HazMon

A Decision-Support System to Predict and Monitor the Evolution and Effects of Natural Hazards

Revision 5a

The Charles Stark Draper Laboratory, Inc.

Table of Contents

Table of Figures

HazMon: A Decision-Support System to Predict and Monitor the Evolution and Effects of Natural Hazards

Abstract

To successfully and efficiently mitigate of the potential impact on lives and infrastructure of dynamically evolving natural or man-made hazards, decision-makers require accurate and timely predictions of when, when, and with what intensity the hazards will strike. Monitoring and prediction of hazard evolution requires the efficient utilization of: environmental sensors based in space, on the ground, in the air, and at sea; computation resources to execute prediction models; and computation infrastructure to move data from sensor to computation resources, and ultimately to the mitigation agent (decision maker).

Currently, mitigation agents rely, in large part, on locally owned and controlled sensor and computation resources for hazard prediction. This “stove-piped” method of operation is generically organized among specific Federal and local agencies, on a hazard-by-hazard basis. This paper describes a “HazMon” (Hazard Monitoring) system-of-systems, which provides an architectural context for negotiated coordination of resources among mitigation agents as one way to improve the quality of hazard prediction. The relationship between the quality of prediction and the effectiveness of mitigation is explored. The results support the development of justifiable quality requirements that can then be used to determine technology gaps in resources.

As a result of this research, specific modeling techniques are established to evaluate future enabling technologies with respect to their potential for improving the quality and timeliness of actionable data provided to decision-makers.

1 Introduction

There are various governmental and non-governmental organizations in the US that are responsible for dealing with the impacts of natural hazards. Their charters range from providing assistance after-the-fact to before-the-fact mitigation of the potential impact on lives and property. Presently, hazard mitigation organizations operate as “stove-piped” systems, shown in Figure 1-1. Each organization maintains its own resources for mitigating the hazard—vehicles, supplies and logistical support systems—as well as dedicated environmental monitoring and computation resources (computer platforms and models) used to predict the evolution of the relevant types of hazards. The organization’s dedicated resources may be augmented by generally available environmental information, such as NOAA weather forecasts and climatic information in archived data sets.

The objective of this study to develop a conceptual system design that provides a model to determine how environmental monitoring and prediction technology could be enhanced, or newly developed, so as to improve the ability of these organizations to execute their mission: *save lives and save money* (minimize the costs associated with mitigation). This study was funded by NASA’s Earth Science Technology Office (ESTO) as part of its continued development and evaluation of advanced-concept scenarios that help identify future Earth Science Enterprise technology needs.

The primary assumption on which the study is based is: higher quality predictions of hazard evolution—more accurate and more timely—improve the effectiveness of the mitigation process. Further assumptions are that: 1) providing more relevant sensor data and better computation and modeling can improve the quality of hazard prediction, and 2) sharing relevant environmental monitoring and prediction resources among the various mitigation organizations can provide both higher operational reliability and significant cost benefits overall. Figure 1-2 depicts a HazMon system for managing the sharing of resources across mitigation organizations.

The characterization of such a system is the substance of Section 2, “Operational Concept and Functional Architecture”, while potential limits of the primary assumption are the substance of Section 3, “Assessing Technology Needs”.

Section 4, “Preliminary Requirements Assessment”, presents a detailed hazard-by-hazard analysis of the HazMon timeline, with the goal of identifying “technology gaps”, which present investment opportunities for significantly reducing overall HazMon delay times. Section 5, “Potential Technology Areas for Further Investment”, identifies the technology gaps that can be usefully addressed by EST investment. Section 6 presents “Conclusions”. Detailed definitions of HazMon terminology, and simulation results of the assessment described in Section 4 are found in Appendices A and B, respectively.

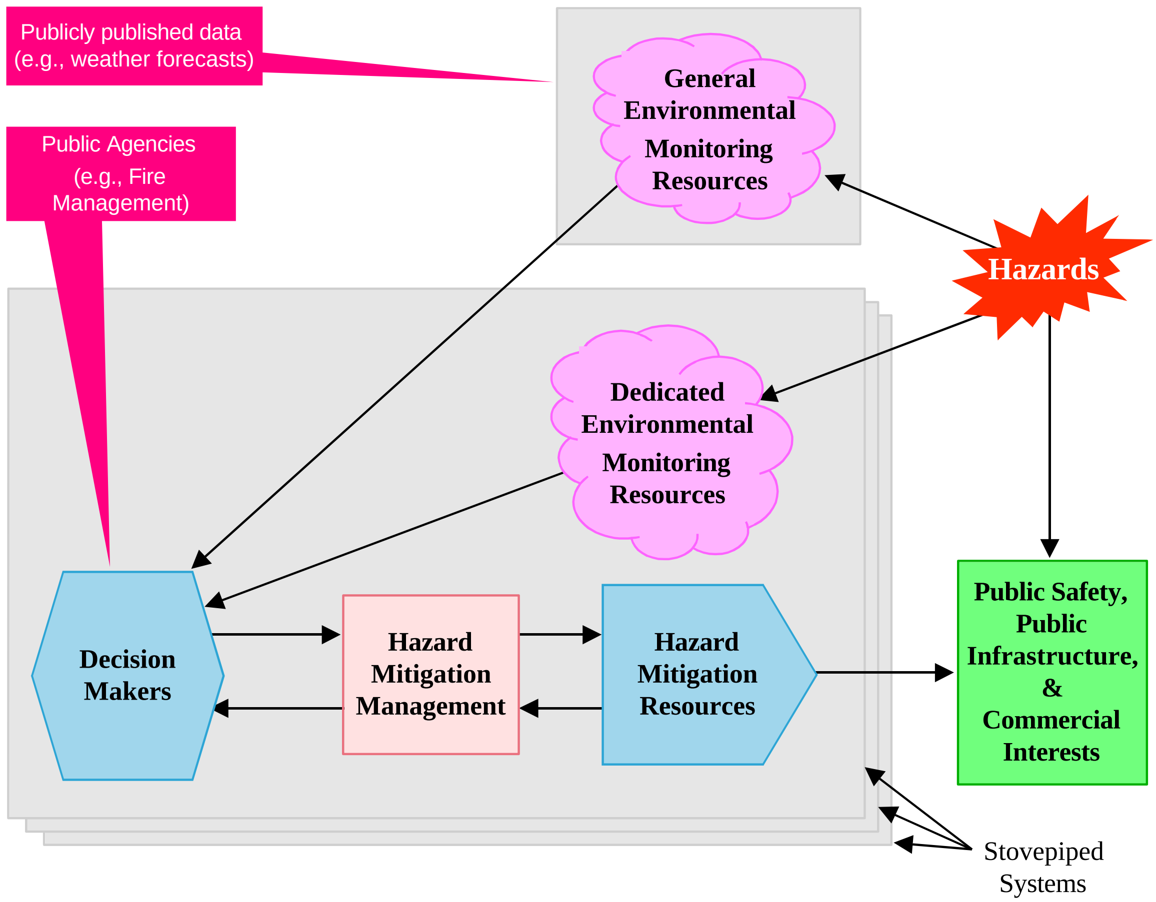

Fig. 1-1 – Dedicated Monitoring Resources in Current Hazard Mitigation Systems

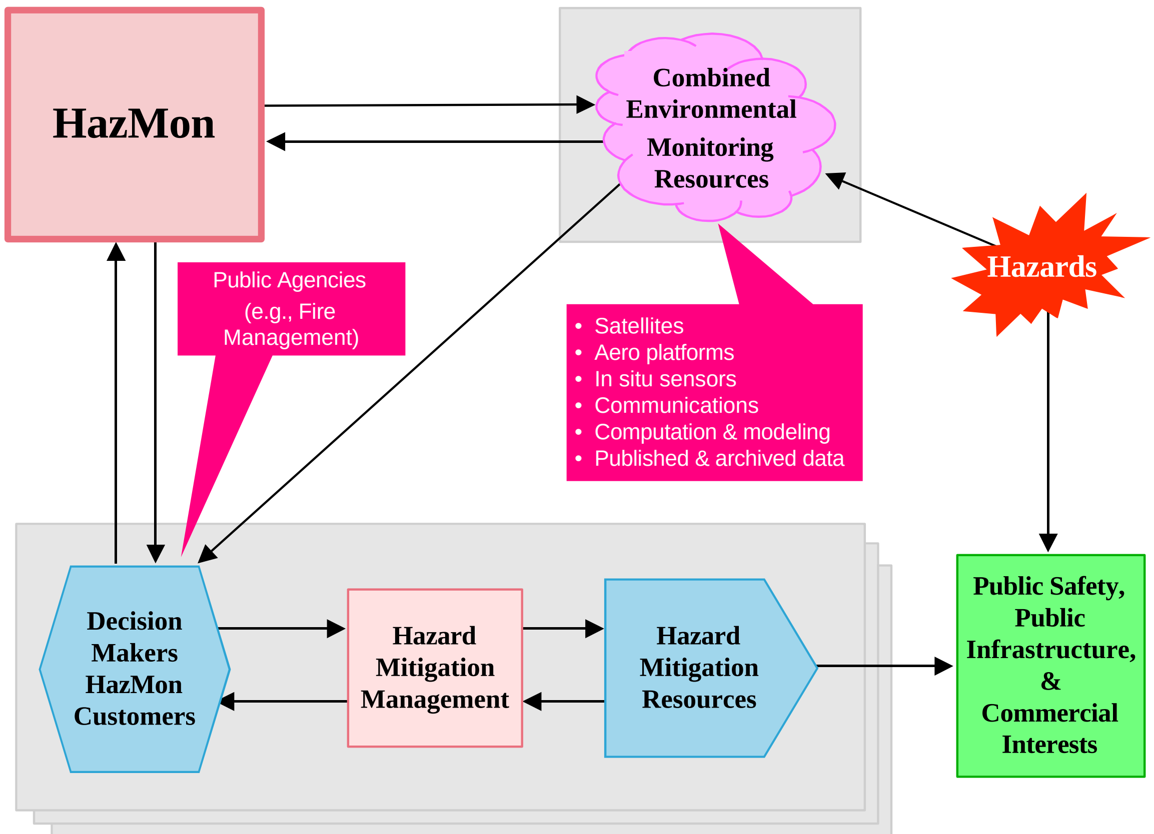

Fig. 1-2 – Shared Monitoring Resources in Proposed HazMon System

2 HazMon – Operational Concept and Functional Architecture

The concepts and terminology used in this section are characterized in the extended Glossary found in Appendix A.

2.1 Overview

The goal of the HazMon system is to provide decision makers with timely and accurate information about rapidly-evolving hazardous environmental situations in order to enable formulation of responses to mitigate the social and economic impact of such situations, i.e., save lives and save money. These situations (hazards) include severe storms, forest fires, floods, volcanic eruptions, tornados, oil spills, chemical spills, and possibly nuclear, biological, and chemical attacks. They may be precipitated by natural events or human actions. HazMon is intended to monitor situations and to predict their evolution, but not to predict their occurrence. Figure 2-1 indicates where the HazMon system fits into the big picture.

HazMon utilizes a network of sensors and computers, connected via a variety of communications links to accomplish its goal. While some of these sensing, computing, and communications *resources* may be owned and directly controlled by HazMon itself, the majority will belong to systems that are owned and operated by other organizations, public and private. The functionality of these external resources will be accessible by HazMon via prior arrangement with the owners of those systems.

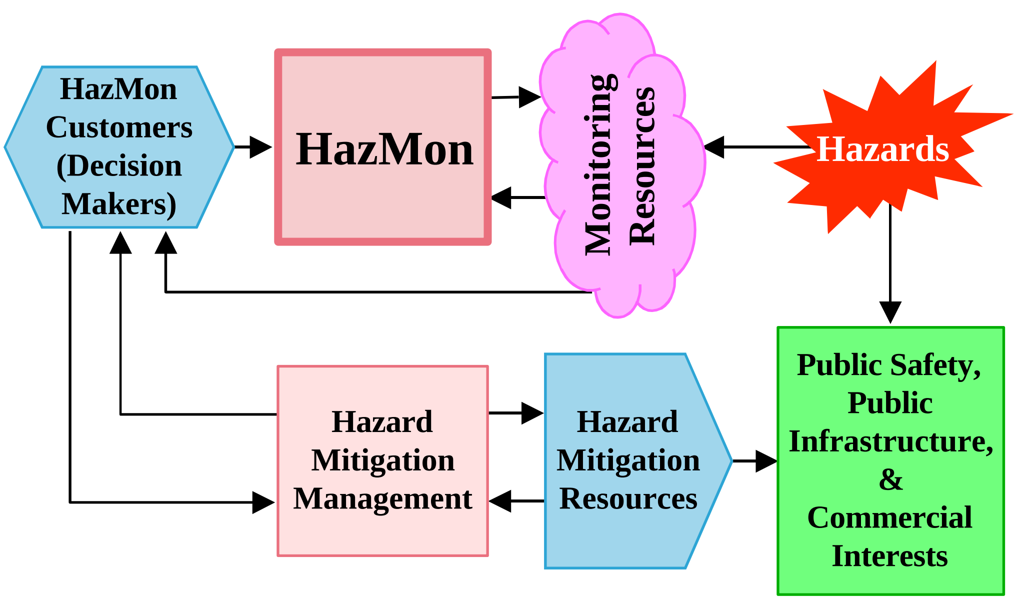

Fig. 2-1 – The Big Picture

Primary *customers* are envisioned to be federal and state organizations such as FEMA, NASA, USGS, state emergency response agencies, and DOD. Judicious design of the HazMon interface protocols will provide the ability to support private entities, and international organizations in the future.

HazMon conducts its operations in response to customer requests, where a request consists of a situation to be monitored, the desired format for delivering the results, and "cost" constraints. HazMon provides an automated, real-time negotiation framework for managing multiple customer requests, and for acquiring the resource functionality needed to support those requests. HazMon attempts to optimize the use of resources to provide the best "product" for each customer, given that customer's priority and cost constraints.

In many cases, potential customers will already have their own hazard monitoring capability. By collaborating with HazMon as resource providers, these organizations will benefit from optimal access to the functionality of a broad range of additional resources available to HazMon from other resource providers. The general relationship between Customers and Resources is shown in Figure 2-2, with HazMon acting as coordinator, broker, and optimizer.

Some customers play a role in mitigation of hazard effects, as indicated in Figure 2-1. HazMon does not directly interact with mitigation systems, but the potential benefits associated with the mitigation of identified hazards serve as inputs to HazMon resource assignment planning.

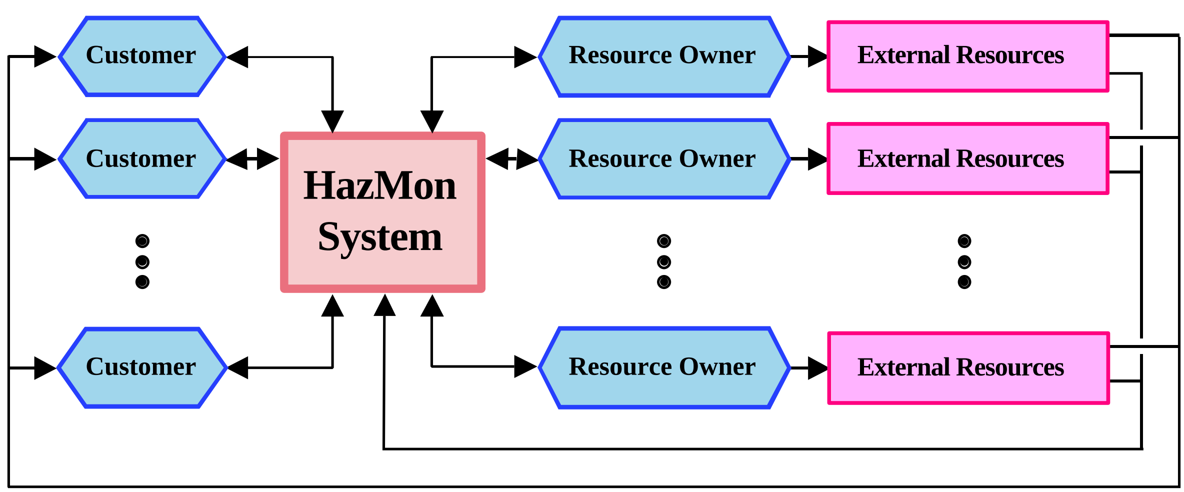

Fig. 2-2– Overview of Customer and Resource Relationship to HazMon

2.2 Scope and Objectives

This section is intended to characterize the overall scope and objectives of the HazMon system and the functionality required to achieve those objectives. Functionality needs to address a host of issues beyond the purely technical:

- HazMon needs to foster and support inter-organizational cooperation by providing the ability to accommodate policy into the operational behavior at a fundamental level. The ability to specify policy must be both dynamic and extensible.

- It must be possible to change policy parameters in real-time, so that in situations of severe resource contention, priorities can be controlled directly by critical decision makers.

- HazMon needs to manage and buffer the flow of “value” among participating organizations. Service to a customer creates a debit in that customers account. Utilization of a resource creates a credit in the owner’s account.

- HazMon needs to specify interfaces in a manner that supports seamless transition from legacy resources to next generation resource deployments. The interface protocols need to be *forward* compatible.

- When the occurrence of multiple hazards results in contention for resources by various customers, the system must be able to optimize the use of those resources to produce the greatest good. The HazMon planning algorithms will need to factor in the potential social costs of each hazard in order to assign resources in a way that is likely to result in the lowest overall social impact.

- Value management, the accounting of resource costs and avoided social cost in the allocation of resources, especially in situations of resource contention, must be *perceived as fair* to all of the organizations participating in HazMon operation.

2.2.1 Scope

The HazMon system is intended to monitor the status of rapidly-evolving hazardous environmental situations and to predict their evolution, but not to predict their occurrence. HazMon is *job oriented*—where a job is a discrete request to look for the occurrence, and track the evolution of, a particular phenomenon in a localized area, over a specified time interval. HazMon is not intended to look for the onset of all possible hazards in all places at all times.

A HazMon job may be requested in response to the actual occurrence of a hazard, such as an oil spill, or in response to the potential occurrence of a hazard based on the existence of precursor conditions. The possible occurrence of hazard may be imminent, as with a tornado, or the precursor conditions could indicate a general increase in the probability of the hazard over a longer time frame, such as dry, hot weather in the Arizona increasing the possibility of forest fires during the summer. Even though, as in the latter case, the HazMon job may be active for several months, it is still a discrete task.[1]

Although hazard mitigation analysis and related trade-offs are outside the scope of HazMon, the estimated benefits of the selected mitigation strategy for each hazard need to be considered by HazMon for optimization of resource allocation.

2.2.1.1 Hazards Addressed

The hazards for which HazMon can be meaningfully tasked must evolve over a sufficiently long enough time interval that a response to mitigate adverse impact is possible. Thus hazards such as earthquakes are ruled out because the onset is unpredictable and the phenomenon is over before any real-time reaction can be mounted. Tornadoes occupy the shortest time frame for which monitoring is realistic. HazMon can be tasked to watch for tornadoes based on precursor conditions. The appearance of these conditions is sufficient to issue preliminary general warnings. A tornado, if it actually occurs, lasts for tens of minutes—long enough for people to react to a specific warning.

Severe Storms	Hurricanes, tornadoes, thunderstorms, blizzards
Fires	Forest fires, smoke plumes and ash from forest fires and urban fires
Floods	Flash floods, rising water in rivers and reservoirs
Volcanic eruptions	Lava flow, ash, smoke
Droughts	Effect of crops and livestock
Pestilence	Biological impacts on crop output
Chemical, Radiological, and Biological Dispersions	Oil spills, chemical spills, radiation releases, and releases of harmful biological agents—triggered by equipment failures, human error, or malicious intent

[1] In the extreme, a collection of long time-frame discrete tasks could produce the effect of constantly monitoring, everywhere, for the occurrence of a particular phenomenon. From a functional perspective the system should be able to support such long-term tasking, but it is not the modus operandi envisioned for the foreseeable future.

2.2.1.2 Area of Coverage

Initially the system will cover the continental United States and those adjacent areas, e.g., the oceans, where phenomena may originate. Subsequently the system would be extended to cover Alaska, Hawaii, and US possessions.

2.2.1.3 Customers and Resource Owners

HazMon customers are those organizations responsible for the mitigation of the impacts of hazards. Resource owners are those organizations that provide sensors, communications links and computation facilities for use by HazMon. An organization may assume either or both of these roles. An initial instantiation of HazMon would likely include some or all of the following organizations:

Organization	Customer	Resource Owner
Federal Emergency Management Agency (FEMA)	X	
Department of Transportation (DOT)	X	
Department of Defense (DOD)	X	X
Department of Commerce (DOC)	X	
Department of the Interior (DOI)	X	
United States Geological Survey (USGS)	X	X
United States Department of Agriculture (USDA)	X	
National Aeronautics and Space Administration (NASA)		X
NOAA: NWS, Office of Climate, Water, and Weather Services National Hurricane Center - Tropical Prediction Center International Oceans Program NPOESS	X	X

2.2.1.4 Researchers

Although HazMon is likely to get much of its data from the same sensors used by researchers, it is not a vehicle for the support of research activities. HazMon needs to produce finished products in real-time. Researchers generally need access to large volumes of raw data for analysis, but do not require it in real time. However, products produced by HazMon can be provided to an archiving facility that will make them accessible to researchers.

2.3 Functional Capabilities

HazMon functions as a coordinator, broker, and optimizer among various systems. It operates by establishing a cooperative “market” among customers and resource owners. Customers specify constraints on the amount they are willing to pay for a particular job, and the HazMon system attempts to spread the benefit among all of its customers in an optimal manner. Organizations are admitted to the market by prior arrangement. These arrangements establish the operational policy and the protocols for interaction with HazMon. The internal functions of HazMon, and the functional interactions between HazMon and other organizations (in the roles of both customer and resource owner) are depicted in Figure 2-3, and described in the following sections.

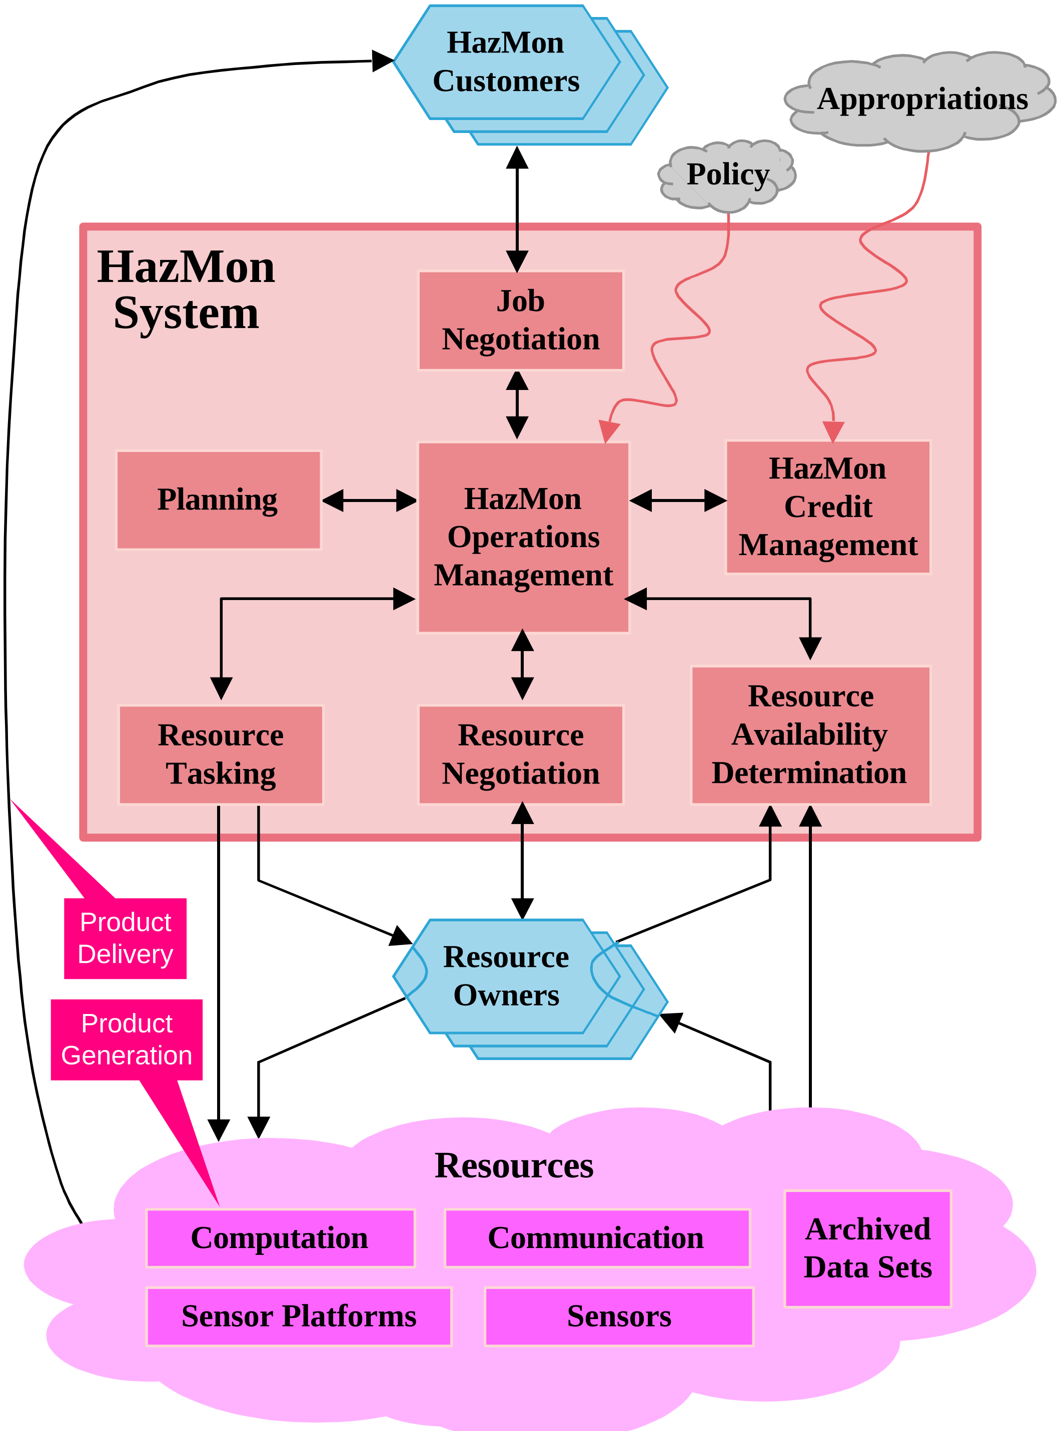

Fig. 2-3 – HazMon Functional Interactions

2.3.1 Job Negotiation

A job negotiation is initiated with a customer request for service. The request specifies the desired observations and predictions to be performed, the format of the results, the quality of the results, and a characterization of related costs—social costs that could be avoided as a result of the job and the cost of resources to perform the job.

