

Generalising Event Trees Using Bayesian Networks with a Case Study of Train Derailment

George Bearfield¹ and William Marsh²

¹ Atkins Rail, Euston Tower, 286 Euston Road, London, NW1 3AT, UK
george.bearfield@atkinsglobal.com

² Department of Computer Science, Queen Mary University of London, Mile End Road,
London, E1 4NS, UK
william@dcs.qmul.ac.uk

Abstract. Event trees are a popular technique for modelling accidents in system safety analyses. Bayesian networks are a probabilistic modelling technique representing influences between uncertain variables. Although popular in expert systems, Bayesian networks are not used widely for safety. Using a train derailment case study, we show how an event tree can be viewed as a Bayesian network, making it clearer when one event affects a later one. Since this effect needs to be understood to construct an event tree correctly, we argue that the two notations should be used together. We then show how the Bayesian Network enables the factors that influence the outcome of events to be represented explicitly. In the case study, this allowed the train derailment model to be generalised and applied in more circumstances. Although the resulting model is no longer just an event tree, the familiar event tree notation remains useful.

1 Introduction

Event trees are used in quantified risk analysis to analyse possible accidents occurring as a consequence of hazardous events in a system. Event trees are often used together with fault trees, which analyse the causes of the hazardous event that initiates the accident sequence. Their origin goes back at least to the WASH-1400 reactor safety study in 1975 [1].

The most serious accident may be quite improbable, so an accurate assessment of the risk requires the probabilities of possible accident scenarios to be determined. The analysis of accidents must consider both the state of the system and of its environment when the hazardous event occurs. The analysis is made more difficult when the environment of a system is complex or variable.

Event trees model an accident as a sequence of events: this is an intuitive approach but it does not explicitly represent the state of the system and its environment, which influences the evolution of events. In this paper, we propose to address this limitation of event trees by using Bayesian Networks (BNs). We have applied this approach to a case study, adapting an existing event tree modelling a train derailment accident. The original author of the event tree was able to explain the system and environmental factors that had been considered when preparing the event tree, but

which could not be included explicitly in it. Using a BN, these factors can be made explicit in the accident model, which can still be viewed as an event tree but is now more general with a single BN-based model taking the place of a set of related event trees.

We argue that the event tree and BN are complementary: an event tree can be translated into a BN allowing two views of the accident model, each view showing different properties of the model most clearly. The generalised model, with system and environmental factors that influence the events made explicit, is a BN but it can still be viewed using the event tree notation.

Event trees are supported by many software packages but are sufficiently simple to be created with standard tools such as a spreadsheet. Perhaps because of this, the notation used by different authors varies. Since we wish to translate between event trees and BNs, the first step, in Section 2, is a precise description of an event tree.

In Section 3, we introduce BNs and show how to translate an event tree into a BN. We first give a ‘generic’ translation based only on the number of events in the tree and then we give rules for simplifying the resulting BN. Section 4 introduces the case study and uses it to show that the combination of event trees and BNs allows a more general model of possible accidents. Conclusions and related work are in Section 5.

2 Event Trees

In this section, we give an informal but precise description of event trees, which will be the basis for the translation of event trees to BNs.

2.1 Events and Outcomes

The evolution of the system following the hazardous occurrence is divided into discrete *events*, starting from the *initiating event*. Each event has a finite set of outcomes; commonly there are just two outcomes – the event happens or does not happen – but a greater number of outcomes can be distinguished.

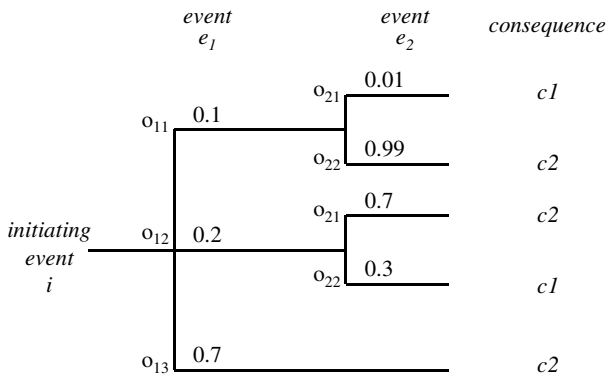


Fig. 1. An example event tree. There are two events: event e_1 has three possible outcomes o_{11} , o_{12} and o_{13} whereas e_2 has only two outcomes o_{21} and o_{22} . Two different consequences are distinguished c_1 and c_2 ; c_1 results both from the event sequence $i \rightarrow o_{11} \rightarrow o_{21}$ and from the event sequence $i \rightarrow o_{12} \rightarrow o_{22}$.

