

www.palgrave.com/journals

On sourcing and stocking policies in a two-echelon, multiple location, repairable parts supply chain

Izack Cohen[1]*, Morris A. Cohen[2] and Elad Landau[1]

[1]*The Technion, Haifa 321000, Israel; and* [2]*The Wharton School, Philadelphia, PA, USA*

This research develops policies to minimize spare part purchases and repair costs for maintaining a fleet of mission-critical systems that operate from multiple forward (base) locations within a two-echelon repairable supply chain with a central depot. We take a tactical planning perspective to support periodic decisions for spare part purchases and repair sourcing, where the repair capabilities of the various locations are overlapping. We consider three policy classes: a central policy, where all repairs are sourced to a central depot; a local policy, whereby failures are repaired at forward locations; and a mixed policy, where a fraction of the parts is repaired at the bases and the remainder is repaired at the depot. Parts are classified based on their repair cost and lead time. For each part class, we suggest a solution that is based on threshold policies or on the use of a heuristic solution algorithm that extends the industry standard of marginal analysis to determine spare parts positioning by including repair fraction sourcing. A validation study shows that the suggested heuristic performs well compared to an exhaustive search (an average 0.2% difference in cost). An extensive numerical study demonstrates that the algorithm achieves costs which are lower by about 7–12% on average, compared to common, rule-based sourcing policies.

Journal of the Operational Research Society (2017) **68(6),** 617–629. doi:10.1057/s41274-016-0108-0; published online 17 November 2016

Keywords: supply chain; repairable parts; sourcing; two-echelon

1. Introduction

This paper introduces a model for two-echelon repairable supply chains for the after-sales support of mission critical products. The objective of the model is to optimize deployment of two classes of resources, ie, spare parts inventory and repair capacity within a two-echelon service support supply chain consisting of a central depot and multiple forward (base) locations. A tactical planning perspective is adopted to support periodic decisions for spare part purchases and repair sourcing. The main focus of the analysis is to account for the heightened level of flexibility for repair sourcing that has been enabled by recent developments in product design, maintenance information support systems, and the tooling used to complete part repairs. These developments allow for easier fault isolation and repair and thus have enabled overlapping repair capabilities within the network of service supply chain facilities, making it possible to source repairs to different repair facilities located at either the central depot or at a forward base – see Muller *et al* (2008) for a discussion of the flexibility provided by new maintenance technologies or a description of one example of a commercial software solution offered by PTC at http://www.ptc.com/service-lifecycle-management/service-parts-information. The PTC software enables "the creation, management and delivery of content for servicing equipment." and "provides up-to-date technical, parts and 2D/3D graphical information to the global service network." The existence of such repair capacity overlap enables managers in these environments to plan for the sourcing of repairs across multiple locations as a part of their tactical resource deployment planning process.

The repair sourcing decision must be made jointly with the selection of target inventory stocking levels at all locations in the service support network. Thus, our model, whose objective is to minimize stocking and repair costs by jointly deciding on repairs sourcing, spare parts purchasing, and parts allocation to different locations, explores the tradeoff between repair cost, inventory cost, and the speed of repair. The model also explicitly captures the impact of risk pooling which is afforded by positioning spare parts and repair capacity at the central depot.

Our analysis contributes to the extensive literature on multi-echelon repairable inventory management by introducing the repair sourcing decision (ie, the planned fraction of repairs to be conducted at each location), in the context of a stochastic planning model that minimizes total cost (inventory plus repair) subject to an overall service constraint. The analysis is based on the derivation of new closed-form structural results. Policy analysis is conducted both analytically and through an extensive numerical experiment leading to managerial insights associated with optimal stocking and repair sourcing policies.

The structural results support the definition of an effective (near optimal, fast) heuristic algorithm, which utilizes the

*Correspondence: *Izack Cohen, The Technion, Haifa 321000, Israel.*
E-mail: izik68@tx.technion.ac.il

structural results and which can be incorporated into existing stocking algorithms that are standard in the after-sales service support industry.

The problem considered here extends the standard multi-echelon model (Feeney and Sherbrooke, 1966) that takes repair sourcing decisions as inputs. There is now an opportunity to improve existing planning models and solution practices by making repair sourcing decisions periodically within the tactical resource planning process (eg, annual, bi-annual budget planning cycles) and jointly with inventory stocking decisions. Our model is motivated by industry situations which we have witnessed, in which changing repair sourcing (eg, performing more in-house repairs or outsourcing more repairs) has proved to be beneficial in response to changes in the installed base, reliability of parts due to engineering change, changes in costs and lead times of repairs, increases in the cost of acquiring parts, and adjustments to the utilization profile of the installed systems. In the implementations that we are familiar with, such changes have led to periodic re-estimation of model parameters and re-solving the model. Additionally, in today's performance-based environment, where outsourcing of after-sales has become more common (Guajardo *et al*, 2012), the importance of updating repair allocation policies periodically, on the basis of total cost tradeoffs and in response to changes in support requirements, is considerable.

Our model also contributes to the literature on Level of Repair Analysis (LORA) which focuses on locating repair capacity and the sourcing of part repairs. We, in particular, contribute to the growing literature which addresses the joint solution of stocking and sourcing for repairable parts.

Our model recommends both the target fraction of failed parts to source from each repair location and the target inventory level for each part at that location. These periodic planning decisions are not a function of the system's real-time status, but rather are made to minimize expected system costs. In contrast, day-to-day control, used for operational decisions, such as where to repair the part that has failed today or how to re-allocate inventory across locations, considers current system status (eg, the load at repair facilities, actual stocking, and pipeline inventory levels) and can also take into account the realized complexity of any part failure. Thus, actual (real time) decisions for sourcing repairs or for re-allocating inventory across locations are not included in the model. Rather, the output of our model can be used to set parameters which support resource allocation and budget planning and also can guide such real-time decisions.

2. Literature review

As noted, our analysis builds on two streams of research literature, ie, Multi-Echelon (ME) models and Level of Repair Analysis (LORA). In this section, we briefly review relevant literature from the two streams and focus on examples that have combined both.

The standard ME model, used to quantify the tradeoff between repairable stock levels and system availability is the Metric inventory model (introduced by Feeney and Sherbrooke, 1966; Sherbrooke, 1968). This model was originally developed for the U.S. Air Force to minimize expected inventory costs subject to meeting a probabilistic service goal. Variants of ME models have been implemented extensively in the aerospace and defense, high-technology, and medical equipment industries throughout the world and are supported by commercial software systems (see for example http://www.ptc.com/service-lifecycle-management/service-parts-management). There have been extensions that consider stocking decisions for multiple indentures (Muckstadt, 1973), lateral transshipments (Lee, 1987), cannibalization (Sherbrooke, 2004; Gaver *et al*, 1993), non-backordering and emergency shipments (Cohen *et al*, 1988), and location-dependent lead times (Wang *et al*, 2000) as well as decisions made under budget constraints (Rustenburg *et al*, 2000). These models relate to stocking decisions for both the initial supply and planning phase as well as during the exploitation phase of a fleet (van Houtum, and Kranenburg, 2015) where a stationary model is periodically updated. The main ME model assumption, which we relax, is that repair sourcing is given as an input. We also consider different repair costs at different repair locations.

LORA models determine, for each part-location in an echelon, if the part should be discarded, repaired, or moved to a higher echelon. LORA research has focused on problem formulations and on the development of heuristic solution algorithms (eg, Saranga and Kumar, 2006 that used genetic algorithms). Basten *et al* (2009) formulated an IP model and solved a realistic example. A number of papers have considered the joint optimization of repair capacities, inventory investments, and repair sourcing. Alfredsson (1997) used an IP model to decide on repair allocations and target stocking levels. Rappold and Van Roo (2009), used a stochastic integer program to jointly determine repair facility location and inventory allocation. Basten *et al* (2012, 2015) introduced an IP LORA model and considered the joint inventory-allocation problem which was solved through both sequential and iterative methods. The analysis of the joint problem has been extended to consider more detailed models for a single location that allow for expediting repairs to an alternative source (Basten *et al*, 2014). The stream of papers by Basten *et al*, are the closest to our research. Their problem formulation is more general than ours and we note (as pointed out in Basten *et al*, 2011), issues such as repair complexity through multiple failure modes, multiple repair locations, finite repair resources, and the outsourcing of repairs can be accommodated within the mathematical programming formulation of LORA models through appropriate changes to input data and re-definition of decision variables. We take a different solution approach. In particular, we develop structural results for a stochastic, steady-state model, leading to a modification of the standard greedy heuristic used in practice and thus the proposed algorithm is highly scalable. Our model directly models full and partial repair sourcing to different locations through sourcing fraction decisions, and thus supports the movement to more flexible repair strategies.