The HazMon negotiation function passes the request parameters to the planning function for a feasibility assessment, i.e., are appropriate resources available for the job that meet the requirements and constraints of the job specification? If resources are available, HazMon replies positively and may offer to provide higher quality outputs at the same cost, or lower quality outputs at a reduced cost. If the desired resources are not available HazMon will indicate the cost necessary to acquire them and may suggest alternatives using available resources. If the job is feasible, the customer may authorize execution of the original request, authorize one of the alternatives, or modify the request and submit it for another feasibility check. If the job is not feasible as specified, the customer may authorize one of the alternatives, or modify the request and submit it for another feasibility check.

If there is little contention for resources the negotiation remains bilateral between HazMon and the customer. In times of high resource contention, or where specific resources are not advertised as being available, HazMon will need to negotiate for them with resource owners.

2.3.2 Resource Availability Determination

Resource status indicates when and under what conditions various capabilities of a resource will be available to support HazMon. Resources may be unavailable or unusable by HazMon for a variety of reasons, e.g., failure, maintenance, preemption by the owner. For instance, sensors deployed on satellites or aircraft may not be in the right place at the desired time. Furthermore, a resource may be initially committed to another activity, but could be preempted to support HazMon based on job priority or negotiated "price".

Before HazMon can determine availability status, it first needs to know the existence of potentially available resources. There are two ways that HazMon gains this information:

1. It knows about some resources a priori and can query the resource owner (or even the resource itself) about availability.

2. A resource owner (or the resource itself) advertises resource existence and availability status to HazMon.

2.3.3 Resource Negotiation

Resource negotiation is strongly affected by policy considerations. The resource negotiation protocols need to reflect this situation. For example, resource owners generally have top priority for the use of their own resources, but under critical conditions high-level decision-makers may need the ability to override that priority.

Costing resource use is also heavily driven by policy issues—issues that have to be pre-negotiated with resource owners. When there is no contention for a resource is it subjected to a baseline cost, or is it free? What happens when another job desires that resource? If two jobs can use the same resource in the same way do they split the cost, or does the second user ride for free? In situations of contention how does the system avoid a bidding war?

2.3.4 Resource Tasking

Resource tasking is responsible for formatting and transmitting command/requests determined by the planning algorithm to resources. With legacy resources, these requests will initially be passed to the resource owners, who will have to use their local interfaces to actually task the resource. Subsequently, it may be possible to build wrappers for the resource (middleware) that interface with the resources' native application programming interfaces (API) and transform them into HazMon-compatible protocol messages.

One objective of the HazMon development philosophy is that collaboration with HazMon would be sufficiently attractive to resource providers that their future systems would directly implement the HazMon protocols.

2.3.5 Product Generation

HazMon itself does not generate the products for the customer. Products are actually generated by the appropriate algorithms on the computation resources owned by resource providers and made available to HazMon.

2.3.6 Product Delivery

HazMon does not directly deliver products to customers; instead it directs computation and communication resources to deliver those products. The customer is responsible for further dissemination of those products. If the computation that produces the final product happens to be on a computer that is owned by the customer, then the delivery is implicit.

The HazMon system should be able to accommodate both push and pull delivery mechanisms.

2.3.7 Planning

HazMon and all of its external resources are actually part of a larger system that includes hazard mitigation management and emergency response resources. These resources are the machinery, hardware, and personnel that act on the environment and public infrastructure, and move the public, in order to mitigate the potential damage resulting from a hazard. They include such things as snowplows, buses, sandbags, and the National Guard. In an emergency, these resources need to be managed, in much the same way that the HazMon resources need to be managed. This function is depicted in Figure 2-4 as HazMit, or hazard mitigation.

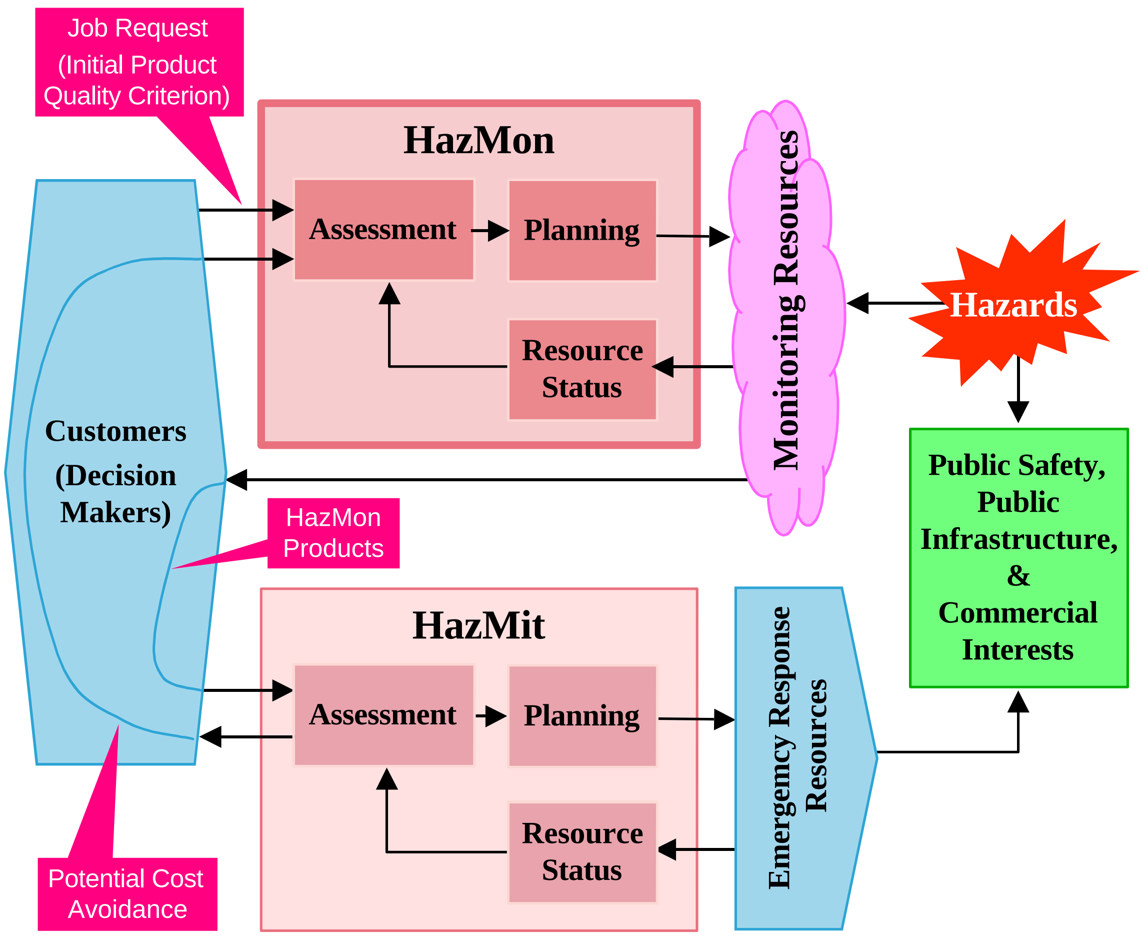

Fig. 2-4 – HazMon / HazMit Data Flow

HazMon *provides* inputs to HazMit (via customers) in the form of HazMon products. But it also *receives* inputs from HazMit in the form of the projected benefits of mitigating the effects of each of the hazards under considerations. HazMon tries to determine how the overall application of resources will provide the greatest overall benefit, so it needs to know the potential benefit of products for each job. But the potential benefit is inextricably related to how the hazard is evolving. This “chicken and egg” problem implies that there is continuous feedback process between HazMon and HazMit, driven by the evolution of hazards.

2.3.7.1 Cost Parameters

The ultimate objective of the HazMon system is to (help decision makers) minimize the social and economic cost associated with natural and man-made disasters. Thus "benefits" within HazMon are measured in avoided cost.

2.3.7.2 Optimization

When HazMon is performing multiple jobs, it will attempt to allocate resources to those jobs so as to meet the requirements of all of the jobs at the lowest aggregate cost. For example, sharing of the outputs of a sensor among multiple jobs can reduce the sensor "cost" for each of those jobs.

2.3.8 Credit Management and Accounting

Customers engage the HazMon system with the understanding that they will incur a cost. Resource owners make their resources available to HazMon with the understanding that they get something in return— either access to other resources at some time in the future, or services as a HazMon customer. HazMon must provide an internal banking system in order manage a fair exchange of "value". HazMon will use this bank to account for services provided and resources used. All forms of value must be translated into a common unit of exchange.

Defining a unit of currency is a significant issue. System implementers will be required to establish and negotiate terms of exchange rates between organizations and HazMon. Alternately, rates may need to be established between resource value as seen by an organization, and product value as seen by a customer, and HazMon. System implementers will need to resolve the following questions: Are the exchange rates fixed, or do they vary over time? If I am a customer who supplies no resources, how do I pay? Inter agency money transfers? Does the Federal budget allocate some amount of HazMon services?

Customers of the HazMon system pay for delivered products and resource owners get compensated for the use of their resources. The credit bank keeps track of the flow of "value" in the system.

2.3.9 HazMon Operations Management (OM)

This is the function that coordinates the interactions of the other HazMon functions, and could be envisioned as the primary User Interface to HazMon. As seen in Figure 2-3, the OM function would focus on managing the other HazMon functions (processes) to produce and distribute products. As a real-time (or near-real-time) system, HazMon must provide a user-centric focus on efficiency and effectiveness of the other processes, and would therefore provide tools for measurement and analysis of internal processes (a

"System Health" capability). OM is likely to be the physical "person-in-the-loop" location, the point where operators, managers, planners, decision-makers reside to actually operate and get results from HazMon. The HazMon OM function (and hence the overall performance of HazMon) is driven by Policy (see Section 2.5 below).

2.4 Interactions with Other Organizations

The HazMon system must be able to interact with other organizations in a variety of ways. Organizations may obtain services from HazMon or they may provide resources to HazMon. The initial interaction establishes the rules for collaboration between HazMon and the organization. Subsequent interactions utilize those rules to support real-time collaboration. The character of the interaction depends on the role played by the organization: customer, resource owner, or resource interface. An organization can play any or all of these roles.

2.4.1 Customer Role

When an organization requests services from HazMon it assumes the role of a customer, although an organization may be a resource provider as well as a customer.

2.4.2 Resource Owner Role

An organization that actually controls a set of resources and is fiscally responsible for acquisition, operation, and maintenance of the resources, is in the role of owner. In this role the owner is responsible for and has the authority to negotiate for the use of and assigns the cost of the resource(s). The resource owner is responsible for publishing the availability of and providing for the real-time status of the resource(s).

2.4.3 Resource Interface Role

An organization may act as a broker for resources while not actually acting in the role of customer or owner. An example of an organization acting in this role is a "third party" which may propose the following (for example): *The use of resource XXX from the YYY facility is to be shared between the Customer and YYY based on the "bucket" concept where all negotiated benefits accumulated as a result of the use of resource XXX collected by YYY are balanced against the costs of these activities as identified and tracked by mutual agreement of the two parties. The deficit of cost over revenue is to be split and should there be a surplus it is to be distributed to the YYY. In the future the "bucket" mechanism may be expanded to the system level.* [2]

[2] This example, modified for HazMon, is taken from the "RESOURCE INTERFACE BETWEEN AHC AND FAIRVIEW HEALTH SERVICES, JANUARY 27, 1999" (http://www.ahc.umn.edu/AHCFVResource.html)

2.5 Policy

HazMon organizational and operational policy will define the rules that regulate how the system will manage interfaces, information and resources to achieve its objectives. One of the policy's primary purposes is to document and publish to the user community the available information assets (resources) and how HazMon will respond to resource contention or changes to operational (run-time) priorities.

Specific policy procedures will include the actions necessary to observe systems and networks for signs of resource contention and/or unexpected behavior, including intrusion. Observation can take the form of monitoring, inspecting, and auditing. From these procedures, HazMon Operational managers will determine the operational steps they need to take to comply with published policy. These steps will thereby uphold the operational status and security of HazMon information and networked systems.

3 Assessing Technology Needs

The overall objective of HazMon is to increase the benefits of hazard mitigation efforts—save lives, save money—in an optimal manner. HazMon provides cost benefits by optimizing the use of hazard monitoring and prediction resources. By bringing more resources to bear on a given situation than is now possible, it can provide improved quality of predictions. In this respect it earns a place among other possible technology improvements that could improve the quality of predictions—more accurate sensors, faster computation, better hazard propagation models, and faster communications, including the ability to move data from space-based assets to the ground with negligible delay.

3.1 Quantization of Mitigation Actions

Given sufficient time and money we can certainly develop technologies to make predictions more accurate and timely—to the limit imposed by Planck's constant. But this relationship is not very useful in setting accuracy targets. What we really need to know is the relationship between the quality of hazard prediction and the effectiveness of the hazard mitigation activity—where *quality* is a function of the *accuracy* of the predicted location, time, and intensity of hazard impact, and *timeliness* of that prediction. Furthermore, because hazards evolve and are mitigated in different ways, this relationship needs to be determined for each type of hazard that needs to be addressed.

For each type of hazard we need to specify a quality target in order to determine whether current technologies meet that target, or development needs to be done. If the relationship between prediction quality and the effectiveness of mitigation were ever increasing, even if a maximum effectiveness were approached asymptotically, then selecting a quality target would be an arbitrary decision. With lives in the balance, the debate inevitably moves from the realm of engineering to the realm of politics. Diminishing returns notwithstanding, what is the value of a life?

As it turns out, the nature and character of real world of mitigation activities allows us to determine, in principle, when quality is good enough:

1. *Accuracy is good enough when a delta increase no longer provides actionable information to a mitigator*

   i.e., when the improved accuracy would not change the mitigation response. For example, knowing where hurricanes will first strike land within a mile is probably good enough, because improving that to a half mile is not likely to provide additional usable information.

2. *Timeliness is good enough when earlier availability of a prediction would have no material effect on mitigation response*

   If it requires three hours to generate the first prediction of where a hurricane will hit in five days, then reducing the prediction lag to one hour will have no material effect on the mitigation response.

These situations stem from the fact that mitigation actions are quantized. In Figure 3-1 the boxes depict a notional line (say a series of blocks on a street in Kansas) being crossed by a notional point hazard (say a tornado). The Gaussian curve gives the probability of the hazard crossing at any point. If we desire $n\sigma$ confidence that no one gets hurt, we would have to evacuate all of the blocks inside the gray cloud. We could achieve the same numerical result by evacuating only the people living between the dotted lines, but that isn't physically how evacuation works. In reality, evacuations occur by quantized geographical areas: towns, or blocks; but (for natural hazards) not generally by building or apartment within a building. If the $n\sigma$ range covers any part of a mitigation area, the whole area is evacuated.

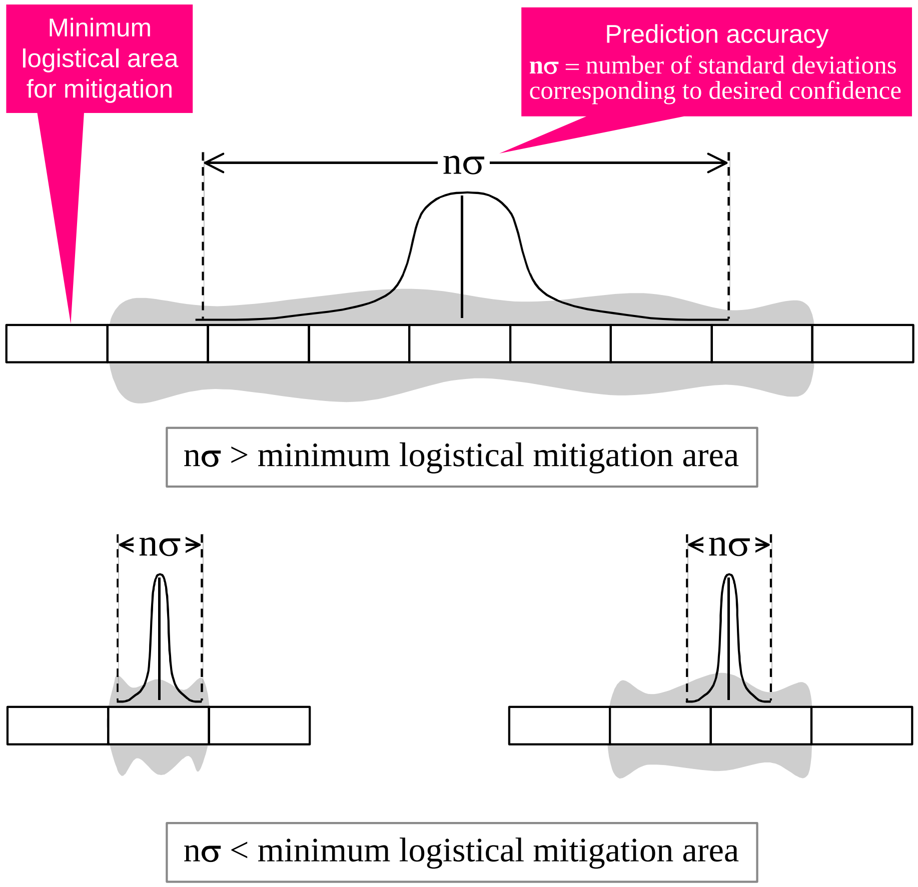

Fig. 3-1 – Spatial Quantization of Mitigation

Now assume that the prediction accuracy is increased as shown in the lower part of Figure 3-1. The curve is much narrower for the same nσ confidence range. If the range were smaller than a single mitigation area, then we would only have to evacuate one or two areas to have the same mitigation effectiveness. No matter how much smaller we made the range by achieving better accuracy, we would still have to evacuate one or two areas, so we get no additional benefit from the additional accuracy. But that additional accuracy incurs a cost that detracts from our overall objective of saving lives *and* saving money.

3.2 General Cost/Benefit Relationships

Figure 3-2 depicts the relationship between cost/benefit and cost/utility curves. The curve on the left shows the general relationship between cost and benefit. For a given situation (say a particular technology or a particular operation) benefit increases monotonically with cost, and the rate of increase slows monotonically, possibly, but not necessarily, reaching zero. The curve represents an efficient frontier of operation, providing the maximum benefit for any given cost. Operating anywhere in the gray area below the curve provides less than the maximum benefit associated with a given cost.

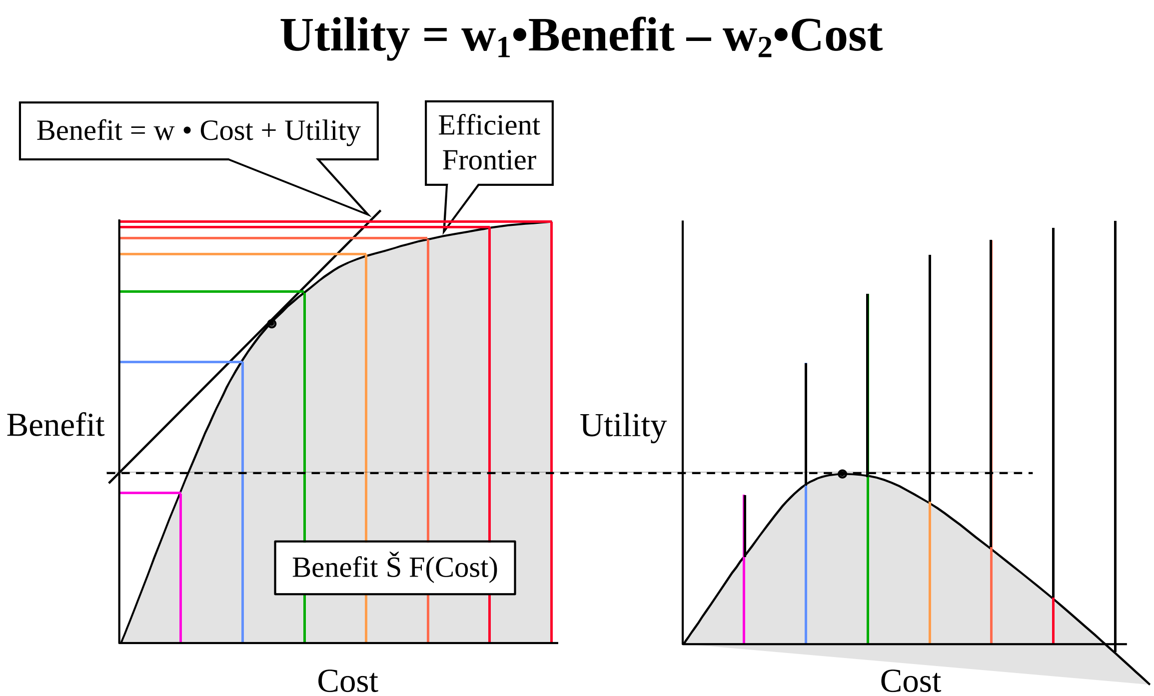

Fig. 3-2 – General Cost/Benefit Optimization

Low cost provides low benefit and higher costs provide more benefits. Picking an operating point requires that we specify the relative importance (to us) of cost and benefit. The utility function characterizes the net gain—the difference between the weighted benefit (expressed in units of cost) and the weighted cost. If we rearrange the terms of the utility function to show benefit as a function of cost, and express the relative weights $\mathbf{w}_1$ and $\mathbf{w}_2$ as **w** the slope of the line, then utility is the value where the line intercepts the benefits axis. Maximum utility (for the given weights) occurs when this line is tangent to the cost/benefit curve.

The curve on the right shows the relationship between utility and cost. It was plotted by taking each vertical benefit line and subtracting (from the top) the related horizontal cost line. Since benefits eventually increase more slowly than costs, the utility curve will eventually go negative, at the point where the costs outweigh the benefits.

3.3 Optimizing Investment in HazMon Resources

How good does hazard monitoring and prediction have to be? In Section 3.1 we saw that beyond a certain point, increasing the quality of predictions does not improve the effectiveness of mitigation activities. The relationship of prediction quality (**Q**) to effectiveness of mitigation **E**(**Q**) is shown in Figure 3-3 (second curve from the left). **E**(**Q**) attains a value of one at the point where prediction is accurate to within a quantum of mitigation activity, and it stays there.

