

Attack tree analysis

D.R. Wiseman

Isograph Ltd., Warrington, United Kingdom

ABSTRACT: Fault tree analysis provides a means of analyzing the causes and probability of a hazardous event. Typically, the hazard of interest will be a system failure or a catastrophic event, caused by random events and component failures. However, in today's increasingly interconnected world, system hazards are more likely than ever to originate from deliberate attacks, such as hacking and malware. This paper will demonstrate how a modified form of fault tree analysis, henceforth referred to as attack tree analysis, may be used to predict the frequency of a threat due to attacks on a system and the failure of defensive measures. The paper will also demonstrate how attack tree analysis can take into account the impact of a successful attack on factors such as cost and safety, as well as the cost and difficulty incurred by the attacker.

1 INTRODUCTION TO ATTACK TREES

Fault Tree Analysis (FTA) is a deductive analysis technique used to determine all possible causes of a hazard. A typical fault tree consists of logic gates and basic events, where the logic gates indicate how failures and events interact to cause a hazard, and basic events represent the most fundamental events to be modelled in the analysis. Though a basic event may be used to represent a variety of different types of event, they are most commonly used to represent failures, such as component failures and component failure modes. It is unusual for an event to represent a deliberate attack on a system.

Attack Tree Analysis (ATA) uses the principles of FTA to provide a concise, easy to understand method of modelling threats to the security of a system. In an attack tree, basic events may represent deliberate attacks against a system, and the failure of security measures that have been put in place to prevent attacks from succeeding. Logic gates indicate how attacks and security failures combine to result in a successful attack.

Similarly to a FTA, the two main outputs from an ATA are the minimal cut sets, which indicate the combinations of attack and security failures that would result in an attack being successful, and the frequency of a successful attack. The consequences of a successful attack may also be included in the model. Unlike a FTA, an ATA must also account for the difficulty to the attacker. An attack might be likely to succeed in theory, but in reality the method

of attack selected may depend on factors such as the skills and equipment required for the attack.

2 INITIATORS AND ENABLERS

The basic events in an attack tree fall into two categories – initiators and enablers. An initiator is an event which must fail last in a sequence if it is to cause a hazard. An enabler, on the other hand, must fail anywhere but last in the sequence if it is to contribute to a hazard. A familiar example would be that of a fire. A fire will cause a safety hazard if it occurs after the failure of a fire suppression system. If the fire occurs while the suppression system is still functional, the hazard is mitigated. Thus, the fire is treated as an initiator, and the failure of the fire suppression system as an enabler.

In the context of an attack tree, initiator events are used to represent attacks, and enablers to represent security measures. For example, a typical initiator would be a hacking attack, and a typical enabler would be a security patch being out of date.

3 CONSTRUCTION

An attack tree is constructed using a top-down methodology, similar to that used for fault trees. The hazard, or TOP event, is selected to represent the success of an attack, and will determine the complexity of the tree. Next, the direct causes of the

TOP event must be identified. This process of cause identification is continued throughout the intermediate levels of the tree, until the basic events are reached.

The TOP event and intermediate levels of the tree are represented by logic gates. These gates determine how causal events combine to result in a hazard. The logic gates that can be used include OR (occurrence of at least one event will result in a hazard), AND (all events must occur in order to cause a hazard) and VOTE gates (m out of n events must occur in order to cause a hazard, where n is the number of inputs to the gate and m is the vote number).




Figure 1: OR, AND and VOTE gate symbols

4 ANALYSIS

Like a FTA, an ATA breaks down into two, distinct stages: the qualitative analysis, wherein the sets of events which will lead to a successful attack are determined, and the quantitative analysis, wherein the frequency of a successful attack is calculated.

4.1 Qualitative Analysis

The qualitative part of an ATA, often described as the minimal cut set analysis, uses Boolean algebra to determine the minimal cut sets (MCS) for the tree. The MCS are the minimum sets of events which, if they occur together, will result in a successful attack. During MCS analysis, each logic gate is replaced by the appropriate Boolean operator ('+' for OR logic, '.' for AND logic). A process of substitution is then used to produce a Boolean expression to represent the TOP event, which is then simplified to give a list of MCS. This procedure does not require any quantitative data. That is, a MCS analysis can be performed without any frequency or probability data (Vesely & Goldberg 1981). Note that the attack tree must be constructed such that each cut set only contains one initiator.

