DAILY AIRTANKER DEPLOYMENT WITH INTERACTING BASES

Kazi Saiful Islam and David L. Martell[1]

ABSTRACT.—Each day regional fire duty officers must decide how many airtankers to acquire and where to deploy them to minimize initial attack response time. We consider the case of fire arrival rates that vary over the day and several bases with overlapping attack zones. We describe how we extended Larson's steady state hypercube queueing model to account for constraints on airtanker strike range and time-dependent system behavior and how our model can be used for daily airtanker deployment planning in the province of Ontario.

INTRODUCTION

Fire plays a significant role in many Canadian forest ecosystems and an average of 1847 fires burned an average of 22710 sq km per year in the province of Ontario during 1976-1988 (Martell 1994). The Ontario Ministry of Natural Resources (OMNR), the agency responsible for forest fire management in Ontario, spends roughly $35 million on the basic fixed cost of its fire management organization and an average of approximately $35 million on Emergency Fire Fighting (EFF) each year. From a national perspective, Canadian forest fire management agencies spend a combined total of roughly $288 million per year.

Airtankers play a very significant role in forest fire management in Canada and in this paper we describe our development of a decision support system that can be used to help decide where to deploy airtankers each day so as to minimize their initial attack response time. We begin by describing the daily airtanker deployment problem and some of the earlier research on the subject. We then present our daily airtanker deployment model and some preliminary results of its application to a small test problem. We conclude with a brief discussion of our future research and implementation plans.

THE INITIAL ATTACK SYSTEM

The OMNR's initial attack system objective is to contain fires while they are small and thus limit the spread and damage they cause. The boreal forest region of Ontario, the portion of the province in which most fire activity takes place, is blessed with an abundance of water in small ponds and streams, and lakes from which amphibious airtankers can pick up water. Fires are placed in the initial attack queue as they are reported and fire suppression resources are dispatched as soon as they become available. Initial attack crews usually travel to fires by helicopter or truck, set portable power pumps at water sources near the fire, lay hose up to the fire, and establish a hose line around the fire as quickly as possible. Large amphibious airtankers based at airports fly out to the fire area, skim low over the surface of a nearby lake, scoop water into their onboard tanks, and drop that water on the fire. The cruising speed of airtankers is such that they usually arrive at the fire and drop several loads before the initial attack crew arrives at the fire site ready to begin their control action on the ground. Airtankers can therefore deliver relatively small fires to fire crews and increase the likelihood they will be controlled at a small size.

THE DAILY AIRTANKER DEPLOYMENT PROBLEM

The OMNR owns 9 CL-215 airtankers and can borrow additional airtankers from other agencies to satisfy the increased demand that arises during "fire flaps" when many fires arrive during short intervals of time and threaten

[1] Kazi Saiful Islam, Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON, M5S 1A4, Canada. E-mail: kazi@smokey.forestry.utoronto.ca. David L. Martell, Faculty of Forestry, University of Toronto, 33 Willcocks Street, Toronto, ON, Canada, M5S 3B3. E-mail: martell@smokey.forestry.utoronto.ca

to overwhelm the initial attack system. Each day fire managers must decide how to deploy the available airtankers at initial attack bases across the province. They attempt to deploy the airtankers close to areas where fires are expected to occur and thereby reduce initial attack response times. Our objective is to develop an airtanker deployment model that fire managers can use to help resolve daily airtanker deployment decisions.

The initial attack airtanker system can be viewed as a spatially distributed queueing system with fires as customers and airtankers as servers. One important distinguishing feature of the initial attack system is the fact that unlike many queueing systems in which customers arrive at stationary servers, airtankers travel to their customers. The fire arrival process is a non-homogeneous Poisson process and airtanker service times are long relative to the diurnal variation in the fire arrival rate. Since the system never really reaches steady state it is not reasonable to partition the day into a small number of time intervals and predict the steady state performance of the system during each interval as is commonly done in urban emergency response system planning.

Daily fire occurrence can be modeled as a time-dependent Poisson process with the fire arrival rate determined by current and past weather in each area. The OMNR partitions the province of Ontario into many small independent sectors and uses historical fire weather and fire occurrence data to predict daily fire occurrence by sector expressed in terms of the expected number of fires per day. Historical fire report time data is then used to model time-dependent fire arrival rates throughout the day. Let $f(t)$ denote the probability distribution of the arrival time of fires in a particular sector. If the expected number of fires arriving in that sector during the day is λ_{Day}, then λ_t, the fire arrival rate at time t can be expressed as $\lambda_t = f(t)\ \lambda_{Day}$.

