Go Back

The Use of System Dynamics in a Large-Scale System Engineering Effort: Case Study of the Deepwater Project

R. Evan Ellis
JAYCOR, Inc.
1410 Spring Hill Drive, Suite 300
McLean, VA 22102
Tel: 703-847-4015
Fax: 703-847-4131
eellis@jaycor.com

Abstract

This paper describes the use of system dynamics modeling and influence diagramming in the design and phased implementation of the U.S. Coast Guard Integrated Deepwater System. From October 1998 through April 2000, Science Applications International Corporation led a multi-company team in a large-scale system engineering effort to design the next generation system of air, sea, information, and logistics support capabilities for Coast Guard operations more than 50 nautical miles from U.S. shores. The SAIC team used system dynamics modeling in close coordination with subject matter experts, design engineers, and decision makers to support not only the development, but also the evaluation, optimization, and implementation of its design. The team leveraged the Vensim SD software package for modeling, expert-in-the-loop sensitivity analysis and optimization. This system engineering effort pushed the boundaries of the system dynamics methodology by using system dynamics techniques and software for agent-based discrete event simulation with GIS-style visualization.

Keywords

Deepwater, optimization, system engineering, discrete event simulation.

Introduction: The Integrated Deepwater System (IDS) Program

This is a paper about the use of system dynamics in what is, arguably, one of the largest formal system analysis and system engineering efforts undertaken in the United States.

From October 1998 through May 2001, the U.S. Coast Guard conducted a major competitive procurement program to design the next generation system for conducting its missions in waters more than 50 miles from U.S. shores. This program was referred to as the “Integrated Deepwater System (IDS),” or simply the “Deepwater Program.”

The Deepwater program was unprecedented for the Coast Guard in its size, scope, and nature, and represents a major landmark for all government procurement. In the Deepwater program, the Coast Guard chose to think of the resources available to it to perform missions, including persons, assets, infrastructure, and procedures, as “system of systems.” Correspondingly, the Coast Guard structured Deepwater to produce the best “system of systems” for successfully accomplishing its projected range of future missions.

Under the program, the Coast Guard provided funding to three industry teams to produce system designs, implementation plans, and supporting analysis for the future Integrated Deepwater System. Each team was given a common, detailed set of reference material on current and projected mission attributes and mission demands. Each was asked to design the system to best meet projected demands over the for a forty-year implementation and support period from 2002 through 2042. The Coast Guard further gave each team wide latitude with respect to the types and numbers of air, sea, and support assets and infrastructure to include.

Each of the three industry teams was also permitted to propose changes to existing Coast Guard assets and infrastructure. Indeed, each team was required to specify how existing Coast Guard cutters, aircraft, air stations, ship stations, support facilities, and other infrastructure would be altered in the implementation of their system design. Deepwater required each to specify detailed plans for when and how certain assets would be upgraded or refitted with new technology, as appropriate, how new technologies would be implemented, how technologies would be assimilated through new training and logistics support in accordance with the implementation timetable, and a broad variety of other details.

Deepwater represented a particular watershed in systems analysis and systems engineering, in that, never before has industry had the opportunity and latitude to deliberately design and optimize a system with such broad impact on the safety, security, and livelihood of Americans. The missions supported by the U.S. Coast Guard Integrated Deepwater System include operations in support of U.S national defense, search and rescue, drug enforcement, alien migrant interdiction and rescue, enforcement of fisheries laws and international treaties, foreign vessel inspection, international ice patrol, monitoring of the commercial offloading of petroleum products, and numerous others. Moreover, even when Department of Defense procurements are included, industry has never before had the opportunity to design and implement a system of this magnitude for the entire duration of its operational lifecycle. Specifically, while various contractors have separately designed and fielded individual classes of cutters, aircraft, logistics systems, and C4ISR systems, no effort prior to Deepwater has given contractor teams had the opportunity to design, build, field, and provide support for this entire range of assets, at once, as a system of systems.

Of the three industry teams funded by the Coast Guard under the Deepwater program, one team was led by Science Applications International Corporation (SAIC), one was led by Lockheed-Martin Corporation, and one was led by Avondale shipyard.

From the outset of the program, the SAIC-led team chose to place the highest-possible quality system analysis and systems engineering at the core of its program effort. During the proposal phase of the project, the SAIC-led team conducted a broad internal survey of advanced modeling tools, techniques, and methodologies. The team chose to use a system dynamics-based methodology as its primary instrument for high level system analysis, evaluation, and design choice.

During the competitive procurement, the methodology was used on a daily basis by key decision-makers across the SAIC team and provided a series of concrete system and process references linking the team together. The methodology played a substantial role in the quantitative and qualitative discussions surrounding design decisions, implementation timetables, and the representation of these analyses to the Coast Guard. The balance of this paper describes this system dynamics-based methodology and how it was used in the Deepwater program by the SAIC-led Deepwater program team.

The Central Positioning of the System Dynamics Methodology in the SAIC Team

System dynamics methodology played a central role in the systems engineering effort of the SAIC-led Deepwater project. This centrality was due, in part, to the significant level of support given to system dynamics modeling and analysis by all levels of program management. This support was reflected not only in words and resources, but also in the placement of system dynamics trained analysts in key project positions, and in the manner by which the methodology was woven throughout the entire systems engineering effort.

