

The Art of Systems Science

Which fork in the road should you take? Systems scientists help provide the rational bases for answering that as well as questions involving incredible complexity.

A classic problem in systems science is embodied in the quandary faced by a traveling salesman: what is the best route to n cities that passes through each city just once? After a few dozen cities, finding that most efficient route gets extremely complicated. In fact, the computational effort increases exponentially as the number of cities increases. Finding a practical approach to solve that problem is one facet of systems science.

Similarly, n interacting systems—facilities, components, machines, processes, and people—are involved in operating the National Ignition Facility (NIF), managing a battlefield, or encapsulating plutonium waste. Systems science tools can also be used to optimize the performance of n interacting complex systems.

At Lawrence Livermore, over 20 systems scientists are working as in-house consultants, applying analytical techniques to projects large



and small. Often, their task is to organize and analyze data to facilitate informed decision making. The information they supply helps managers choose solutions that are the safest, most timely, most productive, or most cost-effective.

The systems scientists work in the Decision Sciences Group led by Tom Edmunds and the Systems Research Group led by Cyndee Annese. The groups, which overlap considerably in expertise and project involvement, are managed as a team by the two group leaders and Annette MacIntyre, a deputy division leader in the Engineering Directorate.

Group members have graduate degrees in statistics, operations research, physics, engineering, economics, and mathematics. They have expertise in decision analysis, computer science, industrial engineering, simulation modeling, and systems engineering. The problems they tackle may involve designing systems to operate in the most effective way; deciding how to allocate scarce human resources, money, equipment, or facilities; or assessing the risks of system options and operations. Systems scientists are typically involved in a variety of projects.

Plethora of Uses

For example, they perform reliability, availability, and maintainability (RAM) analyses look at the big picture of an overall operation and thus assure that the production goals of a facility can be met. The scientists review available design documents, solicit information from experts, talk to vendors and examine their catalogs, and study industrial failure rate data. They examine the interactions among various component systems. In the process, they identify how long it will take to make repairs, how many spare parts will be needed, and so on. Engineers can then incorporate these data in the

final design specifications of the facility.

RAM data can be fed into an operating model for the facility, such as a discrete-event simulation that can be used to optimize plant operation (see the box below for more information). It doesn't matter to the model whether the end product is laser shots, shoes, cars, or laundry detergent. The important thing for analyzing performance is that the data going into the model are the best available and that reasonable mathematical representations of the various processes are used. Says Edmunds, "There is an art to structuring and implementing models that can only be developed through experience." The same tools can be applied for a diversity of purposes and problems.

One goal of systems science is to quantify tradeoffs, according to mathematician Mike Axelrod in the Decision Sciences Group, because "you

can't always get everything you want at the same time." Quantifiable risks and uncertainties surround each goal. In the Yucca Mountain project, for example, the performance of proposed methods for isolating radioactive nuclear waste must be estimated for tens of thousands of years into the future. Inevitably, uncertainties abound for such long time frames. Systems scientists estimate these uncertainties in a formally traceable manner.

Systems scientists also respond to Department of Energy regulations that require periodic completion of hazard and safety analysis reports for each DOE facility. Experts in probabilistic risk assessment recently completed a number of these reports for buildings at Livermore. They analyzed likely risks and worked with building staff to define preventive and mitigating measures, using system safety tools. One tool was fault-tree analysis, which is deductive

Another Kind of Simulation

Systems science modeling often uses discrete-event simulation, which looks at processes (such as manufacturing) that are made up of a series of individual events. In contrast, simulations of physical and chemical processes examine the state of a process at every instant in time. Examples of the latter include continuous simulations of fluid dynamics and high-explosive detonations. Such simulations comprise the majority of modeling work at Livermore.

