

Design of Engineering Systems under Uncertainty via Real Options and Heuristic Optimization

Rania Hassan

Aeronautics & Astronautics and Engineering Systems
rhassan@mit.edu

Richard de Neufville

Engineering Systems and Civil and Environmental Engineering
ardent@mit.edu

Massachusetts Institute of Technology, Cambridge, MA (U.S.A.)

Abstract – This paper presents a practical procedure for using real options valuation in the design optimization of complex engineering systems. Recognition of future uncertainty in both design requirements and the operating environment is the point of departure – a significant shift away from traditional design practice that posits known values for key technical and economic factors. The process leads to the identification of an initial design and a strategy for implementing future expansions according to the way the future unfolds -- in contrast to the usual real options analyses that define a price for the option. Optimization of the design with the recognition of uncertainty leads to significant (greater than 10%) improvements in system performance. The optimization results in multiple desirable attributes, not only the maximization of expected value but also other considerations that may be useful to project managers, such as maximum possible loss or gain and the robustness of the design. A Value-at-Risk diagram conveniently displays these criteria. The optimization itself is extraordinarily complex, compared to standard options analyses, because the reality of design means that the performance of future states is not path independent (because the system evolves in response to its environment) so that the number of possible combinations is astronomical (over 10 to the power of 60 for the simple example problem). A Genetic Algorithm provides a practical solution to such problems, as demonstrated by a generic problem concerning the development of an offshore oil pipeline network.

Keywords: Real Options, Design under Uncertainty, Flexibility, System Architecture, Optimization, Genetic Algorithms, Petroleum Production, Network Optimization, Evolvability, Adaptation.

Background

Traditionally, engineers design large-scale systems to requirements specified in advance. In the United States, for example, the Federal government develops major military and civilian systems through a well-codified “systems engineering” process that explicitly involves a top-down definition of requirements for the design. Internationally, major commercial enterprises define both technical and economic requirements for their engineers to follow. In the extractive industries such as mining and petroleum, management commonly requires engineers to design facilities according to best estimates of the quantity and quality of the material (such as the size and grade of the copper ore), and a sanctioned price of the refined product and discount rate to be used in the evaluation. This process leaves little room for the engineering team to envisage, let alone provide for, alternative scenarios.

Engineering professionals widely accept, indeed endorse, the process of designing to specifications. From personal careers, we know that the engineering paradigm understands its domain to be the strictly technical matters. It is content to leave questions of value or purpose of a system to the managers, generals, or political leaders that specify the nature of the system the engineers are designing. Whether the system is an automobile production plant, a fleet of fighter aircraft, or transportation expressways, engineering practice generally leaves the specification of its purpose, size, and performance to the prospective owners and users. Only recently have designers begun to think systematically about how they should factor the values and assumptions of the stakeholders into the design as variables (see Bartolomei, 2006 for example).

Contrary to this paradigm, as experienced designers know and others can easily imagine, the reality is that system “requirements” change, often radically. The technology, market needs and economic situation easily vary extensively during the years – decades – over which the design, deployment and operation of a system takes place. For example, between the late 1980s when Globalstar and Iridium laid down the requirements for their satellite communications systems, and the late 1990s when they became operational, both the technology and market for cell phones had evolved drastically. The result was that these systems, which met their technical requirements brilliantly, were financial disasters (see de Weck et al, 2004). In general, we can assume that requirements for any major system are likely to change over its lifetime. Put another way, designers should assume that no one is capable of specifying requirements for major systems accurately: technology, markets, needs and the world change beyond our ability to specify precisely.

The design procedure presented here recognizes the unavoidable uncertainty in design requirements. Its point of departure is the fact that we know in advance that important design factors are likely to change over the lifetime of the system. The procedure accepts this reality and deals with it proactively by creating a flexible design with real options that project managers can adapt effectively over time, either to avoid adverse situations, or to take advantage of attractive opportunities. By avoiding bad outcomes and seizing good ones, the procedure leads to significant improvements in performance.

Concept of the Procedure

The basic concept is to introduce flexibility into the design of a system, to enable the project managers to adjust the design to the circumstances that prevail as they implement a system. This flexibility provides designers with “rights but not obligations” to develop their facilities in particular ways. They can, for example, expand or alter the design if that suits

their needs. Flexibility provides “options” in the strict technical sense of the word – these are not merely alternatives, they are capabilities to react easily in a number of ways that would not be possible unless designers make intentional choices in the initial stages of the design. Specifically, the flexibility built into the system creates “real options”, that is “options” that are “real” in that they refer to physical projects rather than the more familiar financial options that deal with company shares, foreign exchange and the like. Numerous textbooks describe real options, for example Trigeorgis (1996). In short, the procedure creates “real options”.

This concept contrasts sharply with traditional design that optimizes a system and then launches it with no intention of revising the design in mid-stream. This traditional approach creates a “rigid” design compared to the “flexible” design that the procedure explained hereafter produces. The case study presented in this paper demonstrates why flexible designs are superior to rigid design on many dimensions.

The output of the procedure is an optimal design. The idea is to identify the best initial design, the one that can evolve most favorably over time. An economic analysis is the means used to achieve this objective. In that sense, the procedure values the real options and performs a “real options analysis”. However, this analysis differs conceptually from common real options analyses performed by economic analysts who are primarily interested in the value of the option, from the perspective of buying or trading it. By contrasting the value of the optimal flexible design with the best flexibility, with the value of the rigid design, the procedure does imply a value for the real option embedded in the flexible design. However, this value is a secondary aspect of the exercise. Ultimately, designers are not interested in whether this value is precisely accurate: they need to know if they are making the right kinds of design decisions. The fact that a flexible design concept performs 10 to 15% better than the rigid concept is the key issue, not whether the improved performance can be known precisely.

The overall object is to present a simple, but effective and efficient approach to the design of large-scale, complex engineered systems under uncertainty. The intended approach should enable system architects to design better projects, ones that deal with both sides of uncertainty, that is, that both hedge risk and exploit opportunity. Its success is to be judged by the extent to which it provides insight into ways to improve design effectively.

