Considering value of information when using CFD in design

by

John Satprim Misra

A thesis submitted to the graduate faculty

in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Major: Mechanical Engineering

Program of Study Committee:
Kenneth Mark Bryden, Major Professor
Arne Hallam
Xinwei Wang

Iowa State University

Ames, Iowa

2009

TABLE OF CONTENTS

ABSTRACT

This thesis presents an approach to find lower resolution CFD models that can accurately lead a designer to a correct decision at a lower computational cost. High-fidelity CFD models often contain too much information and come at a higher computational cost, limiting the designs a designer can test and how much optimization can be performed on the design. Lower model resolution is commonly used to reduce computational time. However there are no clear guidelines on how much model accuracy is required. Instead experience and intuition are used to select an appropriate lower resolution model. This thesis presents an alternative to this ad hoc method by considering the added value of the addition information provided by increasing accurate and more computationally expensive models. In the algorithm presented here for selecting the correct model resolution, the designer, who should be most familiar with the model, creates some quantifiable metrics for model components that he/she identifies as key characteristics. The designer uses the individual component metrics to create a selection utility function. The selection utility function is used to validate the accuracy of low-resolution model by comparing the magnitude of utility that leads to the correct decision. The low-resolution model can then be used to test multiple designs and/or for optimization at substantially lower computational time, giving the designer more flexibility in the design process for the model and other similar models.

CHAPTER 1 Introduction

In the engineering design landscape, the use of computational simulation in the design selection process is becoming an increasingly valuable tool. These computational simulations are numerical representations of a particular aspect (or group of aspects) of a physical product. For example, consider the design of a fan. A finite element model may be used to examine the stress in the fan blade; a computational fluid dynamics model may be used to examine the blade pitch, airflow, and power needed; and an integral model may be used to examine conservation of mass and energy. In each case, instead of building a physical prototype, the engineer can create multiple computational models to guide the design process. These computational models can often provide more detailed answers in less time and at less cost than a physical model.

However, in many of these cases, high-fidelity models can take days or weeks to return a single result on a proposed design, limiting the engineer's ability to explore the design landscape. In these cases, the use of computational modeling is very limited and generally the design falls back on the practice of using physical models or thumb rules, using computational modeling only to confirm the proposed design. The resulting exploration of the design space is limited. Computational models are needed that are robust enough to permit the timely exploration of multiple creative designs and accurate enough to guide the design process.

Based on this, we see that the challenge of using computational models is balancing the information available from the model with the time and cost needed to develop and use the model. Detailed, high-fidelity models may provide a very accurate answer about a

particular detail, but at a high cost in terms of resources, computational time, and time to create and validate the model. Conversely, a simple model can be run interactively on a desktop or laptop and can be created quickly, but may not yield sufficient information to support a design decision. A set of models is needed that can provide sufficient information without excessive cost. One significant challenge in the development of these sets of models is that there are few rule sets, thumb rules, or guidelines available to guide the development of the model sets needed in engineering design.

This thesis explores using the economic principle of value of information as a means to achieve the needed balance in the development of the simulations and models used in engineering design. Recently, the use of value-of-information principles in design has been examined by a several groups of researchers [12, 13]. Value of information is an economic tool that evaluates the benefit of collecting additional information in a decision-making context.

One of the challenges of using value of information with CFD models in design is the development of a set of simpler models. Several approaches are available including orthogonal decomposition, models with lower-resolution grids, accelerated convergence using neural networks, and neural networks. In many of these cases, the focus is on the development of metamodels in place of a high-order or more complex model. Barton [1] has defined a metamodel as a lower-order model that represents (or maps) a more complete model. This definition can be expanded to say that metamodels are a set of models that represent a more complex set of models. In metamodeling, a simplified model of a more detailed computer simulation is built and a number of computational experiments are completed to create an approximation that can be used for untried inputs. Finding the optimal

resolution for a simplified model, or metamodel, can lead an engineer to optimal design selection in not only that problem, but also in other similar design problems.

This thesis examines the use of CFD models using lower-resolution grids as metamodels for a validated and verified CFD model with a detailed grid. In this schema, a standard process is used to create a validated and verified three-dimensional CFD model with a detailed grid. Following this, lower-resolution CFD models are used to model the behaviors and results of the high-resolution model. The choice of lower-resolution model is then based on the information needed for the design process and the time and resources needed for the model.

This thesis is organized as follows. In Chapter 2, the background of value of information and its use as a way to approach engineering design are introduced. The model description for the experiment of a third-world stove used as the example in this work is presented in Chapter 3. The setup for the experiment is provided in Chapter 4. Chapter 5 discusses the results of the experiment and validates the selection of a more robust low-fidelity model. Chapter 6 summarizes the results, discusses how this approach could be approach to other design processes, and suggests avenues for further work.

CHAPTER 2 Background

Detailed engineering models (e.g., computational fluid dynamics and finite element models) are often developed to support the design of energy systems. The process of model development, verification, and validation is a time-consuming but essential process. There are many proposed ways to validate a model, including animation (graphical representation), comparisons to other models, extreme condition tests, historical data validation, facial validity (a person knowledgeable about the real system believes the model and its behavior to be reasonable) [2], white-box validation (micro-scale/component accuracy), and black-box validation (macro-scale/overall accuracy) [3]. In the minimum validation and verification process, the first step is specifying the amount of accuracy required of a model's output variables of interest for its intended application [2]. Next, the assumptions and theories underlying the model should be tested. A face validity check should be done at each model iteration. Then comparisons should be made, if possible, between the model and system behavior data for at least two sets of experimental conditions.

After validation and verification, the simulation model can then be used for a variety of design tasks including optimization. However, detailed models are rarely used interactively or for optimization in the design process because of the time and resources needed to run the models and the complexity of using the models. Many models are slow and time consuming to use and provide far more information than is needed for a particular decision. Other models fail to address the question at hand. As a result, an ad hoc process of model development and extension using reduced-order models is often used.

One way to approach this problem of ad hoc model development and use is to consider the value and fidelity of the information needed as well as the cost of gaining the information. In the example examined in this thesis, an improved plancha stove design model has been validated using a computational fluid dynamics solver [4]. The details of this model and design problem are discussed in greater detail in Chapter 3. In this energy system, the stoves are biomass cookstoves used in lower-income Central American households. The primary purpose of these stoves is to increase efficiency, reduce cooking fuel costs, and reduce the health impact of household cooking. The original stove design had large temperature variations on the cooking surface that limited the stove's effectiveness for cooking and reduced its overall efficiency. Adding baffles to the flow of the flue gas under the cooking surface can improve the stove's performance, but the location and size of baffles for an effective design are unknown. After going through the time-consuming validation process using high-order models in a commercial CFD software package, Star-CD$^{\text{TM}}$, in combination with evolutionary algorithms, an ideal baffle layout can be found for a particular size, shape, and stove application. In the optimization process, 95% of the computational time was spent on the CFD solver for the higher-order models. Once the high-order model has been validated and an ideal answer is known, the value of using reduced-order models can be investigated.

2.1 - Value of information

The basic key consideration when selecting a decision in any analysis is to assess the decision's impact: how does the decision influence future decisions, and how does the decision change under different scenarios? Engineers must consider these questions when deciding how to expend resources to create and analyze information. This information may

provide value by producing a better final design. Though the benefit is uncertain until the resources are spent, the analysis of all possible sources of data is often a waste of resources [5]. Value of information operates under the growth of knowledge assumption, which states that as more information is gathered, less information needs to be acquired [6]. It is easy to show that the return does not grow monotonically with increasing amounts of information [7]. Thus, there must be some point at which the return no longer exceeds the cost of collecting more information. Value-of-information metrics measure the cost of collecting information versus the return. Another way to define value of information is as quantitative measure of the value of knowing the outcome of the information source(s) prior to making a decision. It can be quantified as the difference in value or maximum expected utilities achievable with or without knowing the information sources in a decision-making problem.

2.2 -Value of information properties

There are two types of value of information that are discussed in the literature: the value of perfect information, and the value of imperfect information, which is often referred to just as the value of information. The value of perfect information almost never exists but is easy to calculate and can be a useful tool for the designer. The value of perfect information is calculated by assuming that the key uncertainties can be resolved absolutely before making the decision. If the value of the perfect information is less than the cost to obtain the information, then the analysis can stop.

The probability associated with incomplete data in value of information is linked to operational reliability of data acquisition. The reliability is often judged by the designer. The greater the perceived reliability of the information, the faster the value of information approaches zero [8]. The value-of-information approach emphasizes the reliability of

information sources in a consistent sensitivity analysis to see how value of information changes by varying input parameters including reliability. However, it must be noted that aggregated probabilities of incomplete and imperfect data have the potential to erode all value in data acquisition [9], whereby compounding the data also compounds the uncertainties and could create a scenario that the compounded uncertainties mitigate the reliability of the compounded data in the decision-making process.

In summary, the value of information decreases as more information is gathered, as the uncertainty in information increases, as the cost of acquiring more information increases, as the reliability of previous information acquired increases, and as its impact on a decision's decreases [10]. The value of information increases with the increasing probability of making a wrong decision in absence of new information and with the increasing cost of making a wrong decision absent of new information [11].

2.3 - Applications of value of information

Value of information has been applied and combined with many different analysis methods in different fields. In the financial trading sector, value of information has been used to show a non-trivial, non-monotonic dependence on agents' returns on amount of information possessed [6]. In this model, a double stock market with cumulative information delivered to traders was presented. The results showed that in the long run, non-informed traders performed better than averagely informed traders. That is, the information had a negative effect on averagely informed traders. Only the most informed traders were able to obtain above-average returns, comparable to those achieved through insider trading. The results suggest an explanation for the below-average performance of actively managed funds. The possible cause is that managers are neither perfectly informed (insiders) or completely

uninformed. The completely uninformed cannot be exploited by the better informed traders as their decisions are random, but traders with average forecasting ability can be exploited by the better informed traders.

In the oil-drilling industry, value of information has been used to introduce a new risk and opportunity management method and justify the continuation or discontinuation of data acquisition [9]. In this industry, value is measured monetarily. The two main value-of-information applications by the oil industry are value assurance and value creation. Value assurance can readily identify an immediate target when the value of data acquisition is known. Value creation accounts for sensitivity analyses to identify key risks and probabilities that may be managed to enhance value, which can open new opportunities and facilitate business growth as applied to portfolio management. Engineers identify the impact of the available data on the various combinations of chance scenarios and conscious options. They then compare the optimal set of options to see if further data acquisition is necessary. In a common example, engineers analyze the expected benefit of a series of actions (to drill or not to drill) and outcomes (dry well or very productive well) based on a set of assumptions gathered from preliminary seismic data. They then compare the benefit of these actions to the expected benefit of doing a full high-cost appraisal of the possible new well to decide whether to do the appraisal and whether to drill at the location in question.

