Multifidelity Modeling for Uncertainty Quantification and Optimization in Design of Complex Systems

Karen Willcox
Joint work with Doug Allaire, Andrew March, Leo Ng

7th Research Consortium for Multidisciplinary System Design Workshop
Purdue University
July 19, 2012

Co-Authors

Doug Allaire

Andrew March

Leo Ng

Outline

- Building multifidelity models
  - Surrogate and reduced-order models
- Using multifidelity models
  - Multifidelity model management
- Conclusions, challenges and outlook

Uncertainty quantification and optimization of large-scale complex systems

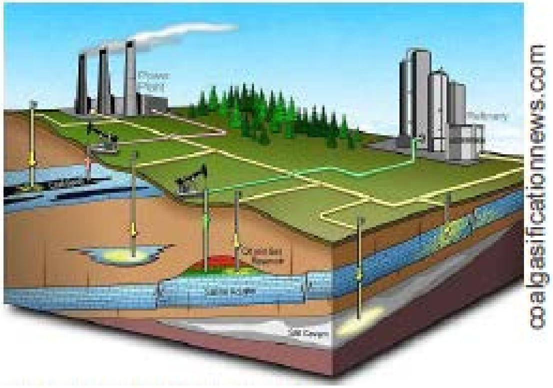

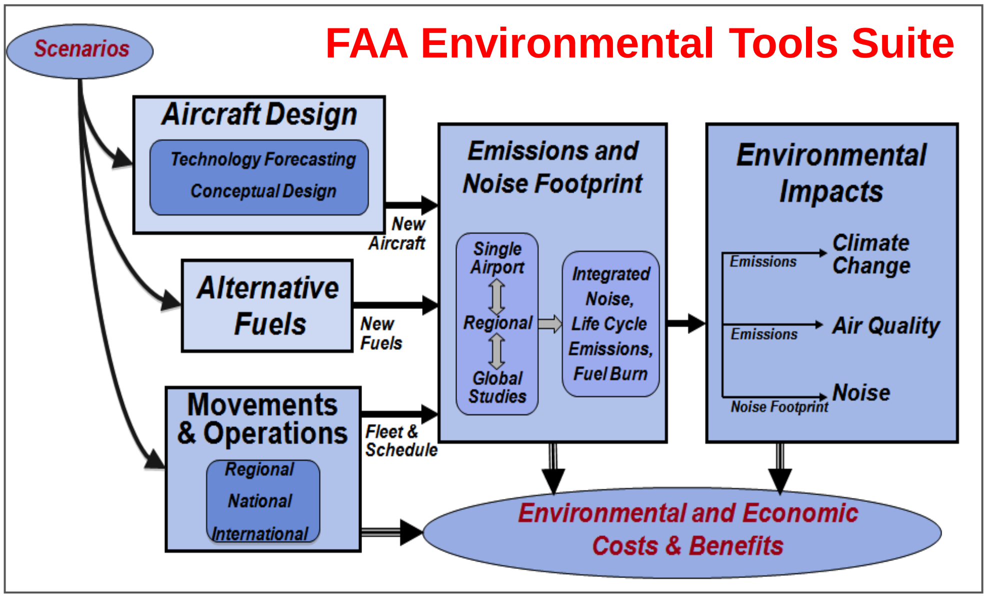

Main message

- Seemingly intractable challenges of identification, prediction and decision—all under uncertainty—for large-scale complex systems can be overcome if

  we use approaches that are **teleological**† and **structure-exploiting**

† *of or pertaining to teleology, the philosophical doctrine that final causes, design, and purpose exist in nature*

From Ancient Greek τέλος (telos, "purpose") + λόγος (logos, "word, speech, discourse")

(http://en.wiktionary.org)

Multifidelity modeling

Often have available several physical and/or numerical models that describe a system of interest.

- Models may stem from different resolutions, different assumptions, surrogates, approximate models, etc.

Multifidelity approaches: How should we best use all available models and data in concert to achieve:

- Better decision-making (optimization, control, design, policy-making)
- Better understanding of modeling limitations $\rightarrow$ guidance for model development

Multifidelity modeling: Ingredients

- Multifidelity model construction
  - Building surrogate, hierarchical or competing models → exploiting structure
- Quantification of uncertainty and model fidelity
  - How good is a model *for a given purpose*
- Multifidelity model management
  - Which model to use when
  - Balancing computational cost with result quality
  - Convergence guarantees
  - Model-model and model-data fusion
  - Model adaptation

Surrogate modeling

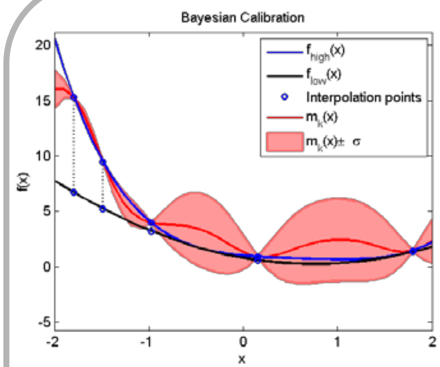

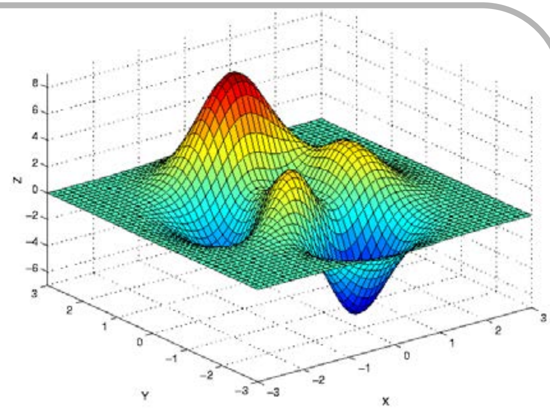

Data-fit models

from Choi et al.

