2008-10

þÿ Optimizing the US Navy s Com

Logistics Force

Brown, Gerald G.

Monterey, California: Naval Postgraduate School.

Optimizing the US Navy's Combat Logistics Force

Gerald G. Brown, W. Matthew Carlyle

Operations Research Department, Naval Postgraduate School, Monterey, California 93943

Abstract: We study how changes to the composition and employment of the US Navy combat logistic force (CLF) influence our ability to supply our navy worldwide. The CLF consists of about 30 special transport ships that carry ship and aircraft fuel, ordnance, dry stores, and food, and deliver these to client combatant ships underway, making it possible for our naval forces to operate at sea for extended periods. We have modeled CLF operations to evaluate a number of transforming initiatives that simplify its operation while supporting an even larger number of client ships for a greater variety of missions. Our input is an employment schedule for navy battle groups of ships operating worldwide, extending over a planning horizon of 90–180 days. We show how we use optimization to advise how to sustain these ships. We have used this model to evaluate new CLF ship designs, advise what number of ships in a new ship class would be needed, test concepts for forward at-sea logistics bases in lieu of conventional ports, demonstrate the effects of changes to operating policy, and generally try to show whether and how the CLF can support planned naval operations. © 2008 Wiley Periodicals, Inc.* Naval Research Logistics 55: 800–810, 2008

Keywords: naval logistics; optimization; combat logistics force

My logisticians are a humorless lot. . . they know if my campaign fails, they are the first ones I will slay.

<div align="right">Alexander</div>

1. WHAT IS THE COMBAT LOGISTICS FORCE, AND WHY MODEL IT?

Our combat logistics force (CLF) is a fleet of transport ships that sustains US Navy combatant ships at sea worldwide. Each transport carries some combination of four basic commodities: ship fuel, aircraft fuel, ordnance, and dry stores and food. An underway replenishment (UNREP) rendezvous of a transport with a client ship (or simultaneously with more than one) demands superb seamanship to approach, rig transfer lines and hoses, convey commodities and, perhaps, personnel, unrig, and depart company. Transports also use helicopters for vertical replenishment (VERTREP) of client ships that may be some distance away. The special hardware and procedures for these operations have been developed and improved since the 1930s, and permit our navy today to operate continuously for extended periods at sea without returning to any port.

Correspondence to: W.M. Carlyle (mcarlyle@nps.edu)

The CLF is being transformed to a fleet with fewer different types of transports, and no more total transports, but it expects to have to serve more clients for a greater variety of deployments in the next decade. Conventional planning has relied on steady-state, average-rate-of-consumption models and rules-of-thumb to assess CLF ability to re-supply our fleet operations (for example, see Appendix A in CNA [8]). Details matter, and we want to determine whether or not, and how, the new CLF can actually support its anticipated missions.

We describe how we have modeled CLF operations, and what we have learned. We present a representative model that plans transport resupply operations with daily resolution, optimizing utilization over any deployment scenario of reasonable length (up to, say, 180 days). Our model takes into account changing daily consumption of each of four basic commodities for each client ship, navigational issues such as slow passages through canals, and the possibility of several client ships, or groups of ships, running low on the same commodity at the same time. Given a deployment scenario and a current configuration of the CLF fleet of transports, a solution to our model is a face-valid logistical plan for the CLF that minimizes shortfalls of any commodity for any customer, highlights unavoidable low-inventory events, and maximizes the utilization of transports by maximizing the total volume delivered over the scenario.

2. CLF CUSTOMERS: US NAVY COMBATANT SHIPS

Most US Navy deployments are groups of ships assembled with a particular mission. Some frequent examples are as follow.

- A carrier strike group consists of a nuclear-powered aircraft carrier (CVN), a guided-missile cruiser (CG), two guided-missile destroyers (DDG), and a fast combat replenishment ship. Accompanying attack submarines are completely autonomous.
- An expeditionary strike group transports a Marine expeditionary unit on an amphibious assault ship (LHA or LHD), with a dock landing ship (LSD), amphibious transport dock (LPD), a CG and two DDGs.
- A surface strike group consists of ships equipped with missiles and missile defense weapons, such as a CG and two DDGs.
- A Littoral Combat Squadron will employ a new class of small ships where larger ships cannot safely navigate, engaging in anti-surface warfare, mine counter measures, intelligence, surveillance and reconnaissance, homeland defense and maritime interdiction, and special operations forces support. For logistic planning purposes, we treat this new, small ship as a frigate (FFG).

To us, any such deployed group of ships in company is a CLF customer. We call such a customer a *battle group* (BG).

3. SCENARIOS: TRANSIT AND BATTLE PLANS

For our purposes, a scenario is set of exogenous, scripted deployment plans that tells us, day by day, where each battle group will be and what it will be doing. This daily fidelity is in keeping with navy practice, where each ship captain transmits a daily situation report including position and state. Our scenarios are filtered from larger, much more detailed contingency plans for a wide variety of missions.

A typical scenario consists of multiple BGs, and for each BG specifies last-minute in-port preparations and/or predeployment workup training in preparation for deployment, a high-speed transit to an area where we will show our military presence, a surge into combat operations to achieve a given objective, a sustainment phase to hold that objective, and perhaps a postcombat period where we stand guard and provide humanitarian assistance when diplomacy and other nonmilitary measures unfold.

Of course, we do not believe that any one of our planning scenarios will ever be followed verbatim. Rather, we use a wide variety of these to see how we will need to support such operations, region-by-region, mission-by-mission, worldwide.

4. DEFINING DEMAND: LOGISTICS PLANNING FACTORS

Standard deployment scenarios used by navy planners include all the client ships, where they will be day by day, and what they will be doing, but do not tell us what logistics support they will need. That's our problem to estimate.