Procedure Steps

The procedure for optimizing the design of complex systems under uncertainty is an extension of existing procedures. It assumes that the designers have models for valuing any design, that is, for estimating its benefits and costs in some metric. This is essential for any optimization process: without a way to measure performance, the concept of optimization lacks meaning. The proposed process builds upon this capability by adding an analysis of the performance of alternative flexible designs, which thereby determines an optimal performance in the uncertain future.

The procedure consists of the following four steps:

1. The specification of the key sources of uncertainty and description of their distribution of possible future states over time. For example, the design of a copper extraction process might specify that the key uncertainties are the price of copper and the grade of ore. The description of the future states would then be probability distributions

over the relevant ranges. Normally, uncertainties grow over time, but they can also decrease, for example as a technology or situation matures and becomes less variable.

2. The definition, for sake of manageable calculation, of a limited number of future states over several stages or design periods in the future. For example, one might assume that the price of copper might move to one of three possible states in the first stage, three more in the next and so on. A stage in this context would be a suitable planning period. That depends on the rate of development of the technology or the market. A stage might be six months for the development of computer systems, and 5 years for major infrastructure projects. In terms of conventional options analysis, this step is similar to that of setting out a binomial or trinomial tree. At each node, that is, at each combination of time and state of the uncertain variables (such as copper price and ore quality), the project managers can develop the project in the optimal way.
3. The valuation of all possible aggregate combinations of designs over the stages and states. The process is comparable, in terms of financial options, to calculating the best time and conditions to exercise an option. Financial analysts usually do this using some variant of dynamic programming. This simple process is inappropriate in the case of design, however, because the value of being at any node (that is, combination of time and states of events) is path dependent. The fact is that project managers respond to circumstances, and will have properly built different facilities according to the way the uncertainties unfold. [In terms of a binomial tree, the project would look different if the prices have been “down and then up” in contrast to “up and then down”.] (Wang 2005, and Wang and de Neufville. 2005 and 2006). This means that dynamic programming would not be applicable. Moreover, the number of possible designs easily swamps most optimization procedures, which is why this procedure suggests the use of a Genetic Algorithm, as the next section describes.
4. The choice of the optimal initial design, and the specification of the optimal development strategies, according to the criteria appropriate to the decision-making unit. The most obvious criterion would be the maximization of expected value. However, many organizations will also be interested in other criteria such as: possible losses, possible maximum gains, the robustness of the design (that is, the reduction in the range of value delivered), the initial capital expenditure as opposed to capital investment in latter stages, etc. Designers can explore these criteria through the Value-at-Risk diagram, as explained further in the case study. (see de Neufville et al, 2006)

Optimization for System Design via a Genetic Algorithm

The value functions for most general designs of infrastructure systems make it inappropriate to use the most common optimization tools, such as variations on linear or dynamic programming. This is because they involve:

- non-convex feasible regions, due to the routine existence of economies of scale in extractive, process and manufacturing industries, which thus preclude linear programming variants;
- path dependence over time, as explained above, that excludes dynamic programming schemes; and
- astronomical size, when any reasonable account is made of the possible joint variations in Location (L), Size (S) and Timing (T) of the construction of any facility (the total number of possibilities being on the order of $[(L \text{ exp } S) \text{ exp } T]$). In the simple case study reported in this paper, the total number is on the order of 10 to the power of 60!

Genetic Algorithms (GA) provide an effective means to solve very large, non-convex, path dependent problems. The GA operates on the Darwinian principle of survival of the fittest in its search process to arrive at optimal solutions. GAs are part of evolutionary computing, a rapidly growing area of soft computing techniques. They are inspired by the principles of genetics and evolution and were developed by John Holland and his colleagues at the University of Michigan (see Holland, 1975). GAs are well suited to complex design problems because they can handle both discrete and continuous variables, and objective and constraint functions are not required to be differentiable or continuous. (See Goldberg, 1989; and Hassan, 2004)

Case Study: Development of Offshore Oil Pipeline Networks

Overall Perspective: To illustrate the proposed procedure, we applied it to a case study considering the design of a pipeline network and production plan for an offshore oil field. The field is to be developed over three stages or planning periods (S1, S2, and S3). The situation is hypothetical and the numbers used are stand-ins to permit calculation.

The hypothetical offshore oil field includes six wells, A through F. Figure 1 shows their geographic distribution and predicted reservoir volumes, which decrease counter clockwise. The concept is to extract the oil by linking the fields to production facilities that are either fixed platforms or FPSO's (Floating, Production, Storage, and Offloading Platforms). The case assumes that a fixed platform would be built at the beginning of a project and cannot be moved or modified to increase or decrease production capacity. Conversely, an FPSO can be moved and its capacity increased or decreased between stages. Fixed platforms and FPSOs are assumed to come in discrete units that can handle up to 100 million barrel per stage.




Figure 1. Layout of hypothetical offshore oil field

For simplicity, the case assumes that pipelines come in three sizes: small, medium and large. Their maximum allowable flow rate depends on both the size of the pipeline and the distance between the oil well and the production facility. It is greater for large pipelines over shorter distance, as Table 1 indicates.

Table 1. Assumed maximum allowable production rates of pipelines as function of size and length

pipeline length, (Km)	Maximum allowable pipeline flow rate, (million barrel/stage)		
	small Ø	medium Ø	large Ø
0 (production facility on top of well)	100	150	200
less than 10	70	120	170
between 10 and 30	50	100	150
more than 30	30	80	130

The objective is to find the best pipeline network architecture and production plan. The base case assumes that net present value is the criterion for performance, as is reasonably standard in practice. When considering uncertain futures, managers should invoke modified and supplemental criteria, such as expected net present value, maximum capital expenditure (CAPEX) and value-at-risk. The analysis section of the case describes these criteria in detail.

