JOURNAL OF MULTI-CRITERIA DECISION ANALYSIS
J. Multi-Crit. Decis. Anal. **10**: 87–99 (2001)
DOI: 10.1002/mcda.291

Hanford Waste Blending and the Value of Research: Stochastic Optimization as a Policy Tool

TIMOTHY LAWRENCE JOHNSON[a,*] and URMILA M. DIWEKAR[b]
[a] *Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA*
[b] *Department of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA*

ABSTRACT

A new approach to stochastic, combinatorial optimization is presented through its application to a contemporary policy problem: cleanup of radioactive wastes stored underground at the US Government's Hanford, WA, nuclear fuels processing site. Current plans call for the tank contents to be selectively combined prior to their immobilization in glass; such blending of wastes reduces the amount of extra material required for vitrification, therefore, decreasing the costs of processing and disposal. Uncertainty in the tank contents, the error inherent in the glass property models governing vitrification, and the computationally intensive nature of the problem, however, render determination of an optimal tank-blend assignment a challenge to existing optimization techniques. Previous studies have focused exclusively on minimization of processing and disposal costs, ignoring such policy-related dimensions as the value of reducing select sources of uncertainty. In addition, the stochastic framework employed by these studies could not guarantee that the glass property requirements (in the guise of model constraints) were met on more than a probabilistic basis. This paper presents a novel stochastic annealing-nonlinear programming framework that incorporates variance—a proxy for the opportunity costs of reducing uncertainty—as an attribute in its objective function. Compared with conventional mathematical programming algorithms, the new optimization framework is seen to be more robust, flexible, and efficient. The algorithm also facilitates analysis of the trade-off between minimizing processing and disposal costs, and reducing the expenses of achieving this savings. Specifically, the prediction error of the glass property models is found to be a more significant source of uncertainty than variation in tank component mass fraction estimates, and constraint violations are traced to specific requirements of the glass property models. Copyright © 2001 John Wiley & Sons, Ltd.

KEY WORDS: nuclear waste remediation; process modeling; stochastic optimization; uncertainty

1. INTRODUCTION: OPTIMIZATION AS A POLICY TOOL

The use of optimization methods in policy analysis and research, while considerable in its potential, has not been fully appreciated in practice. The complexities of policy-oriented problems—especially the multiplicity of conflicting objectives and the need to act with incomplete information that characterize many situations—combine with the computational and analytical demands of optimization techniques to discourage their use. Yet, these tools can play a valuable role, even in contexts where optimization itself may be of secondary importance. Recently developed algorithms, for instance, facilitate the examination of uncertainty. Questions important to policy analysis, such as how conservative decision-makers should be with respect to risks or where limited resources should be allocated in order to reduce uncertainty, can be framed in an optimization context using these methods.

This paper describes a new optimization algorithm through its application to a complex policy problem: the treatment and disposal of wastes generated in the production of nuclear fuels at the US Government's Hanford, WA, nuclear fuels site. Specifically, the optimization framework of previous analyses is extended to incorporate variance as an attribute in the objective function. The augmented algorithm produces results that are more robust than those of

* Correspondence to: Department of Engineering and Public Policy, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA. E-mail: tjohnson@andrew.cmu.edu

Received 29 December 1999
Accepted 22 April 2001

traditional optimization techniques, with greater efficiency and flexibility. The new framework also facilitates evaluation of important policy dimensions that previous analyses of the Hanford remediation effort have examined only on an informal basis. The methods developed here are applicable to similar problems in combinatorial optimization, and represent a fundamental extension of stochastic programming techniques to real world situations.

The paper first discusses the larger context of the Hanford, WA, waste remediation effort, emphasizing the need to find an optimal blend of the stored wastes prior to their immobilization in glass. The heart of the paper—the development of a robust stochastic annealing-nonlinear programming (STA–NLP) algorithm—follows. Results pertaining to both waste remediation and the efficacy of the new optimization framework are then presented. The paper continues with a discussion of the analysis, and concludes with an assessment of the augmented framework's strengths.

2. HANFORD TANK WASTE: AN UNCERTAIN MIX

The 40-year rush to stay ahead of the Soviet Union in the nuclear arms race, coupled with the demands of military secrecy during the Cold War, rendered environmental concerns secondary to the requirements of national security (US Congress, 1991; Gerber, 1998). The US Department of Energy's Hanford, WA, site is one of many government facilities now confronting the emerging effects of this legacy. Buried beneath Hanford, WA, are 177 storage tanks, with capacities ranging from 50000 to 1 million gallons, holding waste from nearly two-thirds of the weapons-grade nuclear fuel produced in the US following World War II (Gephart and Lundgren, 1995). Mixed into this chemical waste is high-level radioactive material, including strontium, cesium, and plutonium isotopes—the latter with a half-life of 10000 years.

The short-term focus of the original fuel production did not consider the eventual disposal of its concomitant waste materials. The storage tanks, for instance, were originally built with an expected useful life of no more than a few decades; moreover, poor documentation was kept, as chemicals were dumped into the tanks, removed for reprocessing, evaporated to increase storage capacity, and juggled from one reservoir to the next (Gephart and Lundgren, 1995). The long-term consequences of this situation would not be quite as serious if the waste could be maintained underground indefinitely. Known and suspected leaks, the volatile nature of reactions among the tank contents, and uncertainty about the future of the Hanford, WA, site, however, preclude this option (US General Accounting Office, 1994; Probst *et al*., 1996).

The current, though evolving, remediation strategy consists of a multi-stage process, beginning with characterization and preprocessing of the tank contents, separation of low- and high-level radioactive components, and conversion of both waste streams into glass 'logs' for permanent disposal in a national repository (Gephart and Lundgren, 1995; *Nuclear News*, 1997). The blending problem, on which this paper focuses, applies to the glass formation process—known as 'vitrification'—and seeks combinations of tank wastes that minimize the amount of glass produced.

