Received: 5 January 2018 | Revised: 23 January 2018 | Accepted: 24 January 2018
DOI: 10.1002/nav.21780

RESEARCH ARTICLE

WILEY

# Scheduling combat logistics force replenishments at sea for the US Navy

**Gerald G. Brown** | **Walter C. DeGrange** | **Wilson L. Price** | **Anton A. Rowe**

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

**Correspondence**
Gerald G. Brown, Operations Research Department, Naval Postgraduate School, Monterey, CA 93943.
Email: ggbrown@nps.edu

**Funding information**
Office of Naval Research; Military Sealift Command

**Abstract**

The Replenishment at Sea Planner (RASP) is saving the U.S. Navy millions of dollars a year by reducing fuel consumption of its Combat Logistics Force (CLF). CLF shuttle supply ships deploy from ports to rendezvous with underway U.S. combatants and those of coalition partners. The overwhelming commodity transferred is fuel, ship-to-ship by hoses, while other important packaged goods and spare parts are high-lined, or helicoptered between ships. The U.S. Navy is organized in large areas of responsibility called numbered fleets, and within each of these a scheduler must promulgate a daily forecast of CLF shuttle operations. The operational planning horizon extends out several weeks, or as far into the future as we can forecast demand. We solve RASP with integer linear optimization and a purpose-built heuristic. RASP plans Replenishment-at-Sea (RAS) events with 4-hour (Navy watch) time fidelity. For five years, RASP has served two purposes: (1) it helps schedulers generate a daily schedule and animates it using Google Earth, and (2) it automates reports command-to-ship messages that are essential to keep this complex logistics system operating.

**KEYWORDS**
decision support, fuel conservation, naval logistics

"Oh God, thy sea is so great and my boat is so small,"

Breton Fisherman's Prayer that graced the Oval Office desk of John F. Kennedy.

We might irreverently paraphrase for our purposes "my boat is so *slow*."

## 1 | THE COMBAT LOGISTICS FORCE

Scheduling at-sea resupply of combatant ships of the U.S. Navy and its coalition partners is a key factor in maintaining a high state of readiness and, in time of conflict, operational effectiveness. The Navy Combat Logistics Force (CLF) consists worldwide of about 30 special transport ships of several types that carry ship and aircraft fuel, ordnance, dry stores and food, as well as mail and repair parts. CLF *shuttle* ships deploy from port to deliver supplies via replenishment-at-sea (RAS) rendezvous with client combatant ships underway, thus making it possible for our naval forces to operate at sea for extended periods. This is a complex operation given the numbers of ships involved, their wide dispersion, and the necessary frequency of replenishments.

This article describes this unique problem, an optimization-based decision support system to solve it, the system's road to adoption and the benefits achieved by the system's use for 5 years.

## 2 | US NAVY 5TH AND 7TH FLEETS

The Navy has divided the world's oceans into administrative divisions (see Figure 1): the US Fleet Forces Command covers the Atlantic Ocean, 3rd Fleet the Eastern Pacific, 5th

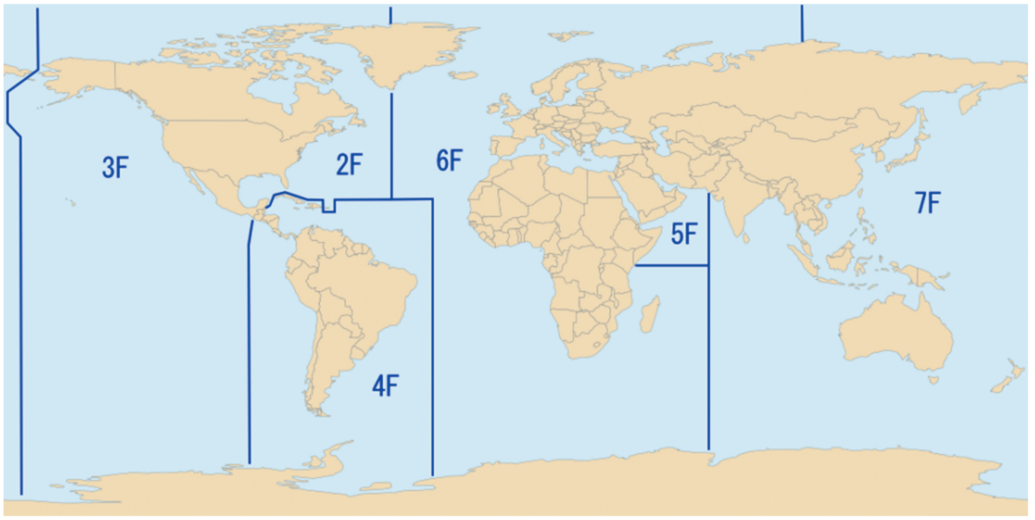


**FIGURE 1** US Navy numbered fleet areas of responsibility worldwide. 7th Fleet extends from the International Dateline in the Pacific Ocean to the Arabian Gulf and Indian Ocean western border of India, from the Kuril Islands in the north, to Antarctica, an ocean area more than 14 times that of the continental United States. The majority of world seaborne commerce traverses this area. 7th Fleet assures "free trade" in this area with a force of about 40-to-50 combatant ships [Color figure can be viewed at wileyonlinelibrary.com]

Fleet the Red Sea-Indian Ocean – Arabian Gulf, 6th Fleet the Mediterranean, and 7th Fleet the Western Pacific. The last three Fleets are of most concern due to the large areas they cover and their distances from the continental U.S. Each Fleet has an office detachment that schedules its CLF ships to resupply combatants in its area.

Worldwide operations require the U.S. Navy and its logistics support ships to traverse vast ocean areas, particularly in the 7th Fleet. Combat logistics shuttle ships must have the endurance to support operating forces across this expanse. Imagine for a moment starting an automotive trip from, say, Los Angeles to Denver, a distance of just over 1000 statute miles (about 1600 kilometers, or 884 nautical miles), and being limited to a sustained (ie, no fuel or rest stop) underway speed of 21 statute miles per hour (33 kilometers per hour, or 18 knots). Your non-stop trip will require 49 hours, one-way. This is representative of the distances and speed at which Combat Logistic Force supply shuttles can operate. (18 knots is actually faster than we would prefer to travel because minimal fuel consumption speed is even slower.)

Why is the U.S. Navy dispersed worldwide? In July, 1944, the U.S. invited all its allies to Bretton Woods, New Hampshire, to make a surprising announcement. The U.S. committed to defend "free trade" globally: our Navy would protect all maritime trade among allies. Although the resulting agreement (Bretton Woods Agreements Act 1945) is rather terse about this commitment, the conference was quite explicit, and the seven decades since have been the most peaceful on the high seas in human history (eg, Chambers, 1999).

For example, the majority of global seaborne commerce traverses the 7th Fleet area of responsibility (AOR), where the Navy and coalition partners patrol with between 40 and 50 combatant ships. These combatants often operate together in strike groups, so at any given time there are 3-to-6 of these larger strike groups, and about 10-to-30 combatants operating independently or in small groups, for a total of 13-to-36 moving customers. 8-to-10 CLF supply shuttle ships are loaded at ports located all around this AOR. A shuttle ship may visit more than one customer on a deployment from port, may return to a different port, and may anticipate making many such sorties over a planning horizon of 360 days, or longer. (5-th fleet has 30-to-40 combatants, 2-to-4 larger strike groups, 15-to-25 independent or small group combatants served by 6-to-8 shuttles, and a planning horizon of 30-to-45 days.)

## 3 | CLF SCHEDULING

Intended operational schedules for combatants are published well into the future. In some cases, the scheduler may not know which particular combatant will be in a particular area performing some particular mission, but we do know some combatant will be there and likely what sort of combatant (ie, what class of ship) this will be. For instance, we can expect to have persistent presence off the Horn of Africa for piracy interdiction. And, of course, contingencies may require changes. Regardless, these intended notional schedules are exogenous inputs to the CLF scheduler over which the scheduler has no control. With this incomplete information, preliminary logistics schedules are created to confirm the intended operations can be supported with available CLF resources.

Closer in, the scheduler receives daily reports from each combatant with its state (fuel and other supplies), and we have some idea from that ship's schedule the rate at which she will continue to consume these. Commodities include marine diesel fuel used in propulsion plants, jet propulsion aviation fuel suitable for shipboard storage and handling, dry, frozen and chilled goods, ultra-high-temperature milk, water, and ordnance (with respective mnemonics DFM, JP5, DRY, FRZ, CHL, UHT, H2O, and ORDN). By weight, DFM and JP5 completely dominate deliveries, but the other categories require much greater per-weight storage, handling, and transfer care. U.S. combatants can make their own fresh water, but not all coalition ships can do so.

Typically, a rolling schedule is revised daily taking into account changing operational assignments and states of combatants. Deciding when, where, and by which CLF ship each combatant customer should be served superficially resembles a multiple traveling salesman problem with moving customers and many operational side constraints, including ports (ie, supply depots) with varying costs, availability of commodities, and hours of operation. Side constraints such as those accounting for demands that vary with customer employment over time complicate matters.

## 3.1 | Manual scheduling

Previous to this project, planning was manual, and the principal use of computer systems support was for recording and managing data. The schedulers manually collected data concerning each combatant customer, her predicted future position, time and place of the last replenishment and the current levels of onboard commodities and ordinance. Then, a specific supply ship was manually assigned to resupply one or more specific combatants at specified times and locations. The scheduler's preliminary routine plan was reviewed by a parallel logistics team to suggest remedies for any conflicts or inconsistencies, with each logistics team member having different expertise and responsibility.