In supply-chain management, value-of-information analysis has been used to generate revenue by reducing information-gathering costs and redirecting inventory flow [7]. In this case, a framework is introduced to determine the impact of marginal information on the cost performance of a supplier. The bullwhip effect is a common problem in supply chain management. Suppliers forecast demand for their product and the size of the inventory buffer

needed to deal with variations of product demand. The farther upstream a supplier is, the larger the variations they see and the larger the needed buffer. The larger buffer is the larger the carrying cost. The value of information has been found to increase with the level of marginal information (i.e., a base scenario where the supplier's only information is the history of orders it has received compared to a scenario where the order history, the customer's replenishment policy – what level of remaining stock dictates re-ordering– and the average lead time to fulfill the replenishment order are available). With the customer's additional information, the supplier can reduce error in their forecast and can reduce the buffer size and the associated carrying costs.

In the most similar engineering design application, value of information has been used for simple modeling problems and for refining models [12]. As a starting point, it is theorized that it is valuable to explicitly represent imprecision in available characterization of uncertainties by using imprecise improbabilities, or a bounded interval, of some attribute of the model. Uncertainty is the gap between certainty and the designer's present state. Reducible uncertainty represents the gap between the present state of information and the state of precise information (the actual state). Irreducible uncertainty arises from random processes and accounts for the gap between certainty and the state of precise information. If the error bounds of the information are available, the designer can measure the improved confidence in decisions as he increases the model's refinement.

Using a simple pressure vessel, this application of value of information emphasizes that model refinement should be based both on the accuracy of the model and on the decision's outcome. In this scenario, a range for the strength of material used to manufacture a pressure vessel is predicted, and with additional information, the range can be reduced. An

improvement potential is introduced as the value-of-information metric. A function utility is introduced as the outcome associated with both ends of the range. The improvement potential compares the upper range's maximum expected utility of the entire design space to the current lower-bound range utility. When the improvement potential reaches zero, narrowing the range of strength information is no longer needed.

Further emphasis is put on accuracy and on the decision's outcome in a shock simulation model, where reducing cell size is addressed in addition to choosing the size of a statistical volume element that represents a small section of material through which a shock is propagated. Again a range is given for the upper bound and lower bound of cell length and the size of the statistical volume element. Combinations of three values selected between each range create nine different options. A utility is again introduced and an improvement potential compares the maximum utility of the upper bound to a specific lower bound utility. It is found that though a designer's goal should be to have an improvement potential of zero, he/she must keep in mind that refining a simulation model improves accuracy but increases costs. A perfect model may not be needed to make a good decision. The designer's objective in this approach is to determine the level of accuracy (or refinement) needed to make an appropriate decision.

This work can be extended to using CFD models in design. As discussed in Chapter 1, CFD models are complex and time consuming both to build and then to operate. And yet, they can provide a wealth of detailed information that may or may not be relevant to the design process and decisions that need to be made. The idea that a designer can find an optimal model refinement for the use of CFD in an engineering problem is more fully addressed in Chapter 3.

2.4 - Creation of value-of-information utility measurement

One of the key aspects of using value of information in building an information space is developing measurement of the utility or metric of the design or decision. Today's CFD modeling community's research environment places great importance on evaluation involving quantifiable metrics that can be assessed and judged with clarity and accuracy. But in engineering design, it is impractical to fully characterize a model due to constraints such as time and computational ability. In the design phase, only partial characterization of the design problem is available and much of this understanding is of a subjective nature rather than a quantitative nature [12]. Thus, the creation of the value-of-information utility falls on the shoulders of the engineering designer [13]. The engineering designer is most familiar with the design question at hand. And although the final decision is unknown, she/he is most able to create a metric that identifies the important components that are to be used as gauges for the engineering design.

For the financial trading sector case, the value-of-information metric used by the financial analyst had components that included the information levels of future dividends and the expected return on the dividends paid to the investors after each period. For the oil-well drilling case, the value-of-information metric used by the geological engineers had components that included reliability of information, probability of the correct decision, and monetary return on investment. For the supply-chain management case, the value-of-information metric used by the supplier reflected the reliability of the information resources and their cost reduction implications. In the simulation design case, the value-of-information metric used by the engineering designer had a cost formula portraying the bill of materials

cost to manufacture the pressure vessel, the reliability of the simulation model, and the bounds of the outcome of the engineering designer's decision.

When considering the use of CFD modeling in design, areas that need to be considered include grid or mesh resolution, the properties of state defining the fluid, and similar to the pressure vessel and shock problem, the validity or reliability of the simulation model, the simulation wall convergence time, and the outcome of the engineering designer's decision.

CHAPTER 3 Description of the CFD Model and the Design Question

CFD models have the capacity to provide a wide range of detailed design information that may or may not be related to the design process and the decisions that need to be made. However, as discussed earlier, it is time consuming to build, validate, and verify a CFD model. Once the model has been constructed, using it is also time-consuming. Part of the problems that arises is that the validation of the model and the use of the model for design and decision making is often integrated or coupled so that validated models are thought to be ready to use in the design process. But the relationship between the model outputs and the critical design questions is often not clear. The CFD model may only provide qualitative information; it may provide information that is too detailed for the stage of the design; or the accuracy of the model may not be sufficient to answer the particular design question. This thesis explores one aspect of this question – following model validation, how can lower-resolution CFD models that act as metamodels in the design process be designed based on a coarser grid? The proposed selection and development criteria for these models are based on the value of the information provide by the models within the context of the engineering design process.

The design problem used in this study is the design of an improved wood cookstove used in lower-income households in Central America. As shown in Figure 1, the improved wood cookstove being designed is a freestanding cookstove with a small combustion chamber that heats a large flat cooking surface (plancha). The cooking is used both as a grill for cooking tortillas (a flat unleavened bread) and for cooking beans and rice. Based on this, the ideal temperature distribution is a single hot spot for quickly boiling water and an even
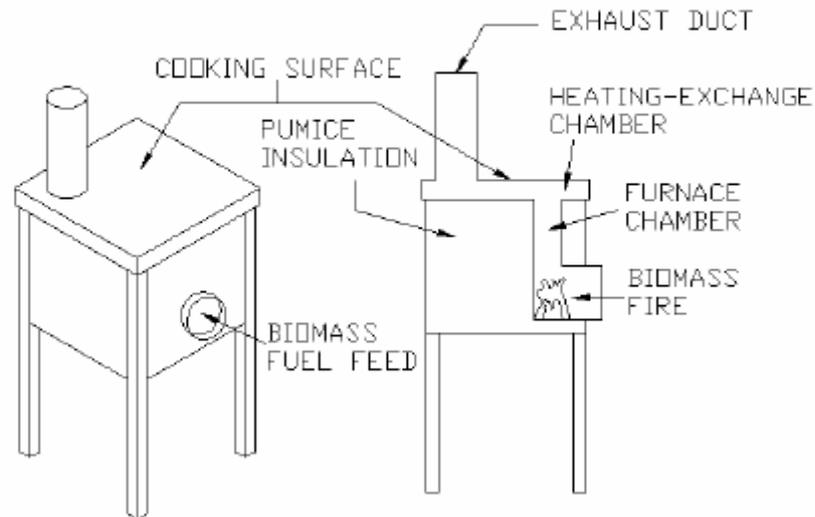
Figure 1- Improved wood cookstove

temperature over the rest of the stove's surface for cooking tortillas or simmering foods. The

primary source of information for this design is based on CFD analysis [4]. This stove was

developed to increase efficiency, reduce cooking fuel costs, and reduce the health impact of

household cooking [16-19]. The proposed stove design had a large temperature variation on

the cooking surface that limited the stove's effectiveness for cooking. Using CFD baffles

added to the flow of the flue gas under the cooking surface can enhance the stove's

performance. A CFD model was used for the design because there is no simple, low-cost

solver that could predict the surface temperature of the stove and a cut-and-try approach was

excessively expensive and time consuming. Without baffles, the cooking surface had

temperatures as cool as 150°C in some regions and as high as 600°C in others. The user sets

the overall surface temperature by varying the fuel feed rate. However, the relative spatial

temperature variation is generally unaffected by the fuel feed rate.

3.1 - Stove model

As shown in Figure 1, the cookstove consists of an elbow-shaped combustion chamber approximately 10 cm × 10 cm into which small wood sticks are fed. Hot gases leave the combustion chamber, travel under the cooking surface, and exit through a chimney. The small size of the combustion chamber relative to the size of the cooking surface is unusual and creates the large spatial variation in the temperature of the cooking surface. For this study, this stove was modeled using commercial CFD software, Star-CD^{TM}. The geometry modeled is a simple rectangular prism representing the 54 cm × 54 cm × 2.5 cm heat exchange chamber coupled to the furnace chamber and exhaust duct. Boundary conditions were determined from in-field measurements. The remaining surfaces of the model were assumed to be adiabatic to simulate the pumice insulation used in the stove construction. Experimental studies in the field with various baffle configurations found no evidence of unsteady flows. Additionally, transient flow computational studies found no oscillating or unsteady flows with various baffle configurations. Based on this, a steady-state CFD model utilizing the Reynold-averaged flow equations is used in the fitness evaluation. Combustion within the combustion chamber was not modeled. Instead, the boundary conditions at the inlet of the heat transfer chamber during typical cooking evolutions were measured. These were a velocity of 3.88 m/s and a temperature of 977 K. Density changes in the air from changing the temperatures were included. Turbulence was modeled using the K–epsilon model with an intensity of 0.1 and an entrance length of 4.8 cm. Resistance to heat transfer from the cooking surface was modeled using a heat transfer coefficient of 20 W/m^2 K, a thermal conductivity of 30 W/m K, and a surface thickness of 1.6 cm. The surface profile is

in excellent agreement with the experimental data collected in Nicaragua. A full description of the computational model and validation process is given in [20].

3.2 - Baffles

In this engineering design problem, baffles are inserted into the flow path to obtain the desired temperature distribution on the cooking surface. These baffles are perpendicular to the surface and parallel to one of the edges of the cooking surface. By changing the flow of hot gases underneath the cooking surface, the baffles change the cooking surface spatial temperature distribution.

In the physical stove, thin metal plates are welded to the underside of the cooking stove to guide the flue gases. These baffles are combined to create an effective design to maintain an even temperature cooking surface. The criteria for the baffles were that they had to be perpendicular to the stove surface, parallel to the heating wall chambers, and no design could include more than three baffles to ensure that the stoves were easy to manufacture by hand. The models were built in Star-CD^{TM} using rectilinear grids. The baffles in Star-CD^{TM} were modeled as two-dimensional adiabatic baffle cells.

