LA-UR- 99-522

Approved for public release; distribution is unlimited.

CONF-990201--

Title: EVALUATION OF HIGH-LEVEL WASTE PRETREATMENT PROCESSES WITH AN APPROXIMATE REASONING MODEL

Author(s): Terry F. Bott, TSA-11
Stephen W. Eisenhawer, TSA-11
Stephen F. Agnew, CST-4

Submitted to: WM'99 Conference
February 28 - March 4, 1999
Tucson, AZ

RECEIVED
APR 13 1999
OSTI

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

EVALUATION OF HIGH-LEVEL WASTE PRETREATMENT PROCESSES WITH AN APPROXIMATE REASONING MODEL

Terry. F. Bott, Stephen W. Eisenhawer and Stephen F. Agnew
Los Alamos National Laboratory
Los Alamos, New Mexico 87545
(505) 667-9207

ABSTRACT

The development of an approximate-reasoning (AR)-based model to analyze pretreatment options for high-level waste is presented. AR methods are used to emulate the processes used by experts in arriving at a judgment. In this paper, we first consider two specific issues in applying AR to the analysis of pretreatment options. We examine how to combine quantitative and qualitative evidence to infer the acceptability of a process result using the example of cesium content in low-level waste. We then demonstrate the use of simple physical models to structure expert elicitation and to produce inferences consistent with a problem involving waste particle size effects.

INTRODUCTION

Removing high-level waste from storage tanks for long-term disposal typically involves four phases: retrieval, pretreatment, immobilization of the low-level activity waste (LAW) and high-level activity waste (HLW) and disposal. Many options exist for performing each of these phases, and therefore, there are many potential alternatives for the overall process. The individual phases interact strongly with one another, so it is not possible to optimize each phase separately, which complicates the design process considerably. An efficient design requires matching the phases to one another using an iterative approach. Detailed simulation of a process often is used in the search for an optimization design. However, this is not a practical approach at the current stage of the remediation process design for a number of reasons. First, there is considerable uncertainty about the characteristics of the waste in individual tanks. This strongly affects the choices for retrieval and pretreatment and the schedule in which the tanks will be retrieved. Second, the flowsheets for the individual phases are still subject to rapid and large changes. Finally, many of the process evaluation criteria are qualitative in nature and are not necessarily closely linked to the simulation results.

In an earlier paper (1), we described a prototype of a tool called the Pretreatment Analysis Tool (PAT). The purpose of this paper is not a detailed description of the PAT. However, some background on its theory and application will make the specific cases we discuss more understandable. The purpose of the tool is the rapid assessment of the effect of pretreatment process changes on the overall removal process. The PAT is based on the theory of approximate reasoning (AR) (2). The PAT uses AR to emulate the evaluation of pretreatment options by a group of experts. The algorithm captures the knowledge and reasoning processes of experts

through a series of linked rule bases. This is more efficient and practical than attempting to assemble a group of experts on a frequent basis.

The requirement for expert evaluation of pretreatment options arises in the following way. A designer developing a pretreatment process might wish to concentrate his analysis within a small region in the parameter space. For example, a specific concern might be associated with waste pH or the decontamination factor for a specific isotope. However, optimization of a single parameter frequently leads to tradeoffs in other process characteristics that can be highly unfavorable. A designer would be less likely to make this mistake if he were able to consult a panel with the expertise to assess these tradeoffs. In the PAT, we explicitly consider tradeoffs associated with the acceptability of costs, waste loading, and decontamination factors as functions of the pretreatment option. The AR algorithms provide expert feedback to facilitate the exploration of the design parameter space and quickly identify pretreatment parameters that result in unacceptable outcomes.

In this paper, we consider two specific technical issues encountered during the implementation of AR in the PAT. The first of these involves integrating qualitative judgments with numerical data to generate an evaluation of the acceptability of a process. We illustrate our solution to this problem with an example that considers the site's sensitivity to the residual activity in the LAW. The second issue involves using approximate physical models to provide a logical structure to organize expert opinion. For this example, we present a simple model for particle-size distillation during retrieval and show how the model simplifies the elicitation process.

OVERVIEW OF AN APPROXIMATE REASONING MODEL OF PRETREATMENT

AR models are designed to emulate the deliberation processes used by subject matter experts in making a decision and are a special type of expert system. The primary characteristics of the AR model presented here are as follows.

- The relationship between data and models is defined explicitly by an inductive logic structure.
- The inference process is implemented using formal logical implication.
- All of the relevant data are transformed into natural language expressions called linguistic variables.
- A linguistic variable is represented in the model using fuzzy sets. The degrees of membership in fuzzy sets are used directly in the logical implication operations.
- Judgments about the importance or quality of data are also treated as linguistic variables. These judgments are incorporated explicitly into the inference model.
- Uncertainty is represented using standard Monte Carlo techniques for random variables. Ambiguity is treated using fuzzy sets.

Additional applications of AR in the areas of safety analysis and the evaluation of system reliability are described in Eisenhawer (3,4) and Bott (5).

