Less, But Better: Optimizing Naval Manpower Efficiency

Emily Frost	Audra Podany	Andrew Martin
Booz Allen Hamilton	Booz Allen Hamilton	Booz Allen Hamilton
Norfolk, VA	Washington, DC	Arlington, VA
Frost_Emily@bah.com	Podany_Audra@bah.com	Martin_Andrew2@bah.com

ABSTRACT

Advancements in algorithms and computing power now allow for analysis that was previously performed at long intervals to be done more quickly, more often, and with more flexibility. The Naval Manpower Requirements System, employed by the Naval Manpower Analysis Center (NAVMAC), is one system built for such legacy analyses. Its current manual processes and heuristic-based algorithm result in sub-optimal, and infrequent updates of, its output: manpower requirements documents. These shortcomings leave the Fleet subject to minimized manpower forecasts and contribute to suboptimal manning across the board, from recruiting to detailing to retention.

Our solution combines a new end-to-end software platform with a modern optimization approach. The platform enables analysts to seamlessly upload requirements for naval activities and to optimize their manpower determinations. This system uses a mixed-integer linear programming approach, powered by Python and best-in-class optimization solvers, to return the provably optimal result each time for the given inputs. Additional parametric and sensitivity analysis tools provide increased confidence to those results and overcome uncertainty in the inputs. The formulation allows NAVMAC to more rapidly provide improved forecasts for Fleet manpower, while still being flexible enough to support changing business rules and requirements – thereby significantly improving the Navy's ability to predictively manage manpower across the manning lifecycle.

ABOUT THE AUTHORS

Emily Frost is an associate within the Analytics division of the Strategic Innovation Group at Booz Allen Hamilton. Her background as a surface warfare officer in the U.S. Navy informs her drive to support clients in adopting data- and analytics-driven decision-making. Her areas of interest include operations research, analytics strategy, and machine learning. Ms. Frost holds degrees in mathematics and English literature from the United States Naval Academy, and a master's degree in operations research from the Massachusetts Institute of Technology.

Audra Podany is an associate within the Digital division of the Strategic Innovation Group at Booz Allen Hamilton. Her primary focus and passion lie in tackling large, complex, digital transformation projects for both commercial and government clients. She has a strong background in business process improvement, data analysis, and operations research. Ms. Podany graduated from the Massachusetts Institute of Technology with a bachelor's degree in management science and minor in mathematics.

Andrew Martin is a senior consultant within the Analytics division of the Strategic Innovation Group at Booz Allen Hamilton. His areas of expertise include optimization, operations research, and energy analytics. He also has experience in nuclear safety oversight and commercial building energy management. Mr. Martin holds degrees in civil engineering and business from Purdue University.

Less, But Better: Optimizing Naval Manpower Efficiency

Emily Frost
Booz Allen Hamilton
Norfolk, VA
Frost_Emily@bah.com

Audra Podany
Booz Allen Hamilton
Washington, DC
Podany_Audra@bah.com

Andrew Martin
Booz Allen Hamilton
Arlington, VA
Martin_Andrew2@bah.com

INTRODUCTION

As part of the Navy's digital transformation, the objective of this effort is to develop an end-to-end digital platform that leverages modern technologies and development methods to replace legacy systems and heuristic-based algorithms for determining Naval manpower requirements. The solution aims to improve the robustness and defensibility of manpower requirements for Congress and other senior decision-makers. Furthermore, the new system, combined with process reengineering, will drastically reduce the time and effort needed to determine Fleet manpower requirements by automating manual processes, streamlining workflows, and digitizing outputs. This paper will focus on the new optimization approach and structure, rather than providing details on the entire workflow development and software stack of the new manpower determination system.

APPROACH

Before explaining our optimization approach, the current Naval Manpower Requirements System (NMRS) must be understood so that the distinction is made apparent. First, it operates on a divisional level, with each organizational unit considered independently. In each division, it uses directed requirements to create manpower requirements (which can be thought of as eventual people assigned to the ship) that are associated with "buckets" of available hours for that manpower requirement to complete watch and work. If there are more hours to assign than remaining bucket space, then NMRS will "grow" another position requirement and continue. Once finished, Naval Manpower Analysis Center (NAVMAC) analysts will review the results, and if the results do not meet their expectations, they adjust the inputs and try again. Typically, this process involves adding additional directed requirements to artificially influence the original position composition, above and beyond input requirements generated by stakeholders such as the Fleet and higher headquarters.