Cost and Revenue Models: The models used in this case study are simplistic and hypothetical, but sufficient to capture the essential elements of the problem. The cost model includes the cost of the pipelines, production facility, and operations. Revenues are the product of the volume extracted in each stage and the corresponding oil price. The analysis uses a discount rate (r) of 15% to calculate present values.

The cost of building the small, medium and large pipelines is 40, 80 and 120 million US\$ per kilometer respectively. A pipeline built in earlier stages of the project can be used in subsequent stages without additional capital cost. However, this is only physically feasible if the pipeline serves a production facility that does not move. A pipeline can carry up to its full capacity per stage (as shown in Table 1) or less if the volume remaining in a well is less than the maximum allowable flow rate. The cost of a fixed platform equals the leasing price of an FPSO per stage, that is, US\$ 50 million.

The case study also assumes that operating costs vary linearly with the amount of oil produced. Specifically, this cost is US\$1 per barrel. Equation (1) then defines the present value of the total cost for a given design:

$$PV(\text{Cost}) = \sum_{i=1}^3 \frac{(\text{Production Facility} + \text{Pipelines} + \text{Operations})_i}{(1+r)^{i-1}} \quad (1)$$

Revenues are the quotient of the oil price/barrel multiplied by the sum of the discounted extracted volumes from all three stages, as in Equation (2).

$$PV(Revenue) = \sum_{i=1}^3 \frac{Volume_i \times Oil\ Price_i}{(1+r)^{i-1}} \quad (2)$$

where $Volume_i = (V_A + V_B + V_C + V_D + V_E + V_F)_i$

Base Case – Conventional Optimization of a Rigid Architecture: The base case assumes a deterministic situation, as is commonly done in practice. Specifically, it posits that the price of oil is constant over all three stages at US\$ 30/barrel; and that volume in each well is precisely as predicted and shown in Figure 1.

If there is no uncertainty about the future, the optimum design can be fully specified in advance. This is a rigid design because the plan undergoes no modification once it has been defined. For example, the plan will determine the location and the capacity of the production facility and pipeline sizes. These are built at the beginning of the project and cannot be changed at later stages. The same rigidity applies to the production rates, which the plan specifies before the project starts and implements over the three stages.

The design objective is to specify the location of the production facilities, the pipeline sizes, and the production plan that maximizes net present value as in Equation (3).

$$\text{Maximize } NPV = PV(Revenue) - PV(Cost) \quad (3)$$

Although this is a simple description of the problem, it involves many design variables as Table 2 indicates. Most importantly, even this simple problem implies a huge design space: about 268 million possible combinations could be considered.

Table 2. Design variables describing a rigid architecture

Variable	Description	Discrete values
X_1	platform location *	16 possibilities [†]
X_2	pipeline size to well A ‡	N, S, M, or L [§]
X_3	producing well A in Stage II	yes or no **
X_4	producing well A in Stage III	yes or no **
	⋮	
X_{17}	pipeline size to well F ‡	N, S, M, or L [§]
X_{18}	producing well F in Stage II	yes or no **
X_{19}	producing well F in Stage III	yes or no **

* same location in all stages and built in Stage I

† carefully chosen on the grid shown in Figure 1

‡ built in Stage I, production rate = the minimum of pipeline capacity or the oil volume in well

§ N: No pipeline, S: Small, M: Medium, L: Large

** if yes, production rate = the minimum of pipeline capacity or the oil volume *remaining* in well

For this base, deterministic case, the GA optimizer selected small pipeline sizes for all oil wells and used them at their maximum capacity for all wells in Stage I, for wells A, B, C and D in Stage II, and for well A in Stage III. Figure 2 shows the facilities in use in each stage. This plan implies present values of revenue and cost of US\$ 1828 and 1006 million respectively, for a net present of US\$ 822 million.




Figure 2. Optimal rigid design for deterministic base case

Application of Procedure

Step 1 – Specification of Key Uncertainties: The first task in optimizing a system that will perform in an uncertain future is to identify, as best one can, what these futures might be. Good design should deal with the prospective situation as clearly as possible. It is worth noting in this context that considering the uncertainties as part of the design at the beginning is demonstrably superior to the common practice of designing to the most likely or expected situation, and then considering uncertainties (opportunities and risks) afterward. This follows from a basic proposition of probability analysis, known as Jensen’s Inequality shown in Equation (4), which states that the expected value of a metric is not equal to the value calculated from average values. To assume otherwise is to fall into the “Flaw of Averages” (see examples in de Neufville et al, 2006).

$$EV(f(x)) \neq f(EV(x)) \quad (4)$$

As even casual observers recognize, the price of oil fluctuates substantially. We thus took it as the primary uncertain parameter for this case study. Additionally, our complete analysis considered the quantity of oil to be uncertain. Indeed, in practice the quantity and quality of oil and other extracted assets (gas, copper, coal, etc.) cannot be known precisely in advance. More forcefully, initial estimates are often very wide of the real situation. The analysis with uncertainty in the quantity of oil is not reported in this paper, as it parallels and echoes the results of the analysis of uncertainty in price. Interested readers can get these results from the authors.

Step 2 – Definition of Future States: This case study used a simple discrete distribution with symmetric transition probabilities, as Figure 3 shows. The range of oil prices and the transition probabilities are based on neither predictions nor historical data, but are rather assumed for simplicity.




Figure 3. Evolution of future oil prices (possible prices in boxes, transition probabilities along the arrows)

Readers accustomed to multi-stage binomial matrices of probabilities may well ask why the analysis should not use a more detailed analysis with more stages and states. The answer is that the traditional binomial representation implies that the value of being in any state is independent of the path for getting there as Figure 4 suggests. This is simply not the case for

engineering projects, whose evolution does indeed depend on the way uncertainties evolve, since managers do respond to opportunities and threats as they develop.

