

2010

Open architecture, inventory pooling and maintenance modules

Geraldo Ferrer

Naval Postgraduate School, gferrer@nps.edu

Follow this and additional works at: <http://digitalcommons.unl.edu/usnavyresearch>

Ferrer, Geraldo, "Open architecture, inventory pooling and maintenance modules" (2010). *U.S. Navy Research*. 32.
<http://digitalcommons.unl.edu/usnavyresearch/32>

This Article is brought to you for free and open access by the U.S. Department of Defense at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in U.S. Navy Research by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.



Open architecture, inventory pooling and maintenance modules

Geraldo Ferrer*

Naval Postgraduate School, IN-234 555 Dyer Rd., Monterey, CA 93943, USA

ARTICLE INFO

Article history:

Received 22 May 2009

Accepted 21 July 2010

Available online 27 July 2010

Keywords:

Open architecture
Inventory management
Pooling effects
Distribution network
Ardalan heuristics
Case study

ABSTRACT

The adoption of open architecture has several economic implications in the life of an asset, including developmental, production, storage, training and maintenance costs. This research responds to an inquiry by the Program Executive Officer—Integrated Weapons System (US Department of the Navy) regarding the value of open architecture (OA) in the design of complex assets. With this intent, we evaluate how the inventory allocation of spare engines for the F-16 operations in the continental United States would be affected with and without the adoption of open architecture, focusing on the benefits of inventory pooling to meet the demand of many users from a small number of storage sites. We use a distance-constrained version of the Ardalan heuristic for solving the facility location problem, responding to practical limitations exposed by the model. This article shows that open architecture may provide substantial supply chain cost reduction, and simplification of the distribution network when combined with proper inventory storage policies.

Published by Elsevier B.V.

1. Introduction

The combined use of commonality and modularity in product design has allowed automobiles, aircraft, computers and a host of other machines (including most military systems) to be reusable beyond their first lifecycle and to be given many more years of operation. This versatility substantially impacts the availability and maintenance cost of many durable assets. Modularity enables the division of the product development effort among many specialists (firms or individuals), ensuring the development of the most advanced and competitive systems. Modularity facilitates the separation of component-wear phenomena as the system ages, enabling maintenance professionals to locate and repair damaged modules without affecting the integrity of other modules in the system.

Commonality, however, presents a disadvantage that many engineers will recognize: the adoption of common design in a competitive environment hinders creativity and innovation in product development; suppliers of high-technology products would prefer to develop their own designs than to share them with competitors. The design team would rather showcase its capabilities, especially in the development of advanced systems or in the adoption of new technologies. Hence, while modularity remains a powerful product-development philosophy that brings agility and cost reduction to product design, the adoption of common designs for complex products may be not the best approach to system acquisition—especially in circumstances

requiring the development of advanced technologies. In these scenarios, the traditional “commonality” must be enhanced with the adoption of “open architecture” features—allowing modules from competing sources to be used in the same system, without constraining the creativity and innovation from the designers involved in the development of the module.

Open architecture provides the opportunity to introduce product aggregation, one of the three aggregation (or *pooling*) approaches to managing and improving supply-chain performance, along with time aggregation and place aggregation. Product aggregation is intended to reduce product variety without compromising the functionality required by the user.

The purpose of this case study is to evaluate open architecture as the design philosophy for the acquisition of complex systems with advanced technologies. This is done by analyzing the case of the F-16 spare engines, showcasing the cost benefits that the US Air Force might be enjoying today had the aircraft engine suppliers been required to adopt open architecture. This study assumes that a complex system (such as the Joint Strike Fighter, or other weapon systems acquired by the uniformed services of the US Department of Defense) is a combination of hardware and software components that may be acquired from multiple developers or suppliers. This study proposes that the adoption of open architecture in the acquisition of these systems can substantially reduce the costs of these programs.

Next section describes the problem that motivated this study, and Section 3 explains how open architecture affects product development and life-cycle management. Section 4 describes current inventory management policy and allocation of F-16 spare engines at Air Force bases in the contiguous United States. Section 5 introduces a brief literature review of the methodology used to

* Tel.: +1 831 656 3290; fax: +1 831 656 3407.
E-mail address: gferrer@nps.edu

rationalize F-16 spare engine allocation. Section 6 analyzes the case under three scenarios: the benchmark scenario, which is based on current policy, a scenario with limited inventory pooling, and a third scenario with open architecture of engine design. Section 7 presents the conclusions and suggestions for future research.

2. Motivation: F-16 alternative engines

Modularity facilitates the development of new systems using modules that were previously designed and developed for other systems, providing major time and cost savings product development initiative that can exploit these benefits. Moreover, because of commonality, high-value modules in a system may be recovered at the end of the system's life and used in another product—a process often called cannibalization. In the case of complex models with high engineering content, such as aircraft turbines, it may be desirable to have multiple suppliers offering competing designs, which could be accomplished with the adoption of open architecture. The [Defense Acquisition University \(2006\)](#) defines open architecture as follows:

The confluence of business and technical practices yielding modular, interoperable systems that adhere to open standards with published interfaces. This approach significantly increases opportunities for innovation and competition, enables reuse of components, facilitates rapid technology insertion, and reduces maintenance constraints.

Modularity and commonality are the two aspects in product design that support the adoption of an open architecture. They facilitate the execution of an agile product development program with a wide-reaching product line that meets the requirements of multiple users with different needs. The renewed emphasis on open architecture allows strategic resource allocation, facilitating the acquisition of better assets with lower costs. A current example of this design approach is the F-35 Lightning II, Joint Strike Fighter, a multi-role aircraft currently in production for the uniformed services of the US Department of Defense (DoD) and for many of the US allies. The [Federation of American Scientists \(2005a\)](#) indicates that among its strengths, “JSF ... will capitalize on commonality and modularity to maximize affordability.”

In practice, previous developments and acquisitions of weapon systems by the DoD usually did not have this focus. For instance, Pratt & Whitney (P&W) and General Electric Aircraft Engines (GEAE) produce engines for the F-16 aircraft used by the US Air Force and a few foreign military forces. The P&W F100-PW-200 aircraft engine was originally selected over GEAEs as the sole source engine for the F-16. The original F-16 was designed as a lightweight, air-to-air day-fighter. Air-to-ground responsibilities transformed the first production F-16s into multi-role fighters. The first operational F-16A was delivered in January 1979 to the 388th Tactical Fighter Wing at Hill Air Force Base, Utah. The delivery of 2200+ aircraft to the US Air Force continued until March 2001 ([Federation of American Scientists, 2005b](#)).

The decision to choose an alternate fighter engine for the F-16 led to the development of the General Electric Aviation Engine's F110 series. With the implementation of the Alternative Fighter Engine competition for the F-16 in 1985, GEAE fielded the F110-GE-100 version to compete with Pratt & Whitney's F100-PW-220 engine. Throughout the production of the F-16, the performance requirements for both suppliers were identical, but the engines delivered were not interchangeable. In fact, the airframe manufacturer, Lockheed-Martin, had to deliver structurally different frames to use the different engines. For example, aircraft with

production numbers ending in zero are designed and built with significantly larger air intake to accept the GEAE F110 series engine. Aircraft with production numbers ending in two are designed and built with smaller air intake to use the P&W F100 series engine. Each engine type (GEAE or P&W) uses different control software (with implications in the cockpit controls and pilot training), requiring unique airframe interface. With the exception of the engine, the airframe interface and the control software, aircraft of the same generation would otherwise be identical.

