

Publication III

Jirka Poropudas and Kai Virtanen. 2007. Analyzing air combat simulation results with dynamic Bayesian networks. In: Shane G. Henderson, Bahar Biller, Ming-Hua Hsieh, John Shortle, Jeffrey D. Tew, and Russell R. Barton (editors). Proceedings of the 2007 Winter Simulation Conference (WSC 2007). Washington DC, USA. 9-12 December 2007. Pages 1370-1377. ISBN 978-1-4244-1306-5.

© 2007 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.

ANALYZING AIR COMBAT SIMULATION RESULTS WITH DYNAMIC BAYESIAN NETWORKS

Jirka Poropudas
Kai Virtanen

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

ABSTRACT

In this paper, air combat simulation data is reconstructed into a dynamic Bayesian network. It gives a compact probabilistic model that describes the progress of air combat and allows efficient computing for study of different courses of the combat. This capability is used in what-if type analysis that investigates the effect of different air combat situations on the air combat evolution and outcome. The utilization of the dynamic Bayesian network is illustrated by analyzing simulation results produced with a discrete event air combat simulation model called X-Brawler.

1 INTRODUCTION

The most cost-efficient and flexible method for the analysis of air combat (AC) is often constructive simulation (e.g., Feuchter 2000; Hill, Miller, and McIntyre 2001; Stevens and Lewis 1992). Commonly used AC simulation models are based on the discrete event simulation methodology (Law and Kelton 2000) that offers flexibility for detailed modeling of the vehicles and hardware, the uncertainties inherent in AC, and pilot decision making (Glærum 1999; Lazarus 1997; Virtanen, Raivio, and Hämäläinen 1999). Simulation allows running a great number of replications that is necessary for producing statistically reliable results.

AC simulation data can be used for estimating descriptive statistics, empirical distributions of simulation output, or regression models that describe the dependence between simulation parameters and the simulation output (Law and Kelton 2000). In general, these statistical models are used to compare different aircraft and hardware configurations, or tactics. The models can also be extended to include the game setting of AC (Poropudas and Virtanen 2006). Unfortunately, these approaches stunt the AC into static event without time evolution and give no information on the progress of the AC during simulations. Thus, such models do not admit running what-if type analysis that study how

the occurrence of given situations affect the course and outcome of the AC.

In this kind of analysis, one has to present the relevant features of the simulation data and the time evolution of the AC scenario in an easily interpretable manner. For example, it is not realistic to assume that an analyst would watch the animations from all the simulation runs and fathom the general trends or phenomena proposed by the simulations.

To model the progress of AC, it is possible to present the AC state by a random variable whose statistical distribution is estimated at different time instants. However, the brute force estimation of these distributions is ineffective. If one wishes to run a what-if type analysis, the screening of the simulation data can be time consuming and has to be repeated for each individual situation.

In this paper, the above mentioned analysis is carried out using a dynamic Bayesian network (DBN) (e.g., Neapolitan 2004). The DBN presents the joint distribution of random variables using of a network graph and probability tables. The structure and probabilities of the DBN are estimated from AC simulation data produced with a discrete event air combat simulation model X-Brawler (Anon. 2002). The estimated DBN accumulates the information and dependencies inherent in the simulation data and provides an effective, practical, and time-saving method for analyzing different what if-type scenarios and courses of AC.

The paper is organized as follows. First, the AC scenario under consideration and the X-Brawler simulation model are introduced. Then, the estimation of the structure as well as probability tables of the DBN model are presented and the DBN based AC analysis is demonstrated with examples. Finally, conclusions are made and future research topics are proposed.

2 SIMULATING THE AIR COMBAT SCENARIO

In this paper, a symmetric 1 vs. 1 AC scenario that starts with identical blue and red aircraft facing each other at long

range is analyzed. The simulated pilots follow identical tactics and their objective is to shoot down the opponent. The purpose of the analysis is to study how the occurrence of a given AC situation affects the evolution of the scenario and the final outcome of the AC. Therefore, the AC state and outcome variables need to be defined.

