Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Modeling of Transport Risk for Hazardous Materials

Erhan Erkut, Vedat Verter,

MODELING OF TRANSPORT RISK FOR HAZARDOUS MATERIALS

ERHAN ERKUT

University of Alberta, Edmonton, Alberta, Canada

VEDAT VERTER

McGill University, Montreal, Quebec, Canada

(Received June 1995; revision received June 1996; accepted April 1997)

The transport of hazardous materials is an important strategic and tactical decision problem. Risks associated with this activity make transport planning difficult. Although most existing analytical approaches for hazardous materials transport account for risk, there is no agreement among researchers on how to model the associated risks. This paper provides an overview of the prevailing models, and addresses the question "Does it matter how we quantify transport risk?" Our empirical analysis on the U.S. road network suggests that different risk models usually select different "optimal" paths for a hazmat shipment between a given origin-destination pair. Furthermore, the optimal path for one model could perform very poorly under another model. This suggests that researchers and practitioners must pay considerable attention to the modeling of risks in hazardous materials transport.

The transportation of hazardous materials (hazmats) is an important problem in industrialized societies, due to the pervasiveness of these materials. Hazardous materials, or dangerous goods, include explosives, gases, flammable liquids and solids, oxidizing substances, poisonous and infectious substances, radioactive materials, corrosive substances, and hazardous wastes. For most members of industrial societies, life without hazmats is inconceivable. Unfortunately, most hazmats are not used at their point of production, and they are transported over considerable distances. It is estimated that between 250,000 and 500,000 hazmat shipments take place in the United States every day, resulting in an annual total shipment volume of between 1.5 and 4 billion tons. To put this number into perspective, according to the U.S. Department of Commerce (1994), roughly every fifth truck on U.S. highways is a hazmat truck.

What differentiates shipments of hazmats from shipments of other materials is the risk associated with an accidental release of these materials during transportation. Hazmats can be extremely harmful to the environment and to human health, since exposure to their toxic chemical ingredients could lead to the injury or death of plants, animals, and humans. This risk is recognized by society, and in many instances strict regulations govern the movement of hazmats. Consequently, hazmat carriers have better accident records than other carriers. Nevertheless, even if they are a rare occurrence, accidents do happen during the transportation of hazmats. For example, in 1979 a train carrying toxic chemicals was derailed in Mississauga, Ontario, and chlorine leaking from damaged tank cars forced the evacuation of 200,000 people. In a particularly gruesome accident in 1982, 2,700 fatalities were reported due to a gasoline truck explosion in a tunnel in Afghanistan. Major gasoline shipments by truck were prohibited in Germany after a 1987 six-fatality accident involving a gasoline truck and an ice cream parlor. Although the average annual number of fatalities due to hazmat transport accidents is dwarfed by the annual number of fatalities due to regular traffic accidents (10 vs. 35,000 in the United States), hazmat transport accidents receive special attention by the media, which sensitizes the public, and consequently the regulators, to the danger of transporting hazmats.

The modeling of transport problems is a popular application area in OR. In most transport planning models, the objective is to move products from origins to destinations at minimal cost. However, for hazmat shipments, a cost-minimizing objective is usually not suitable. The risk associated with hazmats makes these problems more complicated (and more interesting) than many other transport problems. Consideration of risk in hazmat transport models is not merely an academic exercise. The Hazardous Materials Transportation Uniform Safety Act of 1990 requires the U.S. Department of Transportation to determine standards for designating the routes to be used for hazmat shipments. Clearly, a consideration of societal risks is important in such strategic decisions.

In our opinion, OR can play a significant role in the strategic and tactical decision processes surrounding

Subject classifications: Transportation: travel route choice, risk.
Area of review: OR Chronicle.

0030-364X/98/4605-0625 $05.00

hazmat transport. This is not a revelation; there is already a considerable amount of OR literature in this area. However, there is no agreement among researchers on the proper representation of the associated transport risks. The goal of this paper is to provide some insight into hazmat transport modeling by considering different ways of modeling transport risk and attempting to answer the question "Does it matter how we model transport risk?" We begin our paper by discussing the application of the traditional definition of risk to hazmat transport in an OR context. Then we present alternative risk models from the literature and provide a bicriterion approach for hazmat route selection. In the empirical part of the paper, we compare the hazmat routes selected by different risk models and discuss similarities and dissimilarities between these routes.

1. MODELING OF RISK FOR HAZMAT TRANSPORT

Although risk is a popular term in the media, and a popular topic with many authors, there is no universally accepted definition of risk. Most people would agree that risk has to do with the probability and the consequence of an undesirable event. Although some authors define risk as only one of these terms (i.e., probability or consequence), it is more common to define risk as the product of both the probability of and the consequence of the undesirable event (Covello and Merkhofer 1993). Note that this is an "expected consequence" definition, and it is the definition that we refer to as "traditional risk" in this paper (primarily for the reason that it is the definition used in the U.S. Department of Transportation 1989 guidelines for transporting hazmats, which have influenced many researchers in this area). We emphasize that, depending on the circumstances, it might make sense to use other definitions of risk. In fact, the primary objective of this paper is to explore the different ways in which risk can be modeled in the context of hazmat transport.

In the case of hazmat transport, an undesirable event is an accident that results in the release of a hazardous substance. This is usually called an *incident*. Although there can be many undesirable consequences of an incident (such as damage to wildlife, economic losses, and injuries), almost all the literature in this area is concerned with fatalities. Hence, it is common to assume that the undesirable consequence is proportional to the size of the population in the neighborhood of the incident, where the size of the neighborhood depends on the substance carried. Furthermore, the probability of an incident occurring depends on the substance carried and the road type.

Clearly, the risk associated with transporting a hazardous material depends not only on the substance being transported but also on the road network characteristics, such as road type and population, along the chosen route. Hence, we develop a traditional risk model for hazmat

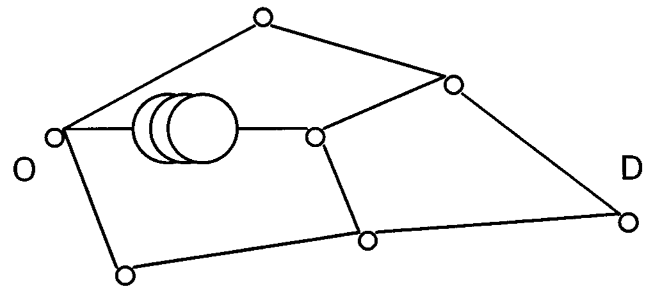

Figure 1. Depiction of hazmat transport between an origin (O) and a destination (D) as the movement of a danger circle on a network.

transport by beginning with its building blocks, namely the unit road segment risk and the edge risk.

1.1. Unit Road Segment Risk

Using the traditional risk definition, the risk of transporting hazmat B over a unit road segment A (such as a one-mile stretch) can be written as:

$$R_{AB} = p_{AB} C_{AB}, \tag{1}$$

where

p_{AB} = probability of an incident on the unit road segment A for hazmat B, and

C_{AB} = population along the unit road segment A within the neighborhood associated with hazmat B.

For most hazmats, estimates of incident probabilities are between 0.1 and 0.8 per million miles. The consequence depends on population density near the road segment and can be estimated by using the impact area for the hazmat under consideration. The impact area can be represented as a circle and the radius of the circle can be anywhere from 0 to 7 miles. (CANUTEC 1992 suggests an initial evacuation radius of 0.5 to 1 mile for most hazmat fires.) If we view the impact area as a "danger circle", we can visualize the hazmat transport activity as the movement of this danger circle along a path in a road network from an origin to a destination, as illustrated in Figure 1. Clearly, the movement of the danger circle carves out a band, on both sides of a road, which is the region of possible impacts.

The use of this danger circle in estimating the consequence of a hazmat incident requires some justification. This is, in fact, a simplification necessitated by data limitations. To illustrate this, consider the modeling of the effect of a ruptured container of pressurized ammonia at a given location and time. To accurately estimate the number of fatalities that could result from this rupture, we need to first estimate the concentration of the gas as a function of the rate and type of release, the distance from the container, the existing winds, and the topography. Depending on the velocity and the density of the gas, three different dispersion models must be applied to estimate the concentration at different distances (Glickman and Raj 1991). We then need to estimate the probability of human fatality as a function of ammonia concentration. Finally, multiplying these probabilities with the number of people residing at different distances from the container and aggregating, we can find an expected number of fatalities as a result of this incident.

It is easy to see that this estimation process requires a lot of information. It requires an accurate knowledge of meteorological conditions and topography, the effect of ammonia on humans, and the location of individuals at the time of the release. Although it might be possible to estimate the risk for a specific incident (by fixing the model parameters, such as the rate of release and wind direction), it is quite unrealistic to expect the generation of this information for all links of a transport network for route planning purposes. Furthermore, the human dose-response relationship for many toxic chemicals is not known, making it difficult to estimate under what conditions fatalities will occur. Hence, although theoretically possible, it is practically impossible to produce accurate estimates of incident consequences. This is why we need to use the concept of "danger circle" in the above risk model. This can be viewed as a worst-case scenario, where we assume that every individual residing in the danger circle will be subject to the same undesirable consequence (death, in the worst case) regardless of the distance to the incident, meteorological conditions, topography, etc.

If we use the danger circle approximation, we must estimate the population within a certain distance of the road. This is difficult since the exact population in an area depends on the time of the day, as considerable population shifts occur due to travel to-and-from work. Since most population estimates are based on census counts, they might not provide accurate estimates for daytime population. Hence, if a hazmat shipment is scheduled to pass through a region during the day, then the population-at-risk estimate of the danger circle will be inaccurate. This is not a limitation imposed by the danger circle approximation, but a limitation imposed by the availability of data. Nevertheless, it is worth mentioning here since it could introduce inaccuracies to models based on the danger circle concept or, in general, to models that require location-specific population estimates. (If daytime population data are available, or they can be estimated accurately, they should be used instead of census counts when dealing with daytime hazmat shipments. If hazmat shipments can occur during the day or night, then it may be best to use a weighted combination, or the maximum, of the daytime and nighttime populations.)

Having argued for the necessity of the danger-circle approximation, we now turn to the use of the above model for unit road segment risk in route evaluation and selection models. Such models are network optimization models, where roads are represented as edges of the network. In the context of hazmat routing it is desirable for an edge to be relatively uniform in its two important attributes: incident probability and population density around the road. For example, a long stretch of a highway that goes through a series of population centers and farmland should not be represented as a single edge, but as a series of edges. Thus, a network to be used for hazmat routing is usually different from a network to be used for other transport planning purposes. This difference is quite important since it limits the portability of network databases between different transport applications.

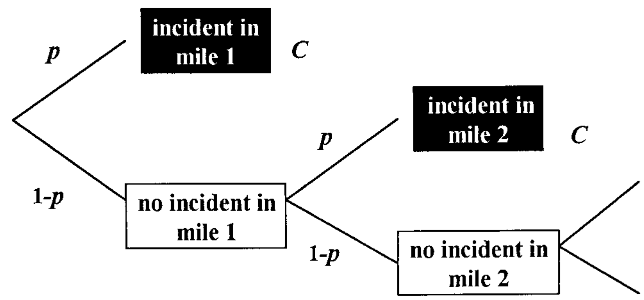

Figure 2. Partial probability tree displaying possible outcomes of a hazmat transport along an edge, where p = incident probability per mile and C = population impacted in the case of an incident.