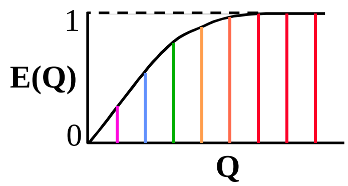

Fig. 3-3 – Effectiveness of Mitigation as a Function of Prediction Quality

We can construct a utility function to cover the entire scope of hazard monitoring and mitigation activities:

$$A = P \bullet E(Q) \bullet E(M) - C(Q) - C(M)$$

where:

A = Actual Savings (expressed in units of cost)
P = Potential Savings (expressed in units of cost)
Q = Quality of Prediction
E(Q) = Effectiveness of Mitigation Actions as a Function of Q
C(Q) = Cost of Q
M = Mitigation Action
E(M) = Effectiveness of Mitigation Action itself due to factors other than Q
C(M) = Cost of M

We want to explore the sensitivity of A (actual savings) to Q (prediction quality). To do this, we assume a specific mitigation strategy, so the utility function reduces to

$$A = P \bullet E(Q) - C(Q)$$

Figure 3-4 shows the intermediate relationships that result in the utility function (the rightmost curve). The curve of **C**(**Q**) vs.**Q** is just a cost benefit function with the axes reversed. Note that there is no limit on either quality or the cost of quality. The curve of **E**(**Q**) vs.**Q** indicates that effectiveness is maximized at a value of one for a given quality and stays there. The curve of **P• E**(**Q**) vs.**Q** is just the previous curve scaled by the potential savings. A hazard effects what is in its path. **P** is derived from how much damage and loss of life would be expected in the absence of any mitigation at all. The curve of **A** vs.**Q** is the utility function itself. Note that utility peaks before **E**(**Q**) reaches a value of one. To save all of the lives that could be saved (given the particular mitigation strategy) requires that we spend enough to get predictions of maximum useful quality. Although this reduces actual savings (in dollars), that reduction is limited because **E**(**Q**) reaches a maximum.

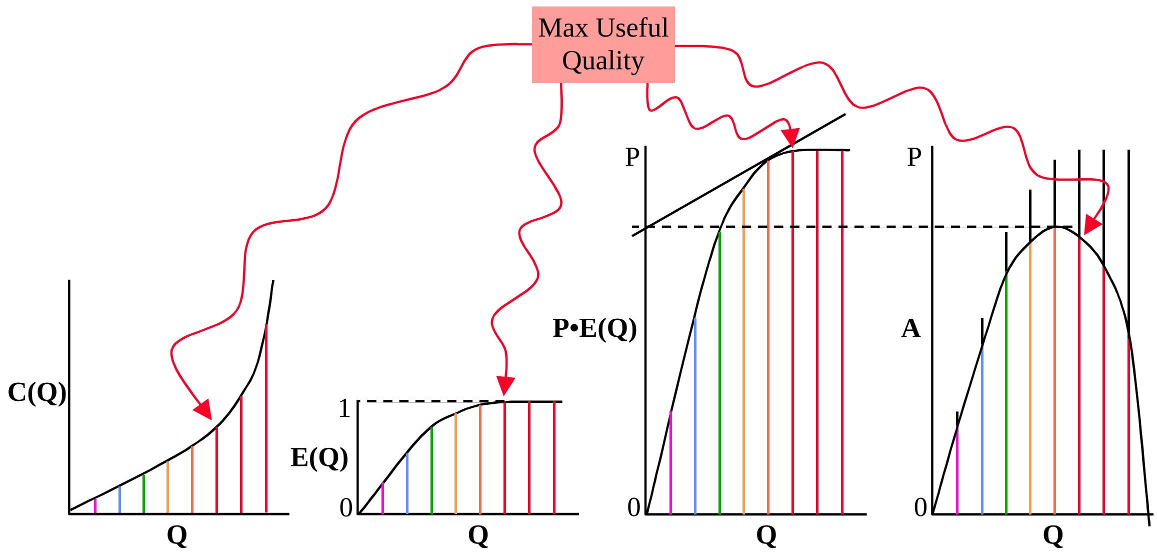

Fig. 3-4 – HazMon Cost/Benefit Relationships

If **E**(**Q**) was simply asymptotic and never actually reached a value of one, it would be necessary to spend ever more to improve prediction quality in order to save additional lives. Given that **E**(**Q**) does reach one, the expenditures necessary to save maximum lives are limited. Furthermore, the **Q** for which **E**(**Q**) = 1 can be determined from the characteristic quantization of the mitigation strategy applicable to any given hazard.

3.4 Timeliness and Accuracy

Prediction quality (i.e., the outputs from model resources) depends on accuracy (and resolution) and timeliness (initial and subsequent update rate) of input data; specifically, the following relationships are key:

- Higher accuracy has an inverse impact on timeliness, i.e., more time is required to obtain the desired sensing or modeling accuracy.
- Yet, higher accuracy is achieved by obtaining higher resolution sensor data

To realize the benefit of high-resolution data, a system must either:

Require more time to transmit, with the associated negative impact on timeliness

Or

Require significantly higher communications bandwidth, with an associated higher system cost to transmit. This impact is seen if utilizing:

- Higher fidelity models

- Higher order models, or
- More (iterative) model recursion

Additionally it is observed that:

> Models which require more time to compute (to achieve the desired accuracy) have an inverse impact on timeliness (i.e., longer compute time, less responsive outputs to Customers)

Or

> In order to improve timeliness, the system must require significantly faster computation with an associated higher computation cost

The basic timeline for defining characteristic times associated with HazMon is shown in Figure 3-5 below. Significantly, what is illustrated is the breakout of system-level delay times, which, when summed, are the minimum times required for HazMon to achieve the required accuracy to provide actionable information to Mitigation-side decision-makers. More detailed hazard-by-hazard analysis of this timeline, with the goal of determining the potential “technology gap”, is discussed in Section 4.

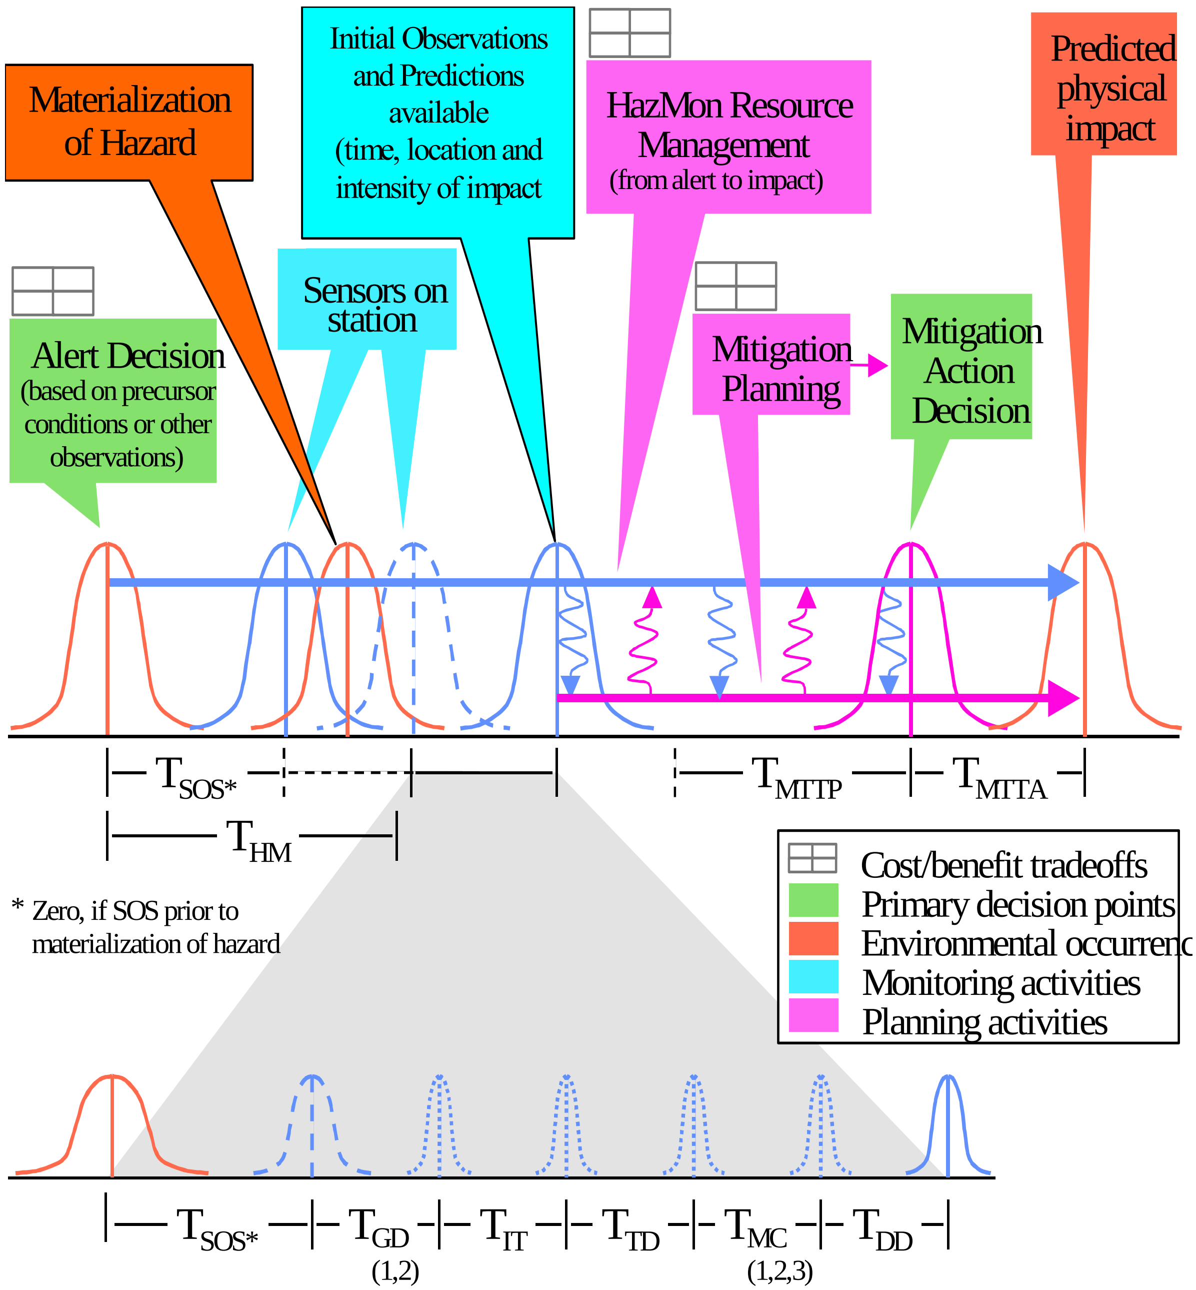

Fig. 3-5 – Characteristic Times

where

T_{SOS} = Time to sensor on station
T_{HM} = Time to hazard materialization
T_{MTTP} = Minimum time to plan
T_{MTTA} = Minimum time to act
$T_{GD(1,2)}$ = Time to gather data
(1) = Time required for actual data acquisition by sensor platform
(2) = Time to Plan, re-plan, re-locate sensor
T_{IT} = Time to initiate data transmission
T_{TD} = Time to transmit data
$T_{MC(1,2,3)}$= Time for modeling calculations
(1) = Time required to provide data to modeler
(2) = Time for database query & data access
(3) = Time to process, integrate, fuse data
T_{DD} = Time for data dissemination

Also note that

T_{SOS} & $T_{GD(1,2)}$
Depend on the type sensors used to monitor a particular hazard, their proximity to the hazard, and the ability to move them appropriately

T_{IT}
Depends on the communication mechanisms used by the sensors (satellite-based sensors may need to wait for ground station access)

T_{TD}
Depends on the volume of data required to get the necessary prediction quality for a particular hazard

$T_{MC(1,2,3)}$
Depends on the complexity of the model for the particular hazard, the total amount of data required, and the update rate of more recent sensor observations

T_{TD}
Depends on the locations and situations of the mitigators

4 Preliminary Requirements Analysis

A fundamental question in defining HazMon requirements is, "How good is good enough"? In other words, what are the required (or even desired) accuracies or minimum delay times which if realized would drive HazMon technology development? Presently we do not have such *a priori* requirements; if we did, specific technology gaps would be relatively straightforward to identify. But the observable world of hazard monitoring (sensing), data/information processing and delivery, and decision-making tools does not readily provide such "good-enough" specifications. Clearly, more accuracy, less processing time, and faster delivery are desirable, but how much, and at what cost?

This section identifies two means by which to answer the question "How good is good enough?" The first means is taken from the AIST Program, through the categories of Needs and Goals. The second means is a modeling approach. The modeling approach identifies each component in the overall system, identifies their relationships, prepares mathematical models for the components, and conducts a sensitivity analysis according to a set of measurable objectives. The approaches can complement each other in the effort to identify and validate a set of requirements.

4.1 AIST Program Needs and Goals

One look at this issue comes from the AIST Program Needs Documentation, in the following listing of AIST Needs & Goals:

AIST Needs Categories

Data Collection and Handling
Transmission and Dissemination
Data & Information Production
Search, Access, Analysis & Display
Systems Management

AIST Goals

Improve mission performance through automation and autonomy
Enable distributed heterogeneous sensorwebs
Enable seamless, ubiquitous communications networks
Improve transmission efficiency of large data volumes
Improve access/retrieval and scalability of the storage and management of large data volumes
Improve performance, flexibility and adaptability of data processing and networking
Improve organization and search of scientific data
Improve extraction and fusion of scientific data
Improve analysis of scientific data
Improve data access through metadata interoperability
Improve system interoperability and use of standards
Improve system management and operations

Improve reliability of hardware/software
Reduce life cycle cost of ground and space operations and processing

This information should not be taken totally out of context, because lists of potential technical approaches follow each goal; the significant point is that engineering analysis seems necessary to determine a measurable metric to associate with each goal, and an associated cost as well. From the perspective of an acquisitions manager deciding what technology developments to fund, such cost-benefit trade-offs are crucial, to ensure that expenditures to "Improve extraction and fusion of scientific data" actually improve the performance of HazMon. That is, lives and property are saved as a result of such "improvements". Again, the question to resolve remains, "How good is good enough"?

4.2 A Modeling Approach to Requirements Analysis

To answer "How good is good enough?" we assembled a simple model of HazMon systems to enable sensitivity analyses. The objective of the sensitivity analyses is to perform cost / benefit assessments. The results of a cost / benefit analysis would clearly identify the benefits (monetary and loss of life savings) compared to the costs associated with mitigation and losses due to damage. Cost / benefit ratios that are functions of timeliness and accuracy can readily identify which improvements in timeliness and accuracy are key to improving hazard mitigation and thus should be targeted for study. This section describes our chosen cost/benefit analysis method, describes a functional hazard scenario, describes the modeling techniques, and presents representative results from an analysis. These results are preliminary in that the model is incomplete. However, even preliminary quantitative results are significant to expose first-order technology gaps.

The data used as input to the models in this report is a combination of representative US Census data and hypothetical data. Hypothetical data is used where there was insufficient time to identify and collect real data. Real models and data need to be collected and validated, before results can be used to select technologies that will improve cost/benefit ratios. However, the results provided in this report indicate the types of trends that might be expected.

4.2.1 Attributes of a Cost / Benefit Analysis

Hazards can entail both monetary costs and loss of life. Monetary costs are based on mitigation expenses and damage to property. The probability of loss of life is a function of the hazard itself, the density of the population and the success of mitigation procedures. Mitigation is executed in an area where the hazard is expected to occur, and therefore its effectiveness is a function of the ability to predict the future state of a hazard. Mitigation effectiveness is related to "When should an area be mitigated?" Prediction of the future state of the hazard depends not only on hazard modeling algorithms but also the data, which comes from sensors, which are inputs to these algorithms.

Due to sensor, estimation and prediction errors, a portion of the mitigation area may reside outside of the area where the hazard actually occurs, and a portion of the hazard area may receive no mitigation. The larger these errors become, the larger the resulting monetary costs and the probability of loss of life. In order to reduce the likelihood of missing part of the hazard area, mitigation procedures incorporate safety factors. As uncertainty increases, safety factors increase and mitigation costs increase correspondingly.

Katz and Murphy present a prototype decision-making model in their book, *Economic Value of Weather and Climate Forecasts*[3]. This model associates a cost with a pair consisting of an event and an action. Applying that notion to this HazMon study, Figure 4.1 illustrates two kinds of costs as a function of a hazard occurring (H) or not and whether mitigation procedures are taken (M) or not. The cells are color-coded so that illustrations in the following subsections can be easily related to cells in these decision-making matrices. Purple represents an area where both mitigation and the hazard occurs. Red represents an area where a hazard occurs but mitigation doesn't, and blue represents the area where mitigation is executed but the hazard doesn't impact.

	$\overline{M}$	M
$\overline{H}$	0	C_3
H	C_2	C_1

Figure 4.1a. Cost Decisions

	$\overline{M}$	M
$\overline{H}$	0	P_3
H	P_2	P_1

Figure 4.1b. Probability of Loss of Life Decisions

Fig. 4.1 Decision-Making Matrices

The following subsections describe means of determining the trends of costs and probable loss of life, as a function of available time. The contents of cells in the decision-making matrices are then functions of time.

4.2.2 Error Analysis

If we predict the state of a hazard exactly, then mitigation occurs exactly in the area where the hazard occurs and therefore the costs C_2 and C_3 as well as the probabilities P_2

[3] Katz, Richard W., Murphy, Allan H., Economic Value of Weather and Climate Forecasts, Cambridge University Press, 1997

and P_3 are zero. (The only way of reducing costs, in this situation, is by changing mitigation strategy and performance.) However, since errors in measurement, estimation and prediction do exist, C_2, C_3, P_2 and P_3 will be non-zero. The following paragraphs examine how errors, as variations between the actual and predicted state, can be handled. These variations are examined within the context of a linear-first, order system. Although actual measurement and prediction errors are non-linear and higher order, this kind of analysis is sufficient to quantitatively bound the problem.

Consider the equation for a first-order, linear system with state, *x*, and external disturbances, *u* in Eqn. 4.1. In this model, A represents the dynamics of the environment and B represents the effects of external disturbances.

(Eqn. 4.1) $$\dot{x} = Ax + Bu$$

To further bound the analysis to just the dynamics of the hazard in the environment without the effect of external forces, consider the homogeneous form of the equation in Eqn. 4.2. The bounds of the hazard system can be considered arbitrarily, so there is no loss of generality that comes with this simplification.

(Eqn. 4.2) $$\dot{x} = Ax$$

The variation, δx, between the actual state and the predicted estimate of the state appears in Eqn.4.3. $\hat{x}$ is the predicted estimate and *x* is the actual state.

(Eqn.4.3) $$\hat{x} = x + \delta x$$

The homogeneous equation for the variation, δx , follows as shown in Eqn.4.4.

(Eqn.4.4) $$\delta\dot{x} = A\delta x$$

The well-known solution to this equation is:

(Eqn.4.5) $$\delta x(t) = \delta x(t_0)e^{-At} + K$$

We are interested in the effect of variations on timeliness, area and intensity. So, define the components of the state of interest to be as shown in Eqn.4.6.

(Eqn.4.6) $$\delta x = \begin{bmatrix} \delta T \\ \delta A \\ \delta I \end{bmatrix}$$

Furthermore, the variation in location of the area is a function of the variation in timeliness and the characteristic speed, V, of the hazard. This relationship is illustrated in Eqn.4.7.

(Eqn.4.7) $\delta l = V\delta T$

These considerations are the basis for the modeling that follows.

4.2.3 Effects of Delays

Delays contribute to growth of errors. For example, equation 4-5 represents the growth of error in the predicted state. If a delay, ΔT, is introduced, the resulting state error is represented by equation 4-8.

(Eqn.4.8) $\delta x(t) = \delta x(t_0)e^{-At+\Delta T} + K$

4.2.4 A Hazard Scenario Model

For the sake of a simple visualization of the growth and traversal of a hazard to and through an area of interest, consider the contents of Figure 4.2. At the top of the figure, the orange rectangle represents the path of the hazard. The red squares represent the area of the hazard at various times. The black dot within the red squares represents the location of the hazard while the arrow represents its speed. The blue square represents the mitigation area.

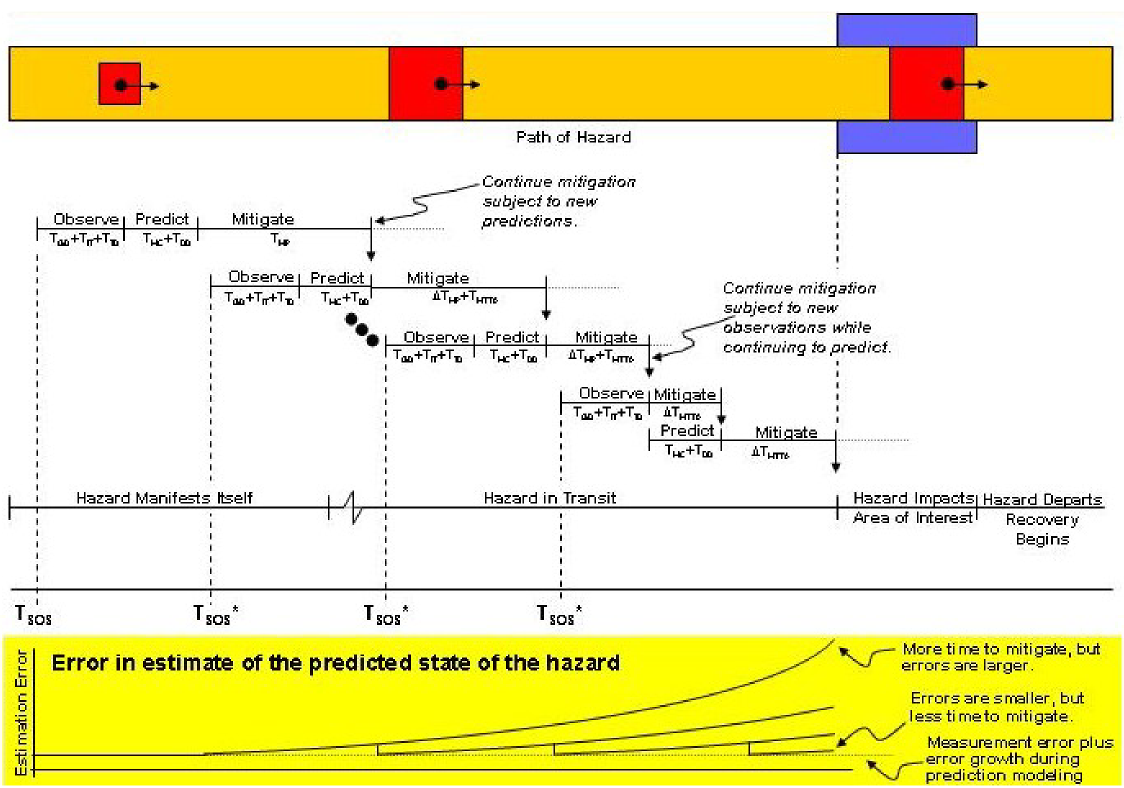

Fig. 4.2 Hazard Scenario Model

The timeline in the middle of Figure 4.2 illustrates cycles of observation, prediction and mitigation. Times associated with these activities are those that are defined in Section 3 of this document.

The plot at the bottom of Figure 4.2 illustrates growth of the error in predictions of future states. Note that the best estimate that a predictor can achieve is equivalent to the measurement error plus the effects of unobserved changes in the hazard area and speed

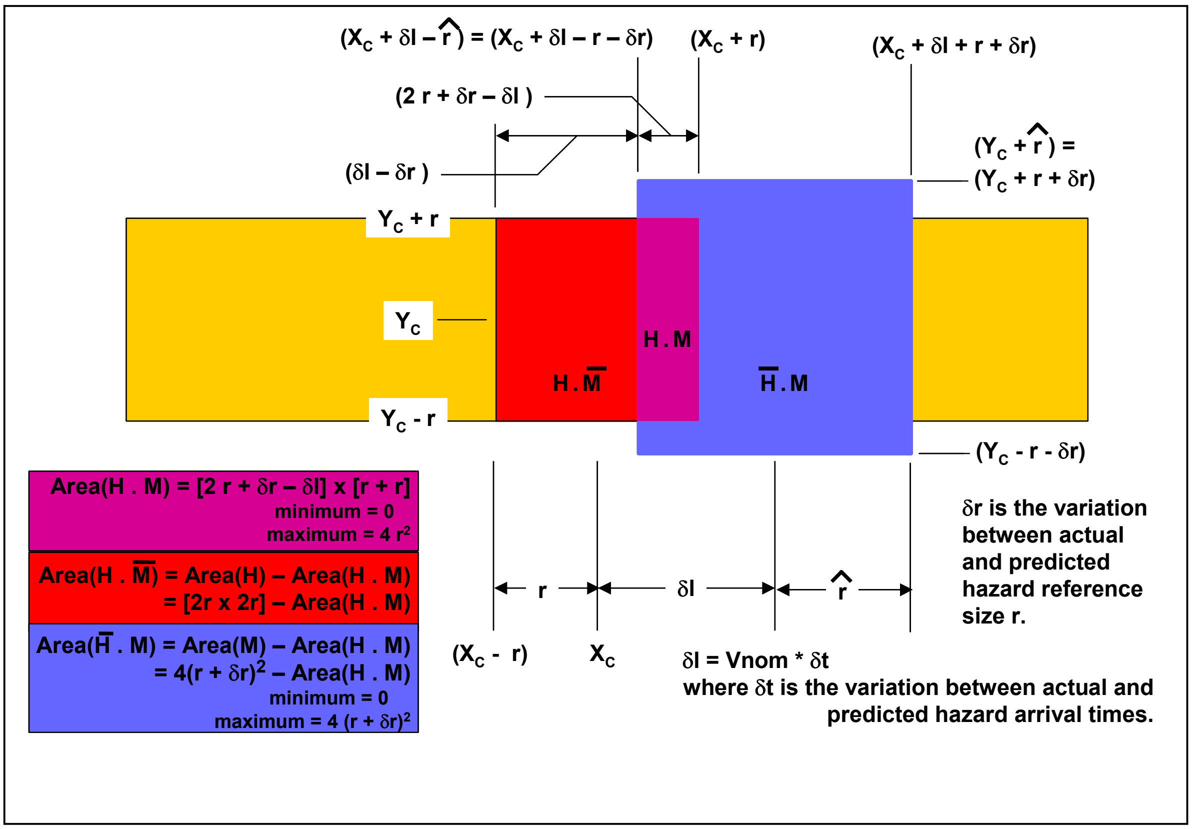

during the time it takes to achieve the prediction. The error in prediction always growths with time. The earlier that predictions are made the more time for mitigation, but the larger the errors in the predicted state of the hazard at the time of expected impact.

