

A Framework for Early and Approximate Uncertainty Quantification of Large System Simulation Models

Magnus Eek¹, Johan Karlén¹, and Johan Ölvander²

¹Saab Aeronautics, Linköping, Sweden, {magnus.eek, johan.karlen}@saabgroup.com

²Machine Design, IEI, Linköping University, Linköping, Sweden, johan.olvander@liu.se

Abstract

Uncertainty Quantification (UQ) is vital to ensure credibility in simulation results and to justify model-based design decisions – especially in early development phases when system level measurement data for traditional model validation purposes are scarce. Central UQ challenges in industrial applications are computational cost and availability of information and resources for uncertainty characterization. In an attempt to meet these challenges, this paper proposes a framework for early and approximate UQ intended for large simulation models of dynamical systems. A Modelica simulation model of an aircraft environmental control system including a liquid cooling circuit is used to evaluate the industrial applicability of the proposed framework.

Keywords: Uncertainty quantification, aleatory uncertainty, epistemic uncertainty, model validation, aircraft system simulation models, Modelica

1 Introduction

In early phases of aircraft system design, a common situation is that only a simulation model of the system to be developed is available. That is, no physical test rig of the system exists, implying that no system level measurement data for model validation purposes are available. Nevertheless, for taking justified design decisions, an approximate estimation of the uncertainties in the simulation results is needed. Historically, this situation often ends up with a subjective experience-based assessment of simulation results uncertainty. The usefulness of subjective experience-based assessments should not be underestimated, but a more rational response to the need for well-founded estimations of uncertainties would be to apply a structured method for Uncertainty Quantification (UQ). UQ is here understood as the process of identifying, quantifying, and assessing the impact of uncertainty sources embedded along the

development and usage of simulation models. A case study carried out at Saab Aeronautics revealed a large gap between this part of industry and academia regarding the application of UQ (Eek, Kharrazi, Gavel and Ölvander, 2015). There are several reasons why UQ is not widely applied in aircraft system design; one is the always tight project budgets and time schedules and another the fact that many engineers working with Modeling and Simulation (M&S) are not aware of structured methods for UQ or do not know where to begin.

An overview of UQ methods in an industrial context together with a discussion on how model size and information availability affect the complexity of a UQ effort is provided by de Rocquigny, Devictor and Tarantola (2008). The UQ community, with substantial activity in high-consequence areas like nuclear power and climate, but also Computational Fluid Dynamics (CFD) in general, has come far but many techniques found in the literature seem to be suited for fast executing models or models with only a few uncertain parameters. This is rarely the case in M&S of aircraft vehicle systems, and a central challenge to enable UQ is therefore to reduce the dimensionality of the problem, i.e. to reduce the number of uncertain parameters that need to be considered. Another challenge is to find a suitable level of detail when performing UQ. That is, to simplify the UQ as far as possible without affecting the UQ results in such a way that they become irrelevant. In the work presented in this paper, a simulation model of a cooling system in the Saab Gripen fighter is used as the experimental frame for development and evaluation of an approximate and computationally relatively cheap UQ method.

The outline of the paper is as follows. First, the industrial application example is described, followed by an introduction to an established method for UQ. A few methods to reduce the computational cost as well as the engineering workload of UQ are then elaborated. This is followed by a proposal for a framework for early

and approximate UQ, together with experimental results. A discussion and conclusions are presented in the final section.

2 Industrial Application Example

Air supply, cabin pressurization, and avionics cooling are basic requirements for most aircraft, usually facilitated by an Environmental Control System (ECS) taking engine bleed air which is conditioned (decreased in pressure and temperature and dried) prior to distribution. For the Saab Gripen E fighter, a setup with two ECS is under consideration; a Primary ECS (P-ECS) and a Secondary ECS (S-ECS). The main task of the P-ECS is to provide conditioned air for cooling of

the avionics equipment, cabin thermal control and pressurization, pressurization of the fuel and anti-g systems, and to provide conditioned air to the On-Board Oxygen Generating System (OBOGS). Via a liquid cooling system, hereinafter referred to as Liquid Loop (LL), the S-ECS facilitates cooling of tactical systems. A more detailed description of the Saab Gripen ECS including system schematics is given by (Carlsson, Andersson, Gavel and Ölvander, 2012a).

To support system design and to estimate cooling performance, simulation models of the systems are developed and integrated into a total S-ECS model as shown in the schematic description in Figure 1.

