

Agent-Based Modeling as a Tool for Manpower and Personnel Management

Manuscript #334

Authors:

Iavor Trifonov
Oliver Bandte
Eric Bonabeau
Paolo Gaudiano

Icosystem Corporation
10 Fawcett Street
Cambridge, MA 02138, USA
{oliver,iavor,eric,paolo}@icosystem.com

Corresponding Author:

Paolo Gaudiano
Icosystem Corporation
10 Fawcett Street
Cambridge, MA 02138, USA
paolo@icosystem.com
+1-617-520-1070

Agent-Based Modeling as a Tool for Manpower and Personnel Management

Iavor Trifonov, Oliver Bandte, Eric Bonabeau, Paolo Gaudiano
Icosystem Corporation
10 Fawcett Street, Cambridge, MA 02138, USA
{oliver,iavor,eric,paolo}@icosystem.com

1 Abstract

We have developed an agent-based model of the US Navy's Manpower and Personnel (M&P) systems, and used the model as a tool to analyze and design M&P policies. The model captures the dynamics of sailor recruitment, training and retention, as well as their performance during missions. Our model makes it possible to gain a deep understanding of the dynamics of the entire M&P systems. We expect our tool to offer several benefits to the Navy, including the ability to design new policies for existing ships or new ships; the ability to understand the impact of shipboard technologies to increase automation; and the ability to study the impact of various interventions on sailor retention. The model also promises to be useful for personnel management in the commercial sector.

2 Overview

The work described in this paper, which was supported from a Phase I SBIR from the Office of Naval Research (contract N000144-04-M-0167), provides a quantitative framework for analysis and policy design for Manpower & Personnel (M&P) systems of the United States Navy. An improved understanding of M&P systems is critical in addressing current and future needs of the Navy (see GAO report 03-520): the number of responsibilities a sailor has on a ship is increasing while the amount of time available to complete those tasks is fixed; continuous overloading of sailors with watch standing, operations or maintenance duties can lead to degraded performance of both the battlegroup mission and ship operations; in addition, sailor burnout increases, which leads to a decrease in personnel retention. Understanding the effects of mission requirements on the behavior of a sailor is of great importance, especially given tighter budgets, longer deployments and possibly lower levels of recruitment (Moore et al., 2002). Further, being able to understand how the behavior of a sailor can impact mission readiness is of great importance in determining whether a mission is likely to succeed. The use of computational models that can integrate traditional Operations Research (OR) approaches (e.g., Holder, 2005) with the sciences of Complex Adaptive Systems (CAS) represents a potentially significant contribution to support the process of shipboard M&P system planning.

Even if the effect of single policies on an individual was understood, being able to understand their compounding effect on the crew, battlegroup, or Navy has been largely left to qualitative measures. The lack of powerful computational tools hinders the Navy's ability to react to new processes, technology or changes in M&P systems. Understanding

the effects of planned military operations on the sailor community as a whole is crucial to performing robust shipboard M&P planning in the long run. Furthermore, a quantitative approach to understanding the dynamics of M&P is crucial to designing new ships with new capabilities, or retrofitting existing ships with technologies for automation.

With the support of an ONR SBIR, we used a CAS technique called Agent-Based Modeling, or ABM (Axelrod, 1997; Bonabeau, 2002) to develop a quantitative model of Shipboard Manning and Personnel Behavior. The development of this model was driven by two goals. First, to demonstrate that the dynamics of activities on the ship is integral to the M&P process in the Navy. Specifically, we have shown that what happens on the ship has a strong impact on the sailor community as a whole, demonstrating that models that do not capture these dynamics are incomplete.

The second goal was to demonstrate that ABM is a methodology that can easily capture the intricacies of a ship's dynamics while at the same time transparently scaling to the aggregate level where policy design and decision making take place. We demonstrated that disaggregating sailors by their characteristics and modeling the interactions between them is natural to ABM and important to the problem at hand.

Our work demonstrates that it is possible to leverage quantitative computational techniques to gain a better understanding of personnel management. The complexity of personnel management is a general problem that afflicts many large organizations, both Government and private. Our findings thus promise to have a significant impact by helping to improve policies and reduce costs for the Navy, for other Government agencies and for large corporations in the private sector.

3 Background: the Navy M&P system

Manpower & Personnel is the process of managing servicemen and workers in the Navy. On a large scale, this includes the determination of manpower requirements, the budget allocation for fulfilling these requirements (manpower programming), the provision of people to fulfill the requirements (personnel planning), and the distribution of the available personnel. Figure 1 is an illustration of the M&P system.