The events form a sequence in time: a tree of possible outcomes for all the events is constructed and the *consequence* or loss evaluated for each path through the tree. Some paths may be judged to lead to the same consequence. Fig. 1 shows an example event tree.

2.2 Probabilities and Consequence

The event tree specifies a logical combination of the event outcomes for each consequence. For the event tree in Fig. 1, the logical formulae for the consequences are $c_1 = (o_{11} \wedge o_{21}) \vee (o_{12} \wedge o_{22})$ and $c_2 = (o_{11} \wedge o_{22}) \vee (o_{12} \wedge o_{21}) \vee o_{13}$.

The probability of each consequence is calculated from the event probabilities, determined from data or experience. For example in Fig. 1, the probability of outcome o_{11} of e_1 event is 0.1. However, the probability of an outcome may depend on the outcomes of events earlier on the path: in Fig. 1 the probability of outcome o_{21} of event e_2 depends on the outcome of event e_1 . The probabilities labelling the branches of the tree for e_2 are therefore conditional probabilities, in this example: $p(o_{21} | o_{11})$, $p(o_{22} | o_{11})$, $p(o_{21} | o_{12})$, and $p(o_{22} | o_{12})$.

The probabilities of the two consequences are calculated by multiplying the probabilities along each path and then adding the probabilities of paths leading to the same consequence. The calculation for Fig 1 is shown below.

<i>Consequence</i>	<i>Calculation</i>	<i>Result</i>
C_1	$0.1 \times 0.01 + 0.2 \times 0.3$	0.061
C_2	$0.1 \times 0.99 + 0.2 \times 0.7 + 0.7$	0.939

It is notable that the logical formulae for the consequences do not carry any information about how the outcome of one event is influenced by earlier events or even of how the events are ordered in time. The logical formulae are sufficient for combining the probabilities of event outcomes to give the consequence probabilities. On the other hand, understanding how the outcome of one event is influenced by earlier events is crucial for judging the event probabilities and the event tree shows only part of the information used during its construction:

- The time ordering of events shows the set of earlier events on which a probability may be conditioned; later events cannot influence the outcome of earlier events.
- However, some earlier events may have no influence and the event tree does not show what subset of the earlier events actually conditions each probabilities. Indeed, we have seen cases where inexperienced users of event trees are unaware that the probabilities attached to branches in an event tree are *conditional* probabilities at all.

In the example of Fig. 1, when event e_1 has outcome o_{13} the tree does not branch for the possible outcomes of event e_2 . We refer to this as a *don't care* condition. There is more than one reason why the event tree may contain such a condition:

- Only one of the outcomes of e_2 is possible following the outcome of the earlier event.
- Both outcomes of e_2 are possible, but the consequence is the same for both.

It is important to note that the event tree does not distinguish between these reasons – there is no need to do so to calculate the consequence probabilities.

3 Translating an Event Tree to a Bayesian Network

In this section we first introduce BNs and describe a ‘generic’ representation of an event tree as a BN before showing how it can be simplified for a specific event tree.

3.1 Bayesian Networks

A BN [2] is a graph with a set of probability tables. The nodes of the graph represent uncertain variables and the arcs represent the causal relationships between the variables. The arcs are directed from ‘parent’ to ‘child’ with, conventionally, the parent as the cause and the child the effect. There is a probability table for each node, providing the probabilities of each state of the variable, for each combination of the states of parent variables. The model of cause is probabilistic rather than deterministic and this makes it possible to include factors that influence the frequency of events, but do not determine their occurrence.

Although the underlying theory (Bayesian probability) has been around for a long time, executing realistic models was only first made possible in the late 1980s using new algorithms. Methods for building large-scale BNs are even more recent [3] but it is only such work that has made it possible to apply BNs to the problems of systems engineering.

The RADAR group at QMUL, in collaboration with Agena Ltd, has built applications based on BNs that have shown the technology to be effective. Several such applications are for dependability assessment, notably the TRACS tool [4] used to assess vehicle reliability by QinetiQ (on behalf of the MOD) and a tool used by Philips to manage software quality [5].

3.2 A Generic Translation from ET to BN

Any event tree with three events e_1 , e_2 , and e_3 can be represented by the BN shown in Fig. 2. Two types of arc complete the network:

- Consequence arcs (shown as dotted lines in Fig. 2) connect each event node to the consequence node. This relationship is deterministic: the probability table for the consequence node encodes the logical relationship between the events and the consequences. (An example is shown in Fig. 5.)
- Causal arcs (shown as solid lines in Fig. 2) connect each event node to all events later in time. We say that e_1 influences the probability of (or, equivalently, is a causal factor for) event e_2 .

