Synthesis and Validation of Control Software For A

Vehicular Electric Power Distribution Testbed

Robert Rogersten KTH Royal Institute of Technology, Stockholm, Sweden

Huan Xu University of Maryland, College Park, Maryland, USA

Necmiye Ozay University of Michigan, Ann Arbor, Michigan, USA

Ufuk Topcu University of Pennsylvania, Philadelphia, Pennsylvania, USA

Richard M. Murray California Institute of Technology, Pasadena, California, USA

Modern aircraft increasing rely on electric power, resulting in high safety-criticality and complexity in their electric power generation and distribution systems. Motivated by the resulting rapid increase in the costs and duration of the design cycles for such systems, we investigate the use of formal specification and automated, correct-byconstruction control protocols synthesis for primary distribution in vehicular electric power networks. We discuss a design workflow that aims to transition from the traditional "design+verify" approach to a "specify+synthesize" approach. We give an overview of a subset of the recent advances in the synthesis of reactive control protocols. We apply these techniques in the context of reconfiguration of the networks in reaction to the changes in their operating environment. We also validate these automatically synthesized control protocols on high-fidelity simulation models and on an academic-scale hardware testbed.

I. Introduction

Next generation aircraft are moving away from hydraulically and pneumatically powered systems into electrically powered systems [1]. As dependence increases on electric power, however, the electric power generation and distribution systems become more critical to the safe operation of aircraft [2]. Because of this increased reliance on electric power, these systems on next-generation aircraft need to be highly reliable and fault-tolerant. In current practice the design of an aircraft electric power system is constructed in an ad hoc manner, and is either borrowed from legacy designs or created "by hand" using designer experience and knowledge. The entire process from running simulations to software testing (i.e., testing control logic) and hardware testing is both time consuming and costly, as mistakes are found throughout. Moreover, unexpected system failures at the hardware level require returning to the design phase to make changes. As the design stage progresses, the more expensive changes must be in re-design.

The difficulty in design of large-scale, complex systems partly lies in the lack of a formal structure to verify the correctness of a system. We propose a formal design methodology for an electric power system that integrates the use of formal methods [3, 4] in order to guarantee correctness. The overall design flow is shown in Fig. 1. The first step in the methodology is translation of specifications. System requirements, including safety and performance properties and customer requests, are typically given in English, text-based form. In order to apply formal methods to establish the correctness of a system, specifications must be translated into a formal specification language (e.g., linear temporal logic is used in this paper) [5, 6], that is mathematically-based and unambiguous. While the details of the translation are not covered within the scope of his paper, this process is critical to overall system design.

Once abstracted and specified formally, we then proceed to the control synthesis layer in the methodology. In this step, we take the abstract model and formal requirements and automatically synthesize a control protocol. The Python-based Temporal Logic Planning toolbox, TuLiP [7], is used to construct a controller that is guaranteed to be correct with respect to the system requirements. If no such controller exists, then specifications or the model can be modified. Instead of constructing a system by hand and then verifying its correctness (i.e., "design and verify"), we "spec-

ify and synthesize" a control protocol. TuLiP has been used to synthesize controllers in past work on aircraft electric power systems [8] and vehicle management systems [9, 10].

Once a control protocol is synthesized, the next step in the methodology is the simulation and hardware tests layer. From the abstract model, a simulation model can be constructed in a tool such as Simulink [11]. Here, (relatively) high-fidelity simulations can be performed, i.e., the behavior of the system can be tested by injecting faults or failures. Because specifications may arise from legacy designs or other customers, in this step of the methodology we can adjust the types of components used in the model. For example, different batteries may have different voltage ratings that may not be able to satisfy all requirements given. Thus, the simulation layer provides information on how good the abstract model and specifications are, and whether or not those need to be modified. These models can also be used for "testing" of design artifacts (e.g., controllers that are synthesized using the abstract models).

Previous work has discussed the implementation of software on hardware [12]. The hardware test step introduces the physical aspects of the system into the design stage. Thus, controllers can be tested for their correctness on a physical system. If any undesired behaviors arise that would not necessarily violate specifications, but are not considered reasonable or "optimal", the hardware implementation provides information back to the specification and abstract model level.

Because of the growing complexity of electric systems, in particular, and embedded systems in general, the use of ad hoc techniques for design is becoming more difficult and time-consuming. The advantages of a formal methodology, such as the one demonstrated in this paper, is the ability for systematic exploration of the design space, as well as the ability to formally analyze and guarantee correctness. In this paper, we demonstrate this design flow for an electric power system and its academic-scale hardware implementation. This demonstration serves only the purpose of proof-ofconcept. Extending the tools used in our study, modifying them to align with the needs of particular application areas, and transitioning them are among the important challenges that are faced. They are, however, beyond the scope of the current paper.



Fig. 1 Methodology of design flow for an aircraft electric power system. An abstract model and formal requirements are used to synthesize a controller. Then control protocols can be tested in high-fidelity simulations and implemented on a hardware platform.

II. Background

An electric power system provides power to buses and subsystem loads. In more-electric aircraft, these loads include lighting, heating, and safety- and mission-critical subsytems (such as avionics, de-icing, and flight actuation) [1]. Fig. 2 shows a single-line diagram for an electric power system [13]. Each of the two engines power a high-voltage and low-voltage AC generator. Two additional generators are mounted on an auxiliary power unit and can be used to supply power in case of emergencies. The primary loads can be considered as safety- or mission-critical. Primary loads include avionics, communication systems, and window heating. The electric power system of an aircraft often contains transformers and rectifier units that broadly divide the system into four categories, namely high-voltage AC, high-voltage DC, low-voltage AC, and low-voltage DC. The power is distributed from the generators to the loads in series connection through buses, transformers, rectifier units and electronically controlled switches called contactors.



Fig. 2 A single-line diagram for an electric power system topology adapted from a Honeywell Patent for a more-electric aircraft [13] by Richard Poisson of United Technologies Aerospace Systems.

