Optimal Approximations of Coupling in Multidisciplinary Models

Ricardo Baptista*, Youssef Marzouk†, Karen Willcox‡,
Massachusetts Institute of Technology, Cambridge, MA

and Benjamin Peherstorfer§
University of Wisconsin-Madison, Madison, WI

This paper presents a methodology for identifying important discipline couplings in multicomponent engineering systems. Coupling among disciplines contributes significantly to the computational cost of analyzing a system, and can become particularly burdensome when coupled analyses are embedded within a design or optimization loop. In many cases, disciplines may be weakly coupled, so that some of the coupling or interaction terms can be neglected without significantly impacting the accuracy of the system output. Typical practice derives such approximations in an ad hoc manner using expert opinion and domain experience. This work proposes a new approach that formulates an optimization problem to find a model that optimally balances accuracy of the model outputs with the sparsity of the discipline couplings. An adaptive sequential Monte Carlo sampling-based technique is used to efficiently search the combinatorial model space of different discipline couplings. An algorithm for selecting an optimal model is presented and illustrated in a fire detection satellite model and a turbine engine cycle analysis model.

Nomenclature

A_{sa}	Satellite: Area of the solar array	d	Number of discipline couplings
BPR	Turbine: Bypass ratio	D_{KL}	Kullback-Leibler divergence
C_d	Satellite: Drag coefficient	$\mathbf{f}$	Model output variables (QoI)
CPR	Turbine: Compressor pressure ratio	$\mathcal{F}$	Function for output variables

*Graduate Student, Center for Computational Engineering, rsb@mit.edu, AIAA Student Member
†Associate Professor, Department of Aeronautics and Astronautics, ymarz@mit.edu, AIAA Associate Fellow
‡Professor, Department of Aeronautics and Astronautics, kwillcox@mit.edu, AIAA Associate Fellow
§Assistant Professor, Department of Mechanical Engineering, peherstorfer@wisc.edu

$\tilde{\mathbf{f}}$	Linear approximation of model outputs
F_n	Turbine: Net thrust
F_s	Satellite: Average solar flux
FPR	Turbine: Fan pressure ratio
H	Satellite: Altitude
$h(M)$	Model objective function
I	Satellite: Moment of inertia
K	Number of disciplines
L	Number of particles in SMC
L_a	Satellite: Aerodynamic torque moment arm
L_{sp}	Satellite: Radiation torque moment arm
M	Discipline couplings in the decoupled model
m	Number of input variables
$\mathcal{M}$	Model space
M_0	Discipline couplings in the reference model
n	Number of state variables
OPR	Turbine: Overall pressure ratio
p	Number of output variables
$\mathcal{P}(M)$	Function to count removed couplings
P_{ACS}	Satellite: Power of attitude control system
P_{other}	Satellite: Power other than P_{ACS}
P_{tot}	Satellite: Total power
$P_t(M)$	Probability of model M at step t
q	Satellite: Reflectance factor
$q_{\mathbf{B}}$	Proposal distribution to sample models
$\mathcal{R}_i$	Residual function for state variable i
R_D	Satellite: Residual dipole of spacecraft
$\mathcal{S}_k$	Input state variables to discipline k
$\Delta t_{eclipse}$	Satellite: Eclipse period
Δt_{orbit}	Satellite: Orbit period
T_4	Turbine: Engine burner temperature
$TSFC$	Turbine: Thrust-specific fuel consumption
v	Satellite: Satellite velocity
W	Turbine: Mass-flow rate
$w_t^{(l)}$	Weight of model l at step t
$\mathbf{x}$	Model input variables
$\mathbf{y}$	Model state variables
$\mathbf{y}(\mu_{\mathbf{x}})$	First-order mean of state variables

Symbols

α_t	Update to tempering parameter
δ	Threshold for minimum sample diversity
γ	Metropolis-Hastings acceptance probability
κ_t	Metropolis-Hastings transition kernel
Λ	Set of λ parameters
λ	Parameter in combinatorial optimization
$\mu_{\mathbf{x}}$	Mean of input variables
$\mu_{\tilde{\mathbf{f}}}$	Mean of linearized model output
$\pi_{\mathbf{f}}$	Probability density of outputs
$\pi_{\mathbf{x}}$	Probability density of inputs
$\rho(t)$	Tempering parameter at step t
$\Sigma_{\tilde{\mathbf{f}}}$	Covariance of linearized model output
τ_{tot}	Satellite: Total torque
θ	Satellite: Deviation of moment axis
θ_{slew}	Satellite: Maximum slewing angle
ζ_t	Particle diversity at step t

I. Introduction

Multidisciplinary analysis and optimization (MDAO) couples multiple computational models to represent complex interactions in the design of engineering systems. With the increasing number of disciplines and improved fidelity in multidisciplinary models, the coupling among disciplines can contribute significantly to the computational cost of analyzing these systems. This coupling may include both one-directional (feed-forward) coupling, and bi-directional (feedback) coupling that requires iterative numerical methods to compute model outputs. The coupled MDAO problem may be formulated in a variety of different ways [1], and various MDAO architectures have been developed to manage discipline coupling for large-scale problems (see Ref. [2] for an overview). Monolithic architectures solve the system using a single optimization problem, while distributed approaches partition the model into subproblems, each involving a smaller number of variables [3–7]. Typical practice derives these discipline couplings using expert opinion and domain experience. Given the significant impact of couplings on the computational tractability of evaluating a model, it is of interest to systematically identify *which* disciplines should be coupled in a model and where couplings may be approximated. This paper addresses this open challenge to yield *optimal approximations* to coupling in multidisciplinary models.

The application of MDAO originated in structural optimization and aircraft wing design [8, 9]. It has since been extended to many different engineering systems, such as the design of complete aircraft configurations [10, 11], internal combustion engines [12], wind turbines [13], and spacecraft [14]. Multidisciplinary models of such systems often demonstrate varying degrees of coupling. Couplings can be categorized as being "strong" or "weak" based on the response of a discipline output to a change in a coupling variable [15]. As a result, for certain quantities of interest (QoI) in a model, neglecting weak interactions negligibly impacts the accuracy of the system outputs. Simultaneously, decoupling discipline interactions can reduce the number of feedback loops and present substantial computational savings when using the model for analysis and/or design optimization. Therefore, an important challenge in the field of MDAO has been to identify a decoupled model that best trades off the sparsity of the discipline couplings with the accuracy of the model in representing the engineering system's outputs. One approach to do this uses weighted design structure matrices to decompose the model and rearrange the disciplines to minimize feedback. This method estimates the strength of the couplings based on the discipline connectivity [16] or the sensitivity of the model outputs to the coupling at each iteration of a multidisciplinary optimization process [17]. However, existing methods leave open the question of identifying which discipline couplings are most important in a model for characterizing system outputs over a range of input variables.

In this work, we consider optimal approximations to coupling in the context of multidisciplinary models with uncertain inputs. Using such models in optimization or uncertainty quantification (UQ) is challenging, particularly in the nonlinear setting; both types of analysis require an accurate characterization of the uncertainty in the model outputs. To this end, users often rely on traditional and computationally intensive approaches such as Monte Carlo sampling [18]. As a result, identifying an approximate coupling that reduces the dimensionality of the information exchanged between disciplines and that eliminates large feedback loops, while minimizing information loss in the output uncertainty, presents an important source of computational savings for MDAO. This paper addresses this challenge by formulating the search for a decoupled model as a combinatorial optimization problem and leveraging recent advances in sequential Monte Carlo algorithms [19, 20] to solve this problem. To the best of our knowledge, it is the first attempt in the literature to identify a sequence of optimal discipline couplings that trade off the accuracy of the model output uncertainty with the sparsity of the discipline couplings.

The remainder of this paper is organized as follows. Sections II and III introduce the mathematical formulation and methodology to efficiently search for an optimal approximation to the coupling of a multidisciplinary model. Section IV presents results for two engineering systems: a fire detection satellite and a turbine engine cycle analysis model. Finally Section V concludes the paper and presents future research directions.

II. Background and Problem Formulation

This section introduces the notation and background for multidisciplinary model analysis in Sections A and B. The definition and the problem formulation for model coupling approximations are presented in Sections C and D, respectively.

A. Multidisciplinary Models

We consider a numerical multidisciplinary model to be a mapping from disciplinary input variables to disciplinary output variables, mediated by a set of internal state variables. The output variables represent quantities of interest (QoI) that depend on the input and state variables of the model.

For a model in this work, we represent a vector of m input variables by $\mathbf{x} = [x_1, \ldots, x_m]^T \in \mathcal{X} \subseteq \mathbb{R}^m$, a vector of n state variables by $\mathbf{y} = [y_1, \ldots, y_n]^T \in \mathcal{Y} \subseteq \mathbb{R}^n$, and a vector of p output variables by $\mathbf{f} = [f_1, \ldots, f_p]^T \in \mathbb{R}^p$.

The state variables in the model are defined implicitly by the governing residual equations

$$\mathcal{R}_i(\mathbf{x}, \mathbf{y}) = 0, \quad i = 1, \ldots, n, \tag{1}$$

where in most generality each function $\mathcal{R}_i$ depends on all input and state variables, and consists of algebraic equations, differential equations, and/or other operations. Letting $\mathcal{R} = [\mathcal{R}_1, \ldots, \mathcal{R}_n]^T$, the coupled system of equations, $\mathcal{R}(\mathbf{x}, \mathbf{y}) = \mathbf{0}$, defines the mapping of input to state variables, $\mathcal{R}\colon \mathcal{X} \to \mathcal{Y}$. The function $\mathcal{F}\colon \mathcal{X} \times \mathcal{Y} \to \mathbb{R}^p$ computes the output variables explicitly as

$$\mathbf{f} = \mathcal{F}(\mathbf{x}, \mathbf{y}). \tag{2}$$

In a multidisciplinary problem, the n residual equations in equation (1) are partitioned into disjoint groups, each representing a particular engineering discipline or subsystem of the model. Each discipline computes a set of internal state variables and coupling variables that are passed to other disciplines for modeling system interactions. The coupling variables supplied by a given discipline are defined as a transformation of its internal state variables. In this work, we will include these transformations in the set of residual equations, $\mathcal{R}$, and include both the internal and coupling variables in the vector of state variables, $\mathbf{y}$. A graphical representation of a model with two disciplines and its coupling variables (y_1 and y_2) is given in Figure 1.

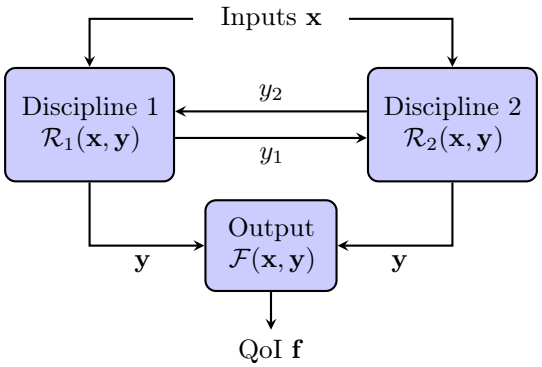

Fig. 1 Graphical representation of a multidisciplinary model

The dependence of the residual equations on coupling variables from other disciplines defines the overall discipline coupling in the model. This coupling can include one-directional (feed-forward) coupling, or bi-directional (feedback) coupling between any two disciplines. For instance, the dependence of disciplines 1 and 2 on y_2 and y_1, respectively, results in the feedback coupling loop for the model in Figure 1. Solving these systems of feedback coupled residual equations requires using an iterative numerical method to determine

the states and corresponding output variables. Common iterative methods include fixed point iteration or gradient-based algorithms for solving nonlinear equations, such as Newton's method.

B. Uncertainty in Multidisciplinary Models

In many multidisciplinary models, the input variables are not known exactly. As an example, when analyzing a rocket during launch, it is common to consider wind gusts as a source of uncertainty in the multidisciplinary model [21]. In such cases, the input variables, $\mathbf{x}$, are represented as random variables endowed with a probability density function (PDF) $\pi_{\mathbf{x}} \colon \mathcal{X} \to \mathbb{R}_+$.

With uncertain inputs, the state and output variables of the model are also random variables where the outputs, $\mathbf{f}$, have an induced joint PDF, $\pi_{\mathbf{f}} \colon \mathbb{R}^p \to \mathbb{R}_+$. Forward uncertainty quantification (UQ) propagates the input uncertainty through the coupled model to characterize the uncertainty in the output variables, which is defined by their distribution. For engineering applications, this typically includes determining the mean and variance of the output variables, although sometimes other properties of the output distribution are also of interest, including higher-order moments and the PDF itself.

For linear residual and output equations, the mean and variance in the output variables can be quantified analytically in terms of the model equations and the mean and variance of the inputs. However, for general nonlinear equations, forward UQ is more challenging and is commonly performed with various probabilistic techniques such as sampling-based methods, localized Taylor-series model expansions, functional approximations (e.g., polynomial chaos expansions), etc [22].

One of the traditional and robust algorithms for quantifying uncertainty is based on Monte Carlo simulation. This sampling-based method comprises three steps:

1. Draw N realizations, $\mathbf{x}^{(j)}$, of the input variables from $\pi_{\mathbf{x}}$ that define samples $j = 1, \ldots, N$.
2. For each sample, solve the coupled residual equations, $\mathcal{R}(\mathbf{x}^{(j)}, \mathbf{y}^{(j)}) = \mathbf{0}$, for $\mathbf{y}^{(j)}$ and evaluate the output $\mathbf{f}^{(j)} = \mathcal{F}(\mathbf{x}^{(j)}, \mathbf{y}^{(j)})$.
3. Use the output samples, $\mathbf{f}^{(j)}$, to estimate the mean, variance, or other attributes of the PDF, $\pi_{\mathbf{f}}$.

It is well known that the standard deviation of a Monte Carlo estimator converges at a rate of $1/\sqrt{N}$ with an increasing number of samples, N [22]. As a result, to accurately estimate output statistics, a large number of samples and corresponding solves of the coupled residual equations are required with this method. For expensive simulations under limited computational resources, this motivates the use of alternative techniques that reduce the computational cost of solving the multidisciplinary model. One approach for efficiently

propagating uncertainty in multidisciplinary systems is a decomposition-based method that combines Monte Carlo sampling of each discipline or domain with importance sampling to analyze the uncertainty in feed-forward systems [23]. For multidisciplinary systems with feedback, recent methods include a likelihood-based approach to decouple feedback loops and reduce the model to a feed-forward system [24], dimension reduction techniques to represent certain coupling variables with spectral expansions that depend on a small number of uncertain parameters [25–27], and the use of adaptive surrogate models for individual disciplines [28] or to approximate the coupling variables and reduce the number of model evaluations [29].