There are several other relevant papers, which relate to stocking and sourcing decisions. Sleptchenko *et al* (2005) applied a priority queuing model to determine stocking/sourcing priority decisions in a finite repair capacity situation, where high-priority jobs interrupt low-priority jobs. They assumed fixed repair allocation/sourcing and did not consider repair cost differences. A different type of model was developed by Caggiano *et al*, (2006) to dynamically decide on repair and inventory allocations for operational day-to-day purposes.

The main contribution of this paper is the development of new closed-form functions, for expected backorders and their derivatives with respect to both sourcing and stocking decisions which support a model formulation that directly accounts for the flexible sourcing of repairs. Beginning with Simon (1971) and extended by Graves (1985), Axsater (1990) and Rustenburg *et al* (2003), prior exact solutions are based on numerical approaches for computing the backorder probability distribution from which performance metrics, such as expected backorders, can be computed. The closed-form results developed as a part of our analysis enable the development of a heuristic solution algorithm that solves the joint sourcing and stocking problem and leads to the definition of threshold-driven sourcing policies.

3. The model

In this section, we introduce the model, which includes a two-echelon network consisting of multiple (N) bases and a central depot (ie, $N+1$ locations). Our approach to dealing with the joint stocking and sourcing problem starts with a discussion of the model's assumptions and its dynamics in Section 3.1. In Section 3.2, we formulate a steady-state, nonlinear, stochastic optimization problem to minimize the stocking and repair costs subject to a fleet availability constraint by deciding on the amount to stock for each part-location combination (based on an order-up-to target stocking level policy) and target repair sourcing among the different locations. To solve this problem, we develop, in Section 3.3, closed-form expressions for the expected number of backorders and their derivatives (or differences), with respect to sourcing and stocking decisions, assuming that repair costs are equal and in Section 3.4, we use these expressions to characterize the optimal sourcing and stocking policies. Finally, in Section 3.5, we establish convexity that enables us to define threshold-driven sourcing policies through analysis of the Karkush–Kuhn–Tucker (KKT) conditions. In Section 4, we extend to the general case where repair costs may be different across bases and introduce a greedy algorithm to solve the joint stocking and sourcing stocking problem.

3.1. The model's assumptions and dynamics

Repairable parts inventory is typically managed at the subsystem or Line Replaceable Unit (LRU) level and we consider multiple LRUs. We consider a system (eg, aircraft) to be available if all of its LRUs (hereafter referred to also as "parts") are in working order (that is, a single part's failure brings the system down), and we assume that each part appears once in the system's Bill of Materials (this assumption can be easily relaxed). Figure 1 illustrates the flow of parts within the repairable parts supply chain. The sequence of events is as follows: A failed part generates a demand for a good unit. The base replaces the failed part with a good part from its good inventory if one is available; otherwise, the demand is backordered at the base. The base sources the failed part repair with probability r_i, to its local repair site; otherwise, the repair is sourced to the depot repair facility (transshipments between bases are not allowed). A part, which is sourced for depot repair, is replaced with a good unit that the depot sends

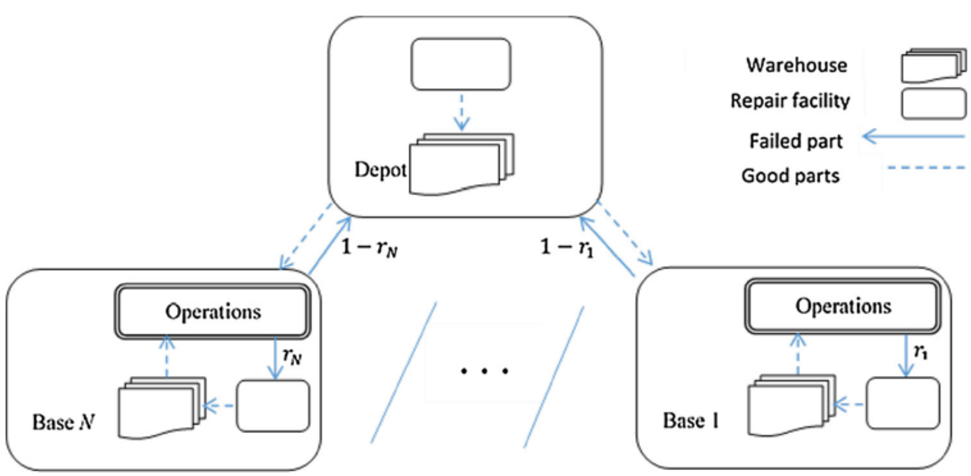

Figure 1 Repairable parts flows.

to the base, if available; otherwise, the demansd is backordered at the depot.

We make the standard assumption of independent failures across the system. M_i systems are supported by base i. A base must satisfy a predetermined service level for its systems, which defines a minimal number of operationally ready systems. The objective is to minimize the expected steady-state cost per period. A period corresponds to the planning horizon over which an organization wishes to manage its repair and purchase budget (it can range from one year to the expected end of life of the fleet of systems) and is determined by the time scale chosen for the demand and lead time parameters. Application of the model can support budget planning decisions for an upcoming period since purchase and repair sourcing decisions determine the repair and purchase budgets (eg, how much budget should be allocated to purchase new parts in order to achieve optimal target stocking levels and what repair capacity is necessary to satisfy demand for repair services based on the fractions of repairs to be allocated to the different locations). These costs, derived from our model solution, would be an input to typical budget planning process, especially those involving negotiation between the customer and the service provider. Typically, model parameter estimates would be updated for the upcoming planning period with the timing set to cover the period duration (eg, estimate demand per year for a 1-year horizon). The resulting steady-state solution would then be used to explore the cost and service tradeoffs for that period. We note that expected parts purchase and repair costs can be adjusted to account for the current inventory in the system as well as the repair pipeline status at the beginning of the planning period.

Our use of purchase costs rather than holding costs is consistent with what we have seen in practice in the context of a budget planning process where the budget for new part purchases and for repair capacity must be set periodically. We assume that the initial investment to set up the repair capability and inventory (eg, for the purchase of equipment, as well as for training and construction of the infrastructure) has been made, and that both the bases and the depot have overlapping repair capabilities. That is, each repair location, either a base or the depot, can repair a part for maximum flexibility. We note that the model can be applied to situations with partial overlap in repair capability, ie, a situation in which only some of the parts can be repaired at the bases. We note that use of partial overlap (through an upper bound on the repair fraction) can be used in the model to account for the mix of repair complexity that will be encountered as parts fail. The optimal solution to our model provides a recommended repair fraction, r_i, and the real-time allocation of repairs to bases could be guided by the result, ie, never assign a part to a location whose repair fraction is zero, always do so when the value is 1 and use an intermediate value to guide real-time allocations.

Our experience with practice also indicates that in today's environment, with outsourcing options and systems which are designed for efficient fault detection, isolation, and maintainability, the difference between defining a base level repair and a depot level repair is often based on historical reasons. Thus, maintenance managers can control repair sourcing by changing the allocation of repairs between the bases and the depot. As noted, this heightened level of flexibility for repair sourcing is now possible in modern systems due to recent developments in product design, maintenance information support systems and the tooling used to complete part repairs. Typically, a repair starts with finding and isolating the problem and deciding about the needed repair, which necessitates expertise and test equipment. Technology and the design of modern systems have led to more efficient fault detection and isolation and repair guidelines (which can be based on online step-by-step technical manuals, on-board sensors, real-time data analysis, or mobile test devices) which enhance the expertise available at all locations and thus support the overlap of repair capabilities.

