

Publication IV

Jirka Poropudas and Kai Virtanen. 2009. Influence diagrams in analysis of discrete event simulation data. In: Manuel D. Rossetti, Raymond R. Hill, Björn Johansson, Ann Dunkin, and Ricki G. Ingalls (editors). Proceedings of the 2009 Winter Simulation Conference (WSC 2009). Austin, Texas, USA. 13-16 December 2009. Pages 696-708. ISBN 978-1-4244-5771-7.

© 2009 Institute of Electrical and Electronics Engineers (IEEE)

Reprinted, with permission, from IEEE.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of Aalto University's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Proceedings of the 2009 Winter Simulation Conference

M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, eds.

INFLUENCE DIAGRAMS IN ANALYSIS OF DISCRETE EVENT SIMULATION DATA

Jirka Poropudas
Kai Virtanen

Systems Analysis Laboratory
Helsinki University of Technology
P.O. Box 1100, FIN - 02015 HUT, FINLAND

ABSTRACT

In this paper, influence diagrams (IDs) are used as simulation metamodels to aid simulation based decision making. A decision problem under consideration is studied using discrete event simulation with decision alternatives as simulation parameters. The simulation data are used to construct an ID that presents the changes in simulation state with chance nodes. The decision alternatives and objectives of the decision problem are included in the ID as decision and utility nodes. The solution of the ID gives the optimal decision alternatives, i.e., the values of the simulation parameters that, e.g., maximize the expected value of the utility function measuring the attainment of the objectives. Furthermore, the constructed ID enables the analysis of the consequences of the decision alternatives and performing effective what-if analyses. The paper illustrates the construction and analysis of IDs with two examples from the field of military aviation.

1 INTRODUCTION

Discrete event simulation is a flexible and widely used methodology for the analysis of complex real world systems (Law 2006). In simulation analysis, a simulation experiment is conducted by running simulations with alternative values of simulation parameters that determine the properties of the studied system. The values of the simulation parameters and the corresponding output are recorded to produce the simulation data that are used to estimate, e.g., descriptive statistics, empirical distributions of the simulation output, or statistical models that describe the dependence between the simulation parameters and the simulation output (Law 2006). Furthermore, in order to support decision making, the simulation data can be used for finding the simulation parameters that describe optimal system settings or configurations. The optimization is carried out either directly through simulation based optimization (e.g., Fu, Glover, and April 2005, Ólafsson and Kim 2002, Swisher and Jacobson 1999) or by using simulation metamodels (e.g., Cheng and Currie 2004, Friedman 1996, Barton 1998) which refer to simplified and analytical models representing the simulation model. Such approaches are used to determine the simulation parameters that result in the desired simulation outcome or produce optimal simulation output with respect to a given optimization criterion. In general, the optimization criteria of such analyses depend on the final outcome of the simulation or measure the system performance using time averages. Thus, they omit the time evolution of the simulation and give no explicit information on the progress of the simulation with different values of the simulation parameters.

In this paper, simulation based decision making refers to the use of simulation in the analysis and solution of decision problems under uncertainty. The definition of the decision problem consists of available decision alternatives and an objective describing the preferences of the decision maker. The alternatives are compared using a suitable simulation model that represents the uncertain, random, or uncontrollable factors that may be an integral part of the decision problem. The alternatives are fed to the simulation model as simulation parameters and the resulting simulation output is collected. Then, based on the simulation output, one can solve the optimal decision alternative or rank the alternatives based on a given optimization criterion associated with the objective of the decision problem.

The simulation data can be used, in addition to the optimization of the simulation output, to better understand what actually happens during the simulation. In such case, the data should be presented in a form that describes the evolution of the simulation, i.e., how the simulation state changes in time. For increased analytical capability, the presentation should also include the interdependencies between the simulation states at different time instants – preferably in a probabilistic

manner in order to represent the uncertainties related to simulation. Poropudas and Virtanen (2007) presents a method for analyzing the evolution of discrete event simulations with dynamic Bayesian networks (DBNs, e.g., Neapolitan 2004, Jensen and Nielsen 2007, Pearl 1986). In the presented approach, chance nodes of a DBN and their interconnecting arcs are used to describe the joint probability distribution of the random variables that presents the evolution of the simulation state during simulation. The analysis of such networks has been found more effective and less time consuming than the brute force analysis of raw simulation data (Poropudas and Virtanen 2007). Unfortunately, DBNs are only descriptive and offer no direct way for optimizing the simulation output.

In this paper, the above mentioned analysis is extended to simulation based decision making by constructing influence diagrams (IDs, Howard and Matheson 1984, Howard and Matheson 2005) from simulation data. The IDs are probabilistic models that are used for structuring and solving decision problems under uncertainty (e.g., Diehl and Haines 2004, Virtanen, Raivio, and Hämäläinen 1999). In the construction of an ID, a DBN is first built by combining simulation data with expert knowledge (Neapolitan 2004, Poropudas and Virtanen 2007) to present the evolution of the simulation state. Then, the DBN is extended into an ID by including the simulation parameters into the model to represent the parameters' effect on the simulation output and the changes in the simulation state. In practice, the parameters are represented by decision nodes that have a discrete set of states corresponding to the available decision alternatives.

To enable the comparison of decision alternatives, the preferences of the decision maker are included into the ID by adding a utility node and a corresponding utility function (von Neumann and Morgenstern 1944). The utility function is defined as a function of the simulation output that measures the attainment of the objectives of the decision problem. Such an approach is versatile because different utility functions offer almost unlimited possibilities for modeling the preferences of the decision maker. The constructed ID is then solved to find the optimal decision alternatives that maximize the expected value of the utility function. The ID is also used for studying the evolution of the simulation and conducting what-if analyses. Thus, the ID methodology offers an approach to simulation metamodeling that makes the simulation analysis of decision problems more transparent and gives valuable information about the consequences of decision alternatives.