In the high-fidelity model, the heat-exchange chamber is represented by a $160 \times 160 \times 30$ rectangular grid. Two entry regions were attached to the bottom and top of the grid to model the chimney and the inlet of the furnace heat transfer chamber (Figure 2). An outlet boundary condition was applied to the chimney with an assumption that all flow exits at the chimney and there were no air losses around the stove. The inlet boundary condition was applied to the entrance of the heat transfer chamber. A wall condition was applied to the top of the stove that initialized the temperature as ambient 293 K.
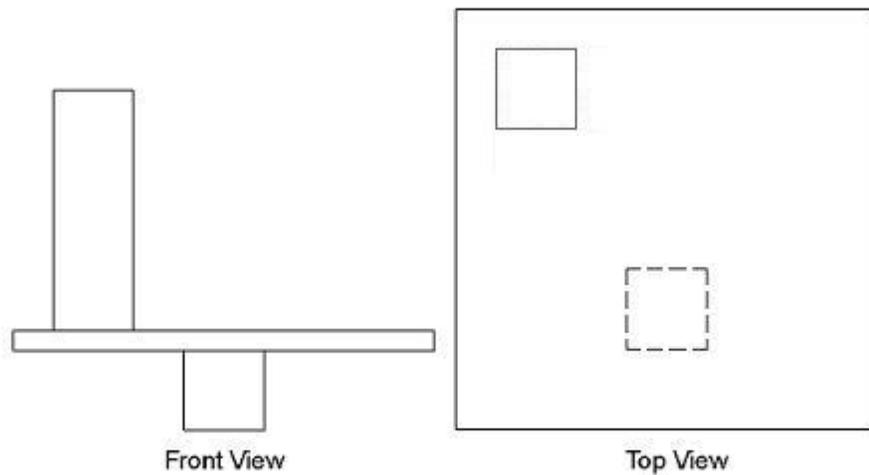
Figure 2 - Computational volume for the wood cookstove analysis

Before adding baffles, the high-fidelity model temperature profile looks like Figure 3.

Clearly, the temperature profile is varied. Near the heat source, the temperature is near 900

K. Circular patterns emerge as radial distance from the heat source is increased. The middle

of the stove has temperatures near 700 K, a 200 K drop, and near the edges the temperature

profile falls to the range of 400-600 K. This wide temperature variation renders most of the

stove surface unusable for proper tortilla cooking.

Baffle cells were added to guide the hot flue gases in the high-fidelity model. Figure 4

shows the top view of an addition of a single baffle extending across the middle of the stove

surface, and Figure 5 shows the resulting temperature profile of the high-fidelity model. The

profile seen in Figure 5 differs from Figure 3 in that there are two localized temperature

gradients, extending on both sides of the newly inserted baffle. More baffles and/or a better

arrangement could then reduce cool spots and narrow temperature variation in the model.
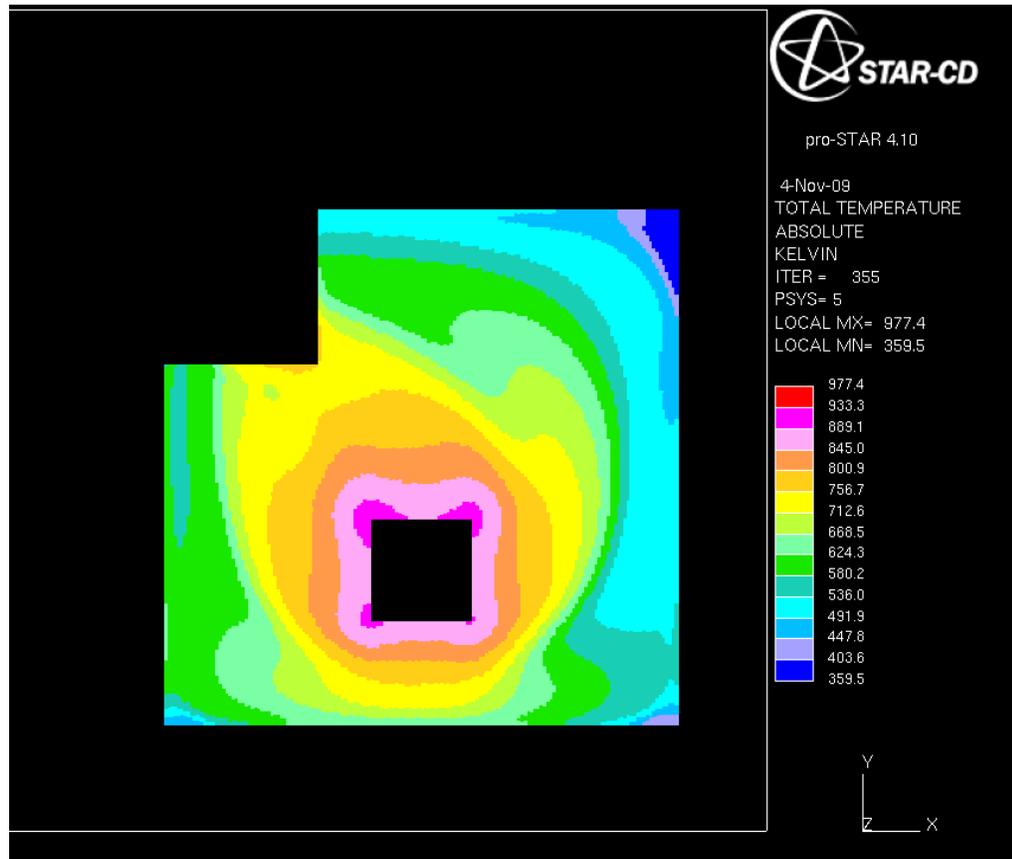
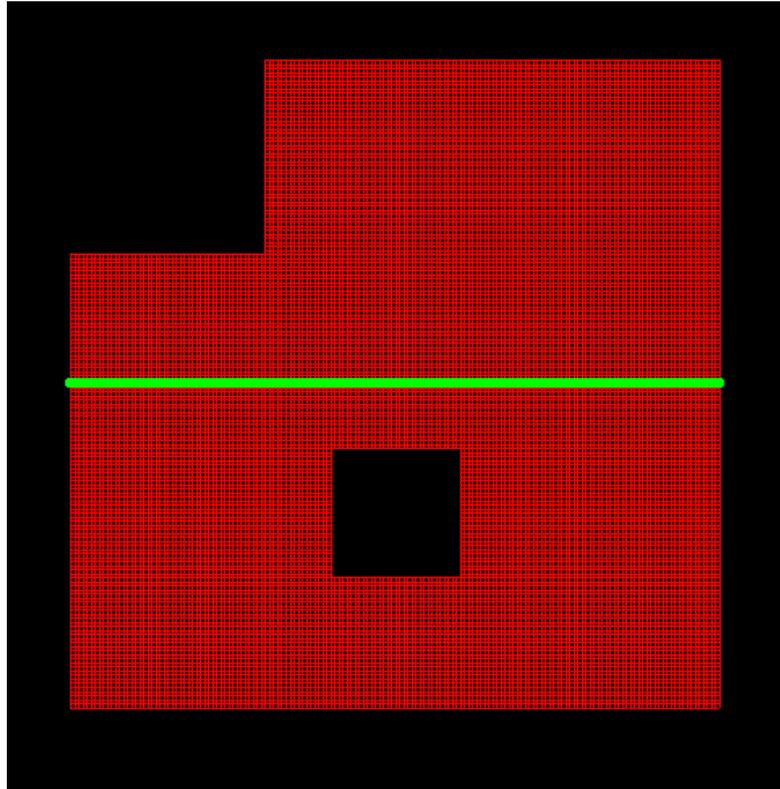**Figure 3 - High-fidelity model temperature profile with no baffles**

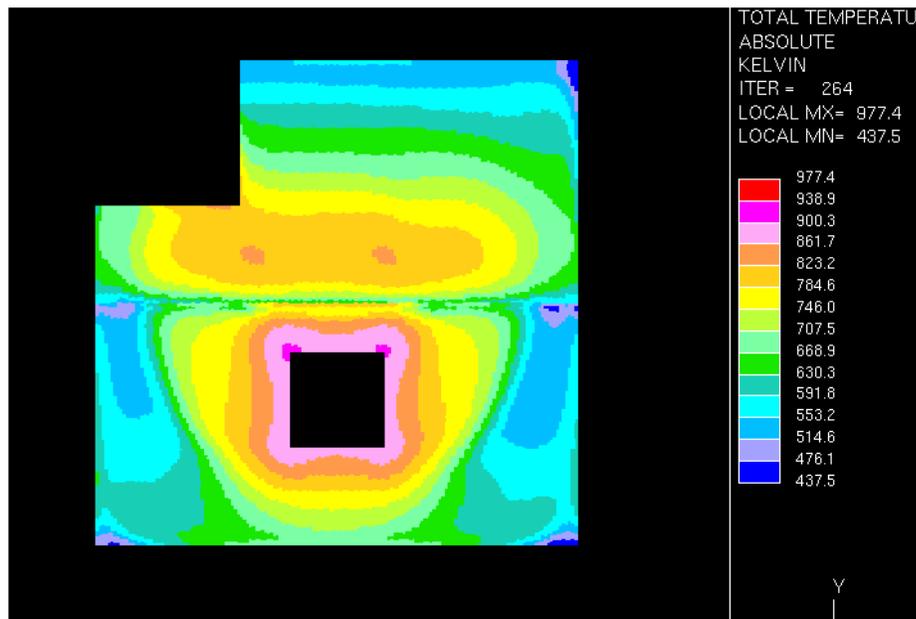**Figure 4 - Single baffle placement in high fidelity model**

Figure 5 - Single baffle placement temperature profile of high-fidelity model

CHAPTER 4 Design of the Experiment

The goal of this computational experiment is to determine if using a lower-resolution model could accurately lead to a correct design decision based on choosing between ten different design choices based on the results of varying resolution CFD models. At the highest model fidelity choices are straightforward, though time consuming. Each high-fidelity CFD model run for this relatively simple problem takes approximately three hours. Depending on the optimization and design algorithms chosen, in a practical design and optimization problem, 100 to maybe 1000 runs would be performed. In addition, most practical energy systems are much more complicated and require significantly more time and computational resources. Because of this, reducing the grid resolution of CFD models during the design phase is a common solution. However, this reduction in grid resolution is done in an ad hoc fashion without specific guidelines or goals. In this experiment, we are seeking to start to develop an understanding that can guide the process of moving CFD models from high-resolution models to practical design tools.

4.1 - Overall design of the experiment

Ten baffle layout designs were chosen based on a range from poor to very good and quantitatively verified by usable area calculations. Twenty different low-resolution models were created using five different planar (x,y) resolutions and four different depth (z) resolutions. The planar grid resolutions examined were 20×20, 40×40, 60×60, and 80×80 cells. The depth resolutions used were 3, 6, 9, 12, and 15 cells. The high-fidelity model had been validated earlier and was used as the source of absolute information. The lower-resolution models were compared with the high-fidelity models. The metrics used to

characterize the model were the nodal temperature variation, the standard deviation of the usable area, the usable area error associated with each resolution, and a selection utility function based on the nodal temperature variation and standard deviation of the usable area. The nodal temperature variation compared how accurately the lower model's nodal temperature values match the higher node location's temperature value for the high-fidelity model. The nodal temperature values were interpolated to match the location of the high-fidelity model's nodes. The usable areas were calculated by comparing the number of nodes that fell within a range of 37.5°C of the average non-interpolated nodal temperature. The standard deviation of the usable area examined the variation in usable area expected from the lower-order models. The usable area error examined the resolution effects in the planar and depth scales and how they affected the calculated usable area of the lower-order models. The selection utility metric was created by the designer to quantifiably lead to the correct answer by summing a utility associated with the nodal temperature variation and the usable area standard deviation.