Figure 1: Overview of the Navy M&P system.

The M&P process is typically perceived at an aggregate level, representing the aggregate flow of people within the personnel planning stage. During personnel planning people are recruited, trained, and ‘managed’ in communities by the so-called *community managers*. These communities serve as a pool of servicemen for fulfilling actual billets in the process of personnel distribution.

Currently, however, the actual fleet is perceived by the M&P process as an aggregate bucket-sink for the resources that have been acquired and are dispatched during personnel planning and distribution. Personnel is represented as an aggregate inflow of people ready for distribution, and two aggregate outflows – one for people who leave the Navy, another for people who remain in the Navy after serving a period in the fleet.

In reality, being at sea and serving in the fleet are integral parts of the M&P process, as every serviceman’s career consists of a string of sequential assignments on shore and at sea. When thinking about the M&P process at an aggregate level, however, the separation emerges. One of the reasons for this is the fact that the majority of the policy levers at the disposal of decision makers within the M&P process are located outside of the fleet ‘sector’. Policies governing the recruiting, training, retention, and overall community management of personnel affect the events at sea only indirectly and in a way that cannot be currently represented at the aggregate level.

In particular, at the ship level an XO has to match a list of diverse billets to a list of diverse and usually insufficient personnel, whose availability and distribution are a direct result of the M&P process. The particular skill set available to the XO and the way it is allocated to serve the ship, given the requirements, drives the dynamics of the events on the ship. These events constitute the performance of the ship, its readiness, the achievable workload, the ability to complete a mission – the things that the M&P is ultimately supposed to support.

What happens on the ship, however, is not directly a subject to the policies implemented by the M&P decision makers. They impact the life on the ship indirectly by making a particular set of sailors available to the XO. Furthermore, the M&P policy levers are applied on a different scale relative to a single ship, which further prevents the feedback of information directly from the fleet to the M&P process.

Understanding the fleet sector better may have two implications for the decision makers in the M&P process. One could attempt to examine the reasons behind the outflow of people and the loss of personnel as a direct consequence of people’s experience on the ship. On the other hand, a more detailed model of the fleet may be used to provide a more sophisticated metric for the ability of the M&P process to maintain an optimal inventory of people and skill sets at the disposal of the fleet. Using such a model of a ship (strike force, fleet), one would be able to tell how well a particular ship performed (in terms of achieved workload, fatigue/overload of the crew, readiness of the ship and so on) given the available and required resources. As the available resources come directly from the M&P process, knowing how the configuration of the pool of servicemen affects the fleet as a whole provides a benchmark for the calibration of the M&P process.

To accomplish that, one has to introduce more detail both within the M&P process and the fleet, so that the different skill sets are exposed on the ship and the available policy levers in the M&P process are made relevant to the diversity of the personnel in terms of

ratings and pay grades. Agent-based models provide the ability to model individual entities and observe global system behavior as a function of the interactions between them. Modeling sailor behavior, interactions, and decision making both on shore and at sea is a natural application of ABM to a complex system such as the Navy M&P process.

4 Modeling approach

For our purposes it is sufficient to model an individual ship and its interaction with the available supply of sailors through the M&P process. The representation of the M&P process is simplified as this project focuses primarily on the dynamics of a sailor's life on the ship. Nevertheless, we include an on-shore pool of servicemen, which is managed through recruitment, retirement, detailing, and training policies. The presence of a simplified M&P process is necessary to the extent that the impact of ship dynamics on the M&P process can become evident.

The career of a sailor is modeled as a sequence of assignments on shore and at sea. To capture that, our model includes two main components: a pool of sailors on shore and a ship with a crew at sea. At the end of each mission the ship will return to shore and gather a new crew from the available pool for the next mission, while the crew that just returned will join the pool for the duration of their next on-shore assignment.

Figure 2 shows an overview of the model structure.




Figure 2: Overall diagram of the Navy M&P system we simulated in Phase I.

The input to this model is a definition of a ship and an initial pool of available sailors, as well as a list of sailor types that could exist. The output of the model is comprised of time series data for variables such as: the sailors' daily distribution of time while on board a ship, ship department/division staff availability and utilization, on-shore pool composition, ship crew composition, sailor shortage composition, level of sailor

frustration per type and watch station, and so on. Different metrics for the performance of the system may be developed from these variables.