Our model is appropriate when the following conditions hold: (1) The systems are modern, designed for efficient fault isolation; (2) the actual repair duration is only a small fraction of the total repair lead time, which is typical (eg, Sherbrooke, 2004 p. 23); (3) bases and the depot have comparable repair capabilities when they overlap. Under such conditions, it is reasonable to assume that the repair lead time and costs are primarily characteristics of the repair organization (eg, the cost of personnel, the estimated load of the repair location due to other commitments, historical information about lead times and costs) and thus will not be influenced by changing the repair sourcing fractions. We are therefore assuming that changing sourcing fractions does not imply changing the mix of repair complexities at a repair location. The assumption that the expected repair lead time and cost are fixed at each repair location is appropriate for the tactical perspective of our model in which sourcing decisions are made up front, based on commitments, from various possible repair sources (eg, a central depot and bases), to achieve repair lead times and costs.

We note that a re-formulation of the model so that the unit of analysis is a part with a specific failure mode and an associated level of repair complexity is possible. Doing so would allow costs and lead times to vary with repair sourcing for a part since different failure modes could have different parameters and different repair source allocations. This approach, however, introduces a problem in terms of the stocking decision which is made at the part level and not at the part/failure mode level. Thus, capturing the impact of part stocking which affects availability for multiple failure modes is not captured by the standard repairable inventory model that our formulation is based on. As a result, introduction of failure modes into our model would make it intractable. We therefore assume that the sourcing decisions derived by our model, which would be based on the commitments for cost and lead time, would be independent of realized repair complexities.

We note that committing to lead times and costs in maintenance repair contracts, however, is standard practice. This is also supported by research; for example, Arts *et al* (2014) use a deterministic expedited repair lead time that reflects a repair time that was agreed upon with an external repair provider.

We use the standard repairable inventory model assumptions: Poisson demand, backordering of excess demand at each stocking location, target stocking policy at each location, and ample repair capacity based on a steady-state, continuous review formulation.

For each part there are $2N+1$ decisions to be made within the planning period ie, the TSLs (Target Stocking Levels) at the depot (S_0) and at the bases $(S_1,\ldots,S_N)$ and the fraction of repairs to be carried out at a base repair facility $(r_1,\ldots,r_N)$. We assume that all parts are repaired either at a base or at the depot and none are condemned. Actual repair lead times can be generally distributed and system performance is a function of only the expected repair lead times (by Palm's theorem) and are assumed to be given (eg, each repair location commits to a standard for repair lead time).

3.2. Model formulation

We use the following notation throughout the paper:

i	Location: 0 for depot; 1,…, N for bases
j	An LRU
$S_{i,j}$	TSL for j at i
$\breve{S}_j = (S_{0,j},\ldots,S_{N,j})$	Vector of base TSLs for j
$R_{i,j}$	Number of j parts in repair
$r_{i,j}$	j repair fraction sourced to i
$\breve{r}_j = (r_{1,j},\ldots,r_{N,j})$	Vector of repair fractions for j
M_i	The number of systems supported from i
$\lambda_{i,j}$	Average number of demands per period per part per base
$C_{i,j}$	Part j repair cost at i
$L_{i,j}$	Repair lead time of j at i
$TT_{i,j}$	Transport time of j from the depot to Base i
$BO_{i,j}$	Expected backorders of j at i
p_j	Per unit part j purchase cost
$\tilde{A}, A_i$	Overall availability goal and base availability, respectively

Following Sherbrooke (2004), base availability is defined as $A_i=\prod_{j=1}^{K}\left(1-\frac{\mathrm{BO}_{i,j}}{M_i}\right)$, and the overall availability constraint can be written as $\frac{\sum_{i=1}^{N}A_iM_i}{\sum_{i=1}^{N}M_i}\geq\tilde{A}$. In many cases, requiring each of the bases to achieve equal availability is reasonable from a practical perspective and also simplifies our analysis. In such cases, the overall availability constraint can be replaced by N constraints, one for each base i, of the form: $A_i\geq\tilde{A}$. To simplify the analysis, we approximate the equal availability service target by requiring equal target fill rates at the bases, ie, $P\left(R_{1,j}>S_{1,j}-1\right)=\ldots=P\left(R_{N,j}>S_{N,j}-1\right)$ for each LRU j. This approximation can be relaxed at the cost of a more complicated analysis and is reasonable since availability and fill rate are both determined by the stocking and sourcing policy and both are common measures of customer service.

The objective function is the summation of repair and inventory-related costs over all bases and LRUs. We have assumed that fixed costs are absorbed and included in the repair costs, and thus we do not consider the possibility of economies of scale or finite capacity at the repair facilities. The objective is expressed as $\sum_{j=1}^{K}\sum_{i=1}^{N}\left\{\lambda_{i,j}\cdot\left[r_{i,j}\cdot c_{i,j}+(1-r_{i,j})\cdot c_{0,j}\right]+\left(S_{i,j}+S_{0,j}\right)\cdot p_j\right\}$.

This formulation implies that the initial inventory is zero and thus, ignores holding costs. The zero inventory assumption can be relaxed by considering prior stock purchase investments, which are treated as a sunk cost. For example, if we introduce S_j as the existing total inventory then the formulation can be presented as $\sum_{j=1}^{K}\sum_{i=1}^{N}\left\{\lambda_{i,j}\cdot\left[r_{i,j}\cdot c_{i,j}+(1-r_{i,j})\cdot c_{0,j}\right]+\left(S_{i,j}+S_{0,j}-S_j\right)^{+}\cdot p_j\right\}$. In case one wants to account also for the disposal of excess inventory, in cases where it exists, we can add the term $\left(S_{i,j}+S_{0,j}-S_j\right)^{-}\cdot p_j$, to the formulation, where p_j is the per unit selling price of excess inventory. To keep the formulation simple, we assume zero initial inventory going forward. (Interested readers are referred to van Houtum and Kranenburg (2015), Section 2.8 for a discussion about inventory planning during the exploitation phase.)

The overall multiple parts problem can be decomposed into K single-LRU problems with an overall availability constraint that takes into account interactions between base availability, for LRU j as follows:

$$\min_{\breve{S}_j,\breve{r}_j}\sum_{i=1}^{N}\left\{\lambda_{i,j}\cdot\left[r_{i,j}\cdot c_{i,j}+(1-r_{i,j})\cdot c_{0,j}\right]+S_{i,j}\cdot p_j\right\}+S_{0,j}\cdot p_j \qquad (P1)$$

such that $\sum_{j=1}^{K}\mathrm{BO}_{i,j}\leq\Theta$, where

$$\Theta=-M_i\cdot\log\left(\frac{\tilde{A}\cdot\sum_{l=1}^{N}M_l-\sum_{l=1(l\neq i)}^{N}\left[\prod_{j=1}^{K}\left(1-\frac{\mathrm{BO}_{l,j}}{M_l}\right)\right]M_l}{M_i}\right),$$

$S_{i,j}\geq 0$ and integer for $i=0,\ldots,N;\ j=1,\ldots,K,$
$0\leq r_{i,j}\leq 1$ for $i=1,\ldots,N;\ j=1,\ldots,K.$

The following is a summary of our solution approach which is documented in the subsequent sections and the appendices.

(a) We develop closed-form expressions for expected backorders as a function of stocking and sourcing policy at both the depot and the bases.
(b) We then develop closed-form expressions for the partial derivative of these functions with respect to the stocking fraction decisions for any given stocking and sourcing

policy. This includes variation in backorders at one location due to changes in sourcing policy at another location.

(c) We then develop a closed-form expression for the total change in expected backorders due to a change in sourcing fraction at each location, based on the partial derivatives.

(d) Optimal sourcing policies for a given stocking policy are then developed based on an analysis of the Karush–Kuhn–Tucker conditions under the assumption of equal repair costs at the bases and the depot.