We call this representation *generic* since the nodes and arcs depend only on the number of events. However, assuming that the BN is only used to determine the consequence probabilities (i.e. just as the event tree), some of the arcs may not be necessary allowing the BN to be simplified. In the next two sections we give rules for eliminating unnecessary arcs.

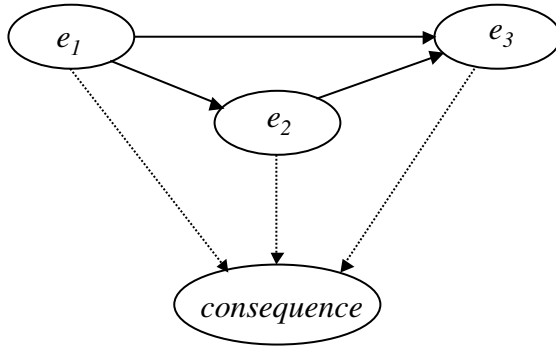


Fig. 2. Generic BN representation of an event tree. Nodes e_1 , e_2 , and e_3 represent the events; each node has a state for each outcome. The node *consequence* has a state for each of the consequences in the event tree.

3.3 Eliminating Consequence Arcs

The consequence arc from an event can be eliminated if the logical formulae for the consequences do not refer to any outcome of the event. Fig. 3 shows an example: the logical expression for c_1 is $(o_{11} \wedge o_{21}) \vee (o_{12} \wedge o_{21})$ but this can be simplified to o_{21} ; since this expression (and the similar expression for c_2) includes only the outcomes of the e_2 event, the BN node e_1 is not needed as a parent of the *consequence* node. The set of consequence arcs is not determined by the branching structure of the event tree but by the assignment of consequences to each of the paths through the tree.

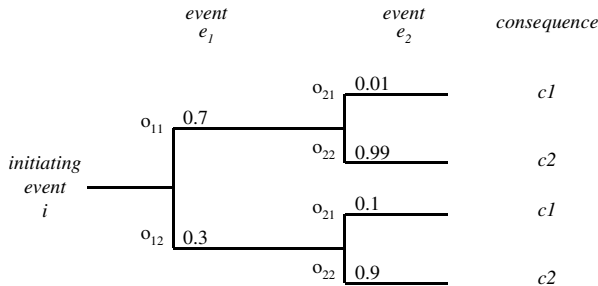


Fig. 3. Example of an event tree allowing a consequence arc to be eliminated, since e_2 determines the consequence whatever the outcome of the first event: the first event influences the relative probability of the two outcome of e_2 but does not change the consequence

3.4 Eliminating Causal Arcs

A causal arc to an event e_t from an earlier event e_f can be eliminated if and only if the probabilities labelling branches for event e_t do not depend on the outcome of event e_f . We can see this in the event tree: are the probabilities labelling an outcome o_{xy} the same on all branches for this outcome or do they differ? An example of this is shown

in Fig. 4, where both branches for o_{21} have probability 0.1 and both branches for o_{22} have probability 0.9:

$$p(e_2 = o_{21} | e_1 = o_{11}) = p(e_2 = o_{21} | e_1 = o_{12}) = 0.1$$

$$p(e_2 = o_{22} | e_1 = o_{11}) = p(e_2 = o_{22} | e_1 = o_{12}) = 0.9$$

Because the probabilities of the outcome of event e_2 do not depend on the outcome of event e_1 no causal arc is needed from e_1 to e_2 . More generally, if for all outcomes of e_i the probability $p(e_i | \dots, e_f, \dots)$ does not depend on the outcome of e_f (given the outcome of the other events) then the two events are ‘conditionally independent’ and the arc from e_f to e_i is not needed.

The complete BNs, including the probability tables, for the event trees in Figs 4 & 5, showing the two types of elimination, are given in Fig. 5.

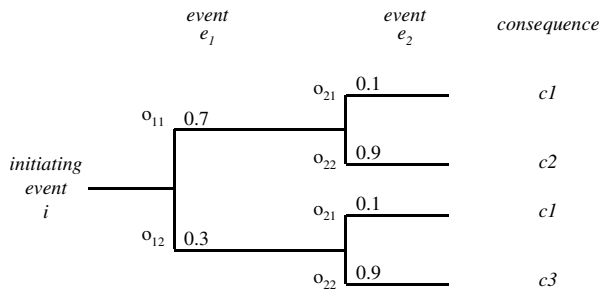


Fig 4. Example of an event tree allowing a causal arc to be eliminated: the probabilities of the two outcomes of event e_2 are the same whatever the outcome of event e_1 . Note that this figure and Fig. 3 have the same shape but differ in the pattern of probabilities and consequences.