The AR approach used in the PAT can handle both quantitative and qualitative parameters in the evaluation. Some quantitative pretreatment parameters evaluated in the PAT include

- HLW volume,
- LAW volume,
- HLW/LAW weight ratio,
- ^{137}Cs loading in LAW,
- ^{90}Sr loading in LAW,
- TRU loading in LAW, and
- processing rate or schedule.

Additional parameters are needed to represent the interface between retrieval and pretreatment properly, including the volume and activity of in-tank waste remaining after retrieval and the particle size distribution of the waste delivered by retrieval to pretreatment. This latter parameter is important because of the effect of waste particle size on the efficiency of certain pretreatment options and because it affects the likelihood of transfer-line plugging. Quantitative inputs to the PAT include both simple numerical values and mathematical expressions that use several input parameters. These quantitative inputs may be stochastic. Uncertainty is included by using probability density functions or intervals to describe the stochastic behavior. The qualitative inputs include such concepts as the site sensitivity to radionuclide content in the LAW.

The basic output of the PAT is a qualitative measure called "the acceptability" of the pretreatment process. "The acceptability" is a linguistic variable that uses natural language expressions to describe how well a specific combination of retrieval, pretreatment, and disposal options meets the specifications for final waste form, waste removal rate, residual waste, and operability of the process equipment. The acceptability is a function of the waste characteristics, the site characteristics, and the various process parameters.

The PAT algorithm is based on inductive logic models for the parameters needed to represent the pretreatment process. Each of the models evaluates the acceptability of some aspect of the pretreatment process. The individual acceptability evaluations are combined using a series of forward-chaining logical implications to infer the overall acceptability of a particular pretreatment option. This evaluation is based not only on how well the option meets the pretreatment phase design goals but on how it affects the retrieval and immobilization phases as well.

EVALUATING A COMBINATION OF NUMERICAL AND QUALITATIVE INPUT DATA

In designing the inductive logic models for PAT, it became clear that a consistent methodology for combining generic quantitative data with site-specific qualitative judgments of the data would

be necessary. For example, one metric for evaluating the performance of a pretreatment process is the activity level for specific isotopes in the LAW stream. If the loading is too high, the pretreatment option is unacceptable. This site sensitivity is a qualitative judgment based on expert elicitation and may vary among the sites because what is and what is not acceptable is site-dependent. For instance, a particular site may be especially sensitive to ^{137}Cs because of groundwater conditions at the LAW storage site. Thus, the evaluation of the acceptability of the loading for a LAW stream depends on the well-defined measurement of activity in the LAW and the qualitative tolerance of each site for a particular radionuclide. We show here how to combine these two very different types of data in a reproducible and defensible manner.

Logic structures specify how the input data combine to produce acceptability values. Such an inductive logic structure for evaluating the acceptability of the radionuclide levels in the LAW stream is shown in Fig. 1. The primary inputs appear at the left of the diagram. These primary elements are combined pairwise to generate the initial acceptability measures for each species of interest . The species-specific acceptability measures then are combined as well. This forward-chaining continues until the final desired acceptability measure is generated.

In the PAT, the radionuclides of concern are ^{137}Cs, ^{90}Sr, and transuranic isotopes (TRU). The acceptability logic submodule for each of the different radionuclides is identical in the current version of the PAT, so we use the cesium loading acceptability evaluation as an illustrative example. The two inputs to the evaluation are the ^{137}Cs loading in the LAW and the site sensitivity to LAW cesium content. The LAW loading is a numerical value that describes one aspect of the performance of a particular pretreatment option. This is a quantitative input to the AR model and is treated objectively in the evaluation. As noted above, each site will have its own qualitative evaluation of how sensitive it is to the waste loading for a given radionuclide. The isotope loading and the site sensitivity are considered together to produce an acceptability evaluation for the pretreatment option with respect to LAW loading. The relationship between activity and site sensitivity is defined in the PAT by a series of "if–then" rules given in Table I. These rules are of the form

If the Loading is L_i and the Sensitivity is S_j then the Acceptability is C_k (1)

For example, the rule in the upper left corner of the rule base in Table I is

If the ^{137}Cs loading is very low and the site sensitivity for ^{137}Cs is high then the ^{137}Cs acceptability is preferred.

These statements are in the form of a logical implication. If we make the additional statement

The Loading is L_i and the Sensitivity is S_j (2)

then it follows immediately that the acceptability is C_k. The combination of two propositions in the form of Statements 1 and 2 is the *modus ponens* tautology. It is the standard form of inferential reasoning used in AR models.