It must be stressed, however, that a complete analysis of the role vitrification will play in Hanford's waste disposal strategy requires consideration of the entire remediation context: its technical *and* political aspects. Significant obstacles may render improvements in the vitrification process secondary in concern when viewed in this context. The objectives of this paper are, therefore, twofold: (1) to describe a new, robust approach to stochastic optimization, while (2) illustrating how a complex policy application like the Hanford blending problem benefits from its use. The importance of increasing the probability that vitrification will succeed and extending the analytical framework to include a more realistic set of objectives make Hanford's blending problem an important application—as well as a significant methodological challenge. (Note that Hanford is being used as a contextual example; results are, therefore, illustrative of those that would be obtained in an actual decision-making environment.)

3. OPTIMIZING VITRIFICATION: THE ADVANTAGES OF BLENDING

Selective blending of the tank wastes prior to vitrification reduces the amount of glass produced only because the tank contents differ (Hrma and

Bailey, 1995; Narayan *et al.*, 1996). A nonhomogeneous mixture of gases, liquids, slurries, and solids, the tank contents include salts, complex organic compounds, various metals, low- and high-level nuclear waste and water (Gephart and Lundgren, 1995). In order to convert these substances into glass, oxides collectively referred to as 'frit' (e.g. SiO_2, B_2O_3, Na_2O, Li_2O, CaO, and MgO) must be added to the waste as it melts.

Blending takes advantage of the fact that the frit constituents are present to varying degrees in each tank; a selective combination of wastes, therefore, reduces the need to add frit during vitrification by matching tanks with complementary frit requirements (Narayan *et al.*, 1996). In addition, blending decreases the *proportion* of so-called 'limiting' components in the combined waste streams. The presence of these components adversely affects the vitrification process; sodium, for instance, reduces glass durability, while aluminum increases both the melt temperature of the material and its viscosity (Hrma and Bailey, 1995). Hence, selective blending of the tank contents increases the probability that vitrification will succeed, reduces frit requirements (i.e. achieves a greater waste-to-frit mass ratio), and minimizes the volume of glass produced.

Blending involves a two-stage decision process: assignment of individual tanks to a given blend and determination of frit requirements. The latter decision depends on the contents of each tank, and is governed by both analytical and empirical glass property models derived for the Hanford tank wastes (Mendel *et al.*, 1980; Hrma and Piepel, 1992). The resulting constraints pertain to the vitrification *process*—rather than characteristics of the subsequent glass—and include bounds on the waste component mass fractions, crystallinity requirements, solubility limits, and attributes of the molten glass, including its liquidus temperature, viscosity and electrical conductivity (Jantzen and Brown, 1993; Hoza, 1994; Narayan *et al.*, 1996). Appendix A presents the glass property models in their guise as deterministic constraints.

This nested structure is typical of problems in combinatorial optimization, where a number of discrete decisions (e.g. the tank-blend assignments) must be made prior to the optimization of some function of continuous characteristics (e.g. the frit mass of each blend)—a dependency that prevents these decisions from being decoupled (Painton and Diwekar, 1995; Chaudhuri and Diwekar, 1998). Solution techniques must cycle between the discrete and continuous decision levels, until a satisfactory 'optimum' is reached. This computational complexity prevents the use of a heuristic, back of the envelope', approach; a formal optimization strategy is, therefore, required. The determination of an optimal set of tank-blends, however, is not a trivial task and, like other problems in stochastic optimization, is one that challenges current techniques.

4. COMPUTATIONAL COMPLEXITY: A CHALLENGE TO EXISTING OPTIMIZATION METHODS

Optimization techniques seek to minimize (or maximize) the value of an objective function, subject to both equality and inequality constraints on its constituent decision variables (see, for example, Ravindran *et al.*, 1987). The most general optimization procedures restrict analysis to a deterministic set of linear objective and constraint equations. Research has focused on extending these methods to more realistic applications involving nonlinear functions of the decision variables, multiple layers of discrete and continuous variables, and probabilistic equations. For several reasons, however, applications like the blending problem remain difficult to solve in a computationally tractable framework without resorting to the excessive use of simplifying assumptions. Such obstacles have rendered traditional mathematical programming techniques ineffective as a means of finding an optimal combination of the Hanford tanks.

- First, standard stochastic optimization methods have difficulty incorporating probabilistic events and parametric uncertainties in a way that recognizes their full importance. Significant uncertainties, for example, exist in both the contents of each Hanford tank and in the prediction error of the empirical glass property models. Stochastic programming techniques are restricted in practice to a few tractable functional forms (including the use of Gaussian distributions), and are therefore limited to relatively simple problems (Painton and Diwekar, 1995; Birge, 1997; Chaudhuri and Diwekar, 1998).

  A common approach to combinatorial optimization, for instance, generates expected values of uncertain parameters or variables by recursively sampling over their respective distributions, propagating these values through the analysis without regard to the underlying variation (Birge, 1997; Birge and Louveaux, 1997; Chaudhuri and Diwekar, 1998). In such a framework, constraints are met only in a probabilistic sense. Hence, while *expected* values of the decision variables satisfy model constraints, there is no guarantee that *actual* (i.e. sample) values will produce an acceptable—or even a feasible—solution. Robustness to variation acquires additional relevance when framed in a policy context like that of Hanford's waste remediation effort, where constraint violation implies the possible failure of vitrification. The following section describes an STA algorithm that explicitly addresses this aspect of the blending problem—an issue that previous analyses have ignored (e.g. Hoza, 1994; Narayan *et al*., 1996).

