

International Conference on Computational Science, ICCS 2012

Dynamic Data Driven Methods for Self-aware Aerospace Vehicles

D. Allaire^a, G. Biros^c, J. Chambers^b, O. Ghattas^c, D. Kordonowy^b, K. Willcox^a

^aDepartment of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA

^bAurora Flight Sciences, Cambridge, MA

^cInstitute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX

Abstract

A self-aware aerospace vehicle can dynamically adapt the way it performs missions by gathering information about itself and its surroundings and responding intelligently. Achieving this DDDAS paradigm enables a revolutionary new generation of self-aware aerospace vehicles that can perform missions that are impossible using current design, flight, and mission planning paradigms. To make self-aware aerospace vehicles a reality, fundamentally new algorithms are needed that drive decision-making through dynamic response to uncertain data, while incorporating information from multiple modeling sources and multiple sensor fidelities.

In this work, the specific challenge of a vehicle that can dynamically and autonomously sense, plan, and act is considered. The challenge is to achieve each of these tasks in real time—executing online models and exploiting dynamic data streams—while also accounting for uncertainty. We employ a *multifidelity* approach to inference, prediction and planning—an approach that incorporates information from multiple modeling sources, multiple sensor data sources, and multiple fidelities.

Keywords:

DDDAS, multifidelity, statistical inference

1. Introduction

Seeing, feeling, and thinking aerospace vehicles can adapt the way they perform missions by gathering information about themselves and their surroundings and responding intelligently. The vehicle will respond to events and to degradation over time in the same way as a self-aware organism—sprinting when healthy and under favorable conditions, and slowing down as it ages and degrades. The ability to sense system anomalies allows a vehicle to rely more heavily on healthy systems to complete missions. This self-aware aerospace vehicle can perform missions that are impossible using current design, flight, and mission planning paradigms.

A self-aware aerospace vehicle has a number of benefits. The ability to sense will allow decisions to be made mid-mission. For example, a launch vehicle approaching maximum dynamic pressure can use information regarding current system health to compensate underperforming systems with healthier systems. Under off-nominal conditions or health states, modifications to the mission or flight planning can be implemented to perform the mission based on responsive adaptation by the vehicle to its state. For an aging aircraft that must operate for twenty, thirty, or even fifty

Email address: dallaire@mit.edu (D. Allaire)

years, the ability to tailor or restructure its everyday flight to minimize wear, fatigue, or environmental degradation adds years to its life and reduces the maintenance required to keep the aircraft flight-worthy. Additionally, the aircraft will fly to its capability, performing missions beyond its traditional design envelope. Figure 1 shows the results of an initial study to estimate the benefit of a condition-aware airframe [1]. It was found that the airframe could operate at 130% of the designed performance at the beginning of its life and could extend its lifetime by 400% without changing any other feature. This was achieved by replacing traditional damage tolerant design with real-time health and capability assessment of the airframe. This capability extended to the entire air vehicle could revolutionize the performance of aerospace vehicles.

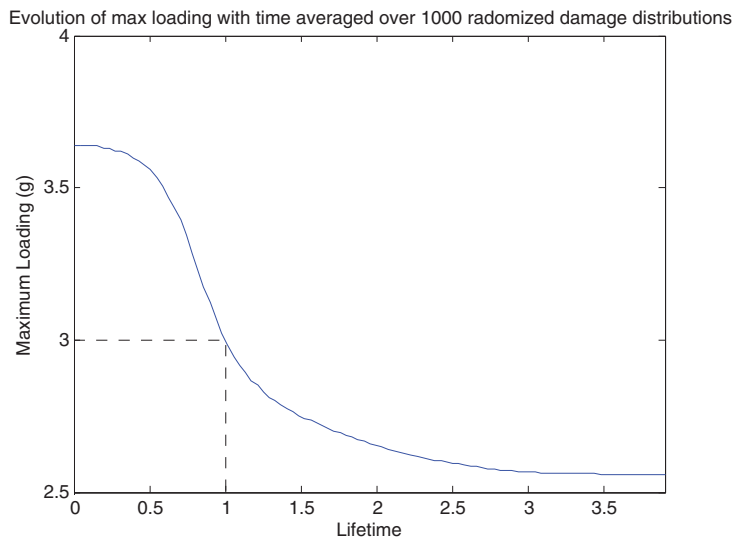


Figure 1: Condition-Aware Capability. The potential benefit of condition-awareness in terms of maneuver loading over the vehicle lifetime for an airframe is shown in blue compared with the traditional design shown by the black box [1].

