SAND2019-2458PE

LDRD

Laboratory Directed Research and Development

Exploration of multifidelity UQ strategies for network applications

Gianluca Geraci, Jonathan Crussell, Laura Swiler and Bert Debusschere

SECURE LDRD Grand Challenge
External Advisory Board
March 5th, 2019

Multifidelity Uncertainty Quantification

UQ context at a glance:

- High-dimensionality, non-linearity and bifurcations/discontinuities
- Large set of modeling choices available (network topology, operative conditions, etc.)

Natural candidate:

- **Sampling**-based (MC-like) approaches because they are **non-intrusive**, **robust** and **flexible**...
- **Drawback:** Slow convergence $\mathcal{O}(N^{-1/2}) \rightarrow$ many realizations to build reliable statistics

Goal of MF UQ:

Reducing the computational cost of obtaining MC reliable statistics by combining several models

Pivotal idea:

- Simplified (**low-fidelity**) models are **inaccurate** but **computationally inexpensive**
  $\Rightarrow$ **low-variance** estimates
- **High-fidelity** models are **costly**, but **accurate**
  $\Rightarrow$ **low-bias** estimates

Let consider a **random variable Q**, we want to compute **its expected value $\mathbb{E}[Q]$** (or some high-order moment):

$$\hat{Q}_N^{\mathrm{MC}} = \frac{1}{N}\sum_{i=1}^{N} Q^{(i)}$$

Let's use MC to compute the value $\pi = \dfrac{\#\mathrm{Hit}}{N}$

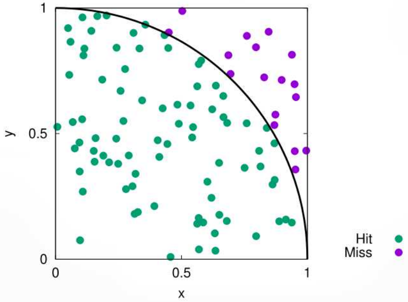

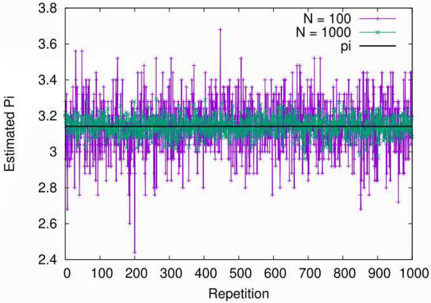

Numerical problems **cannot be resolved with infinite accuracy** (discretization error), the MC estimator for a specific **M**th level

$$\hat{Q}^{MC}_{\mathsf{M},N} \stackrel{\text{def}}{=} \frac{1}{N}\sum_{i=1}^{N} Q_{\mathsf{M}}^{(i)}$$

Let's use MC to compute the value $\pi = \dfrac{\#\text{Hit}}{N}$

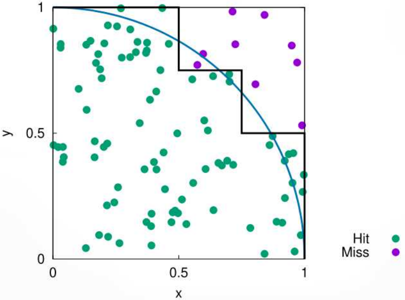

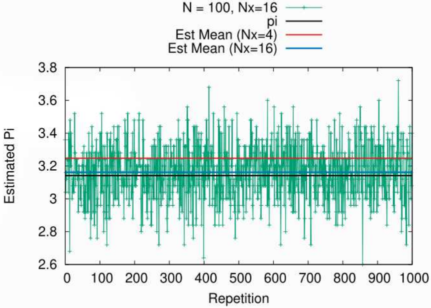

Two sources of error in the **Mean Square Error**:

$$\mathbb{E}\left[(\hat{Q}_{M,N}^{MC} - \mathbb{E}[Q])^2\right] = \mathbb{V}\mathbf{ar}\left(\hat{\mathbf{Q}}_{\mathbf{M},\mathbf{N}}^{\mathbf{MC}}\right) + \left(\mathbb{E}\left[\mathbf{Q}_{\mathbf{M}} - \mathbf{Q}\right]\right)^{\mathbf{2}}$$