The time step for the model is the duration of a single ship mission (3 months), while the time horizon is determined arbitrarily. A smaller time step that would allow for modeling events during a mission is unnecessary for our purposes at this stage, as the structure of the ship is static. Every day on board is the same – no random events and effects are modeled – therefore everything that happens during a mission can be calculated upfront. However, a finer level of detail could be easily added to the model.

4.1 Model details

Agent-based models capture the behavior of the systems they emulate by simulating agents and their interactions (Bonabeau, 2002). In this section we describe all of the agents that are used in our model of the Navy M&P process.

4.1.1 The Sailor Agent

The sailor agent captures the properties and the rules of behavior of a Navy sailor. Sailor agents can be either on the ship at sea, or on shore. They are brought in the system through the recruitment strategy of the on-shore pool agent and can leave the system either by retiring or by declining to reenlist.

The defining characteristics of the sailor agent are its rating and pay grade. Higher pay grades mean higher experience for a sailor, while the rating refers to the particular profession (skill set) for which the sailor specializes. Ratings determine the billets to which sailors can be assigned while at sea and therefore determine the watch stations and divisions they serve on. The rating and the pay grade can be adjusted through training, which is determined by the on-shore pool agent.

While on the ship, each sailor agent is assigned to a particular billet and serves on a set of watch stations. This determines the amount of departmental work that that sailor is able to perform, apart from the required watch standing. When the amount of departmental work completed by a sailor differs from the expected amount, the sailor experiences *frustration* by an amount proportional to the difference between expected and actual work completed during the mission.

The frustration experienced by sailors is a key variable in the model. The amount of frustration experienced by each sailor is recorded for each mission, and its memory lasts over time (we model this as a weighted sum of past frustration, with more emphasis on the most recent mission). The accumulated level of frustration determines the probability of a sailor quitting the Navy at the end of their enlistment period.

4.1.2 The On-Shore Pool

The on-shore sailor pool is essentially a bucket of available sailors with an inflow of new recruits and a sink for retiring conscripts. However, sailors that are spending time in the pool are going through training and pay grade adjustment processes.

At the beginning of every mission, a ship crew is compiled from the on-shore pool. The composition of the crew is based on what is required (the ship billets) and what is

available (what is in the pool). The actual matching of available sailors to billets obeys a set of rules that are a simplified representation of the detailing process (see below).

4.1.3 Recruitment

The flow of sailors into the on-shore pool is a simplified representation of the recruitment process. The recruitment policies are based on two assumptions: new sailors can only be hired with pay grade one, and only with initial ratings. Initial ratings would be ratings that do not require in-house training (seamen – SN and firemen – FN). Any ratings that are not designated as initial require a sailor to follow a chain of advancement through training before they can become a specialist with that rating.

Beyond these two assumptions however, the recruitment policies can differ greatly. Typical recruitment strategies include:

1. Recruit the same number of sailors that left last time step (due to retirement as well as “quitting”).
2. Same as before, but recruit a minimum of X sailors even if nobody left.
3. Recruit a constant quota each time step.
4. In conjunction with any of the above, a cap may be applied to the total pool size.

The ability to select different recruiting strategies makes it possible to test various scenarios and understand the impact that each policy has on the size and stability of the sailor pool.

4.1.4 Detailing

The detailing process matches the demanded positions on the ship with what is available in the pool. To simplify matters, each billet on the ship is assigned a priority and the billets are filled in order of decreasing priority. This resolves the competition between billets that require similar types of sailors (rating and pay grade) and allows for controlled experiments with the detailing policy by simply rearranging the priorities.

If multiple sailors are found to be eligible for the same billet, the model chooses at random. If no sailor is found, the search process moves on to the next billet.

For our initial work we used this simplified detailing strategy. However, it would be simple to extend the model to allow for simulation of additional detailing strategies that reflect more accurately current practices. Furthermore, it is possible with this model to test the impact of other proposed detailing strategy (e.g., Holder, 2005).

4.1.5 Training and Advancement

The training process is ongoing and advances the sailors in ratings as well as pay grade. In our model, pay grade 6 is the highest pay grade. Advancement requires completion of a certain number of years in service or a certain number of missions.

4.1.6 Retention

By default, sailors in our model retire after 20 years of service, while each enlistment period lasts for five years. After each enlistment period the sailor has an option to re-enlist or not, which depends probabilistically on the personal level of frustration as

explained earlier. The retention policy may be experimented with by adjusting the sailor sensitivity to frustration. This is in effect is a simplified representation of retention strategies such as offering sailors reenlistment bonuses or other perks.