The paper is organized as follows. The construction of IDs from simulation data is presented in Section 2. In Section 3, potential ways for utilizing the IDs are outlined. The utilization of IDs is illustrated in Section 4 with two examples of simulation analysis involving the maintenance of military aircraft and a 1 vs. 1 air combat engagement. Finally, conclusions are made and future research topics are proposed in Section 5.

2 CONSTRUCTION OF INFLUENCE DIAGRAMS FROM SIMULATION DATA

2.1 Influence Diagrams

IDs (Howard and Matheson 1984, Howard and Matheson 2005) are models that give graphical and analytical means for outlining and solving decision problems under uncertainty. IDs can be used to model the dependence relations between available decision alternatives and the objectives of the decision problem. Moreover, if the decision problem involves factors that are beyond the decision maker's control, they too can be incorporated in the ID. The objectives are described using a utility function that includes both the preferences of the decision maker as well as her attitude towards risk (von Neumann and Morgenstern 1944). The IDs are useful for structuring decision problems and, more importantly, they can be solved with effective algorithms to obtain the combination of decision alternatives that is most preferred by the decision maker (Shachter 1986, Jensen and Nielsen 2007).

The definition of an ID is based on a directed acyclical graph over decision, chance, and utility nodes (see, Fig. 1). Each decision node of the ID, denoted by $d(j)$, $j \in \{1, \dots, m\}$, has a discrete set of states, denoted by $D(j)$, presenting the available decision alternatives at the given node. The decision nodes are ordered so that the decisions are made in sequence and the previous decisions are recalled when making further decisions. The arcs of the ID are used to model both the dependencies between nodes and the flow of information. Mainly, if there is an arc from a chance node to a decision node, the state of the chance node is known at the time of the decision.

Each chance node $x(i)$, $i \in \{1, \dots, n\}$, of the ID has a discrete set of states, denoted by $X(i)$, and a probability table that gives the conditional probabilities of these states when the states of the parents of the node, i.e., the nodes with an arc pointing to the node of interest, are known. The chance nodes are used to model the framework and environment of the decision problem by representing the joint distribution of random variables, i.e., the probabilities $P(x(1) = x_1, \dots, x(n) = x_n)$ for all possible combinations of states x_1, \dots, x_n , where each $x_i \in X(i)$. The random variables present uncertain factors that affect the outcome of the decision problem but can not be directly controlled by the decision maker. Yet, the probability distributions of these variables may depend on the chosen decision alternatives. In other words, if there is an arc from a decision node to a chance node, the probability distribution of the chance node depends on the state of the decision node.




Figure 1: Example of an ID with two decision nodes, $d(j)$, $j \in \{1,2\}$; three chance nodes, $x(i)$, $i \in \{1,2,3\}$; and utility node, u . Decision $d(1)$ is made without knowledge of the chance nodes’ states. At the time of decision $d(2)$, the decision maker knows the state of chance node $x(2)$ and remembers the decision made at $d(1)$. The state of the utility node depends on the decision made at $d(2)$ and the state of chance node $x(3)$.

A utility node u is used to describe the objective of the decision maker. In this paper, only cases with one objective, i.e., a single utility node, are discussed. IDs could also be used for analyzing decision problems with multiple objectives represented by several utility nodes (Diehl and Haimes 2004). The utility node has a set of states, denoted by U , that represent different outcomes of the decision problem and depend on the states of the decision and chance nodes with an arc pointing to u . The decision maker’s preference between these outcomes is modeled using a utility function $f(\cdot) : U \rightarrow [0, 1]$ where the least and most desirable outcomes are scaled to values 0 and 1, respectively. If the decision maker, whose preferences are represented by the utility function, is prepared to accept the axiomatic definition of rationality as given by utility theory (see, e.g., Keeney and Raiffa 1993), her objective is set as the maximization of the expected value of the utility function (von Neumann and Morgenstern 1944). Thus, when making a decision, the decision maker selects the state of a decision node that maximizes her expected utility.

It should be noted that technically there is no other distinction between decision nodes and chance nodes except that the decision maker can choose the states of the decision nodes. On the other hand, if the chance nodes of the ID are considered independently from the rest of the diagram, they form a Bayesian network (BN, also known as Bayesian belief network, see, e.g., Pearl 1986, Jensen and Nielsen 2007). When a BN is divided into disjoint time slices, where each time slice consists of random variables realized at the same time instant, it is called a dynamic Bayesian network. By connecting the time slices with arcs that represent the dependence between random variables realized at different times, the DBNs can be applied for modeling time dependent phenomena. In this paper, the chance nodes of an ID form a DBN that represents the time evolution of the simulation and gives the probability distribution of the simulation state at different times.

There are many methods for solving an ID and finding the optimal decision alternatives for all decision nodes (see, e.g., Jensen and Nielsen 2007). Generally, the combination of decision alternatives that maximizes the expected utility is sought. On the other hand, the ID provides the probability distribution for the states of the utility node conditional on a given combination of decision alternatives. Thus, other criteria could also be used for comparing the decision combinations. Regardless of the method used and the optimization criterion, the optimal solution of the ID is found in feedback form that gives the optimal decision alternative for each decision node as a function of its parents’ states. In practice, the solution of the ID is obtained with computationally effective algorithms (e.g., Shachter 1986, Shachter 1988, Jensen and Nielsen 2007) that are already implemented in available software such as GeNIe (Decision Systems Laboratory 2009) and Hugin (Andersen, Olesen, and Jensen 1990).