During the 1990's, calls were made for intelligent flight [2] and collaborative artificial intelligence [3] to create systems that were naturally able to emulate human decision-making to ease the interaction of operators from autonomous tasks. However, while complex control systems have been created to allow operation of highly dynamic aerospace vehicles, and vehicle health management architectures have been developed to increase their reliability and safety, the vision of an intelligent vehicle that can sense, plan, and act [4] remains beyond current systems or architectures. Current vehicle health management architectures developed by NASA [5] and the Department of Defense [6] determine the health state of a vehicle. However, the architectures remain uncoupled from the potential capability of the vehicle to make on-the-fly decisions using harvested information; instead these systems are used to increase mission reliability and vehicle availability only. The fundamental flight laws of the aerospace vehicle cannot be modified by current systems, which are instead set a priori during the design. This design paradigm, where sensory input checks, but does not form the basis for capability, is suboptimal. Instead, vehicle performance is limited at all times according to extreme environments and physical degradation in order to avoid this difficult yet rewarding dynamic data-driven problem. Unlike such approaches, our approach not only senses vehicle state, but also couples it to vehicle performance parameters.

In the next section, we describe an adaptive structural response model, adapted using vehicle-specific dynamic data. In Section 3 we describe development of a resource allocation procedure for dynamic online management of multifidelity structural response models and sensor data. We conclude in Section 4 with a discussion of ongoing and future work.

2. Adaptive Structural Response Model

An adaptive structural response model provides a representation of the physical aerospace vehicle that changes as the vehicle degrades and predicts the capability of the vehicle to perform actions. Figure 2 depicts the components of the model. In order to adapt the model based on real-time degradation, various types and quantities of sensors are placed on the vehicle that continuously monitor aspects of its health, loading, and environment. These sensors include distributed strain sensors to measure local strain/stress fields throughout the airframe, structural health monitoring sensors to measure changes in the structure, temperature sensors to determine environmental temperatures that can alter the structural stiffness and strength of the vehicle, distributed pressure sensors to measure local loading on the structure, and inertial sensors to measure body accelerations and rotations that induce loading in the structure. The locations of these sensors are determined by physical access constraints and the criticality of the location to the overall structure capability to all actions that the vehicle would perform to complete a mission. The rates at which these sensed variables are measured depends not only upon the physical characteristics of the sensor and storage space within the data acquisition unit onboard the vehicle, but also on the rate that the variable will change and the impact of that variable rate of change on the overall capability. Because the overall capability is determined by the sum of localized capabilities, sensor measurement rates must also take into consideration the criticality of the measurement location. Therefore, the measurement rates can be determined by the desired fidelity of overall structural capability for any given action.

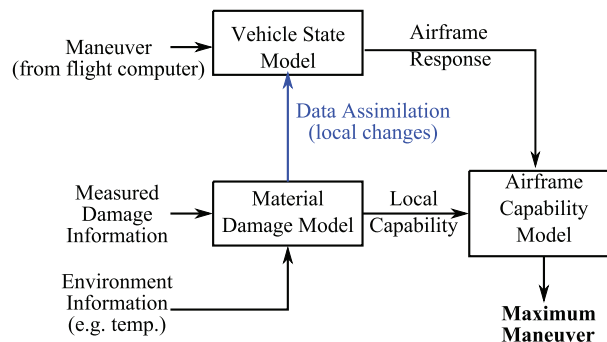


Figure 2: Adaptive structural response model is updated as the vehicle degrades and predicts the capability of the vehicle to perform actions.

The various sensors inform a model of the vehicle that is adapted through our multifidelity inference methods. The prediction and planning modules interrogate the updated model to predict the response of the vehicle, which can be divided between static response (i.e., the ability of the vehicle to survive the performance of an action) and the dynamic response (i.e., the change in the vehicle geometry through the action). The structural response model is based on a finite element modeling approach. We derive projection-based reduced-order models for rapid online estimation of structural response. In the context of the finite element model, the static response is the stress or strain at elements compared with their corresponding strength, and the dynamic response is the stiffness and mass of the elements, which can change with degradation. The sensor data described then adapts that model based on degrading events, which are monitored.