4.2 Quantitative

Once the MCS analysis is complete, the quantitative analysis may be performed. This part of the ATA requires that frequency data is provided for all initiator events and probability data for all enabler events in the tree. First, the frequency of each cut set must be calculated. This is done by multiplying the fre-

quency of the cut set initiator by the probability of each of the enabler events in the set. The frequency of a successful attack is then determined by combining the cut set frequencies using the Cross Product method (Vesely & Goldberg 1981). Here, the addition rule of probability is applied to the cut sets in order to determine the TOP event frequency.

4.3 Approximation Methods

The analysis of a large, complex attack tree can be a time-consuming process, even for a computer program. Thus it may be necessary to apply an approximation method in order to complete the analysis in a practical amount of time. One such method is the Upper Bound approximation. Here, the analyst may specify the maximum number of cross product terms of each order that should be calculated. If this limit is reached, the calculation will end with the last set of upper bound terms to fall below the limit. Another example is that of the Esary-Proschan approximation (Esary & Proschan 1963).

Both of the methods discussed are upper bound methods, so any deviation from the true result will always tend towards the pessimistic, meaning that the ATA will not underestimate risk. Note that these methods are most accurate when the basic event probabilities are small.

4.4 Consequences and Risk

Any successful attack on a system is likely to have consequences. In ATA, consequences may be allocated to the TOP event. Each consequence may have an associated weight (a numerical value indicating the severity of the attack) which may then be used to calculate the risk due to an attack. Risk may then be calculated by multiplying the frequency of a successful attack by the weight of the associated consequence.

Consequences fall into a variety of categories, including safety, financial and security. ATA may be used to determine the total risk to a system due to consequences across a range of categories. These risk values are then compared to limiting values, set internally by the analyst, or externally by a regulatory body. This can prove useful at the design stage, when trying to build a system that is robust against attack.

As well as determining the risk due to an attack, it is also possible to determine the sensitivity of system risk to changes in initiator frequency and enabler probability. These sensitivity values provide an indication of how risk might be mitigated most effectively. Risk sensitivity is calculated using Equation 1:

$$I_i^{BB} = \frac{\delta P}{\delta p_i} \quad (1)$$

where I_i^{BB} is the risk sensitivity for the event i , P is the risk due to a given consequence, and p_i is the frequency or probability of the event i . Reducing the frequency or probability of an event with a high risk sensitivity will have a large, positive impact on the system risk.

4.5 Indicators

ATA provides an estimate of how frequently an attack will succeed if attempted. However, the analysis must also account for the obstacles involved for the attacker. For example, hacking into a computer may require a high level of expertise, while breaking into a bank vault may require very expensive equipment. Indicators are numerical values that are used to specify the extent of the problems that the attacker must overcome in order to successfully attack a system. Typical indicator categories include cost, difficulty and equipment. Indicator values are applied at the basic event level.

In order to determine the indicator values for a successful attack, the analyst must choose how the indicator values are to be calculated for the intermediate gates in the attack tree. For example, the analyst may determine that the lowest cost should be selected for OR logic (i.e. if the attacker has a choice of methods, they will select the cheapest), whereas cost values should be summed in the case of AND logic (e.g. the attacker must pay all costs in order for the attack to be successful).

5 EXAMPLE

In August 2015, Charlie Miller and Chris Valasek published a paper describing how an attacker could gain access to on-board systems in a car by hacking into the entertainment system using wireless technology (Miller & Valasek 2015). Using this method, an attacker could gain access to door locks and headlights, and even critical systems such as steering and brakes. Furthermore, research performed at the University of California, San Diego, has demonstrated how such attacks could also be performed via a normal mobile phone, if the target vehicle has an on-board diagnostics (OBD) dongle installed (Foster & Koscher 2015). The following is a simple, hypothetical example demonstrating how to build and analyze an attack tree to determine the frequency and difficulty of an attacker successfully accessing a vehicle's on-board computer. (Note that the quantitative data presented here are not taken from a specific source, but have instead been generated for the purposes of this example.)

