Information Report FF-X-66

January 1978

AIRPRO

AN AIR TANKER PRODUCTIVITY COMPUTER SIMULATION MODEL

THE EQUATIONS

(Summary)

by

A. J. Simard and G. A. Young

Forest Fire Research Institute
Canadian Forestry Service
Department of Fisheries and the Environment
240 Bank Street
Ottawa, Ontario
KlG 3Z6

TABLE OF CONTENTS

AIRPRO

AN AIR TANKER PRODUCTIVITY COMPUTER SIMULATION MODEL

THE EQUATIONS

(Summary)

A. J. Simard and G. A. Young

INTRODUCTION

AIRPRO is a simulation model designed for computer implementation. Its purpose is to simulate the use of air tankers in wildland fire suppression operations. The model can be used to analyze a wide variety of questions with respect to air tanker systems, including: dispatch, resource and tactic selection, productivity, effectiveness, fleet size and composition, and, to a limited extent, allocation. In addition, because the model incorporates ground suppression and fire growth in some detail, it can also be used to analyze a variety of fire management questions external to air tanker systems.

The first part of this report outlines the specifications used to guide overall model development. The second part summarizes the relationships and interactions defined by the 300 primary equations which constitute the model. For those interested in more detailed information than is presented here, a complete description of the model is available from the Forest Fire Research Institute[1].

---

[1] Information Report FF-X-66 (Documentation).

I. SPECIFICATIONS

1. Goal Definition

A. The problem - there is no procedure available to quantitatively determine air tanker productivity or effectiveness and hence, the appropriate role of air tankers in wildland fire management systems.

B. Use of the model - by researchers, to provide data to aid fire management agencies with air tanker system presuppression planning.

C. Goals - that the model be able to:

1) quantitatively measure air tanker productivity and effectiveness;

2) determine the optimum combination of resources and tactics to employ in specific fire suppression operations;

3) summarize the above for an agency over one or more fire seasons.

2. Requirements

A. Quality

1) Errors: the cumulative effect of all errors should be insignificant, when summarized over a large set of data. The cumulative effect of errors should also be sufficiently small to permit the drawing of conclusions based on individual fires.

2) Validity: assumptions and relationships which perturb the overall system by more than 5% will be verified.

3) Scope: the model should be applicable to all fire suppression environments in North America where air tankers are used.

4) Resolution: the model should be able to detect statistically significant system responses on the order of 10%.

B. Analysis

1) Verification:

- output generated by the environmental, fire, and ground suppression components will be compared with observed results on the individual fire reports;

- output generated by the air tanker component will be compared with previous air tanker research, to the extent possible.

2) Experimentation:

- the model will test all reasonable combinations of air tanker resources and tactics on a set of 3,000 historical fires which occurred in the province of New Brunswick;

- the results will be analyzed and inferences will be made, if warranted.

C. Implementation

1) Documentation: the model and computer program will be fully documented. There will also be a report on the analysis of the data, the results, and their interpretation.

2) Communication: the findings will be summarized and written reports presented to system managers. In addition, formal and informal verbal presentations will be made.

3) Application: this is the prerogative of system managers. The authors will provide whatever assistance is necessary.

4) Evaluation: the overall modeling effort will be evaluated, using procedures outlined elsewhere.[2]

---

[2] See FF-X-66 (Documentation) for more information.

D. Resources

1) Time: three months will be required for planning, two years for model development, six months for analysis and one year for implementation.

2) Cost: non salary expenses will include $30,000 for model development and analysis and $10,000 for data analysis and preparation.

3) Manpower: the project will require three persons; one project supervisor and two programming full-time assistants.

4) Equipment: a large computer with remote access, time sharing, and interactive debugging will be required.

3. Model Description

A. General

1) Class: operational model - to provide research data.

2) Size: medium.

3) Rigor: as appropriate to the specific component.

- mental models (tactic selection, suppression, etc.);
- correlative models (mop-up, cost, etc.);
- mechanistic models (distances, sunset time, etc.).

B. Design

1) Organization: modular - in general, one subroutine for each major component or function.

2) Aggregation: mixed - the air tanker component will be disassociated while the remaining components will tend toward integration.

