

Introduction

The problem of procuring and allocating electronic equipment for Naval ships is an extremely vital one, involving an annual budget of fifty to one hundred million dollars, and, more important, the safe and efficient operation of the entire Navy. Haphazard or inefficient allocation plans are therefore expensive and dangerous, and the importance of a mathematical systems approach is obvious. On the other hand, the problem is complex and involved; there are hundreds of different ship classes, many available equipments, and all kinds of complex relationships within and between the ships and the equipments. The problem has been fully described and discussed by J. W. Smith (see [1]). As a result of the efforts of Smith and others, including the authors, a systematic analytical technique for obtaining acceptable allocation plans has been developed; a brief description of the method is given in [1]. The technique may also be applicable to the procurement problem.

Computationally, the technique uses the standard linear programming model for the assignment problem; see [2]. The crux of the problem is therefore the evaluation of the coefficients in the linear program; i.e., finding values for the assignment of a given equipment to a given ship. We will here concentrate on this aspect of the problem. A brief description of the method has already been given in [3].

The technique used here has been described and discussed in [2] from the general, theoretical viewpoint, and for a greatly simplified version of the Naval Electronics Problem. The purpose of this paper is to present the detailed application of that technique to the Naval Electronics Problem in all its actual, real-life complexity. This paper is intended primarily, not for those whose main interest lies in the general, theoretical aspects of the allocation problem, but rather for those who are concerned with Naval Electronics, and for those who are interested in the working out of a technique such as this in a particular case, with all the messy detail with which particular cases tend to be burdened.

This paper discusses one aspect of the Naval Electronics Problem; the whole problem has been worked on by many people in many places and at many times. Among those who contributed significantly to the research described here are J. E. I. Heller, John P. Mayberry, Norman Shapiro, and George Suzuki. The most credit, though, must go to Jack W. Smith, whose ideas originated the project, and whose leadership kept it going throughout.

Description of the Problem

It is useful to view the list of possible assignments as being arranged in tabular form. The rows of the table represent electronic equipments, the columns represent “positions,” i.e., places where the equipments will be put. Thus the individual squares of the table represent assignments of electronic equipments to positions; we wish to fill in the table with the values of the various assignments represented by the squares.

We will make use of the rather commonplace idea that the parts and the positions may be grouped into equivalence classes in considering an allocation problem. The parts—in our problem the electronic equipments—will be grouped into models; two equipments are said to be of the same model if they are built to the same specifications and have identical performance characteristics. Thus if we are given a position and two equipments of the same model, then we have no preference as between assigning the one or the other equipment to the position. For convenience, we will define the void model, 0, to consist of no equipment at all.

Because of the complex interrelationships between positions, the grouping of the positions is not so quickly achieved. The basic notions are those of priority, state, and goodness. *Priority* refers to the intrinsic importance of the mission which a position is to fulfill. For example, a position on an aircraft carrier which is meant to contain an air search radar equipment has higher priority than a similar position on a destroyer. Similarly, a sonar position on an aircraft carrier may have lower priority than an air search radar position on the same ship. *State* refers to the model that is already installed in the position in question. *Goodness* refers, roughly speaking, to the suitability of the various models to the position in question. Thus a large model might be more suitable to an aircraft carrier than to a destroyer, or a certain model of a radar might be more suitable to a position requiring medium- or short-range equipment than to one requiring long-range equipment.

If we were to go according to only one of these factors, the solution to our problem would be simple enough. Thus if we were to decide to consider priority only, then we would fill the position of highest priority first, the next position on the priority list next, and so on. If we were to consider state only, we would first fill all the positions with nothing installed, then all the positions with the most inferior equipment installed, and so on. But matters aren't so simple. High-priority positions may have fairly good equipments already installed, there may be equipments that are especially suitable only for low priority positions, etc. All kinds of combinations of the three basic factors can and do occur, in varying

proportions, in real-life problems. It is precisely the interrelationships between the factors of priority, state, and goodness that complicate the solution to the naval electronics problem. Since this is the crux of the problem, we must ask the reader to bear with us while we examine more closely these concepts that we have introduced.

Priority is determined by a directive, called the MIP (Material Improvement Plan), that is issued by CNO (the Chief of Naval Operations). Items in the MIP are pairs consisting of a ship class (e.g., CVA aircraft carrier, DMS destroyer) and a requirement (e.g., sonar, medium-range surface search radar). The MIP ranks these items priority-wise. Two positions that correspond to the same item have essentially similar missions, and so are considered to have the same priority. The relative priority of other positions is given by the relative priority of the corresponding MIP items.

In practice, the MIP can be divided into a relatively small number of classes. Within each class, CNO has determined which item has priority; but this is always a much finer distinction than the large distinction made between the classes. The set of all positions corresponding to MIP items that belong to one of these large classes is called a *priority group*. So much for priority.

The *state* of a position is determined by which model of equipment is already installed in the position (before the contemplated allocation). Nothing more need be said about the state at the present time. The set of all positions of a given state is called a *state group*.

