How to Optimize Air-to-Ground Weapon Purchases[1]

Matthew S. Goldberg and David M. Goldberg

The military service branches use modeling and simulation to determine their requirements for nonnuclear weapons fired from aircraft at targets on the ground. The annual modeling exercise is complicated, and solutions take a long time to compute. In this article, a father-son team demonstrate ways to reduce computation time that still offer high-quality solutions to the complex problem of determining requirements for conventional air-to-ground weapons.

Introduction

Determining requirements for U.S. military aircraft weapons involves numerous combinations of delivery aircraft, weapon types, and targets, which makes the modeling problem high-dimensional and slow to converge to a solution.

Two distinct problems need to be solved (see Figure 1). The first problem, *aircraft weapon budgeting*, is at the strategic level. The objective is to

We developed two greedy heuristics that enable us to accelerate the solution to the aircraft weapon budgeting problem, with minimal loss in the quality of the solution.

[1] Based on M. S. Goldberg and D. M. Goldberg, "Optimizing the Purchases of Military Air-to-Ground Weapons," *Military Operations Research* 24, no. 4 (2019): 37–52, https://www.jstor.org/stable/26853512.

determine the optimal inventory of air-to-ground weapons for fighting in a particular wartime scenario; the next several years' budgets are devoted to procuring that inventory. The second problem is *weapon target assignment*, which occurs at the tactical level in the theater of military operations. As part of their daily air tasking orders, air operations commanders match weapons in their local inventories with potential targets to destroy the highest-value set of targets (see U.S. Department of Defense, Joint Chiefs of Staff 2019). The ability to hit targets is constrained by the number of sorties (missions) available in the wartime theater for each aircraft that can feasibly deliver each type of weapon.

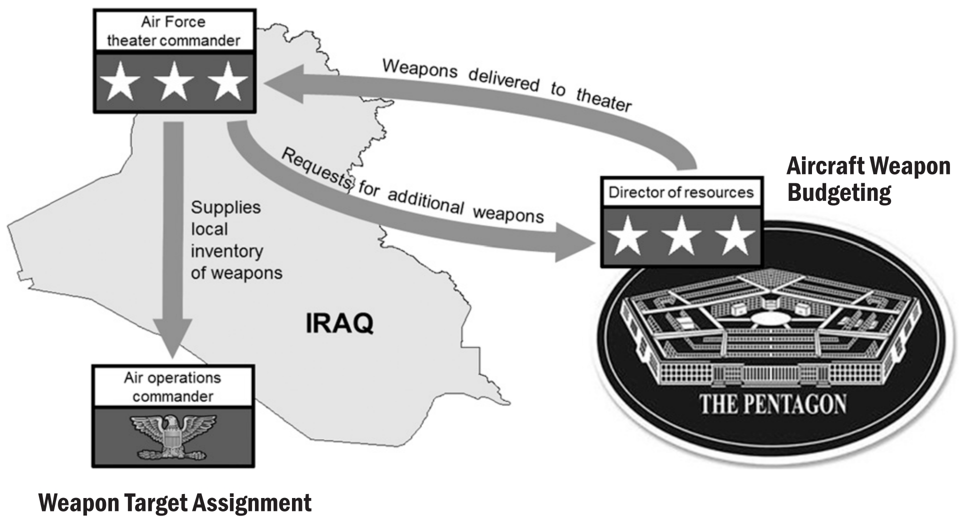

Figure 1. Relationship Between Aircraft Weapon Budgeting and Weapon Target Assignment

The military branches generally use nonlinear programming to solve these problems, and that is our approach as well. In our original paper, we made three major contributions:

1. We extended the work of Boger and Washburn (1985) by deriving an expression for the expected number of targets that would be destroyed when multiple aircraft-weapon combinations attack a given target population, and when some "dead" targets appear "live" and act as decoys.

2. We demonstrated the applicability of two heuristics that identify the preferred weapons to maximize a utility function defined over various types of targets destroyed.

3. We investigated the ability of those two heuristics to reduce the dimensionality of the problem and accelerate the solution.

