Risk Analysis of Changing and Evolving Systems using CORAS

Mass Soldal Lund[a], Bjørnar Solhaug[a], and Ketil Stølen[a,b]

[a]SINTEF ICT, P.O. box 124 Blindern, 0314 Oslo, Norway
[b]Department of Informatics, University of Oslo, Norway
{mass.s.lund,bjornar.solhaug,ketil.stolen}@sintef.no

Abstract. Risk analysis is the identification and documentation of risks with respect to an organisation or a target system. Established risk analysis methods and guidelines typically focus on a particular system configuration at a particular point in time. The resulting risk picture is then valid only at that point in time and under the assumptions made when it was derived. However, systems and their environments tend to change and evolve over time. In order to appropriately handle change, risk analysis must be supported with specialised techniques and guidelines for modelling, analysing and reasoning about changing risks. In this paper we introduce general techniques and guidelines for managing risk in changing systems, and then instantiate these in the CORAS approach to model-driven risk analysis. The approach is demonstrated by a practical example based on a case study from the Air Traffic Management (ATM) domain.

Keywords: Risk management, risk analysis, change management, Air Traffic Management, security

1 Introduction

Risk management is coordinated activities to direct and control an organisation with regard to risk [15]. Risk analysis is a core part of risk management, and involves the identification and documentation of risks with respect to the organisation or the target system in question. When deriving the risk picture for the target of analysis, established risk analysis methods and guidelines typically focus on a particular system configuration at a particular point in time. The resulting risk picture is then valid only at that point in time and under the assumptions made when it was derived. However, systems and their environments tend to change and evolve over time.

In order to appropriately handle change, each of the risk management activities must be supported with specialised techniques and guidelines. Figure 1 is adapted from the ISO 31000 risk management standard [15] and illustrates the seven activities of the risk management process. The five activities in the middle constitute the core activities of a risk analysis, and are described as follows:

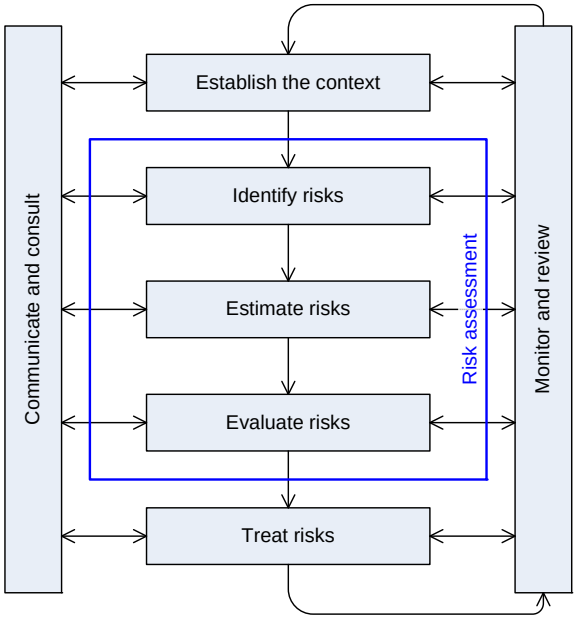

Fig. 1. Risk management process

- *Establish the context* is to define the external and internal parameters to be accounted for when managing risk, and to set the scope and risk criteria for the risk management policy.
- *Risk identification* is to find, recognise and describe risks.
- *Risk estimation* is to comprehend the nature of risk and to determine the risk level.
- *Risk evaluation* is to compare the risk estimation results with the risk criteria to determine whether the risk and its magnitude are acceptable or tolerable.
- *Risk treatment* is the process of modifying the risk.

The remaining two activities are continuous activities of the overall risk management process, and are described as follows:

- *Communicate and consult* are the continual and iterative processes an organisation conducts to provide, share or obtain information, and to engage in dialogue with stakeholders about risk management.
- *Monitoring* involves the continuous checking, supervising and critically observing the risk status in order to identify changes from the performance level required or expected, whereas *review* focuses on the activity undertaken to determine the suitability, adequacy and effectiveness of the subject matter necessary to achieve established objectives.

When targeting changing and evolving systems the main challenge is to ensure that the analysis results are kept valid under change. A straightforward way

to ensure this is to conduct a full risk analysis from scratch when faced with a potentially risk relevant change. Needless to say, such a strategy is not to prefer as it is time and resource consuming, and as it often implies conducting exactly the same analysis again to the extent that the risk picture is persistent. Instead, a customised analysis method for changing systems should provide guidelines and techniques for how to systematically trace the relevant changes from the target to the risks, and thereby for how to update only the part of the risk picture that is affected by the changes.

How to adequately deal with changes in a risk analysis depends, however, on the nature of the changes in each specific case. We have for this reason identified three different perspectives on change, each with its specific methodological needs. These are the *maintenance* perspective, the *before-after* perspective, and the *continuous evolution* perspective.

The approach of this paper is to take as starting point the risk management principles and guidelines of the ISO 31000 standard and generalise these to the setting of changing and evolving systems. This standard is quite general, and most of the established risk management methodologies can be understood as instantiations of the activities it prescribes. The objective of such an approach is to understand and explain how to deal with changing risks without restricting to a specific methodology.

ISO 31000 comes with no guidance on risk analysis techniques. While a risk analysis method provides methodological advice on how to carry out the various activities of risk management, a risk analysis technique is more narrow in the sense that it addresses only some aspects of the risk analysis process. A risk analysis method typically makes use of one or more risk analysis techniques. Risk modelling refers to techniques that are used to support the activities of risk identification and risk estimation. In this paper we take the risk graphs of [3], which can be understood as a common abstraction of different risk modelling techniques, as the starting point for generalising risk modelling to the setting of changing systems. The idea is, as is also demonstrated in the paper, that the extensions we make to risk graphs carry over to other risk modelling techniques.

In order to demonstrate and validate our approach, we present a case study from the domain of Air Traffic Management (ATM). The specific risk analysis method applied in this case study is CORAS [20] used as an instantiation of the general approach to risk analysis of changing systems. CORAS is a model-driven approach to risk analysis that consists of a method, a language and a tool to support the risk analysis process. The ATM case study thus demonstrates a concrete instantiation of the approach in this paper in both the CORAS method and the CORAS language.

The structure of the paper is as follows. In Sect. 2 we make a classification of change by introducing the perspectives of *maintenance*, *before-after* and *continuous evolution*. As the methodological needs are strongly situation dependent, we discuss these in relation to the mentioned perspectives before we present our approach to risk analysis of changing and evolving systems in Sect. 3. Section 4 introduces the formal foundation for our approach by presenting the syntax and

semantics of risks graphs and their generalisation to the setting of changing and evolving risks. The section moreover presents a calculus with rules for reasoning about risk graphs, and a set of guidelines for how to do consistency checking of risk graphs. In Sect. 5 we instantiate our approach with CORAS, thus generalising CORAS to the setting of changing and evolving systems. In Sect. 6 we demonstrate our approach by applying this generalisation of CORAS for changing and evolving systems to the practical example of the ATM case study. Finally, we present related work in Sect. 7 before we conclude in Sect. 8.

2 Classification of Change

The appropriate way of handling change when managing, modelling, analysing and documenting risks depends greatly on the context and on the kind of changes we face [19]. For example, do the changes result from maintenance or from bigger, planned changes? Do the changes comprise a transition from one stable state of the target of analysis to another stable state, or do they reflect the continuous evolution of a target designed to change over time? Do the changes occur in the target system or in the environment of the target? The answers to such questions have implications on how to methodologically approach the problem.

In the following we present three main classes of change by introducing three respective perspectives on change: *maintenance*, *before-after*, and *continuous evolution*. After this presentation, the remainder of the paper focuses on the *before-after* perspective. This is the perspective of the case study from which the practical example is provided, and is perhaps the perspective that most clearly illustrates the need for specialised methods when dealing with risk analysis of changing and evolving systems.

2.1 The Maintenance Perspective

The changes we address from the *maintenance* perspective are those that accumulate over time and gradually make any previous risk analysis results invalid. The challenge for the risk analysts can be described by the following example scenario:

> The risk analysts conducted a risk analysis three years ago and are now requested by the same client to analyse the same target system anew so as to update the risk picture and reflect any changes to the target or its environment, thus restoring the validity of the analysis.

Typical changes in this setting can be bug fixes and security patches, an increase in network traffic, and increase in the number of attacks. For such changes, the risk picture remains more or less the same, but risk values may have changes such that previously acceptable risks are now unacceptable, or vice versa. The objective is then to maintain the documentation of the previous risk analysis by conducting an update.

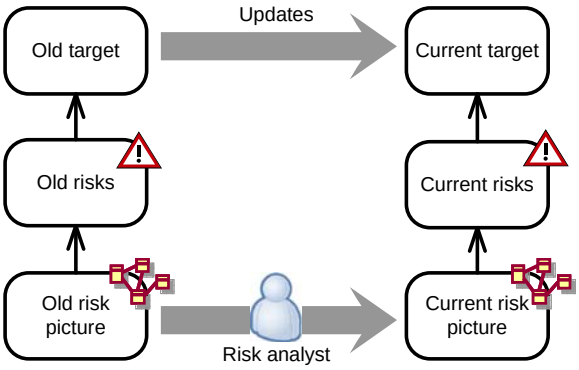

Fig. 2. The maintenance perspective

The illustration in Fig. 2 shows the principles by which the risk analysts conduct the analysis from the *maintenance* perspective. Assuming that we have available the documentation of the previous risk analysis, as well as the description of the old and current target of analysis, we start by identifying the changes that have occurred. We then use the relevant changes as input when deriving the current risk picture based on the previous risk analysis results.

Methodologically, the main challenge is how to reuse the old risk analysis results and not conducting a full risk analysis from scratch. This requires a systematic way of identifying the updates of the target and which parts of the risk picture that are affected such that the risk picture can be updated without addressing and spending effort on the unaffected parts.

2.2 The Before-After Perspective

The changes we address from the *before-after* perspective are those that are planned or anticipated, and perhaps of a radical kind. The challenge for the risk analysts can be described by the following scenario:

> The risk analysts are asked to predict the effect on the risk picture of implementing the changes to the target of analysis.

Examples of changes in the *before-after* perspective are the rolling out of a new system, or making major organisational changes such as implementing a merger agreement between two companies. We must thus understand the current risk picture and how this will change to the future risk picture as a consequence of implementing the changes to the target of analysis. As the planned or anticipated process of change may involve risks in itself, the *before-after* perspective may require a risk analysis of this process.

The illustration in Fig. 3 shows the principles by which the risk analysts conduct the analysis from the *before-after* perspective. Assuming that we have available the description of the current target of analysis as well as the change

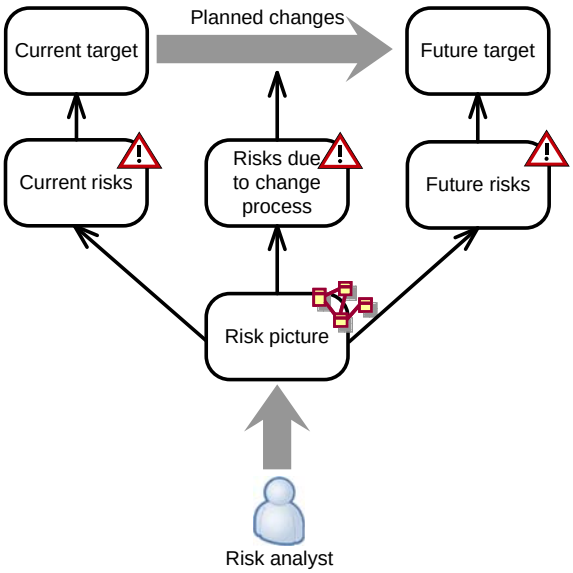

Fig. 3. The before-after perspective

process to bring the target from the current to the future state, we can work out a coherent risk picture for the future target and for the change process.

Methodologically, the main challenges are how to obtain and present a risk picture that describes the current and the future risks, as well as the impact of the change process itself. This requires an approach to making a description of the target "as-is" and a description of the target "to-be", describing the process of change in sufficient detail, identifying and documenting the current and future risks without doing double work, and identifying risks due to the change process.

2.3 The Continuous Evolution Perspective

The changes we address from the *continuous evolution* perspective are predictable and gradual evolutions that can be described as functions of time. Such predictions can be based on well-founded forecasts or planned developments. The challenge for the risk analysts can be described by the following scenario:

> The risk analysts are asked to predict the future evolutions of risk, which mandates that they conduct a risk analysis that establish a dynamic risk picture that reflects the expected evolution of the target.

Examples of predictable changes are the slow increase in the number of components working in parallel or gradually including more sites in a system. Examples of well-founded forecasts are the expected steady increase of end-users, adversary attacks, and annual turnover.

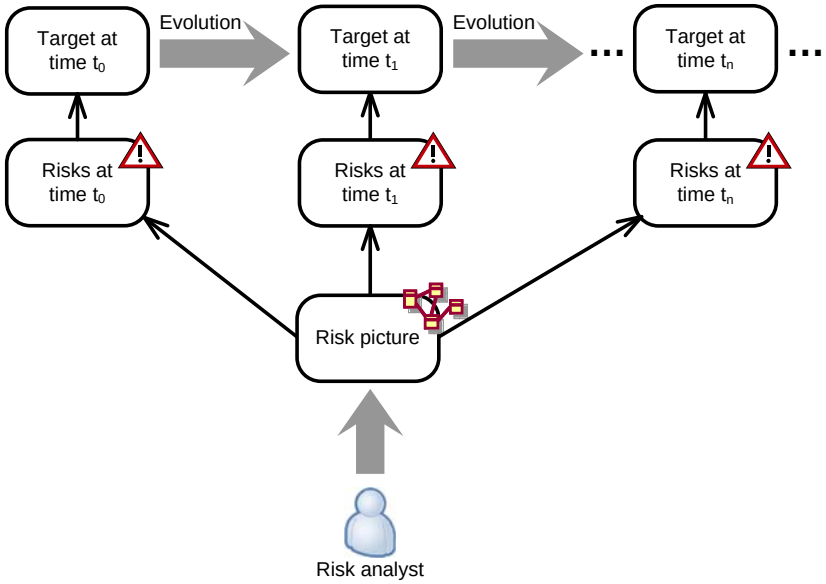

Fig. 4. The continuous evolution perspective

The illustration in Fig. 4 shows the principles by which the risk analysts conduct the analysis from the *continuous evolution* perspective. Assuming that we have available a description of the target of analysis as a function of time, such that we can derive the target at any given point in time, we use this as input to the risk analysis. Understanding how the target and its environment evolve, we seek to work out a risk picture as a function of time that shows how risks evolve.

