N89 - 15565

Vulnerability-Attention Analysis for Space-Related Activities

Dan Hays and Sung Yong Lee
Johnson Research Center, University of Alabama in Huntsville
John Wolfsberger, NASA - Marshall Space Flight Center

Abstract. Techniques for representing and analyzing trouble spots in structures and processes are discussed. Identification of vulnerable areas usually depends more on particular and often detailed knowledge than on algorithmic or mathematical procedures. In some cases, machine inference can facilitate the identification. The analysis scheme proposed first establishes the geometry of the process, then marks areas that are conditionally vulnerable. This provides a basis for advice on the kinds of human attention or machine sensing and control that can make the risks tolerable.

Introduction

This paper outlines the main elements of a scheme for analyzing and representing the vulnerability of a structure or process, and for indicating kinds of intelligent attention that could avoid or repair problems at the vulnerable points or regions.

We want to provide tools, mostly computer-based, for depicting trouble spots and for noting their causes and effects. Distinctively, the analysis is explicit about agents, either human or machine-resident, that may be involved in noticing and doing something about potentially harmful situations. The more usual approach is to focus mainly on the devices and processes themselves. However, vulnerability and attention are closely related. Generally, unattended processes are more likely to develop problems, perhaps serious ones. When problems develop in such situations, further problems often ensue. Conversely, attention itself, whether coming from humans or from sensor arrangements, may be open to certain problems such as overload or faulty coordination.

Viewed as an artificial intelligence problem, vulnerability-attention analysis is more a matter of knowledge representation than machine inference, though in handling causal patterns it seems likely that automated techniques could probably save some time. Again, it seems that things go wrong in similar ways for similar structures, so that if structural similarity can be established, then an inference system might suggest looking for certain kinds of problems.

Outline of the Analysis

The steps of the analysis are as follows:

1. Describe the structure or process.
2. Identify vulnerable parts, given specified circumstances.

3. Analyze the causal antecedents of the problems.

4. Trace effects.

5. Based on this analysis, recommend kinds of intelligence and attention that can be applied to avoid or correct the problems.

6. Analyze the allocation of intelligent resources relative to availability, involvement in routine operations, and so on.

There are three domains of information in the analysis:

- the structure or process itself,
- the broader and less well defined realm of factors that could impact the focal structure or process, and
- the sentient resources and their organization.

The structure or process domain—which describes and characterizes the machine, the manufacturing sequence, the managerial procedure—may be the best understood of the three. It is also likely to be somewhat idealized. The second domain, covering various causal factors that can affect the key process or structure, may also be idealized, though conditions of ordinary use are likely to be well enough understood. Objectively, the number of potentially influential external factors is almost always greater than internal sources. Balancing this diversity of possible external causes of problems is the fact that many systems almost always function within a small range of environments. Information about the third domain, that of the sentient resources, is likely to be more variable. If they are humans and only interact occasionally, or if these humans are the designers and testers of the system, they will often be taken for granted, or at least not subject to scrutiny and analysis as part of the system. Generally it seems to be the case that humans who might be involved in detecting and guarding against vulnerabilities will not be considered so systematically as the machines, unless they are involved in operational steps of machine-based processes, in which case they are likely to be treated as components of a mechanized system and paced as such. If the sentient resources are themselves mechanical, they will be treated as part of the physical structure. When the attention of humans is non-routine, for example when some sort of managerial supervision is involved, or when attention is needed only occasionally, as when maintenance or repair is required, the human resources are not likely to be as well planned, since these matters are often not so predictable. In some situations, the human resources may be slighted to no ill effect, when some persons are clever enough to juggle many complex processes. An unjustified reason for skimping on human attention to risky processes is because the accident has not happened yet, or simply to cut costs.

We believe that all humans associated with a system will have an epistemic or knowledge-related role, whether or not this is planned by system designers. They may also have an action role in the process itself. Attempting to understand and represent the dynamics of knowledge in relation to machine and procedure is a major goal of our investigation.

About Vulnerability

Though vulnerability may be thought of as being a property of objects (*intrinsic vulnerability*), it has to do both with the structure or process and with its environment. Trans-situational vulnerability of a system and its parts can be characterized, a sort of "others things equal" vulnerability. But it is usually more informative to describe *contingent vulnerability,* where susceptibility to problems changes with outside factors, history, and so on..