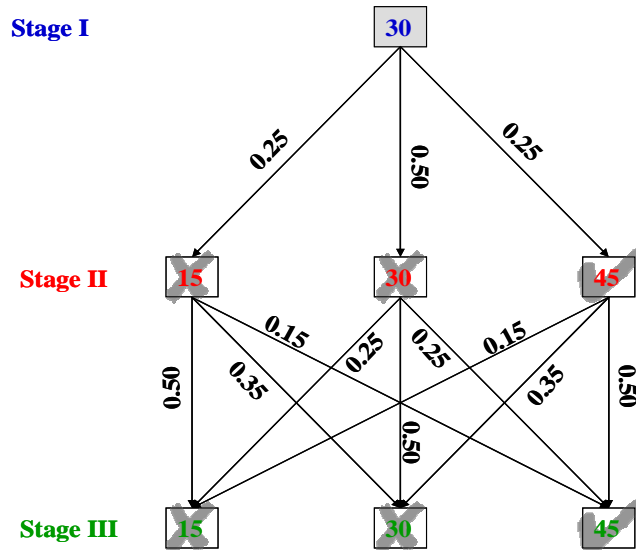


Figure 4. Path independence in financial options valuation

The pervasiveness of path dependence explains why there are nine states at the end of Stage III, instead of three. This is because the value of having a US\$ 30/barrel of oil in the last stage is very different if the preceding price in Stage II were only US\$ 15 (and the owners has held back production and investment) or if it were US\$ 45 (so that the owners had accelerated pumping and investment). A realistic representation of the situation is thus as in Figure 5.




Figure 5. Designing Real Options in the system (path dependency illustrated)

Figure 5 makes evident that the value of the investment in the oil risk is risky. The actual revenues depend on the future prices of oil. The project could bring in considerably greater or less net present value than estimated in the deterministic analysis. Using the distribution in Figure 3 to evaluate the base case, the deterministically optimal rigid design shown in Figure 2, leads to nine different net present values. Their probability of occurrence depends on the two-stage path along the nine branches of the discrete tree in Figure 5.

The risk associated with any project is conveniently displayed in a Value-at-Risk (VAR) diagram as in Figure 6. The VAR is simply a plot of the cumulative density function of the net present value distribution. For this case study, in which revenues are a linear function of price, the expected value associated with the distribution is the same as that calculated for the deterministic case. Because the distribution of the price of oil is assumed symmetric around the base price of US\$ 30 as shown in Figure 3, and because the cost and revenues functions are linear, the VAR is also symmetric, and the expected outcomes range from -56% to $+56\%$ of the mean of US \$ 822 million. The wide range of outcomes can also be characterized in terms of the standard deviation of the distribution around the mean, which in this case is US\$ 285 million or, about 30% of the average. As shown in Figure 6, the VAR diagram provides a means to present the performance of the project not only from the perspective of an expected value, but also considering the range of uncertainty.

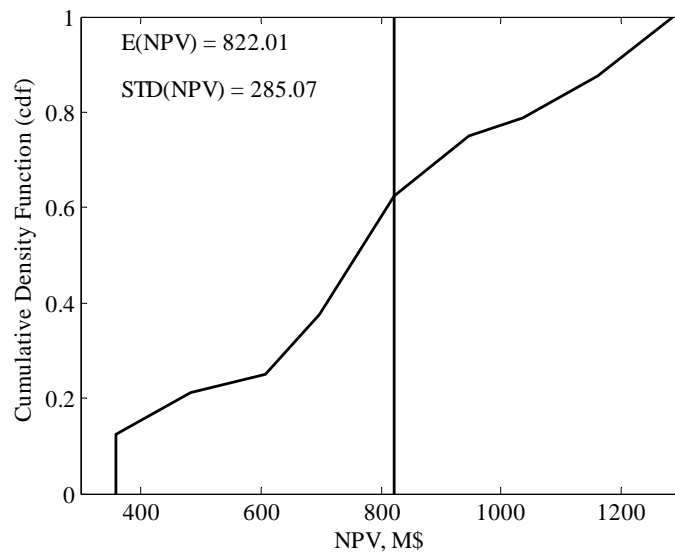


Figure 6. Value-at-risk (cumulative density function of NPV) for rigid optimal design

Step 3 – Valuation of Possible Designs: Once designers recognize the risk and opportunities, as in Step 2 by setting out representative scenarios, they can explore the possibilities for optimizing under these uncertainties. The objective is to define system architectures whose configuration allow them to adapt depending on scenario materialization. For the case study, designing for flexibility can be explained using the schematic in Figure 5. Stage I should be designed bearing in mind that Stage II could take three different forms depending on oil price. Likewise, each of the forms of Stage II should be designed bearing in mind that Stage III could take three different forms based on oil price. Therefore, a flexible architecture for the case study will be composed of a single form for Stage I, 3 forms for Stage II and 9 forms for Stage III.

Note that there are nine forms for Stage III although there are only three possible oil prices in this case. This is because of the “path dependency” issue. For example, if the oil price goes to US\$ 15 /barrel or 45/barrel in Stage II, this stage will most probably look differently for the two prices. Likewise, if the price goes to US\$ 30/barrel in Stage III, there will be three different forms of this stage depending on what was the form of Stage II.

In the statement of the optimization problem for a rigid architecture, three sets of design variables are needed to represent the decisions for each of the three stages, as Table 2 shows. However, the optimization problem statement for a flexible architecture must contain 13 sets of design variables (Table 3). One set is required for each of the 13 forms (= 1 +3 +9) associated with the three stages shown in Figure 5.