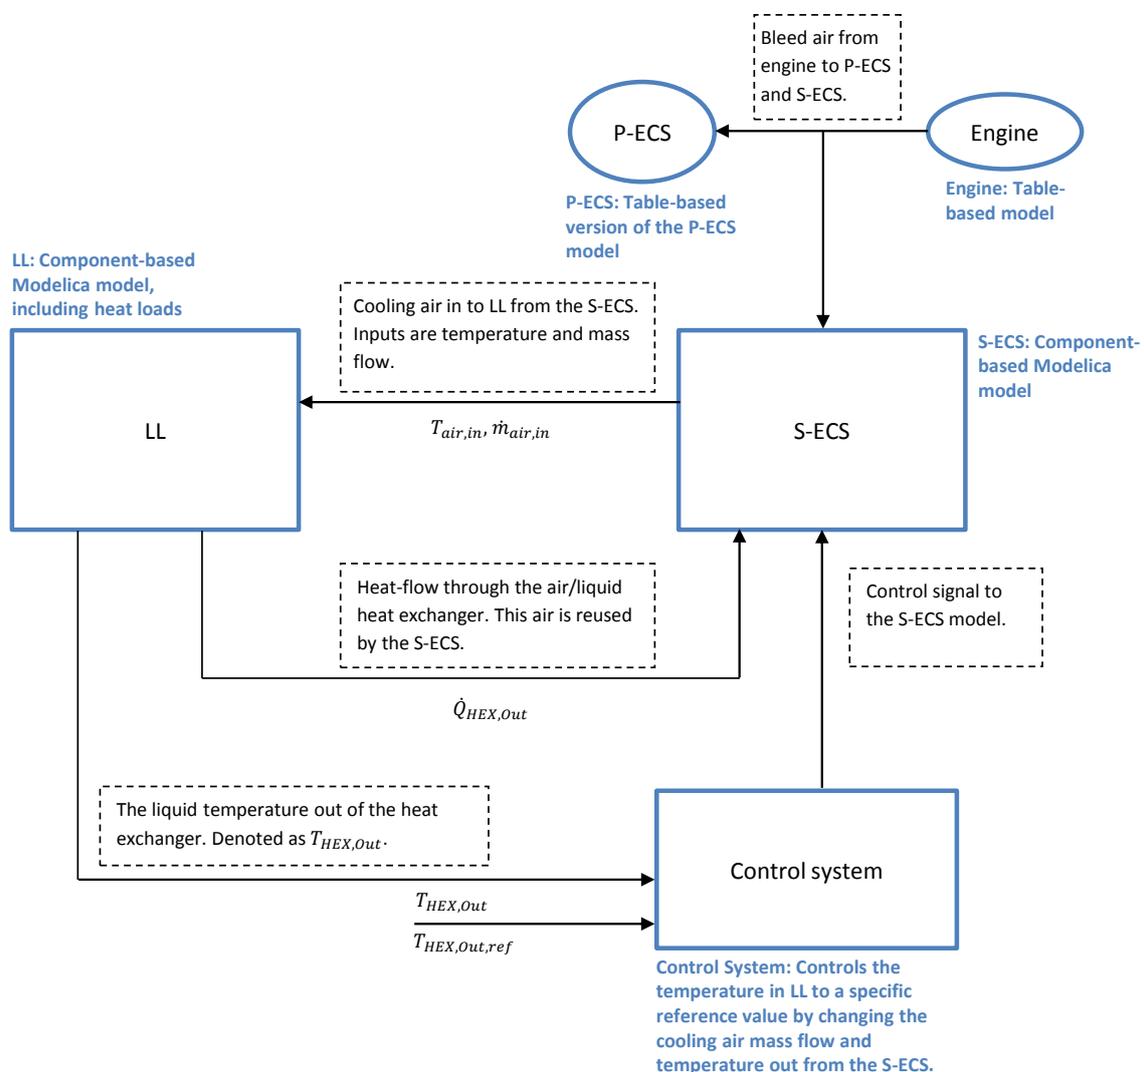


Figure 1. Schematic description of the total S-ECS simulation model.

The simulation model described above is used to estimate cooling performance, both in steady-state operation at specific points in the flight envelope and for full-length flight missions including transient conditions. The simulation

models are developed in the equation-based object-oriented language Modelica¹, using the

¹ <https://modelica.org/>

Dymola tool². Normally, simulation and analysis of results are performed directly in Dymola, but in the development of the UQ method presented in this paper, scripting and post-processing are done in MATLAB³, see Figure 2.

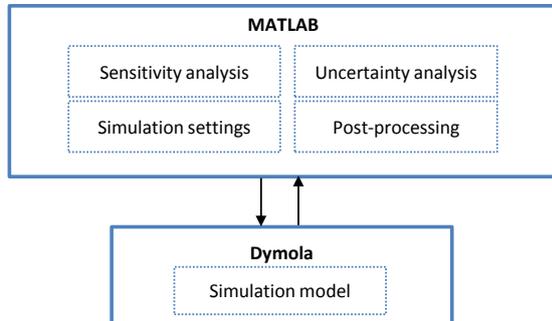


Figure 2. Simulation environment used for development and testing of UQ method.

3 Theoretical Background

Roy and Oberkampf (2011) describe a comprehensive framework for UQ where uncertainties are extensively propagated through a simulation model in order to obtain a detailed uncertainty response. This framework is generally considered to give a well-founded UQ and consists of six steps, as described below.

1. Identify all sources of uncertainty
2. Characterize uncertainties
3. Estimate the uncertainty due to numerical approximations
4. Propagate input uncertainties through the model
5. Estimate model form uncertainty
6. Estimate total uncertainty

The uncertainties can be divided into three categories in order to simplify and structure the identification process. These are input uncertainties, numerical approximations, and model form uncertainties. The input uncertainties can for example stem from model input signals, model parameters, boundary conditions, or initial conditions.

In the first step of the UQ framework proposed by Roy and Oberkampf (2011), all inputs and parameters in a simulation model should be considered as uncertain unless there is strong evidence that their uncertainty does not affect the system response.

In the second step, uncertainties are generally categorized into two different types,

either epistemic or aleatory (Helton, 1996, Helton and Davis, 2003, Oberkampf and Roy, 2012). Epistemic or reducible uncertainties are normally characterized with an interval, a minimum and maximum value, with no probability assigned to possible values. Epistemic uncertainties can be said to derive from lack of knowledge. Aleatory or irreducible uncertainties are stochastic uncertainties that are generally characterized with a Probability Density Function (PDF) or its corresponding Cumulative Distribution Function (CDF). Uncertainties characterized as aleatory therefore contain more information than epistemic ones. Mixed epistemic and aleatory uncertainty may be represented by an imprecise PDF, which here refers to when the parameters of the PDF, for example the mean and standard deviation, are given as intervals or PDFs themselves.

In the third step, the numerical uncertainties due to the solver are estimated. Numerical approximations can be treated as epistemic and added to the system response uncertainty.

When the identification and characterization have been done, the input uncertainties are propagated through the simulation model using a sampling technique. The idea of the sampling is to repeatedly take a random set of uncertainties and simulate the model. These simulations will give response deviations and hence the simulation results' uncertainty due to input uncertainties is obtained. Sampling techniques commonly used in industrial applications are Monte Carlo Sampling (MCS) and Latin Hypercube Sampling (LHS). The latter is a stratified sampling technique in which the range of each uncertain variable is divided into l intervals with equal probability. Only one value is drawn from each interval, but in a random order. This method ensures that all intervals of the uncertain variable are represented, but not more than once. This is in contrast to brute force MCS that is totally random with no guarantee of representing the entire variability of the uncertainties (Saltelli, Ratto, Andres, Campolongo, Cariboni, Gatelli et al., 2008).

Roy and Oberkampf (2011) suggest that each combination of the epistemic uncertainties is propagated with LHS in an outer loop while all the aleatory uncertainties are propagated with MCS in an inner loop. For m epistemic uncertainties, the recommended number of LHS intervals is m^3 . The number of simulations needed in order to propagate input uncertainties can then be calculated as

² <http://www.3ds.com/products-services/catia/products/dymola>

³ www.mathworks.com/products/matlab/

$$n = m^3 \cdot s \quad (1)$$

where n is the total number of simulations needed, m is the number of epistemic uncertainties, and s is the number of MCS samples of the aleatory uncertainties.

When the propagation of input uncertainties has been done, it is time to estimate the model form uncertainties. These stem from the selected model structure and assumptions and approximations of underlying equations made when developing the model. These uncertainties can be identified and characterized by comparing the simulation results with experimental measurement data in a probabilistic manner.

The total uncertainty in the simulation response consists of uncertainties due to inputs, numerical approximations, and model form. Roy and Oberkampf (2011) describe how these uncertainties can be combined to form a *probability box* for the system response (also known as a *p-box*), which is an interval-valued CDF.