Eccles [13] provides a classic document with lessons learned from World War II, and advises "all logistics planning is based on usage factors, which are average figures computed in many various ways." Today, we call such consumption estimates *logistics planning factors*. For purposes of estimating demand, we specify the employment state of each ship as, for example, in port, in transit, operating on station, or in combat. We aggregate demand into four categories: food and dry stores (STOR), ship fuel (DFM), aviation fuel (JP5), and ordnance (ORDN). The number of personnel aboard determines consumption of STOR, regardless of ship employment. DFM (distillate fuel, marine) consumption is fairly easy to estimate from ship engineering publications (e.g., see Brown et al. [6]) and employment state. Ship power consumption is stated in kilowatts, with a basic "hotel load" required to support the ship and her crew and systems, and propulsion plant demand as a function of speed. Some JP5 is consumed by helicopters, but the overwhelming volume is required by carrier aircraft. Although some ORDN is consumed in, for example, gunnery training, most weight consists of air-dropped munitions delivered in combat.

5. CLF SHIPS: FLEET COMPOSITION

The combat logistics force is being consolidated to just three ship types, with 30 total ships (see Fig. 1).

- The TAO187 (Henry J. Kaiser) class was introduced in 1986 as the first US Navy UNREP ship designed for operation by civil service mariners. With a crew of about 82 civilians and 21 navy personnel, it can carry about 180,000 barrels of fuel oil, and 271 tons of cargo lube oil, dry stores, and refrigerated containers, at about 20 knots.
- The T-AOE6 (Supply) Class was introduced in 1994, and with a crew of about 176 civilians and 59 navy personnel, it can carry 156,000 barrels of fuel oil, 1,800 tons of ordnance, 250 tons of dry stores, and 400 tons of refrigerated stores, at speeds exceeding 26 knots. This ship also features some self-defense armament.
- The T-AKE1 (Lewis and Clark) joined the fleet in 2006, and with a crew of about 123 civilians and 49 navy personnel, it can carry 18,000 barrels of fuel oil,
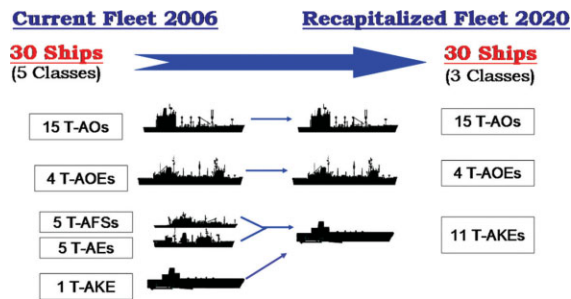
Figure 1. Combat logistics force ship classes. The CLF is shifting from five to just three classes of support ship, but expects to serve a fleet growing from 284 to 310 ships. Navy ship class names offer telltales of function: "T" means Military Sealift Command transport, commanded and crewed by civilian mariners accompanied by some uniformed navy crew, "A" auxillary, "O" fuel oil, "E" explosive ordnance, "F" refrigerated, and "K" general cargo. The respective crew complement and steaming speed for each T-AO is 103 and 20 knots, T-AOE 235 and 26 knots, and T-AKE 172 and 20 knots. The T-AOE is faster, has defensive armaments, and is favored as a station ship. The slower T-AO and T-AKE may pair up as station ships, but are more often used as shuttle ships. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

5,900 tons of ordnance, dry stores, and refrigerated stores at up to 20 knots. Customarily, each storage hold is designed to carry either ordnance or dry stores, but T-AKE storage holds can be converted between the two.

Other ship classes (e.g., T-AE, T-AFS, T-AOE(X)) are either scheduled to leave active service for the reserve fleet, or are still on the drawing boards.

Generally, each ship type is characterized by size, speed and flexibility. Cost increases with each of these attributes, especially speed [e.g., [16]].

CLF ships fill one of two roles: a *station ship* accompanies a BG group of customer ships, acting as a local storage facility for the BG, while a *shuttle ship* transits between BGs and replenishment ports. A station ship may not be fast enough to keep up with its BG customers at their top speed, but it needs to be fast enough to periodically rendezvous with and resupply these customers. The standard term for a shuttle ship UNREP of a station ship, or serving an entire BG through several UNREPS, is a *consolidation*, or CONSOL.

6. SEA ROUTES: NAVIGATING THE WORLD'S OCEANS

Because where we will need to navigate our CLF ships worldwide is an output of our planning, rather than an input, we need an "automated sea routes" model. We could find no such tool, so we built our own. A key feature of our sea routes model is that it takes a basic set of worldwide waypoints and feasible transits between them, and integrates each route followed by the BGs in a given scenario. This provides a navigable network for transports to follow in service of the BGs.

This sea routes network construction proceeds in four phases.

1. A a node is defined for each "port" (i.e., a forward logistics site or an at-sea, pre-positioned modular cargo delivery system ship) and for each of a number of at-sea waypoints frequently used by ships navigating worldwide (e.g., Gibraltar gate). A "fast arc" connects each adjacent pair of these nodes between which full-speed transit is feasible, and nodes are positioned such that each fast arc is navigable along a great circle route.