There are three levels of design challenges in an electric power system: the primary distribution problem, the secondary distribution problem, and the load-shedding problem. In primary distribution, the main concern is in providing power from generators to buses. Generators must be able to supply power to buses connected to safety-critical loads. In the secondary distribution problem, the design question is how much power should be allocated to system loads by buses. Finally, the load-shedding problem is, if a power failure or emergency situation were to arise, what loads should be shed, and in what order, so that safety-critical loads can still be powered and the aircraft can land safely.

In this paper we consider the primary distribution problem. The overall goal is to design a controller that can react to component failures by changing the topology so that new ways for power delivery are created. Moreover, the controller must ensure that that safety-critical buses and loads are always powered. The state of the system (i.e., the status of all contactors and health of components) is estimated by current and voltage sensor measurements. If a fault is detected the controller reacts and reconfigures the contactors so that the system still satisfies all requirements. This controller can take actions depending on the system and environmental conditions during operation. The control logic not only accounts for a static configuration of contactors given a fault, but also determines the correct sequence of contactor switches in order to guarantee all specifications are still satisfied.

Typical electric power system specifications are categorized in terms of safety, reliability, and performance. On aircraft with variable-frequency generators, a mismatch in frequency and voltage can lead to, for example, overcurrents and fires. A safety specification, therefore, would be to disallow any configuration of components in which more than one generator provides power to a bus. A typical reliability specification requires that the system must be able to account for a certain number or subset of failures. The total number of allowable, simultaneous failures is known as a reliability level. Every component has a probability of failure, determined from past operational data. Assuming independence of failures on components, the maximum number of components that can fail (i.e., the environmental conditions for which the controller must account) is determined the reliability specification [14]. Finally, consider performance specifications that effect the overall quality of the flight. A standard performance specification would disallow buses connected to critical loads to be unpowered for a length of time greater than some pre-determined time bound. This ensures proper operation of loads necessary for safe flight. While this paper examines a limited subset of electric power system specifications, more general specifications have been used [15, 16].

III. Modeling, Specifications, and Synthesis

At the core of the design methodology we advocate in this paper are models and specifications in mathematically based languages and the corresponding algorithms that automate the synthesis of software-based control protocols from these models and specifications. We now give an overview of these building blocks as tailored to the discussion in the subsequent sections.

A. Modeling and Specifications

The initial step in any model-based flow is determining the level of fidelity to be used in the design of the control protocols. Partly for aligning with the industrial practice and partly for leveraging the currently available synthesis tools, we use purely discrete (and finite) models for the evolution of the configuration of a power distribution network, for example, as shown in Fig. 2.

Roughly speaking, the variables of interest can be grouped as those under direct control (we will call these as "control" variables) and those that can change without the control of the system (we will call those as "environment" variables). As a modeling convention, we will consider that the environment variables evolve adversarially (specific meaning of "adversarial" will be concretized later in this section) against the system. Typically, controlled variables are the statuses of the contactors. They open and close with directives from the controller. Examples of environment variables include the health statuses of the generators and rectifier units which typically take binary values (i.e., healthy vs unhealthy).

Let now x be the set of variables (including both controlled and environment). Then, the evolution of the system can be described by sequences of valuations x_t (we will also call these valuations as the "states" of the system) of x at the time steps t = 0, 1, 2, ... Let M denote all sequences $x_0x_1x_2...$ that can be generated by the system. M can be considered as a model of the system. Note that, besides this abstract representation of the model, we can equivalently use finite-state, nondeterministic transition system in order to represent all possible behaviors of the system [17, 18].

We will characterize the correctness of the system in terms of the properties satisfied by the sequences in M. To formalize this notion, we use temporal logic based languages [5, 17]. Roughly,

temporal logic allows to unambiguously specify and reason about infinite sequences of states. We specifically employ linear temporal logic (LTL) to describe system behavior. An LTL formula is built up from a set of atomic propositions and two kinds of operators: logical connectives and temporal modal operators. An atomic proposition is a statement over the system variables that has a unique truth value (*True* or *False*) for a given valuation of x. For example, let g and c denote the health status of a particular generator and the status of a particular contactor, respectively. Then, given a configuration of the system, the truth value of "g = healthy", "c = open", and "g = health and c = closed" can be determined and all these statements are atomic propositions. In other words, atomic propositions are the lowest level of building blocks for specifying the system behavior and logical connectives, including negation (\neg), disjunction (\lor), conjunction (\land), and implication (\rightarrow), and temporal operation, including always (\Box), eventually (\diamondsuit), until (\mathcal{U}), and next (\bigcirc), connect these building blocks to create more sophisticated specifications of the system. For example, given atomic propositions p and q, we can write

- invariance (a specific form of safety) properties as $\Box p$,

- guarantee or reachability properties $\Diamond p$,
- progress or recurrent properties as $\Box \diamondsuit p$,
- response properties as $\Box(p \implies \Diamond q)$, and
- next-step response properties as $\Box(p \implies \bigcirc \diamondsuit q)$.

In this paper, we use LTL as the specification language for convenience. Depending on the underlying model and properties of interest, one may consider other, potentially more suitable specification languages including timed temporal logic [19], probabilistic temporal logic [20], and branching-time logics [17]. For further details on the range of specification languages we refer the reader to [5, 17] and for details of and a complete treatment for primary distribution in vehicular electric power networks to our earlier work [18].

B. Synthesis of Reactive Control Protocols

The overall goal of the design problem is synthesizing a control protocol that, when implemented on the electric power system, ensures that the controlled system satisfies its specifications. The correctness of the system is though not merely a function of the controlled variables. It needs to be interpreted in conjunction with the environment variables. For example, the generator from which each bus shall be powered is constrained by the health statuses of the generators, which cannot be controlled by the system. Hence, the control protocol needs to react to the changes in both the controlled variables and environment variables.

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 and may not be possible to construct a protocol that ensures the satisfaction of the system requirements. For example, if all the 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. Hence, such behaviors of the environment shall be disregarded in the protocol design. An essential component of the protocol synthesis problem is the environment assumptions that specify what environment behaviors the controller shall correctly react to. Consequently, the overall goal is to design a protocol that determines how the 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.