4 Methods for Reduction of UQ Workload and Computational Cost

In the following subsections, a few methods to reduce the engineering workload as well as the computational cost required for UQ are proposed and described in the context of the industrial application example from section 2.

Some important conditions, which have great impact on the time needed for the UQ, both in terms of engineering workload and computational cost, can be stated:

- Limitations in the availability of information required for uncertainty characterization.
- Computational expense of the simulation model.
- The dimensionality of the UQ problem, i.e. the number of uncertainties.
- Number of nominal simulations required. For example, a model-based performance evaluation of a system to be developed may require a large number of nominal simulations, which in turn should be evaluated in terms of UQ.

The above combination of characteristics restricts the choice of UQ method. In early system development phases, the availability of information required for characterization of uncertainties is typically very limited.

Additionally, in early phases the system to be developed may not be completely defined and sometimes several possible system layouts are simulated. Thus, for UQ to be considered in early phases, the computational cost of UQ has to be kept as low as possible.

The comprehensive framework described in section 3 is mathematically stringent but requires many simulations when the number of epistemic uncertainties is large. A simulation of one single stationary operational point using the total model of the Gripen E, shown in Figure 1, takes about 20 minutes to run and can therefore be considered computationally expensive. To estimate the stationary cooling performance of the complete flight envelop (spanned by altitude and Mach number), typically about 60 points are needed. The computational cost combined with the high number of identified uncertainties (over 200 in the LL model alone) makes the UQ method described in Chapter 3 unfeasible. It is therefore of great importance to reduce the computational cost in order to enable a UQ.

4.1 Approximate Characterization of Uncertainties

To mitigate the problem of computational cost, one possibility is to consider a simplified characterization of the uncertainties in order to get rid of the inner sampling loop. A simplification of the characterization would decrease the accuracy of the UQ, but could in practical situations be necessary in order to enable a UQ. This simplification could be made by assuming all uncertainties to be epistemic or by assigning the epistemic uncertainties a uniform probability distribution (Carlsson, Steinkellner, Gavel and Ölvander, 2013). Both these alternatives are described in the two following subsections.

Generating surrogate models of the different sub-models seen in Figure 1 was initially considered. Both grey-box and neural-network approaches were considered. However, a surrogate model with the possibility to describe the liquid temperature over time in the total simulation model accurately, with all required inputs such as heat loads, altitude, speed, g-loads, thrust, climate (humidity and air temperature) and the model uncertainties on top of this, was considered to be too complex and time-consuming to derive. In addition, early attempts to create a surrogate model of only the LL model of the Gripen E, sufficiently representative for use in UQ, have shown that this is also a challenging task. At this stage, the research was directed towards the methods described in the following subsections.

4.1.1 Approximate characterization using uniform distributions

Approximating all uncertainties using uniform distributions is a significant simplification that can be pragmatically justified with the argument that a uniform PDF, with respect to the width of the resulting output distribution, is the most conservative distribution available. Thus, in an approximate UQ context, a uniform PDF may be seen as a compromise between a normal distribution and a fully epistemic interval.

An evaluation of this approach was made with a less computationally expensive simulation

model of the liquid loop (LL) in the Gripen Demo test aircraft. This model is similar to the LL model of Gripen E in terms of modeling technique, but far less complex. All the identified uncertainties (25) were approximately characterized with a uniform distribution and were sampled with 350 LHS through the simulation model for one flight mission. The simulation results were then plotted against measurement data with an estimated experimental uncertainty. The result can be seen in Figure 3 where the black lines represent the LHS.

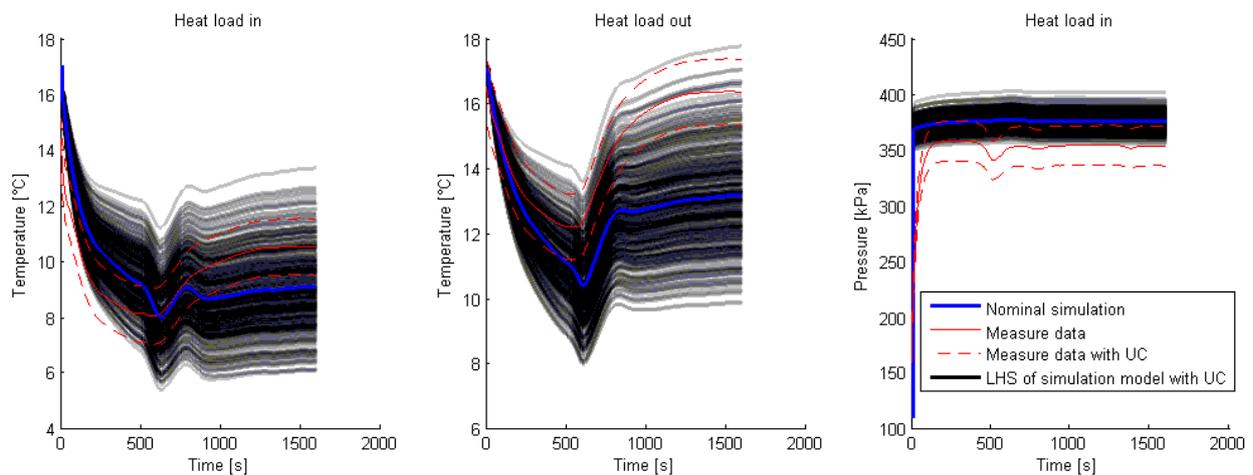


Figure 3. Propagation of the identified uncertainties (25 uncertainties) with 350 LHS through the Gripen Demo simulation model, compared to measurement data of the system.

From Figure 3 it can be seen that using the simplified characterization, the measured temperatures are captured within the range of simulation results. However, regarding the pressure level at the heat load inlet port, the identified and approximately characterized uncertainties do not produce simulation results that fully overlap the experimental data from flight testing. It should be noted that the nominal simulation results differ from the nominal measurement data both near steady-state (in the end of the simulation) and in the qualitative dynamic behaviour. This may be due to both simulation and experimental uncertainty, which both may be larger than estimated.

The method of approximating the uncertainty characterization using uniform distributions may be a suitable alternative for some applications. However, for the total S-ECS simulation model shown in Figure 1 this simplification is still not sufficient in terms of computational efficiency.

4.1.2 Approximate characterization using intervals

A method that can significantly reduce the number of simulations needed in order to determine the output uncertainty is to approximate all uncertainties as epistemic and apply an interval-based minimum/maximum analysis. Since epistemic uncertainties give an epistemic output, one theoretically needs only two simulations in order to find the output interval. If it is possible to find the uncertainty settings which give the output interval, one can drastically reduce the number of simulations needed of the total simulation model.