From the outset, the SAIC-led industry team recognized system engineering and system analysis as a key component in its Deepwater project effort. This recognition was reflected not only in its methodological choices, but also in its project organization and the structure of meetings.

The SAIC-led team represented a heterogeneous collection of contractors with expertise and manufacturing capability in key areas relevant to Coast Guard needs. These project partners included the shipbuilders Bath Iron Works and Marinette Marine Corporation, the aircraft manufacturers Sikorsky and CASA, the Naval Engineering firm Gibbs & Cox, the logistics services corporation AMSEC, the C4ISR contractors Rockwell-Collins and Fuentes Systems, and the domain experts SOZA & Company and Clark-Atlanta University.

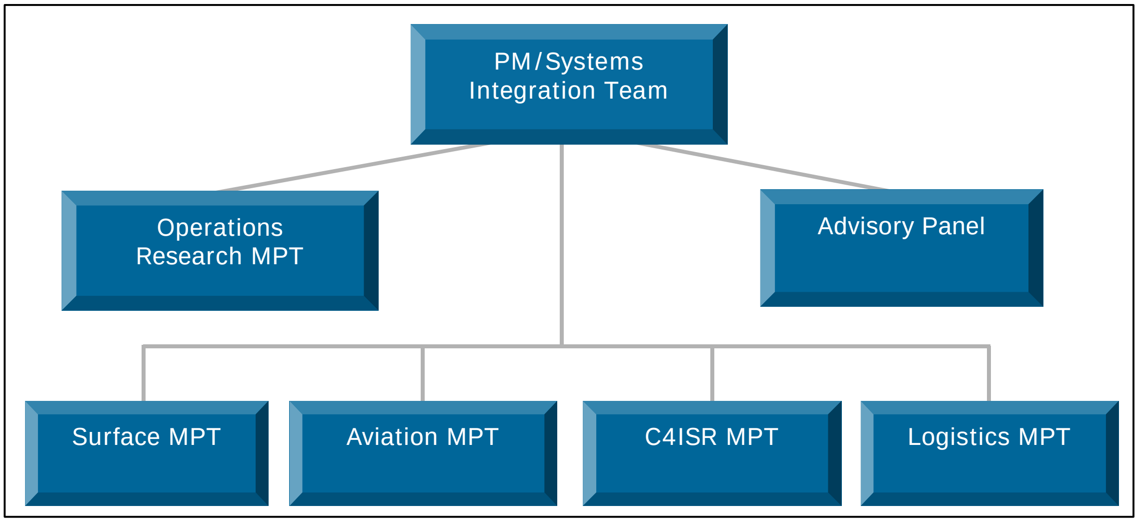

Figure 1

The inter-company project team was organized functionally as a series of Matrix Product Teams (MPT), with multiple companies represented on each functional team as appropriate to their domain of expertise. As depicted in Figure 1, the "Program Management" or "Systems Integration Team" coordinated the project. This team was supported by a group of senior advisors, including a number of retired Coast Guard officers, and an Operations Research team. Detailed design work on the project was performed by four functionally oriented MPTs under the systems integration team (SIT): the surface MPT, the aviation MPT, the C4ISR MPT, and the logistics MPT. The head of each MPT was represented on the systems integration team and provided direct liaison to it.

With respect to system engineering, all high-level activities were coordinated by the system integration team. These included the tracking of requirements, the resolution of interface issues between categories of assets, and final decisions on the allocation of specific capabilities to meet mission demands.

Within the system engineering effort, all system-level modeling and quantitative analysis, including the detailed implementation of the system dynamics-based design methodology, was performed by the operations research team. The operations research team obtained data from and provided analytical support to each of the other functional teams. At the same time, it acted as the analytical arm of the systems integration team. The operations research team not only performed analyses at the request of the systems integration team, but also was co-located with the systems integration team, had representation on the systems integration team, and discussed issues daily with it. The central role of the operations research team, and its special relationship with the key project decision-makers on the system integration team was a critical enabler of the central role that the system dynamics-based methodology played in the Deepwater system engineering effort of the SAIC-led team.

The System Dynamics-Based Methodology of the IDS Project Team

From the beginning of the project, the use of system dynamics in the Deepwater system engineering effort required a combination of influence diagramming, explicit modeling at multiple levels of analysis, and other techniques. While the system dynamics-based methodology employed was consistent with the "phased approach" advocated by James Lyneis and others, [1] the specific and detailed analytical requirements of the project required the team to employ a number of methods atypical of traditional system dynamics.[2]

---

[1] The general approach followed was the "phased approach," as discussed by James Lyneis. Se "System dynamics for business strategy: a phased approach." *System Dynamics Review*. Vol. 15, No. 1. Spring 1999. pp. 37-70. Intellectual grounding for the integration of multiple levels of influence diagramming and modeling can also be found in Geoff Coyle, "The practice of system dynamics: milestones, lessons and ideas from 30 years of experience." *System Dynamics Review*. Vol. 14, No. 4. Winter 1998. pp. 343-366.

