Open access • Proceedings Article • DOI:10.1109/ACC.2001.946006

Dynamic network flow optimization models for air vehicle resource allocation

— Source link

Kendall E. Nygard, Phillip Chandler, Meir Pachter

Institutions: North Dakota State University, Air Force Research Laboratory, Air Force Institute of Technology

Published on: 25 Jun 2001 - American Control Conference

Topics: Multi-swarm optimization, Swarm behaviour, Flow network, Resource allocation and Optimization problem

Dynamic Network Flow Optimization Models for Air Vehicle Resource Allocation

Kendall E. Nygard, Professor
Department of Computer Science and Operations Research
North Dakota State University
Fargo, ND 58105-5164

Phillip R. Chandler
Flight Control Division
Air Force Research Lab (AFRL/VACA)
Wright-Patterson AFB, OH 45433-7531

Meir Pachter, Professor
Department of Electrical and Computer Engineering
Air Force Institute of Technology (AFIT/ENG)
Wright-Patterson AFB, OH 45433

Abstract

A weapon system consisting of a swarm of air vehicles whose mission is to search for, classify, attack, and perform battle damage assessment, is considered. It is assumed that the target field information is communicated to all the elements of the swarm as it becomes available. A network flow optimization problem is posed whose readily obtained solution yields the optimum resource allocation among the air vehicles in the swarm. Hence, the periodic reapplication of the centralized optimization algorithm yields the benefit of cooperative feedback control.

I. Introduction

Autonomous wide area search weapon systems are small powered munitions, each with a turbojet engine and sufficient fuel to fly for short periods. They are deployed in swarms from a larger aircraft flying relatively high. Each is equipped with a solid-state sensor for searching and target identification, and a small warhead. There is interest in developing ways in which multiple munitions can work cooperatively to more efficiently and effectively carry out search, target classification, attack, and battle damage assessment.

In this paper we describe a time-phased network optimization model designed to produce task assignments for the powered munitions each time it is run. The model is intended to run simultaneously and independently at discrete points in time on all of the munitions, and explicitly assign each to choose to strike, assist in classifying a target, or continue searching.

II. Scenario

We begin with a set of N vehicles, deployed simultaneously, each with a life-span of 30 minutes[2]. We index them i = 1, 2, …., N. Targets that might be found by searching fall into known classes according to the value or "score" associated with destroying them. We index them with j as they are found, so that j = 1, 2, ….. and V_j is the value of target j. We assume that at the outset there is no precise information available about the number of targets and their locations. This information can only be obtained by the vehicles by carrying out searches and classifying what they find using Automatic Target Recognition (ATR) methodologies. Thus, at the time of deployment, all N of the vehicles begin searching, using a pre-established scheme for geographically dividing the work among themselves. We expect that at some point in time, a sensor on one of the vehicles detects a potential target. We assume that the target information, including location, type, and classification probability is immediately shared with all other vehicles within communication range, and that they, in turn, propagate the information to others. Ideally, all the weapons systems will quickly have the new target information in a local database. However, there may be cases in which not all vehicles can be reached, in which case some will be functioning with incomplete information. Target detection may be inconclusive, because a single pass over a target at a specific aspect angle may produce a classification probability that falls short of a required threshold value. In this case, a second pass at a different aspect angle can boost the classification probability, possibly to meet the threshold level. The second pass could be done by the initial spotting weapon system or by another, referred to as cooperative classification. If p_1 and p_2 are the classification probabilities on the first and second look respectively, then $p_1 + p_2 - p_1 * p_2$ ($> p_1, p_2$) is the classification probability on the two looks, assuming independence of the observations. If a vehicle strikes a target, it is not certain that the target is destroyed. We let q_j be the probability that target j is destroyed if attacked by a vehicle.

III. A Network Optimization Model

Network optimization models are typically described in terms of supplies and demands for a commodity, nodes which model transfer points, and arcs that interconnect the nodes and along which flow can take place[1]. There are typically many feasible choices for flow along arcs, and costs or values associated with the flows. Arc can have capacities that limit the flow along them. An optimal solution is the globally least cost (or maximum value) set of flows for which supplies find their way through the network to meet the demands. To model weapon system allocation, we treat the individual vehicles as discrete supplies of single units, tasks being carried out as flows on arcs through the network, and ultimate disposition of the vehicles as demands. Thus, the flows are 0 or 1. We assume that each vehicle operates independently, and makes decisions when new information is received. These decisions are determined by the solution of the network optimization model. The receipt of new target information is an event that triggers the formulation and solving of a fresh optimization problem that reflects current conditions, thus achieving feedback action. At any point in time, the database onboard each vehicle contains a *target* set, consisting of indexes, types and locations for targets that have been classified above the probability threshold. There is also a *speculative* set, consisting of indexes, types and locations for potential targets that have been detected, but are classified below the probability threshold and thus require an additional look before striking. Figure 1 below illustrates the network model at a particular point in time.