Methodologically, the main challenges are how to identify evolving risks and present them in a dynamic risk picture. This requires us to generalise the target description to capture evolution and to identify and likewise generalise the risks affected by evolution.

3 Our Approach

As explained in Sect. 2 this paper focuses on the *before-after* perspective. Confining to this perspective, we present in this section our approach to risk analysis of changing systems by adequately generalising the risk management process of ISO 31000. Our main concern is the five activities in the middle of Fig. 1 as this comprise the actual risk analysis activities. We moreover focus on what is needed for handling change, referring to the ISO 31000 standard [15] for the principles and guidelines for risk management in the traditional setting. The reader is also referred to the practical example in Sect. 6 for further details about the various activities and tasks.

In the following we go through each of the five activities of the risk analysis process. Subsequently, in Sect. 4, we introduce language support for this process by a generalisation of risk graphs to a risk graph notation with the expressiveness to explicitly model changing risks in the *before-after* perspective.

3.1 Context Establishment

The activity of establishing the context involves making the premises for the subsequent risk analysis. Establishing the context should result in a target description, which is the documentation of all the information that serves as the input to and basis for the risk analysis, and a set of risk evaluation criteria. Any information about the target of analysis that is relevant for the risk analysis and its outcome needs to be included in the target description. This means that any risk relevant change in the target of analysis must be reflected by changes to the target description.

Establishing the context of the analysis also includes articulating the goals and objectives of the analysis and deciding its focus and scope. In particular, we need to determine precisely what the target of analysis is and what the assets that need to be protected are. The risk analysis is conducted with respect to the identified assets, and it is only by precisely understanding what the assets are that we can conduct a risk analysis that meets the overall goals and objectives of the analysis.

Establishing the Target Description. The target description includes the documentation of the target of analysis, the focus and scope of the analysis, the environment of the target, the assumptions of the analysis, the parties and assets of the analysis, and the context of the analysis. The UML [22] class diagram of Fig. 5 gives an overview of the elements of a target description.

In a risk analysis, the notions of party, asset and risk are closely related. A *party* is an organisation, company, person, group or other body on whose behalf

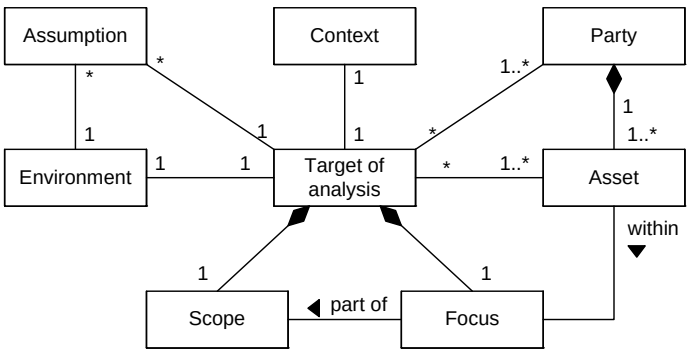

Fig. 5. Target description

the risk analysis is conducted. An *asset* is something to which a party assigns value and hence for which the party requires protection. A *risk* is the likelihood of an unwanted incident and its consequence for a specific asset. This means that if there is no party, it makes no sense to speak of assets. And without assets there can moreover be no risks.

The *target of the analysis* is the system, organisation, enterprise, or the like that is the subject of a risk analysis. The *focus of the analysis* is the main issue or central area of attention. The focus is within the *scope of the analysis*, which the extent or range of the analysis. The scope defines the border of the analysis, i.e. what is held inside and what is held outside of the analysis. The *environment* of the target is the surrounding things of relevance that may affect or interact with the target; in the most general case the environment is the rest of the world. The *assumptions* are something we take as granted or accept as true, although it may not be so; the results of a risk analysis are valid only under the assumptions. The *context* of the analysis is the premises for and the background of the analysis. This includes the purposes of the analysis and to whom the analysis is addressed.

For a given target of analysis we assume that we can use a traditional risk analysis method to conduct the context establishment while not taking into account changes. The additional task of establishing the changes in the context includes making a description of the target of analysis when the changes have been taken into account. This extended target description should include both a description of the changes and the result of implementing them, although it should be possible to deduce the latter from the current target description and the description of the changes. Changes that concern the target of analysis can be new or different work processes, the introduction of new services or applications, changes in users or roles, etc. These may imply changes the risk picture, and therefore imply the need for a new risk analysis of parts of the target. There may, however, also be changes in parties, changes in assets or asset priorities, changes in the environment or in the assumptions, changes in the focus or scope, and so on, that must be documented in the description of the changing target of analysis.

Establishing the Risk Evaluation Criteria. The risk evaluation criteria are a specification of the risk levels that the parties of the risk analysis are willing to accept. The criteria will later be used to evaluate the significance of risk, and should reflect the values, objectives and resources of the parties in question.

When we are deciding the risk evaluation criteria we need to take into account not only the views of the parties, but also the nature of the assets, the types of consequences and how they should me measured and described. Specifically, we need for each asset to define a consequence scale where each consequence value describes a level of impact of an unwanted incident on an asset in terms of harm or reduced asset value. We furthermore need to define a likelihood scale, the values of which will be used to describe the frequency or probability of unwanted incidents and threat scenarios to occur.

Recall that a risk is the likelihood of an unwanted incident and its consequence for a specific asset. The risk level is the level or value of a risk as derived from its likelihood and consequence. The risk level of each combination of a likelihood and a consequence is calculated by a risk function. Since it is only the party of a given asset that can determine the severity of a risk, it is the party that must determine an adequate risk function. Essentially, the risk evaluation criteria are a specification of the level at which risks become unacceptable.

Some changes are of a kind that does not affect the assets or other values, objectives or resources of the parties. In this case, there is no need to reconsider the risk evaluation criteria. For other changes, the value or priorities of assets may change, new assets may arise, the parties may become more or less risk averse, and so forth. In that case we need a new iteration on establishing and documenting the risk evaluation criteria.

3.2 Risk Identification

The risk identification involves identifying and documenting unwanted incidents with respect to the identified assets, as well as identifying vulnerabilities and sources of risk.

The risk identification should involve people with appropriate expert knowledge about the target of analysis. The activity of extracting the relevant information relies on techniques and tools for identifying risk relevant information, for structuring the information in a sensible way, and for adequately documenting the information. While the documentation of the risks that are identified should serve as a means for reporting the findings to the relevant stakeholders, it should at the same time facilitate the subsequent estimation and evaluation. In this section we focus on the methodological guidelines for risk identification of changing and evolving systems. In Sect. 4 we discuss more closely the required risk modelling techniques.

As mentioned above, it is the target description that serves as the input to and basis for the subsequent risk analysis. The objective of the risk identification is to identify and document the changing risks given the description of the changing target. A main principle is that to the extent that we have identified and documented the risks for the target of analysis before changes, we only address the parts of the target that are affected by the change when identifying the changing risks.

This means that when considering the target description without the changes, the risk identification and the risk documentation are conducted according to traditional risk analysis methods. When this is completed we need to update the resulting risk documentation according to the changes. This is conducted by making a walkthrough of the current target description and risk documentation. The risks that are persistent under change can immediately be included in the new documentation with no further investigation. The risks that may be affected by change need to be considered again: Previous scenarios, incidents, etc. may change, new may arise, and others may disappear.

The methodological guidelines for risk identification of a changing target of analysis are summarised as follows:

1. Identify and document risks by using as input the target description before changes have been taken into account.
2. Establish and document the traceability between the target description before change and the risk documentation resulting from the previous step.
3. Based on the traceability and the description of the changed target, identify the parts of the risk documentation that are persistent under change.
4. Conduct the risk identification of the changed target only with respect to the parts of the target and the risks that are affected by the change.

In order to conduct the activity of risk identification of a changing target of analysis, there is thus a need for techniques for identification and modelling of risks that change, as well as techniques for establishing and modelling traceability between target description and risk models, both which are topics of Sect. 4.

3.3 Risk Estimation

The objective of the risk estimation is to establish an understanding of the severity of identified risks, and to provide the basis for the subsequent risk evaluation and risk treatment. By considering the causes and sources of risk, the risk estimation amounts to estimating and documenting the likelihoods and consequences of the identified unwanted incidents. It is the likelihoods of unwanted incidents and their consequences for assets that constitute risks, and by making estimates of the likelihoods and consequences we can understand which risks are the most in need of treatment and which risks are less relevant.

Given the documentation of identified risks from the previous activity, including the documentation of the changing risks, the risk estimation of a changing and evolving target is quite similar to traditional risk analysis: The estimation is conducted by a walkthrough of the risk documentation addressing each of the relevant elements in turn. To the extent that risks are persistent under the changes, the estimation is not repeated.

The estimates need to be continuously documented, which means that there must be adequate support for including the estimates in the risk models. In order to conduct the activity of risk analysis and the documentation of the results, there is hence a need for techniques for making estimates of likelihoods and consequences of changing risks, as well as modelling support for documenting the results. As we shall see, such techniques and support is provided by the approach to the modelling of changing risk presented in Sect. 4.

3.4 Risk Evaluation

The objective of the risk evaluation is to determine which of the identified risks that need treatment, and to make a basis for prioritising the treatment options. Basically, the risk evaluation amounts to estimating the risk levels based on the

likelihood and consequence estimates, and to compare the results with the risk evaluation criteria.

The risk evaluation of a changing and evolving target is conducted in the same way as risk evaluation of a traditional risk analysis. Given the risk documentation of the changing risks with the risk estimates, the risk evaluation is conducted by calculating the risk level of each pair of an unwanted incident and asset harmed by the incident. The calculation is straightforwardly done by using the risk function defined during the context establishment. For changing systems, the criteria may of course be different before and after the changes in question.

3.5 Risk Treatment

The risk treatment succeeds the risk identification, estimation and evaluation activities, and the objective is to identify and select a set of treatment options for the risks that are not acceptable according to the risk evaluation criteria. A treatment option is thus an appropriate measure to reduce risk level, and the implementation of the selected treatments should bring the risk level down to an acceptable level.

For changes that are planned or predicted, it may be that we are only concerned about the future risks and in ensuring that implementation of the changes results in a system with an acceptable risk level. For changes that are planned to be implemented over a longer time span it may, however, be that we also need to identify treatments for current unacceptable risks and treatments that are consecutively implemented in the future in order to maintain an acceptable risk level under the period of change.

4 Foundations

Risk analysis involves the process of understanding the nature of risks and determining the level of risk [15]. Risk modelling refers to techniques that are used to aid the process of identifying, documenting and estimating likelihoods and consequences of unwanted incidents. In this section we present an approach to risk modelling referred to as risk graphs [3]. Once we have introduced the syntax and semantics of risk graphs, we generalise these to enable modelling and reasoning about changing risks. We then provide a calculus for analysing likelihoods in risk graphs and a means for relating risk graphs to target models.

4.1 Risk Graphs

A risk model is a structured way of representing an unwanted incident and its causes and consequences by means of graphs, trees or block diagrams [24]. We introduce risk graphs as an aid for structuring events and scenarios leading to incidents and estimating likelihoods of incidents. There exist several modelling techniques that can be used for such structuring of scenarios and incidents, and for the reasoning about likelihoods of incidents, for example fault trees [13],

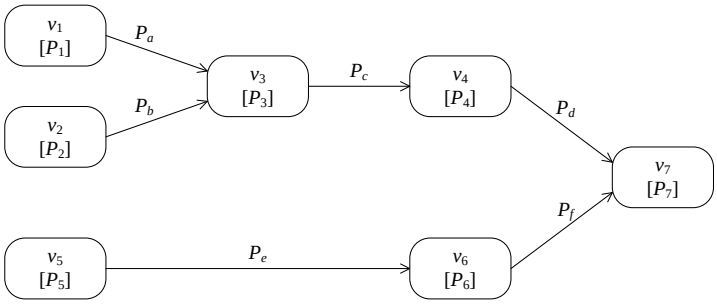

Fig. 6. Risk graph

event trees [14], attack trees [25], cause-consequence diagrams [21], Bayesian networks [2] and CORAS threat diagrams (see Sect. 5). Risk graphs can be understood as a common abstraction of these modelling techniques [3]. By giving formal semantics to risk graphs, we thereby also provide a risk model semantics that can be used to explain and reason about several established approaches to risk modelling.

A risk graph consists of vertices (representing threat scenarios) and a finite set of directed relations (representing the "leads to" relationship) between them. An example risk graph is shown in Fig. 6. Each vertex in a risk graph is assigned a set of likelihood values representing the estimated likelihood for the scenario to occur. The assignment of several likelihood values, typically a likelihood interval, represents underspecification of the likelihood estimate. A relation from vertex v to vertex v' means that v may lead to v'. Also the relations can be assigned likelihood sets. These are conditional likelihoods that specify the likelihood for a scenario leading to another scenario when the former occurs. One threat scenario may lead to several other threat scenarios, so the probabilities on the relations leading from a threat scenario may add up to more than 1. A risk graph is furthermore allowed to be incomplete in the sense that a given threat scenario may lead to more scenarios than what is accounted for in the risk graph. The probabilities of the relations leading from a threat scenario may for this reason also add up to less than 1.