2. A "slow arc" connects any node pair with a fixed transit time (e.g., a canal or restricted passage). This
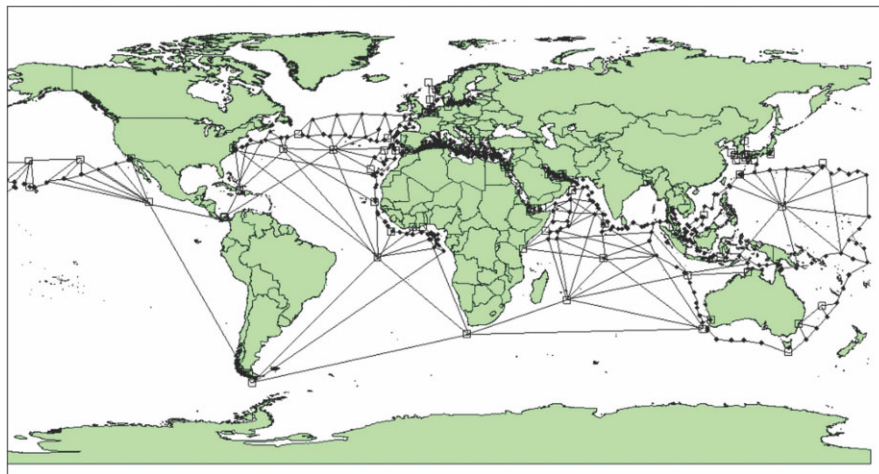

Figure 2. Worldwide sea route network. This particular network is for a case with 13 battle group customers operating over a 90-day planning horizon. Figure from Doyle [12]. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

base network is scenario independent, and has 102 nodes, 32 of which are ports, connected by 187 fast arcs and 11 slow arcs.

3. For any given scenario, we then trace the planned tracks of each BG, adding any daily position not collocated with one already in the network. Fast arcs connect these successive daily positions. A typical scenario might have 13 BG customers in transit for 90 days. After this step, we typically have about 600 nodes and 700 arcs.

4. Wherever a fast arc intersects some other fast arc, we induce a waypoint, and add the induced fast arcs. This brings us up to about 900 nodes and 64,000 arcs (see Fig. 2).

Each node may be assigned an administrative restriction that limits passage to only a subset of CLF ships. Each port node may admit only certain classes of CLF ship, or certain individual ships. These exclusionary devices can be used to keep CLF ships from leaving particular areas of operation, and to represent varying levels of access to sea space and forward supply ports.

Because of these restrictions, we must define a sea route network for each CLF shuttle ship that is physically or administratively distinct from others. There are about 16 of these shuttle ships for most scenarios (the CLF ships serve as station ships for the BGs). Each CLF ship class has its transit speed (vice its maximum speed), expressed in nautical miles per day (e.g., T-AO 384, T-AKE 308, T-AOE 600), and we use these to compute transit time from the great circle distance of each fast arc.

Finally, we apply the Floyd-Warshall algorithm [[1], see pp. 147-150] to generate the shortest point-to-point navigation times between any two locations in our sea route network, worldwide.

7. PLANNING CLF EMPLOYMENT SCHEDULES: AN INTEGER LINEAR PROGRAM

We use an integer linear program to plan optimal employment of CLF ships to minimize policy penalties accruing from any commodity shortage. An exact formulation follows.

7.1. Indices [~cardinality]

$s \in S$	Shuttle ship [~25]
$p \in P$	Port available to load shuttle ships [~35]
$bg \in BG$	Battle group [~13] (alias bx, by)
$d \in D$	Day [~181] (alias dx, dy, dh)
$c \in C$	Commodity (DFM, JP5, STOR, ORDN) [~4]
$\hat{c} \subseteq C$	Dry commodity subject to load fraction restrictions (STOR, ORDN) (alias $\hat{c}$)

7.2. Provided Data [units]

$spdSHUTTLE_s$	Speed of shuttle ship s [nm/day]
$inptTAT$	Time to reload shuttle ship in port [days]
$portok4s_{s,p}$	Binary indicator that shuttle ship s can reload at port p [binary]
$legdays_{s,bg,d,p}$	Shuttle ship s transit time at speed $spdSHUTTLE_s$ from bg position on day d to port p following given sea routes and/or BG tracks [days]
$useBG_{bg,d,c}$	Consumption by bg during day d of commodity c [c-units]
$mxload_{bg,c}$	Maximum capacity of bg to carry commodity c [c-units]
$safety_c$	Minimum desired fraction of $mxload_{bg,c}$ to be held at all times [fraction]
$extremis_c$	Extreme minimum desired fraction of $mxload_{bg,c}$ to be held at all times, $extremis_c \leq safety_c$ [fraction].
$hitOK_{bg,d}$	Logical indicator if bg can CONSOL on day d [binary]
$capacity_{s,c}$	Shuttle ship s capacity to deliver commodity c [c-units]
$mnfrac_{\hat{c}}, mxfrac_{\hat{c}}$	Minimum, maximum fraction of T-AKE dry capacity that must be loaded with dry commodity $\hat{c}$ [fraction]
$safety\_penalty_c$	Penalty per deficit unit of desired storage below safety-stock held by any BG [penalty per c-unit]
$extremis\_factor$	Multiplier (>1, e.g. 10) for penalty per deficit unit of desired storage below extremis held by any BG [dimensionless]
$negative\_factor$	Multiplier ($>extremis\_factor$, e.g. 1000) for penalty per deficit unit of desired storage below zero held by any BG [dimensionless]

7.3. Derived Data

$mxconsol_{s,bg,c}$	Maximum delivery shuttle ship s can make to bg on any day of commodity c [c-units]. This is defined as: $\min\{mxload_{bg,c}, capacity_{s,c}\}$.
$mxconsol2_{s,bg,bx,c}$	Maximum delivery shuttle ship s can make in one voyage to CONSOL bg and bx with commodity c [c-units]. This is defined as: $\min\{mxload_{bg,c} + mxload_{bx,c}, capacity_{s,c}\}$.