Determining port call schedules is a key concern. These require diplomatic clearance and require prepositioning combatant customer-specific cargo, such as repair parts.

The schedulers may not have had an explicit objective function, as operations researchers (OR's) usually do, nor were they likely to refer to solutions as being "optimal" or otherwise. We venture to call such a plan a "satisficing" solution if the customers, the combatant ships, are satisfied with the service level they receive.

In the manual scheduling mode, schedulers make assignments, a customer at a time, typically using rules of thumb that have worked in the past. The quality and feasibility of the overall schedule is assessed by a larger group during a holistic review where many pairs of eyes can try to find a problem with the proposed plan or approve it as satisfactory. The quality of a schedule generally depends on the experience level of the scheduler, and a number of factors may collude to render a less than optimal, sometimes less than satisfactory schedule:

- individual ship assignments may be based on rules of thumb that are outdated or actually don't apply in more complicated cases;
- scheduling by hand is time-consuming and schedulers may not have a chance to explore multiple alternatives, generally proceeding with the first option they find;
- assignments may conflict with each other in ways not immediately evident to a scheduler, such as capacity constraint violations. Individual assignments may look good, but collectively a plan may become unworkable; and
- last-minute changes in requirements or delivery locations may offer short notice and few easy corrective options. Manual rescheduling is challenging, even daunting, especially for new-to-the-job surface schedulers.

Current US Navy personnel policies for the management of the schedulers do not mitigate the effects of any of these factors: naval personnel are typically assigned to logistical commands for periods of 1 or 2 years, and receive only on-the-job training once they arrive. If overlap with outgoing personnel is not adequate, and it often is not, this results in a significant recurring loss of expertise.

## 3.2 | Past experience with CLF and other marine shippers is tantalizing bait

When Navy leadership thought there was an opportunity to achieve significant cost savings (SEE TEXT BOX) in the scheduling of CLF ships, we were asked to help (Wray 2010). This should not be interpreted as a rebuke of the manual schedulers. They had never been asked to consider costs and they had been producing otherwise satisfactory results for many years.

REDUCING THE COST OF CLF OPERATIONS
Reckoning the cost of CLF shuttle ships and their cargos is too complicated for manual scheduling. Each CLF shuttle may be sent to a number of ports with differing port charges, fuel and commodity availabilities, and prices. A large cost component is the fuel burned by the CLF shuttles themselves; to estimate this for any given resupply plan, one must know the distance each CLF vessel must travel as well as the speed maintained on each leg of the trip. Fuel consumption is reduced by travelling at the slowest speed consistent with the supply needs of the combatant customers.

The typical remedy available to adjust or correct a schedule is to order the CLF shuttles to *travel faster*, which is costly because fuel consumption is a super-linear function of underway speed. Avoiding this recourse is one of the largest available sources of savings.

We agreed to model this situation, and provide a user interface and systems to talk to other systems. We were anticipating applying modeling and optimization as we have done for some time for CLF (Brown & Carlyle 2008). However, our prior work has been strategic capacity assessment, not operational planning, and we had not before had to worry about time fidelity finer than a day, or CLF shuttle ship fuel consumption, or a rolling time horizon cluttered with arbitrary fixed future events.

## 4 | REPLENISHMENT AT SEA PLANNER (RASP)

We had to develop a clean-sheet model dubbed Replenishment at Sea Planner (RASP) with time fidelity originally requested to be hours, but later relaxed to be 4-hour Navy watches. In peacetime, we prefer daylight RAS evolutions for the safety of our crews (ie, forenoon 08-12, afternoon 12-16, and dog watch 16-20, although operations are also possible during first 20-00, middle 00-04 and morning watch 04-08). There may also be shift premium cost considerations for civilian merchant mariner shuttle crews on weekends and holidays. We had to consider fine geographic detail, operational business rules, and recognition of the full costs of all actions taken by CLF ships, *including their own-ship fuel consumption*. The team that has developed RASP is composed of experienced modelers and seafarers with exhaustive knowledge of and experience with Navy operations.

### 4.1 | Prior work in ship scheduling and routing

About 50,000 civilian merchant ships carry approximately 90% of world trade, and cost almost $400 billion annually (ICS, 2018). Understandably, models to improve routing and scheduling of these ships have received a lot of attention in the open literature. A sequence of survey papers reviews this literature (Ronen, 1983, 1993; Christiansen, Fagerholt, & Ronen, 2004; Christiansen, Fagerholt, Nygreen, & Ronen, 2013; Meng, Wang, Andersson, & Thun, 2014). Christiansen et al. (2013) alone refer to 132 citations. These surveys offer a taxonomy of problem characteristics and cite papers accommodating various subsets of these. Some models concentrate on ship activities, while others focus on cargo supply chains. For ship scheduling, a modest variety of models apply, including linear integer programs that track binary flows (ship movements) through a space-time network (such as the network we present). Another modeling option is to generate candidate voyages through a space-time network and select one such voyage for each ship scheduled (as is done here). A third alternative is constructing voyages heuristically while adhering to whatever side constraints apply (as our heuristic does). Any of these models may be generated a priori (as we do), or on-the-fly during solution. Some are omniscient over a full planning horizon (as ours is), or intentionally myopic over a shorter rolling time horizon. Some papers employ off-the-shelf optimization packages to solve the linear integer models directly (as we have done), some employ decomposition to generate a sequence of subproblems to solve, and others suggest search heuristics (eg, Korsvik, Fagerholt, &Laporte, 2010, 2011). The reader is cautioned that only a scarce few of these research papers include real-world application.

Examples of prior military cargo ship scheduling and routing include Dantzig and Fulkerson (1954) introducing a network linear program to assess the ability of Navy tankers to deliver fuel, and Brown and Carlyle (2008) presenting a network integer linear program to assess logistics capacity worldwide for the US Navy simultaneously responding to multiple crises. The distinguishing differences between routing and scheduling civilian merchant ships and our CLF include: (1) the bulk of our cargos are fungible between loading and delivery location, rather than source–destination specific; (2) loading and (especially) delivery locations are not fixed, and suggesting these is key advice from our models; (3) customers are moving, often faster than the maximum speed of our supply ships; (4) our ship routes are not pre-cataloged in standard shipping guides as are those for commercial vessels; (5) a wide diversity of commodities are conveyed, ranging from frozen food to flammable fuels to explosive munitions; (6) simultaneous pipeline, highline, and helicopter deliveries take place; and (7) we must respond quickly to the vagaries of an extended, daily, real-world application influencing operations over an Earth hemisphere.

### 4.2 | Replenishment at sea planner (RASP) modeling

The following describes the components of RASP.

#### 4.2.1 | Indices [∼cardinality]

$g \in G$ Combatant customer strike group [∼30]
Each group is composed of one or more combatants traveling in company.

$d \in D$ Day of planning horizon, a contiguous ordinal set [∼45–360] (alias $\delta$)

$t \in T$ Time period, a contiguous ordinal set (alias $\tau$)

May be a day, or some fraction of a day. $|T| = k|D|$ for some integer $k \geq 1$.
For example, with six four-hour naval watches per day, $k = 6$.

$d(t)$ Day of time period $t$, $d(1) = 1$, $d(k) = 1$, $d(k+1) = 2, \ldots$

$t \in T_d$ Set of time periods during planning day $d$

$s \in S$ Shuttle ship [~10]

$c \in C$ Commodity group (DFM, JP5, DRY, FRZ, CHL, UHT, H2O, ORDN) [~8]

$p \in P$ Port [~25]

$a \in A$ Set of potential actions for a shuttle ship at any location. $A = \{POS, LOG\}$, indicating, respectively, that a shuttle must simply be in a given location on the start of a time period, or that the shuttle ship will have an opportunity for a logistic event with some strike group $g$.

### 4.2.2 | Given data [units]

These are exogenous inputs over which we have absolutely no control.

$lat_{g,d}, lon_{g,d}$ Coordinates of combatant strike group $g$ at start of day $d$ [degrees]

$RAS\_OK_{g,t} = 1$ if RAS permissible during time period $t$, 0 otherwise [binary]

$window_g$ Minimum number of days between RAS events (from any shuttle) for combatant customer $g$ [days]

$s\_lat_s, s\_lon_s$ Initial coordinates of shuttle ship $s$ [degrees]

$min\_speed_s, max\_speed_s$ Minimum, maximum speed of shuttle ship $s$ [knots]

$fuel_s(speed)$ Fuel consumption as a function of shuttle speed. Standard navy fuel consumption tables are in gallons per hour or barrels per day versus speed in knots. The function here maps knots to [fuel units]

$inptTAT_s$ Inport turn-around time to reload shuttle ship $s$ [days]

$p\_lat_p, p\_lon_p$ Coordinates of port $p$ [degrees]

$x\_lat_{s,t,a}, x\_lon_{s,t,a}$ Coordinates fixed by the scheduler for shuttle $s$ to occupy at the start of timeperiod $t$ to perform action $a$ [degrees].