The Syntax of Risk Graphs. Formally a risk graph is a set D of elements e. An element is a vertex v or a relation $v \rightarrow v'$. Let $P \subseteq [0,1]$ denote a probability set. We then write $v(P)$ to indicate that the probability set P is assigned to v. Similarly, we write $v \xrightarrow{P} v'$ to indicate that the probability set P is assigned to the relation from v to v'. If no probability set is explicitly assigned, we assume the probability set assigned to the element to be $[0,1]$. Using this textual notation, the risk graph shown in Fig. 6 can be represented by

$$D = \{v_1(P_1), v_2(P_2), v_3(P_3), v_4(P_4), v_5(P_5), v_6(P_6), v_7(P_7),$$
$$v_1 \xrightarrow{P_a} v_3, v_2 \xrightarrow{P_b} v_3, v_3 \xrightarrow{P_c} v_4, v_4 \xrightarrow{P_d} v_7, v_5 \xrightarrow{P_e} v_6, v_6 \xrightarrow{P_f} v_7\}$$

The Semantics of Risk Graphs. Risk graphs are used for the purpose of documenting and reasoning about risks, particularly the documentation and analysis of threat scenarios and unwanted incidents and their likelihoods. The approach of [3] assumes that scenarios and their probabilities are represented by a probability space [4] on traces of events. We let $\mathcal{H}$ denote the set of all traces (both finite and infinite) and $\mathcal{H}_{\mathbb{N}}$ the set of all finite traces. A probability space is a triple $(\mathcal{H}, \mathcal{F}, \mu)$. $\mathcal{H}$ is the sample space, i.e. the set of possible outcomes, which in our case is the set of all traces. $\mathcal{F}$ is the set of measurable subsets of the sample space, and μ is a measure that assigns a probability to each element in $\mathcal{F}$. The semantics of a risk graph is statements about the probabilities of the trace sets that represent vertices or the composition of vertices. In other words, the semantics is a set of statements about the measure μ.

For composition of vertices, $v \sqcap\!| \, v'$ denotes the occurrence of both v and v' where the former occurs before the latter. We let $v \sqcup v'$ denote the occurrence of at least one of v and v'. A vertex is atomic if it is not of the form $v \sqcap\!| \, v'$ or $v \sqcup v'$. We use lower case v as the naming convention for arbitrary vertices, and upper case V as the naming convention for the set of finite traces representing the vertex v.

When defining the semantics of risk graphs we use the auxiliary function $tr(\_)$ that yields a set of finite traces from an atomic or combined vertex. Intuitively, $tr(v)$ is the set of all possible traces leading up to and through the vertex v, without continuing further. The function is defined by

$$\begin{aligned} tr(v) &\stackrel{\text{def}}{=} \mathcal{H}_{\mathbb{N}} \gtrsim V \text{ when } v \text{ is an atomic vertex} \\ tr(v \sqcap\!| \, v') &\stackrel{\text{def}}{=} tr(v) \gtrsim tr(v') \\ tr(v \sqcup v') &\stackrel{\text{def}}{=} tr(v) \cup tr(v') \end{aligned}$$

where $\gtrsim$ is the operator for sequential composition of trace sets, for example weak sequencing in UML sequence diagrams [10]. Notice that the definition of the composition $v \sqcap\!| \, v'$ does not require v to occur immediately before v'. The definition implies that $tr(v \sqcap\!| \, v')$ includes traces from v to v' via finite detours.

A probability interval P assigned to v, denoted $v(P)$, means that the likelihood of going through v is a value $p \in P$, independent of what happens before or after v. The semantics of a vertex is defined by

$$[\![v(P)]\!] \stackrel{\text{def}}{=} \mu_c(tr(v)) \in P$$

where the expression $\mu_c(S)$ denotes the probability of any continuation of the trace set $S \subseteq \mathcal{H}$, and is defined as

$$\mu_c \stackrel{\text{def}}{=} \mu(S \gtrsim \mathcal{H})$$

A probability interval P assigned to a relation $v \rightarrow v'$ means that the likelihood of v' occurring after an occurrence of v is a value in P. This likelihood is referred to as the conditional likelihood. The semantics of a relation is defined by

$$[\![v \xrightarrow{P} v']\!] \stackrel{\text{def}}{=} \mu_c(tr(v \sqcap\!| \, v')) \in \mu_c(tr(v)) \cdot P$$

$$
\begin{aligned}
[p_i, p_j] \cdot [p_k, p_l] &\stackrel{\text{def}}{=} [p_i \cdot p_k, p_j \cdot p_l] \\
[p_i, p_j] \cdot p &\stackrel{\text{def}}{=} [p_i \cdot p, p_j \cdot p] \\
[p_i, p_j] + [p_k, p_l] &\stackrel{\text{def}}{=} [p_i + p_k, p_j + p_l] \\
[p_i, p_j] + p &\stackrel{\text{def}}{=} [p_i + p, p_j + p] \\
[p_i, p_j] - [p_k, p_l] &\stackrel{\text{def}}{=} [p_i - p_k, p_j - p_l] \\
[p_i, p_j] - p &\stackrel{\text{def}}{=} [p_i - p, p_j - p] \\
p - [p_k, p_l] &\stackrel{\text{def}}{=} [p - p_k, p - p_l]
\end{aligned}
$$

Fig. 7. Interval arithmetic

Our definitions of interval arithmetic in the setting of risk graphs are given in Fig. 7.

The semantics $[\![D]\!]$ of a risk graph is the conjunction of the expressions defined by the elements in D, formally defined as

$$[\![D]\!] \stackrel{\text{def}}{=} \bigwedge_{e \in D} [\![e]\!]$$

A risk graph is said to be correct (with respect to the world or a specification of the relevant part of the world) if each of the conjuncts of $[\![D]\!]$ is true. We say that D is inconsistent if it is possible to deduce *False* from $[\![D]\!]$.

4.2 Risk Graphs for Changing Risks

In order to support the modelling of changing risks we need to generalise risk graphs to allow the simultaneous modelling of risks both before and after the implementation of some given changes. For this purpose we extend the risk graph notation to three kinds of vertices and three kinds of relations, namely *before*, *after* and *before-after*. When an element (vertex or relation) is of kind *before* it represents risk information before the changes, when it is of kind *after* it represents risk information after the changes, and when it is of kind *before-after* it represents risk information that holds both before and after the changes.

The Syntax of Risk Graphs with Change. A risk graph with change is represented by a pair (D_b, D_a) of sets of elements, the former consisting of the vertices and relations of kind *before* and the latter consisting of vertices and relations of kind *after*. Table 1 gives an overview of the language constructs and the naming conventions we use for referring to them. The symbols written in sans serif and the arrows denote specific language constructs, whereas v denotes an arbitrary vertex of any kind. Table 2 gives an overview of the various ways of specifying likelihoods. Recall that any of the likelihoods can be undefined, in which case they are completely underspecified.

Figure 8 shows an example of the visual representation of a risk graph with change. Solid lines, like on the vertex v_1 and the relation from v_1 to v_3, indicates

Table 1. Naming conventions

Variable	Diagram construct
v	Vertex *before-after*
vb	Vertex *before*
va	Vertex *after*
v	Vertex
$v \rightarrow v'$	Relation *before-after*
$v \rightarrow_{\mathsf{b}} v'$	Relation *before*
$v \rightarrow_{\mathsf{a}} v'$	Relation *after*

Table 2. Denoting likelihoods

Likelihood spec.	Interpretation
$\mathsf{v}(P\ P')$	v occurs with likelihood P *before*, and v occurs with likelihood P' *after*
$\mathsf{vb}(P)$	vb occurs with likelihood P *before*
$\mathsf{va}(P)$	va occurs with likelihood P *after*
$v \xrightarrow{P\ P'} v'$	v leads to v' with conditional likelihood P *before*, and v leads to v' with conditional likelihood P' *after*
$v \xrightarrow{P}_{\mathsf{b}} v'$	v leads to v' with conditional likelihood P *before*
$v \xrightarrow{P}_{\mathsf{a}} v'$	v leads to v' with conditional likelihood P *after*

elements that only exists *before*, while dashed lines indicates elements that exists *after*. The vertices with a white shadow, like v_2, are those that exist both *before* and *after*, while those with black shadows, like v_5, exist only *after*. The dashed relations with a single probability set, like the relation from v_5 to v_6, exist only *after*, while those with double probability sets, like the relation from v_3 to v_4, exist both *before* and *after*.

Since we are operating with vertices and relations of kind *before-after* as language element of their own, we also allow the representation of risk graphs

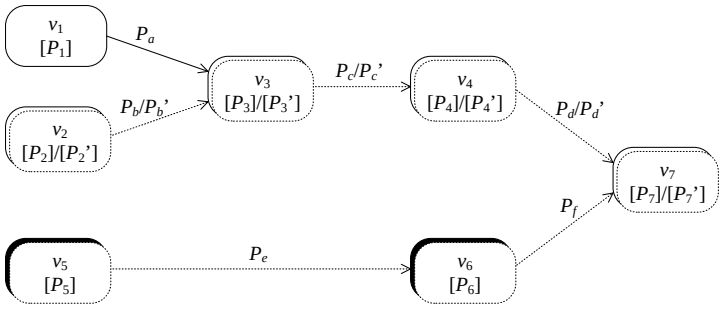

Fig. 8. Risk graph for changing risks

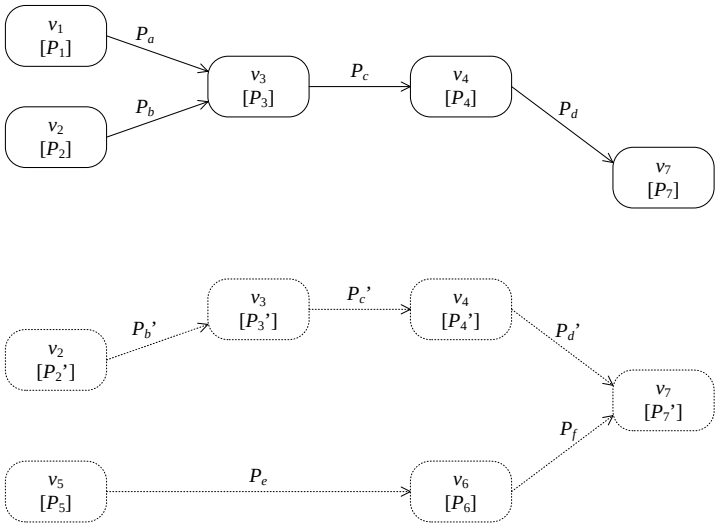

Fig. 9. Two views on risk graphs for changing risk

with change as a single set D of vertices and relations, where each element is of one of the kinds *before*, *after* or *before-after*. This single set of elements is then syntactic sugar for the equivalent representation of a pair of sets of elements. For such a combined representation D we use the functions *before*(_) and *after*(_) to filter the combined risk graph with respect to the elements of kind *before* and *after*, respectively. The following define the function *before*(_) for singleton sets of elements.

$$\begin{aligned}
\mathit{before}(\{\mathsf{v}(P\ P')\}) &\stackrel{\text{def}}{=} \{\mathsf{vb}(P)\}\\
\mathit{before}(\{\mathsf{vb}(P)\}) &\stackrel{\text{def}}{=} \{\mathsf{vb}(P)\}\\
\mathit{before}(\{\mathsf{va}(P)\}) &\stackrel{\text{def}}{=} \emptyset\\
\mathit{before}(\{v \xrightarrow{P\ P'} v'\}) &\stackrel{\text{def}}{=} \{v \xrightarrow{P}_{\mathsf{b}} v'\}\\
\mathit{before}(\{v \xrightarrow{P}_{\mathsf{b}} v'\}) &\stackrel{\text{def}}{=} \{v \xrightarrow{P}_{\mathsf{b}} v'\}\\
\mathit{before}(\{v \xrightarrow{P}_{\mathsf{a}} v'\}) &\stackrel{\text{def}}{=} \emptyset
\end{aligned}$$

The filtering of a risk graph with change D with respect to the *before* elements is then defined as

$$\mathit{before}(D) \stackrel{\text{def}}{=} \bigcup_{e \in D} \mathit{before}(\{e\})$$

The definition of the function *after*(_) is symmetric. For a risk graph with change D of elements of the three different kinds, the representation as a pair of elements of kind *before* and elements of kind *after* is then given by (*before*(D), *after*(D)). Figure 9 shows the graphical representation of the two risk graph *before*(D) (top) and *after*(D) (bottom) where D is the risk graph shown in Fig. 8.

The Semantics of Risk Graphs with Change. Given the syntax of risk graphs with change as defined above, we can define the semantics as a straightforward generalisation of the semantics of regular risk graphs. The semantics $[\![(D_b, D_a)]\!]$ of a risk graph with change is defined as

$$[\![(D_b, D_a)]\!] \stackrel{\text{def}}{=} [\![D_b]\!] \wedge [\![D_a]\!]$$

For a combined representation D of a risk graph with change, the semantics is defined as

$$[\![D]\!] \stackrel{\text{def}}{=} [\![\mathit{before}(D), \mathit{after}(D)]\!]$$

4.3 Reasoning about Likelihoods in Risk Graphs

In this section we introduce rules for calculating probabilities of vertices in risk graphs, and we provide guidelines for consistency checking probabilities that are assigned to risk graphs.

The first rule is referred to as the *relation rule*, and captures the conditional likelihood semantics of a risk graph relation. For a vertex v that leads to v', the vertex $v \sqcap\!| v'$ denotes the occurrences of v' that happen after an occurrence of v.

Rule 1 (Relation). *If there is a direct relation from v to v', we have:*

$$\frac{v(P) \quad v \xrightarrow{P'} v'}{(v \sqcap\!| v')(P \cdot P')}$$

The second rule is referred to as the *mutual exclusive vertices rule*, and yields the probability of either v or v' occurring when the two vertices are mutually exclusive:

Rule 2 (Mutually exclusive vertices). *If the vertices v and v' are mutually exclusive, we have:*

$$\frac{v(P) \quad v'(P')}{(v \sqcup v')(P + P')}$$

The third rule is referred to as the *statistically independent vertices rule*, and yields the probability of either v or v' occurring when the two vertices are statistically independent.