In addition, for T-AKE shuttle ships and dry commodities $\hat{c}$ sharing dry storage, and subject to limits on the minimum and maximum fractions of dry capacity that must be carried in every T-AKE load, this is restricted to:

$$\min\left\{ mxload_{bg,c}, \min[mxfrac_{\hat{c}}, \right.$$
$$\left. 1 - \sum_{\tilde{c}\neq\hat{c}} mnfrac_{\tilde{c}}] * capacity_{s,\hat{c}} \right\}$$

or, the maximum permitted T-AKE load of dry commodity $\hat{c}$, or the amount of commodity $\hat{c}$ that can be loaded after the minimum loads of other dry commodities $\tilde{c} \neq \hat{c}$ sharing dry storage are loaded.

$cycledays_{s,bg,d,bx,dx}$ If shuttle ship s departs bg on day d to reload at some port p, the minimum number of days before a rendezvous with BG bx on day dx is

$$\min\left\{ \infty, \min_{\substack{p|\\portok4s_{s,p}}}\left[\min_{\substack{dx \geq legdays_{s,bg,d,p}\\+inptTAT\\+legdays_{s,bx,dx,p}}} \begin{array}{c}(legdays_{s,bg,d,p} + inptTAT\\ +legdays_{s,bx,dx,p})\end{array}\right]\right\}$$

Note that this admits a cycle with slack time (or, "shuttle waiting time") $dx - d - cycledays_{s,bg,d,bx,dx} \geq 0$, and that because of the relative motion of a shuttle ship and a BG over navigable sea routes, and their daily proximity to ports and to each other, there will be cases in which planning for a shuttle to wait for this amount of time is better than restricting plans to have no such slack.

$directdays_{s,bg,d,bx,dx}$ The number of steaming days for shuttle s to transit from the position of bg on day d directly to the position of bx on subsequent day dx (i.e., without reloading in any port) (policy limits may govern the minimum or maximum days allowed between these planned events).

7.4. Decision Variables

$HIT_{s,bg,d}$ Binary indicator of shuttle s CONSOL visit to bg on day d (depends on $hitOK_{bg,d}$)

$HIT2_{s,bg,d,bx,dx}$ Binary indicator of shuttle s CONSOL visit to bg on day d, followed by a second CONSOL visit to bx on day dx before returning to port (depends on $hitOK_{bg,d}$)

$CONSOL_{s,bg,d,c}$ Amount of shuttle s delivery to bg on day d of commodity c [c-units]

$CONSOL12_{s,bg,d,bx,dx,c}$, $CONSOL22_{s,bg,d,bx,dx,c}$ Amount of shuttle s deliveries of commodity c to respectively, bg on day d ("12") and bx on subsequent day dx ("22") (i.e., without reloading at some intermediate port). [c-units]

$SHORTAGE_{bg,d,c}$ amount of inventory deficiency of c for bg, at end of day d [c-units]

$EXTREMIS_{bg,d,c}$ amount of extreme deficiency of c for bg, at end of day d [c-units]

$NEGINV_{bg,d,c}$ magnitude of negative inventory of c for bg at end of day d, has this [c-units]

7.5. Formulation

$$\text{s.t.} \quad \sum_{s,dh\leq d} CONSOL_{s,bg,dh,c}$$
$$+ \sum_{s,dh\leq d,bx,dx} CONSOL12_{s,bg,dh,bx,dx,c}$$
$$+ \sum_{s,dh\leq d,bx,dx} CONSOL22_{s,bx,dx,bg,dh,c}$$
$$\leq \sum_{dh\leq d} useBG_{bg,dh,c} \quad \forall bg,d,c \quad (1)$$

$$\sum_{s,dh\leq d} CONSOL_{s,bg,dh,c} + \sum_{s,dh\leq d,bx,dx} CONSOL12_{s,bg,dh,bx,dx,c}$$
$$+ \sum_{s,dh\leq d,bx,dx} CONSOL22_{s,bx,dx,bg,dh,c} + SHORTAGE_{bg,d,c}$$
$$+ EXTREMIS_{bg,d,c} + NEGINV_{bg,d,c}$$
$$\leq \sum_{dh\leq d} useBG_{bg,dh,c} - (1 - safety_c)mxload_{bg,c} \quad \forall bg,d,c$$
$$(2)$$

$$CONSOL_{s,bg,d,c} \leq mxconsol_{s,bg,c}HIT_{s,bg,d} \quad \forall s,bg,d,c$$
$$(3)$$

$$CONSOL12_{s,bg,d,bx,dx,c} + CONSOL22_{s,bg,d,bx,dx,c}$$
$$\leq mxconsol2_{s,bg,bx,c}HIT2_{s,bg,d,bx,dx} \quad \forall s,bg,d,bx,dx,c$$
$$(4)$$

$$HIT_{s,bg,d} + \sum_{\substack{by,dy|\\d-dy\geq directdays_{s,by,dy,bg,d}}} HIT2_{s,by,dy,bg,d}$$
$$+ HIT_{s,bx,dx} + \sum_{\substack{by,dy|\\dy-dx\geq directdays_{s,bx,dx,by,dy}}} HIT2_{s,bx,dx,by,dy} \leq 1$$
$$\forall s,bg,d,bx,dx|dx - d < cycledays_{s,bg,d,bx,dx} \quad (5)$$

$$\sum_{bg} HIT_{s,bg,d} + \sum_{\substack{bx,by \\ dx \le d \le dy}} HIT2_{s,bx,dx,by,dy} \le 1 \quad \forall s,d \quad (6)$$

$$HIT_{s,bg,d} \in \{0,1\} \quad \forall s,bg,d \quad (7)$$

$$HIT2_{s,bg,d,bx,dx} \in \{0,1\} \quad \forall s,bg,d,bx,dx \quad (8)$$

$$0 \le CONSOL_{s,bg,d,c} \le mxconsol_{s,bg,c} \quad \forall s,bg,d,c \quad (9)$$

