Risk-Based Prioritization of Research for Aviation Security Using Logic-Evolved Decision Analysis

S. W. Eisenhawer; Los Alamos National Laboratory; Los Alamos, New Mexico, USA

T. F. Bott; Los Alamos National Laboratory; Los Alamos, New Mexico, USA

M. R. Sorokach; NASA-Langley Research Center; Hampton, Virginia, USA

F. P. Jones; NASA-Langley Research Center; Hampton, Virginia, USA

J. R. Foggia; Consultant: Chicago, Illinois, USA

Keywords: aviation security, risk analysis, prioritization, decision analysis

Abstract

The National Aeronautics and Space Administration is developing advanced technologies to reduce terrorist risk for the air transportation system. Decision support tools are needed to help allocate assets to the most promising research. An approach to rank ordering technologies (using logic-evolved decision analysis), with risk reduction as the metric, is presented. The development of a spanning set of scenarios using a logic-gate tree is described. Baseline risk for these scenarios is evaluated with an approximate reasoning model. Illustrative risk and risk-reduction results are presented.

Introduction

An important element of the national infrastructure is the aviation system. Major efforts are being directed toward improving procedures and providing technical measures to reduce the risk of terrorist attacks on the system. Of necessity, many of these efforts are short term and are intended to address the current risk in the aftermath of the events of September 11, 2001, using existing technology. Improving aviation security also will require the introduction of new technology. As part of its Aviation Safety and Security Program (ASSP), the National Aeronautics and Space Administration (NASA) is developing technologies for application in this 10- to 20-year timeframe. These technologies are directed to all three phases of the aviation system (aircraft, airport, and national airspace) and may have the potential to improve security and system capacity.

By their very nature, many of the NASA technologies under consideration are quite immature. Inevitably, some of them will prove to be more suitable than others for a variety of reasons, and there is a strong incentive to direct the limited available funding toward the most promising concepts. How can the prioritization of this research be done? The first step in answering this question is to decide on a metric to be used to rank order the technologies. Although many metrics are available, such as cost and technical risk, the "risk reduction" associated with terrorist attacks was chosen as the initial metric for the technology rank ordering. Several steps are required to perform the rank ordering:

- determine a spanning set of terrorist-attack scenarios,
- develop an inferential model to infer risk for a scenario,
- evaluate each scenario for baseline risk and the change caused by an introduction of candidate technologies, and
- rank order the technologies using appropriate measures of risk reduction.

Taken together, these steps represent the essential phases of the decision process that interests NASA in the development of a research-portfolio-prioritization scheme for its aviation security initiative.

The methodology chosen for this problem was logic-evolved decision (LED) analysis (ref. 1), which was developed at Los Alamos National Laboratory. LED analysis uses linked, formal logic models to represent the basic functions of a decision analysis tool. Each logic model is a directed graph called a logic-gate tree. The process trees are developed deductively using general-purpose tree-construction software called LED TOOLS (ref. 2). Two process trees are needed for the aviation security prioritization analysis—a possibility tree that provides the basis for the spanning set of attack scenarios and an inference tree that defines how the risk metric is to be inferred for the baseline case and with the assumed introduction of new security technologies.

LED models are particularly well suited to decision problems that require the integration of expert knowledge that resides with a large and diverse set of subject-matter experts. The use of approximate reasoning techniques within LED allows accurate reflection of the inferential processes that are used to make judgments when many of the factors are highly subjective, as is the case here. Additionally, the results (including uncertainty) are expressed in an easily understandable form using this methodology. In this paper, we show how LED analysis is being applied to prioritize research for aviation security. We present the logic models that were used to represent the attack scenarios and the inferential process for risk, and we describe the interactions with experts that were used to develop these models. Also presented are illustrative risk and risk-reduction rankings for candidate technologies associated with the security of large, commercial passenger aircraft.

Possibility Tree for Aircraft-Attack Scenarios

When terrorist attacks against United States aviation are considered, the overall system conveniently can be divided into three target segments: aircraft, airports, and the national airspace. Aircraft, in turn, may be categorized in many ways. The division used here is in accordance with the Federal Aviation Administration (FAA) regulations, which are based primarily on operational intent and specify particular requirements that are associated with aircraft specifications, operations (including security requirements), maintenance, and crew training. In this paper, we consider attacks associated with aircraft that operate under Part 121 of the FAA regulations and that carry passengers and cargo. Typical aircraft include large passenger aircraft (e.g., Boeing 747s, 767s, 757s, and 737s and similar Airbus types), as well as regional jets (e.g., Canadair RJs). All cargo operations (e.g., Federal Express and United Parcel Service) with these, as well as smaller aircraft, also are Part 121 operations. However, the size of the aircraft is not the critical parameter; identical aircraft will operate under Part 91 when they are owned by a business or private individual. Business jets (e.g., Boeing Business Jet, Gulfstream IV, and Citation X), along with typical general aviation aircraft (e.g., Cessna C172), also operate under Part 91.

