

Early Risk Quantification Strategy for Design Space Reduction Decisions in Set-Based Design

J.B. Van Houten^{1,*}, A.A. Kana², D.J. Singer¹ and M.D. Collette¹

ABSTRACT

Perceptions of feasibility in design spaces are susceptible to change if new and conflicting information becomes available later. Design space reduction decisions made in set-based design can amplify vulnerability to new information if remaining design spaces and present perceptions are unable to adapt. This paper considers different ways new information can alter perceptions of feasibility for complex design problems and introduces an early, probabilistic strategy for quantifying the risk of eliminating potential design solutions based on the vulnerability of remaining design spaces to new information. Emergent designs of a set-based design process gauging this risk are evaluated against one neglecting it for an analogous design problem. Early results indicate that the probabilistic model is able to effectively delay design decisions and prevent lock-in while design space understanding is still growing.

KEY WORDS

Set-based design; Space reductions; Information; Fragility; Risk quantification.

INTRODUCTION

Design decisions made within the web of interdependencies and requirements ingrained in the marine design process produce complex knowledge structures. While different methods have been proposed to characterize the knowledge generation accompanying these decisions (Braha and Reich, 2003; Hatchuel and Weil, 2009; Shields, 2017; Goodrum, 2020), each one seeks to track and better understand the emergence of (or lack thereof) design solutions. Decisions made in set-based design (SBD) build up these knowledge structures gradually, but they also leave design spaces vulnerable to emergent design failures if the information supporting them changes. Providing designers with a tool to understand the potential impacts of new information on a reduced design space after eliminating potential design solutions from consideration would assist them in making more informed space reduction decisions.

Using iteration to make decisions and generate knowledge is an understood reality of many complex design problems (Wynn and Eckert, 2017). Different studies have investigated how enhancing the allocation of resources (Smith and Eppinger, 1997) or communicative pathways (Mihm and Loch, 2006; Parraguez et al., 2015) between iterative tasks can promote more efficient information flow. As these strategies are improved upon to assist with iterative design decisions, they can fixate a designer's knowledge on one decision path, restricting the solutions that can be attained through others (Page, 2006). Examples of this fixation are shown in Van Houten et al. (2022) where viable solutions within a discipline's design space are significantly limited by the path chosen. In some cases, designers can lose influence if a finite set of absorbing paths constrain the knowledge structures generated from these temporal decision processes (Niese et al., 2015; Kana, 2017).

¹ Naval Architecture and Marine Engineering, University of Michigan, Ann Arbor, USA

² Maritime Transport and Technology, Delft University of Technology, Delft, the Netherlands; ORCID: 0000-0002-9600-8669

* Corresponding Author: joeyvan@umich.edu; ORCID: 0000-0001-6569-8524

A consequence of becoming overly fixated on a particular decision path is leaving a design susceptible to emergent design failures. Dong (2017) discusses the prevalence of this problem in product development when companies introduce innovative technologies into their product's existing functional architecture. He argues that integration issues arise before the establishment of their product's physical architecture and should instead be attributed to the solution principles the design team committed to during development. He and others (Shields and Singer, 2017; Goodrum, 2020) insist that understanding emergent design failures requires a shift in viewing them from a product-centric to a knowledge-centric perspective. As Goodrum (2020) explains, a design decision is a commitment to a knowledge structure, and how those decisions affect future design activities will vary depending on how new knowledge integrates with existing knowledge.

Set-based design (SBD) protects against emergent design failures stemming from path fixation by having design decisions focus on eliminating undesirable regions rather than making premature commitments to hard-set characteristics. By delaying commitments and keeping variable sets open, SBD decisions create low-risk knowledge structures (Shields and Singer, 2017) and allow designers to maintain influence over the design problem while their understanding of it grows (Bernstein, 1998; Singer et al., 2009). Advantages of SBD include basing the earliest and most critical design decisions on acquired data, promoting institutional learning within the design environment, encouraging concurrence in the design and manufacturing process, and supporting a search for more globally optimal designs (Ward et al., 1995). These advantages have fueled US Navy interest in making ship design and analysis tools compatible with SBD methods (Doerry, 2012) and applying SBD to various projects such as the Ship to Shore Connector (Mebane et al., 2011), Amphibious Combat Vehicle (Burrow et al., 2014), and Small Surface Combatant (Garner et al., 2015). Despite the advantages, it is still either infrequently applied to problems in industry or generally confined to introductory design stages (Toche et al., 2020). Singer et al. (2009) claim SBD's biggest obstacle in naval design coincides with current government acquisition policies conforming to point-based methodologies. Other hurdles are summarized in McKenney and Singer (2014) and Gumina (2019) and involve having to manage misconceptions *about implementation* and lacking a regimented process *for implementation*.

The SBD implementation process is multifaceted and has disciplines individually explore areas of their design spaces to accumulate information, form perceptions of preferred and nonpreferred areas from this information, and propose space reductions from these perceptions (Bernstein, 1998). A representation of an example design space is depicted in Figure 8 of Andrews (2018) where *space reductions* refer to reducing the range of potential design solutions being left open. A Design Integration Manager (DIM) will then consider the space reductions proposed and the information supporting them to finalize a conceptually robust set of space reductions across all disciplines (Singer et al., 2009). Each of these later steps are directly tied to the information gathered at the beginning, so effective decision-making in SBD necessitates robust information. Gembarski et al. (2021) evaluates the robustness of information in decision-making by using Bayesian probabilities to model uncertainties that originate from a scarcity of information. Sypniewski (2019) takes a different approach and assesses how the inherent biases of information that has already been gathered can lead to inadequate characterization of a design space and misinformed decisions. As the robustness of information pertains to decisions made during SBD specifically, research is limited. Doerry (2015) presents a method for measuring the diversity of information in a design space to increase the likelihood of viable solutions being found later; however, this method intends to insure reduction decisions against uncertain information rather than understand the uncertainty permissible for those decisions to remain advisable.

The purpose of this paper is to present a new approach for quantifying the risk of design space reduction decisions in SBD by considering the potential for new information to alter perceptions of feasibility and incite emergent design failures. Although in a much broader sense, the aim of this research is to assess how future information can undo the foundation of knowledge already established through previous design decisions. In the following sections, a brief background on SBD and a design space's fragility (or its vulnerability to new and conflicting information) will first be provided. Next, details of a SBD simulation will be presented for proposing reasonable space reductions. After explaining how the simulation works, an early framework built for assessing the fragility of design spaces and quantifying the risk of space reduction decisions will be explained. The developed fragility framework plugs in at the very end of the space reduction process. Finally, emergent design spaces for simulations performed on a problem analogous to complex design with and without the framework will be observed and discussed.

SET-BASED DESIGN

SBD is a convergent design approach that seeks a final solution through the gradual elimination of design spaces rather than cycles of rework and refinement synonymous with most iterative approaches. Bernstein (1998) describes the *ideal* way SBD should be performed with illustrative help from Figure 1 developed by Dr. William Finch. In the early stages of SBD, disciplines individually explore areas of their design spaces and expand their ranges of potential design solutions. From a marine design perspective, these disciplines may consist of (but not be limited to) a weights division negotiating lightship and deadweight tonnage allotments along with center of gravity positioning, a stability division considering allowable beam and vertical center of gravity pairings, and a structural division contemplating various plate thickness and stiffener sizing schemes. As potential solution spaces are identified by each discipline, they work together to identify areas of overlap between their interdependent design spaces that satisfy all requirements of the design problem. For example, the weights division may have its own displacement and trim requirements to satisfy, but the vertical center of gravity of a load case cannot prevent the stability division from satisfying intact or damage stability requirements, and the lightship allotment must be sufficient for the structural division to satisfy material yielding requirements within a specific safety factor. As a solid understanding of potential design solutions and trade-offs is formed, nonpreferred areas of each discipline's design space are eliminated until a final solution remains.