3.5 Handling ‘Don’t Care’ Conditions

The event trees in Figs. 3 and 4 are both complete: a path exists for all possible combinations of outcomes of the two events. An event tree that is complete in this way includes all the probabilities needed to complete the node probability tables for the event nodes. However, this is not the case when there are *don’t care* conditions in the event tree. In this section we show how the rules described above can be adapted for *don’t care* conditions.

Consider the *don’t care* branch in the event tree of Fig. 1: suppose that it is instead split into the two outcomes of event e_2 , the first given probability α and the other $1-\alpha$. Any probability α could be used: since the two branches both lead to the same consequence (or set of consequences) the value chosen has no effect on the consequence probabilities. We are free to choose α to simplify the BN as far as possible, so we choose α to create conditional independence whenever this is possible.

This procedure produces the fewest causal arcs but it does not distinguish between the two reasons given at end of section 0 why a *don’t care* condition may occur. This is satisfactory because the distinction doesn’t affect the calculation of the consequence probabilities in either the event tree or the BN. However, by assuming that event outcomes are conditionally independent except when the probabilities shown in the event tree force the opposite conclusion we may have ignored causal

relationships between events that really exist. If we use the BN model of the event sequence for other calculations we may need to add the causal arcs modelling these causal relationships to the BN. We could do this by determining the probabilities of the outcomes of *don't care* conditions and adding extra branches into the event tree. The resulting BN has some interesting properties but we do not need it to calculate consequence probabilities.

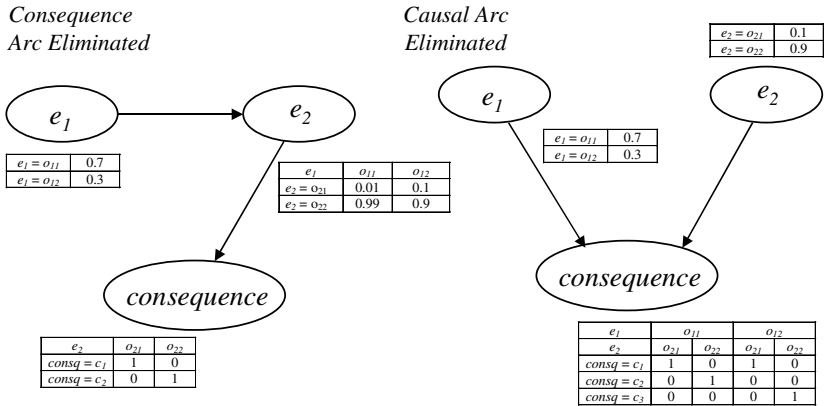


Fig. 5. Complete BNs for event trees of Figs. 4 & 5, showing the two types of arc elimination

3.6 Using a Hierarchy of Nodes for Consequence

Rather than having a single BN ‘consequence’ node with a probability table determined from the logical relationship between events and consequences, it is possible to represent this relationship using a hierarchy of nodes, determined from the event tree structure. A node can be introduced for each vertical line (representing a branch or decision point) in the event tree provided that more than two sequences lead from the branch. The parents of this node are the node representing the event and the nodes from the decision points to the right. Using a hierarchy of nodes has two potential advantages:

- more efficient propagation of the BN
- clearer representation (for the risk analyst) of the logical relationship between events and consequences.

We do not consider the efficiency of propagation further in this paper. In section 4.2, we assess whether the clarity of the model improves using this translation for a realistic event tree.

4 Why Use Bayesian Networks to Model Event Sequences

The previous section showed how to construct a BN equivalent to an event tree; however, if the two models are equivalent what purpose does the BN serve? We examine this using a case study of train derailment, which is introduced in section 4.1.

In the following sections, we first argue that an event tree and a BN provide complementary views of the relationship between events. Secondly, we show how an event tree expressed using a BN can usefully be generalised by making the factors influencing the evolution of events explicit, producing a more widely applicable model of the accident.

4.1 Case Study: Train Derailment

A ‘Derailment Study’ was carried out in 2001 as part of development studies for a proposed upgrade to an urban railway. The objective of the study was to quantify the risks to passengers and staff arising from derailment. This required the consequences of derailment to be analysed and event trees were constructed for this. Other models were used to analyse the frequency of derailment and, given the accident sequences, the likely toll of injuries. Since the ultimate aim was to ensure that risks were tolerable, some conservative assumptions were made.