The experiments were all run consecutively on similar processors over a period of two weeks. The output nodal values for temperature and velocity were stored in vector arrays and the convergence time was saved to .post files in Star-CD^{TM}. The vector arrays and .post files were converted to text files that could be read into Excel.

The nodal values for the low-resolution model were first used to calculate the usable area for the stove. The nodal values for temperature of the low-resolution model were interpolated in Excel to match the nodal placements of the high-resolution model. The nodes where chimney and flue gas entrance would be were then removed to match the high-fidelity model. The temperature error L_2 calculations were then made from these nodal values.

Normally, the designer would not know where the results of his design would fall in the design space, but the ability to define a good answer and a poor answer allows both ends of the spectrum to be used to check how the final decision is affected by using the lower-order models of differing accuracies versus their high-resolution counterparts.

The temperature error L_2 values of all ten designs per resolution were compared graphically against their computational wall time for convergence. Trend lines were fit across fixed planar and depth resolution L_2 values to examine if a correlation existed. To quantify the usable area standard deviation values, each resolution model compared the usable area of the ten different designs to their high-fidelity model usable area counterpart get a standard deviation value for all twenty models. The twenty standard deviation values were compared graphically to their average computational time and a trend line was fit to examine if a correlation existed. The usable area error values were compared graphically against their computational time across fixed planar or depth dimensions. Trend lines were added to each design to examine the effects of fixing a dimension of resolution in the lower-order models. The selection utility values were compared graphically to their average computational time. Trend lines were fit across fixed planar and fixed depth dimensions to examine the effect of a fixed dimensional resolution on the selection utility model.

4.2 - Baffle designs chosen for study

Ten different baffle layouts were selected. The ten designs were initially chosen based on the quantitative usable area range as judged by the designer to range from poor to very good. Normally a designer may create more than ten designs, but these designs were chosen because they addressed most of the range of the design landscape. Usable area calculations as specified in Section 4.5 were performed on the ten designs in the high-fidelity

model to quantitatively verify that the ten designs fell within the range. The optimal baffle layout was based on the best design in an earlier design study [4], and quantitatively had the highest usable area. An obviously poor design was created by the designer to represent the poor end of the spectrum, and its quantitative usable area was the lowest. The remaining designs were created to have quantitative usable areas that fell between the poor design's usable area and the optimal design's usable area. The baffle depth of each model was chosen to be two-thirds of an inch because it fell in the predicted range of optimum baffle depth [4].

4.3 - The high-resolution model

Based on earlier studies [4], a high-resolution model using a $160 \times 160 \times 30$ cell grid for the stove was used as the high-resolution model. The model chosen was validated and verified using standard CFD metrics and experimental data. As suggested in [1], to ensure valid models, another high-resolution model was run. In addition, all the designs were run at $160 \times 160 \times 60$ cells. Both models were validated using experimental results as discussed in [4]. Flow features and temperatures remained the same and the model was numerically stable in both cases. A 1-2% difference in the usable area (defined in Section 4.5) was observed. Thus, it was determined that it would be acceptable to use the $160 \times 160 \times 30$ cell models results as the "true" baseline.

This high-resolution model was run for each baffle design chosen for this study. The resulting temperature values at each grid point were used as the "true" state of each baffle design. To compare nodal temperatures for the lower-resolution grids, the values were interpolated across the stovetop surface to match the 160×160 cell plane resolution surface. Each high-resolution model took 2.5 to 3 hours to converge, which ranged from 15 to 1800 times longer than the lower-resolution models.

4.4 - Metamodels examined

As discussed earlier, the metamodels used in this study were created by reducing the grid resolution of the high-resolution CFD model. Grid resolution was independently varied in the planar (x, y) direction and the depth direction. Because the physical stove was square in the planar direction, the grid resolution was maintained as a square in the x-and y-directions (parallel to the cooking surface).

The planar (x, y) grid resolutions examined were 20×20, 40×40, 60×60 and 80×80 cells. This is approximately one, four, nine, and sixteen cells per square inch, respectively, in the plane resolution of each model. The depth (z) resolutions used were three, six, nine, twelve, and fifteen cells per inch. As is common in CFD models, there are more cells per inch in the z- direction than the x- and y-directions because of the importance of flow in this dimension. It should be noted that only three cells would not be used in the depth resolution of a CFD model because three cells is not sufficient to resolve even the general details of the flow (and certainly not the most important details). However, models using three cells in the depth direction were included in this study to show how increasing the depth resolution to six cells greatly improved the accuracy of the information, reduced the error rate, and allowed a better understanding of the value of the information.

4.5 - The utility metric

The performance of each of the lower resolution models on each of the proposed designs was determined using a set of computational experiments. Three performance areas were examined in each case. These were the accuracy of the temperature profile, the accuracy of the usable cooking area, and the accuracy of selecting the best design. Temperature profile and usable cooking area were chosen because they are the raw data that

would be used by the design engineer to choose a particular baffle design. The accuracy of selection was examined because the designer would generally make these choices on a relative basis, comparing the relative performance of various designs. As noted earlier, the goal of the design is to maximize the surface area of the stove available for cooking by creating a uniform temperature on the surface of the stove.

Temperature profile

The first question about the low-order models is how the temperature values at each location of the stove surface compare with those of the high-resolution model. To compare these values, the temperature values computed by the metamodels were interpolated to match the high-resolution model measurement locations. Then these temperatures were compared to their high-resolution counterpart as follows

$$L^p = \frac{\sum \sqrt{\left(T_{ij}^{n \times n} - T_{ij}^{160 \times 160}\right)^2}}{nodes_{160 \times 160}}$$

(1)

The temperature error L_2 value is in units of Kelvin. A lower temperature error L_2 value would indicate a model that has temperatures that are closer to the high-order model temperatures and a more accurate model. A higher temperature error L_2 value would indicate that a model has temperatures that are farther from the high-order model temperatures and is a less accurate model.

Usable area

The quality of the design was determined by the percentage of usable space on the stove. The area of a stove was considered usable if the temperature calculated in it was within 37.5°C of the average temperature of the surface. The acceptable temperature variation was chosen based on earlier studies in which stove users identified preferred stove characteristics. The area directly above the combustion chamber was not included in the calculation of the usable area, as it would be used for boiling water and not for cooking. The area occupied by the chimney was also not included in the usable area (Figure 6). For a given design, the usable area was divided by the available surface area to obtain the percentage of usable surface area.
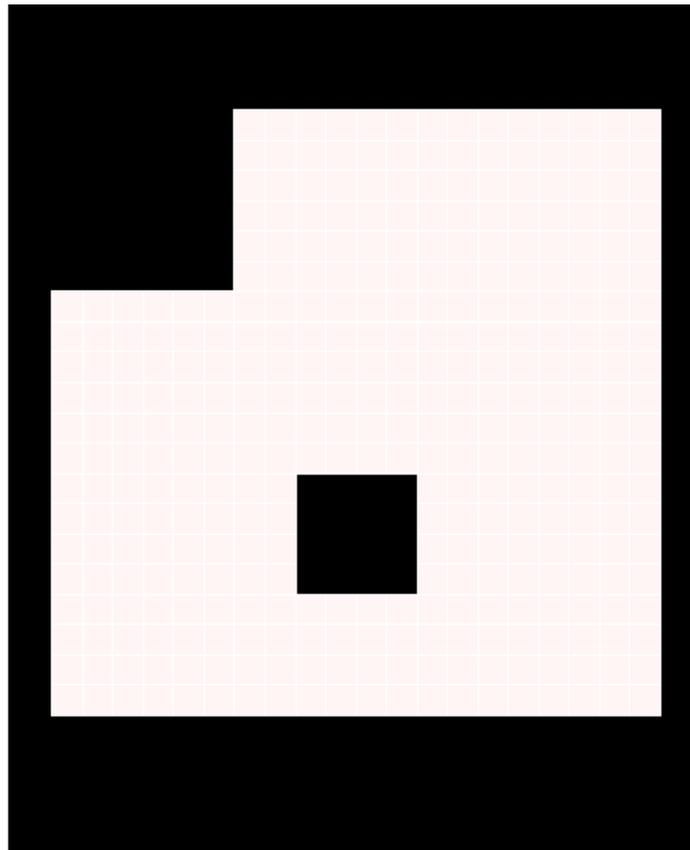
Figure 6 -Usable area of stove

Two metrics, usable area standard deviation and usable area error, were used to compute the difference in usable area calculated for the lower-resolution model as compared to the usable area estimated for the high-resolution model. The standard deviation across the ten designs in the design space was calculated to quantify the variance in the usable area calculation.

$$\sigma = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (x_i - \mu)^2},$$

(2)

where μ = high-resolution model usable area value and x_i= low-resolution metamodel usable area value. Using this metric, a large standard deviation for a given resolution indicates a less accurate model. A lower standard deviation area error value would indicate that a model was closer to a high-resolution model and a more accurate model.

To compare error associated with increasing resolution in the depth versus increasing resolution in the plane, the following equation was used:

$$A_{error,n \times n} = A_{usable,160 \times 160} - A_{usable,n \times n}$$

(3)

The difference between the high-resolution model usable area and the low-resolution model usable area was calculated and compared as one resolution (depth/plane) was held constant while the other resolution (plane/depth) was increased. If the models were found to favor one resolution over the other, further investigation was performed to see how that resolution affected the robustness of a model.

Selection utility

The last and most important question of the low-order models is whether the model leads the designer to the correct answer. To quantifiably answer this, we will evaluate the utility of each resolution by using the temperature error L_2 value and the usable area standard deviation value σ and comparing their combined predictive ability to lead the designer to a good design. As stated in Section 2.4, the creation of the value-of-information metric must come from the designer [13], who should have the most familiarity with the problem and can make a usable utility equation. The following equation was created:

$$Utility = \left|\left(\frac{200K - L_2}{200K}\right)\right|\omega_{L_2} + \left|\left(\frac{50\% - \sigma}{50\%}\right)\right|\omega_\sigma$$

(4)

where $\omega_{L2} = 0.5$ and $\omega_\sigma = 0.5$. Equation 4 combines temperature error L_2 values and the usable area variance σ values by making each component non-dimensional. The selection of 200 K for temperature error and 50% for variance of usable area as comparative values was made based on the worst-expected-case scenario, where the designer felt that the worst performing models would not have values exceeding 200 K for L_2 and 50% for σ, respectively. We assume that a "good" low-order model is represented equally by the L_2 value and the usable area standard deviation value. The weight given to each value is 0.5.