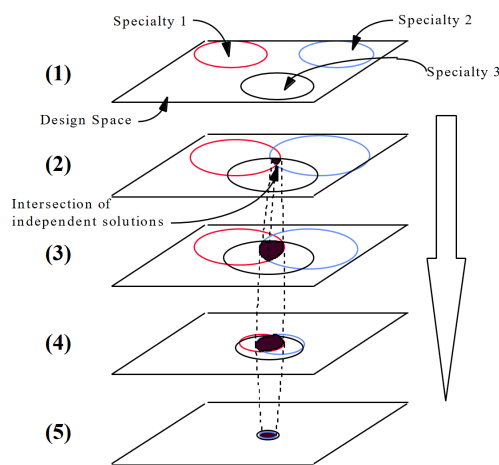


Figure 1: Ideal convergence of the SBD process through gradual elimination of nonpreferred areas (Bernstein, 1998)

Through this process, a major principle of SBD is delaying decisions until the consequences of those decisions are understood (Ward et al., 1995; Singer et al., 2009). During discussions with managers utilizing “set-based concurrent engineering” at Toyota, Ward et al. (1995) learned that a critical aspect of their job is to discourage engineers from making important design decisions too soon. They believe it is necessary to delay decisions to ensure all the requirements of the customer are met while also ensuring that the design is manufacturable. Bernstein (1998) and Singer et al. (2009) discuss the benefits of delaying design decisions from the perspectives of accrued knowledge, committed costs, and stakeholder influence. They explain that knowledge of a design is gathered with time as designers run analyses to build their understanding of the characteristics and requirements driving the process. By delaying decisions through a set-based approach, designers can increase the influence maintained and decrease the costs incurred until the information and existing knowledge supporting these decisions is more robust.

The difficulty of eventually making these reduction decisions is that design spaces cannot be understood absolutely. Different disciplines often manage large design spaces that cannot be explored completely while tolerating analyses with varying degrees of uncertainty. Moreover, it is common for changes in design requirements as well as the fidelity or underlying assumptions of analyses to be introduced throughout the design process that shift preferred and nonpreferred areas. Shields and Singer (2017) assert that space reduction decisions create low-risk knowledge structures while also acknowledging that SBD relies on considerable knowledge generation and decision-making to work effectively. In their words, “Only making decisions when the supporting knowledge is well-understood and is unlikely to change leaves stable knowledge to be

further developed” (Shields and Singer, 2017). Each space reduction decision in SBD is supported by information that is incomplete, uncertain, and susceptible to change. If designers do not account for this uncertainty of information, their reduction decisions may lead to exceedingly fragile design spaces, or design spaces whose perceived feasibility is vulnerable to new and conflicting information. In instances when new information exposes fragile design spaces, designers using a SBD approach cannot simply rely on backtracking and reopening design spaces either, because their design timelines are limited by the considerable time already spent exploring those design spaces in the first place.

Design Space Fragility and Reduction Decisions

To help visualize a design space’s fragility, Figure 2 has been created to mirror the third layer in Bernstein’s explanation of SBD. In Figure 2, the perceived feasible regions of each discipline are located within the red, blue, and black circles. The green regions signify perceived feasible areas of the design space for one discipline, the yellow regions signify the same perceived feasibility for two regions, and the orange region signifies the same perceived feasibility for all three regions. Suppose the fragility is being assessed from the red discipline’s perspective. One source of fragility is attributed to learning new information that alters the perceived feasible space of the red discipline itself, as depicted by the dashed red circle in Figure 3. As it pertains to the red discipline, the pink region captures newly perceived feasible space, and the grey region captures newly perceived infeasible space. If new information shifts the perceived feasible space of the red discipline such that the grey region outweighs the pink region, then that originally perceived design space would have been very vulnerable to the new information. Another source of fragility is attributed to learning new information that alters the perceived feasible space of an interdependent discipline, as depicted by the dashed blue circle in Figure 4. As it pertains to the blue discipline, the pink region captures newly perceived, shared feasible space, and the grey region captures perceived, shared infeasible space. If new information shifts the shared feasible space such that the grey region outweighs the pink region, then that originally perceived design space would have also been very vulnerable to the new information. In both scenarios, new information’s effect on perceived feasibility is determinant of a design space’s fragility.




Figure 2: Overlapping regions of perceived feasible spaces for three disciplines of a design problem




Figure 3: Fragility attributed to design change of main discipline




Figure 4: Fragility attributed to design change of interdependent discipline

While a design space’s fragility directly corresponds to its vulnerability to new information, that vulnerability can be amplified by the particular space reductions that have previously been made. In both Figure 3 and Figure 4, the DIM may have already decided to eliminate portions of the pink region. If that is the case, disciplines would be left without newly perceived feasible space, meaning that the grey region would further outweigh the pink region. Designers want to avoid space reduction decisions that lead to exceedingly fragile design spaces, yet they must make reductions to keep the design process moving. At each space reduction cycle, every design space is susceptible to increases in design space fragility that can be further exacerbated by previous reductions. And even though the figures have portrayed instances of increasing fragility in the context of new information and previous space reductions, the size of a feasible design space alone does not dictate its fragility. For example, if a multidimensional design space is very large, but one of its input variables only has two feasible values, then that design space is vulnerable to new information from its own discipline and/or other disciplines finding

those two values to be infeasible. By effect, there are varying levels of risk for space reduction decisions due to the varying levels of fragility that result from prior reductions and new information.

Originating Sources of New Information

In the development of solutions to complex design problems, designers are compelled to explore and gain an understanding of their own discipline’s design space, integrate the understanding and preferences of designers from interdependent disciplines with their own, and endure changing design requirements and maturing analyses throughout the entire process. Bearing each of these challenges in mind, three different sources of new information are worth considering when characterizing the fragility of a design space: (1) newly explored design points of a directly affected discipline, (2) newly explored design points of an interdependent discipline, and (3) new or updated design requirements or analyses. In this work, only the first originator for new information will be considered, but it is important to keep the other two in mind for future improvements to the fragility assessment process.

To observe how new information originating from newly explored design points of a directly affected discipline can impact perceptions of design space behavior, consider Figures 5 to 6. In these figures, green points represent tested designs that are feasible, red points represent tested designs that are infeasible, green regions represent perceived feasible spaces, red regions represent perceived infeasible spaces, and yellow regions represent spaces of mixed feasibility. With the information from design points presently available in Figure 5, clear regions of feasibility have been formed for the discipline; designers of *this discipline* are perceiving smaller values of *Variable 1* to be feasible and larger values of *Variable 1* to be infeasible. However, those perceptions shift in Figure 6 when new information originating from newly tested design points becomes available. Larger values of *Variable 1* are still perceived as infeasible, but designers have also learned they may have less area to work with for smaller values of *Variable 1* than they previously thought. Before learning this new information, suppose the decision is made to eliminate some of the smallest values of *Variable 1* because, in contrast to this discipline, other disciplines prefer large values of *Variable 1* to small values. Designers of this discipline may be inclined to approve the space reduction thinking they still have plenty of feasible space with which to work. Later, they would regret to learn that the space reduction decision has limited far more feasible solutions remaining for them than they originally anticipated.




Figure 5: Perceptions of feasibility before sampling new points within primary design space

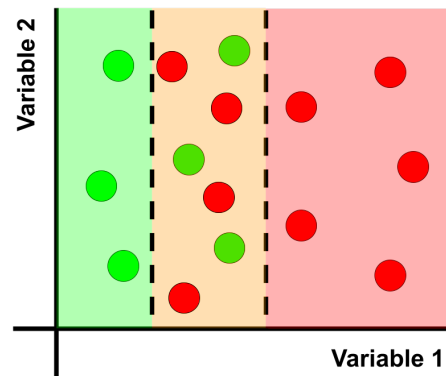


Figure 6: Perceptions of feasibility after sampling new points within primary design space

The intent of a fragility framework will be to protect design spaces against scenarios like the one described. DIMs may be capable of taking proposed space reductions from disciplines and carefully assessing the impact those reductions would have on other disciplines with the information *at hand*, but they lack a tool for understanding the consequences of those reductions if the perceptions formed from that information *changes*. Before introducing an early strategy for quantifying this space reduction risk, a SBD simulation to which a fragility framework can tie in must be developed.