Airtanker service time includes the time required to fly to and from the fire as well as the on-scene fire fighting time. It is reasonable to assume the on-scene fire fighting time is exponentially distributed and that it varies by daily weather and forest vegetation, which vary by sector. Since the service time includes both the on-scene fire fighting time plus the two-way travel time it clearly cannot be exponential, but later we describe a heuristic calibration procedure that makes it possible to approximate system behavior by assuming the service time is exponential with calibrated parameters that are based on both the travel and on-scene fire fighting times.

PREVIOUS RESEARCH

Fortin (1989) studied amphibious airtanker operations in the province of Quebec, which owned 21 airtankers during the 1980's. She analyzed Quebec's historical airtanker data and found that an Erlang distribution could be used to model airtanker service time - the travel time plus on-scene fire fighting time. Given the large number of airtankers relative to the fire load she modeled Quebec's airtanker system as an $M(t)/E_k/\infty$ queueing system. The OMNR has a larger initial attack fire load than Quebec so its 9 airtanker fleet size renders the infinite server assumption inappropriate for Ontario.

Martell and Tithecott (1991) partitioned the Northwest region of Ontario into 4 sectors, each of which enveloped a single airtanker base, and modeled initial attack airtanker operations in each sector as an independent queueing system. The duty officer provided daily subjective assessments of the expected number of people and lightning caused fires in each zone by subjectively modifying predictions produced by statistical models and they solved the M(t)/M/s queueing system equations numerically for each zone independently. The results were then presented to the fire managers in the form of computer-generated graphs. The fire managers involved expressed concern that the model did not capture the important interaction between airtanker bases.

Larson (1974) developed the hypercube queueing model (HQM) which models the interaction of distinguishable servers in a spatially distributed queueing system. The HQM is a descriptive model that provides steady state probabilities of the system states and many related performance measures given the airtanker deployment. The model is based on an M/M/*s* queueing system with distinguishable servers, but airtanker service time has a general distribution since it includes a very significant travel time component in addition to the on-scene fire fighting time. The HQM model assumes airtankers have sufficient fuel capacity to cover any, including the most remote, fire in the region.

The HQM is a steady state model and its application to a system in which fire arrival rates vary over time is contingent on the assumption that one can partition the day into discrete time intervals and use the HQM to produce steady state results for each such interval. In the case of airtankers fighting forest fires, peak congestion follows the peak fire arrival rate and system performance deteriorates. When the system becomes congested the best HQM performance is achieved by moving all airtankers to a centrally located base. Urban fire trucks can "gas up" on the way to a fire, but airtankers operate in remote areas and since they can seldom refuel en route to or from a fire, they have a maximum strike range beyond which they cannot fly, so the use of a single central base would leave many fires "uncovered" by an airtanker.

Table 1 illustrates the more significant differences between the HQM and an airtanker system. Given these discrepancies, we decided to extend the HQM by restricting airtankers to designated strike ranges and modeling the time-dependent behavior of the system.

Table 1—Comparison of the hypercube queueing model and an airtanker system

Larson's HQM	Airtanker System
Distinguishable servers	Distinguishable servers
One initial attack queue	One queue for each airtanker
No limit on strike range	Constrained attack radius
Exponential service time (needs to be calibrated)	Exponential service time (needs to be calibrated)
Constant arrival rate	Time-dependent arrival rate
Steady state solution	Time-dependent solution

A DAILY AIRTANKER DEPLOYMENT MODEL

Our daily airtanker deployment model is an extension of Larson's (1974) basic HQM developed by introducing: 1) time-dependent fire arrival rates and 2) constraints on the initial attack range of airtankers. The service time includes an on-scene component and a travel-time component that depend upon estimates of the probabilities that each airtanker is dispatched to a fire at each fire node. The model provides transient system state probabilities and many related performance measures such as the virtual response time to a fire that arrives in the system at time t given the airtanker deployment. We then extended a heuristic developed by Berman et al. (1987) to find "good" locations for s airtankers for a specified deployment time interval.