The model is demand driven, with the large rectangular node on the right exerting a demand pull of N units (labeled with a supply of –N), so that each of the LOCAAS nodes on the left (with supply of +1 unit each) must flow through the network to meet the demand. In the middle layer, the top M nodes represent all of the targets that have been identified with the required minimum classification probability at this point in time and thus are ready to be attacked. An arc exists from a specific vehicle node to a target node if and only if it is a feasible vehicle /target pair. At a minimum, the feasibility requirement would mean that there is enough fuel remaining to strike the target if tasked to do so. Other feasibility conditions could also enter in, if, for example, there were differences in the onboard weapons that precluded certain vehicle /target combinations, or if the available attack angles were unsuitable. The bottom R nodes of the middle layer represent all of the potential targets that have been identified, but do not meet the minimum classification probability. We call them *speculatives*. The minimum feasibility requirement for an arc to connect a vehicle /speculative pair is sufficient fuel for the vehicle unit to assume a position in which it can deploy its sensor to assist in elevating the classification probability beyond threshold. The lower tier models alternatives for battle damage assessment for targets that have been struck. Finally, each node in the vehicle set on the left has a direct arc to the far right node labeled sink, modeling the option of continuing to search. The capacities on the arcs from the target and speculative sets are fixed at 1 Because of the integrality property, the flow values are constrained to be either 0 or 1. This enforces a condition that at most one vehicle can attack any given target, avoiding the need to model the nonlinear scoring situation that would occur if multiple weapons could attack a single target. Each unit of flow along an arc has a "cost" which is an expected future value. The optimal solution maximizes total value. The values are listed in Table 1.

The network optimization model can be expressed in closed form as follows:

$$Z = \max \sum_{i,j \in I, i \neq j} c(i,j)x(i,j) \quad (1)$$

Subject to:

$$\sum_{j \in I, i \neq j} x(i,j) - \sum_{k \in I, k \neq i} x(k,i) = 0 \quad i \in I \quad (2)$$

$$\sum_{(i,j) \in A} x(i,j) \leq b(i,j) \quad \{\ (i,j) \mid i,j \in I, i \neq j\ \} \quad (3)$$

$$x(i,j) \geq 0 \quad \{\ (i,j) \mid i,j \in I, i \neq j\ \} \quad (4)$$

This particular model is a capacitated transshipment problem (CTP), a special case of a linear programming problem. Constraint set 2 enforces a condition that flow-in must equal flow-out for all nodes. Constraint set 3 mandates that flows on arcs must not exceed specified upper bounds. Restricting these capacities to a value of one on the arcs leading to the sink, along with the integrality property, induces binary values for the decision variables x(i,j).

IV. Coordinated Time-Phased Operation

At any point in time, each vehicle is in exactly one of the four possible modes, striking, classifying, searching or BDA. Striking means that the vehicle is in the process of attacking the target. Classifying means that the vehicle is searching for a known speculative, and upon finding it, will alter the classification probability (possibly, but not necessarily above threshold). BDA is carrying out battle damage assessment. Searching means the vehicle is carrying out a pre-established search process, looking for detections. Any of the vehicles that are at work at the point in time (i.e., have not been expended in striking a target) can cause an *event* that triggers a decision point. The possible events are given as follows:

1. A vehicle in search mode makes a detection, and classifies a target above threshold probability.
2. A vehicle in search mode makes a detection, and classifies a speculative below threshold probability.
3. A vehicle in classify mode makes a detection, and classifies a target above threshold probability.
4. A vehicle in classify mode makes a detection, and classifies a speculative below threshold probability.
5. A vehicle in strike mode transforms a target into a BDA objective.
6. A vehicle in BDA mode transforms a BDA objective into a target.
7. A vehicle in BDA mode removes a BDA objective from the problem.