2.2 Construction of Influence Diagrams

When analyzing a decision problem using simulation and an ID, the ID is constructed by combining expert knowledge with automated learning where the probability tables of chance nodes and the structure of the network are estimated from data (e.g., Neapolitan 2004). First, the decision problem at hand is presented by identifying the decision alternatives and the objective. In simulation, the values of simulation parameters correspond to the decision alternatives and simulation output used to measure the attainment of the objective. The simulation state is defined in such a manner that it gives a comprehensive and effective presentation for the decision problem. In practice, the simulation state at each time instant $t \in \{t_0, t_1, \dots, t_n\}$

is represented by a discrete random variable, $x(t)$, that is associated with a chance node in the ID. (If necessary, continuous simulation state variables may be discretized resulting in discrete valued variables.) The time evolution of the simulation state and interdependence between the states at different times are modeled with the conditional probability tables of the chance nodes. This presentation for the simulation state is in fact a DBN (Poropudas and Virtanen 2007).

The evolution of the simulation is affected by the values of the simulation parameters that represent factors that may or may not be controlled by the decision maker. To model this dependence, the simulation parameters are included in the ID. If the factors presented by the simulation parameters are controllable, they are included in the ID as decision nodes. If the parameters can not be controlled, they are added as chance nodes representing external random variables. In such cases, expert knowledge is needed to define the prior distributions of the external variables, i.e., the probabilities of different simulation parameter values.

When the decision and chance nodes of the ID have been determined, the initial structure of the ID is drawn by a subject matter expert. In other words, the nodes are connected with arcs to delineate their interdependence. The conditional probability tables of the chance nodes are estimated using automated learning where the initial structure of the diagram and the pre-processed simulation data are fed to a software such as GeNIe or Hugin. The estimation of probabilities is based on the observed frequencies, i.e., the probability of a combination of states depends on the number of times it has been observed (Neapolitan 2004). The final structure of the diagram is found by comparing various structure alternatives and the best structure is selected, e.g., by using a scoring rule (Neapolitan 2004, Poropudas and Virtanen 2007). The scoring rules are based on trade-off between the size of the model and the fit between the probabilities given by the model and those estimated from the simulation data. If there exists notable lack-of-fit, additional arcs may be added to the initial structure in order to include the observed dependencies. When constructing IDs for continuous time phenomena, the time instants represented by the chance nodes are selected so that the maximal absolute error between the probabilities estimated from the simulation data and those given by the ID is minimized. The optimization problem and its solution are presented in detail in Poropudas and Virtanen (2009).

Finally, a utility node is added to the model. In practice, the dependence between the state of the utility node and the states of the decision and chance nodes is represented with arcs. All combinations of the states of the utility node's parents are associated with a state of the utility node. Each of these states represents a potential outcome of the decision problem that results from the selected decision alternatives and the realized states of the chance nodes. Thus, the ID can be used to calculate the probability distribution of the states of the utility for each combination of the decision alternatives. The decision maker's preference between the states of the utility node is measured using a utility function that can be elicited by, e.g., proposing lotteries between different states and finding the subjective probabilities that render the alternatives equally desirable (Keeney and Raiffa 1993).

3 ANALYSIS OF SIMULATION DATA WITH INFLUENCE DIAGRAMS

The constructed IDs are studied in order to better understand the decision problem and solve the optimal decision alternatives. In this paper, the IDs are used for the following purposes: studying time evolution and tracking of the probability of a given simulation state, conducting what-if analyses, and solving optimal decision alternatives as well as studying the decision alternative's effects on the evolution of simulation. It should be noted that such analyses could also be conducted by brute force estimation and comparison of probability distributions calculated directly from the simulation data. Nevertheless, the utilization of IDs makes the analysis more effective and less time consuming as there is no need for repeated screening of the simulation data (Poropudas and Virtanen 2007). Furthermore, the available software tools can be utilized both to calculate probability distributions and solve the optimal decision alternatives.

Time Evolution and Tracking. The estimated ID gives the joint distribution of the random variables describing the simulation state at times $t \in \{t_0, \dots, t_n\}$. In other words, the probability of the simulation progressing through the states x_0, \dots, x_n at the respective time instants is $P(x(t_0) = x_0, \dots, x(t_n) = x_n)$. The joint distribution can be used to calculate the probability distribution $P(x(t_i) = x)$ over the simulation states $x \in X(i)$ at fixed time instant t_i . The changes in this distribution can be studied visually to better understand the evolution of the simulation (Poropudas and Virtanen 2007). On the other hand, one can simply inspect the probability of the simulation being in a fixed state x at different times $t \in \{t_0, \dots, t_n\}$. These probabilities, $P(x(t_i) = x)$, can be tracked to analyze the likelihood of a given simulation event taking place, i.e., the simulation being in the state x , during the simulation. To compare the available decision alternatives, the ID is used to calculate the conditional probability distributions $P(x(t_0) = x_0, \dots, x(t_n) = x_n | d(j) = d_j)$ and $P(x(t) = x | d(j) = d_j)$ for each decision alternative $d_j \in D(j)$. The conditional probability distributions show the consequences of a given decision alternative by describing the progress of the simulation when that alternative is selected. In cases where the simulation state has a reasonable numerical interpretation, all the above mentioned distributions can be used to calculate descriptive statistics such

as the expected value of the simulation state at time t , $E(x(t))$, that shows the evolution of the average state of the simulated system.

What-if Analysis. The chance nodes of the ID represent the dependencies between simulation states at specific times as well as between simulation parameters and simulation states. The IDs are used to conduct what-if analysis where the model is fed evidence, i.e., the state of the simulation at a given time t or the value of an uncontrollable simulation parameter. Then, the distributions of other variables are updated accordingly to study the effects of this observation, $P(x(t_0) = x_0, \dots, x(t_n) = x_n | x(t) = x)$. The updated distributions show how the occurrence of a given simulation state at given time affects both the progress and the final state of the simulation. The same kind of analysis can also be conducted conditional on decision alternatives available in the node $d(j)$. Then, the conditional distributions $P(x(t_0) = x_0, \dots, x(t_n) = x_n | x(t) = x, d(j) = d_j)$ are calculated to see how the simulation progresses when the alternative d_j has been chosen and the simulation is observed in state x at time t . When applicable, it is possible to use all these conditional probability distributions to calculate the respective conditional expected values of the simulation state, e.g., $E(x(t_n) | x(t) = x)$.