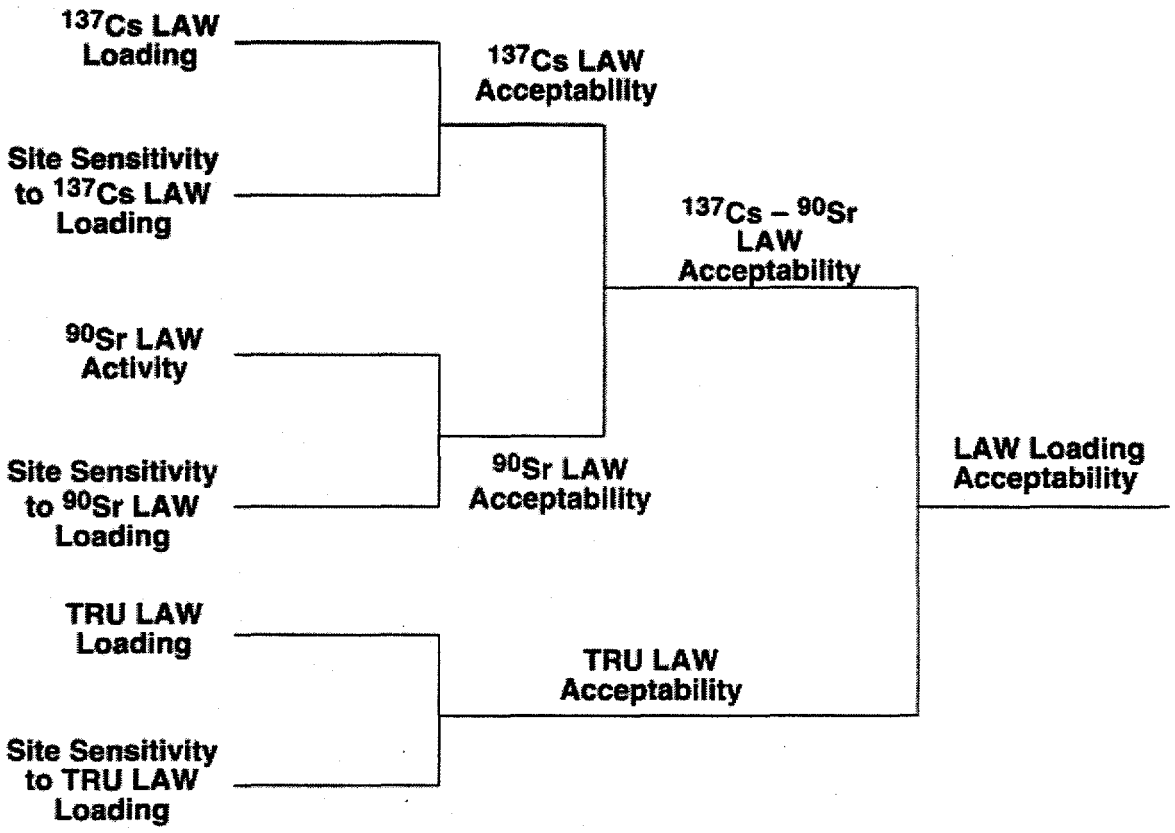

Figure 1. Inductive Structure for Evaluating Acceptability of Low Activity Waste Radionuclide Loadings

Table I. Rule Base for Inferring Site-Specific ^{137}Cs LAW Acceptability from ^{137}Cs LAW Activity and Site Sensitivity

		Site Sensitivity		
		High	**Average**	**Low**
^{137}Cs LAW Loading	**Very Low**	*Preferred*	*Ideal*	*Ideal*
	Low	*Acceptable*	*Preferred*	*Ideal*
	Quite Low	*Undesirable*	*Acceptable*	*Preferred*
	Quite High	*Unacceptable*	*Tolerable*	*Acceptable*
	High	*Unacceptable*	*Unacceptable*	*Tolerable*

All of the inputs and outputs in the rule bases are treated linguistically. The ^{137}Cs loading is described using a set of descriptors {Very Low, Low, Quite Low, Quite High, High}. Each of these descriptors is itself a fuzzy set. The set of descriptors used to represent a variable is called the universe of discourse (UOD) for that variable. The UOD for the site sensitivity is {Low,

Medium, High}, and the acceptability is represented by {Ideal, Preferred, Acceptable, Tolerable, Unacceptable}.

Primary inputs, which may be numerical in nature, must be translated into their correct linguistic variable form so that they can be assigned membership in sets within the UOD for that variable. The translation of numerical input to linguistic variable uses a set of membership functions such as those shown in Fig. 2. There is a membership function for each element in the class of fuzzy sets that span the UOD for residual ^{137}Cs activity. The degree of membership in each set is determined by the membership functions. These functions map a numerical value of the input variable into membership values in the fuzzy sets. Here the membership functions are meant to represent a consensus among experts from different sites on how the natural language descriptors should be related to numerical values for activity. Because the sets used in the PAT are fuzzy, a particular ^{137}Cs activity can have membership in more than one set. To illustrate the use of the membership functions, consider a ^{137}Cs LAW loading of 200 Ci/mL. From Fig. 2, a residual activity of 200 corresponds to memberships of 0.75 in the "Quite High" set, 0.25 in the "High" set, and 0 in the other sets. We denote the degrees of membership in the class of fuzzy sets by a vector, μ. Here $\mu = [0, 0, 0, 0.75, 0.25]$.