Rule 3 (Statistically independent vertices). *If vertices v and v' are statistically independent, we have:*

$$\frac{v(P) \quad v'(P')}{(v \sqcup v')(P + P' - P \cdot P')}$$

Table 3. Guidelines for consistency checking likelihoods

How to check consistency of likelihoods in risk graphs
Exact values in complete diagrams
Assigned value: $v(p)$
Calculated value: $v(p')$
Consistency check: $p = p'$
Exact values in incomplete diagrams
Assigned value: $v(p)$
Calculated value: $v(p')$
Consistency check: $p \geq p'$
Intervals in complete diagrams
Assigned interval: $v([p_i, p_j])$
Calculated interval: $v([p'_i, p'_j])$
Consistency check: $[p'_i, p'_j] \subseteq [p_i, p_j]$ or, equivalently, $p_i \leq p'_i$ and $p_j \geq p'_j$
Intervals in incomplete diagrams
Assigned interval: $v([p_i, p_j])$
Calculated interval: $v([p'_i, p'_j])$
Consistency check: $p_j \geq p'_j$

As a small example of probability calculation consider the risk graph in Fig. 6 and assume we want to calculate the probability of v_3 from v_1 and v_2. By Rule 1 we calculate $(v_1 \sqcap\!| \, v_3)(P_1 \cdot P_a)$ and $(v_2 \sqcap\!| \, v_3)(P_2 \cdot P_b)$. Assuming that v_1 and v_2, as well as $v_1 \sqcap\!| \, v_3$ and $v_2 \sqcap\!| \, v_3$, are statistically independent, we use Rule 3 to calculate $((v_1 \sqcap\!| \, v_3) \sqcup (v_2 \sqcap\!| \, v_3))(P_1 \cdot P_a + P_2 \cdot P_b - P_1 \cdot P_a \cdot P_2 \cdot P_b)$.

Assuming that the likelihood estimates in Fig. 6 are correct, there is still one issue to consider before we can conclude about the likelihood of the vertex v_3. The issue is whether or not the risk graph is complete. If the risk graph is complete, the graph shows all the possible ways in which v_3 may occur. In that case we have that $v_3 = (v_1 \sqcap\!| \, v_3) \sqcup (v_2 \sqcap\!| \, v_3)$ and that $P'_3 = P_1 \cdot P_a + P_2 \cdot P_b - P_1 \cdot P_a \cdot P_2 \cdot P_b$ is the correct likelihood of this vertex. If the risk graph is incomplete, there may be further scenarios that can lead to v_3. In that case we only know that P'_3 is the lower bound of the probability of v_3.

Consistency checking of risk models is important, as it is a useful means for detecting errors or misunderstandings of the risk estimates that are documented during a risk analysis. The basis for the consistency checking is the likelihood values that are already assigned to the vertices and relations of a risk graph. The guidelines for consistency checking depend on whether the risk graph in question is complete, and whether the likelihoods are given as exact probabilities or as probability intervals. The guidelines are given in Table 3.

As an example of consistency checking, consider the risk graph in Fig. 6, assuming first that the graph is complete. By the above example, we know that the probability of the vertex v_3 is $P'_3 = P_1 \cdot P_a + P_2 \cdot P_b - P_1 \cdot P_a \cdot P_2 \cdot P_b$ given the vertices and relations that lead to this vertex. The assigned probability P_3 must therefore equal the calculated probability P'_3 in order to be consistent with the

preceding probability estimates if we are working with exact values. If the P's are intervals, we must have that $P_3' \subseteq P_3$ for the risk graph to be consistent.

Discarding the assumption of the completeness of the graph gives the consistency requirement that the assigned probability P_3 must be greater than or equal to the calculated probability P_3', i.e. that $P_3 \geq P_3'$, if we have exact values. On the other hand, if P_3 and P_3' are intervals $P_3 = [p_i, p_j]$ and $P_3' = [p_i', p_j']$, the requirement is that $p_j \geq p_j'$.

4.4 Relating Risk Model to Target Description

Risk analysis of changing systems requires means for identifying the parts of a risk picture that are affected by changes to a specific part of the target (and therefore need to be reassessed), as well as identifying the parts of the risk picture that are not affected (and therefore valid also after the changes). Thus we need techniques for identifying and documenting the relation between the target description and the risk models in a way that gives us traceability between target elements (elements of the target description) and risk model elements.

Two key artifacts in a risk analysis are the target model and the risk model. The target model is the core part of the overall target description and documents the events, scenarios and actors that are the subject for the risk analysis. Given these two artifacts we introduce a third artifact in order to establish the relation between the former two: a trace model.

The trace model is of a table format that allows the tracing from target model elements to risk model elements, and vice versa. Initially we can think of the trace model as a set of pairs (u_{id}, v_{id}) of target model identifiers u_{id} and risk model identifiers v_{id} representing the rows of the table. This of course require that each of the elements have a unique identifier. In the following we assume that we already have a target model and a risk model of elements with unique identifiers, since obtaining such models by indexing the elements is a trivial task.

From a pragmatic point of view, there are two obvious shortcomings of the table format given above. To make efficient use of the trace model it should convey information about the relations in an intuitive way; the use of possibly tool generated indexes for the model elements is not intuitively informative. Furthermore, in many cases several target model elements are logically understood as a whole. Without some means of grouping several rows of the table into one compound relation, such structures of the target model will be obscured.

To mitigate this we introduce a third column in the table for tagging the target model element/risk model element pairs. The grouping of pairs is then conducted by inserting the same tag on several rows. The name of the tag should be chosen by the end-user, and should be a unique name that conveys intuitive information about the grouping. More formally, the trace model is now a set of tuples (u_{id}, v_{id}, t) of a target model identifier, a risk model identifier and a tag.

We extend the risk graph notation with a language construct for explicitly specifying the relation to the target model. The construct is used for annotating risk graphs with the tags of the trace model. We understand this construct as a mere visualisation of the trace model in the risk graphs, and not as part of

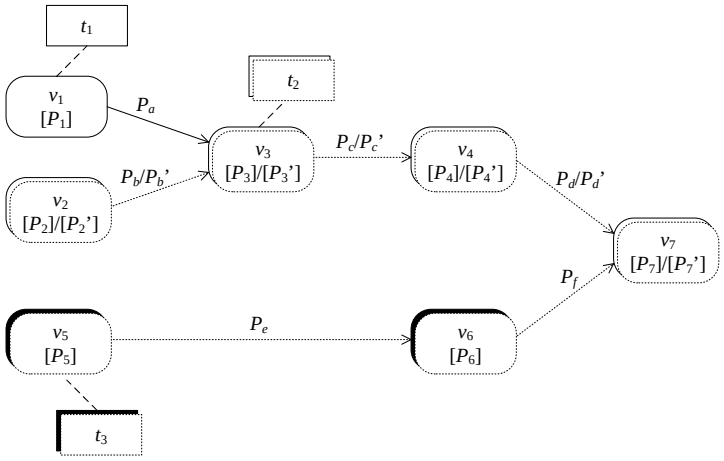

Fig. 10. Relating target model and risk model

the semantics of risk graphs. An example of the visualisation of a trace model in a risk graph with change is shown in Fig. 10. As with the other elements of risk models with change, the target model relations can be specified as existing *before* the change only, (for example t_1), *after* the change only (for example t_3), or both *before* and *after* the change (for example t_2).

5 Instantiation of CORAS

In a CORAS risk analysis, threat diagrams are used intensively to facilitate risk identification and risk estimation. The diagrams are furthermore used as a part of the documentation and reporting of the analysis results. Figure 11 depicts an example of a threat diagram. In fact, this threat diagram shows the scenarios that are modelled by means of the risk graph in Fig. 10. Threat diagrams describe how threats may exploit vulnerabilities to initiate threat scenarios, how threat scenarios may lead to unwanted incidents or other threat scenarios, and the assets harmed by the unwanted incidents. The language constructs are threats (deliberate, accidental and non-human), vulnerabilities, threat scenarios, unwanted incidents and assets. Only threat scenarios and unwanted incidents may be assigned likelihoods.

There are furthermore three kinds of relations in threat diagrams, namely initiates relations, leads-to relations and impacts relations. An initiates relation has a threat as source and a threat scenario or unwanted incidents as target. It can be annotated with a likelihood that describes the likelihood for the threat to initiate the related scenario or incident. A leads-to relation has a threat scenario or unwanted incident as both source and target. It can be annotated with a conditional likelihood. An impacts relation has an unwanted incident as source

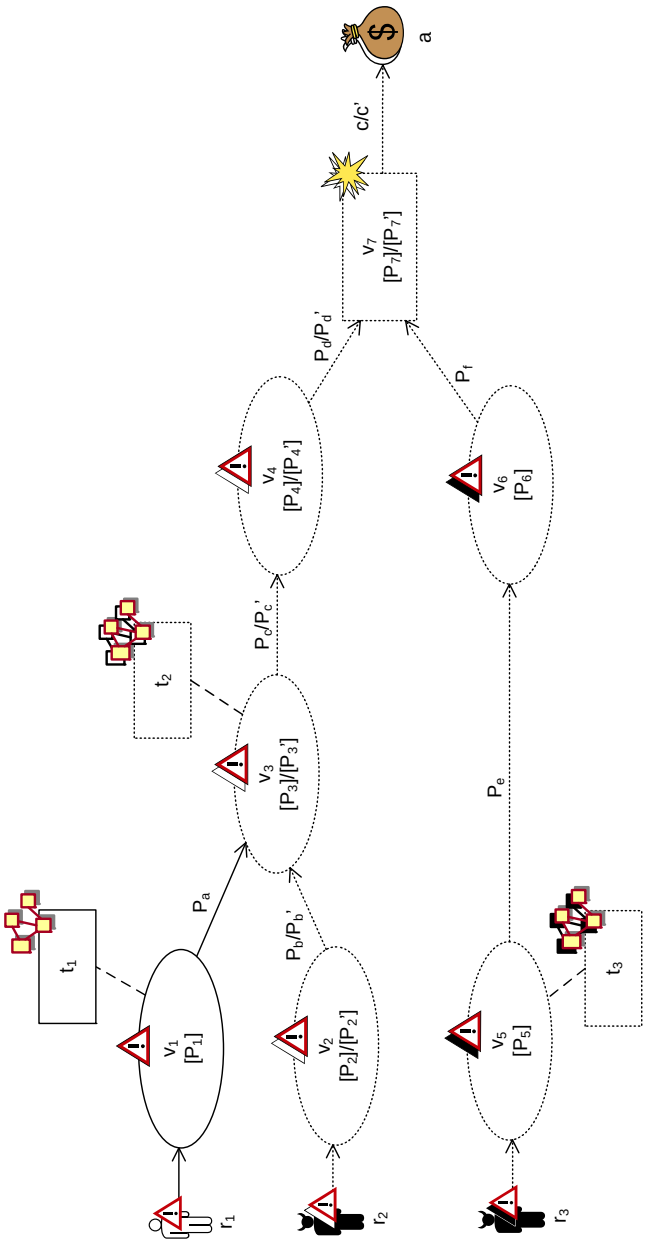

Fig. 11. Instantiation of risk graphs in CORAS

and an asset as target, and can be annotated with a consequence value that describes the harm of the incident on the asset when the incident occurs.

While all scenarios and relations in Fig. 10 are present in Fig. 11, there are some significant differences between the two diagrams. The threat diagram explicitly shows the initiating threats, distinguishes v_7 from the other scenarios as an unwanted incident, and explicitly shows the asset that is harmed.

The differences between threat diagrams and risk graphs are summarised as follows:

- Initiate relations and leads-to relations in threat diagrams can be annotated with vulnerabilities, while the relations in risk graphs cannot.
- Threat diagrams distinguish between four kinds of vertices, namely threats, threat scenarios, unwanted incidents and assets, while risk graphs only have scenarios.
- Threat diagrams distinguish between three kinds of relations, namely initiates relations, leads-to relations and impacts relations, while risk graphs only have leads-to relations.

Given the differences between threat diagrams and risk graphs, the techniques for reasoning about likelihoods nevertheless carry over to the CORAS instantiation. The vulnerabilities are mere annotations on relations, and can be ignored in the formal representation of the diagrams. Moreover, the various vertices and relations of threat diagrams can be interpreted as special instances of the risk graph vertex and relation:

- An unwanted incident of a threat diagram is interpreted as a scenario of a risk graph.
- A set of threats $r_1, \ldots, r_n$ with initiates relations to the same threat scenario s is interpreted as follows: The threat scenario s is decomposed into n parts, where each resulting sub-scenario $s_j, j \in \{1, \ldots, n\}$, corresponds to the part of s that is initiated by threat r_j. The threat r_j with initiates relation of likelihood l_j to sub-scenario s_j is then combined into the risk graph scenario *Threat* r_j *initiates* s_j and the scenario is assigned likelihood l_j.
- An impacts relation from unwanted incident u to asset a with consequence c in a threat diagram is interpreted as follows: The impacts relation is interpreted as a risk graph relation with likelihood 1; the asset a is interpreted as the risk graph scenario *Incident* u *harms asset* a *with consequence* c.

With this interpretation, we refer to Sect. 4.3 for the techniques for reasoning about likelihoods in CORAS threat diagrams. However, notice that Rule 1 (Relation) applies to the CORAS leads-to relations only and that Rule 2 (Mutually exclusive vertices) and Rule 3 (Independent vertices) apply to the CORAS threat scenarios and unwanted incidents. In order to allow all likelihood reasoning to be conducted directly in CORAS diagrams, we introduce a separate rule for the initiates relation. We let r denote a threat, v denote a vertex (threat scenario or unwanted incident), $r \rightarrow v$ denote the initiates relation from threat r to vertex v, and $r \sqcap\!| \, v$ denote the occurrences of vertex v that are initiated by the threat r.