$$0 \le CONSOL12_{s,bg,d,bx,dx,c} \le mxconsol_{s,bg,c}$$
$$\forall s,bg,d,bx,dx,c \quad (10)$$

$$0 \le CONSOL22_{s,bg,d,bx,dx,c} \le mxconsol_{s,bx,c}$$
$$\forall s,bg,d,bx,dx,c \quad (11)$$

$$0 \le SHORTAGE_{bg,d,c}$$
$$\le (safety_c - extremis_c) * mxload_{bg,c} \quad \forall bg,d,c \quad (12)$$

$$0 \le EXTREMIS_{bg,d,c}$$
$$\le extremis_c * mxload_{bg,c} \quad \forall bg,d,c \quad (13)$$

$$0 \le NEGINV_{bg,d,c} \quad \forall bg,d,c \quad (14)$$

$$\underset{\substack{HIT,HIT2, \\ CONSOL,CONSOL12,CONSOL22, \\ SHORTAGE,EXTREMIS,NEGINV}}{\text{MIN}} \sum_{s,bg,d,c} -0.1 safety\_penalty_c$$

$$* CONSOL_{s,bg,d,c} + \sum_{\substack{s,bg,d,bx,dx| \\ dx-d \ge directdays_{s,bg,d,bx,dx}}} -0.1 safety\_penalty_c$$

$$* (CONSOL12_{s,bg,d,bx,dx,c} + CONSOL22_{s,bg,d,bx,dx,c})$$

$$+ \sum_{bg,d,c} safety\_penalty_c * SHORTAGE_{bg,d,c}$$

$$+ \sum_{bg,d,c} extremis\_factor * safety\_penalty_c$$

$$* EXTREMIS_{bg,d,c}$$

$$+ \sum_{bg,d,c} negative\_factor * safety\_penalty_c * NEGINV_{bg,d,c}$$

$$(15)$$

8. DISCUSSION

Inequalities (1) limit day-by-day cumulative CONSOL volumes of each commodity to the cumulative usage of each BG through the end of that day. We assume that on the first day, each BG is full to capacity with every commodity. Thereafter, daily use is deducted, and replenishments are accounted from shuttle CONSOLs. Elastic inequalities (2) reckon cumulative inventory state of each commodity at the end of each planning day, and compare this to the cumulative usage less desired safety-stock level at the end of that day, representing any shortage, extreme shortage, or negative inventory required to reconcile this state. Each inequality (3) limits the CONSOL volume transferred from a shuttle ship, to a BG, on some given day, to be zero unless a replenishment event takes place. Similarly, each inequality (4) controls the

successive CONSOL volumes transferred from a shuttle to a BG on some given day, followed by a second CONSOL on a BG on some given later day, to be zero unless a replenishment event takes place for that shuttle on BG on that day, followed by the second BG on the second day. Constraints (5) restrict successive shuttle rendezvous with battle groups so that each such visit is followed by sufficient time to cycle to a port for re-supply. Each constraint (6) permits a shuttle to engage in at most one activity on a given day. Variable domains are stated by constraints (7)–(14). The objective (15) expresses a penalty with a component for any shortage below safety-stock, and extreme shortage below minimum stock, and any negative inventory, less rewards for commodity volume delivered; the rewards here are ten percent of the safety stock shortage penalties, and attract maximal delivered volumes, rather than merely deliveries to avoid shortages.

Our model can schedule a single shuttle ship sortie from port to make two separate CONSOL visits, perhaps to two different battle groups. It turns out that allowing multiple CONSOLs does not help in our large scenarios, and it increases the solve times significantly, so we do not use this feature in those scenarios. For a small 30-day scenario in the Arabian Gulf, with three battle groups being supported by one T-AKE, we find that allowing multiple CONSOLs per sortie increases the overall minimum inventory levels seen, but that the difference is modest. We include this feature to allow the model to represent a situation that could occur in the real world, but we conjecture that it is useful in a few circumstances where many battle groups are close to each other for extended periods, and have low to moderate consumption of the relevant commodities. The two-CONSOL sorties can be toggled on and off, therefore, for completeness, we display in our formulation the fully general model we have implemented.

9. PRIOR MODELS OPTIMIZING NAVAL LOGISTICS

Dantzig and Fulkerson [9] present what we would now call a pure network formulation to minimize the number of identical tanker ships required to deliver a fixed slate (schedule) of naval fuel shipments. Each shipment is characterized by load day and location, unload day and location, and times for loaded outbound, and empty return voyages to the next load location.

Brown, et al. [5] plan crude oil tanker voyages with a set partition model selecting for each vessel the best employment schedule among all those feasible over a 60–90 day planning horizon. The sorties carry full loads of crude oil on various-sized tankers from Middle East terminals to refiners in Europe and the Americas.

Bausch, et al. [2] schedule lighter and barge operations conveying multiple distilled and refined petroleum products among coastal port facilities. Tides and complex operating and safety rules govern hours of operation, each vessel has distinct cargo capacities and operating rules, and vessels must adjust their speeds to arrive at destinations when they can dock and transfer cargo.

10.   RESULTS: CLF INSIGHTS

The integer linear optimization model (or, if you prefer, mixed-integer program, or MIP) presented here is the latest in a series that has evolved to answer a number of questions. Following are abstracts of our results.

Borden [3] presents seminal work on the history of CLF, recounts prior analyses and, anticipating decisions to procure the T-AKE, shapes the fundamental questions:

- How many T-AKEs will be enough?
- What is the optimal T-AKE load of ordnance and dry stores for its convertible-storage holds?
- How should we operate the T-AKEs?