4.1.7 The Ship Agent

The ship is defined by a list of billets. A billet is a position on the ship, which requires a sailor with a particular rating and a range of possible pay grades. As mentioned earlier, a billet has an associated assignment priority. This is used during the detailing process and helps determine which positions on the ship should be filled first.

The ship billets are organized in departments and divisions. Departments and their respective divisions are responsible for getting a certain amount of work done, as reflected by the billets. Each division is expected to perform an amount of work equal to the sum of the expected work from each billet. Each sailor is expected to perform departmental work during their free time, including on-watch, off-duty time. This amount is calculated upfront for each billet, assuming that each billet is manned by a sailor.

During the simulation, however, not all billets are necessarily filled by sailors. This increases the expected work per manned billets, because the division is still responsible for the same amount of work. When a sailor does not have enough available time to complete all the expected work, frustration arises representing the pressure of incomplete tasks and burn out.

The administrative organization of the ship in departments, divisions, and their respective billets is fed to the model as input. For our purposes we designed a simplified ship which consists of four departments – Operations, Engineering, Weapons, and Supply. The Operations and Supply departments have each two divisions, Deck and Navigation for the Operations, Food Services and Administrations for the Supply department. Figure 3 illustrates this organization.

Each listed billet has a title (for differentiation in the model) and requires a particular rating and a range of pay grades. All billets are numbered from 1 to 32 as an identifier for the Watch-Station-Bill as well as in the model. This numbering reflects the assignment priorities as well.

In addition to the administrative organization of the ship, sailors are also organized in the so called Watch-Station-Quarter bill. This bill provides a designated watch station for every sailor that is on board. Sailor man their specified watch station whenever they happen to be on watch during the day.

The difference between the watch station organization and the administrative organization is that the watch station distribution of the crew is performed for every mission, while the administrative distribution of billets is static. Since only some billets are filled by an actual sailor, the Watch-Station-Quarter bill is different for every mission.

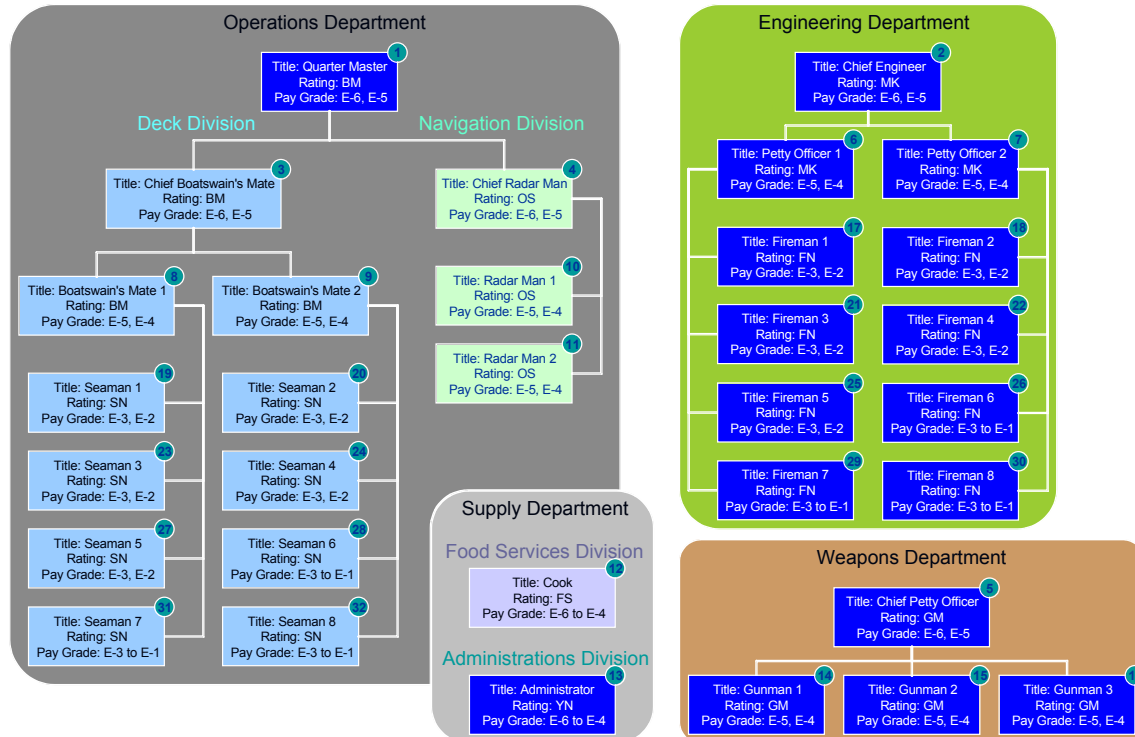


Figure 3: Administrative organization – Departments, Divisions and Billets.