Because of the uncertainty in waste properties and the performance of a particular pretreatment process, it will normally be the case that the ^{137}Cs LAW activity should be treated as a random variable. In this case, the degrees of membership also will be random. Monte Carlo simulation techniques are used to find the associated probability density functions (PDFs) for each element

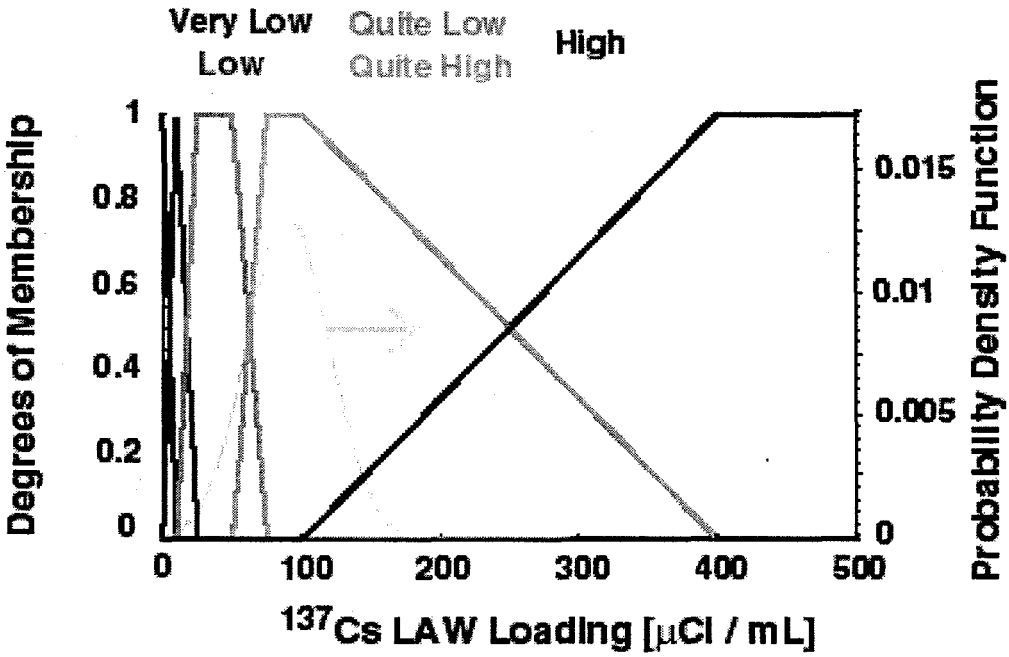

Figure 2. ^{137}Cs LAW Activity Membership Functions

in the μ vector. Suppose that the activity is described by the PDF shown in Fig. 2. Each Monte Carlo trial generates a separate value for the residual activity and therefore separate degrees of membership in the residual activity sets. With a sufficient number of samples, PDFs and the associated quantiles for each residual activity can be estimated. For this example, the median degrees of membership are μ(median) = [0, 0.01, 0.18, 0.74, 0.08] The natural language expression for this result might be "The ^{137}CS LAW loading acceptability is rather high." Degrees of membership for site sensitivity to ^{137}Cs residual activity are assigned directly by an expert familiar with the site. As noted earlier, the UOD for site sensitivity is{Low, Medium, High}, and for this example, we use a degrees of membership vector μ = [0, 0.75, 0.25]. This set membership distribution represents a sensitivity that is meant to be somewhat above the average for all sites. The corresponding natural language expression might be "The site sensitivity for ^{137}Cs LAW activity is fairly high."

With both the ^{137}Cs loading and site sensitivity expressed as degrees of membership in the appropriate fuzzy sets, the rule base defined in Table I is used to infer the memberships in LLWLAW loading acceptability for the site. The degrees of membership for the acceptability sets are calculated using a max-min rule that operates on the antecedent degrees of membership for each rule in the rule base (Ross 1995). The median results from the Monte Carlo simulation are μ (median) = [0, 0.01, 0.18, 0.74, 0.25]. The results of the AR algorithm are interpreted by examining these degrees of membership for the site acceptability. Comparing this vector with the one obtained above based on the numerical value of the cesium activity, μ(median) = [0, 0.01, 0.18, 0.74, 0.25], shows that consideration of site sensitivity has shifted the numerically larger degrees of membership toward the Tolerable and Unacceptable sets. The corresponding natural language expression for the result of the simulation might be "The site-specific acceptability of the ^{137}Cs LAW loading is barely tolerable." This shows the ability of an AR model to make expert-like qualitative judgments.

The binary inference we have just examined (two inputs, one output) is the basic building block of the PAT model. Referring to Fig. 1, the same process is carried out for the ^{90}Sr and TRU LAW activities. The site-specific acceptabilities then are combined using the rule base in Table II. This rule base defines what inferences can be drawn when the separate acceptabilities are considered together. In this case, we have chosen to infer a combined cesium and strontium acceptability first. This acceptability then is combined with the TRU acceptability by applying the rule base a second time. An overall acceptability evaluation is performed by combining a series of logic submodules that consider other tradeoffs of concern. Examples of other submodules include those designed to infer the acceptability of the decontamination factors for the radionuclides and overall cost acceptability. The LAW loading acceptabilities inferred here also can appear in these other submodules.