Table 3. Design variables describing a flexible architecture

variable	description	discrete values
X ₁	FPSO location in S I , o.p.* = 30	16 possibilities
X ₂ : X ₇	pipeline sizes to wells A to F	N, S, M, or L
X ₈	FPSO location in S II , o.p. = 15	16 possibilities
X ₉ : X ₁₄	pipeline sizes to wells A to F	N, S, M, or L
X ₁₅	FPSO location in S III , o.p. = 15	16 possibilities
X ₁₆ : X ₂₁	pipeline sizes to wells A to F	N, S, M, or L
X ₂₂	FPSO location in S III , o.p. = 30	16 possibilities
X ₂₃ : X ₂₈	pipeline sizes to wells A to F	N, S, M, or L
X ₂₉	FPSO location in S III , o.p. = 45	16 possibilities
X ₃₀ : X ₃₅	pipeline sizes to wells A to F	N, S, M, or L
X ₃₆	FPSO location in S II , o.p. = 30	16 possibilities
X ₃₇ : X ₄₂	pipeline sizes to wells A to F	N, S, M, or L
X ₄₃	FPSO location in S III , o.p. = 15	16 possibilities
X ₄₄ : X ₄₉	pipeline sizes to wells A to F	N, S, M, or L
X ₅₀	FPSO location in S III , o.p. = 30	16 possibilities
X ₅₁ : X ₅₆	pipeline sizes to wells A to F	N, S, M, or L
X ₅₇	FPSO location in S III , o.p. = 45	16 possibilities
X ₅₈ : X ₆₃	pipeline sizes to wells A to F	N, S, M, or L
X ₆₄	FPSO location in S II , o.p. = 45	16 possibilities
X ₆₅ : X ₇₀	pipeline sizes to wells A to F	N, S, M, or L
X ₇₁	FPSO location in S III , o.p. = 15	16 possibilities
X ₇₂ : X ₇₇	pipeline sizes to wells A to F	N, S, M, or L
X ₇₈	FPSO location in S III , o.p. = 30	16 possibilities
X ₇₉ : X ₈₄	pipeline sizes to wells A to F	N, S, M, or L
X ₈₅	FPSO location in S III , o.p. = 45	16 possibilities
X ₈₆ : X ₉₁	pipeline sizes to wells A to F	N, S, M, or L

* oil price, for rest of acronyms, see footnotes of Table 2

The large number of design variables needed to formulate optimization problems for designing for flexibility is typical for engineering projects. This is because of the path dependency. In that sense, working with real options in the design of engineering systems is distinctively different from valuing financial options. This situation leads to two major points:

- Optimization problems involving real options analysis are typically very expensive compared to those involving financial options; and
- Methods commonly used to evaluate financial options, such as the Black-Scholes formula, do not apply to evaluation of the value of flexibility in the design of complex engineering systems.

The analysis used a Genetic Algorithm to optimize the design for flexibility. The GA is a population-based search, i.e. it moves from a set of solutions to another with likely improvements in each iteration. At the end of the optimization, the GA therefore not only provides an optimal design (or pseudo optimal because there is no proof of optimality as the GA is not a calculus-based search) but also a set of good designs with high performance values.

Step 4 – Choice of Optimal Design: To illustrate the choice of design, this section presents not only the optimal solution, as defined by the optimization process, but also few others. These alternatives provide different risk profiles that might be attractive to the owners and operators of the system, as the subsequent discussion illustrates.

Figures 7, 8, and 9 show three alternative solutions in detail. The top row in each Figure contains three subplots. The middle subplot is the design of Stage I. The one on the left summarizes the performance of the alternative in terms of the present values of cost and revenue, the expected net present value and the associated standard deviation. The right hand subplot compares the VAR distribution of the flexible architecture (dashed line) to that of the rigid architecture under uncertain oil prices (solid line). The second row in Figures 7, 8 and 9 includes three subplots describing the three forms of Stage II at the oil prices of US\$ 15, 30 and 45/barrel. The bottom of the plot consists of three 3×2 matrices. The second columns of the matrices include the three forms of Stage III for the oil prices of 15, 30 and 45 organized vertically from top to down. The left hand, middle and right hand matrices are associated with the three Stage II forms when the oil price is 15, 30 and 45 respectively. The text subplots in the left columns in the matrices summarize the data (cost, revenue, net present value and probability of occurrence) corresponding to the respective tree path. For example, the top text subplot in the left hand matrix summarizes the data of the following path: oil price is equal to 30, 15 and 15 in Stages I, II and III respectively. In each graph, the solid lines represent the pipelines built in Stage I that could be used in Stages II and III if the FPSO locations do not change and the pipeline sizes stay the same. The pipelines built in Stage II are marked as dash-dotted lines and could be used in Stage III under the same conditions. Finally, pipelines built in Stage III are marked as dashed lines

The first alternative, in Figure 7, illustrates the concept of flexibility or real options “in” design as defined by de Neufville et al (2006). A real option “in” the system exists thanks to some technical arrangement that enables the project managers to alter the performance of the system. Real options “in” a system differ from real options “on” the system in that the latter have nothing to do with the design of the system. A real option “on” the system might be the option to delay its construction for example. This flexibility is an option, and it is real in that it deals with physical products rather than contracts, but it treats the engineering system as a black box. A real option “in” the system, by contrast, results from

specific design decisions. In the case of Figure 7, the design of Stage I is suboptimal from a cost perspective because the cheaper design for Stage I would be to place the FPSO on top of one of the wells being produced, B or C, with a pipeline to the other well. However, such a decision would affect the subsequent stages negatively. This is because it is cheaper overall to keep the FPSO at its suboptimal position in Stage I, when considering the cost of the pipelines extended in Stages II and III to the remaining oil wells. This alternative demonstrates the idea that in a flexible architecture the design of the early stages of a project should consider the subsequent effects on later stages.