The concept of *goodness* is in reality associated with pairs consisting of a model and a position, rather than just with positions. Thus the corresponding concept for a position alone is that of a *goodness vector*, i.e., the set of all goodnesses obtained when we fix the position and let the model vary over all possible models. Two positions have the same goodness vector if each model is equally suitable to both positions. The set of all positions whose goodness vector is the same as that of a given position is called a *goodness group*.

Goodness groups are determined in the following way: Insofar as the position is concerned, suitability of a model to a position is determined by the kind of ship on which the position is located, and the requirement which it is to fulfill. When we say “kind of ship” we do not mean ship-class, which is a rather specific concept, but a much broader concept such as “space-critical ship” or “weight-critical ship” or their negatives. Whereas there may be fifty or more ship classes involved in a problem, there are usually no more than about four or five “kinds” of ships. “Requirement” is used here in the same way it was used before; examples of a requirement are sonar and medium-range surface search radar. A

goodness group may now be defined as the set of all positions fulfilling the same requirement and located on ships of the same “kind.”

Summing up, the positions have been grouped in three independent ways: into priority groups, into state groups, and into goodness groups. It is important to note that the classification of a position by one of these methods in general has no effect on its classification by another method.

We emphasize that the naval electronics problem need not be regarded as one huge allocation problem, but can be broken up into a number of smaller independent problems. There are usually many different models capable of filling a given electronic requirement. Conversely, a given model is often capable of filling more than one requirement. Thus, when considering the allocation problem for a given requirement, it is usually necessary to consider some models that are suitable for other requirements as well; these new requirements will then have to be considered at the same time as the original requirement. These in turn may be capable of being filled by other models, which necessitates considering still more requirements. The process eventually terminates; in fact, electronic equipment in use by the Navy is sufficiently specific so that usually at most three or four requirements can be strung together in the manner indicated above. These three or four requirements are, in fact, usually closely related anyway, so that from the administrative point of view, it is also convenient to consider them at once. The entire allocation problem for all electronic requirements can be conveniently separated into a number of smaller problems, each dealing with a “string” of requirements of the type described, together with all models appropriate for any of the requirements in the string. In the rest of this paper, attention will be focussed on a single problem dealing with only one string of requirements.

In Figure 1 we give a small example of what our structure might look like in a specific case; but Figure 1 is only meant to indicate some of the possible complex interrelationships, and should not be interpreted as giving the structure of a typical problem. Although a typical problem need not exhibit more complexity, it is usually a good deal larger. Following is a list giving the approximate order of magnitude for the number of groups and elements of each kind that might appear in a typical realistic allocation problem:

Number of available equipments:	500
Number of models:	8
Number of positions:	800
Number of MIP items:	40
Number of priority groups:	5
Number of state groups:	5
Number of goodness groups:	5

Our aim is to find the value—to be thought of intuitively as a “military worth”—of assigning a given equipment to a given position. We will assume that this value depends only on the model of the equipment, and on the priority (see “Priority Groups and the MIP”), state, and goodness groups of the position. Thus we are interested in determining all the values $v(M_i, P, G, M_j)$, where this expression stands for the value of assigning an equipment of Model M_i to a position in priority group P and goodness group G which has an equipment of model M_j already installed.

Brief Review of [2]

Let us review briefly the ideas introduced in [2], and the conclusions reached there. The central idea of [2] is that significant quantitative information about the values we are seeking can be obtained from qualitative decisions made by Naval Personnel. As an example, let us consider the MIP.

The MIP is provided by CNO for guidance in making up procurement and allocation programs. On an intuitive basis, we may conclude that “all other factors being equal,” the assignment of an equipment to a position higher on the MIP should have higher value than a similar assignment to a position lower on the MIP. In accordance with the terminology previously introduced, this can be formalized as follows:

If P_1 and P_2 are priority groups such that P_1 is higher on the MIP than P_2 , then for every model M and goodness group G , we have

$$v(M, P_1, G, 0) \geq v(M, P_2, G, 0) \quad (1)$$

with strict inequality holding for nonvoid M . The MIP can thus be regarded as a set of inequalities involving the values we are seeking. As such, it does already give us quantitative information regarding the values. On the other hand, it is clear that much more information is needed in order even to approximate numerical values.

Basically, the MIP consists of a set of decisions made by CNO in “small” or “token” allocation problems. The inequality (1) is tantamount to a decision by CNO that he would rather assign an equipment of M to a position in P_1 and G that has nothing installed than to a position in P_2 and G that has nothing installed (unless M is void). The “token” allocation decisions that constitute the MIP are of a very special kind—the positions involved in a given decision are always contained in the same goodness group and state group, and the equipments involved are of the same model. The inequalities that these decisions yield are not sufficient

to determine our values numerically, even approximately. It is natural to try to extend the scope of the “token” decisions. Additional decisions will yield additional inequalities, and by means of these we will be able numerically to approximate the assignment values.