The outcome of the AC at time t is described by a variable O_t whose set of possible values, denoted by θ , is:

- *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
- *neutral*, both sides are alive

The definition of AC state variable depends on the analyzed scenario and properties of the used simulation model. Now, the AC is simulated using a discrete event air combat simulation model called X-Brawler. The X-Brawler combines realistic hardware models with a hierarchical, value-driven pilot decision making logic that is based on the detailed and possibly imperfect situation awareness (Anon. 2002, Lazarus 1997). To model the actual AC flight tactics and the chain of command, the pilot decision making is hierarchical, i.e., the flight leader can direct the flight by giving out tactical orders.

The pilots' basic skills and behavior have been hard-coded into the simulation model. Additionally, the pilot decision making can be customized and refined using simulation input files, i.e., so called RULES files (Anon. 2002). These files contain if-then type statements to portray different tactics and to dictate the desired pilot behavior in different AC situations called *phases*. The actual pilot decisions are made using the value driven decision model that links all the aforementioned factors together (Lazarus 1997).

The definition of the AC state needed in the analysis is based on the ten phases of the decision making logic. The AC state at time t consists of the pair of random variables (B_t, R_t) that represent the phases of the two pilots. The sets of possible values for B_t and R_t , denoted \mathcal{B} and \mathcal{R} , include the following phases

- *fly-in*, flying to combat
- *approach*, maneuvering towards the opponent
- *support missile*, relaying radar information to launched missile
- *retreat*, trying to create distance to the opponent
- *point*, flying directly towards the opponent
- *WVR*, engaged in combat within visual range
- *decide*, making a decision between *retreat* and *WVR*
- *evade*, trying to avoid opponent's missile
- *fly-out*, returning to air base

- *killed*, aircraft has been shot down

The RULES file also includes a set of rules that determine the transitions between phases and depend on the pilot's assessment of the AC situation. The possible transitions are illustrated in Figure 1. The pilots start in the *fly-in* phase and the AC ends in the *fly-out or killed* phase.




Figure 1: The phase diagram of the RULES file describing the pilot decision making.

The AC state distribution can be estimated from the simulation data at different time instants and analyzed to understand the evolution of the AC. Furthermore, the data can be screened for replications that satisfy some desired condition in order to evaluate the effect of this condition on the AC. However, this approach is quite inflexible and the simulation data needs to be re-screened for all different conditions. In the following, this is avoided by fitting a DBN to the simulation data to allow more effective analysis.

3 CONSTRUCTION OF THE DYNAMIC BAYESIAN NETWORK

3.1 Bayesian Networks

Bayesian networks (BNs) present graphically the joint distribution of several random variables (e.g., Jensen 2001). The definition of a BN includes a directed acyclical graph and a probability distribution table for each node of the graph. The nodes correspond to the random variables and the arcs present the dependencies between the variables. For example, a BN model for the joint distribution of variables A_1, A_2 , and A_3 is presented in Figure 2. To present the joint distribution numerically, one has to define the distribution for A_1 , and the conditional distributions for A_2 and A_3 , i.e., how A_2 and A_3 are distributed when the value of their predecessor is known.




Figure 2: Simple Bayesian network.

Dynamic Bayesian networks (DBN) are BNs with a special structure as DBNs can be divided into disjoint *time slices*. Each time slice consists of random variables mea-

sured at the same time and there are connections between the slices in order to model phenomena that evolve in discrete time. For example, the model in Figure 2 could describe a random variable at three time instants in which case each variable would constitute its own time slice.

BNs can be constructed based either upon expert knowledge or upon learning where the probability distributions of the variables and the structures of the networks are estimated from data (e.g., Neapolitan 2004). The estimation of distributions is based on frequency calculations, i.e., the probability of a combination of values depends on the number of times it has been observed. It should also be noted that if some values are not observed for a random variable, the corresponding probabilities are zero and the dimensions of the conditional probability table can be reduced. The structure estimation is based on comparing different structure alternatives and selecting the best one, e.g., using a scoring rule (Neapolitan 2004).

The BNs are used to analyze the joint distribution of variables and their dependencies with computationally effective algorithms for distribution calculations. The BNs can also be fed evidence, i.e., observations on the values of variables. In practise, the distributions of the variables are altered so that the probability of the observed value is set as one. Then, the distributions of the other variables are updated or the most likely value combination is solved. This kind of analysis displays, what happens to the other variables' distributions if such a observation is made.