- Nonlinearities and nonconvexities in the space defined by the objective and constraint functions generate a second computational challenge. Both sets of equations in the blending problem contain second-order terms. Previous analyses (e.g. Hoza, 1994; Narayan *et al*., 1996) have employed mathematical optimization techniques, such as mixed integer NLP (MINLP). These methods are derivative-based search algorithms, and are thus prone to becoming 'trapped' at a local minimum (Ravindran *et al*., 1987). Evaluation of several initial points is, therefore, required to ensure that a 'global' solution is found, increasing the computational burden. Simulated annealing—a more contemporary optimization algorithm that is not a derivative-based search technique—forms the basis of the new framework described below.
- Like related applications in combinatorial optimization, the blending problem requires an iterative evaluation of continuous and discrete variables, resulting in a combinatorial 'explosion'. Dividing N tanks into B blends of T tanks each yields

$$N!/B!(T!)^B \qquad (1)$$

potential solutions. A subset of 21 of the 177 Hanford tanks, for instance, partitioned into three blends of seven tanks each produces 66.5 million tank-blend combinations. Conventional optimization techniques running on a high-end workstation (a DEC Alpha 400) take an estimated 60 days to identify a solution; the algorithm described below, in contrast, requires less than 24 h.
- Consideration of the multiple objectives typical of realistic applications poses a fourth source of problem complexity. Previous efforts to address the blending problem (e.g. Hoza, 1994; Narayan *et al*., 1996), for instance, have focused solely on the cost of vitrification (i.e. minimization of frit, which is equivalent to minimizing glass volume and, hence, disposal costs). While these efforts have included a representation of the different sources of uncertainty inherent in the blending problem, they have not recognized reduction of this uncertainty as an important objective in itself. Significant policy dimensions related to the vitrification process have thus been ignored. The augmented framework described in the next section facilitates a comparative analysis of the resulting trade-offs.

5. OPTIMIZATION WITH VARIANCE REDUCTION: A NEW FRAMEWORK

A realistic assessment of applications similar to the Hanford blending problem requires the use and further development of improved techniques, especially as they relate to combinatorial optimization in situations of uncertainty. This section describes a new approach to stochastic optimization as applied to the blending problem; a detailed discussion of the model's policy implications and the progressively complex objective functions from which they are derived follows.

The new framework employs a coupled STA–NLP algorithm; the sequential association of optimization techniques reflects the two-stage decision structure typical of combinatorial problems. STA is a probabilistic extension of the simulated annealing (SA) algorithm (see Painton and Diwekar, 1995 for STA, Van Laarhoven and Aarts, 1987 for SA). Figure 1 illustrates the augmented STA–NLP model as applied to the blending problem; a similar nested structure would fit other combinatorial optimization applications. The iterative framework consists of three 'loops':

- a *discrete variable* decision loop that selects those tank-blend configurations that minimize the expected frit mass *and* its associated sample variance;

- a *continuous variable* decision loop that determines the amount of frit required by each of the resulting blend combinations, subject to a series of glass-property constraints; and
- a *sampling* loop that calculates expected values of the waste component mass fractions.

This section describes the elements of the STA–NLP framework, beginning with those at the centre of Figure 1.

- *Sampling loop:* the sampling loop utilizes a Latin hyper-cube (LHC) algorithm (see Iman and Shortencarrier, 1984) to calculate expected values of the waste component mass fractions for each tank over distributions provided by Hanford scientists (see Appendix B for a representative tank's composition). LHS, compared with traditional Monte Carlo methods, produces more uniform samples and is less computationally intensive. The waste component mass fraction sampling distributions represent a combination of expert judgement, physical inspection, and analytical calculations (Mendel *et al.*, 1980; Hrma and Piepel, 1992; Hopkins *et al.*, 1994). The difficulty in sampling from a nonhomogeneous mixture, the high cost of probing the tanks, component reactions, and the poorly documented history of the tank contents result in large estimated mass fraction variances (var_{samp}). Normal distributions characterize this uncertainty; determination of sample size is discussed below. [Note that the mass fractions are assumed to be uncorrelated—an assumption that is not realistic. Reactions among the components, for instance, introduce but one source of dependence (Hopkins *et al.*, 1994). Limited knowledge precludes a more accurate assessment, although the correlation structure would not be difficult to incorporate with the availability of trustworthy estimates.]

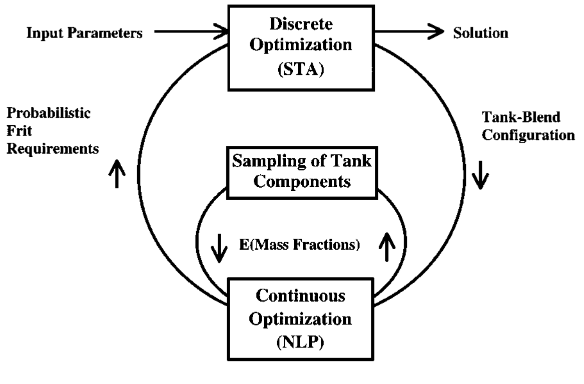

Figure 1. The STA–NLP framework as applied to the blending problem.

- *Continuous decision loop (NLP):* the continuous variable decision loop determines the frit mass required by each blend in a given tank configuration based on the expected waste component mass fractions. The constraint functions (Appendix A) of the NLP model dictate frit requirements (i.e. the NLP loop 'adds' frit to each blend until all of its constraints are satisfied). As the waste mass fractions are expected values, the constraints are met only on a probabilistic basis and the resulting frit requirements are expected values. Constraints may be violated when the *sample* waste mass fractions are used with frit requirements based on the sample averages (as is the case in actuality). Minimizing the frit mass sample variance (var_{frit})—in addition to its expected value—therefore, increases the robustness of the algorithm. Note the parallel to Taguchi's robust design methodology, which seeks to maximize a 'signal-to-noise' ratio of a quantity's expected value to a measure of its variation (see, for example, Kacker, 1985).
The width of the constraint bounds reflects an additional source of uncertainty: the prediction error inherent in the empirical glass property models. Greater error produces more conservative bounds. As an example, the conductivity constraints are written:

$$\text{lower limit} + U_{\text{cond}} < \Sigma_i c_a^i p^{(i)} + \Sigma_i \Sigma_j c_b^{ij} p^{(i)} p^{(j)} < \text{upper limit} - U_{\text{cond}} \qquad (2)$$

where the c_i are coefficients from the empirical glass property model and $p^{(i)}$ represents the mass fraction of the ith glass component. The prediction error of the corresponding empirical conductivity model fixes the uncertainty term, U_{cond}. Note that larger values of U_{cond} produce tighter bounds, and more restrictive frit mass requirements. Relaxing the constraint width (i.e. [lower limit + U_{cond}] − [upper limit − U_{cond}]) is, therefore, desirable.