Of course, it does not necessarily have to be CNO that reaches the decisions on these “token” problems. We will postulate the existence of a Naval Authority (called “the Board”) authorized to exercise judgment and arrive at decisions on allocation problems. In practice, the Board might consist of either one or several members.

Let us now turn to the conclusions reached in [2]. Starting out from a situation in which positions were distinguished only by their Priority and their State, we reached the conclusion that there are functions q and g such that

$$v(M_i, P, M_j) = q(P)(g(M_i) - g(M_j)), \quad (2)$$

where $v(M_i, P, M_j)$ is the value of assigning an equipment of model M_i to a position in priority group P which has an equipment of model M_j already installed; $q(P)$ is the *priority rating* of P , and $g(M)$ is the *goodness rating* of M . We also showed that the goodness and priority ratings could be fairly closely determined numerically using only the qualitative decisions of the Board in token allocation problems. Finally, the goodness rating g is such that if M_1, \dots, M_n ranges over all models, then

$$\max_i g(M_i) = 1. \quad (3)$$

Priority Groups and the MIP

The *value vector* associated with a given position b is the set of values obtained by assigning equipments of different models to b . A similar definition can be made for any function that associates a number with each pair consisting of a model and a position. A *value group* is the set of positions having a given value vector. A value group is the intersection of a priority group, a state group, and a goodness group. The priority group of a position depends only on the priority “class” of the MIP item to which the position corresponds, not on the MIP item itself. The formula for value introduced previously involves the tacit assumption that the distinction in priority between MIP items in the same priority group is small as compared to the distinction between priority groups. In the typical case presented, it cuts the number of values that must be ascertained from 1000 to 400; actually, as will be seen later, the saving is much greater than this.

Somehow, though, the distinction between members of the same priority class should be taken into account. This will be done as follows: Define a *subvalue* group to be the set of positions in a value group that correspond to a given MIP item. If we solve the linear program under the above assumption (i.e., that the value vector is constant throughout a value group rather than just throughout a subvalue group), we will in general obtain some “split” value groups,¹ i.e., value groups whose positions are assigned equipments of various models (including possibly voids), instead of equipments of just one model. Since for a given model, all positions of a given value group V have the same value, the question then arises, which positions should be assigned which models. This question is answered by reference to the distinctions between activities within the priority group of which V is a subset: We simply assign the better models to the positions that are higher in the MIP list. If any subvalue group is then still split by the allocation, then the choice of assignments within the subvalue group is arbitrary. The process is entirely in accord with the intuitive meaning of the MIP list: the more important activities get the better equipments. (If we could, we would apply the method throughout the whole problem: the only trouble is that to define “better equipments” uniquely, we must have the value of a model constant throughout the positions under consideration, and of course that is not so unless we restrict ourselves to a single value group.)

The Value Function

The difference between the mathematical structure of the problem presented in [2] and that presented here is that here the positions are distinguished by their goodness group as well as by their priority and state groups. However, we *can* apply the results obtained in [2] if we restrict our attention to one goodness group. Put differently, in using (2), we must take the goodness group G of a position into account, and allow the priority and goodness ratings to vary with goodness group as well as with priority group and model. We thus obtain

$$v(M_i, P, G, M_j) = q(P, G)(g(M_i, G) - g(M_j, G)) \quad (4)$$

and

$$\max_i g(M_i, G) = 1, \quad (5)$$

1. The number of split value groups never exceeds the number of models.

where M_i , P , and M_j are as before and G is the goodness group of the position to which the assignment is being made.

For reasons that will become apparent later, let us henceforth call the $g(M_i, G)$ *relative* goodnesses rather than just goodnesses. The relative goodness vector $g(G) = (g(M_1, G), \dots, g(M_n, G))$ can be determined by the method given in section 5 of [2]. The $q(P, G)$ may also be determined as in [2]. However, the evaluation of the $q(P, G)$, unlike that of the relative goodness ratings, can be considerably simplified.

To determine the $q(P, G)$ in accordance with the methods of section 6 of [2], we would ask the Board the following type of question: Given two positions, one in priority group P_i with an equipment of model M_i already installed, the other in priority group P_j with an equipment of model M_j already installed, both in the same goodness group G , and an available model M better² in G than either M_i or M_j , to which position would you assign M ?

In our situation, we are presented with the possibility of an alternative type of question: Given two positions, one in goodness group G_i with an equipment of model M_i already installed, the other in goodness group G_j with an equipment of model M_j already installed, both in the same priority group P , and an available model M better in G_i than M_i and better in G_j than M_j , to which position would you assign M ?