The adoption of two engine suppliers for the F-16 fighter aircraft was intended to eliminate the monopoly held by Pratt & Whitney as the sole-source engine supplier for that aircraft. However, allowing the newcomer (GEAE) to design a product that was not interchangeable with the existing engine did not eliminate some of the monopoly effects in the long-term, and created costly logistics constraints.

Similar to the F-16 acquisition experience in the 1980s and the 1990s, the ongoing acquisition process of the Joint Strike Fighter includes the development of two competing power plants: the Pratt & Whitney F135, and the GEAE F136, developed in partnership with Rolls-Royce. On its website, the Federation of American Scientists states that the F-35 propulsion systems will be “physically and functionally interchangeable in both the aircraft and support systems.” According to the Joint Strike Fighter Program Office, “the F135 and F136 teams are working closely to develop common propulsion system components” ([F-35 JSF Program, 2007](#)).

In this study, we analyze current usage data of P&W and GEAE spare engines held in various bases in the continental United States to support the F-16 operations to identify substantial cost reduction from the pooling effects that could be achieved with the use of better inventory allocation (place aggregation), as well as the adoption of open architecture (product aggregation). One important caveat exists, however: considering the limited amount of usable data available about the acquisition and use of these aircraft, the reader is cautioned that this analysis is not a critique of the acquisition of the F-16 aircraft or its engines. Rather, it intends to discuss how it would have benefited had it adopted open architecture.

3. Open architecture as a design approach to simplify the supply chain

This section presents open architecture and how it generally benefits product design. The concept stems from the development approach used by many software houses, in which sub-routines (modules) are developed by individual designers having only two major constraints: the functionality (i.e., the sub-routine does what is expected to do) and the standardized interface with the main program (i.e., the sub-routine has seamless integration with other modules in the software).

[Nelson \(2007\)](#) indicates that open architecture principles have been around since at least 1981, when IBM developed its personal computer. The design of the IBM-PC was a major breakthrough in that it was made of a set of physical modules that could be replaced by similar modules of different design, make or performance, as long as they satisfied a limited set of interface requirements and fulfilled the expected functions. For example, a hard disk drive of a given capacity and make could be upgraded by another hard disk of different make and greater capacity, as long as it satisfied a simple set of interface constraints. By contrast, one is not usually able to replace the engine of an automobile by one from a different maker, even if the two have similar performance, size or functionality.

The open architecture design philosophy was extremely successful for desktop computers, and it still describes most desktop computers built almost 30 years later. In contrast, proprietary designs have led to expensive and less successful products in the computer industry—such as the computer Amiga that preceded the IBM-PC, the short-lived Unix desktop, and various generations of the Macintosh desktop. With the exception of the IBM-PC, the adoption of open architecture in computer hardware design is limited. Space and weight limitations have restricted the use of open architecture in the design of laptop computers. Hence, internal components developed for one particular laptop usually cannot be used in a different model or brand. Open architecture benefits have been usually restricted to the interfaces with external accessories and, in some cases, to memory units.

It is important not to confuse open architecture with “open source” (Coar, 2006). Software developed under an open source philosophy is copyright-free and can be modified and extended by any other software writer, as exemplified by the Linux operating system and the Mozilla web browser. Nonetheless, to enable continued expansion, open source software usually adopts open architecture as the means to ensure a compatible interface between the works of multiple authors.

In 2006, the US Navy released the *Open Architecture Contract Notebook*, explicating the open architecture guidelines to be adopted by Acquisition Officers (PEO-IWS, 2006). Specifically, it is recommended that contracts include this statement: “The Contractor will be required to define, document, and follow an open systems approach for using modular design, standards-based interfaces, and widely supported consensus-based standards” (p. 7). While these recommendations usually target software design, they can be quite useful in the design and acquisition of all complex hardware. The adoption of open architecture principles in hardware design provide some of the same benefits found in software design, in addition to the following:

- 1 Simplified maintenance: the modularity found in open architecture products makes it easier to remove, replace and repair damaged modules with minimal impact to the whole system.
- 2 Simplified logistics: open architecture enables the use of modules by different makes, or even different generations, if they maintain the same interface standards.
- 3 Reduced acquisition cost: open architecture allows a true competition between potential suppliers in all phases of the lifecycle of the product, requiring just that each potential supplier adopt the standard module interfaces.

These benefits become more critical when we realize that many complex assets depend on the successful integration of multiple hardware and software modules. The determination of standard interfaces between modules allows substantial savings in the operation and maintenance of weapon systems, as illustrated by the F-16 aircraft engine. In some cases, an engine or engine components (modules) may be removed and transported to Tinker Air Force Base (AFB), the maintenance depot for F-16 engines, without any maintenance action by local technicians. In other cases, local technicians may be capable of performing the required maintenance action locally and returning the engine to operable or Ready for Issue (RFI) status.

The managers of active duty air force bases aggressively track the status of engine changes. They expect turnaround of less than 24 hours from each engine change operation, which requires keeping a certain inventory of spare engines readily available. This culture seems to contrast with Air National Guard (ANG) and

Air Force Reserve (AFR) units, in which the F-16 aircraft are used less intensively. Guard and Reserve units typically have fewer assigned aircraft and, therefore, have a lower spare engine stock. However, given that their primary mission is the defense of the national air space, they too can benefit from having a short engine maintenance turnaround.

The US Air Force uses the F-16 in 30 bases of various sizes, including Active Duty (AD), Air Force Reserve (AFR) and Air National Guard (ANG). Each base has its own stock of spare engines to meet demand. Some of the bases use aircraft with GE engines; others use aircraft with P&W engines. As explained earlier, bases do not use engines of different make in their fleets because they are not interchangeable in any way. The most notable differences associated with the two power plants are as follows:

- 1 Airframes are structurally different, with a distinct engine bay for each engine make.
- 2 Engines have different durability and reliability, leading to distinct preventive maintenance needs.
- 3 Repair parts, maintenance jigs and tools are different.
- 4 The software that controls engine performance and interprets the pilot’s command from the cockpit is different.
- 5 Aircraft using different engines respond differently to the pilot’s commands. This mandates a non-trivial period of adaptation when a pilot switches from one aircraft type to the other.

In short, because of different design choices made by the engine manufacturers, we have effectively two distinct aircraft types in service under the codename F-16. This creates undesirable limitations in the way aircraft and engines are used and maintained. The open architecture design approach would effectively eliminate many of the differences between the two engines, without constraining the creativity and flexibility of the design engineer.