Heuristic to Solve the Daily Airtanker Deployment Problem

We deal with a fire region that contains several airports (nodes G_A) and a large set of potential fire locations (nodes G_F) where fires may occur. The daily airtanker deployment problem is a very complex continuous time stochastic control problem with control variables that indicate which airport each airtanker should proceed to after it completes the on-scene service of each fire. That problem is beyond the scope of our current analysis, and we deal with a much simpler problem that entails deciding where to base each airtanker for some specified time interval which we refer to as the deployment time interval. We assume each time an airtanker completes serving a fire it returns to the airport at which it was deployed at the start of the deployment time interval. It cannot be deployed at some other airport nor can it be dispatched on a multiple fire sortie that might take it to more than one fire before it returns to its base. The deployment time interval will be some time interval less than or equal to the entire day. We envisage duty officers initially selecting a deployment interval that begins at the start of a day and extends to the mid-point or the end of the day, and then selecting new deployment intervals that begin after the duty officer has observed the system and decided he or she wishes to revise the deployment on a rolling planning horizon basis.

Let $\boldsymbol{X} = (X^1, X^2, ..., X^s)$ denote the locations (airports) where the s airtankers are deployed at the start of the deployment time interval. Assume airtanker I is constrained to an initial attack radius a_i specified by the duty officer. An airtanker certainly cannot fly beyond the range imposed by its fuel tank capacity but its actual range is flexible as the duty officer can make a trade off between the amount of fuel loaded into the tank, the flying distance to the fire and the amount of flying time devoted to cycling between a landable lake close to the fire and dropping water on the fire. Duty officers assess regional fire loads and tend to limit the initial attack radius and concentrate their airtanker resources on high priority areas as anticipated fire loads increase so they are not "caught short" with all airtankers far from their bases when a high priority fire is reported. The initial attack strike range is therefore a decision variable that must be less than or equal to the maximum physical distance imposed by the fuel tank capacity and our model is designed to allow duty officers to evaluate the initial attack radius as well as the deployment base for each airtanker.

Airtanker initial attack zones may overlap. Let N_i denote the set of potential fire locations (nodes) "covered" by airtanker i (within radius a_i). Therefore $N_i \equiv \{[j] \mid d(X^i, j) \leq a_i, \forall j \in G_F\}$. Let $\overline{S}_i(X^i)$ denote the mean service time of airtanker i deployed at base X^i. Since we assume airtankers always return to their bases between fires it includes the fire suppression time plus two-way travel time.

$$\overline{S}_i(X^i) = \sum_{j \in N_i} \overline{w}_{ij} \left(\overline{U}_{ij} + \frac{2}{v_i} d(X^i, j) \right) \quad for\ i = 1, ..., s$$

where U_{ij} is the average fire suppression time for airtanker i fighting a fire at node j and $\overline{w}_{ij}$ is the probability (averaged over the deployment interval) that the next fire that arrives at node $j \in N_i$ will be assigned to airtanker i.

Our dispatch strategy is as follows. As soon as a fire is reported, the nearest available (idle) airtanker that covers that node will be dispatched to that fire, but if no covering airtankers are available, the fire is assigned to the queueing subsystem of the covering airtanker that has the shortest queue length and is served by that airtanker on a FCFS basis. The behavior of the system is described by a set of difference equations representing an M(t)/M/s system with distinguishable servers that can be solved numerically. Our state variable is an $s+1$ dimensional vector with the first element denoting the number of fires in the system. The next s elements are the number of fires in each airtanker's queueing subsystem. Table 2 shows the first 42 possible system states. State 17, for example, has 3 fires in the system. One is being served by airtanker 2, a second is waiting in airtanker 2's queue and the third fire is being served by airtanker 3; airtanker 1 is idle.

Table 2.—The first 42 states if the M(t)/M(s) system with distinguishable servers

state index	state symbol
0	0:000
1	1:100
2	1:010
3	1:001
4	2:200
5	2:110
6	2:101
7	2:020
8	2:011
9	2:002
10	3:300
11	3:210
12	3:201
13	3:120
14	3:111
15	3:102

16	3:030
17	3:021
18	3:012
19	3:003
20	4:400
21	4:310
22	4:301
23	4:220
24	4:211
25	4:202
26	4:130
27	4:121
28	4:112
29	4:103
30	4:040
31	4:031
32	4:022
33	4:013
34	4:004
35	5:500
36	5:410
37	5:401
38	5:320
39	5:311
40	5:302
41	5:230

Let $E[R_{X,t}]$ denote the expected virtual response time at time t, $E[W^{(q)}_{X,t}]$ the expected virtual waiting time in the initial attack queue at time t, and $E[T_{X,t}]$ the expected travel time to a fire that arrives in the system at time t. Then $E[R_{X,t}] = E[W^{(q)}_{X,t}] + E[T_{X,t}]$. Let $AR_X(\Delta)$ denote the average over the deployment time interval $\Delta \equiv [\Delta_L, \Delta_U]$, of $E[R_{X,t}]$, the expected virtual response time. Then $AR_X(\Delta) = \frac{1}{\Delta}\int_{\Delta_L}^{\Delta_U} E[R_{X,t}]\,dt$.

