Proceedings of the 2004 Winter Simulation Conference
R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

STREAM OPTION MANAGER

William P. Niedringhaus
Michael J. White
Patrick L. Jones

The MITRE Corporation
Center for Advanced Aviation System Development
7515 Colshire Drive
McLean, VA, 22102, U.S.A.

ABSTRACT

Stream Option Manager (SOM) is a set of mathematical tools developed at The MITRE Corporation's Center for Advanced Aviation System Development (CAASD). While still under development, it holds the potential to provide modelers an automated capability for separation, sequencing, spacing and proper exit strategy of aircraft in en route airspace. SOM addresses all of these services in a single, integrated linear programming solution. SOM computes modified flight paths meeting all constraints while minimally deviating from preferred trajectories. A prototype implementation of SOM is being evaluated in several applications. These include (a) an automated generator of sector complexity metrics e.g., by measuring the number and magnitude of separation actions needed to resolve a given traffic situation, (b) a means of performing system-wide analyses of revised conflict resolution requirements, e.g., revised separation standards, and (c) an automated controller-surrogate for simulation studies. The SOM metrics will be described, and example applications will be presented.

1 INTRODUCTION

Over many years of analysis of the complexity of air traffic control (ATC) scenarios, the same few metrics keep resurfacing: counts of aircraft, conflicts, transitions, crossings, and the like. Some specific examples of this traditional type of metric appear in Andrews (1997) and Alliott (1997); a summary of ATC metrics issues is presented in (Air Traffic Control Quarterly, 1997). A consensus is that these metrics capture only some aspects of sector complexity.

In contrast to other complexity measuring algorithms, SOM is distinctive in that it attempts to measure not *how hard a complex problem looks* (the traditional approach) but *how hard it is to solve.* Since hard-looking problems may turn out to be unexpectedly easy, and easy-looking ones hard, SOM arguably gets us a step closer to the core issues involved in measuring the complexity of dense ATC scenarios.

(Niedringhaus, 1998) describes an early version of the use of SOM for measuring complexity. It was tested only on simple scenarios (maximum twelve aircraft over thirty minutes). This paper documents our recent progress in upgrading the methodology to apply to realistic scenarios involving hundreds of aircraft over many hours.

SOM automatically resolves complex problems using an algorithm documented in (Niedringhaus 1995). It uses the x, y, and z positions of a set of aircraft at certain future times as variables in a linear program. Using linear inequalities, SOM assures the aircraft are separated properly over a period of time, obey in-trail constraints, obey their performance limits, and avoid certain regions of airspace (such as moving thunderstorms), while maximizing forward progress.

To evaluate new ATC concepts, large numbers of complex ATC scenarios must be tested, including many that are computer-generated—more than could be tested with humans in the loop. Thus there is a need for a fully-automated methodology to evaluate scenario complexity.

Evaluating complex problems as they evolve is difficult for traditional metrics, especially for automatically-generated "what if" scenarios. For instance, an hour into an artificial scenario (perhaps modeling that all aircraft fly their preferred routes), there may be a lot of separation problems. But in the real world, ATC would be acting all along, a factor that traditional metrics cannot model. SOM does model this. It uses its automated problem resolutions not only to calculate metrics but also to feed them back into the scenario as it advances. The intent is that if SOM finds complexity an hour into a complex scenario, it is genuine, and not simply the effect of an absence of any sort of (modeled) ATC intervention over the hour.

2 OVERVIEW OF THE METRICS

The metrics are based directly on SOM's resolutions, which are outputs of a linear program defined in (Niedringhaus, 1995). Values include delay induced, number of maneuvers needed, and the effectiveness of the maneuvers. Each is a snapshot in time, which applies over a given time horizon. Each measures a different aspect of density/complexity.

Weighted sums are used. The weights are intended to reflect common sense—that is, large or prompt maneuvers indicate a higher contribution to complexity than small or deferrable maneuvers. Thus a five-mile horizontal maneuver might be weighted more heavily if it must occur immediately, versus 10 minutes in the future. Or, among two maneuvers, 10 minutes in the future, the one with greater magnitude in miles would contribute more to complexity.