1.2. Edge Risk

Noting that (1) is valid for a unit road segment, we now move on to the computation of risk along an edge of a hazmat routing network. Due to the construction of the network, it is safe to assume that an edge is a collection of n unit road segments (e.g., miles) each with the same incident probability p and population in (the danger circle) C. (For simplicity, we suppress the indices for the hazmat type and road type.) The vehicle will either have an incident in the first mile, or it will make it safely to the second mile. If it makes it safely to the second mile, it will either have an incident in the second mile, or it will not, and so on. We assume that the trip ends if an incident occurs. The tree in Figure 2 displays all possible outcomes of a trip along this edge.

The expected consequence (risk) associated with this edge would be expressed as follows:

$$pC + (1-p)pC + (1-p)^2pC + \cdots + (1-p)^{n-1}pC. \tag{2}$$

Note that, given p, n, and C, we can easily compute (2) and attach this value to the edge as an edge attribute (impedance) to be used in route selection problems. However, it is possible to simplify (2) considerably by using an approximation:

$$p^s \cong 0, \quad \text{for } s > 1. \tag{3}$$

This approximation is very reasonable since p is at most on the order of one-in-a-million-miles. With this approximation, the risk associated with a trip on this edge becomes:

$$(pn)C. \tag{4}$$

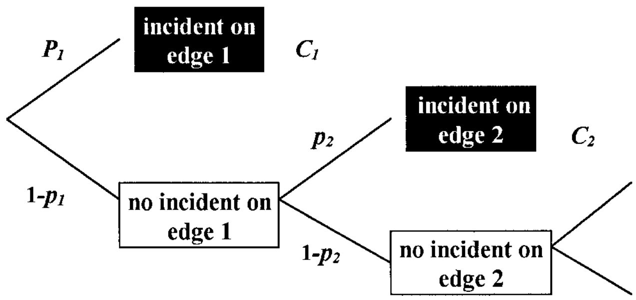

Figure 3. Partial probability tree displaying possible outcomes of a hazmat transport along a path, where p_i = incident probability along the ith edge of the path and C_i = population impacted by an incident on the ith edge.

Defining

p_i = probability of an incident on edge i = (prob. of an incident on a unit segment of edge i)(no. of unit segments in edge i), and

C_i = the number of people within the danger circle along edge i, we can express the risk of one trip along edge i as

$$R_i = p_i C_i. \tag{5}$$

Note that we need uniform incident probability and uniform population density along an edge to express the edge risk in this simple form. (If the network is designed in a way that these two attributes are not uniform on the edges, then the computation of edge risk becomes more complicated, and it may become impossible to use other definitions of risk, such as perceived risk, which is introduced in the next section.) This limitation on edge definition is why the construction of a hazmat transport network is more complicated than the construction of a transport network for cost minimization to be used in the case of ordinary (nonhazardous) materials.

Although the approximation in (3) simplifies the calculation of edge risks, it is not crucial since we can use (2) directly to calculate edge risks. However, it becomes crucial when we simplify the path-risk minimization problem to that of a shortest path problem, as we will see next.

1.3. Path Risk

Having expressed the risk associated with travel on a unit road segment and on an edge in (1) and (5), we now address the following question: How would we express the risk for an entire path between an origin and a destination? A path is a collection of edges, and we can view the travel on the path as a probabilistic experiment, as displayed in Figure 3.

The expected consequence (risk) associated with this trip would be expressed as follows:

$$p_1C_1 + (1 - p_1)p_2C_2 + (1 - p_1)(1 - p_2)p_3C_3 + \ldots. \tag{6}$$

We can view $(1 - p_1)(1 - p_2) \cdots (1 - p_{k-1})p_kC_k$ as the edge attribute (impedance) of the kth edge. Note that the

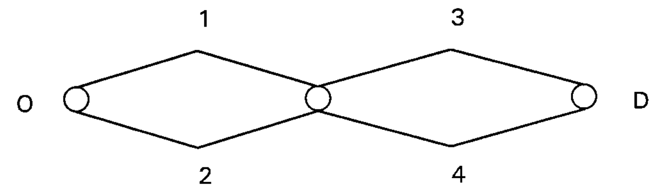

Figure 4. Example demonstrating path-dependency of edge attributes in hazmat transport problems.

edge attributes are path-dependent. To clarify this, consider a very small path selection problem with four edges (Figure 4).

Defining $x_j = 1$ if edge j is on the path, 0 otherwise, we can write the following path selection problem with a risk-minimization objective:

$$\begin{aligned}\min\ & (p_1C_1)x_1 + (p_2C_2)x_2 \\ & + [(1 - p_1)x_1 + (1 - p_2)x_2](p_3C_3)x_3 \\ & + [(1 - p_1)x_1 + (1 - p_2)x_2](p_4C_4)x_4,\end{aligned} \tag{7}$$

$$\text{subject to} \quad x_1 + x_2 = 1, \tag{8}$$

$$x_1 + x_2 = x_3 + x_4, \tag{9}$$

$$x_3 + x_4 = 1, \tag{10}$$

$$x_j = (0, 1) \quad \text{for } j = 1, \ldots, 4. \tag{11}$$

Note that the objective function contains products of the decision variables. This is necessary since the incident probabilities (and hence the risks) of edges 3 and 4 depend on whether the vehicle takes edge 1 or edge 2 in the first part of the trip. However, it is also undesirable since it results in a nonlinear integer programming (NIP) problem. This problem can be rewritten as an linear integer (LIP) programming problem using a substitution of variables. However, the resulting LIP will have considerably more variables than the NIP, and it does not offer a particularly attractive way to solve the route selection problem. On the other hand, an assumption similar to (3) (namely $p_ip_j \cong 0$, for all i, j) is sufficient to achieve a considerable reduction in problem difficulty. Using this approximation (and the constraint set), the objective function can be rewritten as:

$$\min \sum_{j=1}^{4} (p_jC_j)x_j. \tag{12}$$

Thus, the risk minimization problem can be (approximately) solved by using a shortest path algorithm. Note that (12) is merely the sum of expected values of incident consequences (expected number of people impacted per trip).

As pointed out in Erkut and Verter (1995a), almost all the papers on hazmat transport use this approximation, without always explicitly stating it. (Sometimes this approximation is stated as $(1 - p_i) \cong 1$, for all i.) We would like to focus on this approximation since it is so crucial in making the risk minimization problem tractable. How large an error is introduced into the solution by this approximation? By ignoring the higher-order terms from the analysis, we are overestimating incident probabilities. In

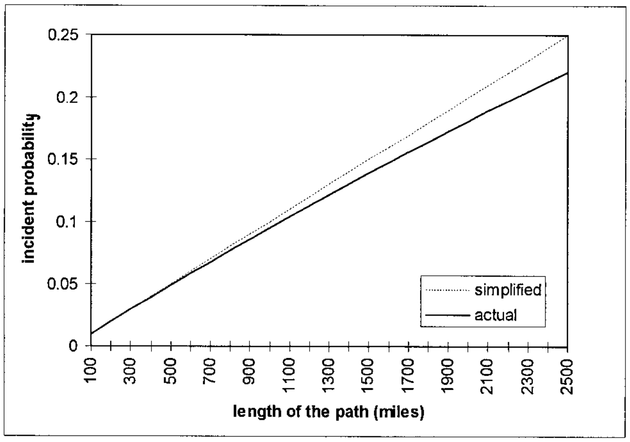

Figure 5a. Plot of actual incident probability and approximated (simplified) incident probability as a function of distance.

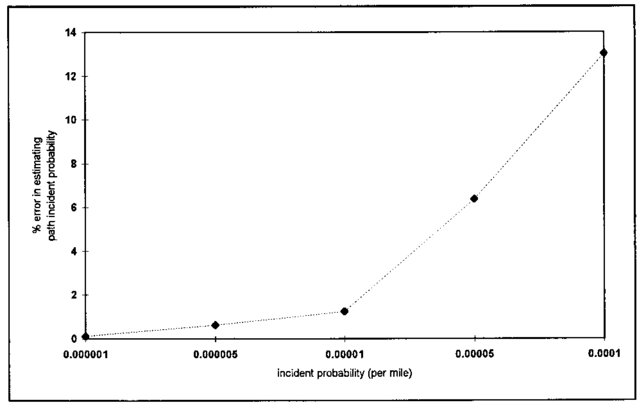

Figure 5b. Plot of percentage error in estimation of path incident probability (for 2,500-mile path) as a function of per-mile incident probability.

fact, the further the trip stretches the larger the overestimation. Does the overestimation ever become sufficiently large to invalidate the results found using this approximation?

We attempt to answer this question using a numerical example. Consider a path that consists of 2,500 unit road segments (miles). Assume $p_i = p$ for all i. Figure 5a shows a plot of the actual and the approximated incident probabilities on this path as a function of the trip length. The incident probability used is 10^{-4} per mile, and the distance, d, varies from 0 to 2,500 miles. The approximated probability is a linear function of distance, namely pd, and the actual probability is a concave function of d.

Figure 5a suggests that the errors introduced may be significant. However, actual incident probabilities are much smaller than 10^{-4} per mile. In fact, when we use a probability of 10^{-6} per mile, the two functions become indistinguishable when plotted (which is why we used 10^{-4} per mile in Figure 5a). With a unit incident probability of 10^{-6} the actual probability of an incident on a 2,500-mile path is 0.00249688 while the linear approximation yields 0.0025 (a difference of only 0.125%). Furthermore, the largest unit road segment incident probability error (which is made on the last mile of the trip) is only 0.25%. Hence, it seems that the errors introduced by the approximation are negligible for all relevant trip lengths and incident probabilities. Figure 5b shows percentage errors in estimating the incident probability on a 2,500-mile path as a function of per-mile incident rates. Note that for incident probabilities below 10^{-5}, the percentage error is under 2%.

In the above paragraph we concentrated on errors introduced by our approximation of incident probabilities and found them to be insignificant. Yet there is another source of error in the edge incident probability calculations: the discretization of the incident process. Note that in Figure 2, the hazmat transport activity is viewed as a probabilistic experiment, where there are two outcomes at each stage. The probabilities are attached to unit road segments and the experiment is repeated as many times as there are unit road segments on the edge. However, in reality the transportation of hazmats is a continuous activity, and the incident probabilities should be interpreted as incident "rates." The discrete probabilistic model estimates the probability of having one incident along a 2,500-mile path as: $p + (1 - p)p + (1 - p)^2 p + \cdots + (1 - p)^{2499} p$. Note that this is a finite geometric series that is equal to $1 - (1 - p)^{2,500}$. Suppose that we wish to increase the accuracy of the probability estimate through a finer discretization of the trip that uses shorter (and more) road segments. If we increase the number of road segments by a factor of k (to nk), the incident probability on each segment becomes (p/k), and the geometric series becomes $1 - (1 - p/k)^{nk}$. To compute the exact probability, we need to take the limit of this series as k goes to infinity. This limit is $1 - e^{-pn}$, which can be expressed as $pn - (pn)^2/2! + (pn)^3/3! + \ldots$ by using a Taylor's series expansion. Clearly, the first term is the only significant one, and higher-order terms are near zero. Using the first two terms, the "exact" probability is computed to be 0.00249688, which is the same as the discrete approximation to eight significant digits. Thus, we can conclude that, even for very long trips, the discretization does not introduce significant errors into the calculation of incident probabilities since the relevant incident probabilities are very small (around one incident per million miles).