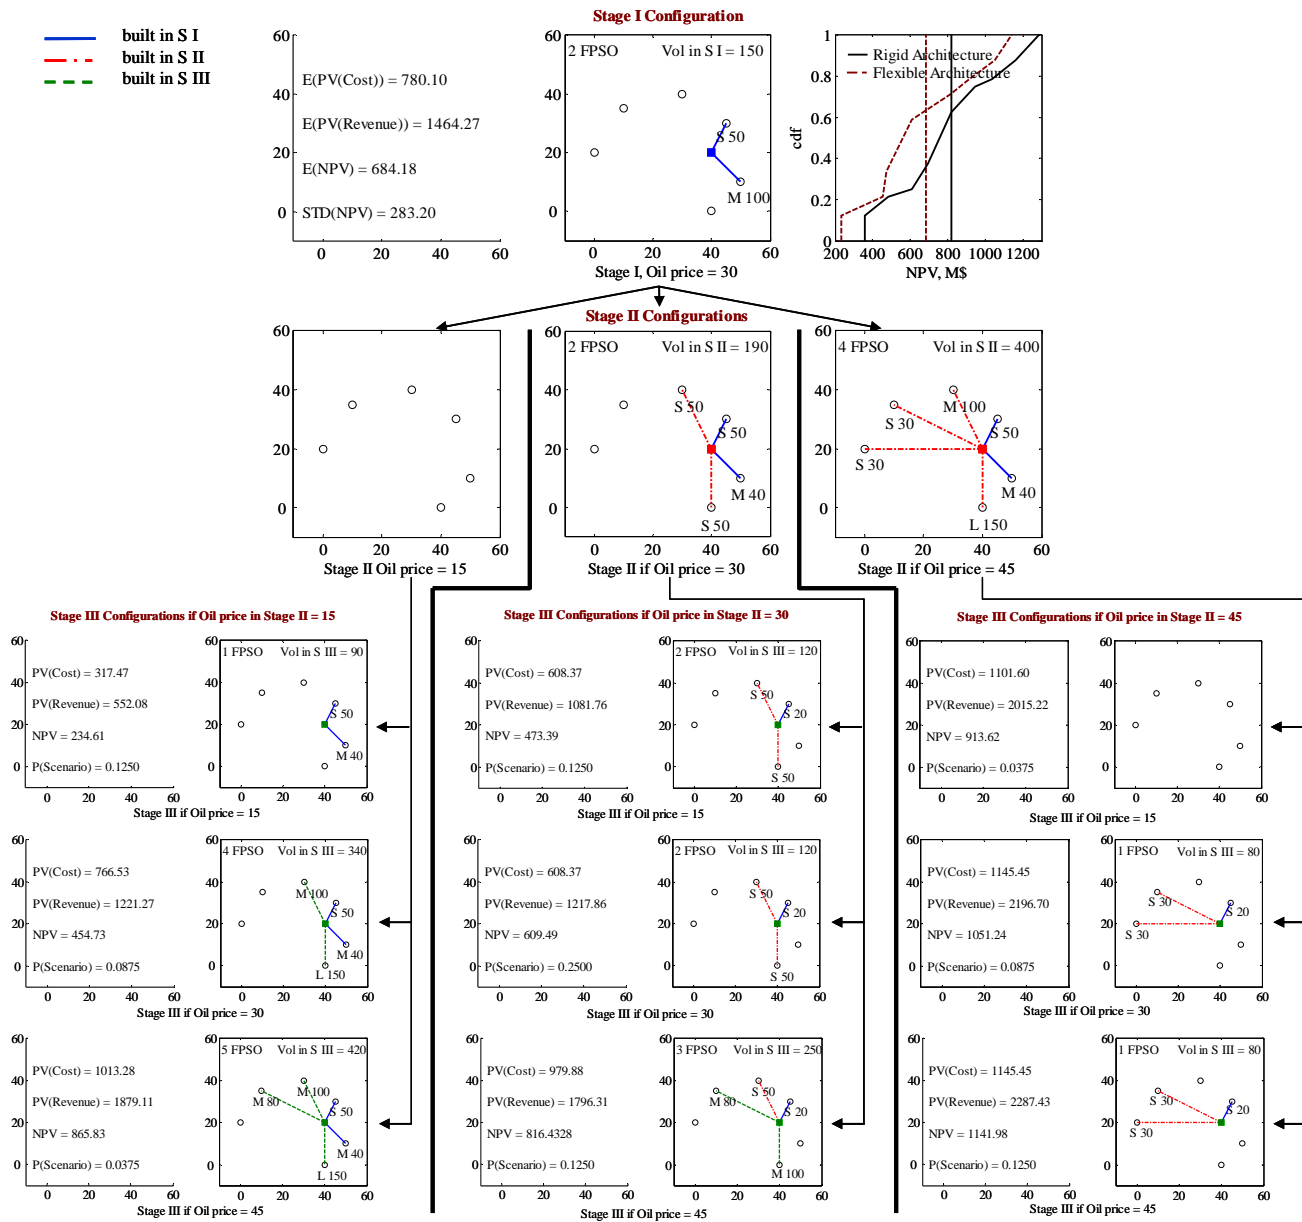


Figure 7. A demonstration of flexibility in system design

The second alternative, in Figure 8, is the optimal flexible architecture generated by the GA, that is, the one with the highest expected net present value. This design not only has an expected net present value 13% greater, but also has several other advantages compared to the rigid architecture:

- it requires less upfront capital expenditure (CAPEX),
- its expected present value of cost is, and
- its standard deviation of net present value is 43% less.

This optimal flexible architecture not only improves the value of the project while decreasing the cost, but also is robust in that it produces a tight range of “better outcomes” by shifting the VAR distribution to the right of that of the rigid architecture.

The third alternative, in Figure 9, presents a low cost, near optimal flexible architecture. The GA generated this alternative as part of its evolutionary analysis. Indeed, one of the benefits of the GA is that it can provide a large number of attractive designs that have interesting features. For example, this particular architecture has the merit of only requiring 68% of the expected present value of the cost of the rigid architecture, while delivering about a 9% gain in expected net present value. This is possible because this flexible alternative only exploits one oil well in Stage I, thereby eliminating the cost of extending expensive pipelines early in the project. This reduction in CAPEX might make this alternative a favorite for some investors. This alternative illustrates that a narrow focus on one dimension of merit – such as expected net present value – may hide attractive alternatives.

Finally, the full analysis also examined some other designs not presented in this paper. These were rigid designs designed for “robustness” in that they reduced the standard deviation of the expected performance (see Hassan and de Neufville, 2006). They thus would be more likely to deliver the results promised. For some designers, this is a most important feature.

Summary of Results: Table 4 compares the results of the optimal rigid, flexible, and robust designs (not presented in this paper, see Hassan and de Neufville, 2006). It demonstrates that flexible architectures provide a range of advantages over rigid and robust architectures. Flexible architectures:

- improve the expected value of the project,
- tighten the range of outcomes such that the architectures are also robust on top of being flexible, and
- decrease upfront CAPEX cost as well as the expected present value of the cost.