Figure 1 shows the attack-scenario possibility tree for a Part 121 passenger and cargo aircraft (Part 121 PC). The top node in the tree is a terrorist attack involving "a Part 121 air operation handling passengers and cargo." The gate associated with this node is an EXCLUSIVE OR (EOR). The inputs to this gate, which appear below and indented to the right by convention, are "The attack targets passengers/crew." (indicated as Node PC in figure 1), "The attack targets the aircraft." (Node AC), and "The attack uses the aircraft as an enabling system." (Node AES). Each of these nodes defines the starting point for a distinct set of attack scenarios to be developed in turn using deduction. The gate type for nodes PC and AC is a CONTINUATION (CO), which is used to facilitate the natural-language development of the logic. Node AC has a single input "The attack is", an EOR with the inputs "The attack is on the airframe." (Node AF), and "The attack is on critical on-board systems." (Node OBS). Each of these nodes is shown as an open diamond indicating that further development of these scenarios can be seen by double-clicking in LED TOOLS.

Figure 2 shows the continued development of a single airframe attack involving a man-portable missile that commonly is referred to as a Man PAD (Node AF-2). Man PAD attacks against a passenger aircraft on departure and a successful attack against a cargo aircraft on approach (in Iraq) have been reported. Node AF-2 is defined as a PROCESS (PO) gate, and the inputs describe a sequence of steps to be carried out in mounting this attack. Note that in a process gate, the order of the inputs from top to bottom indicates the process sequence. The first input to this gate describes the type of portable missile used: "*The missile is*". This gate is a TAXONOMY gate and is useful for defining an ordered set of alternatives. In this case, three alternative types of missiles are shown. All three appear as solid circles, indicating that they are terminal nodes. Two of these ("*a Stinger*" and "*another Man PAD*") appear as blue (dark) icons. Blue is a visual indication that a node has been excluded—that is, it will not be considered when the possibility tree is evaluated. The next input in the process is “*The attacker acquires the weapon system.*” (Node AC).

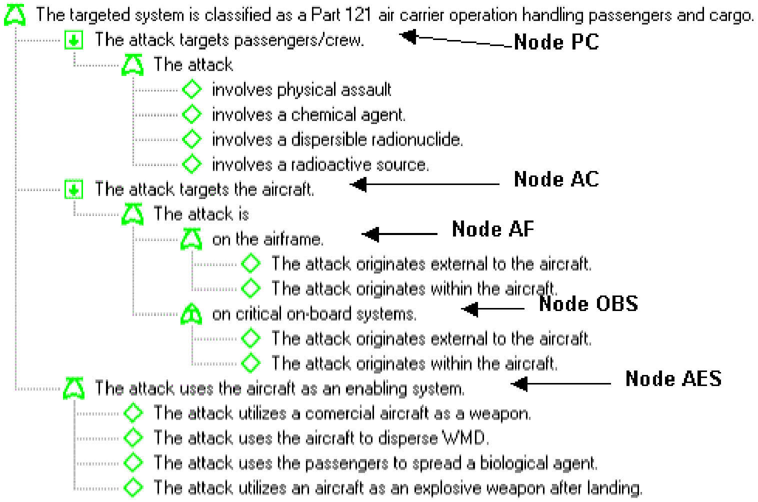

```
The targeted system is classified as a Part 121 air carrier operation handling passengers and cargo.
    The attack targets passengers/crew.                          Node PC
        The attack
            involves physical assault
            involves a chemical agent.
            involves a dispersible radionuclide.
            involves a radioactive source.
    The attack targets the aircraft.                             Node AC
        The attack is
            on the airframe.                                     Node AF
                The attack originates external to the aircraft.
                The attack originates within the aircraft.
            on critical on-board systems.                        Node OBS
                The attack originates external to the aircraft.
                The attack originates within the aircraft.
    The attack uses the aircraft as an enabling system.          Node AES
        The attack utilizes a comercial aircraft as a weapon.
        The attack uses the aircraft to disperse WMD.
        The attack uses the passengers to spread a biological agent.
        The attack utilizes an aircraft as an explosive weapon after landing.
```

Figure 1 – High-Level Development of Part 121 PC Attack Possibility Tree

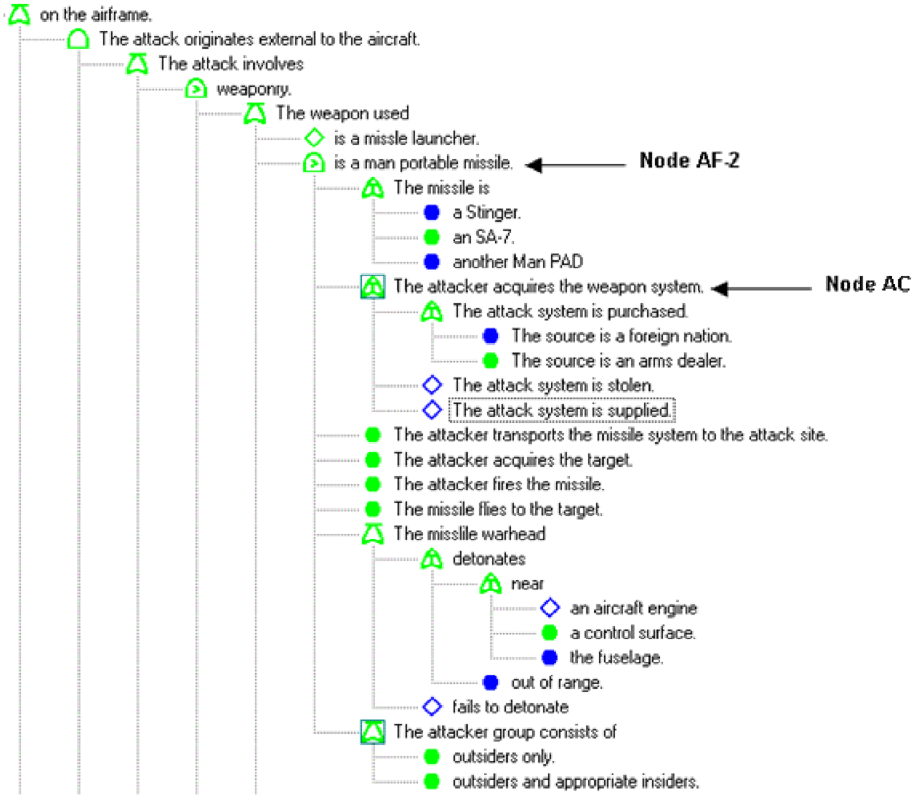

```
on the airframe.
    The attack originates external to the aircraft.
        The attack involves
            weaponry.
                The weapon used
                    is a missle launcher.
                    is a man portable missile.                   Node AF-2
                        The missile is
                            a Stinger.
                            an SA-7.
                            another Man PAD
                        The attacker acquires the weapon system.  Node AC
                            The attack system is purchased.
                                The source is a foreign nation.
                                The source is an arms dealer.
                            The attack system is stolen.
                            The attack system is supplied.
                        The attacker transports the missile system to the attack site.
                        The attacker acquires the target.
                        The attacker fires the missile.
                        The missile flies to the target.
                        The misslile warhead
                            detonates
                                near
                                    an aircraft engine
                                    a control surface.
                                    the fuselage.
                                out of range.
                            fails to detonate
                        The attacker group consists of
                            outsiders only.
                            outsiders and appropriate insiders.
```