An approach to find this upper and lower output response is to use constrained optimization (Carlsson, Gavel and Ölvander, 2012c). When applying an optimization approach to computationally expensive models, it is advantageous to utilize Response Surface Methodology (RSM) in which the output response to the input uncertainties are described by a response surface (Khuri and Mukhopadhyay, 2010). An optimization approach using RSM in order to find the

uncertainty parameter settings that result in a minimum/maximum response can be complex if the uncertainties change nonlinearly and have strong interactions (Swiler, Mayes and Paez, 2009).

The total simulation model shown in Figure 1 consists of several connected sub-models, which can be run much faster when simulated as standalones. This opens up for new possibilities for interval-based UQ using RSM. In this work, a linear response surface has been used to approximate the output dependency of the input uncertainties. To clarify, the response surface is not a surrogate model of the simulation model itself, but rather a surrogate model of the relation between input and output uncertainties, i.e. used as a sensitivity measure. In fact, this can be seen as a way of performing global Sensitivity Analysis (SA) (Saltelli and Annoni, 2010). The response for a simulation can be written as

$$y = y_{nom} + \sum_{i=1}^n \Delta x_i \beta_i \quad (2)$$

where Δx_i denotes a change in input i from its nominal value and β_i denotes a constant, y represents the output affected by uncertainties and y_{nom} represents the nominal output. The constants β_i can be estimated using least square estimation and tell us how the uncertainties affect the response and can therefore be used to find the uncertainty settings that give the minimum and maximum response. In this paper, LHS is used to provide the input data to the least square estimation of β_i .

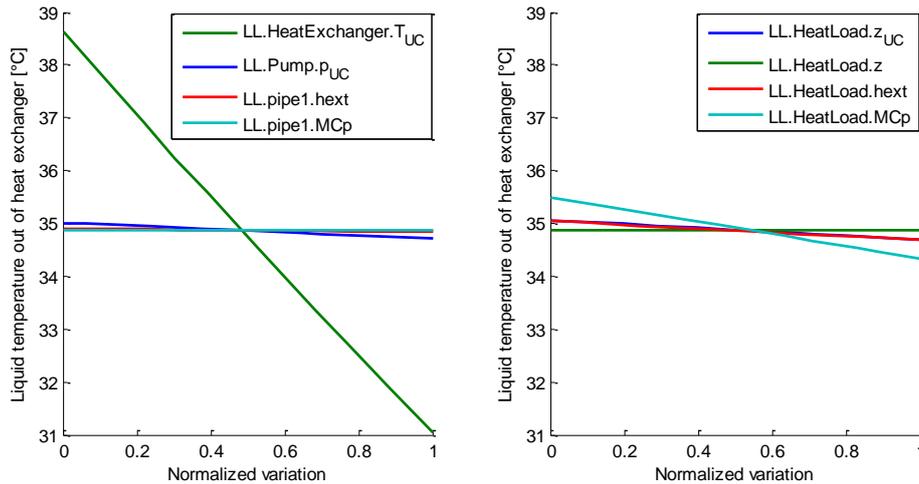


Figure 4. Graphs showing how the uncertainties affect the liquid temperature out of the heat exchanger when they are varied one-at-a-time 10 times over their intervals with LHS.

Such a study does not take interaction effects between uncertainties into consideration. It does, however, indicate that a linear approximation of how the uncertainties affect the response can be

This is where the structure of the simulation model is utilized. The sub-models can be simulated as standalones in order to derive a response surface. The estimated constants β_i in the response surface can be used to find the settings of the uncertainties that give the minimum and maximum cooling performance with respect to each specific sub-model. These uncertainty parameter settings are from here on referred to as *minimum/maximum output parameter setting*.

We have so far assumed a linear dependency between the input and output uncertainties (in this case, the output is the uncertainty of the cooling liquid temperature at the heat exchanger outlet). But can this assumption be justified? A small change in an uncertainty can be assumed to have a small impact on the temperature, and this change can therefore be considered as linear. This assumption can be examined by varying the uncertainties one-at-a-time (OAT) while studying the temperature response; a simple form of local SA as described by Saltelli et al. (2008). This was done for the Gripen Demo LL model which was simulated as a standalone. The model was simulated for 440 seconds, during which time the uncertainties were varied OAT 10 times over their intervals with LHS while the other uncertainties were fixed at their nominal values during the simulation. It can be seen from Figure 4 that the uncertainties seem to affect the liquid temperature out of the heat exchanger linearly.

justified. In order to examine whether interactions between the uncertainties have to be considered, one could compare the result when propagating uncertainties through the response

surface with an uncertainty propagation of the original model. This was done for the Gripen Demo LL model. A response surface according to equation (2) was derived with 26 simulations ($n = m + 1$). The uncertainties were then extensively propagated with MCS through the response surface and the resulting distribution of the heat exchanger outlet temperature was compared to the results of a MCS of the original model, see Figure 5. In this comparison, the uncertainties were characterized as uniform PDFs and propagated with 12000 MCS. The estimated β_i from the response surface were used to determine the min/max parameter output settings by examining the sign of β_i . These settings were then used to simulate the original model twice (with the minimum and maximum output parameter settings respectively), and the results can be seen as the green lines in Figure 5.

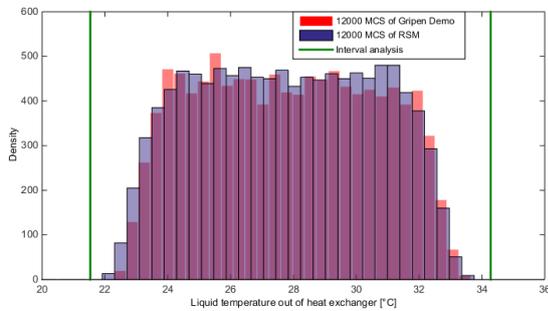


Figure 5. A comparison of the heat exchanger outlet temperature after 2000 seconds' simulation of the original Gripen Demo LL model (transparent red) and a linear response surface (transparent blue). The green lines are the results of the interval-based method. The purple area comes from the overlap of the blue and the red distribution.

Qualitatively, the linear response surface describes the output from the original Gripen Demo LL model relatively well. Figure 6 shows the time-series data obtained from two interval-based simulations and an extensive MCS respectively. Both Figure 5 and Figure 6 show that the results from the two interval-based simulations enclose the results from 12000 MCS, hence indicating that we have found the minimum/maximum output due to the identified uncertainties.