In this work, we focus on approximate coupling as a way to reduce the computational cost of finding a solution to the multidisciplinary model. We note that this could also be combined with existing techniques to accelerate forward uncertainty quantification.

C. Model Coupling Approximation

In practice, many disciplines are often weakly coupled and the residual equations in each discipline are most sensitive to only a subset of the coupling variables [15]. Our approach of approximate coupling is to trade off the number of number of coupling variables between disciplines with the accuracy in the output QoI.

In a multidisciplinary model with K disciplines, we define $\mathcal{S}_k \subseteq \{y_1, \ldots, y_n\}$ as the set of coupling variables that are arguments to the residual functions of the k-th discipline (e.g., $\mathcal{S}_1 = \{y_2\}$ and $\mathcal{S}_2 = \{y_1\}$ for the example in Figure 1). Collecting the coupling variables associated with each discipline, we define a *model* M as $M = [\mathcal{S}_1, \ldots, \mathcal{S}_K]$. In this context, the size of the model $d(M)$ is the total number of discipline couplings it contains, i.e., $d(M) = \sum_{k=1}^{K} |\mathcal{S}_k|$, where $|\mathcal{S}_k|$ denotes the cardinality of $\mathcal{S}_k$. Given a multidisciplinary model M, our goal is to identify the coupling variables (i.e., components y_i of $\mathbf{y}$) that have a small effect on the overall model output. Each of these variables as inputs to a certain discipline (e.g., $y_i \in \mathcal{S}_k$) is then fixed to a nominal value, such as its mean value, and removed from S_k. The corresponding disciplines are thereby only coupled through deterministic values for the coupling variables. By fixing these random variables, these coupling variables are no longer unknown discipline inputs and do not require iteration to find their converged values. We refer to this process as *decoupling*.

The operation of replacing the dependence of residual equation, $\mathcal{R}_2$, on the coupling variable, y_1, with the fixed constant input, $\overline{y_1}$, is seen graphically in Figures 2 and 3. By decoupling variable y_1 from discipline 1 to discipline 2, the original feedback coupled multidisciplinary model is converted into a feed-forward model. In this case, instead of having to iteratively solve the coupled residual equations to find $\mathbf{y}$, the state variables

can be more cheaply determined by solving the one-way coupled model.

For the remainder of this paper, we denote the set of discipline couplings in our reference multidisciplinary model as M_0, and use M to represent a model that results from decoupling one or more coupling variables.

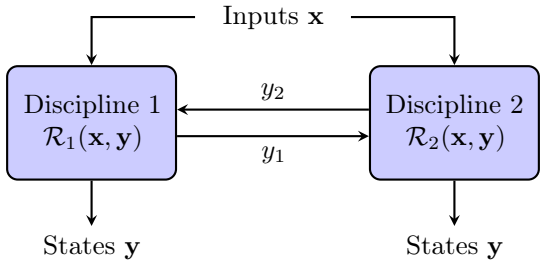

Fig. 2 Feedback coupled two discipline model, M_0

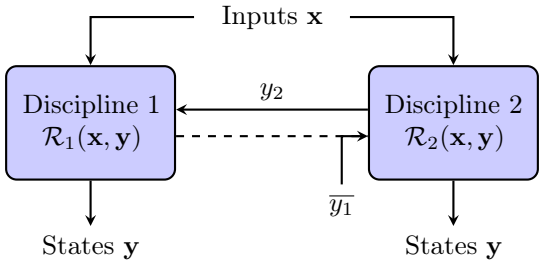

Fig. 3 One-way coupled two discipline model, M

In the context of forward UQ, the error introduced by decoupling variables between disciplines is defined as the accuracy of the PDF of the output variables of model M relative to the PDF of the output variables of the reference model M_0. One widely used measure for comparing PDFs is the Kullback-Leibler (KL) divergence. Here we consider the divergence $D_{KL}(\pi_{\mathbf{f}_{M_0}}||\pi_{\mathbf{f}_M})$ between the PDF of the output variables in the decoupled model, $\pi_{\mathbf{f}_M}$, and the PDF of the outputs in the reference model, $\pi_{\mathbf{f}_{M_0}}$. The KL divergence has useful information theoretic interpretations [30]; in particular, it provides an indication of the information lost when using $\pi_{\mathbf{f}_M}$ to approximate the distribution of the output variables with the discipline couplings in model M. The reader is referred to Ref. [30] for more properties of the KL divergence.

D. Search for Optimal Coupling

We search for the maximum number of discipline inputs that can be decoupled while minimizing the information lost from decoupling. These competing objectives are used to formulate a combinatorial optimization problem to find an optimal model by exploring the space of possible decoupled models, which we denote by $\mathcal{M}$. The optimal decoupled model is denoted by $M^*(\lambda)$ and is found by solving

$$M^*(\lambda) = \underset{M \in \mathcal{M}}{\arg\min}\ D_{KL}(\pi_{\mathbf{f}_{M_0}}||\pi_{\mathbf{f}_M}) - \lambda \mathcal{P}(M). \tag{3}$$

The two parts of the objective function in (3) are the KL divergence that measures the accuracy of the PDFs for the output variables in the decoupled model relative to the reference model, and the function $\mathcal{P}\colon \mathcal{M} \to \mathbb{N}$ that represents the sparsity of the model couplings. In this work, $\mathcal{P}(M) = d(M_0) - d(M)$, i.e., the function

counts the number of removed discipline couplings in M with respect to the reference model M_0, so that maximizing $\mathcal{P}(M)$ tends to favor sparse models with fewer couplings. Finally, these two parts are combined in the objective function by using a weighting parameter, λ, that controls the relative importance given to accuracy vs. sparsity.

We note that other metrics can also be used for the sparsity penalty, $\mathcal{P}(M)$, in the optimization problem. For example, by using additional knowledge about the contribution of each discipline coupling to the total cost of running the model, we can assign a different weight to each coupling that favors removing computationally expensive discipline couplings.

For a model with K disciplines, the number of discipline couplings, d, can grow as $d = K^2 - K$ for the case of a fully coupled reference model with single univariate coupling variables between all disciplines. The decoupling problem then associates a binary variable with each discipline coupling: the coupling can remain active or become inactive. Thus the cardinality of the model space $\mathcal{M}$ grows exponentially with d: $|\mathcal{M}| = 2^d$. Using a binary representation for each coupling variable, a model can be equivalently expressed as a binary vector in $\mathcal{M} = \{0, 1\}^d$, where a value of one denotes that a discipline coupling is active and a value of zero denotes an inactive coupling. For $K = 2$, this space is given by the four models in Figure 4, while for $K = 3$ disciplines there are $2^6 = 64$ possible models. As a result, in the case of many disciplines it is not feasible to compare all models by enumeration. Instead, the next section discusses tractable approaches for evaluating the objective function and exploring this high-dimensional binary space over different models.

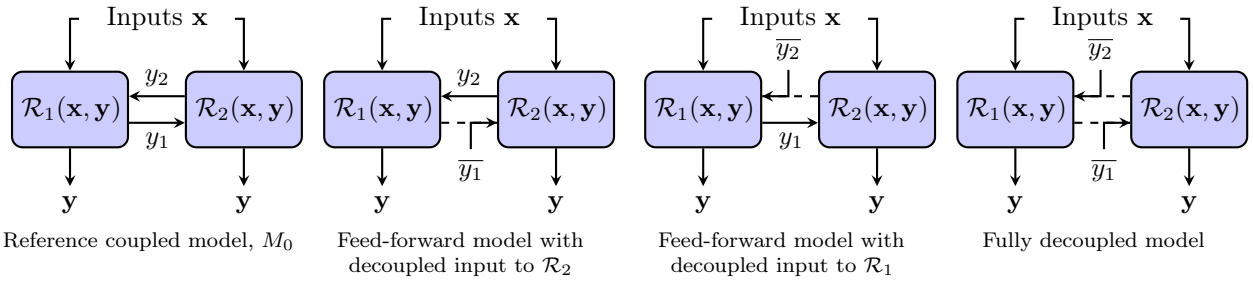

Fig. 4 Space of possible decoupled models, $\mathcal{M}$, for $K = 2$ disciplines

III. Approximate Coupling via Optimization

In Section A we describe the model selection procedure for finding optimal decoupled models. Section B introduces the use of model linearizations for estimating the KL divergence, and Section C describes a sequential Monte Carlo algorithm to solve the combinatorial optimization posed in equation (3).

A. Model Selection Procedure

In an engineering setting, it is of interest to select the model coupling approximation based on a combination of accuracy requirements and the availability of computational resources. To find a model that satisfies these criteria, we suggest solving the optimization problem in equation (3) for different values of λ in a set Λ to identify a sequence of optimally decoupled models, $M^*(\lambda)$. An optimal model for each application is selected from this sequence by trading-off accuracy in the output variable PDFs for a reduced number of discipline couplings and thus reduced computational cost for each sample evaluation.

To efficiently find the optimal model for each value of λ, we use a sequential Monte Carlo (SMC) algorithm for the combinatorial optimization, as described in Section C. To approximate the KL divergence term in equation (3), Chapter 3 of Ref. [31] presents model linearizations as an inexpensive and often sufficiently accurate approach to characterize the model output uncertainty and it is the approach we adopt in this work. The complete procedure for optimal model selection is summarized in Algorithm 1. The next two subsections provide details on Line 2 and Line 3, respectively, of Algorithm 1.

Algorithm 1: Optimal Model Selection

Input : Reference model: M_0, Set of λ parameters: Λ, Model sparsity function: $\mathcal{P}(M)$

1 **for** $\lambda \in \Lambda$ **do**

2 Setup objective function, $h(M;\lambda) = D_{KL}(\pi_{\mathbf{f}_{M_0}} || \pi_{\mathbf{f}_M}) - \lambda \mathcal{P}(M)$, using model linearizations;

3 Determine $M^*(\lambda) = \arg\min_{M \in \mathcal{M}} h(M;\lambda)$ using sequential Monte Carlo algorithm;

4 **end**

5 Select optimal model by comparing the accuracy and sparsity of $M^*(\lambda)$ for each $\lambda \in \Lambda$;

6 Return decoupled model, M^*;

B. Model Linearization

Our linearization of a multidisciplinary model is given by the first order Taylor series approximation of the model outputs in equation (2). This linearization is found by evaluating the partial derivatives of the function $\mathcal{F}$ with respect to input and state variables, which are denoted by $\partial_{\mathbf{x}}\mathcal{F}$ and $\partial_{\mathbf{y}}\mathcal{F}$, respectively. Both of these derivatives are evaluated at a specific linearization point for the input and state variables.

One natural linearization point is the mean of the state variables. However, finding the mean requires characterizing the distribution for these variables *a-priori*. Instead, an approximation commonly used in the literature is to linearize the models around the state variables corresponding to the mean values of the input variables [24]. The mean value of $\mathbf{x}$ is denoted by $\mu_{\mathbf{x}} = \mathbb{E}[\mathbf{x}]$ where $\mathbb{E}$ is the expectation operator with respect to the PDF $\pi_{\mathbf{x}}$. The state variables, $\mathbf{y}(\mu_{\mathbf{x}})$, corresponding to $\mu_{\mathbf{x}}$ satisfy the residual equations,

$\mathcal{R}(\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})) = \mathbf{0}$, and these states are referred to as a first-order mean for $\mathbf{y}$ [32].

Using the Taylor series expansion about the first-order mean, the linearized approximation of the output variables is given by

$$\mathbf{f} = \mathcal{F}(\mathbf{x}, \mathbf{y}) \approx \mathcal{F}(\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})) + \partial_{\mathbf{x}}\mathcal{F}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} (\mathbf{x} - \mu_{\mathbf{x}}) + \partial_{\mathbf{y}}\mathcal{F}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} (\mathbf{y} - \mathbf{y}(\mu_{\mathbf{x}})). \tag{4}$$

Similarly, the linearized approximation of the coupled residual equations for the model about the same linearization point is given by

$$\mathcal{R}(\mathbf{x}, \mathbf{y}) \approx \mathcal{R}(\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})) + \partial_{\mathbf{x}}\mathcal{R}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} (\mathbf{x} - \mu_{\mathbf{x}}) + \partial_{\mathbf{y}}\mathcal{R}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} (\mathbf{y} - \mathbf{y}(\mu_{\mathbf{x}})), \tag{5}$$

where $\partial_{\mathbf{x}}\mathcal{R}$ and $\partial_{\mathbf{y}}\mathcal{R}$ denote the partial derivatives of the residuals with respect to the input and state variables, respectively. Rearranging the approximation in equation (5) for the state variables, and noting that $\mathcal{R}(\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})) = \mathbf{0}$, the solution for the state variables that satisfy $\mathcal{R}(\mathbf{x}, \mathbf{y}) = \mathbf{0}$ is given by

$$\mathbf{y} = \mathbf{y}(\mu_{\mathbf{x}}) - (\partial_{\mathbf{y}}\mathcal{R}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})})^{-1} \partial_{\mathbf{x}}\mathcal{R}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} (\mathbf{x} - \mu_{\mathbf{x}}). \tag{6}$$

Here we note that the matrix of partial derivatives, $(\partial_{\mathbf{y}}\mathcal{R})^{-1}\partial_{\mathbf{x}}\mathcal{R}$ in equation (6), contains the sensitivity of the state variables to perturbations in the inputs. These sensitivities to the input variables are often computed as part of an adjoint analysis for single or multiple discipline models [33, 34]. Furthermore, the sensitivities are typically found without assembling and inverting the full matrix of partial derivatives for the residual equations with respect to the state variables. Instead, the values in the matrix are determined by solving a linear system using an iterative numerical method [2].