An event triggering a decision point results in the following sequence of actions:

1. The vehicle responsible for the event stores the information in its own database, and communicates it to the others. The information is an ID, timestamp; the event type (1., 2., or 3. above); the target or speculative location; its type; and the vehicle location. In turn, those vehicles receiving the information relay it to others in which they are in contact, and add their own ID, timestamp, and location to the message.
2. Upon receiving the event information, each vehicle stores the new information in their local database, evaluates the value functions, parameterizes and autonomously solves the network optimization problem.
3. The solution is stored in the local database. Since all vehicles are solving the same problem, all now know the modes of all the others.
4. If a vehicle is in attack or classify mode, it completes the task underway, then switches to the mode specified by the model solution. If it is in search mode, it immediately switches to the newly assigned mode.

Elapsed time between events is a function of threat density, number of vehicles in the swarm and their locational distribution, and ability to detect targets and speculatives (dependent on such things as terrain, weather, ATR procedure employed). To get a feel for the dynamics of the approach, Figure 2 depicts a timeline example with 6 events, their times and their effects.

Key strengths of the approach are as follows:

1. Solutions are globally optimal, yet computed locally and independently. Communication of solutions is implicit, accomplished by all vehicles running the same model in parallel.
2. Requires very little preplanning, and is driven by using new information and dynamically reassigning roles to each vehicle.
3. Bandwidth requirements are very low. Model results are communicated implicitly through identical models on the distributed aircraft.
4. If communications fail between any vehicle, model synchronization may be not be global. Yet

feasible and potentially high quality solutions can still be computed locally, albeit with information that is not completely current.
5. There is a low computational burden on each vehicle, tractable even with modest computers.
6. The model explicitly considers deferred strike decisions as an available option.
7. The model explicitly considers deferred cooperative classification.

Potential weaknesses of the approach are as follows:

1. There is a large burden on being able to accurately specify cost functions.
2. Synchronization may be incomplete in situations when not all vehicles are in communication range or acknowledgements are not received. In these cases, solutions are not truly global, and, in some cases, could result in redundant actions such as two vehicles attacking the same target.
3. Model solutions may be delayed in circumstances where threats are being rapidly detected, since the weapons complete strike and classify tasks before solving a new model.
4. The approach assumes that battle damage assessment is perfect and instantaneous. Adding explicit battle assessment is straightforward, but requires development of the corresponding cost function component.

V. Model Solver

The model is a special case of the well-known capacitated transshipment problem (CTP). From the perspective of linear programming, solution of the model has several advantageous computational characteristics. The most important is that the coefficient matrix of any basis matrix is totally unimodular. From a computational view, this implies that there exists an optimal solution that is all integer. Furthermore, the basis can graphically be represented as a spanning tree, allowing basis change calculations to be carried out very rapidly on this specialized data structure in integer arithmetic. Solutions to such problems require little memory, and for the anticipated small problems generated in vehicle situations, would solve to optimality in negligible time –no more than a second or two, even on a modest processor.

VI. Conclusions

In this paper, an approach is developed to quantitatively address the tactical deployment of a swarm of air vehicles such that the benefit of cooperative operations is achieved. The relevant mathematical paradigm entails the real time solution of a network flow problem. This results in an integer programming problem of special structure, which is easily solved.

VII. References

1. Ford, L. R Jr. and D. R. Fulkerson, “Flows in Networks”, Princeton University Press, Princeton, N.J., 1962
2. Murphy, R. A., “An Approximate Algorithm For a Weapon Target Assignment Stochastic Program,” Approximation and Complexity in Numerical Optimization: Continuous and Discrete Problems, editor: P. M. Pardalso, Klewer Academic, 1999.

C(i,j)	=	Expected value of vehicle i attacking target j	
	=	(Probability of destroying target j) * (Value of target j)	
	=	$q_j * V_j$	i = 1, 2, …., N j = 1, 2, …., M
C(i,k)	=	Expected value of vehicle i assisting in classifying speculative k	
	=	(Probability of elevating classification probability above threshold) * (Probability that speculative will be attacked in the future if threshold is met) * (Value of speculative k)	
	=	$e_i * (a_i \mid p_i >= .95) * V_k$	i = 1, 2, …., N k = 1, 2, …., R
C(i,g)	=	Expected value of vehicle i carrying out BDA on target g	i = 1, 2, …., N g = 1, 2, …., G
	=	(Probability that target g is judged not dead) * (Probability that target g will again be struck) * (Value of target g)	
C(i,s)	=	Expected value of vehicle i continuing to search	
	=	(Probability that the vehicle will ultimately attack a target j) * (Probability of destroying target j) * (Value of target j) + (Probability that the vehicle will ultimately assist in elevating the classification probability of one or more speculatives above threshold) (Probability that speculatives elevated will be attacked) * (Value of speculatives attacked)	i = 1, 2, …., N
C(j,s)	=	0	j = 1, 2, …., M
C(k,s)	=	0	k = 1, 2, …., K
C(g,s)	=	0	g = 1, 2, …., G