4. Case description: F-16 spare engine storage

This section describes the inventory management of the F-16 spare engines, as practiced by US Air Force bases using this aircraft. The expected demand for spare engines was 656 P&W engines and 773 GEAE engines in 2008, reflecting a negative trend of approximately 5.8% per year since 2001. Demand originates in 13 bases using Pratt & Whitney-powered aircraft, and 18 bases using General Electric-powered aircraft in the Continental United States. In general, these bases hold a total pre-positioned inventory of 159 spare engines, turning the inventory fewer than 9 times per year. Table 1 shows the historic demand in each base, next to the International Civil Aviation Organization (ICAO) code of the respective airfield. Based on a simple linear regression of the 8-year demand in each of the 31 bases, the forecasted demand for year 2008 is also shown in the table. As we can observe, approximately half of the forecasted demand is fragmented across 24 Air National Guard (ANG) bases, and the remainder is distributed in four Active Duty (AD) bases; a small demand is generated in two Air Force Reserve (AFR) bases. Each base has different capabilities to provide engine maintenance, with all the complexity that such maintenance entails. In general, the ADs have the support personnel and equipment to give some service, while the ANGs and AFRs have limited maintenance infrastructure. Notice that four ANG bases are co-located with AD bases (Andrews AFB (DC), Kelly AFB (NM), Buckley AFB (CO) and Kirtland AFB (NM)), where valuable synergies regarding engine maintenance can be expected.

Table 1

Historic demand of F100 and F110 engines.

Sources: Historic data adapted from Henderson and Higer (2007) and <http://www.f-16.net>. Forecast by the author.

Pratt & Whitney (F100 engine)										
Base	ICAO	'01	'02	'03	'04	'05	'06	'07	2008 forecast	
									Mean	Std err
ANG-Burlington VT	BTV	52	39	51	48	43	31	40	34.9	6.5
ANG-Duluth MN	DLH	49	52	56	38	30	27	33	23.0	7.1
ANG-Des Moines IA	DSM	40	37	32	28	12	25	13	8.9	5.7
ANG-Ellington TX	EFD	38	53	42	38	39	33	37	33.4	5.7
ANG-Fresno CA	FAT	57	35	41	45	36	34	55	41.4	10.3
ANG-Ft Smith AR	FSM	34	33	27	36	39	29	5	17.1	10.2
ANG-Ft Wayne IN	FWA	38	37	33	47	23	28	29	25.7	7.3
Hill AFB-Depot UT	HIF	56	71	64	54	67	67	67	67.7	6.5
Nellis AFB NV	LSV	88	57	92	90	97	66	60	69.9	17.7
Luke AFB AZ	LUF	365	355	344	327	338	273	250	248.1	19.1
Tinker AFB OK	TIK	0	0	0	0	0	0	0	0.0	0.0
ANG-Toledo OH	TOL	44	35	24	24	13	18	7	1.3	4.4
ANG-Tulsa OK	TUL	43	34	21	22	21	18	9	4.9	4.5
ANG-Tucson AZ	TUS	209	160	176	153	126	102	107	80.1	14.1
General Electric (F110 engine)										
Base	ICAO	'01	'02	'03	'04	'05	'06	'07	2008 forecast	
									Mean	Std Err
ANG-Albuquerque NM	ABQ	48	40	45	39	27	38	37	31.3	5.6
ANG-Eggharbor NJ	ACY	44	53	51	53	38	20	22	19.4	9.5
ANG-Andrews DC	ADW	33	39	26	21	34	29	26	25.0	6.0
ANG-Buckley CO	BKF	40	33	48	39	33	28	29	27.4	6.0
Cannon AFB NM	CVS	175	153	126	102	93	99	65	48.9	11.3
ANG-Sioux SD	FSD	42	60	49	42	43	45	42	41.0	6.6
ANG-Gt Falls MT	GTF	33	45	47	55	46	53	29	44.4	10.6
Hill AFB UT	HIF	245	201	225	236	202	230	220	217.0	17.8
AFR-Homestead FL	HST	40	31	41	32	34	26	33	28.4	4.7
ANG-Montgomery AL	MGM	39	30	50	46	26	45	34	37.3	9.8
ANG-Madison WI	MSN	45	33	33	33	28	25	36	26.4	5.6
ANG-Selfridge MI	MTC	65	45	47	51	37	43	33	30.1	6.6
AFR-Ft Worth TX	NFW	47	40	41	20	38	26	27	21.1	7.5
ANG-Richmond VA	RIC	36	42	37	36	20	25	6	8.7	7.0
ANG-Springfield OH	SGH	53	64	63	57	43	54	56	51.3	7.2
ANG-Kelly TX	SKF	61	49	46	56	50	49	46	45.1	5.0
ANG-Springfield IL	SPI	33	19	24	27	23	31	34	31.0	5.7
ANG-Syracuse NY	SYR	0	0	31	38	35	33	39	38.5	3.3
Tinker AFB OK	TIK	0	0	0	0	0	0	0	0.0	0.0

To prevent shortage, which would affect the readiness of the respective base, a base-stock inventory management policy is adopted such that a prescribed level of inventory is kept at each base. The Oklahoma City Air Logistics Center, located at Tinker AFB, provides “supply chain management, including acquisition, repair, storage, distribution, disposal and the technical and engineering services for the center’s assigned engines,” which include major maintenance activities for the F100 and F110 engines (Tinker AFB, 2007). This depot is conveniently located in the center of the country, but only seven bases are within 1-day driving range (approximately 550 miles)—an important consideration since managers expect to maintain the base stock at all times. Traveling time to other bases is as long as 3 days. Hence, the lead time for an order placed from each base is typically between 5 and 7 business days, depending on the distance to the customer and provided that Tinker has the engine in stock ready for issue.

The US Air Force propulsion requirements determine the spare engine inventory level, adding a safety stock based on the demand variability and on the service level associated with the user’s priority. This service level depends on the primary assignment for

each location: either combat (80% service level), or training (70% service level). All F-16 users in this study are considered combat units—except those located at Luke AFB (AZ), a training base (Henderson and Higer, 2007). The demand variability is caused by two random variables that regulate the queuing system at the Tinker AFB maintenance depot:

- 1 Number of hours flown per month: this drives the actual demand seen at the depot.
- 2 Maintenance service time: this drives the waiting time until an engine can be serviced.

The asset utilization randomness drives the need for a safety stock. Considering the forecasted demand for 2008, and the unit prices of \$3.27M (P&W) and \$2.95M (GEAE), the recommended safety stock in all bases is worth \$123.9M, as shown in Table 2. This base stock policy meets the forecasted demand according to the current practice of each base keeping its own inventory. The difference between the base stock and the safety stock ($86 - 40 = 46$) is the sum of the expected lead-time demand in each site.

Table 2
Inventory distribution according to category and make

Inventory status: each base stores its own replacement engines			
	P&W	GEAE	Total
Annual demand	656	773	1429
Base stock	38	48	86
Safety stock	18	22	40
Safety stock value	\$58.9	\$65.0	\$123.9

5. Warehouse location and related literature

As discussed in the previous section, the USAF currently manages the inventory of spare engines without taking advantage of any pooling benefit. An alternative approach would be to hold the inventory in select bases within acceptable distance. The selection of sites that should hold inventory is characterized as the *p*-median problem; this section presents some of its early literature.