$g\_lat_{g,t}, g\_lon_{g,t}$ Coordinates of strike group $g$ at the start of time period $t$ [degrees]

$fuel\_cost$ Cost of shuttle own-fuel, shuttle diesel fuel marine (DFM) [\$/fuel unit]

$port\_cost_{s,p}$ Cost of a visit by shuttle $s$ to port $p$ [\$]

### 4.2.3 | Minimum-distance transit routes

To assess the feasibility of a logistics plan, RASP must know what routes CLF ships will take between sequential at-sea replenishment locations and to and from resupply ports. Given input is simply latitude and longitude of these locations, with no indication of how best to transit between them, it is up to RASP, as an internal matter, to determine how long it will take to transit from each point to the next within a particular AOR. Prior work by Brown and Carlyle (2008, p. 802) designed a global, static sea route network of navigable node coordinates and legs connecting adjacent nodes. This static route network proved too coarse and unwieldy for an operational environment, where finer geographic resolution is needed and destinations change frequently. What RASP required was an on-demand routing engine, capable of connecting any two points within an AOR and returning an estimated transit distance. This did not turn out to be trivial, because obstacles to navigation such as land masses, and perhaps diplomatically or militarily excluded areas, must be avoided. Using a data base (not illustrated here) containing spherical polygon vertices defining these obstacles, as well as port and canal entrance locations, all AOR positions are automatically connected via minimum-distance great-circle ocean routes avoiding obstacles using a purpose-built algorithm (Washburn & Brown, 2016; Brown & Washburn, 2017). (This has also been retrofitted into the CLF model of Brown & Carlyle, 2008.) Figure 2 illustrates all minimum-distance sea routes from one location of interest to all others.

With minimum-distance sea routes calculated from any given location in an AOR to any other, we can begin to assemble candidate CLF voyages, correctly assessing the feasibility and fuel cost for each.

### 4.2.4 | Candidate voyage generation

For each shuttle $s$ define an acyclic time-based directed graph with a node for every candidate event (a coordinate and time). There is an initial node $s\_lat_s, s\_lon_s$ for $t = 1$, a node for every port coordinate in every time period $p\_lat_p, p\_lon_p$, a node for every combatant strike group time period coordinate $g\_lat_{g,t}, g\_lon_{g,t}$ for which $RAS\_OK_{g,t} = 1$, and a node for every coordinate fixed ("pinned") by the scheduler $x\_lat_{s,t,a}, x\_lon_{s,t,a}$. Passing through the nodes by time period $t$, define an arc from every node to every subsequent node in a following time period that is adjacent (ie, can be reached) *if the shuttle travels within the allowable speed interval* $[min\_speed_s, max\_speed_s]$, not port-to-port, accounting for any port visit delay of $inptTAT_s$ days, and not if the arc would make two RAS visits to combatant strike group $g$ within *window* days. If the scheduler has fixed events (pinned location-times) that violate shuttle speed restrictions, inport turn-around times, or visits a strike group $g$ too frequently, this is permitted *but diagnosed*. Manual restrictions by the scheduler can extinguish voyage paths, and this is permitted *but diagnosed*.

Whenever there is a pinned event in the directed graph, this can best be exploited by back-casting (on this acyclic graph) the earliest and latest times by which any preceding location must be reached while still permitting the subsequent


**FIGURE 2** Minimum-distance sea routes avoiding obstacles from one source location of interest to all others. Coordinate N22 37′55″E178 38′55″ is on the "chop line" between 3rd Fleet to the east, and 7th Fleet to the west, where responsibility for a combatant customer transitions between these fleets [Color figure can be viewed at wileyonlinelibrary.com]

pinned event to be achieved. The forward, stack-based enumeration of paths can backtrack much earlier when encountering a location-time event from which a subsequent pinned event cannot be reached. This trades work by a linear-time back-cast labeling for an exponential forward-reaching enumeration. The back-casting can be terminated when it reaches a prior pinned event, and at this point it may be possible to diagnose inconsistencies between the successively pinned events.

The (time-forward) path enumeration originates at a location and time where a shuttle is planned to be (a pinned event), and proceeds forward in time until the time of the next pinned event. For most daily planning cycles, the path enumeration is restricted to discover candidate paths that might fill in gaps between pinned events.

Due to the relatively slow speeds of our shuttle ships, and the large distances to traverse, we can selectively generate all achievable voyages with stack-based enumeration of all paths through this directed graph over a planning horizon of 30 days or more (eg, see fig. 5 on page 46 of Brown, Carlyle, Dell, & Brau, 2013).

As the voyages are enumerated, their leg-by-leg own-fuel consumption cost is computed from *fuel_cost* and the $fuel_s(speed)$ function, and port charges $port\_cost_{s,p}$ are assessed. (Demurrage costs would also apply if we were to use in-charter supply shuttles.)

Strike groups are supposed to maintain at all times a given fraction of their fuel capacity called SAFETY stock. Because any manual restriction by the scheduler is permissible, we must allow for fuel levels that violate this, and even fall to a lower EXTREMIS fraction of capacity, and perhaps even NEGATIVE levels. This last is a modeling artifice introduced to make anything the scheduler restricts admissible, and to diagnose the times and magnitudes of induced shortages.

For a 30-day planning horizon in a large AOR, each shuttle may anticipate one of as many as 10,000 achievable voyages, though *commitments fixed by the scheduler and widely separated coordinates typically result in far fewer candidate voyages*. The result is:

#### 4.2.5 | Generated indices [∼cardinality]

$v \in V_s$ Voyages for shuttle $s$ [∼10,000]

$\{g,t\} \in GT_v$ For voyage $v$, two-tuples of combatant strike group $g$ time period $t$ RAS rendezvous

$\{t,p\} \in TP_v$ For voyage $v$, two-tuples of time periods with port $p$ visits

#### 4.2.6 | Generated data [units]

$v\_cost_v$ Voyage costs (excluding loaded commodity costs) [$]
$fuel\_burned_{v,t}$ Own-fuel burned by shuttle $s$ on voyage $v \in V_s$ during time period $t$ [fuel units]

Given this generated data, we are prepared to suggest an optimized fleet schedule.

#### 4.2.7 | Additional index

$\ell \in L$ Commodity level (eg, SAFETY, EXTREMIS, NEGATIVE), an ordinal set

#### 4.2.8 | Additional given data [~units]

RASP maintains a catalog of combatant customer ships with their commodity capacities and consumption rates for a variety of employment activities. These are automatically aggregated for a combatant customer strike group of ships.

$g\_uses_{g,t,c}$ Consumption by $g$ during time period $t$ of commodity $c$ [$c$-units]
$g\_mxload_{g,c}$ Maximum capacity of $g$ to carry commodity $c$ [$c$-units]
$g\_starting\_c_{g,c}$ $g$ Inventory at start of planning horizon of commodity $c$ [fraction of $mxload_{g,c}$]
$g\_limit\_c_{g,c,\ell}$ Commodity limit triggering a shortage violation (ie, safety stock) [$c$-units] [fraction of $mxload_{g,c}$]
$g\_penalty\_c_{g,c,\ell}$ Positive multiplicative penalty for a shortage violation [$/$c$-unit violation]
$g\_priority_g$ Weight assigned to RAS volume delivered to strike group $g$ [scalar]
$npv_t$ Net present value discount term [fraction] We sometimes refer to this as the "fog of future planning" discount.
$s\_capacity_{s,c}$ Shuttle ship $s$ capacity for commodity $c$ [$c$-units]
$s\_init\_load_{s,c}$ Shuttle ship $s$ initial inventory of commodity $c$ [fraction of $s\_capacity_{s,c}$]
$pier_s$ Pier capacity used by shuttle $s$ [pier capacity]
$pier\_cap_p$ Port capacity [pier capacity]
$c\_cost_{p,c}$ Commodity $c$ cost at port $p$ [$/$c$-unit]
$c\_priority_c$ Priority of commodity $c$ [scalar]
$reward_c$ Reward for delivery of commodity $c$ [$/$c$-unit]

#### 4.2.9 | Decision variables [units]

$VOYAGE_v$ Binary indicator that shuttle voyage $v$ is selected.
$VISIT_{g,d}$ Binary indicator that at least one shuttle visits $g$ on day $d$
$LOAD_{s,t,p,c}$ Amount of commodity $c$ loaded by shuttle $s$ at start of time period $t$ at port $p$ [$c$-units]
$HOLD_{s,t,c}$ Shuttle $s$ commodity $c$ contents at start of time period $t$ [$c$-units]
$RAS_{s,g,t,c}$ Amount of shuttle $s$ delivery to $g$ during time period $t$ of commodity $c$ [$c$-units]
$VIOLATION_{g,t,c,\ell}$ Amount of inventory deficiency of $c$ for $g$, at start of planningperiod $t$below level $\ell$ [$c$-units]