We now, equipped with LTL as a specification language, formally state the reactive synthesis problem. Let E and P be sets of environment and controlled variables, respectively. Let $s = (e, p) \in$ $dom(E) \times dom(P)$ be a state of the system. Consider a LTL specification φ of assume-guarantee form

$$\varphi = \varphi_e \to \varphi_s,\tag{1}$$

where, roughly speaking, φ_e is the conjunction of LTL specifications that characterizes the assumptions on the environment and φ_s is the conjunction of LTL specifications that characterizes the system requirements. The synthesis problem is then concerned with constructing a strategy, i.e.,

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 the controlled plant: the environment attempts to falsify the specification in (1) and the controlled plant tries to satisfy it.

For general LTL, it is known that the synthesis problem has a doubly exponential complexity [21]. For a subset of LTL, namely generalized reactivity (1) (GR(1)), Piterman et al., have shown that it can be solved in polynomial time (polynomial in the number of valuations of the variables in *E* and *P*) [22]. GR(1) specifications restrict φ_e and φ_s to take the following form, for $\alpha \in \{e, s\}$,

$$\varphi_{\alpha} := \varphi_{\mathrm{init}}^{\alpha} \, \wedge \bigwedge_{i \in I_{1}^{\alpha}} \Box \varphi_{1,i}^{\alpha} \wedge \bigwedge_{i \in I_{2}^{\alpha}} \Box \Diamond \varphi_{2,i}^{\alpha},$$

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. Given a GR(1) specification, the digital design synthesis tool implemented in JTLV (a framework for developing temporal verification algorithm) [23] generates a 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 [7]. For the examples discussed in this paper, we use TuLiP.

C. A Closer Look at the Synthesized Controllers

Fig. 3 shows different views of the resulting controller automaton for a toy example. This controller has four states. The top left corner of Fig. 3 shows the output from TuLiP which roughly lists each of the states, the corresponding configurations (i.e., the status of the two generators and three contactors in this example) and the states to which the system may transition (as a function of the environment move) from the current one. The big box on the right hand side pictures the configurations of the corresponding network in each of the four states. For example, state 0 (read from the text in the top, left box) corresponds to a configuration where both generators are healthy,



Fig. 3 An automatically generated controller from TuLiP (top, left) and its translation into a Matlab function. The automaton is synthesized for a two-generator and three-contactor case. The generator status variables are rgen and lgen, and the contactor status variables are c1, c2, and c3. Each state has successors, which define to which state the controller can transition depending on the current controlled and environment state. In addition, no-successor states exists.

two of the contactors are closed, and the one connecting the two buses is open. If one of the generators (rgen) become unhealthy, then the system transitions to state 2, contactor c2 opens, and contractor c3 closes in reaction to this change in the generator health. Finally, the left bottom corner of Fig. 3 shows (part of) the control automaton written into a Matlab function which is used to drive the testbed as discussed in the next section.

IV. An Aircraft Electric Power Testbed

We now discuss an end-to-end implementation of the "specify+synthesize" design flow on an academic-scale electric power testbed we had developed in our recent work [12]. We begin with an overview of the testbed and its basic functionality. Fig. 4 shows the physical layout of the testbed (left) and its single-line diagram (right). It was build to mimic some of the characteristics of the primary electric power distribution systems on aircraft. It contains transformers that supply power to an AC systems and rectifier units that separate the DC part of the system from the AC part. We refer to these transformers as generators because we are only interested in the role as voltage sources.



Fig. 4 *Left*: A photo of the physical layout of the electric power system testbed. *Right*: Singleline diagram of the power system testbed. Contactors are represented by double bars. The AC and DC sides of the system are separated by rectifier units (RU).

The generators and rectifier units are crucial for the safe operation of an aircraft. Therefore, the testbed focuses on the failure of these components.

A. The Structure of the Hardware Testbed

The topology of the hardware testbed is shown in Fig. 4 (right); it contains two generator sources and two rectifier units. The generator sources are modeled by transformers with a secondary side voltage of 24 VAC. The DC section is connected to the AC section through two rectifier units. The rectifier units contain a diode-rectifier bridge and DC bus capacitor to achieve a low ripple in the DC-side voltage. It also contains a variable DC voltage regulator that is tuned to 2.5 VDC. Consequently, the testbed has two different voltage levels: 24 VAC and 2.5 VDC.

The single-line diagram also contains four buses, specifically, two AC buses and two DC buses. Multiple lamps attached to these buses are considered as the primary loads in the distribution network. The main design goal is to keep these loads powered even in the presence of failures in the generators or rectifier units. A detailed circuit schematic of the testbed hardware is shown in Fig. 5.

The contactors in electric power distribution networks on aircraft are designed to switch three-



Fig. 5 Circuit schematic of the testbed. The topology is the same as in Fig. 4 (right). The relays are represented by the numbered boxes. The numbered arrows denote voltage sensing connections from Fig. 8.

phase electric power. Their functionality are replicated by simpler relays in the testbed. In particular, a commercially available relay board [25] that provides a set of computer-controlled relays that can communicate with programming languages supporting serial communications for example Matlab, is used. The relays on the board are numbered and range from 0 to 7. The same numbering convention is used in Fig. 4 (right) and Fig. 5, and throughout the paper. We also remark that for the results presented in this paper, the contactors c1 and c2 shown in Fig. 5 were not used and left closed. However they can be used to isolate the AC and DC subsystems and to test the AC subsystem separately.

The hardware testbed is also equipped with switches/plugs for injecting faults and with sensors that monitor the health status of components. These elements are discussed in detail in Section IV D.

B. The Specifications

The testbed mimics only a small fraction of the functionality that exists in the primary distribution networks on aircraft. Therefore, only a subset of the typical specifications are well-defined for the testbed. We now discuss these specifications, both their descriptions in English and their translations into temporal logic statements.