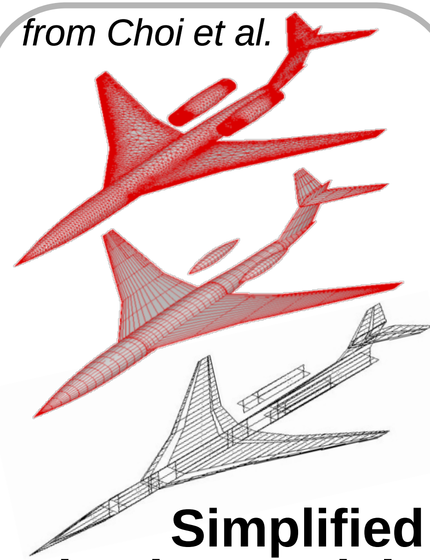

Simplified physics models

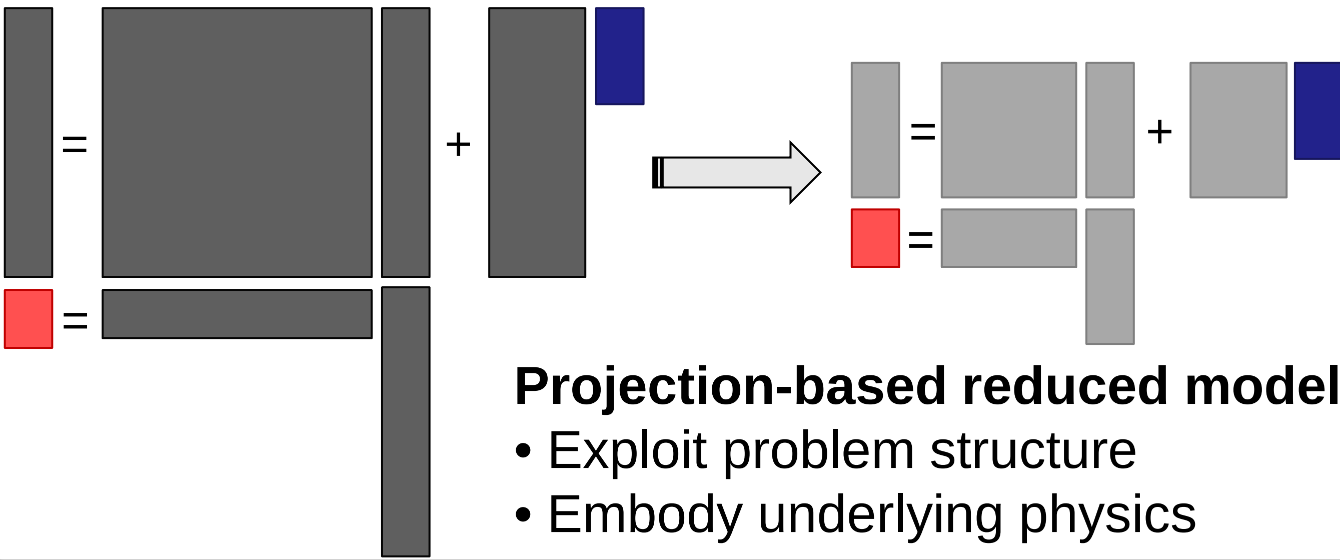

Projection-based reduced models

- Exploit problem structure
- Embody underlying physics

Multifidelity modeling: State of the art

- Focus and progress on deriving surrogates:
  - Projection-based model reduction methods (e.g., Krylov-based, POD, balanced truncation, reduced basis, etc.)
  - Recent breakthroughs in model reduction for parametrically varying and nonlinear systems: (Discrete) Empirical Interpolation Method *(Barrault et al., 2004; Chaturantabut & Sorensen, 2010)*
  - Data fit models (e.g., Gaussian process/Kriging)
- Multifidelity strategies for deterministic optimization problems *(Alexandrov, Booker, Dennis, Lewis et al., 1997,2001)*
  - Otherwise, less focus on how to use surrogates (beyond just replacing high-fidelity simulations)
- Many open questions in quantification of uncertainty and multifidelity model management

Multifidelity philosophy: Use cheap models as much as possible; use adaptation of low-fidelity models

Example: Optimization under uncertainty

$$\min_{x} f\big(x, s(x)\big)$$
$$\text{s.t.} \quad g\big(x, s(x)\big) \leq 0$$
$$h\big(x, s(x)\big) = 0$$