Figure 2 – Sub-Tree for Man PAD Attack on an Aircraft

This input is also a taxonomy gate. In this case, a box surrounds the icon. This box indicates that the node is a replicant, a sub-tree that may be used in multiple places in the tree. Each of the successive inputs describes another part of the process. The last input describes the attacker.

The attack scenario possibility tree is a logical equation written in natural-language form. One class of solutions is the set of paths. Each path is a unique attack scenario. For example, the Man PAD scenario is as follows.

AF-2a: *The targeted system is classified as a Part 121 air carrier operation handling passengers and cargo. The attack targets the aircraft. The attack is on the airframe. The attack originates external to the aircraft. The attack involves weaponry. The weapon used is a man portable missile. The missile is an SA-7. The attacker acquires the weapon system. The attack system is purchased. The source is an arms dealer. The attacker transports the missile system to the attack site. The attacker acquires the target. The attacker fires the missile. The missile flies to the target. The missile warhead detonates near a control surface. The attacker group consists of outsiders only. The weapon is located off airport.*

This scenario is unedited and appears exactly as shown from solving the logical equation. The degree of detail in a scenario is controlled by the exclusion operation noted above and a companion termination operator. Note that the description of the attacker is quite simple: "*The attacker group consists of outsiders only.*" A companion scenario (AF-2b) is identical, except that the attack group involves outsiders and appropriate insiders. In the actual analysis, these nodes are expressed in considerably more detail.

The possibility tree provides a mechanism for describing a very large number of scenarios in compact form. If all of the nodes in the possibility tree that are considered here were included in the evaluation, more than 14.2 million attack scenarios would exist; many of these attacks are variations of a single basic theme. For evaluating the value of security technologies, the goal was to obtain a spanning set of scenarios. For the analysis of Part 121 PC aircraft, we used 96 scenarios in the spanning set: 48 basic scenarios multiplied by 2 attacker types. A summary of the spanning set is given in table 1.

Table 1 – Summary of Attack Scenarios in Spanning Set

Type of Attack	**Number of Scenarios**	**Example**
PC: Attack on Passengers and Crew	8	Dispersion of chemical agent in passenger compartment
AF: Attack on Airframe	40	Missile attack with man-portable system
OBS: Attack on Critical Onboard Systems	40	Jamming or spoofing of navigational aids
AES: Use of aircraft as an enabling system for weapons-of-mass-destruction (WMD) attack	8	Variations of 9/11 World-Trade-Center attack

Inference Model for Attack Risk

As noted earlier, to perform a risk-based prioritization of the research technologies, it is necessary to estimate the baseline or current risk. Thus, our starting point is to rank order the spanning set of scenarios. Our inferential model for terrorist risk incorporates a game theoretic perspective (ref. 3). The game to be played is asymmetric. Specific attackers will choose to attempt only a particular subset of attack scenarios associated with particular targets and to employ specific attack modes. They will attempt to allocate their assets to inflict the maximum amount of terror. Conversely, the defenders must try to protect all of the targets for which they bear responsibility against all attack scenarios. They will attempt to minimize their risk. This model is an example of an extensive game that will be played with imperfect knowledge by all players. We cast ourselves as the defender in this game and therefore wish to minimize the attack risk. To achieve this goal, we must consider the counter-objective of the attackers—that is, to estimate risk, we must make an estimate of the optimum strategy for the attackers.

Figure 3 shows a segment of the inference tree developed for this analysis. The top of this segment is "*Change in Risk with New PM Suite*". Here, "PM suite" refers to a set of preventive/mitigative capabilities associated with a research technology. The gate type here is an INFERENTIAL AND (IA) used exclusively for inference trees and represents the operation of inferring the output node from the inputs. The manner in which this inference actually is performed will be discussed in the following paragraphs. As expected, the change in risk is inferred from the current risk (Node RC) and the future risk with the PM suite in place (Node RF). The logic used to evaluate these two risk variables is the same, and we consider the current risk in more detail here.