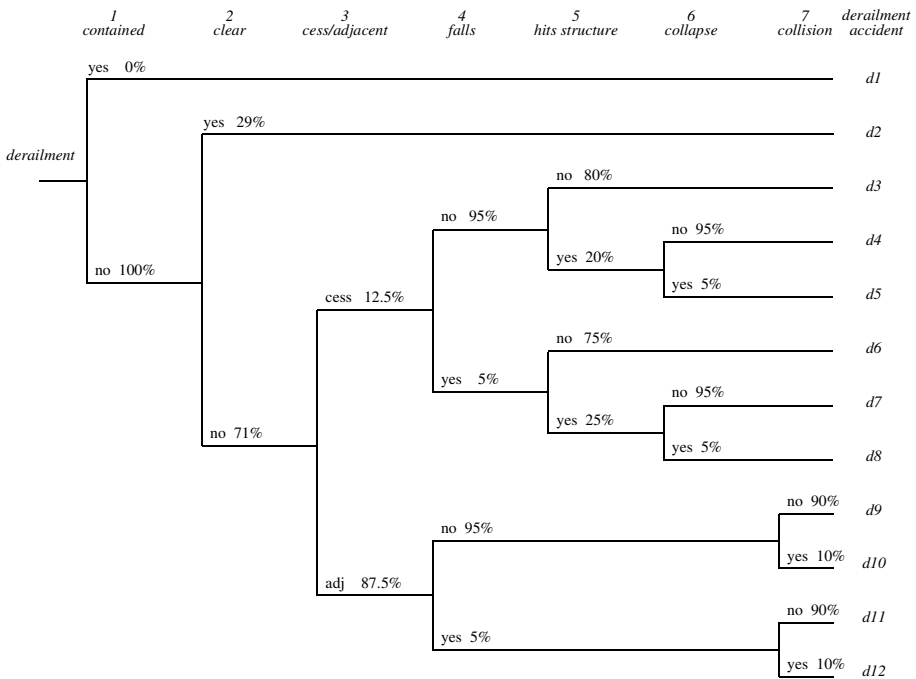


Fig. 6. An event tree from the ‘Derailment Study’ covering derailment in open track areas. The structure of the event tree, and the event probabilities, were adapted from a network-wide model by considering factors specific to the local circumstances.

The analysis used separate event trees for six different infrastructure areas, each with different characteristics including open track, in tunnels and on bridges. Here, we consider only derailments on areas of open track, which is track not in tunnels or carried on bridges. The analysis drew on a version of the ‘Safety Risk Model’ (SRM)

[6], which analyses the risk arising from different hazards using historic accident data and expert judgement for the UK rail network as a whole. The event trees for the derailment study used the structure of the SRM but had to be tailored to the local circumstances: for example, the maximum speed limit is 30 miles per hour, the trains are electric multiple units with third-rail electrification. The original author of the derailment study was available and assisted the authors with the case study.

The event tree for open track derailment is shown in Fig. 6. The events, all of which have only two outcomes, are described in Table 1. Twelve consequences or ‘derailment accidents’ are distinguished: for example ‘d2’ is ‘minor derailment within clearance’ and ‘d7’ is a ‘major derailment to cess, striking line-side structure’. Given the frequency of the initiating ‘derailment’ event, the frequency of each accident can be calculated. The ‘equivalent fatalities’ for each accident are estimated by a separate method, which is not relevant here.

Table 1. Derailment Events

<i>Event</i>		<i>Description</i>
1	Derailment containments controls the train.	An extra raised ‘containment’ rail, if fitted, limits movement sideways.
2	The train maintains clearance.	The train remains within the lateral limits and does not overlap adjacent lines or obstruct beyond the edge of the track area.
3	Derails to cess or adjacent line.	The train can derail to either side of the track: derailing to the ‘cess’, or outside, may lead to a collision with a structure beside the line, while derailing to the ‘adjacent’ side brings a risk of colliding with another train.
4	One or more carriages fall on their side.	The carriages may remain upright or fall over.
5	Train hits a line-side structure.	The train hits a structure beside the line.
6	The train structure collapses.	Collision with a line-side structure causes the train structure to collapse.
7	Secondary collision with a passenger train.	A following or on-coming train collides with the derailed train.

4.2 Causality in the Event Sequence

Fig. 7 shows the BN generated for the event tree, using the algorithm described in section 3. Comparing the two notations – the BN of Fig. 7 and the event tree of Fig. 6 – we see that:

1. The logical combination of events leading to each accident is most clearly shown in the event tree.
2. The occurrence of conditional probabilities – arising from dependence between the events – is shown more clearly in the BN.