Design variables	x
Uncertain parameters	ξ
Model outputs	$y(x, \xi)$
Statistics of model	$s(x)$

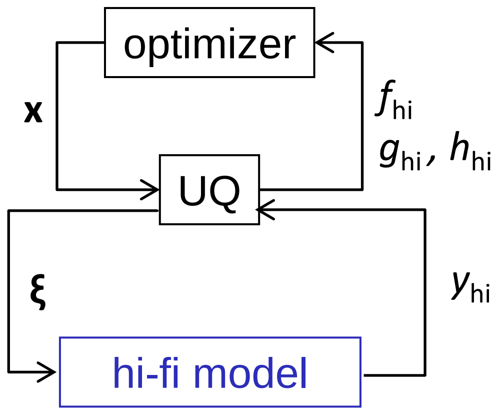

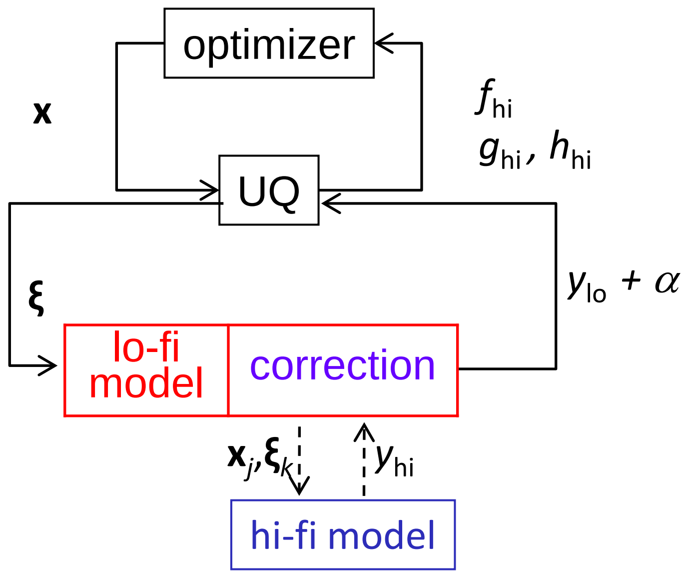

Adaptive corrections: Exploit model local accuracy

- Computed using occasional recourse to the high-fidelity model
- Constructed so that surrogate has desirable properties (e.g., for convergence)

Multifidelity philosophy: Maintain guarantees of convergence with respect to highest-fidelity models

High-fidelity model: $f_{\text{high}}(\mathbf{x})$

Surrogate: $m_k(\mathbf{x}) = f_{\text{low}}(\mathbf{x}) + \alpha_k(\mathbf{x})$

Trust-Region Algorithm for Iteration k

1. Compute a step, $\mathbf{s}_k$, by solving the trust-region subproblem,

$$\min_{\mathbf{s}_k} \quad m_k(\mathbf{x}_k + \mathbf{s}_k)$$
$$\text{s.t.} \quad \|\mathbf{s}_k\| \le \Delta_k.$$

2. Evaluate $f_{\text{high}}(\mathbf{x}_k + \mathbf{s}_k)$.

3. Compute the ratio of actual improvement to predicted improvement,

$$\rho_k = \frac{f_{\text{high}}(\mathbf{x}_k) - f_{\text{high}}(\mathbf{x}_k + \mathbf{s}_k)}{m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)}.$$

4. Accept or reject the trial point according to ρ_k,

$$\mathbf{x}_{k+1} = \begin{cases} \mathbf{x}_k + \mathbf{s}_k & \text{if } \rho_k > 0 \\ \mathbf{x}_k & \text{otherwise.} \end{cases}$$

5. Update the trust region size according to ρ_k,

$$\Delta_{k+1} = \begin{cases} \gamma_1 \Delta_k & \text{if } \rho_k \le \eta_1 \\ \Delta_k & \text{if } \eta_1 < \rho_k < \eta_2 \\ \gamma_2 \Delta_k & \text{if } \rho_k \ge \eta_2. \end{cases}$$

Trust-Region Model Management

(Alexandov, Lewis, et al., 1997, 2001)

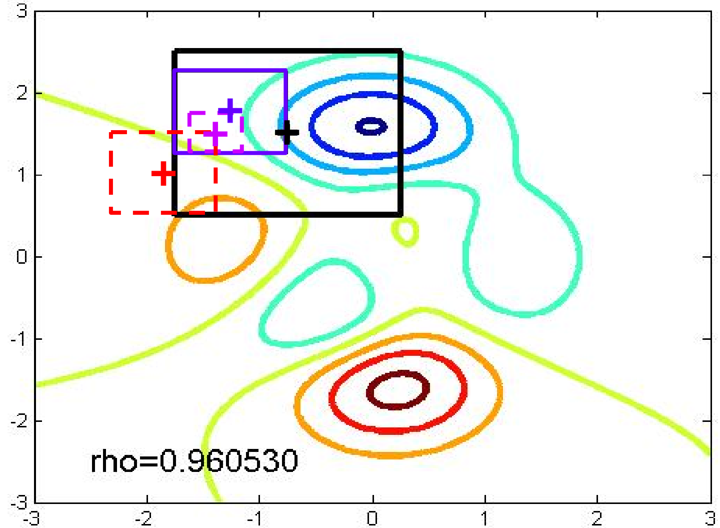

- Provably convergent if surrogate is at least first-order consistent at center of trust region or "fully linear" in gradient-free case *(Conn et al., 2001)*
- Achieved through adaptive corrections or adaptive calibration