[2] These requirements were largely driven by three sources. First, the Deepwater contract itself called for a number of very analytically detailed deliverables. Second, the system integration team defined the system engineering effort in a manner that required specific and detailed sensitivity analysis and analytical support

The system dynamics-based methodology used by the team represented a pragmatic adaptation of influence diagramming, modeling, and other system-dynamics based techniques to meet team needs as those needs unfolded over the two plus years of the program. The system integration team was open to apply system dynamics to new problems and questions as they arose. Reciprocally, the operations research team was willing to adopt its approach to the level of analytical detail and level of grounding in empirical evidence mandated by the issue at hand. Over time, this interplay resulted in the creation of a number of system dynamics models and influence diagrams at varying levels of detail. Some of these models influenced key decisions through the manner in which they illustrated fundamental structures of interaction, such as operational cycles. Others influenced decisions by providing quantitative feedback for contemplated system options, and by allowing the reasons for differential outcomes to be readily traced. Finally, some made use of advanced features of current generation system dynamics software, including optimization and sensitivity analysis, to allow real-time, expert-in-the-loop tradeoff analysis.

Over the course of the Deepwater program, the operations research team produced twelve models that were ultimately included in a report in some fashion, plus countless influence diagrams. The smallest of these models had fewer than ten variables, and was made and used over the course of several days. The largest of these models had in excess of 5.5 million variables (when arrays are considered), took over a year to fully implement and validate, and was used for a variety of advanced trade studies.

The system dynamics modeling choices made by the operations research team were impacted by its position and responsibilities in the project as a whole. The operations research team was simultaneously supporting the relatively broad and generalized analytical needs of the system integration team, while at the same time, attempting to provide technically credible analysis to detailed design engineers on the other teams. In some cases, this mandated separate models to accommodate each user. In other cases, it required a model both broad and detailed to accommodate multiple users while satisfying the analytical requirements for engineering-level detail in each of their disciplines.

In the opening days of the Deepwater program, the operations research team employed influence diagramming and relatively simple models to explore fundamental dynamics surrounding how the Coast Guard did, and could, perform its missions. An early influence diagram, for example, served as a basis of team discussion for understanding and optimizing the fundamental Coast Guard operational cycle.

As depicted in Figure 2, the dominant feedback loop in the diagram is the sequence “Detect, Classify, Respond, Recover,” corresponding to the traditional Coast Guard framework for thinking about mission performance. Negative feedback loops from “Respond” through “Classify,” and from “Respond” through “Detect” to “Classify” capture the effects of mission performance on subsequent activity. Minor negative

---

for particular targeted questions. Third, each of the asset-centered teams had specific questions and ways of thinking about the problem that had to be accommodated.

feedback loops from “Respond” directly to “Classify” and from “Respond” to “Detect” then to “Classify” captures how committing assets to certain missions may make them unavailable to perform patrol surveillance functions in support of others.

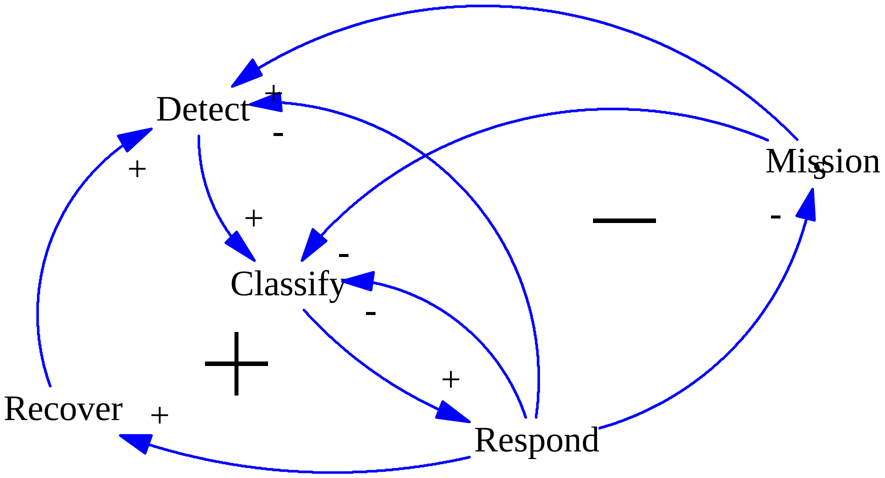

Figure 2 – Basic Operational Cycle

This diagram and the surrounding discussion, and a series of more detailed diagrams derived from it, paved the way for team consideration of how various types of assets could be understood as contributing to the design as a whole. The operational cycle was analyzed with respect to the particulars of each of the Deepwater Coast Guard missions. The particulars of the SAR mission were, for example, quite different from the particulars of drug interdiction, yet each could be understood within the context of the basic operational cycle. The team also examined the operational cycle from the perspective of each category of asset in the Deepwater system: C4ISR, logistics assets, aviation, and surface assets. This multi-perspective analysis of the operational cycle ultimately led the team to identify fundamental mission and system drivers and constraints that served as the framework for subsequent system discussions.

Even in the early days of Deepwater, modeling was performed at multiple levels of detail. One relatively small model, for example, was done in coordination with the surface asset team, and focused on the operational availability of Coast Guard cutters. The discussion surrounding the model focused the team on the relationship between physical and doctrinal determinants of cutter availability. The application of the model highlighted a layered system of availability bottlenecks, and thus impacted system-engineering decision-making early in the design process. The traditional approach to increasing cutter availability involved increasing storage space for potable water, perishable goods, and other supplies. The model led to the discovery, however, that the impact of such design actions would be limited, at a certain point, by crew-based needs to return to port, including leave and training. It is significant to note that this insight did not come from the isolated execution of a “validated” model, but rather, to experienced operators meaningfully interacting with, and objecting to, a straw-man engineering-based representation of cutter availability.