Table 4. List of acronyms

Acronym	Meaning
ACC	Area Control Center
ADS-B	Automatic Dependent Surveillance-Broadcast
AMAN	Arrival Manager
AOIS	Aeronautical Operational Information System
ATCO	Air Traffic Controller
ATM	Air Traffic Management
ATS	Air Traffic System
CWP	Controller Working Position
FDPS	Flight Data Processing System
OPS Room	Operation room
SUP	Supervisor

Rule 4 (Initiates). *If there is an initiates relation from threat r to vertex v, we have:*

$$\frac{r \xrightarrow{P} v}{(r \sqcap\!| v)(P)}$$

6 Practical Example

In this section we present a practical example of a risk analysis of a changing system under the *before-after* perspective. As the running example we use an Air Traffic Management (ATM) risk analysis case study conducted within the SecureChange project.[1]

European Air Traffic Management is currently undergoing huge changes with introduction of new information systems and decision support systems, as well as the reorganisation of ATM services, as part of the ATM 2000+ strategic agenda [6] and the Single European Sky ATM Research (SESAR) [26]. The case study presented in the following focuses in particular on the introduction of two new systems: Arrival Manager (AMAN) and Automatic Dependent Surveillance-Broadcast (ADS-B). The results of the case study are fully documented in a SecureChange deliverable [17]; in this paper only a selection of the target and the analysis is presented. A list of acronyms from the ATM domain used in the example is given in Table 4.

6.1 Context Establishment

The context establishment includes articulating the overall goals and objectives of the risk analysis, and deciding its focus and scope. This includes making a

[1] The case study was conducted in close interaction with ATM personnel with expert knowledge about the target of analysis and the planned changes. For more information on the SecureChange project, see `http://www.securechange.eu/`.

description of the target of analysis, identifying the assets and deciding the risk evaluation criteria.

Goals and Objectives. An important part of Air Traffic Management is the services provided by ground-based Air Traffic Controllers (ATCOs) located at Area Control Centers (ACCs). One of the main responsibilities of Air Traffic Controllers is to maintain horizontal and vertical separation among aircrafts and between aircrafts and possible obstacles. They must ensure an orderly and expeditious air traffic flow by issuing instructions and information to aircrafts, and by providing flight context information to pilots, such as routes to waypoints and weather conditions.

An important characteristic of the ATM domain of today is that there are limited interactions with the external world, and therefore also limited security problems in relation to information flow to and from the environment. A further characteristic is that humans are at the center of the decision and work processes and the role of automated decision support systems and tools is limited.

The planned and ongoing changes raise new security issues and security concerns with immediate impact on safety issues. The overall objective of this risk analysis is to understand, document and assess security risks of ATM with particular focus on the arrival management process with the involved activities, tasks, roles, components and interactions. The party of the analysis is the ATM service provider.

Target Description. The target of the analysis is a specific Area Control Center and the activities of the Air Traffic Controllers in the arrival management process. The Area Control Center is a ground-based center with the responsibility of managing the traffic of a given airspace. The actual traffic management is conducted from the Operation room (OPS Room), which is the operational environment of the Air Traffic Controllers. The Air Traffic Controllers have different roles, some of which have their own Controller Working Position (CWP). The Controller Working Positions makes a range of tools for surveillance, communication and planning available to the Air Traffic Controllers.

In the following we first document the target of analysis before the introduction of AMAN and ADS-B. Thereafter we describe the planned changes, before we document the target description where the introduction of AMAN and ADS-B is reflected. The selected part of the target is documented by use of UML structured classifiers and activity diagrams.

Before Changes. The structured classifier of Fig. 12 shows the structure and communication links of the ATM components, while the structured classifier of Fig. 13 shows the Operation room as consisting of a number of ACC islands that are connected to the ACC network. Each ACC island consists of a number of Controller Working Positions, each of which are operated by exactly one Air Traffic Controller, and is divided into sector teams with the responsibility of assigned sectors of the airspace. The Operation room furthermore has a number

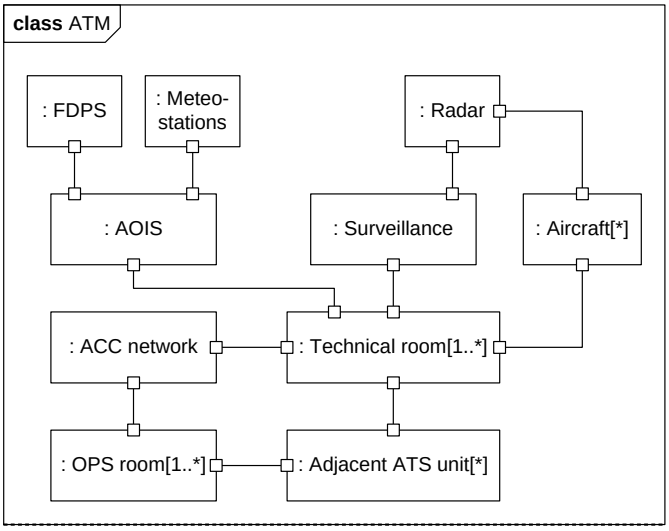

Fig. 12. ATM before changes

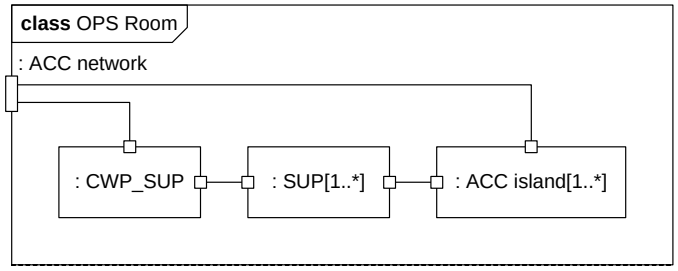

Fig. 13. Operation room before changes

of Supervisors (SUPs) that communicate with the ACC islands, and that also are connected to the ACC Network via the Controller Working Position CWP_SUP. The UML activity diagram of Fig. 14 gives a high-level overview of the various tasks of the arrival management process.

Planned Changes. The changes we are addressing are, in particular, the introduction of the Arrival Manager (AMAN) in the managing of air traffic, as well as the introduction of the Automatic Dependent Surveillance-Broadcast (ADS-B). The AMAN is a decision support tool for the automation of Air Traffic Controllers' tasks in the arrival management, such as the computation of arrival sequences for aircrafts approaching an airport. The introduction of the AMAN affects the Controller Working Positions, as well as the Area Control Center as a whole. The main foreseen changes from an operational and organisational point of view are the automation of tasks (i.e. the usage of the AMAN for the com-

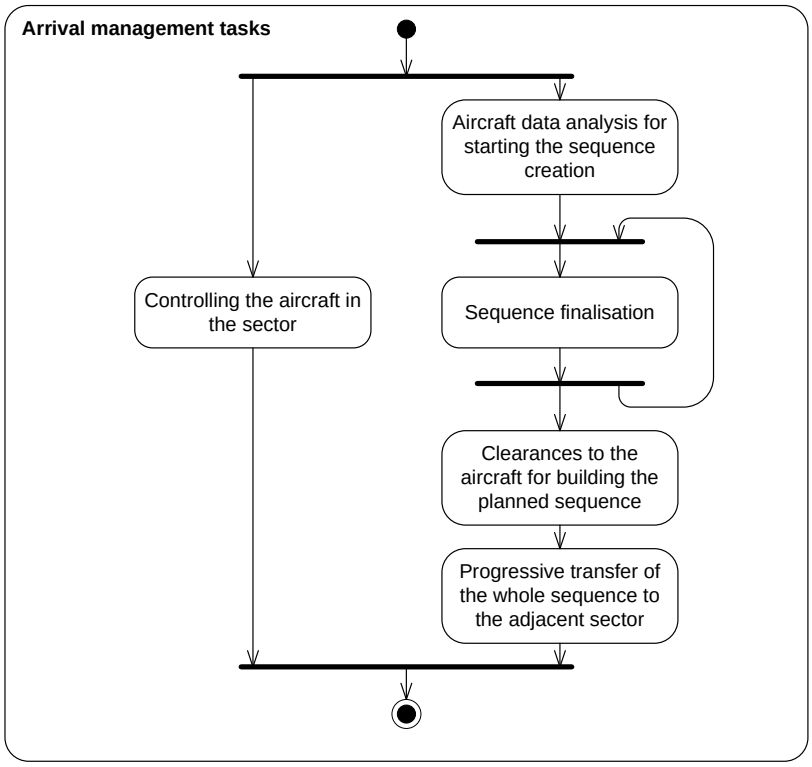

Fig. 14. Arrival management tasks before changes

putation of the arrival sequence) that currently are carried out by Air Traffic Controllers.

The introduction of the ADS-B is actually independent of the AMAN, but is introduced during the same time frame. ADS-B is a cooperative GPS-based surveillance technique for air traffic control where the aircrafts constantly broadcasts their position to the ground and to other aircrafts.

After Changes. In order to highlight the changes in the diagrams, we use grey shading to indicate elements that are introduced. The UML structured classifier of Fig. 15 shows the structure and communication links of the ATM components after the changes. At this level we only see the introduction of the ADS-B. The diagram of Fig. 16 shows the internal structure of the Operation room after the introduction of the AMAN. The AMAN is connected to the ACC network, and thereby also to the ACC islands and the Controller Working Positions. The UML activity diagram of Fig. 17 gives a high-level overview of the arrival management tasks after the changes.

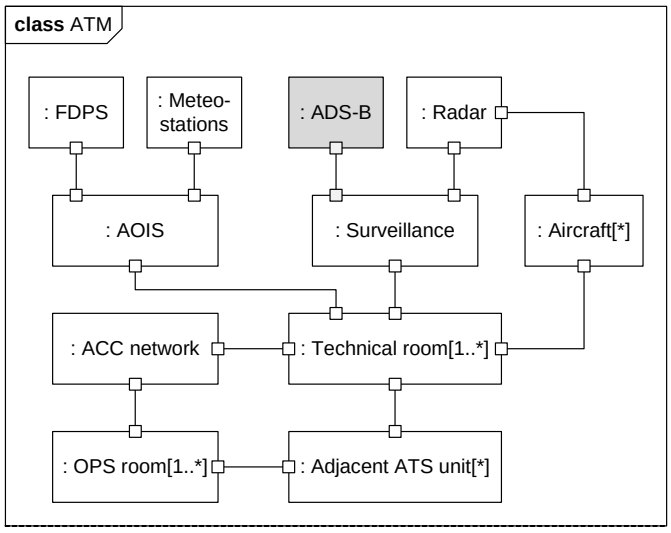

Fig. 15. ATM after changes

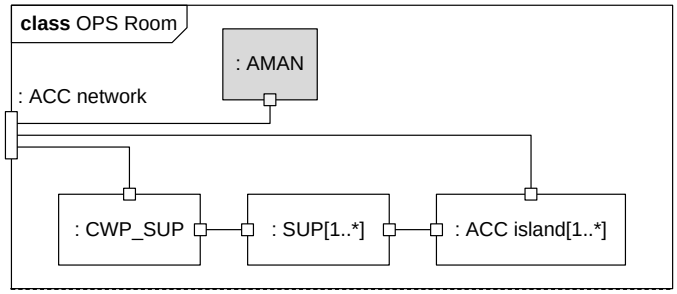

Fig. 16. Operation room after changes

Assets. The purpose of asset identification is to identify the parts, aspects or properties of the target with respect to which the risk analysis will be conducted. An asset is something to which a party assigns value, and hence for which the party requires protection.

In this analysis, the party is the ATM service provider who owns the Area Control Center in question. The risk analysis addresses security issues. Before the changes, the focus is on the security property *Information Provision*: The provisioning of information regarding queue management sensitive data by specific actors (or systems) must be guaranteed 24 hours a day, 7 days a week, taking into account the kind of data shared, their confidentiality level and the different actors involved.

As explained above the ATM of today has little interaction with the external world, but the planned changes will raise new security issues. Therefore, in the

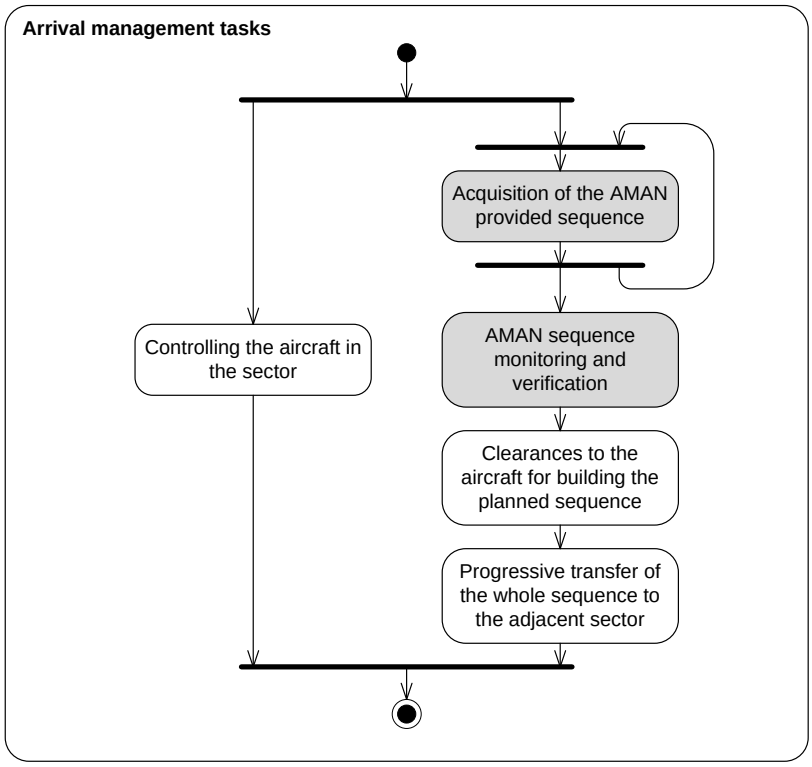

Fig. 17. Arrival management tasks after changes

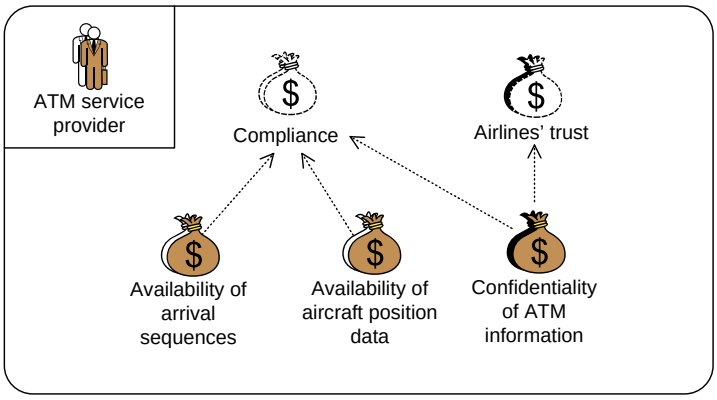

Fig. 18. Assets before and after the changes

analysis of the situation after the changes, the focus is extended to also include the security property *Information protection*: Unauthorised actors (or systems) are not allowed to access confidential queue management information.