(e) This solution is used to characterize three classes of repair sourcing policies; ie, central at the depot, local at the bases, and mixed at both the depot and the bases under the assumption of equal repair costs and an equal availability goal at all of the bases, for a fixed stocking policy. The optimal sourcing policy is shown to be dependent on the relative (base vs. depot) lead times values. The solutions are characterized by a threshold value for the lead times.

(f) We then analyze a more general version of the problem where the assumption of equal repair costs and a fixed stocking policy is relaxed. This yields the joint solution to the stocking and sourcing problem and the solution now depends on the relative value of repair lead times and repair costs at the depot vs. the bases. Four exhaustive cases are analyzed. For some of the cases, the solution heuristic must be utilized.

(g) Finally, we conduct an extensive numerical study (over 8 million problem instances) which includes those cases where there are no closed-form analytical results. The gap between the solution derived by application of our results and algorithm are compared to solutions which were generated through complete enumeration for a subset of cases. Managerial insights were then developed.

3.3. Closed-form expressions

The closed-form expressions for expected backorders and their derivatives presented below are used to solve Problem (P1), to gain insights about the structure of optimal policies, and to develop a heuristic algorithm for solving the general joint stocking and sourcing problem. As the mathematics involved is straightforward but tedious, we present here only those results that are needed for the derivation of the stocking algorithm to solve (P1). The proofs are contained in the online appendix. For conciseness, we do not include results regarding the first differences with respect to $\breve{S}$ as their use is limited to the heuristic algorithm.

We begin with a closed-form expression for backorders at a base. For readability, we suppress the subscript j. For the remainder of the paper, unless noted otherwise, we present results that pertain to a single LRU.

Proposition 1 *For any stocking (S_0, S_i) and sourcing $\breve{r}$ policy, the expected number of backorders at Base i is*

$$BO_i\left(S_0, S_i, \breve{r}\right) = \lambda_i \cdot \rho_i - S_i + e^{-\{\lambda_i\rho_i\}} \sum_{k=0}^{S_i-1} \frac{\{\lambda_i\rho_i\}^k}{k!}(S_i - k),$$

$$\textit{where} \quad \rho_i = r_i[L_i - (TT_i + L_0)] + (TT_i + L_0) + \frac{(1-r_i)}{\sum_{j=1}^{N}\lambda_j(1-r_j)} \cdot \left\{ -S_0 + e^{-\sum_{j=1}^{N}\lambda_j(1-r_j)\cdot L_0} \cdot \sum_{n=0}^{S_0-1} \frac{\left[\sum_{j=1}^{N}\lambda_j(1-r_j)\cdot L_0\right]^n}{n!} \cdot (S_0 - n) \right\}.$$

In general, the multi-echelon stocking problem, to determine TSLs is not jointly convex in S_0 and the $S_i's$ and therefore, solutions to the overall problem, (P1), require a heuristic algorithm, except for those special cases where structural results can be derived.

Next, we present the closed-form expression for the derivatives of expected backorders with respect to the repair sourcing fraction (the proofs of Propositions 1–3 are added in an online supplement.

Proposition 2 *For stocking and sourcing policy $(S_0, S_i, \breve{r})$, the partial derivative of expected backorders at Base i, with respect to r_i is,*

$$\frac{\partial BO_i(S_0, S_i, \breve{r})}{\partial r_i} = \lambda_i \left[(L_i - TT_i - L_0) + \frac{\sum_{k=1\,(k\neq i)}^{N} \lambda_k(1-r_k)}{\left[\sum_{k=1}^{N}\lambda_k(1-r_k)\right]^2} \cdot S_0 \cdot P(R_0 \geq S_0 + 1) + L_0 \cdot P(R_0 \leq S_0 - 1) \right] \cdot P(R_i > S_i - 1).$$

Proposition 3 *For stocking and sourcing policy $(S_0, S_k, \breve{r})$, the partial derivative of the expected backorder at Base k, with respect to the allocation fraction r_i (whereas $k \neq i$) is,*

$$\frac{\partial BO_k(S_0, S_k, \breve{r})}{\partial r_i} = -S_0 \cdot \frac{\lambda_k(1-r_k)\lambda_i}{\left[\sum_{j=1}^{N}\lambda_j(1-r_j)\right]^2} \cdot P(R_0 > S_0) \cdot P(R_k > S_k - 1) < 0$$

$$\forall k \text{ where } k \neq i.s.$$

Note that changing r_i affects backorders at all of the bases that use the depot.

Using Propositions 2 and 3, we introduce Proposition 4, which enables calculation of the overall change in the expected number of backorders associated with changing the sourcing fraction r_i:

$$\Delta BO = \frac{\partial BO_i(S_0, S_i, \breve{r})}{\partial r_i} + \sum_{k=1\,(k\neq i)}^{N} \frac{\partial BO_k(S_0, S_k, \breve{r})}{\partial r_i}.$$

Proposition 4 *For a stocking and sourcing policy* $(\breve{S}, \breve{r})$*, the total change in the expected number of backorders due to changing sourcing* r_i *is,*

$$\Delta BO = \Bigg\{ \lambda_i (L_i - (TT_i + L_0 \cdot P(R_0 > S_0 - 1))) + \frac{\lambda_i \cdot \left[\sum_{l=1(l \neq i)}^{N} \lambda_l \cdot (1 - r_l)\right]}{\left[\sum_{l=1}^{N} \lambda_l \cdot (1 - r_l)\right]^2} \cdot S_0 \cdot P(R_0 > S_0) \Bigg\} \cdot P(R_i > S_i - 1) - \sum_{n=1(n \neq i)}^{N} \frac{\lambda_n \cdot (1 - r_n) \cdot \lambda_i}{\left[\sum_{l=1}^{N} \lambda_l \cdot (1 - r_l)\right]^2} \cdot S_0 \cdot P(R_0 > S_0) \cdot P(R_n > S_n - 1).$$

If service levels at all the bases are constrained to be equal, then the overall change becomes,

$$\Delta BO = \lambda_i (L_i - (TT_i + L_0 \cdot P(R_0 > S_0 - 1))) \cdot P(R_i > S_i - 1).$$

Proof of Proposition 4 The proof follows from summation of the expressions in Propositions 2 and 3.

ΔBO can be positive, negative, or zero. Increasing r_i reduces the number of backorders at all of the other bases that use the depot, so the term $\sum_{k=1\,(k \neq i)}^{N} \frac{\partial BO_k(S_0, S_k, \breve{r})}{\partial r_i}$ is always negative. This can be intuitively understood as taking some of the load from the depot and allowing it to provide better service to the other bases.

In the next subsection, we use Karush–Kuhn–Tucker (KKT) conditions to analyze optimal sourcing policies for a given stocking policy under the assumption (to be relaxed in a subsequent section), of equal repair costs at the bases and the depot (which means that it is sufficient to minimize total backorders). For the KKT conditions to be necessary and sufficient for optimality, one needs to show convexity of the sum of backorders with respect to $\breve{r}$. In other words, we need to prove that the Hessian of the sum of backorders with respect to $\breve{r}$ is positive semi-definite. Thus, it is necessary to develop closed-form expressions for $\frac{\partial^2 \left(\sum_{l=1}^{N} BO_l(S_0, S_l, \breve{r})\right)}{\partial^2 r_i}$ and $\frac{\partial^2 \left(\sum_{l=1}^{N} BO_l(S_0, S_l, \breve{r})\right)}{\partial r_i \partial r_j}$. These complicated expressions are presented in Part A of the Appendix. We also note that the expected number of backorders at any base i is a decreasing, convex function of either S_0 or S_i separately (the proof of this result is omitted since its use is limited to the developed heuristic algorithm).