- *Discrete decision loop (STA):* governing the process is a stochastic variant of the simulated annealing algorithm. Analogous to metallurgical annealing, STA is an iterative procedure that compares the current objective function value with that from a randomly perturbed tank-blend configuration, accepting lower values on a probabilistic basis that becomes progressively restrictive (Painton and Diwekar, 1995; Chaudhuri and Diwekar, 1998). Probabilistic acceptance of a given solution helps the STA algorithm avoid termination at a local optimum, as derivative-based optimization methods (e.g. MINLP) are prone to do. The STA procedure ceases when a stopping criterion related to its rate of convergence is satisfied.

Finally, the STA algorithm includes sample size as an additional decision variable, one that receives greater weight as the optimum is reached. The use of fewer samples when the algorithm is far from a minimum and 'exploring' the surface of the objective function gives STA its computational advantage. (Note that the sample size term is not shown in the objective functions described in the next section. See Painton and Diwekar (1995) and Chaudhuri and Diwekar (1998) for details). The following section examines augmentation of the objective function to include, among other attributes, the uncertainty induced by the sampling procedure.

6. VARIANCE AS AN ATTRIBUTE: THE ANALYSIS OF UNCERTAINTY

Sources of uncertainty in the blending problem have important technical implications, and reflect significant aspects of the policy-making process surrounding Hanford's remediation efforts. Expansion of the objective from minimization of frit to include different sources of variation represents an important methodological development, one that capitalizes on the STA–NLP framework to make stochastic optimization a more robust mathematical technique and a more useful policy tool. This section illustrates the STA–NLP framework's advantages through progressive extensions to the blending problem's objective function. The base analysis is presented first, and results accompany the description of each extension. The following section discusses the corresponding implications.

6.1. Base objective: minimization of frit mass

A subset of 12 tanks, divided evenly into three blends, forms the basis of this analysis. The evaluation presented below utilizes the STA–NLP framework of Figure 1. All computer code was written in FORTRAN, with prepared routines modified for the NLP optimization and LHS sampling loops (Iman and Shortencarrier, 1984); data was taken from Hopkins *et al.* (1994) and Narayan *et al.* (1996).

Initial remediation efforts at the Hanford, WA, site will focus on a limited number of storage tanks; the criticality of a tank's condition (its position on a 'watch list'), and the compatibility of its contents with the demands of vitrification will govern the selection process (Gephart and Lundgren, 1995). Evolving waste management plans have established an initial period over which the vitrification process will be refined (*Nuclear News*, 1996, 1997). Hence, both practical realities and computational convenience justify the use of a subset of tanks in this *illustrative* analysis.

The blending scheme planned for the Hanford tank waste sits on a continuum of strategies between vitrification without blending and a 'total blend' of all tank materials into one waste stream (Narayan *et al.*, 1996). The total blend alternative, in theory, requires the least amount of frit; the impossibility of combining all tank wastes into one batch, however, precludes its use in practice. The no-blend and total solutions are of interest, nevertheless, as they provide bounds against which changes in discrete-blend frit requirements may be compared.

A deterministic analysis of the blending problem yields a further basis for comparison. 'Deterministic' in this sense implies use of the STA–NLP framework with nonstochastic constraint bounds and without the sampling loop (i.e. the modes of the empirical mass fraction distributions are substituted for sample means generated from these distributions). Table I presents the frit requirements from these preliminary solution schemes, based on the base-case objective:

$$\text{minimize}\{\text{frit mass}\} \quad (3)$$

Note that the difference between the deterministic and stochastic solutions (1101 kg) is commonly referred to as the *value of the stochastic solution.*

Figure 2 present histograms of the frit mass requirements and corresponding proportion of constraint violations when the *individual* waste

Table I. Frit requirements as determined by basic solution techniques

Solution method	Required frit mass (kg)
Worst case (no blending)	13 410
Best case (one blend of all tanks)	9839
Deterministic solution (STA–NLP without sampling)	11 161
Single attribute stochastic solution (STA–NLP with sampling)	10 060

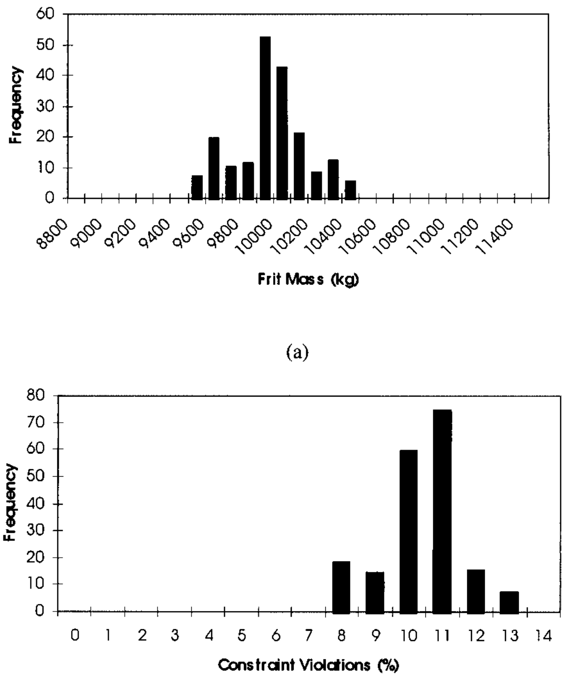

Figure 2. Distribution of (a) frit masses, and (b) constraint violations for the base-case objective: minimize frit mass.

mass fraction sample values are used with the tank-blend configuration derived from their expected value.