CREATING A SBD SIMULATION

Before discussing the logic behind a framework intended to evaluate the fragility of design spaces, a SBD simulation that proposes reasonable space reductions is needed. With the simulation, experiments can be run that compare emergent design spaces when there are fragility checks in place compared to when there are not. The decision was made to automate the simulation for a couple of reasons. For one, automating the simulation removes the impact that human inconsistencies would have on the emergent design spaces by ensuring the same criteria are used to explore design spaces and propose space reductions every time. Additionally, automating the simulation speeds up the process and cuts back on the time it would take for a DIM to evaluate the present state of the data and formulate their next exploration or reduction decision. The simulation is not meant to be a perfect replication of how SBD activities are performed and reductions are made because SBD is fundamentally a human-centric process that is driven by knowledgeable designers. A simplified depiction of how the SBD simulation works is shown in Figure 7, while a more detailed depiction of the simulation and the actual Python code can be viewed via the link in the Data Access Statement at the end of this paper. Different parts of the simulation fall under the groupings of problem setup, exploration, or space reduction.




Figure 7: Simplified flowchart of the automated SBD simulation with a fragility framework plug-in

Problem Setup

The problem setup portion initializes the design problem and prepares various aspects of the code for the time-based simulation. The user has the option to set certain user inputs that dictate characteristics of the design problem, exploration tendencies of the algorithm, and space reduction tendencies of the algorithm. The code is intended to work for any design problem as long as the user defines the input variables, output variables, initial input rules, initial output rules, and analyses of each discipline. To ensure exploration does not continue indefinitely, a bounded timeline is established, and the user must define how long the analyses of each discipline take to run during exploration. This portion of the code also sets an exponential-based function that encourages a minimum amount of each discipline’s design space to be reduced relative to the time elapsed. An exponential function was chosen to promote exploration without required space reductions early on while saving more of the required space reductions for later once more information has been gathered. In an actual SBD process, there usually is not any sort of timeline established to keep space reductions on pace. However, for an automated SBD process without human involvement, it is important to ensure all of the space reductions are not delayed until the end.

Exploration

The exploration portion is where each discipline is free to sample designs within their remaining design space. Exploration occurs whenever there is time remaining, there are no space reductions being proposed by any discipline, and there are no space reductions being forced because of the time that has elapsed compared to the size of each discipline’s remaining design space. Based on guidelines set by the user in advance of the simulation, a specific amount of time will be dedicated to

sampling new points at each exploration cycle. The location of new points chosen for sampling are randomized throughout each discipline's remaining design space. Each discipline will then calculate the output values of the new points according to their specific analyses and check whether the calculated output values satisfy the current set of output requirements. Design time is only reduced during exploration as it is assumed that the time to run analyses is significantly greater than the time to propose space reductions or evaluate fragility.

For infeasible points, the algorithm is programmed to determine the extent to which each point fails. This failure calculation finds the difference between the calculated amount and the nearest threshold amount of each failing requirement (normalizing the difference to ensure no one requirement bears more influence on the failure amount than the others), and then calculates the root-mean-square deviation of the normalized differences. After going through this process, each infeasible point will have a failure amount assigned to it where values closer to zero indicate that the point is right on the threshold of passing, while larger values indicate that the point is very far off from passing. These calculated failure amounts come in handy when the SBD simulation goes through its space reduction and fragility assessment processes. Equation 1 shows an example calculation of the failure amount (FA) involving three different requirements ($y_1 > 0.2$, $0.3 < y_2 < 0.6$, $y_1 + y_2 < 0.8$), three different calculated amounts ($y_1 = 0.1$, $y_2 = 0.25$, $y_1 + y_2 = 0.35$), and three different ranges of output values ($y_1 \in [0.05, 1.2]$, $y_2 \in [0, 0.9]$, $y_1 + y_2 \in [0.05, 2.1]$).

$$FA = \sqrt{\frac{\left(\frac{0.1-0.2}{1.2-0.05}\right)^2 + \left(\frac{0.25-0.3}{0.9-0}\right)^2 + \left(\frac{0}{2.1-0.05}\right)^2}{3}} = 0.103 \quad (1)$$

For feasible points, the algorithm is programmed to determine the extent to which each point passes. This passing calculation is done in a similar process to that of the failure calculation with a couple of slight modifications. First off, the difference is now taken between the calculated output amount and the nearest threshold amount of each *passing* requirement. Secondly, rather than taking the root mean square of each requirement's normalized difference between the calculated value and threshold, the passing calculation simply takes the minimum normalized difference to gauge how close an output value is to failing the nearest requirement. By effect, passing amount values closer to zero indicate that the point is right on the threshold of failing, while larger values indicate that the point is very far off from failing. These calculated passing amounts also come in handy when the SBD simulation goes through its space reduction and fragility assessment processes. Equation 2 shows an example calculation of the passing amount (PA) involving the same three requirements ($y_1 > 0.2$, $0.3 < y_2 < 0.6$, $y_1 + y_2 < 0.8$), three new calculated amounts ($y_1 = 0.4$, $y_2 = 0.35$, $y_1 + y_2 = 0.75$), and the same ranges of output values ($y_1 \in [0.05, 1.2]$, $y_2 \in [0, 0.9]$, $y_1 + y_2 \in [0.05, 2.1]$).

$$PA = \min\left(\frac{|0.4 - 0.2|}{1.2 - 0.05}, \frac{|0.35 - 0.3|}{0.9 - 0}, \frac{|0.75 - 0.8|}{2.1 - 0.05}\right) = 0.024 \quad (2)$$

Space Reduction

The space reduction portion is where each discipline can propose new input rule(s) that reduce design spaces if adopted by the DIM. The process each discipline goes through to propose these rule(s) involves a series of steps. It starts out by sorting all the failure amounts of previously explored points remaining in a discipline's available design space by magnitude. The algorithm then labels a percentage of the points with the highest failure amounts as "bad" while all of the remaining points are labeled as "good". Next, a decision tree classifier (DTC) is built in the design space based on these "good" and "bad" points. The goal of the DTC is to create boundaries within the design space such that all the "good" and "bad" points are grouped together as succinctly as possible. After the decision tree is formed, the boundaries of the decision tree grouping the highest fraction of "bad" points together are extracted to act as the threshold of a newly proposed input rule. A benefit of the decision tree is that each one of its boundaries are defined by a first-order equation only involving one input variable. This benefit reflects what is often seen in actual SBD as human designers tend to define simple rules for elimination rather than complicated equations cutting through the design space. The new input rule is finally checked against a set of user-

defined criteria to ensure it is supported by adequate information before a discipline can formally propose it.

If no rules are proposed by any discipline, then the simulation will ask whether a reduction should be forced for any discipline. If the answer to that question is yes (due to substantial space remaining relative to time remaining), then one of the user-defined criterion will be relaxed, and the disciplines will reassess if they have any reductions to propose. If the answer to that question is no, then the algorithm will continue to the exploration part of the simulation.

If one or more disciplines do have a new rule to propose, then the DIM becomes responsible for thoughtfully merging these requests based on available information and the impact each rule would have on other disciplines. This is a difficult part of the SBD process to reproduce because a preferred reduction for one discipline may not be preferred by another; human DIMs must often think critically when finalizing requests based on infeasibility and dominance. In the simulation, this merging process is handled by having each discipline directly affected by the input rule form an opinion on it that is represented by a value between 0 and 1 (where a value of 0 indicates the discipline is completely opposed to the rule, while a value of 1 indicates the discipline is completely in favor of the rule.) The opinions are quantified through each discipline's answers to two different questions:

1. Is the proposed space reduction removing clearly infeasible designs in your discipline's design space?
2. If the proposed space reduction is enacted, how less likely is it that feasible designs exist in your discipline's remaining design space?