3.4. Characterization of optimal sourcing policies

We define three classes of sourcing policies: a central policy in which repairs are performed at the central depot, a local policy in which repairs are performed at the bases, and a mixed policy where repairs are shared between the bases and the depot. The nature of optimal sourcing policies is characterized, based on analysis of the KKT conditions. For tractability, we assume the following: (1) Bases maintain the same availability goal, which is approximated by equal fill rates, ie, $P(R_i > S_i - 1) = P(R_l > S_l - 1)$ for all $i \neq l$; (2) repair lead times in all bases are either equal to, shorter than, or longer than the depot lead time, ie, $L_i \geq L_0 + TT$ or $L_i < L_0 + TT$ for $i = 1, \ldots, N$; and (3) all orders up to levels S_0 and S_i are fixed. The first two assumptions are reasonable and generally hold in practice and the third will be relaxed in the next section when we solve the joint problem. Recall that we assume, for now, equal repair costs at the depot and at the bases—this assumption also is relaxed subsequently.

Under the assumptions noted above, it is sufficient to find a sourcing policy that minimizes the sum of backorders across all the bases, thus, the problem to be solved is

$$\min_{r} \sum_{i=1}^{N} BO_i\left(S_0, S_i, \breve{r}\right) \qquad (P2)$$

such that, for all $i = 1, \ldots, N{:}\, 0 \leq r_i \leq 1$.

For KKT conditions to be not only necessary but also sufficient for optimality, the objective of Problem (P2) has to be convex (note that its domain is convex). We have shown that the Hessian

$$\nabla^2 f\left(\breve{r}\right) = \begin{pmatrix} \frac{\partial^2 \left(\sum_{i=1}^{N} BO_i(S_0, S_i, \breve{r})\right)}{\partial^2 r_1} & \cdots & \frac{\partial^2 \left(\sum_{i=1}^{N} BO_i(S_0, S_i, \breve{r})\right)}{\partial r_1 \partial r_N} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 \left(\sum_{i=1}^{N} BO_i(S_0, S_i, \breve{r})\right)}{\partial r_N \partial r_1} & \cdots & \frac{\partial^2 \left(\sum_{i=1}^{N} BO_i(S_0, S_i, \breve{r})\right)}{\partial^2 r_N} \end{pmatrix}$$

is positive semi-definite with respect to $\breve{r}$. Details about the KKT conditions and their analysis are provided in Part B of the appendix. Based on this analysis we present Theorem 1, which characterizes optimal sourcing policies for three classes of relationships between repair lead times at the bases and at the depot: (1) $L_i > L_0 + TT$, (2) $L_i = L_0 + TT$, and (3) $L_i < L_0 + TT$.

Theorem 1 *If repair costs at all locations are equal, then the optimal sourcing policy, which minimizes the sum of backorders* $\sum_{i=1}^{N} BO_i\left(S_0, S_i, \breve{r}\right)$ *is*

1. *central, where all repairs are sourced to the depot for* $L_i > L_0 + TT$,
2. *local, mixed, or central where all repairs are sourced to the bases for* $L_i = L_0 + TT$ *and* $S_0 = 0$*or central for* $L_i = L_0 + TT$*and* $S_0 > 0$, *and*
3. *when* $L_i < L_0 + TT$*the optimal sourcing policy is either local, or central or mixed. When it is mixed then sourcing fractions can be found numerically by*

solving the equations $L_i = TT_i + L_0 \cdot P(R_0 > S_0 - 1)$ *for* $i = 1, \ldots, N$.

Proof of Theorem 1 Part (1) is obvious, and Part (2) follows the solution of the KKT conditions shown in the appendix. The proof for Part (3) builds on the KKT solutions by which, in an optimal policy, r_i is either 0 or 1 for all the bases, or is mixed. Thus, we use the condition on the Lagrangian (Λ) derivative $\frac{\partial \Lambda(\breve{r})}{\partial r_i}$, and assume equal service levels at the bases. After some mathematical manipulations, we get the following equation system: $\frac{\partial \Lambda(\breve{r})}{\partial r_i} = \lambda_i(L_i - (\mathrm{TT}_i + L_0 \cdot P(R_0 > S_0 - 1))) \cdot P(R_i > S_i - 1) = 0$ for $i = 1, \ldots, N$ that can be solved numerically by solving $L_i = \mathrm{TT}_i + L_0 \cdot P(R_0 > S_0 - 1)$ for $i = 1, \ldots, N$.

3.5. Threshold policies for the case of faster bases ($L_i < L_0 + \mathrm{TT}$)

Let us analyze Part (3) of Theorem 1, which illustrates a fundamental tradeoff between efficiency and risk pooling. From a risk-pooling perspective, central stocking and sourcing is preferred, and in contrast, when considering only repair lead times, a local policy is optimal. We are able to develop thresholds for base repair lead times, which determine which case applies. In this context, the situation that we have in mind is that of an organization that repairs its parts locally but can switch sourcing of repairs to a central depot (recall that we are still assuming equal repair costs at all locations). In such a situation, it is imperative to challenge in-house local repair facilities by notifying them that achievement of a repair lead time goal is necessary if they wish to remain attractive as a repair source (in fact, we faced this exact situation in a recent consulting project in the commercial airline industry).

We start with the case of identical bases, that is $L = L_i$ for $i = 1, \ldots, N$. Under this assumption, Lemma 1 defines a threshold policy, where the threshold is denoted as $\bar{L}$. (The proofs of Lemmas 1 and 2 use the closed-form expressions that were developed earlier along with the KKT conditions. The detailed proofs are in Part C of the appendix.)

Lemma 1 *For N identical bases, equal repair costs, and a total stocking quantity* S', *a local policy is preferable to a central policy if* $L < \bar{L}$. *The threshold* $\bar{L}$ *is the solution to the equation*

$$\mathrm{BO}_1(0, S_1, 1, \ldots, 1) + \mathrm{BO}_2(0, S_2, 1, \ldots, 1) + \ldots + \mathrm{BO}_N(0, S_N, 1, \ldots, 1) = \mathrm{BO}_1(S'_0, S'_1, 0, \ldots, 0) + \ldots + \mathrm{BO}_N(S'_0, S'_N, 0, \ldots, 0),$$

where S'_i, *for* $i = 0, 1, \ldots, N$, *are optimal stocking levels under a central policy. If* $S_i = \frac{S'}{N}$ *is an integer, then it is the optimal stocking allocation for each of the bases, under a local policy. Otherwise,* S_i *is not an integer, and thus the stocking targets are determined as follows:* $\lfloor \frac{S'}{N} \rfloor$ *is allocated to each base, and there is a remainder; add one additional part to each base until all of the remainder is allocated.*

We extend Lemma 1 to the case of non-identical bases (with respect to lead times) in Lemma 2.

Lemma 2 *For N non-identical bases, equal repair costs, and a total stocking quantity* S':

(1) *a central policy is not optimal when* $L_i \leq \bar{L}_i$ *holds for at least one of the bases, and*
(2) *a local policy is optimal when* $L_i < \bar{L}_i \quad \forall i = 1, \ldots, n$.

The thresholds $\bar{L}_i$ *are determined as a solution to the equations* $\bar{L}_i = \mathrm{TT}_i + L_0 \cdot P(R_0 > S'_0 - 1)$, *where* S'_0 *is the optimal stocking target at the depot under a central policy.*

Lemmas 1 and 2 enable finding thresholds such that if base repair times are lower than these thresholds, central policies are not optimal. The opposite also holds; ie, if base repair times are longer than the thresholds, local policies are not optimal.

4. Analysis of the Problem and a Heuristic Solution Algorithm

In this section, we relax the assumptions that the bases and the depot have equal repair costs and that the stocking policy is fixed (ie, by analyzing the partial derivatives for the backorders with respect to the depot and base TSLs). This enables finding solutions that take into account the tradeoff between lead times, inventory costs, and repair costs. We consider 4 cases, based on relative (ie, base vs. depot) values for repair lead times and unit repair costs. Recall that the original problem is non-convex, thus its general solution must be based on heuristics apart from instances for which we found structural properties that enable finding the optimal stocking and sourcing policy directly. The heuristic solution algorithm that was developed can solve the most general version of the problem and is a variant of the greedy algorithm, used in practice, to solve the multi-echelon stocking problem and can be extended to multiple LRUs.

We describe below each one of the four cases, provide a motivation for their existence, and present the argument for the suggested solution. Then, we present Table 1, which summarizes the results.