The largest of the Deepwater models began its life early in the Deepwater program, and came to play a variety of roles. The decision to construct the detailed model, largely in contradiction to the traditional practices of system dynamics,[3] was based on a widely-shared sentiment among team members that key attributes of system performance could only be validly addressed through a discrete simulation of individual assets, with individual capabilities, performing specific missions, in specific contexts (eg. visibility, sea state, size and speed of target vessel, etc.). Moreover, the cannon of complexity theory suggested that the interaction between individual effects could significantly impact system-level outcomes.[4] Certain system-level behavioral characteristics of importance to the team could not be reliably established by aggregating the results of detailed simulations and representing them as general tendencies in a system level simulation. The detailed model thus became a key vehicle for relating the impact of detailed technical design decisions in one area to system performance as a whole.

The detailed model also allowed the project team to obtain high-level insight into tradeoffs between performance and cost. By modeling to the asset level, the team was able to use the model to perform first-order activity-based costing of system alternatives. Specifically, the detailed model attached capital costs, personnel costs, and operations costs to each assets, including fuel consumption and maintenance costs according to the simulated mission use profile, as well as manpower costs according to simulated asset crewing requirements and facility staffing requirements.

In the process of providing these benefits, the detailed model broke new ground in the integration of system dynamics with techniques conventionally found only in object-oriented programming, agent-based modeling and discrete event simulation.

While the detailed model may have sacrificed a certain level of clarity in its system portrayal and ease of validation, it proved itself valuable on a number of fronts.[5] The

---

[3] Most of the discipline emphasizes the need to keep models relatively simple to maximize their analytical utility in understanding the object of study. The writings of Forrester (1961) and Coyle (1999) both make this point. Perhaps the only prominent, albeit qualified, exception in the literature can be found in Lyneis (1999).

[4] See, for example, Joshua M. Epstein and Robert Axtell, *Growing Artificial Societies*. Washington D.C. Brookings Institution Press, 1996. See also Murray Gell-Mann. "The Simple and the Complex." In *Complexity, Global Politics, and National Security.*, Ed. David S. Alberts and Thomas J. Czerwinski. National Defense University, June 1997. pp. 3-28. See also See also M. Mitchell Waldrop (1996). *Complexity: the Emerging Science At the Edge of Order and Chaos*. New York. Simon and Schuster, 1992.

[5] An excellent treatment of the value and inherent drawbacks of large system dynamics models can be found in Lyneis (1999). Lyneis notes that such models can be difficult to understand, and may tend to be treated as "black boxes" because of their complexity. At the same time, he notes that large models perform many valuable roles. Some of these include (1) providing assurance that all structural elements involved in producing principal dynamic behaviors are included, (2) supporting accurate cost-benefit analysis, (3) providing enough detail to permit analysis at the level that real decisions are executed, and (4) building confidence in the simulation results through the inclusion of detail.

model helped the team to establish that, while certain design decisions could exert small impacts on system performance as a whole, over short time periods, these decisions might often go unrecognized due to the stochasticity of mission demand, asset location, weather, and other considerations. Furthermore, the detailed model exposed counterintuitive relationships that would, arguably, have been overlooked by a higher-level model without a priori specification of the phenomenon.

The detailed model showed, for example, that increasing the range or endurance of assets would not always increase system performance, and in many cases would decrease it. Increased endurance, as reflected in the simulation, caused assets to accept missions, on average, at greater distances from their original position. Particularly when the system was already saturated by mission demand, increasing asset range increased the relative amount of time per mission dedicated to transit, and thus decreased the number of missions that could be performed per asset. More importantly, the representation of the details of each operation allowed the operations research team to establish that the number of missions in which the extra range made the difference between success and failure was insignificant, compared to the net increase in time per mission.

The detailed model also allowed the team to establish the relative lack of importance of cutter speed to asset performance. The model allowed team subject matter experts to see that, where response time was critical to mission success, typically only a preliminary airborne response could make a difference. Where a preliminary airborne response was not required, no reasonably contemplated increase in cutter speed made a substantial difference. It was particularly interesting that these results proved counterintuitive to current and former Coast Guard personnel due to an "instrument bias" effect in their own experience. Specifically, every cutter skipper could vividly recall a handful of situations in which he had been engaged in a chase with some target vessel in visual range, and wanted or needed more speed. The model suggested, however, that virtually the ***only*** time speed actually did make a difference to individual outcomes was when the cutter was engaged in a chase with the target object close enough to be in visual range. The persons focusing on the value of cutter speed based on visual chase incidents had not fully taken into consideration how few of their missions actually involved such pursuits.

Finally, the detailed model allowed the team to gain insight into cost-performance tradeoffs involving very dissimilar system options and complex interdependencies between cost drivers. The team used the detailed model, for example, to compare system alternatives with higher manpower costs to others with high fuel and maintenance costs to others with high capital costs. Some options that appeared to produce reasonable increases in system performance, for example, turned out to be unreasonable because they required retention of a large and costly logistics infrastructure to support a small number of assets. In other cases, the higher operational tempo from a downsized force led to added fuel and maintenance requirements that negated the capital and manpower savings achieved by the reduction.