Our objective is to identify a location vector X^* in the fire region that will minimize $AR_X(\Delta)$.

The Mean Service Time Calibration

The airtanker service time includes the time to fly to and from the fire plus the on site service time. Since the flying time to a fire depends upon the deployment strategy, the service time we use to identify a good deployment strategy is an endogenous variable that depends upon the deployment strategy. We modified the mean service time calibration suggested by Berman et al. (1987) for our airtanker model. The steps are as follows.

Step 0: Estimate the initial mean service time for all airtankers based on a simple assumption concerning the assignment of fires to airtankers (an airtanker fights all the fires within its attack radius)

Step 1: (i) Run the model using the initial estimate of the mean service time for all airtankers

(ii) Estimate $\overline{w}_{ij}$, the probability (averaged over the deployment interval) that the next fire that arrives at node $j \in N_i$ will be assigned to airtanker i

Step 2: Calculate a revised estimate of the mean service time based on the current estimate of $\overline{w}_{i\,j}$

Step 3: Calculate e_i, the difference between the initial and revised estimates of the mean service times for all airtankers

Step 4: Stop when e_i is small enough for all *i*.

The Airtanker Deployment Heuristic

We use the following heuristic procedure to identify a good deployment strategy for the airtankers.

Step 1: i) Choose the deployment time interval $\Delta \equiv [\Delta_L, \Delta_U]$

ii) Specify an initial airtanker deployment $X = (X^1, X^2, X^3, \ldots, X^s)$

iii) Set the initial estimate of the average of the expected virtual response time over the deployment interval Δ equal to some large number, $AR_X^0(\Delta) = M$

Step 2: Solve the modified HQM and perform the mean service time calibration to produce

1) calibrated mean service time

b) $\overline{w}_{ij}$, the probability (averaged over the deployment interval) that the next fire that arrives at node $j \in N_i$ will be assigned to airtanker *i*

Step 3: Evaluate a revised estimate of the average of the expected virtual response time over the deployment time interval , $AR_X^R(\Delta)$

Step 4. If $\left| AR_X^R(\Delta) - AR_X^0(\Delta) \right| < z$,

Then

Stop, ***X*** is the chosen deployment

Otherwise

Set $AR_X^0(\Delta) = AR_X^R(\Delta)$

GO TO Step 5 to revise the airtanker deployment

Step 5: Myopically revise the location of each airtanker (one at a time) to improve the rate at which it serves the fires that have been assigned to it, and solve the modified HQM again

a) Revised airtanker location

$X_R = (X_R^1, X_R^2, \ldots, X_R^s)$

b) Set $X = X_R$

Return to Step 2 to solve modified HQM and perform the mean service time calibration.

PRELIMINARY RESULTS

We applied our model to the very simple test problem with 3 airtankers and 19 fire nodes illustrated in Figure 2 which shows the fire nodes and fire arrival rate at each node. Figure 3 describes the performance of the system for the initial deployment of the 3 airtankers at nodes 1, 3 and 8. The solution is described in Table 3, the first column of which shows the initial and interim deployment strategies evaluated at each iteration. The second column contains the time averaged initial attack response time and the third column contains the recommended deployment strategy generated by the deployment revision procedure. We tried two different initial deployment strategies for this test problem. The results presented in Table 3 indicate they converge to two different deployment strategies with almost identical performance measures.

TABLE 3. —Preliminary results with 3 airtankers and 19 fire nodes

Initial Deployment (X^1, X^2, X^3)	$AR_X(\Delta)$	Recommended Deployment (X^1, X^2, X^3)
	hours	
1, 3, 8	0.612	15, 6, 12
15, 6, 12	0.525	18, 7, 11
18, 7, 11	0.481	13, 7, 11
13, 7, 11	0.473	13, 7, 11
(Stop)		
3, 9, 19	0.572	6, 8, 18
6, 8, 18	0.472	6, 8, 18
(Stop)		

FUTURE RESEARCH AND IMPLEMENTATION

We have made many simplifying assumptions in order to develop a tractable model and many aspects of the system merit further investigation. Forest fires grow as they wait in the queue so service time actually depends upon waiting time, but we have ignored that important characteristic of initial attack system. It may well prove unnecessary to account for such behavior on most days but the issues should be investigated to determine the extent to which our assumption is acceptable.