The *p*-median problem can be described as a process for identifying *p* facilities to serve customers in *n* nodes of a network ($p < n$). Finding a limited number of Air Force bases to store spare engines for all bases in the Continental United States would be a *p*-median problem with special characteristics, which we describe later. The heuristics to solve the *p*-median problem usually strive to find the minimum transportation time, given the respective demands. As Reese (2005) described in his bibliography, there are four primary approaches to the facility location problem, and the *p*-median is one of them, along with the *p*-center, the incapacitated facility location and the quadratic assignment problems. The *p*-median problem is a location–allocation problem, because it finds the supply nodes and allocates the customers that should be served by each of them.

The unconstrained *p*-median problem seems to have been first described by Hakimi (1965). Toregas et al. (1971) formulated the constrained problem as a mixed integer program, and many researchers have attempted to solve both the constrained and unconstrained problems using either approximate heuristics, or exact branch-and-bound, or Lagrangean relaxation algorithms. A large number of heuristics were developed to solve specific variations of the problem; we focus on four of them.

Kuehn and Hamburger (1963) proposed a greedy heuristics and tested it using a simple data set that became standard to test other solutions. This was later described as an *add* heuristics, because it assigns facilities to locations successively, as they improve market coverage and reduce transportation cost. Teitz and Bart (1968) proposed a *replacement* heuristics: from an initial arbitrary allocation of *p* facilities, the remaining locations are successively tested as replacements for elements in the original set if they improve performance, or they are disregarded if they don't. This heuristics does not have distance constraint capability; rather, the manager has to find the solution with increasing number of nodes until all facilities are served from a node within the distance constraint. Khumawala (1972) proposed two greedy heuristics that he called “delta” and “omega”, both of them incorporating distance constraints. Delta is an *elimination* heuristics that starts with the whole pool of locations, and gradually eliminates, one at a time, those that cost the least to shut down, until none can be eliminated without violating the distance constraint. Omega is an *add* heuristic that gradually selects from the pool of locations, one at a time, those that provide greatest cost reduction, considering the facilities already open. Facilities are added until the constraint is met for all customers. For reviews, see Vidal and Goetschalckx (1997), Beamon (1998),

Erenguc et al. (1999), Sarmiento and Nagi (1999), Reese (2005) and Nagi and Salhi (2007).

Ardalan (1988) compared the performance of the Kuehn and Hamburger (1963) heuristics with the ones originally proposed by Khumawala, also considered efficient and simple to implement. They were compared against the optimal solution under several location–allocation scenarios, with customer locations varying from 10 to 40 nodes from which 3 or 4 supply sources would be allocated using either heuristics; a total of 14 different scenarios. Each problem was replicated 50 times with each heuristics. Ardalan found that the proposed heuristics performed better than Khumawala's in almost every situation examined, except in one case, the selection of 4 supply nodes from 10 customer locations. Given the impact of this study, the Kuehn–Hamburger heuristics has since been known as the Ardalan heuristics, being included in several operations management textbooks.

Rahman and Smith (1991) and Chaudhry et al. (1995) compared Teitz and Bart (1968) and the Ardalan (1988) heuristics, with and without maximum distance constraints, using a variety of data sets. Although Ardalan generally provided worse solutions than Teitz–Bart, it should be noted that Teitz–Bart had to be applied several times with different random seeds before the best objective value was eventually found. Summarizing the comparative studies, we have the following:

- Ardalan's heuristics is superior to Teitz–Bart's in problems with fewer customers and distance constraints.
- Ardalan's is more powerful than Khumawala's in problems with 10–40 customers and 3 or 4 nodes, except when more nodes are desired for fewer customers.
- Teitz–Bart's is generally superior to Ardalan's in problems with 30–75 customers without distance constraint.

These comparisons are not enough to indicate that any particular heuristic is *the best*. Since Teitz–Bart's is cumbersome for treating problems with distance constraints, the most appropriate selection in these cases should be one of Khumawala's heuristics or Ardalan's heuristic.

6. Analysis: three approaches to spare engines storage

The following section presents a centralized storage solution for the spare engine inventory problem, a lower bound that ignores distance and time constraints. Then it uses the Ardalan heuristic to select the bases that should hold the inventory of each engine type separately, or both engine types together. We select Ardalan's method for two reasons: it is the easiest to incorporate distance constraint using MS Excel, and it is better than Khumawala's in problems with lower customer-to-node ratio. All approaches in this section meet an 85% service level in all bases at a lower cost than what is currently practiced by the Air Force.

6.1. Centralized storage

It is a well-known statistical fact that when two independent random variables are added, the resulting variable is proportionally less variable, provided that the variables are not correlated. Recent demand data for the F-16 replacement engines (Table 1) show that the changes in demand in the bases do not follow any particular pattern that would suggest a correlation. Demand increased in some bases, decreased in others, and oscillated randomly in most of them. A direct verification of the two-way correlation between the 30 bases shows that most of them fit in

the range from -0.2 to $+0.5$. For all we know, the variations in demand were not driven by a common policy that affected all bases in the same way each year. Rather, local policies affected the changes separately and thus, there is prior expectation that the coefficient of variation of the joint demand of multiple bases is lower than the coefficient of variation of the demand in individual bases. Hence, to manage the demand variability of two or more bases, it is possible to hold lower aggregate inventory in a single facility than it is to hold each inventory separately.

This simple observation has powerful applications that are often ignored. For instance, Tinker AFB does not have any assigned F-16 aircraft. However, considering its status as the maintenance depot and its central location, it is conceivable to store all F-16 engines at the depot, regardless of make, and ship them directly to the respective base when needed. Pooling this demand under a single inventory would reduce the safety stock, generating substantial savings. Under this policy, the total inventory of Pratt & Whitney engines necessary to satisfy demand (with the same degree of confidence in satisfying the demand in each base) would drop from 38 (see Table 2) engines to 18 engines, as shown in Table 3. Likewise, GEAE engine inventory would drop from 48 to 20. These inventory reductions would be credited exclusively to the pooling effect; in other words, to achieve the same service level, a centralized (or *pooled*) inventory requires a smaller safety stock than a distributed (or *pre-positioned*) inventory. In addition to adopting a centralized inventory, if the engines were designed using an open architecture, we would be able to reduce the inventory further, from 38 to 35 engines. The safety stock would reduce by 30%, from 10 to 7 units. Open architecture would require that P&W and GEAE engines could be used interchangeably in any airframe.

There are a few weaknesses associated with central storage solution. Considering that it is necessary to ensure that the drive time to receive the engine when ordered from the field is less than 1 day, inventory consolidation at only one central location might not meet the operational needs: it would impose up to a 3-day traveling time between the inventory and some users, compromising their readiness. However, according to users in the field, one business day (550 miles driving distance or less) is an acceptable traveling time for a replacement engine.

Proponents of pre-positioning will point to availability (or *readiness*) as one of its greatest benefits. However, just as the centralization of the inventory in a single location is inefficient, pre-positioning is costly and may expose the user to potentially lower inventory availability (due to increased demand variability) unless there is an additional investment in safety stock. The following section shows an alternative approach that does not

centralize the inventory in a single location, but in a few selected locations. The solution combines some of the advantages of pre-positioned inventories with the risk-pooling benefits of an aggregate storage plan.