In the asset identification these security properties are covered by the two corresponding concepts of confidentiality and availability. More precisely we define the assets of the analysis to be *Availability of arrival sequences* and *Availability of aircraft position data* before the changes, with the addition of *Confidentiality of ATM information*, after the changes. In addition we define the indirect assets *Compliance* and *Airlines' trust*. Indirect assets are assets that, with respect to the target and scope at hand, is harmed only via harm to other assets. Assets that are not indirect we often refer to as direct assets.

We use CORAS asset diagrams to document the assets and the relations between them. The asset diagram showing assets both *before* and *after* is given in Fig. 18. *Availability of arrival sequences* and *Availability of aircraft position data*, as well as the indirect asset *Compliance*, shown with white shadows, are relevant both *before* and *after*, while *Confidentiality of ATM information* and the indirect asset *Airlines' trust*, shown with black shadow, are relevant only *after*. The arrows of the diagram specify harm relations. One asset is related to another if harm to the former may lead to harm to the latter. In this case, harm to all direct assets may lead to harm to the indirect asset *Compliance*, but only harm no *Confidentiality of ATM information* may lead to harm to *Airlines' trust*.

Risk Evaluation Criteria. The risk evaluation criteria define the level of risk that the party, i.e. the ATM service provider, is willing to accept for the given target of analysis. Basically, the criteria are a mapping from risk levels to the decision of either accepting the risk or evaluating the risk further for possible treatment.

In order to speak of risk levels, we need first to define consequence and likelihood scales and a risk function. The consequence and likelihood scales are partly based on requirements and advisory material provided by EUROCONTROL [5, 7], and are given in Tables 5 through 7. The risk function is a mapping from pairs of consequence and likelihood to risk levels and is documented by the risk matrix shown in Fig. 19. We use three risk levels, namely low (light grey), medium and high (dark grey). With these definitions we define the risk evaluation criteria as follows:

- *High risk*: Unacceptable and must be treated.
- *Medium risk*: Must be evaluated for possible treatment.
- *Low risk*: Must be monitored.

In this analysis, the scales, risk function and risk evaluation criteria apply both before and after the changes, while in the general this is not always case as both scales, risk functions and risk evaluation criteria may change in the transition from *before* to *after*.

6.2 Risk Identification

Risk identification is conducted as a structured brainstorming involving personnel with first hand knowledge about the target of analysis. By conducting a

Table 5. Consequence scale for availability

Consequence	Description
Catastrophic	Catastrophic accident
Major	Abrupt manoeuvre required
Moderate	Recovery from large reduction in separation
Minor	Increasing workload of ATCOs or pilots
Insignificant	No hazardous effect on operations

Table 6. Consequence scale for confidentiality

Consequence	Description
Catastrophic	Loss of data that can be utilised in terror
Major	Data loss of legal implications
Moderate	Distortion of air company competition
Minor	Loss of aircraft information data (apart from aircraft position data)
Insignificant	Loss of publicly available data

Table 7. Likelihood scale

Likelihood	Description
Certain	A very high number of similar occurrences already on record; has occurred a very high number of times at the same location/time
Likely	A significant number of similar occurrences already on record; has occurred a significant number of times at the same location
Possible	Several similar occurrences on record; has occurred more than once at the same location
Unlikely	Only very few similar incidents on record when considering a large traffic volume or no records on a small traffic volume
Rare	Has never occurred yet throughout the total lifetime of the system

		Consequence				
		Insignificant	Minor	Moderate	Major	Catastrophic
Likelihood	Rare					
	Unlikely					
	Possible					
	Likely					
	Certain					

Fig. 19. Risk function

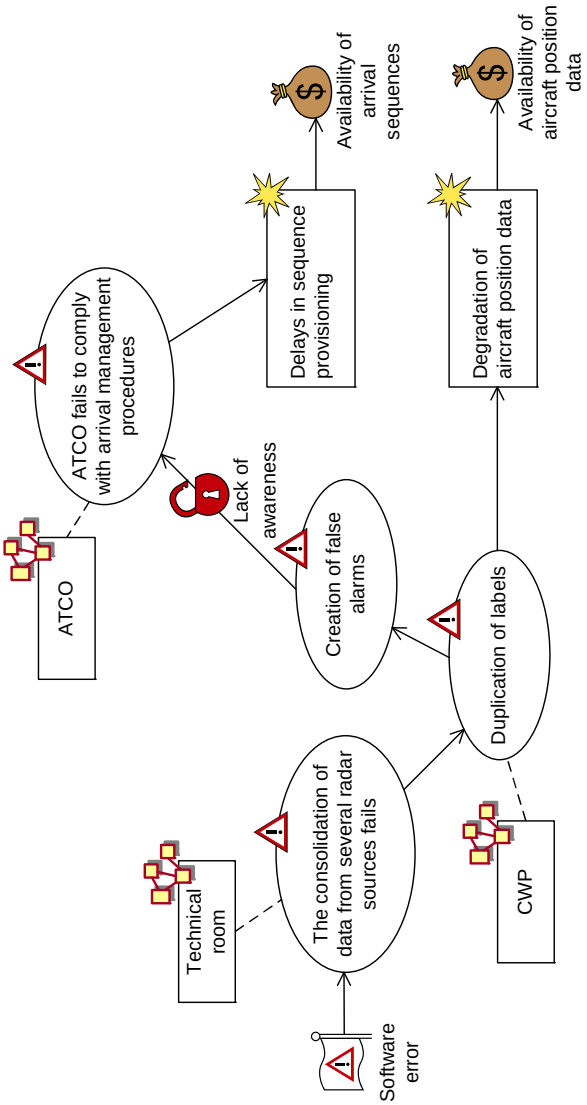

Fig. 20. Threat diagram before changes

walkthrough of the target description, the risk are identified by systematically identifying unwanted incidents, threats, threat scenarios and vulnerabilities. The results are documented on-the-fly by means of CORAS threat diagrams.

While the various parts of the threat diagrams are modelled and documented, the relations to the target of analysis are identified and documented at the same time. The approach to risk identification is to first identify and document risks for the target of analysis before the changes. Once this is completed, we

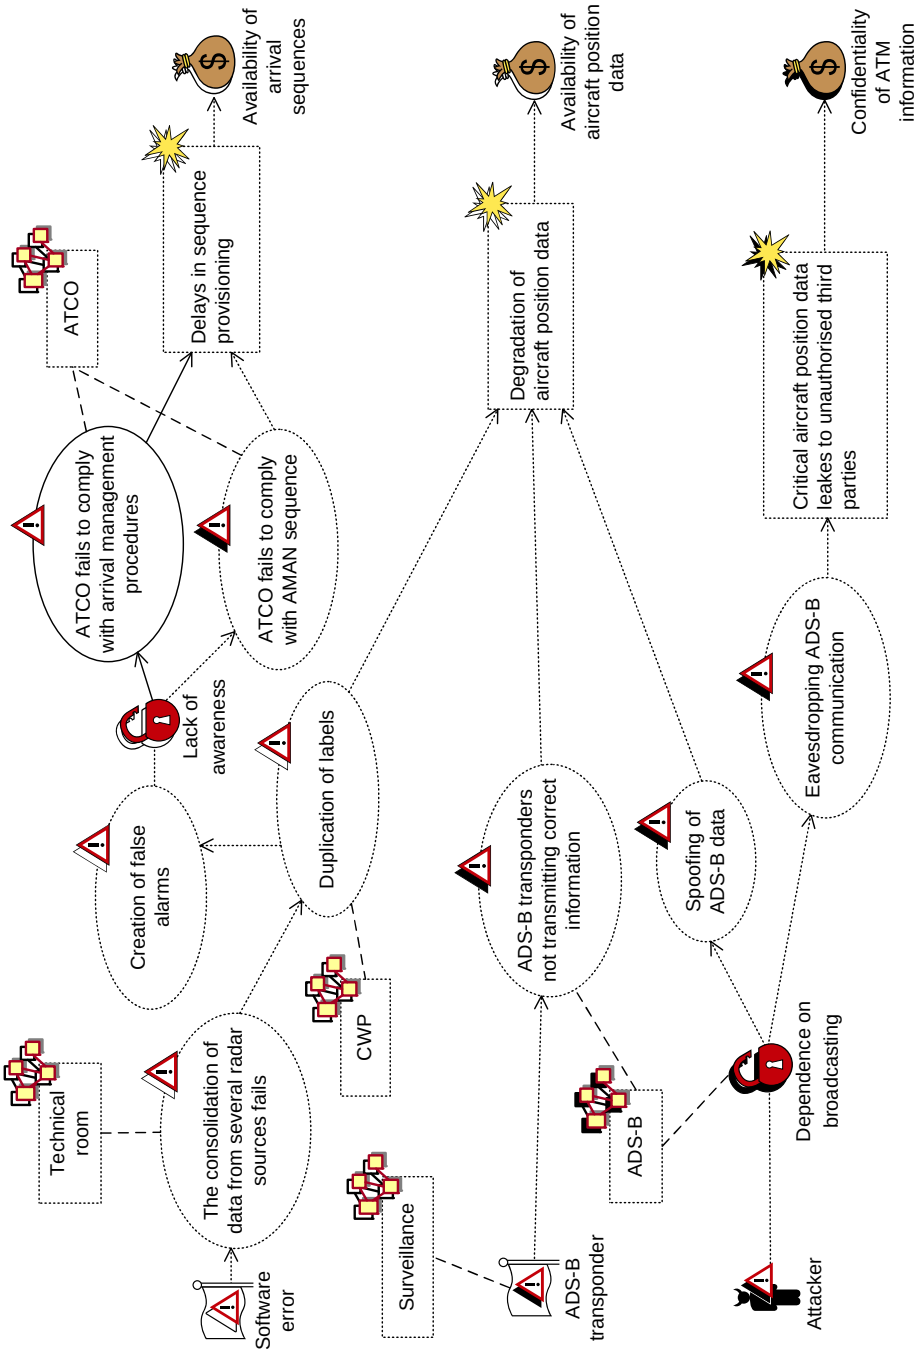

Fig. 21. Threat diagram before and after changes

proceed by identifying and documenting the risks after the changes. Based on the documented relations to the target description, i.e. the trace model, we identify the parts of the threat scenarios that are not affected by the changes and therefore do not have to be addressed again from scratch.

Before Change. The threat diagram of Fig. 20 documents unwanted incidents that may arise due to duplication of labels on the Controller Working Position interface. A label depicts an aircraft with its position data, and is derived from radar data. When several radar sources are used, the label is generated by automatically consolidating (merging) the data from the various sources. In some cases, software errors may yield duplicated labels that may lead Air Traffic Controllers to believe there are two aircrafts. The duplication may also lead to false near miss alarms. This may lead to two unwanted incidents *Delays in sequence provisioning* and *Degradation of aircraft position data*.

After Change. Figure 21 is the result of the risk identification where the changes have been taken into account. The unwanted incidents documented in Fig. 20 persist under the changes and we therefore find them as "two-layered" *before-after* elements in Fig. 21. However, the incidents *Delays in sequence provisioning* may be caused by the threat scenario *ATCO fails to comply with arrival management procedures* only before the changes and by the threat scenario *ATCO fails to comply with AMAN sequence* only after the changes.

Due to the introduction of ADS-B as a means for surveillance, there are also further threats and threat scenarios that are relevant for the unwanted incident *Degradation of aircraft position data* after the changes. This is documented by the threats *ADS-B transponder* and *Attacker*, and the threat scenarios *ADS-B transponders not transmitting correct information* and *Spoofing of ADS-B data*. The threat *Attacker* also initiates a new threat scenario *Eavesdropping ADS-B communication* which leads to a new unwanted incident *Critical aircraft position data leaks to unauthorised third parties* that may cause harm to the new asset *Confidentiality of ATM information*.

6.3 Risk Estimation

The risk estimation basically amounts to estimating likelihoods and consequences for the unwanted incidents. Usually, we also estimate likelihoods for threat scenarios in order to get a better basis for estimating the likelihood of unwanted incidents and to understand the most important sources of risks.

To the extent that threat scenarios and unwanted incidents before changes are completely unaffected by the changes, the risk estimates need not be conducted twice. However, when scenarios and incidents are affected by the changes, the value of these must be reassessed. Likewise, we must estimate likelihoods and consequences for scenarios and incidents that only are relevant after the changes.