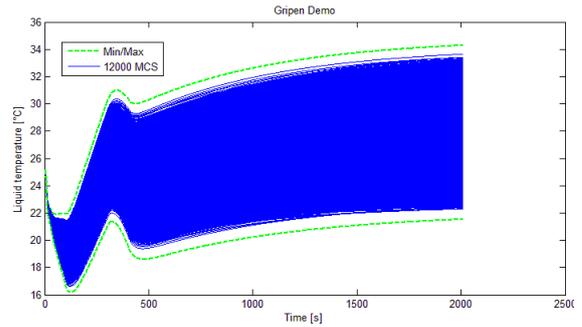


Figure 6. A comparison between 12000 MCS of the Gripen Demo LL model (blue) and two simulations of the same model using the minimum/maximum output parameter settings (red).

Another approach to find the min/max parameter output setting could be a local SA. However, by using the response surface one could use statistical methods as the coefficient of determination to examine the fit of the linear approximation and increase the credibility of the estimated min/max output parameter settings as described by Saltelli and Annoni (2010).

It is important to be aware that an epistemic approximation of all the uncertainties makes the UQ less informative since the information about aleatory uncertainties is lost. This might not be a problem if the uncertainties are epistemic or if the most significant uncertainties are epistemic. More about how to utilize SA in a UQ context is given in the following section.

4.2 Dimensionality Reduction of the UQ Problem

The credibility of the UQ relies on a good characterization of the uncertainties. According to Carlsson et al. (2013), the characterization workload for this type of simulation model can be extremely time-consuming. That characterization of uncertainties in general is time-consuming and expensive is also recognized by Saltelli et al. (2008) and Helton, Johnson, Sallaberry and Storlie (2006).

Treating all uncertainties as epistemic reduces the characterization workload, but the uncertainties still have to be assigned with well-founded intervals. The proposed characterization approach is to find the most important uncertainties, which affect the results the most, and devote the most resources to assigning these uncertainties with the most accurate intervals possible. The less important uncertainties with respect to the output of interest can be roughly assigned with "big enough intervals" or where justified ignored.

4.2.1 Reduction of uncertainties using SA

In combination with an initial rough uncertainty characterization, the importance of the uncertainties can be estimated using SA. Local SA using an OAT approach is easy to implement and has been proven to be suitable for complex 1-D simulation models of dynamical systems (Steinkellner, 2011). By comparing the sensitivity of a selected output with respect to each input uncertainty to the sum of the total sensitivity, a simple sensitivity metric can be calculated as

$$S_i = \frac{|y_{nom} - y(\theta_i)|}{\sum_{i=1}^n |y_{nom} - y(\theta_i)|} \quad (3)$$

where y_{nom} denotes is the nominal output, S_i is the sensitivity of the selected output to perturbations in the uncertain parameter θ_i , and n is the total number of uncertain parameters. The goal of this SA is to find the most influential uncertainties, which should be characterized more carefully, and to find insignificant uncertainties that can be safely ignored. This reduces the number of uncertainties, which is positive in terms of both computational cost and work effort related to detailed uncertainty characterization.

It should be noted that ignoring apparently insignificant uncertainties will make the uncertainty analysis less precise. Therefore, care has to be taken when setting a threshold for which uncertainties to ignore. A suitable threshold for this type of application could be to ignore the uncertainties that together account for less than 1% of the total sensitivity. This can be written as $\sum S_i < 1\%$ where the least significant uncertainties are added together. Such a threshold reduces the number of uncertainties for consideration from 259 to 39 for the Gripen E LL model, implying a significant reduction of the characterization workload.

As an alternative to the local SA OAT approach, a simple form of global SA may be carried out by using regression techniques, as described in section 4.1.2, equation (2).

4.2.2 Reduction of uncertainties using aggregation

In industrial applications, a major challenge when characterizing input uncertainties is lack of information. Obtaining credible information for uncertainty characterization may require extensive physical testing or development of alternative models of higher fidelity. Subject matter experts (SMEs) can be used, but in practice persons with relevant knowledge of a specific detail of an application may be very hard to find.

A method shown to be useful in the context of component-based 1-D system simulation models is to aggregate a number of input uncertainties into a smaller number of output uncertainties (Carlsson, Gavel and Ölvander, 2012b, Carlsson et al., 2012c). To clarify using an example of a pipe component, in which the differential pressure is calculated, the uncertain parameters (e.g. diameter, surface roughness, and pressure loss coefficients) are kept at their nominal values, and a new parameter is introduced to describe the total uncertainty of the differential pressure over the pipe. In this way, the number of uncertainties to be characterized and propagated is reduced. A nice feature of this method is that it can be used to describe not only input uncertainty but also model form uncertainty at component level.

5 Proposed Framework for Early and Approximate Uncertainty Quantification

Based on the comprehensive framework proposed by Roy and Oberkampf (2011), in combination with a number of methodological simplifications described in section 4, a method is proposed for early and approximate UQ applicable for large 1-D dynamic simulation models. The method consists of a number of steps as illustrated in Figure 7, and is briefly described in the subsections below. A prerequisite for carrying out a study according to the proposed workflow is that a main output of interest is defined and that a qualitative analysis is carried out of how each sub-model affects this output.

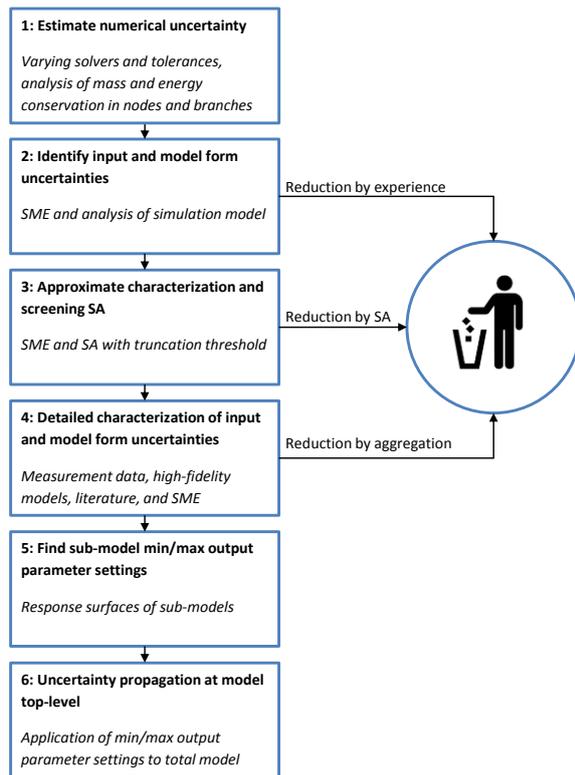


Figure 7. Proposed workflow of an approximate UQ for use in early system development phases, “SME” is an abbreviation of Subject Matter Expert. The number of uncertainties for consideration is reduced in three steps.