To answer the first question, the area of the remaining design space that the space reduction would remove (the *eliminated design space*) is assessed. This area of the design space is discretized, and then a Gaussian process regressor (GPR) is trained with the available data from all of the explored points. The x -training data of the GPR are the input locations of the explored points, while the y -training data are the difference between each point's passing and failing amount. The trained GPR is then used to predict the difference between the passing and failing amounts at each discretized (unexplored) point in the eliminated design space. A negative value indicates the discretized point is predicted to be infeasible, while a positive value indicates the discretized point is predicted to be feasible. As the question is concerned with discerning *clearly infeasible* areas of the eliminated design space, the predicted difference between the passing and failing amounts is permitted to fluctuate between three standard deviations. The fraction of this range staying *below* zero for each discretized point is calculated, and the average value of those fractions acts as the metric answering the first question. Large values of this metric reflect *assuredness* that the eliminated design space is clearly infeasible.

The answer to this first question may be sufficient for validating proposed spaces reductions early in the SBD process when designers are purposefully delaying commitments to hard-set specifications and working with large infeasible spaces. If the answer to this question is a resounding yes for all disciplines involved, then the DIM can move forward with the space reduction for the universally infeasible design space. However, later on in the SBD process when these infeasible spaces have diminished, dominance-based decisions may be required that consequentially cut away at some of the feasible spaces of various disciplines. In these cases, the second question becomes more important to ask to ensure that a dominance-based reduction decision does not end up severely limiting any one discipline from producing a feasible design.

To answer the second question, the area of the design space that would remain after the space reduction (the *reduced design space*) is assessed in relation to the original design space before the space reduction (the *non-reduced design space*). The same trained GPR is now used to predict the difference between the passing and failing amounts at each discretized (unexplored) point in the reduced and non-reduced design spaces. However, now the question is concerned with discerning *likely feasible* spaces to ensure the space reduction does not significantly reduce the potential of finding a feasible design. In this case, the normal distribution about each discretized point's predicted value is determined, and the average fraction of those distributions falling *above* zero is calculated for both the reduced and non-reduced design spaces. The ratio of the reduced average to the non-reduced average acts as the metric answering the second question, where large values reflect *little reduction* in the *likeliness* of finding feasible designs in the reduced design space.

With metrics produced that quantify a discipline's answers to both questions, the influence that each metric should have on a discipline's overall opinion of a space reduction can be determined. Equation 3 is used to quantify this opinion (OP)

where m_1 and m_2 are the metrics for the first and second question, and w_1 and w_2 are the weights assigned to each metric. Because the second question only comes into play if the answer to the first question is not a resounding yes, the weight of the second metric is dictated by the value of the first metric. If m_1 is high because clearly infeasible spaces would be eliminated, then w_2 should be low because those spaces are not needed regardless. On the other hand, if m_1 is low because clearly infeasible spaces are *not* being eliminated, then w_2 should be high to account for how much the space reduction would actually hinder the remaining space in the opinion formulation. To reflect this behavior while not limiting the relationship between w_2 and m_1 to a linearly inverse correlation, a user-specified quadratic Bezier curve between the two variables is adopted. Once w_2 is determined, w_1 is calculated through the equation $w_1 = 1 - w_2$ to ensure the overall value on the opinion stays between 0 and 1.

$$OP = w_1 * m_1 + w_2 * m_2 \quad (3)$$

Once the opinions are formed, the DIM can finally decide how much influence each opinion should have when finalizing the universal set of input rule(s) for the newest space reduction. In the simulation, this decision is made by establishing a threshold which permits a discipline to veto an input rule based on the value of their opinion in relation to the opinion of the discipline proposing the new input rule(s). Early on in the SBD process when space reductions do not necessarily need to be forced by the DIM, this threshold is low to allow for more vetoing of rules and less dominance-based reduction decisions. Later on in the SBD process when space reductions are becoming more forced by the DIM, this threshold is high to prevent more vetoing of rules and allow for more dominance-based reduction decisions.

At the end of this space reduction process, the DIM will have decided on a universal set of space reductions by which all disciplines must abide. Again, it is not meant to be a perfect representation of how space reduction decisions are made in SBD. Rather, it is meant to consistently produce reasonable space reductions for both infeasibility and dominance-based decisions so that the fragility of those decisions can be studied.

FRAGILITY FRAMEWORK

Traditionally in SBD, the space reduction decision process ends with the universal set of reductions instituted by the DIM. At this point, designers have explored their own design spaces to form perceptions and propose space reductions, and the DIM has merged them together with the information available through infeasibility or dominance. As discussed though, this process, which only considers present information, leaves reduced design spaces vulnerable to new information.

The intent of a fragility framework is to gauge the vulnerabilities of each discipline's design space to new information before committing to any space reductions. To accomplish this goal, a developed framework will require components that address various complexities inherent to the space reduction process. Table 1 summarizes those space reduction complexities and corresponding fragility framework requirements. In this work, a Probabilistic Fragility Model (PFM) is introduced for fragility assessment. The PFM is still a work in progress and does not address every framework requirement outlined in the table. Still, it addresses many complexities inherent to SBD's space reduction process and has the potential to be expanded further in future work. After discussing the underlying logic behind the PFM, it is incorporated as the final step in the SBD simulation's space reduction process.

Probabilistic Fragility Model

The main idea behind the PFM is to characterize a discipline's present understanding of a design space with straightforward probabilities of feasibility and infeasibility and then to quantify its vulnerability based on how likely that understanding is to change. There are three main parts to the PFM which include forming perceptions of feasibility in the design spaces, determining potentials for regret and windfall from those perceptions, and using metrics to compare those potentials between the reduced and non-reduced design spaces.

In the first step, designers need to form perceptions of feasibility throughout their design space by leveraging data from their explored points thus far. To meet this requirement, the same GPR from the space reduction process is used to predict the difference between the passing and failing amounts (pass-fail) for all unexplored areas of the design space. With these predictions, designers can form perceptions for feasible and infeasible areas depending on if the pass-fail value is positive or negative. Designers will also have an idea of how much different areas are passing or failing depending on its magnitude. Figure 8 depicts this process for the remaining areas of a design space involving two input variables (x_1 and x_2). On the left-hand side of the figure, pass-fail amounts are formed for three explored points. Data from those explored points train a GPR, and then the trained GPR forms predictions for the discretized areas of the design space.

Table 1: Complexities that exist when making space reduction decisions with uncertain information and the corresponding fragility framework requirements addressing these complexities

Space Reduction Complexity	Framework Requirement
Space reductions are focused on eliminating undesirable solutions from a ranging design space. The desirability of solutions are rooted in perceptions of feasibility formed by running <i>discrete</i> design points through the analyses established by each discipline.	The framework needs to form initial perceptions of feasibility with presently available information. A technique for converting information from explored points and their output values into perceptions of feasibility <i>throughout</i> each discipline’s design space is required.
Perceptions of feasibility are uncertain because they are formed with incomplete information within a discipline’s design space. Information from newly analyzed design points <i>within a design space</i> could alter perceptions.	Formed perceptions of feasibility for unexplored areas of the design space are not definitive. The framework should account for the possibility of new design points being tested with feasibility that is contradictory to expectations.
Perceptions of feasibility are uncertain because of the <i>interdependencies</i> that exist through shared variables between disciplines. Vulnerabilities of one design space to new information could directly or indirectly amplify the vulnerabilities of other design spaces.	The framework must include a cross-discipline component that ties the individualistic fragilities of each discipline together such that the vulnerabilities tracked across interconnected design spaces are representative of their dependencies on each other.
Perceptions of feasibility are uncertain because they are formed with output information that is susceptible to change. New information originating from <i>changes to design requirements or analyses</i> could alter perceptions.	The location of calculated output values within the objective space must not be treated as definitive. Instead, the framework should account for the possibility of output values and requirements shifting in relation to each other.
A design space may be fragile when considering all input variables together (i.e. x_1, x_2, x_3) and when considering various <i>combinations</i> of input variables (i.e. x_1, x_2).	The framework cannot only measure the fragility of a design space as a whole. It must be flexible enough to also identify component-based fragilities.
The number of ways new information can alter perceptions of feasibility within a design space is <i>unbounded and unknown</i> until the information is made available. The risk of a space reduction in context of itself is unlimited.	Comparing the fragility of a reduced design space to a non-reduced design space and determining what new information a discipline <i>can</i> handle rather than it would <i>have to</i> handle will narrow the DIM’s scope and allow space reduction risk to be quantified.




