

Dynamic State Estimation in Distributed Aircraft Electric Control Systems via Adaptive Submodularity

Quentin Mailliet[§], Huan Xu[†], Necmiye Ozay[¶] and Richard M. Murray[‡]

Abstract— We consider the problem of estimating the discrete state of an aircraft electric system under a distributed control architecture through active sensing. The main idea is to use a set of controllable switches to reconfigure the system in order to gather more information about the unknown state. By adaptively making a sequence of reconfiguration decisions with uncertain outcome, then correlating measurements and prior information to make the next decision, we aim to reduce the uncertainty. A greedy strategy is developed that maximizes the one-step expected uncertainty reduction. By exploiting recent results on adaptive submodularity, we give theoretical guarantees on the worst-case performance of the greedy strategy. We apply the proposed method in a fault detection scenario where the discrete state captures possible faults in various circuit components. In addition, simple abstraction rules are proposed to alleviate state space explosion and to scale up the strategy. Finally, the efficiency of the proposed method is demonstrated empirically on different circuits.

I. INTRODUCTION AND MOTIVATION

The increasing focus on more-electric aircraft within the aerospace industry signifies progress in the direction of more energy efficient vehicles. Electric systems are replacing pneumatic, mechanical, and hydraulic subsystems, thus reducing weight, easing maintenance, and improving aircraft controllability and configurability [14]. Yet as more subsystems rely on electric power, the flight-criticality of an aircraft becomes more dependent on the electric power system as well. Because the state of the system is determined solely from sensor measurements, the problem of state estimation from sensor readings is crucial to the safety of the entire aircraft.

Estimation of electric power systems using optimization-based techniques is a well-established area [1], [4], [15]. A large body of work exists on diagnostics of electric power systems focusing on AC systems [5], as well as large vehicle systems. [12] examines the diagnostics for the international space station, [10] for an aircraft electric system, and [7] for a marine vehicle power system. For a DC system, [8] uses an optimization-based approach to estimate fault states.

Previous work in electric power system state estimation has focused on static, centralized estimation problems with continuous states. We perform discrete state estimation using active control of switches within a distributed control architecture. The system reconfigures itself through a set of controllable contactors (i.e., electrically controlled switches).

Once reconfigured, new sensor measurements are taken to gain more information about the unknown state. We adaptively sequence switching actions by use of a greedy strategy that maximizes the one-step expected uncertainty reduction. By exploiting recent results in adaptive submodularity [11], [6], we provide theoretical bounds for the worst-case performance of the greedy strategy. Such dynamic state estimation techniques have been proposed in the context of Markov jump linear systems [3], information gathering in robotics [13], [20], active hypothesis testing [16], and active learning [9]. To the best of our knowledge, these ideas have not been applied before in electric power system state estimation and fault diagnosis problems.

Recently, correct-by-construction control synthesis has been applied to power allocation and distribution in aircraft electric power systems [21], [17], [22]. A critical assumption in these papers is that the high-level reactive control protocol has an accurate knowledge of the system states, including fault states, so that it can reroute the power accordingly. An expensive, hence undesirable, solution to achieve accurate state estimates is to equip the system with a large number of sensors. Software, however, is cheaper and more amenable to change than hardware.

The goal of the current paper is to obtain high-accuracy state estimates with a limited number of sensors by utilizing software-based dynamic estimation strategies. We are particularly interested in detecting and localizing faults in the system. As it is common to use discrete models for fault diagnosis [19], continuous values of voltage and current, as well as health statuses of components in the system, are discretized before performing state estimation. Additionally, a discrete framework is well-suited for combining the proposed estimation strategy with control synthesis results described in [22] in future work.

II. PROBLEM SETUP

A. General problem description

Consider an aircraft electric power system topology, which can be represented by a graph data structure $G = (\mathcal{N}, \mathcal{E})$. Figure 1 shows a representative single-line diagram. The set \mathcal{N} of nodes in the graph contains the following components: generators (\mathcal{G}), rectifier units (\mathcal{R}), and voltage sensors (\mathcal{S}). The set \mathcal{E} of edges contains all contactors (and solid wire links) between components. The status of contactors $\mathcal{C} \subseteq \mathcal{E}$ can either be *open* or *closed*. A node corresponding to a rectifier unit has no outgoing edges on the AC side and no incoming edges on the DC side to reflect the fact that

[§]Mines ParisTech (Ecole des Mines de Paris), Paris, 75005 France, correspondence: quentin.mailliet@mines-paristech.fr

[†]University of Maryland, College Park, MD 20742

[¶]University of Michigan, Ann Arbor, MI 48109

[‡]California Institute of Technology, Pasadena, CA 91106

they contain a diode. The rest of the edges in the graph are bidirectional.