4.1.8 The Watch Station Agent

Watch stations are implemented by the Watch Station Agent. Watch stations represent the positions on the ship in terms of its operations (e.g. helm, radar, engine, etc.). Similarly to the billets, different watch stations require different specialization of the sailors that will be manning them. To represent that fact, watch stations are defined with a set of eligible billets. This eligibility pool designates what positions (billets) are capable and/or expected to man the particular watch station.

The fact that the watch station eligibility pools are defined in terms of billets implies that they are dynamic, since different billets are filled by actual sailors every time. An empty billet cannot be assigned to man a watch station in the Watch-Station-Quarter bill.

Watch stations need to be manned 24 hours of the day. This requires a schedule for watch standing. For simplicity we model an even split of watch standing among all sailors that are assigned to a particular watch station. In order to be able to man a watch station with a different set of sailors for every watch period, there need to be at least two sets of different sailors (manned billets) that can stand watch every other period (1 in 2). This restriction is used to determine if the ship can or cannot sail. If it is impossible to assign at least twice the minimum required sailors per watch station per period, the ship's operations cannot be performed properly, the ship stays in port, and no mission is performed.

An important aspect of watch standing is that the watch station eligibility pools overlap: a given billet may be eligible for manning more than one watch station. This means that the

same billets may be assigned on different watch stations during different missions. In the model, the watch station workload is divided equally among sailors. Hence, when a certain billet cannot be filled, sailors from other billets need to share the additional watch station workload – even if the sailor comes from a different department from that of the missing sailor(s). This creates significant interactions within the model: departmental work is performed during the day-work time for a sailor as well as during their on-watch, off-duty time and the expected work per division is based on a full crew. When a billet is not filled, some watch stations get fewer people for assignment. This means that the people who are actually assigned will get less on-watch, off-duty time. Therefore all of them will have less time for departmental work than expected and since people from different departments are assigned to serve on the same watch stations, departments that are otherwise fully staffed suffer from the fact that other departments are understaffed.

It is important to understand that the watch station organization of the crew is completely separate from the administrative organization. This means that people who belong to different departments may happen to serve on the same watch station. This independence allows for people from different departments and divisions to interact while on watch. This mechanism is essential to the propagation of the effects of understaffing from one division to another across the administrative organization of the ship.

5 Summary of results

In this section we summarize some key results obtained with our model. The M&P process strives to maintain the right mix and amount of sailors so that “the right person can be delivered at the right place, at the right time.” This is a very complex problem given the diversity of sailor specialties, the budget restrictions, recruitment problems, training requirements, etc. Here we offer some observations in light of these problems.

First of all, we need to be able to maintain a constant size of the pool at a value that is in accord with budget restrictions as well as the requirements of the Navy. As explained earlier, in this model people may leave the Navy for two different reasons – they have either reached their retirement age, or they decided to quit due to frustration. The retirement age is a parameter that is relatively easy to control for policy makers, while the frustration of sailors is a product of the dynamics of the entire system.

Sailors who retire have been in the Navy for a while, and they have usually reached high standing in terms of pay grade and rating. On the other hand, people are only recruited with low pay grades and initial ratings. This poses the problem that the people who leave the system and those who enter it are not directly interchangeable. There is a pipeline that each new sailor needs to go through, in which he or she can advance to higher pay grades and ratings in order to be able to fill the more specialized billets. During that process, however, sailors are going on missions and may become frustrated and leave. As a consequence, we need to be able to recruit enough people so that a sufficient amount of them reach the higher pay grades and ratings before they quit due to frustration.

We begin with an unlimited pool size. Figure 4 shows the pool composition for a simulation where the pool is initialized with 100 random sailors. The recruitment policy is to simply hire as many people as were lost due to retirement and frustration during the last time step.




Figure 4: Pool composition for simple rehiring recruitment policy

The graph in Figure 4 displays the total number of sailors in each rating-pay grade combination. These amounts are displayed on top of each other and add up to the total size of the on-shore pool for every time step. Figure 4 shows a constant pool size, which is expected given that exactly the same amount of people are hired as had left. At the same time we observe a cyclical pattern in the bands of different sailor types. The amplitudes in the cycles of some sailor types are quite dramatic. At times, some of the bands are completely interrupted. This means that at times particular sailor types are not available at all.