As discussed in Section III, the formal specifications we consider have an assume-guarantee form, i.e., they contain both assumptions on the possible environment behavior and guarantees on the system behavior. The environment assumptions include the following.

- At least one of the generators is always healthy.
- At least one of the rectifier units is always healthy.

The guarantees on the system behavior include the following.

- No AC bus can be powered from two different AC sources simultaneous at any time.
- AC and DC buses are powered at all times.

The synthesis problem can then be states to constructing a reactive control logic that ensures the realizability of the temporal logic specification $\varphi_e \to \varphi_s$, where

$$\varphi_e = \Box(((gen_1 = healthy)) \lor (gen_2 = healthy)) \land ((ru_1 = healthy)) \lor (ru_2 = healthy))), \quad (2)$$

$$\varphi_s = \Box \neg ((c_0 = closed) \land (c_7 = closed) \land (c_3 = closed)) \land \bigwedge_{i \in \{1,2,3,4\}} \Box (bus_i = powered), \quad (3)$$

and gen_1 , gen_2 , ru_1 , and ru_2 are the health statuses of the two generators and two rectifier units, respectively. The contactors c_0 and c_7 are next to the generators in the topology shown in Fig. 5, and c_3 is between the AC buses. Therefore, contactors c_0 , c_7 , and c_3 can never be closed at the same time, which otherwise would lead to paralleling two AC sources. The buses bus_1 , bus_2 , bus_3 , and bus_4 can be considered to be in an electrical connection to the loads. The final part of φ_s ensures that a bus can never be unpowered, given that the environment assumptions hold. However, measurements are taken at discrete time intervals. A continuous implementation have to allow a certain unpowered time. In addition to the centralized control protocol that realizes the global specification $\varphi_e \to \varphi_s$, we synthesize a distributed reactive control protocol following the theory in [10]. More specifically, we decompose the global specification into local specifications for the AC and DC parts of the system is such a way that, if the local specifications are realizable separately, then they can be implemented together and ensure the correctness of the global specification (under additional mild technical assumptions discussed in [10]). For example, the relatively simple global specification in (2)-(3) are decomposed into $\varphi_{e,AC} \to \varphi_{s,AC}$ and $\varphi_{e,DC} \to \varphi_{s,DC}$ for the AC and DC parts respectively, where

$$\begin{aligned} \varphi_{e,AC} &= \Box(((gen_1 = healthy) \lor (gen_2 = healthy)), \\ \varphi_{s,AC} &= \Box \neg ((c_0 = closed) \land (c_7 = closed) \land (c_3 = closed)) \land \bigwedge_{i \in \{1,2\}} \Box(bus_i = powered), \\ \varphi_{e,DC} &= \Box(((ru_1 = healthy) \lor (ru_2 = healthy)), \text{ and} \\ \varphi_{s,DC} &= \Box \neg ((c_5 = closed) \land (c_4 = closed) \land (c_6 = closed)) \land \bigwedge_{i \in \{3,4\}} \Box(bus_i = powered). \end{aligned}$$
(4)

C. The Synthesized Controller Automata

We use TuLiP to synthesize the control protocols for the global and the distributed specifications for the AC and DC parts. The centralized controller realizing the global specifications has 16 states (i.e., one state for each of the possible environment configurations in this case). On the other hand, each of the automata for the AC and DC parts contains 4 states.

Fig. 6 shows part of the output from TuLiP for the centralized controller. Each entry in the list has two lines that correspond to one state in the automaton. The valuations of the environment variables (i.e., the health statuses of the generators and rectifier units) and the controlled variables (i.e., the statuses of the contactors) are in the first line. The second line lists the possible transitions from the current state. Out of these possible transitions, the one that is implemented as the transition in the controlled variables is picked based on the transition in the environment variables. For example, if, from state 0 in which all generators and rectifier units are healthy, the environment transitions to a configuration in which rgen = 0 and rru = 0, then the controller transitions to state 4 and the contactor statuses switch to the values listed under state 4.

By its construction, as long as the environment satisfies its assumptions then, the controller can execute indefinitely and the contactors take actions such that the system requirements are satisfied.

```
State 0 with rank 0 -> <rgen:1, rru:1, lru:1, lgen:1, c3:0, c0:1, c7:1, c4:0, c5:1, c6:1>
With successors : 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 0
...
State 2 with rank 0 -> <rgen:0, rru:0, lru:0, lgen:1, c3:0, c0:0, c7:0, c4:0, c5:0, c6:0>
With no successors.
...
State 4 with rank 0 -> <rgen:0, rru:0, lru:1, lgen:1, c3:1, c0:0, c7:1, c4:1, c5:0, c6:1>
With successors : 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 0
...
State 6 with rank 0 -> <rgen:0, rru:1, lru:0, lgen:1, c3:1, c0:0, c7:1, c4:1, c5:1, c6:0>
With successors : 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 0
...
```

Fig. 6 TuLiP auto-generated automaton. All states are not represented in the figure. The environment variables are lgen, rgen, lru1, and rru. The controlled variables are c0, c7, c3, c5, c6, and c4.

```
global state;
while 1
    lgen = readlgen();
    rgen = readrgen();
    lru = readlru();
    rcu = readrru();
    [c0, c7, c3, c5, c6, c4] = controller(lgen, rgen, lru, rru);
    write2contactors(c0, c7, c3, c5, c6, c4);
```

Fig. 7 Code that implements a control cycle on the hardware testbed. In each control cycle the sensors are first read for the generators, then the sensors for the rectifier units. After that, the mscript generated by TuLiP is used to decide on the contactor statuses. Finally, the state of contactors are set.

However, if the environment assumptions are violated, then the controller may end up in a state with no outgoing transition, referred to as the "no-successor" state in Fig. 6. For example, if both generators or both rectifier units are unhealthy (as in state 2) the controller will enter a no-successor state.