Multifidelity philosophy: Use cheap models as much as possible

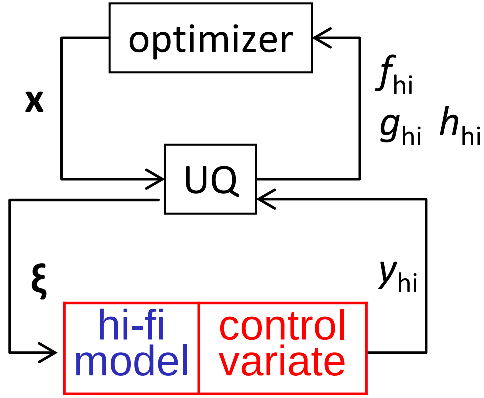

Control variates: Exploit model correlation

- Estimate correlation between high- and low-fidelity models
  → reduce high-fidelity samples needed at optimization iterations

Monte Carlo estimate $\hat{s}(x)$ for statistics at given design point

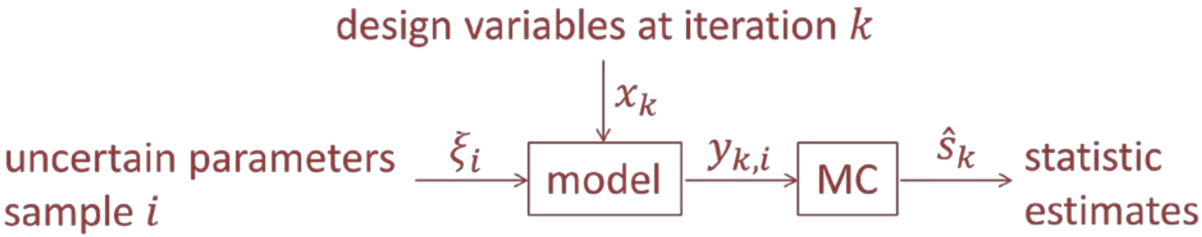

Single fidelity

Control variate estimator of s:

$$\hat{s} = \bar{s} + \beta\left(t_{ref} - \bar{t}\right)$$

- $\bar{s}$: hi-fi model MC estimate
- β: control parameter
- $\bar{t}$: lo-fi model MC estimate

Multifidelity

Multifidelity philosophy: Use high-fidelity models to complement rather than supplant low-fidelity results

Model fusion: Bayesian update (~Kalman filter)