3.2 Dynamic Bayesian Network for the AC Scenario

To estimate the DBN model for the AC scenario defined in Section 2, the AC scenario is simulated 10000 times and a suitable DBN model is fitted to the data. The DBN is estimated using learning where the pre-processed data is fed to a BN-software such as Hugin (Andersen, Olesen, and Jensen 1990) to enhance the DBN structure and to estimate the probability tables (Neapolitan 2004).

In the DBN, the AC state is presented by a pair of discrete random variables, $B_t \in \mathcal{B}$ and $R_t \in \mathcal{R}$, that represent the phase of the blue and red pilot at time $t \in \{t_0, t_1, \dots, t_n\}$. The outcome O_t is function of B_t and R_t that is used to follow the kills and losses during the AC. The time slices are set as $\{0, 25, 50, \dots, 300\}$ seconds so that the variable O_{300} is the final outcome of the AC.

The DBN structure is defined so that a single time slice consists of the variables B_t , R_t and O_t where O_t is a function of B_t and R_t (see, Figure 3). To model the evolution of the AC, it is assumed that the variables B_t and R_t depend on their previous values. The dependencies are denoted by the black arrows in Figure 3. This structure is further refined by adding some arcs to fully represent the dependencies proposed by the simulation data (Neapolitan 2004). These additional arcs are denoted by white arrows in Figure 3.

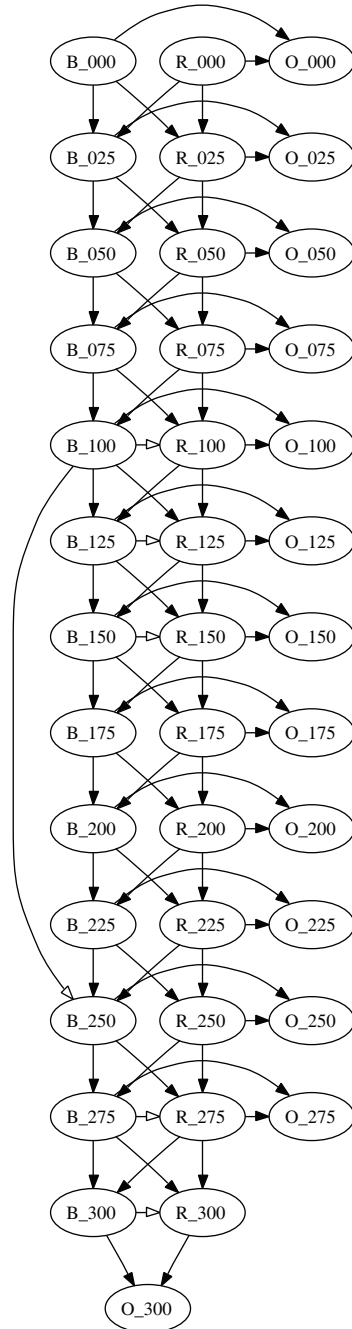


Figure 3: DBN model for the AC scenario. The black arrows present the assumed dependencies and the white arrows are added based on the simulation data.

After the probability tables of the variables are estimated from the simulation data, the DBN gives a presentation for the joint distribution of $(B_0, \dots, B_{300}, R_0, \dots, R_{300}, O_0, \dots, O_{300})$ that represents the evolution of the AC simulations. The main phase probabilities are presented as functions of time in Figures 4–8. The solid and dashed lines are the blue and red phase probabilities estimated directly from the simulation data. The stars and circles mark the corresponding probabilities given by the DBN. The fitting of the DBN is almost perfect pointwise as the stars and circles match the lines. Only inconsistency is found in Figure 6 where there exists a sharp peak in the probability curve. The DBN can not represent the peak as it does not include any variables in this time interval because the time slices are distributed uniformly on the time axis. The problem could be solved without increasing the size of the DBN by selecting the time instants more deliberately.




Figure 4: Probabilities for the *approach* phase. Blue: solid line and stars. Red: dashed line and circles.




Figure 5: Probabilities for the *point* phase. Blue: solid line and stars. Red: dashed line and circles.




Figure 6: Probabilities for the *WVR* phase. Blue: solid line and stars. Red: dashed line and circles.