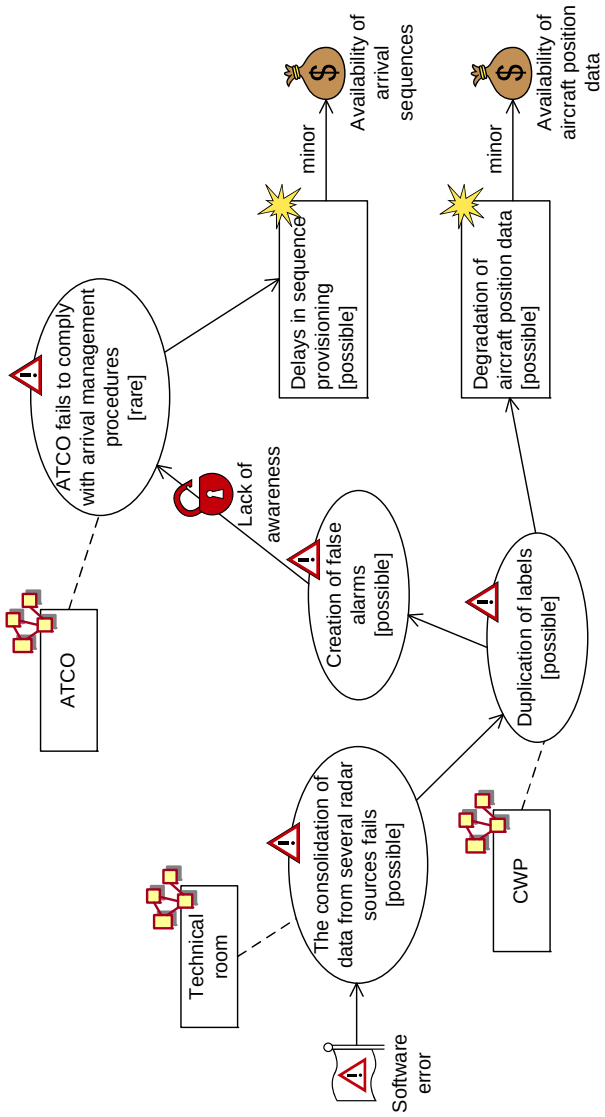

Fig. 22. Risk estimation before changes

Before Change. The estimation of likelihoods and consequences before the changes is documented in Fig. 22.

After Change. The threat diagram of Fig. 23 documents the estimation of likelihoods and consequences when the changes are taken into account. The issues in relation to ADS-B are relevant only after the changes and are therefore only assigned *after* values. The likelihood of the threat scenario *Creation of false*

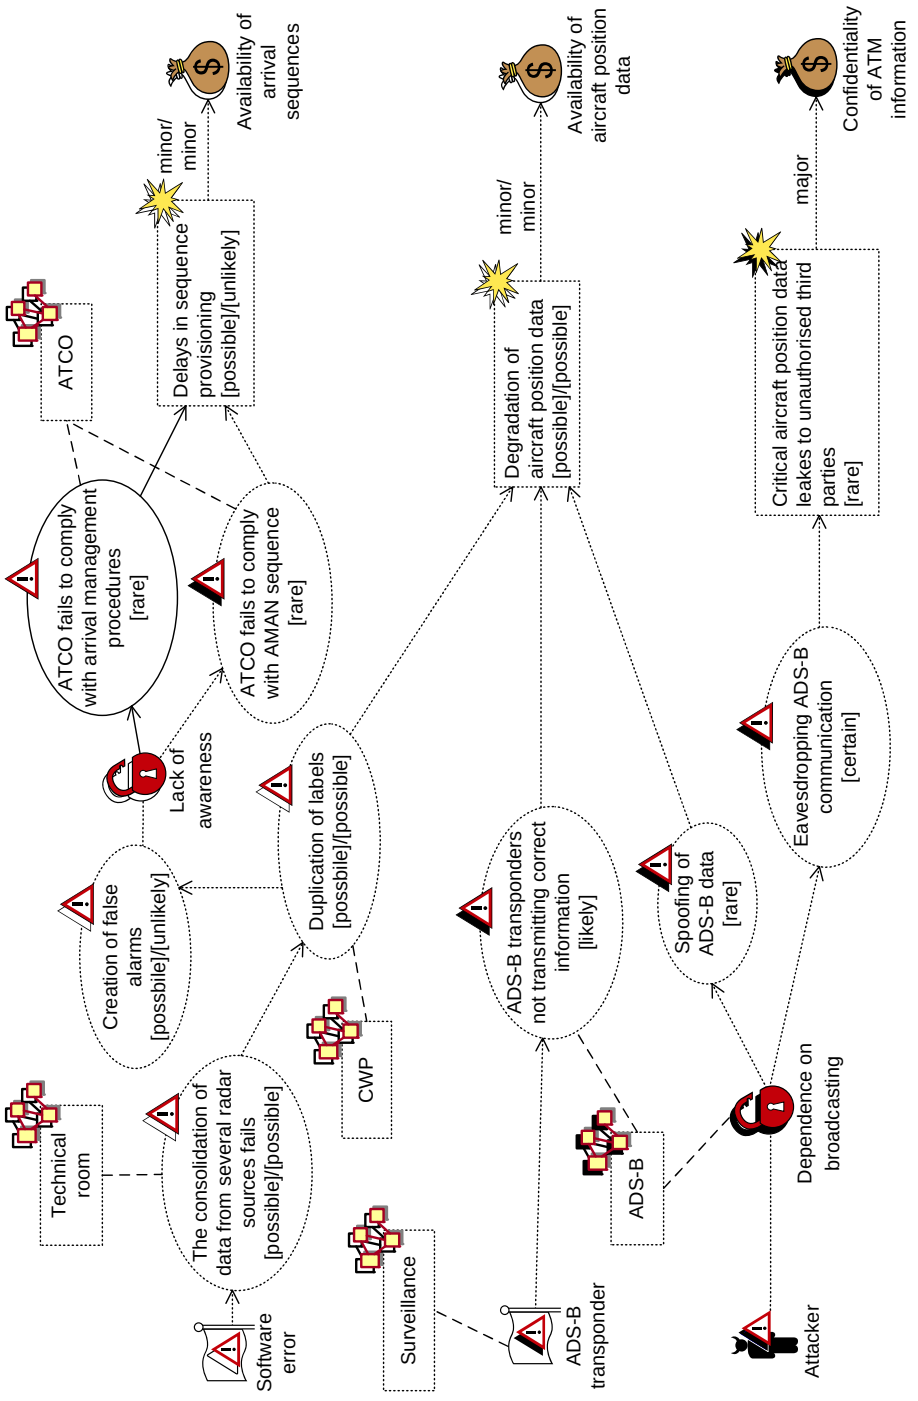

Fig. 23. Risk estimation before and after changes

alarms and the unwanted incident *Delays in sequence provisioning* furthermore changes under the changes to the target of analysis. The threat scenario *ATCO fails to comply with arrival management procedures* is not relevant after the changes and are thus only assigned a *before* value.

6.4 Risk Evaluation

The purpose of risk evaluation is to decide which of the identified risks are acceptable and which must be evaluated for treatments. But before we evaluate the risks, we also need to estimate risks to the indirect assets, which until this point have not been dealt with. Indirect assets are assets that are harmed only through harm to other assets. This activity therefore consists of following the harm relations from direct to indirect assets documented in the asset diagram (recall Fig. 18) and estimating how the harm caused by unwanted incidents propagate. The result of this is documented in Fig. 24.

A risk is the likelihood of an unwanted incident and its consequence for a specific asset. From the risk estimation documented in Figs. 23 and 24, we get seven risks when coupling unwanted incidents and assets.

- R1:*Delays in sequence provisioning* (risk *before-after* toward *Availability of arrival sequences*)
- R2:*Degradation of aircraft position data* (risk *before-after* toward *Availability of aircraft position data*)
- R3:*Critical aircraft position data leaks to unauthorised third part* (risk *after* toward *Confidentiality of ATM information*)
- R4:*Delays in sequence provisioning* (risk *before-after* toward *Compliance*)

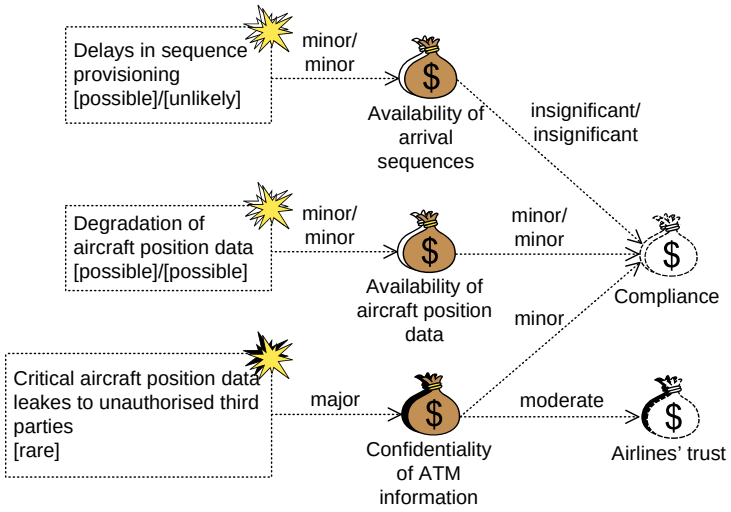

Fig. 24. Risk estimation for indirect assets

		Consequence				
		Insignificant	Minor	Moderate	Major	Catastrophic
Likelihood	Rare		**R6**	**R7**	**R3**	
	Unlikely	**R4**	**R1**			
	Possible	*R4*	*R1*, ***R2***, ***R5***			
	Likely					
	Certain					

Fig. 25. Risk evaluation

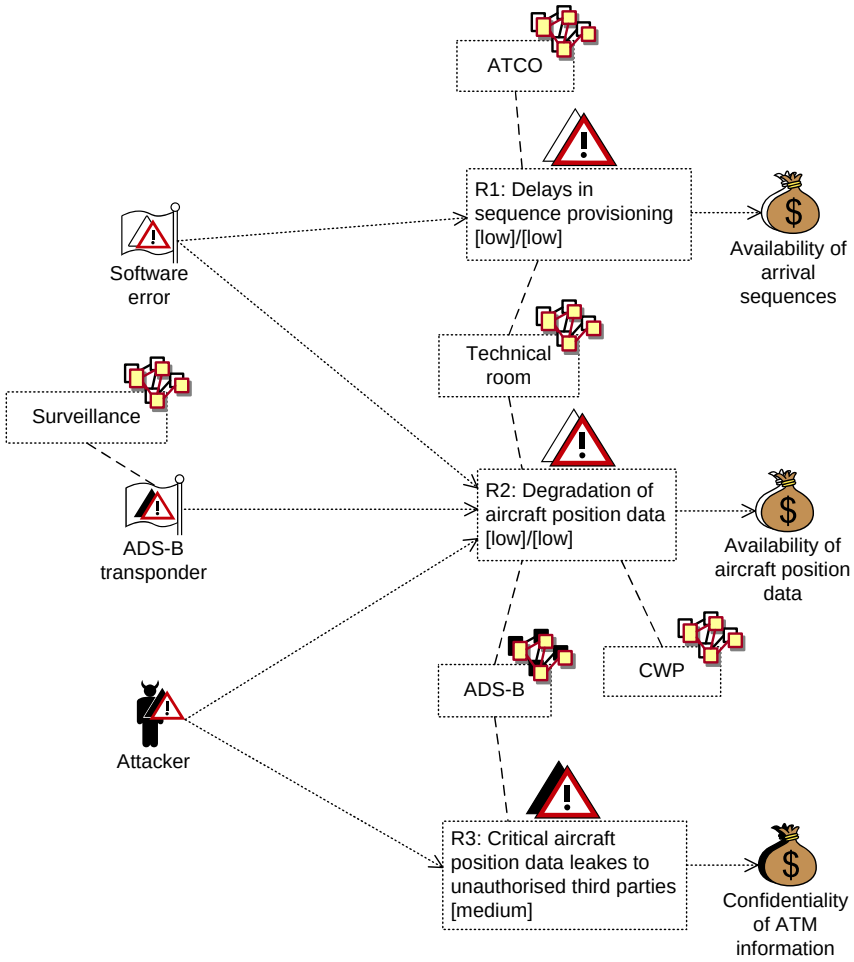

Fig. 26. Risk diagram before and after changes – direct assets

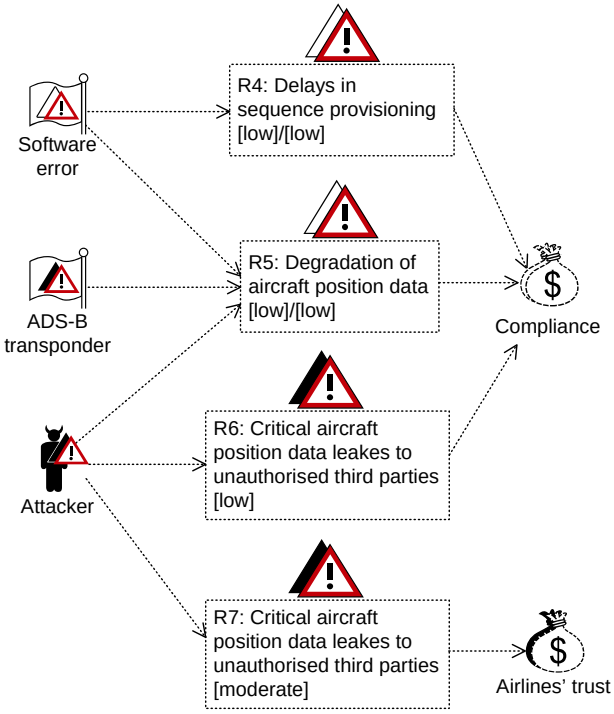

Fig. 27. Risk diagram before and after changes – indirect assets

- R5:*Degradation of aircraft position data* (risk *before-after* toward *Compliance*)
- R6:*Critical aircraft position data leaks to unauthorised third part* (risk *after* toward *Compliance*)
- R7:*Critical aircraft position data leaks to unauthorised third part* (risk *after* toward *Airlines' trust*)

The first step of the risk evaluation is to calculate the risk levels. These are obtained by plotting the risks in the risk function defined during the context establishment (recall Fig. 19) using the estimated likelihoods and consequences. The risk evaluation is conducted separately for the risks before the changes and the risk after the changes. Figure 25 shows the risks plotted in the risk function; we use *italic* to indicate risks *before* and **bold** to indicate risks *after*.