Combine similar models

Trust models with lower variance

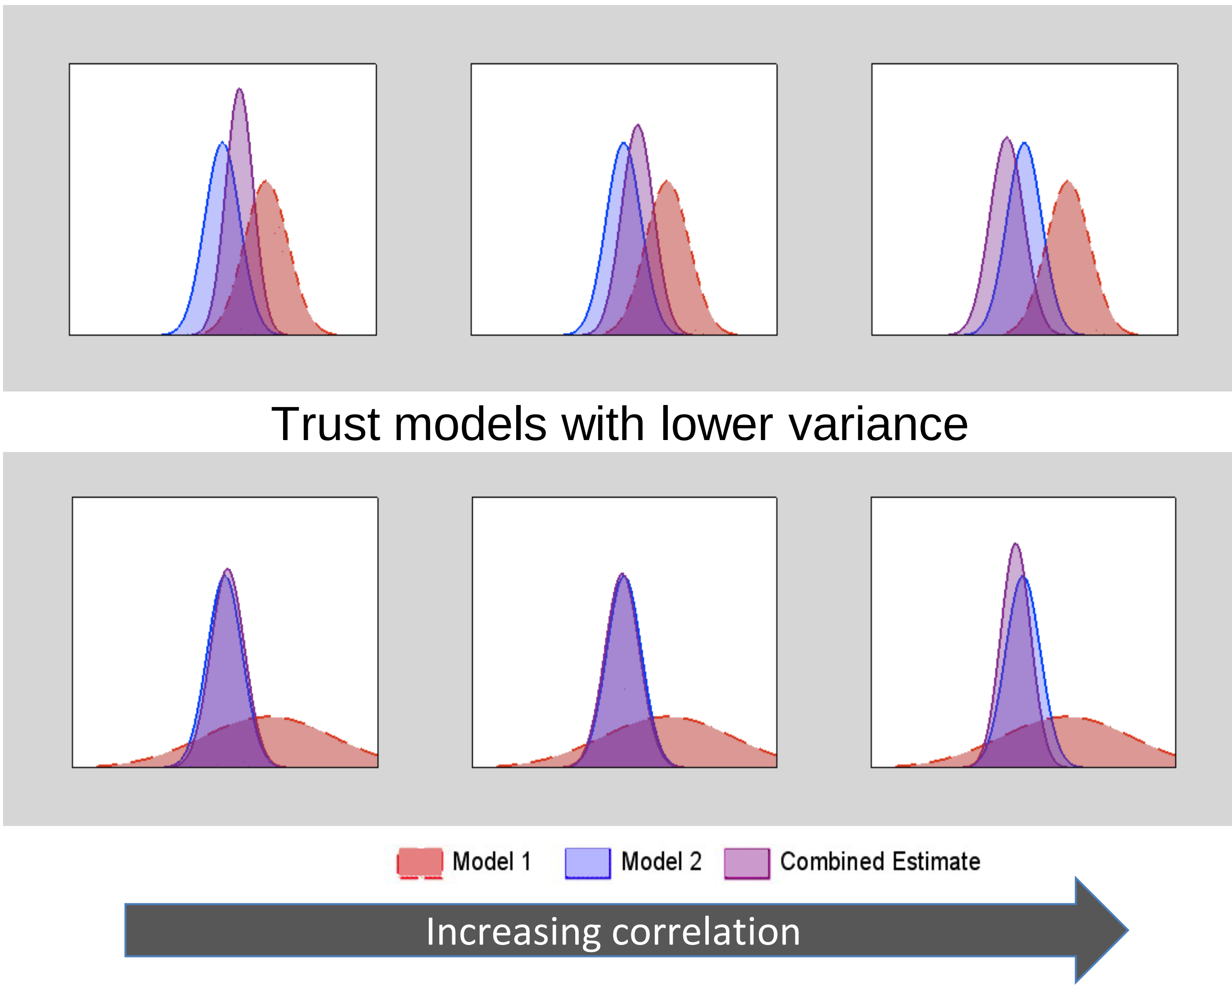

Design under uncertainty example: Acoustic horn

Decision variables: horn geometry, *b*

Uncertainty: wavenumber, wall impedances

Output of interest: reflection coefficient, s_r

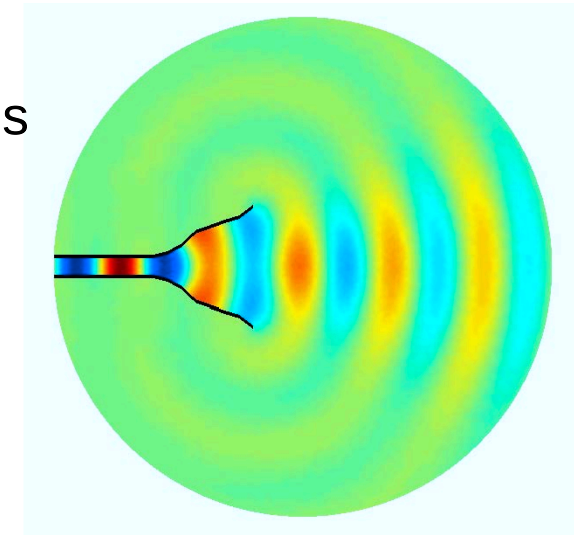

$$\min_b \mathbb{E}[s_r] + \sqrt{\mathbb{V}\mathrm{ar}[s_r]}$$

Multifidelity models:

Finite element model (35,895 states)

Reduced basis model (30 states)

Multifidelity approach:

Control variates

	Equivalent number of hi-fi evaluations	
Regular MC	44,343	
Multifidelity MC	6,979	(-84%)

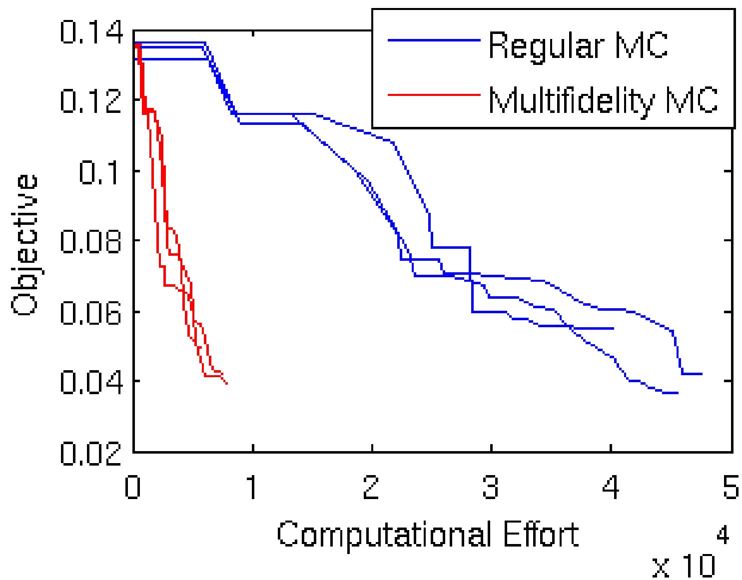

Multidisciplinary design example: Aircraft wing (with black-box codes)