5.1 Estimate Numerical Uncertainty

To avoid unwanted effects of numerical uncertainty in later stages of the workflow, the numerical uncertainty is here estimated in the first step. In contrast to models of relatively high physical fidelity as often used in computational fluid dynamics (CFD) or computational solid mechanics (CSM), for 1-D simulation models of dynamical systems, the numerical uncertainty is typically insignificant. However, this has to be justified, which is the aim of the first step. This may be done by repeated simulations with varying solvers and solver tolerances, and by checking consistency of mass and energy conservation in nodes and branches of the model.

5.2 Identify Input and Model Form Uncertainties

In the second step, all uncertainty sources in addition to numerical uncertainty are identified. Regarding input uncertainty (uncertainties in for example input signals, parameters and boundary conditions), this is in accordance with the framework proposed by Roy and Oberkampf (2011). However, as the proposed approximate framework includes an estimation of model form

uncertainties by analysis at component or sub-model level rather than using measurement data from extensive repeated physical testing at system-level, there is a significant difference between the two frameworks. Estimation by analysis is less stringent but is the only applicable alternative found when extensive measurement data on system-level are not available – which is typically the case in early system development phases. For representation of model form uncertainty, the aggregation method described in section 4.2.2 may be used. Already at the second step, the number of uncertainties for later consideration can be reduced by experience.

5.3 Approximate Characterization and Screening SA

To further reduce the number of uncertainties, the third step includes an approximate experience-based characterization of the identified uncertainties in combination with a screening SA. Section 4.2.1 provides a simple example of how this can be done by the use of local SA. However, if global SA is applied and implemented with RSM, this step may be merged with step five below.

5.4 Detailed Characterization of Input and Model Form Uncertainties

The remaining uncertainties, and especially the ones found to be the most significant, are carefully characterized in step 4. To reduce the computational cost, all uncertainties are characterized using intervals. As shown in Figure 7, also this step may involve reduction of the number of uncertainties for consideration. A method that can be applied at this step is uncertainty aggregation as briefly described in section 4.2.2.

5.5 Find Sub-Model Min/Max Output Parameter Settings

With the characterization done, the minimum/maximum output parameter settings are found in step 5 using a linear response surface for each sub-model, as described in section 4.1.2. It should be noted that for some cases the minimum/maximum output parameter settings can be derived without applying RSM, e.g. for simpler sub-models or when there is significant experience of the specific sub-model.

5.6 Uncertainty Propagation at Model Top-Level

Finally, in step 6 the different minimum/maximum output parameter settings of the sub-models are

combined in order to obtain the resulting uncertainty of the total simulation model. For the industrial application example used in this paper, the minimum/maximum output parameter settings found are combined and two simulations for each operating point are carried out to obtain estimations of the best and worst cooling performance, respectively.

6 Results

To evaluate the industrial applicability of the proposed workflow for approximate UQ, it has been applied on the total S-ECS simulation model. The simulation models as well as the inputs used in the work described by this paper have been adjusted so as not to reveal confidential information about the systems and their performance. The outcome of interest from the simulations performed in this work is thus not the simulation results in themselves, but rather how the proposed UQ method performs on an industrial grade simulation model of a complex system.

The minimum/maximum output parameter settings for the Gripen E LL were found with a linear response surface by simulating the sub-model as a standalone. The inputs to the LL were the cooling air mass flow and temperature. The LL output of interest, affecting the cooling performance on a system level, was the liquid temperature out of the heat exchanger. An increase in the heat exchanger outlet temperature relative to the reference cooling temperature indicates deficient cooling performance. The inputs were the same throughout the different simulations when deriving the response surface. Only the uncertainties were changed between the different simulations using LHS and $1 + n$

samples, where $n = 259$ uncertainties (including all identified uncertainties). These uncertainty settings were then used when simulating the total simulation model.

Regarding identification and characterization as well as response surface generation, it should be noted that this study is limited to the uncertainties in the LL sub-model. Adding uncertainties from other sub-models will not increase the number of simulations needed of the total simulation model. The only increase in computational cost is related to finding the minimum/maximum output parameter settings using RSM also for the other sub-models (P-ECS, S-ECS, and Engine). The computational cost of adding uncertainties from other sub-models is therefore considered relatively low.

In the industrial application studied, a complicating factor is that the heat exchange with ambient air, and thereby the LL cooling performance, is affected by the actual flight case. Flying at low altitude in combination with high speed implies high skin temperature, which in turn heats up the components and the liquid in the LL. Similarly, the combination of high altitude and low speed cools the LL. These characteristics are visualized in Figure 8, which also shows the 25°C curve. It should be noted that the equipment bay temperatures, which directly affect the LL equipment, differ from the skin temperature, both in steady-state and transient conditions; however, this is accounted for in the LL model. To summarize the problem; when finding minimum/maximum output parameter settings for uncertain parameters related to the LL equipment heat exchange with ambient, the operational point in the flight envelope has to be considered.

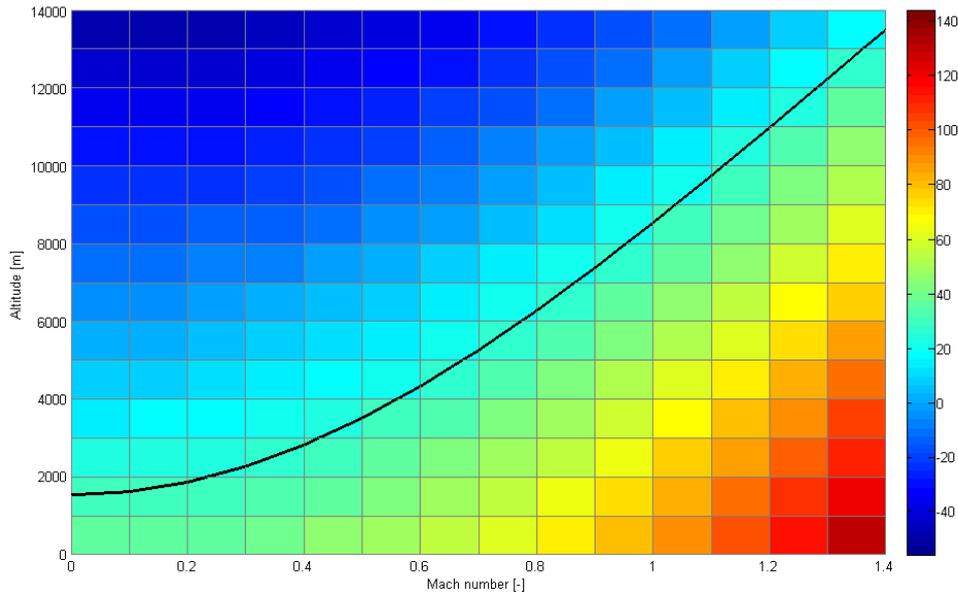


Figure 8. Flight envelope with colours representing skin temperature (recovery temperature) in [°C] for International Standard Atmosphere (ISA) + 20°C. The black curve indicates a temperature of 25°C, which is the assumed control reference temperature out of the LL heat exchanger.