Figure 8: Forming perceptions of feasibility for unexplored areas of the design space

In the next step, designers need to consider the consequences of their formed perceptions of feasibility being incorrect. This requirement leads to the introduction of regret and windfall in a design space. Suppose the sampled design space in Figure 9 is considering the space reduction depicted by the black box. The space reduction would eliminate portions of the design space perceived as feasible (top-left) as well as portions of the design space perceived as infeasible (top-right). Now suppose new information comes along that throws off those perceptions of feasibility as depicted by the left-hand design space in Figure 10. This new information would cause designers to regret the space reduction if they are left *with infeasible* space that was expected to be feasible or left *without feasible* space that was expected to be infeasible (instances of regret). In contrast, the new information would benefit designers if they are left *with feasible* space that was expected to be infeasible or left *without infeasible* space that was expected to be feasible (instances of windfall). Before committing to a space reduction, the PFM considers these *potentials* for windfall and regret for the reduced design space in context of the non-reduced design space (right-hand side of Figure 10). This logic allows designers to consider the consequences of moving forward with a space reduction compared to forgoing the space reduction.

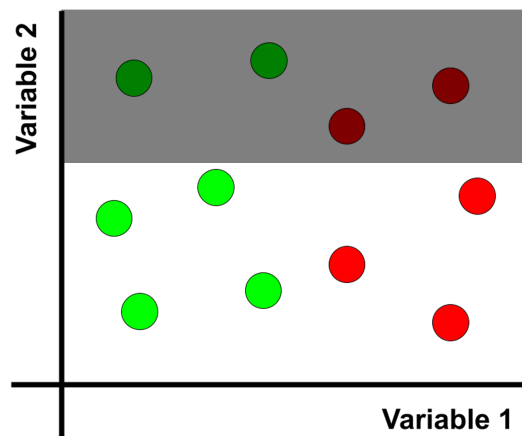


Figure 9: Design space considering a proposed space reduction (signified by the grey box)

Considering the consequences of a space reduction alone is not enough as there are different likelihoods of these consequences actually coming to fruition. Fortunately, probabilities of feasibility can be calculated from the predicted pass-fail values of the GPR to account for the regret and windfall likelihoods. Figure 11 depicts this process for the PFM. Once all




Figure 10: Instances of regret and windfall for the reduced design space (left) and non-reduced design space (right)

discretized areas of the design spaces are labeled as feasible or infeasible based on their predicted pass-fail value, the value itself and the accompanying standard deviation are used to create a normalized probability distribution on each prediction. The probability of feasibility or infeasibility is determined based on the portion of a point's probability distribution maintaining the predicted positive or negative pass-fail value. Finally, the potentials for regret or windfall are taken from the complementary probabilities of feasibility for each discretized point. Whether the complementary probability results in a potential for regret or windfall will depend on where the discretized point falls in the non-reduced or reduced design space.

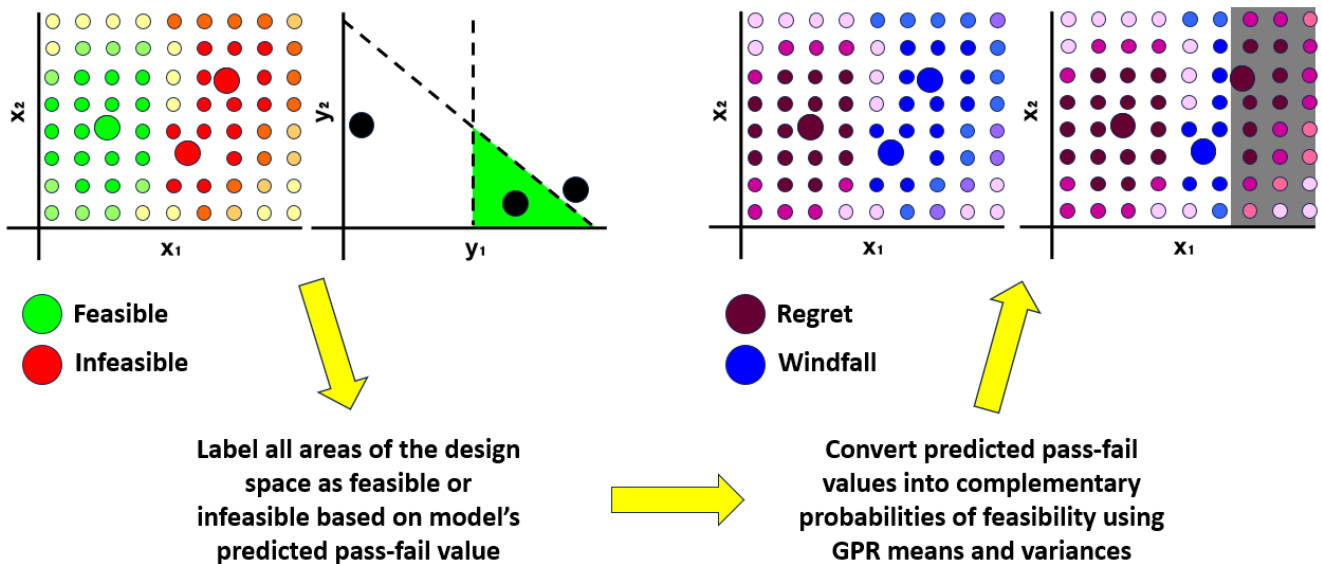


Figure 11: Converting perceptions of feasibility into potentials for regret and windfall for a non-reduced and reduced design space

The last step in the PFM involves gauging regret and windfall potential for the entirety of the reduced design space from each discretized point. A straightforward method to do so is summing up each discretized point's regret and windfall potentials. However, to actually understand the *risk* of a space reduction from these potentials, further context is needed. This context is provided by calculating the same summation for the non-reduced design space as is done for the reduced design space and then calculating the *added* potentials for regret and windfall (shown in equations 4 and 5). Working with these added potentials for regret and windfall allow a DIM to understand the risk of moving forward with a space reduction in context of leaving the design space untouched. High risk space reductions correspond to large positive values of added potential for regret and large negative values of added potential for windfall, which would reflect a reduced design space taking on more potential for regret and giving up potential for windfall. In equations 4 and 5, Δ_{reg} and Δ_{wind} are the added potentials for regret and windfall, $p_{reg,red}$ and $p_{wind,red}$ are a discretized point's probability of regret and windfall in the reduced design space, and $p_{reg,nonred}$ and $p_{wind,nonred}$ are the same probabilities for the non-reduced design space.

$$\Delta_{reg} = \frac{\sum_{i=1}^n p_{reg,red}(x)}{\sum_{i=1}^n p_{reg,nonred}(x)} - 1 \quad (4)$$

$$\Delta_{wind} = \frac{\sum_{i=1}^n p_{wind,red}(x)}{\sum_{i=1}^n p_{wind,nonred}(x)} - 1 \quad (5)$$

Incorporating Fragility in SBD Simulation

In the SBD simulation, fragility checks occur immediately after the DIM has merged proposed space reductions into a universal set, as depicted by the red box in Figure 7. Up to this point in the simulation, all space-reduction-related decisions are solely supported by present information that is assumed not to change. The fragility framework intends to protect design spaces against this assumption and consider the effect that the space reductions would have on the remaining design spaces if new information were to alter the perceptions formed.