Decision variables: wing geometry, structural members

Disciplines: aerodynamics, structures

Outputs of interest: weight, lift-to-drag ratio

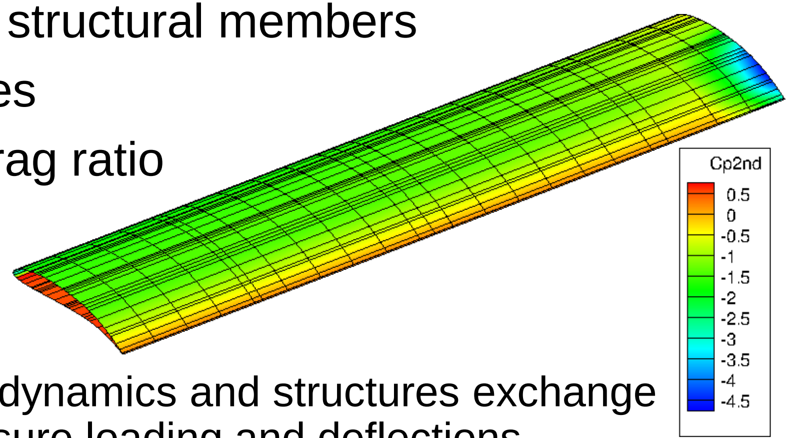

Aerodynamics and structures exchange pressure loading and deflections, requiring an iterative solve for each analysis.

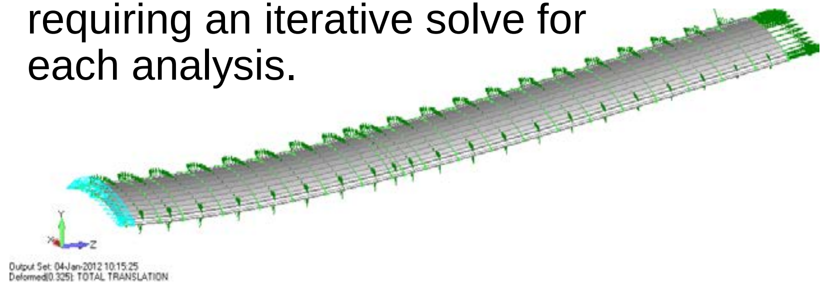

Multifidelity models:

Structures: Nastran (commercial finite element code; MSC)
Beam model

Aerodynamics: Panair (panel code for inviscid flows; NASA)
FRICTION (skin friction and form factors; W. Mason)
AVL (vortex-lattice model; M. Drela)
Kriging surrogate

Multidisciplinary design example: Aircraft wing

Multifidelity approach:

- Trust region model management
  - Derivative free framework *(Conn et al., 2009)*

- Adaptive calibration of surrogates
  - Radial basis function calibration to provide fully linear models *(Wild et al., 2009)*
  - Calibration applied to correction function (difference between high- and low-fidelity models) *(Kennedy & O'Hagan, 2001)*

Low-Fidelity Model	Nastran Evals.	Panair Evals.	Time* (days)
None	7,425	7,425	4.73
AVL/Beam Model	5,412	5,412	3.45
Kriging Surrogate	3,232	3,232	2.06

• Time corresponds to average of 30s per Panair evaluation, 25s per Nastran evaluation, and serial analysis of designs within a discipline.

Exploiting multidisciplinary structure

System
(aerostructural wing design):

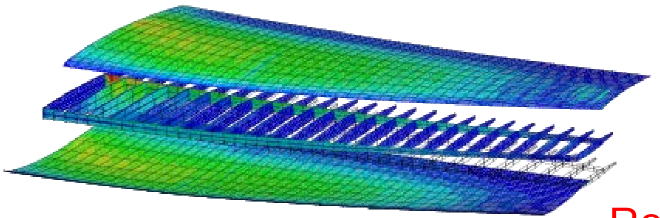

Aerodynamics:

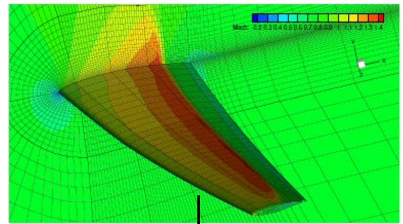

$r_{Cp}(\tilde{\mathbf{x}}^1, \mathbf{t}_{Def})$

Targets: deflection configuration →

← Response: pressure distribution

$$\min_{\mathbf{x}\in\Re^n, \tilde{\mathbf{x}}^1\in\Re^n, \tilde{\mathbf{x}}^2\in\Re^n, \mathbf{t}\in\Re^\nu} w(\mathbf{x})\frac{D(\mathbf{x},\mathbf{t})}{L(\mathbf{x},\mathbf{t})}$$

$$\begin{aligned} s.t. \quad & \sigma_{vM}(\mathbf{x},\mathbf{t}) \le \sigma_{yield} \\ & r_{Cp}(\tilde{\mathbf{x}}^1, \mathbf{t}_{Def}) = \mathbf{t}_{Cp} \\ & r_{Def}(\tilde{\mathbf{x}}^2, \mathbf{t}_{Cp}) = \mathbf{t}_{Def} \\ & \mathbf{x} = \tilde{\mathbf{x}}^1 = \tilde{\mathbf{x}}^2 \end{aligned}$$

Targets: pressure distribution configuration →

← Response: deflection

Structures:

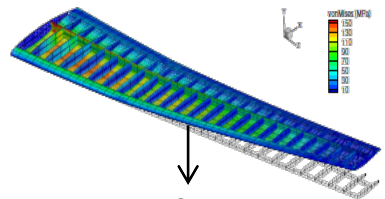

$r_{Def}(\tilde{\mathbf{x}}^2, \mathbf{t}_{Cp})$

Images from: Kenway, Kennedy, and Martins, "A CAD-free approach to high-fidelity aerostructural optimization." AIAA 2010-9231 (MAO 2010).