- **Sampling error:** replacing the expected value by a (finite) sample average, *i.e.*

$$\mathbb{V}\mathbf{ar}\left(\hat{\mathbf{Q}}_{\mathbf{M},\mathbf{N}}^{\mathbf{MC}}\right) = \frac{\mathbb{V}\mathbf{ar}(\mathbf{Q})}{\mathbf{N}}$$

- **Model fidelity (e.g. discretization):** finite accuracy

Accurate estimation $\Rightarrow$ **Large number** of samples evaluated for the **high fidelity** model

$$\mathbb{E}[Q_M] - \hat{Q}_{M,N}^{MC} \sim \sqrt{\frac{\mathbb{V}\mathbf{ar}(\mathbf{Q}_{\mathbf{M}})}{\mathbf{N}}}\mathcal{N}(0,1)$$

Pivotal idea:

- **High-fidelity** models are **costly**, but **accurate**
  - **low-bias** estimates
- Simplified (**low-fidelity**) models are **inaccurate** but **cheap**
  - **low-variance** estimates

Single Fidelity

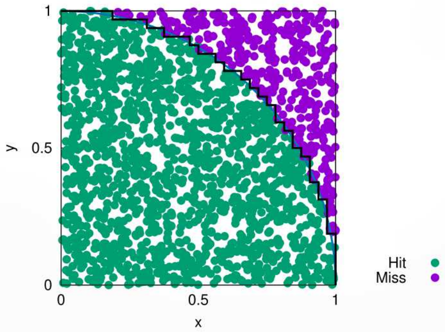

Multi Fidelity

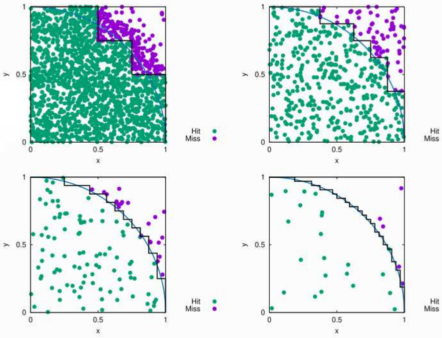

Approximate Control Variates

Optimal Control Variate
M low-fidelity models with known expected value

Let's consider M **low-fidelity models with known mean**. The **Optimal Control Variate (OCV)** is generated by adding **M unbiased terms to the MC estimator**

$$\hat{Q}^{\mathrm{CV}} = \hat{\mathbf{Q}} + \sum_{\mathbf{i}=\mathbf{1}}^{\mathbf{M}} \alpha_{\mathbf{i}} \left(\hat{\mathbf{Q}}_{\mathbf{i}} - \mu_{\mathbf{i}} \right)$$

- $\hat{Q}_i$ **MC estimator** for the ith low-fidelity model
- μ_i **known expected value** for the ith low-fidelity model
- $\underline{\alpha} = [\alpha_i, \ldots, \alpha_M]^{\mathrm{T}}$ **set of weights** (to be determined)

Let's define

- The covariance matrix among all the low-fidelity models: $\mathbf{C} \in \mathbb{R}^{M \times M}$
- The vector of covariances between the high-fidelity Q and each low-fidelity Q_i: $\mathbf{c} \in \mathbb{R}^{M}$
- $\bar{\mathbf{c}} = \mathbf{c} / \mathbb{V}ar(Q)$, where ρ_i is the correlation coefficient (Q, Q_i)