After calculating the added potential for regret and added potential for windfall, the amount of risk a DIM is willing to take on must be determined. In this work, a maximum risk threshold that increases exponentially with project time is established. The idea behind choosing an exponential threshold is to undertake little risk in space reduction decisions early on when conflicting information can be more common and influential on design space perceptions, and to increase the amount of risk endured later when conflicting information is less common and design spaces are more locked in.

With the exponential risk threshold set, the actual risk experienced from a space reduction in the simulation is set as the added potential for regret subtracted by the added potential for windfall. This risk setup allows the DIM to weigh each discipline's shift in potential to be hurt and helped by new information following a space reduction. In any given space reduction cycle, disciplines are permitted to keep proposing space reductions that can be added by the DIM as long as the risk does not exceed the threshold for any one discipline. If the threshold is exceeded, the input rule that put the risk over the limit is temporarily banned from being proposed again until design spaces have been explored further.

With that, the SBD simulation is established with the option to include a probabilistic fragility check before fully committing to each space reduction decision. After first validating the pre-fragility portion of the simulation with a simple SBD problem, emergent design spaces of the design problem with and without the PFM are ready to be observed.

VALIDATING SIMULATION WITH A SBD PROBLEM

Investigating the impact of this new fragility framework on emergent design spaces requires a simple design problem. Testing the framework on a simple problem directs focus on the framework and makes interpretation of results in its early stages

more straightforward. After first describing the problem, it will be used to validate the efficacy of the SBD simulation without yet incorporating the PFM. Once the simulation is validated, it will then be used to assess the vulnerabilities of emergent design spaces for SBD simulations with and without the PFM.

SBD Problem

A design problem has been created for testing the developed fragility framework. As shown in Figure 12, the design problem involves three different disciplines having some shared input variables and unique output variables. The input variables are analogous to the different ship characteristics that a discipline has influence over, while the output variables are analogous to the different ship performance characteristics with which a discipline is concerned. Circling back to the marine design disciplines used as an example in the *SET-BASED DESIGN* section, Discipline 1 could represent the Weights division, Discipline 2 could represent the Stability division, and Discipline 3 could represent the Structural division. It is important to note that only a simple design problem is being investigated right now, and there are numerous other disciplines and sub-disciplines (e.g. arrangements, powering, seakeeping, maneuvering) that would tie into more complex problems.




Figure 12: Input and output variables for three disciplines of SBD problem

Each one of these disciplines have analyses that calculate the output variables from the input variables to provide insight on how a potential design solution will perform. For this design problem, arbitrary mathematical equations act as the analyses for each discipline as shown in equations 6 to 10. These equations are analogous to the different parametric models or design programs used to evaluate performance metrics of a potential design solution.

Discipline 1:

$$y_1 = 0.8x_1^2 + 2x_2^2 - x_3 \quad (6)$$

Discipline 2:

$$y_2 = 1.25x_5 - 12.5x_3^3 + 6.25x_3^2 \quad (7)$$

$$y_3 = (x_4^3 + x_5)^2 \quad (8)$$

Discipline 3:

$$y_4 = 2x_5 + 0.2 \sin(25x_6) - x_1^{\frac{1}{5}} \quad (9)$$

$$y_5 = x_1^{\frac{1}{3}} - \cos(3x_5) \quad (10)$$

Each of the input and output variables have requirements that must be satisfied. The bounds on all the input variables are normalized such that they must fall between 0 and 1. The bounds on the output variables are unique and described as follows: $0 \leq y_1 \leq 0.4$ or $1.2 \leq y_1 \leq 1.6$, $0.5 \leq y_2 \leq 0.7$, $0.2 \leq y_3 \leq 0.5$, $0 \leq y_4 \leq 0.5$, $0.8 \leq y_5 \leq 1.6$. These

requirements are analogous to the different stakeholder or industry-set design requirements that the design must satisfy. In marine design, the bounds on the input variables could be normalized length, beam, depth, etc. ranges, while the bounds on the output variables could be standardized displacement, wind-righting arm, yielding stress, etc. ranges.

The equations and required bounds of the design problem produce the feasible spaces depicted in Figures 13 to 15. In each of the figures, red points represent discrete design solutions not meeting output requirements, while green points represent discrete design solutions meeting output requirements. The feasible boundaries are unknown to the designers of each discipline, so they must form their perceptions of these feasible boundaries solely from the discrete points they decide to test. The equations and bounds are formulated in such a way that the feasible spaces have complex boundaries for designers of the SBD simulation to learn.

Validating SBD Simulation Approach

The code simulating SBD decisions needs to be validated to ensure reasonable space reductions are being proposed from the present information. The aim of the fragility framework is to evaluate the suitability of a space reduction decision based on the potential for perceptions of feasibility to be altered by *new information* rather than present information. Even so, if unreasonable space reductions that neglect infeasible areas and surround feasible areas are proposed, then new information could not make perceptions any worse, and the framework would never find a reduced design space to be fragile.




Figure 13: Feasible bounds of Discipline 1 depicted with 2,000 sampled points




Figure 14: Feasible bounds of Discipline 2 depicted with 2,000 sampled points




Figure 15: Feasible bounds of Discipline 3 depicted with 2,000 sampled points

To validate the SBD code, 200 runs of the SBD problem without the fragility framework were executed. Table A1 shows all of the user inputs selected for these simulations. Short analysis run times ([2, 3, 4] iterations) relative to a long project timeline (1000 iterations) were chosen to ensure ample time is given to build up information for proposing reasonable space reductions. Each discipline was also given the goal of reducing their designs space's down to at least 5% of their initial size.

Figures 16 to 18 show locations of remaining solutions in each discipline's design space at the end of all 200 runs with the tan-colored points. The more opaque points show where remaining designs are most commonly found in each discipline at the end of a run, while the more transparent points show where remaining designs are less commonly found or not found at all. The surfaces show the feasible bounds of the design problem unknown to designers in the simulation. As can be seen in each figure, the remaining solutions are most commonly found within or around each discipline's feasible areas. This finding shows that remaining designs are being narrowed down to these feasible areas without actual knowledge of them, indicating that the SBD code is proposing reasonable space reductions.




Figure 16: Locations of remaining design solutions for Discipline 1 at the end of the simulation




Figure 17: Locations of remaining design solutions for Discipline 2 at the end of the simulation




Figure 18: Locations of remaining design solutions for Discipline 3 at the end of the simulation

ASSESSING FRAGILITY IN SBD PROBLEM

With the SBD code validated, experiments can be run comparing the emergent design spaces of SBD simulations including the PFM to those that do not. As the proposed method does not yet consider the potential of analyses or requirements to change, random design changes do not need to be introduced to the simulation. The PFM currently only considers the impact that new information arising from newly tested design points in a primary discipline would have on a remaining design space. The results focus on evaluating the fragility of design spaces following space reductions by tracking remaining designs as the simulation progresses.

Experimental Setup

Experiments are run for two different scenarios. In the first scenario, emergent design spaces with and without the PFM are compared to each other when a large amount of time has been allotted to the design problem relative to analysis run times. These test cases will have more time to generate information and presumably form more stable perceptions of design space behavior before proposing various space reductions. In the second scenario, emergent design spaces with and without the PFM are compared to each other when little time has been allotted to the design problem relative to analysis run times. These test cases will have less time to generate information and presumably form less stable perceptions of design space behavior before proposing various space reductions.