Instead, we formulate a mixed-integer linear programming (MILP) problem to minimize cost of Navy manpower requirements across a naval ship subject to all categories of work, watch, and directed requirements. The formulation optimizes manpower requirements across an entire ship, as opposed to filling requirements by division. Global optimization, as we refer to it, results in a more optimal answer than divisional optimization because input requirements can often be fulfilled by individuals who reside in different departments. In the previous approach, these requirements were forced to fall entirely within individual divisions. With global optimization, they can be distributed across the ship to take advantage of any slack (unused hours) that may appear in individual departments. Furthermore, by structuring our approach as a MILP, we obtain a solution all at once, rather than by iteratively building up positions. Thus, we can be certain that the solution is indeed a global minimum and that as much slack has been wrung out of the system as possible.

Before decomposing the components of the MILP, it is prudent to note what this approach is not solving. NAVMAC's charter is to provide the minimum manpower required to man a ship without being constrained by available Sailors and without consideration of available funding. In doing so, Ship Manpower Documents (SMDs) provide a true baseline of what manpower is actually required in order to deploy a ship during wartime. Also, although SMDs can provide a useful input for planning Navy recruitment, retention, and training as a whole, this project does not attempt to directly contribute or optimize in those areas.

Additionally, the MILP does not attempt to address any inaccuracies present in the data inputs. NAVMAC has long been aware that the data inputs do not fully capture all work and all watches required by the modern Navy; however, as mentioned above, the legacy NMRS toolset does not provide the capability and flexibility required to handle complete and up-to-date data inputs. The MILP takes the current data inputs as givens, leaving the problem of ensuring and improving their accuracy as a separate scope of work.

Finally, this approach does not attempt to transform NAVMAC's current methodology into a scheduling problem. Work inputs are in the form of average hours per week and do not necessarily reflect conditions onboard the ship at any particular point in time. Once an SMD has been produced and the ship has been manned by a resource sponsor, personnel may get sick or go on leave, equipment may need maintenance more or less frequently than expected, or work may need to be staged in certain ways. It is up to the leaders on the ship to determine who does what when in accordance with their availability and the current situation onboard.

STRUCTURE

As a mixed-integer linear program, the goal is to select values for certain integer and continuous variables so as to minimize the value of an objective function while respecting the constraints.

Objective Function

The objective function for the optimization is to minimize the dot product of the position variables and their respective costs. Mathematically, we can represent the objective function as:

$$\boldsymbol{min} \sum_{l \in L} \boldsymbol{cost_l} \times \boldsymbol{positions_l} \qquad (1)$$

with *L* representing the set of all labor types. We define a labor type to be a specific combination of a rating, paygrade, and zero or more Navy Enlisted Classifications (NECs) that identify specific skillsets. In general, the cost of a labor type increases with increasing paygrade and number of NECs. Therefore, the model attempts to minimize not only the total number of positions on the ship, but also their seniority and necessary training. The result does not preclude additional or more-qualified individuals from staffing the ship, but rather establishes a rigorous baseline for when those staffing decisions are made.

Variables

The basic model uses four types of variables: position, watch, work, and directed. The approach frames the labor capacity problem in terms of an average workweek, in which each manpower requirement (mapping to positions in our formulation) has a given number of labor hours available. The labor requirements come in the form of watch stations, work (preventive and corrective maintenance, own unit support, etc.), and directed requirements (those skills which have been determined by senior leadership to be required on a ship).

The position variables are nonnegative integers that represent the number of each labor type on the ship. For example, one labor type might be an FCA2 with the SPY-1D(V) NEC – in other words, a second-class petty officer (2) rated to work on the Aegis Fire Control System (FCA) and specifically trained to work on the SPY radar system. In effect, the position variables summarize the breakdown provided by the watch, work, and directed variables.