D. Determining the Statuses of the Generators, Rectifier Units and Buses

Note that the execution of the reactive protocols synthesized in Section IV C requires the knowledge of the valuations of the environment variables at each execution step. The lines of code in Fig. 7 illustrate the control cycle workflow, which consists of three steps: read environment variables, run control logic, and assign values to controlled variables. The four environment variables are represented in Fig. 7 and are labeled as lgen, rgen, lru, and rru. The controlled variables are c0, c7, c3, c5, c6, and c4. The health statuses of the generators and rectifier units are not directly monitored. Their values are deduced from certain voltage measurements at appropriate locations in the power distribution network. Let V_0 be a pre-specified positive constant and T_r be the time that elapses while the corresponding sensor reading takes place. Then, a generator or a rectifier unit will be registered unhealthy if the magnitude of its voltage reading stays below V_0 for a time T_{sense} that elapse so it includes T_r . Time T_{sense} therefore has to be greater than or equal to T_r . For example, the function readrgen() in Fig. 7 check if the right generator is above or below the threshold V_0 .

Each of the rectifier units in the testbed consists of a single-phase diode rectifier followed by a capacitor and a voltage regulator. The capacitor is connected to the DC bus in order to reduce the voltage ripple at the input of the voltage regulator. The specification used in controller synthesis requires that all buses are powered at all times. However, there exists time T during which the bus is unpowered that needs to elapse when the controller take actions. See Section IVE for the details on how the constant T is chosen based on the characteristics of the testbed estimated empirically. This situation does not necessarily mean that relevant components are influenced by that the voltage is below V_0 during time T. For example, consider a rectifier unit connected to an AC bus. It contains a capacitor which charges to the peak voltage each half cycle of the AC voltage sine curve and then discharges at a slower rate through the load while the rectified voltage drops before the beginning of next half cycle. Therefore, duration, call T_{RU} , of time that takes for the capacitor voltage to drop below an acceptable value depends on the capacitance of the capacitor and the amount of current drawn by the load. If T_{RU} is strictly larger than T, then it can be guaranteed that the DC voltage stays above a pre-specified threshold provided that the corresponding rectifier unit is healthy. Furthermore, the time over which a generator has to remain healthy during each control cycle is not arbitrarily small because it needs to be healthy for at least a time, call T_r , during that the sensor is read. Otherwise, we would violate the environment assumptions. The time T_r is enough to charge the capacitor to the peak voltage. Therefore, the capacitances and the current drawn by the DC loads in the testbed are arranged so that T_r is large enough to charge the capacitors in the rectifier units to their peak voltage.

For proper operation of the controller, the sensors shall provide complete and consistent information. To this end, their placement, functionality and accuracy play crucial roles in design. Analog-to-digital (A/D) inputs on the relay board are used to monitor the system conditions; the input connections range from 1 to 8, as shown Fig. 8. The system can have four threshold values because it has four sensors. The A/D inputs on the relay board can read 0 to 5 VDC. The first two sensors will have a threshold value of $V_0 = 5V_{AC}/256$, and the other two sensors will have a threshold value of $V_0 = 5V_{DC}/256$. We check the voltage on an A/D input with an accuracy of 8 bits therefore, V_{AC} and V_{DC} are scaled in the range of 0-256. The first lines of the code in Fig. 7 read the voltages from each sensor and check if the voltage measurement is above or below the threshold values, V_{AC} and V_{DC} . The status of each environment variable can then be assigned as healthy or unhealthy accordingly.

The voltage sensing connections are represented by the numbered arrows in Fig. 8, which correspond to the numbered arrows in Fig. 5. The transformers that act as generator sources can be unplugged in order to simulate a generator failure. The A/D inputs cannot handle 24 VAC; therefore, voltage sensing for generator failures on 24 VAC is handled using additional relays. The relays connect a 3.6 V circuit to a battery when triggered by the



Fig. 8 Sensing configuration for the testbed. The numbered arrows denote voltage sensing connections from Fig. 5.

voltage from the transformers. Therefore, the threshold value V'_{AC} for the system is set whenever the additional relays are not triggered anymore. The threshold value V_{AC} , which is read from the A/D inputs, is set to 100, approximately 2 V using an 8-bit resolution. The voltage V'_{AC} is set by the relay manufacturer but is usually a low value compared to when they are triggered. The relays used in the testbed have a minimum turn-off voltage of 3.6 VAC and a maximum turn-on voltage of 18 VAC.

The rectifier units are connected to a switch which can be used to generate a fault in the DC subsystem. The voltage sensors of the rectifier units are directly connected to the A/D inputs of the relay board because the voltage is tuned to 2.5 VDC using the variable DC voltage regulator on the rectifier units. When the status of a switch that injects a fault on the DC subsystem is changed, there will be no potential difference between ground and the wire connected to the sensor; therefore, V_{DC} can be chosen anywhere between 0 and 128 and is usually set to 100 (i.e., equal to V_{AC}).

E. Testbed Characteristics

We now describe the characteristics of the hardware testbed. The characteristics depend on the relay delay time, T_d , and control cycle times, T_c and T'_c . The relay delay time is the time delay between the time a command to actuate the relay is written on the relay board and the time the action (i.e., relay opening or closing) is completed. The control cycle times are defined as

$$T_c = 4T_r + T_I + T_w$$

$$T'_c = 4T_r + T_I,$$
(5)

where T_r is the time taken to read the health status from one environment variable, T_I is the time taken to run the Matlab script generated from TuLiP, and T_w is the time taken to write information to the board. We also have to consider the control cycle time T'_c because, if the system remains the same, writing information to the board is not necessary in that iteration.