The time between events in a discrete-event simulation may be a few seconds or many hours. What matters to the simulation—and what triggers a new event—is a change in some part of the system. Simulations may be run for laser system operations, manufacturing systems, economic markets, or the entire nuclear weapons complex. Analyses seek to predict the feasibility of a design or the performance of existing or proposed systems before their implementation. Design analyses and tradeoff studies can be performed inexpensively via simulation, and the practical implications of proposed systems can be identified and examined. Experiments with simulation models enable engineers to test ideas and proposed operating policies and to suggest alternatives. Key issues typically include throughput, resource utilization, reliability, availability, maintainability, and scheduling.

A large number of simulations have been run to plan the operation of the National Ignition Facility. What happens in terms of downtime when one or more of NIF's 192 lasers is off line? How frequently will parts need to be changed out? What happens when spare parts are not available as planned? Which spare parts matter the most to avoid downtime? What are staffing requirements for the system? These are the kinds of parameters that can be understood and optimized using a discrete-event simulation model.

and helps analyze problems from the top down to evaluate many modes of failure both qualitatively and probabilistically.

Systems science statisticians also developed a model for forecasting electric energy demand at the Laboratory. This forecast must be accurate because Livermore contracts for bulk electric power. Usage over or under the forecast results in extra costs.

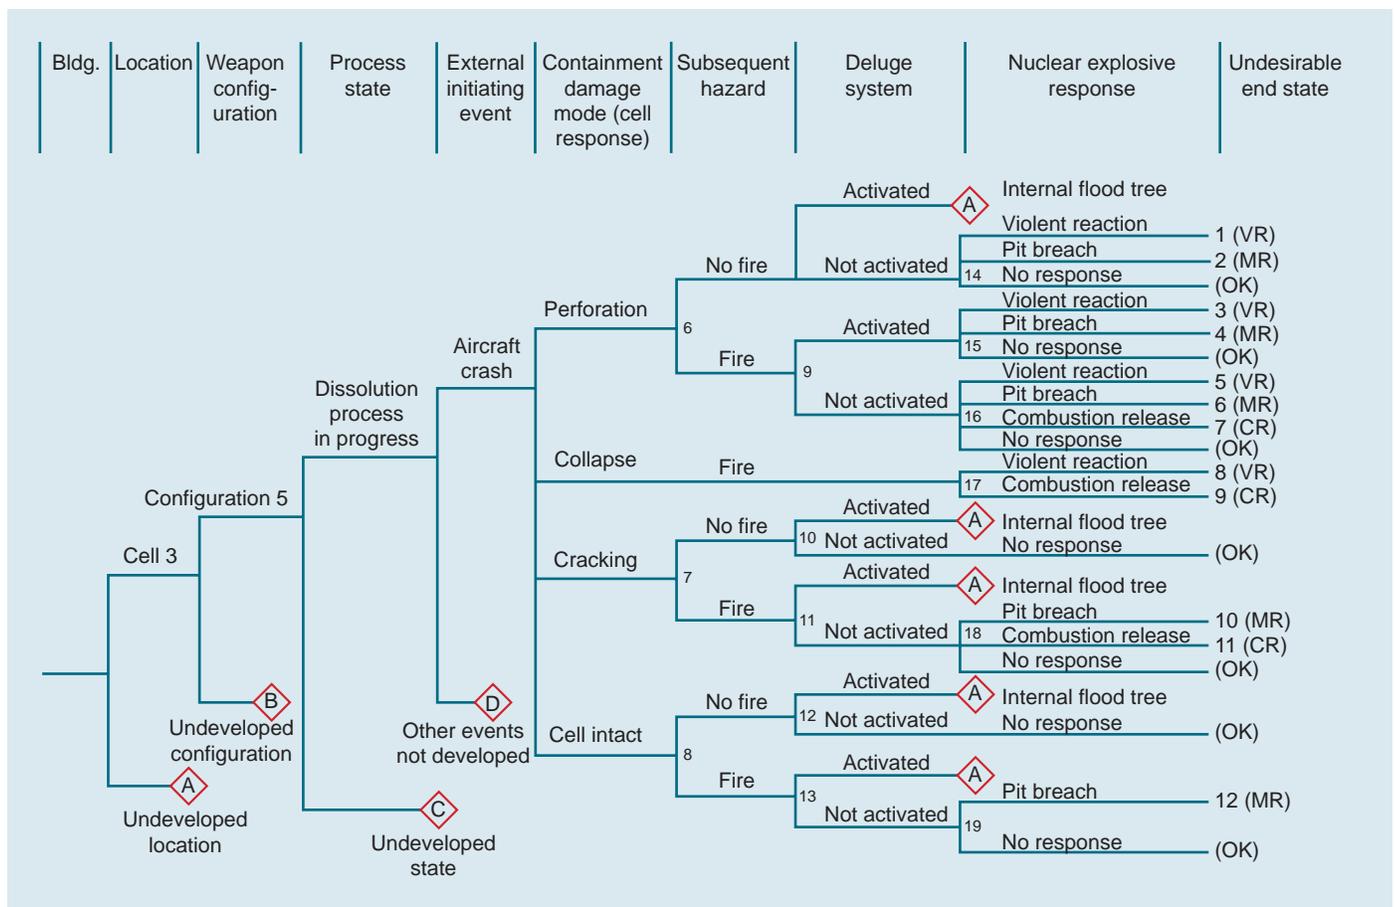
Livermore statisticians have developed a statistical sampling method for property management at Livermore. DOE requires a periodic audit of the records of more than 50,000 items in capital and attractive property inventory. The statisticians used statistical