#### 4.2.10 | Simplified formulation

$$\text{s.t.}\quad HOLD_{s,t,c} - \sum_{v\in V_s|c=\text{`DFM'}} fuel\_burned_{v,t} VOYAGE_v - \sum_{g\in G} RAS_{s,g,t,c} + \sum_{p\in P} LOAD_{s,t,p,c} \overset{\circ}{=} HOLD_{s,t+1|t<\|T\|} \quad \forall s \in S, t \in T, c \in C \tag{1}$$

$$\sum_{\substack{s\in S,\\ \tau<t}} RAS_{s,g,\tau,c} \overset{\circ}{\le} \sum_{\tau<t} g\_uses_{g,\tau,c} + [g\_mxload_{g,c}(1 - g\_starting\_c_{g,t,c})]_{t=1} \quad \forall g \in G, t \in T, c \in C \tag{2}$$

$$\sum_{\substack{s\in S,\\ \tau<t}} RAS_{s,g,t,c} + \sum_{\ell\in L} VIOLATION_{g,t,c,\ell} + g\_starting\_c_{g,c} mxload_{g,c} \ge \sum_{\tau<t} g\_uses_{g,\tau,c} - g\_mxload_{g,c}(1 - g\_limit\_c_{g,c,\text{`SAFETY'}}) \quad \forall g \in G, t \in T, c \in C \tag{3}$$

$$RAS_{s,g,t,c} \le \min\{g\_mxload_{g,c}, s\_capacity_{s,c}\} \sum_{v\in V_s|\{g,t\}\in V_s} VOYAGE_v \quad \forall s \in S, \forall g \in G, t \in T, c \in C \tag{4}$$

$$\sum_{v\in V_s} VOYAGE_v \le 1 \quad \forall s \in S \tag{5}$$

$$\sum_{v\in V_s|\{g,t\}\in GT_v \wedge t\in T_d} VOYAGE_v \le VISIT_{g,d} \quad \forall g \in G, d \in D \tag{6}$$

$$\sum_{d-window_g \le \delta \le d} VISIT_{g,\delta} \overset{\circ}{\le} 1 \quad \forall g \in G, d \in D \tag{7}$$

$$LOAD_{s,t,p,c} \le \sum_{v\in V_s|\{t,p\}\in TP_v} s\_capacity_{s,c} VOYAGE_v \quad \forall s \in S, t \in T, p \in P, c \in C \tag{8}$$

$$\sum_{\substack{s\in S,\\ v\in V_s|\{t,p\}\in TP_v}} pier_s VOYAGE_v \le pier\_cap_p \quad \forall t \in T, p \in P \tag{9}$$

$$VOYAGE_v \in \{0,1\} \quad \forall s \in S, v \in V_s$$

$$VISIT_{g,d} \in \{0,1\} \quad \forall g \in G, d \in D$$

$$0 \le LOAD_{s,t,p,c} \le s\_capacity_{s,c} \quad \forall s \in S, t \in T, p \in P, c \in C$$

$$0 \le HOLD_{s,t,c} \le s\_capacity_{s,c} \quad \forall s \in S, t \in T, c \in C$$

$$HOLD_{s,1,c} = s\_init\_load_{s,c} s\_capacity_{s,c} \quad \forall s \in S, c \in C$$

$$0 \le RAS_{s,g,t,c} \le \min\left\{g\_mxload_{g,c}, s\_capacity_{s,c}\right\}$$
$$\forall s \in S, g \in G, t \in T, c \in C$$
$$0 \le VIOLATION_{g,t,c,\ell} \le g\_mxload_{g,c} * (g\_limit\_c_{g,c,\ell}[-g\_limit\_c_{g,c,\ell-1}]_{\ell>1})$$
$$\forall g \in G, t \in T, c \in C, \ell \in L \quad (10)$$

$$\underset{\substack{VOYAGE,VISIT,\\ LOAD,HOLD,\\ RAS,VIOLATION}}{\mathrm{MIN}} \sum_{\substack{s\in S,\\ v\in V_s}} v\_cost_v VOYAGE_v + \sum_{\substack{s\in S,t\in T,\\ p\in P,c\in C}} c\_cost_{p,c} LOAD_{s,t,p,c} + \sum_{\substack{g\in G,t\in T,\\ c\in C,\ell\in L}} npv_t g\_penalty\_c_{g,c,c\_level} VIOLATION_{g,t,c,\ell} - \sum_{\substack{s\in S,g\in G,\\ t\in T,c\in C}} g\_priority_g * c\_priority_c reward_c RAS_{s,g,t,c} + \text{elastic penalties} \quad (11)$$

#### 4.2.11 | Discussion

Each equality (1) accounts for a shuttle cargo contents from one period to the next, for every time period in the planning horizon (the model can schedule a single shuttle ship sortie from port to make many separate RAS visits, perhaps to different combatant strike groups). Each inequality (2) limits time period by time period cumulative RAS volumes of each commodity to the cumulative usage of each strike group through the end of that time period. We assume that at the start of the first planning day, each strike group contains some stated initial load quantity. Thereafter, daily use is deducted and shuttle RAS volumes are added. Each elastic inequality (3) determines the difference between the cumulative inventory state of each commodity at the end of each time period and the cumulative usage less desired safety-stock level at the end of that time period, representing any shortage, extreme shortage, or negative inventory required to reconcile this state. Each inequality (4) limits the RAS volume transferred from a shuttle ship to a strike group on some given time period to be zero unless a replenishment event takes place on a selected shuttle voyage. Each inequality (5) allows at most one voyage to be selected for each shuttle. Each inequality (6) forces a visit event for any selected voyage containing that visit. Each inequality (7) allows each combatant strike group to be interrupted by at most one RAS event per a specified minimum time epoch. Each constraint (8) limits loading of each shuttle to locations and times selected by a voyage for that shuttle. Each constraint (9) ensures simultaneous shuttle port calls do not exceed the capacity of a port during a time period. Variable domains are stated by (10). The objective (11) expresses the fuel cost of selected shuttle voyages, in-port charges for shuttle port calls, fixed costs of shuttle operation, differential costs of commodities by loading port, a penalty with a component for any combatant shortage below safety-stock, extreme shortage below minimum stock, and any negative inventory, less a reward for commodity volume delivered (this reward offers a mechanism to concentrate schedule attention by commodity and by strike group).

### 4.3 | Optional elastic features

On occasion a scenario may exhibit infeasibility due to data errors or scheduler inputs that exceed any feasible course of action. In such cases, we employ elastic logical variables to help diagnose and repair the infeasibility. Constraints (1), (2), and (7) are instrumented by addition of the following terms to their respective right-hand sides:

$-S\_INV\_ARTIFICIAL_{s,t,c} + S\_INV\_SURPLUS_{s,t,c}$ for each $\overset{\circ}{=}$ constraint (1),

$+G\_INVENTORY\_SURPLUS_{g,t,c}$ for each $\overset{\circ}{\le}$ constraint (2), and

$+G\_VISIT\_SURPLUS_{g,d}$ for each $\overset{\circ}{\le}$ constraint (7).

The objective function is augmented respectively with penalty terms:

$$\sum_{s\in S,t\in T,c\in C} npv_t penalty_c (S\_INV\_ARTIFICIAL_{s,t,c} + S\_INV\_SURPLUS_{s,t,c}),$$
$$\sum_{g\in G,t\in T,c\in C} npv_t penalty_c\ G\_INV\_SURPLUS_{g,t,c}), \text{ and}$$
$$\sum_{g\in G,d\in D} npv_{\min\{t(d)\}} visit\_penalty\ G\_VISIT\_SURPLUS_{g,d}$$

(where *visit_penalty* is used by the scheduler to discourage too-frequent RAS interruption of strike groups).

In practice, the last of these options is almost always essential: It is very hard for a scheduler to envision situations where shuttles must visit some strike group more frequently than desired.

### 4.4 | Practical considerations

This is called a “simplified” formulation because considerable notational clutter has been suppressed. For example, a strike group may not be present for some subset of the planning horizon, may not demand all commodity types, and may make a port call to resupply pier-side (we know when this will happen, but do not control it). Similarly, a shuttle may not be available for some subset of the planning horizon, may not be able to load at some ports, may not carry all commodity types, and may be able to adjust storage between commodity categories. The operational system accommodates all this detail at the expense of a lot of additional, distracting notation. The resulting model is more complicated to view, but no harder to solve.

In a consolidation ("console") event, two shuttles rendezvous at sea to transfer cargo. This might include rush-order repair parts or mail destined for a particular customer ship, picked up by one shuttle and relay-transferred to another bound for that customer. These console cargos typically show up at the last minute, and so we do not try to automatically schedule console events, depending upon the scheduler to pin these rendezvous for us.

*A scheduler will not use a planning system that cannot be completely controlled manually*. Accordingly, any event and location can be pinned by the scheduler ($x\_lat_{s,t,a}, x\_lon_{s,t,a}$). There is considerable machinery to *alert the scheduler when such a restriction is not feasible* (say, requiring a shuttle to make 100 knots underway), and the graphical user interface uses Google Earth to animate schedules. Scheduler restrictions can be beneficial by greatly reducing the number of admissible voyages for a shuttle. We have not needed to add a mechanism to retain persistence in schedules suggested from one calendar day to the next (ie, reduce excessive changes in schedules already promulgated) (eg, Brown, Dell, & Farmer, 1996 and Brown, Dell, & Wood, 1997). We were pleasantly surprised that carrying over forecast daily locations from one planning day to the next, coupled with the slow speed of our shuttles and combatant customers compared to the span of the AOR, retains stability intrinsically.

An important feature of our planning systems is that all data is visible and controllable by the planner (scheduler). There are no hidden "IRKs" (independent rheostat knobs) intended to influence solutions behind the scenes.

Finally, although RASP superficially appears to be a vehicle routing problem (VRP), it does not bear close resemblance to those addressed in a large VRP literature. In particular, we do not seek Hamiltonian cycles for our shuttles. Due to all the essential side constraints in our model, we have not had much success trying to apply suggestions from the VRP literature (eg, Lawler, Lenstra, Rinnooy Kan, & Shmoys, 1987; Toth & Vigo, 2002) to guide us to quickly achieving solutions with a certificate of good quality. Packing constraints (5) (that are usually, but not always, satisfied exactly) interact with constraints (6) and (7) and lend significant integer linear program pre-solve reductions and model-tightening cuts (see, eg, Garfinkel & Nemhauser, 1972, Chapter 5 and section 8.3; Nemhauser & Wolsey, 1988, sections II.1 and II.2). We have had success with solvers such as CPLEX (IBM, 2017). If we're careful with our candidate voyage generation, we achieve good-quality solutions within minutes on a portable workstation.

### 4.5 | Initial testing with historical data

Initial testing with historical operational data showed that there was an opportunity for significant cost savings. The U.S. Navy is making large investments in research into new hull coatings, changes to engineering configurations, hydro form designs and similar measures (eg, Karafiath, McCallum, & Hendrix, 2004; OSD, 2016). In these areas, the Navy regards an improvement in fuel efficiency of less than one percent as a major breakthrough. Our early results revealed that improvement of more than an order of magnitude better than this is attainable.