Elements in the set of generators $\mathcal{G} \subseteq \mathcal{N}$ and rectifier units $\mathcal{R} \subseteq \mathcal{N}$ are uncontrollable, and can take values of *unhealthy* (i.e., the component is online but outputting a voltage not in admissible range), *healthy* (i.e., the component is online and outputting the correct voltage), or *offline* (i.e., no power output, open circuit). Measurements read from the sensors $\mathcal{S} \subseteq \mathcal{N}$ will depend on the status of generators, rectifier units, and contactors. We say that there is a *live path* between two components if there exists a simple path in the graph G that connects the two nodes corresponding to these components, there is no offline component along the path including end nodes, and the contactors along this path are all closed. The readings of a sensor $s \in \mathcal{S}$ can then take the following values (i) *improper voltage*: if there is a live path between s and some $g \in \mathcal{G}$ (not offline by definition of live path), and either g or some rectifier $r \in \mathcal{R}$ along such a path is unhealthy; (ii) *admissible voltage*: for all $g \in \mathcal{G}$ that have a live path to s , both g and r along such paths are healthy; or (iii) *no voltage*: there is no live path between s and any generator $g \in \mathcal{G}$.

On top of the circuit and sensing topology is a distributed control architecture with a dynamic state estimation mechanism. We assume that one of the embedded controllers is responsible for dynamic state estimation, hereafter referred to as the fault detection controller. The fault detection controller is able to control a subset $\mathcal{C} \setminus \mathcal{C}'$ of contactors (e.g., those labelled with blue in Fig. 1).

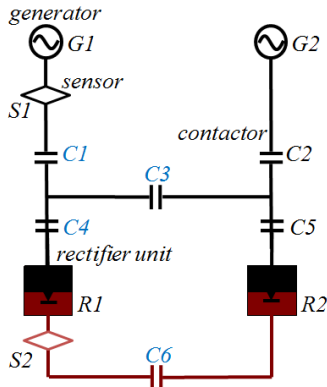


Fig. 1. A single-line diagram of a simple circuit with AC components (in black) and DC components (in red).

The state x of the system is defined as a valuation on all components $n \in \mathcal{G} \cup \mathcal{R}$ and uncontrollable contactors $e \in \mathcal{C}' \subseteq \mathcal{C}$. We define Ω as the set of all states, i.e all the different valuations of the components. The state x is unknown and hence modeled as a random variable X that can only be determined by sensor measurements mapped back to a set of possible states in which the circuit may be. The overall goal is to design a strategy the fault detection controller runs to adaptively estimate the discrete state of the circuit by taking “actions” (i.e., closing and opening controllable contactors), and then reading voltage sensor

measurements.

B. Mathematical formulation

In this section we introduce the relevant notations used throughout the rest of the paper and mathematically formulate the estimation problem.

The state X of the circuit is modeled as a random variable. Data on component types and reliability levels can be used to build a probability measure $\mathbb{P}[x]$ on Ω . At the beginning of the state estimation process the system is in the (unknown) state $x_0 \in \Omega$. We assume that faults in the system are independent, and that x_0 remains fixed during the estimation process. This is a reasonable assumption because the timescale of the estimation process is meant to be much smaller than the failure rates of the components and the timescales of the other controllers in the system.

For the controllable subset of contactors, there exists a set \mathcal{V} of actions v that can be performed and a set \mathcal{Y} of measurements y that can be observed. For an action $v \in \mathcal{V}$, $y = \mu(v, x) \in (Y)$ is the unique outcome of performing action v if the system is in the state x . The actions $\{v_0, \dots, v_t\}$ performed and outcomes $\{y_0, \dots, y_t\}$ observed up until step t are represented by the partial realization $\psi_t = \{(v_i, y_i)\}_{i \in \{0, \dots, t\}}$. Given two partial realizations ψ_t and $\psi_{t'}$, we say that ψ_t is a subrealization of $\psi_{t'}$ if $\psi_t \subseteq \psi_{t'}$. At each step t , the probability measure $\mathbb{P}[x]$ can be updated by conditioning it on ψ_t to obtain $\mathbb{P}[x | \psi_t]$.

We are interested in an estimation process adaptively eliminating “invalid” states to get to the actual state x_0 . We define $D(y, v)$, with $y = \mu(v, x_0)$, to be the set of states $x \in \Omega$ that are indistinguishable from x_0 under the action v . Formally, $D(\mu(v, x_0), v) = \{x \in \Omega \mid \mu(v, x) = \mu(v, x_0)\}$. We further extend this concept by defining $h(v_{0:t}, x_0)$, the set of states that produce the same set of outcomes $\{\mu(v_0, x_0), \dots, \mu(v_t, x_0)\}$ as x_0 under the same set of actions $\{v_0, \dots, v_t\}$. In the remainder of the paper, we use S_t as a shorthand for $h(v_{0:t}, x_0)$. If, at step t , we perform a new action $v' \notin \psi_t$, there exists a recursive relation between the two sets of states:

$$h(v_{0:t} \cup \{v'\}, x_0) = h(v_{0:t}, x_0) \cap D(\mu(v', x_0), v'), \quad (1)$$

which leads immediately to

$$S_t = \bigcap_{i \in \{0, \dots, t\}} D(\mu(v_i, x_0), v_i). \quad (2)$$

As only intersections are taken, the order of actions v_i does not matter.

To represent the uncertainty in the state estimate, we define an objective function $f : 2^{\mathcal{V} \times \mathcal{Y}} \times \Omega \rightarrow \mathbb{R}_+$ that maps the the set of actions $A \subseteq \mathcal{V}$ under state x_0 to reward $f(A, x_0)$. A strategy π is a function from partial realizations to actions such that $\pi(\psi_t)$ is the action v_{t+1} taken by π when observing ψ_t . We denote $\check{\mathcal{V}}(\pi, x_0) \subseteq \mathcal{V}$ the set of all the actions performed under the strategy π , the state of the system being x_0 . In the general case, $\check{\mathcal{V}}(\pi, x_0) \neq \mathcal{V}$.