The multifidelity nature of our DDDAS process enters through our use of a variety of structural response models and a variety of update frequencies. During times when resources are low, modifications to flight limits are based on low-fidelity modeling. In times when more resources are available, more detailed analysis is performed to more precisely assess the impact of damage on material limits and overall structural/vehicle state. For example, during the offline phase (e.g., on the ground between flights), we use our highest fidelity structural models (finite element models). These high-fidelity models are updated with damage information that has occurred during the previous flight. During the online (flight) phase, we use reduced-order models that permit rapid execution. Vehicle state information is used to update the low-fidelity models with time-sensitive inference methods. The reduced-order model modifies estimated flight limits based on the dynamically sensed health of structure.

Open research challenges include the development of multifidelity structural response models, including an understanding of how to incorporate damage effects in the lower-fidelity models. Uncertainties in the sensed health of the structure must also be accounted for. We will approach these challenges by considering a finite element model coupled to data- and physics-based failure models. High fidelity models will use the historical sensor data captured during previous missions to update the vehicle specific structural model. Such models will reanalyze the structure and update material allowables. The modified allowables will be uploaded to the on-board reduced-order models. These models will utilize real-time sensor data to calculate the instantaneous structural state of the vehicle and update the flight envelopes, accounting for system uncertainties (e.g., possible false sensor readings, conflicting sensor information, etc.). The up-to-the-second flight envelope can permit a central avionics computer to access a database of safely achievable maneuvers in order to carry out the assigned mission.

3. DDDAS Methods

The distinguishing feature of our proposed research is a multifidelity approach to achieving the DDDAS paradigm. By “multifidelity” we mean an approach that draws upon information from multiple modeling options and multiple sensor data sources. These modeling and sensing options each have different levels of fidelity and different resource requirements. For example, a low-fidelity model may give rapid but low-confidence estimates that can be used to provide some indication of a load limit. This estimate could be refined by higher-fidelity models and additional data in order to provide a more confident prediction before the decision must be executed.

At the core of our evolving DDDAS process is a decision-making hub that controls model execution, data gathering, and model adaptation to support the inference, prediction and planning tasks. Our research is developing the resource allocation procedure that supports this decision-making. The decisions to be made include: when and what to measure from sensors, when and which structural response models to use, and what are the current quantities of interest. For example, if an event has occurred that threatens the integrity of the wing just prior to a pull-up maneuver, the load limit of the wing is a quantity of interest. Depending on the time until the maneuver and available computational resources (both of which may be uncertain), the resource allocation procedure makes decisions on how to gather information from models and data most efficiently to determine the load limit with an acceptable level of confidence prior to the maneuver.

More formally, consider the problem of estimating a quantity of interest q at a given set of conditions \mathbf{x} , with a specified level of confidence and within a specified amount of time. For this, assume that we have a set of models $\mathcal{M} = \{M_1, \dots, M_k\}$ we may use to estimate the quantity of interest, and that we will only choose one model to make the estimate. These models are multifidelity in the sense that some may be better than others at estimating the quantity of interest. In general, we may measure our confidence in our estimate of q by some dispersion measure $\mathcal{D}(Q(\mathbf{x}))$, where Q represents the probability distribution of q . To note the dependence on the model used to estimate the quantity of interest, we write the dispersion measure as $\mathcal{D}(Q_{M_i}(\mathbf{x}))$, where $i \in \{1, \dots, k\}$. The smaller the dispersion measure, the higher we consider the fidelity (with respect to the quantity of interest) of the model used to estimate it. Assuming there is some amount of time \mathcal{T} and some amount of available computational resources \mathcal{R} , we would like to exercise the highest fidelity model that does not require more than the available computational resources or take longer than the allowable time to execute. Formally, we wish to solve the following:

$$M^* = \arg \min_{M \in \mathcal{M}} \mathcal{D}(Q_M(\mathbf{x}))$$

$$\text{subject to } t_M \leq \mathcal{T}, r_M \leq \mathcal{R},$$

where t_M and r_M are respectively, the time and computational resources required to execute model M . The M^* we find, and how we find it, in general depends on the dispersion measure we use.