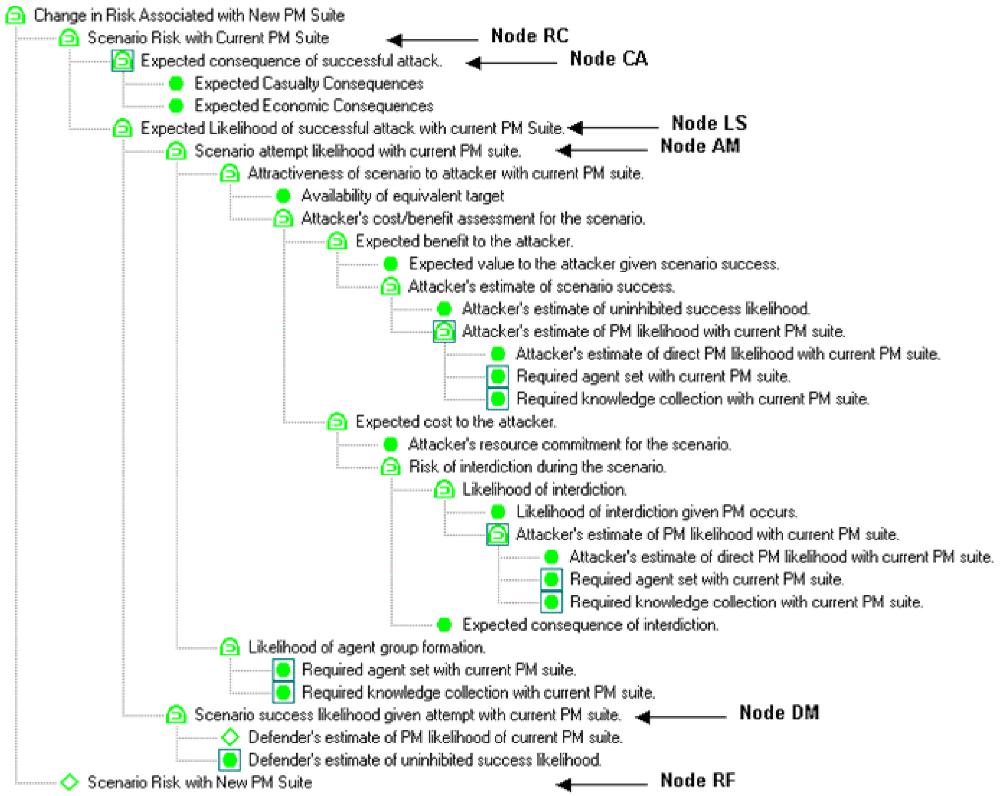
Change in Risk Associated with New PM Suite
Scenario Risk with Current PM Suite
Node RC
Expected consequence of successful attack.
Node CA
Expected Casualty Consequences
Expected Economic Consequences
Expected Likelihood of successful attack with current PM Suite.
Node LS
Scenario attempt likelihood with current PM suite.
Node AM
Attractiveness of scenario to attacker with current PM suite.
Availability of equivalent target
Attacker's cost/benefit assessment for the scenario.
Expected benefit to the attacker.
Expected value to the attacker given scenario success.
Attacker's estimate of scenario success.
Attacker's estimate of uninhibited success likelihood.
Attacker's estimate of PM likelihood with current PM suite.
Attacker's estimate of direct PM likelihood with current PM suite.
Required agent set with current PM suite.
Required knowledge collection with current PM suite.
Expected cost to the attacker.
Attacker's resource commitment for the scenario.
Risk of interdiction during the scenario.
Likelihood of interdiction.
Likelihood of interdiction given PM occurs.
Attacker's estimate of PM likelihood with current PM suite.
Attacker's estimate of direct PM likelihood with current PM suite.
Required agent set with current PM suite.
Required knowledge collection with current PM suite.
Expected consequence of interdiction.
Likelihood of agent group formation.
Required agent set with current PM suite.
Required knowledge collection with current PM suite.
Scenario success likelihood given attempt with current PM suite.
Node DM
Defender's estimate of PM likelihood of current PM suite.
Defender's estimate of uninhibited success likelihood.
Scenario Risk with New PM Suite
Node RF

Figure 3 – Inference Model for Attack Risk

The risk to the defender associated with a scenario is inferred from the aggregate consequence resulting from the attack and the likelihood that the attack is successful. Current risk is inferred from the expected aggregate consequences (CA) of a successful attack and the likelihood of a successful attack (LS). The aggregate consequences are evaluated based on casualty and economic consequences. In principle, other impacts such as psychosocial consequences also could be considered. As with the possibility tree, we use natural-language expressions in the deduction of the inferential structure. The objective is to deduce those factors determining risk and how they logically combine. The terminal nodes in this model are the input variables that are needed to evaluate the risk. A simplified representation for the attacker's decision process, called the Attacker Model, is shown in figure 3, starting with the node "Scenario attempt likelihood with current PM suite." (Node AM). The basic inferential model is a cost/benefit analysis from the attacker's perspective. It is evident that the attacker's actual decision model is inaccessible to us, and we can only try here to approximate the analysis of a rational opponent. The attacker's cost estimate depends on the inherent difficulty of the attack, as well as the risk of interdiction. Similarly, in evaluating benefit, a rational attacker takes into account the perceived gain resulting from a successful attack and the likelihood that the attack is in fact successful. This latter node is inferred from the uninhibited success likelihood—the chances of success even with no defense—and the attacker's estimate that the defender can prevent the attack. The attacker also must assemble the appropriate agent group to perform the attack. In our simple model, we picture some general pool of agents available to the attacker and shrink this pool, depending on the knowledge required to carry out the attack. The corresponding defender model (DM) is a simple vulnerability estimate that depends again on the uninhibited success likelihood and the likelihood that the attack can be prevented. The defenders' estimate of the chances of an attack succeeding in the absence of any preventive actions may well be different from what they believe to be the attackers' estimate. The prevention likelihood will depend on the preventive measures in place and their effectiveness, information that is directly available to the defender.