sampling theory to satisfy the DOE requirements without having to check every record. They reduced sampling costs by 90 percent, well worth the investment. The sampling method can be used for years to come, and DOE is considering applying this technique to property management throughout the DOE complex.

Systems scientists have assessed the risk associated with transporting spent nuclear fuel. They looked at a wide range of accidents and developed a probability distribution of the radiation that might be released from each one. For another project, they analyzed the hazards of assembling and

disassembling nuclear weapons, developing a systematic and traceable method to assess the hazards from each part of the operation. This project required a marriage of traditional hazard-risk analysis and time-and-motion studies. To define appropriate controls, the team used database techniques, talked to experts in the field, developed simulations of the processes, and used specialized statistical methods that handle sparse data, rare events, and uncertainty.

Because of the diverse ways that systems science can be applied to Laboratory projects, systems scientists are involved in numerous projects at



An event tree developed as part of a probabilistic risk assessment for the processing of a nuclear weapon. This part of the tree shows what might happen if an aircraft were to crash into the building while processing was under way.

any given time. A few of the larger ones are discussed here.

Analysis for NIF

The National Ignition Facility (NIF) is a cornerstone of the Department of Energy’s Stockpile Stewardship Program to assure the safety and reliability of our nuclear arsenal without nuclear testing. It is also the largest single project under way at Livermore. NIF’s 192 laser beams will be used to heat and compress a tiny deuterium–tritium fuel pellet with the goal of achieving thermonuclear ignition for the first time in a laboratory.

The NIF Conceptual Design Report specifies a goal of 90-percent facility availability for operating three shifts round-the-clock, that is, having no more than 10-percent downtime (lost opportunities) due to unscheduled maintenance and unexpected equipment failure. The goal for laser system reliability is 80 percent. To meet those goals, design engineers must know what the unavailability and unreliability allowances are for each NIF system. They need answers to such questions as: If the quality of a component is lower to save money, what effects will the reduced quality have on reliability and availability? Where would redundancies improve performance and be most cost effective? What preventive measures can be incorporated into the design to reduce maintenance and downtime?