We constructed a model that suggests an overall schedule that minimizes total costs and ensures the timely delivery and complete resupply of all operational vessels. The introduction of a cost-minimization objective was new and differs from that of the manual schedulers who implicitly shared the objective of maintaining customer satisfaction but had no formal numerical metric.

## 5 | INITIAL IMPLEMENTATION EXPERIENCE WITH 5TH FLEET (FROM BAHRAIN)

In November, 2010, we finished, tested and delivered RASP to the 5th Fleet schedulers, located in Bahrain. We had visited to study the existing manual system and how it was employed. Relevant parameters were obtained, business rules explicitly written out and the requirements of combatant ships taken into account. The model was complicated but it ran comfortably on a portable workstation, and the human interface, screens and reports, were state of the art.

We expected rapid adoption of this model, but when presented it the 5th Fleet schedulers and their leadership, RASP encountered heavy push-back. The schedulers told us that RASP would increase workload because it would require additional data entry—with scheduler time already strained by the manual process in place.

To address this concern, we proposed an acceptable *quid pro quo*: RASP would be modified so that all required daily reports and messages would be produced by the enhanced system thus requiring just a single data-entry point. This proposal was very well received. Their existing system already required schedule data to be entered in three different places. RASP would have been a fourth.

Once a version that generates operational reports and messages was ready, RASP was deployed to 5th fleet using a portable workstation. Commercial solvers used by RASP were not authorized for use directly on the local classified network, thus necessitating the deployment via dedicated hardware. Unfortunately, schedulers did not adopt use of the standalone workstation. While RASP introduced the possibility of a single point of data entry, this point of entry was a remote one. The only way authorized to transfer changes between the workstation and the secure network required burning a CD and requesting information technology (IT) assistance to transfer the contents onto the network. Scheduling changes, often required multiple times a day, were far too frequent to make this viable for operational use. In the end, schedulers rejected the stand-alone workstation and additional IT headaches. They preferred redundant data entry at the time.

### 5.1 | If you want to reduce costs, you need visibility of these costs

Moreover, preliminary testing using the stand-alone workstation revealed we wanted a lot more data than was already being entered into any computer data base at 5th Fleet. We needed port costs, port-specific commodity costs and availability, combatant customer forecast locations and states, and any fixed future events already promulgated in a schedule. We could not schedule operations as efficiently as possible without knowing these details. The schedulers did not want to have to create and tend a database with these essential details, preferring instead to rely on situational awareness and ongoing message traffic that is not transcribed into any form we can feed into on a computer. DeGrange (2012) recounts this early experience.

The salient mitigation action taken in October 2011 during the trial implementation of RASP by the 5th Fleet scheduling office was to offer direct support to the schedulers. Our OR team was reinforced by the addition of a reserve officer, Adrian Zavala, with the technical background, job experience and soft skills necessary to understand the potential of RASP and the human relations challenges involved in having it adopted. After a brief training period with the RASP software, he traveled to Bahrain and undertook required data quality improvements and on-the-job training of the scheduler as well as the thankless task of daily transitional data entry. His overall contribution was, and continues to be invaluable.

Meanwhile, from July 2011 to August 2012 other members of our OR team developed a heuristic adaption of the original solver to run directly on their standard-issue networked (classified) computers, thereby eliminating the need for burning CDs and the stand-alone workstation. Additional development efforts also continued to win official approval that every automated report and message adhered *exactly* to legacy formats, including specific colors used for highlights, identical font sizes, and specific print settings.

Even so, although dialog was active and the promise of RASP reducing workload kept them engaged, getting 5th Fleet to actually make the jump to the new system remained elusive. The process of discovery and the elaboration and development of additional desired features and specifications seemed without end.

### 5.2 | RASP adoption and daily use

The final move to acceptance of RASP in its daily scheduling system by 5th Fleet came unexpectedly and suddenly. The 5th Fleet scheduling office committed to base their daily planning operation on RASP March 29, 2013, and continues to use RASP today for daily scheduling and other planning excursions.

Each day, the scheduler creates a new EXCEL workbook and uses an automatic feature to propagate this from the prior day's final schedule with anticipated updating of positions and states of shuttles and combatant customers. This presents a pre-completed, fixed schedule for the entire planning horizon. The scheduler checks the newly appeared last planning horizon day, reviews message traffic for updates and changes to operational schedules, and either revises forecast fixed events by time and location, or frees these to allow RASP to fill back in automatically. On a good day, very little has to be changed by the scheduler or RASP. On a day with changes, the easiest thing to do is un-pin the obsolete legacy schedule features, and either try by hand to adjust things, or let RASP do so.

On a really bad day, there may be no option but to ask the operational commander to make changes to some combatants' schedules. This might involve asking a combatant to make a port call to resupply, or have a transiting strike group change intended course to rendezvous with a shuttle. Sometimes challenges arise, such as after the 2011 Fukushima Japan tsunami, or anytime there is an emergent demand for emergency relief or a show of force. When large numbers of combatants are ordered to make best speed to arrive on station, we may have to abandon our *delivery boy* service and set up a *gas station*, asking customers to come visit us. RASP can do this.

When the scheduler is satisfied, the aforementioned operational reports are produced automatically. Figure 3 is a representative 2-week superposition of schedule activities. Figure 4 shows how we can track scheduled fuel consumption by each shuttle over the planning horizon.

It would be most unusual to conduct an extensive revision of a complete schedule just a day old. The daily scheduling drill may be challenging, but in the end we hope not too many messages are required changing plans already promulgated to and likely in execution by our shuttles.

Soon after adoption, and for months thereafter, 5th Fleet schedulers realized an "extra 2 hours per day to consider alternatives" (per Frank Miller, Lieutenant Commander, USN Supply Corps, direct boss of 5th fleet surface schedulers at the time).

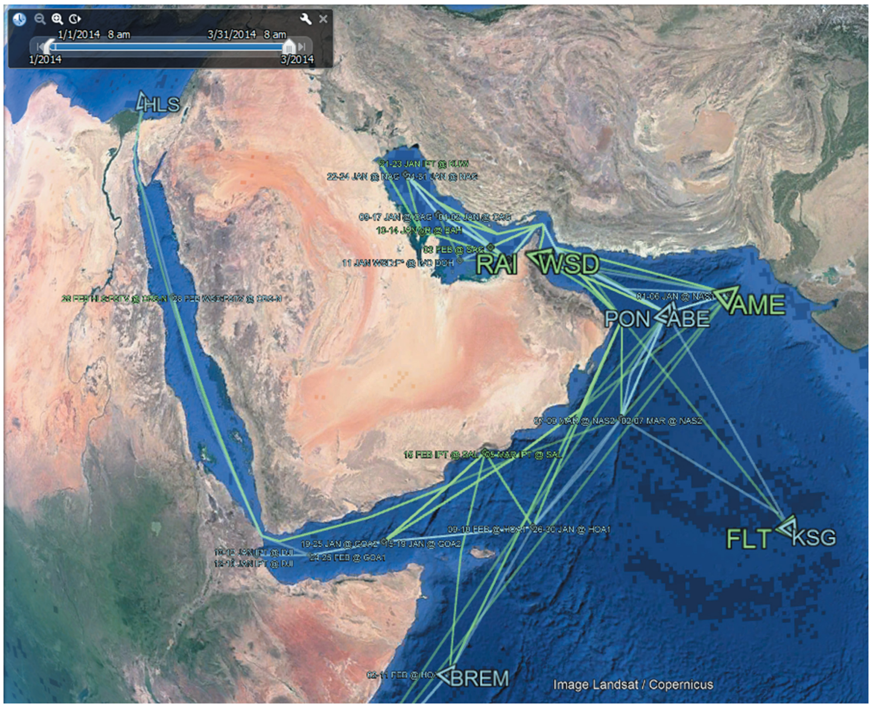


**FIGURE 3** Superposed tracks of ships operating in 5-th Fleet over a 2-week planning horizon, with combatant customer tracks in blue, CLF shuttle ship tracks in green, and thicker tracks used more frequently. The three-letter acronyms identify combatant customers such as the carrier strike group ABE (USS Abraham Lincoln CVN 72 and her escorts). The slider bar at the upper left can be used to freeze the state of all ships, look forward or back at any instant of the planning horizon, or animate motion. Port congestion and restrictions are a key concern in 5-th Fleet [Color figure can be viewed at wileyonlinelibrary.com]

## 6 | SCHEDULING 7TH FLEET (FROM SINGAPORE): RAPID SUCCESS

In parallel to the 5th Fleet effort, we discovered in September 2012 that the 7th Fleet scheduling office (in Singapore) was interested in RASP, so we started a second RASP project in that office, while remaining active in 5th Fleet.

Given our 5th Fleet experience, for 7th Fleet we anticipated local differences in daily reporting and messaging requirements and formats. Such variation is understandable, given the vast domain here. We again matched (this time in advance of a request to do so) the exact format and colors of displays, reports, and messages. Figure 5 compares the master schedule report format between 5-th and 7-th Fleets.