This article summarizes the last two of these contributions. We developed two greedy heuristics that enable us to accelerate the solution to the aircraft weapon budgeting problem, with minimal loss in the quality of the solution.

Current Practice

The military branches often estimate their requirements for aircraft weapons in two stages. They do not impose an explicit budget constraint in the optimization model during the first stage, but only constraints based on the number of sorties available in the wartime theater. The solution to that problem generates requirements for weapons, to which they apply cost factors to estimate the budget needed to fully purchase them. The budget amount is treated as an output (byproduct) of the optimization rather than an input (constraint). Next, they run budgetary excursions in which they incrementally reduce the budget below 100 percent to see how the optimal solution changes and how much wartime capability is sacrificed. The budgetary excursions also highlight cases in which the optimal inventory is too expensive to purchase in the next year's budget (netting out current inventories), and has to be spread out over several years' budgets in the future.

Greedy Heuristics

We wondered whether greedy heuristics would be useful in either partially characterizing or fully computing a solution. A simplified version of our problem is similar to the fractional (or continuous) knapsack problem first studied in a classic paper by Dantzig (1957), who developed a "bang-per-buck" criterion for entering variables into the solution.

As a first step, we reduced the number of aircraft-weapon-target combinations considered in aircraft weapon budgeting by omitting those we deemed infeasible. The infeasible combinations included aircraft that cannot carry certain types of weapons and aircraft-weapon combinations that are not effective against certain targets. We next developed two heuristics that further reduce the number of aircraft-weapon-target combinations to be considered.

Either heuristic, or a hybrid of the two, greatly reduces time to produce a solution for the particular circumstances. We identified two cases in which the smaller, faster-running model produced the same number of expected kills as the full model: (1) the number of sorties available to deliver the weapons in the scenario being modeled is unlimited or (2) the weapon procurement budget is unlimited. These cases—called "edge cases"—although unlikely, demonstrate the gain in computational speed by pre-screening the aircraft-weapon-target combinations to a reduced number. Our research question was whether the improvement in computational speed outside of the edge cases comes at the price of some degradation in the quality in the solution—the utility provided by the expected numbers of targets destroyed of each type. Details about these heuristics follow.

Least Cost to Kill

Goldberg (1991) supplied mathematical conditions under which the least cost-to-kill (LCTK) criterion identified the optimum combination for destroying a single type of target when the budget constraint is binding but the sortie constraints are not. The cost-to-kill ratio divides the cost to purchase a weapon by its pure probability of kill against the target type, without adjustment for decoys. That ratio is minimized or, equivalently, its inverse ratio—"bang-per-buck"—is maximized. The LCTK heuristic eliminates choice variables from the optimization problem that would never appear in positive quantities in the optimal solution under the specified circumstances. However, the LCTK heuristic only partially characterizes the solution to the nonlinear program. The heuristic determines the aircraft-weapon combination to use against a particular target type but not the allocation of budget dollars to destroying each type.

Expected Kills per Sortie

We complemented Goldberg's previous research by exploring a greedy heuristic in the converse situation, when the sortie constraints are binding but the budget constraint is not. The expected kills-per-sortie (EKS) heuristic considers only combinations that destroy the most targets per sortie. This criterion, like the LCTK criterion, favors weapons with a high probability of kill; but rather than contrasting that probability against procurement cost, it favors aircraft with a high load factor so that more weapons can be delivered on a single sortie.

In an illustrative analysis of the EKS heuristic, we considered an aircraft that can carry either four 500-pound bombs or two 1,000-pound bombs on a single sortie. A sortie against unprotected or immobile targets might destroy an average of 1.2 targets if loaded with four 500-pound bombs but only 1.0 targets if loaded with two 1,000-pound bombs. However, a sortie against heavily defended or moving targets might destroy an average of 1.6 targets if loaded with four 500-pound bombs but 1.8 targets if loaded with two 1,000-pound bombs. The EKS heuristic would steer the model solution toward the preferred weapon type in both situations.