Because of NIF’s enormous size and importance, systems scientists approached this RAM study in two steps. They first allocated unavailability and unreliability from the top down, looking at the interaction of major functions to determine where the most downtime might occur. Then from the bottom up, they assessed individual component performance to determine overall system performance. The bottom-up estimates were used to

	Availability		Reliability	
	Estimated lost opportunities		Estimated ruined shots	
Conventional facilities	0.9915	(8.1%)	0.9985	(0.6%)
Laser system	0.9894	(10.1%)	0.8878	(53.3%)
Beam transport	0.9956	(4.2%)	0.9852	(6.7%)
Computer control	0.9448	(54.0%)	0.9941	(2.7%)
Optical components	0.9911	(8.5%)	0.9852	(6.7%)
Laser control	0.9920	(7.6%)	0.9493	(23.3%)
Target area	0.9921	(7.5%)	0.9852	(6.7%)
Goals	0.90		0.80	

The first-cut allocation of availability and reliability among the eight major systems at Livermore’s National Ignition Facility. The goal was 90-percent availability and 80-percent reliability.

monitor how evolving designs were addressing RAM performance goals and to check on the plausibility of the top-down allocations. As design of NIF progressed, designs were modified and allocations were iterated to meet system performance goals. If goals could not be achieved, availability could be increased by on-site backups, on-hand spares, and on-call maintenance. The table above shows the breakout of how availability and reliability were allocated throughout the major NIF systems in the first-cut allocation.

Because no laser facility truly compares to NIF, Bayesian statistical techniques were used to help handle the problem of sparse data and uncertainty. Systems scientist Alan Sicherman, who headed the top-down study, was pleased to note, “The top-down and bottom-up assessments were close to each other, which was pretty remarkable.”

To integrate the RAM studies over all systems and to assess whether the total facility goals could be met, Laboratory scientists developed a discrete-event simulation model known as NIFSim. The NIFSim model used subsystem availability estimates from the RAM assessment as input to produce estimates of the availability of NIF as a function of the failure rates of its components and subsystems under various scenarios. In addition to failure-

rate data, the NIFSim model reflected planned maintenance, personnel allocations, and operating modes such as all 192 or fewer beams, shots every 4 or every 8 hours, and periodic high-power (2- to 20-megajoule) shots.

NIFSim was used to check plans for facility operations to maintain a desired shot schedule. It showed what happens under different assumptions, such as performing maintenance only on planned maintenance days or having spare parts readily available. NIFSim predicted that at the completion of final design, NIF can be available for approximately 766 shots per year for an 8-hour shot cycle. For a 4-hour shot cycle, 1,360 shots are possible. The importance of spares availability and planned maintenance was illustrated by the simulation showing that if no spare parts are available, only 15 shots per year are possible.

NIFSim can be customized to evaluate various scenarios of shot schedules, spare part availability, maintenance deferral, staffing levels, and operating practices related to unexpected equipment failures.

Systems scientists are helping to plan for the production of certain spare parts known as line replaceable units, or LRUs, for NIF. These are limited-life components and include some laser mirrors, spatial filter lenses,