Substituting the linearization for the state variables in equation (4), a linear approximation to the output QoIs in the multidisciplinary model with respect to the input parameters, $\mathbf{x}$, is given by

$$\begin{aligned} \mathbf{f} \approx \tilde{\mathbf{f}}(\mathbf{x}) &= \mathcal{F}(\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})) + \left[\partial_{\mathbf{x}}\mathcal{F}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} - \partial_{\mathbf{y}}\mathcal{F}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} (\partial_{\mathbf{y}}\mathcal{R}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})})^{-1} \partial_{\mathbf{x}}\mathcal{R}|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})}\right] (\mathbf{x} - \mu_{\mathbf{x}}) \\ &= \mathcal{F}(\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})) + \left.\frac{d\mathcal{F}}{d\mathbf{x}}\right|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} (\mathbf{x} - \mu_{\mathbf{x}}), \end{aligned} \tag{7}$$

where $\tilde{\mathbf{f}}(\mathbf{x})$ represents the linear approximation and $\frac{d\mathcal{F}}{dx}$ denotes the matrix of total derivatives of $\mathcal{F}$ with respect to the input variables. With this linear relationship in the inputs, $\mathbf{x}$, the mean and variance of

the linearized outputs are propagated analytically from the mean and variance of the input variables. In particular, if the input variables are normally distributed with covariance matrix, $\Sigma_{\mathbf{x}}$, the linearized output variables are also normally distributed with mean, $\mu_{\tilde{\mathbf{f}}} = \mathbb{E}[\tilde{\mathbf{f}}]$, and covariance, $\Sigma_{\tilde{\mathbf{f}}} = \mathbb{E}[(\tilde{\mathbf{f}} - \mu_{\tilde{\mathbf{f}}})(\tilde{\mathbf{f}} - \mu_{\tilde{\mathbf{f}}})^T]$, where

$$\mu_{\tilde{\mathbf{f}}} = \mathcal{F}(\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})), \tag{8}$$

$$\Sigma_{\tilde{\mathbf{f}}} = \left[\frac{d\mathcal{F}}{d\mathbf{x}}\bigg|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} \right] \Sigma_x \left[\frac{d\mathcal{F}}{d\mathbf{x}}\bigg|_{\mu_{\mathbf{x}}, \mathbf{y}(\mu_{\mathbf{x}})} \right]^T. \tag{9}$$

In the case of normally distributed inputs, to compute the KL divergence for the output variables of each decoupled model, $M \in \mathcal{M}$, relative to M_0, we compute the mean and covariance of the linearized model outputs using (8) and (9). The mean and covariance for the linearized output variables of model M_0 are denoted by $\mu_{\tilde{\mathbf{f}}_{M_0}}$ and $\Sigma_{\tilde{\mathbf{f}}_{M_0}}$, respectively. Similarly, the mean and covariance for the decoupled model M are denoted by $\mu_{\tilde{\mathbf{f}}_M}$ and $\Sigma_{\tilde{\mathbf{f}}_M}$, respectively. Using the closed form expression for the KL divergence between two Gaussian distributions, we estimate the KL divergence between the PDFs for the outputs of models M_0 and M with the equation

$$\begin{aligned} D_{KL}(\pi_{\mathbf{f}_{M_0}} || \pi_{\mathbf{f}_M}) &\approx D_{KL}(\pi_{\tilde{\mathbf{f}}_{M_0}} || \pi_{\tilde{\mathbf{f}}_M}) \\ &= \frac{1}{2} \left\{ \mathrm{Tr}(\Sigma_{\tilde{\mathbf{f}}_M}^{-1} \Sigma_{\tilde{\mathbf{f}}_{M_0}}) + (\mu_{\tilde{\mathbf{f}}_M} - \mu_{\tilde{\mathbf{f}}_{M_0}})^T \Sigma_{\tilde{\mathbf{f}}_M}^{-1} (\mu_{\tilde{\mathbf{f}}_M} - \mu_{\tilde{\mathbf{f}}_{M_0}}) - p + \ln \left(\frac{|\Sigma_{\tilde{\mathbf{f}}_M}|}{|\Sigma_{\tilde{\mathbf{f}}_{M_0}}|} \right) \right\}, \end{aligned} \tag{10}$$

where $\mathrm{Tr}(\cdot)$ denotes the matrix trace operator, and $\ln(|\cdot|)$ is the log-determinant of a matrix. We note that the derivatives in equation (7) are typically available when computing system outputs or performing design optimization. This leads to a small incremental cost for evaluating the KL divergence in equation (10).

For numerical stability when computing the KL divergence with equation (10), we require that the covariance matrices of the reference and decoupled model are invertible and are well-conditioned. One necessary condition for the covariance to be full-rank and invertible is that the dimension of the random inputs, m, is at least the dimension of the output variables, p. For rank-deficient covariance matrices, adding Gaussian noise with small variance to the output QoIs ensures the PDF for the outputs is fully supported and the covariance is invertible. In practice, the selected QoIs should also have similar orders of magnitude for the covariance matrix to be numerically well-conditioned and for the objective function to closely weight the QoIs equally. One recommended approach is to scale QoIs to have similar orders of magnitude or to work with normalize physical variables (e.g., normalize absolute fuelburn by gross weight for an aircraft design problem).

Lastly, it is important to note that each decoupled model, $M \in \mathcal{M}$, results in a different first-order mean for the state variables. As a result, to compute the mean and covariance of the output variables, it is necessary to solve the nonlinear residual equations once for each model M at the mean input variables. The resulting first-order mean of $\mathbf{y}$ is used as the linearization point to evaluate the partial derivatives of the residual and output functions in the model linearization.

Nevertheless, in some cases there are efficiencies that can be exploited to avoid a system evaluation for every possible model considered. One choice that can significantly reduce the cost of evaluating the model linearizations and the KL divergence is the value selected for decoupled state variables, $\bar{y}_i$, that are fixed inputs to other disciplines. In particular, if the first-order mean value for the state variable y_i from M_0 is also used as the fixed input, $\bar{y}_i$, the solution of the unknown state variables in a decoupled model are still given by the first-order mean values for the states in the reference model. In this case, it is not necessary to re-evaluate the system's outputs after decoupling certain connections. As a result, the linearization point for the derivative evaluations also remains the same for a decoupled model, M, and only the dependence of the residual equations on the state variables change. To account for this change in variable dependence, it is sufficient to mask the effect of the decoupled inputs in the $\partial_{\mathbf{y}}\mathcal{R}$ matrix and recompute the linearization for model M using the updated matrix in equation (7).

C. Sequential Monte Carlo

To find an optimal decoupled model that best balances the accuracy and sparsity of an approximate model coupling, we solve the combinatorial optimization problem in equation (3). Common algorithms for combinatorial optimization over large binary spaces find approximate solutions by using local neighborhood searches, depth-first branch and bound searches, or randomized methods such as simulated annealing and genetic algorithms. For a comparison of these methods, the reader is referred to Ref. [35].

To effectively explore the model space, there is growing numerical evidence that global particle methods based on sequential Monte Carlo, which track a population of possible solutions, are robust and can often outperform heuristics and local search methods for binary optimization problems [36, 37]. This is particularly evident in strongly multi-modal objective functions where local search methods can become trapped in certain modes of the model space. Therefore, in this work we chose to use SMC to perform the combinatorial optimization.

SMC is described below and summarized in Algorithm 2. More details on the algorithm's implementation

can be found in the Appendix. The main steps of SMC are:

1. Generate a set of weighted particles representing possible decoupled models
2. Update the weights based on the value of the objective function
3. Propose new particles to explore the space while seeking the optimum
4. Repeat steps **2-3**

To find the optimal solution using SMC, we first define a sequence of probability distributions on the model space, which we denote by $P_t \colon \mathcal{M} \to [0,1]$ for each index $t \in \mathbb{N}_0$. These distributions progress from a distribution that is easy to sample from, such as the uniform weighting over the set of all possible decoupled models (i.e., $P_0(M) = \mathcal{U}(M)$), to the final target distribution of interest that concentrates most mass over the models that minimize the objective function. To smoothly move towards the target distribution, the goal of an efficient particle method in this context is to learn the correlations and properties of the model space in order to find the global minimizers without enumerating all models.

Particle methods approximate the distribution at each step by a finite set of $L \in \mathbb{N}$ weighted particles that each represent a particular model. We denote particle l at step t of the algorithm by $M_t^{(l)} \in \mathcal{M}$ for $l = 1, \ldots, L$. Each of these particles is also associated with a weight that we denote by $w_t^{(l)} \in [0,1]$. The collection of particles and weights at each step is given by $(\mathbf{M}_t, \mathbf{w}_t) = \{M_t^{(l)}, w_t^{(l)}\}_{l=1}^{L}$. The weights are set to $1/L$ when initializing the system at $t = 0$. At each iteration of the algorithm, the weights are updated based on the subsequent distributions using the ratio

$$u_{t+1}^{(l)} = w_t^{(l)} \frac{P_{t+1}(M_t^{(l)})}{P_t(M_t^{(l)})}, \tag{11}$$

followed by the normalization step, $w_{t+1}^{(l)} = u_{t+1}^{(l)} / \sum_l u_{t+1}^{(l)}$ for $l = 1, \ldots, L$. Here we note that it is sufficient to evaluate the models at the unnormalized versions of P_t and P_{t+1} given that the weights only require the ratios of these probability mass functions. As described in a recent study by Schäfer et al. [37], one common and successful technique to construct these distributions over the model space is to use a tempered family. This family of distributions assigns to each decoupled model, $M \in \mathcal{M}$, the probability mass

$$P_t(M) \propto \exp(-\rho(t) h(M)), \tag{12}$$

where $h \colon \mathcal{M} \to \mathbb{R}$ is the value of the objective function for each model in the combinatorial optimization problem, and $\rho \colon \mathbb{N}_0 \to \mathbb{R}_+$ is a monotonically increasing tempering parameter that depends on step t.

For the combinatorial optimization problem posed in equation (3), the objective function is given by $h(M) = D_{KL}(\pi_{\tilde{\mathbf{f}}_{M_0}} || \pi_{\tilde{\mathbf{f}}_M}) - \lambda \mathcal{P}(M)$. Therefore, as t and $\rho(t)$ increase at each iteration, the tempered family concentrates more mass on the set of models that minimize the value for the objective function by assigning higher weights to these models. However, by just repeatedly reweighing the initial set of particles using equation (11), the weights will become uneven and eventually lead to particle degeneracy with a poor sample approximation to each distribution.

As a result, the main ingredient that differentiates sequential Monte Carlo samplers from vanilla importance sampling is the series of alternating reweighting and proposal steps that resample and move the particles in order to explore the space outside of the current set of models. At each iteration, the algorithm measures a series of metrics including the effective sample size (ESS) to measure the uniformity of the weights $\mathbf{w}_t$ (i.e., relative optimality of the models at iteration t), and the particle diversity to measure the number of unique models in $\mathbf{M}_t$. The ESS and particle diversity are denoted by $\text{ESS}(\mathbf{w}_t)$ and $\zeta(\mathbf{M}_t)$, respectively, and defined in equation (13). Based on the ESS, SMC resamples the particles from its current empirical distribution and moves them by using an adapted proposal distribution that samples new particles in the model space. This combination of steps ensures there is a smooth transition between all distributions, P_t.

$$\text{ESS}(\mathbf{w}_t) = \frac{1}{\sum_{l=1}^{L} (w_t^{(l)})^2}, \qquad \zeta(\mathbf{M}_t) = \frac{1}{L} |\{M_t^{(l)} : l \in 1, \ldots, L\}| \tag{13}$$

The resample and move steps are repeated until the sample diversity drops below a threshold $\delta > 0$, indicating that most of the mass in P_t is concentrated on a few decoupled models. At the end, the algorithm returns the maximizer of the final distribution within this small set of particles as the optimal decoupled model. A summary of the algorithm to find the optimal model coupling using SMC is presented in Algorithm 2. More details on the proposal mechanism, the resample and move steps and how to update $\rho(t)$ (i.e., Lines 7-10 of the algorithm) can be found in the Appendix. The reader is also referred to Ref. [37] for specific details and limitations on each function in the SMC algorithm for solving combinatorial optimization problems.

In practice, the number of particles at each iteration of the SMC algorithm, L, and the minimum particle diversity, δ, are chosen relative to the cardinality of the model space. While the number of particles should be small enough to ensure that the model selection is computationally feasible, it should also be large enough to approximately represent the distributions, P_t. In our numerical examples, we chose a minimum particle diversity of $\delta = 0.1L$ to terminate the algorithm and indicate that most probability mass is concentrated on a small number of unique particles.

Finally, as empirically demonstrated in Ref. [37], this algorithm can successfully perform optimization on binary spaces of 250 dimensions—which, in the present context, corresponds to multidisciplinary models with up to 250 coupling variables. We note that for models that have more coupling, the algorithm can be first applied to groups of coupling variables to find a coarse approximation to the optimal decoupled model, followed by refining the retained coupling in each group.

Algorithm 2: SMC for Optimal Model Coupling

Input : Model Objective Function: $h(M)$, Minimum Particle Diversity: δ, Number of Particles: L

Initialize counter: t $= 0$;
Sample models: $M_t^{(l)} \overset{iid}{\sim} \mathcal{U}(M)$ for $l = 1, \ldots, L$;
Initialize weights: $w_t^{(l)} = 1/L$ for $l = 1, \ldots, L$;
Compute particle diversity: $\zeta_t = \zeta(\mathbf{M}_t)$;
while $\zeta_t > \delta$ **do**
    Update counter: $t \leftarrow t + 1$;
    Fit proposal: $q_t \leftarrow$ `Fit_Proposal`$(\mathbf{M}_t, \mathbf{w}_t)$;
    Resample: $\mathbf{M}_t \leftarrow$ `Re_sample`$(\mathbf{M}_t, \mathbf{w}_t)$;
    Move models: $\mathbf{M}_t \leftarrow$ `Move`$(q_t, \mathbf{M}_t, \mathbf{w}_t)$;
    Update $\rho(t)$ to maintain ESS$(\mathbf{w}_t)$;
    Update weights: $\mathbf{w}_t \leftarrow$ `Importance_Weights`$(h(M), \rho(t), \mathbf{M}_t, \mathbf{w}_t)$;
    Compute particle diversity: $\zeta_t = \zeta(\mathbf{M}_t)$;
end
Return $M^* = \arg\min_{M \in \mathbf{M}_t} h(M)$, $h(M^*)$

IV. Numerical Results

In this section, we present results for applying the proposed model selection algorithm to two aerospace engineering examples. These are a three-discipline model of a satellite used to detect forest fires in Section A that was presented in an uncertainty quantification context by Sankararaman et al., [24] and a model for turbine engine cycle analysis in Section B that was developed by Hearn et al. [38]. We conclude in Section C by referencing a publicly available implementation of the model selection algorithm for future analyses.

A. Fire Detection Satellite Model

We begin by studying the satellite model with a nominal set of parameters for the input uncertainty distribution in Section 1. The effect of different input uncertainty on the optimal models is analyzed in Section 2.

1. Nominal Input Uncertainty

To analyze the performance of a fire detection satellite under uncertain operating conditions, we consider a model that comprises three disciplines: orbit analysis, attitude control, and power analysis. As seen in Figure 5, the model features both feed-forward and feedback coupling variables to exchange information between the disciplines.