Approximate Reasoning Model for Risk

The inferential model identifies the basic factors and their relationships. At this point in the analysis, a decision must be made as to how the risk metric is to be measured using these factors. One approach is to represent the inferential process as a series of linked numerical functions. This approach is a natural choice when much of the input data are quantitative and when the data relationships easily can be cast in functional form. Consideration of the problem at hand suggests that these conditions are not met. Much of the data is qualitative in nature, and the nature of the attacker/defender extensive game is primarily one of perception. An analysis methodology that is well suited to this set of inferential specifications is approximate reasoning (AR) (ref. 4). AR is intended to emulate the type of expert judgment used by subject-matter experts. Variables in AR are linguistic—natural-language expressions that can take on a set of discrete natural-language values. A one-to-one correspondence exists between the natural language used in the inference tree and these linguistic variables.

In AR, inferences with linguistic variables are made using formal logical implication defined in a series of linked, forward-chaining rule bases. A detailed discussion of AR is beyond the scope of this paper; we offer a short example herein to demonstrate the compatibility of AR with the type of expert-based inferential reasoning we wish to emulate. From figure 3, we conclude that risk, R, can be inferred from the aggregate consequences, C_a, and from the likelihood of a successful attack, L_s, as

$$C_a \wedge L_s \Rightarrow R \quad , \tag{1}$$

where $\Rightarrow$ is the implication operator. A linguistic variable takes on the values from its universe of discourse (UOD), which is the set of words used to describe it. Here we use $R \in$ {Very Low, Low, High, Very High}, $C_a \in$ {Negligible, Moderate, High, Catastrophic}, and $L_s \in$ {Very Unlikely, Unlikely, Likely, Nearly Certain}. The rule base that implements the logical implication is shown in table 2. The shaded entry in the rule base corresponds to the implication "If the Aggregate Consequences are High and the Attack Success Likelihood is Unlikely Then the Risk is Low". Each linguistic variable associated with figure 3 requires a UOD, and a rule base is associated with each INFERENTIAL AND gate in this model. The UODs and the rule bases are chosen to be consonant with the judgments of the subject-matter experts for the problem: in this case, intelligence officers, commercial pilots, etc.

Table 2 – Rule Base for Inferring Risk from Aggregate Consequences and Likelihood of a Successful Attack

Aggregate Consequences	Catastrophic	*Low*	*High*	*Very High*	*Very High*
	High	*Low*	*Low*	*High*	*Very High*
	Moderate	*Very Low*	*Low*	*Low*	*High*
	Negligible	*Very Low*	*Very Low*	*Very Low*	*Low*
		Very Unlikely	Unlikely	Likely	Nearly Certain
		Likelihood of Successful Attack			

To perform evaluations with this AR model, we treat each element in a UOD as a fuzzy set. The ability to express ambiguous or hedged expressions of a linguistic variable (e.g., "The aggregate consequences are moderate to high, but somewhat closer to moderate") is encoded using the degree of membership vector, $\mu(C_a)$ = [Negligible: 0, Moderate: 0.8, High: 0.2, Catastrophic: 0]. The concept of likelihood appears many times in the risk model. Likelihood is a natural-language expression associated with outcome uncertainty. We interpret a degree of membership for L_s (e.g., $\mu(L_s)$ = {Very Unlikely: 0.2, Unlikely: 0.4, Likely: 0.3, Nearly Certain: 0] as an expression of the possibility that the likelihood is described by each of these descriptors. Note that the sum of these memberships is not 1.0—that is, our expression of outcome uncertainty is imprecise and non-probabilistic. A natural-language equivalent might be “the likelihood of a successful attack is in the range from very unlikely to likely”. Logical implication with fuzzy/possibilistic inputs to a rule base is performed using the min-max operator (ref. 5). For the two vectors given here, the corresponding risk vector is $\mu(R)$ = [Very Low: 0.2, Low: 0.4, High: 0.2, Nearly Certain: 0], which can be expressed variously as “the risk is low,” “the risk is very low to high,” etc. The application of linguistic variables, implication rule bases, and fuzzy sets in AR provides the set of capabilities needed to evaluate risk for the aircraft-attack problem.

Illustrative Attack-Scenario Risk Ranking

We applied the inference model of figure 3 using the AR implementation discussed previously. This application required determining the PM measures currently in place for each scenario and estimating their effectiveness. In addition, many of the inferential inputs associated with the attacker model required expert elicitation from intelligence experts. The use of this information and the resultant risk ranking of the scenarios cannot be discussed here. Instead, we present illustrative results to show how the rank ordering can be used.