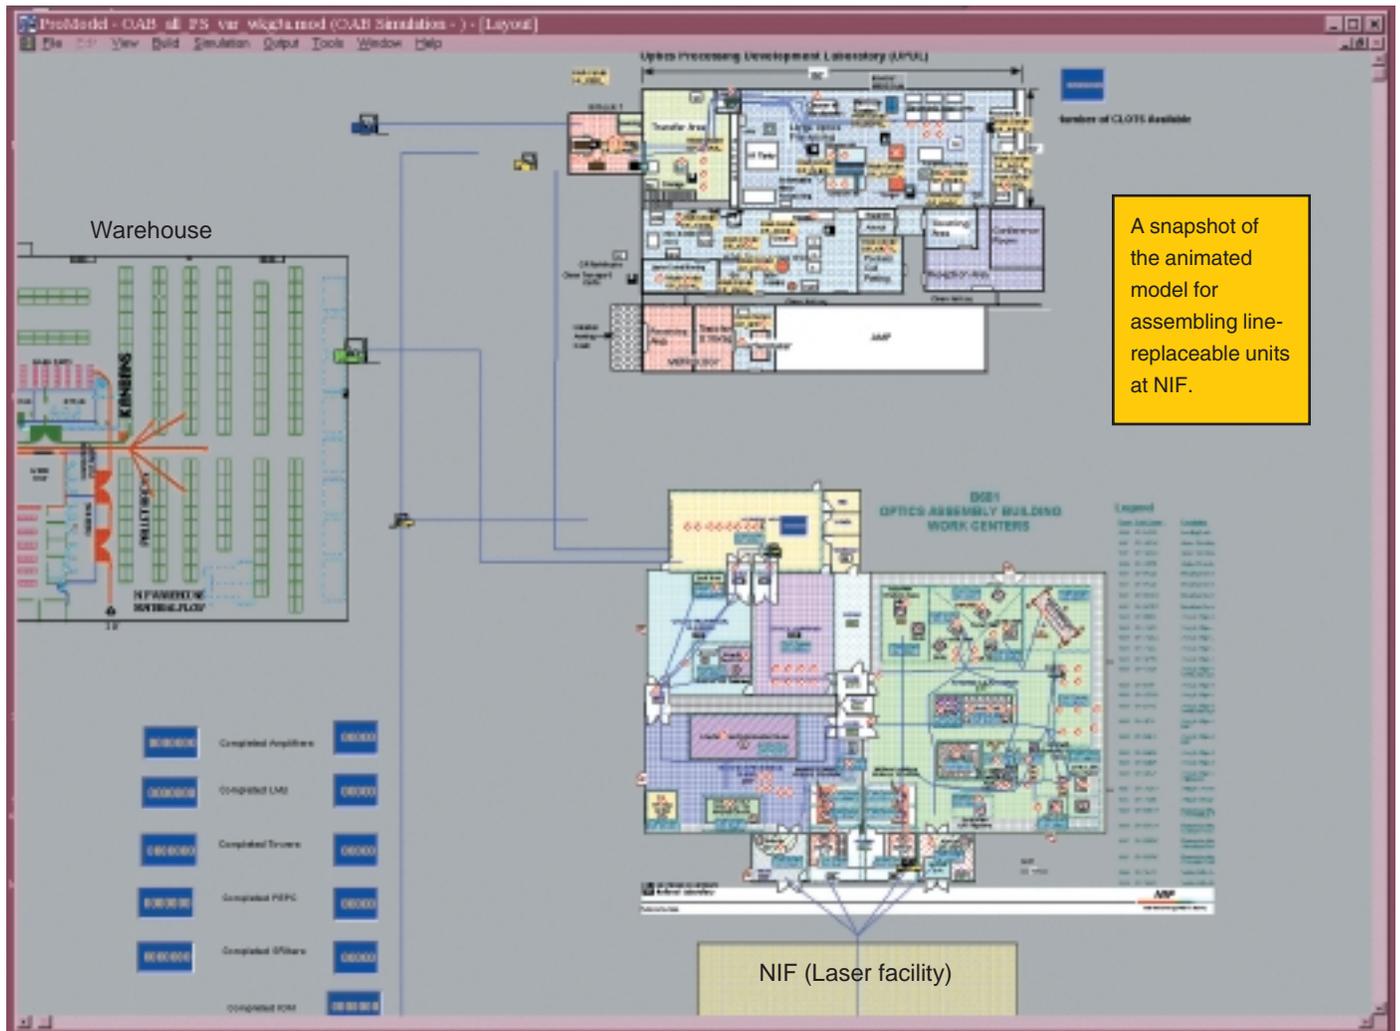
transport spatial filter diagnostic and alignment towers, cavity spatial filter towers, amplifier slab cassettes, plasma electrode Pockels cells, and flashlamp cassettes. Assembly of many LRUs from smaller components will take place at Livermore. Many of the LRUs are as large as a refrigerator and cannot be stored easily. Instead, they will be manufactured “just in time.”

For just-in-time manufacturing to work effectively, planners must work backward from the time new parts are to be installed during the first five years as

well as from the start time of later refurbishments. What’s more, they must consider random failures so they can minimize operational downtime. The model for assembly, installation, and refurbishment of NIF LRUs analyzes and verifies production capabilities under varying constraints. It is a decision-analysis tool for identifying risk, allocating resources, scheduling staff, and assuring that the activation schedule can be met within given time, resource, and budget constraints. As a cost- and time-saving device, perhaps its most important use is for estimating

where resources may be most strategically allocated.