Table 2 highlights the differences made in the simulation between these two scenarios, while all other design parameters selected for the simulations match up with those shown in Table A1. The first and third test cases do not include any sort of fragility checks, but the second and fourth test cases include the PFM. Each test case is executed over 200 runs, and the averages of those runs are observed when examining the emergent design spaces.

Table 2: Independent variables for various test cases of the SBD simulation

Design Parameter	User Input	Test Case 1	Test Case 2	Test Case 3	Test Case 4
Project timeline (iterations)	iters_max	200	200	1000	1000
Fragility check	fragility	False	True	False	True
Starting Fragility Threshold	fragility_shift	0.0	0.2	0.0	0.2

Results and Discussion

While executing the runs of each test case, the total design space remaining and the remaining number of feasible solutions are tracked across each discipline. Figures 19 to 21 display these results as various percentages. In each figure, the “Total Space” curves show the average percentage of the remaining design space, the “Feasible Space” curves show the average percentage of the remaining feasible designs, and the “Feasible-to-Remaining” curves show the ratio of the remaining feasible designs to the remaining design space, all over the elapsed project time.

One immediate takeaway from each of the figures is that including fragility checks with the PFM does not result in a higher ratio of feasible designs to space remaining by the end of the simulations. When comparing test cases of the same project timeline in each figure, the ratio of feasible-to-remaining designs is generally the same or slightly higher for test cases that do not include fragility checks. For Test Case 1 and Test Case 2 in Discipline 3, the ratio of feasible-to-remaining designs is much higher (roughly 25%) for simulations without the PFM. While these results may seem deterring, one explanation for them is there being more total space remaining at the end of the simulations with the fragility checks. In Discipline 3 specifically, there is on average about 10% more designs remaining in Test Case 2 at the end of the simulations than every other test case. More test cases remaining can result in a lower ratio of feasible-to-remaining designs. This occurrence is confirmed by the fact that despite its much lower ratio, Test Case 2 has more remaining feasible designs than Test Case 1 at

the end of the simulation for Discipline 3. Regardless, the behavior this work is more concerned with studying occurs for the emergent design spaces rather than the final design spaces.

Another takeaway is that having more time to explore each discipline's design space does lead to a higher understanding of the space. Across each discipline, Test Cases 1 and 2 (having an 80% shorter project time) consistently retain fewer feasible designs than Test Cases 3 and 4. As Test Cases 3 and 4 have more time to explore areas of their design space before proposing space reductions, they can be more careful about eliminating feasible solutions. It is worth noting that the feasible-to-remaining design space ratio of Test Case 1 does rapidly catch back up to Test Cases 3 and 4 towards the end of the simulations. While it would require further investigation to confirm this explanation, the smaller dispersion of designs in Figure 18 than Figure 16 and Figure 17 hints that this behavior may be attributed to the actual equations and requirements established for Discipline 3 in the SBD problem. The feasible solutions for Discipline 3 can vary over the entire range of its unique variable (x_6) but has distinct feasible regions for its shared variables (x_1 and x_5). This coincidence gives designers of Discipline 3 some more freedom to still find feasible designs in their design space whether or not any sort of premature design lock-in has occurred for its shared variables.

Circling back to the fragility aspect of the simulations, the total space remaining results show that simulations including the PFM support a more gradual reduction of design spaces with less lock-in than simulations neglecting fragility checks. For most of Discipline 1's reduction time and for all of Discipline 2 and 3's reduction time, simulations including fragility checks maintain a larger remaining design space than simulations without them. The PFM is effectively delaying the space reduction process and forcing disciplines to really consider the potential consequences of a space reduction before they commit to it. Those consequences are being realized at the "notches" in Test Case 1 and 3's total space curves just past the 40% elapsed timeline, primarily in Disciplines 2 and 3.

The rapid decline in total space remaining and apparent change in pace at each notch suggests the disciplines are eliminating infeasible areas of their design spaces without much hesitation and then finding themselves locked in on the design solutions that remain. While design changes are not introduced in the experiments conducted for this work, the design spaces of Test Cases 1 and 3 find themselves very vulnerable at this point to new information stemming from slight changes to requirements or analyses. Whereas the design spaces of Test Cases 2 and 4, who do consider the consequences stemming from new and conflicting information, are more prepared to handle design changes. Test Case 4 in Discipline 2 does see a similar notch in its total space remaining curve at the same elapsed project time. However, it tries to correct for this rapid space reduction to a greater extent than Test Cases 1 and 3, almost to the point of meeting back up with the total space remaining curve of Test Case 2 at 70% of its elapsed timeline. For what Test Case 2 sacrifices in extra space retained at the end of SBD process, it makes up for in flexibility to handle new information.

As a whole, the results are encouraging and support introducing a step for fragility assessment to support DIMs making space reduction decisions. Using the PFM to make these fragility checks is a promising first attempt, but it is by no means perfect. The PFM meets many of the framework requirements laid out in Table 1, but it is still lacking in a few areas. Namely, the PFM has no such network component that considers vulnerabilities of interdependent disciplines, it does not account for the possibility of calculated output values or design requirements shifting, and it is not yet built to identify any component-based fragilities. Furthermore, the PFM has only been tested for space reduction decisions of a simple design problem involving just three interconnected disciplines. To really justify the PFM's incorporation into the space reduction process, it needs to be proven against more comprehensive design problems while tracking metrics such as the diversity of remaining design spaces to better substantiate claims of rapid convergence and lock-in. Future work will focus on addressing each of these shortcomings.




Figure 19: Size of Discipline 1's total design space, feasible design space, and feasible design space relative to remaining design space over the elapsed project time




Figure 20: Size of Discipline 2's total design space, feasible design space, and feasible design space relative to remaining design space over the elapsed project time

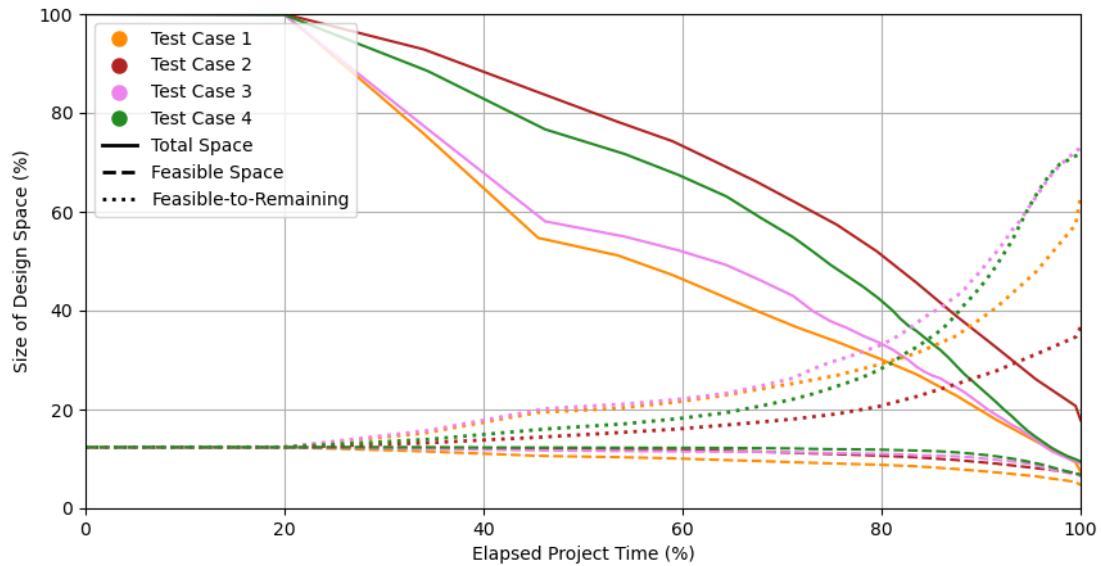


Figure 21: Size of Discipline 3’s total design space, feasible design space, and feasible design space relative to remaining design space over the elapsed project time

CONCLUSIONS

The intent of this work is to introduce a framework to help DIMs make more informed space reduction decisions in SBD by considering the vulnerabilities of remaining design spaces to new information. The framework, or Probabilistic Fragility Model (PFM), uses present perceptions of feasibility formed from sampled points in a design space to gauge the potential and likelihood for those perceptions to be altered by new information. An automated SBD simulation is built to observe the emergent design spaces of a space reduction process including the PFM against one that does not for a simple design problem involving three interdependent disciplines. When tracking their emergent design spaces, initial results indicate that the framework could be a useful tool for delaying space reduction decisions and preventing designers from fixating on certain design solutions while new knowledge is still integrating with existing knowledge. Such a framework could become a critical final step to ensuring space reduction decisions are made with both present and future information in mind.