Figure 4 shows illustrative results for the spanning set of scenarios. The results are shown for the attack with and without an insider(s) as a member of the attack group. The height of the bars is designed to reflect the memberships in the risk UOD. A height of 54 would correspond to a risk with full membership in the “very high” risk set and similarly for “high” (18), “low” (6), and “very low” (1.5). Although strictly notional, the figure shows how including an insider in the attacker group can increase the risk. This increase occurs because many of the PM measures in place for certain scenarios are aimed primarily at preventing intruders or passengers from carrying out the attack. For our model scenario, Node AF-2, this is not the case. Including an insider does not increase the likelihood of a successful attack.

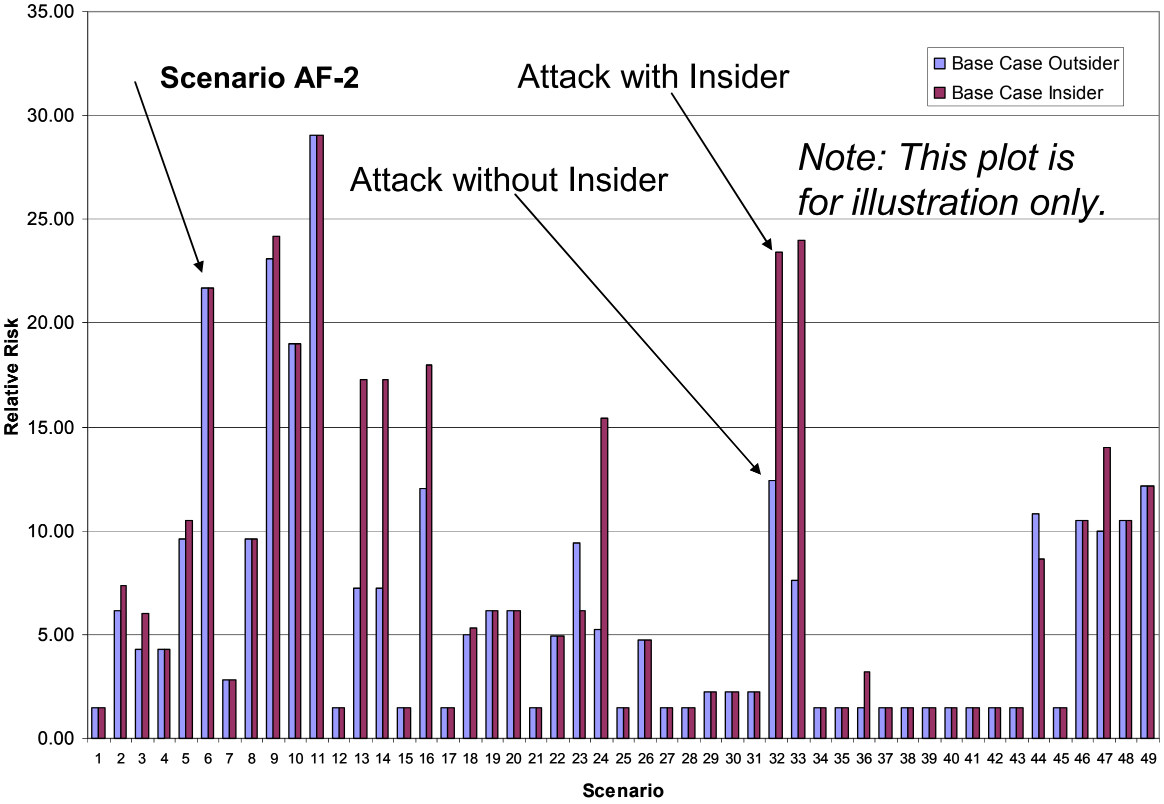

Figure 4 – Illustrative Risk Rankings for the Spanning Set of Scenarios

Technology-Dependent Risk Reduction

NASA Aircraft Security Technology Portfolio: The objective of the Aviation Security (AvSEC) initiative within NASA's AvSSP is to increase the resiliency of the Air Transport System (ATS) from threats and hostile acts and to identify and inform users of potential risks. The AvSEC initiative seeks to reduce the risk to the ATS by limiting vulnerabilities, preventing vulnerabilities from being exploited, and mitigating the consequences. The project is focused on developing technologies for aircraft and onboard systems. Existing technologies and ongoing efforts also will be leveraged to reduce the risk to other components in the ATS, such as the National Airspace System (NAS). NASA is considering technologies that will address all aspects of the aviation industry, including commercial, business, and general aviation. Technologies will be developed with the recognition that meeting security requirements is one consideration in an integrated system that must also meet safety, capacity, and mobility requirements.

Eight specific subprojects within AvSEC deal with PM technologies that address risk reduction for Part 121 PC aircraft: lightweight fire- and explosion-resistant (LWFR) materials to reduce the damage to the airframe from an attack, protected asset flight systems (PAFS) to prevent intrusion of aircraft into restricted airspace, damage-adaptive controls (DAC) to provide automated response to damage to control surfaces and engines from an attack, vehicle recovery (VR) to allow for the safe landing of aircraft in association with PAFS intervention, electromagmetic emission (EME) protection to prevent interference from external fields with flight-critical systems, safe and secure information flow (SASIF) to ensure that voice and data communications (including navigational signals) between aircraft and the ground are not compromised, fuel tank inerting (FTI) to limit the damage from fires and explosions, and biological and chemical sensing (BIOS) to detect WMD attacks using aircraft.