The result that flexibility reduces cost deserves special attention. Designers and financial analysts often assume that flexibility and options always cost money. This is simply not true in engineering design – in contrast to the financial world. In the market place, since options have value, traders can only obtain them at a price. In engineering, however, designers create real options by being insightful, and have neither need nor obligation to buy them. When they create flexible architectures, they create the possibility of reacting appropriately to future design needs, thus avoiding the costs of both inappropriate elements that might not be needed, and premature investments that are more economically deferred (see for example Hassan et al, 2005). Good designers can both create flexibility and reduce costs simultaneously.

The option value of any alternative can be deduced from the results in Table 4. As a working estimate, it is simply the difference between the value of the alternative and the traditional rigid design. Thus the option value of the alternative with the highest option value (Flexible II) is simply $(929 - 822) = 107$. In this case, the best flexible design:

- represents about a 13% improvement over the traditional deterministic approach in net present value terms, while
- being more robust and more likely to perform according to expectations (lower standard deviation of returns),
- more than doubling the minimum possible value of the project (from 358 to 788), and
- reduces the expected total costs (from 1006 to 969).

The only advantage of the traditional design in this case is that it is more of a gamble and has a small probability of a slightly higher upside potential (1286 compared to 1148).




Figure 8. Flexible architecture with highest expected net present value (optimal architecture)

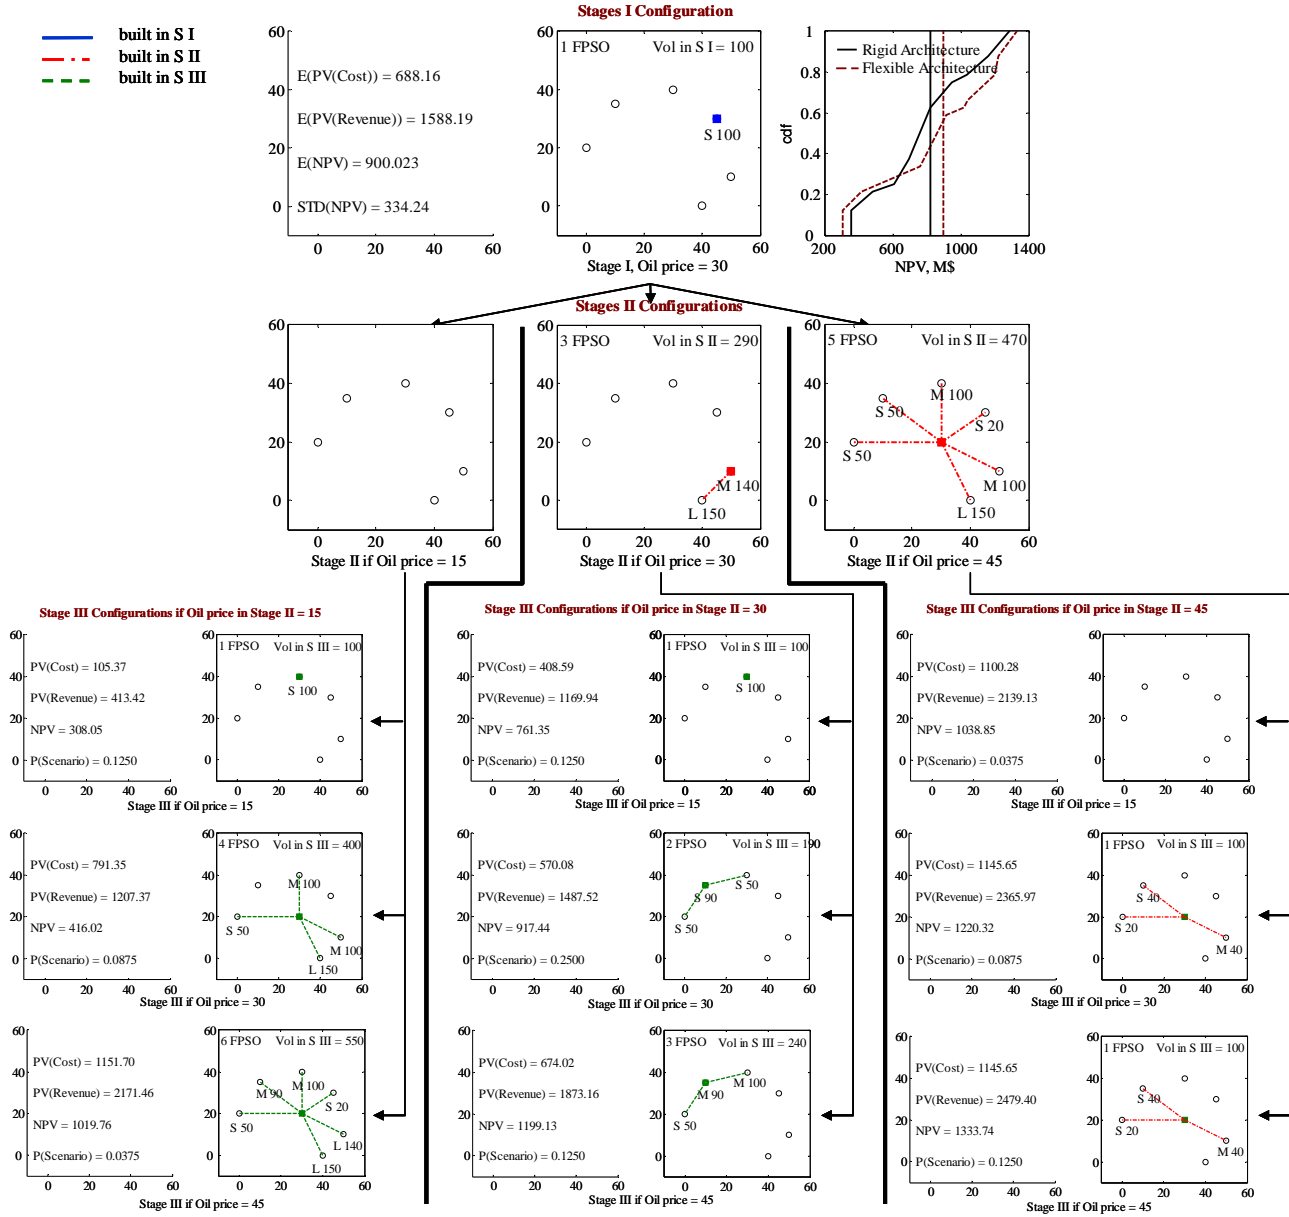


Figure 9. Low cost, near-optimal flexible architecture

Table 4. Summary of performance metrics of optimal rigid, robust and flexible pipeline architectures