SOM computes the following set of nine complexity metrics:

1. **Problem**—a weighted sum of density and number/timing/miss-distance of conflicts—measures how hard the "remaining" problem is. As time passes, SOM commits to immediate maneuvers, but otherwise keeps taking fresh looks.
2. **Maneuver—**a weighted sum of the number, timing and degree of SOM maneuvers—interprets scenario as "hard" if many maneuvers are required, especially if they are immediate and large.
3. **Delay**—the delay induced by the maneuvers—interprets a hard scenario as one that requires delaying the aircraft.
4. **Separation**—SOM's degree of success in separating the aircraft—interprets the scenario as hard if not all aircraft pairs are fully separated.
5. **In-Trail**—SOM's degree of success in meeting miles-in-trail restrictions—interprets the scenario as hard if not all restrictions can be met.
6. **Return-to-Nominal**—SOM's degree of success in returning aircraft to their nominal positions—interprets the scenario as hard if aircraft are forced far off their nominal.
7. **Collinearity**—the degree to which maneuvers are collinear—interprets scenarios as hard if maneuvers are complex.
8. **Persistence**—the degree to which maneuvers planned are retained as time passes—interprets scenarios as hard if the strategy must change as the horizon advances.
9. **Sector**—the degree to which aircraft resolutions respect sector boundaries—interprets scenarios as hard if aircraft are forced outside their sector temporarily, or exit at an unexpected place.

3 CALIBRATION OF THE METRICS

The value for each metric processed a further step: it is mapped onto the range [0, 1), with the intent that:

- A value near 0 represents negligible complexity
- A value near 1 represents intolerably-high complexity (the value of 1 is never assigned; it is approached asymptotically with increasing complexity)
- A value of 0.5 is intended to connote a level which might keep a controller busy but not burdened.

We have selected values for weighting but have not yet calibrated our weights with experienced controllers, which we expect to do at a future time. Thus, our results are more useful as comparing scenarios relative to each other, rather than assigning an absolute measure to any given ATC scenario.

It may be just as well that we have not taken time to calibrate the weights yet, since SOM's resolutions are steadily being improved. With each improvement, metrics for a given scenario typically go down. One that generated, say, an 0.5 might yield only an 0.4. The scenarios are not getting easier, SOM is getting smarter. Fortunately, our experimental results indicate that relative complexity among scenarios appears to be preserved across SOM enhancements.

4 THE PASSAGE OF TIME

To measure evolution of a scenario as time passes, the following steps are iterated:

a. Resolve the scenario at time Sim Time, looking out into the future to time Horizon (which might be 10, 20 or 40 minutes later), and ignoring aircraft positions beyond then
b. Compute the SOM metrics based on the resolution
c. Advance the clock by a time-increment (perhaps 2, 5 or 10 minutes) to a new Sim Time, feeding back any prompt resolution maneuvers requiring changes in aircraft positions at the new Sim Time.

The overall complexity of the scenario tends to rise and fall as time passes. One may use the average over time to evaluate complexity over a period of time longer than SOM's horizon. For instance, in our experiments in Indianapolis Center, we ran SOM for five hours of data, using a 30- or 40-minute horizon.

5 "OPTIMISTIC" AND "PESSIMISTIC" PARAMETER SETTINGS

It is useful to run the same scenario through SOM with different parameters, e.g. using looser or tighter bounds on performable maneuvers, required separation, position uncertainty growth with respect to time, and the like. With "optimistic" settings, such as:

- uncertainties add up stochastically (i.e., sometimes cancel out)
- aircraft may make any reasonably-performable maneuver,

SOM may be able to fully resolve a complex scenario. But it might fail to do so with more "pessimistic" settings, such as:

- uncertainties always add up adversely, and
- aircraft are limited to their most-preferred set of maneuvers.

In either case, SOM provides useful information. In the optimistic case, SOM suggests a minimum magnitude of actions needed to solve the problem. In the pessimistic case, SOM suggests where and how possible solutions might break down. It is sometimes revealing to see how "pessimistic" the parameters need to be to force SOM to fail.

"Failure" does not cause SOM's linear program to become infeasible (Niedringhaus 1995). Instead, variables representing the deficiency in desired versus achieved separation or miles-in-trail spacing become positive instead of zero. For example, SOM may not be able to separate all aircraft pairs by a desired distance of five miles, but it may be able to so separate all but one pair, and that one pair is separated by only four miles. This information is part of the linear program output and can be fed into the metrics computation.