6.2. Robustness: minimizing variance

The variance in frit requirements, var_{frit}, is a measure of the STA–NLP algorithm's robustness. The magnitude of var_{frit}, for instance, directly affects the probability that the NLP/glass property constraints are met when *actual* (i.e. sample) values of the waste component mass fractions are used in place of their sample mean. Hence, there is a desire to keep this source of variation as low as possible. Including variance as an attribute produces the following objective:

$$\text{minimize}\{\text{frit mass} + w_1 * \text{var}_{\text{frit}}\} \tag{4}$$

Note that 'frit mass' in Equation (4) is an expected value, and that var_{frit} has been scaled so that both terms have the same order of magnitude. The variance of frit mass is used instead of its standard deviation. While portfolio theory optimization frameworks employ the latter, quality control models like the loss function—which, like the blending problem, are characterized by a nonlinear domain—feature variance (again, compare this approach with Taguchi's robust design methodology (Kacker, 1985)).

The optimization framework illustrated here and extended in the following section, unlike formal multi-attribute decision analysis, is *qualitative* in nature. Specific meaning, for instance, cannot be attached to w_1 (or the w_i in general; see Equation (8)). The highly nonconvex, nonlinear, and discrete character of the blending problem precludes the assessment of 'weights' customary to multi-attribute optimization algorithms. The parsimonious choice of an additive objective function in Equation (4), as well as the selection of units and scaling factors for its terms, determines the trade-offs produced by variation of w_1. Attention, therefore, should focus not on the w_1 term, but on the relative changes in frit mass, its variance, and the number of constraint violations that parametric adjustments of w_1 produce. The scale of w_1 values explored was selected iteratively in order to observe the complete range of the criteria of interest—in this case, constraint violations (which decreased from a maximum of eleven percent to zero; see Table II). Future work aimed at

Table II. The balance between expected value and variance minimization

w_1	Frit mass (kg)	$\sqrt{(\text{var}_{\text{frit}})}$ (kg)	% constraints violated
0.50	10 255	293	11
1.0	10 075	190	6
2.0	10 647	138	2
4.0	11 558	118	0

exploring the entire objective space and obtaining the complete Pareto set of weights is discussed in a later section of this paper.

An illustration clarifies this caveat. The decrease in NLP constraint violations (produced by using waste component mass fraction sample values rather than their means, for which the constraints are *always* met) can be examined as a function of increasing frit mass (Table II).

As shown in Table II, an increase in expected frit mass of approximately 4% yields an 80% reduction in constraint violations—an important factor for decision making at Hanford. Again, the corresponding increase in w_1 (from 0.5 to 2.0) that produces this result does not have a meaningful interpretation. Nor is the trade-off between frit mass and its variance constant over the range of w_1. Instead, the value of this framework lies in its ability to explore trade-offs in terms of relative changes between different factors relevant to the blending problem (the compromise between increasing frit mass and decreasing constraint violations). The following sub-section builds on this flexibility.

6.3. Reducing uncertainty: minimizing the time devoted to research

The ability to incorporate additional terms into the objective function illustrates stochastic optimization's potential as a policy tool. Better characterization of the Hanford tank wastes and glass property models, for instance, would result in lower frit requirements. The decrease in frit mass that a reduction in uncertainty yields, however, must be weighed against the opportunity costs of pursuing the objective. The extensions introduced here facilitate this analysis: an examination of the trade-offs inherent in allocating scarce resources to reducing uncertainty. Such extensions are generalizable to similar situations, which are ubiquitous, especially in nuclear waste management where the long-lived nature of the waste creates large uncertainties.

The analysis rests on a key assumption: that time spent on research increases understanding and, therefore, decreases variation in quantitative estimates derived from this knowledge. Research activities introduce their own costs and risks; hence, time spent learning and experimenting therefore needs to be minimized. While reducing uncertainty is profitable, the time required to achieve a reduction tempers the benefit. An augmented objective captures this trade-off:

$$\text{minimize}\{\text{processing and disposal costs } \textit{and} \text{ time devoted to reducing uncertainty}\} \quad (5)$$

As before (Equation (4)), processing and disposal costs are represented by the expected frit mass and its associated variance. As illustrated below, the sampling variance of the tank waste component mass fractions and the uncertainty in the empirical glass property models (through its effect on constraint width) serve as proxies for resources devoted to reducing uncertainty.

The expanded blending objective therefore attempts to minimize frit mass, but—beyond finding an optimal tank-blend assignment—limits the extent to which improved waste characterization and more accurate glass property models contribute to this goal. Research efforts, for instance, could aim at easing the constraint bounds via improvements in the glass property models' prediction error; as the constraints govern frit requirements, less conservative limits in an optimization framework translate into the need for a smaller safety margin and, therefore, less frit. Proportional relaxation of the constraints, however, carries an increasing penalty: the time and opportunity costs of related research activities.

To understand how the augmented blending objective captures this trade-off in mathematical terms, note that the type of investigation relevant to the problem will exhibit diminishing marginal returns as uncertainty declines nonlinearly with time spent on research. For characterization of the tank waste components, an exponential relationship between sampling variance and time provides an adequate first-order functional approximation of this nonlinear dependence:

$$\text{uncertainty in waste composition} \Leftrightarrow \text{var}_{\text{samp}} \propto \exp(-\text{time}) \quad (6)$$

or

$$\text{time} \propto -\ln(\text{var}_{\text{samp}})$$

A similar relationship holds for the constraint width term. Note, however, that the width of the constraint bounds varies *inversely* with the prediction error of the empirical glass property models (U_{cond} in Equation (2)):

$$\begin{aligned}\text{time} &\propto -\ln(\text{prediction error})\\ &\propto -\ln(\text{constraint width})^{-1}\\ &= \ln(\text{constraint width}) \quad (7)\end{aligned}$$

Once again, minimization of resources devoted to reducing uncertainty, *taken by itself*, is captured in this model by seeking tank-blend combinations with larger input sampling variances and prediction errors (i.e. narrower constraint bounds). Excessive values, however, are simultaneously penalized through their detrimental effect on the expected frit mass and its associated sample variance. The optimum reflects a balance in this trade-off: a low frit mass and $\mathrm{var_{frit}}$, with moderate values of $\mathrm{var_{samp}}$ and the constraint widths. Combining these arguments, the augmented blending objective becomes:

$$\text{minimize}\{\text{frit mass} + w_1 * \mathrm{var_{frit}} - w_2 * \ln \Sigma\, \mathrm{var_{samp}} + w_3 * \ln \Sigma \text{ constraint width}\} \quad (8)$$

The caveats discussed in the previous section with respect to w_1 and Equation (4) apply to Equation (8): the w_i do not have a meaningful interpretation, and they should not be taken as a measure of 'priority'. Rather, the w_i are simply a means of assessing trade-offs between the conflicting goals, in this instance, of decreasing the costs of vitrification and minimizing the unnecessary reduction of uncertainty. Table III presents results of a parametric analysis of changes in the w_i, similar to that presented earlier in Table II. Table IV presents a qualitative summary of these results, the implications of which are discussed in the following section.