6.2. Regional storage

Some of the risk pooling benefits can be observed through regional inventory storage. We identify this as the *p-median* problem. A word of caution: a transportation model based on a mixed-integer program would not be a useful approach to find the storage points in this problem for two main reasons: (1) The problem is fairly complex to be analyzed using software typically available to most managers (MS Excel and Solver). (2) Most important, if every customer is also a potential sourcing point, and the number of storage points is pre-determined but not pre-identified, the solution process would encounter discontinuities in the objective function. This would prevent us from finding the optimal solution, even for a small problem with just 13 customers.

The heuristics popularized by Ardalan (1988) requires the development of a table of distances between potential inventory locations and users, as well as the assignment of weights to help prioritize the decisions. It is a greedy procedure that sequentially identifies the locations that are closest to the most demanding users until all warehouses are identified. In this problem, the table of distances was created with *Yahoo! Maps* (<http://maps.yahoo.com>). As recommended by the heuristic, an Ardalan table is created as the product between the user's demand (d_i), the table of distances (x_{ij}) and a weight associated with that delivery (w_j). Because of existing resources at active duty bases, it is usually more desirable to store engines there than at ANG or AFR. Moreover, it is more desirable to store engines at Tinker AFB than at any other AD base because it is the depot that provides major maintenance support for the F-16 engines. Hence, this modified Ardalan matrix assigns weights to the sources (w_j) that act as a "source penalty," rather than to customers, which would indicate their priority levels. Since Tinker is the ideal source, its weight is 1. Other AD bases received a weight of 1.1, while the AFR and ANG sites received a weight of 1.5. To create a distribution network for k customers, this process generates a square matrix with k rows and columns in which the value of each cell (a_{ij}) is determined by the expression:

$$a_{ij} = d_i x_{ij} w_j$$

This matrix is the root of the procedure to identify a set of storage locations that require low transportation time to the respective users and provide the benefit of inventory aggregation. The following steps identify an efficient set of storage locations:

Step 1: Let $s=1$. This variable is the number of storage points at the end of this round.

Step 2: Generate the value $A_j = \sum_i a_{ij}$, the weighted sum of all shipments from storage point j to each customer i .

Step 3: Identify $A_m = \text{Min}\{A_1, \dots, A_{k-s+1}\}$. Column m defines the least onerous storage location in round s .

Step 4: Move column m to the end of the matrix.

Step 5: For each cell (i,j) that satisfies $j \leq k-s$, let $a_{ij} = \text{Min}\{a_{ij}, a_{mj}\}$.

Step 6: If a stopping point is reached, stop. Otherwise, let $s=s+1$ and repeat steps 2–5.

The stopping point could be, for instance, a pre-established number of storage facilities or some capacity limitation. In this case, since this is a distance-constrained problem, we added storage points until all users were served by inventories within a

Table 3
Central storage of spare engines at Tinker AFB.

Inventory status: all replacement engines stored at Tinker AFB			
	P&W	GEAE	Total
Annual demand	656	773	1429
Base stock	18	20	38
Safety stock	5	5	10
\$ Safety stock	\$16.4	\$14.8	\$31.1
Inventory status: all replacement engines stored at Tinker AFB. OA allows using engines of different makes in any F-16 airframe			
All makes, common interface			
Annual demand	1429		
Base stock	35		
Safety stock	7		
\$ Safety stock	\$21.71		

1-day drive (approximately a 550-mile distance). However, as any heuristics, some exception may be necessary to ensure that it finds a solution that is efficient (low cost) and effective (meets all practical constraints). Consequently, we added another twist: each low-cost location found with the heuristic is selected as a new storage point only if it increases the network coverage, i.e., one or both conditions are satisfied: (1) the low-cost location is not within range from any of the existing storage locations or (2) the low-cost location is within range from a customer that cannot be served by any of the existing storage locations. If these

conditions are not satisfied, that low-cost location is not contributing with the inventory-pooling objective, and the next low-cost location is selected in its place.

The heuristics is illustrated with the allocation of the Pratt & Whitney spare engines, in Table 4. All 13 customers are served from 6 locations, shown in Fig. 1. Nellis AFB (LSV), located within a 1-day drive distance from two previously assigned storage locations (LUF and HIF), does not improve network coverage, so it is not selected in rounds 5 and 6. The next lower-cost location in each round, DLH (ANG-Duluth in Minnesota) and FWA (ANG-Ft Wayne in Indiana) are selected instead. Notice that despite the inventory pooling efforts, three locations (Hill AFB-Depot in Utah, ANG-Duluth and ANG-Burlington in Vermont) store inventory for just their needs, because of their distance to other bases using the same engine type.

Once the low cost storage locations are identified, each user *i* is assigned to the storage location *j* that satisfies the equation $a_{ij} = \text{Min}\{a_{i,k-s+1}, a_{i,k}\}$. Each storage location holds the inventory to meet the demand for all users assigned to it, in addition to a safety stock based on the aggregate demand variability in the bases supplied by that location, and on the lead-time for that location to resupply from the central depot in Tinker.

Locations holding inventory for multiple users benefit from the pooling effect already discussed. For example, Luke AFB (LUF) holds inventory for its needs and for three other bases. The standard error of the forecasted annual demand (260 business days) from each of these bases range between 10.3 and 19.1. However, the standard error of the aggregate demand is just 31.4. Considering a lead-time of 6 business days from TIK (the depot) to LUF, the safety stock to meet the demand variability of all four bases is just 4.9 units. Also, the aggregate expected demand from the four users is 440 units per year (or 10.2 units during the lead-time), which leads to a base stock of 16 (=4.9+10.2) units. Detailed information about the inventory allocations appears in Table 5.

Using the heuristic to assign storage points for GEAE spare engines, eight storage locations are sequentially identified (Table 6). In this case, no exceptions are necessary, since every allocation suggested by the heuristic expands the network coverage—adding at least one base to within the 1-day delivery threshold. In this analysis, eight bases are needed to hold the inventory for 19 bases using GEAE engines (Fig. 2). Again, two bases (AFR-Homestead FL and ANG-Montgomery AL) hold

Table 4
Heuristic application to the Pratt & Whitney spare engines allocation problem. Lowest choice location is shown in parentheses, if different from selected site.

Heuristic round	Lowest cost location	Selected distribution sites	Users within range from all sites
s=1	LUF	LUF	4
s=2	TIK	LUF, TIK	8
s=3	BTV	LUF, TIK, BTV	9
s=4	HIF	LUF, TIK, BTV, HIF	10
s=5	DLH (LSV)	LUF, TIK, BTV, HIF, DLH	11
s=6	FWA (LSV)	LUF, TIK, BTV, HIF, DLH, FWA	13 (all)

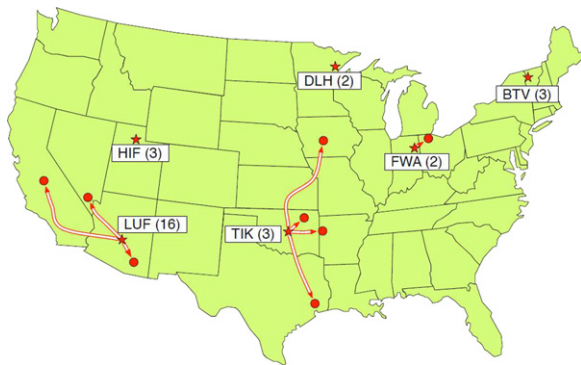


Fig. 1. Regional storage of Pratt & Whitney spare engines. All maps: storage locations identified as a star with base stock in parentheses. Users identified as a circle. Typical roads are shown.