The state space of our model grows very quickly as the number of servers increases. Larson (1975) developed approximation procedures for solving his steady state HQM model. It may prove necessary to develop approximate solution methods for our time-dependent constrained attack range airtanker HQM.

We have used a calibrated exponential service time distribution that merits further investigation. We are currently developing a simulation model that can be used to investigate some of these issues and to “validate” our HQM. We also plan to field test our model in Ontario during the 1998 fire season.

LITERATURE CITED

BERMAN, O., R.C. LARSON, AND PARKAN, C. 1978. The stochastic queue p-median problem. *Transportation Science* 21: 207-216.

FORTIN, J. 1989. Daily deployment of airtankers for initial attack on forest fires in the province of Quebec. M.Sc.F. thesis. Toronto, ON: University of Toronto, Graduate Department of Forestry.

LARSON, R.C. 1974. A hypercube model for facility location and redistricting in urban emergency services. *Computers and Operations Research* 1(1): 67-95.

LARSON, R.C. 1975. Approximating the performance of urban emergency service systems. *Operations Research*, 23: 845-868.

MARTELL, D.L. 1994. The impact of fire on timber supply in Ontario. *The Forestry Chronicle*. 70(2): 164-173.

MARTELL, D.L., AND A. TITHECOTT. 1991. Development of daily airtanker deployment models. In *Proceedings of the 1991 Symposium on Systems Analysis in Forest Resources*, comp. Buford, M.A. Charleston, SC. March 3-6, 1991. Asheville, NC: U.S. Department of Agriculture, Southern Research Station.

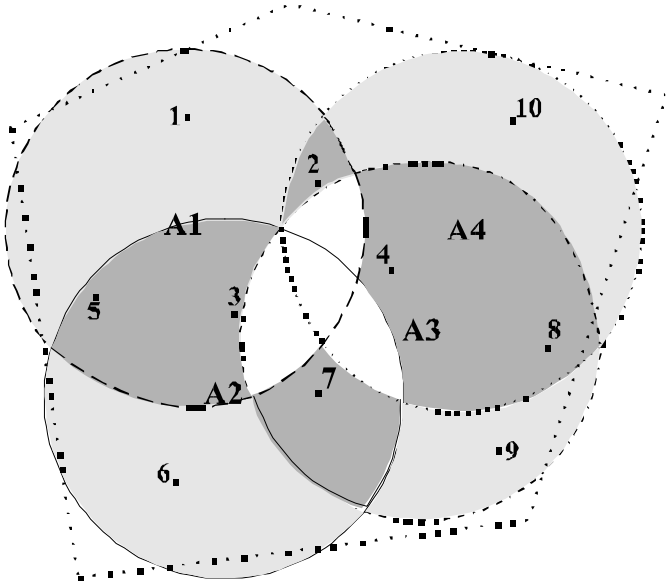

Figure 1.—Airtanker bases, initial attack strike range boundaries and fire nodes. The airtanker bases are nodes A1, A2, A3 and A4 and the fire nodes are nodes 1, 2, ..., 10. Some fires may lie within range of one or more airtankers while others may not be covered by any airtankers.