Solution of Optimal Decisions. The ID is solved to see how the decision alternatives affect the distribution of the utility node. The alternatives can be compared to find the one that maximizes the expected utility $E(f(\cdot) | d(j) = d_j)$. On the other hand, the ID retains the information necessary for further analysis by, e.g., studying different utility functions.

Consequences of Decisions. In addition, the ID shows the progress of the simulation for a given decision alternative which increases one's understanding over the decision problem and reveals the actual consequences of the decisions in addition to the resulting expected value of the utility. By representing the evolution of the simulation, the probability distributions produced by the ID show how or why the alternative performs better than others. The ID also gives probabilities for simulation events consequent on the decision alternatives that can be easily compared. For example, one can recognize alternatives that lead to high values of the expected utility but may result in the occurrence of unwanted events during the simulation with an unacceptably high probability. Thus, the analysis of the ID may give additional insight to decision problems in comparison to other approaches used in simulation based optimization (see, e.g., Cheng and Currie 2004, Fu, Glover, and April 2005, Ólafsson and Kim 2002, Swisher and Jacobson 1999).

4 EXAMPLE ANALYSES

Next, two examples on the construction and utilization of IDs in the analysis of simulation data are presented. In the first example, an ID is constructed to study a discrete event simulation model for the maintenance of a military aircraft fleet (Mattila, Virtanen, and Raivio 2008). The ID is used to describe the dependence between the number of personnel in a maintenance unit and the availability of aircraft. In the second example, a discrete event air combat simulation model called X-Brawler is studied (L-3 Communications Analytics Corporation 2002). Here, an ID is constructed for comparing the effectiveness of decision alternatives in 1 vs. 1 air combat. The comparison is carried out based on several alternative objectives in order to see how the optimal decision alternative depends on the objectives. On the other hand, the ID is used to describe the evolution of air combat and study the consequences of the available decision alternatives.

4.1 Aircraft Maintenance Simulation

In this example, the maintenance of military aircraft is studied using a discrete event simulation model (Mattila et al. 2008). The simulation model is presented in detail in Mattila et al. (2008). The example is based on a simplified model with only one air command. The example presents a scenario where a fleet of 10 aircraft is serviced by a maintenance unit within a single air command. The maintenance unit is responsible for turnaround and preflight inspections, minor periodic maintenance, and minor failure repairs such as simple component changes. In the example, an ID is constructed from the simulation data to describe the dependence between the number of personnel in the maintenance unit and the availability of aircraft, i.e., the number of mission capable aircraft, that is used to measure the performance of the maintenance unit.

For demonstration purposes, the number of personnel, denoted by $z \in Z = \{2, 3, 4\}$, is studied as an uncontrollable simulation parameter. In other words, it is taken as a factor that is included in the simulation model only as a parameter and added to the ID as a chance node. The simulation state $x(t)$ at time t , $x(t) \in X(t) = \{0, 1, \dots, 10\}$, is set as the availability of the aircraft at the beginning of the day t . At the beginning of the simulation there are nine available aircraft, i.e., $x(0) = 9$. The state of the simulation is checked daily for 20 days, $t \in \{0, 1, \dots, 20\}$. To model the interdependencies between the simulation states as well as the dependence between the simulation parameter and the simulation states, the structure of the ID is defined as shown in Fig. 2.

To gather the necessary simulation data, the simulation is replicated 10000 times for each value of z . The simulation data is fed into the GeNle software to estimate the conditional probability tables for the chance nodes $x(t)$. The software




Figure 2: The ID for the aircraft maintenance example. The progress of the simulation is described with chance nodes $x(t)$. The simulation state $x(t)$ depends only on the previous state and the number of personnel in the maintenance unit, $z \in Z = \{2, 3, 4\}$.

is also used to confirm that the structure of the ID is satisfying and there are no statistically significant dependencies that would require additional arcs. The prior distribution of the external variable z can not be estimated from simulation data. Therefore, it is assessed by an expert and defined as $P(z = 2) = P(z = 3) = 0.20$ and $P(z = 4) = 0.60$. In other words, the most probable alternative is that the maintenance unit operates at full capacity and the other two alternatives are equally probable.


To analyze the evolution of the simulation, the constructed ID is used to calculate the probability distribution of the number of available aircraft $P(x(t) = x|z)$ for all time instants t and values of z . The conditional expected values for these distributions $E(x(t)|z)$ for each value of z are presented in Fig. 3(a) where the continuous lines present the values estimated from the simulation data and circles with the dotted lines represent the values given by the ID. The red curve presents the expected value of availability given by the ID that incorporates the expert knowledge about the distribution of z to the simulation data. In other words, the curve combines the expert knowledge with the simulation data giving a comprehensive overall picture of the simulated scenario.

The dependence between the performance of the maintenance unit and the number of personnel is studied in a what-if type analysis, see Fig. 3(a). There is a slight difference between the curves representing values $z = 3$ and $z = 4$, but in effect they act in a similar manner to the overall expected value of availability. On the other hand, the maintenance unit with only two mechanics, i.e., $z = 2$, struggles to service the fleet and the expected value of availability dips heavily in the second half of the simulation.