The **variance of the OCV estimator** (optimal weights $\underline{\alpha}^{\star} = -\mathbf{C}^{-1}\mathbf{c}$)

$$\mathbb{V}ar\left(\hat{Q}^{\mathrm{CV}}\right) = \mathbb{V}\mathbf{ar}\left(\hat{\mathbf{Q}}\right)(\mathbf{1} - \mathbf{R}^{\mathbf{2}}_{\mathbf{OCV}}) = \mathbb{V}\mathbf{ar}\left(\hat{\mathbf{Q}}\right)\left(\mathbf{1} - \bar{\mathbf{c}}^{\mathrm{T}}\mathbf{C}^{-\mathbf{1}}\bar{\mathbf{c}}\right), \quad 0 \leq R^2_{OCV} \leq 1.$$

NOTES:

1. For a single low-fidelity model: $R^2_{OCV-1} = \rho_1^2$
2. For all estimators in literature (MLMC, MFMC, *etc.*): $R^2 \leq \rho_1^2 \leq R^2_{OCV}$

For complex engineering models the **expected values of the M low-fidelity models are unknown** *a priori*

- Let's consider N_i LF evaluations: $N_i = \lfloor r_i N \rfloor$

The generic **Approximate Control Variate** is defined as

$$\tilde{Q}\left(\underline{\alpha}, \underline{\mathbf{z}}\right) = \hat{\mathbf{Q}}(\mathbf{z}) + \sum_{\mathbf{i}=\mathbf{1}}^{\mathbf{M}} \alpha_{\mathbf{i}} \Delta_i(\mathbf{z}_i)$$

The **optimal weights and variance** can be obtained as

$$\underline{\alpha}^{\mathrm{ACV}} = -\mathbb{C}ov\left[\underline{\Delta}, \underline{\Delta}\right]^{-1} \mathbb{C}ov\left[\underline{\Delta}, \hat{Q}\right]$$

$$\mathbb{V}ar\left(\tilde{Q}\left(\underline{\alpha}^{ACV}\right)\right) = \mathbb{V}\mathbf{ar}\left(\hat{\mathbf{Q}}\right)\left(\mathbf{1} - \mathbf{R}^{\mathbf{2}}_{\mathrm{ACV}}\right).$$

NOTES:

1. For a single low-fidelity model: $R^2_{ACV-1} = \frac{\mathbf{r_1 - 1}}{\mathbf{r_1}} \rho_1^2$
2. We can build provably optimal estimators: $\rho_1^2 \leq R^2_{ACV} \leq R^2_{OCV}$

Numerical Experiments

Network Configuration

- 1 client - 1 server (possible to extend to multiple clients)
- 100 Requests

Uncertain Parameters

- `DataRate` $\sim \mathcal{U}(5, 500)Mbps$
- `Delay` $\sim \mathcal{U}(1, 3)ms$

Fidelity definition

- HF: `ResponseSize` $16MB$ – runtime $20min$
- LF: `ResponseSize` $1MB$ – runtime $50s$
- LF*: `ResponseSize` $500B$ and 10 Requests – runtime $0.15s$

	$\mathcal{C}$
HF	1
LF	0.0417
LF*	0.000125

TABLE: Normalized Cost

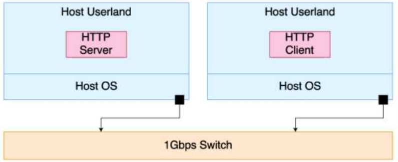

FIGURE: Network Configuration

Requests/second (Expected Value)

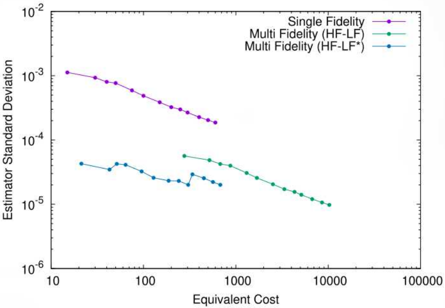

Figure: Estimators Standard Deviation.