7. DISCUSSION: THE IMPLICATIONS OF UNCERTAINTY

The results from the previous section have implications for stochastic optimization in general and the blending problem in particular. The importance of attending to matters of robustness, for instance, is apparent in Table II: as reduction in frit variance is emphasized (i.e. as w_1 increases), the proportion of constraint violations decreases to zero and the frit masses become clustered more tightly around their mean. The expected frit mass, however, is uniformly higher with fewer constraint violations—a compromise that illustrates the balance between reducing the volume of immobilized waste and increasing the probability that vitrification succeeds. The augmented STA–NLP algorithm facilitates such an analysis.

Beyond providing a framework in which similar trade-offs may be assessed, however, policy-makers desire answers to specific questions. Note that the most important question concerning the blending problem is not minimization of frit mass per se; indeed, consideration of the entire context

Table III. Parametric results of the trade-off in reducing sources of variation

w_1 ($\mathrm{var_{frit}}$)	w_2 ($\mathrm{var_{samp}}$)	w_3 (constant width)	E[frit mass] (kg)	$\sqrt{(\mathrm{var_{frit}})}$ (kg)	% constraints violated
0	0	0	10 255	293	11
1	1	1	10 932	214	8
1	1	3	11 061	192	9
1	3	1	9931	478	5
1	3	3	9971	337	5
3	1	1	10 815	184	2
3	1	3	10 050	175	3
3	3	1	12 008	245	2
3	3	3	11 217	230	3

Table IV. A qualitative summary of the trade-off in reducing sources of variation

Focus of research	E[frit mass]	Result $\mathrm{var_{frit}}$	% constraint violations
Robustness/minimization of frit variance	Increases	Decreases	Decreases
Minimize time devoted to tank characterization	Increases	Increases	Increases
Minimize time devoted to improving glass property models	Does not change	Decreases	Increases

of Hanford's remediation effort and the politics of radioactive waste disposal may decrease the priority of reducing frit mass—especially on the order of the savings seen above (compare the values in Tables I and III). (Expanding the problem scale by including a larger subset of tanks, however, would increase the importance of lowering frit mass; a greater number of tanks would also take better advantage of blending, and result in more impressive reductions of frit.)

More important are questions concerning uncertainty: to what extent is imperfect information acceptable, and where should scarce resources be allocated to leverage the impact of this narrow part of Hanford's waste remediation effort on the whole of its strategy? Not all sources of uncertainty, after all, are significant. In pursuing answers to such questions, stochastic optimization functions more as an exploratory tool than as a means of providing 'one best' solution.

The preceding analysis illustrates this capacity. An examination of the constraints, for instance, reveals that the crystallinity requirements are most consistently violated, with the P_2O_5 solubility limit and the component bound on Al_2O_3 frequently exceeded as well (see also Hopkins *et al.*, 1994). Resources would be profitable allocated to reducing the error in the corresponding glass property models ahead of additional waste pretreatment efforts designed to mitigate the effects of these limiting components.

Perhaps more significant is the ability to determine what sources of uncertainty need to be reduced and which, in contrast, may be tolerated. The relationship, however, between the required frit mass, its variance, and constraint violations is complicated. As described, the constraint width terms enter the objective function as penalties; considered in isolation of their effects on frit mass, larger values are desired (i.e. the devotion of resources to reducing uncertainty is minimized). The 'benefit' of greater uncertainty in the tank waste distributions and glass property models, however, is balanced by its detrimental effect on the expected frit mass and its variance.

Results from the preceding section illustrate this relationship. As the sampling variance term increases (i.e. characterization the tank wastes is less complete), variation in frit mass increases and constraint violations become more numerous. This effect is not surprising: a change in the variance of the waste component sampling distributions leads to a proportionate shift in the frit variance and a similar impact on both the average frit mass and extent of constraint violations.

Compared with these changes, however, the variance in frit mass *decreases*, while the percentage of constraint violations *increases* with the constraint width uncertainty (see Equation (2), and compare parts (a) and (b) of Figures 3 and 4, which illustrate the effect of increasing w_3 in Equation (8)); greater uncertainty in the glass property models translates into narrower constraint bounds, and a smaller range across which frit requirements may vary without consequence. This impact on process robustness leads to the conclusion that improvements in the glass property models should come before efforts to reduce uncertainty in the tank waste composition. The presence of nonlinearities in the glass property models (constraints)—which inflate the effects of variance—provides one explanation for the pattern of these results.

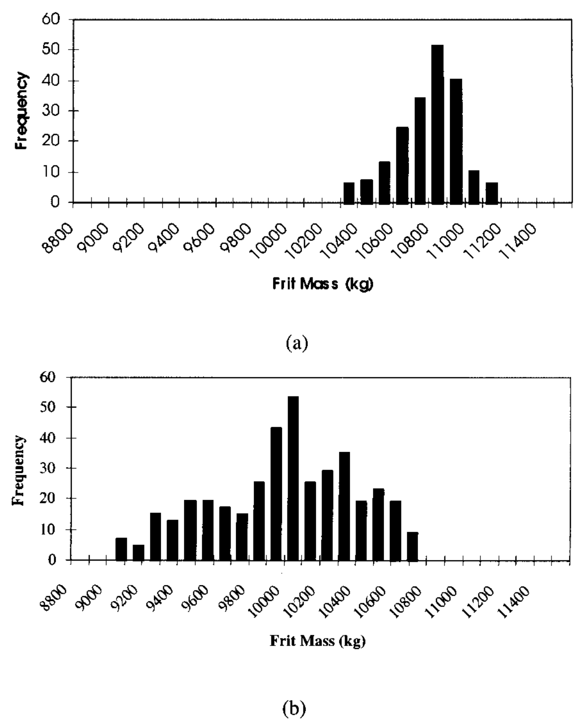

Figure 3. Distribution of frit masses for different constraint width objective coefficients; (a) $w_3 = 1.0$, (b) $w_3 = 3.0$.