As discussed in Section IV D, we reason the definition of when a bus becomes unpowered based on these timing characteristics. Consider the code listed in Fig. 7, which shows that the controller reads the health status from each environment variable in a specified order. Therefore, we have to include first T_c , and then part of T'_c , from the previous control cycle in the limit T. We approximate the time to read the health status from generators and rectifier units as $T_r \approx T'_c/4$ because the time T_I is negligible compared to T_r . Therefore, a reasonable estimate of time T can be calculated with

$$T \approx \max\left(T_d\right) + \max\left(T_c\right) + \frac{4-n}{4}\max\left(T_c'\right),\tag{6}$$

where $n \in \{1, 2, 3, 4\}$ is a number that denotes the order of when the faulty environment variable is read in the code. Table 1 summarizes the variables that characterize the testbed. The relay delay time T_d can be found from the board specifications and should be less than 20 ms. We also get a relay delay time from the additional relays that measure the AC voltage. However, we assume that both the delay time from the board relays and the additional relays never exceed 20 ms. The control cycle times and the time it takes to run the code are estimated empirically. We calculated $T_r \approx T'_c/4 = 58.5 \text{ ms}$ and used Eq. (5) to calculate $\max(T_w) = \max(T_c) - \min(T'_c) = 166.7 \text{ ms}$ and

Variable	Max Value	Description
T	$587.9\mathrm{ms}$	Time limit for the bus to stay unpowered
T_d	$20\mathrm{ms}$	The relay delay time
T_c	$333.3\mathrm{ms}$	Control cycle time $(T_r + T_I + T_w)$.
		The mean and minimum values are $303.7\mathrm{ms}$ and $282.5\mathrm{ms}$, respectively.
T_c'	$234.1\mathrm{ms}$	Control cycle time without relay changes $(T_r + T_I)$.
		The mean and minimum values are $187.5\mathrm{ms}$ and $166.6\mathrm{ms},$ respectively.
T_I	$1\mathrm{ms}$	Time it takes to run the control logic generated from TuLiP .
T_r	$58.5\mathrm{ms}$	Time it takes to read information from one sensors
T_w	$166.7\mathrm{ms}$	Time it takes to write information to all relays that need to take actions

Table 1 Summary of the variables that characterize the testbed. The values for T_c , T'_c , and T_I were calculated from 20, 250, and 400 measurements, respectively. The times were calculated on a Macbook Pro with 2.3 Ghz Intel Core i7 Processor.

 $mean(T_w) = mean(T_c) - mean(T'_c) = 116.2 \,\mathrm{ms}.$

For an example calculation, consider a configuration of two generators and two rectifier units, such as shown in the topology of Fig. 4, where one generator is read first in the code (n = 1) and the other generator is read second (n = 2). The rectifier units are read third and fourth, respectively, in the code. The maximum unpowered time for the left AC bus on the hardware testbed can be calculated with Eq. (6). Thus, $T \approx \max(T_d) + \max(T_c) + \frac{3}{4}\max(T'_c) = 587.9$ ms. The unpowered times for the right AC bus, left DC bus, and right DC bus are calculated in the same way as 470.4 ms, 411.8 ms, and 353.3 ms, respectively.

Thus, it can be concluded that the unpowered time depends on where the fault is injected. The components connected to the right DC bus are least affected in the case of a fault, whereas the components connected to the left AC bus are most affected in the case of a fault.

V. Simulation Models for the Testbed

As discussed earlier, correctness of an automatically synthesized control software should be interpreted with respect to the abstract models and specifications used in synthesis. Therefore, before control software is implemented and tested on actual hardware, it is useful to develop high-fidelity simulation models to explore potential shortcomings of the abstract models used in synthesis as well as to test continuous time properties not precisely captured by LTL specifications. This section details the simulation models for the hardware testbed including potential control architectures.

In this work, we used Matlab Simulink [11], a graphical tool with a wide variety of built-in

functions that can be assembled into complete systems, and, in particular the SimPowerSystems toolbox [24], which is a physical modeling tool for electric power systems. With SimPowerSystems, models for an entire electric power system can be built just as it would be assembled from physical components. The constituent blocks are linked together with ideal conductors and may be linear, nonlinear, continuous, or discrete. It is also easy to integrate TuLiP controllers into Simulink models as TuLiP has the ability to export controllers in the form of a Matlab script that can be used as a Simulink block.

The SimPowerSystems models used in this study are built in accordance with the hardware. The generator units are connected to be 180° out of phase in order to create a shortcut when paralleled. The rectifier units in the Simulink model are built from a transformer, diode bridge, and capacitor to smooth out the ripple from the AC-to-DC conversion. Generators and rectifier units are equipped with fault injection inputs and fault sensors. The delays in the relay opening and closing times are modeled using saturated integrators to capture the formation of the electromagnetic field when the relays are actuated. Fig. 9 shows the topology when a centralized control architecture is used. In the this model, the embedded Matlab function block, BPCU, runs the control logic script generated from TuLiP. There are several adjustable parameters in the model that are initialized with a configuration script. The relay delay time, T_d , is set to 20 ms. The time T_r it takes to read a sensor value is modeled with a delay between the sensing and the control command times. Because sensors are sequentially read as indicated in Section IVE, we set the delays from the fault sensors to kT_r , where $k \in \{0, 1, 2, 3\}$ is the order in which the sensor is read (k = 0) is the first sensor and k = 3 is the fourth). The time, T_I for running the control logic, the mean time of $4T_r$, and the mean time of T_w to write the information to the relays are lumped into a sampling time T_s of the BPCU block; therefore, we let $T_s = 4T_r + T_I + T_w = T_c$. The mean value of T_c is chosen according to Table 1. To reflect the variability in timing, a uniform random value is added to the sensor reading delays $T_r \approx T'_c/4$, so that the overall control cycle time T'_c ranged between its maximum and minimum value given in Table 1. The configuration script is also used to define different scenarios that involved different combinations of fault conditions.