The watch variables show how much of a given watch station requirement is fulfilled by a given labor type. These variables are sparse: rather than the full Cartesian product of watch stations and labor types, watch variables are only created where a given labor type could stand a given watch station in accordance with the business rules discussed below. Watch variables are also either integer or continuous depending on a property of the watch station indices called condition. Specifically, elevated conditions of readiness such as General Quarters (COND I) or Amphibious Operations (COND IA) are treated as integer variables because there are not hours associated with those watches: the conditions are stood up for as long or as short as the commanding officer deems necessary. Normal steaming conditions, such as COND III, are treated as continuous variables because they have known hours per week associated with them.

Work and directed variables are very similar to watch and illustrate how the various work and directed requirements are fulfilled by the labor types, except that work variables are always continuous, and directed variables are always integers. For example, a work input requirement might specify an average of 50 hours per week of equipment maintenance that must be performed by an FCA2 or higher. Those 50 hours need not all be performed by the same position and can be distributed across position variables in a continuous manner. Directed requirements, on the other hand, are comprised of skills like a command master chief. Each ship must have one individual rated as a CMC, and therefore these requirements are treated as integer variables.

Constraints

There are two types of constraints in the model. The first ensures that all of the requirements are met. For example, the sum of all of the watch variables corresponding to a given watch station must equal the requirement for that watch station. The same must also be true for the work and directed variables. So, if work type 2 requires an average of 50 hours of weekly labor and could be performed only by a BMC or a BMCS, then the following would become a constraint:

$$\boldsymbol{work[2, BMC] + work[2, BMCS] = 168} \qquad (2)$$

The second type of constraints ensures that there are sufficient positions of each labor type to fulfill the tasks that they are assigned in the watch, work, and directed variables. These constraints must be fulfilled both in terms of position counts and in terms of hours.

As a reminder, a directed requirement is a mandate that a specific labor type must appear on the ship regardless of whether there is any watch or work available for him to actually perform. So, the number of positions of a given labor type must equal or exceed the sum of the directed requirements variables for that labor type. For example,

$$\boldsymbol{position[BMC] \geq directed[1, BMC] + directed[3, BMC] + directed[4, BMC]} \qquad (3)$$

where a given labor type happens to be eligible to fill three different directed requirements.

Likewise, some of the watch station requirements, depending on their condition, function similarly to directed requirements. However, conditions of readiness are wholly disjoint. In other words, an individual in a single designated positon could fill a Condition I watch station, a Condition IA watch station, and a directed requirement. During such conditions only watch stations are manned and other work is not performed; it is similarly important to note that only one condition of readiness pertains at a time.

However, in Condition III, work is performed, and both work and watch are measured in terms of hours. So, for each labor type, the number of positions is scaled by the number of hours in a standard workweek, and this product must be greater than the sums of watch and work that that labor type is assigned. An example might look like

$$\begin{aligned}\boldsymbol{position\cancel{ersonnel}[BMC] * 67} & \\ & \boldsymbol{\geq watch[1, BMC] + watch[3, BMC] + work[2, BMC]} \\ & \boldsymbol{+ work[3, BMC] + work[4, BMC]}\end{aligned} \qquad (4)$$

if a BMC happens to only be able to stand those two watches and perform only those three types of work.

Assignment and Regrouping

The above model succeeds in its purpose of determining the most efficient manpower requirement mix for the ship, but because the ship is still administratively organized into divisions (identified by org codes) in the real world, every individual needs to be assigned to a particular division. Such assignment should respect the org codes of the input requirements that he is fulfilling as much as possible. To do so, a secondary optimization is performed, based on the existing model.

First, the existing position variables are fixed so that their values cannot change. Then new position, watch, and work variables are created that are indexed by combinations of labor types and their feasible org codes, rather than just labor types. So if the original optimization says that four BM3s are required, and if they can feasibly be assigned to either the 040CAA or 040DAA org codes, then two new variables will be created for each of the originals, and the values of these new variables must sum to the values of the originals. An example might look like

$$\boldsymbol{position}[\boldsymbol{BM3}] = \boldsymbol{position}[\boldsymbol{BM3}_{\boldsymbol{040CAA}}] + \boldsymbol{position}[\boldsymbol{BM3}_{\boldsymbol{040DAA}}] \quad (5)$$

Similarly, for work, there would be many new constraints in the form of

$$\boldsymbol{work}[\boldsymbol{B}, \boldsymbol{BM3}] = \boldsymbol{work}[\boldsymbol{B}, \boldsymbol{BM3}_{\boldsymbol{040CAA}}] + \boldsymbol{work}[\boldsymbol{B}, \boldsymbol{BM3}_{\boldsymbol{040DAA}}] \quad (6)$$

Additional constraints are also added, similar to before, to ensure that there are sufficient positions for the watch and work that is assigned to them.