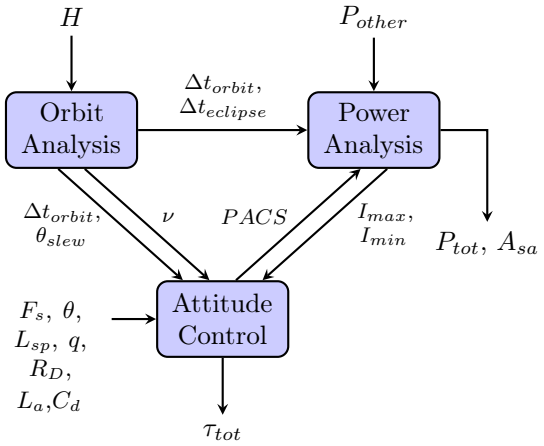

Fig. 5 Fire detection satellite model from Ref. [24]

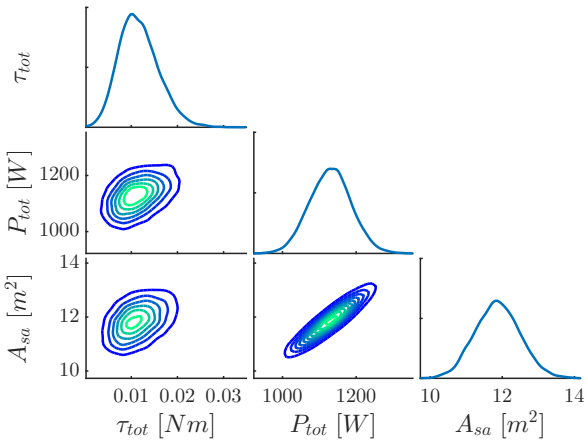

Fig. 6 Joint PDF of the reference model

To represent the uncertain conditions, the model includes nine Gaussian random variables that are described by their mean and standard deviation parameters in Table 1. After propagating the input uncertainty through the disciplines, the seven state variables connecting the disciplines are also uncertain and are distributed with potentially non-Gaussian probability distributions. These state variables are: the orbit period (Δt_{orbit}), eclipse period ($\Delta t_{eclipse}$), satellite velocity (v), maximum slewing angle (θ_{slew}), the power of the attitude control system (P_{ACS}), and the moments of inertia (I_{max} and I_{min}). In addition, the orbit period is an input to two disciplines (orbit analysis and attitude control) resulting in a total of $d = 8$ coupling variables for this satellite model as seen in Figure 5. These coupling variables with the discipline residual equations in Ref. [24] are used to compute the three output QoIs: total torque (τ_{tot}), total power (P_{tot}), and the area of the solar array (A_{sa}). The joint PDF of these output variables, $\pi_{\mathbf{f}}$, is displayed in Figure 6.

By defining the model displayed in Figure 5 as the reference model, M_0, our algorithm for optimal model selection considers possible decoupled models, M, that have a smaller number of discipline couplings by fixing a subset of the coupling variables of each discipline to their first-order mean value from M_0. In order to empirically validate the use of model linearizations to estimate the KL divergence of the decoupled models, we compare the Gaussian distribution for the linearized outputs of the reference model to the uncertainty in

Table 1 Random variable parameters in the satellite model from Ref. [24]

Random Variable	Symbol	Mean	Standard Deviation
Altitude	H	18.0×10^6 m	1.0×10^6 m
Power other than attitude control	P_{other}	1.0×10^3 W	50.0 W
Average solar flux	F_s	1.4×10^3 W/m^2	20.0 W/m^2
Deviation of moment axis	θ	15.0°	1.0°
moment arm for radiation torque	L_{sp}	2.0 m	0.4 m
Reflectance factor	q	0.5	1.0
Residual dipole of spacecraft	R_D	5.0 Am^2	1.0 Am^2
Moment arm for aerodynamic torque	L_a	2.0 m	0.4 m
Drag coefficient	C_d	1.0	0.3

the output variables of the nonlinear coupled system based on 10^4 Monte Carlo samples. As seen in Figure 7, the joint empirical PDF for the output variables, $\pi_{\mathbf{f}}$, is closely approximated by the multivariate Gaussian distribution resulting from the linearization in equation (7).

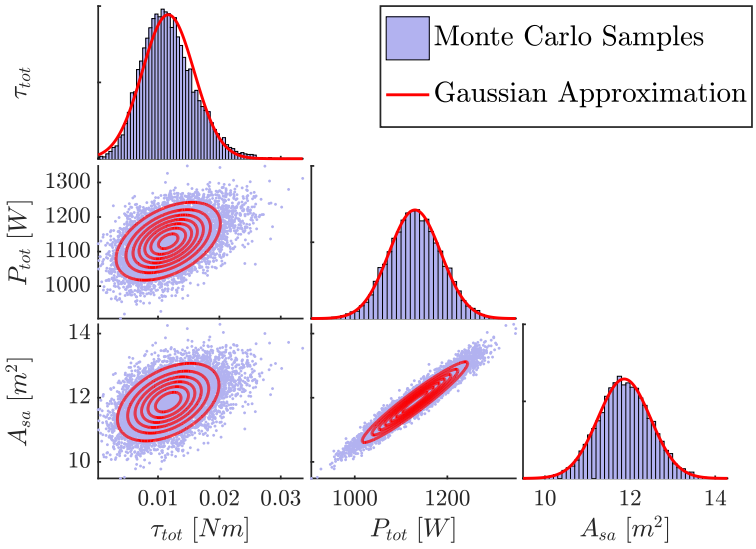

Fig. 7 Comparison of Gaussian approximation and Monte Carlo samples

Using these linearizations for the output variables, the SMC algorithm explores the model space using $L = 100$ particles to determine the optimal decoupled model, $M^*(\lambda)$, that solves the combinatorial optimization problem posed in equation (3). For each value of λ, the algorithm identifies a subset of discipline couplings that best trade off the accuracy of the output PDFs, that is estimated from the KL divergence, $D_{KL}(\pi_{\tilde{\mathbf{f}}_{M_0}} || \pi_{\tilde{\mathbf{f}}_{M^*}})$, and the added sparsity in the discipline couplings, as measured by $\mathcal{P}(M^*)$.

With increasing values for λ, the objective in the optimization problem adds a greater penalty to models with less sparsity in the discipline couplings. This results in a set of increasingly sparser models that

remove additional couplings at the expense of accuracy in the output PDF. These decoupled models and the corresponding output joint PDFs for four different values of λ in $\Lambda = \{10^{-4}, 10^{-3}, 0.1, 1\}$ are presented in Figures 8 to 15. The decoupled connections between each set of disciplines are indicated by the dashed lines in these figures.

For the fire detection satellite model, the algorithm identifies that the feedback coupling variables between the attitude control and power disciplines for the moment of inertia (I_{max} and I_{min}) could be fixed to their mean values (Figure 8). This decoupling has a small effect on the accuracy of the joint output PDF, as seen in Figure 9. With increasing values for λ, the slewing angle, the satellite velocity, and the orbit period are also found to weakly contribute to the uncertainty in the total torque and the overall power requirement for the attitude control subsystem, as seen in Figures 10 and 11. The subsequent sparser model for $\lambda = 0.1$ fixes these discipline coupling variables along with the state inputs to the power discipline for the orbit and eclipse period. Finally, we note that for $\lambda = 1$, the optimal model decouples all disciplines and as a result only propagates the first-order mean values for all coupling variables. In this case, the model outputs are computed independently for each discipline and one can take advantage of parallel computation to do so. The joint PDF for the output variables of the model for $\lambda = 1$ is given in Figure 15.

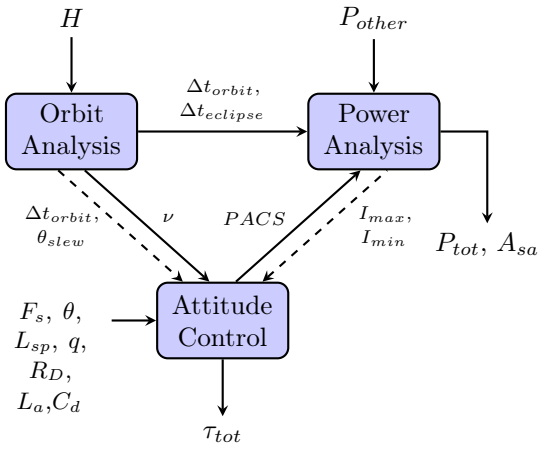

Fig. 8 Optimal model coupling for $\lambda = 10^{-4}$

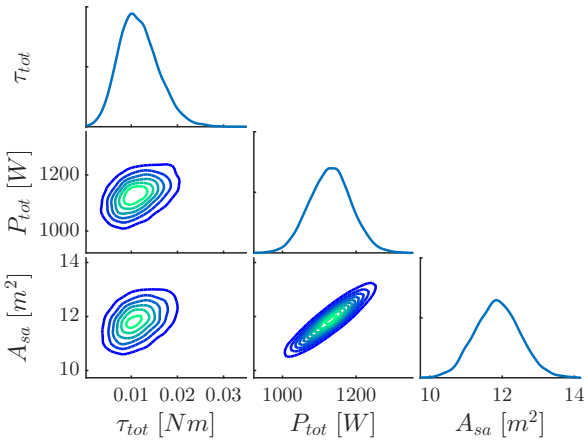

Fig. 9 Joint PDF for $\lambda = 10^{-4}$

In Table 2 we present the values for the KL divergence between the Gaussian distribution for the linearized outputs of the reference model, M_0, and the optimally decoupled models, M^*, that we denote as the linearized KL divergence. The number of active coupling variables, representing the increasing sparsity of each decoupled model with larger values of λ, is also presented in the table below.

From the table and figures above, we observe that the optimal models identified by the model selection

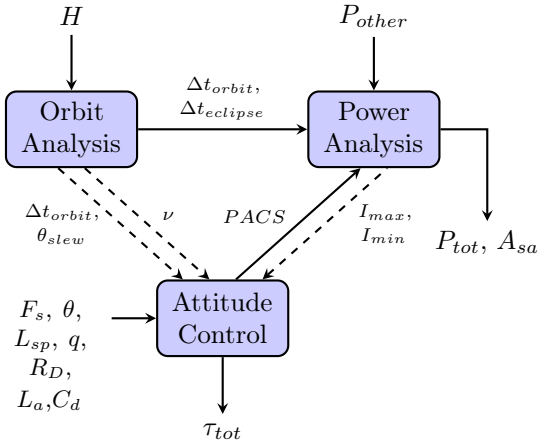

Fig. 10 **Optimal model coupling for** $\lambda = 10^{-3}$

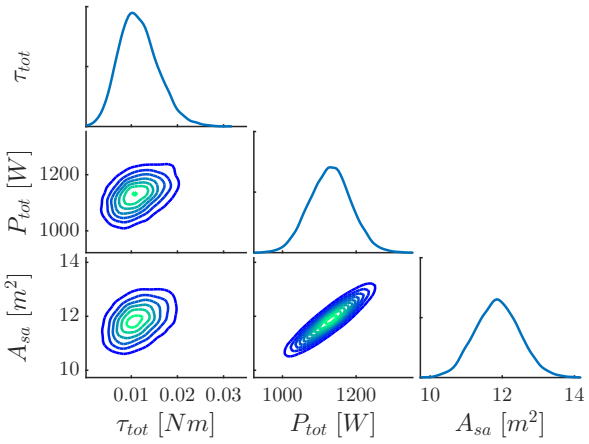

Fig. 11 **Joint PDF for** $\lambda = 10^{-3}$

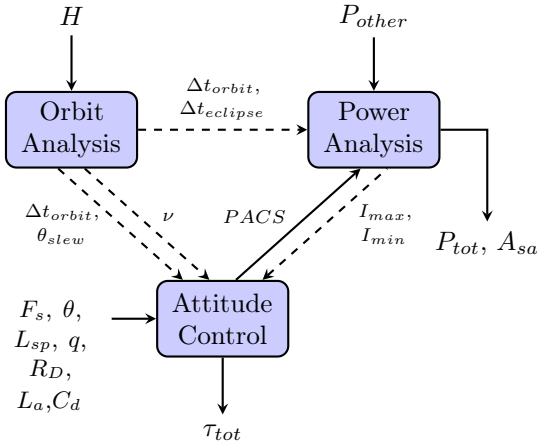

Fig. 12 **Optimal model coupling for** $\lambda = 0.1$

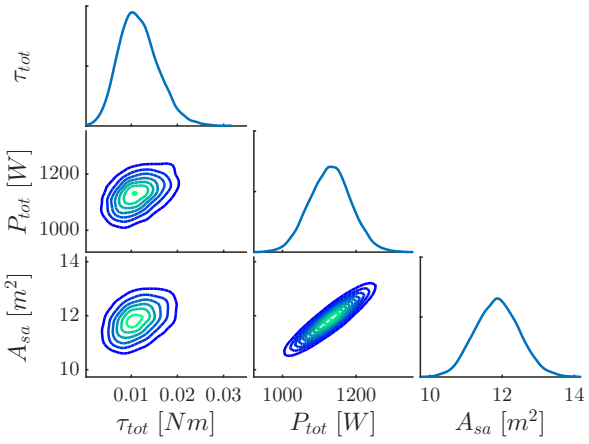

Fig. 13 **Joint PDF for** $\lambda = 0.1$

Table 2 **Linearized KL divergence for decoupled satellite models with nominal inputs**

	$\lambda = 10^{-4}$	$\lambda = 10^{-3}$	$\lambda = 0.1$	$\lambda = 1$
Linearized KL Divergence	7.30×10^{-16}	7.43×10^{-4}	4.20×10^{-3}	1.33×10^{-1}
Active Coupling Variables	4/8	3/8	1/8	0/8

algorithm corresponding to $\lambda = 10^{-4}, 10^{-3}, 0.1$ result in similar distributions for the model outputs as the reference model based on the linearized KL divergence. Furthermore, while the final optimal model for $\lambda = 1$ approximately captures the marginal uncertainty in the outputs, this fully decoupled model does not represent, as accurately, the correlations between the output variables, as seen in the joint marginal PDF for P_{tot} and τ_{tot} in Figure 15. Nevertheless, given that all decoupled models feature only feed-forward connections, these models with approximate coupling can be used to cheaply propagate the uncertainty in

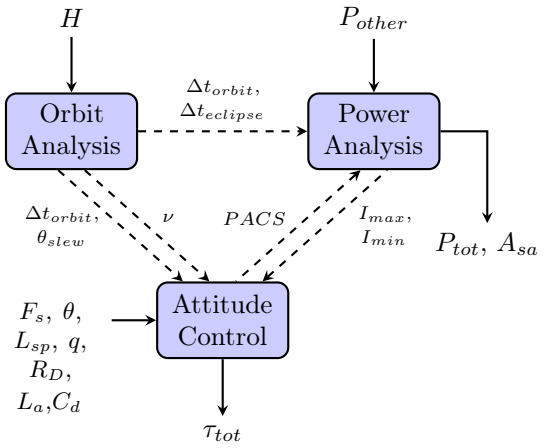

Fig. 14 Optimal model coupling for $\lambda = 1$

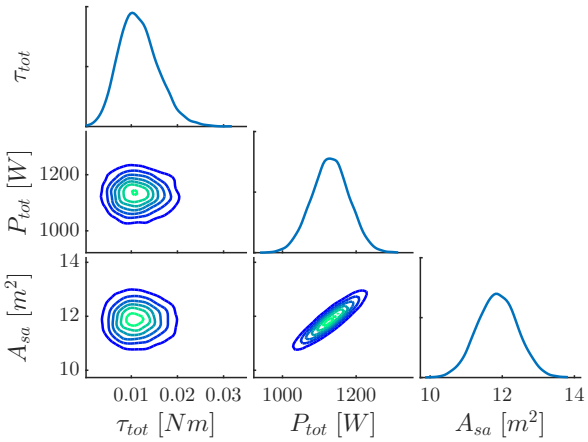

Fig. 15 Joint PDF for $\lambda = 1$

the multidisciplinary model without having to iteratively solve for the outputs that correspond to each set of input variables. This results in a substantial computational savings for performing forward UQ. In particular, if an accuracy of about 10^{-1} in KL is sufficient in the output uncertainty, the disciplines can be analyzed independently using the model corresponding to $\lambda = 1$.