5.1 Attack Tree Construction

The frequency with which an attacker can gain access to critical systems via the on-board computer is to be determined. The TOP event is thus defined as 'Hacker gains access to on-board computer', henceforth referred to as ACCESS. ACCESS will occur if a wireless attack is attempted and the on-board computer is not secure. The on-board computer will be vulnerable if the entertainment system does not have the latest security patch installed, or if an OBD dongle is available for wireless connection. An OBD dongle will be available for connection if the driver has connected a dongle, and the dongle security patch is not up to date. The resulting attack tree is shown in Figure 2.

The ATTACK event is an initiator, representing the frequency with which a wireless attack is expected to take place. For the purposes of this example, let us assume that in the United States, the frequency of a hacking attack against a vehicle is $1 \times 10^{-12} \text{ hour}^{-1}$.

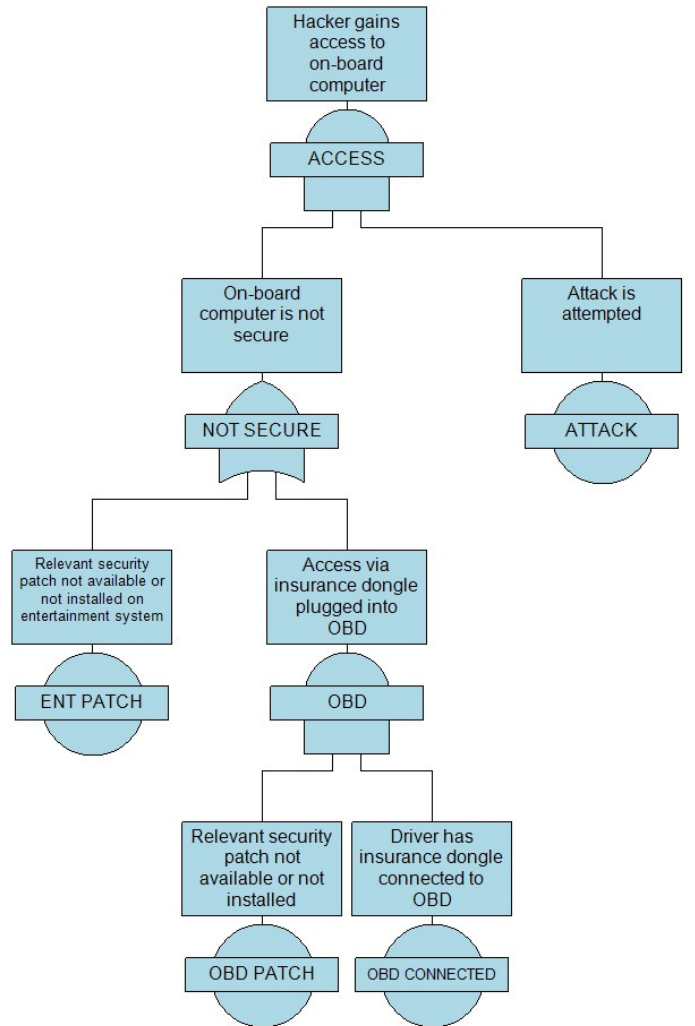


Figure 2: Simple attack tree representing the frequency of a success attempt to hack into a car's on-board computer.

The other events in the tree are enablers, representing the probability of security failures in the system. Approximately 2% of vehicles in the United States have a OBD dongle installed. The probability of se-

curity patches being out of date is estimated to be 0.25. The frequency and probability data of the attack tree events are given in Table 1.

Table 1. Quantitative data for the attack tree shown in Figure 2.

Event	Frequency hour ⁻¹	Probability	Cost \$	Difficulty	Equipment
Attack	1×10^{-12}				
Ent patch		0.25	1000	10	100
OBD patch		0.25	100	1	10
OBD Connected		0.02			

The cost, difficulty, and equipment required for an attack will differ depending on the method of access. Attacking via an un-patched entertainment system will require a laptop or PC with specialized software installed, whereas attacking via an unsecured OBD dongle requires a mobile phone and a downloadable app. The cost of a laptop installed with the require software is estimated to be \$1000, and that of a mobile phone to be \$100.

While the cost of an attack may be expressed easily in units of currency, more abstract indicators such as the difficulty and the sophistication of the equipment required must be quantified using a dimensionless scale. In this example, the sophistication of the equipment used by the attacker is expressed on a scale of 1 to 100. The equipment required for an attack via the entertainment system is more sophisticated than that required for an attack via an OBD dongle. Hence, the analyst has allocated an equipment indicator value of 100 for the ENT PATCH event, and 10 for the OBD PATCH. A similar principle has been followed when allocating the difficulty indicator values. The cost, difficulty and equipment indicators for each event are given in Table 1.