Fortunately, although we did not make full use of the organizational change tools at our disposal in Bahrain, we didn't miss the opportunity to do so in Singapore, this time more quickly. In Singapore we gathered key information we had not initially taken into account in Bahrain:

- Early RASP versions had commingled two planning functions in a fashion not aligned with the organization. One organization schedules the CLF ships, while a second plans supplies of commodities.
- Coalition combatant customers, non-US ships about which we have little data, are also supplied via CLF ships. We are therefore obliged to use guesswork to plan their re-supply.
- We had focused on the ship-scheduling function where there is a considerable opportunity for saving money. However, the key incentive for schedulers is to reduce the time required to produce the schedule, and not to save money. While RASP might save millions in fuel, the most important immediate impact for schedulers has been a reported time savings of a couple hours per day.

Adrian Zavala, already involved in the Bahrain work, engaged simultaneously with Singapore. He undertook the temporary increase in his overall workload, and identified desirable embellishments of RASP. We followed his advice closely, and hastened to enhance workflow functionality to automate as much detail as possible, and relax many of the data requirements that seemed to present obstacles to acceptance. We realized that relaxing such details greatly moderated

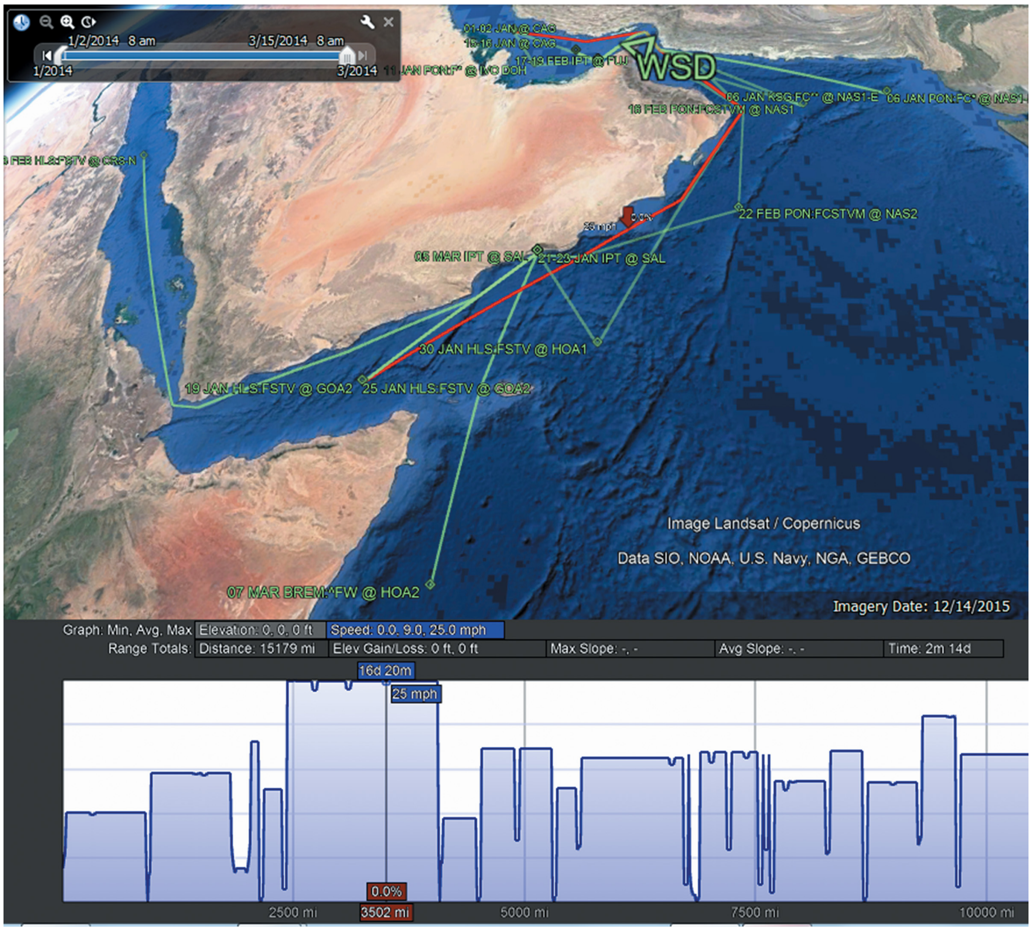


**FIGURE 4** Sixty-day trace over the planning horizon from left to right of the underway speed profile of USNS Walter S. Diehl, T-AO 193, a CLF shuttle ship. Most operations are at reasonably fuel-efficient speeds, but one deployment highlighted in red is at maximum speed 25 mph (21 knots). Presumably there was reason to order this high-speed voyage, motivated perhaps by a change in customer combatant plans, a casualty (mechanical failure) on another shuttle, or perhaps an earlier scheduling oversight [Color figure can be viewed at wileyonlinelibrary.com]

RASP's immediate effectiveness, but that this was a necessary course of action to gain initial acceptance and trust.

7-th Fleet scheduling is complicated by long distances and customers transiting at speeds higher than we can achieve with many of our shuttles. RASP can consider many possible rendezvous locations with fast-moving customers traversing long distances, with the goal of reducing CLF fuel consumption.

RASP went live for 7th Fleet on July 31, 2013.

## 7 | SUSTAINED OPERATIONS AND CONTINUED IMPROVEMENTS

The RASP "Quick Solve" heuristic, as implemented at present, helps schedulers produce schedules that achieve significant savings. However, running a "Full Solve" optimization on a schedule lacking pinned future events (fortunately, this doesn't happen day-by-day) may take hours to complete and requires specialized software. This would provide better schedules, especially in situations when CLF resources are stressed. The much faster Quick Solve option is the go-to operational tool, running directly on the secure scheduler network and taking only minutes to solve, even if it does not furnish a true optimal solution.

### 7.1 | Seeking a fast, approximate schedule

During initial heuristic development we considered several schemes to approximately solve the integer linear program, but finally realized that classic methods primarily focused on the routing problem could not accommodate the myriad side constraints. We concluded it would be a good idea to find out how the schedulers had long been solving these problems by hand. It turns out their heuristic is effective, and mirroring it was a good way to gain their confidence and acceptance. The heuristic details are elaborate and situation-specific and therefore are of scant interest here. Suffice to say we had our best success sorting strike groups by a combination of priority

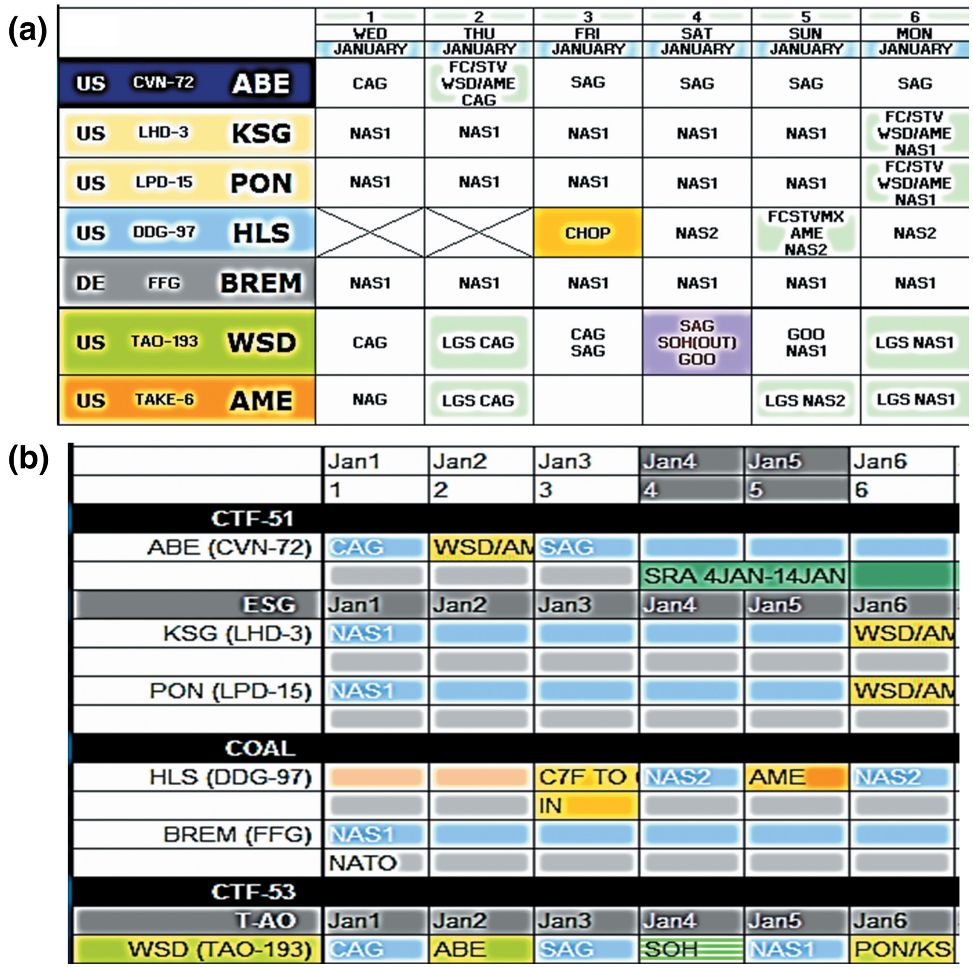


**FIGURE 5** a) 5-th and b) 7-th Fleet views of the same schedule. Each Fleet has its own preference. 7-th prefers a sparser Gantt-like schedule, while 5-th prefers a detailed text grid. RASP produces both, allowing each fleet to view its schedule as the other would. This has opened the door to standardize and adopt the better of each [Color figure can be viewed at wileyonlinelibrary.com]

and demand volume (as schedulers were already doing before our arrival). The heuristic schedules the highest-volume customers first, followed by the others in decreasing priority order. Because there are many customer ships of the same class, it is important to have a tie-breaker for ones with similar volume. Fortunately, this additional prioritization preference is provided directly from how the fleets choose to display the ships in their reports. The more important customers are always on top of others (Rowe, 2016). Along the way, 1-opt opportunities are exploited among shuttles and customers to make minor exchanges that are otherwise eye-catching oversights.