When we implement the distributed logic on the hardware testbed it is still centralized in that

only one relay board is connected to one computer; that is, all sensors and relays are connected to the same computer, which leads to the same timing characteristics regardless a centralized or distributed logic is used. With Simulink, it is possible to mimic the behavior of distributed control architecture with two relay boards controlled by different automaton running on two different computers. The distributed Simulink model is shown in Fig. 10 where AC and DC subsystems are sensed and controlled by different embedded Matlab function blocks. There are several advantages of this distributed architecture. First, it increases the robustness of the system. For instance, even if the computer running the control logic for the DC subsystem fails, or both rectifier units fail, the AC subsystem will continue operating and providing power to the AC buses. Second, it reduces the control cycle time by reducing the number of sensors each controller reads from. This effect would be particularly noticeable for the hardware used in the testbed as the largest contribution to the control cycle times is the total time it takes to read data from the board, i.e., $4T_r$ for the centralized case and $2T_r$ for the distributed case. Finally, in the distributed architecture it is possible to introduce and study the effects of asynchrony by choosing different sampling times for the two different controllers in the Simulink model.

VI. Controller Tests

We now discuss some implementations of the controllers by running tests on Simulink and on the hardware. The following examples also illustrate the differences between the high-fidelity simulations and testbed characteristics.

A. An Example Control Test on Hardware

Fig. 11 shows the voltage measurement for a centralized 16-state controller. The measurement was taken on the AC bus when the power cord to the transformer which is read by the sensor for n = 2 in Eq. (6) was unplugged. The power cord was unplugged at t = 2.83 s, which is denoted by the first vertical line in Fig. 11. The second vertical line from the left indicates when the controller reacted and powered up the bus through another path, which occurred at t = 3.1 s. After that, we plugged the power cord back in again. At time t = 3.73 s the controller react on that the power cord was back, which was accompanied by a discernible change in the sine curve. Once the transformer



Fig. 9 An example SimPowerSystems model which corresponds to the single-line diagram in Fig. 4. The embedded Matlab function called BPCU (bus power control unit) controls the system with a 16-state TuLiP automaton. There are two AC loads connected to the AC subsystem and two DC loads in the bottom connected to the DC subsystem. In addition, there is a Matlab function that can be used for fault injection at a specific or random time.

was plugged in again after a fault, the time during which the bus had been without power is not noticeable because the controller sends simultaneous commands to the two relays.

The measured unpowered bus times are listed in Table 2 for n = 2. The maximum value is $T_{max} = 414.9 \text{ ms.}$ As calculated in Section IV E, time T = 470.35 ms; therefore, $T_{max} < T$. We used a digital storage oscilloscope (Rigol DS1052E 50MHz) for the measurements. The measurement data were imported into Matlab to plot sinusoidal curves (e.g., Fig. 11) and analyzed the signal to estimate the unpowered times.



Fig. 10 An example SimPowerSystems model which corresponds to the single-line diagram in Fig. 4. The model has two embedded Matlab functions called BPCU; each of them runs on a four-state TuLiP automaton. There are two AC loads connected to the AC subsystem and two DC loads in the bottom connected to the DC subsystem. In addition, there is a Matlab function that can be used for fault injection at a specific or random time.

	Bus-unpowered time
Mean	$333.9\mathrm{ms}$
Max	$414.9\mathrm{ms}$
Min	$232.7\mathrm{ms}$

Table 2 Time for which bus was unpowered after a fault had been injected on the hardware testbed. These values were calculated using measurements from 10 fault injections.

B. An Example Control Test on Simulink

Fig. 12 illustrates the bus voltage measurements of the Simulink model when a fault was injected on a generator which is read by the sensor for n = 2. Note the similarities with the hardware measurements based on the unpowered time and change in the sine curve when the faulty generator



Fig. 11 Bus voltage measurement when a generator was switched off and then turned back on. The first vertical line indicates the fault, the second vertical line is when the controller reacts, and the third line is when the generator was turned back on.

	Bus-unpowered tim	ıe
Mean	$269.7\mathrm{ms}$	
Max	$379.0\mathrm{ms}$	
Min	$146.0\mathrm{ms}$	

Table 3 Time for which bus was unpowered after a fault had been injected in the Simulink model. These values were calculated using measurements from 10 fault injections.

was switched on again.

The measured unpowered bus times are listed in Table 3; the maximum value is $T_{max} = 333.0$ ms. Thus, we can verify with Eq. (6) that $T_{max} < T$.

C. Comparison Between Simulation Results

Fig. 11 and Fig. 12 show the similarities in the AC voltages measured in the Simulink based simulations and in the hardware tests. Fig. 13 illustrates the measured voltage on the DC bus when a rectifier fault was injected. The same unpowered behavior are seen in both figures. However, no change could be detected in the voltage when the rectifier unit became healthy in the Simulink based simulation. Partly, because of the ideal behavior of the components; e.g., contactor delays.

Table 2 and Table 3 show that the unpowered time is slightly lower in the Simulink based simulation compared to that in the hardware testbed.



Fig. 12 Bus voltage measurement in Simulink when a generator was switched off and then turned back on.

Bus-unpowered time
$204.4\mathrm{ms}$
$274.0\mathrm{ms}$
$121.0\mathrm{ms}$

Table 4 Time for which bus was unpowered after a fault had been injected in the distributed logic Simulink model. These values were calculated using measurements from 10 fault injections.

Table 4 lists the unpowered times of the distributed logic in Simulink. Note the decrease in the unpowered times compared to the values shown in Table 2 and Table 3. An interesting observation from the executions of the centralized and distributed controllers (synthesized to realize the local specifications discussed in Section IV B) is that if the centralized controller senses that both rectifier units are unhealthy (i.e., the environment assumption on the DC side is violated), the entire controller stops working because a no-successor state has been reached. On the other hand, in the case of the distributed controllers, the AC part continues executing and its own requirements are still fulfilled whereas the DC part stalls at a no-successor state with no guarantees on the satisfaction of its requirements.

VII. Conclusions, Limitations and Extensions

We demonstrated a formalized workflow for the design of control protocols for primary distribution in electric power systems on more-electric aircraft. The steps of the workflow include



Fig. 13 Bus voltage measurement on the testbed when a rectifier unit was switched off and then turned back on. Figure in (a) is for when the rectifier was turned off and on twice. Figure (b) is for when a fault was injected at 2.83 ms in the Simulink model.