Finally, we analyze the pointwise errors of the three model outputs in each decoupled model, M^*, relative to the reference model, M_0. The pointwise error is given by the difference in the values of the output variables in the decoupled and reference models for the same input sample. The errors for 10^4 input samples are displayed in Figures 16 to 19. With an increase in λ, the sparser models result in progressively greater pointwise errors, which are seen in the greater spread of the output values between both models. Although the set of discipline couplings corresponding to $\lambda = 0.1$ and 1 have relatively accurate joint PDFs for the model output uncertainty, their larger pointwise errors makes them less adequate for computations that require accurate pointwise approximations, such as multidisciplinary optimization. Nevertheless, the models associated with $\lambda = 10^{-4}$ and 10^{-3} (Figures 8 and 10) have accurate PDFs for the output variables and low pointwise errors relative to the reference model, M_0.

2. Sensitivity to Input Uncertainty Assumptions

In comparison to methods that partition and reorder disciplines to reduce feedback loops based only on the graphical structure of the model [39], the algorithm presented in this work accounts for the system's physics to identify an optimal approximation of the system coupling. In particular, this coupling is dependent on the sensitivities of the outputs to each input variable and the scaling by the input uncertainty (i.e., the covariance

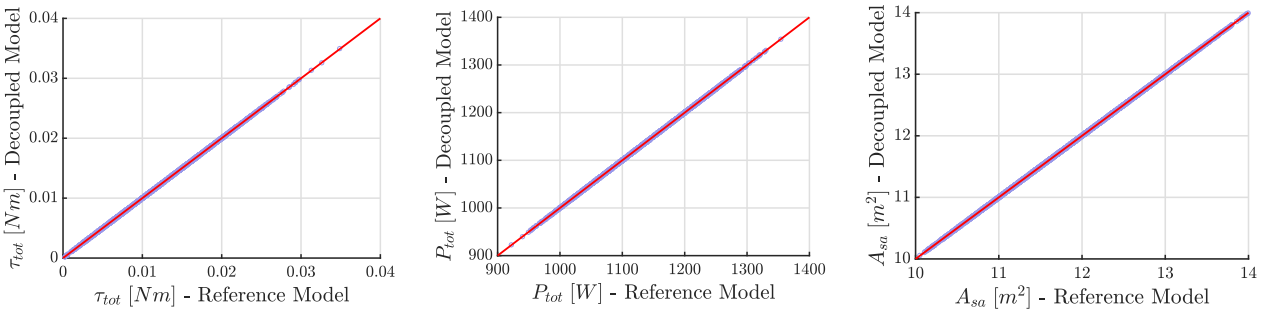

Fig. 16 Comparison of QoI with the optimal decoupled model for $\lambda = 10^{-4}$

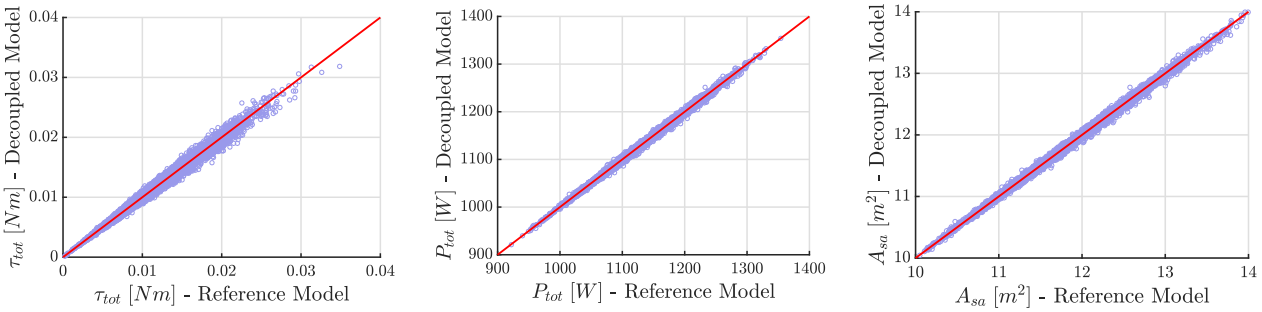

Fig. 17 Comparison of QoI with the optimal decoupled model for $\lambda = 10^{-3}$

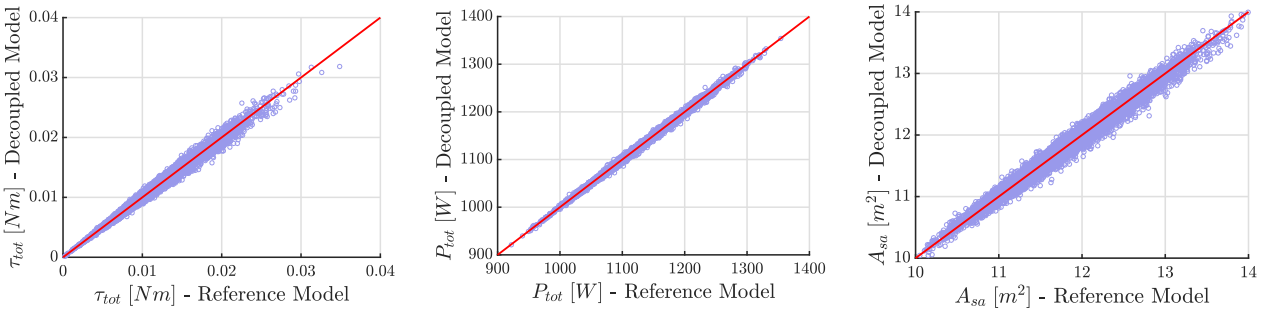

Fig. 18 Comparison of QoI with the optimal decoupled model for $\lambda = 0.1$

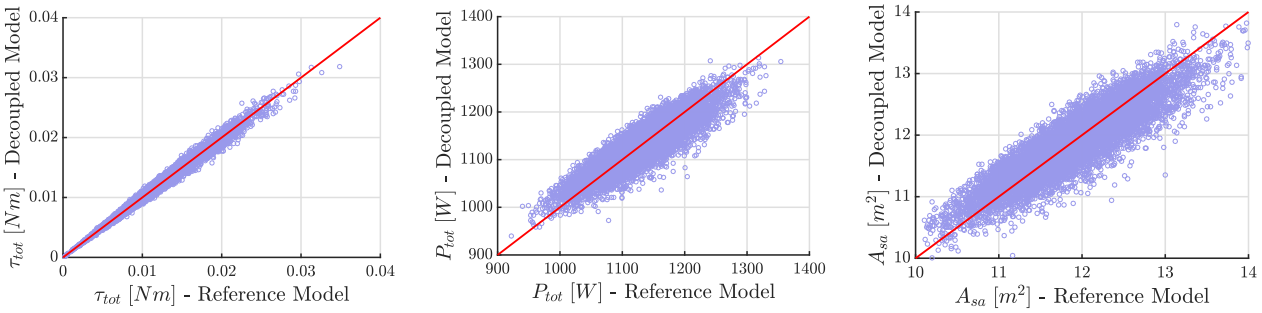

Fig. 19 Comparison of QoI with the optimal decoupled model for $\lambda = 1$

matrix, $\Sigma_{\mathbf{x}}$, for normally distributed input variables). Therefore, when changing the parameters describing the uncertainty in the input variables, the optimal decoupled model should be reassessed to minimize the

discrepancy in the estimates for the output uncertainty.

For the satellite model, we consider the impact of different input parameters by introducing uncertainty in the satellite's slewing time period (Δt_{slew}), which was held constant in the above analysis. This variable is an input to the attitude control discipline and together with the coupling variable, θ_{slew}, affects the slewing torque and the resulting total torque on the satellite. We note that with a constant value for the Δt_{slew} input, the effect of slew on the total torque calculation can be accurately represented with only the mean value for the slew angle, $\bar{\theta}_{slew}$. This allows for decoupling the θ_{slew} variable in the optimal models identified above in Section 1.

By describing the slewing time-period as a normal random variable with a shifted mean, $\Delta t_{slew} \sim \mathcal{N}(5, 0.5^2)$, the joint PDF for the output variables based on 10^4 Monte Carlo samples is displayed in Figure 20. Using the model linearization approach, the optimal decoupled models identified by the model selection algorithm and the joint PDFs of the outputs for $\lambda = 10^{-4}$ and 10^{-3} are presented in Figures 21-24. These models differ from the optimal models identified in Section 1 due to the necessity of maintaining the θ_{slew} coupling variable in this case to accurately compute the total satellite torque. The models for $\lambda = 0.1$ and 1 match the earlier optimal models in Figures 12 and 14, respectively. This is due to the greater weight on having sparse discipline couplings in the objective function with these larger values for the λ parameter.

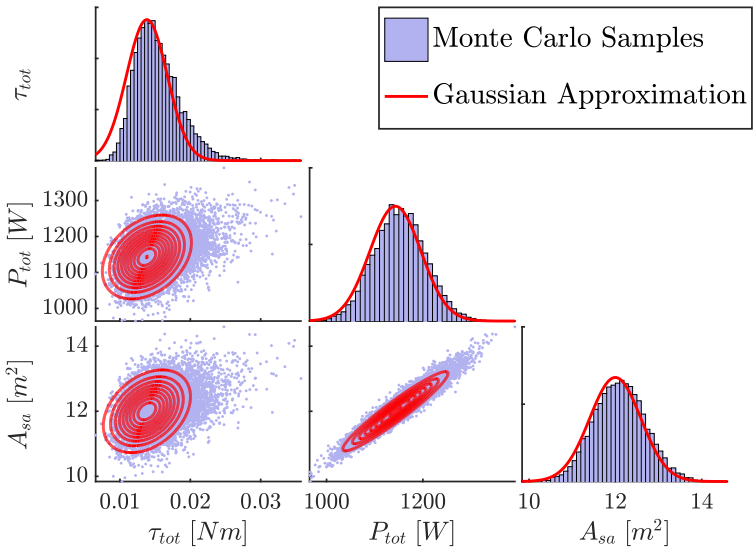

Fig. 20 Gaussian approximation and Monte Carlo samples with different inputs

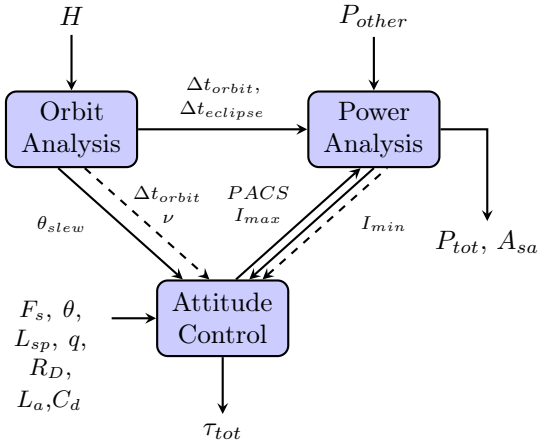

Fig. 21 Optimal model coupling for $\lambda = 10^{-4}$ with input $\Delta t_{slew} \sim \mathcal{N}(5, 0.5^2)$

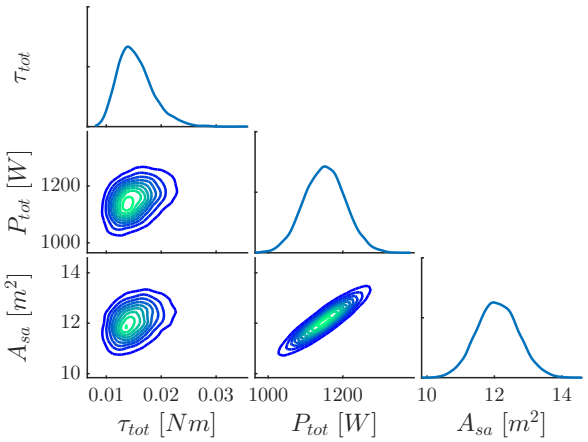

Fig. 22 Joint PDF for $\lambda = 10^{-4}$ with input $\Delta t_{slew} \sim \mathcal{N}(5, 0.5^2)$

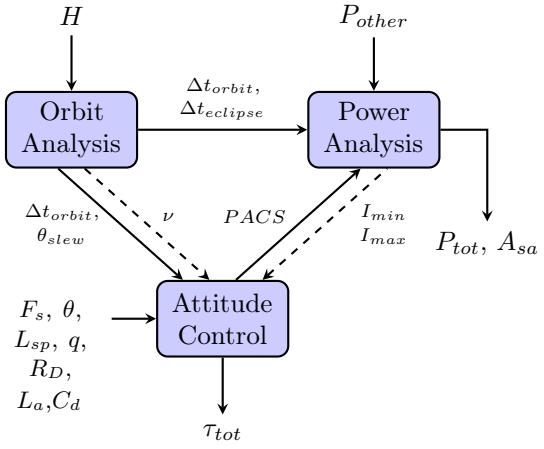

Fig. 23 Optimal model coupling for $\lambda = 10^{-3}$ with input $\Delta t_{slew} \sim \mathcal{N}(5, 0.5^2)$

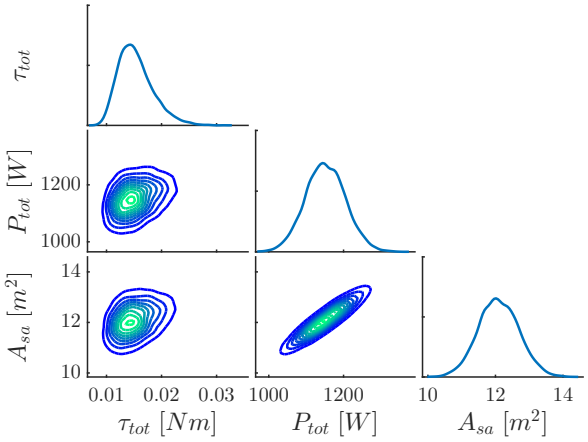

Fig. 24 Joint PDF for $\lambda = 10^{-3}$ with input $\Delta t_{slew} \sim \mathcal{N}(5, 0.5^2)$

Table 3 presents the values for the linearized KL divergence between the Gaussian distribution for the outputs of the reference model, M_0, and these optimally decoupled models, M^*. Furthermore, Table 4 evaluates the linearized KL divergence when using the optimal decoupled models from Figures 8-14 to estimate the outputs with the different input condition for Δt_{slew}. The discrepancy, particularly for $\lambda = 10^{-4}$, demonstrates that models determined with the nominal input uncertainty conditions in Section 1 are suboptimal when using different inputs. This emphasizes the importance of accounting for the inputs when identifying optimal decoupled models that accurately propagate the uncertainty to the model's outputs.