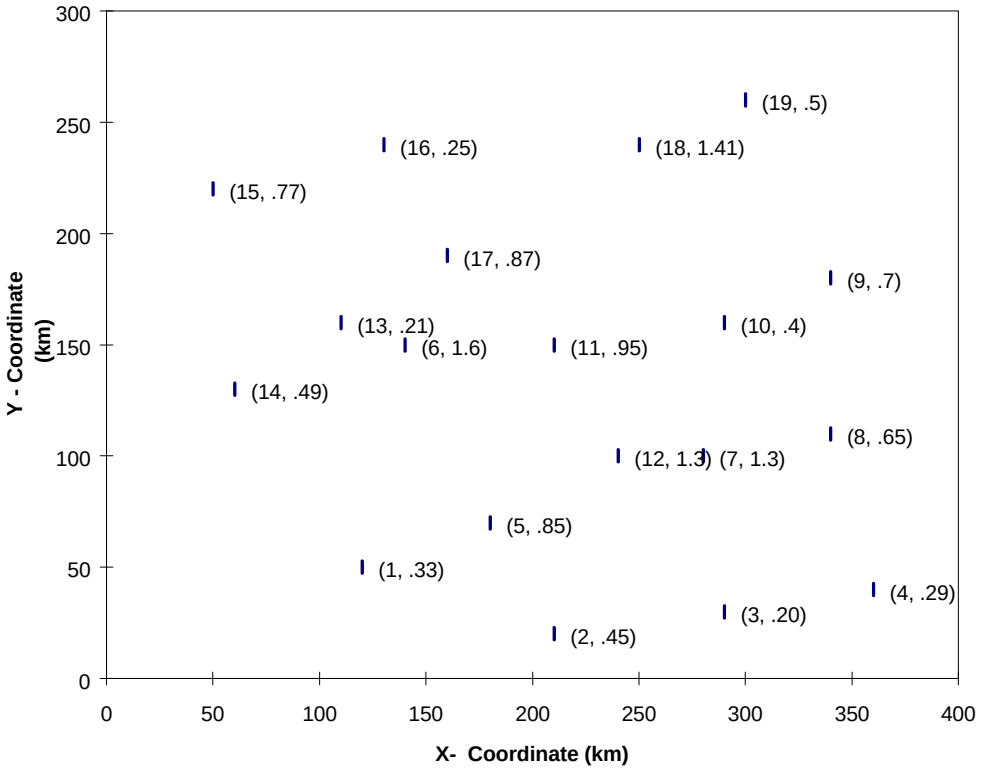

Figure 2.—Map of fire and airtanker nodes for a small test problem. The first number inside the parentheses is the node number and the second is average fire arrival rate at that node expressed in terms of fires per hour.

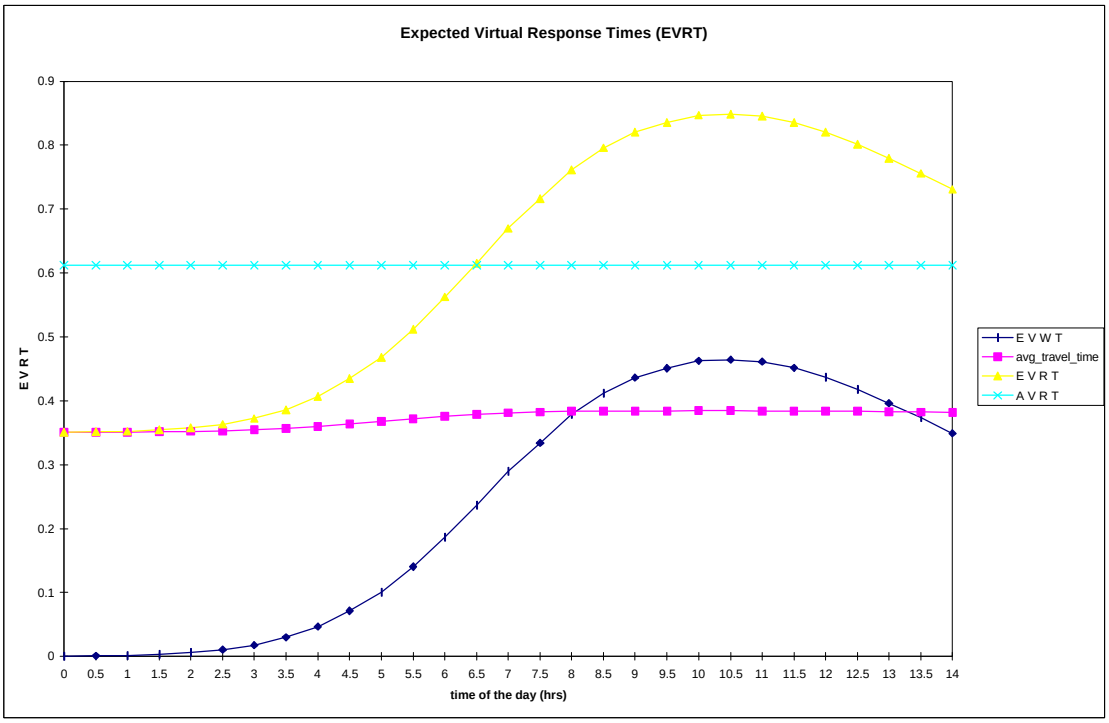

Figure 3.—Predicted airtanker system performance with airtankers deployed at nodes 1, 3 and 8. EVWT is the expected virtual waiting time, EVRT the expected virtual response time, and AVRT the average over the deployment interval, of the expected virtual response time.