The fault detection controller is assigned a budget $k \ll |\mathcal{V}|$, indicating the number of steps within which the estimation process should terminate. The system is initially

in the state x_0 , which is fixed and unknown, and the controlled contactors are in some initial configuration v_0 . Initial configuration v_0 and the corresponding measurement y_0 constitute ψ_0 . Then, for $i = 1, \dots, k$, we consider the following process:

$$v_i = \pi(\psi_{i-1}) \quad (3a)$$

$$y_i = \mu(v_i, x_0) \quad (3b)$$

$$\psi_i = \psi_{i-1} \cup (v_i, y_i). \quad (3c)$$

Equations (3a) - (3c) represent the decision making, measurement, and update in the estimation process, respectively.

The goal is to reduce the uncertainty of X represented by the probability distribution $\mathbb{P}[x]$ through performing k actions. To that end, the following reward function is considered:

$$f(v_{0:k}, x_0) = -\mathbb{P}[S_k] = -\sum_{x \in S_k} \mathbb{P}[x]. \quad (4)$$

The behavior driven by the maximization of f is to remove as much probability mass from Ω as possible in k steps. It is also worth noting that when the underlying probability distribution on Ω is uniform, f is just proportional to the size of S_k and so maximizing f is equivalent to minimizing the number of indistinguishable states.

The goal of estimation is to find the strategy that allows the “best expected estimate” for the state, i.e, the strategy π^* such that

$$\pi^* \in \arg \max_{\pi} \mathbb{E}[f(\tilde{\mathcal{V}}(\pi, X), X)], \quad (5)$$

subject to $|\tilde{\mathcal{V}}(\pi, x)| \leq k$ for all x , and with expectation taken with respect to $\mathbb{P}[x]$.

III. STRATEGY

In this section, we describe the algorithm used to solve the state estimation problem and give performance guarantees on the worst-case execution.

A. Greedy strategy

The optimal strategy for the fault detection controller would plan ahead for k steps. Complexity, however, scales up exponentially with k . To address the problem efficiently we develop a greedy strategy that selects, at each step, the action maximizing the expected one-step gain in uncertainty reduction. At step t , the greedy strategy uses the available information ψ_t to compute the probability measure $\mathbb{P}[x | \psi_t]$ on the set S_t , using a Bayesian update:

$$\mathbb{P}[x | \psi_t] = \frac{\mathbb{P}[\psi_t | x] \mathbb{P}[x]}{\mathbb{P}[\psi_t]}, \quad \forall x \in \Omega. \quad (6)$$

As the measurement process is deterministic, for a given $x \in \Omega$ we have $\mathbb{P}[x | \psi_t] = \mathbb{1}_{\{x \in S_t\}}$, meaning that $\mathbb{P}[x | \psi_t] = 1$ if x belongs to S_t , and $\mathbb{P}[x | \psi_t] = 0$ otherwise. From (6) we then get:

$$\mathbb{P}[x | \psi_t] = \begin{cases} \frac{\mathbb{P}[x]}{\mathbb{P}[\psi_t]} & \forall x \in S_t \\ 0 & \text{elsewhere} \end{cases} \quad (7)$$

The term $\mathbb{P}[\psi_t]$ is the same for all x . It is a normalization coefficient that can be computed using $\sum_{x \in S_t} \mathbb{P}[x | \psi_t] = 1$ to obtain

$$\mathbb{P}[\psi_t] = \sum_{x \in S_t} \mathbb{P}[x]. \quad (8)$$

At each step t , the strategy consists of choosing the next action v_{t+1} that maximizes the gain in uncertainty reduction. Our measure of uncertainty comes from the value of the function f , established in Eq. (4), and therefore the benefit is expressed in terms of the change in f as we choose the action v . Consistent with our goal, we choose to maximize in mean the benefit at each step, the expectation taken with respect to the updated probability measure $\mathbb{P}[x | \psi_t]$. We obtain the greedy strategy:

$$v_{t+1} \in \arg \max_{v \in \mathcal{V}} \mathbb{E}[f(v_{0:t} \cup \{v\}, X) - f(v_{0:t}, X) | \psi_t]. \quad (9)$$

B. Worst-case performance guarantees

Greedy strategies in general can perform arbitrarily bad [2]. However, by exploiting recent results on adaptive submodularity, we give a lower bound on the performance of the proposed strategy. For a brief overview of adaptive submodularity and related definitions, see Appendix A. We next show that the function f defined in Eq. (4), is adaptive monotone and adaptive submodular (Def. 2 and 3).

Proposition 1: The function f defined in Eq. (4) is adaptive monotone.