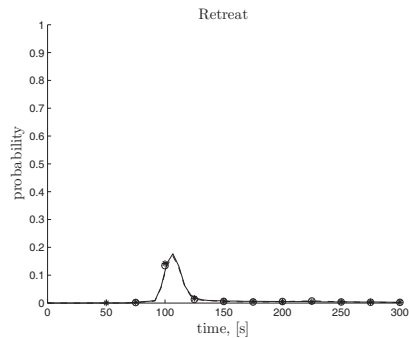


Figure 7: Probabilities for the *retreat* phase. Blue: solid line and stars. Red: dashed line and circles.

4 THE UTILIZATION OF THE DYNAMIC BAYESIAN NETWORK

In this Section, the utilization of the DBN is illustrated with example analyses of the AC scenario defined in Section 2. The examples demonstrate the evolution of the AC, the effect of observations on the AC and its outcome, and the effect of being in a given phase at different times. The analyses are based on the probability distributions given by the DBN and the distributions estimated from the simulation data are presented for comparison.

The evolution of phase probabilities in time is studied to understand the course of the simulated AC scenario. These distributions can be estimated directly from the simulation data and this analysis does not necessitate the DBN model. The advantage of the DBN lies in the computing time – the calculation of the conditional DBN probabilities is very fast compared to the re-screening of the original data to find the appropriate simulation runs. Most importantly, once




Figure 8: Probabilities for the *fly-out* phase. Blue: solid line and stars. Red: dashed line and circles.




Figure 10: Probabilities for the *approach* phase conditional on blue being in phase WVR at time 125.




Figure 9: The Distribution of the AC outcome O_t .

the DBN model is estimated there is no need for further processing of the simulation data.

The phase distributions are presented in Figures 4–8 to show the progress of the AC in the simulations. At the beginning of the combat, the pilots observe each other and start the approach maneuvering (Figure 4). Around the time $t = 80$, they point their aircraft towards the opponent (Figure 5) and make the decision to either engage WVR combat (Figure 6) or retreat (Figure 7). After this, there is a about 25% chance of the pilot returning to the air base (Figure 8). The outcome distribution of the AC is plotted in Figure 9. The outcome stays in the *neutral* state until the time $t = 110$ when the kill probabilities increase rapidly. The probability curves level off after time $t = 150$ and approach their final values: $P(\text{blue adv.}) = 0.253$, $P(\text{red adv.}) = 0.244$, $P(\text{mut. disadv.}) = 0.492$, and $P(\text{neutral}) = 0.012$.

In addition to the analysis presented above, the DBN can be fed evidence, i.e., observations on the values of the AC state variables at different times. The evidence allows the updating of the probability distributions and the effect of evidence is propagated through the DBN by calculating the conditional joint distribution of other variables. Naturally,

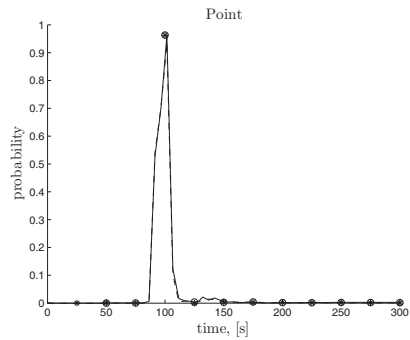


Figure 11: Probabilities for the *point* phase conditional on blue being in phase WVR at time 125.

one can also concentrate on the distribution of single variable such as the AC outcome.

To demonstrate the study of an observation on the blue phase at some time and its effect upon the course and outcome of the AC scenario, it is assumed that blue is found to be in WVR phase at time $t = 125$. Additionally, all the blue phases are analyzed to see their effect on the distribution of AC outcome.

The observation that blue is in WVR phase at time $t = 125$ results in the conditional phase probabilities presented in Figures 10–13. As before, the lines have been directly estimated from the simulation data while the stars and circles denote the probabilities given by the DBN model. The figures show that the DBN gives almost exactly the same probabilities as the original simulation data without the need for re-screening of the data.

Figures 10 and 11 show how the condition "blue in WVR at $t = 125$ " affects the beginning of the AC: the *approach* phase is unaffected, but the peak probability of *point* increases from 0.805 to 0.963, i.e., in the given scenario




Figure 12: Probabilities for the *WVR* phase conditional on blue being in phase *WVR* at time 125.