The ideal low-order model utility would have a value of unity, 1, because its temperature component, L_2, and usable area component, σ, would both be 0. The lowest low-order model utility value is 0. Using the absolute value of both components in the utility function addresses the issue of either component exceeding their respective worst-expected-case values. For models with values slightly above 200 K for L_2 or 50% usable area for σ, (i.e. 205 K or 52%) the utility value low will still be low. If the model has very large

temperature fluctuations above 200 K or usable area standard deviation values above 50%,

(i.e. 300 K or 75%), the designer should remove the model from the design space.

       To find a utility value, the designer needs to compare the magnitude of the utility

value that leads the designer to the correct decision. This can be done by correlating trends

found from the utility plots or by verification of some goal/selection as pre-determined by the

designer.

CHAPTER 5 Results

Using the protocols discussed in Chapter 4, computational runs were completed (ten designs × twenty lower-resolution models/design). The results of these runs were analyzed using the metrics described in Section 4.5 and the results plotted. These plots are presented in this chapter and a number of conclusions are drawn.
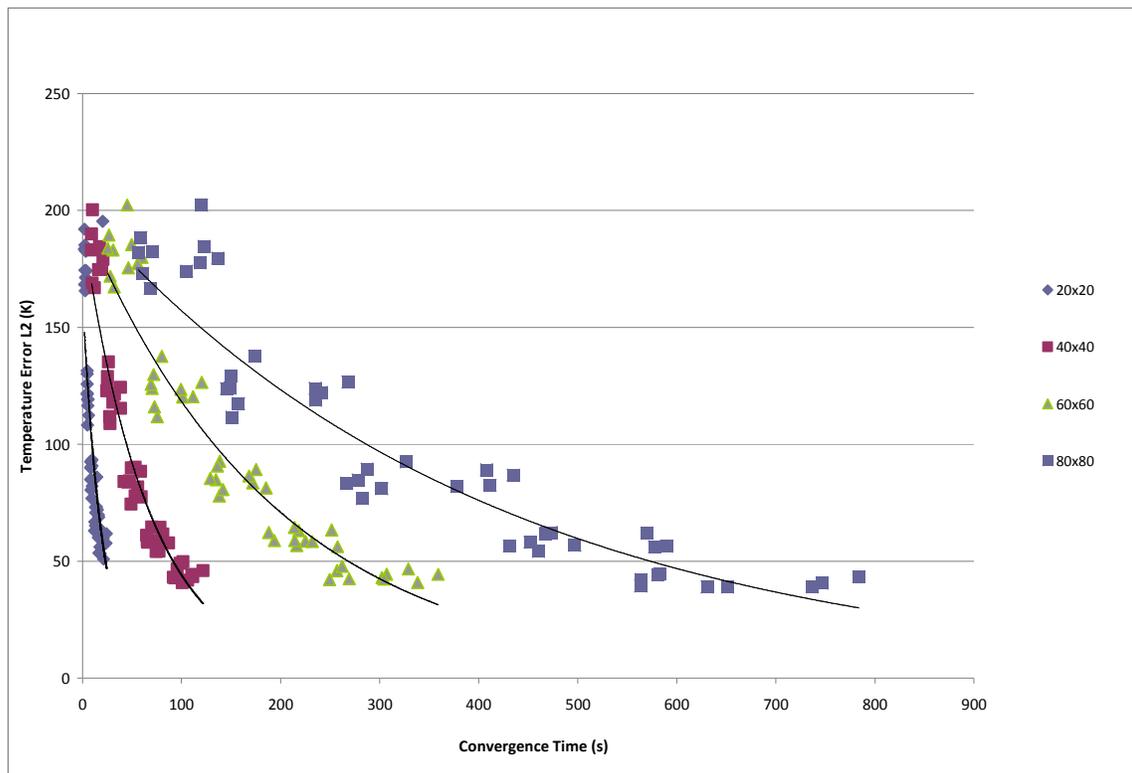
5.1 - Temperature

As would be expected, the accuracy of the metamodels in predicting the temperature is strongly dependent on which planar resolution or depth resolution is being considered. Figure 7 shows the temperature error as a function of computational time for various grid-resolutions in the depth axis. As shown, the errors are strongly a function of the grid resolution in the direction of the depth and to a lesser degree the resolution of the grid in the xy plane. In fact, there is very little crossover between the various depth grid resolutions. In addition, as the computational time increases at any fixed-depth grid resolution, the error is about the same. This implies that at a fixed-depth grid resolution, the model accuracy appears to be significantly less dependent on the xy plane grid resolution than on the grid resolution of the depth. Also, the range of the error band decreases as the depth resolution increases. This means that as the fixed-depth grid resolution increases, the range of the temperature variation of the design space narrows. Thus, at low fixed-depth resolution, the average temperature values are farther off from the true value and the range of the accuracy band in the design space is wider. Conversely, at high fixed-depth resolution, the accuracy temperature values are much closer to the true value and the range of the accuracy band in the design space is narrower.
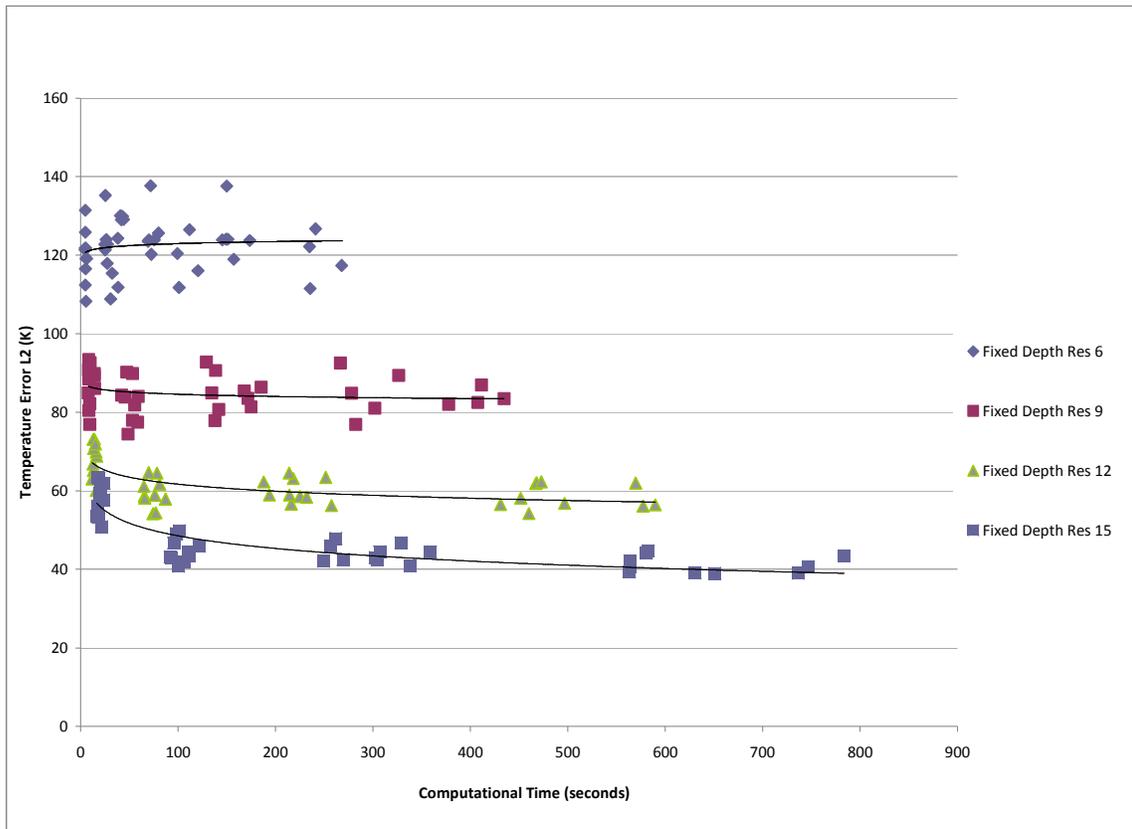
Figure 7 - Temperature error as a function of computational time for fixed-depth resolutions

Figure 8 shows the temperature error as a function of computational time for various grid resolutions in the xy-plane. As shown, the error decreases with increasing depth resolution. This is consistent with observations made in the fixed-depth grid resolution graph. Second, as the fixed-plane resolution is increased, the cost band for the design space increases without significant increases in accuracy. At the lowest fixed-plane resolution, the range of the cost band is about 20 seconds. At the highest xy-plane resolution, the range of the cost band is about 650 seconds. This means that nearly the same error values occur regardless of the fixed-plane resolution. This implies that the accuracy of the model for predicting temperature profiles is independent of the grid resolution in the xy plane.

**Figure 8 - Temperature error as a function of computational time for fixed-plane grid resolutions**

5.2 - Usable cooking area

As discussed in Section 4.5, the usable cooking area of the stove was determined by calculating the percentage of the available cooking surface within 37.5°C of the average surface temperature. The acceptable temperature variation was chosen based on earlier studies in which stove users identified preferred stove characteristics. The usable area was used in the high-fidelity models to quantitatively rank designs from poor to very good. Therefore, it is important to identify how the calculated usable area differed between the high-fidelity models and the lower-resolution models. If the lower-resolution models have some inaccuracy of usable areas in comparison to the high-fidelity model, it is important to quantify that error and understand how it affects the designer's decision. The standard

deviation for ten designs shows the variance in usable area calculated for lower-resolution models. It is an easily replicable approach that also renders a quantitative metric that can be used in the selection utility metric. The area error metric quantifies the difference between the high-fidelity model's usable area value to the lower-resolution model's usable area value. It is used to identify the role that fixed resolutions played on usable area if any existed and how it affects which resolutions a designer uses.