The ID also enables further what-if analysis where the state of the simulation is observed at a given time and the consequences of this observation are calculated. For example, the ID is fed evidence that $x(10) = 8$, i.e., it is assumed that there are eight aircraft available at the beginning of day 10. The probabilities of the evidence for different values of z are $P(x(10) = 8|z = 2) = 0.189$, $P(x(10) = 8|z = 3) = 0.229$, and $P(x(10) = 8|z = 4) = 0.246$. Thus, the probability of the observed simulation state increases along the number of personnel. The resulting conditional expected values of availability are presented in Fig. 3(b). In the case of two mechanics, the expected value of availability dips considerably at the latter stages of the simulation, despite there being eight available aircraft at time $t = 10$.

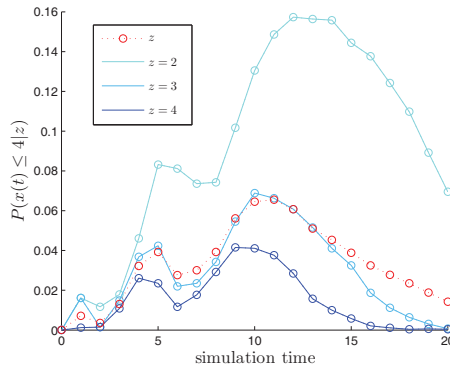
In addition to the estimates of the expected availability that have already been studied in Mattila et al. (2008), the ID gives the probabilities of states $P(x(t) = x)$ and the conditional probabilities of states $P(x(t) = x|z)$ in an easily handleable and compact form. For example, if the number of available aircraft falls to four or less, i.e., $x(t) \leq 4$, no more than two pairs of aircraft can be deployed to carry out missions. The ID can be used to effectively calculate the probability of the occurrence of such an event for different sizes of personnel, as presented in Fig. 3(c). The figure shows that, in the case of two mechanics, the probability peaks at $P(x(12) \leq 4|z = 2) = 0.157$ and the performance capacity of the fleet may be found inadequate at this time.

To summarize the findings, the analysis of the ID constructed from simulation data shows that the availability of the aircraft is heavily influenced by the number of personnel. Based on the prior probability distribution of the number of personnel, the expected value of availability remains mostly above seven aircraft and the probability of less than four available aircraft is low. There is no clear difference between the expected value of availability and the conditional expected values with three or four mechanics, even though the maintenance unit with four mechanics performs best. On the other hand, if there are only two mechanics in the maintenance unit, the performance degrades noticeably. The expected value of availability dips heavily later in the simulation. This is true even for the case where there are eight available aircraft on the tenth day.



(a) The expected value of availability $E(x(t))$ and the conditional expected values of availability $E(x(t)|z)$ for different numbers of personnel.

(b) The conditional expected value of availability $E(x(t)|x(10)=8)$ and the conditional expected values of availability $E(x(t)|x(10)=8, z)$ for different numbers of personnel.



(c) The probability of "at most 4 aircraft available", $P(x(t) \leq 4)$ and the conditional probabilities $P(x(t) \leq 4|z)$ for different numbers of personnel.

Figure 3: Performance of the maintenance unit. The continuous lines depict the estimates calculated from the simulation data. The circles and dotted lines present the corresponding estimates given by the ID. The red curve gives the corresponding estimates when the prior distribution of z is included in the ID.

In fact, with only two mechanics there is a remarkably high probability of less than four available aircraft on the latter half of the simulation.

The utilization of the ID gives several advantages to the analysis of the simulation data. The fit between the ID and the simulation data is almost perfect as the continuous and dotted blue curves in Fig. 3 are indistinguishable from each other. This refers to the validity of the model and it would appear that the ID incorporates the essential information included in the simulation data. In what-if analyses, the ID can be used to calculate conditional probabilities effortlessly and accurately. The ID also shows the time dependence of the simulation state in a compact manner. Furthermore, by using the ID, expert knowledge can be seamlessly joined with the simulation data for more extensive analysis that includes factors that are only modeled as parameters in the simulation model.

4.2 Air Combat Simulation

In the second example, an air combat scenario is analyzed using a discrete event air combat simulation model called X-Brawler (L-3 Communications Analytics Corporation 2002). The scenario presents a 1 vs. 1 air combat where blue chooses between two alternative tactics at the beginning of the engagement. The scenario is studied using generic unclassified models for blue and red aircraft as well as for missiles, radars, and other hardware.

The decision problem analyzed in the scenario concerns the maneuvering of blue and is described as follows. The initial geometry of the scenario is advantageous for blue as it approaches red from behind near the launch area of the medium range missiles. In other words, at the start of the scenario blue is preparing to launch its missile towards red. On the other hand, blue has only one semi-active medium range missile at its disposal while red carries two such missiles. Therefore, blue will be at a considerable disadvantage if its missile misses the target. In addition to the medium range missiles, both aircraft are equipped with short range missiles that are used if the engagement leads to combat within visual range.

Suppose now that blue has two decision alternatives: an aggressive and defensive tactic. In the aggressive alternative, blue launches its missile and continues to approach towards the opponent in order to engage within visual range combat and use its short range missiles if the initial missile is unsuccessful. If the defensive alternative is chosen, blue launches its missile and performs a drag maneuver, i.e., flies away from the opponent, before continuing the engagement. This maneuvering should allow blue to avoid the first missile that is possibly launched by red.

To analyze the decision problem, the ID presented in Fig. 4 is constructed. In the ID, decision node d corresponds to the decision of blue, where the set of available decision alternatives is $D = \{\text{aggressive, defensive}\}$. To model the evolution of the simulation, the state of the air combat at time t is described by a random variable $x(t)$ that corresponds to a chance node of the ID. The set of states for variable is $X(t) = \{\text{neutral, blue advantage, red advantage, mutual disadvantage}\}$ where the definitions of the states are

- *neutral*, both sides are alive
- *blue advantage*, blue is alive and red has been killed
- *red advantage*, blue has been killed and red is alive
- *mutual disadvantage*, both sides have been killed

In other words, the variable $x(t)$ reflects the outcome of the simulated air combat up to the time t .