Proof: Given an action $v \in \mathcal{V}$ and partial realization ψ_t at step t , we need to show the expected marginal benefit $\Delta(v | \psi_t)$ (see Def. 1) is nonnegative. For the cost function f , $\Delta(v | \psi_t)$ can be written as:

$$\Delta(v | \psi_t) = \mathbb{E}[f(v_{0:t}, X) | \psi_t] - \mathbb{E}[f(v_{0:t} \cup \{v\}, X) | \psi_t]. \quad (10)$$

By Eq. (7), we get

$$\Delta(v | \psi_t) = \sum_{x \in h(v_{0:t}, x_0)} \mathbb{P}[x | \psi_t] \phi(x), \quad (11)$$

with

$$\phi(x) = \sum_{\tilde{x} \in h(v_{0:t}, x)} \mathbb{P}[\tilde{x}] - \sum_{\tilde{x} \in h(v_{0:t} \cup \{v\}, x)} \mathbb{P}[\tilde{x}]. \quad (12)$$

By Eq. (1), $h(v_{0:t} \cup \{v\}, x)$ is a subset of $h(v_{0:t}, x)$ for every $x \in \Omega$. Thus, $\phi(x) \geq 0$, all the terms in the sum in Eq. (11) are non-negative, and $\Delta(v | \psi_t) \geq 0$. ■

Proposition 2: The function f defined in Eq. (4) is adaptive submodular.

Proof: Given in Appendix B. ■

Theorem 1: For any true state $x_0 \in \Omega$, the uncertainty reduction achieved in k steps by the greedy strategy given in Algorithm 1 is no worse than $(1 - 1/e)$ of what can be achieved in k steps by any other strategy, including the best possible strategy.

Proof: Follows directly from Propositions 1 and 2 and Theorem 2 given in Appendix A. ■

IV. IMPLEMENTATION

In this section, we give implementation details on the dynamic estimator employing the greedy strategy on some typical aircraft electric power system topologies. In order to reduce online computation, the inverse mapping from sensor measurements to compatible states of the circuit is conducted offline. Additionally, we propose some abstraction rules to reduce the size of the circuit as well as computation time.

A. Implementation details

The overall estimation process is summarized in Algorithm 1.

Algorithm 1 Adaptive greedy strategy

Input: Probability measure $\mathbb{P}[x]$ on Ω , number of actions to perform k . The system is in the state $x_0 \in \Omega$, fixed and unknown, and the controlled contactors are in some configuration v_0 .

Output: Partial realization ψ_k and the set S_k of compatible states after k actions are taken based on the strategy π_{greedy}

- 1: Take the measurement $y_0 = \mu(v_0, x_0)$.
 - 2: $\psi_0 = \{(v_0, y_0)\}$
 - 3: **for** $t \in \{1, \dots, k\}$ **do**
 - 4: $v_t = \pi_{greedy}(\psi_{t-1})$
 - 5: Perform action v_t
 - 6: Take the measurement $y_t = \mu(v_t, x_0)$
 - 7: $\psi_t = \psi_{t-1} \cup \{v_t, y_t\}$
 - 8: $S_t = S_{t-1} \cap D(y_t, v_t)$
 - 9: Compute $\mathbb{P}[x | \psi_t]$ (Bayesian update)
 - 10: **end for**
 - 11: **return** $(\psi_k, S_k, \mathbb{P}[x | \psi_k])$
-

In this algorithm, some items can be precomputed to improve run time. In particular, the inverse mapping from sensor measurements to compatible states does not have a closed form expression and the computation of the inverse map involves searching for paths on the graph $G = (\mathcal{N}, \mathcal{E})$ representing the circuit topology. Therefore, for all measurements $y \in \mathcal{Y}$ and all actions $v \in \mathcal{V}$, the sets $D(y, v)$ of states consistent with the action-measurement pairs (v, y) are computed offline to achieve a faster implementation. This collection is then accessed on the fly to significantly reduce the computation time as it is the most costly part of the algorithm.

Assumptions about the components and circuit can be easily incorporated in our framework. In particular, because these circuits are designed to achieve certain reliability levels, one common assumption is that at least one generator and one rectifier unit are online (delivering correct or improper voltage). These assumptions render certain states impossible, which are removed from the initial state set Ω .

B. Model reduction via abstraction

Although the greedy strategy provides an efficient way (with performance guarantees) to solve the dynamic state estimation problem, the offline computation for complex topologies can be very demanding as the number of possible states is exponential in the number of components whose states are being estimated. In this section, we give a set of rules that can be recursively applied to reduce the size of

the circuit by clustering certain components together into metacomponents.

Components (generators, rectifier units, contactors) are connected through their ports to form the circuit, and sensors are placed on some of these ports. The main reduction idea is that when two uncontrolled components are connected together and there is no sensor on their internal connecting port, some of the individual states of the components may become indistinguishable from what can be measured with the available sensors. Therefore, they can be treated as a single basic component, called a metacomponent, having the same global overall behavior. It is then possible to hierarchically estimate the system state, first by estimating the state of the metacomponent, and then mapping this state to possible states of individual components forming the metacomponent. When running the greedy algorithm on the reduced circuit, the probabilities of metacomponent states should be adjusted accordingly to ensure a lossless abstraction.

The rules we use to simplify the circuits are summarized in Fig. 2. Figure 2(a), for example, shows how the combination of generator and contactor can be abstracted into a single “generator” metacomponent. For the original combination of components, the contactor can either be open (o) or closed (c), and the generator can either be healthy (h), unhealthy (u), or offline (o). Thus, the set Ω has six possible states, represented as a tuple of contactor status and generator health: $x_1 = (c, h)$, $x_2 = (c, u)$, $x_3 = (c, o)$, $x_4 = (o, h)$, $x_5 = (o, u)$, and $x_6 = (o, o)$. The “generator” metacomponent, however, has three possible states, corresponding to healthy, unhealthy, and offline: $\tilde{x}_1 = h$, $\tilde{x}_2 = u$, and $\tilde{x}_3 = o$. These metacomponent states can be mapped back to the corresponding original components, such that $\tilde{x}_1 = \{x_1\}$, $\tilde{x}_2 = \{x_2\}$, and $\tilde{x}_3 = \{x_3, x_4, x_5, x_6\}$.