To distinguish these two types of questions, we will call the former a priority question, and the latter a *top anchor* question, for reasons which will soon be clear. According as to how a top anchor question is answered, we can decide which of

$$\frac{q(P, G_i)}{q(P, G_j)} \begin{cases} > \\ \approx \\ < \end{cases} \begin{cases} g(M, G_j) - g(M_j, G_j) \\ \\ g(M, G_i) - g(M_i, G_i) \end{cases} \quad (6)$$

holds. Let us fix P , G_i , and G_j . For each M , M_i , and M_j , the right side of (6) can be numerically evaluated (approximately) by the methods of section 5 of [2]. Thus by varying M , M_i , and M_j , we will obtain a number of numerical estimates for the (fixed) right side of (6), both from above and from below. Taken together, these estimates approximately determine the ratio $q(P, G_i)/q(P, G_j)$. The process is similar to that used for priority questions in section 6 of [2].

Let $p(P)$ be the maximum of the $q(P, G)$ over all goodness groups G . Define

$$t(P, G) = q(P, G)/p(P). \quad (7)$$

2. Formally, "having higher relative goodness."

If the maximum of the $q(P, G)$ is taken on for $G = G(P)$ then we deduce from (7) that

$$t(P, G(P)) = 1. \quad (8)$$

Clearly

$$q(P, G_i)/q(P, G_j) = t(P, G_i)/t(P, G_j). \quad (9)$$

We have already seen how the left side of (9) may be determined. Hence we can also determine the right side of (9). (8) gives us the value of one of the $t(P, G)$. Since we now know the ratios, we can determine the values of all the $t(P, G)$ from the single value given by (8).

Our determination of the values of the $t(P, G)$ was based entirely on the answers to top-anchor questions, and on their interpretation in the form (6). Since the right side of (6) does not involve P in any way, the only way in which the values of the $t(P, G)$ could depend on P is if the answers to the top anchor questions depended on P . As a matter of actual³ fact, though, the answers to the top-anchor questions are independent of P . We conclude that $t(P, G)$ is independent of P , so that we may write

$$t(P, G) = t(G). \quad (10)$$

Combining (7) with (10) we obtain

$$q(P, G) = p(P)t(G); \quad (11)$$

$p(P)$ is called the *priority rating* of P ; $t(G)$ is called the *top anchor* of G .

From (11) and (4) we deduce

$$v(M_i, P, G, M_j) = p(P)t(G)(g(M_i, G) - g(M_j, G)). \quad (12)$$

The saving effected by (11) is considerable. In our typical example, (11) allows us to determine five $t(G)$ and five $p(P)$ —a total of ten parameters—instead of about twenty-five $q(P, G)$ we might have had to determine in the absence of such a formula.

Intuitively, (11) and (12) have the following meaning: The $g(M, G)$ are supposed to be a measure of goodness, i.e., a measure of the suitability of a piece of equipment to its position. We have arbitrarily set the highest $g(M, G)$ in a given G to be equal to 1. If this were to be a measure of *absolute* goodness, such a measure would be manifestly unrealistic. The reader can convince himself of this by considering what would happen if a new model were to be developed suitable for only one of the goodness

3. I.e., experimental.

groups, say G . The new model would have the highest possible goodness in G , namely 1; the old models would have goodnesses in G correspondingly smaller than 1; whereas in goodness groups other than G , the old models would retain their former goodnesses. Obviously either *all* goodnesses should be depressed with the introduction of the new model, or else the old goodnesses should be left undisturbed and the new model be given a goodness higher than 1; both are contrary to the stipulation that the highest goodness in a given goodness group be equal to 1. Obviously, therefore, the $g(M, G)$ should be considered only as *relative* goodnesses, i.e., they should be considered as giving valid comparisons *within* a goodness group only. To obtain comparisons valid *between* the goodness groups, each of the relative goodness vectors should be multiplied by a scalar factor, which compares the state of development in one goodness group with that in the others. Since this factor anchors the top goodness in any goodness group, it is called the top anchor. The *absolute goodness* in G of a model M is given by the product $t(G)g(M, G)$. An *absolute improvement* is the difference between two absolute goodnesses; and as in [2], to obtain the value of an assignment, we must multiply the absolute improvement it accomplishes by the priority rating of the position which it fills. This is formula (12).

We remark that (12) is not only a “reasonable” formula, but is actually the *unique correct formula*, assuming that linear programming is at all applicable and that there *is* a formula for value in terms of the priority, relative goodness, and top anchor ratings. More precisely, we are making the following assumptions:

1. The problem may be set up as a linear program.
2. The value of assigning an equipment to a position depends only on the model of the equipment and the priority, state, and goodness groups⁴ of the position.
3. Absolute goodness is independent of Priority Group.

Precise mathematical statements and justification of all but one of these assumptions may be found in [2]. The additional assumption is

4. Although we have used the vague idea of “suitability” to give an intuitive meaning for goodness, the formal definition of goodness group involves only the well-defined (in any particular problem) notions of “kind” of ship and requirement. Therefore any consideration affecting assignment value that depends only on these two notions can be formally incorporated into the goodness concept. An example is the “Utilization Factor” mentioned in [1]. Even if it is assumed that this is relevant, it should not appear explicitly in our formulae, because it depends only on the goodness group, and is therefore incorporated in the goodness ratings. Other examples might be such factors as reparability and obsolescence. Like goodness, they depend only on the requirement and “kind” of ship with which the position is associated (and of course on the model of the equipment to be installed).

embodied in (10), and is justified by the experimental conclusion that the answers to the top-anchor questions do not depend on P .