Finally, we note that by further decreasing the mean and variance of the time period to $\Delta t_{slew} \sim \mathcal{N}(1, 0.01^2)$, the importance of the θ_{slew} coupling variable becomes more pronounced. In particular, the optimal decoupled models for $\lambda = 10^{-4}, 10^{-3}, 10^{-2}$ and $\lambda = 1$ with this input uncertainty are displayed in Figures 25 and 26,

Table 3 Linearized KL divergence with optimal decoupled satellite models found based on an uncertain Δt_{slew} input

	$\lambda = 10^{-4}$	$\lambda = 10^{-3}$	$\lambda = 10^{-2}$	$\lambda = 1$
Linearized KL Divergence	0	2.80×10^{-4}	1.04×10^{-2}	7.75×10^{-2}
Active Coupling Variables	5/8	4/8	1/8	0/8

Table 4 Linearized KL divergence with decoupled satellite models found based on the inputs in Section 1

	$\lambda = 10^{-4}$	$\lambda = 10^{-3}$	$\lambda = 10^{-2}$	$\lambda = 1$
Linearized KL Divergence	68.5×10^{-4}	68.5×10^{-4}	1.04×10^{-2}	7.75×10^{-2}
Active Coupling Variables	5/8	4/8	1/8	0/8

respectively. Even for larger values of λ, these optimal models retain both the θ_{slew} and $PACS$ coupling variables for accurate output estimation, as compared to the increasingly decoupled models that are found based on the nominal inputs in Section 1.

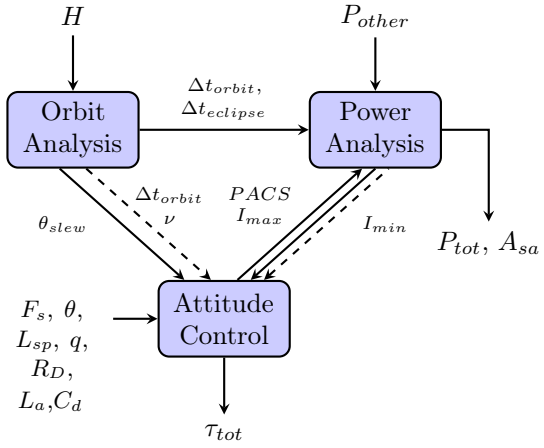

Fig. 25 Optimal model coupling for $\lambda = 10^{-4}, 10^{-3}, 10^{-2}$ with input $\Delta t_{slew} \sim \mathcal{N}(1, 0.01^2)$

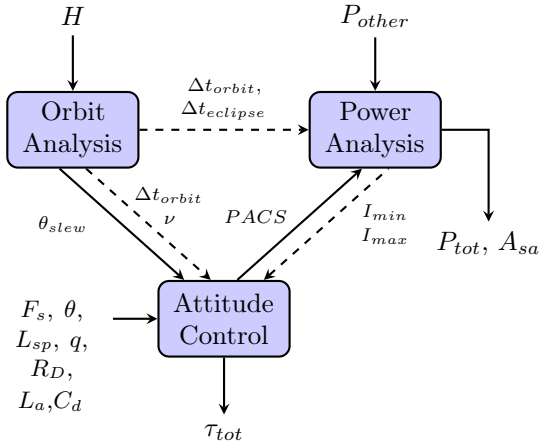

Fig. 26 Optimal model coupling for $\lambda = 1$ with input $\Delta t_{slew} \sim \mathcal{N}(1, 0.01^2)$

B. Turbine Engine Model

We use here the model for turbine engine cycle analysis that was presented in Ref. [38] and recently used in a coupled propulsive-aerodynamic model to study boundary layer ingestion [40]. The reference model for the system, M_0, is displayed in Figure 27 and consists of 13 disciplines for 12 engine components and a performance function, which are connected by 22 coupling variables. These disciplines model the engine with a core air-stream passing through a fan, compressor, burner and turbine as well as a second stream of air that bypasses the engine core. For each set of inputs that include the fan pressure ratio (FPR), compressor

pressure ratio (CPR), bypass ratio (BPR), and mass-flow rate (W), the model computes 4 output quantities of interest: thrust-specific fuel consumption ($TSFC$), net thrust (F_n), overall pressure ratio (OPR), and engine burner temperature (T_4).

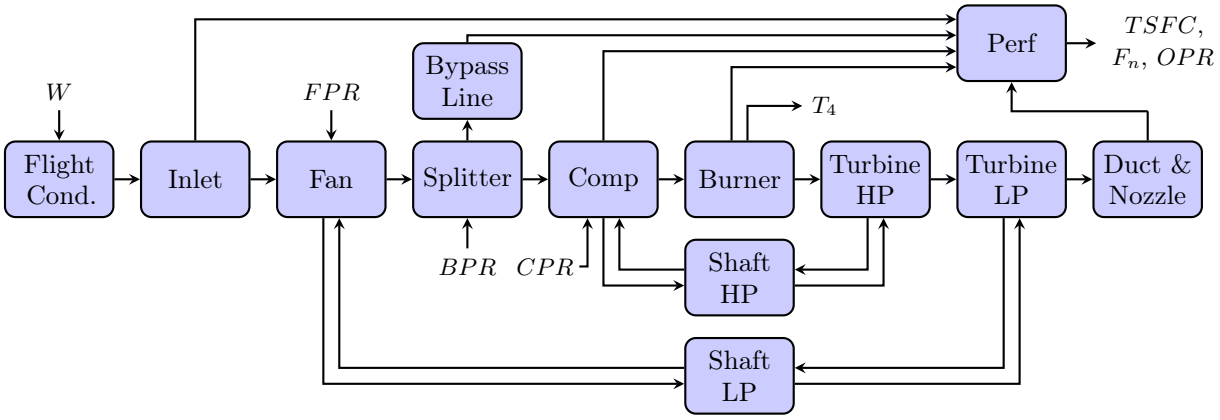

Fig. 27 Turbine engine cycle model from Ref. [38]

To analyze various operating conditions, the model inputs are represented with Gaussian distributed random variables whose parameters are listed in Table 5.

Table 5 Random variable parameters in the turbine engine cycle model

Random Variable	Symbol	Mean	Standard Deviation
Fan pressure ratio	FPR	1.5	0.01
Compressor pressure ratio	CPR	15.0	0.10
Bypass ratio	BPR	2.0	0.10
Mass-flow rate	W	1000.0 lbm/s	$\sqrt{10.0}$ lbm/s

By propagating 10^3 Monte Carlo samples of these inputs through the nonlinear system, the non-Gaussian PDF of the output variables, $\pi_{\mathbf{f}}$, is empirically characterized for the selected QoIs in Figure 28. In the optimal model selection algorithm, we use the linearizations of these output variables with respect to the inputs when fixing different subsets of the discipline couplings to their estimated mean values. We then compare the resulting multivariate Gaussian distribution for the output of each linearized decoupled model, $\pi_{\tilde{\mathbf{f}}_M}$, to the output PDF of the linearized reference model, $\pi_{\tilde{\mathbf{f}}_{M_0}}$.

Applying the model selection procedure described above with $L = 1000$ particles, the SMC algorithm identifies decoupled models in the space of size $|\mathcal{M}| = 2^{22}$ for four different values of λ in $\Lambda = \{0.1, 1, 10, 10^3\}$. These models are presented in Figures 29 to 32 with the decoupled connections indicated by dashed lines. Similarly to the satellite model, we observe that with an increase in λ, the optimal solutions to the combinatorial optimization problem have greater sparsity in the number of discipline couplings. This results in a set

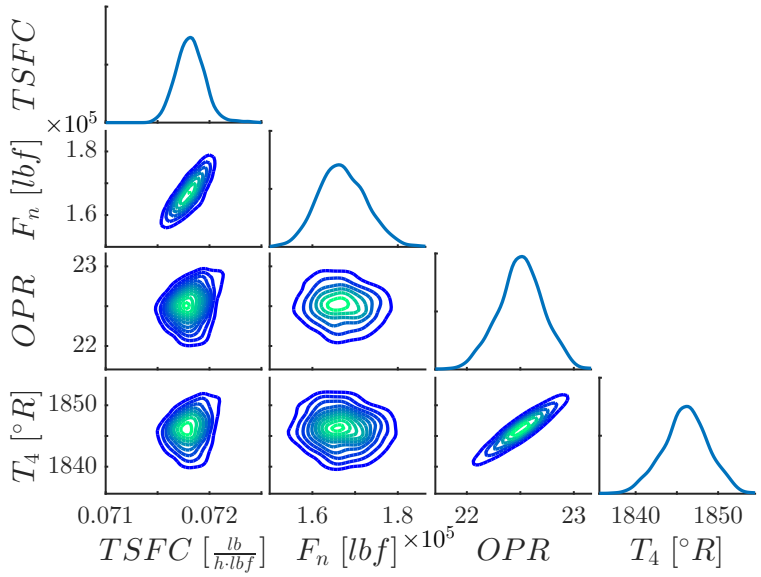

Fig. 28 Joint PDF of the reference model

of models that sequentially remove weak connections to increase the sparsity of couplings in exchange for accuracy in the output PDFs. For this turbine engine case study with $\lambda = 0.1$, and $\lambda = 1$, this results in neglecting the effect of the feedback from the low-pressure turbine on the fan and the variability of inlet ram drag on system performance.

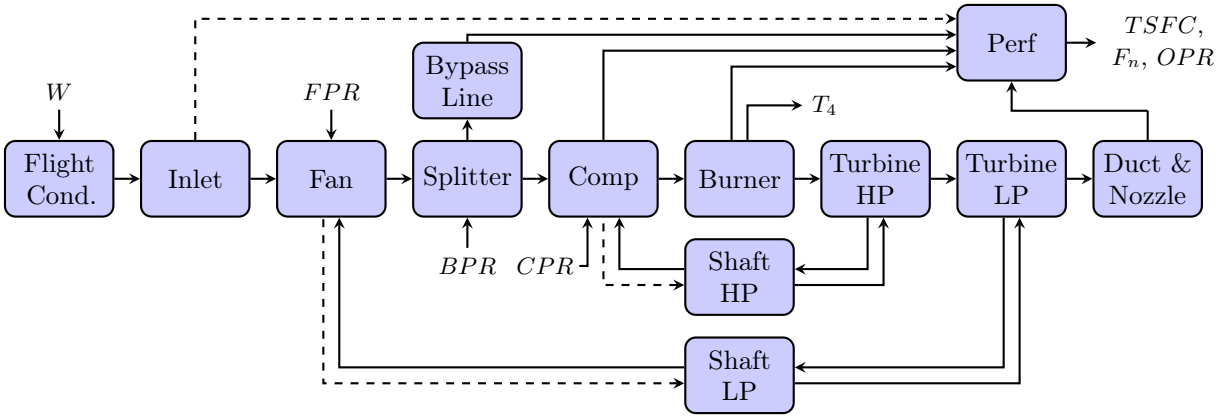

Fig. 29 Optimal model coupling for $\lambda = 0.1$

With a further increase in λ and a greater penalty on low sparsity models, the effect of the bypass line and the feedback from the high-pressure turbine are also neglected in the optimal models. The accuracy of the system outputs corresponding to the decoupled model for each value of λ is displayed in the joint PDF of the output QoIs in Figures 33 to 36. We observe that while the distribution for the output $TSFC$ variable is closely captured with the decoupled models for $\lambda = 0.1$, and 1, the correlations of this variable are not as accurately represented when decoupling connections from the bypass line and feedback from the high-pressure