Case 1 Both repair lead time and cost are higher at the bases than at the central depot, $L_i > L_0 + \mathrm{TT}_i$ and $c_i > c_0$. This case describes a situation in which a central repair facility, either a depot or an external subcontractor, is faster and more efficient in performing repairs. Such a situation may occur, for example, if the depot specializes in a complex system (eg, engine repairs) and/or serves multiple customers such that its repair volume is large enough to justify investing in efficient repair equipment and training personnel to a level that allows for quick and cost-effective maintenance. In Theorem 1 (Part 1), we proved that for slower bases central repair is optimal. Since central repair also minimizes repair costs by repairing everything at the cheaper location (the depot) it is clearly the optimal policy. Thus, we set $\breve{r}_i = 0 \quad i = 1, \ldots, N$ and determine the associated stocking policy.

Case 2 Repair lead time is shorter at the central depot but the repair cost is higher there, ie, $L_i > L_0 + \mathrm{TT}_i$ and $c_i < c_0$. This case describes a classical "pay more for quicker service" situation. It occurs when central repairs are outsourced to an original equipment manufacturer (OEM) who provides the best repair option (from the perspective of repair quality and lead time), but is more expensive.

Ultimately, this tradeoff between cost and time has to be resolved on a case-by-case basis. Our analysis is directed to applications where strategic and tactical inventory and repair resource deployment decisions must be made and where the scale can be enormous (ie, many thousands of stock-part combinations). Our algorithm (whose development is detailed at the end of this section) is based on the closed-form expressions for the derivatives of the backorders with respect to repair fractions and stocking levels, and thus can deal with large-scale joint stocking and sourcing problems in multiple location settings.

Case 3 Both repair lead times and repair costs are lower at the bases than at the central depot, $L_i < L_0 + \mathrm{TT}_i$ and $c_i < c_0$. We have seen this case in situations where the part's technology is simpler (ie, structural parts) or when there are efficient fault identification procedures and equipment. In such circumstances, bases with cheaper workforce costs and less demand (compared to a central depot) have a chance to conform to this case. We build on the results of Theorem 1 Part (3) that analyzes such a situation for equal repair costs. In addition, we use Lemmas 1 and 2 to state that the optimal policy is local under the following conditions: (1) $L < \bar{L}$ for identical bases, or (2) $L_i < \bar{L}_i$ for all the bases when the bases are non-identical.

We can rule out central policy as optimal when $L_i < \bar{L}_i$ exists for at least one of the bases.

As the difference $c_0 - c_i$ increases, a local policy becomes more attractive. For the cases, which are not covered by the thresholds, one should use the solution algorithm.

Case 4 Repair lead times are shorter at the bases, but the repair costs there are higher, ie, $L_i < L_0 + \mathrm{TT}$ and $c_i > c_0$. In general, here all 3 policies—local, central, and mixed—may be optimal. We use Lemma 1 to rule out a local policy in certain cases. Particularly, an optimal policy is either mixed or central if $L > \bar{L}$. We note that as $c_i - c_0$ increases, a central policy becomes more attractive compared to mixed policies. The same conclusion does not apply to non-identical cases and in such cases, the solution algorithm must be used.

Table 1 summarizes these results, which provide guidance for jointly setting stocking and sourcing policies for the 4 cases, subject to an equal service level requirement at all of the bases. Cases for which the assumption of equal base service does not hold are not covered in Table 1 since we lack structural insights for them, but they are covered by the proposed solution algorithm.

As noted, there are instances, notably for Cases 3 and 4, for which the solution algorithm has to be used. The standard solution algorithm for solving ME models is a greedy heuristic based on a marginal analysis that evaluates the benefit of stocking one more item at the bases or at the depot. The algorithm developed in this paper chooses, at each iteration, a decision that optimizes the marginal benefit defined as the change in the total number of backorders due to a decision (eg, to either change a repair fraction or to change a stocking level at a location) divided by the incremental change in cost due to that decision. The algorithm may be used when the optimal policy cannot be determined analytically and when mathematical programming approaches are not feasible due to problem scale. The algorithm can handle multiple bases and multiple parts without assuming equal service levels at the bases, and it utilizes the structural and analytical closed-form expressions that were developed for the derivatives. Below, we summarize the solution algorithm (more details, including the starting conditions, are presented in Part D of the appendix):

Table 1 Repair sourcing and stocking policies

Values	$L_i < L_0 + \mathrm{TT}$	$L_i > L_0 + TT$
$c_i < c_0$	If $L < \bar{L}$(identical bases) or $L_i < \bar{L}_i$, a local policy is optimal. Otherwise, use the solution algorithm (Case 3)	Use the solution algorithm (Case 2)
$c_i > c_0$	If $L > \bar{L}$(identical bases), a mixed or a central policy is optimal. Otherwise, use the solution algorithm (Case 4)	Central policy is optimal (Case 1)

(1) Start with zero stock at the bases and depot and set the repair allocation fraction so that the objective function of (P1) achieves its lowest value (eg, if $c_i > c_0$ set $r_i = 0$). Calculate the availability.
(2) As long as the availability goal is not satisfied then there are three decision options: to increase the stock at the bases, to increase the stock at the depot, or to allocate more repairs to the depot/bases. Each of these options increases cost and decreases the sum of backorders. The algorithm chooses at each iteration, the decision that leads to the largest decrease in backorders per unit of cost.
(3) When the availability goal is exceeded then the algorithm checks if adjusting the allocation fractions to make it binding improves the solution.

We note that this algorithm can be easily enhanced in several ways, but this is left for future research.

5. A numerical study

An extensive numerical study complements the analysis and provides insights for those instances for which there are no closed-form analytical results. The study also examines the quality of the suggested solution algorithm in terms of the gap between recommended and optimal solutions.

5.1. Design of experiment

The experiment examines a supply chain with two bases and a single depot, and includes a full range of realistic parameter combinations. Overall, 8.64 million individual single-item problem instances were solved, each representing a different type of part. We conducted the calculations on a 12-core computer cluster using a Linux operating system. The suggested solution algorithm, coded in Matlab, solved each one of the problem instances in an average time of 0.016 s. For a representative sample of 2000 of the instances, we also conducted an exhaustive search procedure. The search procedure explores, for each instance, all relevant combinations of the decision variables (base and depot TSLs and repair allocations), and chooses the one that minimizes cost. For this experiment, the repair allocation values were discretized using a 0.01 step size. The solution algorithm performed quite well when compared to the exhaustive search procedure (an average of 0.2% difference in cost, a maximal error of 4.9%). Performance of the algorithm could be improved by applying backtracking procedures (eg, convexification, see Sherbrooke, 2004, p. 53), and parallel searches with different sourcing fractions but based on these results we chose not to explore this option.

The parameters for the experiment were set according to what we have observed in real-world aerospace and defense environments. The unit purchase cost p was normalized to 1 and repair costs, c_i, c_0, were set to vary from 5 to 50% of the purchase cost. For higher repair costs, it may be reasonable to condemn the part upon failure. Repair lead times L_i, L_0, varied between 20 and 110 days and thus represent a wide spectrum of LRU repair lead times. The transfer time between depot and base was set to 7 days and the average number of yearly repair demands that each base generates for the particular LRU ranged from 10 to 100. In the subsequent graphs, the time unit is years. The availability goal values ranged from 90 to 99.5 %, in order to illustrate interactions between sourcing, stocking, and pooling.

5.2. Results

We present specific results from which one can extract managerial insights that can lead to sourcing or stocking policy recommendations. For more complex scenarios, it is necessary to use the proposed solution algorithm. Section 5.2.1 overviews the results for all of the four cases, and then we focus on Case 3, in which the bases have shorter repair lead times than the depot (in Section 5.2.2). Section 5.2.3 presents benefits from using the joint stocking and sourcing algorithm compared to using simple benchmark policies.