To summarize, with an incident probability of 10^{-6} incidents per mile and a trip length of 2,500 miles, the continuous and the discrete probabilistic models both compute an incident probability of 0.00249688 for the trip. The linear approximation results in a slight overestimation, computing a probability of 0.0025. Note that using the linear model amounts to taking into account only the first term of the Taylor's series resulting from the continuous model and ignoring the higher-order terms. Thus, the error is limited to the sum (with alternating signs) of the higher-order terms of the Taylor series expansion. Given the magnitude of the incident probabilities, this seems reasonable. Hence, we can express the (traditional) risk of transporting a prespecified hazmat along path P as follows:

$$\mathrm{TR}(P) = \sum_{i \in P} p_i C_i. \tag{13}$$

2. ALTERNATIVE RISK MODELS FOR HAZMAT TRANSPORT:

Having discussed the approximations necessary for the use of the traditional risk model in hazmat-path selection models, we now turn to the validity of this model in the context of hazmat transport. One objective of modeling is to accurately represent the underlying problem environment. While risk is an issue in hazmat transport, due to the disutility imposed on society, it is not clear that the traditional definition of risk provides a proper representation of this disutility. The multiplication of a probability with a population figure, and the resulting "expected value," might not mean much to the population at large. A simpler concept may be "population exposure," the total number of people exposed to risks during a transport activity. Another relevant concept is simply "incident probability." (Note that these two concepts can be thought of as extreme points of a "risk spectrum" that includes the traditional risk model.)

It could be argued that the population exposure model is one proper way to model the public's *perceived* risk. This model has been used by ReVelle et al. (1991) in a study commissioned by the U.S. Department of Energy. It may be particularly relevant if the hazmat being carried imposes an exposure risk, rather than an incident risk, on the population near the route (for example, radiation from high-level nuclear waste). In this scenario, we can think of p as the probability of exposure for people living near the road, and we can assume that it is equal to one (i.e., exposure to minimal radiation with certainty). Although exposure to minute amounts of radiation for very short periods of time is not likely to pose health risks, public perceptions surrounding nuclear risks are quite different from the assessments of scientists. It is reasonable to expect that public opposition to shipments will be proportional to the size of the population exposed. Hence, selecting routes that minimize exposure may, in turn, result in the minimization of public opposition.

At the other end of the spectrum, if we ignore the variation in population density, or assume that all population densities are equal to some constant (in the danger circle), then the traditional risk model simplifies to the incident probability model. This version may be suitable if the hazmat carried has a very small danger radius. In this case, a relevant objective may be to minimize the incident probability in order to minimize risks imposed on drivers on the road, as well as incident-related costs. This model has been adopted by Saccomanno and Chan (1985). In defense of this approach, one may suggest that each hazmat incident receives a considerable amount of negative publicity independent of its consequences, and hence the likelihood of an incident should be minimized to mitigate public opposition.

In addition to these three closely related models, alternative models have been suggested in the literature to more precisely model the risk in hazmat transport. The traditional "expected consequence" representation of risk might be deemed inappropriate for hazmat logistics since this representation assumes a risk-neutral public. Most people, however, would judge a low-probability–high-consequence event as more undesirable than a high-probability–low-consequence event even if the expected consequences of the two events are equal. This risk-averse attitude is not reflected by the traditional representation of risk, although it is prevalent in the decisions involving hazmats. Abkowitz et al. (1992) suggested modeling perceived risk (for a unit road segment), PR, by using a risk preference parameter q in the following form:

$$\mathrm{PR} = pC^q. \tag{14}$$

Using this approach, the perceived risk associated with one trip along path P, PR(P), can be expressed as follows:

$$\mathrm{PR}(P) = \sum_{i \in P} p_i (C_i)^q. \tag{15}$$

A risk-averse attitude regarding the consequences of potentially hazardous activities can be represented by using $q > 1$. When $q > 1$, the edge impedances in the route selection problem would be increased in proportion to the population along the edges, and edges with larger population densities would become less attractive. Consequently, the model would tend to select edges that go through sparsely populated regions.

The above model can also be thought of as a simple disutility function for perceived societal risk. It provides us with a pragmatic way of accounting for risk aversion; for a given value of q, the perceived risk of an action can be quantified and used as data in an optimization model. Note that this model reduces to the traditional risk model when $q = 1$, which is sometimes also referred to as the "technical risk." Clearly, it is possible to develop other models for modeling the disutility of risk. However, this approach has not been popular with researchers who deal with hazmat route-selection problems. In fact, we know of only two papers that use a proper risk disutility function (from a multiattribute utility theory perspective) for hazmat route selection (Kalelkar and Brooks 1978, McCord and Leu 1995).

The traditional representation of risk is also criticized for not being appropriate in the case of multiple hazmat shipments between an origin and a destination. The expected consequence definition assumes that a selected path can be taken as many times as necessary regardless of the number of accidents on that path. However, if a catastrophic accident should occur during the transportation of an extremely hazardous substance, further shipments may be suspended to allow for a reevaluation of the routing

Table I
Summary of the Five Risk Models Suggested in the Literature for Hazmat Transport Risk

Approach	Model	Sample References
Traditional Risk	$\text{TR}(P) = \Sigma_{i \in P} p_i C_i$	Alp 1995, Erkut and Verter 1995b
Population Exposure	$\text{PE}(P) = \Sigma_{i \in P} T_i$	ReVelle et al. 1991, Batta and Chiu 1988
Incident Probability	$\text{IP}(P) = \Sigma_{i \in P} p_i$	Saccomanno and Chan 1985, Abkowitz et al. 1992
Perceived Risk	$\text{PR}(P) = \Sigma_{i \in P} p_i (C_i)^q$	Abkowitz et al. 1992
Conditional Risk	$\text{CR}(P) = \Sigma_{i \in P} p_i C_i / \Sigma_{i \in P} p_i$	Sivakumar et al. 1993a, 1993b, 1995

Note that we use T_i in the second model, which is the total population in the impact region along edge i. $T_i = 2C_i d_i / \pi r$, where d_i is the length of edge i and r is the radius of the danger circle.

policy. Sivakumar et al. (1993a, 1993b, 1995) proposed the use of conditional risk (i.e., expected consequence given the occurrence of the first accident) to deal with this type of routing problem. Using our notation, the conditional risk for path P, CR(P), can be expressed as:

$$\text{CR}(P) = \frac{\sum_{i \in P} p_i C_i}{\sum_{i \in P} p_i}. \tag{16}$$

By minimizing CR(P) we would minimize the expected consequence at the time of the first incident.

As the above discussion shows, researchers have attempted to model the risks associated with hazmat transport in many different ways. We summarize the five different risk models discussed so far in Table I. Traditionally, these models have been justified by intuition. To provide a more structured basis for discussion, we offer three axioms in the next section.

We should point out that there is a fundamental difference between the first four models in Table I and the conditional risk model. The first four models are single-attribute models. (Although the traditional risk model and the perceived risk model contain two attributes, these two attributes are preprocessed, and a new attribute is defined for each edge.) In contrast, the conditional risk model can be viewed as a multiplicative multiattribute model, where the first attribute is traditional risk and the second attribute is incident probability. Unlike the first four models, this model does not result in an additive objective function.

3. AN AXIOMATIC APPROACH TO THE MODELING OF HAZMAT TRANSPORT RISK

We now state two general axioms for general path evaluation and selection models, and one axiom for transport risk models. We suggest that a credible hazmat path selection model should satisfy all three of these axioms. At the end of this section, we point out an exception where a credible multiattribute model (that combines risk and cost, for example) may violate one of these axioms. Figure 6 may help to clarify the ideas behind the first two axioms.

Let **P1** denote the set of all paths between O' and D', and **P2** denote the set of all paths between O and D. Let $P1 \in \mathbf{P1}$ and $P2 \in \mathbf{P2}$ such that $P2 = \{PA, P1, PB\}$ as in Figure 6. Define V to be an evaluation function that operates on paths (such as the distance function).

Axiom 1. *Monotonicity axiom for path evaluation models* (Erkut 1995):

$V(P1) \leqslant V(P2)$.

Axiom 2. *Optimality principle for path selection models:*

$V(P2) = \min_{P \in \mathbf{P2}} V(P) \Rightarrow V(P)1 = \min_{P \in \mathbf{P1}} V(P)$.

For the third axiom, we assume that the evaluation function in question is a measure of risk. Let $V(P) = R(p(P), C(P))$, where $R(\cdot)$ is a risk function, $p(\cdot)$ is the vector of edge incident probabilities, and $C(\cdot)$ is the vector of edge consequences for the hazmat of concern.

Axiom 3. *Monotonicity axiom for risk models:*

$p(P) \leqslant p'(P)$ and $C(P) \leqslant C'(P)$

$\Rightarrow R(p(P), C(P)) \leqslant R(p'(P), C'(P))$.

One immediate implication of the first axiom can be phrased in a nontechnical way as follows: As we add one (or more) link(s) to an existing path, the total impact (value) of the path can get no smaller. This seems like a reasonable requirement for all path evaluation models. The second axiom requires that all subpaths of an optimal path should themselves be optimal. It is merely a restatement of Bellman's optimality principle in the context of path selection models. The third axiom states that path risk is a nondecreasing function of edge incident probabilities and edge consequences. This implies that in a valid risk model, increased probability or consequence on an edge cannot result in reduced path risk.

Among the five models given in Table I, four of them (namely the traditional risk model, the perceived risk

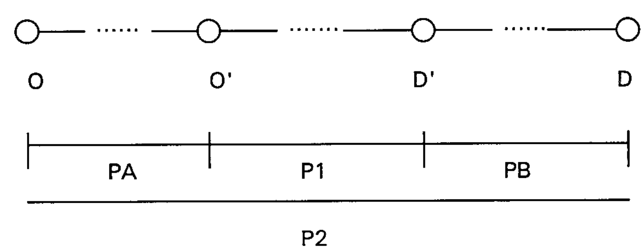

Figure 6. Graphical representation of notation used in axioms 1 and 2.