LIke much of engineering thought, the ordinary concept of vulnerability focuses on the object (device, process, etc.) rather than on its situations of use. However, it directly implicates situational factors much more clearly than do terms like "risk" or "weakness": a structure is thought of as being vulnerable to something and perhaps as being vulnerable in certain ways.

To say that something or someone is vulnerable means that it may receive effects, ordinarily from external sources. The additional connotation is that the effects may be harmful, or may change the recipient's structure significantly. On examination, practically anything in the universe can be affected by something, resulting either in a change of state or configuration *(ultimate vulnerability)*. . Nevertheless the concept is useful since not everything is equally affected by everything else. Thinking of which parts are relatively vulnerable, or of the overall vulnerability of certain systems, reminds us of some of the associated causal scenarios that could result in changes worth noting, and prepares us to deal with unwanted change.

Students in both engineering and psychology classes were asked to identify the vulnerable aspects of various entities (objects, devices, procedures, people and relationships). They did this easily. In many cases, they linked the undesirable results to a part of the entity, for example, a pump that wears out. In other cases the problematic region was seen more globally, as in a software system that crashed readily but from diverse and unpredictable causes rather than from something more localized such as parameter passing between procedures. Frequently, design problems seem implicated. Students seem to characterize ordinary conditions of use when asked to conceptualize vulnerable aspects of machines. Vulnerable behavior of persons is more often thought of as contingent on unusual circumstances (being away or in a new situation, for example), or as depending on apparently volitional but statistically unlikely actions of the participants. Without going further into discussion of these exercises, the implication is that identifying points of vulnerability seems a natural way to think. Listing the vulnerability of parts is convenient for cataloguing things to watch out for both within a system during operation and in its environment. (Such a list could, of course, obscure causal relations, or lead one to think that the part is somehow responsible for the things that could happen to it involving potent outside sources.)

Representation of Causes and Effects

Even a bare listing of risky areas can be helpful to someone who must deal responsibly with a system. A *surface vulnerability description* can help channel attention and avoid surprises, even when causes of potential problems or remedies for them are not thoroughly understood.

Causal analysis of events or conditions leading to problems will make the depiction more thorough and probably easier to conceptualize.

The depiction of causal paths leading to problems in a structure or process is more of a challenge than just describing the basic system (itself not always an easy task if the system is large or has complex relationships). This is so for several reasons:

- In some cases, more than one causal sequence could lead to the same costly result.

- Causal factors and processes could conceivably be very numerous. Frequently they will reside largely in sources outside the basic system, including ones that are not in the ordinary environment of the system.

- Antecedent circumstances may be described loosely. (Our experience is that level of abstraction problems are likely and pernicious.)

- Causes of problems may not be understood. Thus there may be nothing to represent. (We conjecture that heuristically programmed computer advisory systems might *suggest* problem areas that could be explored.)

One aim of the analysis is to provide graphic representations of causal processes. In doing so, we would like to combine the more abstract tradition of engineering analysis which list causes and effects analytically ("A and B causes event C, which in the case of state D also causes E to happen.") with the kinds of depiction of linked parts or pictures of structures and processes that has been more common in recent computation (for example, in the renderings of semi-animated devices in various AI programs). In vulnerability representations that we have been exploring, causal depictions are linked to parts or regions of a basic schematic or diagram. In some cases, the causal sequences are almost entirely internal to the parts of the focal system, so they can be shown as highlights or certain parameter values of the the basic depiction. In most cases, though, a representation of external events, states, entities, etc. needs to be included, if only as verbal labels. Generally it is the case that alternate causal paths, or simply lists of possible sources and kinds and problems, must be represented for a given combination of system part or region and vulnerability type. Thus, a basic kind of interface with a vulnerability representation would be the familiar one of "selecting" a part or region, or a type of failure (or change), then choosing from a display of alternative possibilities and

paths leading to this kind of problem, from a menu, a pop-up, or similar artifice. Because of variations in level of understanding of causal factors, it may not be possible to draw diagrams in all cases. Verbal descriptions are often informative, though for analysis by the system one would probably want more information on connectedness to be included.