3) Scale:

- there are three scales of time measurement - 0.1 to 1.0 hours for individual fires, zero to several days between fires, and one to several years for overall totals;

- space is measured in acres, with 0.1 acres or 10% of the area (whichever is greater) considered the minimum significant difference;

- three system levels are considered - fire suppression, air tanker utilization, and some air tanker utilization subsystems.

C. Technical Characteristics

1) Time simulation: uneven increment between events, with even increment for fire growth and ground suppression.

2) Flow:

- continuous for fire growth and ground suppression;

- discrete for events - crew arrival, change of hour, sunrise, sunset, and air tanker drop.

3) Certainty: deterministic - a large sample size is used to simulate the effect of stochastic elements.

4) Analytical technique: primarily simulation - with real-world observations, empirical data, regression analysis, and mathematical models being used as appropriate.

II. THE MODEL

The model consists of five components:

1. administration (input, output, initialization, tabulation, and control);
2. the environment (distances, sunset, weather, and fuels);
3. the fire (occurrence, behavior, and growth);
4. ground suppression (control, mop-up, and economics);
5. air tanker use (tactic selection, delivery, drop, line holding, and costs).

1. Administration

The administrative component links the model, the computer, and the user. This linkage is shown graphically in Fig. 1. The model is linked to the computer by establishing initial conditions for all system variables. This is done once, at the beginning of the run, for those variables that tabulate totals for the entire run or do not change during the course of the simulation. Initialization also takes place each time a new fire is processed by the model and each time that a new resource or tactic is used to refight the fire. In addition, initialization takes place during the course of fighting a fire, to avoid having to restart suppression when not necessary. The administrative component also links the model to the computer by controlling the sequence in which the computer processes the various technical subroutines.

Administration links the user to the model and the computer by bringing in user supplied fire and system descriptive data. It also tabulates and outputs results for ground suppression and various air tanker resources and tactics, as well as for the overall run. The result sequence also determines the optimum air tanker combination, based on the minimization of cost-plus-loss. Through input parameters, the user may specify a wide variety of options detailing what the model will and will not do, as well as what types of output are desired. From another point of view, the system manager uses the computer to solve a problem, with the model (through a program) being used to control the computer.

Another function performed by the administrative component is keeping track of time. While this is not strictly administrative, it is a control rather than a technical function. Since the model contains both event oriented and continuous processes, two time mechanisms are used: an event calendar and an even increment processor.

Figure 1

Linkage Between the Computer, User and Model

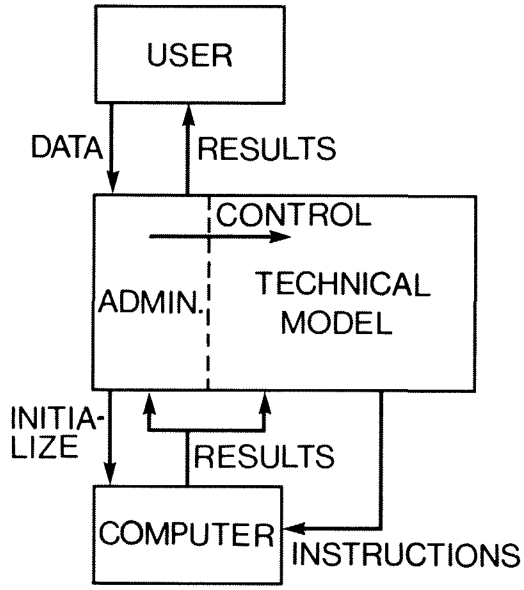

There are five events processed by the model: crew arrival, change of hour, sunrise, sunset, and an air tanker drop. These are processed by means of an event calendar. The calendar is simply a list of the elapsed times to the next occurrence of each event. As each new fire is encountered, the event calendar is initialized relative to the time of detection. As the processing of each event is completed (including calculation of the next occurrence), the calendar is searched to find the event with the smallest elapsed time since detection. The event is then processed by the model. The procedure continues until the fire is controlled by ground forces.