Systems scientists on another NIF project developed a statistical method for measuring “how clean is clean” on metallic surfaces in proximity to optics. The intense energy of the laser beam can drive microscopic particles from laser vessel surfaces onto nearby optics. As the laser beam passes through the optics, energy is deposited in any dust and dirt particles that are present, thus damaging the optics by punching microscopic holes in their surfaces. A cleanliness level for optics has been



established, but to check every optical surface is not practical. The challenge, then, is to establish a process for performing random checks that ensure, to an acceptable level of confidence, the cleanliness of all vessel surfaces interfacing with optics.

Finally, systems scientists have simulated the process for manufacturing KDP (potassium dihydrogen phosphate) optics for NIF. In modeling the manufacturing process, the scientists presented questions and facts that optics producers had to consider to improve efficiency. The developed model is now used for estimating production (how many finished crystals can be produced per month) and for determining process bottlenecks. This information is then used in making decisions about process parameters such as the number of machines of each type needed, the number of shifts for each machine, or whether more than one vendor at a time should be producing optics for NIF to meet its installation and maintenance goals.

Forecasts for the Stockpile

Recently, systems scientists completed development of decision aids to forecast the reliability of the nuclear stockpile and select an optimal set of science-based activities (stockpile surveillance, physics experiments, weapon simulations, and plant production) to enhance confidence in the future performance of the stockpile. Using reliability and process modeling as well as discrete-event simulation and multiattribute utility theory, Livermore systems scientists developed a forecast framework that projects the status of various weapon systems into the future. As shown in the figure below, future reliability depends on the activities that the DOE invests in today.

Systems scientists also developed a database and simulation model

framework of facilities and associated capabilities throughout the DOE weapons complex. The framework focused on readiness and technical base facilities that will be used to develop the new generation of weapons codes and physical experiments to validate the codes.

Although the model still requires more detailed data to realize its full potential, the framework is designed to address issues such as: How well can experimental throughput capacity address the goals and timelines of the overall research campaign? Is available expertise adequate for experimental design, execution, and data interpretation? Is there sufficient funding for ongoing experimentation as well as for the construction and maintenance of facilities? Is the Laboratory hiring and training the right disciplines to maintain design, experimentation, code development, theory, and refurbishment capabilities to continue stockpile certification?

Modeling Waste Processing

Discrete-event simulation modeling is helping to find the best method for processing 13 metric tons of plutonium that are no longer needed for national defense. Livermore and several partners are developing a ceramic material in which the plutonium will be immobilized, a production process for fabricating the material, and a process for placing the finished ceramic-plutonium "pucks" in waste canisters. The plutonium comes from several sources with varying levels of impurities. It must be blended to evenly dilute the impurities before the pucks are fabricated. The challenge is to accomplish the blending operation with as few reblands as possible to minimize handling costs and personnel exposures.

Systems scientists developed an impurity blending model, whose logic

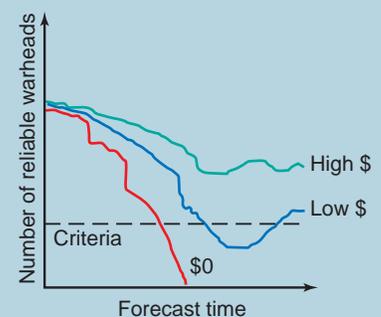
Problem: Forecasting stockpile reliability

The nuclear weapons stockpile must continually be certified as safe and reliable, a challenging task in the absence of nuclear testing. The problem is to select an optimal set of science-based investments (in stockpile surveillance, physics experiments, weapon simulations, and plant production) that maximize confidence in the near- and far-future performance of the stockpile.

Approach: Reliability, process, and cost modeling and discrete-event simulation

This effort requires modeling nuclear weapons complex enterprises in three fundamental areas: (1) determining how weapons expertise, physics experimentation, and computer simulation capability will affect the future ability to detect and assess problems and to design and certify remedial actions and certify the results; (2) modeling weapons production facilities, capacities, and costs; and (3) assessing past and predicting future failure types and rates in the stockpile.