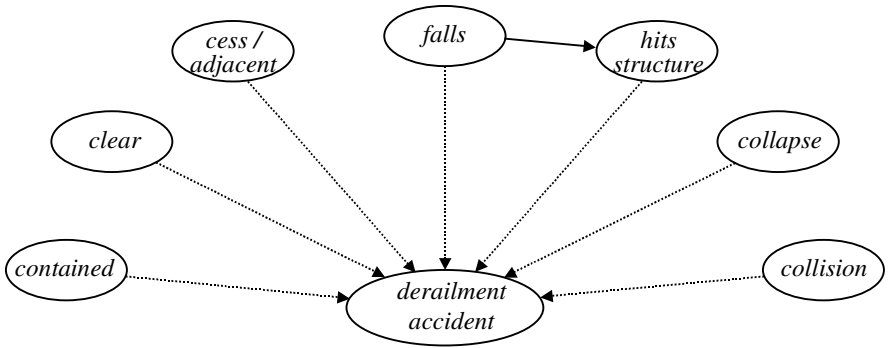


Fig. 7. Equivalent BN for open track derailment

The first point remains true even if the single ‘derailment accident’ node is replaced by a hierarchy of nodes as described in section 0, producing the BN shown in Fig. 8. Although this alternative translation may improve the efficiency of Bayesian propagation, the logical relationship between events and consequence is still more clearly shown in the original event tree.

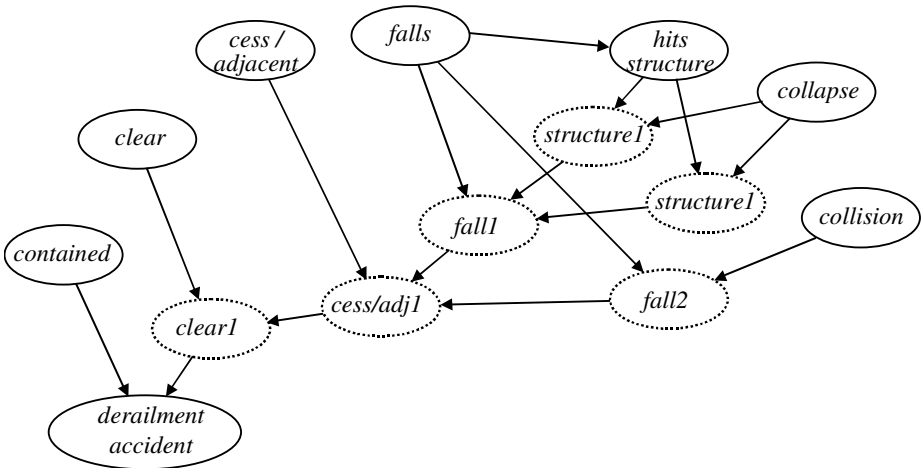


Fig. 8. Alternative translation of open track derailment event tree, using a hierarchy of nodes to encode the logical relationship between events and consequences

It may seem surprising that there is only a single causal arc – from ‘falls’ to ‘hits structure’ between the nodes representing events. This arc occurs because the probability $p(hits | falls = yes) \neq p(hits | falls = no)$. For other events, the probability of each outcome is the same on all the branches. The absence of other causal arcs depends on our treatment of *don't care* conditions. For example, a collision is only possible following a derailment to the adjacent side, but we do not need to represent

this relationship by a causal arc since it is captured by the branching structure of the event tree. Since the two views of the event tree show different information most clearly, we propose to use them together: the BN view is used to ensure that conditional probabilities are handled correctly and the tree view is used for mapping event sequences to consequences. The BN can be shown without the consequence node and arcs, so this part of the BN can be chosen to optimise propagation.

4.3 Generalised Event Trees

As described above, the event tree was originally prepared from a network-wide event tree for derailment accidents. To be applied to an analysis in a specific location, the network-wide model has to be tailored. In this section, we show how a more general model can be represented as a BN, which can be tailored automatically.

The author of the event tree was asked to identify the conditions of the infrastructure and the operation of the railway that influence a derailment accident. Table 2 shows the conditions identified. The causal relationships between these conditions were then elicited together with the probability tables. Fig. 9 shows the resulting BN, with the consequence node and arcs omitted for clarity.

Table 2. Derailment Operating and Infrastructure Conditions

<u>Conditions</u>	<u>Description</u>
Fitted	Whether the derailment containment is fitted: Yes, No
Curvature	The curvature of the track: Severe, Mild, None
Number of tracks	The number of adjacent tracks: 2, 4
Track Speed	The running speed of the track (mph): 0-10, 10-30, 30-60, 60>
Derailment Speed	The speed of the derailment (mph): >15, <15
Lineside Density	The density of objects beside the line: High, Low
Lineside Type	The type of equipment beside the line: Fixed, Anchored
Density of Traffic	The traffic density: High, Low
Peak	The time of day when the incident occurs: Peak, Off peak
Passenger Loading	How full the coaches are: >50%, <50%
Crashworthiness	The crashworthiness of the train: High, Low
Rolling Stock	The type of rolling stock: High Speed Train, EMU