5.2.1 – Usable area variance

For each low-resolution model, a usable area is calculated for ten designs. The usable area for each design in the lower-resolution model is compared to the usable area for each design in the high-fidelity model. The ten differences in usable area are used as specified in Section 4.5 to calculate a standard deviation value for that particular lower resolution. The ten designs' computational times are averaged to a single time quantity for each lower-resolution model. This created twenty standard deviation values from twenty different lower-resolution models. Figure 9 shows standard deviation values as a function of average computational time. As shown, the standard deviation of usable area decreases exponentially with increasing average computational time. At low grid resolution, the standard deviation is high. As grid is increased, the rate of decrease in standard deviation diminishes. This trend would suggest that the designer use a model with a grid that results in a moderate computational time (100-200 seconds) because further time invested does not significantly decrease the standard deviation value. The selection utility metric uses the standard deviation metric to assess its impact on the designer's decision.
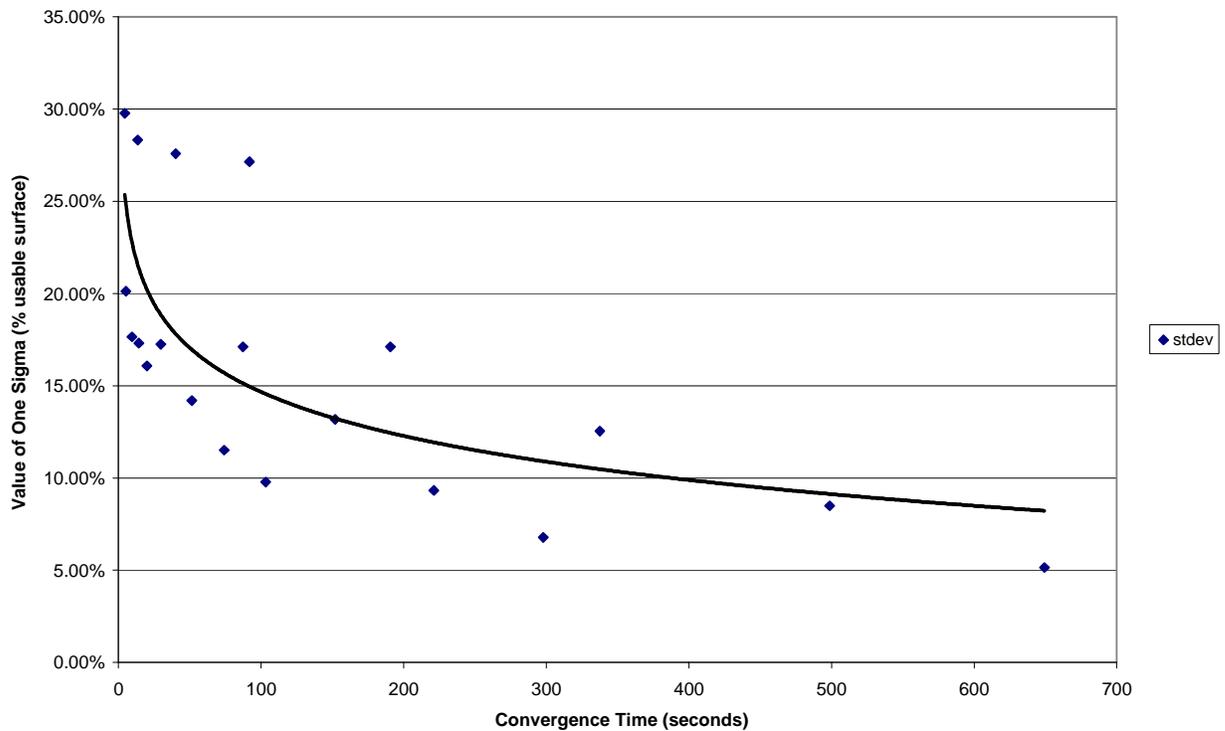
Figure 9 - Usable area standard deviation values chart
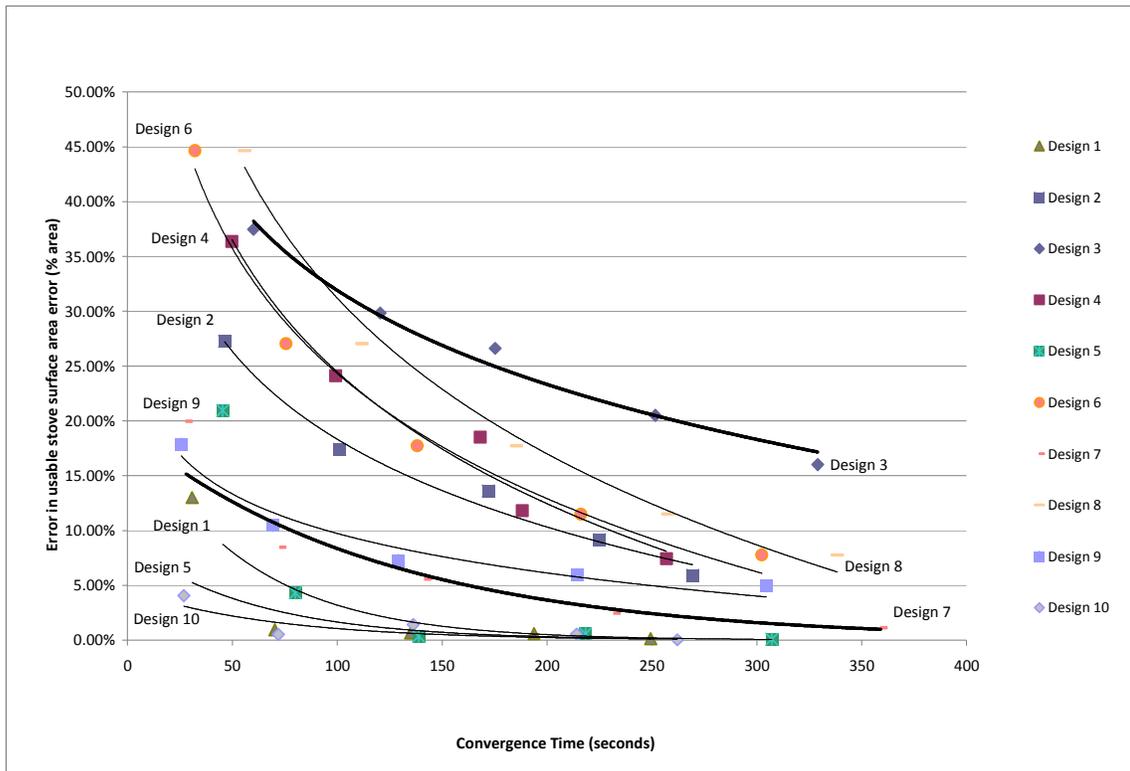
5.2.2 – Usable area error

The area computed by the metamodels is used by the designer to choose between competing designs. In the same manner as the error associated with the temperatures was examined, variations in grid resolution in the depth and planar directions are examined. Figure 10 shows the error in the usable stove surface area as a function of computational time and baffle design for changing grid resolution in the xy plane with a fixed grid resolution of 15 cells in the depth dimension. As shown, the error associated with each design decreases exponentially as the plane resolution increases. At low-plane resolution, the area error is high. As the plane resolution is increased, the rate of decrease in error diminishes.
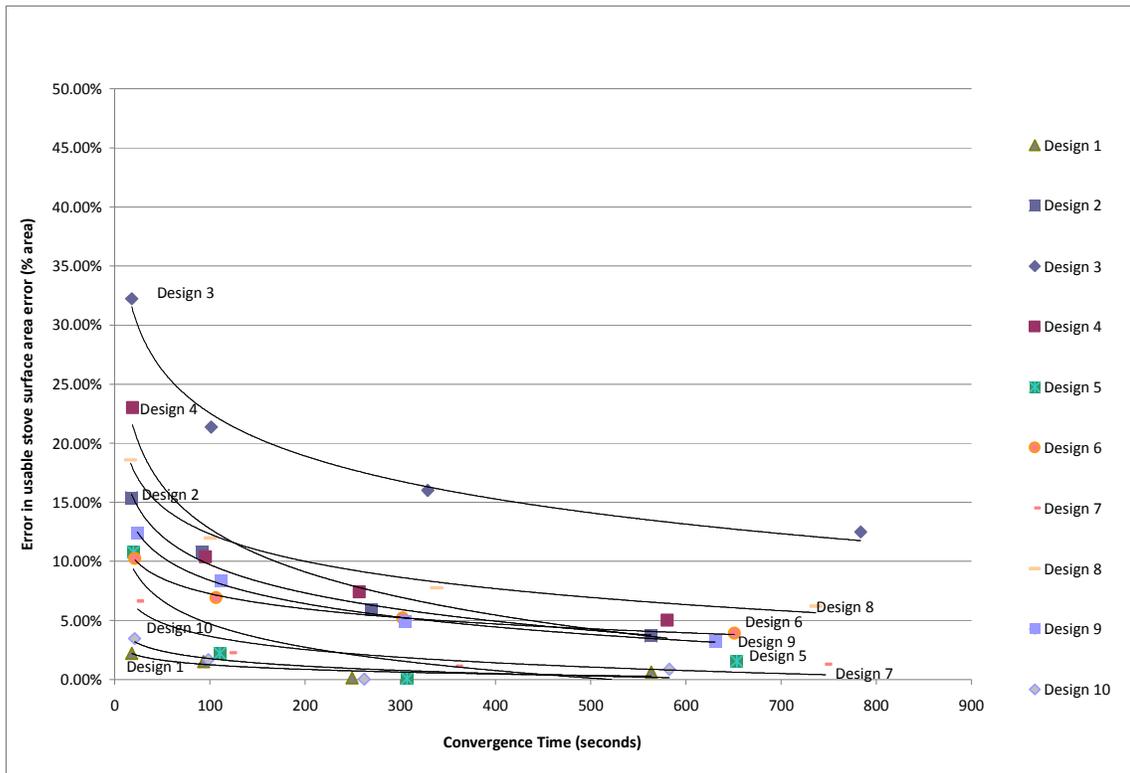
Figure 10 – A$_{error}$ usable area values for fixed grid resolution of 15 cells in the depth dimension

Figure 11 examines the error associated with the computed usable area as a function

of computational time and baffle design for changing grid resolution in the depth plane with

a fixed-grid resolution of 60×60 cells in the xy plane. Once again, as the computational time

increases, the error associated with each result decreases exponentially. At low time cost (a

result of low resolution), the computed usable area has an error of 5% to 45% depending on

the design. As time invested in the computation increases, the upper end of the error band

drops substantially. This would validate part of a value-of-information model, which calls for

diminishing returns as more accurate information is gathered.

Figure 71 -A$_{error}$ usable area values for grids of fixed plane resolution

In both Figures 10 and 11, there is a wide band of area error. At the higher end of the

error band are Designs 3, 4, and 8 in both the fixed grid planar resolution and fixed grid

depth resolutions. These designs are ranked as the top performing designs by the quantitative

measure of usable area in the high-fidelity model and are included the optimal design

(Design 8) as found in [4]. When examining the raw data, the low-resolution models

overestimated the usable area, further separating the top performing designs from the other

designs. For example, the high-fidelity model's usable area for Design 8 was 30%. For the 40

$\times 40 \times 06$ grid model, the usable area for Design 8 was 55%. Conversely, at the bottom end

of the error band of both fixed resolutions were Designs 1, 5, and 10. These designs are

ranked as the lowest-performing designs by the quantitative measure of usable area by the

high-fidelity model. This maintains a more accurate usable area of poor designs in the lower resolution model, so the poor designs still return a low usable area value in the low-resolution model. For example, the high-fidelity model's usable area for Design 1 was 19%, whereas the $40 \times 40 \times 06$ grid model's usable area for Design 1 was 23%. The lower model's usable area would affect the designer's decision by pushing the good designs farther from other designs and maintaining the poor design at the lower end of the usable area spectrum. This may lead the designer to using the lowest order models. However, the average performing designs have error values that cross over (Design 2 crosses Design 4 in Figure 10) or have even higher usable error values than those of the higher performing designs (Design 6 has the highest usable area error in Figure 11) at the lower convergence times (the result of low resolution), which could lead to average designs getting higher usable areas than good designs in the metamodel. This pushes the designer to use a higher resolution model. The usable area error metric gives the designer insight on how the lower-resolution model works for very good or poor designs, but it alone cannot lead the designer to the optimal grid choice because it does not effectively deal with the entire design space and could lead the designer to a sub-optimal design selection.