Table 5
Regional storage of Pratt & Whitney spare engines. Source: The author.

Pratt & Whitney (F-100 engine)						
Base	ICAO	2008 forecast Mean	2008 inventory		Distance from inventory (mile)	
			Base stock	Safety stock		
Luke AFB AZ	LUF	248.1	16	4.9	0	
ANG-Tucson AZ	TUS	80.1	At LUF		145	
Nellis AFB NV	LSV	69.9	At LUF		278	
ANG-Fresno CA	FAT	41.4	At LUF		579	
Tinker AFB OK	TIK	0.0	3	1.8	0	
ANG-Des Moines IA	DSM	8.9	At TIK		545	
ANG-Ft Smith AR	FSM	17.1	At TIK		183	
ANG-Ellington TX	EFD	33.4	At TIK		467	
ANG-Tulsa OK	TUL	4.9	At TIK		119	
ANG-Burlington VT	BTV	34.9	3	1.1	0	
Hill AFB-Depot UT	HIF	67.7	3	1.1	0	
ANG-Duluth MN	DLH	23.0	2	1.1	0	
ANG-Ft Wayne IN	FWA	25.7	2	1.3	0	
ANG-Toledo OH	TOL	1.3	At FWA		98	

Note: Storage locations in bold.

inventory exclusively for their needs because of their distance to other bases using the same engine. Table 7 shows storage points, the size and the distance from each base to the respective inventories.

As the analysis shows, regional storage reduces the safety stock of P&W engines from 18 to 11 engines and the safety stock of GEAE engines from 22 to 15 engines, in contrast to the fully

distributed storage of engines shown in Table 2. By pooling the demand from each base into a limited number of storage points, the coefficient of variation of the forecasted demand is reduced, which leads to lower safety stock requirement and substantial savings. This inventory allocation requires that two users (ANG-Des Moines in Iowa and ANG-Great Falls in Montana) be served by inventory located more than 500 miles away, but no more than 600 miles from the user. Yet, this allocation allows all bases to receive their spare engines within 1 day.

Table 6
Heuristic application to the GEAE spare engines allocation problem.

Heuristic round	Lowest cost location	Selected distribution points	Users within range
s=1	TIK	TIK	4
s=2	HIF	TIK, HIF	6
s=3	ADW	TIK, HIF, ADW	12
s=4	SGH	TIK, HIF, ADW, SGH	14
s=5	HST	TIK, HIF, ADW, SGH, HST	15
s=6	MGM	TIK, HIF, ADW, SGH, HST, MGM	16
s=7	CVS	TIK, HIF, ADW, SGH, HST, MGM, CVS	17
s=8	FSD	TIK, HIF, ADW, SGH, HST, MGM, CVS, FSD	19 (all)

6.3. Regional storage with open architecture benefit

The regional storage performance could be substantially improved if the engines were designed with an open architecture mindset. Without OA, the inventory distribution in Figs. 1 and 2 benefits only from the regional pooling effect, observed when we aggregate the demand for identical items from customers in close proximity. In addition to the benefits associated with simpler design and maintenance of these complex assets, the adoption of open architecture would increase the number of bases in some geographic regions that could be served by the same storage

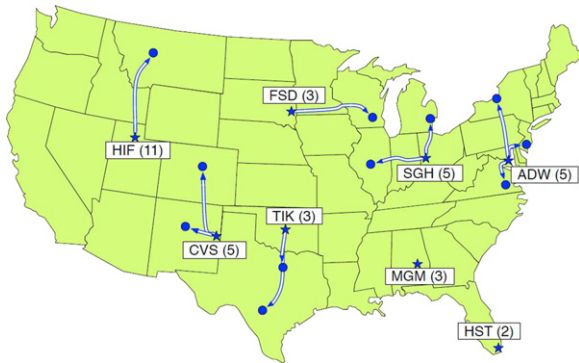


Fig. 2. Regional storage of General Electric spare engines.

Table 8
Heuristic application to the complete spare engines allocation problem.

Heuristic round	Lowest cost location	Selected distribution points	Users within range
s=1	ABQ	ABQ	6
s=2	ADW	ABQ, ADW	15
s=3	HIF	ABQ, ADW, HIF	18
s=4	LUF	ABQ, ADW, HIF, LUF	19
s=5	TIK	ABQ, ADW, HIF, LUF, TIK	24
s=6	FWA	ABQ, ADW, HIF, LUF, TIK, FWA	26
s=7	FSD	ABQ, ADW, HIF, LUF, TIK, FWA, FSD	28
s=8	HST	ABQ, ADW, HIF, LUF, TIK, FWA, FSD, HST	29
s=9	MGM	ABQ, ADW, HIF, LUF, TIK, FWA, FSD, HST, MGM	30 (all)

Table 7
Regional storage of General Electric spare engines.
Source: The author.

General Electric (F-110 engine)					
Base	ICAO	2008 forecast	2008 inventory		Distance from inventory (mile)
		Mean	Base stock	Safety stock	
Tinker AFB OK	TIK	0.0	3	1.2	0
ANG-Kelly TX	SKF	45.1	At TIK		485
AFR-Ft Worth TX	NFW	21.1	At TIK		211
Hill AFB UT	HIF	217.0	11	3.5	0
ANG-Gt Falls MT	GTF	44.4	At HIF		544
ANG-Andrews DC	ADW	25.0	5	2.3	0
ANG-Eggharbo NJ	ACY	19.4	At ADW		168
ANG-Syracuse NY	SYR	38.5	At ADW		386
ANG-Richmond VA	RIC	8.7	At ADW		122
ANG-Springfield OH	SGH	51.3	5	1.8	0
ANG-Selfridge MI	MTC	30.1	At SGH		253
ANG-Springfield IL	SPI	31.0	At SGH		347
AFR-Homestead FL	HST	28.4	2	0.8	0
ANG-Montgomery AL	MGM	37.3	3	1.5	0
Cannon AFB NM	CVS	48.9	5	2.0	0
ANG-Albuquerque NM	ABQ	31.3	At CVS		220
ANG-Buckley CO	BKF	27.4	At CVS		493
ANG-Sioux SD	FSD	41.0	3	1.4	0
ANG-Madison WI	MSN	26.4	At FSD		429

Note: Storage locations in bold.

location, adding another dimension of supply-chain aggregation to reduce the need for safety stock. This pooling effect created by open architecture is called *product aggregation*: different products that are perfect substitutes can be held as a single inventory pool. This aggregation has the same pooling effect as when we pool inventories from different locations.