Fig. 4.3. Areas of Coverage during the Impact of a Hazard

A detailed examination of the area of mitigation relative to the area of hazard impact is illustrated in Figure 4.3. In this figure, red represents the area where the hazard occurs without mitigation. Purple represents the area where the hazard occurs and mitigation is performed. The blue area represents the area where mitigation was performed because the hazard was expected but the hazard in fact did not or has not yet occurred. In this figure, r equals half the length of a side of the hazard area. $\hat{r} \; (= r + \delta r)$ equals half the length of a side of the mitigation area. δr is the error in predicted length of a half-side, r. δt is the error in the predicted time of impact. Also, within Figure 4.3, X_c, Y_c are the coordinates of the location of the center of the hazard impact area.

4.2.5 The Modeling Technique

Modeling a complex system with all of its interactions is a tedious, iterative process. The whole hazard, population/property, mitigation, prediction and sensing system is such a system. A typical approach is to start with the assembly of piecewise linear, first-order models. Once a first-order model is developed and analyzed, higher degrees of fidelity and interactions can be added. (Delays as a result of sensing and mitigation planning are added as the first higher-order degree of fidelity in this report.) Results from a linear, first-order model often reveal valuable insights to system sensitivities to parametric variations. The ultimate goal is to understand the system's sensitivity to variations in accuracies within components of the sensor subsystem. The goal for this study is to begin building a model that enables us to understand the system's sensitivity to variations in model parameters. These model parameters represent characteristic properties of the hazards, property / population, mitigation, prediction and sensor systems.

Figure 4.4 illustrates the general linear, first-order model, as represented in the Laplace domain, that is used for each of the model components. K_x is the steady state value of the model and τ_x is the time constant of the model.

$$\boxed{\frac{Output}{Input} = \frac{K_x}{(1 + \tau_x s)}}$$

Fig. 4.4. General Linear, First-Order Model

The following subsections identify the component models, their relationships, model parameters, and representative results. The final subsection identifies further work that is necessary to complete the first-order model and continue with analysis and higher order modeling.

4.2.5.1 Model Parameters

Six model parameters are identified and described. The models are hazards, property and population, mitigation, prediction, estimation, and sensor models, shown in Figure 4.5 below.

Modeling Parameters

- Hazard Model
  - Reference size
  - Reference speed
  - Reference intensity
- Property Model
  - Average valuation
  - Density
  - Damage time constant
  - Mitigating action time constant
- Population Model
  - Density
  - Loss of life time constant
  - Mitigating action time constant
- Sensor Model
  - δr measurement error
  - δt measurement error
  - δi measurement error
- Estimation Model
  - Currently subsumed in sensor model
- Prediction Model
  - δr growth time constant
  - δt growth time constant
  - δi growth time constant
- Mitigation Model
  - Mitigation planning time constant
  - Mitigation Costs
    - Property
    - Population
  - Scale Factors
    - δr
    - δt
    - δi

Fig. 4.5 Modeling Parameters

4.2.5.1.1 Hazards Model

A range of hazards was considered; hazards are discussed in greater detail in Section 4.2.5.3.1 below, and some characteristic properties shown in Figure 4.8 below.

4.2.5.1.2 Property and Population Models

Property and population are modeled in terms of density. Property is also modeled in terms of valuations that can be mitigated and cannot be mitigated as a function of hazard intensity. Population is also modeled as a probability of loss of life as a function of population density, hazard intensity and hazard duration.

The property and population models have separate time constants for mitigation actions.

4.2.5.1.3 Mitigation Model

There are separate mitigation models for property and population. These mitigation models focus on mitigation activities that can take place after a hazard is first detected, such as deactivating utilities, boarding windows, etc. for property and evacuation for the populace. These mitigation models do not address mitigation activities that can occur prior to detection of a hazard, such as improving building codes.

There is a time constant for mitigation planning. The time constants for mitigation actions are incorporated into the property and population models.

4.2.5.1.4 Prediction Model

The significant aspects of the prediction model are its error growth rate and the total duration of modeling computations plus data dissemination. The error growth rate is modeled by a time constant. The prediction cannot be used until the prediction computations are complete, and since the computations consume non-zero time the error that exists at the end of the computation will increase as the time to predict increases.

4.2.5.1.5 Estimation Model

Drawing from navigation's use of the term "Estimation", estimation is the process of incorporating sensor measurements into an updated prediction of the state of a system. Sensors often do not directly measure the state of the system. Sensors often measure components of the state or measure quantities that map into components of the state. For purposes of simplifying this first cycle of system modeling, we define that the estimation model is perfect and that the sensors measure components of the state.

4.2.5.1.6 Sensors Model

The sensor models are also linear, first-order. The significant aspects of the sensor models are their measurement errors and their durations of time-on-station, time-to-gather-data, time-to-transmit. The error growth rate is modeled by a time constant. The prediction cannot be used until the prediction computations are complete, and since the computations consume non-zero time the error that exists at the end of the computation will increase as the time to predict increases.

4.2.5.2 Relationships Among Components

Figure 4.6 illustrates the relationships among the components. The losses include both monetary losses of property as well as the probability of loss of life. Mitigation costs are the monetary resources that can be consumed by the mitigation process. Other interfaces among the model components should be obvious from the diagram.

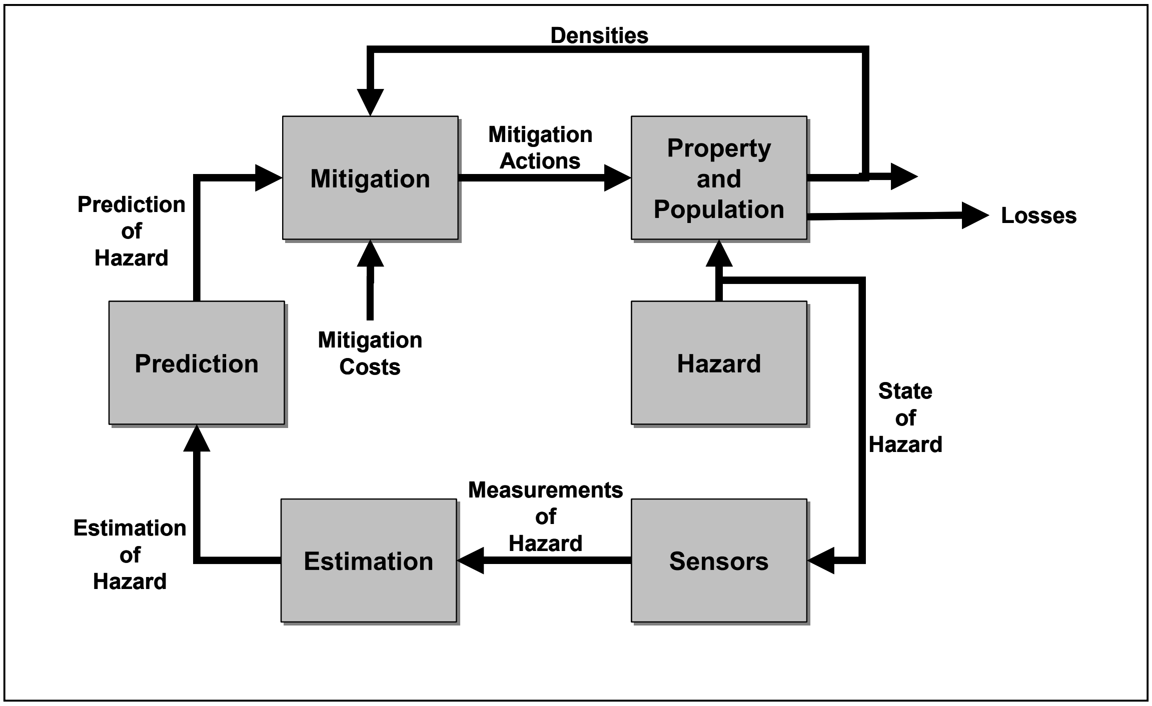

Fig. 4.6 Relationships Among Model Components

4.2.5.3 Model Parameters

Figure 4.7 summarizes the parameters that are implemented in each of the models. The following sections provide references from which values of these parameters were derived. These derived values were used during execution of the model and analysis. Values for these parameters are given in Figure 4.8.

4.2.5.3.1 Hazard Parameters

Representative values of parameters for hazard models are identified in the top portion of Figure 4.8. These include characteristic size, Rnom, and speed, Vnom, parameters.

Hazard Parameters	Generic	Hurricane	Tornado	Volcano	Pestilence	Oil Spill	Wildfire	NBC	Earthquake
Rnom (km)	50	300	5	10	50	50	50	10	200
Vnom m(Km/hr)	20	20	50	10	0.01	0.01	1	20	200
Imax	10	10	10	10	10	10	10	10	10
Tau Loss of Life (LoL)	8	8	8	8	8000	4	1	8	8
Tau Property Damage	8	8	8	8	8	4	1	8	8
% of property value damagable at Intensity=5	0.1	0.1	0.8	0.1	0.1	0.1	0.8	0.01	0.8
% of property density damagable at Intensity=5	0.5	0.5	0.8	0.5	0.5	0.5	0.5	0.1	0.8
Prediction Parameters									
Tau dr growth	5	10	15	20	36	24	12	10	5
Tau dt growth	10	5	10	10	24	24	8	5	10
Mitigation Parameters									
Tau Population	10	10	10	10	2000	20	10	20	10
Population Mitigation Cost Rate	5	5	5	5	5	5	5	5	5
Tau Property	8	8	8	8	48	10	8	10	8
Property Mitigation Cost Rate	50	10	10	5	0.00000022	0	100	0	0
Total Mitigation Delays (hrs)	0	0	0	0	0	0	0	0	0
Sensor Parameters									
dr measurement error	5	10	0.05	5	5	5	1	5	5
dt measurement error	0.2	0.2	0.1	1	1	1	0.25	1	0.1
Total Sensor Delays (hrs)	0	0	0	0	0	0	0	0	0
Environment Parameters									
Area (sq mi)		438	665	1679	2576	600,000	2362	23	789
Area (sq km)		1,134	1,722	4,349	6,672	1,554,000	6,118	60	2,044
Population		232,064	82,068	719,407	11,621	0	20,255	1,541,150	2,890,444
Population Density (Persons/mi^2)	327	530	123	428	5	0	9	67007	3663
Population Density (Persons/km^2)	126	205	48	165	2	0	3	25871	1414
Number of Property Units		99,683	35,163	277,060	26,265,913	1,000,000	11,217	798,144	969,484
Median Property Valuation ($K)	200	95	83	150	0.00221	0.991	268	1000	270
Property Density (Property/mi^2)	518	228	53	165	77,373	2	5	34,702	1,229
Property Density (Property/km^2)	200	88	20	64	30,224	1	2	13,398	1,116
Representative Locations									
State		GA	OH	WA	NE	TX, LA, AL, MS, FL	CO	NY	CA
County		Chatham	Muskingum	Pierce	Custer	Gulf of Mexico	Routt	New York	Orange
Largest City / Town / Landmark		Savannah	Zanesville	Mt Ranier	Broken Bow		Sarvice Creek	New York	Irvine
Reference		1	1	1	2,3,4	5,6,7	1	1	1

References

1 http://quickfacts.census.gov/qfd/
2 http://www.hort.purdue.edu/newcrop/cropmap/nebraska/counties/custer.html
3 http://www.grainline.com/asp/ldp/countyldp.asp
4 Property is corn yield has been 26,265,913 bu / 217,261 acres
5 http://www.mms.gov/eppd/socecon/techsum/gm/29143.doc
6 http://www.epa.gov/gmpo/governance.html
7 Fish and shellfish annually net $991M

Fig. 4.8. Values of Parameters

Time constants for damage, Tau (Property Damage), and probability of loss of life, P(LoL), are currently assigned arbitrarily. Trends in sensitivity to changes in time constant values are more important than the precise values at this point in modeling.

The hazard model provides for specifying the percent of a properties valuation that can be damaged by a given type of hazard. Values for these percentages are currently assigned rather arbitrary values. A further study is necessary to incorporate more realistic data.

4.2.5.3.2 Property / Population Parameters

A brief search of the US Census web site provided some property and population data. Reference 1 in Figure 4-8 identifies the web site. This census site provides the county size, population, number of properties and property valuation data.

Nebraska crop yield data for corn comes from a Purdue University site identified by reference 2. Reference 3 identifies the source for corn commodity prices in Custer County, NE.

4.2.5.3.3 Mitigation Parameters

Time constants for executing mitigation actions use arbitrary values. Again further examination and consensus on realistic values with mitigation authorities need to occur.

Delays due to mitigation planning are shown as zero in Figure 4.8. Values for these delays are varied parametrically from zero to 1 hr to 2 hrs in simulations of the model. These values are currently arbitrary. Realistic data needs to be acquired from mitigation agencies..

4.2.5.3.4 Sensor Parameters

Values of sensor parameters are also considered arbitrary at this point. However, even though they are arbitrary, they are based on representative figures of merit that have been deduced by examination of various sensor systems.

Delays due to sensors are shown as zero in Figure 4.8. In results that are produced, they are parametrically varied from zero to 2 hrs to 4hrs.

4.2.6 Representative Results

This section describes a representative set of results for a generic hazard. Results for each of the hazards are available in Appendix B.

Data from Figure 4.8 is used to produce these results.

Figure 4.9 illustrates trends for costs, losses and probable loss of life. The horizontal axis represents the time available to mitigate once a hazard is detected. The model simulated a 24-hour period.

In general the expectation is that more time available for mitigation leads to lower monetary costs, losses and loss of life. However, with more available time, predictions farther into the future need to be made. These longer-range predictions come with larger errors. Therefore mitigation can occur over the wrong area. Errors in prediction are reduced as the time of impact approaches. More of the impact area will be mitigated and less un-affected areas will be mitigated – however, there is naturally less available time to mitigate. The generic hazard in Figure 4.9 illustrates this phenomenon, described in the following text.

The top frame illustrates (reading from right to left) that for each hazard intensity level other than zero, for a given probable loss of life at 24-hours this probability reduces as less mitigation time is available and then tends to increase again as less mitigation time is available. The reason for the initial reduction in the 24 to a10-hour period is that prediction errors decrease, Therefore more of the area where the hazard will actually impact is being mitigated. In the 10 to 0 hour period, probable loss of life increases. The reason for this is that even though predictions of the impact location and area are getting better, there is less time to conduct mitigation actions.

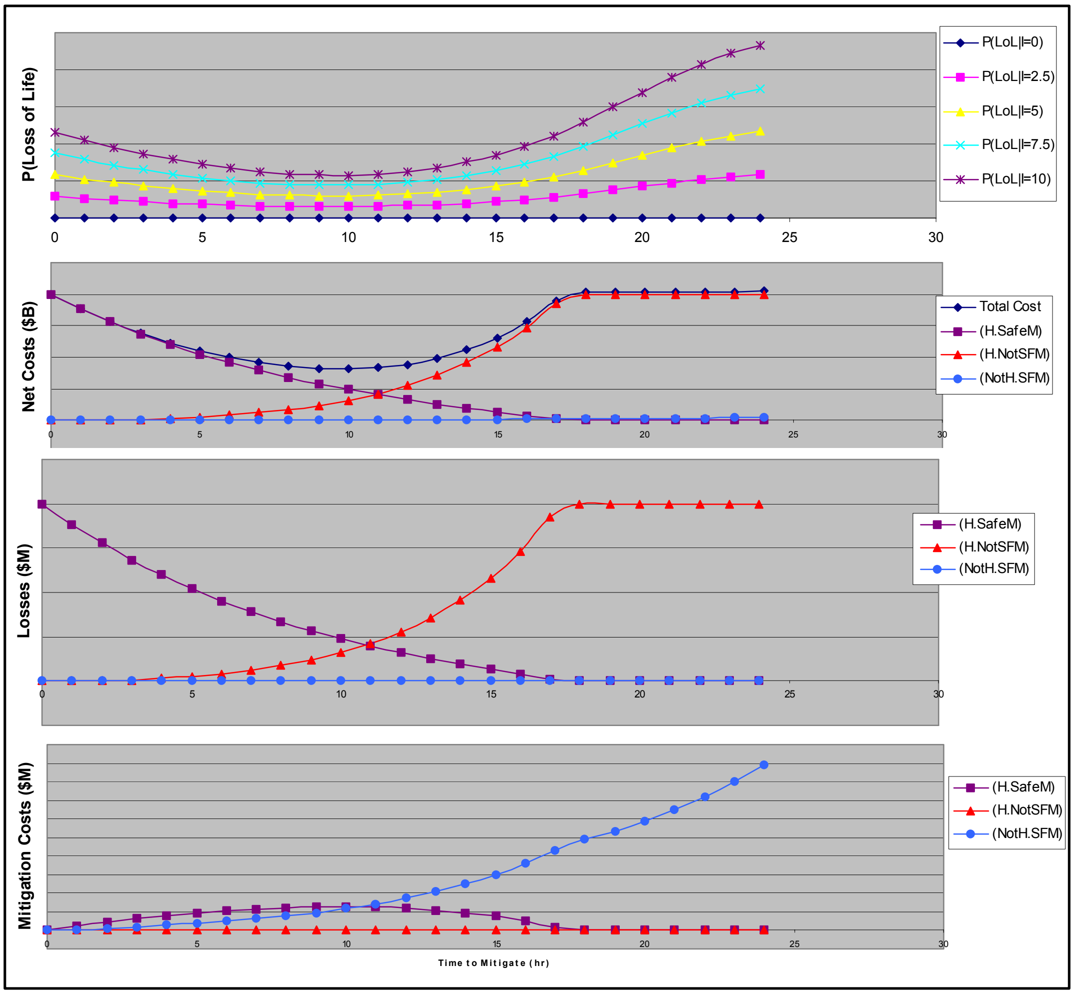

Fig. 4.9. P(LoL) and Costs for a Generic Hazard

The Net Costs, in the second frame, are the sum of the Losses and Mitigation Costs that are plotted in the bottom two frames. In addition to the net total costs, also plotted are net costs for the areas where:

- the hazard impacts and mitigation actions are executed (H.SafeM),
- the hazard impacts and mitigation actions are not executed (H.NotSFM), and
- no hazard impacts but mitigation actions are executed (NotH.SFM) are plotted.

Notice that in reading right to left in the bottom plot frame that mitigation costs are highest for the area where the hazard doesn't impact (the blue line) and then decreases as time to impact gets closer. Again, this is because errors in predicted location are high and then reduce.

In the middle plot frame, notice that losses are high in the hazard area (red line) when actions need to be taken 24-18 hours in advance. This is because mitigation actions are being taken in the wrong area. The purple line indicates that losses tend to increase from 17-0 hours in the area where both the hazard exists and mitigation actions exist. This is a two-fold reason for this increase. First more of the hazard area is being mitigated as available time decreases. This is because errors are decreasing. However, even though more of the hazard area is now known so that it can be mitigated, there is less time to conduct mitigation actions.

[Note: The "Safe" and SF prefixes mean that a Safety Factor on the mitigation area is imposed. For the examples in this report the value of the safety factor, however, is 0. One would expect that as the value of the safety factor is increased monetary losses and probable loss of life will decrease but mitigation costs will increase.]

Figure 4.10 illustrates trends of P(LoL) and Net Costs as total delay (sensor + mitigation) is increased from zero to 3 hrs to 6 hrs. The delay due to sensors represents the times $T_{GD} + T_{IT} + T_{TD}$. These times are described in Section 3. The times due to mitigation represent the times $T_{MC} + T_{DD} + T_{MP} + T_{MTTA}$. Notice that cost and the probable loss of life increase as delays are increased. The inflection at 10 hours could be significant. This point does not appear to change but the time window around it for acting and having an effect on costs decreases as delays increase.

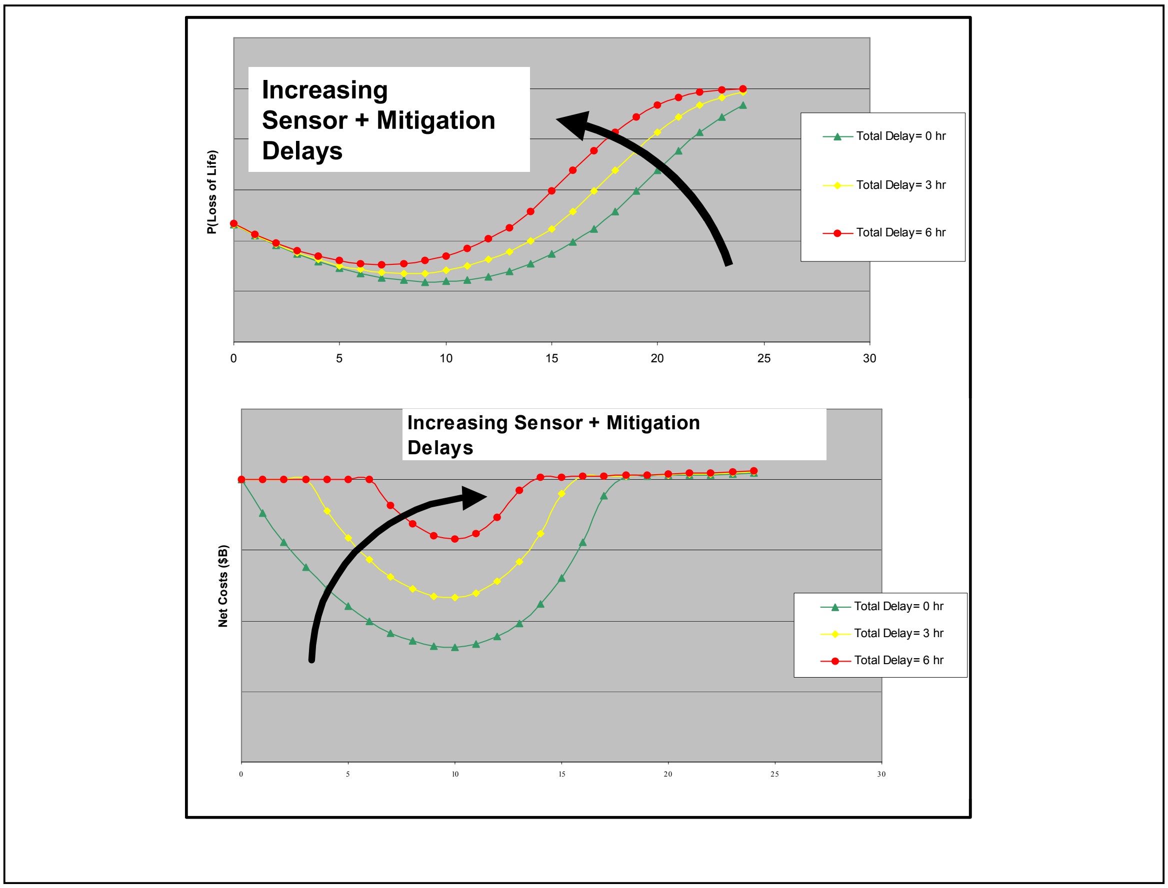

Fig. 4.10 Delay Trends in P(LoL) and Costs for a Generic Hazard

Appendix B contains results for each of the hazards identified in this report.

4.2.7 Further Modeling Work

The value of cost / benefit analyses through modeling is important. However, the models must adequately represent real situations. Contact and interchanges with experts in related fields are necessary to validate the models used in cost / benefit analyses.

Contact and interchange with experts that have experience with hazard modeling, such as weather is necessary. Web sites and documentation provide limited insight. For example, Specifications that have been located on weather models, e.g. MM5, OMEGA, and RAMS, identify resolutions on their prediction capability, such as 50km, 10km, etc. However, their specifications do not identify what kinds or resolutions of sensor data that is necessary to produce these kinds of results. The model should be extended for the purpose of understanding the impact of sensor resolution.