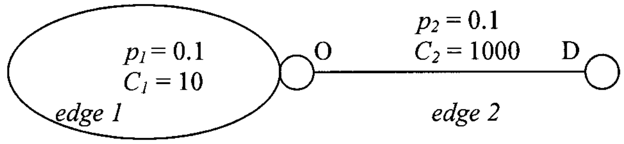

Figure 7. Simple example demonstrating the effect of looping on the total expected consequence of a hazmat trip.

model, the incident probability model, and the population exposure model) satisfy all three of these axioms, while the conditional risk model violates all three of them. The violation of these axioms is not merely a theoretical matter, as it may result in unreasonable solutions to hazmat route selection problems. Erkut (1995) provides illustrative examples where the violation of the axioms stated above by the conditional risk model results in indefensible hazmat route selections. The difficulty with this model is apparent in its functional form. Since the denominator is equal to incident probability, routes with high incident probabilities will tend to have low conditional risks. Consider a hypothetical case where all routes between an origin and a destination have the same expected consequence (the numerator). In this case, this model would act in a suicidal manner and choose the route with the highest incident probability. Furthermore, increasing the incident probability along an edge with zero population (for example, by sabotaging a bridge) will make routes that use this edge more desirable. Since we find the behavior of the conditional risk model unreasonable, we do not consider it in the remainder of this paper, and we suggest that it be used with great care for hazmat route selection, if it is used at all.

In developing the risk models in Table I, we used assumption (3) ($p^s \cong 0$, for $s > 1$). It is possible to develop close relatives (more accurate versions) of these models *without* using (3). However, the versions of the traditional risk model and the perceived risk model *without* assumption (3) would violate all three axioms (see Erkut 1997 for a discussion and numerical examples). In particular, the violation of Axiom 1 gives rise to an interesting phenomenon in the selection of routes, namely looping. Figure 7 provides a very simple example that demonstrates the violation of Axiom 1 by the traditional risk model if assumption (3) is not made.

In this example, the risk of taking *edge* 2 is 100, whereas the risk of taking *edges* 1 and 2 in succession can be computed, using (6), as $(0.1)10 + (0.9)(0.1)1000 = 91$. In fact, this risk can be further reduced to 82.9 if one loops on *edge* 1 twice before tracing *edge* 2. Clearly, each additional loop on *edge* 1 further reduces the risk. This happens since the trip through the loop is reducing the incident probability on *edge* 2, where the consequence is much higher than that of *edge* 1 (the loop).

Boffey and Karkazis (1995) point out that looping may reduce transportation risks if one does not make assumption (3). A looping hazmat route is clearly undesirable. Thus, if one wanted to use a model which is prone to looping, one would have to restrict the feasible set to loopless paths (as in Sivakumar 1993a and Boffey and Karkazis 1995). However, if one makes assumption (3), there is no need to be concerned about looping. Since this assumption reduces the route selection problem to a shortest path problem (and since probabilities and consequences are nonnegative), any route that contains a loop cannot be optimal. This can be viewed as a fringe benefit of using the simplifying assumption (3).

We now discuss an exception where Axiom 2 can be violated by a credible hazmat routing model, if it is a multiattribute model. Evaluation functions that satisfy Axiom 2 are called order-preserving functions. Caraway et al. (1990) show that an additive two-attribute utility function may be non-order-preserving if one of the utility functions is additive and the other is multiplicative. Nembhard and White (1997) build on this result and demonstrate that an additive two-attribute hazmat routing model consisting of a linear cost attribute and an exponential population exposure attribute may be non-order-preserving. To provide some intuition for this interesting behavior, we use an informal example. Suppose an optimal trip from Miami to Washington, DC, takes path P_1, with a given total cost and total population exposure. Now consider extending this route to Boston. It is possible that the Washington-Boston leg adds so much to the population exposure attribute that (given the exponential nature of the population exposure attribute) the optimal Miami-Boston route may follow a different path, say P_2, for the Miami-Washington leg, which costs more than P_1 but exposes fewer people. Note that if the utility function is not order-preserving, then a shortest path algorithm is not guaranteed to find an optimal solution. Nembhard and White (1997) report that non-order-preservation may occur with some regularity on actual road networks, and that ignoring this nature of the model may result in poor route selections on real transport networks. They propose an enumeration algorithm based on artificial intelligence to deal with this property. We finish this section by noting that while the two-attribute utility model used by Nembhard and White violates Axiom 2, it does satisfy Axioms 1 and 3.

4. A BICRITERION APPROACH TO MODELING OF RISK

The models discussed in the previous section use incident probabilities or population exposure or both to quantify risk. One can argue that lower incident probabilities should be preferred to higher ones, and lower population exposures should be preferred to higher ones. Thus, we can view the risk minimization problem as a bicriterion optimization problem: one of minimizing incident probability and population exposure. (These two objectives were included in the multiobjective schemes of List and Turnquist 1994 and Abkowitz et al. 1992.) The efficient frontier

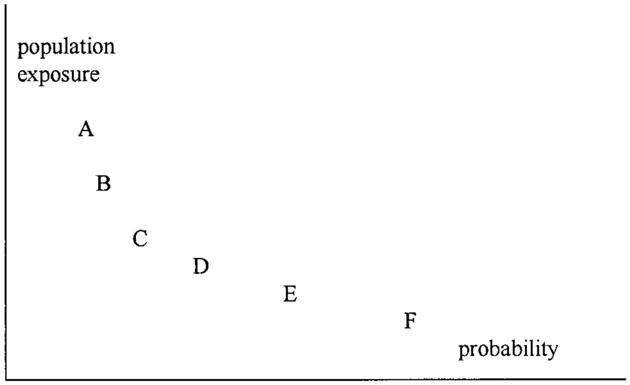

Figure 8. Efficient frontier of a typical bicriterion risk minimization problem.

of the bicriterion risk minimization problem would, in general, have the above form, where each letter denotes a different hazmat route.

The generation of this efficient frontier is not a trivial task. While it is possible to generate the entire efficient frontier by using a bicriterion shortest path algorithm, such as the one proposed by Hansen (1980), the potential number of efficient solutions is not polynomial in the number of network nodes, and attempting to generate all efficient solutions may be quite cumbersome. As a shortcut, a subset of the efficient frontier can be easily generated by using weighted combinations of probability and consequence. The optimization problem to be solved is a shortest path problem, and the weights can be varied parametrically to generate many efficient solutions.

We note that the optimal solutions of the incident probability model and the population exposure model represent the extreme points of the efficient frontier (points A and F, respectively, in Figure 8). However, optimal solutions to the traditional risk and the perceived risk models might not be on the efficient frontier. Note that the two criteria for the bicriterion path selection problem are to minimize the sum of incident probabilities and the sum of population exposures. However, the traditional risk model is not a function of these two sums; it is the sum of pairwise products of the incident probabilities with population exposures. This is why the point minimizing the traditional risk might not be on the efficient frontier. Clearly, this is also true for the points minimizing the perceived risk. In the next section we give a numerical example where this is the case.

We believe that the bicriterion approach to modeling of risk can be useful in hazmat routing decisions. Instead of choosing one of the two extreme models of risk, a decision-maker can use the bicriterion approach to generate a number of efficient alternatives. This would allow the decision-maker to select a compromise solution based on the trade-offs between the two criteria and other attributes (such as cost or length). It is, of course, possible (even advisable) to pose the hazmat routing problem as a multicriteria decision problem, where cost, and possibly other criteria, are considered alongside population exposure and incident probability. As long as each criterion is additive in its edge attributes, the weighting method can be used to generate a subset of efficient points. Although this is a potentially useful area of research, we do not discuss it further here since we focus only on the modeling of risk in this paper and not on the modeling of the entire hazmat routing problem.

5. EMPIRICAL ANALYSIS

In this section, we report the results of our computational analysis with different risk models. The tool we used in this comparison is a professional software package: PC*HazRoute (ALK Associates 1994). (ALK Associates has loaned a copy of PC*HazRoute to E. Erkut for academic research purposes; the authors have no financial or legal association with the company or the software.) This software has a database which contains information about 0.5 million miles of roads in the continental United States. For each road segment, the length and the population density in the neighborhood of the road segment are stored. Accident and release probabilities are determined by the road type. The impact radius can be selected depending on the hazmat under consideration. This decision-support-system finds optimal hazmat transport paths for a given origin-destination pair and a selected objective. We selected the following six objectives for our study:

1. Shortest travel distance,
2. Minimum population exposure,
3. Minimum societal risk,
4. Minimum DoT risk,
5. Minimum accident probability, and
6. Minimum incident probability.

The shortest path might not be a good choice for transporting hazmats. However, we used it as a benchmark to compare the optimal solutions found by the other five criteria. Objective 2 ignores incident probabilities and finds the path that exposes the fewest number of people to the hazmats. Objective 3 is the traditional definition of risk. It uses the following formula to find the risk associated with each edge:

Societal risk =

length of the edge (in miles)
× accident rate (probability) on the edge (per mile)
× conditional release probability given an accident
× population density in the neighborhood of the edge (persons per mile-sq)
× $(\pi)(\text{impact radius})^2$ (in miles-sq).

Thus, the societal risk of an edge is the expected number of people to be impacted in one trip of the hazmat truck on that edge.

Objective 4 is the definition of risk suggested by the U.S. Department of Transportation (1989). This definition is

Table II
Accident and Conditional Release Probabilities Used in the Analysis

Road Type	Accident Rate (per million miles)	Conditional Release Probability
Rural		
Two-lane	2.19	0.086
Multilane undivided	4.49	0.081
Multilane divided	2.15	0.082
Limited access	0.64	0.090
Urban		
Two-lane	8.66	0.069
Multilane undivided	13.92	0.055
Multilane divided	12.47	0.062
Limited access	2.18	0.062

similar to the definition of societal risk with two differences: it ignores conditional release probabilities, and it computes population impacted by using a rectangle instead of a circle. For these two reasons, we believe that this definition is not particularly suitable. However, we include it in the comparison since this is the measure promoted by the DoT. Objective 5 finds the path that minimizes the accident probability, ignoring all other information. Finally, Objective 6 concentrates on incident probabilities and finds the path that minimizes the probability of a hazmat accident involving a release.

In all our computations, we used the default values of the software for accident and release probabilities, and we used an impact radius of one mile. The default values for the probabilities are given in Table II. These values, reported by Harwood et al. (1993), are weighted averages of data from three states (California, Illinois, and Michigan) in the United States.

We used each one of the six objectives for all possible pairs of the following eight cities in the eastern United States: Chicago, Detroit, Boston, Nashville, Atlanta, Houston, New Orleans, Jacksonville. We selected these cities for the following reasons: all of these cities are major population centers, they "span" the eastern United States, and there is considerable variation in the resulting trip lengths. In all, we found six paths for every one of the 28 pairs in our experiment. These paths contain anywhere from 175 edges (Jacksonville-New Orleans) to 1158 edges (Boston-Houston). We stored each of the 168 paths generated by PC*HazRoute and processed these solutions to obtain the results reported in this section.

We were interested in two separate questions in this analysis. First we wanted to answer the question "How similar are the paths found by different risk objectives for a given origin-destination pair?" If it turns out that the different risk models select the same path consistently, then the issue of which risk model to use would be a moot one. If on the other hand, different risk models select very different paths, then this would imply that researchers and practitioners should pay a lot of attention to the modeling of risk in hazmat transport. The second question we dealt with was "How does the optimal solution for one objective perform under the other objectives?" If, for example, all optimal solutions with respect to objective i are always near-optimal with respect to objective j, then we would conclude that using objective i instead of objective j is not such a bad strategy. This would be true even if the paths found under the two objectives are very different. If on the other hand, some of the optimal solutions under objective i perform very poorly under objective j, then we would conclude that the proper selection of the objective is very important indeed.