(i) establishing formal specifications that capture safety and performance requirements and abstract models of the allowable evolution of the underlying system; (ii) automatically synthesizing control protocols from these specifications and models; and validating/testing these protocols on high-fidelity simulations models and a hardware testbed. For the hardware tests, we employed an academic-scale setup we had developed in our recent work to initiate some of the salient features of power networks on aircraft.

On the hardware testbed, we injected faults in the hardware testbed by unplugging the power cords and changing the switches. With this method of fault injection, it is relatively difficult (if not impractical) to switch off a generator and a rectifier unit within the same control cycle. A more accurate approach to generate faults would be using an additional relay board which would enable to systematically study synchronous, correlated, and cascaded failures and their influence on controller performance; with the current method of fault injection, it could be difficult to switch off a generator and a rectifier unit within the same control cycle.

Through the high-fidelity simulations, we showed that the bus unpowered time significantly decreases when we use distributed controllers running with different automata on two different computers and on two relay boards. Therefore, it would be more suitable on the hardware testbed to use a distributed control architecture more like that on an aircraft.

On an aircraft, the controller is an embedded system designated for a specific task. To increase

its reliability and performance, the hardware model could be adapted to run the relay boards through microcontrollers. Embedded code for these microcontrollers can be readily generated using Matlab.

VIII. Acknowledgments

The authors wish to acknowledge the funding from the Industrial Cyber-Physical Systems Center (iCyPhy), and AFOSR (award # FA9550-12-1-0302), and thank Rich Poisson from United Technologies Aerospace Systems for helpful discussions about the development of the hardware testbed.

References

- Rosero, J., Ortega, J., Aldabas, E., and Romeral, L., "Moving towards a more electric aircraft," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 22, No. 3, 2007, pp. 3–9.
- [2] Moir, I. and Seabridge, A., Aircraft Systems: Mechanical, Electrical and Avionics Subsystems Integration. 3rd Edition, John Wiley and Sons, Ltd, 2008.
- [3] Clarke, E. M. and Wing, J. M., "Formal methods: state of the art and future directions," ACM Computing Surveys, Vol. 28, No. 4, 1996, pp. 626–643, doi:10.1145/242223.242257.
- [4] Woodcock, J., Larsen, P. G., Bicarregui, J., and Fitzgerald, J., "Formal methods: Practice and experience," ACM Computing Surveys, Vol. 41, No. 4, 2009, p. 19.
- [5] Manna, Z. and Pnueli, A., The Temporal Logic of Reactive and Concurrent Systems Specification, Springer, 1992.
- [6] Emerson, E. A., "Temporal and modal logic," *Handbook of Theoretical Computer Science*, Vol. 2, 1990, pp. 995–1072.
- [7] Wongpiromsarn, T., Topcu, U., Ozay, N., Xu, H., and Murray, R., "TuLiP: a software toolbox for receding horizon temporal logic planning," in "International Conference on Hybrid Systems: Computation and Control,", 2011.
- [8] Xu, H., Topcu, U., and Murray, R., "A Case Study on Reactive Protocols for Aircraft Electric Power Distribution," in "IEEE Conference on Decision and Control,", 2012.
- [9] Wongpiromsarn, T., Topcu, U., and Murray, R. M., "Formal synthesis of embedded control software for vehicle management systems," in "AIAA Infotech@Aerospace,", 2011.

- [10] Ozay, N., Topcu, U., and Murray, R. M., "Distributed Power Allocation for Vehicle Management Systems," in "IEEE Conference on Decision and Control,", 2011.
- [11] Simulink, version 8.0 (R2012b), The MathWorks Inc.
- [12] Rogersten, R., Xu, H., Ozay, N., Topcu, U., and Murray, R. M., "An Aircraft Electric Power Testbed for Validating Automatically Synthesized Reactive Control Protocols," in "International Conference on Hybrid Systems: Computation and Control,", 2013.
- [13] Michalko, R., "Electrical starting, generation, conversion and distribution system architecture for a more electric vehicle," US Patent number US 7439634, 2008.
- [14] Lyu, M. R. et al., Handbook of Software Reliability Engineering, Vol. 3, IEEE Computer Society Press CA, 1996.
- [15] Wood, A. J. and Wollenberg, B. F., Power Generation, Operation, and Control, John Wiley & Sons, 2012.
- [16] "MIL-STD-704F, Aircraft Electric Power Characteristics," http://www.wbdg.org/ccb/FEDMIL/std704f.pdf, 2004.
- [17] Baier, C. and Katoen, J., Principles of Model Checking, MIT Press, 1999.
- [18] Xu, H., Topcu, U., and Murray, R. M., "Specification and synthesis for aircraft electric power distribution," *IEEE Transactions on Industrial Informatics*. Under review.
- [19] Bengtsson, J., Larsen, K., Larsson, F., Pettersson, P., and Yi, W., UPPAAL a Tool Suite for Automatic Verification of Real-Time Systems, Springer, 1996.
- [20] Hinton, A., Kwiatkowska, M., Norman, G., and Parker, D., "PRISM: A tool for automatic verification of probabilistic systems," in "Tools and Algorithms for the Construction and Analysis of Systems," Springer, pp. 441–444, 2006.
- [21] Pnueli, A., "Applications of temporal logic to the specification and verification of reactive systems: a survey of current trends," in de Bakker, J. W., de Roever, W. P., and Rozenberg, G., eds., "Current Trends in Concurrency. Overviews and Tutorials," Springer-Verlag New York, Inc., pp. 510–584, 1986.
- [22] Piterman, N., Pneuli, A., and Sa'ar, Y., "Synthesis of reactive(1) designs," Verification, Model Checking and Abstract Interpretation, Vol. 3855, 2006, pp. 364–380.
- [23] Pnueli, A., Sa'ar, Y., and Zuck, L., "JTLV a framework for developing Verification Algorithms," in "International Conference on Computer Aided Verification,", 2010.
- [24] SimPowerSystems, version 5.7 (R2012b), The MathWorks Inc.
- [25] The specific relay board used in the testbed are supplied by RelayPROS. URL: www.relaypros.com