Result: Forecasts weapon system and investment decision scenarios



Systems scientists forecast the future reliability of various weapon systems as a function of investments in science-based activities.

flow is shown in the figure below. The model used Monte Carlo simulation techniques to account for uncertainty in the impurity content of incoming cans of plutonium. With this model, the team developed feed-stream schedules and material staging methods that minimize reblend requirements for 13 feed and plant design cases. The model provides the capability for rapid evaluation of new feed scenarios and operating constraints

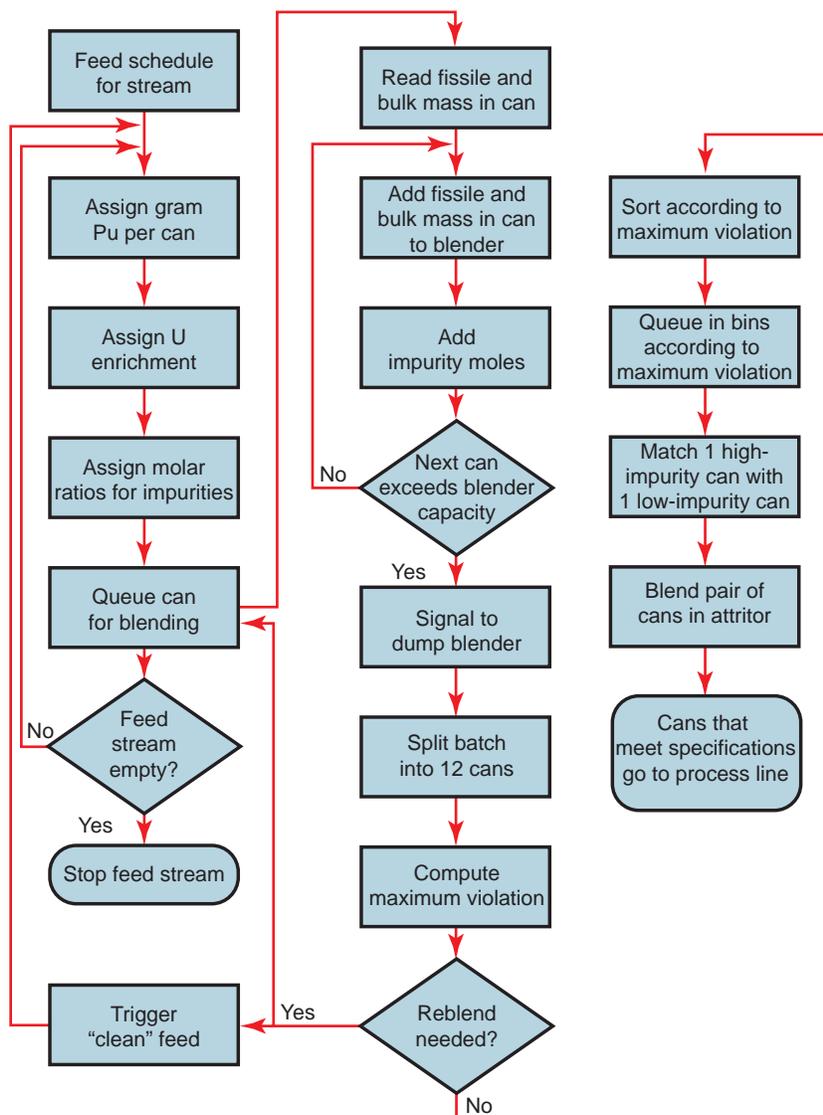
as they arise, including variations in isotopic mixtures.