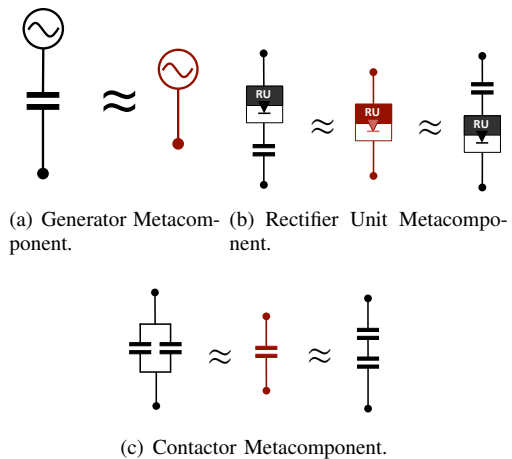


Fig. 2. Metacomponents used for abstraction. In terms of possible external behaviors (i.e., what can be measured from the external ports), two-component circuit units (shown in black) are equivalent to the single component units (shown in red).

V. EXAMPLES

To assess the performance of the greedy strategy, we have systematically tested the dynamic estimator described in Algorithm 1 on multiple circuits typical of those found in electric power systems. For these experiments we have taken a uniform probability distribution over Ω , the reward function f defined in Eq. (4) becoming the size of the feasible set.

In many cases it is not possible to completely eliminate the uncertainty on the state of the system when there is a limited number of sensors. In order to evaluate the performance of the greedy strategy, we compare it with a brute force strategy, which exhaustively tries every action $v \in \mathcal{V}$. Hence, no strategy can gather more information than the brute force strategy. Although the brute force strategy is not practically applicable, as $|\mathcal{V}|$ can be very large, it gives an upper bound on achievable performance, and can be used as a benchmark. Overall test methodology is summarized in Algorithm 2 and has been performed for every or some of the initial configurations $v_0 \in \mathcal{V}$ of the uncontrolled contactors.

Algorithm 2 Test methodology

Input: Initial configuration of the controlled contactors $v_0 \in \mathcal{V}$

- 1: **for** $x_0 \in \Omega$ **do**
- 2: Set the circuit in the state x_0
- 3: Put the controlled contactors in the configuration v
- 4: Run the strategy tested (Greedy or brute force strategy)
- 5: Record the computation time and the value of f at the end for statistics.
- 6: **end for**

A. Tests on a small-size circuit

The test small-size circuit, shown in Fig. 1 is comprised of 12 components. Six components are unknown ($G1, G2, R1, R2, C2, C5$), four contactors are controlled ($C1, C3, C4, C6$), depicted in blue in Fig. 1, and two voltage sensors are available ($S1$ and $S2$). The size of the state-space generated, taking into account the assumptions on faults, is 1600. A more precise description of the actual hardware circuit can be found in [18]. On this particular example there are four controlled contactors, so the brute force strategy performs $|\mathcal{V}| = 2^4 = 16$ actions. Both strategies have been run on the same Intel® Core™ i3-2310M 64 bits CPU 2.10 GHz, 6.00 Gb RAM.

Simulation result: The greedy strategy with a horizon length of $k = 6$ performs as well as the brute force strategy, i.e., the value of the objective function f at the end of the 6 steps using the greedy strategy is the same as after the brute force strategy with 16 steps.

1) *Average execution time:* The average execution time for the greedy strategy is shown in Fig. 3. It takes on the order of milliseconds to compute the next best action to perform. On the other hand, the offline computation takes 30 seconds for this circuit.

2) *Final value of the reward function f :* Using our metric of performance and comparing greedy and brute force strategies, Fig. 4 shows that the greedy strategy performs as well as the brute force. A point at coordinates (n, m)




Fig. 3. Histogram of execution time for the greedy strategy.

signifies that in $m\%$ of the cases, there are n or fewer indistinguishable states after the strategy (greedy or brute force) stops. Starting from 1600 possible states, the greedy strategy reduces the number of candidates to less than 20 in 100% of the cases. In approximately half of the cases, there are four states or fewer that are still indistinguishable after the $k = 6$ steps.

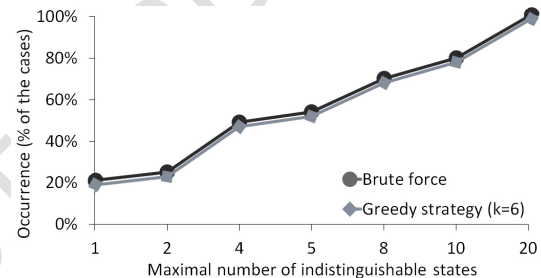


Fig. 4. Performance comparison between greedy and brute-force strategies.

While the greedy strategy shows potential, real-world problems are bigger and more complex. To that end, the lossless abstraction process designed in Subsection IV-B proves useful for scaling up the strategy.