Figure 5 shows the average frustration level for each of the SN (seaman) ratings (pay grades 1, 2, and 3). Similar frustration levels are observed for the other ratings as well. Higher frustration levels tend to plague the initial ratings and pay grades, since these are the types of sailors that are responsible for most of the work and watch standing on board. Even if just a few billets are not manned, the effect is felt across the entire crew.

Using a different hiring policy provides a different picture. Figure 6 shows the resulting pool composition when a minimum of 1 person is hired at each time step. We see a steady growth in the total pool size, although the cyclical pattern remains to some degree. We have the steady growth due to the fact that we always hire at least one new person – regardless of whether anybody left. On the other hand, if more than one person left the previous time step, we would hire that amount – more than one person. This explains the different rates of growth in the pool size. Nevertheless, even though towards the end of the simulation the pool has doubled in size, we still find high levels of frustration among the sailors, as revealed by Figure 7.




Figure 5: Seamen frustration for simple rehiring recruitment strategy




Figure 6: Pool composition for a minimum inflow of 1

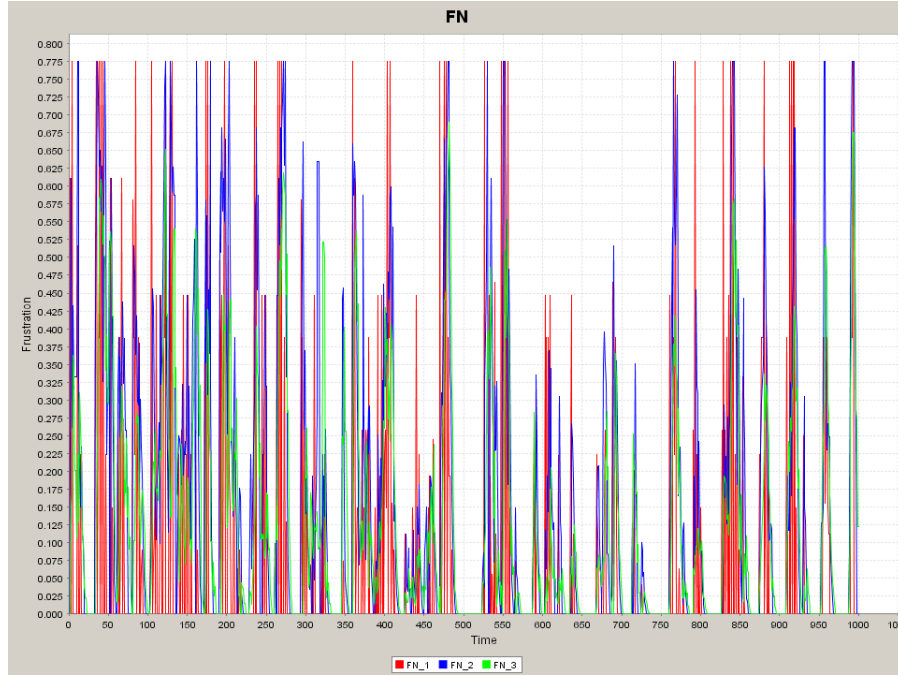


Figure 7: Average fireman frustration

These results demonstrate that the proposed recruitment strategy is unable to achieve a steady flow of people through it, nor low frustration levels.

If we gradually increase the minimum amount of sailors hired at each time step, we start to see the system reaching an equilibrium mode. The pool size stabilizes after an initial growth and the frustration goes down to zero. The higher the minimum amount of people hired at every time step, the larger the equilibrium pool size is and the quicker the initial growth. Figure 8 demonstrates the results of a recruitment policy in which at least 3 new people are hired at every time step.

The fact that an equilibrium pool size is actually reached may not seem intuitive. Hiring a constant positive amount of people should keep adding people to the pool. What happens is that after the pool grows to a certain size it becomes big enough that the number of people that retire or choose not to reenlist is equal to the minimum constant amount of people that are being hired. At this point equilibrium is reached.