"By its nature, the heuristic works in step-by-step fashion, matching the way a surface router thinks and creates schedules by hand. This makes a heuristic solution more easily confirmed by a router [scheduler] as a 'good solution,' the logic more easily tracked and explained. The Quick Solve does not try to achieve savings by taking a *better* approach than routers, it does so by taking the *same* approach, but applying consistent, cost-saving principles to the decisions made. (Rowe, 2016)" And, we are using exact computations to evaluate options the schedulers could only roughly estimate. Carrier Strike Groups are always the biggest customers, and are faithfully attended by CLF. Strike groups with other big-deck ships follow, and so forth. Recall that day-to-day we usually have few major schedule revisions, and Quick Solve completes these intuitively.

The Quick Solve's very fast solution times have eased schedulers' transition to RASP and have inspired features such as a comparative dashboard described below. While we hope to see RASP's "Full Solve" return as an operational solve option once approved for use on the secure network, more cost reductions are within reach using the heuristic. The speed of the heuristic has changed what is possible to do at operational tempo. Some of these impacts are detailed in the following sections.

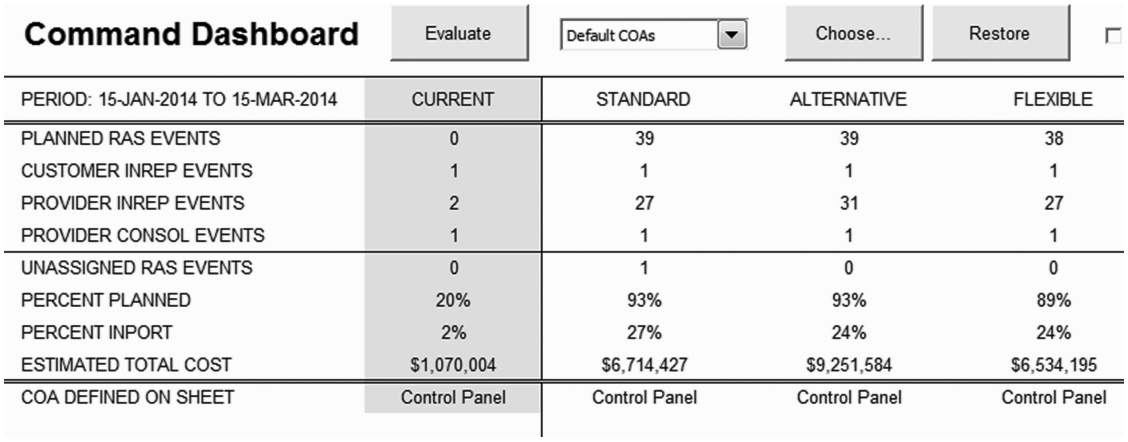

Command Dashboard | Evaluate | Default COAs | Choose... | Restore

| PERIOD: 15-JAN-2014 TO 15-MAR-2014 | CURRENT | STANDARD | ALTERNATIVE | FLEXIBLE |
|---|---|---|---|---|
| PLANNED RAS EVENTS | 0 | 39 | 39 | 38 |
| CUSTOMER INREP EVENTS | 1 | 1 | 1 | 1 |
| PROVIDER INREP EVENTS | 2 | 27 | 31 | 27 |
| PROVIDER CONSOL EVENTS | 1 | 1 | 1 | 1 |
| UNASSIGNED RAS EVENTS | 0 | 1 | 0 | 0 |
| PERCENT PLANNED | 20% | 93% | 93% | 89% |
| PERCENT INPORT | 2% | 27% | 24% | 24% |
| ESTIMATED TOTAL COST | $1,070,004 | $6,714,427 | $9,251,584 | $6,534,195 |
| COA DEFINED ON SHEET | Control Panel | Control Panel | Control Panel | Control Panel |

**FIGURE 6** RASP's Command Dashboard allows schedulers to create, name, and compare competing parallel candidate schedules. The scenarios shown here illustrate what is possible by changing a few top-level solver constraints and penalties. Rows tally the number of RAS events scheduled, combatant port calls, shuttle port calls, and shuttle-to-shuttle "consol" cargo transfers, along with required RAS events not scheduled, percent completion of a schedule, percent in-port shuttle days, and estimated total schedule cost. The 'CURRENT' plan (shown at left in gray) is the common solver starting point. It contains only a few fixed events. Each of the remaining columns summarizes a completed schedule for a named set of scheduling parameters. The 'STANDARD' schedule is the most-constrained to encourage the highest quality of service. In this example, the 'STANDARD' level is not fully achievable and one of the RAS events is not fulfilled. The 'ALTERNATIVE' schedule relaxes constraints to the point where the problematic RAS event becomes assignable; however this comes at a huge increase in cost. The 'FLEXIBLE' schedule relaxes constraints even more, allowing the event to be assigned at a cost similar to the 'STANDARD' plan, but with a lower service quality across the fleet [Color figure can be viewed at wileyonlinelibrary.com]

## 7.2 | A comparative planning dashboard

RASP has been outfitted with a "Command Dashboard" with which a scheduler can develop multiple, parallel schedules and compare the level of service and cost of each. Figure 6 illustrates a seed initial schedule and several solver-completed ones.

Given possible courses of action (COAs), as shown in Figure 6, Navy decision makers can review the details and compromises of each plan before deciding what is best. Alerted ahead of time of a problematic RAS event, as shown in this example, a Navy command might choose to make arrangements for a customer ship to self-replenish in port. The dashboard comparisons can help justify any costs associated with this self-replenishment.

## 7.3 | Shuttle fuel burn rate profiling and awareness

Figure 7 shows an aggregate view of voyage fuel consumption by a T-AOE (a CLF shuttle capable of relatively high underway speed) for a RASP COA. We'd prefer to travel at slower fuel-efficient speeds, but faster speed may be necessary if we have too few shuttles to meet too many RAS evolutions. Speed allows more to be done with a shuttle, but as you can see in this figure, speed exacts a super-linear cost penalty in shuttle fuel consumption (refer back to Figure 4).

## 7.4 | Extending the time horizon

RASP can generate schedules quickly and it completely automates scheduling reports and messages. This has enabled schedulers to look our farther. The only limitation on the duration of planning horizon now is the availability of reliable customer data. RASP has enabled 5-th fleet to extend its typical planning horizon from 45 days to 90 days, while 7-th fleet now plans out its notional schedules typically up to 360 days in advance, and sometimes even farther for long-term planning studies.

## 7.5 | Cross-fleet unified planning

RASP pilot studies have indicated additional cost savings are possible by better collaborating and planning between the fleets, making better use of existing data such as the moment and location when each combatant ship will enter or leave (ie, in Navy vernacular, "cross the chop line" into or out of) an AOR of each fleet as well as the state of her stores of fuel, ordinance and other supplies. CLF ships can then be positioned advantageously so as to serve these arriving and departing combatant customers. Tests with historic operational data on an extended time horizon across fleets have shown that RASP can provide additional cost savings for CLF. The increased problem size and complexity of a cross-fleet, or even global, scheduling problem would have been too daunting to take on before RASP. Now fleets can start to look more beyond their

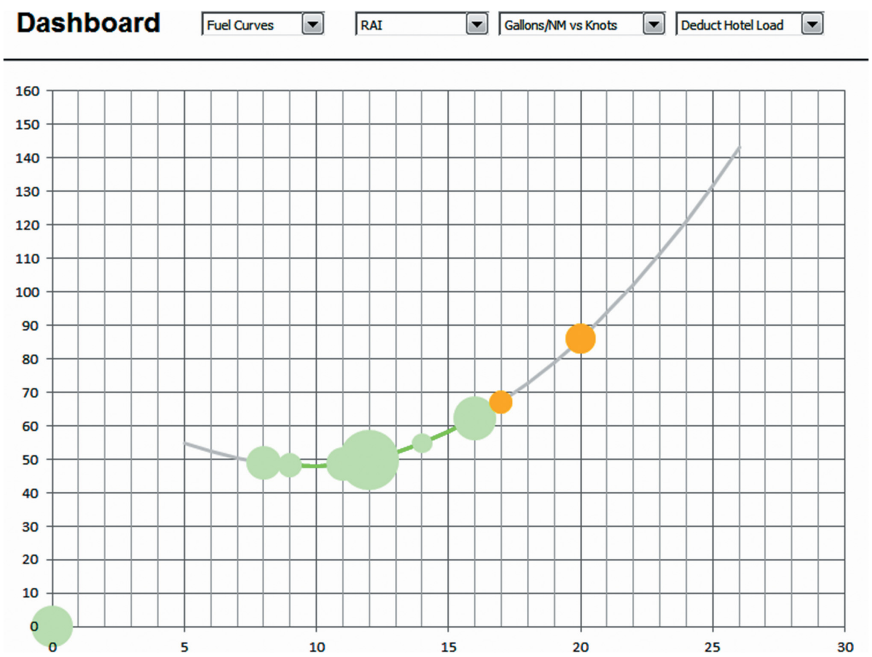


**FIGURE 7** Bubble chart fuel consumption curve: This curve shows fuel consumption (in gallons per nautical mile) as a function of underway speed (in knots) for a CLF Fast Combat Support Ship (T-AOE). Fuel consumption for a chosen voyage plan is shown by the circles along the curve. The area of each circle on the curve is proportionate to the number of hours spent at that speed. A speed of 10 knots is the most economical, but in this plan significant time is spent at higher, more costly speeds [Color figure can be viewed at wileyonlinelibrary.com]

own borders, and multi-fleet operations better planned. The automated Quick Solve and reporting capabilities of RASP puts this augmented problem well within the range of what is possible.