The relationships in the model are causal. For example, a train derailing on a tight curve will be more likely to exceed its clearances while one travelling in a straight line is more likely to maintain its clearances, as its momentum will tend to carry it forward in the direction of travel. The probability table for the event ‘clear’ (whether the train maintains clearance in a derailment) is:

Derailment Speed	> 15 mph			<15mph		
	None	Mild	Severe	None	Mild	Severe
Yes	0.75	0.6	0.29	0.9	0.7	0.4
No	0.25	0.4	0.71	0.1	0.3	0.6

The values 0.29 and 0.71 are taken from the original event tree (Fig. 6), since the circumstance of the original study were ‘Derailment Speed’ > 15 mph and severe track ‘Curvature’. The author of the original event trees judged the other probabilities: although the generalised model requires more such judgements they are similar to those needed to construct an event tree.

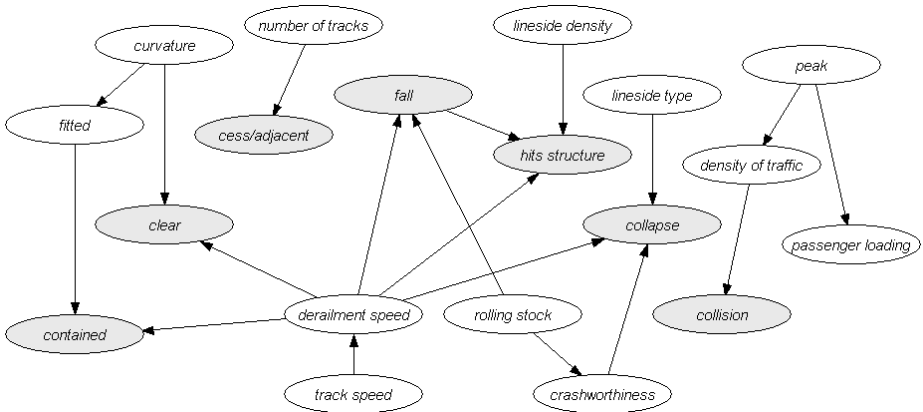


Fig. 9. Derailment BN generalised with the factors that determine the event probabilities. Event nodes are shaded; the consequence node and arcs are not shown.

The generalised model can be used to calculate the accident probabilities in different scenarios. We can compare the scenario in the original study (a dense urban line) with a scenario more typical of an inter-city line:

	Urban Scenario	Inter-city Scenario
Fitted	‘No’	‘No’
Curvature	‘Severe’	‘None’
Number of Tracks	4	2
Derailment Speed	‘> 15’ mph	‘> 15’ mph
Lineside Density	‘High’	‘Low’
Lineside Type	‘Anchored Equipment’	‘Fixed Equipment’
Rolling Stock	‘EMU’	‘High Speed Train’
Density of Traffic	‘High’	‘Low’

These data can be entered into the BN and new event probabilities calculated. The probabilities (relative to the probability of the initial derailment event) of the derailment accidents for the two scenarios are shown in Fig. 10. In the new scenario the less severe accidents are more likely: this results mainly from the absence of curvature. However, following the original study, we have considered only two possible derailment speed ranges and this should be re-examined before drawing any real conclusions. We also note that speed is a factor in the severities (equivalent fatalities) of the accidents, which are estimated using another method.




Fig. 10. Accident probabilities for two scenarios calculated using the BN. The ‘urban’ scenario is identical to the original derailment study giving the same probabilities as the event tree; the hypothetical scenario shows an example of the use of the generalised BN to adapt the accident analysis to different circumstances.

5 Discussion

5.1 Summary

We have shown how a BN can represent an event tree. The translation from BN to event tree is automatic (though we have not yet automated it) and reversible. We argue that the two notations are complementary and should be used together. The event tree shows the logical relationship of events, which is not shown clearly on the BN diagram where it is encoded in a probability table. On the other hand, the BN diagram shows clearly where event probabilities are conditioned on earlier events.

A greater advantage of using a BN is that the accident model can be generalised by including the conditions that influence the evolution of the events in the accident. This generalisation reverses the process used originally to analyse derailments in our case study, where an event tree for a specific location was developed from a network-wide model. The original author of the event tree remarked on the value of analysing causal influences on the events and was lead to re-examine some of the allocated probabilities.