For the case of a real system, the problem described above is complicated by the fact that execution times, resources required, and input configurations are uncertain at the time of such decisions. Thus, our general resource allocation procedure is based on Bayesian statistical decision theory, where minimization of expected loss is the key driver in the decision-making. Let \mathbf{X} be the uncertain input configuration at which we wish to know a quantity of interest q . Also, let R_M and T_M be the uncertain amount of computational resources and time required to execute model M , and define $\theta = [\mathbf{X}, R_M, T_M]$. Let a be some action we may take to estimate q , such as using a model from

the set \mathcal{M} as before, using some combination of these models, using experimental data, etc. Let the set of possible actions be denoted as \mathcal{A} . We again measure our confidence with some dispersion measure $\mathcal{D}(Q(\mathbf{x}))$, but now we have to consider that we are not certain about \mathbf{x} . Thus, we formulate a loss function, taking into account the variability of $\mathcal{D}(Q_M(\mathbf{X}))$, as well as that of R_M and T_M , which we denote as $\mathcal{L}(\theta, a)$. This loss function must take into account losses associated with poor estimates of q as well as violations of the time and resource constraints given in the deterministic formulation and the subsequent losses that may occur as a result. The loss function is thus problem dependent. If $\pi(\theta)$ represents the distribution of θ at the time of decision making, then our goal is to identify a Bayes action $a^\pi \in \mathcal{A}$ such that $a^\pi = \min \mathbb{E}^\pi[\mathcal{L}(\theta, a)]$.

The decision-making problem incorporates global sensitivity analysis to identify how much information may be obtained from a given resource (be it sensor data, high-fidelity models, low-fidelity models, or some combination thereof). The combination of estimated information, computational resources required, and time required for a particular set of modeling or sensing resources factors into the expected loss calculation. Depending on the time available prior to requiring an estimate, different procedures will be employed by the procedure in an attempt to quickly identify actions that are close to optimal. For example, we might use a greedy approach, where a subset of resource options (perhaps based on historical data) are selected to be analyzed and the best is chosen to estimate the load limit. Another possibility is a multifidelity approach to the decision problem, where reflex-like decisions use low-fidelity models/data with low confidence to provide some indication of the load limit (which at times may be all that is required), with subsequent updates by higher fidelity information up to the time permitted to make an estimate.

4. Future Work

Our future research will develop new multifidelity approaches to perform the methodological components of the DDDAS process: inference to estimate vehicle state, prediction of flight limits, and planning to achieve mission objectives. While methods exist to achieve each of these tasks in a deterministic offline setting enabled by high performance computing, we need fundamentally new approaches to move to the dynamic data driven setting while also accounting for the effects of uncertainty. We will develop multifidelity approaches to state estimation that fuse information from multiple data sources and multiple models. We will use a reduced-order modeling framework to achieve inference in real time, together with higher fidelity models in the offline phase. We are also developing machine learning techniques that are scalable on GPU and manycore architectures, such as a new approach for nearest-neighbor interpolation in very high dimensions. We will also develop new methods for rapid online updating of models using dynamic data. Our approach uses a stochastic process model to capture model inadequacy. Dynamic data are used to update the mean and covariance of the stochastic process in a sequential manner using fast online algorithms. This updating will be tied to the resource allocation process to determine what time and computational resources are available for the update.

References

- [1] D. Kordonowy, O. Toupet, Composite Airframe Condition-Aware Maneuverability and Survivability for Unmanned Aerial Vehicles, AIAA-2011-1496, presented at Infotech@Aerospace, St. Louis, Missouri, March 29–31 (2011).
- [2] R. Stengel, Toward intelligent flight control, IEEE Transactions on Systems, Management, and Cybernetics 23 (6) (1993) 1699–1717, doi:10.1109/21.257764.
- [3] B. Grosz, Collaborative systems, American Association for Artificial Intelligence, 0738-4602-1996, pp. 67–85 (1996).
- [4] H. Asada, J.-J. Slotine, Robot Analysis and Control, Wiley, 1986.
- [5] A. Srivastava, R. Mah, M. Claudia, Integrated vehicle health management, Technical Plan, Version 2.03, NASA Aeronautics Research Mission Directorate Aviation Safety Program (2009).
- [6] B. Glass, Integrated system vehicle health management (ISVHM) architecture design: ISHM impact study and architecture design, AFRL-RB-WP-TR-2008-3148.