An explanation for high error in usable area for good designs and low error in usable area for poor designs may lie in the general rule of thumb in CFD modeling that a minimum of 10 cells must be used in the flow for it to make sense. In fact, it is surprising that any applicability is observed in grids with depth resolutions of 3 or 6 cells. The usable area metric is based on the results from a two-dimensional surface, but the CFD model itself is a three-dimensional object. This means the surface cells are affected by the temperature and velocity effects of their neighbors beside *and below* them. As the temperature error graphs

show, changing the depth resolution affects the overall accuracy more than changing the plane resolution.

As stated previously, all the baffles were made to be two-thirds of the stove depth. Thus, the bottom third of the stove cells directly below the baffles are "free" in that there are no baffles guiding fluid flow. For poor depth resolution, this may lead to the "free" cells having unrealistic thermal and velocity values because there are so few points to accurately model the flow. This can be easily seen in the magnitudes of both the temperature and velocity values (Figures 12 and 13). In designs that have baffle locations near the heat source, the "free" cells are observed to have higher and unrealistic temperature and velocity values as the depth resolution is decreased. Generally, poor quality designs have baffles located farther away from the heat source. We postulate that the temperature and velocity values near these "free" cells are lower because the heat source has to travel farther before it reaches the baffled areas. The magnification effect is still the same for a given resolution, but the temperature and velocity values being magnified are lower because they are farther from the heat source. Thus, designs of poor quality would have lower error values because the lower magnified temperature and velocity values would lead to smaller "free" cell vector values. As the depth resolution increases, the magnification effects on the error values of the "free" cells decrease. Figures 12 and 13 show the velocity profile of a baffle placed directly next to the heat source. In Figure 12, the depth resolution is 3 cells. In the bottom third of the model, the "free" cells only get the single velocity value of about 4 m/s. In Figure 13, the depth resolution is 15 cells. In the bottom third of the model, the "free" cells have five times the amount of heat transfer calculations and the velocity values of the fluid have a range of 2 m/s to 4 m/s for the same design. To further validate the hypothesis about the effect of baffle
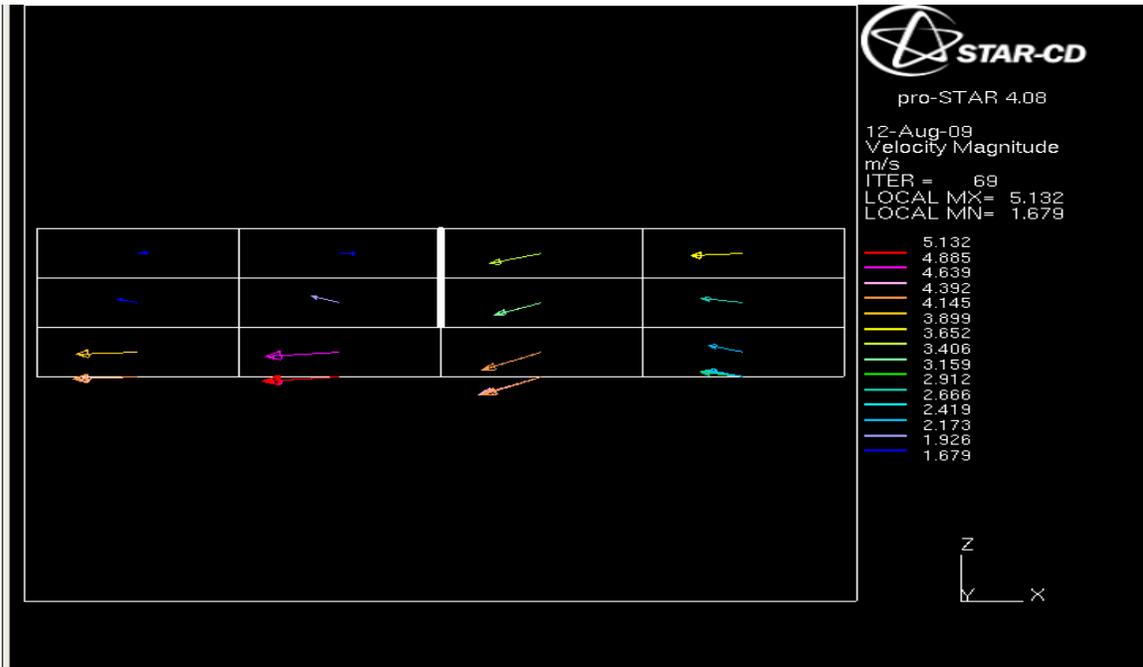
Figure 82 - Velocity magnitude around baffle at 3 cell depth
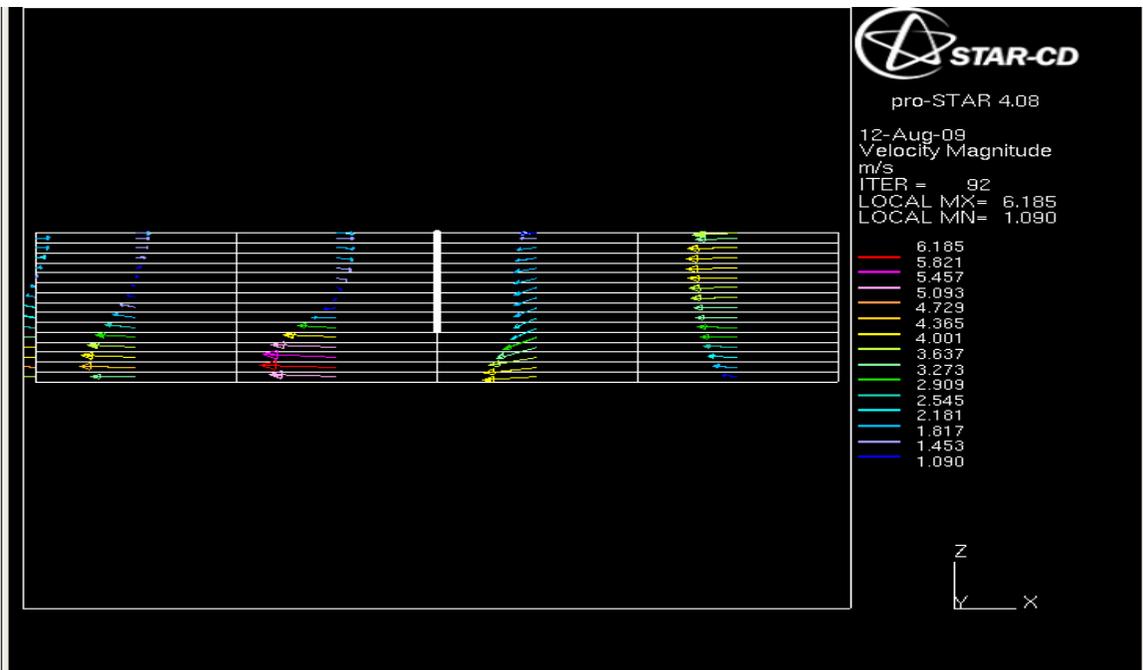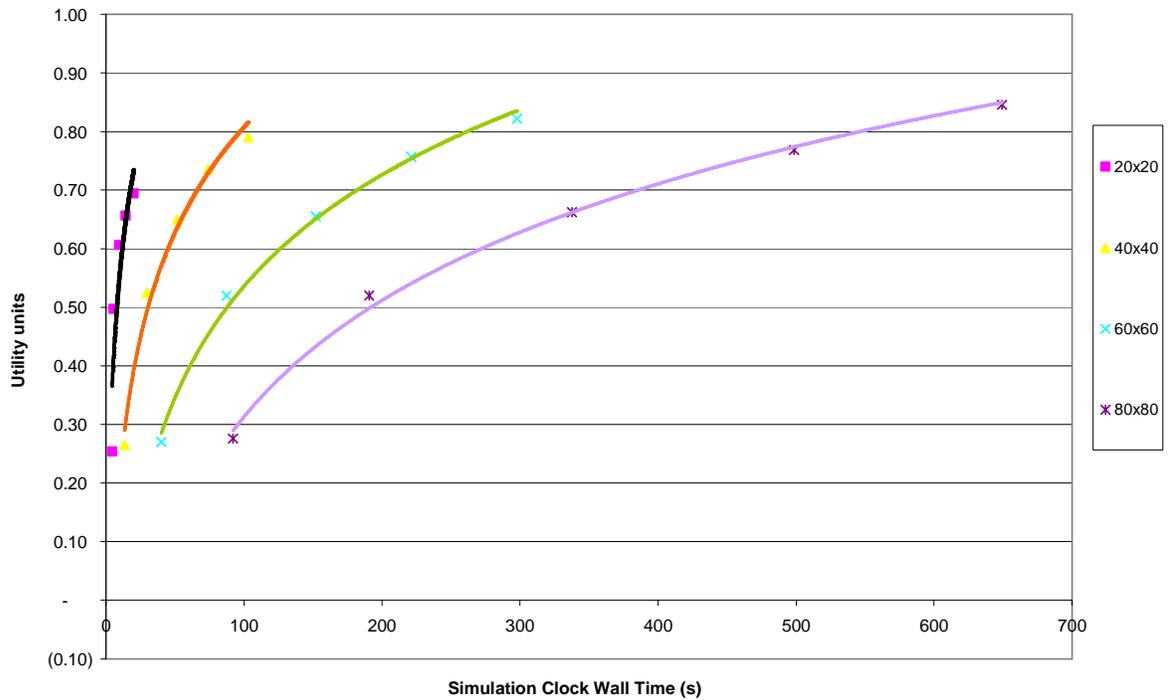
resolution

Figure 93 - Velocity magnitude around baffle at 15 cell depth resolution

placement near the heat source, when checking the layout of the one poor quality design with high overestimation of surface area, the layout does indeed have two baffles directly over the heat source, but is an overall poor design. This verifies the rule of thumb that at least 10 cells should be used in the direction of the flow in a CFD model. Thus, a designer would want to use a higher grid in the depth resolution.