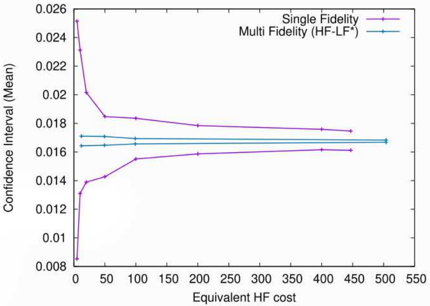

Figure: Confidence Interval convergence

First minimega-NS3 demonstration
Network Configuration: 1 Client - 1 Server

Network Configuration

- 1 client - 1 server (possible to extend to multiple clients)
- 100 Requests

Uncertain Parameters

- `DataRate` $\sim \mathcal{U}(5, 500)Mbps$
- `ResponseSize` $\sim \ln \mathcal{U}(500, 16 \times 10^6)B$

Fidelity definition

- `minimega` – HF: 100 Requests (average over 10 repetitions)
- `ns3` – LF: 10 Requests (`Delay` $50ms$)
- `ns3` – LF*: 1 Requests (`Delay` $5ms$)

	$\mathcal{C}$
HF	1
LF	0.016
LF*	0.002

Table: Normalized Cost

We assume **serial execution for the low-fidelity model**, however we might easily increase the efficiency of LF (`ns3`) by running multiple concurrent evaluations

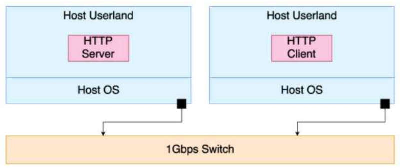

Figure: Network Configuration

FIRST minimega-NS3 DEMONSTRATION
ESTIMATOR STANDARD DEVIATION

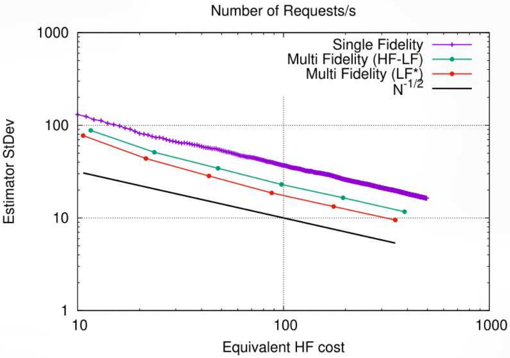

FIGURE: Exp. Value StDev

- The **variance reduction** we obtain w.r.t. MC is

$$\mathbb{V}ar\left(\tilde{Q}\left(\underline{\alpha}^{ACV}\right)\right) = \mathbb{V}ar\left(\hat{Q}\right)\left(1 - \frac{\mathbf{r_1} - \mathbf{1}}{\mathbf{r_1}}\rho_1^2\right)$$

- The **number of low-fidelity simulations** is $N_{LF} = N \times r_1$ where

$$r_1 = \sqrt{\frac{\mathcal{C}_{HF}}{\mathcal{C}_{LF}}\frac{\rho_1^2}{1-\rho_1^2}}$$

- For each HF simulation we need to spend an **extra cost** in LF simulations

$$\text{Eq.Cost}: \quad \mathcal{C}_{tot} = N\left(1 + \mathbf{r_1}\frac{\mathcal{C}_{\mathbf{LF}}}{\mathcal{C}_{\mathbf{HF}}}\right)$$

- For this case

	ρ_1	r_1	$r_1\mathcal{C}_{LF}/\mathcal{C}_{HF}$
LF	0.86	4.69	0.075
LF*	0.90	10.83	0.022

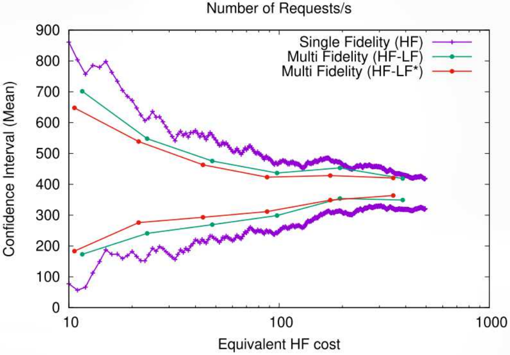

FIGURE: Exp. Value Confidence Interval

- The **variance reduction** we obtain w.r.t. MC is

$$\mathbb{V}ar\left(\tilde{Q}\left(\underline{\alpha}^{ACV}\right)\right) = \mathbb{V}ar\left(\hat{Q}\right)\left(1 - \frac{\mathbf{r_1} - \mathbf{1}}{\mathbf{r_1}}\rho_{\mathbf{1}}^{\mathbf{2}}\right)$$