Finally, a new objective function is created that minimizes the sum of weighted “distance” measures among the org codes. For example, a BM3 assigned to the 040CAA org code who is also assigned work from the 040CAA org code would carry zero cost in this “distance” measure, because the org codes are the same. If that same BM3 were assigned work from the 040DAA org code, a small cost would apply because the org codes differ but still belong to the same 040 overarching organizational level (department, in this case). On the other hand, if he were assigned work from the 079DAA org code, a different department, the associated “distance” cost would be much higher.

Minimizing this cost function accomplishes both assigning all labor types to specific org codes and regrouping the original watch and work variables to reduce how much labor (either watch standing or work tasks) is done outside of one’s own org code, without affecting the overall mix and number of labor types on the ship.

BUSINESS RULES

Allowed Matrices

In order to intelligently create constraints that adhere to all business rules, the business rules are effectively coded into "allowed" matrices. An allowed matrix is a two-dimensional array with labor types as column names and input requirements as indices. Each entry is either true or false to indicate whether that labor type is allowed to fulfill a given input requirement, and the optimization modeling module only creates the watch, work, and directed variables where the entries are true. The allowed matrices also allow for easy visual inspection of the first type of constraint – ensuring that all requirements are met – row by row. For instance, if a given row of an allowed matrix has no true entries, then something has gone wrong with implementing the business rules or a necessary labor type is missing, as that requirement could never be fulfilled. Note: while the allowed matrices show which labor types are allowed to fulfill input requirements, they do not indicate whether or not a particular labor type will actually be selected to partially or completely fulfill that input requirement in the solution. An example allowed matrix for directed requirements is shown in Table 1.

Table 1. Sample Directed Requirements Allowed Matrix

		Labor Types								
		BMSN	BM3	BM2	BM1	BMC	FCSN	FC3	FC2	FC1
Directed requirements	SN	T	T	T	T	T	T	T	T	T
	BM3	F	T	T	T	T	F	F	F	F
	BM1	F	F	F	T	T	F	F	F	F
	FCSN	F	F	F	F	F	T	T	T	T
	FC1	F	F	F	F	F	F	F	F	T
	UTC	F	F	F	F	F	F	F	F	F

Rating and Paygrade Rules

In general, the rating of a particular labor type must match that of the input requirement in order for it to be eligible. However, the input requirements only specify the minimum paygrade needed, and any labor type with a paygrade equal to or greater than that listed in the input requirement is eligible, assuming that it has the same rating. For instance, if a requirement specifies a need for a boatswain's mate third class (BM3), then anyone from a boatswain's mate third class up to a master chief boatswain's mate (BMCM) would also be eligible, and this would be reflected by trues in the allowed matrix.

There are other peculiarities when assembling the allowed matrices that are more nuanced than need to be discussed in detail here. They come into play, for instance, when certain ratings merge at higher paygrades, when dealing with "undesignated" E-3 paygrades, and when an input requirement only requests a generic petty officer, rather than one with a specific rating.

NEC Rules

Beyond rating and paygrade qualifications, many input requirements also list necessary Navy Enlisted Classifications (NECs), which identify additional training and certification requirements. So even if a labor type meets the criteria above, if he does not hold the proper NEC as well, then he is not eligible to fulfill that input requirement. The current formulation allows labor types to hold up to two NECs, and so if two input requirements list two different NECs, it is still possible that they are filled by a single labor type. In reality, a given individual can hold many NECs, but by only using two, rather than three or more, this helps to avoid a combinatorial explosion in the number of potential labor types that the model must solve for, as well as make it easier to find individuals who hold the correct combination of NECs when the ship is ultimately staffed.