6 HUMAN-GENERATED VERSUS COMPUTER-GENERATED RESOLUTIONS

It is legitimate to question whether the difficulty experienced by a computer algorithm in resolving complex ATC problems can shed any light on how hard they may be for a human. After all, the means used by humans differ greatly from those used by SOM.

First note that the nine SOM metrics are just as well defined for human-generated maneuvers as for SOM generated maneuvers. Although there is no linear program in the human case, it is straightforward to calculate the metrics by comparing the before and after trajectories of each aircraft.

Consider next that SOM feeds back only maneuvers that are needed immediately. SOM deliberately avoids immediate maneuvers unless absolutely necessary, thereby allowing uncertainties to resolve themselves favorably so that no maneuver may ultimately be needed. This SOM tendency helps to minimize the amount of feedback that is used to represent ongoing ATC (and hence minimize effects of any human versus computer differences).

Running SOM multiple times, with varying degrees of optimism and pessimism may help SOM to deal with differences in how computers and humans resolve problems. The former makes maximum use of the computer's ability to number-crunch; its output is likely to lead to resolutions with lower metrics than would an actual controller's resolutions (were they evaluated using SOM's metrics). On the other hand, resolutions under pessimistic assumptions may have higher metrics than the human's. After thorough calibration with ATC experts, our hope is that a weighted average for optimistic and pessimistic SOM runs might be developed that would approximate the metric that SOM would assign to a real controller's maneuvers.

7 SECTOR DATA FOR EXPERIMENTS

We obtained five hours of data for four high-altitude sectors which we denote A, B, C and D. Controllers experienced in all four of these sectors have advised us that the A and B are significantly less difficult than C and D.

The values in Figure 1 show the average value of the nine SOM metrics calculated each sector. Each value is averaged over five variations (ranging from "pessimistic" to "optimistic"), and 30 different snapshots in time: every ten minutes over five hours. Thus each bar in Figure 1 is an average of 9 x 5 x 30 = 1350 measurements. The results confirmed the controller's impression that the sectors in the two right columns are significantly more complex than those in the two left columns.

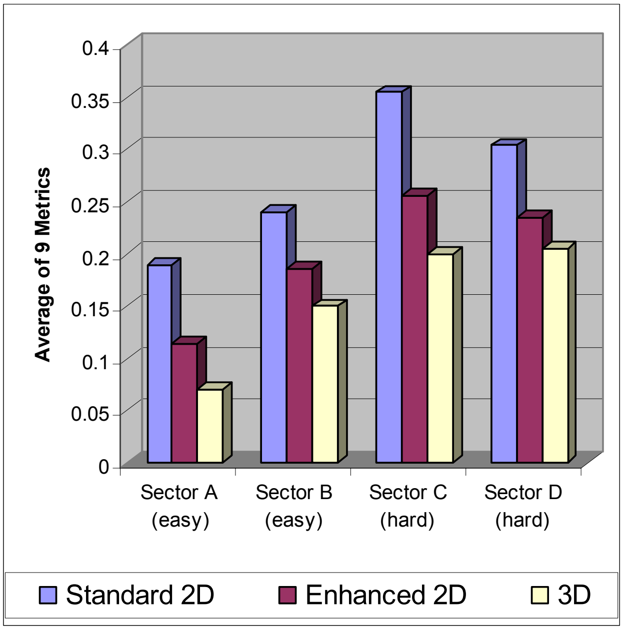

Figure 1: SOM Results for High-Altitude Four Sectors

8 THREE VERSIONS OF SOM (FEBRUARY, APRIL, JUNE 2004)

The different colors represent three different enhancement levels of SOM (February, April and June 2004). In February and April, SOM generated only horizontal (2D) maneuvers. Vertical separation came into play only for aircraft pairs that were already separated vertically; for these pairs SOM (passively) maintained vertical separation. The June experiments included a rudimentary vertical (3D) separation capability. Most separation is still achieved horizontally, but a few conflicts are resolved vertically. Analysis results from the February runs revealed poor results for aircraft that stayed in a sector only briefly (long enough for only one or two SOM maneuvers). By April the algorithm, originally designed to expect three or more maneuver opportunities, was corrected to take reasonable action with only one or two maneuver opportunities.