Figure 4: The ID for the air combat example. Decision node d represents the tactic chosen by blue. The progress of the simulated air combat is modeled with chance nodes $x(t)$. The simulation state $x(t)$ depends only on the previous state and the decision taken by blue. The objective of the decision problem is described by utility node u .

The structure of the ID, i.e., the arcs between the decision nodes and the chance nodes, is determined by a subject matter expert. The structure implies that the simulation state at time t depends on the previous simulation state and the decision taken by blue. In addition, decision d is made before the start of the engagement without any knowledge of the simulation states at latter time instants.

To gather the necessary simulation data for constructing the ID, the scenario is simulated 10000 times for both decision alternatives. The time instants represented by the chance nodes are selected by minimizing the maximal absolute error between the probabilities estimated from the simulation data and those given by the ID. The optimal solution corresponds to time instants $t \in \{0, 122, 135, 146, 157, 202, 237, 500\}$ (in seconds). The GeNIe software is used to estimate the conditional probability tables of $x(t)$ and to determine that no additional arcs are needed to represent the dependencies found in simulation data.

To complete the ID, utility node u is added. In the scenario, the objective of the decision problem depends only on the final state of the air combat. Thus, the set of states for u is defined similarly to chance nodes $x(t)$ of the ID, i.e., $U = X(t)$. The dependence between the utility node and other nodes is defined simply by setting the state of u as equal to the final state of the simulation, i.e., $x(500)$.

The constructed ID is used to calculate the probability distribution of simulation state $x(t)$ for all time instants t . In Figs. 5(a) and 5(b), the probability distributions estimated from the simulation data are presented with the continuous lines and the dotted lines with circles correspond to the probabilities given by the ID. Evidently, the ID gives an accurate representation of the simulation data and there are no noticeable differences between the curves.




Figure 5: The probability distribution of the air combat state as a function of time for tactics $d \in \{\text{aggressive, defensive}\}$. The continuous lines depict the probabilities calculated from the simulation data. The circles and the dotted lines present the corresponding probabilities given by the ID.

Fig. 5(a) shows how the simulated air combat progresses when the decision alternative "aggressive" is chosen. The first observation is that the probability curves remain nearly unchanged before time instant 122 and after 237, i.e., the most important changes in the simulation state take place between these times. There is a sharp peak in the probability of the state "blue advantage" between time instants 135 and 157. This is explained by blue launching its missile first and shooting down red with probability 0.462 before time instant 146. Nevertheless, in this scenario, red has time to launch a missile before being shot down. This leads to a 0.409 probability that blue is also killed before time instant 157 and the air combat ends in the state "mutual disadvantage" with probability 0.416.

In the end, there is only a 0.141 probability for blue winning the engagement. Furthermore, even if blue is not successful in shooting down red with its medium-range missile, there is a 0.397 probability of the air combat resulting in the state "red advantage" where blue is shot down later in the engagement. Overall, there is only a 0.046 probability of both sides surviving the engagement and the simulation ending in the state "neutral".

Fig. 5(b) shows the respective probability curves for the decision alternative "defensive". Unlike the previous case, now there is no peak in the probability of the state "blue advantage". Instead, the probability of "blue advantage" rises sharply to 0.372 by the time instant 157 representing a hit by the blue missile. After this early rise, the probability of the state "blue advantage" continues to climb towards the final value of 0.445. Thus, by performing the drag maneuver, blue is able to avoid the missile that red may have launched before being shot down. On the other hand, there is a 0.477 probability that red will shoot down blue with its second medium range missile later, i.e., the red holds a considerable advantage in the following engagement. Overall, there is a 0.072 probability of both aircraft surviving the combat and only 0.007 chance that both aircraft are shot down.

In addition to describing the evolution of the simulation, the ID is used to calculate the probability distribution of the states of the utility node for both decision alternatives (see, Table 1). The values of the utility function associated with the states of the utility node can be defined in many ways to represent the alternative objectives that the decision maker might have. In the discussed case, the least preferred outcome for blue is clearly "red advantage" and the most preferred "blue advantage". This leaves only the values of the utility function for the states "neutral" and "mutual disadvantage" to

Table 1: The probability distribution of the states of the utility node for each decision alternative.

Tactic, d	State of utility node u			
	<i>blue adv.</i>	<i>red adv.</i>	<i>mut. dis.</i>	<i>neutral</i>
aggressive	0.141	0.397	0.416	0.046
defensive	0.445	0.477	0.007	0.072

be determined. The utility functions are scaled so that the least and most desirable states produce utilities equal to 0 and 1, respectively.

In Table 2, alternative utility functions are presented. Utility function $f_1(\cdot)$ represents a decision maker whose only objective is the survival of blue, i.e., the decision maker receives the highest utility in those states of the utility node where blue has survived the engagement. Conversely, $f_2(\cdot)$ represents a decision maker who is only interested in shooting down red and receives the maximal utility in the states where red has been shot down. Utility function $f_3(\cdot)$ measures the probability of blue winning the air combat. In the utility function $f_4(\cdot)$, the states "neutral" and "mutual disadvantage" are taken equally preferred and their utility is half of the utility received from "blue advantage". Utility function $f_5(\cdot)$ gives an example of a decision maker who is indifferent between the decision alternatives.

Table 2: The comparison of five utility functions. All the utility functions take values 0 and 1 in the states "red advantage" and "blue advantage", respectively. The different values associated with the states "neutral" and "mutual disadvantage" represent alternative objectives. For each utility function, the optimal decision is solved by finding a decision alternative that maximizes the expected value of the utility function.