Reasonable contact and interchange with mitigation agencies are needed in order to provide insight to values of modeling parameters for planning and mitigating property and populations.

Currently, models address the dynamic nature of the location and size of hazards. The intensity is held constant. Future modeling should address the dynamic nature of hazard intensity.

The visual output of results can be improved so that patterns become obvious. Currently, small, yet significant, deviations in costs may not be observable because total costs are being plotted. It might be better to plot costs that deviate from a baseline situation. For example, a baseline situation might be one in which there are no errors and no delays.

Also, the efficient frontier that is identified in Section 3 could be plotted with real data if a spectrum of data is collected and run through the models. Right now decisions on technology gaps are made on an analytical argument. Quantitative results could strengthen these arguments where they matter most.

5 Potential Technology Areas for Further Investigation

As a result of the HazMon research, a number of technology areas were identified as Key Capabilities; that is, areas of existing capability with the potential for significant impact on predicted HazMon performance. These are listed in Table 5-1 below. None of these technologies appear as 'novel' or 'beyond the horizon'. It is significant, however, that these technologies fall into two categories: those which are "self-improving" due to existing, on-going commercial/government investment and development, and those for which development can be accelerated by ESTO leadership and sponsorship.

Table 5-1 Key Capabilities for Improved HazMon Performance

Technology	Category
Platform-driven (sensor) Technologies	
• Getting Platforms on station	Self-improving
– Minimization of Orbit delays	Self-improving
– Employ/enhance/integrate UAV technologies	Self-improving
– Ballistic deployment of low altitude sensors	Self-improving
• Data Communication	
– Commands and Data Uplink/downlink	Self-improving
• Laser-based satellite communications	Self-improving
– On-board processing	Self-improving
• Data Compression	Self-improving
• Edge-finding algorithms	Self-improving
Data Processing Technologies	
• Retrieving external inputs from:	
– Data archives	Self-improving
– "Collaborating" systems (e.g. purchased data)	Self-improving
• Model products (flooding, wind damage, etc.)	Self-improving
• Data Processing	
– Fusion of various internal and external "information" products	Self-improving
– Constant updating of system operational strategy	Self-improving
– Models (flooding, wind damage, etc.)	Self-improving
• Product Dissemination	
– Visualization/presentation of results	Self-improving
– "Broadcasting" delays	Self-improving
• Development of Hazard Information Interchange standards	ESTO sponsored
• Cost-based rules for products and resources	ESTO sponsored
• Rules-based policy management/engine (internal decision-making)	ESTO sponsored
• Resource Optimization/Planning Tools	ESTO sponsored

Technology	Category
• Real-time Policy and Resource negotiation mechanisms	ESTO sponsored
• Real-time dynamic Resource tasking/re-tasking	ESTO sponsored
• Inter-system interfaces and standards	ESTO sponsored
– Protocols	ESTO sponsored
– API wrappers	ESTO sponsored

It is important to note that the required technical capabilities are presently sufficiently mature to allow a HazMon system-of-systems to be "stood up" and operated; there appear to be no "technology gaps" which would *prevent* HazMon from being fielded. Indeed, it seems reasonable to conclude that such a baseline system *should* be established as soon as possible, first to actually measure baseline performance, but just as importantly to give legitimacy to the concept as a whole.

Self-improving technologies include such categories as 'data processing', 'data communications', and 'data dissemination'. These are so-described because current market forces are driving growth in these technologies faster than the projected HazMon system could take advantage. Communications and throughput are generally being driven by the telecommunications and military domains. Also driven by the military domain, data fusion and visualization technologies are being improved almost faster than systems integrators can successfully deal with. The GIS (Graphical Information System) domain, for example, is pioneering significant advances in technology specific to hazard monitoring, information dissemination, and hazard mitigation[4]. With regard to raw computing power, while various estimates of massive volumes of environmental sensor data to be processed do exist, such estimates do not seem consistent with present-day or near-to-medium term projections of the data necessary to provide significant improvements in present-day forecasting. Indeed, data processing rates and throughputs are generally continuing to increase according to Moore's Law[5]. It seems generally reasonable, therefore, to extrapolate that for the foreseeable future, commercially available technologies and capabilities will easily support projected HazMon demands. The AIST program itself has addressed the needs and desires of a potential HazMon user

---

[4] For reference, see the site of ESRI, Inc, a commercial GIS vendor, providing integrated tools with specific application to hazard monitoring and information dissemination: http://www.esri.com/hazards/index.html; the header from this site: "Helping to Build Disaster Resistant Communities: FEMA and ESRI have formed a National Partnership in part aimed at providing multi-hazard maps and information to U.S. residents, business owners, schools, community groups, and local governments via the Internet. The information provided here is intended to assist in building disaster resistant communities across the country by sharing geographic knowledge about local hazards."

[5] Moore's Law: The observation made in 1965 by Gordon Moore, co-founder of Intel, that the number of transistors per square inch on integrated circuits had doubled every year since the integrated circuit was invented. Moore predicted that this trend would continue for the foreseeable future. In subsequent years, the pace slowed down a bit, but data density has doubled approximately every 18 months, and this is the current definition of Moore's Law, which Moore himself has blessed. Most experts, including Moore himself, expect Moore's Law to hold for at least another two decades.

community, not only with a roadmap of technology areas, but Key Investment directions as well. These technologies encompass much of what is key to the successful fielding of HazMon.

With respect to potential areas for ESTO leadership and technical development, it is important to recognize existing ESE partnerships with state and local governments, Federal Emergency Management Agency, Environmental Protection Agency, Department of Defense, National Oceanic and Atmospheric Administration, US Geological Survey, National Institute of Health, Center for Disease Control, the US Army Corps of Engineers, and the Pacific Disaster Center. This may not be technically stimulating to developers and/or implementers of systems, but is extremely important to success of HazMon because the political and organizational aspects of HazMon cannot be overemphasized; technical history is rife with highly qualified systems which did not succeed because of poor organizational development and lack of political support (perceived or real), including funding within the Federal Budget cycle. Table 5-2 then lists potential areas for ESTO leadership and technical development with particular emphasis on realizing the HazMon system-of-systems.

The initial positioning of HazMon as a viable and important tool will depend on the ability of system architects to provide interface to legacy systems while simultaneously adding value as a system-of-systems. Longer term, HazMon can increase its value by developing interface standards to be applied to future systems to reduce the overhead imposed by continuing to integrate stovepiped systems.

6 Conclusions

As a result of this research and modeling, the following conclusions can be made:

1. Monitoring and predicting hazards is as much a problem of planning and coordination as a problem of sensing, predicting, processing and dissemination.

2. A HazMon system, as postulated, is useful even with today's sensing and hazard prediction technology, with the existence of stovepiped systems which can be more usefully combined.

3. First-order modeling has identified which characteristic times are most critical to hazard mitigation, and hence the best targets for technology investment.

4. Existing economic and military development forces will adequately close some of the characteristic time issues without specific development effort by NASA. ESTO should concentrate on those technologies that are not subject to these outside forces.

To summarize:

A concept of operations and a high-level architecture has been developed for a HazMon system that brokers/coordinates a broad set of sensor, communication, and computation resources. A set of hazards for which HazMon is useful has been defined, and a preliminary set of technologies for possible development identified. Modeling techniques and specific models have been generated, which, if expanded, will lead to developing accuracy and timeliness requirements for HazMon predictions, based on the characteristics of the related hazard mitigation mechanisms. Technology needs and potential investment areas are therefore derivable from these accuracy and timeliness requirements. Continued work is necessary to bound the required quality and errors associated with hazard predictions.

Appendix A – Concepts and Terminology

The following terminology is used to characterize the concept of operation and the terms (concepts) used to describe functional architecture of the HazMon system.

A.1 Customer

An organization that engages the HazMon system to provide information about the evolution of an existing or potential environmental hazard.

A.2 Resource Owner

This is the organization that actually owns and controls a set of resources. In order for HazMon to utilize the resources owned by another organization, resource negotiation protocols and resource tasking protocols need to be agreed upon in advance.

A.3 Resources

These are the sensors, computers, application-level algorithms, archived data, and communications links that support HazMon system operation. Some of these may be owned by the HazMon system itself, but the great majority will be owned by other organizations—public and possibly private.

A.3.1 Sensor Resources

Sensors serve as the eyes (and ears, nose and touch) of HazMon by gathering data about the environment. HazMon directs the operation of a sensor when it has been allocated to the HazMon as the result of a negotiation. Directives to sensors owned by external systems are communicated to a pre-designated interface of the owning system via a sensor tasking protocol.

Sensors may be deployed in space (on satellites), in the atmosphere (on manned and unmanned aircraft, and balloons), on the ground, and in oceans, lakes and rivers (tethered, free floating, or on manned and unmanned vehicles.)

A.3.2 Sensor Platforms

Sensors in space and the atmosphere, and some sensors in the water are deployed on vehicles or other platforms such as satellites or balloons. Tasking these sensors may also require tasking of the platforms.

Platforms may host significant computation capabilities to serve all of the sensors aboard.

A.3.3 Computation Resources

These comprise computer hardware and application programs that 1) operate on sensor data to produce HazMon products, and 2) coordinate HazMon operations. Computer resources include large ground-based number crunchers that produce final HazMon products, and platform-based computers that may perform some data reduction and/or data fusion of local sensor data. They do not include dedicated, embedded processors that support only their own sensors. These are considered part of the sensors themselves.

A.3.4 Archived Data Sets

These include historical observations that are useful to various prediction algorithms (wind pattern data, seasonal sea temperature variations, etc.), Digital Terrain Elevation Data (DTED), saved HazMon products, etc.

A.3.5 Communications Resources

These comprise all of the communications links among sensors (sensor platforms) computers, and the customers. Of particular interest are those links, such as satellite ground station, that need to be scheduled.

A.4 Interface Mechanisms

A.4.1 Protocols

Protocols specify the rules and procedures for communicating requests and data between HazMon and Customers, Resource Owners, and Resources. A protocol specifies the message formats, the parameters, and the interaction logic required to carry out such communication. In order to facilitate *forward compatibility* of the HazMon protocols, the protocol messages will be written in XML. Name-value pairs will specify parameters.

A.4.2 Application Wrapper

An application wrapper enables a legacy application to interface with HazMon. It translates the interface of the legacy application into the appropriate HazMon protocols.

A.4.3 Metadata

Metadata describes the syntax and semantics of other data, allowing a user to interpret the content and meaning of that data, e.g., the description of a file format and parameter representation, or the description of the structure and vocabulary of a document. Metadata is most useful when its own syntax and semantics are widely understood. A user can use a metadata description to interpret data whose structure is not known a priori. This is a powerful concept that enables *forward compatibility*—the ability of an "old" application to adapt to "new" data structures, e.g., to adapt to a new protocol.

A.5 HazMon – Customer Interfaces

A.5.1 HazMon Job

A HazMon job consists of those activities necessary to monitor, track, and/or predict the evolution of the specified environmental situation. They include planning the interactions among the sensor, computation and communication resources allocated to the job, and tasking those resources in an appropriate manner. Tasking is effected via a resource tasking protocol between the HazMon system and external systems that provide accessible resources. Sensor resources are tasked to observe various aspects of the environmental situation; communications resources are tasked to move data among sensors and computers; and computer resources are tasked to perform operations that generate the desired output for the customer.

A.5.2 Job Request

The specification, by a customer, of 1) a situation to be monitored, 2) the format and method of delivery of the results, 3) the desired confidence level of the results, 4) priority of the request, and 5) cost constraints. HazMon is intended to monitor situations and to predict their evolution, but not to predict the occurrence of such situations. The situation to be monitored is determined by the customer and is generally characterized by a geographic location and a time at which the situation began, or is expected to likely begin. In some cases the geographic area may be fairly large, such as watching for the occurrence of forest fires in the Southwest during a drought in summer.

A.5.3 Job Negotiation

The process by which the HazMon system acknowledges a job request and indicates its ability to fulfill that request to the confidence level desired by the customer under the cost constraints. The ability to meet a request's specifications will depend on the availability and cost of resources accessible by HazMon, and by the characteristics of existing HazMon jobs and other job requests. If HazMon cannot fulfill all aspects of a request, it will reply with alternative specifications including 1) what can be done for the stated cost, and 2) what cost would be required to meet the specification completely. The customer can accept one of the alternatives, vary the specification's parameters in some other way and resubmit the request, or appeal to a higher-level policy maker for an increase in priority.

In some cases, a new job request of high priority may need to use specific resources previously planned to support existing jobs. In such cases, HazMon will attempt to reassign resources or to access new resources in order to maintain the specified level-of-service for all current jobs. If this cannot be done, HazMon will notify the affected customers of the impact and renegotiate the specifications for the affected jobs.

A.5.4 Job Negotiation Protocol

The protocol that supports interactions between HazMon and Customers.

A.5.5 HazMon Product

The output of the observations and computations associated with a HazMon job. The product is provided in the form specified in the job request. Formats may include text, image, database, and video files, files in customer specific formats, and possibly streaming video. Products are delivered to the customer per the appropriate product delivery protocol.

A product may be the result of extensive computations utilizing data provided by many sensors, or it may simply be unprocessed sensor data.

A.5.6 Product Delivery

HazMon directs products to the specified customer computer. The customer is responsible for further dissemination of those products. If the computation that produces the final product happens to be on a computer that is owned by the customer, then the delivery is implicit.

Products may also be directed to an archive, which would be made available to the research community. Protocols for such a delivery may be different than those for delivery of products to a customer.

A.5.7 Product Delivery Protocol

This is the protocol that *manages* delivery of products from Resources to Customers. The product will actually be moved to the customer via an appropriate data transfer protocol (e.g., FTP).

A.6 HazMon – Resource Owner / Resource Interfaces

A.6.1 Resource Negotiation

When HazMon wishes to employ resources owned by another organization it enters into a real-time interaction with that organization, using protocols agreed upon in advance. HazMon specifies which resources it requires at what times, and what it is willing to "pay" for them. The owner may accept that offer, or may issue a counter-offer. The process iterates until some agreement is reached.

A.6.2 Resource Negotiation Protocol

HazMon can negotiate with other systems, in real-time, for access to their resources. This protocol supports those negotiations.

A.6.3 Resource Tasking

As a result of resource negotiation, HazMon can request that external resources perform certain actions. External resources cannot be controlled directly by HazMon. Tasking requests must be sent via a predetermined interface with the controlling system. This interface needs to be defined in advance with the owner of the system.

A.6.4 Resource Tasking Protocols

The resource tasking protocols supply the mechanism for HazMon to task resources to do a certain job. Depending upon the resource, the interaction may be with the resource owner or with the resource itself.

A.6.5 Resource Status Assessment

Resource status indicates when and under what conditions various capabilities of a resource will be available to support HazMon.

A.6.6 Resource Status Assessment Protocols

The resource status assessment protocols supply the mechanism for HazMon to determine resource availability. These protocols support both polling of resources by HazMon and advertising by resources or resource owners. Depending upon the resource, the interaction may be with the resource owner or with the resource itself.

A.6.7 Resource Advertising

A mechanism by which resource owners (or the resources themselves) inform HazMon of resource availability status.

A.7 HazMon Internal Operations

A.7.1 Planning

Planning is the central activity of the HazMon system. The HazMon Planning function determines which resources are needed to satisfy job requests and how those resources should be tasked.

A.7.1.1 Optimization

HazMon attempts to allocate resources to jobs so as to minimize an aggregate cost function that includes the costs of the hazard impacts and the cost of resources to track those hazards.

A.7.1.2 Hazard Evolution Models

These models utilize sensor inputs to predict the evolution of hazards.

A.7.1.3 Hazard Mitigation Models

Given predictions for a particular hazard, these models are used to predict the benefits of applying a particular mitigation strategy to the hazard—benefits measured in costs avoided and lives saved.

A.7.1.3 Resource Costs

Resource costs represent the cost of utilization of resources. In some cases these cost of resource operation will have been prepaid by government appropriations (sunk costs). In other cases operating budgets may either be insufficient or non-existent.

A.7.1.4 Hazard Impact Costs

The ultimate objective of the HazMon system is to (help decision makers) minimize the social and economic cost associated with natural and man-made disasters.

A.7.2 Credit Management

Customers of the HazMon system "pay" for delivered products and resource owners get compensated for the use of their resources. Credit management handles the flow of "value" among all of the HazMon participants.

A.7.2.1 Credit Bank

The credit bank keeps track of the allocation of "value" in the system.

A.8 Automation vs. Autonomy

In a given system, automation allows certain operations to be performed without a man-in-the-loop. The system may be given a high-level objective by an operator and manage lower-level operations, by itself, in order to meet that objective. For example, an autopilot on an airplane is told to hold altitude and heading. The autopilot uses information from the altimeter, the compass, and the attitude gyroscope to determine the commands to the aircraft's throttle and control surfaces.

A fully autonomous system, on the other hand, is designed to modify its own objectives, and/or to initiate significant actions on its own. For example, an autonomous vehicle with multiple objectives could decide on its own to drop, or modify, some of those objectives, if they could not all be met due to unanticipated circumstances. An autonomous weapons platform could seek out targets and operate against them. These are decisions of the sort that are better made by a human operator, but could be relegated to the system because it must perform in a situation where interaction with a human operator is not possible.

HazMon will necessarily be a highly automated system in that operator control is via the formulation of high-level objectives to track/monitor specific hazards or potential hazards. HazMon itself determines the best mix of resources to utilize to satisfy those requests, and manages the allocation of those resources.

One could conceive of an autonomous HazMon system that might be capable of triggering a response to a hazard on its own, i.e., issue evacuation alerts to a particular geographical area. The consequences of initiating actions of this sort are significant, and policy makers generally want to be in the loop when it comes to issuing such orders. Since HazMon is always accessible to policy makers, there is no benefit to building in the level of autonomy required to support automated hazard responses.

Whether any particular level of action should be automatic or not can be determined from an operator, or decision maker, asking the question: would I rather initiate this action myself, or would I rather not have to think about it?

Appendix B –Modeling Results for Various Hazards

B.1 Hurricane Results

The modeling results for a hurricane are illustrated in Figures B.1-1 and B.1-2.

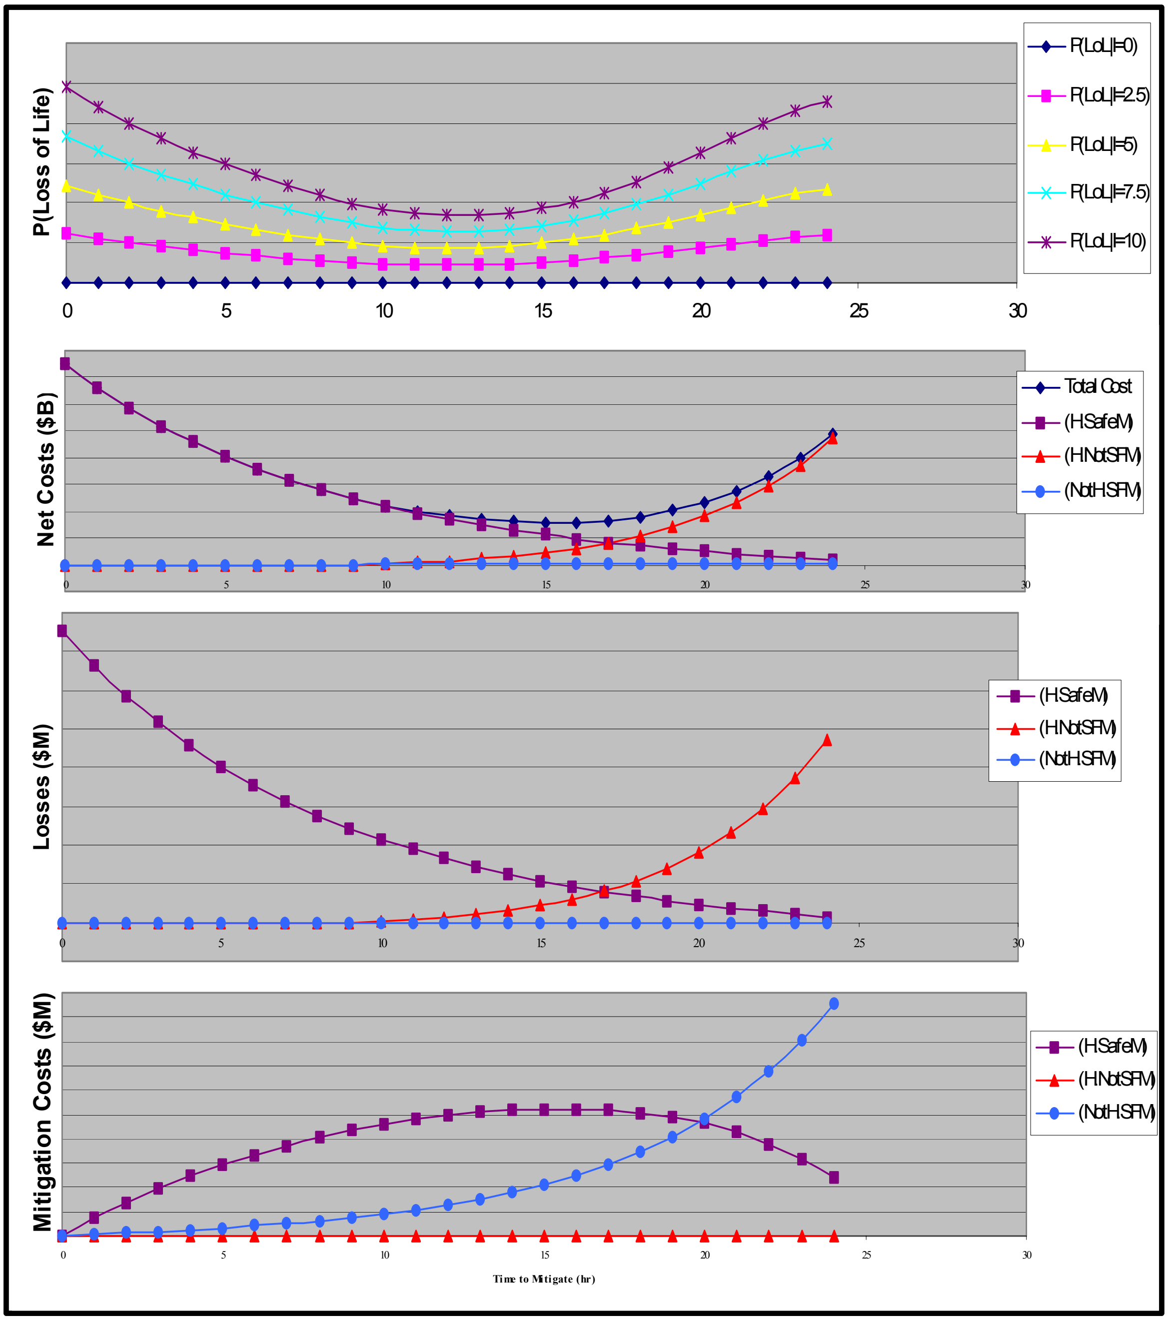

Figure B.1-1 P(LoL) and Costs for a Hurricane

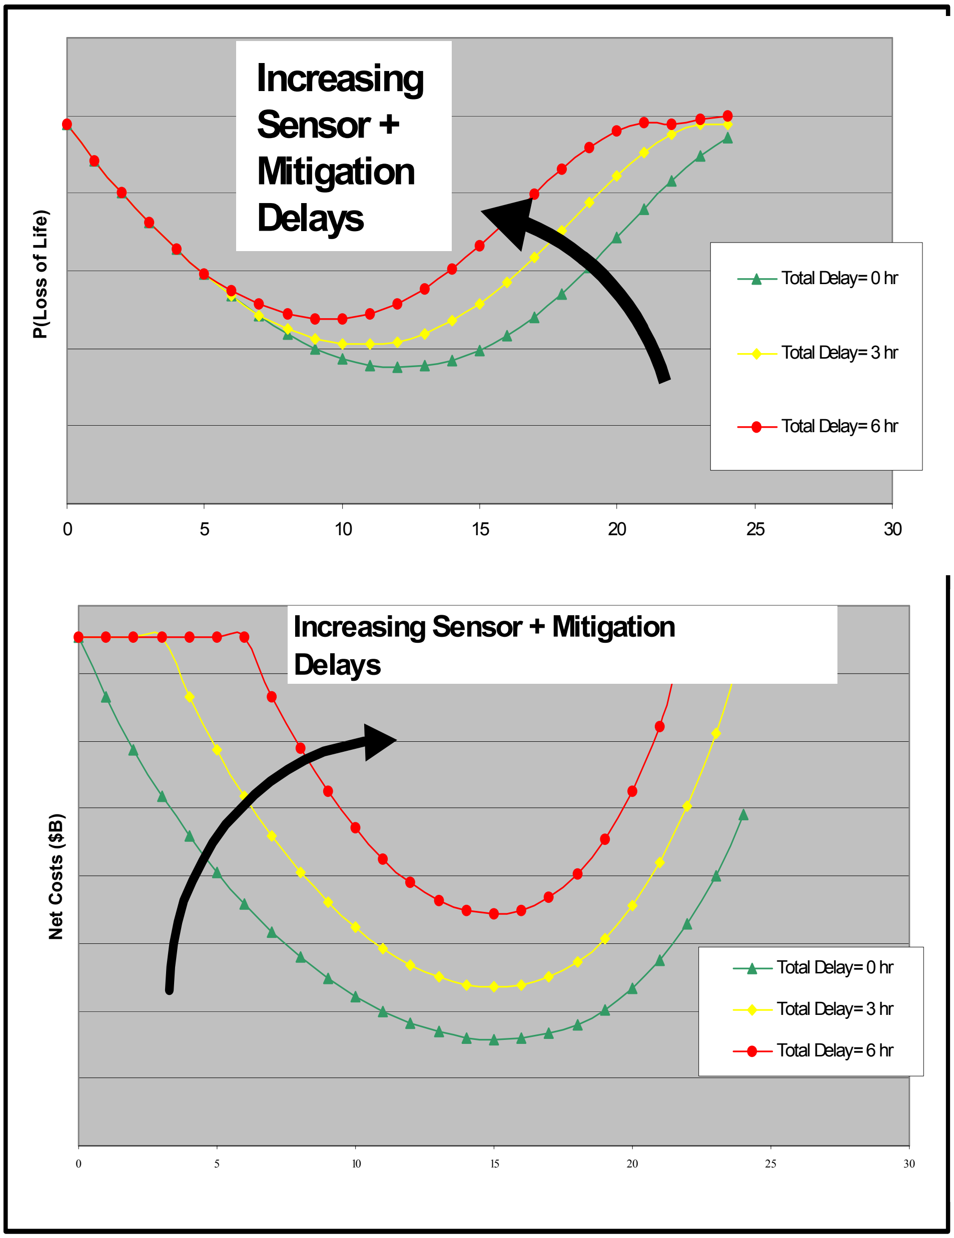

Figure B.1-2 Delay Trends in P(LoL) and Costs for a Hurricane

B.2 Tornado Results

The modeling results for a tornado are illustrated in Figures B.2-1 and B.2-2.