Multidisciplinary feasible (MDF)

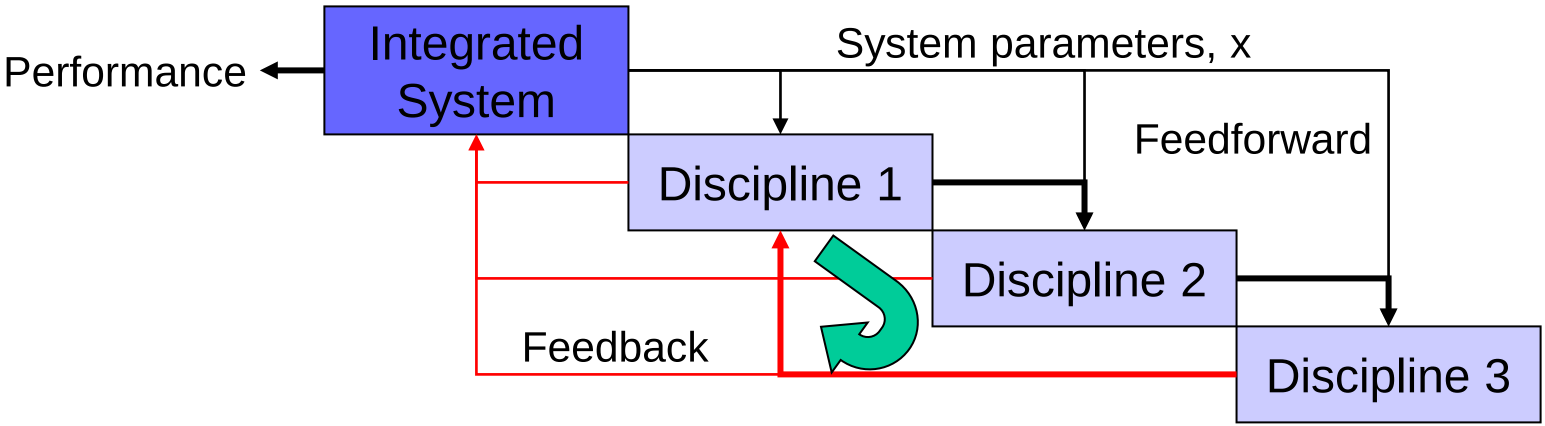

- Feasibility requirements:
  - Internals of each discipline are feasible (i.e., PDEs are solved)
  - Feedback loops are all “closed”
- Each system performance estimate requires an iterative solve
  - Costly when not close to optimum
  - Gradient estimate requires full-system solution for each design variable
- Multifidelity methods and parallelization only at the system level