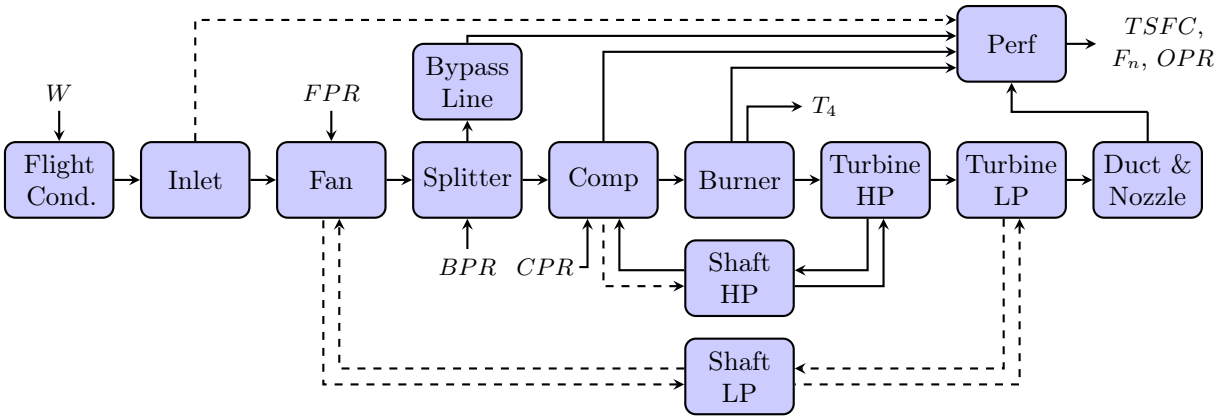

Fig. 30 Optimal model coupling for $\lambda = 1$

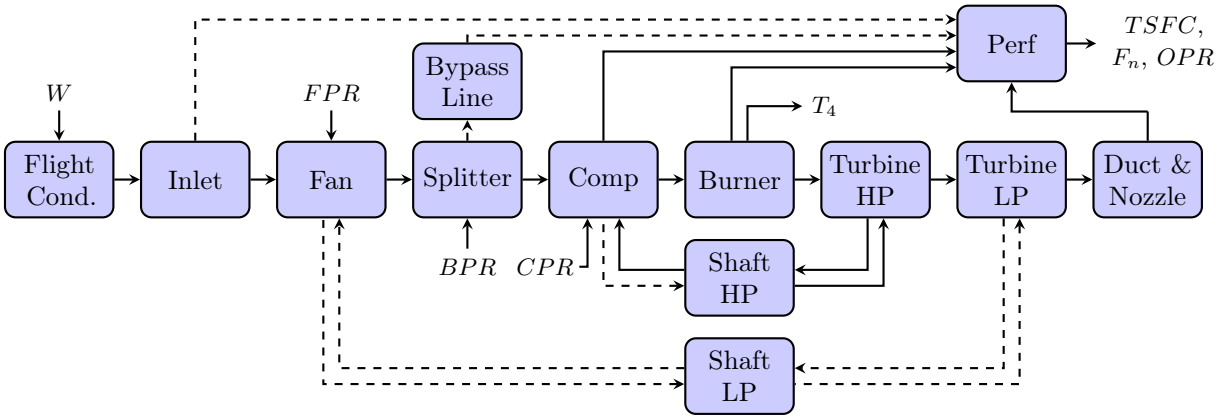

Fig. 31 Optimal model coupling for $\lambda = 10$

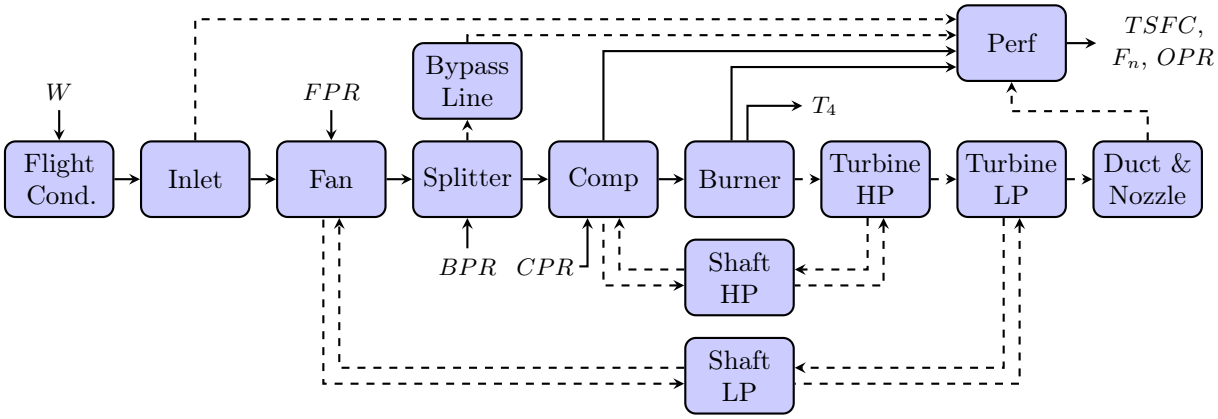

Fig. 32 Optimal model coupling for $\lambda = 10^3$

turbine. This effect on the correlation is expected given the tight balance between the bypass flow and the core flow through the high pressure turbine, that together produce most of the engine's thrust.

Furthermore, while the first three models accurately capture the net thrust, F_n, the uncertainty in this

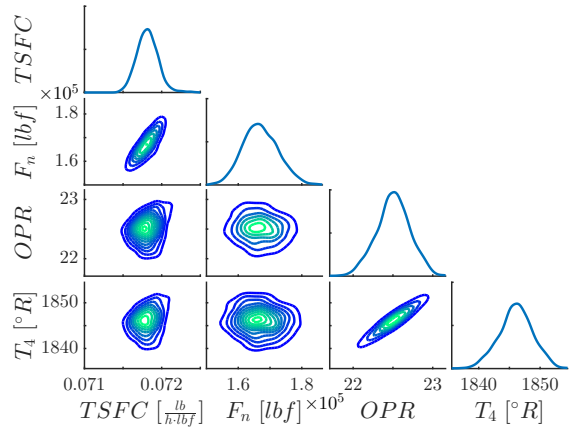

Fig. 33 Joint PDF for $\lambda = 0.1$

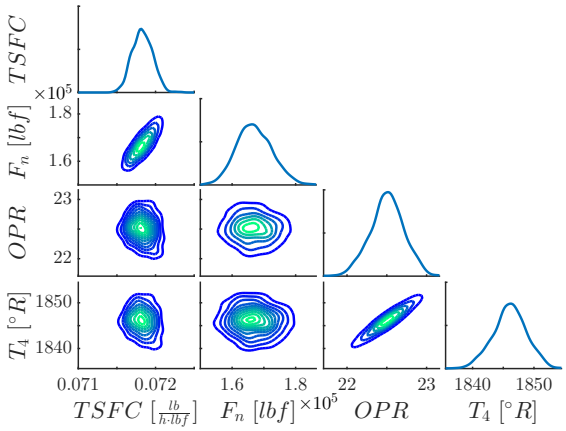

Fig. 34 Joint PDF for $\lambda = 1$

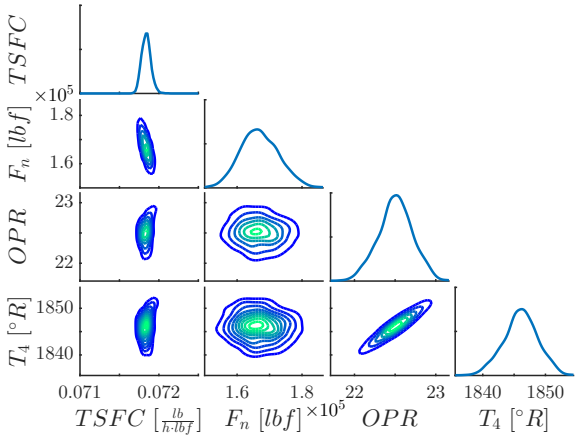

Fig. 35 Joint PDF for $\lambda = 10$

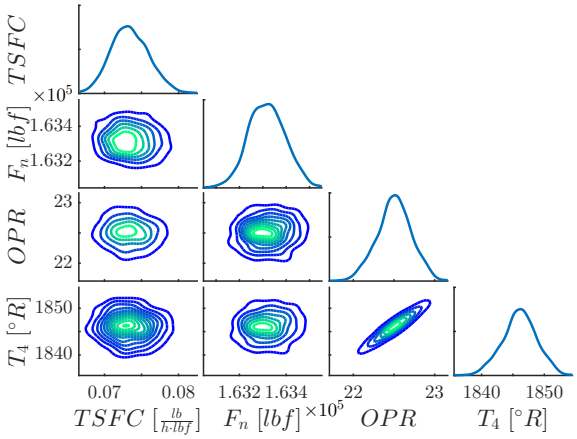

Fig. 36 Joint PDF for $\lambda = 10^3$

variable is poorly represented with the last model when neglecting the contribution of the nozzle. Finally, all four models accurately represent the uncertainty in the overall pressure ratio, OPR, and engine burner temperature, T_4. As a result, depending on the application and available computational resources, users may select an appropriate set of discipline couplings for the turbine engine to model the uncertainty in the system's output variables.

To quantitatively compare the decoupled models, we also present the linearized KL divergence between the approximating Gaussian distributions for the reference and optimally decoupled models of the turbine engine in Table 6. These results are presented along with the number of active coupling variables in each model which demonstrate the increasing sparsity in the decoupled models for larger values of λ.

To further compare the models, the pointwise errors of the three outputs from the performance discipline ($TSFC$, F_n, and OPR) for the decoupled models corresponding to $\lambda = 0.1, 1, 10$ are plotted in Figures 37 to 39. These figures demonstrate the increasing errors in the outputs of the decoupled models relative to the

Table 6 Linearized KL divergence for decoupled turbine engine models

	$\lambda = 0.1$	$\lambda = 1$	$\lambda = 10$	$\lambda = 10^3$
Linearized KL Divergence	6.95×10^{-10}	1.15×10^{0}	1.25×10^{1}	3.73×10^{3}
Active Coupling Variables	19/22	16/22	14/22	7/22

reference model at the same inputs for 10^3 samples with the reduction in the number of discipline couplings. With the exception of the $TSFC$ output variable in the final model (Figure 39), the trends for the output variables are closely captured by the decoupled models associated with $\lambda = 0.1$ and 1. Nevertheless, users with different requirements may still trade off model sparsity and accuracy for performing computations, such as multidisciplinary optimization, with these two decoupled models. While the optimal model for $\lambda = 1$ eliminates one feedback loop from the low pressure turbine leading to lower computational costs than the model for $\lambda = 0.1$, there is a trade off with accuracy, as seen with the larger spread in the pointwise errors of the $TSFC$ variable.

Furthermore, these results for the pointwise error highlight the importance of the bypass line to accurately estimate the $TSFC$ when computing the turbine's performance. We note that this coupling is included in the models corresponding to $\lambda = 0.1$ and 1.

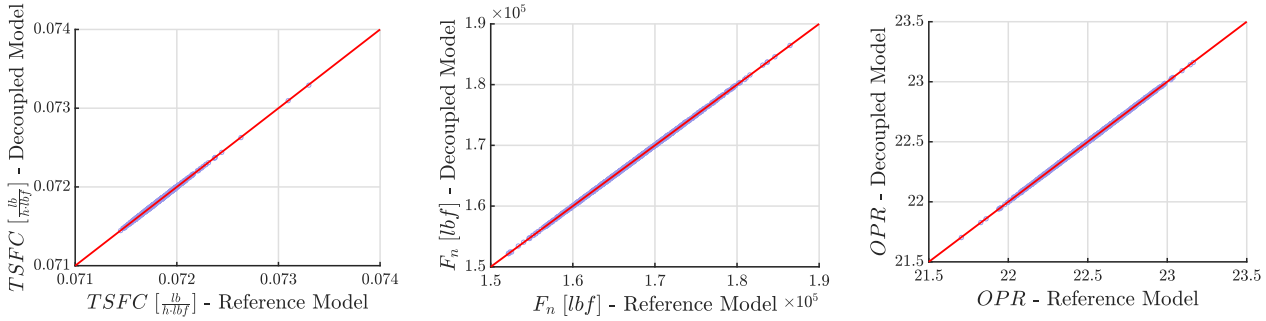

Fig. 37 Comparison of QoI with the optimal decoupled model for $\lambda = 0.1$

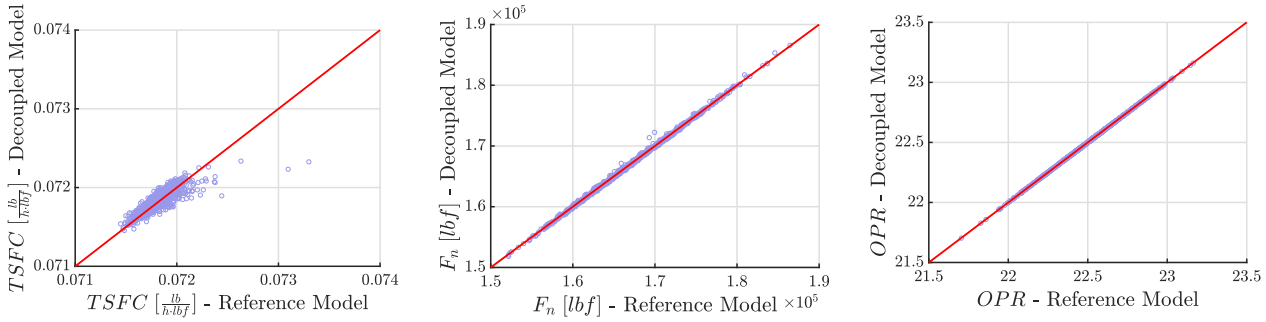

Fig. 38 Comparison of QoI with the optimal decoupled model for $\lambda = 1$

Finally, in addition to producing a sequence of sparsely coupled models with a reduced number of discipline

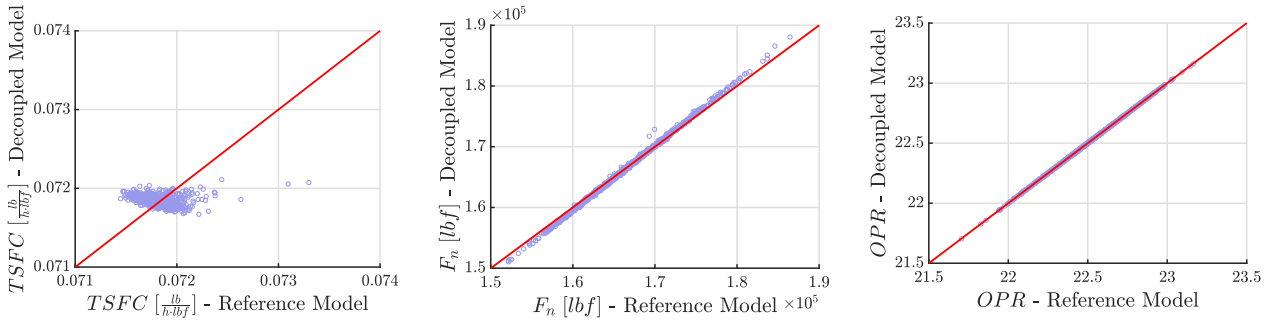

Fig. 39 Comparison of QoI with the optimal decoupled model for $\lambda = 10$

couplings relative to M_0, the model selection algorithm efficiently found all of the optimally decoupled models by visiting less than 5.7% of the model space. This is contrasted with the otherwise intractable process of enumerating all possible models to find the optimum of the combinatorial optimization problem posed in equation (3).

C. Numerical Implementation

A publicly available MATLAB implementation of Algorithm 1 to find the optimal approximate couplings for a multidisciplinary model can be found at `www.github.com/baptistar/coupling`. The application of the algorithm to the fire detection satellite model presented in Section A is publicly available for reference.

V. Conclusion

This paper proposed a method to identify important discipline couplings in multidisciplinary models under uncertainty. Using a combination of model linearizations and sequential Monte Carlo we developed an algorithm to efficiently explore the combinatorial model space of decoupled models and return a model that optimally trades off reduced couplings with accuracy in the characterizations of specified quantities of interest. This algorithm was applied and validated with two engineering problems: a fire detection satellite model and a turbine engine cycle analysis model.