Numerical Determination of Top Anchors and Priority Ratings

To start with, we wish to assure the reader that neither the priority questions nor the top anchor questions will in practice look as forbidding as does the general form given in the previous section. The questions with which the board is confronted will each involve two specific ship-classes, with specific models installed and a specific model available for installation. These models and ship classes will be chosen by the questioner to conform with the specified form, but the board need not and should not concern itself with priority groups, goodness groups, and so on; it should concentrate on answering the specific questions put to it.

Next, we should point out that there are some difficulties and pitfalls involved in the determination of the top anchors and priority ratings. As an example, let us consider the case of the top anchor. The process described in the previous section for obtaining the ratios between top anchors is in practice at best a rather rough one. Even if we were working with precise relative goodness ratings, it would be impossible to obtain precise top anchors. That is because within a certain area, usually something like 5%–10% of the top anchor in question, the decisions are difficult to make, come out as toss-ups, or even involve numerical contradictions of a minor sort. Thus we might get answers implying that a certain ratio is at the same time $>.75$ and $<.73$. This means that it is impossible to determine these ratios with greater precision. It is not only impossible, it is meaningless; to the extent that these parameters “exist” at all, they exist only in this loose, imprecise way. The imprecision in the top anchors is compounded by the imprecision in the relative goodness ratings, on which the determination of the top anchors is based. We may say, though, that given a set of relative goodness ratings, it is possible to determine top anchors with something between one and two significant figures.

One precaution we can take in order to minimize, to the extent possible, the damage caused by imprecise relative goodness ratings is to stay away from ratios on the right side of (6) with small denominators. That is because imprecisions in a small denominator are tremendously blown up in the quotient. In general, some components of the relative goodness vector are more “reliable” than others, as we shall see later; in asking top anchor questions, some attempt should be made to rely more on the more reliable ratios. To facilitate this, an “imprecision interval” should

be calculated for each ratio and placed on the list of ratios that we are using to determine the ratio of our two top anchors. The questioner should then try to stick to questions with as small an interval of imprecision as possible.

Theoretically, to obtain the top anchors, we pick an arbitrary priority group P and then determine all ratios of the form $q(P, G_i)/q(P, G_j)$, i.e., all ratios of the form $t(G_i)/t(G_j)$. From this listing it will be apparent which top anchor is the largest; we set this top anchor = 1, and deduce values for the other top anchors. Unfortunately, it turns out that this can usually not be done, since not all the priority groups intersect all the goodness groups. In fact, very few, if any, do. For the determination of any one ratio, say $t(G_i)/t(G_j)$, we must find a priority group P intersecting both G_i and G_j , so that we can select activities in G_i and G_j respectively that can be compared by means of top anchor questions. Sometimes there is no such priority group. We must then find a “bridging” chain of priority groups P_{i_1}, \dots, P_{i_m} and goodness groups $G_{i_0}, G_{i_1}, \dots, G_{i_m}$, where $i_0 = i$ and $i_m = j$, such that each P_{i_k} intersects both $G_{i_{k-1}}$ and G_{i_k} . Thus we will be able to determine each of the ratios $t(G_{i_{k-1}})/t(G_{i_k})$, and by multiplying them obtain the ratio we seek.

In practice, we choose one goodness group G and try to determine all the ratios $t(G_i)/t(G)$. To the extent possible, no bridging should be used, and all ratios determined independently. That is because the “snowballing” of imprecisions that occurs when bridging is used becomes rapidly disastrous. Care should be taken in the choice of G to minimize the error that will be incurred due to bridging; if possible, G should be chosen so that bridging is entirely avoided.

Everything we have said here about top anchors applies with equal force to priority ratings, which have an almost exactly symmetrical relation with top anchors.

Time Phasing

In discussing the mathematical structure of the problem earlier, we completely ignored its time aspect. In point of fact, equipment is becoming available more or less continuously, and ships are coming in for overhaul more or less continuously. This imposes several additional tasks on our technique for arriving at allocation plans. We will discuss these under the headings of primary, secondary, and tertiary time phasing.

Primary time phasing is simply the job of scheduling, i.e., fixing things so that equipment is available by the time the ship to which it is assigned comes into port for overhaul.