We use CORAS risk diagrams to document the results of calculating the risk levels. These diagrams show the risks together with the threats that initiate them and the assets they harm. The three risks toward direct assets (R1–R3) are documented in Fig. 26, while the risks toward indirect assets (R4–R7) are documented in Fig. 27. By the risk evaluation criteria defined in the context establishment, risks R3 and R7 must be evaluated for possible treatments, while the rest must be monitored.

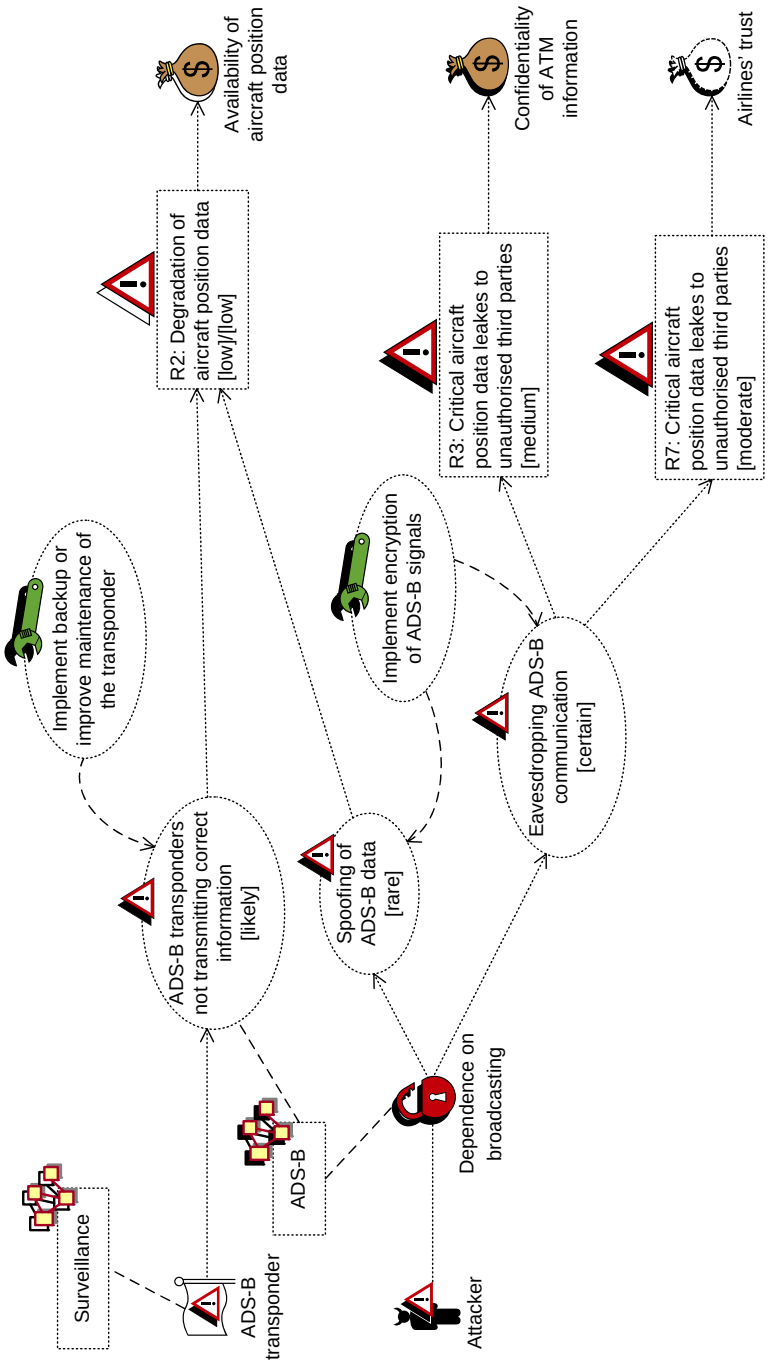

Fig. 28. Treatment diagram after changes

6.5 Risk Treatment

For changes that are planned and/or anticipated, the risk treatment should ensure that the risk level is maintained at an acceptable level through the changes, and that the risk level of the resulting system is acceptable. Whether or not the current risks should be subject to treatment depends on the time frame of the change, as well as the priorities of the parties and other stakeholder. If the planned or anticipated changes are very immediate, it may not make sense to invest in treatments for risks that disappear after the changes.

Identified treatment options are documented using CORAS treatment diagrams, as shown in Fig. 28. The identified treatments focus mainly on the risk picture after the changes and, because only risks R3 and R7 must be evaluated for possible treatments, on the introduction of the ADS-B. One of the treatment options, namely ADS-B encryption, can ensure both confidentiality and authentication, while the second turns out to be relevant only for availability and thus not for R3 and R7.

7 Related Work

The extent to which existing risk analysis methods support the analysis and representation of changing risks depends of course on the relevant perspective of change; recall from Sect. 2 the classification of changes into the *maintenance*, *before-after* and *continuous evolution* perspectives. Somewhat more thorough treatments than provided in this paper of the *maintenance* and *continuous evolution* perspectives in CORAS are given in [17, 20], in particular an approach to the *continuous evolution* perspective inspired by an application of CORAS in risk monitoring [23].

In principle, any risk analysis method can be used to analyse changing risks from the *maintenance* and the *before-after* perspective by staring from scratch and doing complete reassessments. For the management of changing risks to be efficient, however, there should be methodical support for dealing with change in a systematic and focused manner.

Most of the established risk analysis methods provide little or no support for analysing changing and evolving risks. The ISO 31000 risk management standard [15] prescribes change detection and identification for emerging risks, but provides no guidelines for how to do this in practice. A state-of-the-art methodology like OCTAVE [1] recommends reviewing risks and critical assets, but how the risk analysis results should be updated is not specified.

Some approaches have support for associating elements of risk models to parts of the target description, which may facilitate the identification and documentation of risk changes due to target changes. UML based approaches such as misuse cases [28] may utilise built-in mechanisms in the UML for relating elements from different UML diagrams. ProSecO [12] relates risks to elements of a functional model of the target.

With respect to the *before-after* perspective, ProSecO provides some support for modelling the various phases or states of a change process; when the models

change, elements of the model may be transferred to states that indicate the need for additional risk analysis. ProSecO furthermore provides some support for the modelling of the change process by means of state machines. Other academic studies have focused on either maintenance [18, 27] or variants of reassessment [9, 16].

The *continuous evolution* perspective is the most general of the three perspectives on change as it incorporates time into the picture and thereby refers to the instance of the evolving risk picture at any given point in time. Among the existing approaches the support for this perspective is virtually non-existent. Some risk modelling approaches do provide support for updating the values that are annotated to diagrams in the sense that by changing the input values, the derived output values can be automatically updated. These includes fault trees [13], Markov models [11], and Bayesian networks [2, 8]. Still, CORAS is to our knowledge the sole approach to provide means for predicting the evolution of risks related to an evolving target.

8 Conclusion

A traditional risk analysis considers the target of analysis at a particular point in time and in a particular configuration. The result is the documentation of a risk picture valid for that configuration at that point in time. However, systems and their environment tend to change and evolve over time and the risks toward the system may shift. If changes are not captured and reflected in the risk documentation the established risk picture may no longer be valid, with the consequence that we no longer have a full and correct overview of the risks of our target.

Conducting a new risk analysis from scratch when the target or its environment has changed is not an approach to be preferred as risk analyses are time and resource consuming and parts of the risk documentation is likely to still be valid. In order to appropriately handle change, risk analysis methods should be supported with techniques and guidelines that are specialised toward modelling, analysing and reasoning about changing risks. However, the nature of the changes have implications the methodological needs. In this paper we have categorised changes into three perspectives referred to as *maintenance*, *before-after* and *continuous evolution* and characterised the methodological needs for each of them.

The focus in this paper has been on the *before-after* perspective which comprise planned and anticipated changes of a radical and extensive nature. The challenge in this perspective is to obtain and present a risk picture that describes both current and future risks without doing double work. This requires methods and guidelines for distinguishing the parts of the risk documentation that are affected by changes from the parts that are unaffected.

In this paper we handle this by providing modelling support for changing risks and means for relating risk models to target models in such a way that changes in the target models can be traced to the risk models. These are given

as extensions to risk graphs, which may be seen as a common abstraction of a variety of different risk modelling techniques. In this way, the extension may be instantiated into these different techniques. We demonstrate this by providing an instantiation into the CORAS risk modelling language. Further, we demonstrate our approach by a practical example from a risk analysis case study from the Air Traffic Management (ATM) domain.

In difference from other approaches, we provide a systematic approach to change management in the context of risk analysis. This includes methodological guidelines on how to maintain and update risk analysis documentation to reflect both past and future changes to the target without starting from scratch and doing a full reassessment, something that to a large extent is missing in other risk analysis methodologies.

Acknowledgements. The work on which this paper reports has been funded by the European Commission through the projects SecureChange (Contract no. 231101) and NESSoS (Contract no. 256980). Thanks to Massimo Felici, Frank Innerhofer-Oberperfler, Valentino Meduri, Alessandra Tedeschi, and the ATM experts participating in the analysis, for their contributions to the ATM case study.

References

1. Alberts, C.J., Davey, J.: OCTAVE criteria version 2.0. Technical report CMU/SEI-2001-TR-016, Carnegie Mellon University (2004)
2. Ben-Gal, I.: Bayesian networks. In: Ruggeri, F., Kenett, R.S., Faltin, F.W. (eds.) Encyclopedia of Statistics in Quality and Reliability. John Wiley & Sons (2007)
3. Brændeland, G., Refsdal, A., Stølen, K.: Modular analysis and modelling of risk scenarios with dependencies. J. Syst. Softw. 83(10), 1995–2013 (2010)
4. Dudely, R.M.: Real analysis and probability. Cambridge Studies in Advanced Mathematics, Cambridge (2002)
5. EUROCONTROL: EUROCONTROL safety regulatory requirements (ESARR) 4 – Risk assessment and mitigation (2001)
6. EUROCONTROL: Air Traffic Management Strategy for the years 2000+. Volume 1 and Volume 2 (2003)
7. EUROCONTROL: ESARR advisory material/Guidance document (EAM 2/ GUI 5) – Harmonisation of safety occurrences severity and risk assessment (2005)
8. Fenton, N., Neil, M.: Combining evidence in risk analysis using Bayesian networks. Agena White Paper W0704/01, Agena (2004)
9. Goel, S., Chen, V.: Can business process reengineering lead to security vulnerabilities: Analyzing the reengineered process. Int. J. Prod. Econ. 115(1), 104–112 (2008)
10. Haugen, Ø., Husa, K.E., Runde, R.K., Stølen, K.: STAIRS towards formal design with sequence diagrams. Softw. Syst. Modeling 4(4), 355–367 (2005)
11. Howard, R.A.: Dynamic Probabilistic Systems, Volume I: Markov Models. John Wiley & Sons (1971)
12. Innerhofer-Oberperfler, F., Breu, R.: Using an enterprise architecture for IT risk management. In: Information Security South Africa Conference (ISSA'06) (2006)

13. International Electrotechnical Commission: IEC 61025 Fault Tree Analysis (FTA) (1990)
14. International Electrotechnical Commission: IEC 60300-9 Dependability management - Part 3: Application guide - Section 9: Risk analysis of technological systems - Event Tree Analysis (ETA) (1995)
15. International Organization for Standardization: ISO 31000 Risk management – Principles and guidelines (2009)
16. Lee, E., Park, Y. Shin, J.G.: Large engineering project risk management using a Bayesian belief network. Expert Syst. Appl. 36(3), 5880–5887 (2009)
17. Lund, M.S., et al.: SecureChange Deliverable D5.3 Assessment method. Available from `http://www.securechange.eu/sites/default/files/deliverables/D5.3. Assessment Methods.pdf` (2011)
18. Lund, M.S., den Braber, F., Stølen, K.: Maintaining results from security assessments. In: 7th European Conference on Software Maintenance and Reengineering (CSMR'03). pp. 341–350. IEEE Computer Society (2003)
19. Lund, M.S., Solhaug, B., Stølen, K.: Evolution in relation to risk and trust management. Comput. 43(5), 49–55 (2010)
20. Lund, M.S., Solhaug, B., Stølen, K.: Model-Driven Risk Analysis – The CORAS Approach. Springer (2011)
21. Nielsen, D.S.: The cause/consequence diagram method as basis for quantitative accident analysis. Technical report RISO-M-1374, Danish Atomic Energy Commission (1971)
22. Object Management Group: Unified Modeling Language: Superstructure, version 2.1.1 (non-change bar). OMG Document: formal/2007-02-05 (2005)
23. Refsdal, A., Stølen, K.: Employing key indicators to provide a dynamic risk picture with a notion of confidence. In: Trust Management III. Third IFIP WG 11.11 International Conference (IFIPTM'09). pp. 215–233. Springer (2009)
24. Robinson, R.M., Anderson, K., Browning, B., Francis, G., Kanga, M., Millen, T., Tillman, C.: Risk and Reliability – An Introductory Text. R2A, 5th edn. (2001)
25. Schneider, S.: Attack trees: Modeling security threats. Dr. Dobb's J. 24, 21–29 (1999)
26. SESAR Consortium: The ATM Target Concept. SESAR Definition Phase - Deliverable 3 (2007)
27. Sherer, S.A.: Using risk analysis to manage software maintenance. J. Softw. Maint.: Res. Pract. 9(6), 345–364 (1997)
28. Sindre, G., Opdahl, A.L.: Eliciting security requirements by misuse cases. In: 37th International Conference on Technology of Object-Oriented Languages and Systems (TOOLS Pacific'00). pp. 120–131. IEEE Computer Society (2000)