There are two main conclusions that we can draw from this work. First, to identify an approximation to the model coupling with our approach, it is necessary to solve a combinatorial optimization problem. To do so, we presented an algorithm based on sequential Monte Carlo in Section C for solving the optimization problem that is scalable to an increasing number of discipline couplings. This algorithm can be combined with model linearizations to capture a coarse approximation to the output uncertainty, that is often sufficient

for performing model selection in engineering problems without perfectly estimating a high-dimensional distribution, which is not the goal of this work.

Second, from the numerical results we conclude that using a reduced number of discipline couplings can often provide a close approximation to the model outputs while yielding a reduction in the computational cost for each model evaluation. This cost reduction is most evident when neglecting feedback couplings. Nevertheless, the optimal coupling found in all problems is dependent on the input uncertainty, as observed when changing the random inputs to the fire-detection satellite model. Therefore, we recommend reassessing the accuracy of the decoupled model when changing the input variable uncertainty. If necessary, a user should reprocess the optimal model coupling by rerunning the optimization algorithm to identify these model changes.

In addition to the continued application of the algorithm to reveal the natural structure of complex engineering systems under different input uncertainties, future research will address improvements to the measure of model sparsity, $\mathcal{P}(M)$. For example, it may be useful to employ a metric that accounts for the computational cost of each coupling, and that favors feed-forward models and/or parallelism across different disciplines. Lastly, we will explore the use of decoupled models as physically motivated preconditioners to accelerate the iterative solution of reference multidisciplinary models. In particular, connections between decoupling and recent research in nonlinear preconditioning [41, 42] constitute a promising area of future work.

Appendix: Details of SMC Algorithm

In this section, we describe in more detail the Sequential Monte Carlo steps in Algorithm 2. These include fitting the proposal distribution, the resample and move steps, and updating the tempering parameter.

Fitting the Proposal Distribution

To move the models at each iteration, we construct a parametric proposal, $q_{\mathbf{B}} : \mathcal{M} \to [0, 1]$, that approximates the correlation structure of each distribution, P_t. To easily sample new models (i.e., sequentially sample each component of a binary vector $\mathbf{z} \in \{0, 1\}^d = \mathcal{M}$), we consider a proposal distribution that factors as

$$q_{\mathbf{B}}(\mathbf{z}) = q_{\mathbf{B}}(z_1) \prod_{i=2}^{d} q_{\mathbf{B}}(z_i | z_{1:i-1}), \tag{14}$$

where the conditional distribution for each entry, z_i, is represented by a logistic regression model given by equation (15) where $\ell(p) = \log(p) - \log(1 - p)$ and $\mathbf{B} \in \mathbb{R}^{d \times d}$ is a lower triangular matrix that collects all

parameters in the proposal distribution.

$$\ell(q_{\mathbf{B}}(z_i = 1|z_{1:i-1})) = \mathbf{B}_{ii} + \sum_{j=1}^{i-1} \mathbf{B}_{ij} z_j. \tag{15}$$

To determine the unknown parameters, $\mathbf{B}$, we maximize the log-likelihood function for the models and weights at each iteration, $\mathcal{L}(\mathbf{M}_t, \mathbf{w}_t) := \sum_{l=1}^{L} \log(q_{\mathbf{B}}(\mathbf{z} = M_t^{(l)}))$, by solving the first-order condition $\frac{\partial \mathcal{L}(\mathbf{M}_t, \mathbf{w}_t)}{\partial \mathbf{B}} = \mathbf{0}$. In practice, this requires finding a numerical solution with an iterative Newton's method by evaluating the first and second order derivatives of the log-likelihood with respect to the parameters in $\mathbf{B}$. We refer the reader to Ref. [19] for methods on accelerating the convergence of this optimization problem, including how to choose the starting point for the optimization.

Resample and Move Steps

After fitting the proposal, we resample and move the models at each iteration. Following the study in Ref. [37], we use systematic resampling to draw L models from the weighted empirical distribution given by $P_t(M) \approx \sum_{l=1}^{L} w_t^{(l)} \delta_{M_t^{(l)}}(M)$ and replace it by a set of unweighted models. When resampling, models that are closer to the optimum and have larger weights are generated multiple times, while those with smaller weights typically vanish.

To avoid particle degeneracy from repeated weighting and resampling steps, we move the particles at iteration t according to a Markov transition kernel, κ_t. This kernel is chosen to have invariant measure P_t so that by drawing samples $\hat{M}_t^{(l)} \sim \kappa(\cdot|M_t^{(l)})$, the new set of models are still approximately distributed according to P_t. Furthermore, after a sufficient number of move steps, these models are independent of the initial models.

For the sampled models to have invariant measure, P_t, we use a Metropolis-Hastings kernel for κ_t. We sample from this kernel by proposing a new model $\hat{M}_t^{(l)} \sim q_{\mathbf{B}}(\mathbf{z})$ and accepting it with the probability $\gamma(M_t^{(l)}, \hat{M}_t^{(l)})$ in equation (16) or otherwise returning to $M_t^{(l)}$. As empirically demonstrated in Ref. [37], this kernel with adapted parameters is rapidly mixing and results in high average acceptance rates for new models and fast exploration of the model space $\mathcal{M}$. The move step is repeated until the particle diversity defined in equation (13) is above a certain threshold (chosen in practice to be $\zeta \approx 0.95$) or it is no longer changing.

$$\gamma(M_t^{(l)}, \hat{M}_t^{(l)}) := \min\left(1, \frac{P_t(\hat{M}_t^{(l)}) q_{\mathbf{B}}(M_t^{(l)})}{P_t(M_t^{(l)}) q_{\mathbf{B}}(\hat{M}_t^{(l)})}\right) \tag{16}$$

Updating the Tempering Parameter

Lastly, with the weights at each step depending continuously on the parameter $\rho(t)$ and this having a direct influence on the ESS, we monitor the ESS to schedule $\rho(t)$ at each iteration. In particular, the tempering parameter is updated to be $\rho(t+1) = \rho(t) + \alpha_t$ where $\alpha_t \in \mathbb{R}$ is chosen such that $ESS(\mathbf{w}_{t+1}) = 0.9ESS(\mathbf{w}_t)$ in accordance with Ref. [37]. By limiting the decrease in effective sample size with respect to the current distribution, this encourages consecutive distributions to be close and ensures there is a smooth transition between P_t at each iteration.

Acknowledgments

This work has been supported in part by the Air Force Office of Scientific Research (AFOSR) MURI on "Managing multiple information sources of multi-physics systems," Program Officer Jean-Luc Cambier, Award Number FA9550-15-1-0038. The authors also thank T. Hearn and J. Gray for use of the turbine engine cycle analysis model.

References

[1] Cramer, E. J., Dennis, J., Jr, Frank, P. D., Lewis, R. M., and Shubin, G. R., "Problem formulation for multidisciplinary optimization," *SIAM Journal on Optimization*, Vol. 4, No. 4, 1994, pp. 754–776.

[2] Martins, J. R. R. A., and Hwang, J. T., "Review and unification of methods for computing derivatives of multidisciplinary computational models," *AIAA Journal*, Vol. 51, No. 11, 2013, pp. 2582–2599.

[3] Braun, R., Gage, P., Kroo, I., and Sobieski, I., "Implementation and performance issues in collaborative optimization," *6th Symposium on Multidisciplinary Analysis and Optimization*, AIAA, 1996.

[4] Bloebaum, C. L., Hajela, P., and Sobieszczanski-Sobieski, J., "Non-hierarchic system decomposition in structural optimization," *Engineering Optimization*, Vol. 19, No. 3, 1992, pp. 171–186.

[5] Sellar, R., Batill, S., and Renaud, J., "Response surface based, concurrent subspace optimization for multidisciplinary system design," *34th Aerospace Sciences Meeting and Exhibit*, 1996, p. 714.

[6] Sobieszczanski-Sobieski, J., Agte, J. S., and Sandusky, R. R., "Bilevel integrated system synthesis," *AIAA Journal*, Vol. 38, No. 1, 2000, pp. 164–172.

[7] Kim, H. M., Michelena, N. F., Papalambros, P. Y., and Jiang, T., "Target cascading in optimal system design," *Journal of Mechanical Design*, Vol. 125, No. 3, 2003, pp. 474–480.

[8] Haftka, R. T., "Optimization of flexible wing structures subject to strength and induced drag constraints," *AIAA Journal*, Vol. 15, No. 8, 1977, pp. 1101–1106.

[9] Sobieszczanski-Sobieski, J., and Haftka, R. T., "Multidisciplinary aerospace design optimization: survey of recent developments," *Structural Optimization*, Vol. 14, No. 1, 1997, pp. 1–23.

[10] Kroo, I., Altus, S., Braun, R., Gage, P., and Sobieski, I., "Multidisciplinary optimization methods for aircraft preliminary design," *5th Symposium on Multidisciplinary Analysis and Optimization*, AIAA, 1994, p. 4325.

[11] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "High-fidelity aerostructural design optimization of a supersonic business jet," *Journal of Aircraft*, Vol. 41, No. 3, 2004, pp. 523–530.

[12] McAllister, C. D., and Simpson, T. W., "Multidisciplinary robust design optimization of an internal combustion engine," *Journal of Mechanical Design*, Vol. 125, No. 1, 2003, pp. 124–130.

[13] Fuglsang, P., and Madsen, H. A., "Optimization method for wind turbine rotors," *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 80, No. 1, 1999, pp. 191–206.

[14] Braun, R. D., Moore, A. A., and Kroo, I. M., "Collaborative approach to launch vehicle design," *Journal of Spacecraft and Rockets*, Vol. 34, No. 4, 1997, pp. 478–486.

[15] Keyes, D., "Computational Science," *The Princeton Companion to Applied Mathematics*, Princeton University Press, 2015.

[16] Steward, D. V., "The design structure system: A method for managing the design of complex systems," *IEEE transactions on Engineering Management*, , No. 3, 1981, pp. 71–74.

[17] Lu, Z., and Martins, J. R. R. A., "Graph partitioning-based coordination methods for large-scale multidisciplinary design optimization problems," *12th AIAA Aviation Technology Conference*, 2012.

[18] Keane, A., and Nair, P., *Computational Approaches for Aerospace Design: The Pursuit of Excellence*, Wiley, 2005.

[19] Schäfer, C., and Chopin, N., "Sequential Monte Carlo on large binary sampling spaces," *Statistics and Computing*, Vol. 23, No. 2, 2013, pp. 163–184.

[20] Zhou, Y., Johansen, A. M., and Aston, J. A., "Towards automatic model comparison: an adaptive sequential Monte Carlo approach," *Journal of Computational and Graphical Statistics*, 2015.

[21] Brevault, L., Balesdent, M., Bérend, N., and Le Riche, R., "Decoupled multidisciplinary design optimization formulation for interdisciplinary coupling satisfaction under uncertainty," *AIAA Journal*, Vol. 54, No. 1, 2015, pp. 186–205.

[22] Smith, R. C., *Uncertainty Quantification: Theory, Implementation, and Applications*, Vol. 12, SIAM, 2013.

[23] Amaral, S., Allaire, D., and Willcox, K., "A decomposition-based approach to uncertainty analysis of feed-forward multicomponent systems," *International Journal for Numerical Methods in Engineering*, Vol. 100, No. 13, 2014, pp. 982–1005.

[24] Sankararaman, S., and Mahadevan, S., "Likelihood-based approach to multidisciplinary analysis under uncertainty," *Journal of Mechanical Design*, Vol. 134, No. 3, 2012.

[25] Arnst, M., Ghanem, R., Phipps, E., and Red-Horse, J., "Dimension reduction in stochastic modeling of coupled problems," *International Journal for Numerical Methods in Engineering*, Vol. 92, No. 11, 2012, pp. 940–968.

[26] Arnst, M., Soize, C., and Ghanem, R., "Hybrid sampling/spectral method for solving stochastic coupled problems," *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 1, No. 1, 2013, pp. 218–243.

[27] Arnst, M., Ghanem, R., Phipps, E., and Red-Horse, J., "Reduced chaos expansions with random coefficients in reduced-dimensional stochastic modeling of coupled problems," *International Journal for Numerical Methods in Engineering*, Vol. 97, No. 5, 2014, pp. 352–376.

[28] Jin, R., Chen, W., and Simpson, T., "Comparative studies of metamodelling techniques under multiple modelling criteria," *Structural and Multidisciplinary Optimization*, Vol. 23, No. 1, 2001, pp. 1–13.

[29] Chaudhuri, A., Lam, R., and Willcox, K., "Multifidelity Uncertainty Propagation via Adaptive Surrogates in Coupled Multidisciplinary Systems," *AIAA Journal*, 2017. To appear.

[30] MacKay, D. J., *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, 2003.

[31] Baptista, R., "Optimal Approximations of Coupling in Multidisciplinary Models," Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, 2017.

[32] Mahadevan, S., and Haldar, A., "Probability, reliability and statistical method in engineering design," , 2000.

[33] Jameson, A., "Aerodynamic design via control theory," *Journal of Scientific Computing*, Vol. 3, No. 3, 1988, pp. 233–260.

[34] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "A coupled-adjoint sensitivity analysis method for high-fidelity aero-structural design," *Optimization and Engineering*, Vol. 6, No. 1, 2005, pp. 33–62.

[35] Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., and Protasi, M., *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*, Springer Berlin Heidelberg, 2012.

[36] Del Moral, P., Doucet, A., and Jasra, A., "Sequential Monte Carlo samplers," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, Vol. 68, No. 3, 2006, pp. 411–436.

[37] Schäfer, C., "Particle algorithms for optimization on binary spaces," *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, Vol. 23, No. 1, 2013, p. 8.

[38] Hearn, T., Hendricks, E., Chin, J., Gray, J., and Moore, K. T., "Optimization of Turbine Engine Cycle Analysis with Analytic Derivatives," *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2016, p. 4297.

[39] Browning, T. R., "Applying the design structure matrix to system decomposition and integration problems: a review and new directions," *IEEE Transactions on Engineering Management*, Vol. 48, No. 3, 2001, pp. 292–306.

[40] Gray, J., Mader, C. A., Kenway, G. K., and Martins, J. R. R. A., "Approach to Modeling Boundary Layer Ingestion using a Fully Coupled Propulsion-RANS Model," *58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2017, p. 1753.

[41] Cai, X.-C., and Keyes, D. E., "Nonlinearly preconditioned inexact Newton algorithms," *SIAM Journal on Scientific Computing*, Vol. 24, No. 1, 2002, pp. 183–200.

[42] Liu, L., and Keyes, D. E., "Field-split preconditioned inexact Newton algorithms," *SIAM Journal on Scientific Computing*, Vol. 37, No. 3, 2015, pp. A1388–A1409.