Table II. Rule Base for Inferring Two-Species Site-Specific LAW Loading Acceptability from Individual Species LAW Loading Acceptabilities

Site-Specific LAW Loading Acceptability for Component B		Site-Specific LAW Loading Acceptability for Component A				
		I	**P**	**A**	**T**	
	I	*I*	*I*	*P*	*A*	*U*
	P	*I*	*I*	*P*	*A*	*U*
	A	*P*	*P*	*A*	*T*	*U*
	T	*A*	*A*	*T*	*U*	*U*
	U	*U*	*U*	*U*	*U*	*U*
I = Ideal, P = Preferred, A = Acceptable T = Tolerable, U = Unacceptable						

EVALUATING COMPLEX PHENOMENA IN THE PAT

Using a binary rule base is practical as long as the logical inferences can be represented reasonably with a series of forward-chaining, two-antecedent structures. There is a general preference for this two-dimensional structure because of the relative ease in eliciting the necessary information from the subject matter experts. However, in some cases, it is advantageous to construct rule bases of higher dimensions. In this section, we consider such a case and show how a model-based approach can be used to construct the rule base and structure the expert elicitation.

An important parameter in most pretreatment processes is the particle-size distribution of the feed. The efficiency of pretreatment options is sensitive to particle size, and flowsheet variations might be required if this characteristic of the delivered waste changed significantly. This implies a close coupling to the preceding retrieval phase. Any change in the pretreatment option will require a re-evaluation of the particle-size distribution produced in the retrieval process. By the same token, placing too tight a limit on the particle-size distribution to simplify processing could lead to an unacceptable tradeoff relative to transfer-line plugging. Note that the particle-size input is not a simple numerical value as was the case for LAW activity level. The input used to describe particle size is a complicated function of several parameters. To show how such a situation can be evaluated, we consider an AR treatment of the particle-size distillation effect.

Sluicing is planned in some options for retrieval. A phenomena known as particle-size distillation may occur during sluicing. Particle-size distillation occurs when different particle sizes with very different settling rates are suspended in a plenum from which the particles are pumped. If there is a finite settling time for the particles before sweep-out occurs, then smaller particles will be removed preferentially from the plenum and larger particles will settle back onto the undisturbed waste. Thus, the waste that is subject to suspension will become enriched in large particles over time.

It is difficult for experts to directly judge the effect of particle distillation in a given situation on the acceptability of a pretreatment option. They base their judgement on the size distribution of the particles generated by suspension and by the settling fractions for the different sizes of particles. They then must internally evaluate the extent of size distillation that takes place. To facilitate their judgment, we developed a simple model for particle-size distillation in sluicing a tank. The purpose of this model is not high physical fidelity. It is to relieve the experts from directly estimating the effect of distillation on retrieval efficiency.

The model used for particle-size distillation is most easily explained with reference to the diagram in Fig. 3. The tank is modeled with flat waste surfaces and a volume of water added above the waste during sluicing. Sluicing is considered a batch process consisting of a series of steps. Each step includes a suspension phase in which waste is broken up and suspended in the plenum, a settling phase during which the waste particles remain in the plenum subject to settling, and an instantaneous transfer phase or sweep-out in which all the particles still in the plenum are discharged from the system and sent through the transfer lines.

In our model, each sluicing step results in the suspension of a mass M of waste in the supernatant liquid plenum shown. This mass (see Fig. 3) is removed uniformly from the top of the waste. M is the sum of the distilled waste and some undisturbed waste particles called "fresh waste suspended." Whenever any undisturbed waste material is suspended, it is assumed to be composed of two sizes of particles. A mass fraction, f, is composed of large particles and a mass fraction, 1-f, is composed of small particles. During the settling phase of each step, a fraction, X_L, of the large particles and a fraction, X_S, of the small particles settles uniformly from the supernatant liquid onto the surface of the undisturbed waste to form a layer of distilled particles. Normally X_S is much smaller than X_L. The layer is called distilled because small particles have been removed preferentially. Of the unsettled particles that remain in the plenum, a fraction, 1-X_L, of the large particles and, 1-X_S, of the small particles, is swept out through the discharge pipe and a new step begins.