5.3 - Lower-resolution model performance accuracy

Having examined the accuracy of various lower-resolution models as a function of variable, time, and design, the question still remains how much accuracy is needed to assess design performance to pick the best design and at what cost. Selection utility in this study analyzed whether the lower-order models would render the top performing designs as rated by Equation 4. Figure 14 shows selection utility as a function of computational time for various xy plane grid resolutions. As shown, the selection utility shows increased gains as computational time is increased. This is expected since selection utility is derived from the temperature error L_2 value and usable area standard deviation σ, both of which decrease as computational time increases (as noted in Section 4.5, as these two components tended to zero, the selection utility approached the ideal value of unity, 1). This indicates that the varied resolution, the depth resolution, has a direct relationship with the utility, and thus, a direct relationship with the value of the lower-grid resolution.
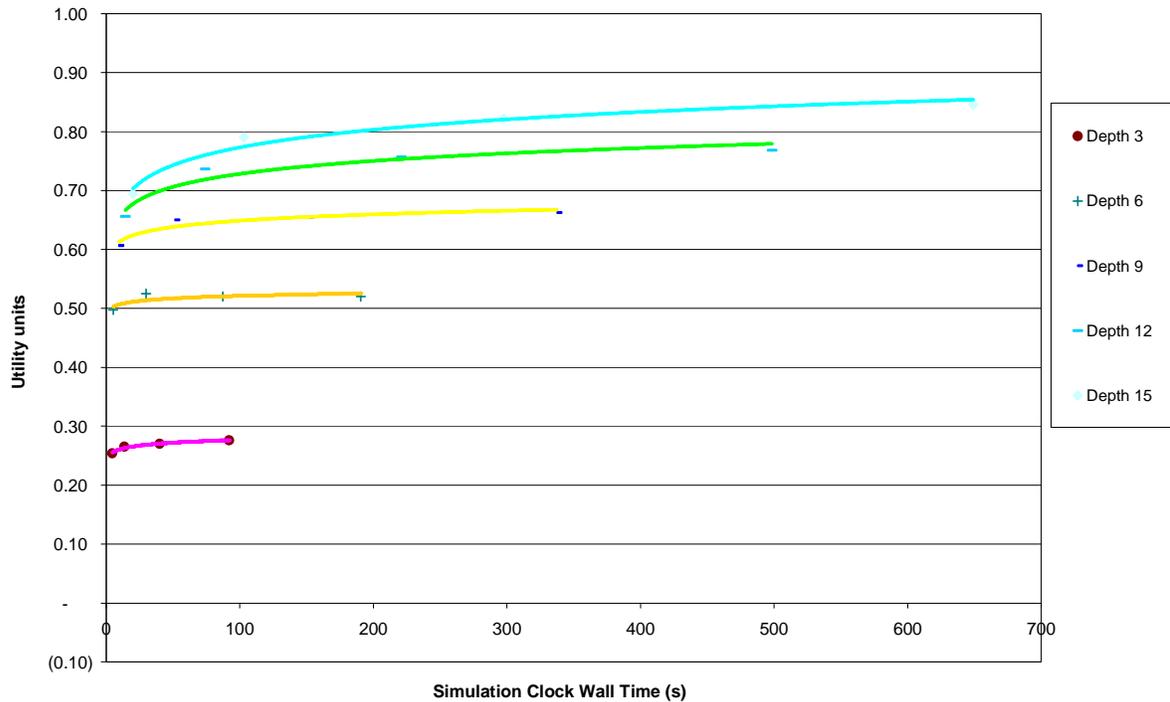
Figure 14 – Utility values for models at fixed plane resolutions

Figure 15 shows selection utility as a function of computational time for various

depth grid resolutions. As shown, there is small gain in selection utility as the computational

time is increased. This is also expected since the utility value is partially derived from the

temperature error L_2 value, which exhibits the very low gains as computational time is

increased in its fixed-depth resolution graphs. The usable area standard deviation σ decreases

exponentially with increasing average computational time, which would explain the small

gains at low computational times and almost none at high computational times.

Using the trends of selection utility as guidelines for the optimal grid resolution, a

high depth resolution should be selected because it has a direct relationship with utility, and a

low plane resolution should be selected to reduce cost since it exhibits much lower returns in utility. From the view of improvement potential from [12], the value of using a $20 \times 20 \times 15$

Figure 10 – Utility values for models at fixed depth resolutions

grid as compared to no information returns a utility value of 0.7 for a cost of 20 seconds. If we chose to invest further in a $40 \times 40 \times 15$ grid, the value increases by 0.1 of an additional investment of 80 seconds. Both grids offer a high selection utility value at low cost, so the next comparison is to find how the selection utility value corresponds to the desired results of the problem. To verify the selection utility value to the low-order model selections, the results comparing the quality of the designs showed that the low order model veritably filtered the top three designs from the rest for both the $20 \times 20 \times 15$ grid and the $40 \times 40 \times 15$

grid. Therefore, with preference being to minimize cost (in this case, computational time), the $20 \times 20 \times 15$ grid would be used.

In summary, the $20 \times 20 \times 15$ grid would be selected to help distinguish the good quality designs (in this case, the upper 30%) from the rest of the tested lower-resolution models. The computational wall time for these models was in the range of 20-25 seconds per run. At 500 times less than the average convergence time of the high-resolution model, using the low-order model here saves the designer time or enables him/her to test a further scope of designs while checking if the layout itself is even worth pursuing. For example, a designer could employ the lower-resolution model in a teaching assistant program that would help identify whether the proposed design would be worth pursuing with further research. It could also be used in conjunction with visual recognition software to teach the computer/student what to look for in the optimal baffle layout.

5.4 - Analysis

After the initial runs found the low-order model dependent on depth resolution and more independent of plane resolution, stove models of $20 \times 20 \times 30$ and $20 \times 20 \times 60$ cell resolution were created and the ten baffle designs were run on each resolution. By doubling the depth resolution from 15 cells to 30 cells, nine of the ten cases saw significant area error reduction, and the tenth had a minor increase. This came at a cost of approximately 60 seconds of computational time. However, by stepping up the depth resolution to 60 cells, the error values for four out of the ten designs had *increases* in area error and the other six had either no or minor area reduction. This occurs because of poor aspect ratio of these cells.

Before, having too few cells in the depth resolution magnified the inaccuracies in both the thermal and velocity values (as discussed in Section 5.2.2). However, the aspect

ratio of the cell must be considered. Each cell in this stove model looks like a very thin plate. The X-Y geometry of the cell is about one centimeter in each direction, whereas the depth of the cell is equivalent to less than $1/20^{th}$ of a centimeter. As a result the error values increased. The way to counteract the effects of the thin plate cells would be to increase the plane resolution of the model, but again, that would greatly increase the convergence time. Further investigation could show that some ideal depth resolution for a 20×20 fixed plane resolution of this model exists, and that it could lie between 15 and 60 cells.

CHAPTER 6 Conclusions and Future Work

In summary, it was found that lower-resolution models could be validated to accurately model high-fidelity CFD models. The process starts with a designer who is familiar with the model and the problem creating some comparable quantifiable metrics that describe the models. The designer uses the metrics to create a selection utility metric, which is used to aid the designer in finding a low-resolution grid that is low cost but returns accurate answers.

In the case of the plancha cookstove, even at moderate temperature and area errors, low-order models could lead the designer to separate good designs from poor designs at a considerably lower computational wall clock time. Depth resolution was observed to play an important role in design selection of the cookstove model, as well as temperature and area accuracy. Depth resolution can be increased at a fixed-plane resolution with little additional expense of computational wall time. Even when combining high depth resolution with low plane resolution, the low-order model still helped lead the designer to the upper 30% of the design space. However, adding too much depth resolution at low plane resolution can lead to errors due to poor aspect ratios, which cuts into the correct design filtration. Thus, some ideal resolution for this model must exist. The proposed resolution for this cookstove model would be $20 \times 20 \times 30$.

Possible future work includes determining the ideal depth resolution for low-plane resolutions, such as 20×20, to minimize temperature and area error. Future work could also explore making a confidence interval in how well a low-order model will estimate the usable area of the stove using the standard deviation values associated with each resolution. In addition, an error magnification characteristic could be investigated for the varying depth

resolutions, especially low-depth resolutions. Also, the concept of design filtration could be

combined with an image/pattern recognizer as a trainer to help find a visual pattern for

sorting the quality of the designs.

REFERENCES

[1] Barton, R., Computationally Inexpensive Metamodel Assessment Strategies, *AIAA Journal* 40:10 (2002) 2053-2060

[2] Sargent, R. G., Verification and Validation of Simulation Models, *Proceedings of the 2007 Winter Simulation Conference* (2007) 124-137

[3] Robinson, S., Simulation Model Verification and Validation: Increasing the Users' Confidence, *Proceedings of the 29th Conference on Winter Simulation* (1997) 53-59

[4] McCorkle, D. and Bryden, K.M., A New Methodology for Optimization of Energy Systems, *Computer Methods in Applied Mechanics and Engineering* 192 (2002) 5021-5036

[5] Ling, J. M., Aughenbaugh, J. M., and Paredis, C., Managing the Collection of Information Under Uncertainty Using Information Economics, *Journal of Mechanical Design* 128 (2006) 980-990

[6] van der Weide, Th.P., van Brommel, P., Measuring the Incremental Information Value of Documents, *Information Sciences* 176 (2006) 91-119

[7] Toth, B., Scala, E., Huber , J., and Kirchler, M., The Value of Information in a Multi-Agent Market Model, *The European Physical Journal B* 55 (2007) 115-120

[8] Chorus, C.G., Arentze, T.A., Molin, E., Timmermans, H., and Van Wee, B., The Value of Travel Information: Decision Strategy-Specific Conceptualizations and Numerical Examples, *Transportation Research Part B: Methodological* 40:6 (2006) 504-519

[9] Koninx, J.P., Value-of-Information – from Cost-Cutting to Value-Creation, *Journal of Petroleum Technology* 53:4 (2001) 84-93

[10] Ketzenberg, M.E., Rosenzweig, E.D., Marucheck, A.E., and Metters, R.D., A Framework for the Value of Information in Inventory Replenishment , *European Journal of Operational Research* 182 (2007) 1230-1250

[11] Leach, P.E., Brown, P.J., and Haskett, W., Value-of-Information Applications in Unconventional Resource Plays, presented at the 2005 SPE Annual Technical Conference and Exhibition, SPE 96879 (2005) 1-11

[12] Panchal, J.H., Paredis, C., Allen, J.K., and Mistree, F., A Value-of-Information Based Approach to Simulation Model Refinement, *Engineering Optimization* 40:3 (2007) 223-251

[13] Fekete, J.D., van Wijk, J.J., Stasko, J.T., and North, C., The Value of Information Visualization, *Information Visualization LNCS* 4950 (2008) 1-18

[14] Proleña–Nicaragua, Alternativas Viables Para Solucionar el Problema de Demanda de Leña en la Región Las Segovias, ed. S.F. Alves-Milho, Managua, Nicaragua, 2000

[15] Instituto Nicaragüense de Energia (INE), Memoria INE 1997, Managua, Nicaragua, 1997

[16] Hong, C.J., Health Aspects of Domestic Use of Biomass Fuels and Coal in China, Shanghai Medical University, Shanghai, China, 1994

[17] Barnes, D., Openshaw, K., Smith, K.R., van der Plas, R., What Makes People Cook with Improved Biomass Stoves? World Bank Technical Paper: Energy Series 242 (1994) 1–39

[18] Malhotra, P., Environmental Implications of the Energy Ladder in Rural India, *Boiling Point* 42 (1999) 3–5

[19] Pandey M., Health Risks Caused by Domestic Smoke, *Boiling Point* 40 (1998) 6–8

[20] Urban G., The EcoStove: A Case Study in Thermal System Optimization Based on Differential Analysis, Master's thesis, Iowa State University, 2001