Secondary time phasing is the problem of taking into account idle warehouse time of available equipment waiting to be installed, and to extend possible minimizing such time. For instance, it may happen that assignment of equipment *a* to position 1 has a slightly higher value than assignment to position 2. As a result, equipment *a* is allocated to position 1, while a considerably inferior model is allocated to position 2. Now it happens that equipment *a* is available in January, the ship on which position 2 is located is coming in for overhaul in February, and the ship on which position 1 is located is not coming in until the following December. Since the value of assigning equipment *a* to position 1 is only slightly higher than assigning it to position 2, it seems clear that we shouldn't let the equipment lie around for almost a year just to wait for the ship on which position 1 is located. The example is admittedly rather extreme, as usually ships of all kinds are coming in fairly often. However, it does serve to illustrate the concept of secondary time phasing.

Tertiary time phasing is the long-range problem. It involves taking into account allocations that will be made at a time in the future which is not being taken care of in the present allocation plan. For example, let us say that we are making up a one-year allocation plan. We may have reason to believe that a new, superior model of equipment will become available for allocation next year. That means that we might want to assign more of the *currently* good models to the less important positions, so as to leave room on the more important positions for the superior model that may be available in the future.

One way to handle the primary and secondary time-phasing problems is first to use the method described in the foregoing section to obtain an un-time-phased allocation plan, then to do the best one can "by hand" to time-phase the resulting plan. In practice, such a scheme works fairly well, because both ships and models of all kinds are coming in fairly frequently. Also, the ships stay in the yards for a period ranging from six weeks to several months, thus giving the time-phaser fairly wide latitude. In some instances, it may be necessary to make comparatively small changes in the allocation plan in order to make it conform to the requirements of primary time-phasing. Actually, in the method now in use in the Bureau of Ships, an un-time-phased allocation plan is first made up "by hand," and it is only afterwards time-phased, also "by hand." The method described in the foregoing sections may then be said to replace the first, but not the second of these processes.

As of now, secondary time-phasing is considered of secondary importance. A blanket rule is made that all equipment becoming available during any given fiscal year will be assigned during that year, but within that restriction not much emphasis is placed on the problem. On the

average, things work out satisfactorily. Of course, the great difficulty in this field is the lack of criteria, the difficulty in comparing the relative importance of assigning an equipment at once and assigning it to an important position.

An alternative method has been developed for handling the primary and secondary time phasing problems—a systematic method that employs the principles of [2] and of the foregoing sections. However, this method does not enjoy the status of the un-time-phased method described in the foregoing sections in the sense that it has been developed only recently and has yet to be applied to and tested on real-life allocation problems. It will be described in detail in a subsequent paper.

Tertiary time-phasing does not really constitute a problem at all. Allocation plans are currently made as far into the future as it is possible to say, even approximately, what will be available and what will be needed. Current practice in the Bureau of Ships is to prepare an allocation plan five years into the future! (This allocation plan is revised every year.) Beyond that, availability of equipment and other factors affecting allocation can only be taken into account in a generalized way. These factors will therefore automatically be considered in the token allocation decisions arrived at by the Board. Thus in a token allocation decision, the Board might very well decide to assign an equipment to the otherwise less important position, because of the tertiary time-phasing kind of reasoning. It might even reverse some pairs in what would otherwise be the “expected” goodness order for a given category. In any case, to the extent that the Board feels that the “tertiary” kind of consideration is significant, its feeling will be reflected in its decision and hence in the values.

Multiple Allowances

A ship S is said to have a multiple allowance in a certain requirement⁵ if the number of positions on S that fill that requirement is greater than 1. Allowances for major electronic equipment on large ships sometimes run as high as 7. Since all the positions filling a given requirement on board a given ship are associated with the same MIP item, they all have the same priority. This may lead to the anomalous situation in which a large ship gets its seventh equipment of a certain kind before a slightly less important ship gets its first.

5. For the meaning of “requirement,” see page 211.

One way to deal with this problem is formally to consider each of the positions in a given multiple allowance as filling a different requirement. (This is actually intuitively justifiable; the first surface search radar on a ship obviously serves a different function from the second, which can only supplement the first.) Each of the positions would then occupy a different position on the MIP. In this ranking, the position with the best equipment already installed should rank highest, with the next best installed second highest, and so on.

Procurement

Up to now, we have been assuming that the number of equipments of each model available at any one time is given us. Actually, of course, in the long run it is only the money that is given us; the problem of how many equipments of each model to buy with the available money must also be solved by the Bureau of Ships. This is the *procurement* problem.

As might be expected, procurement and allocation are in fact intimately related. Procurement effectiveness can only be measured in terms of the allocation doctrine that will be applied following the procurement; conversely, inefficient procurement will greatly hamper any effort to obtain an allocation plan of high value. It would seem that the two problems are inseparable, and should be treated as a combined *procurement-allocation* problem, optimizing procurement and allocation simultaneously.

As a matter of fact, the unified procurement allocation problem can be set up as a linear program, using the same objective function as the one that the allocation problem uses, but replacing the restraints governing model availability by money restraints. See [3]. The problem of finding the coefficients of the objective function is of course the same in this problem as in the problem we have treated.