In order to augment the team's understanding of what was occurring in the detailed model, simulation output was linked to an in-house GIS visualization tool developed specifically for the project. The tool allowed model users to observe the simulated behavior of cutters, helicopters, fixed wing aircraft, mission targets, commercial traffic, and other entities as they interacted in the simulation. Because of the complex and arrayed nature of the detailed model, watching the simulated behavior unfold in physical space was often a more effective way of understanding results at a preliminary level than tracing through causal trees and strip graphs. In understanding the relative lack of importance of cutter speed, for example, it was instructive to watch the replay of cutters "crawling" across hundreds of miles of ocean toward a mission objective, by contrast to helicopters and fixed wing aircraft "zipping" to the scene.

The detailed model and a number of other smaller models were used during most of the initial "concept design" phase of Deepwater to do preliminary tradeoffs on the impact on system effectiveness of including one asset type or another in the design. The design phase involved not only asset decisions, however, but also alternative tactics, doctrine, and concepts of operation associated with each design. In order to reduce the trade space to a manageable size for its preliminary analysis of alternatives, the system integration team chose to construct and evaluate seven candidate designs, each of which was associated with a regionally-specific concept of operation tailored to make best use of the assets included in the design in that region. The seven designs and their characteristics were decided during two special meetings of the system integration team. These meetings featured expanded representation from each of the asset teams, and operations research team. The straw-man system designs were selected not only to reasonably cover the trade space, but also with input from the operations research team concerning what was valid from an experimental design standpoint, and what could reasonably be modeled.

Each of the seven candidate designs was evaluated using the detailed model and several smaller models. Preliminary results from each of the models were presented and discussed in meetings of the operations research team, and subsequently the C4ISR, Logistics, Aviation, and Surface Teams. In each of these reviews, causes for simulated behaviors were examined through the system dynamics models and thoroughly discussed for their broader implications.

The interpretations and revised results from these working team meetings were presented to and discussed within a series of special sessions of the systems integration team. Each system integration team meeting was oriented around a series of design topics, with a discussion of relevant models and findings on any given topic as appropriate. The expanded system integration team meetings involved not only key decision makers, but also subject matter experts including an array of retired Coast Guard personnel from line officers to retired vice admirals with decades of experience on the topics being discussed. The discussion represented a moderated interplay between model insights and qualitative, experience-based commentary. The team used Groupware to capture the rich, multi-level discussion flowing from this interplay, and to poll the group on key points of decision following extended discussion periods.

Over the course of two months, these system integration team meetings employed moderated, model-augmented discussions to progress toward a series of increasingly specific decisions regarding system design and asset choices. It is important to note that the system dynamics models were almost never used in this process to adjudicate between options. Rather, the visual representations and quantitative feedback produced by the models provided a point of departure and key references for the expert discussion. Consistent with Forrester's writings, system dynamics provided a more formalized reference for discussing key interactions than the different "mental models" informally held by each of the subject matter experts.

The decision process and the supporting analysis from modeling results were written up in a series of contractually specified design reports provided to the Coast Guard under the project. A key intermediate product of this process was the specification by the team of its candidate design, including preliminary decisions concerning how many assets of what type would be used to meet the contractually specified mission demands. Each of the working teams then used this "Objective System" as a planning document around which to identify interface issues, to perform more detailed design work, and to conduct trade studies within their respective asset categories.

Following the specification of the objective system, the operations research team shifted its modeling focus to two new categories: detailed trade studies to tradeoffs within asset categories, and development of the implementation plan.

The operations research team worked with asset engineers and subject matter experts to support decisions such as "how much speed was appropriate for a cutter," "how much range was necessary for an aircraft," "what was the optimal basing scheme to meet mission demand," and "what was the best logistics arrangement to make maximum use of the assets in the objective system." In most of these cases, the team used a combination of the detailed model, adapted to address the question at hand, and a custom model, built over the course of several days to address the question.

In several cases, the team used a generic, mid-level model that came to be termed the "intercept" model. This model represented a compromise between the needs for detailed, system-wide discrete simulation, and the need to implement an experimental design that could rapidly explore the decision space. The "intercept" system dynamics model provided a discrete-event simulation of multiple assets executing their operational response cycle against a combination of stationary, time-critical targets and moving targets that had to be intercepted before reaching their objective. A key portion of the structure of the intercept model is shown in Figure 3 below. As with the detailed model, discrete individual assets, their associated states, and decision rules governing asset behavior were represented through model structure. Multiple assets were represented through arrays, with interaction between assets represented by further model structure.