The pool size reached at equilibrium and the stability of personnel levels depend on a variety of parameters. To illustrate, we present two sets of results in Figure 9. The plots in the left column show the pool composition (top) and FN frustration (bottom) when the hiring policy specifies that exactly three people are hired at each time step, regardless of how many may have left. Here it is interesting to see that sailor pool size reaches an equilibrium that is slightly lower than that in Figure 8, when at least 3 people were hired at each step. This slight change in policy results in an equilibrium of about 230 sailors (Figure 9), whereas in the example from Figure 8 the equilibrium is closer to 260 sailors. Hence the net effect is a 30-sailor reduction in pool size without introducing instability.




Figure 8: Pool composition for a minimum inflow of 3

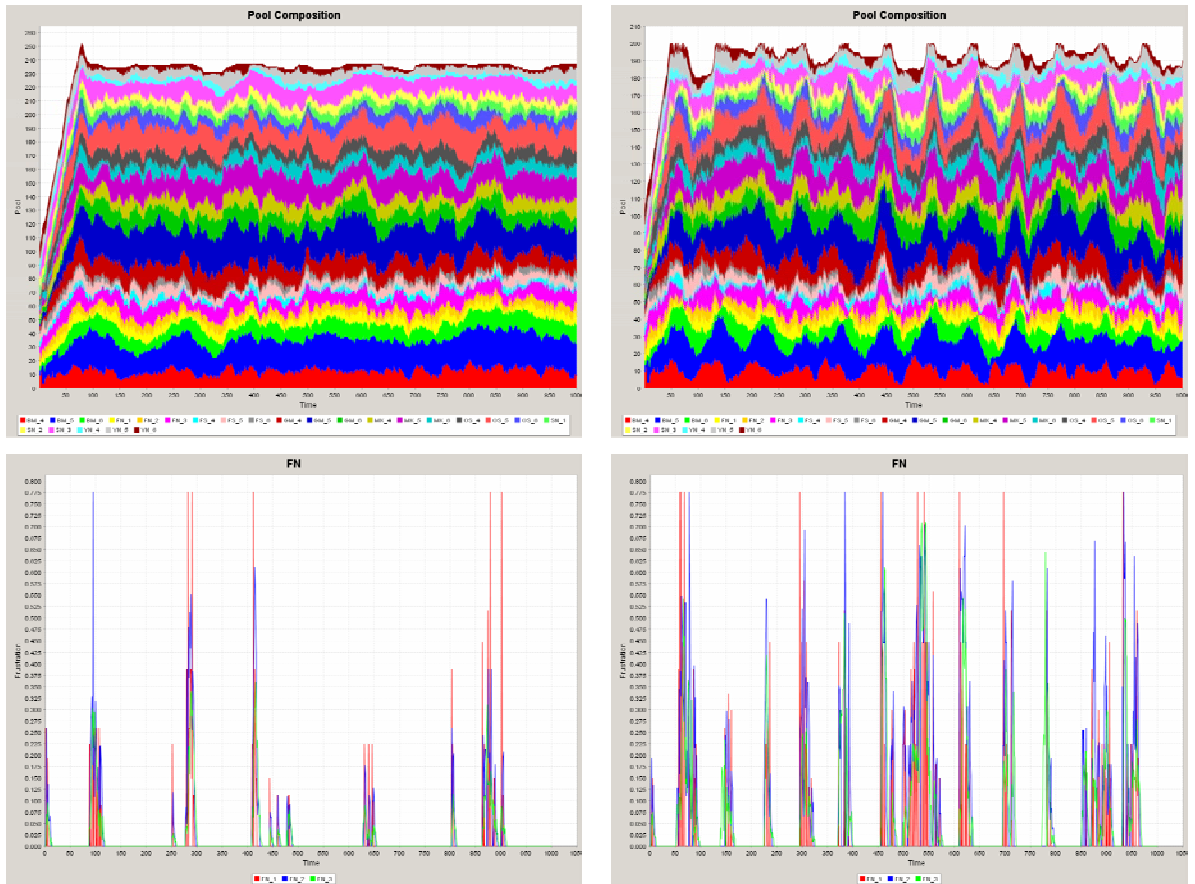


Figure 9: A slight change in recruitment policy has a profound effect on the resulting dynamics of the pool size (top) and frustration levels (bottom).

In the right column of Figure 9 we show the pool size (top) and FN frustration (bottom) when the hiring policy is the same as in the left column (hire exactly three sailors per step), but in addition the total pool size is capped at 200. The results are striking: even though the cap is reduced only by 30 sailors (about 13% lower than the equilibrium level achieved without the cap), the resulting behavior is much less stable, exhibiting significant oscillations and more frequent periods of high frustration.