The same heuristic used earlier to find separate inventory storage can be used to find storage locations for spare engines built using open architecture. The separate inventory allocation required 6 bases for P&W engines and 8 bases for GEAE engines. The joint allocation proceeds as shown in Table 8. In this environment, engines from either manufacturer could be used in any airframe. To meet the demand of all 30 bases from




Fig. 3. Regional storage of spare engines with OA benefit.

locations within a 1-day drive, nine storage points suffice, shown in Fig. 3. Among the selected storage points, there are three Active Duty bases, one Air Force Reserve and five Air National Guard sites. Notice that among the ANG bases, there are two that are co-located with AD bases (ANG-Albuquerque in New Mexico and ANG-Andrews in District of Columbia). These may enjoy some support from this arrangement. Detailed information about the inventory allocation is in Table 9.

Open architecture increases the population density of users that can be served from the same overall inventory pool. In the original allocation, there were 12 bases storing spare engines (two of them storing both types), which amounted to 14 different inventory pools. With this approach, only nine inventory pools are necessary. Note that only two bases (ANG-Montgomery in Alabama and AFR-Homestead in Florida) remain isolated, holding just the engines required for their operation. This is quite an improvement from the previous solution without OA, with five isolated locations. Table 10 summarizes the performance of regionalized storage with and without the benefit of open architecture. Thanks to this added level of aggregation, the total safety stock necessary to absorb the variability of demand in 30 bases is now just 21 units, contrasting quite favorably with the safety stock of 40 engines in the distributed mode based on current policy (Table 2).

As Fig. 3 shows, the use of open architecture greatly improves the distribution network in the Northeast, Southwest, the Midwest and in the central part of the country, where four storage locations (Tinker AFB, Luke AFB, Andrews AFB and ANG-Fort Wayne) serve 20 bases. On the downside, three bases are served by inventory located between 500 and 580 miles away, stressing

Table 9
Regional storage of spare engines with OA benefit.
Source: The author.

Base	ICAO	2008 forecast Mean	2008 inventory		Distance from inventory (mile)
			Base stock	Safety stock	
ANG-Albuquerque NM	ABQ	31.3	5	2.2	0
Cannon AFB NM	CVS	48.9	At ABQ		220
ANG-Buckley CO	BKF	27.4	At ABQ		453
ANG-Andrews DC	ADW	25.0	6	2.6	0
ANG-Richmond VA	RIC	8.7	At ADW		122
ANG-Syracuse NY	SYR	38.5	At ADW		386
ANG-Egg Harbor NJ	ACY	19.4	At ADW		168
ANG-Burlington VT	BTV	34.9	At ADW		523
Hill AFB UT	HIF	284.7	12	2.9	0
ANG-Gt Falls MT	GTF	44.4	At HIF		544
Luke AFB AZ	LUF	248.1	16	4.9	0
ANG-Fresno CA	FAT	41.4	At LUF		579
Nellis AFB NV	LSV	69.9	At LUF		278
ANG-Tucson AZ	TUS	80.1	At LUF		145
Tinker AFB OK	TIK	0.0	4	2.0	0
ANG-Ellington TX	EFD	33.4	At TIK		467
ANG-Ft Smith AR	FSM	17.1	At TIK		183
AFR-Ft Worth TX	NFW	21.1	At TIK		211
ANG-Kelly TX	SKF	45.1	At TIK		485
ANG-Tulsa OK	TUL	4.9	At TIK		119
ANG-Ft Wayne IN	FWA	25.7	7	2.4	0
ANG-Madison WI	MSN	26.4	At FWA		321
ANG-Selfridge MI	MTC	30.1	At FWA		194
ANG-Springfield OH	SGH	51.3	At FWA		141
ANG-Springfield IL	SPI	31.0	At FWA		328
ANG-Toledo OH	TOL	1.3	At FWA		98
ANG-Sioux SD	FSD	41.0	4	1.8	0
ANG-Duluth MN	DLH	23.0	At FSD		396
ANG-Des Moines IA	DSM	8.9	At FSD		292
AFR-Homestead FL	HST	28.4	2	0.8	0
ANG-Montgomery AL	MGM	37.3	3	1.5	0

Note: Storage locations in bold.

the operational constraint in this distribution process. However, most users sit within a 1-day drive from one or more additional storage points. Consequently, the safety stock necessary to manage the demand variability of all 30 bases, which used to be 40 units in the original allocation (see Table 2), is now just 21 units. Considering that the average engine costs the DoD approximately \$3.1M, this reduction accounts as direct savings of \$58.8M—a savings due to the adoption of open architecture in a regionalized inventory distribution.

6.4. Summary of results

The current distribution of spare engines for F-16s was used to illustrate and evaluate the benefits of place and product aggregation. Starting from the status quo, in which the inventory is locally distributed in the hands of each user, and considering that there are two engine makes (Pratt & Whitney and General Electric) that are not interchangeable, we evaluated four alternative distribution models representing different aggregation approaches: with or without open architecture (product aggregation) and centralized or regionalized distribution (place aggregation). As expected, both types of aggregations provided inventory reduction. What perhaps was not expected was the dimension of the safety stock reduction (Table 11).

Table 10
Regional storage of spare engines.

INVENTORY STATUS: replacement engines stored at a limited number of bases			
	P&W	GEAE	Total
Number of storage bases	6	8	12
Annual demand	656	773	1429
Base stock	29	37	66
Safety stock	11	15	26
Safety stock value	\$36.0M	\$44.3M	\$80.3M
Inventory status: replacement engines stored at a limited number of bases. OA allows using engines of different makes in any F-16 airframe			
All makes, common interface			
Number of storage bases	9		
Annual demand	1429		
Base stock	59		
Safety stock	21		
Safety stock value	\$65.1M		

Table 11
Performance of different aggregation approaches.

	Product aggregation	
	Open architecture	Proprietary design
<i>Place aggregation</i>		
Central	Storage points: 1 Safety stock: 7 Safety stock value: \$21.7M Demand-miles: 1412 k	Storage points: 1 Safety stock: 10 Safety stock value: \$31.1M Demand-miles: 1412 k
Regional	Storage points: 9 Safety stock: 21 Safety stock value: \$65.1M Demand-miles: 1.554 k	Storage points: 12 Safety stock: 26 Safety stock value: \$80.3M Demand-miles: 1512 k
Local	Storage points: 30 safety stock: 38 safety stock value: \$117.2M demand-miles: 1412 k	Storage points: 30 safety stock: 40 safety stock value: \$123.9M demand-miles: 1412 k

One important concern is the impact on transportation costs. To facilitate comparison, the analysis included a measure of expected miles driven to each facility, considering that each engine would generally be transported from the main depot to a regional storage base, and then to the user. The baseline measure of 1.41 million miles is the product between the number of units shipped and the distance traveled. This total is the same, whether the storage is centralized or distributed. Pooling the storage into 12 bases (without open architecture) would increase the distance driven—and the transportation cost—to 1.51 million miles. If the storage is pooled into just 9 bases (with open architecture), the distance driven is increased to 1.55 million miles—10% more than the baseline, to obtain a 47% safety stock reduction.

7. Conclusions and future research

An important concern in supply-chain management is the identification of aggregation opportunities that exist in the design, storage and distribution of goods to the final customer. This aggregation, or *pooling*, can take place in three dimensions: time aggregation, place aggregation and product aggregation. The manager should evaluate the trade-offs in each of these aggregation opportunities in order to implement the correct product design, storage and distribution procedures. Time aggregation implies that the inventory is kept to meet the demand over longer or shorter periods of time. Place aggregation implies that the inventory is designed to meet the demand over one or many markets. Finally, product aggregation implies that a product or component is designed to meet the demand associated with one or more applications. When any or all of these aggregations are viable, the company enjoys substantial operational savings due to the reduction of safety stock; in addition, much of the coordination effort may be reduced.