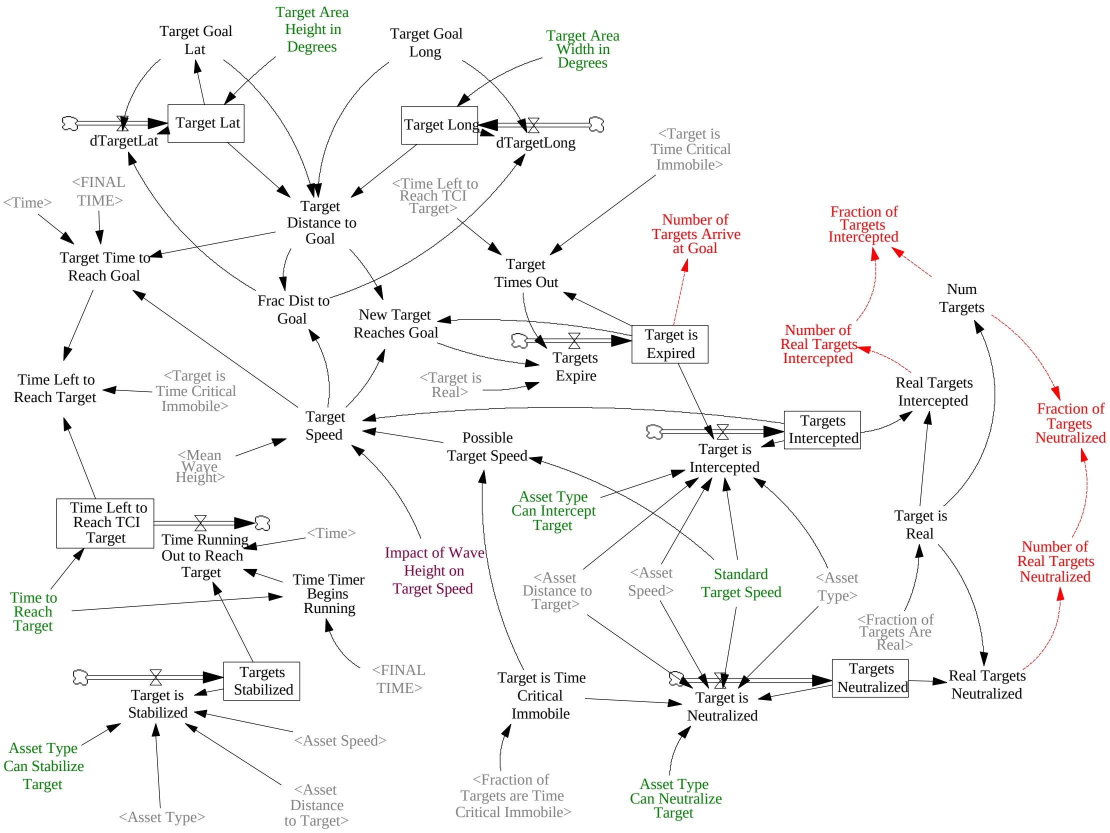

Figure 3 – Partial Representation of the Intercept Model

The team was able to adapt the intercept model to represent a variety of different missions and areas of operation to explore targeted issues such as the impact of C4ISR response time and helicopter endurance on mission performance.

In order to support the development of the implementation plan, the operations research team created a simple asset lifecycle model and linked to a complex fitness function for parametric optimization, as depicted in Figure 4. The lifecycle model was arrayed to include each of the categories of cutters, helicopters, aircraft, logistics facilities, C4ISR facilities, and other assets to exist in the SAIC-team design during the 40-year contract timeframe. In all cases, both new and legacy assets were included. Assets were constructed in the system according to the project budget and implementation plan. Assets aged and were upgraded or retired as called for by the plan. The simulation linked this simulated system to a measure of merit that rolled up values expressed by subject matter experts on the team. These values included not only performance attributes, but conformity with budgets, coordination between platform fielding and the maturity of associated C4ISR technologies, coordination between platform fielding and the availability of an associated logistics support infrastructure, and other items.