There has been a certain amount of opposition to the use of the linear programming technique for determining procurement. Basically, this opposition is probably due to the fact that one of the fundamental reasons for introducing the technique into the field of allocation planning is missing in the case of procurement. This is the combinatorial complexity of allocation plans. Procurement plans are combinatorially relatively simple; they only involve the determination of how many of each of a fairly limited number of models to buy. It may be felt, with some degree of justification, that the whole procurement problem is combinatorially sufficiently simple as almost to be considered a “token” problem in itself.

There is little point in then replacing one “token” problem by a set of different “token” problems. Of course, the difference between the procurement problem and a true “token” problem is that in the latter, we have a clear understanding of the precise implication of our decision insofar as the final allocation is concerned, whereas in the former we do not; nevertheless, it remains true that the procurement problem in itself does not “combinatorially overwhelm” the deciding officer, as does the allocation problem.

Another reason for differentiating between procurement and allocation is that these two events do not really occur simultaneously. Electronic equipment must be ordered an average of two years before it is delivered; thus procurement decisions must be made two years in advance of allocation decisions. A lot can happen in two years. It might not be wise to make procurement depend solely on our estimate of what optimal allocation will be in two years. Some flexibility should be built into a procurement plan that looks that far into the future, as insurance against a change in conditions. On the other hand, how do we measure flexibility? Do current methods arrive at procurement plans that are more flexible? At any rate, it can be seen that the extension of the technique into the field of procurement planning involves conceptual problems that are absent in the case of allocation planning and that have not yet been completely settled.

The procurement-allocation problem we have been considering is only that involving a single string of requirements. Congress makes a single appropriation for the entire navel electronics field, so even if we make the procurements *within* a string by “machine,” we will still have to decide on the distribution of money *between* the strings by “hand.” One way to get around this is to consider the entire Naval Electronics procurement-allocation problem at the same time, comparing the values as between different strings of requirements by the use of “token” procurement-allocation questions, e.g.: If you had \$100,000 available with which you could either buy two surface-search radar equipments of model so-and-so or three radio transmitters of model such-and-such, and if you had these and those positions to fill, which equipments would you buy and how would you allocate them?

We don’t know whether the use of such questions is feasible, or whether we can hope to obtain meaningful answers to them; experimental work has been performed only on the procurement-allocation problem confined to one string of requirements, not on the over-all Naval Electronics procurement-allocation problem. As a final remark, we may say that the unified procurement-allocation problem for all of Naval

Electronics would involve a very large matrix indeed, probably too large to go into the internal memory of any existing computer.⁶

A Brief Account of Experimental Work

At the time the authors began work on this project, they were presented with a specific allocation problem. Involved were some 8 models of radar repeaters to be distributed among some 800 positions divided into 80 MIP items, 6 state groups, and 4 goodness groups. The system of bunching MIP items had not yet been invented, so that each MIP item constituted a separate priority group.

Priority ratings were assigned purely arbitrarily (the lowest item on the MIP was assigned a priority rating of 60, and each succeeding item had a rating that exceeded the previous rating by 1; somewhere along in the middle a gap of about 4 or 5 was left) and the wrong formula for combining priority and goodness ratings was used. Even under such conditions, though, the allocation plan obtained was deemed superior to the one that had been obtained by hand! This would seem to indicate that a *systematic* method using the available information in even the most cursory fashion yields results that are superior to those obtained by the use of haphazard methods. Another advantage of the new method was also noted at that time: several “bloopers” or absurdities in the hand solution had been eliminated in the systematic linear programming solution, despite the use of arbitrary input and incorrect formulae. When the formula that later turned out to be the unique⁷ correct one was adopted, still more acceptable solutions were obtained and absurdities eliminated.

Another problem that was studied, this time from the point of view of procurement as well as allocation, involved 6 models of surface-search radar available for distribution among 700 positions divided into 52 MIP items, 4 state groups, and 5 goodness groups. Work on this problem led to the system of grouping MIP items into priority classes and to the systematic determination of priority ratings, top anchors, and relative goodnesses as described above.

6. A simplified computational technique for dealing with procurement problems involving several distinct allocation problems has recently come to our attention. By making use of Lagrangian multipliers, it allows us to consider each one of the distinct allocation problems separately, so that the whole large procurement problem need not be put on the computer at one time. An application of the technique is given in [4]. In our case, the “procurement problem” is the unified procurement problem for all of Naval Electronics, whereas the “distinct allocation problems” are those determined by the various strings.

7. See the discussion on page 219.

Numerical Methods for Obtaining Relative Goodness Ratings

An example of the calculation of a relative goodness vector involving three models is given in [3]. Often there are more models to work with, and this gives us a more precise determination of the relative goodness vector. It is also true that an improvement order containing toss-up decisions seems not only to yield a relative goodness rating polyhedron of correspondingly smaller dimensionality, but also to reduce the diameter of the polyhedron. Thus a toss-up decision becomes doubly valuable.