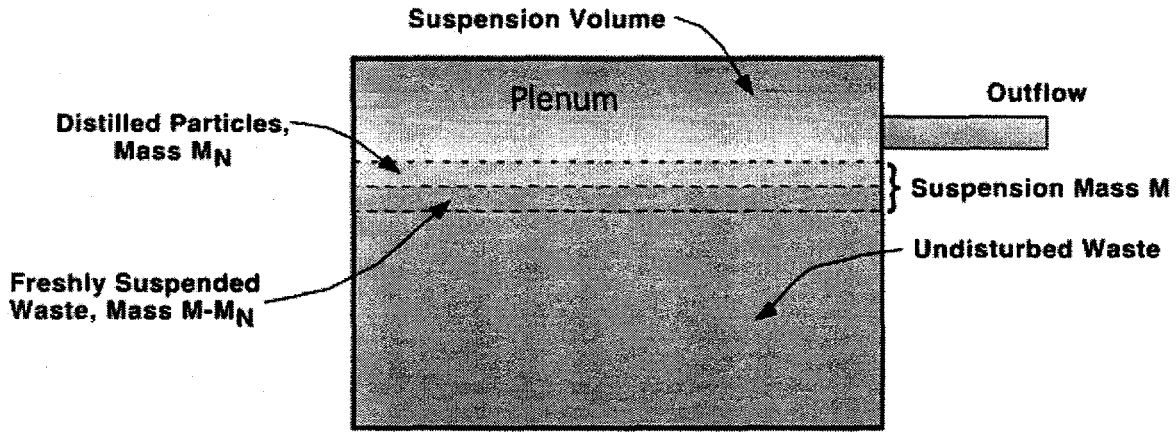

Figure 3. Particle-Size Distillation Model

The picture of the removal process is of an initial transient phase during which a layer of large particles accumulates above the undisturbed waste. Eventually, this layer reaches a steady-state mass that remains constant until all the undisturbed waste is removed. Thereafter, the remaining mass, which is all large particles, is removed according to an exponential decay law because the amount removed at a given step is a fraction of the remaining mass. We will define waste removal efficiency as the ratio of the waste removed to the waste suspended during a sluicing step. In the appendix, we show that this efficiency is

$$\Phi = \frac{1 - X_L - X_S + X_L X_S}{1 - X_L + (X_L - X_S) f} \quad . \tag{3}$$

Equation (3) is a simple approximate model that can be used to provide feedback between the retrieval and pretreatment phases. The behavior modeled here affects both the mass removal rate and the fraction of mass removed that is large particles. The removal rate is of interest because it determines the efficiency of the retrieval process. The fraction of mass in large particles is important because it affects the likelihood of plugging in the transfer lines.

Because the model on which Eq. (3) is based is only approximate, we do not wish to treat Φ as if it were a primary numerical input obtained from a detailed process model. Instead, it is considered to be an intermediate linguistic variable that is to be inferred from X_L, X_S, and f. That is,

$$R(X_L, X_S, f) \Rightarrow \Phi \tag{4}$$

where R is a three-dimensional implication rule base and all of the other symbols are now taken to be linguistic variables. To infer Φ from X_L, X_S, and f, it is necessary to specify R.

The first step in this process is to decide how the antecedents themselves are to be expressed linguistically. All three of the parameters in Eq. (3) and Φ take on values between 0 and 1. Each can be described with the same UOD {Low, Medium, High}. This means that there will be a total of 27 rules in R. Membership in these sets may be either assigned directly by the subject matter experts as with the site-specific sensitivity in the previous example or by using membership functions such as those shown in Fig. 4. Here we use exactly these same membership functions for X_L , f, and Φ. (This is for illustration purposes only.) In practice, the membership functions for each antecedent and the consequent may all be different. Their final forms are arrived at by iteration with the experts. The methodology described here helps facilitate the elicitation and reduce the number of iterations needed. It also ensures that the inferences in R are consistent with the physical model on which they are based.

A rule base, R, that is consistent with Eq. (3) can be developed by evaluating the antecedent triplet (X_L, X_S, f) for selected numerical values of each antecedent. For example, consider the triplet (0.9, 0.5, 0.5). These values were chosen so that X_L has membership 1 in {High}, X_S has