It is advantageous to retain the familiar event tree notation when building the more general accident model. In the case study we were easily able to explain our approach to the author of the derailment event tree: only a short explanation of BNs was needed for this analyst to identify influencing factors and the causal relationships between them. Of course, generalising the accident model in the way we have shown is not automatic. A rigorous elicitation process is needed to understand the influences:

some remained unresolved in our case study, for example the influence of the train weight on the probability of the train falling over in a derailment. The process of judging probabilities for the BN, though time consuming, is similar to that required for building an event tree though potentially many more probabilities are needed.

The validity of the network-wide SRM rests on its use of historic accident data and it is desirable that an accident model for a specific location should have the same basis. At present, the SRM does not include influencing factors although the potential advantages of generalising it have been noted [7]. Clearly, further investigation of the cost-benefit of building such a model is needed.

5.2 Related Work

Others have used BNs to analyse risk. The SCORE project [8] has applied a BN to model accidents in an air-traffic control case study, based on a barrier model of accidents. In [9], an influence diagram is used to model the occurrence of rail breakage, also starting from a barrier model. In both cases the BN replaces the accident model used as a starting point – a barrier model rather than an event tree – rather than providing an alternative view as we have described.

Organisational and management causes of accidents are modelled using BNs in [10] and [8]. Organisational and management causes are examples of ‘influencing factors’ that could be included in our generalised event trees, so both are generalised representations of accidents, but without the connection to an underlying accident model such as an event tree, in the way we have proposed.

The SABINE emergency planning system [11] for accidents in nuclear power plants uses BNs. Part of this system is an accident diagnosis BN, derived from event trees constructed for level 2 PSA. Accident diagnosis requires back propagation from effects to causes and this is prevented by our simple and automatic treatment of *don't care* conditions (section 3.5) which may hide further causal relationships between event outcomes; rather than minimising the number of causal arcs in the BN, we could maximise it, including a causal arc wherever this is possible. We have not followed this approach because diagnosis is not required in our case study.

5.3 Further Work

The derailment study included six separate event trees for different areas of the infrastructure: we are examining how to merge these models. Existing software tools do not allow the event tree and BN views of the accident to be combined conveniently: we would like to investigate how to automate this in practice.

More fundamentally, some of the operating and infrastructure conditions also influence the causes of the initiating event: this is important because such factors introduce correlations between the probability of the initial event and the probabilities of different accident sequences. The present analysis does not capture such correlation and this could lead to an incorrect estimate of the risk. We plan to examine this in future.

Acknowledgement. We are grateful to Colin Howes of Atkins Rail for assistance with the case study, and to the referees' for their comments.

References

1. US Nuclear Regulatory Commission: Reactor Safety Study. WASH-1400, NUREG 75/014 (1975)
2. Jensen, F.V.: *An Introduction to Bayesian Networks*. UCL Press, London (1996).
3. Neil M., Fenton N.: *Building Large Scale Bayesian Networks*. Knowledge Engineering Review. 15(3) (2000) 257-284
4. Neil M., Fenton N., Forey S., Harris R.: *Using Bayesian Belief Networks to Predict the Reliability of Military Vehicles*. IEE Computing and Control Engineering J 12(1) (2001) 11-20.
5. Fenton, N.E., Krause, P., Neil, M.: *Software Measurement: Uncertainty and Causal Modelling*. IEEE Software 10(4) (2002), 116-122.
6. The Safety and Standards Directorate: *Safety Risk Model (SRM)*, Report No. SP-RSK-3.1.3.8 (1999) Rail Safety and Standards Board, UK.
7. Bedford, T., French, S., Quigley, J.: *Statistical Review Of The Safety Risk Model (WP1)*, Report T127 (WP1) (2004) Rail Safety and Standards Board, UK.
8. Neil, M., Malcolm B., Shaw R.: *Modelling an Air Traffic Control Environment Using Bayesian Belief Networks*. In: Proc. 21st Systems Safety Conference, Ottawa, ISBN 0-9721385-2-8.
9. Vatn, J., Svee, H. A.: *Risk Based Approach to Determine Ultrasonic Inspection Frequencies in Railway Applications*. Proceedings of the 22nd ESReDA Seminar, Madrid, Spain May 27-28, 2002
10. Roelen, A.L.C., Wever, R., Cooke, R.M., Lopuhaä, H.P., Hale, A.R., Goossens, L.H.J.: *Aviation Causal Model using Bayesian Belief Nets to Quantify Management Influence. Safety and Reliability (Bedford & van Gelder eds.)*, Swets & Zeitlinger, Lisse, ISBN 90 5809 551 7, (2003) 1315-1320.
11. Zavisca M., Kahlert H., Khatib-Rahbar M., Grindon E., Ang M.: *A Bayesian Network Approach to Accident Management and Estimation of Source Terms for Emergency Planning*. In Proceedings of PSAM7/ESREL '04, Springer-Verlag (2004).