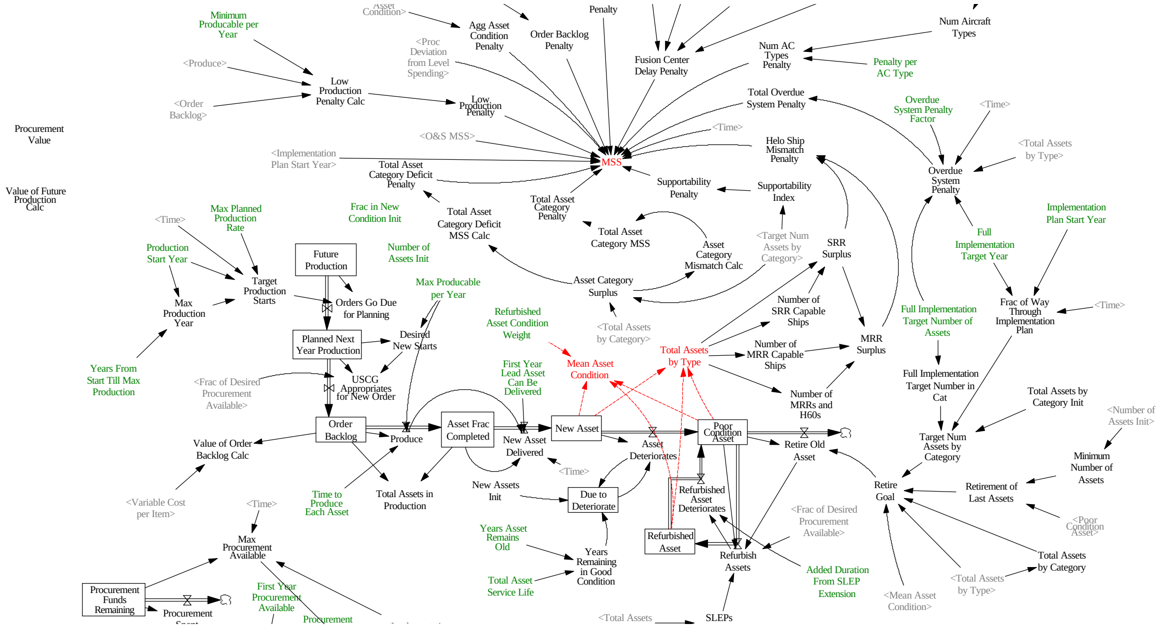

Figure 4 – Implementation Plan Model

As with the system dynamics models supporting the analysis-of-alternatives phase of the project, the team used the asset lifecycle model in an "expert-in-the-loop" workshop format to iteratively develop the most reasonable implementation plan. As in the analysis-of-alternatives phase, the asset lifecycle model was developed by the operations research team, reviewed with each of the asset teams, and finally reviewed and applied by the systems integration team in a series of special sessions.

As with the analysis-of-alternatives phase, the system integration team in the implementation-planning phase was augmented by experienced Coast Guard operators, as well as representatives of key asset contractors who could speak to issues of production capacities, timelines, and issues of technological maturity. In these sessions, the system integration team leveraged the capability of the Vensim™ system dynamics software to perform parametric optimization on the multi-component implementation plan measure of merit. The lifecycle model was run in optimization mode to produce parameters representing a candidate optimization plan. Subject matter experts would review the plan and provide feedback on key considerations that either were not included in the optimization, or appeared to be incorrectly weighted. Members from the operations research team then modified the fitness function in real time to reflect these new issues and re-performed the optimization. For this application, the relatively small size of the model was necessary to execute the optimization algorithm and achieve convergence on a candidate solution multiple times during the each workshop period.

Through expert feedback on each set of optimization results, the team iteratively constructed a fitness function over the course of the meeting that reasonably captured their values. Repeated parametric optimization of the lifecycle model with each new fitness function, in turn produced candidate implementation plans that reflected appropriate compromises between objective requirements and multiple, varied team member values. It is important to note that the final implementation plan produced by the parametric optimization exercise was not incorporated in toto as the implementation plan for the team. Rather, as with model output generated in the analysis-of-alternatives phase, the simulated optimal implementation plans served as the starting point for more detailed discussion and adjustment.

Methodological Issues Raised During the Deepwater Modeling Effort

The experience of the operations research team in Deepwater provides a number of important insights about using system-dynamics to support large system engineering projects. Those insights and associated issues are discussed in the balance of this paper.

Large, Detailed Models May Play a Necessary, Albeit Unconventional Role in Applied System Dynamics. The system dynamics literature, from Forrester through contemporary authors, emphasizes the importance of keeping models small. Simple models, the literature suggests, better support an understanding of the system because they more clearly depict key feedback processes and other behavior drivers.

While elegant simplicity is a laudable goal, the Deepwater project repeatedly demonstrated that large, detailed models might be required within a larger design effort because of the nature of the subject being modeled, and because of the analytical requirements of those being supported by the effort. To argue that system dynamics should not be used under circumstances in which the models cannot be kept simple is to ignore the multiple sources of potential value added from a system-dynamics-based analysis beyond simple, strategic level insights. Even when models become highly detailed, the visual orientation of the methodology enables an invaluable bridge between modelers and subject matter experts, and helps all involved to seat their technical issues in the context of a larger system. Moreover, whatever the model size, the diagnostic capabilities an flexibility of the modern PC-based tools that implement system dynamics permit quantitatively supported, expert-in-the-loop analysis simply not possible with other methodologies. Finally, while large models are certainly more difficult to verify and validate than smaller ones, the visual accessibility to subject matter experts and the causal tracing made possible by the software makes debugging and verification through expert review far easier in system dynamics than in models constructed using conventional programming tools.

System Dynamics Is Part of a Continuous Spectrum of Tools For Studying Complex, Highly Interdependent Systems. System dynamics, with its emphasis on small, continuous flow models, is rarely considered within the larger family of methodologies used to understand complex interdependent systems, particularly such traditionally object-oriented programming applications as discrete event simulation, GIS-based visualization, artificial life, and agent-based modeling. Although, for example, the Santa Fe Institute has begun to include some system dynamics work in its compendium of complex systems research, few references can be found within the system dynamics literature to complexity theory research.

At both the level of modeling and the level of philosophy, system dynamics and complexity research are intimately interconnected. In the “Detailed” and “Intercept” system dynamics models, for example, the continuous-process stock and flow representation of an individual cutter as a "factory" is recognizable as conventional system dynamics. When the representation of cutters and other assets are arrayed and made to interact with each other in each model, however, the system dynamics model begins to resemble an agent-based model, with the behavior of each entity determined by a series of simple rules, but their interaction producing highly complex system-level outcomes. Similarly, the arraying of relatively generic entity structures in the system dynamics model bears a striking resemblance to class objects in an object-oriented program. Stocks defining the cutter structure, such as "Fuel On Board" resembled class variables. Similarly, the equations resolving the values of each stock, flow, and auxiliary variable per time step resemble class methods in an object-oriented programming schema. Although system dynamics is widely seen as "continuous process simulation," the object of study in a system dynamics simulation may appear discrete from a broad enough perspective. In particular, system dynamics models with arrayed structures can be seen to represent discrete entities when the arrayed structure is considered as the unit of analysis.