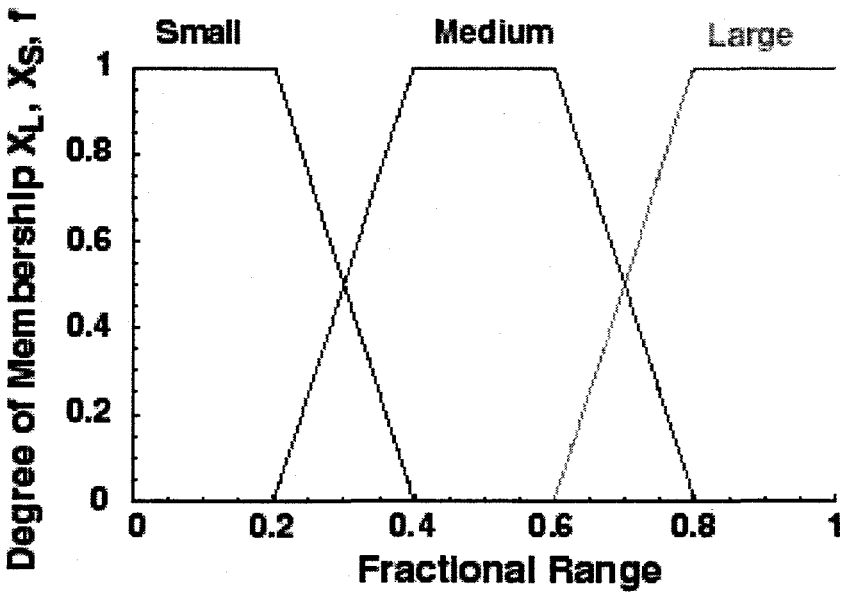

Figure 4. Fraction Membership Function

membership 1 in {Medium}, and f also has membership 1 in {Medium}. If all three antecedents have degrees of membership in only a single fuzzy set, as is the case here, only one rule in R is necessary to infer Φ. The corresponding value of Φ from Eq. (3) is about 0.2. Again referring to Fig. 4, Φ = 0.2 has membership 1 in {Low}. Therefore, the linguistic triplet (High, Medium, Medium) implies that Φ has membership in {Low}. By selecting numerical values for the antecedent triplet that correspond to the 27 distinct set combinations, the complete rule base can be constructed. This is shown in Table III. Each level of this matrix shows the value of the efficiency Φ as a function of the removal fraction values X_L and X_S for a given large particle fraction f. Each level in the matrix is for a different value of f.

This model allows the expert to input his judgments concerning the relative abundance of large- and small-sized particles in the suspended waste and the rates of settling for large and small particles during removal. These estimates are used to infer the effect of size distillation on removal efficiency . The inferential rule base emulates the experts' judgments concerning what constitutes an acceptable level of particle distillation for a given pretreatment process. In practice, other rule bases may be required to completely represent possible tradeoffs associated with sluicing. Additional inputs could include the degree of mixing achieved in a sluicing step and the number of steps required.

SUMMARY AND CONCLUSIONS

A number of system design tradeoffs must be evaluated in the design of a pretreatment process. The degree to which these tradeoffs are resolved needs to be judged before the efficacy of a

Table III. Rule Base for Waste Removal Efficiency

Large Particle Fraction: *Small*				
Large Particle Settling Fraction		**Small Particle Settling Fraction**		
		Large	**Medium**	**Small**
	Large	*Low*	*Medium*	*Medium*
	Medium	*Low*	*Medium*	*High*
	Small	*Low*	*Medium*	*High*
Large Particle Fraction: *Medium*				
Large Particle Settling Fraction		**Small Particle Settling Fraction**		
		Large	**Medium**	**Small**
	Large	*Low*	*Low*	*Low*
	Medium	*Low*	*Medium*	*High*
	Small	*Low*	*Medium*	*High*
Large Particle Fraction: *Large*				
Large Particle Settling Fraction		**Small Particle Settling Fraction**		
		Large	**Medium**	**Small**
	Large	*Low*	*Low*	*Low*
	Medium	*Medium*	*Medium*	*Medium*
	Small	*Medium*	*High*	*High*

particular process can be determined. A prototype PAT was developed to demonstrate how these tradeoff considerations could be performed. The PAT uses AR techniques to provide immediate feedback about the effect of design changes on the overall process. PAT algorithms use fuzzy sets and implication rule bases to emulate the reasoning processes of experts familiar with the system tradeoff issues. The acceptability evaluations for specific aspects of the pretreatment process are aggregated using inductive logic structures to provide an immediate overall evaluation of the acceptability of a proposed design change.

Several issues arose in the course of developing the prototype PAT. Some of the input data used in the evaluation are numerical values; others are highly subjective. Typically, a numerical value and a subjective evaluation of the importance of that value must be combined for an acceptability evaluation. Our approach to combining these inputs uses fuzzy set membership functions to convert the numerical data to linguistic variables that then can be combined with subjective data, also in linguistic form, to produce acceptability measures. The advantage of this approach is that the resultant acceptability judgment is consistent and defensible.

In other situations, input data first must be manipulated using mathematical functions before they are in the proper form to serve as antecedents to infer acceptability measures. Our approach

to this problem has been to develop approximate physical models that define a mathematical relationship between the known quantitative inputs. The inputs derived from the model are taken to be the antecedents for a linguistic variable that can be used directly in an acceptability inference. The dimensionality of the rule base is the number of quantitative inputs for the model. We convert these inputs to linguistic variables whose membership functions are obtained by expert elicitation. A methodology has been demonstrated to ensure that the rule base is consistent with expert opinion and the approximate model used to define the rule base. These two examples provide a methodology for treating a wide and commonly encountered class of problems.

The application of AR as a decision analysis tool in waste management has great promise. The nature and complexity of the tradeoffs that are needed for HLW and other waste processing pose major difficulties for other, more deterministic, methodologies. The application of formal logic and fuzzy sets in the context of a natural language representation of the evaluation problem is well-suited to this class of problems.

REFERENCES

1. S. F. AGNEW, S. W. EISENHAWER, and T. F. BOTT, "Uncertainty and Approximate Reasoning in Waste Pretreatment Planning," International Topical Meeting on Nuclear Waste and Hazardous Waste Management, Spectrum '98, Denver, September 1998.

2. L. ZADEH, "A Theory of Approximate Reasoning," in J. Hayes, D. Michie and L. Mikulich (Eds.), *Machine Intelligence*, Halstead Press, New York, pp. 149–194, 1976.