Utility function	Value of utility function				Expected utility, $E(f_k(\cdot) d)$	
	<i>blue adv.</i>	<i>red adv.</i>	<i>mut. dis.</i>	<i>neutral</i>	<i>aggressive</i>	<i>defensive</i>
$f_1(\cdot)$	1	0	0	1	0.187	< 0.517
$f_2(\cdot)$	1	0	1	0	0.557	> 0.452
$f_3(\cdot)$	1	0	0	0	0.141	< 0.445
$f_4(\cdot)$	1	0	0.5	0.5	0.372	< 0.485
$f_5(\cdot)$	1	0	0.8	0.9	0.515	= 0.515

Table 2 also presents the optimal decision alternatives that are solved by comparing the expected utilities of the decision alternatives. In the case of $f_1(\cdot)$, $f_3(\cdot)$, and $f_4(\cdot)$, the decision maker prefers the defensive tactic because it produces the highest expected utility. In the case of $f_2(\cdot)$, the aggressive tactic is found to be preferable. If the utility function is $f_5(\cdot)$, the decision maker is indifferent between the two alternatives, i.e., both tactics give the same expected utility.

In addition to solving the optimal decision alternative for a given utility function, the ID can be fed evidence of the simulation state to conduct what-if analysis and study the potential consequences of the decision alternatives. For example, if the simulation is in the state "neutral" at time 157, it is known that both aircraft have missed with their first missiles. The probabilities of such evidence are $P(x(157) = \text{neutral} | d = \text{aggressive}) = 0.164$ and $P(x(157) = \text{neutral} | d = \text{defensive}) = 0.616$. Thus, the both aircraft are more likely to survive up to this time instant if the defensive alternative is chosen. In the case of the aggressive tactic, Fig. 6(a) implies that blue remaining within visual range combat is slightly favorable to red who has a 0.382 probability of shooting down blue and surviving, compared to the 0.303 probability for blue. On the other hand, Fig. 6(b) shows that the defensive tactic is disastrous for blue if both sides have survived the first missiles. In this case, red is going to win the air combat with a 0.754 probability, compared to the 0.119 probability for blue. Thus, if blue uses the defensive tactic and both aircraft miss with their first missiles, the increased distance between aircraft is highly disadvantageous for blue due to its lack of additional medium range missiles.

The results provided by the ID can be summarized as follows. The defensive tactic is preferable in the cases where the decision maker values the survival of blue. With this tactic, blue can avoid the first missile of red even though it leaves blue in a disadvantageous position should the blue missile miss its target. On the other hand, if the objective of the decision maker is to shoot down the red aircraft, the aggressive tactic should be selected. Unfortunately, this tactic gives little chance for the survival of blue and is likely to result in both aircraft being shot down.

In this example, the ID constructed from the simulation data was found helpful in many ways. It gives an accurate representation of the simulation data as the continuous and dotted blue curves in Figs. 5 and 6 coincide very well – especially at the time instants included in the ID. The ID also provides an informative presentation for the evolution of the simulated

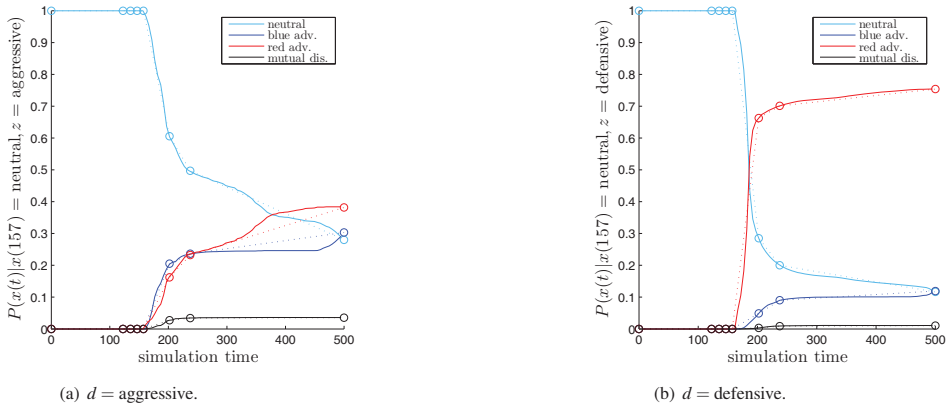


Figure 6: The probability distribution of the air combat state as a function of time for tactics $d \in \{\text{aggressive, defensive}\}$ conditional on $x(157) = \text{neutral}$.

AC. Most importantly, the optimal decision alternatives are easily solved for a variety of utility functions with the help of suitable software. The ID also reveals the consequences of the decision alternatives by enabling various what-if analyses that can be conducted fast and effortlessly.

5 CONCLUSIONS

In this paper, IDs are combined with discrete event simulation in order to solve decision problems. The paper presents both the construction of IDs from simulation data and their utilization in the analysis of decision problems. In such an analysis, a suitable simulation model is used to gather the necessary simulation data for the construction of an ID. Then, the simulation state and its probabilistic changes during the simulation are described by random variables that are represented by the chance nodes of the ID. The decision alternatives related to the simulation parameters are included in the ID as decision nodes and the objective of the decision problem as a utility node with an appropriate utility function. If the effects of uncontrollable factors are studied, the simulation parameters are included in the ID as external variables with prior distributions given by an expert. The construction of the ID involves expert input but is nevertheless straightforward and can be carried out based on a single sweep of the simulation data with the help of suitable software.

IDs can also be seen as a new type of simulation metamodel that is used to analyze the progress of the simulation, effectively solve optimal decision alternatives, and study the consequences of the decision alternatives. The IDs show the time evolution of the simulation and the dependencies between the simulation states at different time instants. Furthermore, the IDs give an effective tool for simulation based decision making. The IDs offer thorough insight into the course of the simulation and show how different events affect its outcome, making simulation based decision making more transparent. Additionally, the IDs present the probability distribution of the states of the utility node and thus alternative utility functions can be compared with little extra effort. It should also be noted that the analysis of the constructed IDs could be extended into multi-objective settings where the decision alternatives are compared with respect to several, possibly conflicting, objectives (Diehl and Haimes 2004).