To answer the first question, we developed several similarity indices. All of our similarity indices are between zero and one, where one indicates perfect similarity and zero indicates perfect dissimilarity. To compute a similarity index for two paths, we process the edge lists of the two paths. Suppose we have two paths, A and B. If these two paths share no common edges, then the similarity index for this pair is zero. At the other extreme, if the two paths are identical, then the index is equal to one. If paths A and B have some common edges, but are not identical, then the similarity index quantifies the similarity between them. We used four indices for similarity measurement. Denoting the length of a path (or a subpath) by $L(\cdot)$, the indices are:

Index 1: $$\frac{L(A \cap B)}{2L(A)} + \frac{L(A \cap B)}{2L(B)};$$

Index 2: $$\sqrt{\frac{L(A \cap B)^2}{L(A)L(B)}};$$

Index 3: $$\frac{L(A \cap B)}{\max\{L(A), L(B)\}};$$

Index 4: $$\frac{L(A \cap B)}{L(A \cup B)}.$$

All indices use $L(A \cap B)$, the total length of the common edges between the two paths. We call this the intersection length. The first index is the arithmetic average of two ratios: the intersection length divided by the length of path A, and the intersection length divided by the length of path B. The second index is the geometric average of the same two ratios. The third index is the ratio of the intersection length divided by the length of the longer of the two paths. Finally, the fourth index is a ratio between the intersection length and the length of the union of the two paths. This fourth index is a weighted variant of a popular tool for similarity measurement in statistics, namely Jaccard's coefficient (Jobson 1992).

The next two tables contain aggregate statistics over the 28 OD-pairs for the first and fourth similarity indices only, since we found that there is a very strong correlation between the first three similarity indices considered. In fact, the first two indices resulted in almost always the same number, and were never off by more than 0.01. The correlation between the first index and the third index was almost perfect, with a correlation coefficient of 0.997. As

Table III
Averages and Standard Deviations (in Parentheses) of Similarity Index Values (Multiplied by 100) Between All Pairs of Criteria Using Index 1 (Upper Triangle) and Index 4 (Lower Triangle) Based on 28 OD-pairs

	Shortest	Popul.	Soc. Risk	DoT Risk	Acc. P.	Inc. P.
Shortest		2 (3)	7 (9)	6 (8)	34 (29)	34 (29)
Population	1 (1)		35 (19)	40 (19)	3 (3)	2 (3)
Soc. Risk	3 (5)	23 (13)		82 (21)	12 (12)	9 (10)
DoT Risk	3 (4)	26 (15)	74 (26)		11 (11)	9 (9)
Accident P.	25 (29)	1 (2)	6 (7)	6 (6)		92 (20)
Incident P.	25 (29)	1 (1)	5 (5)	5 (5)	89 (25)	

well, the numbers for Indices 1 and 3 were always very close to one another. In the instances where they were not the same, Index 3 was slightly smaller. In contrast, Index 4 was quite different from Index 1; using logarithmic regression we found that the relationship between Indices 1 and 4 was approximately as follows: Index 4 = $(\text{Index } 1)^{1.5}$. Index 4 measurements were consistently lower than Index 1 measurements. We provide aggregate statistics for Index 4 (in addition to Index 1) due to the significant differences between Index 4 and the other three indices.

Each table has six rows and columns, one for each criterion in the study. Table III contains averages and standard deviations for Indices 1 and 4, whereas Table IV contains the minimums and maximums for Indices 1 and 4, based on the 28 OD-pairs. All indices are multiplied by 100 for ease of exposition.

These tables indicate that there are very few instances of high similarities between paths selected by different criteria. In fact, of the 15 pairs of criteria considered, there are only two instances of high similarity: between the accident and incident probability minimization, and between societal risk and DoT risk minimization. In addition, there are a few more cases of mild similarity, as indicated by average indices of around 0.34–0.40 for Index 1: between population exposure minimization and the societal risk (and DoT risk) minimization, and between shortest paths and accident (and incident) probability minimization. In all other instances, similarity indices are below 0.12.

In all instances the standard deviations are quite high (relative to the means), indicating considerable deviations around the means. This conclusion is further substantiated in Table IV, where the ranges are quite high for most pairs. Table IV indicates that, even for pairs of criteria with high average similarity indices, there exist pairs of paths that are very dissimilar (which is indicated by low minimum similarity indices). For pairs that do not exhibit a high similarity in Table III, even the maximum similarity indices are below 0.41 for Index 1 and 0.25 for Index 4.

We now move on to the second question of interest: "How does the optimal solution for one objective perform under the other objectives?" To answer this question, we evaluated each of the six optimal paths for an OD-pair using all six criteria. We then normalized the six evaluations under each criterion by dividing all of them by the smallest of the six evaluations under that criterion. Table V contains average and standard deviation information for all OD-pairs, whereas Table VI contains minimum and maximum information. The rows correspond to optimal paths, and the columns correspond to the criteria.

These two tables indicate that the two special cases of the traditional risk model, namely population exposure minimization and incident probability minimization are in conflict. The population exposure minimizing paths usually follow secondary roads and are associated with much higher incident probabilities (4 to 7 times) than the incident probability minimizing paths. In return, incident probability minimizing paths usually follow highways, and expose 2 to 4 times as many people as the population exposure minimizing paths.

These tables also imply that the shortest path is usually a poor choice under most other criteria. For example shortest paths can result in a nine-fold increase in population exposure and a 13-fold increase in societal risk. Even for the two criteria that exhibit a relatively high similarity with the shortest path in Tables I and II, namely accident and incident probabilities, selection of the shortest paths results in an average deterioration of 71% to 84% over the optimal paths with respect to these two criteria. On the

Table IV
Minima and Maxima of Similarity Index Values (Multiplied by 100) Between All Pairs of Criteria Using Index 1 (Upper Triangle) and Index 4 (Lower Triangle) Based on 28 OD-pairs

	Shortest	Popul.	Soc. Risk	DoT Risk	Acc. P.	Inc. P.
Shortest		0–12	0–29	0–29	0–100	0–100
Population	0–6		2–63	2–75	0–12	0–12
Soc. Risk	0–16	1–45		8–100	0–41	0–29
DoT Risk	0–16	1–60	4–100		0–40	0–29
Accident P.	0–100	0–6	0–25	0–24		36–100
Incident P.	0–100	0–6	0–16	0–16	22–100	

Table V
Averages and Standard Deviations (in Parentheses) of Normalized Objective Function Values of Optimal Paths (Rows) Based on 28 OD-pairs

	Length	Popul.	Soc. Risk	DoT Risk	Accid. P.	Incid. P.
Shortest Path	1.00 (0.00)	3.77 (2.19)	4.65 (3.51)	5.09 (4.07)	1.71 (0.37)	1.84 (0.45)
Min. Pop Path	1.72 (0.24)	1.00 (0.00)	1.49 (0.43)	1.46 (0.49)	3.99 (0.50)	5.04 (0.66)
Min. Soc. Risk P.	1.71 (0.31)	1.21 (0.09)	1.00 (0.00)	1.01 (0.02)	2.84 (0.62)	3.48 (0.81)
Min. DoT Risk P.	1.73 (0.33)	1.15 (0.06)	1.01 (0.01)	1.00 (0.00)	3.09 (0.73)	3.83 (0.98)
Min. Acc. P. P.	1.10 (0.06)	2.71 (0.58)	1.84 (0.21)	2.06 (0.28)	1.00 (0.00)	1.00 (0.00)
Min. Inc. P. P.	1.09 (0.06)	2.82 (0.55)	1.94 (0.25)	2.18 (0.32)	1.00 (0.01)	1.00 (0.00)

other hand, the optimal paths with respect to these two criteria are usually not much longer than the shortest paths. For example, the selection of the minimum incident probability path would increase the trip length by an average of 9% (and a maximum of 21%) over the shortest path.

The minimum societal risk paths do not perform very well under several of the other criteria. They are on the average 70% longer than the shortest paths and are associated with incident probabilities that are on the average 3.5 times as high as the minimum incident probability paths. On the positive side, the minimum societal risk paths perform relatively well in terms of population exposure.

It is not surprising that the pairs of criteria that resulted in the highest similarities (namely societal risk vs. DoT risk, and accident vs. incident probability) also exhibit a very high tolerance for one another in Tables V and VI.

6. ILLUSTRATIVE EXAMPLES

In this section, we provide four illustrative examples of path selections using the risk minimization criteria described earlier. We first display the optimal paths for two OD-pairs from the analysis summarized above to provide some spatial insight into the similarity between pairs of risk criteria. For this purpose we selected the Jacksonville-New Orleans pair and the Chicago-Atlanta pair. We use the Boston-Houston pair for the third and fourth examples. In the third example, we generate a number of paths for this trip by using the perceived risk model with different risk-aversion factors. Finally, in the fourth example, we generate a subset of the efficient frontier of the bicriterion risk minimization problem using the weighting method.

For the Jacksonville-New Orleans pair, the shortest path also minimizes the accident and the incident probabilities, and the minimum societal risk path also minimizes the DoT risk. We display the shortest path and the minimum population path in Figure 9a; Index 1 for the two paths shown is 0.12. In Figure 9b we display the shortest path and the minimum societal risk path, with an Index 1 value of 0.29. The value of Index 1 between the minimum population path and the minimum societal risk path is 0.54. This is an example where many similarities exist between paths that are optimal with respect to different objectives.

For the Chicago-Atlanta pair, the minimum societal risk path is the same as the minimum DoT risk path, and the minimum accident probability path is the same as the minimum incident probability path. In Figure 10a we display the shortest path and the minimum societal risk path. In Figure 10b, we display the shortest path, the minimum population path and the minimum accident probability path. The Index 1 between the minimum accident probability path and the shortest path is 0.21. For all other pairs where the optimal paths do not coincide, the values of Index 1 are below 0.05. This is an example where there are few similarities between pairs of optimal paths.