Table 1. Expected future value,“cost”, of each unit flow

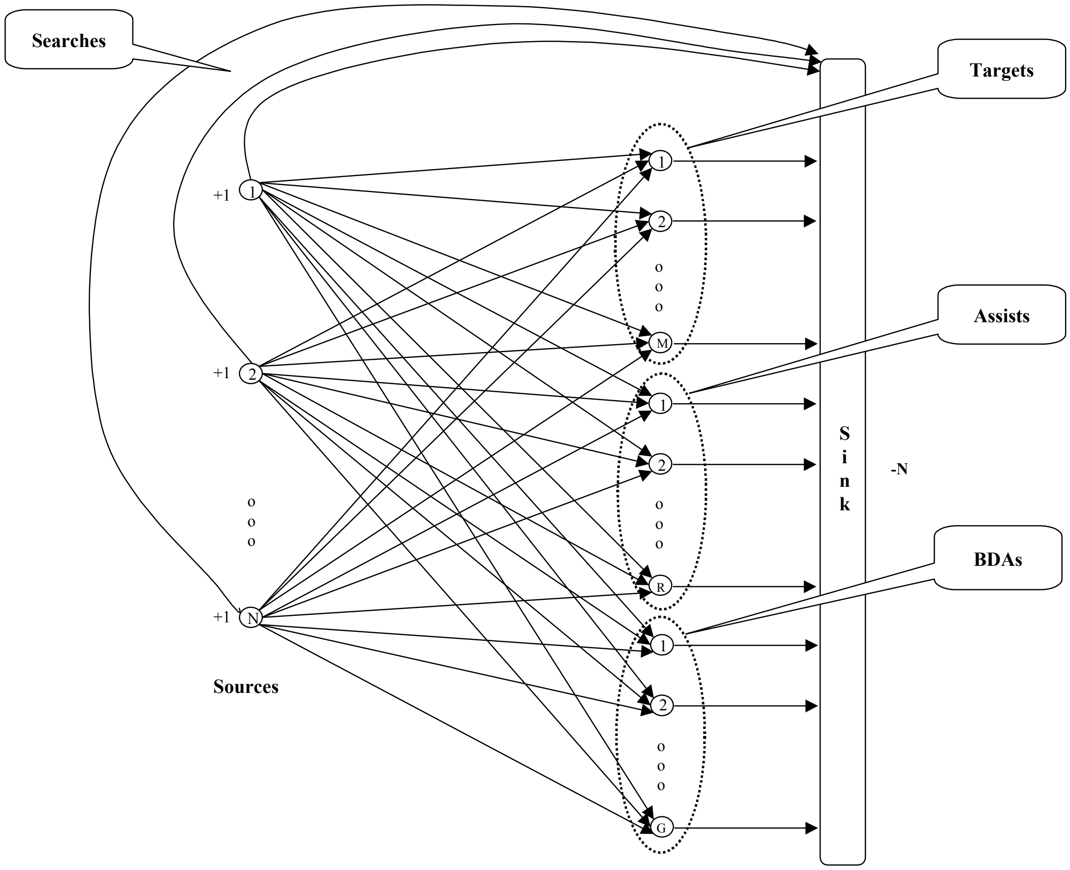

Figure 1. Network optimization model structure at a specific point in time. At this decision point, each of the N available vehicles has options to attack a known target (M choices), assist in cooperatively classifying a target (R choices), Battle Damage Assesement (G choices), or continue searching.

Target 1 detected by Vehicle 2. Classification probability = .98. Solver assigns Target 1 to Vehicle 2, which strikes and destroys Target 1.

Speculative1 detected by Vehicle 1. Classification probability = .83. Solver assigns search mode to all.

Speculative 2 detected by Vehicle 1. Classification probability = .67. Solver assigns Vehicle 3 to classify mode to detect Speculative 1.

Speculative 1 detected by Vehicle 3. Classification probability = .96. Solver assigns Target 2 to Vehicle 1, which strikes and destroys Target 2

Speculative 3 detected by Vehicle 1. Classification probability = .78. Solver assigns speculative 3 to Vehicle 1.

Speculative 3 detected by Vehicle 1. Classification probability = .91. Solver assigns search mode to all.