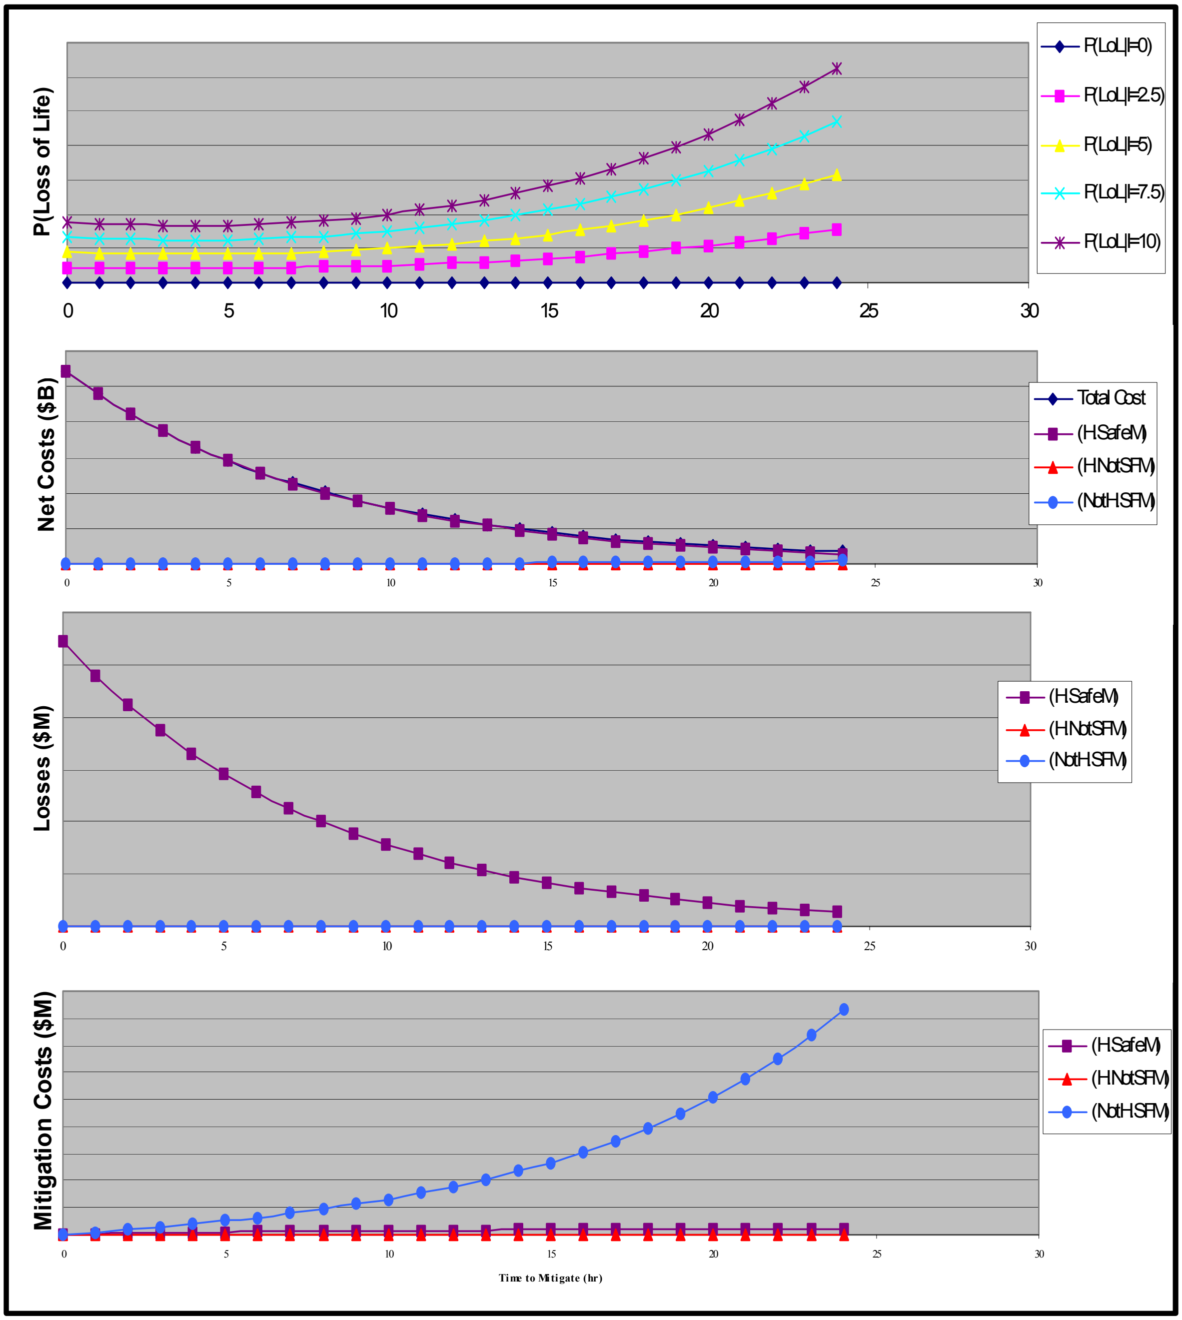

Figure B.2-1 P(LoL) and Costs for a Tornado

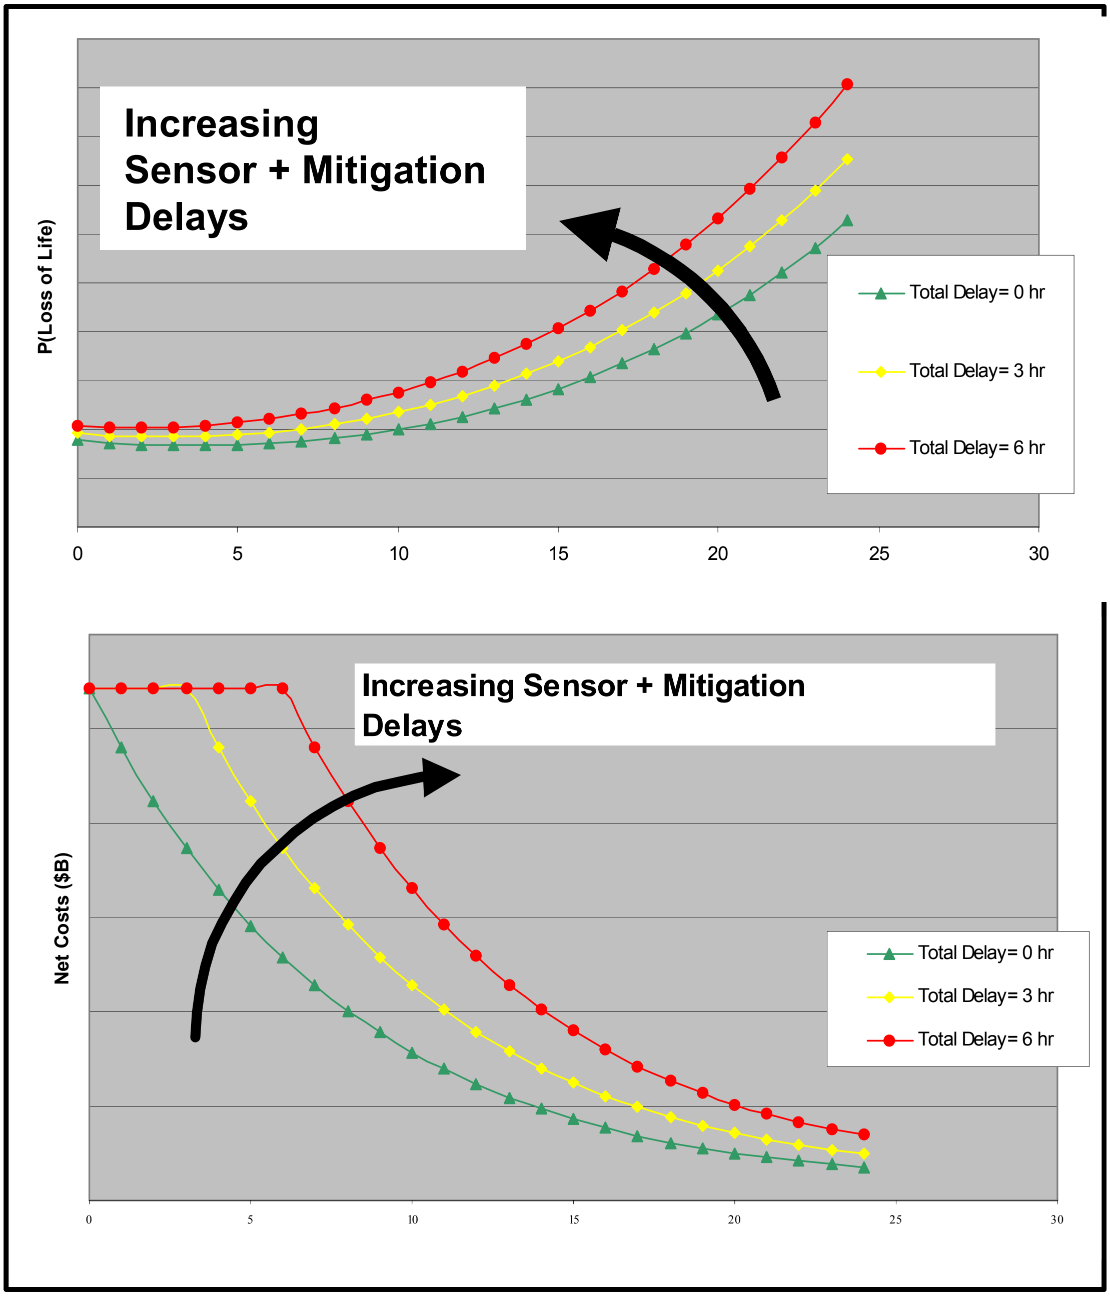

Figure B.2-2 Delay Trends in P(LoL) and Costs for a Tornado

B.3 Volcanic Plume Results

The modeling results for a volcanic plume are illustrated in Figures B.3-1 and B.3-2.

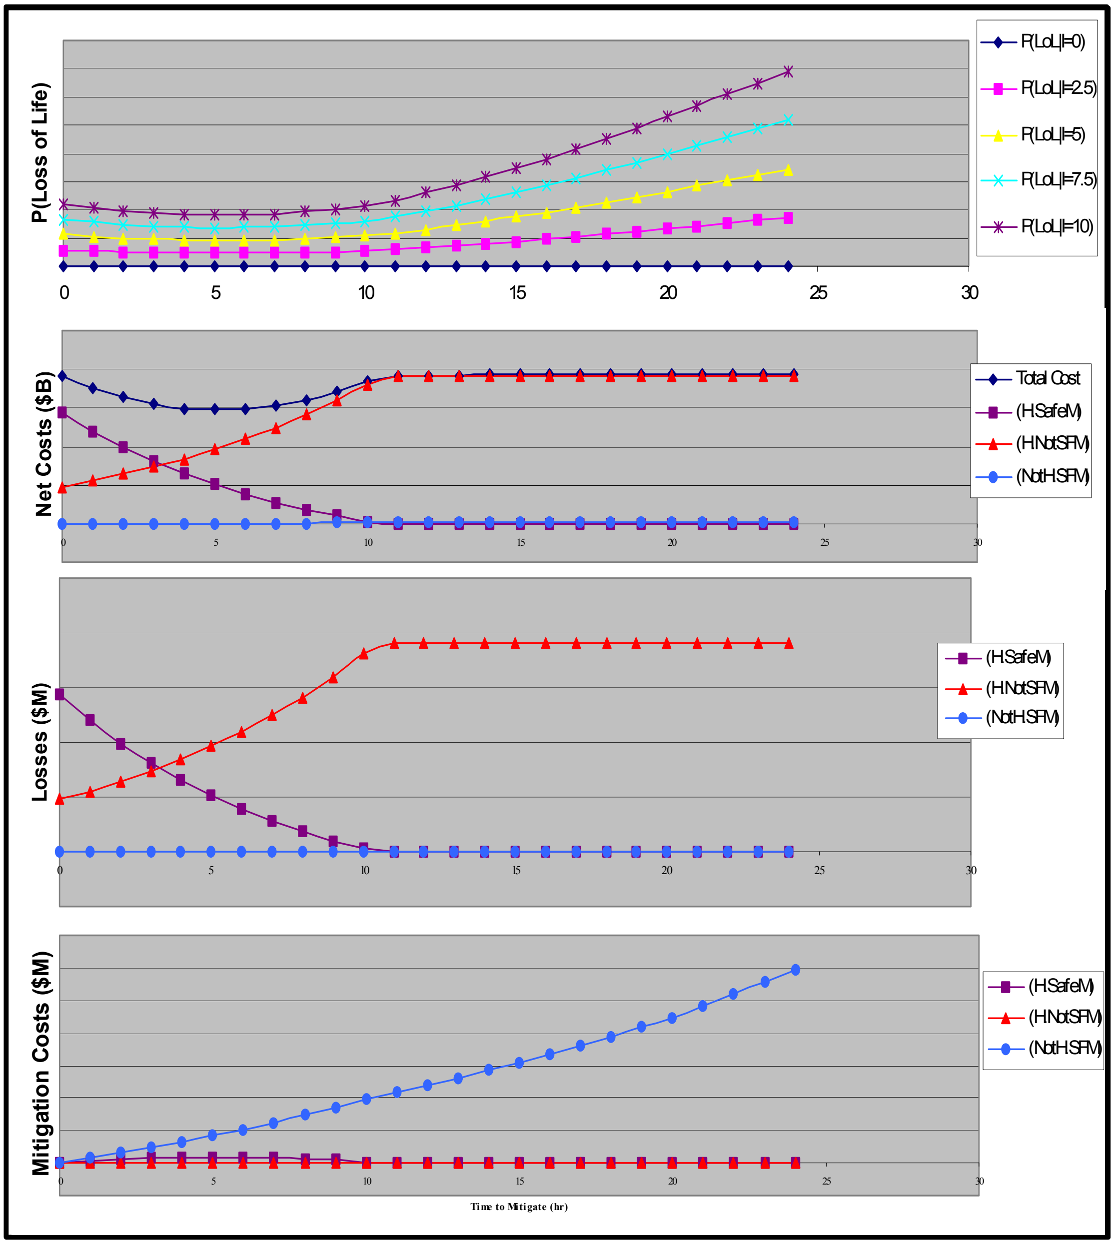

Figure B.3-1 P(LoL) and Costs for a Volcanic Plume

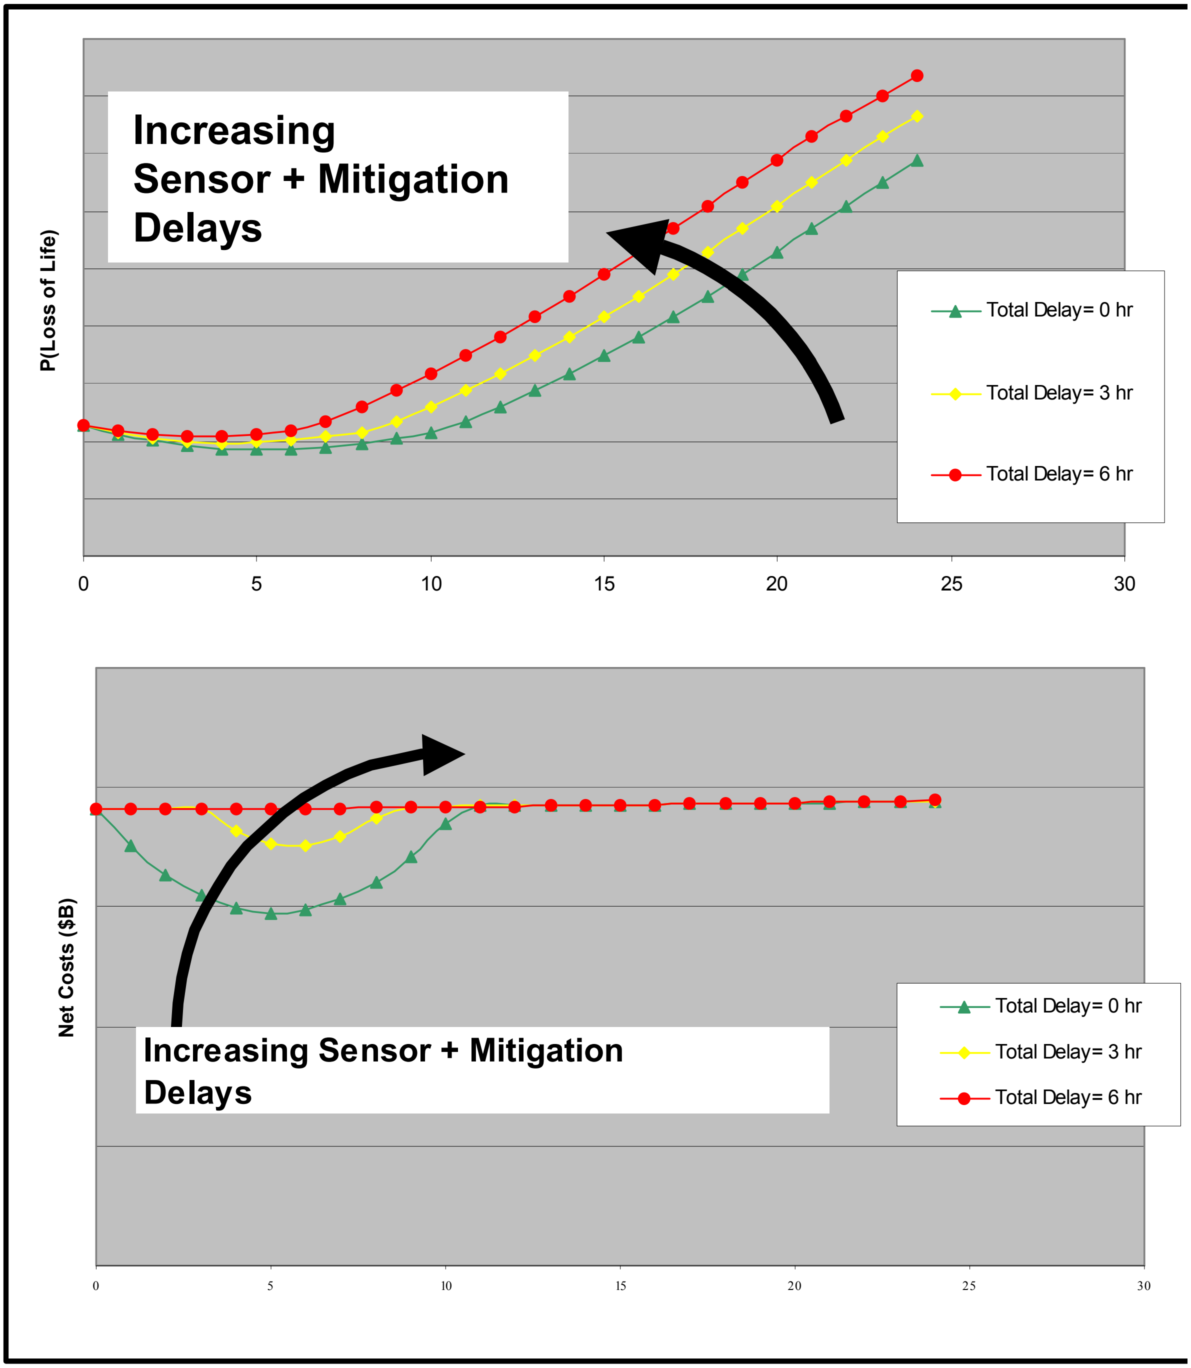

Figure B.3-2 Delay Trends in P(LoL) and Costs for a Volcanic Plume

B.4 Pestilence Results

The modeling results for pestilence are illustrated in Figures B.4-1 and B.4-2.