Decoupling *(Cramer et al., 1984)*

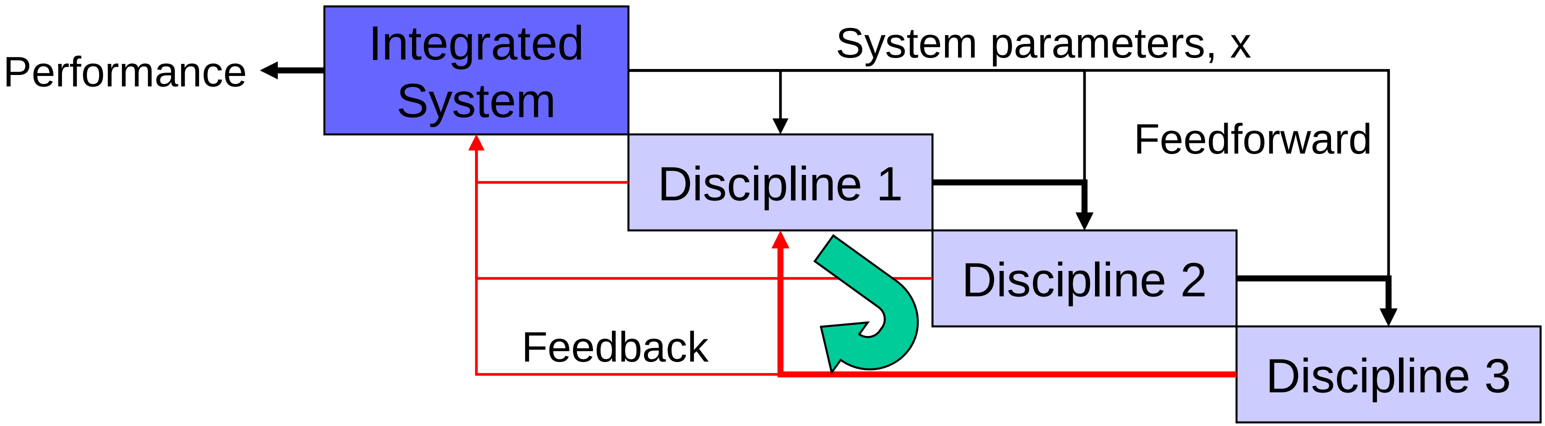

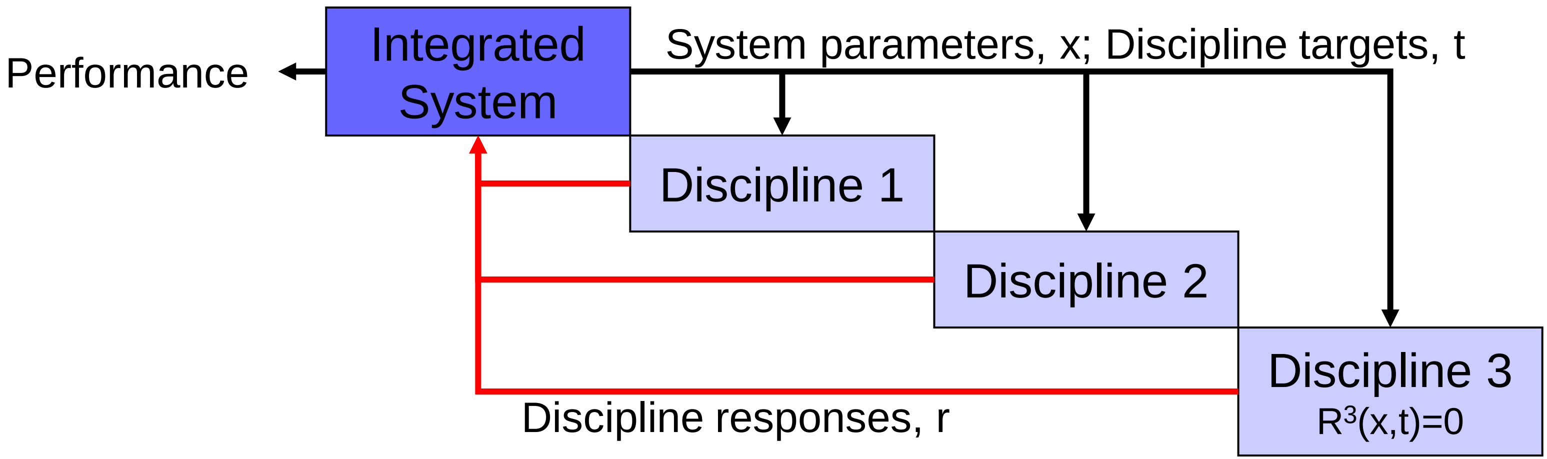

- Individual discipline feasible (IDF): Require **r=t** at convergence
- All-at-once (AAO): Require **r=t, $R^i(x,t)=0$** at convergence

Multifidelity formulations that exploit multidisciplinary problem structure

- MDF formulation
  - Only sees system-level optimization problem
  - Iterative solve for each function evaluation
  - Multifidelity methods and parallelization only at the system level

- IDF formulation
  - Formulate a bi-level optimization problem: system level and disciplinary level
  - Disciplinary optimizations can be done in parallel
  - Disciplinary optimizations can use tailored optimization algorithms (e.g., gradient-based vs. gradient-free)
  - Disciplinary optimizations can exploit discipline-specific multifidelity models
  - Uses Alternating Direction Method of Multipliers to manage the disciplinary interactions

Multidisciplinary design example: Aircraft wing

Multifidelity approach:

- Trust region model management
  - Derivative free framework *(Conn et al., 2009)*

- Adaptive calibration of surrogates
  - Radial basis function calibration to provide fully linear models *(Wild et al., 2009)*
  - Calibration applied to correction function (difference between high- and low-fidelity models) *(Kennedy & O'Hagan, 2001)*

Algorithm	Low-Fidelity Model	Nastran Evals.	Panair Evals.	Time* (days)
Parallel IDF	None	9,073	7,688	2.67
Gradient-free MDF	None	7,425	7,425	4.73
Gradient-free MDF	AVL/Beam Model	5,412	5,412	3.45
Gradient-free MDF	Kriging Surrogate	3,232	3,232	2.06

- Time corresponds to average of 30s per Panair evaluation, 25s per Nastran evaluation, and serial analysis of designs within a discipline.

Conclusions

"All models are wrong, but some are useful."

George Box, 1979

- A formal framework for multifidelity modeling can
  - help us understand when our models are useful
  - provide a new way to think about how to use our wrong-but-useful models for identification, prediction and optimization
- Quantification of uncertainties plays a critical role
  - Many sources of uncertainty in modeling of complex systems
  - Model fidelity ↔ decision task at hand

Acknowledgements

- This work was supported by
  - AFOSR Computational Mathematics Program, AFOSR MURI on Uncertainty Quantification (F. Fahroo)
  - Department of Energy Advanced Scientific Computing Research Program (S. Landsberg)
  - NASA Supersonics Program (N. Alexandrov)
  - Singapore-MIT Alliance Computational Engineering Programme