Performance Exercise

Next, we performed a computational exercise to demonstrate the tradeoff between increased calculation speed versus loss of quality when using these heuristics. We estimated the computational advantage to prescreening the aircraft-weapon-target combinations so that only those satisfying either the LCTK criterion or the EKS criterion enter the problem. Using a realistically sized problem for the U.S. Air Force, we compared results using our heuristics against results obtained with the full optimization model after one hour of calculation time. In our simulations, the heuristics often came up with a solution in just 10 minutes that was superior to the solution the full model came up with in an hour.

Approach

We drew our parameters primarily from Wirths (1989), who provided an unclassified data set derived from the U.S. Air Force's then-current Joint Munitions Effectiveness Manual. Wirths supposed that two aircraft were available to fly sorties limited to 108 and 81, respectively, over the course of a campaign. He considered 24 weapon types and 13 target types. For the target types, he provided target populations, their relative utility values, and decoy rates.

Wirths did not provide data pertaining to weapon costs, load factors, or kill probabilities, so we simulated these parameters to allow us to test the efficacy of our heuristics under a variety of conditions. In all, we generated 10 sets of random parameters to use in our experiments. For each set of parameters, we tested a series of four models: the full aircraft weapon budgeting model (no heuristic), the EKS model, the LCTK model, and a hybrid model having all of the variables in the EKS and LCTK models. All models were run using IBM ILOG CPLEX optimization modeling software on an Intel Core Duo central processing unit at 2.67 gigahertz.

We allowed the full model a maximum runtime of one hour, and each heuristic model was permitted a maximum runtime of 10 minutes. We first ran each model with no budget constraint and recorded the highest budget requested by any of the models. We then tested each model with a budget constrained to, respectively, 90, 80, 70, and 60 percent of the unconstrained requirement.

Results

We provide the results of our demonstration in Table 1, where the highest performing figures at each budget level is in bold. The results show that the speed-quality tradeoff is quite favorable: by applying the heuristics, solutions were calculated much faster and the quality was only occasionally less than that of full models to which we did not make any predetermined exclusions of aircraft-weapon-target combinations. Having the choice of three heuristics allows for useful alternatives to the full budgeting model across the entire range of budget constraints tested.

Table 1. Comparison of Results for 10 Sets of Parameters

	Budget constraint	Full model	EKS model	LCTK model	Hybrid model
Panel A: Average objective values	100%	838.1	**853.7**	736.5	850.1
	90%	723.4	701.5	697.8	**728.4**
	80%	647.3	594.6	**654.2**	648.8
	70%	599.7	537.2	**610.0**	602.6
	60%	489.0	456.5	**506.1**	493.5
Panel B: Average weighted kill percentages	100%	73.4%	**75.1%**	64.8%	75.0%
	90%	63.7%	61.8%	61.6%	**64.1%**
	80%	57.5%	52.7%	**58.5%**	57.8%
	70%	52.7%	47.7%	**53.9%**	53.6%
	60%	43.3%	40.5%	**45.0%**	44.0%
Panel C: Counts of objective values superior to full model	100%	n/a	**10**	0	9
	90%	n/a	2	3	**10**
	80%	n/a	0	**8**	**8**
	70%	n/a	0	**10**	8
	60%	n/a	0	**10**	**10**

When assigning an unlimited budget, we found that the EKS model yielded superior objective values relative to the full model in all 10 trials despite being given a maximum of only one sixth of the solve time. EKS model solutions required an average of only 13.8 seconds to eclipse the objective value that the full model achieved in an hour. Thus, even when run for only a short period of time, the EKS model provided excellent results for estimating military requirements. The EKS model was not nearly as effective when run with restrictions upon the budget. When the budget was restricted to 90 percent, the EKS model achieved superior solutions only twice in 10 trials, and when the budget was even further restricted, the EKS model did not produce any superior solutions (Panel C of Table 1).

Compared to the EKS model, the LCTK model performed particularly well with a budget constraint below 80 percent, providing superior objective values relative to the full model in all 10 trials with budgets of 60 and 70 percent. At a budget of 60 percent, the LCTK model required an average of 31.1 seconds to eclipse the full model's objective values, and at a budget of 70 percent, it required an average of 74.7 seconds. The LCTK model was fairly versatile in that it was effective at various budget levels.