DATA ACCESS STATEMENT

All code for the SBD simulation and fragility framework along with data for each of the test cases can be publicly accessed at https://github.com/Marine-Structures-Design-Lab/DesignSpace_Fragility/releases/tag/IMDC_2024.

CONTRIBUTION STATEMENT

J.B. Van Houten: Conceptualization, Methodology, Software, Investigation, Data curation, Writing – original draft, Visualization, Funding acquisition. **A.A. Kana:** Conceptualization, Resources, Writing - review and editing, Supervision. **D.J. Singer:** Conceptualization, Resources, Writing - review and editing. **M.D. Collette:** Conceptualization, Methodology, Resources, Writing - review and editing, Supervision.

ACKNOWLEDGEMENTS

This work is primarily supported through the Department of Defense (DoD) National Defense Science and Engineering Graduate (NDSEG) Fellowship Program. The program's continued support is greatly appreciated. This work is also supported by the University of Michigan's Rackham Conference Travel Grant.

REFERENCES

- Andrews, D. (2018). The Sophistication of Early Stage Design for Complex Vessels. *International Journal of Maritime Engineering*, Vol 160(SE 18).
- Bernstein, J. I. (1998). Design Methods in the Aerospace Industry: Looking for Evidence of Set-Based Practices. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, USA.
- Braha, D. and Reich, Y. (2003). Topological structures for modeling engineering design processes. *Research in Engineering Design*, 14(4):185–199.
- Burrow, J., Doerry, N., Earnesty, M., Was, J., Myers, J., Banko, J., McConnell, J., Pepper, J., and Tafolla, T. (2014). Concept Exploration of the Amphibious Combat Vehicle. In *SNAME Maritime Convention*, Houston, Texas, USA. SNAME.
- Doerry, N. (2012). A Vision for Ship Design and Analysis Tools. *SNAME (mt) Marine Technology*, pages 8–9.
- Doerry, N. (2015). Measuring Diversity in Set-Based Design. In *ASNE Day*, pages 1–14, Arlington, VA, USA.
- Dong, A. (2017). Functional lock-in and the problem of design transformation. *Research in Engineering Design*, 28(2):203–221.
- Garner, M., Doerry, N., MacKenna, A., Pearce, F., Bassler, C., Hannapel, S., and McCauley, P. (2015). Concept Exploration Methods for the Small Surface Combatant. In *World Maritime Technology Conference*, page D021S004R005, Providence, Rhode Island, USA. SNAME.
- Gembarski, P. C., Plappert, S., and Lachmayer, R. (2021). Making design decisions under uncertainties: probabilistic reasoning and robust product design. *Journal of Intelligent Information Systems*, 57(3):563–581.
- Goodrum, C. J. (2020). *Conceptually Robust Knowledge Generation in Early Stage Complex Design*. PhD thesis, University of Michigan, Ann Arbor, MI, USA.
- Gumina, J. M. (2019). *A Set-Based Approach to Systems Design*. PhD thesis, Naval Postgraduate School, Monterey, CA, USA.
- Hatchuel, A. and Weil, B. (2009). C-K design theory: an advanced formulation. *Research in Engineering Design*, 19(4):181–192.
- Kana, A. A. (2017). Forecasting design and decision paths in ship design using the ship-centric Markov decision process model. *Ocean Engineering*, 137:328–337.
- McKenney, T. and Singer, D. (2014). Set-Based: A concurrent engineering approach with particular application to complex marine products. *SNAME (mt) Marine Technology*, pages 51–55.
- Mebane, W. L., Carlson, C. M., Dowd, C., Singer, D. J., and Buckley, M. E. (2011). Set-Based Design and the Ship to Shore Connector. *Naval Engineers Journal*, 123(3):79–92.
- Mihm, J. and Loch, C. H. (2006). Spiraling out of Control: Problem-Solving Dynamics in Complex Distributed Engineering Projects. In Braha, D., Minai, A. A., and Bar-Yam, Y., editors, *Complex Engineered Systems*, pages 141–157. Springer Berlin Heidelberg, Berlin, Heidelberg. Series Title: Understanding Complex Systems.

Niese, N. D., Kana, A. A., and Singer, D. J. (2015). Ship design evaluation subject to carbon emission policymaking using a Markov decision process framework. *Ocean Engineering*, 106:371–385.

Page, S. E. (2006). Path Dependence. *Quarterly Journal of Political Science*, 1(1):87–115.

Parraguez, P., Eppinger, S. D., and Maier, A. M. (2015). Information Flow Through Stages of Complex Engineering Design Projects: A Dynamic Network Analysis Approach. *IEEE Transactions on Engineering Management*, 62(4):604–617.

Shields, C. P. F. (2017). *Investigating Emergent Design Failures Using a Knowledge-Action-Decision Framework*. PhD thesis, University of Michigan, Ann Arbor, MI, USA.

Shields, C. P. F. and Singer, D. J. (2017). Naval Design, Knowledge-Based Complexity, and Emergent Design Failures. *Naval Engineers Journal*, 129(4):75–86.

Singer, D. J., Doerry, N., and Buckley, M. E. (2009). What Is Set-Based Design? *Naval Engineers Journal*, 121(4):31–43.

Smith, R. P. and Eppinger, S. D. (1997). Identifying Controlling Features of Engineering Design Iteration. *Management Science*, 43(3):276–293.

Sypniewski, M. J. (2019). *A Novel Analysis Framework for Evaluating Predisposition of Design Solutions through the Creation of Hereditary-Amelioration Networks Derived from the Dynamics within an Evolutionary Optimizer*. PhD thesis, University of Michigan, Ann Arbor, MI, USA.

Toche, B., Pellerin, R., and Fortin, C. (2020). Set-based design: a review and new directions. *Design Science*, 6:1–41.

Van Houten, J., Singer, D., and Collette, M. (2022). Balancing Designer Influence with Rework for Design Paths of a Simple Polynomial Model. In *Practical Design of Ships and Other Floating Structures*, volume 1, pages 508–527, Dubrovnik, Croatia.

Ward, A., Liker, J. K., Cristiano, J. J., and Sobek II, D. K. (1995). The second Toyota paradox: how delaying decisions can make better cars faster. *Sloan Management Review*, 36(3):43–61.

Wynn, D. C. and Eckert, C. M. (2017). Perspectives on iteration in design and development. *Research in Engineering Design*, 28(2):153–184.

APPENDIX A

Table A1: User Inputs established in Python for validating SBD simulation

Simulation Parameters	Parameter Values
problem_name	'SBD1'
iters_max	1000
sample	'uniform'
search_factor	100
total_points	10000
run_time	[2, 3, 4]
exp_parameters	array([0.2, 2.2, 1.0, 0.95])
part_params	{'cdf_crit': [0.1, 0.1], 'fail_crit': [0.0, 0.05], 'dist_crit': [0.2, 0.1], 'disc_crit': [0.2, 0.1]}
dtc_kwargs	{'max_depth': 2}
gpr_params	{'length_scale_bounds': (1e-2, 1e3), 'alpha': 0.00001}
bez_point	{'P0': (0.0, 1.0), 'P1': (0.5, 0.8), 'P2': (1.0, 0.0)}