Value Metric (\$ millions)	Rigid	Robust I	Robust II	Flexible I	Flexible II	Flexible III
		STD < 30% E(NPV)	STD < 25% E(NPV)	Flexibility Obvious	Highest E(NPV)	Low Cost
E(NPV)	822	792	685	684	929	900
STD(NPV)	285	233	159	283	165	334
Min Possible Gain	358	407	423	234	788	308
Max Possible Gain	1286	1178	947	1142	1148	1334
E(PV(Cost))	1006	579	439	780	969	688

Computational Cost Considerations

Considerable computational costs are associated with designing under uncertainty. Table 5 compares the effort needed to solve optimization problems for both rigid and flexible pipeline architectures in this case study. It shows that the design space associated with optimizing a flexible architecture is tens of orders of magnitude larger than that of the rigid architecture. This clearly makes designing for flexibility much harder than conventional rigid architecture design, which could be a major hindrance in designing-for-flexibility. However, the GA provides an efficient approach and entails only a single order of magnitude increase in computational cost beyond that needed for optimizing a rigid architecture. For this case study, all the GA runs were done in MATLAB on a desktop with a Pentium 4 processor.

Table 5. Computational cost requirements for rigid and flexible optimization problems

Computational Cost Metric	Rigid	Flexible
Design Variables	19	91
Variable Coding in Bits	28	208
Exhaustive Design Space	268×10^6	411×10^{60}
GA Function Evaluations	14×10^3	17×10^4
GA Computational Time (mins)	< 1	150

Conclusions

This paper provides an efficient procedure for using real options to improve the design of engineering systems under uncertainty. It demonstrates its use in the context of the development of a hypothetical offshore oil field. It also shows that, although the size of the optimization for designing for flexibility objectives are significantly larger than that for traditional design, Genetic Algorithms can find the optimal architectures at relatively low computational cost.

It is important to note that dealing with real options in the design of engineering systems is substantially different from valuing real options financially. Most obviously, the object of the engineering exercise is to identify improved designs, not to determine a price. Moreover, building flexibility “in” engineering systems requires technical knowledge specific to the system under consideration.

Finally, it should be stressed that flexibility is not the enemy of optimality. Designers do not have to give up value in order to achieve flexible designs. On the contrary, this and other case studies indicate that flexible designs often reduce expected cost and CAPEX while improving value compared to traditional rigid designs.

Acknowledgements

The authors thank BP plc and the MITRE Corporation for helping to fund the research that led to this paper. We also gratefully acknowledge our debt to colleagues for their support and advice.

References

- Bartolomei, J., Hastings, D., de Neufville, R. and Rhodes, D. (2006), "Using a Coupled-Design Structure Matrix Framework to Screen for Real Options "In" an Engineering System," *INCOSE 2006*, 10-14 July, Orlando Fl....
- de Neufville, R., Scholtes, S., and Wang, T. (2006) "Real Options by Spread Sheet, Parking Garage Case Example," *ASCE Journal of Infrastructure Systems*, in press for June issue.
http://ardent.mit.edu/real_options/Real_opts_papers/Garage%20Case_Tech_Note%20Draft%20Final%20January.pdf
- de Weck, O., de Neufville, R. and Chaize, M. (2004) "Staged Deployment of Communication Satellite Constellation in Low Earth Orbit," *J. of Aerospace Computing, Information, and Communications*, Vol. 1, No. 3, Mar., pp. 119-131
http://ardent.mit.edu/real_options/Real_opts_papers/JACIC_stageddeploy_aspublished.pdf
- Goldberg, D. (1989), *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, MA, pp. 1-25.
- Hassan, R. (2004) *Genetic Algorithm Approaches for Conceptual Design of Spacecraft Systems Including Multi-Objective Optimization and Design under Uncertainty*, Doctoral Dissertation, Purdue University, May 2004.
- Hassan, R., de Neufville, R., and McKinnon, D. (2005) "Value-at-Risk Analysis for Real Options in Complex Engineered Systems," *IEEE Conference on Systems, Man, and Cybernetics*, Hawaii, October 2005.
http://ardent.mit.edu/real_options/Real_opts_papers/Hassan_deN_IEEE_SMC_2005_final1.pdf
- Hassan, R. and de Neufville, R. (2006) "Foundational Concepts for System Design under Uncertainty," working paper, MIT, Cambridge, MA.
- Holland, J. (1975) *Adaptation in Natural and Artificial Systems: an Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*, U. of Michigan Press, Ann Arbor, MI. Also (1992) MIT Press, Cambridge, MA.
- Trigeorgis, L. (1996) *Real Options: Managerial Flexibility and Strategy in Resource Allocation*, MIT Press, Cambridge, MA.
- Wang, T. (2005) *Real Options "in" Projects and Systems Design – Identification of Options and Solution for Path Dependency*, Doctoral Dissertation, MIT Engineering Systems Division, May 2005.
- Wang, T. and de Neufville, R. (2005) "Real Options "in" Projects", *9th Real Options Conference*, Paris, June 2005.
- Wang, T. and de Neufville, R. (2006) "Identification of Real Options "in" Projects," *4th Annual Conference on Systems Engineering Research*, Los Angeles, CA.