6.1 Stationary cooling performance

The number of evaluated operating points in the flight envelope, where the aircraft altitude is plotted against the speed, was set to 59. Three simulations (minimum, nominal and maximum stationary cooling performance) at each evaluated operating point give a total of 177 simulations. The nominal simulations were utilized to attain the ambient air temperature. This temperature determines which combination of the minimum/maximum output parameters settings to use. In this study the stationary cooling performance is measured in terms of endurance time, i.e. the time in minutes that the system can keep the cooling liquid temperature out of the LL heat exchanger below a predefined temperature level in different points of the flight envelope. Simulating one stationary operating point with the total S-ECS simulation model takes approximately 20 minutes, giving a total simulation time of 59 hours. A typical performance plot for the nominal endurance time, without any uncertainty information, is shown in Figure 9.

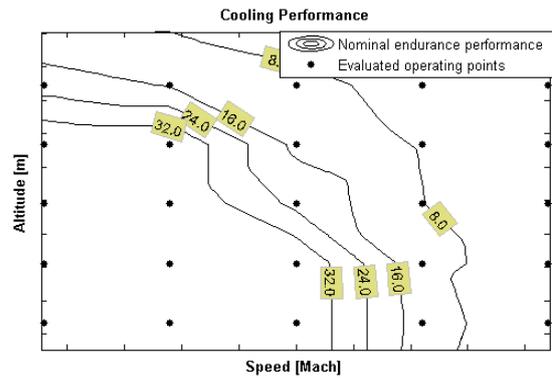


Figure 9. Nominal stationary cooling performance in terms of endurance time in minutes shown in square datatips. The dots are the evaluated operating points and the curves represent interpolated endurance time limits.

In the lower left corner of Figure 9, the nominal endurance time t_{nom} is unlimited with respect to cooling performance (in practical situations there are other limitations, such as available fuel). In the upper right corner of Figure 9, t_{nom} is zero. Figure 10 shows the same nominal data extended with the uncertainties resulting from the approximate UQ.

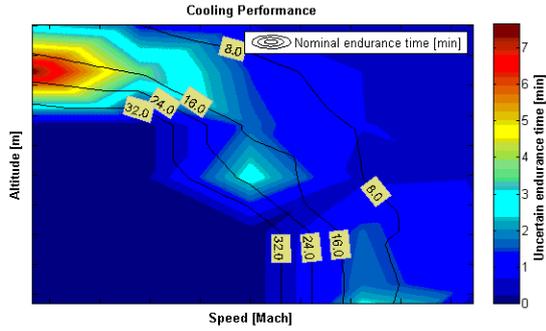


Figure 10. Nominal endurance time in minutes shown in square datatips. The estimated maximum uncertainty of the endurance time is represented by colours; dark blue for “no” uncertainty and dark red for an uncertainty of approximately 8 minutes.

As can be seen from Figure 10, the uncertainty of the endurance time is greatest in the limited region, i.e. where the nominal endurance time is $0 < t_{nom} < \infty$. The reason for this is that in the other parts of the flight envelope shown in Figure 10, the cooling performance of the system is either i) so high that

the S-ECS control system manages to keep the temperature at the set point value regardless of the uncertainties, or ii) so low that the endurance time is zero minutes regardless of the uncertainties.

6.2 Dynamic cooling performance

The nominal dynamic cooling performance for a complete flight mission of 50 minutes was simulated and the proposed approximate UQ framework was applied. Due to the fact that the cooling performance and its uncertainty is dependent on the actual flight case as described in Figure 8 (ambient temperatures heating or cooling the LL), four simulations in addition to the nominal simulation are needed for the approximate UQ.

To speed up the simulations to enable a comparison of the approximate UQ with an extensive LHS, the system was simulated with the less complex Gripen Demo LL model. The result can be seen in Figure 11.

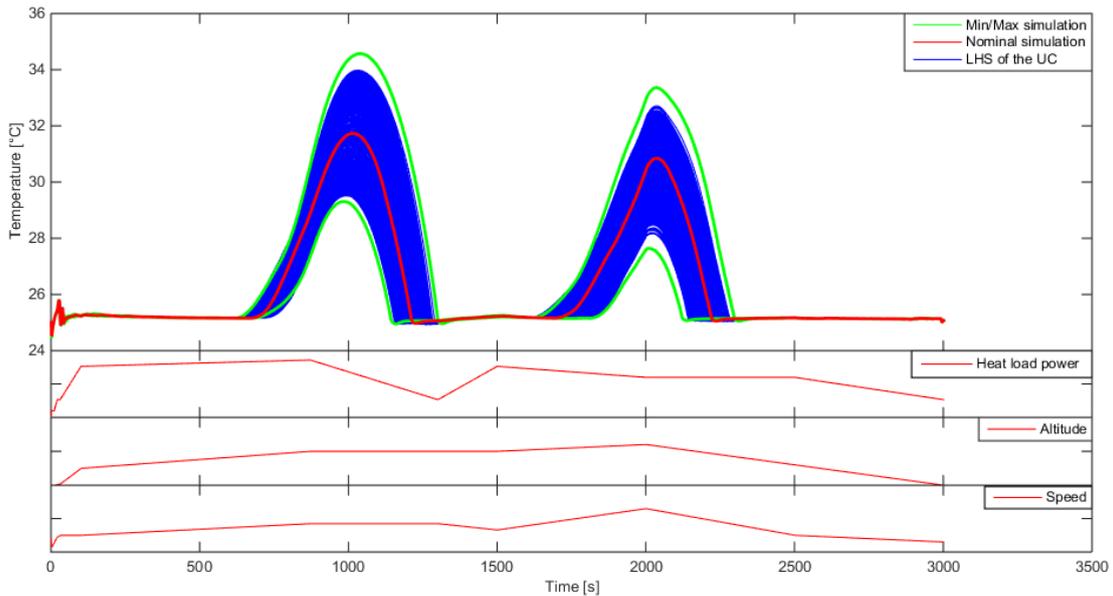


Figure 11. The liquid temperature out of the LL heat exchanger during a flight mission, where the heat loads and the aircraft altitude and speed change over time. The figure includes a comparison of the results from an extensive LHS (25 uncertainties sampled with 800 LH intervals) with the min/max results from the proposed approximate UQ framework.