A now-classic format for causal depiction is Ishikawa diagrams (see Juran & Gryna, 1980, p. 111), where alternative causal "hypotheses" about failure of parts of systems are attached to arrows which point to parts of an industrial process.

When it comes to showing the *effects* of a problematic state or event, their representation will often be more closely tied to the representation of the focal system, since many of the effects may spread within its structure. However, there may also be various effects on the environment. Some are immediate but side-effects, remote consequences, and other unwieldy contingencies may come about. Generally, the more connected a system is to the outside, either physically or epistemically, the more effects will be representable.

Problems with causal depictions and with causal analysis in general should be noted. Shoham (1988) is one who has recently criticized causal diagrams as being oversimplifications. We feel that they may be useful even if they are something of a simplification of what goes on. Even so, tendencies to oversimplify, to assume that the conceptual space is the real world, to be optimistic about one's favorite devices, to look for simplified causal villains, and so on, must be kept in mind.

Comparisons with Traditional Analysis

Vulnerability-attention analysis does not pretend to compete with traditional analyses of failure or fault. It might be thought of as a representation scheme for some of the material uncovered or formulated under standard methodologies. But some similarities and differences are worth noting.

We think that there is an advantage in representing the actual structure or process in some degree of detail. By contrast, note that probably the most highly developed fault analysis methodology, fault tree analysis (Barlow et al., 1975), depicts Boolean combinations of causal factors and events, rooted in descriptions of major conjectured failure states. Working in a causal space, with a manageable algebra, allows risk coefficients to be computed in a fairly straightforward way, and parts of the causal tree to be scrutinized. It is a strong methodology. The kinds of descriptions urged here, which summarize the geometry of the focal system, together with causal depictions associated with points or regions of the system diagram, are not so neat when it comes to managing them mathematically. However, they have the advantage of calling to mind physical relations of adjacency, which may themselves be causally important. For example, when problems occur owing to accidental connection of parts that are not supposed to be connected (e.g., a solder ball or suffusing gas) it is probably easier to think of these

with a process/structure diagram or drawing than with a set of descriptors in an abstract space.

The approach described here is closer in spirit to ordinary failure modes and effects analysis, but would probably lean more toward exploring the connectedness of the entities and events involved more than is sometimes done. (Failure modes analysis is sometimes represented formally; see for example Nielsen, 1975 or Taylor, 1975.)

We would like to be able to reduce some of the labor in identifying cause-effect factors and paths, which is common to all these methods.. One of the questions of vulnerability-attention analysis is whether heuristic analysis, where knowledgeable computer programs interact with subject-matter experts, might reduce some of this labor. It seems that a certain amount of system description, as well as identification of problem-causing paths, depends on particular knowledge of humans. As time goes on it may be possible to incorporate some of this knowledge into computer-based analysis systems or "suggestion systems", in order to reduce some of the tedium of description.

Vulnerability analysis is closer in basic form, though not necessarily in detail, to Ishikawa analysis.

Knowledge Operations

"Attention" is used here roughly as a synonym for "applied intelligence". It serves to point out that knowledge about a system is not useful unless it enters into some real process of noticing, judging, inferring, deciding. adjusting, revising, etc.

Intelligence-in-the-situation requires someone or some knowledgeable machine arrangement.

A variety of persons might be involved with a structure or process at different times. One can distinguish between *pre-attention, on-going attention, and post-attention* relative to an operational phase. Pre-attention consists of efforts to find out possible symptoms of vulnerability, to understand them, and to provide remedies or redesign. On-going attention involves efforts to find symptoms of vulnerability during operations. These might be tipped off by anomalous events, or more directly cued to reliable indicators of specific problems. Knowledge from earlier testing may be useful in this regard. On-going attention may include adjustive or corrective moves. Post-attention evaluates performance, or possibly breakdown, after the fact. Attention at any of these stages could benefit from the knowledge gained at another stage.

These knowledge operations may be *distributed or stratified over agents*. For example, someone who works closely with a machine system will notice small cues that could signal problems. Someone who evaluates statistics of the performance of

many such devices may detect more subtle trends. Managers frequently have massive filtering of information, sometimes constrained by regulation or custom, as well as made difficult by communications overload and slippage, that make the evaluation of what is going on quite difficult.