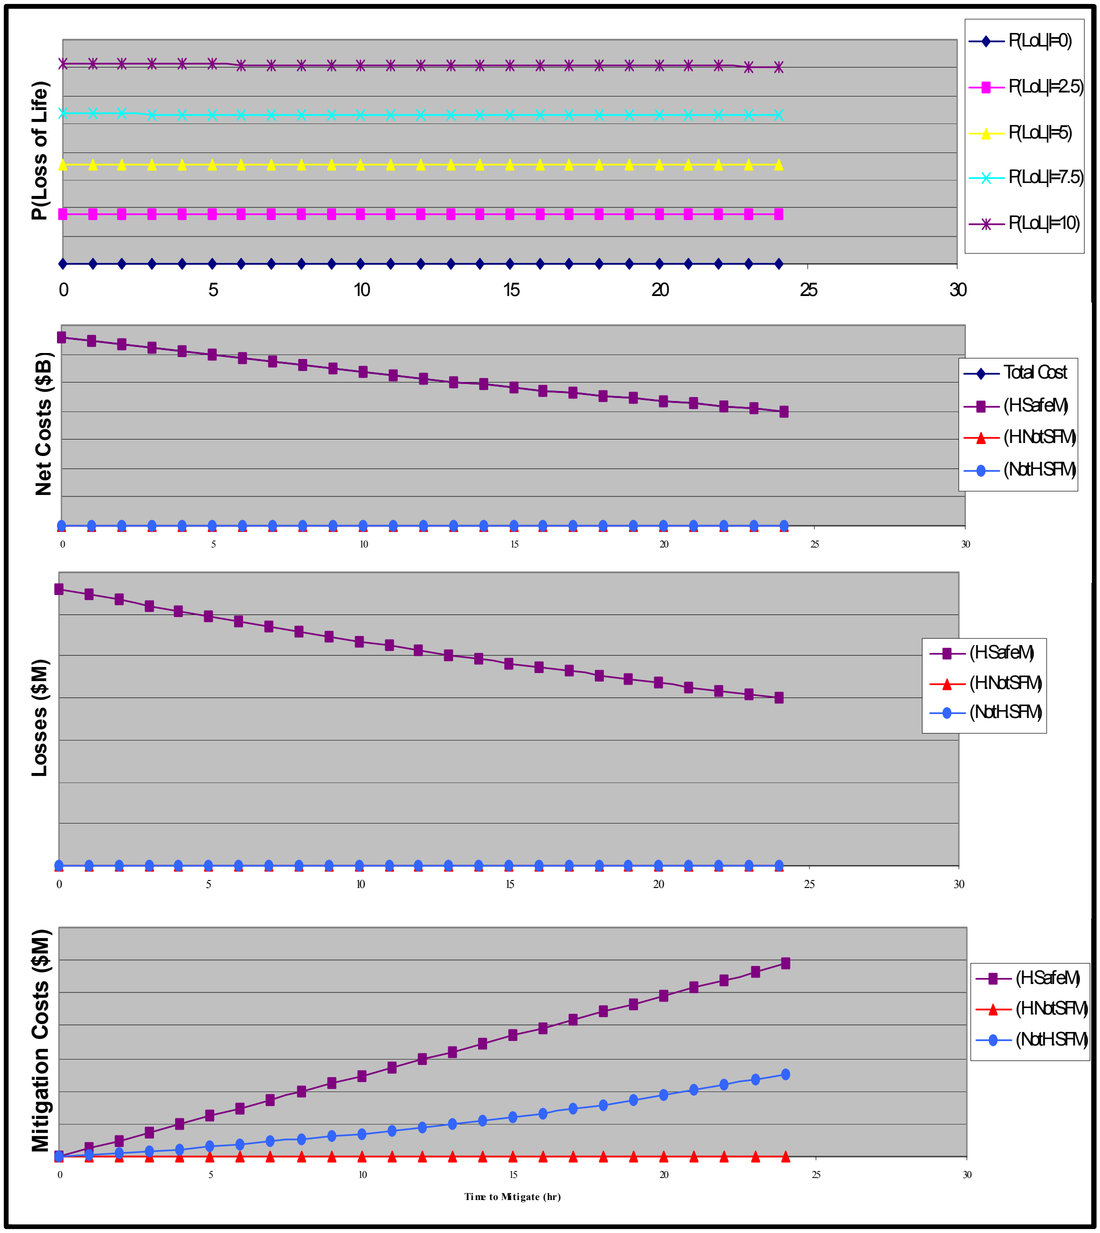

Figure B.4-1 P(LoL) and Costs for Pestilence

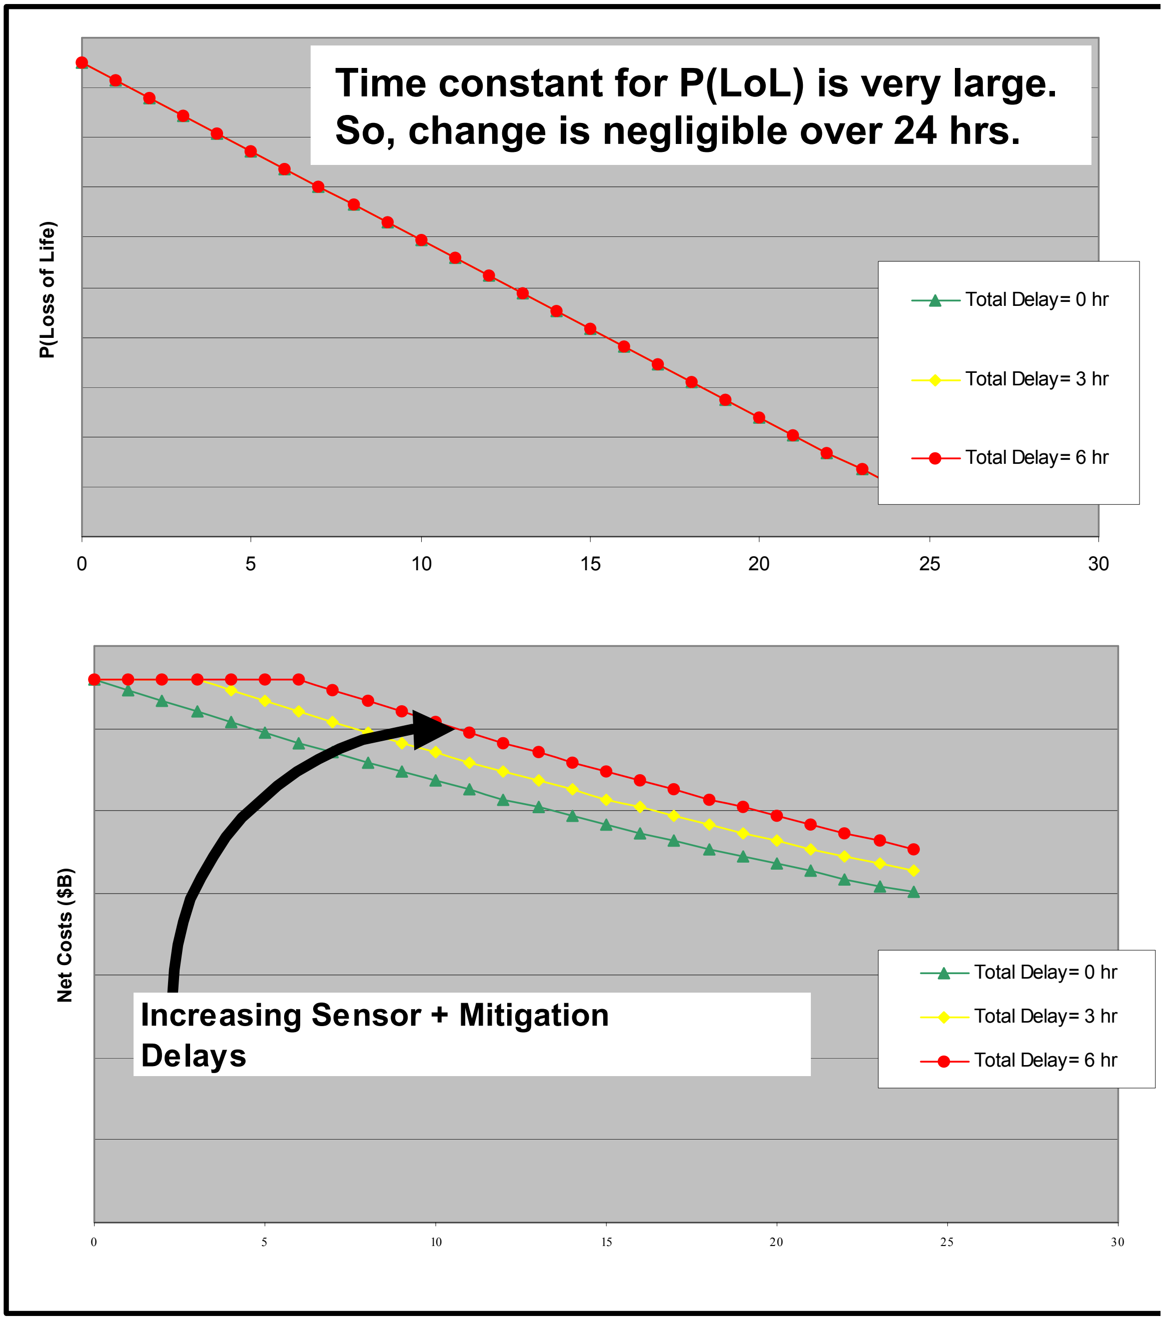

Figure B.4-2 Delay Trends in P(LoL) and Costs for Pestilence

B.5 Oil Spill Results

The modeling results for an oil spill are illustrated in Figures B.5-1 and B.5-2.

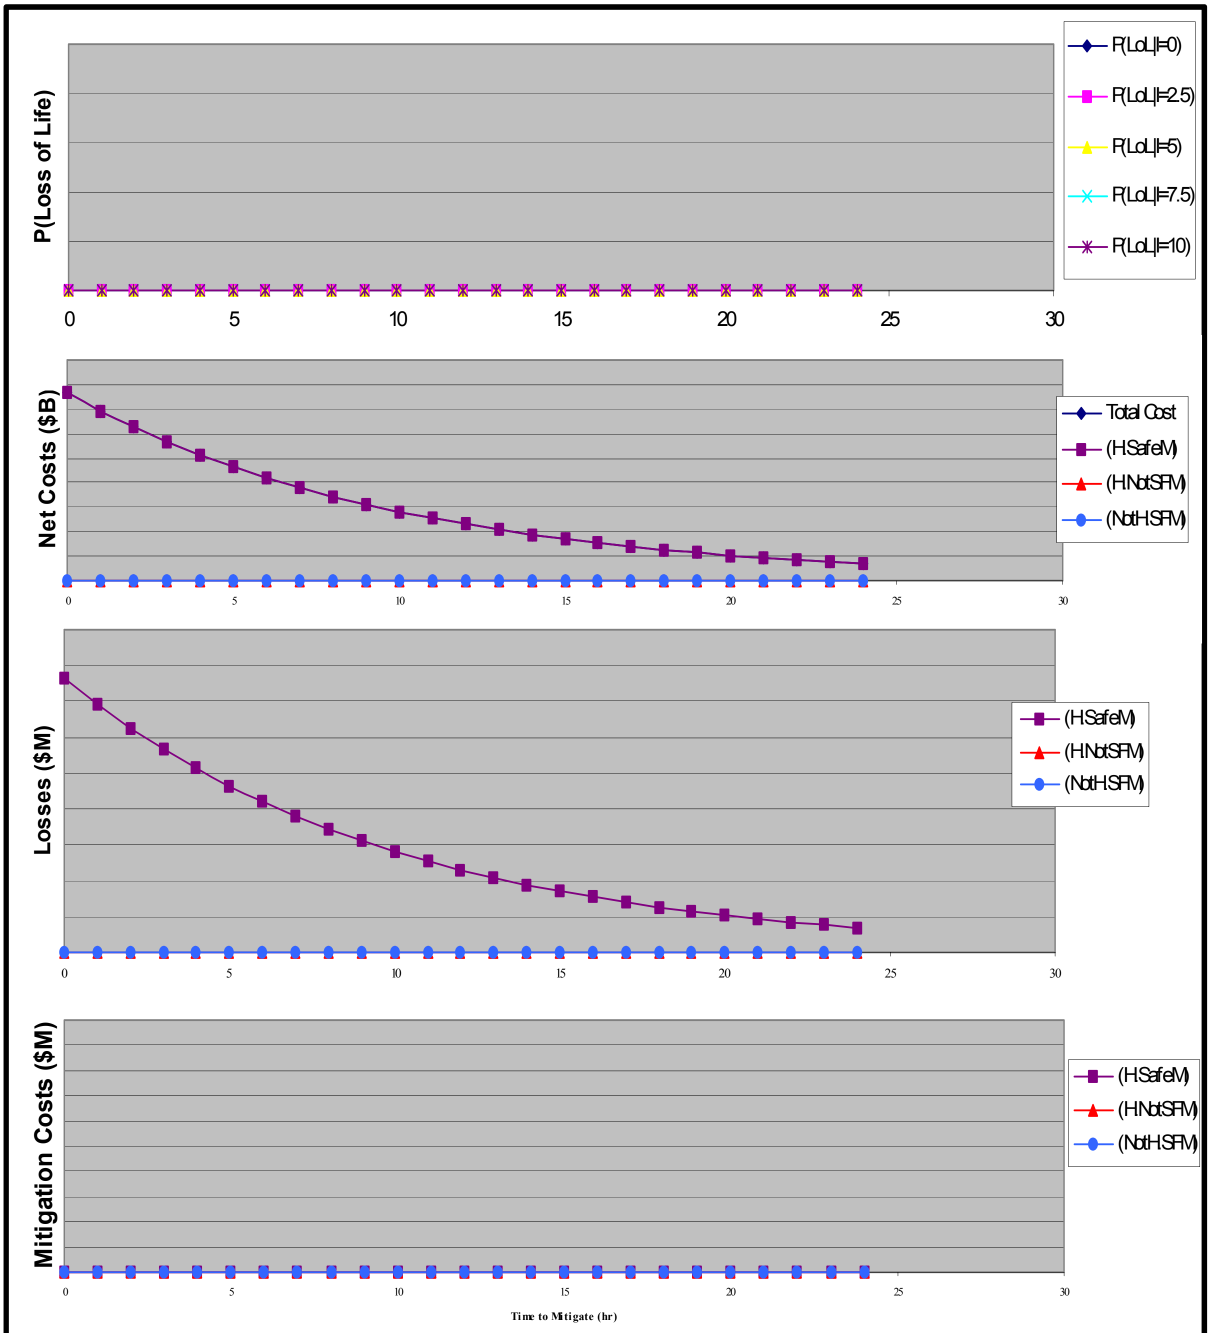

Figure B.5-1 P(LoL) and Costs for an Oil Spill

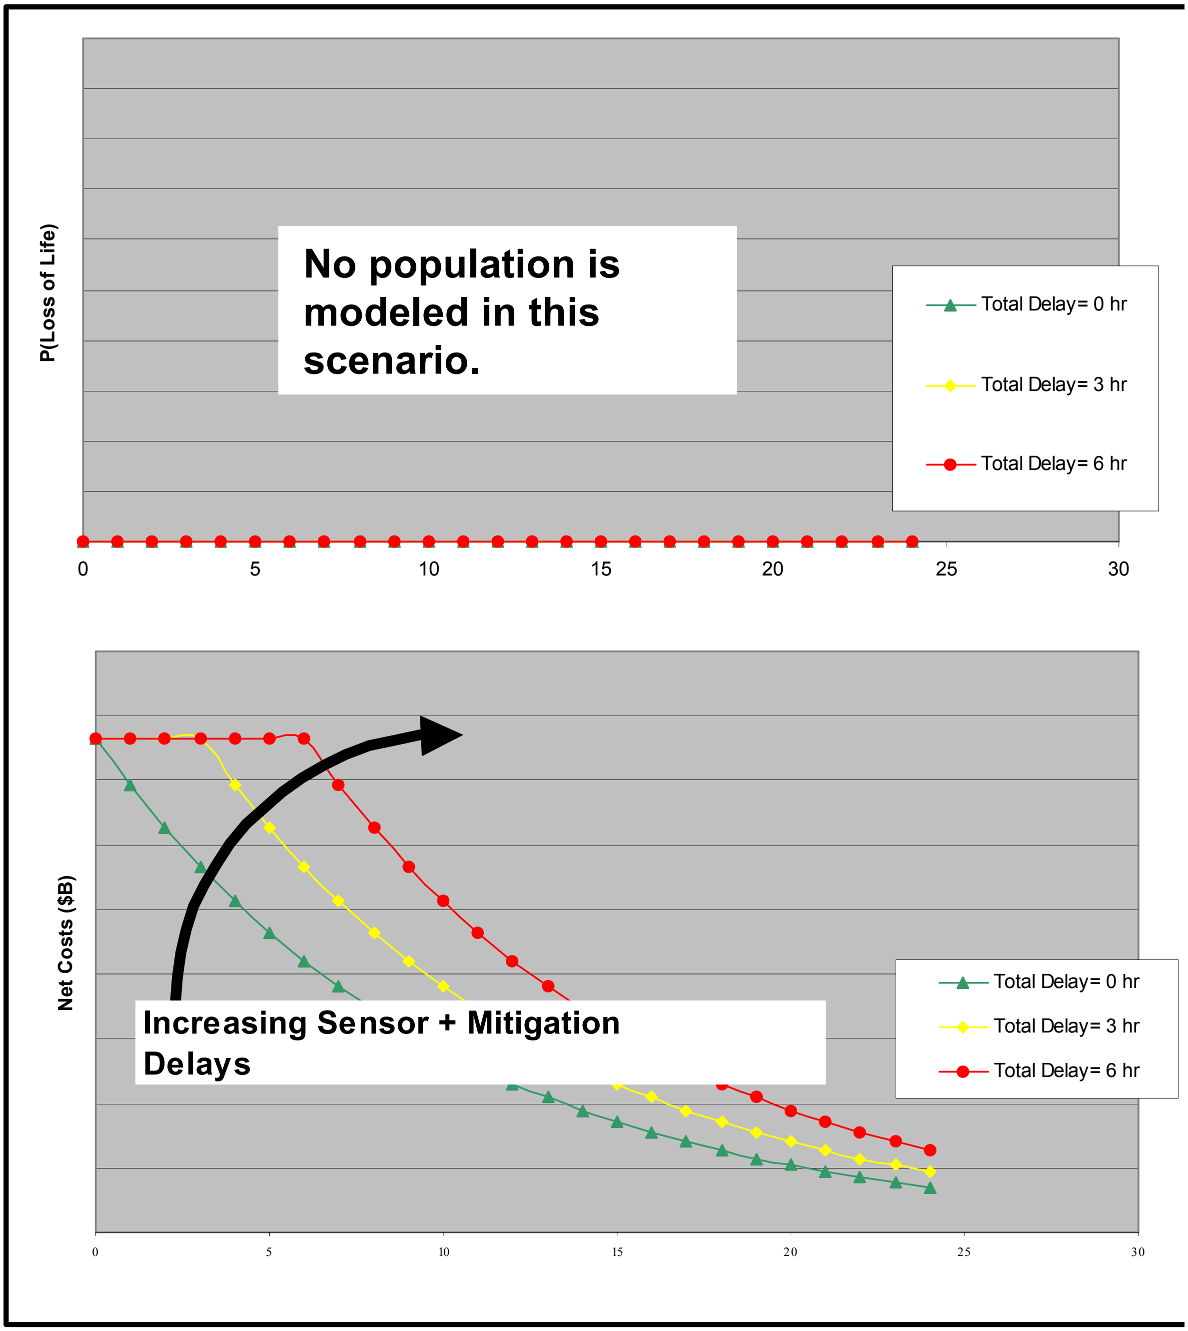

Figure B.5-2 Delay Trends in P(LoL) and Costs for an Oil Spill

B.6 Wildfire Results

The modeling results for a wildfire are illustrated in Figures B.6-1 and B.6-2.

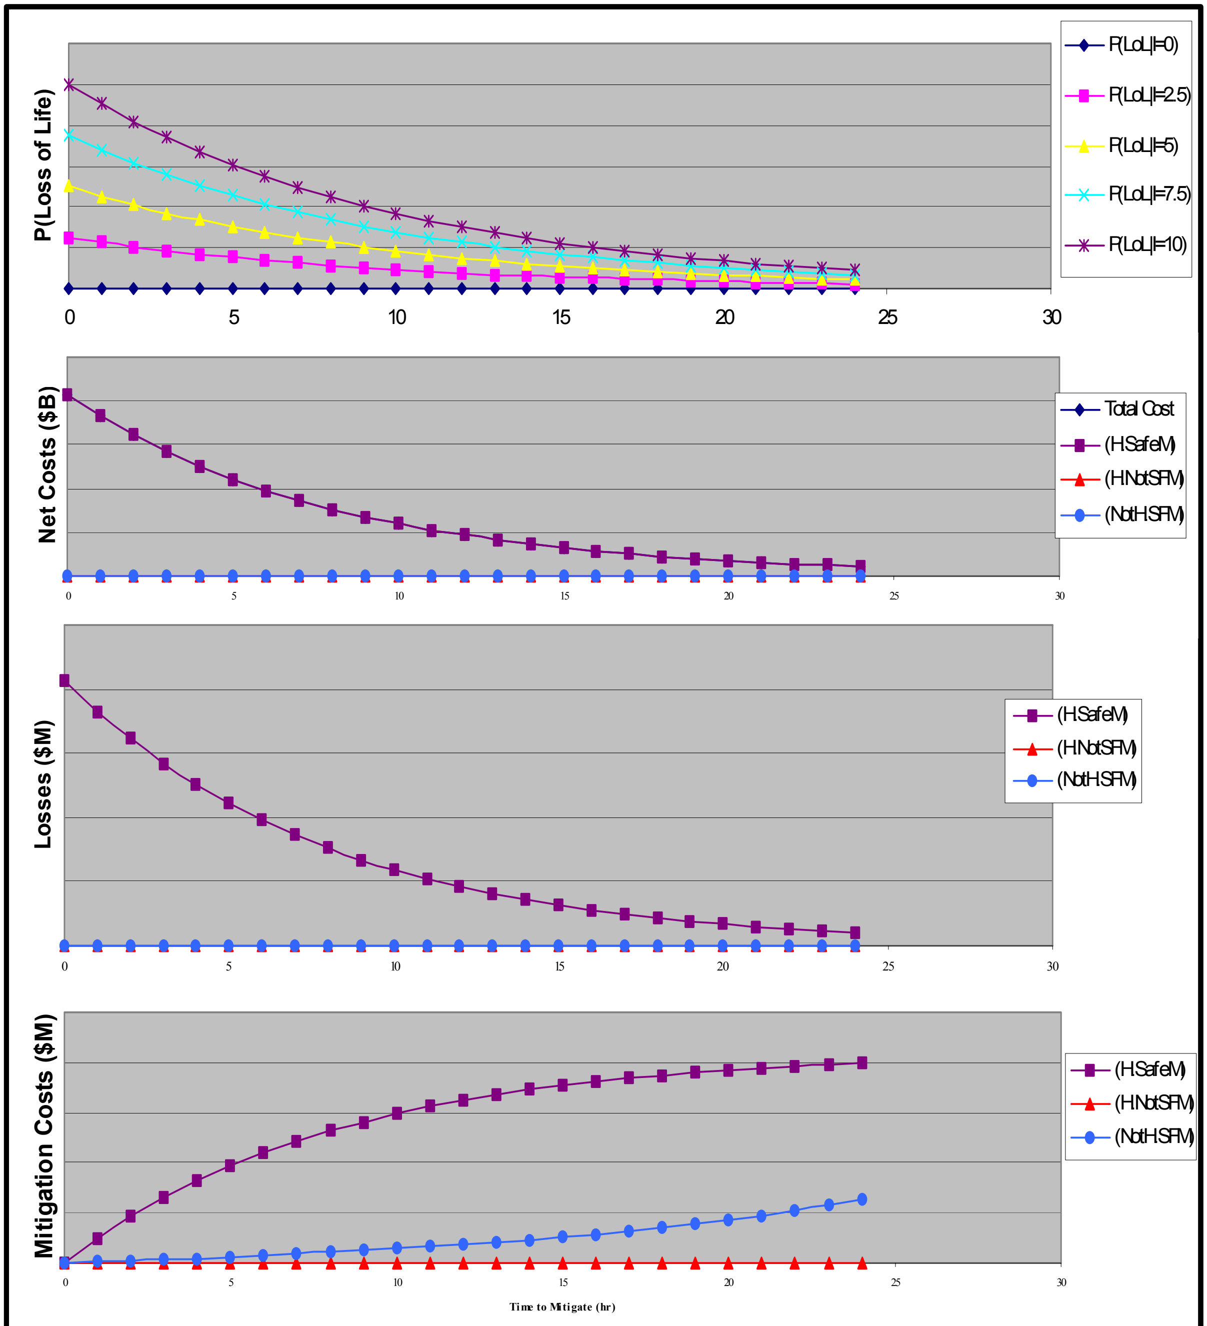

Figure B.6-1 P(LoL) and Costs for a Wildfire

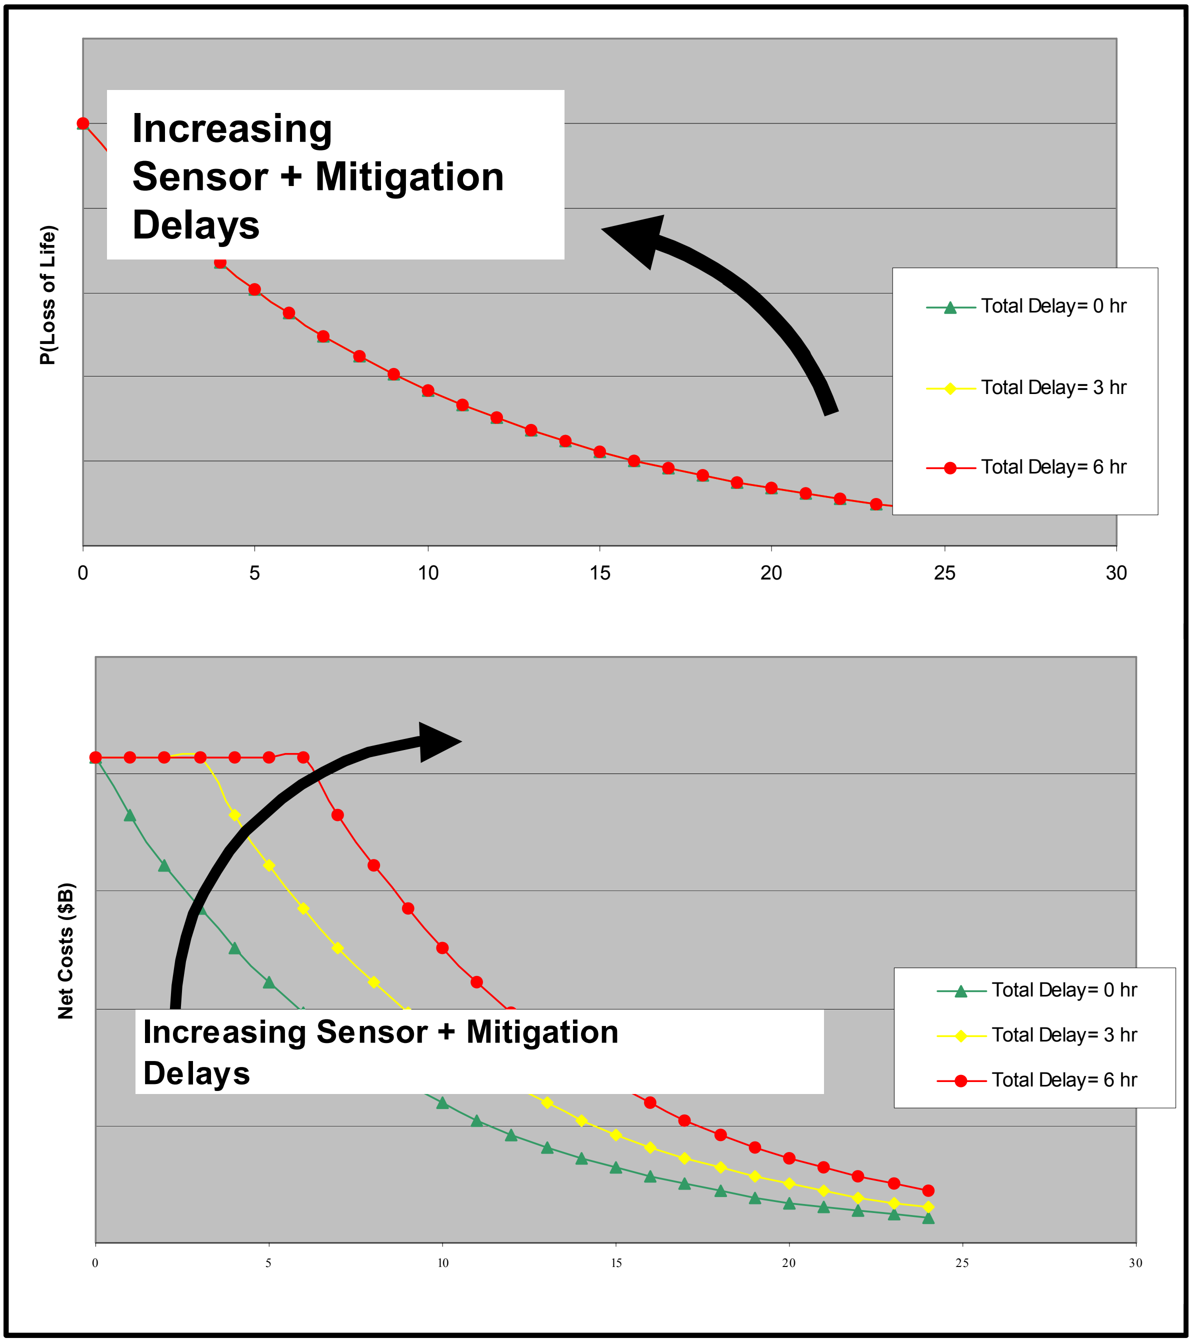

Figure B.6-2 Delay Trends in P(LoL) and Costs for a Wildfire

B.7 NBC Results

The modeling results for an NBC hazard are illustrated in Figures B.7-1 and B.7-2.