The hybrid model was effective regardless of the level of budget. On average, it produced superior solutions relative to the full model across all budget constraints that we tested, and it outperformed the full model in 45 of 50 pairwise comparisons.

In short, when examining military requirements with an unlimited budget, the EKS model provided an expedient alternative to the full model. Conversely, when considering a situation in which the budget is quite restricted, the LCTK model reliably supplied superior solutions in short periods of time. The hybrid model provided excellent solutions regardless of the budget constraint, and it is especially useful when applied in situations with a moderate budget constraint.

The results observed are explained, in large part, by the extent to which each heuristic reduces the size of the problem being solved. In mathematical terms, we use X_{ijk} to denote the number of weapons of type j delivered by aircraft of type i against targets of type k. The full models that we generated using the 10 samples studied averaged X_{ijk} variables numbering 152.4.

As shown in Table 2, the number of variables considered by our heuristic models varied widely, while the number of positive (nonzero) X_{ijk} variables in the solutions they reached show far less variability. The full model expended great computational effort with variable selection, whereas the heuristic models determined many or most of the positive variables in advance.

Table 2. Variables Considered

	Average number of Xijk variables	Average number of positive Xijk variables	Average percentage of Xijk variables selected
Full model	152.4	19.7	12.9%
EKS model	26.0	13.6	52.4%
LCTK model	13.0	9.1	70.0%
Hybrid model	38.8	17.1	44.1%

Conclusion

We used analytical and numerical methods to complete the theoretical understanding of the military's aircraft weapon budgeting problem. First, we extended the work of Boger and Washburn (1985) by deriving an expression that accounts for the possibility that some dead targets may appear live, acting as decoys and drawing additional fire.

Second, we demonstrated the applicability of two greedy heuristics that identify the preferred weapons to maximize a utility function defined over targets destroyed of various types. Goldberg (1991) had provided a formal justification for the LCTK heuristic, in a situation where a binding budget constraint exists but sortie constraints do not. We developed a corresponding criterion using highest EKS for the converse situation in which binding sortie constraints exist, but the budget constraint does not. Although a monetary cost-to-kill ratio does not apply in such a situation because there is no active budget constraint, the scarcity of sorties motivates an alternative criterion. In this latter situation, the preferred weapons are those that economize on scarce sorties by offering the highest EKS.

Finally, we investigated the ability of those two heuristics to reduce the dimensionality and accelerate the solution in a realistically sized problem for the U.S. Air Force. The EKS heuristic correctly preselects the aircraft-weapon combinations that appear in the optimal solution when the procurement budget is fully funded. For our Air Force example, when the procurement budget is set at 60 percent to 80 percent of full funding, the LCTK heuristic achieves a better solution than the full optimization model does in a fraction of the runtime. When the procurement budget is set at about 90 percent of the requirement, the solution may be accelerated by a hybrid approach that includes only the subset of aircraft-weapon combinations that are suggested by either of the two heuristics.

References

Boger, D., and A. R. Washburn. 1985. *Notes from the Stockpile Seminar*. NPS 55-85-014. Naval Postgraduate School: Monterey, CA, https://ia800302.us.archive.org/29/items/notesfromstockpi00boge/notesfromstockpi00boge.pdf.

Dantzig, G. B. 1957. "Discrete-Variable Extremum Problems." *Operations Research* 5, no. 2: 266–277.

Goldberg, M. S. 1991. "Aircraft Weapon Budgeting," *Naval Research Logistics* 38, no. 6: 799–816.

U.S. Department of Defense, Joint Chiefs of Staff. 2019. *Joint Air Operations*. Joint Publication 3-30. https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/jp3_30.pdf.

Wirths, K. P. 1989. "A Nonlinear Programming Model for Optimized Sortie Allocation." Master's thesis in operations research. Naval Postgraduate School. https://calhoun.nps.edu/handle/10945/26102.

About the authors

Matthew Goldberg is Deputy Director of the Cost Analysis and Research Division of IDA's Systems and Analyses Center. He holds a PhD from University of Chicago in economics.

David Goldberg is an assistant professor in the Management Information Systems Department at San Diego State University. He holds a PhD in Business Information Technology from Virginia Tech.