As stated in Section 4.4, the analyst must specify the logic used to determine the indicator values at each gate in the tree. In most cases, it is reasonable to assume that an attacker will prefer the cheapest and easiest path to a successful attack. Thus, the minimum cost and difficulty are selected at each OR gate (given the choice, the attacker chooses the cheaper, easier option), and the summed values are taken at each AND gate (cost and difficulty is additive if the attacker must employ more than one method of attack). The minimum equipment value is taken at each OR gate, and the maximum taken at each AND gate (it is assumed that the more sophisticated equipment can be employed both for complex and simpler and attacks).

5.2 Attack Tree Results

A quantitative analysis of the example attack tree indicates that an attack may be expected to occur with a frequency of 2.538×10^{-13} hour⁻¹, and that the

easiest path of attack is expected to cost \$100, with a difficulty of 1 and equipment requirement of 10.

More detailed conclusions may be drawn by examining the cut sets. The cut sets, along with their calculated frequencies and indicator values, are shown in Table 2. Note that while an attack via an unpatched entertainment system is expected to occur with the greatest frequency, it is also expected to be the more difficult option for an attacker. This information might indicate to the analyst that, while a particular path of attack is more likely to be open, the obstacles to success might prove to be prohibitive.

Table 2. Qualitative data for the attack tree shown in Figure 2.

Cut set	Frequency hour ⁻¹	Cost \$	Difficulty	Equipment
Attack. Ent patch	2.5×10^{-13}	1000	10	100
Attack.OBD patch. OBD Connected	5×10^{-15}	100	1	10

5.3 Risk and Sensitivity

While a successful attack could grant the attacker access to critical on-board systems, this is only possible when the vehicle is stopped, or travelling at low speeds (< 15 kph). For this reason, the safety consequence of ACCESS is judged to be moderate. A moderate safety consequence has a weight of 5 on an dimensionless scale of 1 to 20, where 1 is the least severe and 20 the most. For the purposes of our example, we will assume that a regulatory body has devised this consequence scale, and has placed an upper limit of 1×10^{-12} hour⁻¹ on risk.

The safety risk due to a moderate consequence of a successful attack is 1.269×10^{-12} hour⁻¹. The predicted risk exceeds the limiting value, suggesting that on-board computer security must be improved. Table 3 shows the risk sensitivity results for the example tree.

Table 3. Sensitivity data for the attack tree shown in Figure 3.

Event	Sensitivity
Attack	1.275
Ent patch	5×10^{-12}
OBD connected	1.25×10^{-12}
OBD patch	1×10^{-13}

If we assume that it is not possible for the security analyst to affect the frequency of an attack, then the most effective means of mitigating the risk due to a successful attack is to ensure that the entertainment security patch is up to date.

6 CONCLUSIONS

In this paper, it has been demonstrated that ATA is a useful means of understanding and modeling threats to a system. This analysis technique may be used to predict the frequency with which an attack can be expected to succeed, and also to quantify the obstacles that must be overcome by the attacker. The risk due to attack can also be determined for a range of different risk categories, such as safety and cost. A brief discussion was made of approximation methods that may be used to expedite an analysis. These methods are naturally pessimistic, meaning that the ATA will not underestimate risk. Finally, it was shown that risk sensitivity results may be used to guide the design or redesign of a system to make it more robust against attacks.

7 REFERENCES

- Vesely, W.E., Goldberg, F.F. & Haasl, D.F. 1981, NRC NUREG-0492, VII-15-VII-19.
- Vesely, W.E., Goldberg, F.F. & Haasl, D.F. 1981, NRC NUREG-0492, VII-1.
- Esary, D. & Proschan, F. 1963 Coherent structures with non-identical components. *Technometrics* 5(2), 191-209.
- Miller, C. & Vasalek, C. 2015 Remote Exploitation of an unaltered passenger vehicle. In *Black Hat USA 2015; Proc. intern. symp., Las Vegas, 1-6 August 2015*.
- Foster, I. & Koscher, K. 2015 Exploring controller area networks. In *24th Usenix Security Symposium; Proc. intern. symp., Washington, D.C., 12-14 August 2015*.