Potential Risk Reduction: Each proposed PM suite could be effective against some subset of the spanning set of attack scenarios. The actual effectiveness of the PM suites currently is unknown because much of the research needed to establish the effectiveness has not yet been begun. For this reason, we decided to treat the effectiveness parametrically. The maximum risk reduction occurs when a technology is assumed to be "fully effective" against an attack. We then aggregated the inferred risk reduction for a particular PM suite across the spanning set. Figure 5 shows the preliminary risk reduction determined in this manner for the PM suites. Technologies that are effective against many scenarios or a few relatively high-risk scenarios will have the greatest aggregate risk reduction.

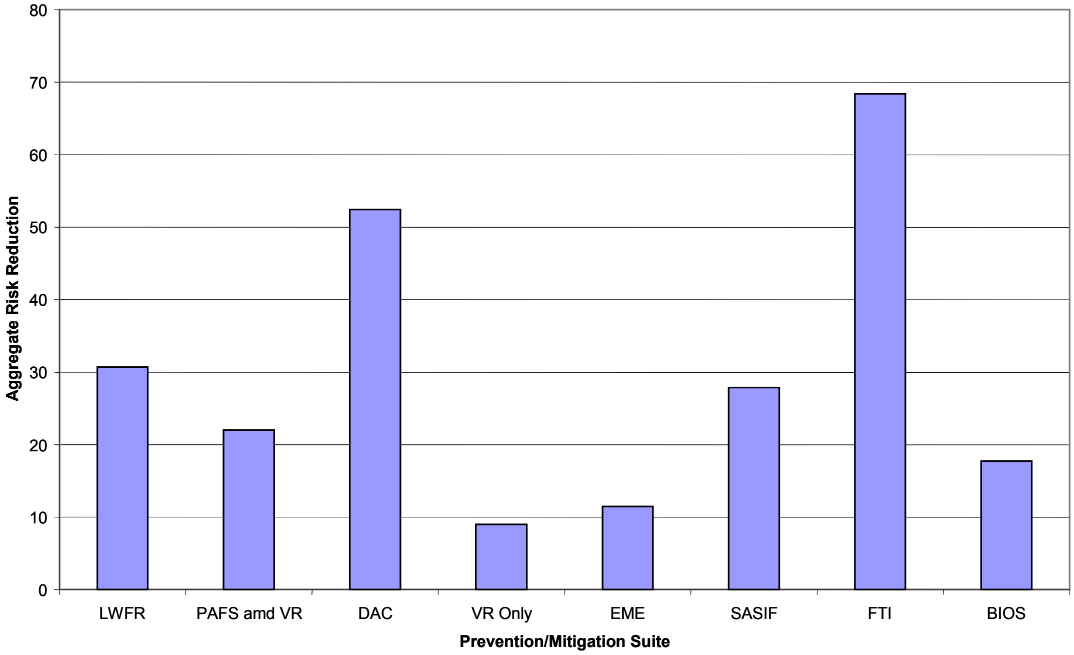

Figure 5 – Aggregate Risk Reduction Potential for Proposed PM Suites

Discussion and Conclusions

The work presented here is the initial phase in an ongoing effort to prioritize NASA security technology research. Currently, we are working on estimating the current risk for Part 121 all-cargo aircraft and Part 91 and Part 135 aircraft, as well as for airports and the national air space. Each of these types of operations introduces new attack scenarios, as well as significant changes in available PM measures. Similarly, the potential risk reduction associated with NASA PM technologies will vary for each set of attack scenarios. In the course of evaluating the risk for these various operations, it has become clear that the risk analysis is an iterative process. As more information becomes available, the spanning sets may change either in the number of scenarios or in the level of detail used to define a particular scenario. In addition, the aviation security system continues to evolve with the introduction of new technologies and procedures. These changes also affect risk. For this reason, the risk reduction results presented here should also be viewed as preliminary. We are working actively with the NASA technologists to understand better how their research reduces risk.

The use of risk reduction as a single rank-ordering metric should be viewed only as a starting point. Other factors clearly play important roles as well. In the near future, cost/benefit analyses will be available for the proposed technologies. The results of these analyses will make it possible to reflect in the rankings the value of technologies that reduce costs, including the costs attributable to delays caused by existing technologies and procedures. Other additions planned for the LED inferential model will consider technical risk and the development of a concept of operations for the technologies.

References

1. T. F. Bott and S. W. Eisenhawer, “A Logic Model Approach to the Conceptual Design of a Scientific/Industrial Complex,” LA-UR-02-4756, ASME-PVP Annual Meeting, Vancouver, 2002, PVP-444, pp. 119-127.
2. T. F. Bott, S. W. Eisenhawer, J. Kingson, and B. P. Key, “A New Graphical Tool for Building Logic-Gate Trees,” LA-UR-03-134, ASME-PVP Annual Meeting, Cleveland, August 2003.
3. S. W. Eisenhawer, T. F. Bott, and D. V. Rao,” Assessing the Risk of Nuclear Terrorism Using Logic Evolved Decision Analysis,” LA-UR-03-3467, to appear in Proc. 2003 ANS Annual Meeting, San Diego, June 2003.
4. R. Lopez de Mantaras, Approximate Reasoning Models, Ellis Howrood, New York (1990).
5. S. W. Eisenhawer, T. F. Bott, and R. E. Smith, “An Approximate Reasoning-Based Method for Screening High-Level-Waste Tanks for Flammable Gas,” *Nuclear Technology*, Vol. 130, June 2000, pp. 351–361.

Biographies

Stephen W. Eisenhawer, Ph.D., Technical Staff Member, Los Alamos National Laboratory, MS K557, Los Alamos, NM, 87545, USA, telephone – (505)-667-2420, facsimile – (505) 667-5531, e-mail – seisenhawer@lanl. gov.