## 8 | INFORMING POLICY CHANGES

RASP has enabled cost-based exploration of many existing navy policies, including the policy on the fleet "reserve fuel level" (safety stock). Using RASP, we have found that this policy has a dramatic influence on CLF costs. Suppose for some combatant customer that $S$ is the safety stock policy, expressed as a percentage of its fuel capacity, meaning that $1-S/100$ is the useable fraction of ship fuel capacity. This dictates that the "hit-rate" at which CLF must visit this combatant is proportional to $H(S) = 1/(1 - S/100)$. Changing the safety stock from $S$ to $S'$ has the following percentage influence on this hit rate: $100\,[H(S') - H(S)]/H(S)$. For example, if the safety stock requirement $S = 60\%$ is relaxed to $S' = 50\%$, this reduces the CLF hit rate by 20%. Regardless of combatant fuel capacity or rate of fuel use, if the combatant operations are independent of those of the CLF ships providing fuel, this 20% reduction applies directly to reduction of CLF fuel consumption. Not only does RASP allow costing of specific safety stock policies, it is easy for schedulers to implement a new policy. Reducing CLF shuttle fuel consumption not only lowers cost to deliver fuel to combatant customers, it means less overall cost to ensure the combatants stay on station. RASP is intended to discover efficiencies and evaluate innovative policies in search of such savings.

RASP has also been used to evaluate changes to CLF fleet composition, commodity storage locations, and port availabilities, both in the context of intentional policy changes and potential fragilities. All-in-all a list of about 100 alternate scenarios has been used to inform and stress-test Navy logistics policies.

## 9 | DEPLOYING TO A SECURE COMPUTING ENVIRONMENT

All operational data is classified, and can only be manipulated on classified computers and networks. Although these systems feature Microsoft Office © and Google Earth © tools, approved software for the fleet schedulers does not include a viable, efficient computational language, mathematical modeling package or large-scale optimization software. In January 2011 we initially provided a portable workstation not connected to a classified network, and asked schedulers to transfer data to this workstation via CD, and copy schedules back to the classified network where other reports had to be completed. This was unsatisfactory.

Faced with no alternative, we converted everything to Visual Basic for Applications (VBA) (Microsoft, 2017). In our experience, interpreted VBA is about two orders of magnitude slower than a compiled language such as C or FORTRAN (eg, INTEL, 2017), but these faster tools are not approved for use by fleet schedulers. We still have the ability to generate and solve our formal optimization model outside the classified bubble, but this is only useful for independently checking the performance of the heuristic. The heuristic schedules are produced fast enough to satisfy schedulers, who can quickly polish any rough details.

Due to all these restrictions, RASP was entirely converted to maximally comply by adopting Microsoft Office with Google Earth animations of schedules added with color highlights when any CLF shuttle ship is scheduled to travel faster than its most fuel-efficient underway speed. These developments have been a big attraction for schedulers and their leadership.

We haven't given up on getting more-powerful mathematical modeling and optimization software approved for use with RASP. This part of our story is not over.

## ACKNOWLEDGMENTS

Portions of this work have been supported by Office of Naval Research and Military Sealift Command. Rob Dell poured over several drafts and made many constructive suggestions. Our faculty has a billet for a military faculty member of Operations Research and manager of our Operational Logistics curriculum and students. The tenants have included Commander Pat Burson (now Captain), Commander Walt DeGrange (an author here, now retired) and Commander Peter Ward. This succession of experienced, seafaring supply experts has marshalled their combined wisdom with that of our students to provide a deep pool of domain expertise. Adrian Zavala is now a civilian contractor supporting RASP worldwide; without his sustained support and the continuity he provides, we would have recurring problems with personnel turnover. Adrian is also a prodigious experimenter, and we are delighted with his continuing probes and insights. There have been and continue to be independent audits of RASP results: ***RASP* works**. We are also pleased that RASP is attracting some favorable press (e.g., DOD 2015 and USN 2015).

## ORCID

*Gerald G. Brown* http://orcid.org/0000-0002-2974-7162

## REFERENCES

Bretton Woods Agreements Act, section 14. (1945). http://avalon.law.yale.edu/20th_century/decad047.asp, (July 23, 1945) accessed August 25, 2017.

Brown, G. G., & Carlyle, W. M. (2008). Optimizing the US Navy's combat logistics force. *Naval Research Logistics*, *55*, 800–810.

Brown, G. G., Carlyle, W. M., Dell, R. F., & Brau, J. W. (2013). Optimizing Intratheater Military Airlift in Iraq and Afghanistan. *Military Operations Research*, *18*(3), 35–52.

Brown, G. G., Dell, R. F., & Farmer, R. (1996). Scheduling coast guard district cutters. *Interfaces*, *26*, 59–72.

Brown, G. G., Dell, R. F., & Wood, R. K. (1997). Optimization and persistence. *Interfaces*, *27*, 15–37.

Brown, G. G., & Washburn, A. R. (2017). Method and system for determining shortest oceanic routes, avoiding obstacles. *U.S. Patent 9,541,401*, January 10.

Chambers, J.W. (Ed.). (1999). *American military history*. New York: Oxford University Press.

Christiansen, M., Fagerholt, K., & Ronen, D. (2004). Ship routing and scheduling: Status and perspectives. *Transportation Science*, *38*(1), 1–18.

Christiansen, M., Fagerholt, K., Nygreen, B., & Ronen, D. (2013). Ship routing and scheduling in the new millennium. *European Journal of Operations Research*, *228*(3), 467–483.

Dantzig, G. B., & Fulkerson, D. R. (1954). Minimizing the number of tankers to meet a fixed schedule. *Naval Research Logistics Quarterly*, *1*, 217–222.

DeGrange, W. (2012). Why won't you use my model? Difficulties in implementing optimal scheduling models, new horizons: The role of operations research in planning DoD's future. In *3RD Annual Military Applications Society Conference*, Monterey, California, March 27–28.

Garfinkel, R., & Nemhauser, G. (1972). *Integer programming*. New York, NY: Wiley.

INTEL Corporation, Visual Studio. (2017). https://software.intel.com/en-us/intel-parallel-studio-xe-compilers-required-microsoft-visual-studio (accessed September 10).

International Business Machines (IBM). (2017). CPLEX. https://www.ibm.com/analytics/data-science/prescriptive-analytics/cplex-optimizer.

International Chamber of Shipping (ICS). (2018). Shipping and World Trade. https://www.ics-shipping.org.

Karafiath, G., McCallum, D., & Hendrix, D. (2004). Designing navy hull forms for fuel economy. *Naval Engineers Journal*, *116*(2), 43–54.

Korsvik, J., Fagerholt, K., & Laporte, G. (2010). A tabu search heuristic for ship routing and scheduling. *Journal of the Operational Research Society*, *61*(4), 594–603.

Korsvik, J., Fagerholt, K., & Laporte, G. (2011). A large neighbourhood search heuristic for ship routing and scheduling with split loads. *Computers & Operations Research*, *38*(2), 474–483.

Lawler, E.L., Lenstra, J.K., Rinnooy Kan, A.H.G., & Shmoys, D.B. (Eds). (1987). *The Traveling Salesman Problem*. New York, NY: Wiley.

Meng, Q., Wang, S., Andersson, H., & Thun, K. (2014). Containership routing and scheduling in liner shipping: Overview and future research directions. *Transportation Science*, *48*(2), 265–280.

Microsoft, Visual Basic for Applications. (2017). https://msdn.microsoft.com/en-us/vba/vba.

Nemhauser, G., & Wolsey, L. (1988). *Integer and combinatorial optimization*. New York, NY: Wiley.

Office of Secretary of Defense (OSD). (2016). http://www.acq.osd.mil/eie/Downloads/OE/FY16%20OE%20Budget%20Certification_07_21_16.pdf.

Ronen, D. (1983). Cargo ships routing and scheduling: Survey of models and problems. *European Journal of Operational Research*, *12*(2), 119–126.

Ronen, D. (1993). Ship scheduling: The last decade. *European Journal of Operational Research*, *71*(3), 325–333.

Toth, P. & Vigo, D. (eds.). (2002). *The Vehicle Routing Problem*. Philadelphia: SIAM.

Rowe, A. A. (2016). RASP Software Design Description, white paper, Operations Research Department, Naval Postgraduate School, 23 pages.

U.S. Department of Defense (DOD). (2014). *Operational Energy Annual Report*. Oct 15 http://energy.defense.gov/Portals/25/Documents/Reports/FY13%20OE%20Annual%20Report.pdf accessed 4 Oct 2015.

U.S. Navy (USN). (2015). Replenishment at Sea Planner, Sealift, Military Sealift Command, Washington, DC, http://www.msc.navy.mil/sealift/2013/October/energy.htm accessed 1 October.

Washburn, A. R., & Brown, G. G. (2016). An exact method for finding shortest routes on a sphere, avoiding obstacles. *Naval Research Logistics*, *63*(5), 374–385.

Wray, R. O. (2010). Military Sealift Command Optimization Initiative, private communication of author Brown with RDML USN, Deputy Commander, Military Sealift Command, Washington, DC, February 12.

**How to cite this article:** Brown GG, DeGrange WC, Price WL, Rowe AA. Scheduling combat logistics force replenishments at sea for the US Navy. *Naval Research Logistics*. 2018;00:000–000. https://doi.org/10.1002/nav.21780