While fire suppression and fire growth are continuous and simultaneous processes, they are handled sequentially by the model. Changing the order of processing changes the simulation results. If suppression is processed first, area burned is reduced, whereas if fire growth is processed first, area burned is increased. Fire growth and suppression are processed each time an event occurs but before the event itself is processed. If the time interval between events is less than a user selected maximum, fire growth and suppression are processed only once. If the time interval between events exceeds the preselected maximum, a series of calls are made, with the event interval being divided into a series of even time increments, each of which is less than the maximum allowable time. Varying the maximum allowable interval between calls permits a reasonably fine calibration of simulation results with observed data.

2. The Environment

Those processes which are external to the fire management system, but which affect it, are grouped together under the environment. There are four environmental components included in the model: distances, sunrise and sunset, weather, and fuels.

Two distances are calculated by the model: fire-to-base and fire-to-lake. The distance from the fire to a lake at least 2.4 km long is part of the data associated with each fire. This distance is increased or decreased, if the required lake length is longer than or shorter than 2.4 km, respectively. The amount of change is calculated from tabular values determined in a previous survey of lake distributions in Canada. Fire-to-base distances are calculated with a series of equations taken from the literature, and which require only the longitude and latitude of any two points on the earth's surface. The average error is only 0.03 km at distances up to 500 km.

Another series of equations, derived from published literature, calculates the time of sunrise and sunset for any location, given the longitude, latitude, and date. The equations have the capability of processing far northern locations, where the sun may be above or below the horizon for 24 hours. The average error is on the order of two minutes or less for all locations of interest in North America.

Meteorological parameters are modeled on two levels: hourly variation, and daily change. Hourly variation is incorporated through tabular data obtained from a previous study of the diurnal variation of the Fire Weather Index (FWI) and Initial Spread Index (ISI) in the province of New Brunswick. Essentially, the index value for each hour is listed as a percentage of the 1600-hour value. An additional grass spread index was included in the model to account for flash fuels not incorporated in the FWI.

Weather and index parameters for the day on which the fire was detected are available as part of the fire data, as is the FWI for the day after detection. Beyond this, the values are reduced each day (to reflect increasing probability of rain) until they reach zero, insuring that no fire "blows up" in the model. No claim is made that the model accurately simulates fire growth or suppression beyond the first two days. In fact, the model continues processing these fires only to provide seasonal system totals. Finally, because meteorological conditions do not change instantly at midnight, the model incorporates a delay function, which applies the daily change gradually during the first 16 hours of each day.

Fuels are incorporated into the model by stratification. Three sets of fuel classification are processed by the model: input fuel types, a standard set, and an agency specific set.

The standard set is used for all generally applicable functions, with several parameter values being associated with each fuel type. The agency specific fuel types are used for the mop-up, cost, and damage regression equations which are specific to each fire management organization.

3. The Fire

To determine the demand for fire management activity, it is necessary to simulate the occurrence, behavior, and growth of wildland fires. The distribution of fire occurrence encountered in the field is generated by using historical data from actual fires.

The model performs a variety of editing and calibration functions before processing the data.

- Minimum values for fire growth and suppression are established.

- Anomalies between recorded free-burning and suppression fire growth rates are reconciled.

- Suppression fire growth is converted to free-burning growth.

- Data is calibrated to a standard time period.

Two aspects of fire behavior are incorporated in the model -- rate of spread and intensity. A theoretical forward rate of spread is calculated directly from the ISI and standard fuel type. Its primary purpose is to provide a mechanism whereby the model can apply diurnal adjustments to the observed spread rates. Similarly, a theoretical fire intensity is calculated from the FWI and adjusted for fuel conditions and season. The theoretical intensity is also used to incorporate diurnal weather variation. Hourly spread and intensity adjustments are made by calculating new theoretical values for each hour and adjusting previous values by the relative change.

Average spread values during the travel and control intervals, obtained from fire reports, are adjusted to initial conditions. The adjustment is such that the application of diurnal meteorological effects during the simulation yields the average value when integrated over time. To do this, the model makes a preliminary estimate of the time that will elapse during the travel and control intervals, including an adjustment for nighttime conditions. It then tabulates the relative difference between adjusted and unadjusted spread during the interval.

The use of average spread values yielded a low area burned for some fires where the reported size was larger than 16 ha. To counteract this, average spread for the problem fires was converted to a "pulse" which was applied immediately at the start of suppression to give these fires a "head start" in the simulation model. With the pulse, simulated growth during the control interval approximated the observed total growth much more closely. Finally, the presence of occasional large discrepancies in observed free-burning and suppression fire spread rates required that an average value be calculated for fire intensity, based on both observations.