In Figures 9 and 10, we observe that some of the paths selected (particularly the minimum population paths) contain a large number of twists and turns. The computed risk on these paths may in fact be understated, since it is well-known that accident probabilities are much higher at intersections (particularly due to left turns of vehicles) than on straight stretches of highways. The software we used does not explicitly model the additional risk at highway intersections, and the user cannot edit the road network database to incorporate such risks. It is possible to model intersections explicitly when designing a software for transportation problems. However, such detailed modeling would increase the size of the database considerably and would

Table VI
Minima and Maxima of Normalized Objective Function Values of Optimal Paths (Rows) Based on 28 OD-pairs

	Length	Popul.	Soc. Risk	DoT Risk	Acc. P.	Inc. P.
Shortest Path	1.00–1.00	1.64–9.09	1.66–13.3	1.73–15.3	1.00–2.26	1.00–2.67
Min. Pop Path	1.42–2.63	1.00–1.00	1.07–2.39	1.03–2.52	3.18–5.57	3.96–7.26
Min. Soc. Risk P.	1.28–2.64	1.05–1.42	1.00–1.00	1.00–1.05	1.17–4.77	1.27–5.98
Min. DoT Risk P.	1.28–3.02	1.05–1.29	1.00–1.02	1.00–1.00	1.17–5.49	1.27–6.93
Min. Acc. P. P.	1.00–1.21	1.89–4.36	1.33–2.23	1.42–2.62	1.00–1.00	1.00–1.01
Min. Inc. P. P.	1.00–1.21	1.91–4.36	1.33–2.54	1.42–2.87	1.00–1.03	1.00–1.00

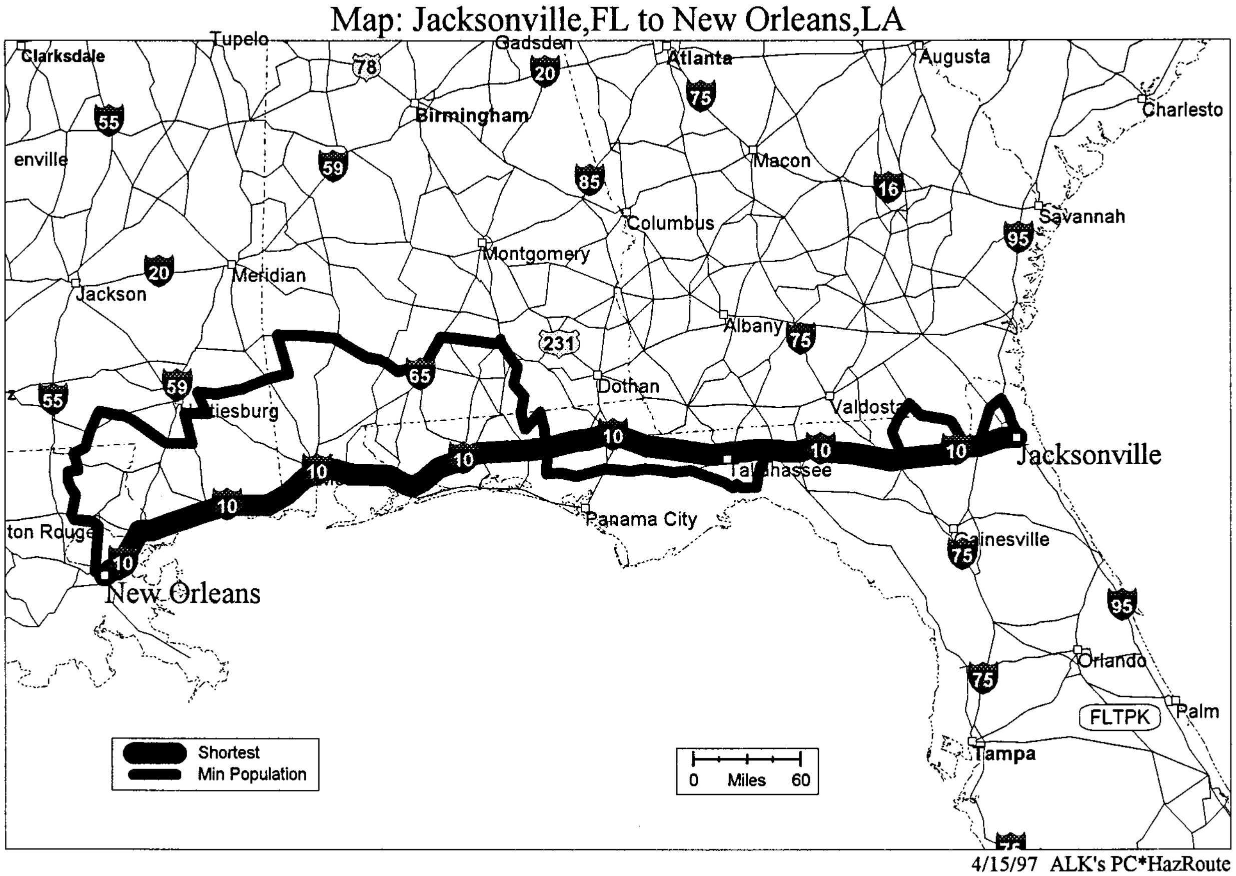

Figure 9a. Shortest path and minimum population path for the Jacksonville-New Orleans pair (Index 1 = 0.12).

slow down the programs. Hence, as usual, there is a cost associated with increased model realism.

We now move on to the Boston-Houston example, where we first demonstrate the use of the perceived risk model in (15). Although this is not one of the default risk models in PC*HazRoute, the "custom risk model" option of the software allows for a limited experiment where the risk-aversion factor (q) can be varied between 1 and 2. In Table VII we report the results of the experiment for the Boston-Houston pair for q values between 1 and 2 in increments of 0.1. For $q = 1$, the path selected is the minimum societal (traditional) risk path. As one would expect, with higher q values the model tends to choose paths with lower population exposure figures, but higher incident probabilities and higher societal risks. However, it is entirely possible for an increase in q to result in higher total population exposure, or a lower total incident probability. Furthermore, the length of the optimal path can increase or decrease as q is changed.

In our final illustrative example, we pose the risk minimization problem as a bicriterion optimization problem and generate a subset of the efficient frontier by minimizing weighted combinations of the incident probability and the population exposure for the Boston-Houston pair. PC*HazRoute has a "weighted risk model" option that allows the user to minimize a weighted combination of several objectives. One can use weights between 0 and 1 (in increments of 0.01) for each criterion. We note here that there is an almost 10 orders-of-magnitude difference between population exposure figures and incident probabilities. The software takes this significant difference into account and scales the weights before the optimization. Table VIII displays the results of 20 runs with different weights. Figure 11 displays the efficient frontier summarized in Table VIII, as well as the solutions to the perceived risk model summarized in Table VII.

As Table VIII indicates, the analysis results in a large number of efficient paths—18 distinct paths in 20 attempts (although some of these paths differ only slightly from one another). This is due to the high density of the U.S. road network, and fewer efficient paths might exist in less dense road networks such as Canada's. In some decision settings, a large number of efficient paths might in fact be desirable, as they make it easier to select a compromise solution.

At one end of the efficient path spectrum, we have the minimum population path (for $w_1 = 0$). As expected, this path has a relatively high incident probability. As w_1 is increased, the incident probability of the efficient path decreases and the population exposed increases. For $w_1 = 1$, the efficient path is identical to the minimum incident probability path. The incident probability of this path is only one-fifth that of the minimum population path, but

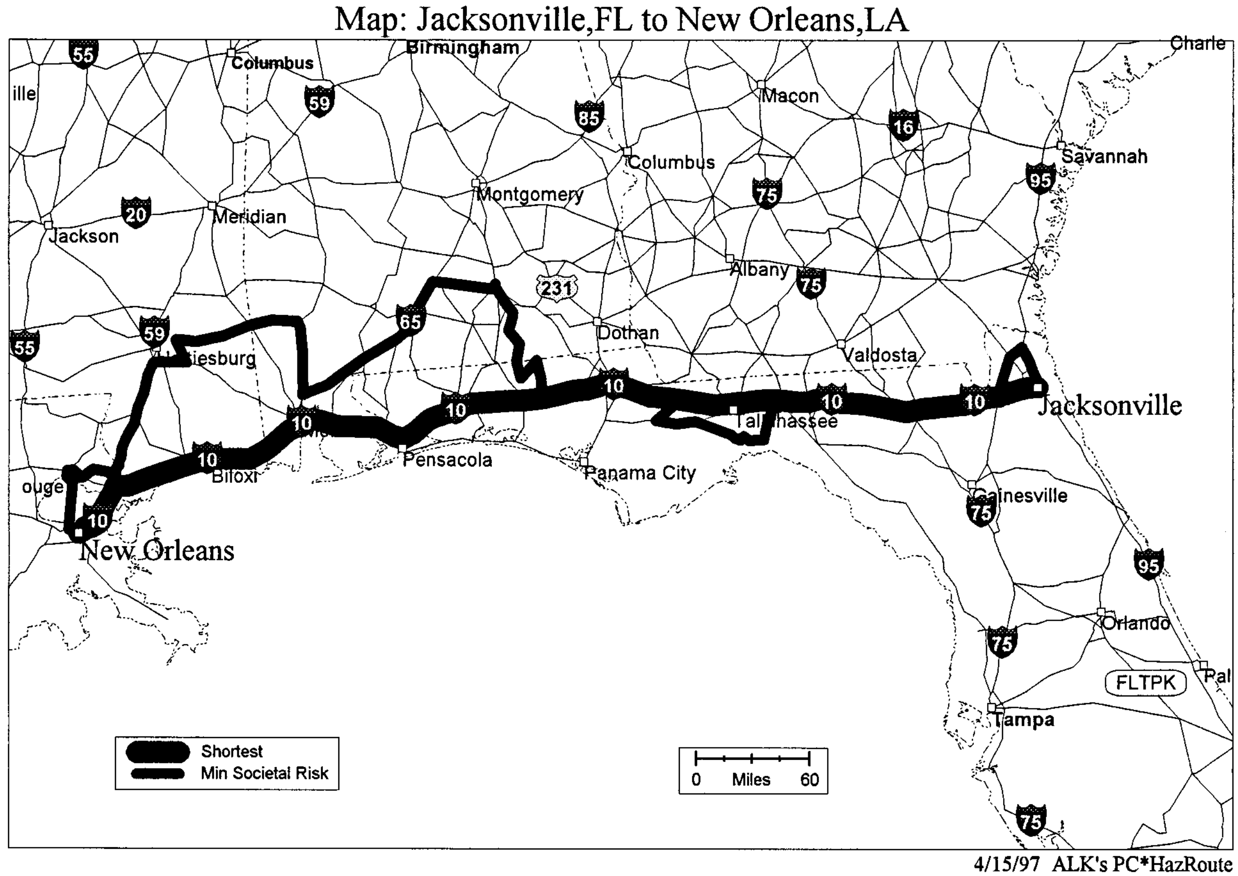

Figure 9b. Shortest path and minimum societal risk path for the Jacksonville-New Orleans pair (Index 1 = 0.29).

the population exposure of this path is over three times higher than for the minimum population path. This demonstrates that the range of objective function values over all efficient solutions can be quite substantial.

In general, as w_1 is increased, the length of the efficient paths decrease. This is due to the high similarity between the shortest path and the minimum incident probability path for many pairs in our analysis. However, in a few instances, a larger w_1 results in a longer path (for example $w_1 = 0.40$, or $w_1 = 0.80$).

Note that the societal risk values of many of the efficient solutions are relatively close to one another. Hence, many of these solutions are almost equally desirable from a societal risk perspective. However, the path that minimizes the societal risk (see Table VII, $q = 1$) is not on the efficient frontier. This path has a societal risk of 0.086, which is considerably lower than the societal risk values in Table VIII. Its incident probability is 3.2×10^{-4}, and it exposes a population of 606,057. Hence, it is dominated by the efficient solutions generated using $w_1 = 0.20$, 0.25, and 0.30. We do not imply here that these three paths should be preferred to the societal risk minimizing path; we state only that these three paths expose a smaller total population and result in a lower total incident probability than the minimum societal risk path. (Recall that the minimum societal risk does not minimize the product of these two sums, but a sum of the pairwise products of incident probabilities and population figures.) It is entirely conceivable that a rational decision-maker may reject the bicriterion model in favor of the traditional risk model.