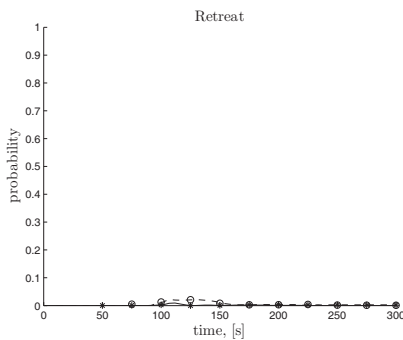


Figure 13: Probabilities for the *retreat* phase conditional on blue being in phase *WVR* at time 125.

WVR combat is almost always entered through the *point* phase.

It is seen in Figure 12 that if blue is in *WVR* at time $t = 125$, the red probability of *WVR* is increased up to 0.714, i.e., the visual range dogfight is a phase that the pilots engage simultaneously. If *WVR* combat is engaged, it is clear that neither side has retreated (Figure 13). Figure 14 shows the AC outcome in the cases where blue is in *WVR* at time $t = 125$. Especially, the final outcome distribution is $P(\text{blue adv.}) = 0.344$, $P(\text{red adv.}) = 0.277$, $P(\text{mut. disadv.}) = 0.360$, and $P(\text{neutral}) = 0.020$. This can be compared to Figure 9 to show how the fact that blue has survived at least until the time $t = 125$ affects the AC outcome. The probability of *blue advantage* is increased and the probabilities of *mutual advantage* and *neutral* are decreased.

Because the DBN enables a fast method for running what-if analysis, the study of different observations is vastly accelerated. For example, in Table 1 one can see what kind of effect the different observations at time $t = 125$ have on the distribution of the final AC outcome. Each line of Table




Figure 14: The distribution of the AC outcome O_t conditional on blue being in phase *WVR* at time 125.

1 contains the following information: the phase observed, the probability of this observation, and the conditional distribution of the AC outcome given this observation. The final line presents the AC outcome distribution without evidence. In 63% of the simulations, blue has already been shot down by this time and the probability of *red advantage* is 0.318. Conversely if blue is in *fly-out* phase, the probability of *blue advantage* is 0.971, i.e., the most probable reason for blue heading back to the air base this early is red being killed. The other phases can be analyzed similarly, but in the case of really unlikely observations such as “blue in *support* phase at time $t = 125$ ” one should keep in mind that such results are based only on a very small fraction of simulation runs so that the distribution estimates may be inaccurate.

The DBN can also be used to track a given state through the entire simulation run. This is illustrated in Table 2 that presents the probability of blue being in *approach* phase at different times. The conditional distribution of the AC outcome is also shown and the final line presents the AC outcome distribution without evidence. One can see that being in the *approach* phase before the time $t = 100$ is almost certain and does not affect the final outcome. On the other hand, if blue is still in the *approach* phase at the time $t = 100$, red has already had time for an attack and the probability of *red advantage* is 0.72. If blue survives this early attack by red, the situation is evened out as the time goes by. While the later rows of Table 2 are based on more unlikely events and fewer simulation runs, they show an interesting trend where the probability of the *neutral* outcome increases almost monotonously in time. In the corresponding simulations, the both pilots have lost track of their opponent and are stuck in the *approach* phase until the end of simulation.

To summarize the examples, the DBN model was applied in analyzing the evolution of the AC scenario as well as the effect of different observations on the evolution and the final AC outcome. Overall, the DBN accelerated the

analysis while the distributions produced by the DBN agreed with the simulation data. With a suitable BN software, the analysis of the DBN is an effective and informative way to comprehend the simulated AC because many AC situations can be considered and compared in little time.

Table 1: The distribution of the AC outcome conditional on the blue phase at time $t = 125$.

B_{125}		Outcome, O_{300}				
Phase	prob.	blue adv.	red adv.	mut. dis.	neutral	
<i>killed</i>	0.634	0.000	0.318	0.682	0.000	
<i>fly-out</i>	0.210	0.971	0.009	0.019	0.001	
<i>WVR</i>	0.110	0.344	0.277	0.360	0.020	
<i>evade</i>	0.018	0.276	0.298	0.419	0.007	
<i>retreat</i>	0.017	0.219	0.204	0.294	0.282	
<i>point</i>	0.010	0.217	0.054	0.453	0.273	
<i>approach</i>	0.003	0.080	0.278	0.016	0.626	
<i>support</i>	0.000	0.434	0.355	0.063	0.148	
-	1.000	0.253	0.244	0.492	0.012	

Table 2: The distribution of the the AC outcome conditional on blue being in the *approach* phase at different times.