He develops from scratch three single- and three multiple-battle group 90-day scenarios variously directed at the Baltics, Arabian Gulf, Philippines, Panama, and Korea, and ranging in scale from a minor contingency to a major theater war. He develops logistics planning factors, a global sea route model, and a MIP to plan CLF CONSOLs.

Borden answers that:

- We need 11 T-AKEs, plus one to allow for maintenance availability and shipyard periods;
- Rather than a single, static, optimal T-AKE load, we simply need to load what is needed each time we load it; and
- T-AKE can not only shuttle to CONSOL faster BGs, but can also join with a T-AO to serve as a station ship pair for a BG, even though this pair is not as fast as its BG customer.

Borden offers a serendipitous insight, discovered by the optimization: forward pre-positioning of a T-AKE to make a timely first CONSOL of an emergent deploying battle group as it speeds by still permits the slower T-AKE to head to a forward logistics site, reload, and follow the faster battle group in time to join it in the destination area of operation. Once there, T-AKE is fast enough to shuttle to and from CONSOL rendezvous with the faster battle group. With multiple battle group deployers, and more than one T-AKE, even more interesting chess moves arise, but all with the same essence: anticipate the position of the BGs, top them off as they pass, then follow them to their area of operation and serve them there.

Borden also evaluates the effects of improving port loading time, decreasing the distance to forward logistics sites and/or at-sea station ships (resupply ships prepositioned in advantaged locations to reload shuttle ships), and either increasing shuttle ship speed or slowing planned BG speed.

We initially distinguished between a station ship and its companion BG ships, with shuttle ship consolidating the station ship, and the station ship, in turn, providing UNREPs to its companions in the BG. This turns out to be more detail than we need, so we now just consider the station ship as organic with its BG, and plan CONSOLs of the BG.

Subsequently, the navy announced plans to purchase 12 T-AKEs. (This is one more than shown in Figure 1 for the 2020 force plan.)

Givens [15] evaluates a proposed new T-AOE(X) ship, each of which would replace both a T-AKE and a T-AO with a faster (and much more expensive) station ship.

Givens refines logistics planning factors by ship type and employment state, including details of UNREP approach, rig and unrig times, and transfer rates. Givens highlights the influence of the minimum inventory levels for combatants that we use to trigger extraordinary efforts to CONSOL. Givens also introduces restrictions on shuttle access to ports, to preclude consideration of silly trans-global transits and limit the sheer size of our planning problems.

Givens concludes that we can support his scenarios without T-AOE(X), but that the cost is an increase in "off station time," when station ships must break company with their customer ships to resupply.

Cardillo [7] examines a 90-day scenario deploying every available naval combatant. He investigates how CLF can best support forces concentrating on one major theater conflict while holding back a secondary contingency, and then turning to deal with the secondary action. His scenarios orchestrate CLF ships moving to one theater, and then transitioning to another. He also anticipates activating and deploying navy reserve-fleet tankers to supplement the active CLF fleet.

Even in the basic scenario, where all ports in both theaters are open to our CLF transports, several BGs get uncomfortably close to running out of fuel in an optimal solution. One BG arrives in the second theater and gets to within one day of running out of DFM, and another gets within three days of a DFM runout. This is a *completely* unacceptable situation for a BG commander, who prefers to maintain at least ten to twelve days of fuel in reserve. The result points to an urgent need for more logistical support in a two-theater scenario, especially during the swing between theaters.

DeGrange [10] models CLF operations with a forward sea logistics base—an at-sea logistics facility made up of ships that can support direct operations inland without amphibious assault or permissive access to nearly port facilities. He also evaluates navy conversion to a single distillate fuel,

in lieu of DFM for ships and JP5 for aircraft. His scenario includes a period of humanitarian assistance following military action, and this puts the T-AKE convertible dry storage compartments to good use.

His results show that, for example, increasing safety-stock requirements from 50 to 60% of full levels requires a significant increase (from seven to nine) in the number of CLF transports required to support sea-basing operations. Converting to a single fuel can reduce the number of T-AOs required from four to three, and also reduces the total number of T-AO CONSOLs.

Doyle [12] distills all these prior contributions, further refines logistics planning factors, and examines the nuances of CLF ship ownership and control. Each CLF ship is customarily assigned to operate in one of the navy's five numbered fleet areas of responsibility worldwide (e.g., Third Fleet extends from the western shores of the Americas to the international dateline). While this may make sense for administrative purposes, during multiple, major worldwide deployments CLF ships may need to leave their home areas. Fleet commanders may be a bit possessive and resist giving up "their" CLF ships. Doyle illustrates quite convincingly how damaging this can be to CLF effectiveness.

Doyle also evaluates an operational model improvement by which a shuttle ship can CONSOL two BGs, on different days of course, in one sortie from a port. Surprisingly, this generalization seldom makes any improvement. When you consider the off-station time required for such less-than-maximal deliveries, they are not attractive. Of course, this is planning. In real-life operations, we will do whatever we need to do to sustain combatants.

11.   IMPLEMENTATION AND COMPUTATION: LEARNING FROM EXPERIENCE

We use the General Algebraic Modeling Language and the CPLEX MIP solver (respective versions 22.3 and 10 at this writing) [14] on WINTEL personal computers. A representative GAMS script totals about 7000 source lines, including an imported scenario data script.

Most of this GAMS code filters data and diagnoses curiosities in the scenario. E.g., a minor data entry error of a single daily BG longitude—one too small to see right away in a graphical interface sea route map—might make that BG move at 100 knots for a couple of days, and this will be tough for CLF to keep up with. Similarly, a unit error for a logistics planning factor, say BBL instead of KBBL for DFM consumption, can wreak havoc if not detected and diagnosed immediately. Even (perhaps especially) with sophisticated data bases, errors happen. We have built in a lot of "tribal wisdom" to save ourselves time, frustration, and (maybe) embarrassment.