This paper analyzes the case of the F-16 spare engine storage to illustrate two types of pooling benefits: by place and by product. Here, product aggregation is achieved with the adoption of open architecture in product design. Place aggregation is achieved with judicious use of the classic *Ardalan (1988)* heuristics subject to a distance constraint, necessary to ensure the desired performance level. We draw several important lessons from this study:

Open architecture is an effective means of product aggregation to facilitate supply-chain improvement for valuable complex assets.

Open architecture can be leveraged by place aggregation when the asset is used by several facilities geographically distributed.

Open architecture provides the greatest inventory reduction benefit when storage can be centralized. If centralization is not desirable, it can still provide substantial benefits, by judicious identification of a regional cluster of users to share joint inventory.

Reduction in the number of storage points generally increases transportation cost. Hence, it is important to evaluate the trade-off between simplified infrastructure and reduced investment in inventories against increases in transportation cost.

The purpose of this study is to evaluate the economic impact of adopting open architecture in the design of complex assets to reduce the lifecycle cost of maintaining those assets. The adoption of open architecture affects several economic components in the life of the asset, including developmental costs, maintenance costs, and inventory-management costs. This article focuses on the benefits of pooling the inventory necessary to meet the demand of many users into a small number of storage sites with product variety reduction obtained with open architecture.

The example showcased in this analysis—distribution of spare engines for the F-16 in continental United States—amply supports open architecture as the design approach to reduce supply chain expenditures without compromising the availability of valuable assets. The relevance here is far beyond the potential savings that the F-16 program could have enjoyed, but is a lesson for future government programs—whether they are weapon systems or other assets supplied by two or more qualified suppliers. Several examples come to mind, among them unmanned aerial vehicles, the space program and high-speed rail equipment.

More than 40 years since Hakimi (1965) first formulated the *p*-median problem, it is still necessary to warn managers in government of the importance to use rigorous techniques to make facility location decisions. Fok (2001) raised this issue, at the same time that he recommended the use of accessible techniques, such as the center of gravity method – for centralized storage – or the Ardalan heuristics – for multi-facility storage, as is our case.

Future studies about the impact of open architecture on complex systems should expand the analysis to incorporate benefits provided by simplified maintenance, as well as to investigate the additional cost and time required to coordinate the developmental efforts to ensure a common interface. On the distribution side, the regular adoption of transshipment (lateral shipment) or the use of multiple storage units should be studied as alternatives to improve inventory pooling, and thus to enhance the value of open architecture in future product development programs.

Acknowledgements

The authors wish to acknowledge two anonymous referees for providing numerous suggestions that greatly improved this study. We are thankful to Retired Rear Admiral Jim Greene, the Naval Postgraduate School Acquisition Research Chair, and to Prof. Keith Snider, director of the Acquisitions Research Program, for securing the necessary research funds to support this work.

References

- Ardalan, A., 1988. A comparison of heuristic methods for service facility locations. *International Journal of Operations and Production Management* 8 (2), 52–58.
- Beamon, B.M., 1998. Supply chain design and analysis: models and methods. *International Journal of Production Economics* 55, 281–294.
- Chaudhry, Sohail S., Choi, In-Chan, Smith, David K., 1995. Facility location with and without maximum distance constraints through the *p*-median problem. *International Journal of Operations and Production Management* 15 (10), 75–81.
- Coar, K., 2006. Open Source Initiative. Available from: <<http://www.opensource.org/docs/osd>>.
- DAU, 2006. Open Architecture, Naval Open Architecture—Terms and Definitions. Available from: <<https://acc.dau.mil/oa>>.
- Erenguc, S.S., Simpson, N.C., Vakharia, A.J., 1999. Integrated production/distribution planning in supply chains: an invited review. *European Journal of Operational Research* 115, 219–236.
- F-35 JSF Program, 2007. F-35 Joint Strike Fighter Program: Technology. Available from: <http://www.jsf.mil/f35/f35_technology.htm>.
- Federation of American Scientists, 2005b. F-16 Fighting Falcon. Available from: <<http://www.fas.org/man/dod-101/sys/ac/f-16.htm>>.
- Federation of American Scientists, 2005a. Joint Strike Fighter. Available from: <<http://www.fas.org/man/dod-101/sys/ac/jsf.htm>>.
- Fok, Wing M., 2001. The Location of Subregional Offices: An Emerging Issue for HR Managers in Government. *Public Personnel Management*. Available from: <<http://www.allbusiness.com/human-resources/employee-benefits/835889-1.html>>.
- Hakimi, S.L., 1965. Optimum distribution of switching centers in a communication network and some related graph theoretic problems. *Operations Research* 13, 462–475.
- Henderson, Robert, Higer, Matthew, 2007. Potential Logistics Cost Savings from Engine Commonality. Naval Postgraduate School, Monterey, CA.
- Khumawala, B.M., July 1972. An efficient algorithm for central facilities location. Paper no. 357, Krannert School of Industrial Administration, Purdue University.
- Kuehn, A.A., Hamburger, M.J., 1963. A heuristic program for locating warehouses. *Management Science* 10, 643–666.
- Nagy, Gábor, Salhi, Saïd, 2007. Location-routing: issues, models and methods. *European Journal of Operational Research* 177, 649–672.
- Nelson, Eric, 2007. Open Architecture Technical Principles and Guidelines (Ver. 1.5.4). International Business Machines Corporation, Armonk, New York.
- PEO-IWS (Program Executive Officer—Integrated Weapons System), 2006. Department of the Navy. Naval Open Architecture Contract Guidebook (Ver. 1.0) (PEO-IWS 7). Available from: <<https://acc.dau.mil/oa>>.
- Rahman, S., Smith, D.K., 1991. A comparison of two heuristic methods for the *p*-median problem with and without maximum distance constraints. *International Journal of Operations and Production Management* 11 (6), 76–84.
- Reese, J., 2005. Methods for solving the *p*-median problem: an annotated bibliography. Manuscript, Trinity University, San Antonio, TX. Available from: <<http://ramanujan.math.trinity.edu/tumath/research/reports/report96.pdf>>.
- Sarmiento, A.M., Nagi, R.A., 1999. Review of integrated analysis of production–distribution systems. *IIE Transactions* 31, 1061–1074.
- Teitz, M.B., Bart, P., 1968. Heuristic methods for estimating the vertex median of a weighted graph. *Operations Research* 12 (3), 955–961.
- Tinker AFB, 2007. 448 Combat Sustainment Wing's mission statement. Available from: <<http://www.tinker.af.mil/units/>>.
- Toregas, C., Swain, R.W., ReVelle, C.S., Bergman, L., 1971. The location of emergency service facilities. *Operations Research* 19 (6), 1363–1373.
- Vidal, C.J., Goetschalckx, M., 1997. Strategic production distribution models: a critical review with emphasis on global supply chain models. *European Journal of Operational Research* 98, 1–18.