Object-Oriented Programming May Be More Appropriate Than System Dynamics Where The Model Is Very Large And Is To Be Used For Black-Box Sensitivity Analysis. The detailed model produced under the Deepwater project arguably pushed the boundary for when it is appropriate to construct a system dynamics model, as opposed to constructing a conventional program. Models implemented using modern PC-based system dynamics software tend to be significantly faster to build, adapt, and de-bug than conventional programs. In part, this is because the system dynamics software circumvents much of the overhead required in conventional programming for building the logical structure of the program, constructing classes, and creating functionality to support data input, output, and behavioral diagnosis. System dynamics models tend to execute far slower than comparable object-oriented programs, however. In part, this is because system dynamics models typically execute in an interpreted form.[6] In part, it is because system dynamics models use a fixed structure of interaction. In a traditional program, if certain logical conditions are not met, entire segments of code can be bypassed. In a system dynamics model, by contrast, the equations specifying all of the variables must be solved at each time step. Thus, in the detailed model, for example, although an aircraft may be parked at an airbase for the entire length of the simulation, its speed, fuel consumption, desired course, sensor readings, mechanical failures, equipment on board, current mission assignment, and other attributes must be calculated at each time step.

As models become larger, the increased amount of time required per run for a system dynamics model becomes an issue, particularly if the models are primarily used for multi-run sensitivity analysis. Models that rely on stochastic behavior are particularly

---

[6] Modern PC-based system dynamics software typically performs a simulation by executing a series of instructions within the software, which then must be implemented by the computer executing the system dynamics software. By contrast, compiled programs are executed directly by the computer. The only current exception in system dynamics software is Vensim™, which allows its models to be compiled and executed directly be the computer.

problematic, because multiple runs are usually required per option being explored in order to obtain results from which valid statistical inferences can be made. The detailed model used by the SAIC-led Deepwater team was particularly challenging in this regard. Even running in a "compiled" mode,[7] the detailed model required 6-10 hours per simulation, with four or more simulations used per input combination. It is also instructive that, when the same model is used repeatedly over a long time frame, the ability to rapidly build and adapt a model, afforded by system dynamics software, is undercut. At the same time, when a model is treated as a black box for performing off-line sensitivity analysis, the ability of a system dynamics model to involve subject matter experts in the review of the model and the diagnosis of results is negated. For the Deepwater project team, the process of incrementally reviewing the detailed model during its construction and validation afforded the team multiple valuable insights. Approximately one year into the program, as the model evolved into a sensitivity analysis tool, a detailed object-oriented simulation would have served the team as well, but for all of the benefits that had accrued to that point.

Optimization is a Useful Venue for the Application of System Dynamics Software. The SAIC Deepwater project team enjoyed a great deal of success in its use of the Vensim™ System Dynamics Software for an iterative optimization drill. The same visual accessibility and software flexibility that allowed other system dynamics models to be adapted and discussed in an expert workshop context, allowed fitness functions to be adjusted and models re-optimized in real-time in a room of subject-matter experts. The system dynamics literature tends to downplay ways in which attributes of the software may facilitate the system learning and analysis process. The use of the Vensim™ optimization functionality in an iterative subject-matter expert workshop context, however, provided significant benefit by facilitating the construction of an implementation plan that simultaneously incorporated insights and constraints from multiple experts. Although different in appearance than traditional model-mediated interactions in system dynamics consultancy and pedagogy, the interactive use of system dynamics models for optimization is consistent with the approach of the discipline in providing qualitative insight through visual referents and quantitative feedback.

Conclusion

The U.S. Coast Guard Integrated Deepwater System program is unprecedented in its size, scope and nature. It represents a major challenge for system engineering, and a major opportunity for system dynamics as an analytical methodology. Perhaps never before has system dynamics been so extensively used, with such close collaboration of key decision makers, in a program of this magnitude. While the victor in the Deepwater design competition has not yet been selected, the SAIC-led team broadly acknowledges that system dynamics made major substantive contributions to its Deepwater system engineering effort.

---

[7] The Vensim™ system dynamics software contains an option to translate the ordered system dynamics equations and the functionality which executes them into C code, then compiles the C code so that the system dynamics model can be executed as a compiled program on the user's machine.

Although the core principles of the system dynamics methodology have been established for thirty years, system dynamics has principally remained a business process re-engineering tool.[8] The successful SAIC experience with system dynamics suggests a far broader range of applications. It also suggests, however, that the broader use of system dynamics may be contingent on a receptiveness to adapt the methodology to the exigencies of the situation, and a willingness to stretch boundaries between system dynamics and other methodologies for exploring complex systems.

---

[8] This is not intended to overlook the substantial volume of system dynamics work in the areas of sustainable development, health care, and the biological sciences. Whatever the importance of these works, however, they have arguably been eclipsed by the sheer volume of work in business strategic planning and management. It is instructive, for example, that the principle commercial practitioners of system dynamics are found in accounting and business consultancy firms, rather than among engineering organizations, product manufacturers, research and development corporations, or government agencies.