There are two significant results from our experiments shown in Figure 1. First, as we steadily enhanced SOM's capabilities in 2004, the two "hard" sectors consistently outscored the two "easy" sectors. Second, as SOM had better tools at it's disposal, it did not have to work so hard, so the metrics consistently declined. The absolute values declined, but the relative ranking stayed constant.

9 NEXT STEPS

We are in the process of adding several enhancements to SOM. First, the logic to decide whether to separate a given pair of aircraft vertically or horizontally is still rudimentary. It picks a sense for some aircraft pairs before it has considered all pairs. Instead, it should evaluate senses for all pairs before deciding on any. Second, SOM ignores the fact that any given sector is benefited by upstream-sector resolutions, and burdened, in turn, by the need to provide such services for downstream sectors. In SOM, for example, aircraft may arrive at a sector already in conflict, too late for SOM to resolve. SOM needs to be operating in the upstream sector as well, so as not to overestimate the workload downstream. SOM is therefore being extended to handle multiple sectors and collect complexity metrics for each.

These enhancements will be completed in summer 2004 and will be applied to further experiments to determine the complexity and difficulty of air traffic control scenarios.

ACKNOWLEDGMENTS

The authors would like to thank Brian Simmons of MITRE/CAASD for recognizing the applicability of SOM for complexity metrics in the current era.

REFERENCES

Andrews, John W., et. al, "Workload Implications of Free Flight Concepts," 1st USA/Europe Air Traffic Management R&D Seminar, Federal Aviation Administration and Eurocontrol, June 1997.

Alliot, Jean-Marc., et. al, "An Experimental Study of ATM Capacity," 1st USA/Europe Air Traffic Management R&D Seminar, Federal Aviation Administration and Eurocontrol, June 1997.

Air Traffic Control Quarterly, Volume 5 #4, 1997 (entire issue devoted to Air Traffic Control metrics).

Niedringhaus, W. P., "Stream Option Manager (SOM): Automated Integration of Aircraft Separation, Merging, Stream Management, and Other Air Traffic Control Problems," IEEE Trans. Systems, Man & Cybernetics, Vol. 25 No. 9, Sept. 1995.

Niedringhaus, W. P., "Solution Complexity Metrics," Proc. AIAA Guidance, Navigation and Control Conference, Boston MA, August 1998

AUTHOR BIOGRAPHIES

WILLIAM P. NIEDRINGHAUS, Ph. D., is a lead engineer at The MITRE Corporation of McLean Virginia. He received his Ph.D. in Electrical Engineering and Computer Science from Princeton University, and his B.S. in Mathematics from Yale. Dr. Niedringhaus has published numerous journal and conference articles in the areas of agent-based modeling, evolution of the airline industry, automated air traffic control, and dynamic load balancing in massively parallel computations. His e-mail address is `<wniedrin@mitre.org>`

MICHAEL J. WHITE is a Lead Engineer at The MITRE Corporation's Center for Advanced Aviation System Development. He joined MITRE in 1983 following earlier aviation research at ARINC Research and the IIT Research Institute. He has a B.S. in Physics from the University of Maryland, and a Master of Engineering Administration from the George Washington University. His research interests center on system-level simulation of the National Airspace System in support of benefits analyses of new technologies and procedures. His email address is `<mwhite@mitre.org>`

PATRICK L. JONES is a Lead Network Systems and Distributed Systems Engineer at The MITRE Corporation. He joined MITRE in 1991 after 12 years in industry with Spatial Data Sciences, Science Applications International Corporation, and Computer Sciences Corporation. He received a B.A. in Geography from SUNY/Buffalo (1979). His research interests include distributed systems, spatial data technologies, modeling and simulation.. His e-mail address is `<plj@mitre.org>`

This work was produced for the U.S. Government under Contract DTFA01-01-C-00001 and is subject to Federal Aviation Administration Acquisition Management System Clause 3.5-13, Rights In Data-General, Alt. III and Alt. IV (Oct., 1996).

The contents of this document reflect the views of the author and The MITRE Corporation and do not necessarily reflect the views of the FAA or the DOT. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty or guarantee, expressed or implied, concerning the content or accuracy of these views.