We have had to remove the Floyd-Warshall shortest-path computations from GAMS to an external C procedure, where it runs many orders of magnitude faster than in GAMS. This external procedure reads an input file from GAMS, computes all shortest paths, and writes an output file for GAMS to use. We compute a check-sum of all the data that influences the shortest paths, and record this signature with the file. Every time we run our GAMS script, it repeats the (fast) computation of this signature checksum and verifies that the shortest-path file it imports is still valid, or complains that something has changed influencing sea route shortest paths, writes a file for the external procedure to use as input, and insists the file be refreshed by the external procedure.

Some scenarios are large and long, and the integer MIPs that result can have several million constraints and several million binary variables. We have trouble solving these models in reasonable time. The old planning adage applies here that *you have to be able to come up with an answer while your client still remembers the question.*

Accordingly, we have introduced some expedient methods to accelerate solution times.

For instance, we sometimes reduce the geographic range and number of ports available to each shuttle ship, and reduce the set of BGs served by each shuttle ship. This works especially well when we have scenarios with two or more geographically remote, and thus essentially physically disjoint missions to deal with at once—our common sense helps the optimization ignore lots of desperate, bad options.

However, for some large-scale, long-term, and thus worldwide scenarios, restricting CLF maneuvers to a subset of customers is not realistic.

Solving a monolithic planning problem, we specify a maximum solve time and a maximum acceptable interval of uncertainty–the relative or absolute difference between the value of the best integer solution found and an upper bound on this value. If we do not satisfy this tolerance in the allotted solve time, we resort to a time-myopic problem cascade; by coincidence, this derives from earlier ship-scheduling work by Brown, et al. [[5]; see p. 341], though the idea has been used since in many other planning contexts.

Our cascade is defined with three terms: days in planning horizon, days in planning window, and days advanced per solve. For the simple example in Fig. 3, these are, respectively, 30, 10, and 5 days. As we advance the planning window through the number of days comprising the planning horizon, we divide the optimization into three components: history, current planning window, and future. In the initial step, the first MIP solves only the current planning window, returning a solution when the integer tolerance is satisfied. The second MIP then advances to the next planning window, fixing the last-determined values for those variables in our history, and relaxes all historic constraints save those that still have influence on the new current planning window. We pay no

Figure 3. Example of a time-myopic problem cascade. We solve a 30-day planning horizon by considering a sliding planning window of 10 days, and advancing this planning window 5 days at a time. As we advance, we fix our history at that last seen in a planning window. For example, in the third integer program, we have 10 days of decisions that are fixed (days 1–10), ten days of decisions available to the model (days 11–20), and ten days that do not yet appear in the model (days 21–30). Figure from Doyle [12].

attention to the variables and data for future time periods until they comprise part of the current planning window. This iteration repeats until the final day of the planning horizon is included in the current planning horizon.

We add another key feature to our cascade: When we cannot solve a current planning window in a specified amount of time, our solver memorizes the last, "goal day" of this solve window, then halves the window length and the number of the days to advance to solve this new, and now smaller, planning window. Despite the reduced planning window, the MIP advances from where we last started and attempts to once again solve the current cascade window. If necessary, this halving continues recursively until a planning window of only two days, and an advance of just one day might fail, in which hypothetical case we would abnormally terminate. When we finally solve a planning window that ends on or after the goal day, we restore the days in the current planning window and the days advanced per solve to their original values until we either reach the end of the planning horizon, or need to repeat the recursive steps when we cannot solve a subsequent current planning window.

In practice, with an initial planning window spanning the full planning horizon, if we allow reasonable maximum solve time and interval of uncertainty, the automatic, emergent planning window reduction rarely activates. But, when it does, we get an alarm, and still reliably get a planning solution. This is a robust solution scheme. And, an alarm is a valuable telltale of model trouble.

When the planning window is a subset of the planning horizon, the variables we fix during the early periods have an impact on the later decision variables as the history period grows. For example, a T-AO replenishment on day 3 of a 30-day scenario may offer the optimal solution for this current planning window. However, establishing this same T-AO replenishment on day 3 of the 181-day scenario could result in a sub-optimal solution for the complete scenario, because the day 3 replenishment does not "anticipate" necessary later requirements in the complete planning horizon. We might see

this same effect in the initial prepositioning of CLF shuttles. We allow the model to locate shuttles based on the initial planning window regardless of later requirements in the planning horizon.

Admittedly, when our cascade planning window is (or becomes) shorter than the planning horizon, this returns a restricted solution, and we forfeit the certificate of solution quality we would otherwise attain.

Nevertheless, a time-myopic optimization may be more realistic than the omniscient, global monolithic one. Our scenario, dependent on known consumption and capacity data, assumes deterministic demand, allowing us to anticipate every future battle group nuance. Such omniscience is arguably "too optimal." The temporal cascade more closely mimics CLF planning that considers recent history and a reasonable forecast of near-term demands to develop shuttle schedules for upcoming sorties. While we concede formal optimality, and admit that the current planning window does not consider possible future spikes in demand, actual experience reveals that a cascade returns feasible solutions that are not far from omniscient-optimal. The cascade offers us a reliable solution strategy and can potentially highlight unusually challenging windows of the planning horizon, through the need to recursively halve the planning window, to alert CLF and individual fleet planners to the need for added shuttle capacity.

For exploratory optimization to merely assemble and filter scenario data for errors—preparatory exercises we admit take much more of our time than making subsequent model plays for the record after we have the problem shaped and debugged—the cascade can solve the largest problems in a minute or two on a personal computer. For important plays where we really need a solution quality certificate, or for cases where we must compare two alternatives and come to an unambiguous choice of the better one, we may spend hours solving a monolith model.