- The **number of low-fidelity simulations** is $N_{LF} = N \times r_1$ where

$$r_1 = \sqrt{\frac{\mathcal{C}_{HF}}{\mathcal{C}_{LF}}\frac{\rho_1^2}{1-\rho_1^2}}$$

- For each HF simulation we need to spend an **extra cost** in LF simulations

$$\text{Eq.Cost}: \quad \mathcal{C}_{tot} = N\left(1 + \mathbf{r_1}\frac{\mathcal{C}_{\mathbf{LF}}}{\mathcal{C}_{\mathbf{HF}}}\right)$$

- For this case

	ρ_1	r_1	$r_1 \mathcal{C}_{LF}/\mathcal{C}_{HF}$
LF	0.86	4.69	0.075
LF*	0.90	10.83	0.022

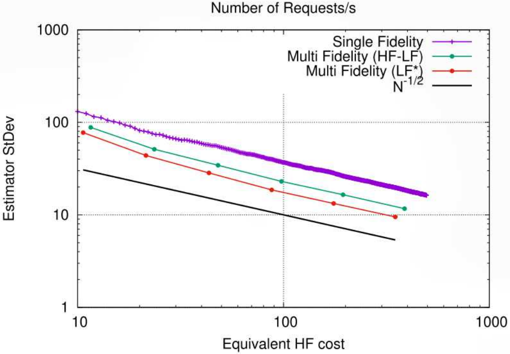

FIGURE: Exp. Value StDev

	ρ_1	r_1	$r_1\mathcal{C}_{LF}/\mathcal{C}_{HF}$
LF	0.86	4.69	0.075
LF*	0.90	10.83	0.022

Example (for LF*)

- Number of **HF runs**: $N = 500$
- Number of **LF* runs**: $r_1 \times N = 5415$
- Equivalent **LF cost**: $r_1 \times N \times \dfrac{\mathcal{C}_{LF}}{\mathcal{C}_{HF}} = 11$
- **Total** estimator **cost** (HF + LF*): $\mathcal{C}_{tot} = 500 + 11 = 511$
- **Variance reduction**: $\left(1 - \dfrac{r_1 - 1}{r_1}\rho_1^2\right) = 0.23$

More than **70% of variance reduction** is obtained by adding **only an equivalent cost of 11 HF runs**

Is it efficient to leverage multiple low-fidelity models at the same time?

	HF	LF	LF*
HF	1	0.86	0.90
LF	0.86	1	0.99
LF*	0.90	0.99	1

TABLE: Correlation matrix

	OCV	ACV
HF+LF	0.26	0.39
HF+LF*	0.19	0.23
HF+LF+LF*	0.08	N/A

TABLE: Variance Reduction, $1 - R^2$

$$\mathbb{V}ar\left(\tilde{Q}\right) = \mathbb{V}ar\left(\hat{Q}\right)\left(1 - R^2\right)$$

NOTE:

- OCV assumes that the LF expected values are known, *i.e.* maximum attainable variance reduction

Concluding Remarks

State-of-the-art

- Multifidelity Uncertainty Quantification proved to be effective for many different applications
- Encouraging preliminary results have been obtained for simple network configurations

Future Directions

- Extension to additional statistics (Tails, risk measures, *etc.*)
- Multifidelity Sensitivity Analysis
- Extension to discrete variables
- Extension to more complex network configurations/topologies
- Exploration of data-driven approaches for LF modelling (model reduction, active directions, *etc.*)
- Exploration of surrogate-based approaches