The approach of vulnerability-attention analysis is knowledge-based rather than algorithmic. However, it may use economical and orderly means of identifying causes and tracing effects. Included in the logic of the analysis is to give each part of the focal system a generic identification, so that more general heuristics can be applied to suggest trouble-spots ("This assembly functions as a valve; valves frequently have certain problems; therefore look for....") In other cases, expert and historical case knowledge would be incorporated, since vulnerability of a system and its parts is actually dependent on conditions, and in some cases on particular causal histories.

Risk analysis is often intuitive, certainly knowledge-based, and can be tedious and difficult to represent. Analysis bearing on the best kinds of attention and intelligence is generally less well understood, possibly because people do a good job of it, by and large. Sometimes they do not, however, so more concern with knowledge operations seems to be of utmost importance, especially with costly and high-risk systems, such as may be found in various parts of the Space program.

Background: the Context-Sensitive Scheduling Problem

The analysis discussed in this paper grew out of work which recast a heuristic scheduler for space activities designed by Floyd and Ford (1986) into an object-oriented form (Hays and Davis, 1988). Davis (1988) reprogrammed the Floyd-Ford scheduler, which in its original form used traditional symbolic programming techniques. Though Davis's version maintained an overall flow of control similar to the original, the treatment of discrete processes as "objects" which pass messages to other objects (in this case, scheduling procedures) that evaluate their suitability for location in a schedule, suggested a partially "decentralized" determination of position which was sensitive to power drain, priority, and other factors.

Yet more radically object-oriented approaches to scheduling could be even more suited to nonhierarchically organized environmental contexts. Some higher-level evaluation or conflict resolution is also needed, of course, to prevent local shortsightedness and to insure that a suitable variety of factors are accounted for. This general kind of decentralized "power" situation for computational entities was discussed in Hays (1977).

Other Applications to Space-Related Activities

Since its beginning, work related to space travel and operations has had to consider risky situations. Precise results have had to be obtained in unusual and often dangerous environments. In many cases, the impossibility of operational attention has meant very careful pre-operational attention leading both to rugged,

protective designs and to detailed attention during construction, testing, flight preparation, and so on (see for example the discussions in Bolger, 1975).

There are many occasions for representing vulnerability and for understanding the optimal application of knowledgeable attention in space-related activities. Scheduling that is more sensitive to context is just one. Testing of devices and procedures, ordinary management for development, design of operational environments with shared machine and human monitoring and decision-evaluation, and various other computer-assisted operations, are all candidates for vulnerability-attention analysis. Deeper understandings of knowledge operations, and the relation of knowledge to external process, should produce efficiencies of analysis and of performance.

Acknowledgements

The work upon which this paper is based has been supported in part by Grant #NAG8-641 from Marshall Space Flight Center, Donnie Ford, Principal Investigator, John Wolfsberger, Project Monitor. Please address correspondence to Dan Hays, 135 Morton Hall, University of. Alabama in Huntsville, Huntsville, AL 35899.

References

Barlow, Richard E., Fussell, Jerry B. & Singpurwalla, Nozer D. (Eds.). *Reliability and Fault Tree Analysis.* Philadelphia: Society for Industrial and Applied Mathematics, 1975.

Bolger, Philip H. (Ed.). *Space Rescue and Safety 1975.* Vol. 41, Science and Technology. American Astronautical Society, 1975.

Floyd, S. & Ford, D. in *Proceedings of the Conference on Artificial Intelligence for Space Applications,* Huntsville, 1986.

Hays, D. "Dominance relations in computing systems." AFIPS Press: *Proceedings of the 1977 National Computer Conference,* 1977, pp. 595-600.

Hays, D. Davis, S., and Wolfsberger, J. "Implementation of a scheduler in LISP and in Ada." *Robotics and Automation Conference,* Huntsville, 1988.

Juran, Joseph M. & Gryna, Frank M. Jr. *Quality Planning and Analysis: from Product Development through Use.* Second Edition. New York: McGraw-Hill, 1980.

Nielsen, Dan. "Use of Cause-Consequence Charts in Practical Systems Analysis." In Barlow, et al., 1975, pp. 849-880.

Shoham, Yoav. *Reasoning about Change.* Cambridge, MA: MIT Press, 1988

Taylor, J. R. "Sequential Effects in Failure Mode Analysis." In Barlow, et al., 1975, pp. 881-894.