Dr. Eisenhawer is a technical staff member in the Decision Applications Division at Los Alamos National Laboratory. Previously, he was Vice President of Systems Engineering at Enhanced Energy Systems in Albuquerque, New Mexico, and a member of the technical staff at Sandia National Laboratories. His technical training is in engineering; he received his B.S. in Mechanical Engineering from Seattle University, his M.S. in Nuclear Engineering from the University of Washington, and his Ph.D. in Nuclear Engineering from the University of Washington in 1977 following doctoral research at the Kernforschungszentrum Karlsruhe in Germany. He is a Registered Professional Engineer and the co-holder of two patents. Together with Dr. T. Bott, he has developed the LED analysis methodology for decision support. They have published approximately 30 peer-reviewed journal and conference papers describing the methodology and its application to a wide variety of practical decision-analysis problems.

Terry F. Bott, Ph.D., Technical Staff Member, Los Alamos National Laboratory, MS K557, Los Alamos, NM, 87545, USA, telephone – (505)-667-9027, facsimile – (505) 667-5531, e-mail – tbott@lanl.gov.

Dr. Bott is a staff member in the Probabilistic Risk Analysis Group at Los Alamos National Laboratory. He has a B.A. in physics from the University of Utah and a Ph.D. in Chemical Engineering from Brigham Young University. He did his doctoral research on fast reactor disassembly calculations. Following graduation, Dr. Bott’s research was in the areas of nuclear reactor safety and risk assessment at Oak Ridge National Laboratory and of nuclear reactor thermal-hydraulics code development at Los Alamos. Recently, Dr. Bott has been involved in design and safety studies for nuclear weapons and other systems. His principal research interests are in the use of formal logic models in decision analysis and knowledge representation.

Michael R. Sorokach, Jr., B.S., Systems Analysis Branch, Aerospace Systems, Concepts & Analysis Competency, MS 348, 8 Langley Blvd, NASA-Langley Research Center, Hampton, VA, 23681-2199, USA, telephone – (757) 864-7143, facsimile – (757) 864-6306, e-mail – Michael.R.Sorokach@nasa.gov.

Mr. Sorokach is an Aerospace Engineer in the Systems Analysis Branch at the NASA Langley Research Center (LaRC). His technical training is in engineering; he received his B.S. in mechanical engineering from Clemson University in 1986. Previously, he was a mechanical engineer in the Model Systems Branch at LaRC. Mr. Sorokach is presently the Level III Program and Vulnerability Assessment Lead for Technology Integration of the AvSEC Initiative within AvSSP. His work is focused on identifying national air-transportation-system vulnerabilities and developing program prioritization approaches, decision support tools, and criteria that can be used to assess the performance of research areas within the AvSEC. He has published one conference paper and authored several system study reports for the AvSSP, submitted an invention disclosure for patent, and received numerous honors and awards throughout his career at LaRC.

Frank P. Jones, B.S., Aviation Safety and Security Program, 8 Langley Blvd., NASA-Langley Research Center, Hampton, VA, 23681-2199, USA, telephone – (757) 864-5271, e-mail – f.p.jones@nasa.gov.

Mr. Jones is the Technical Integration Project Lead of the NASA ArSSP. He has been involved with NASA's safety initiatives since 1997 when, as the Project Manager, he led the formulation of the Aviation Weather Information Project. Mr. Jones has held numerous research and management jobs in his 24 years with NASA: Airspace Operations Systems Manager and Deputy Manager of the Advanced Subsonic Noise Reduction Program at LaRC, NASA Headquarters Office of Aeronautics Flight Controls Program Manager responsible for the Advanced General Aviation Transport Experiment, Fly-by-Light/Power-by-Wire, Propulsion Controlled Aircraft, and Flight Critical Systems projects. Mr. Jones was a flight systems researcher at NASA Dryden Flight Research Center from 1981 to 1988. Between 1989 and 1994, he worked on the Space Station Freedom Program as the Configuration Analysis Manager. Mr. Jones received his B.S. in aeronautical/astronautical engineering from Purdue University in 1983.

John R. Foggia, M.S., 619 Cortland Dr., Lake Zurich, IL, 60047, USA, telephone – (847) 438-3294, e-mail – jfoggia@ameritech.net.

Currently, John Foggia is developing Concepts of Operation for NASA's Aviation Security technology programs as an independent Aviation Engineering contractor for LaRC. He has performed independent consulting work for Los Alamos National Laboratory as an "Aviation Expert," developing aircraft-attack scenarios and LED analysis for aviation security threats. From 1997 to early 2003, Mr. Foggia was President and Chief Pilot of Aviation, Navigation, and Satellite Programs, Inc. (ANSP), directing a company specializing in integrating leading-edge aviation technologies into NAS modernization—including test flight for NASA. From 1990 to 1997, Mr. Foggia was an Airport Manager at the Minneapolis/St. Paul International Airport. Mr. Foggia served as an active-duty Air Force pilot, flying fighter aircraft in the 1980s, including A-10s and F-5Bs, and also an Air Force meteorologist, managing numerous USAF radar and satellite system test and evaluation programs. Mr. Foggia was an MIT Research Scientist in the Fluid Dynamics Engineering doctoral program. He holds an M.S. in atmospheric physics from Texas A&M University and a B.S. in mathematics from Northern Arizona University.