*, pp. 396–407, 2001.
- [30] R. Alur, S. Moarref, and U. Topcu, "Counter-strategy guided refinement of gr (1) temporal logic specifications," in *Proc. IEEE Formal Meth. Comput.-Aided Design*, 2013, pp. 26–33.
- [31] G. Behrmann *et al.*, "Uppaal-tiga: Time for playing games!," in *Computer Aided Verification*. New York, USA: Springer, 2007, pp. 121–125.



Huan Xu received the B.S. degree in mechanical engineering and material science from Harvard University, Cambridge, MA, USA, in 2007, and the M.S. and Ph.D. degrees in mechanical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2008 and 2013, respectively.

Currently, she is an Assistant Professor at the University of Maryland, College Park, MD, USA, with a joint appointment in Aerospace Engineering and the Institute for Systems Research. Current projects include developing real-time flight simulation software

for unmanned aircraft vehicles (UAVs) and the design and implementation of a UAV for the Student Unmanned Aircraft Systems Competition. Her research interests are in control, analysis, and design of unmanned autonomous systems.



Ufuk Topcu received the Ph.D. degree in mechanical engineering from the University of California, Berkeley, CA, USA.

Currently, he is a Research Assistant Professor in the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia, PA, USA, and was a Postdoctoral Scholar at the California Institute of Technology, Pasadena, CA, USA, until 2012. His research is on the analysis, design, and verification of autonomous, networked systems.



Richard M. Murray (F'04) received the B.S. degree in electrical engineering from California Institute of Technology (Caltech), Pasadena, CA, USA, in 1985 and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 1988 and 1991, respectively.

Currently, he is the Thomas E. and Doris Everhart Professor of Control and Dynamical Systems and Bioengineering at Caltech. His research is in the application of feedback and control to networked systems, with applications in biology and autonomy. Current projects include the analysis and design biomolecular feedback circuits; specification, design and synthesis of networked control systems; and novel architectures for control using slow computing.