7. SUMMARY AND CONCLUSIONS

The routing of hazardous materials is an important decision problem that is of interest to hazmat producers and consumers, hazmat carriers, local governments, insurance companies, and the people exposed to the risks from the shipments. Although accident probabilities are quite low for any given trip, the sheer volume of hazmat shipments almost guarantees that there will be some accidents over a sufficiently long period of time. Risks imposed on the residents living near transport routes can be reduced by explicitly taking the risk of the shipments into account in the planning and the selection of routes.

Most risk models in the hazmat transport literature use the concept of a danger zone. The assumption is that residents living inside a circle centered at the incident site, with a given impact radius, will experience the same undesirable consequence, and residents living outside this circle will experience no undesirable consequence. In this paper we argued that, although not very realistic, this assumption

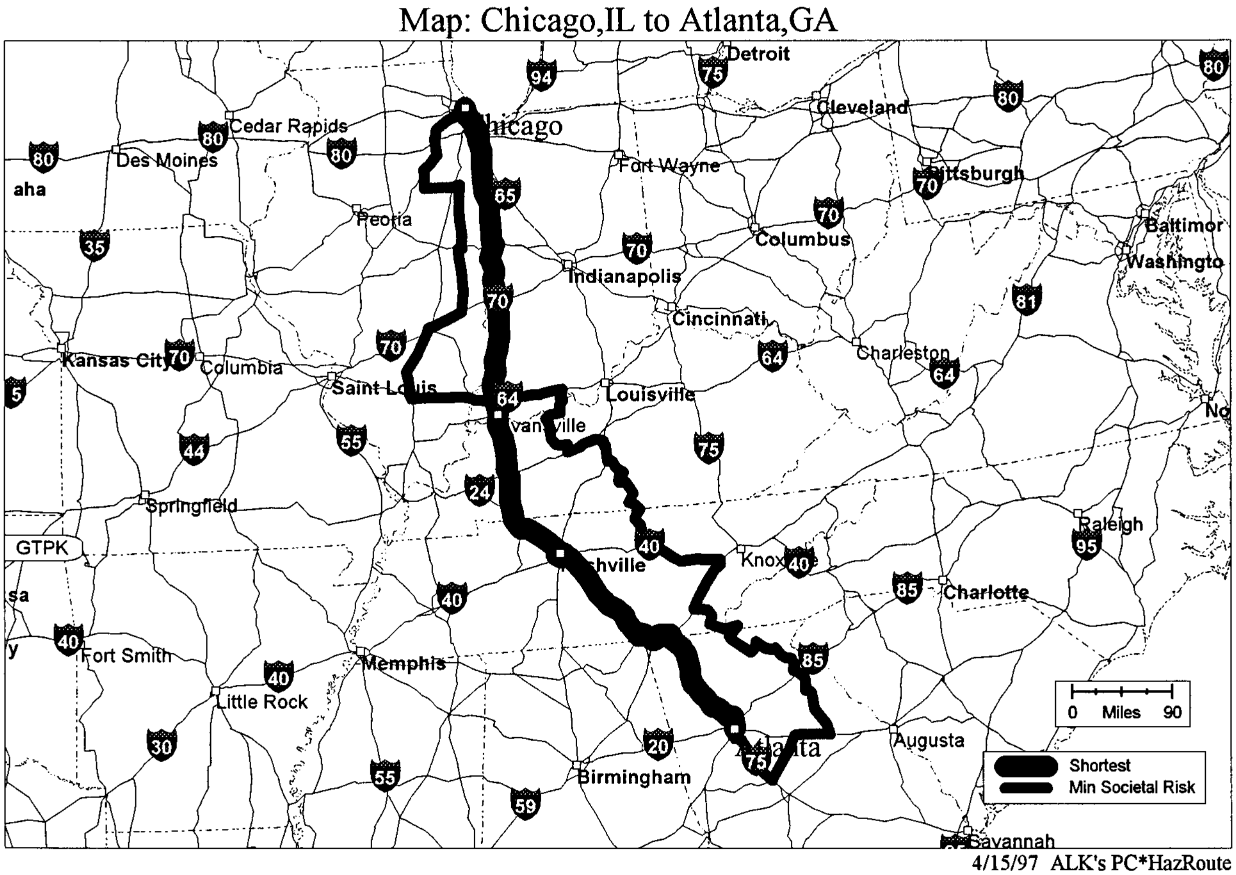

Figure 10a. Shortest path and minimum societal risk path for the Chicago-Atlanta pair.

is inevitable since the data necessary for an accurate assessment of risks to humans simply do not exist for many hazmats. As long as one is interested only in the relative risks of different paths, the danger circle approximation will provide valid comparisons.

All hazmat transport risk models we know of use an approximation to make the route-selection problem tractable. The approximation involves assuming that the products of incident probabilities are equal to zero. We demonstrated that this approximation does not result in significant inaccuracies in the estimation of path incident probabilities. We conclude that this approximation is justified and it is not necessary to develop more sophisticated risk models that do not use this approximation.

Although most researchers agree on the need to include risks in route selection for hazmat transport, they do not agree on how transport risk should be modeled. We discussed five different models: traditional risk, population exposure, incident probability, perceived risk, and conditional risk. We stated three axioms that we believe should be satisfied by a risk model for hazmat transport. We argued against using the conditional risk model, since it violates all three of these axioms. The remaining four models satisfy all three axioms (given the approximation stated in the above paragraph). In fact, each one is reduced to a shortest path problem after a processing of the data to compute edge impedances.

The hazmat route selection problem can also be viewed as a bicriterion shortest path problem, where the objectives are the minimization of total incident probability and total population exposure. This representation of the problem is useful for generating a number of efficient solutions, which can then be presented to a decision-maker for selection based on their performances on the two criteria, or on other criteria such as cost or the traditional definition of risk. We note that the path with the minimum traditional risk may not be on the efficient frontier of the bicriterion problem.

In our empirical analysis, we searched for answers to the following two questions: "How similar are the paths found by different objectives for a given origin-destination pair?" and "How does the optimal solution for one objective perform under the other objectives?" Our analysis was performed using a professional decision-support system for hazmat route selection. We generated optimal paths for 28 pairs of cities in the eastern United States using six different objectives for each pair. We found that the optimal paths with respect to the three fundamental risk models—namely, minimizing the traditional definition of risk, minimizing total incident probability, and minimizing total population exposed—do not exhibit strong similarities. For many of the pairs studied, these three paths were quite different. In general, incident-probability-minimizing paths seem to take fairly short paths and use highways as much

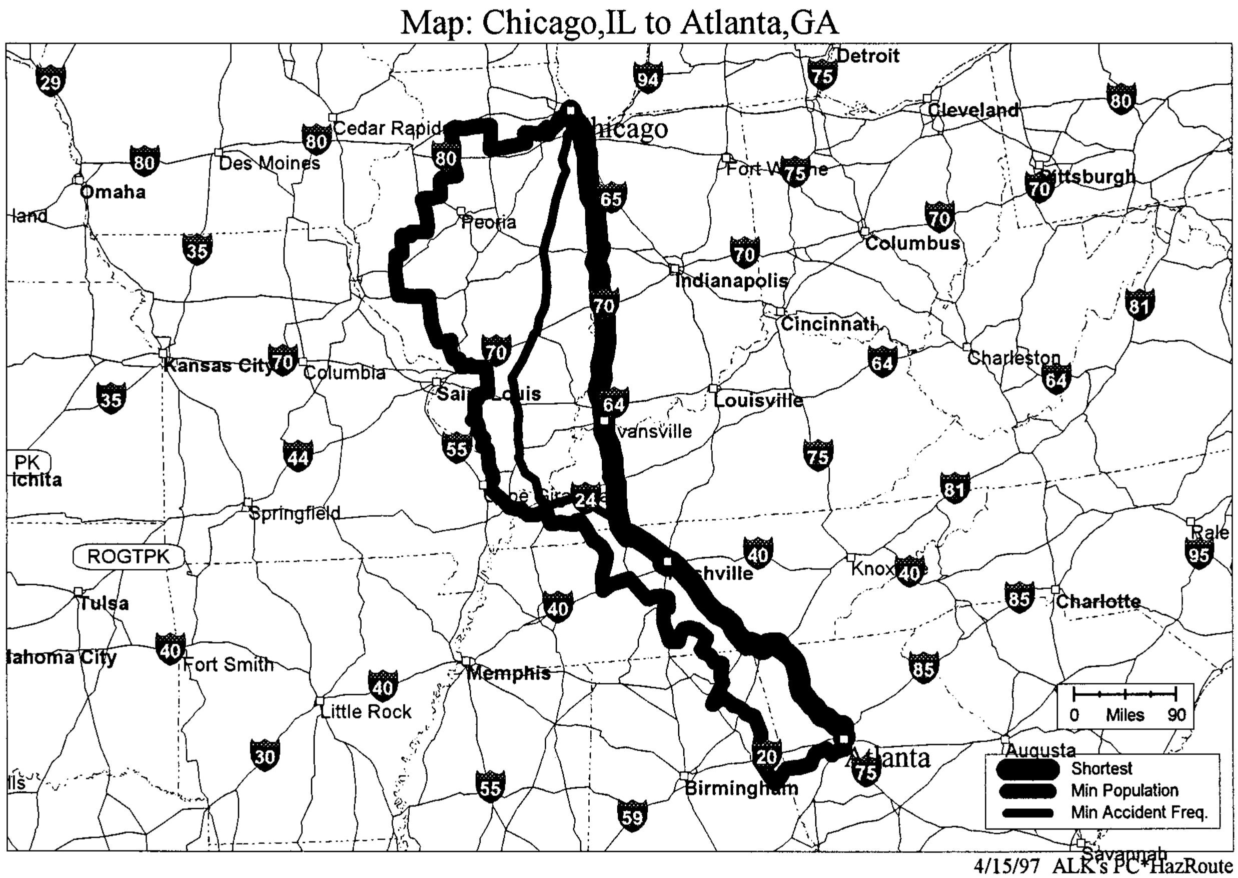

Figure 10b. Shortest path, minimum population path, and minimum accident probability path for the Chicago-Atlanta pair. (Index 1 value between minimum accident probability path and shortest path is 0.21.)

as possible. On the other hand, population-exposure-minimizing paths seem to wind through secondary roads that do not go though densely populated areas and are considerably longer than the shortest paths. There seemed to be a mild similarity between the population-exposure-minimizing paths and risk-minimizing paths. In addition to the spatial dissimilarities between the three optimal paths, the optimal path under one of the criteria was usually not a near-optimal path under the other two criteria. In fact, in most instances there were significant trade-offs among the optimal solutions with respect to these three criteria.