The fire growth model is based on a segmented elipse, with a variable length-to-width ratio. The perimeter is divided into four components (head, two flanks, and rear), each of which is processed separately. The model could be classed as a "point growth" type, with parabolic segments connecting the points.

When a fire is detected (read in from the data set), the model begins by calculating a series of parameters describing a standard elipse (semimajor axis of one). First, the length-to-width ratio is related to wind speed. The point of intersection of the head and one flank is then determined, based on functions derived from the equation for an elipse. The relative length of head and flank are then calculated, using relationships derived from the equation for a parabola, which is used to approximate the eliptical segment connecting each of the four intersections. The relative flank and rear spread rates are calculated next, employing empirical relations obtained from the FWI.

To determine initial fire size and growth rates, the model first calculates the ratio of the rate of perimeter growth to the forward rate of spread. After the intensity at the head is determined, the model calculates flank and rear spread rates, using the observed forward spread rate and the relative values for the standard elipse. The area and perimeter of the fire are then related to each other. The final initializing sequence processes individual flanks. The intensity on each flank is determined by assuming proportionality with the relative spread rates. The length of each arc and chord as well as the arc-to-chord ratio are then calculated, using the observed fire size and the standard elipse parameters. The growth rate at each end of the four chords is also calculated, as is the initial growth rate for each arc.

When an arc is burning freely, growth is a function of arc growth rate and time. During suppression, growth is also proportional to the ratio of the free-burning arc length to total arc length. When one end of an arc is first held, the arc growth rate is modified to reflect the start of suppression and current meteorological conditions. During suppression, the free-burning end moves in response to horizontal and vertical spread vectors. With simple triangulation, new chord lengths, along with resulting new arc-to-chord ratios, are determined. When both ends of an arc are held, only a new arc-to-chord ratio has to be calculated, since the arc continues to grow while the chord length is fixed.

When ground forces have controlled the fire, the final area is determined by assuming an eliptical shape. The head and rear chords are averaged, as are the two flank chords. The averages are used to compute the length-to-width ratio which, in turn, yields the fire area.

4. Ground Suppression

The model begins fighting the fire by setting up a crew arrival schedule. In the case of fires with an observed size of less than 12 ha, the arrival time noted on the fire reports is used directly. For larger fires, crew arrivals are delayed by the model to prevent the initial attack crew from catching the fire too quickly in the simulation. The total line construction rate is allocated to between two and eight crews which arrive at evenly spaced intervals (not exceeding four hours) during suppression.

The rate of line construction on each flank is inversely proportional to the intensity on the flank. The rate of line construction can be split up among crews working in opposite directions with a variable percentage allocated to each direction. The rate of line construction within an air tanker drop is assumed to equal twice the rate at the rear of the fire.

Each flank is processed sequentially with respect to suppression activity. If suppression is taking place on a flank, the amount of time that will elapse until a change of line construction rate occurs is calculated. One of four equations is used, depending on whether the crew is or is not within a drop pattern and whether or not a pattern boundary will be encountered before the flank boundary. Line is then constructed until either a boundary is encountered (pattern or flank) or the end of the time interval occurs. If there is time remaining when a boundary is encountered, a new line construction rate is calculated and the process continues. When the entire perimeter has been controlled, suppression stops and the suppression time is calculated.

Mop-up is modeled in a different manner than the components discussed to this point. Since air tankers are not used for mop-up, the model only needs to estimate the reduction in mop-up time (and hence overall suppression cost) resulting from air tanker suppression. To estimate mop-up time, a series of equations (stratified by overstory species and fuel type) was determined by regression analysis. While several variables are used, they can all be related to either fire area or rate of mop-up. The average R^2 is 0.72 while the average standard error is 110% of the mean. While this predictive accuracy for individual fires is less than desired, it is considered adequate for the present purpose due to the relatively small effect of mop-up time, large sample size, and use of relative rather than absolute results.