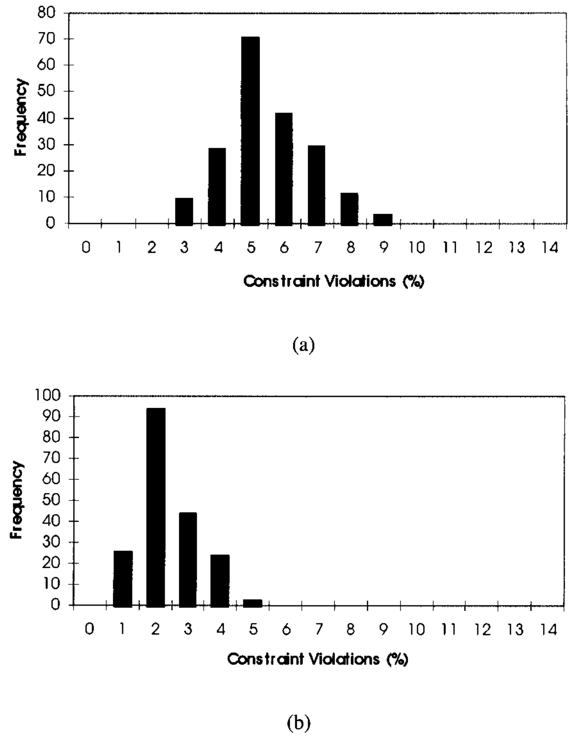

Figure 4. Distribution of constraint violations for different constraint width objective coefficients; (a) $w_3 = 1.0$, (b) $w_3 = 3.0$.

8. CONCLUSIONS: THE STRENGTHS OF THE NEW ALGORITHM

Stochastic optimization deserves to be included in the policy analysis toolbox. The computational and algorithmic complexity of adapting many policy problems to an appropriate framework, however, has limited its application. Developments like the STA–NLP algorithm will lower these technical hurdles, and hopefully increase the practical value of stochastic optimization in the eyes of policy researchers.

The augmented STA–NLP framework, for instance, facilitates a more realistic analysis of questions pertaining to uncertainty than competing optimization strategies, and recognizes the trade-offs that arise when knowledge is incomplete. The analysis presented here, for instance, identified the predictive error of the glass property models as the most significant source of uncertainty—an issue that pervious approaches to the blending problem could not address. Not all sources of uncertainty are consequential; the ability to distinguish between those sources that are important and those that may be tolerated is, therefore, a valuable contribution.

The augmented STA–NLP framework is more robust than equivalent mathematical programming techniques. The inclusion of variance in the objective function helps ensure that constraints are met by actual values of the decision variables—not just their expected values. The new framework is also more flexible than its competitors, as illustrated by the inclusion of other variance-related terms in the objective function. Extension of the blending objective facilitated a structured evaluation of an important policy issue: the trade-off between living with uncertainty now versus the future costs and benefits of increasing knowledge. And, not least important, the STA–NLP algorithm is faster, reaching a solution with a fraction of the computational resources required by more traditional mathematical optimization routines.

The development and application of the STA–NLP algorithm is a work in progress. Current research is focusing first on the enhancement of sampling schemes. A fractal-dimension approach to error characterization, for instance, promises further improvements in algorithmic efficiency. Application of the framework to larger-scale problems is also of interest. An increase in the number of tanks will demonstrate the computational advantage—and the necessity—of using methods similar to those discussed here.

A third focus concerns nonconvexities and formal assessment of the objective function weights. Work is proceeding on the development of an algorithm that utilizes a multi-attribute constraint method to select the complete set of Pareto optimal weights. This approach will facilitate assessment of the objective function over its entire range.

Progress in these developments will continue to lower the technical barriers of stochastic optimization, and move its application further into the realm of policy-making.

ACKNOWLEDGEMENTS

The authors wish to thank Dr Jared Cohon of Carnegie Mellon University for his support, ideas and guidance. This work was sponsored in part by the National Science Foundation under grant CTS-9613561.

REFERENCES

Birge JR. 1997. Stochastic programming computation and applications. *INFORMS Journal on Computing* **9**: 111–132.

Birge JR, Louveaux F. 1997. *Introduction to Stochastic Dynamic Programming*. Springer-Verlag: New York.

Chaudhuri P, Diwekar U. 1998. Pollution prevention design. In *Encyclopedia of Environmental Analysis and Remediation*, Meyers RA (ed.). Wiley: New York; 3772–3791.

Gephart RE, Lundgren RE. 1995. Hanford tank clean up: a guide to understanding the technical issues. Report BNWL-645, Pacific Northwest Laboratory, Richland, WA.

Gerber MS. 1998. Historical generation of Hanford site wastes. Report WHC-SA1-224, Pacific Northwest Laboratory, Richland, WA.

Hopkins DF, Hoza M, Lo Presti CA. 1994. FY94 optimal waste loading models development. Report PVTD-C9-402.04D, Pacific Northwest Laboratory, Richland, WA.

Hoza M. 1994. Multipurpose optimization models for high-level waste vitrification. In *Proceedings of the International Topical Meeting on Nuclear and Hazardous Waste Management—SPECTRUM 94*. American Nuclear Society: La Grange Park, IL; 1072–1077.

Hrma PR, Bailey AW. 1995. High level waste at Hanford: potential for waste loading maximization. Report PNLS-A26-441, Pacific Northwest Laboratory, Richland, WA.