B. Tests on larger circuit

We also test the greedy strategy on a larger circuit representative of more-electric aircraft power distribution systems with multiple generators and demonstrate how abstraction can reduce the offline computation time. The circuit topology is shown in Fig. 5. Contactors controlled by the fault detection controller are depicted in blue.

Applying the lossless abstraction established in IV-B leads to a reduced circuit that eliminates four uncontrolled contactors. Comparing the offline computation for the full and reduced circuit, abstraction reduces the offline computation time from 3795 to 378 seconds, i.e., by a factor of 10.

We have tested the performance of the greedy strategy on this circuit with the methodology described in Algorithm 2. The results obtained on this test set by the greedy strategy (with $k = 6$ actions) were again equivalent to the results for the brute-force strategy (with $|\mathcal{V}| = 32$ actions).

VI. CONCLUSIONS AND FUTURE DIRECTIONS

We proposed a greedy strategy for dynamic state estimation and fault detection in aircraft electric power systems.

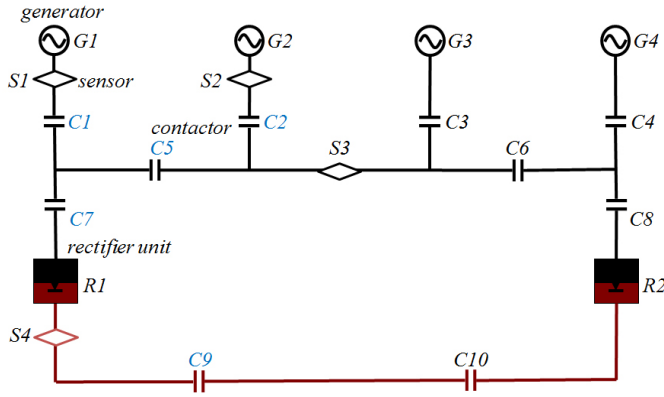


Fig. 5. A single-line diagram of a larger circuit with AC and DC components.

The proposed strategy was shown to have theoretical worst-case performance guarantees. Moreover, the practical performance of the strategy is well above the theoretical lower bound as demonstrated by examples.

The output of the overall estimation process is the set of all states consistent with observations, or the unique feasible state if possible. Future work will integrate these results with partial information games and synthesis of reactive control protocols. We are also interested in incorporating safety requirements while performing estimation. Although one can incorporate such information to disallow potentially unsafe actions, the main difficulty here is to obtain performance bounds when the set of available actions are changing in time. Currently, placement of sensors on the circuit topology is a given. By changing the number and locations of sensors, however, it may be possible to improve state estimation performance. Exploring this design space and trade-offs therein are likewise topics for future work.

APPENDIX

A. Background results in submodularity

We give some definitions and results on adaptive submodularity that follow the exposition provided in [6] and [11]. Notations used here are defined in Subsections II-A and II-B.

Definition 1: Given an objective function f , an action $v \in \mathcal{V}$, and a partial realization ψ_t , $\Delta(v|\psi_t)$ is the *conditional expected marginal benefit* of v conditioned on having observed ψ_t , defined as

$$\Delta(v|\psi_t) \doteq \mathbb{E}[f(v_{0:t} \cup \{v\}, X) - f(v_{0:t}, X) | \psi_t],$$

and the expectation taken with respect to $\mathbb{P}[x|\psi_t]$.

Definition 2: The function $f : 2^{\mathcal{V}} \times \mathcal{X} \times \Omega \rightarrow \mathbb{R}_+$ is *adaptive monotone* with respect to distribution $\mathbb{P}[x]$ if the conditional expected marginal benefit of any action is non-negative. Thus, for all $v \in \mathcal{V}$ and ψ_t with $\mathbb{P}[\psi_t] > 0$,

$$\Delta(v|\psi_t) \geq 0.$$

Definition 3: The function $f : 2^{\mathcal{V}} \times \mathcal{X} \times \Omega \rightarrow \mathbb{R}_+$ is *adaptive submodular* with respect to distribution $\mathbb{P}[x]$ if the

conditional expected marginal benefit of any fixed action v does not increase as more actions are performed and measurements are taken. Thus, f is adaptive submodular with respect to distribution $\mathbb{P}[x]$ if for all $\psi_t, \psi_{t'}$ such that ψ_t is a subrealization of $\psi_{t'}$, and for all $v \in \mathcal{V} \setminus \{v_0, \dots, v_{t'}\}$,

$$\Delta(v|\psi_t) \geq \Delta(v|\psi_{t'}).$$

The adaptive greedy algorithm, a generalization of the greedy algorithm [11], is a strategy that selects the action maximizing the conditional expected marginal benefit, conditioned on outcomes from all previous actions.

Theorem 2 (Theorem 1.14 in [11]): Let π_l^{greedy} be a greedy strategy run for l iterations (so that it selects l actions). Let π_k^* be any policy selecting at most k actions for any realization x . Then,

$$f_{\text{avg}}(\pi_l^{\text{greedy}}) \geq \left(1 - e^{-l/k}\right) f_{\text{avg}}(\pi_k^*),$$

where $f_{\text{avg}}(\pi) \doteq \mathbb{E}[f(\tilde{V}(\pi, X), X)]$ is the expected reward of π .