An example planning scenario includes 13 battle groups served by nine TAO and seven T-AKE shuttle ships over a

90-day planning horizon. Some of the larger battle groups are accompanied by station ships. Five of the battle groups are deployed on day one, and the rest deploy on day 10, 12, etc., until all are underway to their various areas of operation worldwide. A monolithic (i.e., omniscient) solve of this situation generates an optimization model with about 367,000 constraints and about 23,000 variables. A cascade (i.e., myopic) solve using a 30-day planning windows advanced 15 days at a time yields a sequence of problems with between about 37,000 and 75,000 constraints and between 6000 and 9000 variables. The monolith takes several hours to solve on a 2-9 GHz WINTEL laptop, whereas the cascade takes a few minutes. For this particular case, the two solutions compare closely. The CLF fleet conducts a total of 101 CONSOLs. The availability in this scenario of forward (i.e., close to the areas of operation) supply base ports is key: our average cycle times from a CONSOL back to a port and on to the next CONSOL is only just over five days, two of which are spent pierside loading the shuttle ship.

A hard-to-solve MIP typically arises at the end of a succession of plays with a fixed scenario we try to sustain with fewer, and fewer, and finally the absolute minimum number of CLF ships. Given the cost of building any CLF ship, you can see how such a scenario arises. Realistically, these at-the-edge plans are so fragile, we are uncomfortable advising them.

We think our most important model feature is seen in constraints (1) and (2), along with elastic variables for shortage and extremis shortage. A cumulant inventory constraint, combined with an elastic violation device for a shortage, and another for extremis shortage, combine to signal a deficiency, and carry this forward, paying a daily penalty, until this deficiency is remediated. The distinction between shortage and extremis penalties is important: CLF ships are equitably assigned to CONSOL all customers by maximizing deliveries, to preferentially serve needy customers by CONSOL to avoid shortage, and to energetically serve customers in extremis. For more motivation, see Brown, et al. [4].

We use four commodity groups because this is a natural aggregation of at-sea logistics, and forecasting at more than this level of detail with logistic planning factors for each battle group, daily, over a long planning horizon, is daunting at this scale. We have experimented with more commodities, and there is no technical reason we could not accommodate this detail, but the four aggregate groups we now use are an adequate logistical representation of the demands we satisfy at sea.

12. CONCLUSIONS

We have been able to answer some important exigent strategic and operational questions about our Combat Logistics Force. The planning factors and ancillary tools developed to support this work provide a reliable foundation upon which fleet planners can base a comparative analysis. Clearly-stated modeling and data assumptions combine with mathematical optimization to render advice with two distinguishing advantages: optimization has earned its reputation for teasing some surprising insights from the scenarios, and each monolithic solution comes with a quality certificate assuring that no better solution remains undiscovered. These advantages convey an unusual level of confidence, especially in comparison to ubiquitous simulation tools.

We have been studying these problems long enough now to see some of our prescriptions come to life in the CLF fleet. The basic question is always "Can we logistically support this plan?" The plans, and there are a lot of these we constantly tend, express our navy's commitments to support current defense doctrine (e.g., see the latest Quadrennial Defense Review [11].

REFERENCES

[1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin, Network flows: Theory, algorithms, and applications, Prentice Hall, Upper Saddle River, NJ, 1993.

[2] D.O. Bausch, G.G. Brown, and D.R. Ronen, Scheduling short-term marine transport for bulk products, Maritime Policy Management 25 (1998), 335–348.

[3] K.D. Borden, Optimizing the number and employment of combat logistics force shuttle ships, with a case study of the new T-AKE ship, MS Thesis in Operations Research, Naval Postgraduate School, 2001.

[4] G.G. Brown, R.F. Dell, and R.K. Wood, Optimization and Persistence, INTERFACES 27 (1997), 15–37.

[5] G.G. Brown, G.W. Graves, and D.R. Ronen, Scheduling ocean transportation of crude oil, Management Sci 33 (1987), 335–346.

[6] G.G. Brown, J.E. Kline, R.E. Rosenthal, and A.R. Washburn, Steaming on convex hulls, INTERFACES 37 (2007), 342–352.

[7] J.A. Cardillo, Optimizing global operations plans for the combat logistics force, MS Thesis in Operations Research, Naval Postgraduate School, 2004.

[8] Transitioning the combat logistics force in the 21st Century, Center for Naval Analyses (CNA), 1998.

[9] G.B. Dantzig and D.R. Fulkerson, Minimizing the number of tankers to meet a fixed schedule, Naval Res Logist Quart 1 (1954), 217–222.

[10] W.C. DeGrange, Optimizing global combat logistics force support for sea base operations, MS Thesis in Operations Research, Naval Postgraduate School, 2005.

[11] Department of Defense, Quadrennial defense review, 2006, Available at www.defenselink.mil/qdr/ (last accessed 12 December 2007), February 6.

[12] D.E. Doyle, Evaluation of fleet ownership versus global allocation of ships in the combat logistics force, MS Thesis in Operations Research, Naval Postgraduate School, 2006.

[13] H. Eccles, Operational Naval Logistics (NAVPERS 10869), Bureau of Naval Personnel, 1950.

[14] GAMS Development Corporation, 2006, Available at www.gams.com, last accessed 28 January 2007.

[15] R.D. Givens, A comparison of the operational potential and capability of two combat logistics force alternatives, MS Thesis in Operations Research, Naval Postgraduate School, 2002.

[16] R. O'Rourke, Navy ship acquisition: Options for lower-cost ship designs —Issues for congress, Congressional Res Service Report RL32914, 2005.