The ability of our model to exhibit such rich behavior even in its simplified form highlights the importance of using quantitative tools to study the impact of M&P policies: even slight changes in policy may have unexpected consequences.

5.1 Summary of additional results

As part of our Phase I effort, we carried out extensive simulations and parameter sweeps to analyze the behavior of the model. In order to evaluate the impact of parameter changes, we have defined a series of metrics that capture the functionality of the Navy over the course of the entire simulation. Some of the metrics we have devised include:

- **Time at sea:** On what percentage of time steps was the ship able to sail? In other words, when were there enough sailors to fill all critical billets?
- **Ship frustration:** On those steps when the ship actually sailed, what percentage of the total billets were filled? This metric is averaged over all sailors on a given mission, to give a sense of how frustrating a particular mission was.
- **Staffing:** Over the course of the entire simulation, whenever the ship actually sailed, what was the ratio between manned billets versus the total billets posted for the mission?

Using these and other related metrics, we were able to do a full sweep of more than 5,000 different combinations of parameters. Through this analysis we found some interesting results, such as:

- When a cap is imposed on the pool size, the maximum period for retirement has a surprising impact on performance: the longer the retirement period, the worse the performance. The reason for this is that the longer retirement periods mean that the pool contains a higher number of more senior sailors. However, more billets need to be filled by sailors with lower ratings, which are not sufficient in numbers, thereby causing some missions to be skipped and other missions to have high frustration levels. Hence a lot of the newer sailors tend not to re-enlist, which causes further turnover at the lower ranks, and so on. In the meantime the more senior sailors are less frustrated, so they tend to stick around and cause the sailor pool to become more and more top-heavy.
- Surprisingly, a good way to cure this problem is to make the sailors more sensitive to frustration (that is, increase the amount of frustration felt as a result of being overworked). This prevents most sailors from sticking around until retirement, which restores some balance in the sailor pool.
- As with most complex systems, certain parameters exhibit strongly nonlinear behaviors. For instance, it is possible to counteract a tight cap on the maximum pool size by increasing the rate of recruitment. As we increased the recruitment

rate in response to a pool size of 200 and measured the staffing ratio (ratio of manned to total billets), we found that hiring rates of 1 and 2 sailors per time step were unable to overcome the cap, keeping the staffing ratio below 90%, while increasing the hiring rates to 3, 4, 5 or 6 sailors caused the staffing ratio to increase abruptly to nearly 100%.

6 Conclusions and future work

In all, we believe that the results of our work provide strong support for the importance of our quantitative modeling approach as a key tool for analyzing M&P systems and policy design. Even in its simplified form, the model captures a surprising degree of complexity, and exhibits *emergent* behaviors not unlike those seen in the real Navy. Agent-based models are ideally suited to investigate the dynamics of complex systems of this type.

One additional advantage of ABMs is that they tend to scale favorably. The model we have described here is a sound foundation for a more detailed model that captures many more aspects of the real Navy M&P system.

We were recently selected for a Phase II SBIR award to extend this work. Our plan is to extend the model to replicate more faithfully several details of the Navy's M&P process. First, we will increase the details of the behavior of sailors on a single ship. Next, we will replicate an entire battlegroup. Through these enhancements, we will be able to test the impact of various M&P policies. We will also enhance the user interface of the model to make it usable as a stand-alone tool to assist Navy staff in scenario testing and decision making. Finally, we will add a search module that enables users to identify M&P policies that satisfy specific criteria and requirements.

Many of the issues that the Navy faces are likely to be present in any organization with a substantial workforce. We have begun to discuss the applicability of our model within other branches of the military, and with large corporations. We believe that our approach represents a novel methodology that can complement existing approaches to personnel management.

7 References

- Axelrod, R. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration* (Princeton University Press, Princeton, NJ).
- Bonabeau, E. (2002). Agent-based modeling: methods and techniques for simulating human systems, *Proc. Nat. Acad. Sci. USA*, **99**, 7280-7287.
- Holder, A. (2005). Navy personnel planning and the optimal partition. *Operations Research*, **53**(1), 77-89.
- Moore, C. S., Hattiangadi, A. U., Sicilia, G. T. and Gasch, J. L. (2002). Inside the black box: Assessing the Navy's manpower requirements process. *CNA report CRM D0005206.A2*, available on-line through <http://www.cna.org/>.
- United States General Accounting Office (2003). Military Personnel: Navy Actions Needed to Optimize Ship Crew Size and Reduce Total Ownership Costs. *GAO Report 03-520*.