Approach		Outcome, O_{300}				
time [s]	prob.	blue adv.	red adv.	mut. dis.	neutral	
0	0.000	0.253	0.244	0.492	0.012	
25	1.000	0.253	0.244	0.492	0.012	
50	0.999	0.253	0.244	0.492	0.012	
75	0.995	0.253	0.244	0.493	0.010	
100	0.011	0.039	0.724	0.083	0.155	
125	0.003	0.081	0.278	0.016	0.626	
150	0.002	0.078	0.183	0.017	0.723	
175	0.002	0.071	0.073	0.016	0.840	
200	0.002	0.083	0.081	0.019	0.821	
225	0.002	0.071	0.055	0.025	0.849	
250	0.002	0.053	0.068	0.033	0.847	
275	-	-	-	-	-	
300	-	-	-	-	-	
-	-	0.253	0.244	0.492	0.012	

5 CONCLUSIONS

This paper presented the analysis of an AC scenario using a dynamic Bayesian network whose structure and probability tables are estimated from simulation results. The transforming of the simulation data into a DBN is straight-forward and can be carried out based on a single sweep of the data with the help of a suitable software. Once the DBN has been estimated, it makes the analysis of the simulations more effective as performing what-if type analysis does not necessitate the re-screening of the original simulation data. The DBN gives a transparent and tractable model for AC. It offers thorough insight into the course of the simulated AC, and shows how different events affect the outcome of AC.

Such information can not be acquired using the traditional statistical analysis methods.

In the future, the most important target for development was found to be the selection of the time instants in way that gives the best approximation of the simulation data. A DBN could be estimated for m vs. n scenarios in order to analyze the tactics of larger flights. On other hand, the DBN can be utilized to outline AC as a tree that presents the branching of the most likely courses of AC.

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.

Anon. 2002. *The X-Brawler air combat simulator management summary*. Vienna, VA, USA: L-3 Communications Analytics Corporation.

Feuchter, C. A. 2000. *Air force analyst's handbook: On understanding the nature of analysis*. Kirtland, NM, USA: Office of Aerospace Studies, Air Force Material Command.

Glærum, S. 1999. TALUS - an object oriented air combat simulation. In *Proceedings of the 1999 Winter Simulation Conference*, 1160–1167. New York, NY, USA: IEEE Computer Society.

Hill, R. R., J. O. Miller, and G. A. McIntyre. 2001. Simulation analysis: applications of discrete event simulation modeling to military problems. In *Proceedings of the 2001 Winter Simulation Conference*, 780–788. Washington, D.C., USA: IEEE Computer Society.

Jensen, F. V. 2001. *Bayesian networks and decision graphs (information science and statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc.

Law, A. M., and W. D. Kelton. 2000. *Simulation modeling and analysis*. New York, NY, USA: McGraw-Hill Higher Education.

Lazarus, E. 1997. The application of value-driven decision-making in air combat simulation. In *Proceedings of the IEEE International Conference of Systems, Man and Cybernetics, Vol. 3, Orlando, Florida*, 2302–2307.

Neapolitan, R. E. 2004. *Learning Bayesian networks*. Upper Saddle River, NJ, USA: Prentice Hall.

Poropudas, J., and K. Virtanen. 2006. Game theoretic analysis of an air combat simulation model. In *Proceedings of the 12th International Symposium of Dynamic Games and Applications*. The International Society of Dynamic Games.

Stevens, B. L., and F. L. Lewis. 1992. *Aircraft control and simulation*. New York, NY, USA: John Wiley & Sons.

Virtanen, K., T. Raivio, and R. P. Hämäläinen. 1999. Decision theoretical approach to pilot simulation. *Journal of Aircraft* 36 (4): 632–641.

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 and simulation. His e-mail address is [`<Jirka.Poropudas@tkk.fi>`](mailto: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>`](mailto:Kai.Virtanen@tkk.fi).