5.2.1. The combined effect of repair cost and repair lead time differences Figure 2 is a high-level overview of the results for two identical bases as a function of $\Delta L = L_0 + \mathrm{TT} - L_i$ and $\Delta c = c_0 - c_i$. It demonstrates how the optimal policy is driven by the relative values of repair cost and repair lead time and validates our structural analysis and the guidelines presented in Table 1. For example, results indicate that for Case 1, a central policy is optimal. The other cases demonstrate tradeoffs between lead time, repair cost, and pooling, leading to all 3 possible policies.

5.2.2. Tradeoffs between pooling, availability, and demand rate Case 3 is interesting because it captures the tradeoff between the cost and lead time advantage that the bases enjoy when stocking and sourcing are local versus the pooling advantage that a central policy can provide. Figures 3 and 4 present results for Case 3.

Each point in Figure 3 represents an optimal policy, either local or central, for an availability goal and lead time difference combination. The lead time difference is calculated as ΔL. For small differences in repair lead times (eg, $\Delta L/(L_0 + \mathrm{TT}) < 15\%$) the pooling effect dominates, leading to superiority of central policies over local ones. As lead times at the bases become shorter, local policies prevail. As the availability goal increases, pooling becomes more important. Figure 3 demonstrates that the optimal policy, for a fixed ΔL value, may switch from local to central as availability increases (eg, for $\Delta L \cong 0.6$ a central policy becomes optimal for availability goals higher than 0.97). We denote the approximate line that separates central and local policies as

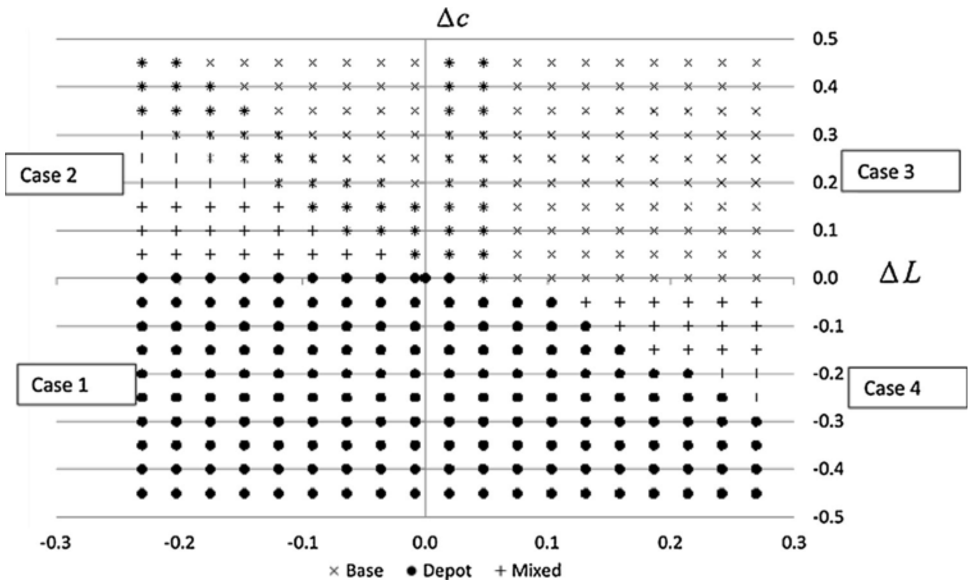

Figure 2 Optimal stocking and sourcing solutions as a function of the difference between repair lead times and repair costs. For Cases 2 and 3, there is sometimes an overlap between base and mixed policies (for different parts).

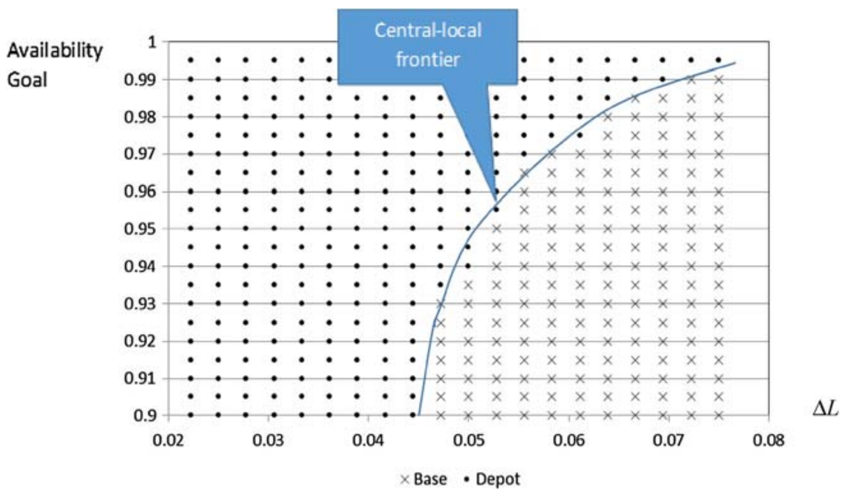

Figure 3 Optimal stocking and sourcing policies for two identical bases. Each point represents an optimal policy for a combination of an availability goal and a lead time difference. The parameters for the figure are $\lambda_1 = \lambda_2 = 60$, $\Delta L = L_0 + \mathrm{TT} - L_i$, $c_0 = c_1 = c_2 = 0.2$.

the central–local frontier. Repair cost differences between the depot and the bases will shift the central–local frontier. For example, an increase of Δc when base repair gets cheaper will shift the central–local frontier to the left, and local policies will become more attractive over a larger portion of the results, compared to Figure 3.

Figure 4 presents optimal stocking and sourcing policies for different $\lambda_1 + \lambda_2$ and ΔL combinations. With an increase in ΔL, local policies become more attractive, as expected. It is interesting to see that as the sum of demands increases, for a large enough ΔL, local policies are favorable over central policies, and for small $\lambda_1 + \lambda_2$ values, central policies are typically optimal. This follows since pooling will be more important in low-demand situations. Increasing Δc will shift the central–local frontier to the left, marking that local policies become dominant over central ones, compared to scenarios with smaller Δc values. (We note that near the central–local frontier, the policies may yield almost similar cost, which sometimes lead to a slight non-monotone behavior due to round offs in our algorithm.)

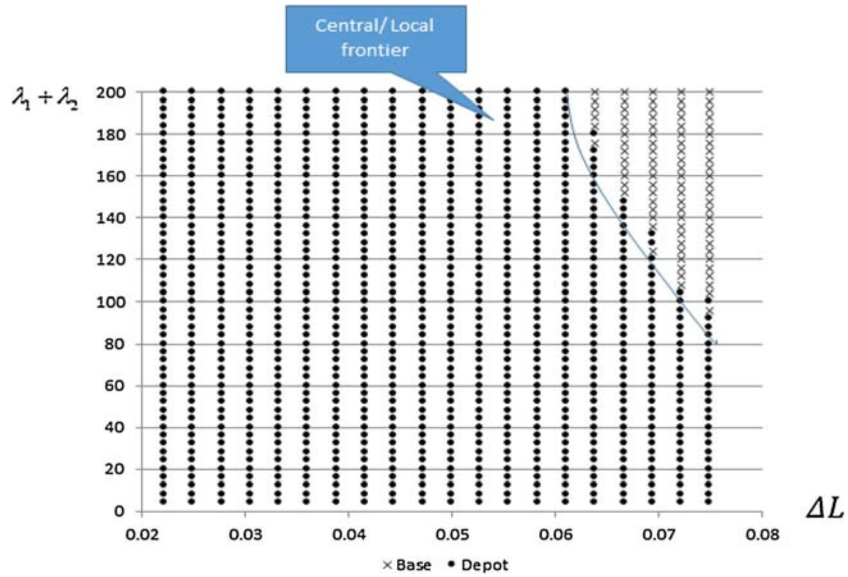

Figure 4 Optimal stocking and sourcing policies for two identical bases. Each point represents an optimal policy for a combination of the sum of demands and a lead time difference. The parameters for the figure are $\Delta L = L_0 + \mathrm{TT} - L_i, \quad c_0 = c_1 = c_2 = 0.2, \tilde{A} = 0.99.$

5.2.3. Possible benefits of the joint optimization of stocking and sourcing policies To illustrate possible benefits from the suggested joint stocking and sourcing policies, we compared the results with a widely practiced class of management policies that base repair sourcing decisions on a single attribute, ie, either the part repair cost or its lead time. Under a time-based policy, all parts with shorter base lead times are repaired locally and under a cost-based policy all parts with lower base repair costs are repaired locally. When summarizing the results over all of the experiments, the suggested approach, based on our algorithm and structural results, outperforms policies that source repairs to the faster location, with total costs being lower, on average, by 11.6 %. The suggested policy was found to be cheaper, on average, by about 7.1 % when compared to policies that source repairs to the lower cost repair location. These results suggest that substantial cost savings could be derived, especially when taking into consideration that for multiple parts the cost saving advantage will be higher.