In particular, by setting $k = l$ we see that the greedy strategy selecting k items step by step obtains at least $(1 - 1/e)$ of the value of the optimal strategy that selects k items step by step.

B. Proof of proposition 2

We first state a lemma that will be useful in the proof.

Lemma 1: The function $b : \mathbb{R}^Y \rightarrow \mathbb{R}$, defined as

$$b(\tau_1, \tau_2, \dots, \tau_Y) = \sum_{i=1}^Y \tau_i - \frac{\sum_{i=1}^Y \tau_i^2}{\sum_{i=1}^Y \tau_i}, \quad (13)$$

is increasing on the positive orthant, i.e., $b(\tau_1, \tau_2, \dots, \tau_Y) \geq b(s_1, s_2, \dots, s_Y)$ if $\tau_i \geq s_i \geq 0$ for all $1 \leq i \leq Y$.

Proof: Note that because b is symmetric, i.e., permutation invariant with respect to its arguments, it is enough to show that it is increasing in one of its arguments. Let $k_1 \doteq \sum_{i=2}^Y \tau_i$ and $k_2 \doteq \sum_{i=2}^Y \tau_i^2$. Define $\tilde{b}(x) \doteq b(x, \tau_2, \dots, \tau_Y) = k_1 + x - \frac{k_2 + x^2}{k_1 + x}$. The partial derivative of b with respect to state x is $\partial \tilde{b} / \partial x = \frac{k_1^2 + k_2}{(k_1 + x)^2}$, which is non-negative by definitions of k_1 and k_2 . ■

Now, we are ready to prove Proposition 2.

Consider two partial realizations ψ_t and $\psi_{t'}$ s.t. $\psi_t \subseteq \psi_{t'}$ and the corresponding sets S_t and $S_{t'}$. Fix an action $v \in \mathcal{V} \setminus v_{0:t'}$. To prove adaptive submodularity, $\Delta(v, \psi_t)$ can be expressed as a function dependent on the size of S_t . We examine the variation of Δ between S_t and $S_{t'}$.

Since the probability measure is non-uniform and can take values in some set $\{p_1, \dots, p_N\}$, we define the subsets of Ω where $\mathbb{P}[x]$ is constant: $F_n = \{x \in \Omega \mid \mathbb{P}[x] = p_n\}$ for $n \in \{1, \dots, N\}$. The collection $F_{1:N}$ is trivially a partition of Ω . It is possible to show that the sets $\{D(y, v^*) \cap F_n \mid y \in \mathcal{Y}, n \in 1 : N\}$ form a partition of Ω and thus a partition of S_t .

Let $\alpha_{n,y} \doteq S_t \cap D(y, v) \cap F_n$. Then for all $x \in \alpha_{n,y}$, we have

$$\mu(v, x) = y \text{ and } \mathbb{P}[x] = p_n. \quad (14)$$

By Eq. (8), we get a new expression for $\mathbb{P}[\psi_t]$:

$$\mathbb{P}[\psi_t] = \sum_{x \in S_t} \mathbb{P}[x] = \sum_{y \in \mathcal{Y}} \sum_{n \in 1:N} p_n |S_t \cap D(y, v) \cap F_n|. \quad (15)$$

Let $\tau_y \doteq \sum_{n \in 1:N} p_n |\alpha_{n,y}|$. Then, conditional probabilities on F_n can be rewritten as

$$\forall x \in F_n, \mathbb{P}[x | \psi_t] = \frac{p_n}{\sum_{y \in \mathcal{Y}} \tau_y}. \quad (16)$$

We then separately compute the two terms in Eq. (10). First term becomes:

$$\mathbb{E}[f(v_{0:t}, X) | \psi_t] = \sum_{x_0 \in S_t} \mathbb{P}[x_0 | \psi_t] \sum_{x \in h(v_{0:t}, x_0)} \mathbb{P}[x]. \quad (17)$$

For $x_0 \in S_t$, $h(v_{0:t}, x_0) = S_t$, we obtain

$$\mathbb{E}[f(v_{0:t}, X) | \psi_t] = \sum_{y \in \mathcal{Y}} \tau_y. \quad (18)$$

For the second term in Eq. (10), we first get

$$\begin{aligned} f(v_{0:t} \cup \{v\}, x) &= \sum_{\tilde{x} \in h(v_{0:t}, x) \cap D(\mu(v, x), x)} \mathbb{P}[\tilde{x}] \\ &= \tau_{\mu(v, x)}. \end{aligned}$$

From Eq. (14) and Eq. (16), we obtain:

$$\begin{aligned} \mathbb{E}[f(v_{0:t} \cup \{v\}, X) | \psi_t] &= \sum_{x \in S_t} f(v_{0:t} \cup \{v\}, x) \mathbb{P}[x | \psi_t] \\ &= \sum_{n \in 1:N} \sum_{y \in \mathcal{Y}} \sum_{x \in \alpha_{n,y}} \tau_y \frac{p_n}{\sum_{z \in \mathcal{Y}} \tau_z} \\ &= \sum_{y \in \mathcal{Y}} \frac{\tau_y}{\sum_{z \in \mathcal{Y}} \tau_z} \sum_{n \in 1:N} p_n |\alpha_{n,y}| \\ &= \sum_{y \in \mathcal{Y}} \frac{\tau_y^2}{\sum_{z \in \mathcal{Y}} \tau_z}. \end{aligned}$$

Finally, putting the two terms of Eq. (10) leads to

$$\Delta(v | \psi_t) = b(\tau_1, \tau_2, \dots, \tau_Y) = \sum_{i=1}^Y \tau_i - \frac{\sum_{i=1}^Y \tau_i^2}{\sum_{i=1}^Y \tau_i}, \quad (19)$$

where $Y \doteq |\mathcal{Y}|$.