3. S. W. EISENHAWER, T. F. BOTT, and R. E. SMITH, "An Approximate Reasoning-Based Method for Screening High-Level Waste Tanks for Flammable Gas," WM'98 Proceedings, Tucson, March 1998.

4. S. W. EISENHAWER, T. F. BOTT, T. M. HELM, and S. T. BOERIGTER, "System Reliability with an Approximate Reasoning Model," Proceedings of the International Conference and Workshop on Reliability and Risk Management, pp. 521–543, San Antonio, September 1998.

5. T. F. BOTT and S. W. EISENHAWER, "Including Uncertainty in Hazard Analysis through Fuzzy Measures," ASME Pressure and Piping Conference, San Diego, July 1998.

6. T. J. ROSS, *Fuzzy Logic with Engineering Applications*, McGraw-Hill, New York, 1995.

APPENDIX

Particle distillation results from the differential settling rates of large and small particles mobilized during sluicing. We will model the retrieval process as a series of steps, each with a suspension, settling, and removal phase. At the beginning of the first step in the removal process, all of the

waste is undisturbed. As a result of the suspension phase of this step, a mass, M, is suspended, consisting of a mass, fM, of large particles and a mass, (1-f)M, of small particles. In the settling phase of the first step, a mass of large particles, fXM, settles uniformly onto the undisturbed waste surface. Then in the removal phase of the step, all the remaining suspended mass of large and small particles is swept out of the plenum. At the end of the first step, the waste consists of a layer of large particles of mass fXM (which is less than the suspension mass M) lying on undisturbed waste. The next sluicing step is assumed to suspend all the large particles in the top layer of the waste plus enough undisturbed material to make a total suspended mass of M.

Let us describe the mass of large particles lying on top of the undisturbed waste at step N by M_N. In step N+1, all the mass, M_N, of large particles lying on the waste, plus a mass, $(M-M_N)$, of the underlying undisturbed waste is suspended because a total mass of M is suspended at each step. Now out of this mass, all the M_N plus a fraction, f, of the freshly suspended undisturbed waste mass $(M-M_N)$ is in the form of large particles. The total mass of suspended large particles is then

$$M_{Large} = M_N + f(M - M_N) \quad . \tag{A-1}$$

Now a fraction, X, of this mass settles on top of the undisturbed waste during the settling phase of the step. Thus, the mass of large particles lying above the undisturbed waste at step N+1 is given in terms of M_N by

$$M_{N+1} = XM_N + Xf(M - M_N) \quad . \tag{A-2}$$

The mass of large particles lying above the undisturbed waste will increase at each step until the amount of large particle mass removed at each step is equal to the amount suspended during that step. This will occur when

$$(1 - X)\left[M_N + f(M - M_N)\right] = f(M - M_N) \quad . \tag{A-3}$$

Solving Eq. (A-3) for MN gives

$$M_N = \frac{Xf}{1 - X + Xf} \quad . \tag{A-4}$$

Thus, the picture of the removal process is an initial transient period during which a layer of large particles accumulates above the undisturbed waste. Eventually, this layer reaches a steady-state mass that remains constant until all the undisturbed waste is removed. The remaining mass, which is all large particles, is removed according to an exponential decay law because the amount removed at a given step is a fraction of the remaining mass.

When steady state is reached, the mass removal for a step is

$$M_{out} = (1 - X)M_{Large} + M_{Small} \quad . \tag{A-5}$$

where M_{Large} and M_{Small} are the suspended mass of large and small particles respectively. The expression of the suspended large particle mass is (A-1), whereas that for the small particles is

$$M_{Small} = (1-f)(M - M_N) \ . \tag{A-6}$$

Substituting Eqs. (A-1) and (A-6) into Eq. (A-5) and simplifying leads to

$$M_{out} = (1 - fX)M - X(1-f)M_N \ . \tag{A-7}$$

Substituting Eq. (A-4) into Eq. (A-7) for M_N and simplifying we get

$$M_{out} = \frac{1-X}{1-X+Xf} M \ , \tag{A-8}$$

showing that the mass out flow is decreased by particle-size distillation below the ideal value of M in which all suspended particles are removed at each step. We define the waste removal efficiency, Φ, as the ratio of the mass suspended to that removed:

$$\Phi = \frac{1-X}{1-X+Xf} \ . \tag{A-9}$$

If we assume that both the small and large particles settle out with fractions X_S and X_L, respectively, then the approach used in deriving the previous formulae can be used to generate new expressions. Now instead of one layer of large particles lying above the undisturbed waste, we have a layer of large and small particles. Two coupled expressions are derived for the masses of large and small particles at steady state analogous to Eq. (A-2). These expressions lead to expression for the waste removal efficiency:

$$\Phi = \frac{1 - X_L - X_S + X_L X_S}{1 - X_L + (X_L - X_S)f} \ . \tag{A-10}$$

This waste removal efficiency is a measure of how much waste is removed during each step in the process compared with the mass suspended by the retrieval process. A perfectly efficient process would have a Φ of 1, and all of the suspended mass would be removed.