The purpose of the model is to minimize cost-plus-loss. The economic consequence of air tanker use must, therefore, be determined. Suppression cost is modeled with agency specific regression equations. The several variables used can be related to three factors: amount of resources, duration of work, and amount of work accomplished. The average R^2 is 0.74 with an average standard error of 120% of the mean value. Although this accuracy precludes analysis on the basis of individual fires, significant results can be obtained with sample sizes as small as five to ten fires.

Two aspects of loss are considered in the model: forest and nonforest damage. Forest damage is modeled with a set of regression equations. Within a specific cover type, the amount of damage is related to the area burned and fire intensity. The average R^2 is 0.88 with an average standard error of 100% of the mean value, suggesting that the damage equations are slightly more accurate than the cost equations. Since both the occurrence and amount of nonforest damage are stochastic, they cannot be predicted with regression equations. Therefore, the model simply associates the observed nonforest damage with the results attributable to suppression without air tankers.

To compare the results of various tactics, the cost-plus-loss with air tankers is subtracted from the cost-plus-loss without air tankers. In addition, the maximum possible saving is determined by subtracting the cost-plus-loss that would result from a fire being held immediately at the time of detection from that realized with unaided ground suppression. The expected saving, with respect to nonforest damage, is determined by multiplying the observed damage by the ratio of the reduced fire area to the fire area with ground suppression only.

5. Air Tanker Use

The first step to be performed by the air tanker component is to select various combinations of resources and tactics for analysis. Since developing dispatch guides is one of the uses to which results obtained from the model will be put, an enumerative procedure, which considers all reasonable combinations had to be employed. Thus, the model selects, in order: air tanker model, number of aircraft, type of retardant, and flank of attack.

With over 2,000 possible combinations per fire, it is clear that some reduction in the number of trials is necessary. Thus, with respect to air tanker model, a trial is not undertaken if the delivery cost exceeds the maximum possible savings. In addition, all secondary models of a specific type are eliminated when the primary model fails to generate positive savings. For number of aircraft, a marginal test is added to the delivery cost test. If the saving with n + 1 aircraft are less than the saving with n aircraft, no further increases in aircraft numbers are tested. The latter test is modified slightly to incorporate the birddog cost. Only a delivery cost test is possible with respect to retardants, while only a marginal test is made with respect to flank of attack. A final series of tests rejects a trial if the air tanker production rate is less than the arc growth rate or the mission cost (one flank) exceeds the maximum possible savings.

The first step in fighting a fire with air tankers is to deliver the retardant. The model begins by calculating circuit time, which is the sum of: loading, takeoff, drop, and landing times. Loading time is a function of the retardant capacity of the air tanker. The remaining three times are based on an average for all aircraft, modified by a maneuverability factor. Maneuverability was assumed to be proportional to the design load limit and inversely proportional to the wing loading, power loading, and control surface loading.

The fire-to-base and fire-to-retardant flying times are determined next, based simply on distance and aircraft speed. The third delivery time of interest is the time between drops. It is obtained by combining various flying and circuit times in accordance with circumstances prevailing at the time of each drop. The first drop involves an initial warm-up period as well as base-to-fire flying time. Water-based operations also require a water pickup. For subsequent drops, the air tankers simply return to the nearest retardant source. An adjustment is made to subsequent drop intervals when more than one air tanker is involved. Finally, when air tanker endurance is about to be exceeded, or sunset is about to occur, the aircraft returns to base for refueling or an overnight layover.

Having delivered retardant, it must now be dropped on the fire. The model begins by determining the depth of retardant required. The depth of retardant required for extinguishment of low intensity fires is related to fire intensity through empirical data available in the literature. For high intensity fires, the model employs a double drop procedure, whereby the first drop reduces the intensity to a level where a second drop can extinguish the fire. The intensity reduction produced by the first drop is developed from theoretical studies for four broad fuel categories available in the literature. The depth required is related to one of ten retardant depth classes used by the model.

The next step is to determine idealized drop patterns in the open. An algorithm was developed in which the maximum pattern length for various depths and release sequences was determined. Pattern lengths for single releases obtained from the literature were used as input data by the algorithm. Essentially, the procedure gradually increases the interval between releases for each tank combination tested (one, two, four, or eight at a time). Individual patterns were combined and total retardant depth determined at a series of points along the overall pattern lengths. The longest release interval which does not result in a break at the desired depth yields the longest pattern length for that depth. These effective line lengths constitute the basic drop pattern input data used by the model.