Table VII
Statistics of Optimal Risk Paths Generated Using the Perceived Risk Model (15) with q Values Varying from 1 to 2 in Increments of 0.1 for the Boston-Houston Pair

q	Population Exposure	Incident Prb. ($\times 10^{-4}$)	Soc. Risk	Length (Miles)
1	606,057	3.20	0.086	2792.1
1.1	593,292	3.30	0.086	2818.8
1.2	566,652	3.70	0.087	2918.3
1.3	567,037	3.70	0.087	2932.3
1.4	557,135	3.90	0.088	2976.8
1.5	548,601	4.00	0.088	2919.8
1.6	542,629	4.00	0.089	2925.0
1.7	540,043	4.30	0.092	3116.2
1.8	549,158	4.80	0.096	3337.2
1.9	612,652	4.80	0.110	3346.5
2.0	612,806	4.80	0.110	3351.5

The risk disutility model (which we did not include in the comparison) can be used to address the issue of risk aversion. A high value of the "risk-aversion parameter" in the model would result in the selection of road segments with low population densities, reducing the undesirable consequences in the case of an incident. However, increasing the risk-aversion parameter does not necessarily result in a reduction in the total number of people placed at risk during the transport. As well, this model requires the selection of a risk-aversion parameter, which may be difficult. A potential use for this model is in the generation of a set of desirable solutions by varying the risk aversion factor parametrically. We caution that one can run into numerical problems when using this model since the population figures are raised to a power (the risk aversion parameter) that can make them very large.

Based on our analysis, we conclude that considerable attention should be paid to the modeling of risk for hazmat transport since the different objectives that are suggested in the literature cannot be used interchangeably. Different models result in different paths, and the models do not tolerate one another very well. Unfortunately, it is

Table VIII
Efficient Solutions Generated by Minimizing a Weighted Combination of Incident Probabilities (multiplied by w_1) and Population Exposure (multiplied by $1 - w_1$) for the Boston-Houston Pair

w_1	Inc. Prob. $\times$ (10^{-4})	Population	Soc. Risk	Length (Miles)
0.00	4.60	491,441	0.11	2,825.7
0.05	4.10	499,438	0.11	2,645.2
0.10	3.90	505,047	0.10	2,579.2
0.15	3.50	524,140	0.10	2,435.7
0.20	3.10	551,590	0.10	2,351.2
0.25	2.90	566,134	0.10	2,283.7
0.30	2.80	574,284	0.10	2,270.0
0.35	2.50	622,188	0.11	2,266.4
0.40	2.40	651,376	0.11	2,310.6
0.45	2.10	715,970	0.10	2,288.7
0.50	2.00	755,735	0.11	2,209.7
0.55	1.80	806,130	0.11	2,176.2
0.60	1.50	932,742	0.13	2,115.3
0.65	1.20	1,080,568	0.13	2,075.8
0.70	1.20	1,101,372	0.13	2,104.4
0.75	1.20	1,124,136	0.13	2,105.9
0.80	1.10	1,211,750	0.13	2,200.8
0.85	1.10	1,211,750	0.13	2,200.8
0.90	0.96	1,499,805	0.17	1,980.8
0.95	0.96	1,499,805	0.17	1,980.8
1.00	0.95	1,639,584	0.19	1,949.3

impossible to suggest one risk model for all instances of hazmat transport; the objectives of the decision-maker and the nature of the hazmat in question are important in the selection of the appropriate model. (For example, one may just use incident probabilities when transporting PCBs, but use both probabilities and population when transporting chlorine.) Instead of using one risk model, which would find one "optimal" path, we recommend the generation of a manageable number of good (efficient) solutions using the different single-criterion risk models, as well as using weighted combinations of the different criteria, for the final decision. Availability of decision-support systems, such as the one we used in this paper, makes this a fairly simple exercise. This would allow a decision-maker to select one of the solutions based on his/her utility function. If this exercise is repeated a number of times with different OD-pairs, it may be possible to make certain inferences about the preferences of the decision-maker, and incorporate these preferences into a decision-maker-specific model for hazmat route selection. However, the real strength of a decision-support system is not in its extraction potential of utility functions, but in its ability to provide a wealth of information to the decision-maker and create an environment for informed decision-making.

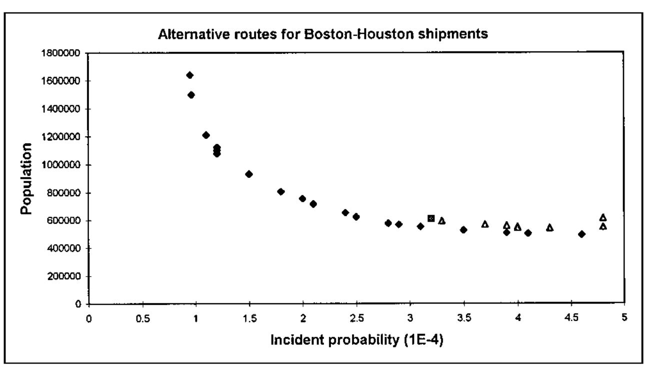

Figure 11. Efficient frontier (diamonds) generated by minimizing a weighted combination of incident probabilities (multiplied by w_1) and population exposure (multiplied by $1 - w_1$), the solutions to the perceived risk model (triangles), and the solution minimizing the societal risk (square), for the Boston-Houston pair.

ACKNOWLEDGMENT

This research has been supported in part by two grants from the Natural Sciences and Engineering Research Council of Canada (OGP 25491, OGP 183631). The support of ALK Associates by allowing the authors to use a copy of PC*HazRoute for empirical analysis is greatly appreciated. We acknowledge the contributions of Rajan Batta, Armann Ingolfsson, Susan Neuman, and two referees to this paper. Specifically, Rajan suggested the concept of the moving danger circle and discussed the conditional risk model with us at length. Armann made numerous comments on an earlier draft and helped us clarify the relationship between the discrete and continuous incident probability models. Susan edited the final draft and improved the flow of the material. Finally, questions and comments by one of the referees helped us to improve the exposition and resulted in the insertion of several clarifying sentences and paragraphs into the text. The authors take the responsibility for all remaining inaccuracies and errors.

REFERENCES

Abkowitz, M., M. Lepofsky, and P. Cheng. 1992. Selecting Criteria for Designating Hazardous Materials Highway Routes. *Transp. Res.* **1333,** 30–35.

ALK Associates. 1994. *ALK's PC*HazRoute*. ALK Associates, Inc., Princeton, NJ.

Alp, E. 1995. Risk-based Transportation Planning Practice: Overall Methodology and A Case Example. *INFOR* **33,** 4–19.

Batta, R. and S. S. Chiu. 1988. Optimal Obnoxious Paths on a Network: Transportation of Hazardous Materials. *Opns. Res.* **36,** 84–92.

Boffey, B. and J. Karkazis. 1995. Linear Versus Nonlinear Models for Hazardous Materials Routing. *INFOR* **33,** 114–117.

CANUTEC. 1992. *Dangerous Goods: Initial Emergency Response Guide*. Transport Canada, Ottawa, Ontario.

Carraway, R. L., T. L. Morin, and H. Moskowitz. 1990. Generalized Dynamic Programming for Multicriteria Optimization. *Eur. J. Opnl. Res.* **44,** 95–104.

Covello, V. T. and M. W. Merkhofer. 1993. *Risk Assessment Methods: Approaches for Assessing Health and Environmental Risks*. Plenum Press, NY.

Erkut, E. 1995. On the Credibility of the Conditional Risk Model for Routing Hazardous Materials. *O. R. Lett*. **18,** 49–52.

Erkut, E. 1997. Revisiting Transport Risk Models for Hazardous Materials. Research Report 97-4, Research Papers in Management Science, Faculty of Business, University of Alberta, Edmonton, Alberta, Canada.

Erkut, E. and V. Verter. 1995a. Hazardous Materials Logistics. In *Facility Location: A Survey of Applications and Methods*. Z. Drezner (ed.), Springer-Verlag, NY.

Erkut, E. and V. Verter. 1995b. A Framework for Hazardous Materials Transport Risk Assessment. *Risk Analysis* **15,** 589–601.

Glickman, T. S. and P. K. Raj. 1993. A Comparison of Theoretical and Actual Consequences in Two Fatal Ammonia Incidents. In *Transportation of Dangerous Goods: Assessing the Risks*. F. F. Saccomanno and K. Cassidy (eds.), Institute for Risk Research, University of Waterloo, Canada.

Hansen, P. 1980. Bicriterion Path Problems. In *Multiple Criteria Decision Making: Theory and Applications*. G. Fandel and T. Gol (eds.), Lecture Notes in Economics and Mathematical Systems, 177, Springer-Verlag, Berlin.

Harwood, D. W., J. G. Viner, and E. R. Russel. 1993. Procedure for Developing Truck Accident and Release Rates for Hazmat Routing. *J. Transp. Eng.* **119,** 189–199.

Jobson, J. D. 1992. *Applied Multivariate Data Analysis, Volume II: Categorical and Multivariate Methods*. Springer-Verlag, NY.

Kalelkar, A. S. and R. E. Brooks. 1978. Use of Multidimensional Utility Functions in Hazardous Shipment Decisions. *Accident Analysis Prevention* **10,** 251–265.

List, G. and M. A. Turnquist. 1994. Routing and Emergency Response Team Siting for High-Level Radioactive Waste Shipments. Technical Report, Department of Civil Engineering, Rensselaer Polytechnic Institute.

McCord, M. R. and A. Y. Leu. 1995. Sensitivity of Hazardous Material Routes to Limited Preference Specification. *INFOR* **33,** 68–83.

Nembhard, D. A. and C. C. White III. 1997. Multiobjective Non-Order-Preserving Path Selection with Application to Hazmat Transport. *Transp. Sci*. **31,** 262–271.

ReVelle, C., J. Cohon, and D. Shobrys. 1991. Simultaneous Siting and Routing in the Disposal of Hazardous Wastes. *Transp. Sci.* **25,** 138–145.

Saccomanno, F. F. and A. Chan. 1985. Economic Evaluation of Routing Strategies for Hazardous Road Shipments. *Transp. Res.* **1020,** 12–18.

Sivakumar, R. A., R. Batta, and M. H. Karwan. 1993a. Establishing Credible Risk Criteria for Transporting Extremely Dangerous Hazardous Materials. In *Transportation of Dangerous Goods: Assessing the Risks*. F. F. Saccomanno and K. Cassidy (eds.), Institute for Risk Research, University of Waterloo, Canada.

Sivakumar, R. A., R. Batta, and M. H. Karwan. 1993b. A Network-based Model for Transporting Extremely Hazardous Materials. *O. R. Lett.* **13,** 85–93.

Sivakumar, R.A., R. Batta, and M. H. Karwan. 1995. A Multiple Route Conditional Risk Model for Transporting Hazardous Materials. *INFOR* **33,** 20–33.

U.S. Department of Commerce. 1994. Truck Inventory and Use Survey. Bureau of the Census, Washington, DC.

U.S. Department of Transport. 1989. Guidelines for Applying Criteria to Designate Routes for Transporting Hazardous Materials. Report No. DOT/RSPA/OHMT-89-02, Federal Highway Administration, Washington, DC.