As shown in Figure 11, the cooling performance is limited in two segments of the flight mission, with an increase in liquid temperature as a consequence. As expected, the min/max uncertainty resulting from the proposed approximated UQ framework encloses the results based on LHS.

7 Discussion and Conclusions

The proposed framework for approximate UQ significantly reduces both computational cost and engineering workload. In comparison with the framework proposed by Roy and Oberkampf (2011), this is at the expense of mathematical stringency as well as UQ accuracy. However, the proposed framework is intended for large system simulation models implying high computational

cost in combination with many uncertainties. In addition, users are often interested in performance predictions in several operational points, both in stationary and transient conditions. This makes computationally expensive UQ methods unfeasible.

Even though the challenge of computational cost is great, in an industrial context where information and resources are limited, the challenge of uncertainty characterization is even greater. This is especially true in early development phases, which is the focus of the proposed framework. Nevertheless, an approximate UQ is more worthwhile than no UQ.

To decrease the computational cost as well as the engineering workload required for uncertainty characterization, reducing the number of uncertainties for consideration is essential. In the proposed framework, the number of uncertainties for consideration is reduced by three means: i) experience, ii) aggregation, and iii) approximate characterization combined with sensitivity analysis.

The interval-based epistemic UQ approach was found to be suitable when analysing closed-loop systems, such as simulation models with a control system. For the industrial application studied it was seen that in many operational points, the uncertainties did not affect the ability of the S-ECS to control liquid temperature to the reference value since the available cooling power was more than enough in these operational points. In a UQ context, these simulations are unnecessary to evaluate with a more comprehensive propagation method, since the liquid temperature will always be close to the reference value. This implies that the most interesting simulations are those where the system starts to have deficient cooling performance.

If a more accurate but still computationally relatively cheap UQ is required to explicitly handle significant uncertainties characterized in detail as aleatory, the proposed framework can be extended; The proposed UQ framework can be used as described to find minimum/maximum output parameter settings for the epistemic uncertainties, and combined with sampling of the aleatory uncertainties. A probability box is thus derived without nested sampling. Instead of performing an extensive sampling of the aleatory uncertainties for each combination of epistemic uncertainties, the aleatory uncertainties are only sampled for the minimum and maximum output parameter settings, respectively.

Finally, it should be noted that if strong nonlinearities exist, the combination of the

minimum and maximum output parameter settings for the individual sub-models does not necessarily result in minimum and maximum response in the total simulation model. Such characteristics could possibly be detected by means of additional random samples at the model top level. Nevertheless, the execution of any UQ effort will benefit from extensive domain experience and system understanding.

Acknowledgements

The research leading to these results has received funding from Saab Aeronautics and the National Aviation Engineering Research Programme (NFFP) jointly driven by the Swedish Armed Forces, the Swedish Defence Materiel Administration (FMV), and the Swedish Governmental Agency for Innovation Systems (VINNOVA), (NFFP6 2013-01211).

References

- Carlsson, M., Andersson, H., Gavel, H., Ölvander, J., Methodology for Development and Validation of Multipurpose Simulation Models, *Proceedings of the 50th AIAA Aerospace Sciences Meeting*, Nashville, TN, USA, 2012a.
- Carlsson, M., Gavel, H., Ölvander, J., Evaluating Model Uncertainty Based on Probabilistic Analysis and Component Output Uncertainty Descriptions, *Proceedings of the ASME 2012 International Mechanical Engineering Congress & Exposition*, Houston, TX, USA, 2012b.
- Carlsson, M., Gavel, H., Ölvander, J., Utilizing Uncertainty Information in Early Model Validation, *Proceedings of the AIAA 2012 Modeling and Simulation Technologies Conference*, Minneapolis, MN, USA, 2012c.
- Carlsson, M., Steinkellner, S., Gavel, H., Ölvander, J., Enabling Uncertainty Quantification of Large Aircraft System Simulation Models, *Proceedings of the Council of European Aerospace Societies (CEAS) 2013 Conference*, Linköping, Sweden, 2013.
- de Rocquigny, E., Devictor, N., Tarantola, S., *Uncertainty in Industrial Practice: A Guide to Quantitative Uncertainty Management*, John Wiley and Sons, 2008.
- Eek, M., Kharrazi, S., Gavel, H., Ölvander, J., Study of Industrially Applied Methods for Verification, Validation & Uncertainty Quantification of Simulator Models, *International Journal of Modeling, Simulation, and Scientific Computing*, 2015.
- Helton, J.C., Probability, conditional probability and complementary cumulative distribution functions in performance assessment for radioactive waste

disposal, *Reliability Engineering & System Safety*, **54**(2–3): p. 145-163, 1996.

Helton, J.C., Davis, F.J., Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems, *Reliability Engineering & System Safety*, **81**(1): p. 23-69, 2003.

Helton, J.C., Johnson, J.D., Sallaberry, C.J., Storlie, C.B., Survey of sampling-based methods for uncertainty and sensitivity analysis, *Reliability Engineering & System Safety*, **91**(10–11): p. 1175-1209, 2006.

Khuri, A.I., Mukhopadhyay, S., Response surface methodology, *Wiley Interdisciplinary Reviews: Computational Statistics*, **2**(2): p. 128-149, 2010.

Oberkampf, W.L., Roy, C.J., *Verification and Validation in Scientific Computing*, Cambridge University Press, Cambridge, UK, 2012.

Roy, C.J., Oberkampf, W.L., A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing, *Computer Methods in Applied Mechanics and Engineering*, **200**(25–28): p. 2131-2144, 2011.

Saltelli, A., Annoni, P., How to avoid a perfunctory sensitivity analysis, *Environmental Modelling & Software*, **25**(12): p. 1508-1517, 2010.

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., et al., *Global Sensitivity Analysis: The Primer*, John Wiley and Sons, 2008.

Steinkellner, S., *Aircraft Vehicle Systems Modeling and Simulation under Uncertainty*, Tekn. Lic. no 1497, Linköping University, Linköping, Sweden, 2011.

Swiler, L.P., Mayes, R.L., Paez, T.L., Epistemic uncertainty quantification tutorial, *Proceedings of the 27th Conference and Exposition on Structural Dynamics (IMAC XXVII)*, Orlando, FL, USA, 2009.