The examples presented in the paper involve military aviation but IDs constructed from simulation data can be undoubtedly applied in the analysis of almost any discrete event simulation model where the simulation state can be defined in a relatively compact manner. In such applications, the IDs are useful for describing the effects of simulation parameters and finding optimal simulation settings. Most importantly, the constructed IDs are valuable in explaining why a given decision alternative should be selected as they describe the consequences of the decision. Thus, the IDs can be used to better explain and rationalize the results of simulation based decision making.

REFERENCES

- Andersen, S. K., K. G. Olesen, and F. V. Jensen. 1990. *Hugin – a shell for building Bayesian belief universes for expert systems*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Barton, R. R. 1998. Simulation metamodels. In *Proceedings of the 1998 Winter Simulation Conference*, ed. D. J. Medeiros, E. F. Watson, J. S. Carson, and M. S. Manivannan, 167–174. Washington, DC: Institute of Electrical and Electronics Engineers, Inc.
- Cheng, R. C., and S. S. Currie. 2004. Optimization by simulation metamodelling methods. In *Proceedings of the 2004 Winter Simulation Conference*, ed. R. C. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 473–478. Washington, DC: Institute of Electrical and Electronics Engineers, Inc.
- Decision Systems Laboratory 2009. GeNie (graphical network interface). Available via <<http://genie.sis.pitt.edu/>> [accessed June 5, 2009].
- Diehl, M., and Y. Y. Haimes. 2004. Influence diagrams with multiple objectives and tradeoff analysis. *IEEE Transactions on Systems, Man, and Cybernetics, Part A* 34 (3): 293–304.
- Friedman, L. W. 1996. *The simulation metamodel*. Norwell, MA, USA: Kluwer Academic Publishers.
- Fu, M. C., F. W. Glover, and J. April. 2005. Simulation optimization: A review, new developments, and applications. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 83–95. Orlando, FL: Institute of Electrical and Electronics Engineers, Inc.
- Howard, R. A., and J. E. Matheson. 1984. Influence diagrams. In *The Principles and Applications of Decision Analysis*, ed. R. A. Howard and J. E. Matheson, Volume 2, 719–762. Strategic Decision Group, Palo Alto, CA.
- Howard, R. A., and J. E. Matheson. 2005. Influence diagrams. *Decision Analysis* 2 (3): 127–143.
- Jensen, F. V., and T. D. Nielsen. 2007. *Bayesian networks and decision graphs*. New York, NY: Springer-Verlag New York, Inc.
- Keeney, R. L., and H. Raiffa. 1993. *Decisions with multiple objectives: preferences and value tradeoffs*. Cambridge, MA: Cambridge University Press.
- Law, A. 2006. *Simulation modeling and analysis*. Fourth ed. New York, NY: McGraw-Hill Science/Engineering/Math.
- Mattila, V. A., K. Virtanen, and T. Raivio. 2008. Improving maintenance decision making in the Finnish air force through simulation. *Interfaces* 38 (3): 187–201.
- L-3 Communications Analytics Corporation 2002. *The X-Brawler air combat simulator management summary*. Vienna, VA: L-3 Communications Analytics Corporation.
- Neapolitan, R. E. 2004. *Learning Bayesian networks*. Upper Saddle River, NJ: Prentice Hall.
- Ólafsson, S., and J. Kim. 2002. Simulation optimization. In *Proceedings of the 2002 Winter Simulation Conference*, ed. E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, 79–84. San Diego, CA: Institute of Electrical and Electronics Engineers, Inc.
- Pearl, J. 1986. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence* 29 (3): 241–288.
- Poropudas, J., and K. Virtanen. 2007. Analyzing air combat simulation results with dynamic Bayesian networks. In *Proceedings of the 2007 Winter Simulation Conference*, ed. S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 1370–1377. Washington, DC: Institute of Electrical and Electronics Engineers, Inc.
- Poropudas, J., and K. Virtanen. 2009. Simulation metamodeling through dynamic Bayesian networks. Manuscript.
- Shachter, R. D. 1986. Evaluating influence diagrams. *Operations Research* 33 (6): 871–882.
- Shachter, R. D. 1988. Probabilistic inference and influence diagrams. *Operations Research* 36 (4): 589–604.
- Swisher, J. R., and S. H. Jacobson. 1999. A survey of ranking, selection, and multiple comparison procedures for discrete-event simulation. In *Proceedings of the 1999 Winter Simulation Conference*, ed. P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, 492–501. Phoenix, AZ: Institute of Electrical and Electronics Engineers, Inc.
- Virtanen, K., T. Raivio, and R. P. Hämäläinen. 1999. Decision theoretical approach to pilot simulation. *Journal of Aircraft* 36 (4): 632–641.
- von Neumann, J., and O. Morgenstern. 1944. *Theory of games and economic behavior*. Princeton University Press.

AUTHOR BIOGRAPHIES

JIRKA POROPUDAS received his M.Sc. degree in systems and operations research from the Helsinki University of Technology (HUT), Espoo, Finland, in 2005. He is currently working on his doctoral thesis at the Systems Analysis Laboratory, HUT. His research interests include statistics, simulation, simulation metamodeling, and statistical analysis of

Poropudas and Virtanen

basketball. His e-mail address is <Jirka.Poropudas@tkk.fi>.

KAI VIRTANEN received the M.Sc. and Dr. Tech. degrees in systems and operations research from the Helsinki University of Technology (HUT), Espoo, Finland, in 1996 and 2005, respectively. He is currently Teaching Research Scientist at the Systems Analysis Laboratory, HUT. His research interests include optimization, decision and game theory with particular attention to aerospace applications as well as discrete-event simulation. His e-mail address is <Kai.Virtanen@tkk.fi>.