We make three additional remarks. First of all, it sometimes happens that where there are n models, more than n independent inequalities arise from the improvement order. In this case, we will obtain a polyhedron of possible relative goodness vectors that is not a simplex. To obtain the center of gravity of such a polyhedron, we break it up into simplexes. The center of gravity of the whole polyhedron is then the average of the centers of gravity of the component simplexes, weighted by their volumes. The center of gravity of a component simplex is found as illustrated in [3], while the volume is easily determined from the determinant of the vertices. In the instances that we have run across, the accompanying calculations have been remarkably quick and painless, largely because of the high incidence of zeros in the determinants involved. In more complicated cases, a high-speed computing machine could profitably be used.

Next, we note that for every polyhedron P of possible relative goodness vectors, there is a smallest⁸ rectangular parallelepiped with sides parallel to the axes that contains P . From the fact that this rectangular parallelepiped is not in general a cube, we may deduce that the precision with which our method determines the components of the relative goodness vector usually varies from component to component. This is why some components of the relative goodness vector are more "reliable" than others, as we noted earlier in this paper.

Our third remark is that occasionally the Board is only willing to give the beginning of an improvement order, but is unwilling to commit itself as to, say, its last third. We then obtain a "truncated" improvement order. This truncated improvement order leads to a polyhedron of possible goodness vectors, and thence to a center of gravity, in exactly the same way that the full improvement order does. However, such a polyhedron will in general be somewhat larger than that obtained with a full improvement order, so that the resulting relative goodness vectors will be less "reliable."

8. Under inclusion.

Concluding Remarks

1. Everything that was said in section 7 of [2] about the advantages and limitations of the type of method described there applies with equal force to this application. In particular, once it is established that the linear programming model is applicable, *any* set of questions may be used that determine the values sufficiently closely. Thus we may with a clear conscience pick out those questions that are most efficient from the computational viewpoint.
2. There are two kinds of consistency problems associated with the work described here. One is the self-consistency problem. Will the token decisions made by any one Board yield consistent numerical answers? As stated above and in section 7 of [2], the applicability of the mathematical model may stand or fall on this question. The available experimental data seems to point to self-consistency; the only inconsistencies encountered were of a very minor nature (see page 220). We concede, though, that this question has not yet received sufficient attention and should be investigated on a more systematic basis.

The second consistency problem is that of consistency as between the numerical implications of decisions made by different naval officers or different Boards. We have stated before, and we state now that this is *not* a requirement that our system need satisfy. An allocation plan and the token decisions leading to it actually constitute military command decisions. Military decisions are always made by the officers appointed for the purpose of making them; they are not less valid because other officers, in the same position, might have made somewhat different decisions. The purpose of the method described here is not to replace the system of military command by a mathematical system; it is to transform a command decision which because of its enormous complexity must needs be made on a largely haphazard basis into a set of less complicated command decisions, each of which can be decided intelligently on the basis of the military judgment and experience of the deciding officers.

In point of fact, Naval Personnel with widely differing backgrounds often exhibit surprising unanimity regarding their decisions. On most of the priority and top-anchor questions that were asked of several officers, there was complete agreement. Some borderline cases might occasionally be called a toss-up by one officer while they would be definitely decided by another one; much more rarely, one officer would decide a problem one way while another one would decide it the other way. The priority ratings and top anchors determined according to the decisions of different Naval Personnel generally differed by no more than about .1 or .15. This

unanimity should actually not be too surprising. The officers draw on their naval experience to answer questions that they consider trivial, but that go a long way toward determining the values we are after. It seems that the naval knowledge on which most officers can agree is sufficient to determine the values quite closely.

3. We might mention that certain of the special assumptions we have made should be checked before each time that the technique is applied to a new string of requirements. As an example, the assumption that the MIP items can be bunched into a small number of priority classes whose members have approximately equal priority ratings should be checked. This can easily be done by asking a complete set of priority questions (i.e., a set of priority questions sufficient to determine the priority rating) for each of a number of activities within a certain priority class. Although the Board itself decided on the grouping involved, it might turn out in some cases that such a grouping is artificial, and that the priority ratings are more adequately represented, for instance, if a few of the ratings are determined individually and some kind of curve fitting or piecewise linear method is used for the others.

References

1. J. W. Smith, "A Plan to Allocate and Procure Electronic Sets by the Use of Linear Programming Techniques and Analytical Methods of Assigning Values to Qualitative Factors," *Nav. Res. Log. Quart.* 3: 151-162, 1956.
2. R. J. Aumann and J. B. Kruskal, "The Coefficients in an Allocation Problem," *Nav. Res. Log. Quart.* 5: 111-123, 1958 [Chapter 11].
3. G. Suzuki, "Procurement and Allocation of Naval Electronic Equipment," *Nav. Res. Log. Quart.* 4: 1-7, 1957.
4. S. G. Allen and T. R. Broida, "Stock Control Policy, Aviation Supply Office," Stanford Research Institute Project No. IU-2109 Final Report, March 1958.