This expression of $\Delta(v | \psi_t)$ in terms of the variables τ_i is similar for the partial realization $\psi_{t'}$; the only change is the set S_t , which is represented in the function b by a different value of the τ_i denoted τ'_i . Since $\psi_t \subseteq \psi_{t'}$ and $S_{t'} \subseteq S_t$, τ_i and τ'_i satisfy $\tau'_i \leq \tau_i$ for all i .

Therefore, adaptive submodularity is equivalent to showing that b is increasing on the positive orthant, and Lemma 1 concludes the proof.

ACKNOWLEDGMENTS

The authors would like to thank Jean-Michel Maillet and Eric Wolff for very useful and enlightening discussions. This work was supported in part by IBM and UTC through the iCyPhy consortium and TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

REFERENCES

- [1] A. Abur and A.G. Exposito. *Power system state estimation: theory and implementation*, volume 24. CRC, 2004.
- [2] J. Bang-Jensen, G. Gutin, and A. Yeo. When the greedy algorithm fails. *Discrete Optimization*, 1(2):121–127, 2004.
- [3] L. Blackmore, S. Rajamanoharan, and B.C. Williams. Active estimation for jump markov linear systems. *Automatic Control, IEEE Transactions on*, 53(10):2223–2236, 2008.
- [4] A. Bose and K.A. Clements. Real-time modeling of power networks. *Proc. of the IEEE*, 75(12):1607–1622, 1987.
- [5] E.M. Davidson, S.D.J. McArthur, and J.R. McDonald. A toolset for applying model-based reasoning techniques to diagnostics for power systems protection. *Power Systems, IEEE Transactions on*, 18(2):680 – 687, may 2003.
- [6] D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42(1):427–486, 2011.
- [7] A.J. Gonzalez, R.A. Morris, F.D. McKenzie, D.J. Carreira, and B.K. Gann. Model-based, real-time control of electrical power systems. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 26(4):470–482, 1996.
- [8] D. Gorinevsky, S. Boyd, and S. Poll. Estimation of faults in dc electrical power system. In *American Control Conference*, pages 4334–4339. IEEE, 2009.
- [9] A. Guillory and J. Bilmes. Average-case active learning with costs. In *Algorithmic Learning Theory*, pages 141–155. Springer, 2009.
- [10] K. Keller, K. Swearingen, J. Sheahan, M. Bailey, J. Dunsdon, K.W. Przytula, and B. Jordan. Aircraft electrical power systems prognostics and health management. In *Aerospace Conference, 2006 IEEE*, pages 12–pp. IEEE, 2006.
- [11] A. Krause and D. Golovin. Submodular function maximization. In *Tractability: Practical Approaches to Hard Problems (to appear)*. Cambridge University Press, 2012.
- [12] L. Liu, K.P. Logan, D.A. Cartes, and S.K. Srivastava. Fault detection, diagnostics, and prognostics: software agent solutions. *Vehicle Technology, IEEE Transactions on*, 56(4):1613–1622, 2007.
- [13] L. Mihaylova, T. Lefebvre, H. Bruyninckx, K. Gadeyne, and J. De Schutter. A comparison of decision making criteria and optimization methods for active robotic sensing. *Numerical Methods and Applications*, pages 316–324, 2003.
- [14] I. Moir and A. Seabridge. *Aircraft Systems: Mechanical, Electrical and Avionics Subsystems Integration*. AIAA Ed. Series, 2001.
- [15] A. Monticelli. Electric power system state estimation. *Proc. of the IEEE*, 88(2):262–282, 2000.
- [16] M. Naghshvar and T. Javidi. Active sequential hypothesis testing. *CoRR*, abs/1203.4626, 2012.
- [17] N. Ozay, U. Topcu, and R. M. Murray. Distributed power allocation for vehicle management systems. In *Proc. IEEE Conference on Decision and Control and European Control Conference*, pages 4841–4848, 2011.
- [18] R. Rogersten, H. Xu, N. Ozay, U. Topcu, and R. M. Murray. An aircraft electric power testbed for validating automatically synthesized reactive control protocols. In *Proc. of the 16th international conference on Hybrid systems: computation and control, HSCC '13*, 2013.
- [19] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D.C. Teneketzis. Failure diagnosis using discrete-event models. *Control Systems Technology, IEEE Transactions on*, 4(2):105–124, 1996.
- [20] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin. *Efficient planning of informative paths for multiple robots*. Carnegie Mellon University, School of Computer Science, Machine Learning Department, 2006.
- [21] T. Wongpiromsarn, U. Topcu, and R. M. Murray. Formal synthesis of embedded control software for vehicle management systems. In *AIAA Infotech@Aerospace*, 2011.
- [22] H. Xu, U. Topcu, and R. M. Murray. A case study on reactive protocols for aircraft electric power distribution. In *Proc. IEEE Conference on Decision and Control*, 2012.