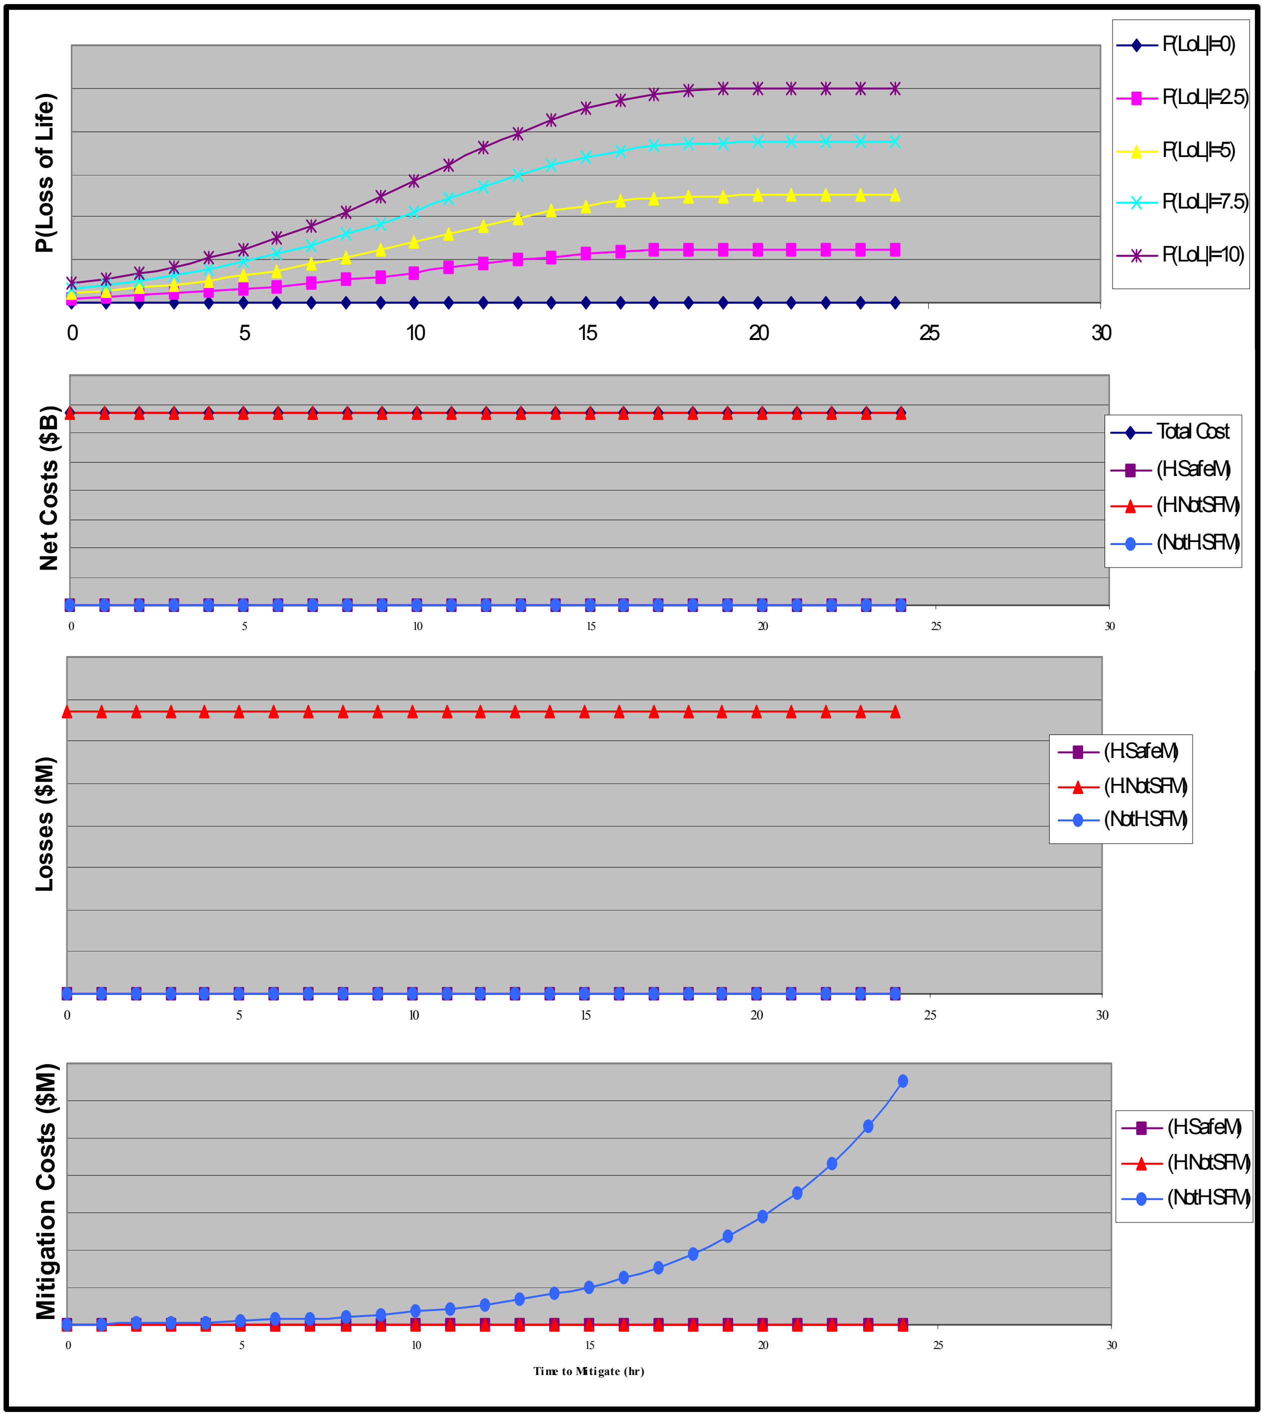

Figure B.7-1 P(LoL) and Costs for an NBC Hazard

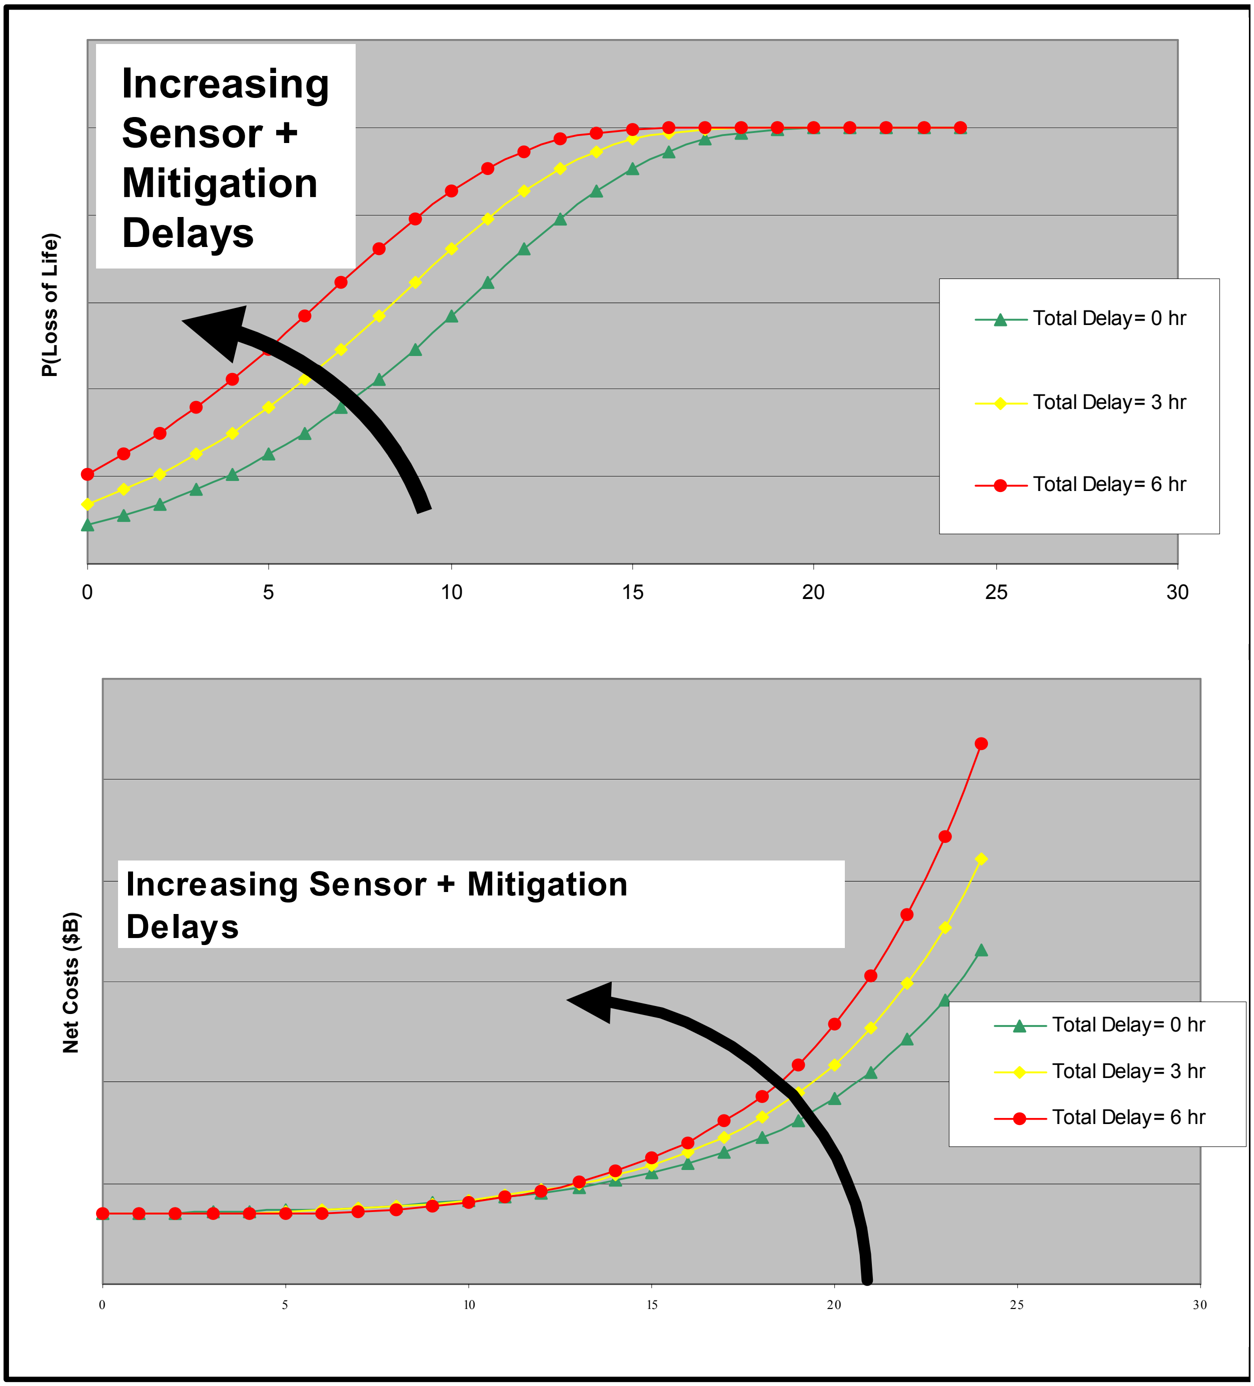

Figure B.7-2 Delay Trends in P(LoL) and Costs for an NBC Hazard

B.8 Earthquake Results

The modeling results for an earthquake are illustrated in Figures B.8-1 and B.8-2.

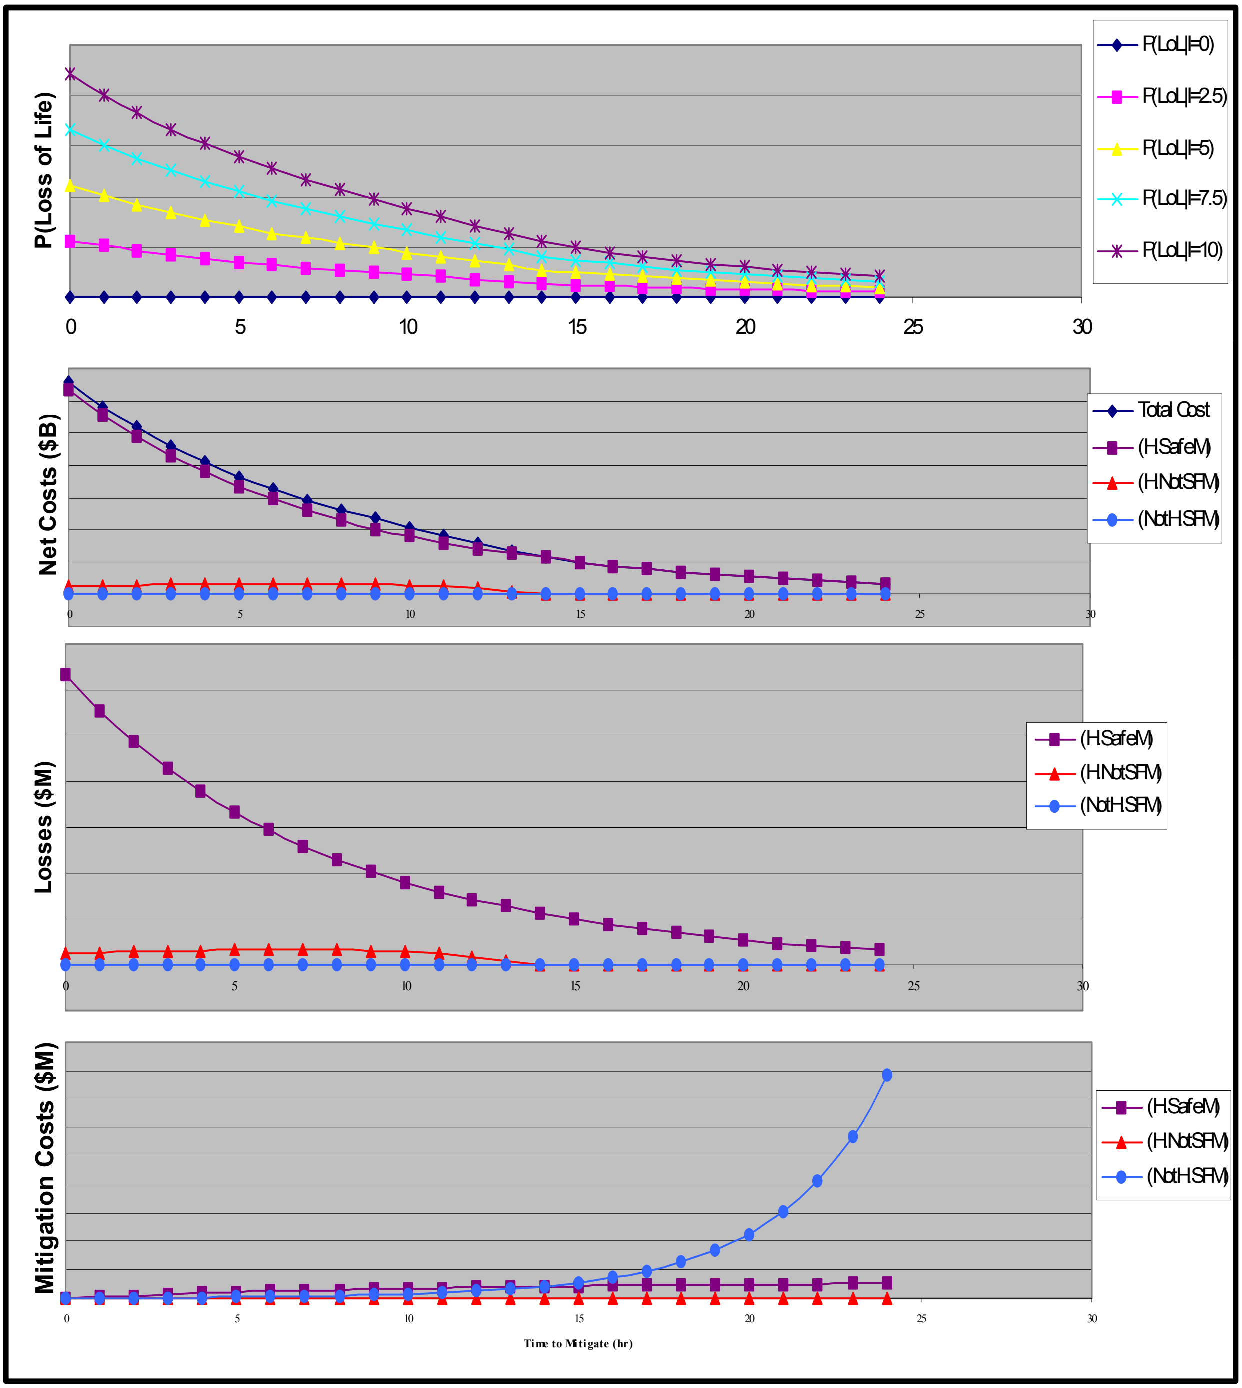

Figure B.8-1 P(LoL) and Costs for an Earthquake

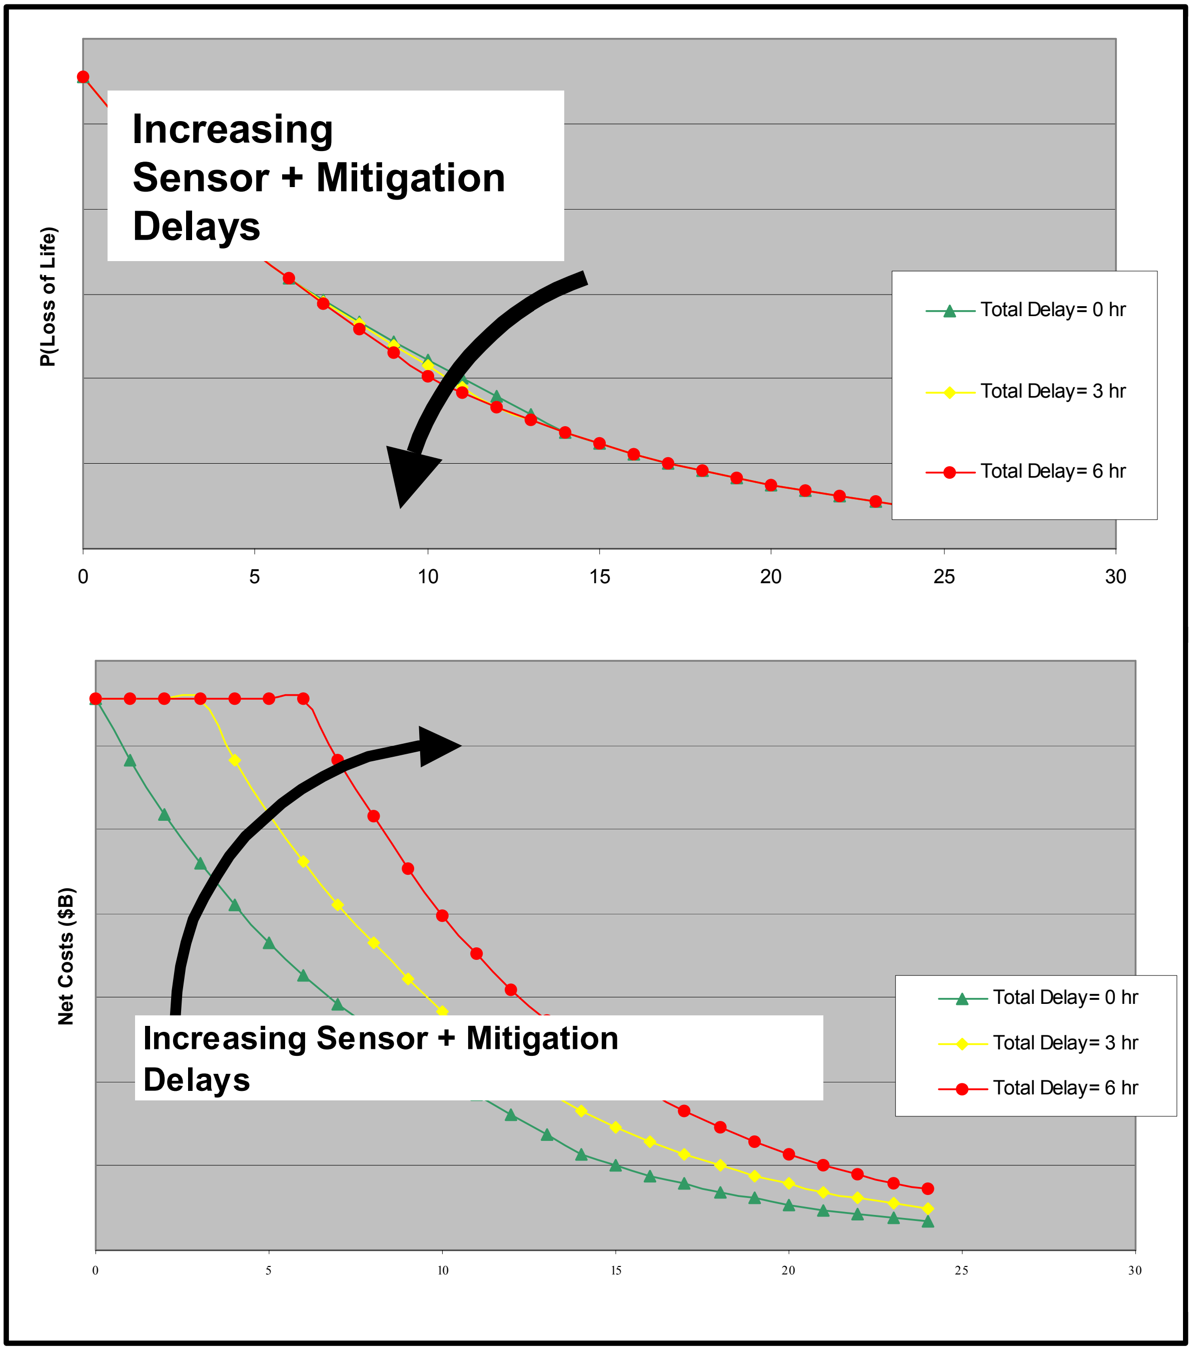

Figure B.8-2 Delay Trends in P(LoL) and Costs for an Earthquake