These insights can have several practical benefits. In implementations of service parts planning systems, we have observed that there is considerable effort devoted to reviewing proposed stocking and sourcing decisions before they are actually implemented by passing them on to the real-time transactional system (which is usually supported by an ERP system). With our results, it is possible to identify the subset of parts where the proposed decisions can easily be verified and hence manual examination is not necessary. This could greatly enhance the productivity of inventory managers who are tasked with a real-time review of planning system recommendations. A second benefit is more strategic. In current competitive maintenance environments where outsourcing is frequent, the computed threshold values can be shared with those employees whose efforts help to determine the repair lead time. These thresholds can serve as part-specific lead time improvement targets that, if achieved, will enable a repair location to retain its work load.

6. Summary

This paper has introduced a model of the multi-echelon repairable supply chain with multiple bases and a single central depot. Our analysis differs from traditional multi-echelon models that only optimize target stocking levels at supply chain locations, as we also consider repair sourcing decisions. The suggested model follows recent trends in the maintenance industry toward more flexible repair sourcing and is appropriate when repair capabilities across the repair locations overlap, ie, parts can be repaired either locally or at the central depot.

We analyze three types of repair sourcing policies: central where repairs are performed in the central depot; local in which all the failures of a base are repaired locally and the stock is allocated between the bases; and a mixed policy that sources a fraction of the repairs to the depot with the remainder being repaired locally.

We develop a heuristic solution algorithm, whose performance was validated through an extensive numerical study, to solve the general joint stocking and sourcing problem to near-optimality. If the bases are required to provide an equal level of service, then we are able to develop structural results that guide the selection of good, and sometimes optimal, stocking, and sourcing policies. We partition the problem into 4 cases based on relative values at the depot and bases for repair lead times and repair costs. When depot repairs are cheaper and faster, the optimal policy is central. An "opposite" situation arises when the bases are cheaper and faster. In such cases, there is a tradeoff between the faster bases' repair times that encourage a local policy, and risk pooling that encourages a central policy. Our analysis leads to a threshold policy for optimal sourcing policies. Solutions for the other two cases, in which the faster location (either the bases or the depot) is the more expensive, are based on a threshold policy and on the use of the suggested solution algorithm.

Future research directions could deal with different types of bases, with some being faster and some being slower when compared to the depot—a scenario that was not included in our model. The algorithm can be applied to multiple parts and could be extended to handle additional factors, such as an overall budget constraint or multiple classes of customers with different service requirements. The suggested model can also be extended to include multiple indenture levels which accounts for parts and sub-assemblies that are used in the repair of LRUs. Another research direction could be to enhance and improve the heuristic algorithm. Such research can apply backtracking procedures, use diverse starting conditions (eg, with r_i equal to 0, 1 or intermediate values), etc.

References

Alfredsson P (1997). Optimization of multi-echelon repairable item inventory systems with simultaneous location of repair facilities. *European Journal of Operational Research* **99**(3):584–595.

Arts J, Basten R, and van Houtum GJ (2014). Optimal and heuristic repairable stocking and expediting in a fluctuating demand environment, Eindhoven University of Technology, Beta Working Paper series 446.

Axsäter S (1990). Modeling emergency lateral transshipments in inventory systems. *Management Science* **36**(11):1329–1338.

Basten RJI, Schutten JMJ and van der Heijden MC (2009). An efficient model formulation for level of repair analysis. *Annals of Operations Research* **172**(1):119–142.

Basten RJI, Schutten JMJ and van der Heijden MC (2011). Practical extensions to a minimum cost flow model for level of repair analysis. *European Journal of Operational Research* **211**(2):333–342.

Basten RJI, van der Heijden MC and Schutten JMJ (2012). Joint optimization of level of repair analysis and spare parts stocks. *European Journal of Operational Research* **222**(3):474–483.

Basten RJI, van der Heijden MC, Schutten JMJ and Kutanoglu E (2015). An approximate approach for the joint problem of level of repair analysis and spare parts stocking. *Annals of Operations Research* **224**(1):121–145.

Caggiano KE, Muckstadt JA and Rappold JA (2006) Integrated real-time capacity and inventory allocation for repairable service parts in two-echelon supply system. *Manufacturing & Service Operations Management* **8**(3):292–319.

Cohen M, Kleindorfer PR and Lee HL (1988). Service-constrained (s,S) inventory systems with priority demand classes and lost sales. *Management Science* **34**(4):482–499.

Feeney GJ and Sherbrooke CC (1966) The (s-1,s) inventory policy under compound poisson demand. *Management Science* **12**(5):391–411.

Gaver D, Isaacson K and Abell J (1993). Estimating aircraft recoverable spares requirements with cannibalization of designated items. *Rand Corporation monograph*.

Graves SC (1985) A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management Science* **31**(10):1247–1256.

Guajardo J, Cohen M, Kim S and Netessine S (2012) Impact of performance based contracting on product reliability: an empirical analysis. *Management Science* **58**(5):961–979.

Lee HL (1987). A multi-echelon inventory model for repairable items with emergency lateral transshipments. *Management Science* **33**(10):1302–1316.

Muckstadt JA (1973). A model for a multi-item, multi-echelon, multi-indenture inventory system. *Management Science* **20**(4):472–481.

Muller A, Marquez AC and Lung B (2008). On the concept of e-maintenance: review and current research. *Reliability Engineering System Safety* **93**(8):1165–1187.

Rappold JA and Van Roo BD (2009). Designing multi-echelon service parts networks with finite repair capacity. *European Journal of Operational Research* **199**(3):781–792.

Rustenburg WD, van Houtum GJ and Zijm WHM (2000). Spare parts management for technical systems: resupply of spare parts under limited budgets. *IIE Transactions* **32**(10):1013–1026.

Rustenburg WD, van Houtum GJ and Zijm WHM (2003) Exact and approximate analysis of multi-echelon, multi-indenture spare parts systems with commonality. In: Shanthikumar GJ, Yao DD and Zijm WHM (eds) *Stochastic modeling and optimization of manufacturing systems and supply chains* (pp. 143–176). Kluwer Academic: Dordrecht.

Saranga H and Kumar UD (2006). Optimization of aircraft maintenance/support infrastructure using genetic algorithms—level of repair analysis. *Annals of Operations Research* **143**(1):91–106.

Sherbrooke CC (1968). Metric: A multi-echelon technique for recoverable item control. *Operations Research* **16**(1):122–141.

Sherbrooke CC (2004). *Optimal Inventory Modeling of Systems. Multi-echelon Techniques*, 2nd ed., Kluwer: Dordrecht.

Simon RM (1971). Stationary properties of a two-echelon inventory model for low demand items. *Operations Research* **19**(3):761–773.

Sleptchenko A, van der Heijden MC and van Harten A (2005). Using repair priorities to reduce stock investment in spare parts networks. *European Journal of Operational Research* **163**(3):733–750.

van Houtum GJ and Kranenburg B (2015). Spare Parts Inventory Control under System Availability Constraints (Vol. 227). Springer: Boston.

Wang Y, Cohen MA and Zheng Y (2000). A two-echelon repairable inventory system with stocking-center-dependent depot replenishment lead times, *Management Science* **46**(11):1441–1453.

Received 4 September 2014; accepted 12 August 2016

Electronic supplementary material The online version of this article (doi:10.1057/s41274-016-0108-0) contains supplementary material, which is available to authorized users.