In an effort to save storage space, only the single and double release pattern lengths are stored. For initial drops involving one or two releases, these data are used directly. For all subsequent drops or initial drops involving more than two releases, additive combinations of the data are used. The model also calculates percentage overlap between drops to aid in relating model results to field application.

Three pattern length adjustments are incorporated in the model: drop accuracy, canopy interception, and retardant viscosity. An average accuracy error is subtracted from the pattern length each time a drop is made. Drop accuracy is governed by five factors: target identification, reaction time, flying errors, wind, and aircraft response. It is assumed that the pilot is aware of the various causes of drop error and that he is attempting to compensate. The model, therefore, assumes errors of compensation rather than total error.

An average target identification error is used by the model. A standard reaction time error and flying errors are related to distance through drop speed. Aircraft response errors are assumed to be inversely proportional to maneuverability. Finally, a pattern drift error is related to wind speed. The expected error is determined with a binomial probability distribution since each individual error can result in either overshooting or undershooting the target.

The canopy interception component is based on extrapolation and interpretation of limited field data. In essence, a family of curves was developed relating the length-of-line held under a canopy to the length held in the open, the percent of maximum useful depth, and the quantity dropped. This is modified by a stand characteristic variable (again based on somewhat limited data) describing age and stocking. Overstory species is incorporated by using available rainfall retention data. Finally, the effect of retardant viscosity on drop patterns, based on data from a series of observed drop patterns, is also included in the model.

Once an adjusted pattern length is determined for every possible drop tactic and release sequence, the model selects that combination which maximizes the rate of production. Partial loads are selected only when they are sufficient to complete a mission.

The sequence of air attack used by the model is the head and right flank (clockwise) followed by the left flank and the rear (counterclockwise). The flanks are attacked one at a time with the air tankers returning to base after each flank is contained.

When a subsequent flank is considered, the model restarts at a point just before the next drop is due and continues from the previous flank. A drop is either tied in with the flank boundary (first drop), a previous drop, or the forward edge of line constructed by ground forces working in the same direction. If the far end of a drop overlaps line constructed by ground crews working in the opposite direction, the length-of-line held is reduced by the amount of overlap. If the far end of a drop spills over onto an adjacent flank, that portion of the perimeter within the drop, but not already controlled by ground forces, is listed as held. The model requires at least three separate drops to completely contain the perimeter, regardless of its length. It is possible with four- or eight-tanked aircraft that a single load can hold the entire perimeter if partial drops are made.

When the fire is controlled, the model accumulates totals for flying time and quantity of retardant dropped. Then, the flying circuit, and retardant cost per load are calculated. From these, the delivery cost, retardant cost, and birddog cost for the mission are determined. The latter three, in turn, yield the total air tanker cost.

Summary

In this report we have outlined the specifications used to guide the overall development of AIRPRO - an air tanker productivity computer simulation model. We have summarized the major relationships and interactions found in the model. In general, the model optimizes air tanker utilization by analyzing all reasonable combinations of resources and tactics and choosing the one which results in the lowest cost-plus-loss. The model consists of five basic components:

- The administrative component links the model, the computer, and the user. It controls input and output functions, initializes the system, tabulates results, and controls the flow of the model.

- The environmental component contains those processes which are external to the fire management system: distance calculation, sunrise and sunset, weather, and fuels.

- The fire component generates fire occurrence with historical data, calculates rate of spread and intensity, and models fire growth with a four-segment eliptical point spread model.

- The ground suppression component fights the fire, including: determining an arrival schedule, calculating line construction rates, building line on each flank, and mopping-up. Suppression costs and losses are also determined with the use of regression equations.

- The air tanker component first selects a resource and tactic combination for analysis. It then delivers retardant to the fire, drops the retardant, locates the drop on one of the four flanks, and finally calculates air tanker costs.

While there will no doubt be changes and improvements made to the model as it evolves from the developmental to the application stage, the initial version summarized here is a powerful tool for analyzing air tanker utilization. It is only through such objective, quantitative analyses that the potential effectiveness of air tankers in fire suppression will be realized. It is also only through such analyses that operational efficiency will be achieved. By contributing to these two objectives, AIRPRO will have achieved its purpose.