RESULTS

The optimization not only produces the total positions needed to man a ship, but also provides a comprehensive breakdown of how each manpower requirement is created. For example, if the output required three boatswain's mates second class (BM2s), the full optimization output would also show exactly which watch stations, work, and directed input requirements those BM2s were satisfying. This type of traceability is key in building trust in the new optimization.

Tools in Use

After originally prototyping the model in Excel, it has since been translated to and improved with the Python programming language. Specifically, the NumPy and Pandas libraries are used to craft the allowed matrices that encode the business rules. Then, the Pyomo modeling library uses these matrices and additional information to craft the necessary variables, objective function, and constraints. However, Pyomo only builds the model and cannot solve it itself. For that, it interfaces with dozens of commercial and open-source solvers. For smaller problems, the open-source GLPK package is sufficient for performing the optimization. For larger problems, however, the powerful commercial Gurobi solver is necessary. It can be accessed via Pyomo, though Gurobi's native Python wrapper runs faster and provides better control.

IMPACT

The integration of the MILP into an end-to-end platform provides numerous benefits and efficiencies to the current processes in place at NAVMAC. The powerful, modern tools supporting this platform allow for SMDs to be generated in a fraction of the time that the process previously required using the legacy system. Additionally, the system allows for multiple optimization runs to be accommodated, saved, and viewed. Analysts can run if-then scenarios after changing input requirements or parameters and see results live.

The use of a MILP establishes an improved level of trust in the manpower requirement outputs. The transition from a pseudo-greedy heuristic to a provably optimal MILP allows NAVMAC analysts to refrain from editing the inputs solely to influence the outputs. Manual workarounds and manual data manipulation are no longer required to generate a reasonable answer. As mentioned previously, the data generated from the optimization provide full traceability of each manpower requirement. This level of detail within the data ensures that the complete picture is available at all times. While that granularity of data is not required for all types of stakeholders, its presence allows for a variety of reports and dashboards to be built on top of the data. Finally, the quality of provability provides a form of "ground truth" that has never previously been available to NAVMAC. For example, if a senior decision-maker were to challenge the manpower requirements determined by NMRS, the analysts at NAVMAC had no way to rigorously justify why those particular numbers had been generated.

A major pain point felt by NAVMAC analysts is the inability to understand the exact execution of business rules. With no formal documentation of the legacy system available, analysts take their best guesses as to how business logic is encoded and how the heuristic performs its calculations. Our approach was designed specifically to supply transparency into the business rules, as well as support future updates to those business rules with minimal rework required. As Navy policies continue to grow and develop, this functionality is key to supporting reliable and consistent outputs.

Next Steps

As discussed above, there are still areas where the manpower requirement determination process could be significantly improved. One such area consists of the input data, while another consists of flexible procedures across multiple ships. Further in the future, this optimization formulation could contribute to dynamic manning of Naval ships.

Ideally, the optimization formulation would be able to take advantage of "live" data – that is, data that is as up-to-date as possible for the input requirements. Opportunity exists in this realm due to the fact that preventive maintenance and corrective maintenance data already reside in electronic databases; no new processes need be created to capture this data. The natural next step would be to construct a pipeline from those databases, and ingest that data directly, thereby eliminating extra manual steps current analysts perform to generate the data.

Further in the future, this data could be used as the basis of predictive modeling – machine learning or similar – to create predictive and corrective maintenance inputs that would accurately forecast future demand. Inputs that accurately reflect future demand would be valuable because the manpower requirements determination process is inherently forward-looking, but currently forced to rely only on backward-looking historical data. Forecasted inputs could still be fed into the MILP to produce the optimally efficient manpower requirement mix *for future work and watch requirements*, not for historical demand.

The optimization formulation, thus far, has been thoroughly tested on only one ship class – LPD-17 – chosen by the end users based on their upcoming needs. Next steps include running additional ship classes through the optimization formulation. The approach taken in building the MILP is meant to generalize to any ship class and as the formulation construction reaches its completion for LPD-17, our hypothesis will be tested as we receive data for other ship classes. Currently in testing are DDG-51, LHD-8, and LSD-41 – two small platforms and two large.

Disclaimer: Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of the Navy or the Department of the Air Force – Air Force Installation Contracting Agency.