Buying Uranium from Russia

Systems science data analysis methods have found their way into a major DOE nonproliferation project that is monitoring the purchase of highly enriched uranium from the Russians. In several formerly closed cities in Russia, highly enriched uranium is

being removed from nuclear weapons, processed to reduce its enrichment level, and shipped to the United States where it is used for fuel in nuclear reactors. Personnel from several DOE laboratories spend from a few weeks to several months at a time in Russia monitoring these processes. Monitors check inventory, measure the reduced level of enrichment, and seal canisters that are sent to the U.S. Monitoring began in 1995 and is expected to continue until 2013.

Currently, four Russian sites require monitoring and typically, six monitoring visits are made to each site per year. From five to ten monitors participate in each of those 24 visits every year, producing reports of what they saw. The Russians also supply records for in-plant processing, interplant transfers, and shipping. They send about 8,000 pages of data per year—all requiring sifting and analyzing for holes or inconsistencies. In 1998, Livermore was assigned to be the repository for these data. A relational database keeps the information organized. From Livermore, the monitoring information travels via a secure network to analysts around the U.S. who examine it and prepare reports for DOE. Together, the Livermore team makes recommendations to DOE about



A logic flow of the simulation model for processing surplus plutonium. Plutonium from several sources has varying levels of impurities that must be diluted evenly. The goal is to perform as few reblends as possible to minimize handling costs and personnel exposures.



Sealed canisters of uranium shipped from Russia.

the effectiveness of the visit and prepares a new set of instructions for the next group of monitors.

Consistency tests for the data include developing tags and seals for chain-of-custody observations, standardizing monitoring forms and trip reports, performing mathematical and statistical analyses, reconciling apparent inconsistencies by seeking Russian clarifications of anomalous observations, and performing a final assessment.

On the Move

Systems science is applicable to a wide range of problems, so systems scientists at Livermore move from one project to another very different one on a regular basis. Whether working on seismic discrimination for treaty verification, allocating resources for contaminant cleanup at Livermore's experimental test site, or developing a model of energy use in China, they take their kit of decision analysis tools to help solve problems.

—Katie Walter

Key Words: decision analysis; discrete-event simulation; National Ignition Facility (NIF); nonproliferation; operations research; reliability, availability, maintainability (RAM); plutonium disposition; risk analysis; statistics; Stockpile Stewardship Program; systems engineering.

For further information contact

Cynthia Annese (925) 422-0264

(annese1@llnl.gov) or Annette MacIntyre

(925) 423-7254 (macintyre1@llnl.gov).

About the Scientists



CYNTHIA ANNESE is group leader of the System Research Group, a systems science team at Livermore that focuses on innovative analysis of information, technical and safety risk assessments, and integrated simulation modeling. She received a B.S. in physics from the University of Idaho at Moscow, and an M.S. and Ph.D. in nuclear engineering from the University of California at Berkeley. Following work as a test engineer for the Fast Flux Test Facility and the Fusion Materials Irradiation Test Facility at the Hanford Reservation in Richland, Washington, she joined the Laboratory in 1985 as a physicist working in laser diagnostics for the Nova laser. Her current assignment is in the Nonproliferation, Arms Control, and International Security Directorate where, in addition to leading the systems science group, she conducts research to characterize design, manufacturing, and testing programs for nuclear weapons in countries of concern to the U.S. government. She also analyzes data to resolve questions about activities involving the international nuclear fuel cycle.



ANNETTE MACINTYRE is a deputy division leader for the Electronics Engineering Technology Division in the Engineering Directorate. She also helps manage the systems science groups, teaming with their group leaders to lead a sizable number of technically diverse scientists and engineers who perform risk, safety, and reliability analyses. She has over 20 years of experience in risk, safety, reliability, and discrete-event simulation modeling in the areas of waste management, economics, public and worker safety, and operation tradeoff studies. Most recently, she has worked on the issues of control reliability and effectiveness in real environments. MacIntyre received a bachelor's degree in German and European Affairs and a master's degree in operations analysis, both from the American University in Washington, D.C. She joined Lawrence Livermore in 1987 as the deputy program manager for reliability, availability, and maintainability in the Laser Isotope Separation program. Prior to that, she was a risk analyst and safety engineer at Rockwell Hanford Operations in Richland, Washington.