Hrma PR, Piepel GF. 1992. Property/composition relationships for Hanford waste vitrification plant glasses—preliminary results through CVS-II Phase 2. Report PHT-D92-03.01/K897, Pacific Northwest Laboratory, Richland, WA.

Iman RL, Shortencarrier MJ. 1984. A FORTRAN 77 program and users guide for the generation of Latin hypercube and random samples for use with computer models. Report SAND83-2365, Sandia National Laboratories, Albuquerque, NM.

Jantzen CM, Brown KG. 1993. Statistical process control of glass manufactured for nuclear waste disposal. *American Ceramic Society Bulletin* **72**: 55–59.

Kacker RS. 1985. Off-line quality control, parameter design, and the Taguchi method. *Journal of Quality Technology* **17**: 176–209.

Mendel JE, Rawest WA, Turcotte RP, McElroy JL. 1980. Physical properties of glass for immobilization of high-level radioactive waste. *Nuclear and Chemical Waste Management* **1**: 17–28.

Narayan V, Diwekar UM, Hoza M. 1996. Synthesizing optimal waste blends. *Industrial and Engineering Chemistry Research* **35**: 3519–3527.

Nuclear News. 1996. DOE issues RFP for Hanford tank waste solidification. *Nuclear News* **39**: 40.

Nuclear News. 1997. DOE selects Hanford tank waste cleanup plan. *Nuclear News* **40**: 49.

Painton L, Diwekar U. 1995. Stochastic annealing for synthesis under uncertainty. *European Journal of Operational Research* **83**: 489–502.

Probst KN, Pilling CA, Dunn KT. 1996. Cleaning up the nuclear weapons complex: Exploring new approaches. Discussion Paper 96-25, Resources for the Future, Washington, DC.

Ravindran A, Phillips DT, Solberg JJ. 1987. *Operations Research: Principles and Practice* (2nd edn). Wiley: New York.

US Congress, Office of Technology Assessment. 1991. *Complex Cleanup: The Environmental Legacy of Nuclear Weapons Production*, OTA-O-484. US Government Printing Office: Washington, DC.

US General Accounting Office. 1994. Nuclear cleanup: completion of standards and effectiveness of land use planning are uncertain. Report GAO/RCED-94-114, US General Accounting Office, Washington, DC.

Van Laarhoven PJM, Aarts EHL. 1987. *Simulated Annealing: Theory and Applications*. Reidel: Dordrecht.

APPENDIX A: THE GLASS-PROPERTY MODELS AS *DETERMINISTIC* CONSTRAINT FUNCTIONS

See Equation (2) in the text for an example of the incorporation of uncertainty.

Note that $p(i)$ refers to the mass fraction of the ith component in the glass; b_i, c_i, and d_i are corresponding coefficients.

A.1. Component bounds

- $0.42 < p(SiO_2) < 0.57$
- $0.05 < p(B_2O_3) < 0.20$
- $0.05 < p(Na_2O) < 0.20$
- $0.01 < p(Li_2O) < 0.07$
- $0.00 < p(CaO) < 0.10$
- $0.00 < p(MgO) < 0.08$
- $0.02 < p(Fe_2O_3) < 0.15$
- $0.00 < p(Al_2O_3) < 0.15$
- $0.00 < p(ZrO_2) < 0.13$
- $0.01 < p(\text{other}) < 0.10$

A.2. Crystallinity constraints

- $p(SiO_2) > 3p(Al_2O_3)$

- $p(\mathrm{MgO}) + p(\mathrm{CaO}) < 0.08$
- $p(\mathrm{Fe_2O_3}) + p(\mathrm{Al_2O_3}) + p(\mathrm{ZrO_2}) + p(\mathrm{other}) < 0.225$
- $p(\mathrm{Al_2O_3}) + p(\mathrm{ZrO_2}) < 0.18$
- $p(\mathrm{MgO}) + p(\mathrm{CaO}) + p(\mathrm{ZrO_2}) < 0.18$

A.3. Solubility constraints

- $p(\mathrm{Cr_2O_3}) < 0.005$
- $p(\mathrm{F}) < 0.017$
- $p(\mathrm{P_2O_5}) < 0.01$
- $p(\mathrm{SO_3}) < 0.005$
- $p(\mathrm{Rh_2O_3} + \mathrm{PdO} + \mathrm{Ru_2O_3}) < 0.025$

The summation in the following constraints is over the total number of chemical components in the glass, including oxides present in both the tank waste and the added frit.

A.4. Viscosity constraints

$\ln(2) < \Sigma_i b^i_a p^{(i)} + \Sigma_i \Sigma_j b^{ij}_b p^{(i)} p^{(j)} < \ln(10)$ (in PaS)

A.5. Conductivity constraints

$\ln(18) < \Sigma_i c^i_a p^{(i)} + \Sigma_i \Sigma_j c^{ij}_b p^{(i)} p^{(j)} < \ln(50)$ (in S/m)

A.6. Dissolution rate for boron

$\Sigma_i d^i_a p^{(i)} + \Sigma_i \Sigma_j d^{ij}_b p^{(i)} p^{(j)} < \ln(28)$ (in g/m^2)

(For elaboration see: Jantzen and Brown, 1993; Narayan *et al*., 1996.)

APPENDIX B: COMPOSITION OF THE HIGH-LEVEL WASTE IN HANFORD TANK B-110

Component	Mass (kg)*	Mass fraction	Mass standard deviation (kg)
Al_2O_3	25 165	0.02	3775
B_2O_3	1076	8.6E-04	140
CaO	14 195	0.011	994
Fe_2O_3	288 285	0.23	11 531
Li_2O	0	0	–
MgO	3378	0.003	135
Na_2O	101 111	0.08	4045
SiO_2	220 305	0.18	8812
ZrO_2	51	4.1E-05	6.2
Other**	603 429	0.48	33 889

* The mass values represent most likely estimates.
** Other includes (e.g.): Bi, NO_2, SO_4, PO_4, NO_3, and P.
(For elaboration see: Hopkins *et al*., 1994.)