



AN OPTIMIZATION FRAMEWORK FOR AIR FORCE LOGISTICS MODELS

Retsef Levi
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

11/04/2014
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)

- FA9550-11-1-0150 - An Optimization Framework for Air Force Logistics Models

Final Report - September, 2014

1 Introduction

In this report we outline the work that has been performed under grant - FA9550-11-1-0150 awarded to us by the *Optimization and Discrete Mathematics Program* in the Mathematics, Information and Life Sciences Directorate of the *Air Force Office of Scientific Research* (AFOSR) during the period 7/1/11-6/30/14 (2nd year of the grant). (This grant was a continuation of award FA9550-08-1-0369.) We shall summarize the body of work that has been performed, list the papers that have been produced and their publication status, and highlight the potential impact on the Air Force (see also appendix).

As the name of the grant suggests, the goal of this project is the creation of a new framework of models and algorithms to address grand future challenges of the US Air Force logistics. In particular, the work under this grant was focused on the following domains:

1. *Maintenance and inventory models with submodular cost functions.*
2. *Maintenance of low observable aircraft.*
3. *Online models for supply chain and logistics models.*
4. *Distribution networks and supply chain management*

The grant have been used to support the work of 7 graduate (6 PhD and 1 master) students, including a recently graduating master student from the U.S. Air Force. Thus far, the grant has partially supported work on 26 papers [8, 21, 13, 9, 2, 20, 6, 10, 11, 12, 23, 27, 14, 7, 19, 24, 16, 17, 15, 25, 26, 3, 28, 1, 18, 4, 5] primarily in the areas mentioned above, as well as in additional areas such as healthcare networks. In addition, we have continued to develop contacts and collaborations with Air Force logistics units and personnel (see list in the appendix). We next discuss some of the major results and contributions related to the work supported by the grant.

2 Major Results and Contributions

Next we describe the major results and contributions during the duration this grant. For each topic we describe the nature of the work and the various papers that were produced:

Models for maintenance of modular systems. Traditionally, the academic and practice work on Air Force logistics have been primarily focused on issues related inventory and spare parts availability. However, in recent years it has been apparent that maintenance resources have increasingly become a major bottleneck that requires careful consideration and optimization. Submodular cost structures are very common in various maintenance settings, including ones that are very typical to the Air Force. Unfortunately, at the same time, they give rise to very hard optimization models. We have generated a series of results and papers that develop both new models and algorithms to find provably near-optimal solutions to core maintenance problems. The algorithms has theoretical worst-case guarantees as well as near-optimal empirical performance.

In joint work with with Jack Muckstadt, Danny Segev and Major Eric Zarybnisky (a former Air Force PhD student that was co-supervised by the PIs) [16, 17], we studied various aspects in the management of *scheduled maintenance activities for modular systems*, such as an aircraft engine. These systems consist of components with *cycle limits* that specify the maximum number of periods of use allowed between subsequent maintenance actions. (For example, a cycle for the starter system in an aircraft engine could be one startup sequence. For components in an aircraft braking system, a cycle could be one landing sequence.) Each component can be used for a certain number of cycles and then must be repaired or replaced due to safety or failure concerns. These cycle limits are determined through a number of methods including physical testing, simulation, and analytical assessment. Although it is possible that components fail prior to their cycle limits, due to the conservative nature of these cycle limits, such events are extremely rare. As a result, it is common to assume that a component is operational until its cycle limit is reached and that after maintenance it again has a full cycle limit. While much of the literature has examined stochastically failing systems, preventative maintenance of *usage limited* components has received less attention.

We have initially focused on a single modular system that consists of components with associated cycle limits. The typical cost structures that arise in practical settings are submodular in the subset of components being maintained. (This reflects the fact that as one component is maintained, it might be beneficial to maintain other components, even if their cycle limit is not due.) The goal is to compute a feasible maintenance schedule that minimizes the cost associated with component maintenance. This model captures some fundamental tradeoffs in the design phase of the system. By making cost tradeoffs early in development, program managers, designers, engineers, and test conductors can better balance the up front costs associated with system design and testing with the long term cost of maintenance. The model provides a framework for design teams to evaluate different design and operations concepts and then evaluate the long term costs. (For example, how to tradeoff the additional cost of extending the cycle-limit of a given component in the system during the development phase versus the long term maintenance cost savings.) However, the optimal policies in these settings tend to be very complex to compute and cannot be operationalized. A natural approach that is often considered in practice is to use cyclic policies that maintain each component at a fixed frequency. The question that arises is the increase in cost associated with using potentially suboptimal cyclic policies. We develop two algorithms to compute provably near-optimal cyclic policies. The *cycle rounding* algorithm computes the cycle-limits iteratively by considering the components in the system in increasing cycle-limits. The algorithm provides a worst-case approximation guarantee of 2. That is, for any input instance of the problem, the algorithm computes a cyclic policy with cost that is at most twice the optimal (over all policies, not necessarily cyclic policies). We also develop a class of *shifted power-of-two* algorithms. The cycle-limits are rounded to power-of-two times a shifting parameter and then an optimal policy with respect to the rounded cycle-limits is computed. (It can be shown that if all cycle-limits are power-of-two's the optimal policy is to maintain each component exactly when it is due.) Based on an innovative cost decomposition scheme and randomized analysis, we show that there is a small set of shifted power-of-two policies, the best of which is guaranteed to have cost at most $1/\ln(2)$ times the optimal cost. This guarantee holds for any submodular increasing cost function. Interestingly, the set of shifted power-of-two policies can be computed based only on the cycle limits independent of any other parameter of the problem, including the cost functions. Moreover, we show that one can choose an a-priori set of shifted power-of-two policies entirely independent of *any* parameter of the problem, and obtain constant worst-case guarantees. The guarantees are improving as the size of the this set grows larger and converge (quickly) to $1/\ln(2)$. In fact, even with just a few policies that are chosen a-priori, one can obtain a provable worst-case guarantee very close to $1/\ln(2)$. The analysis is obtained based on innovative linear programming approach that, for any given predefined number of chosen policies, reveals the worst-case guarantee, as well as tight worst-case instances.

In joint work with Maurice Cheung, Adam El-Machtoub and David Shmoys together with the PI Levi [8], we have studied the extension of the above model with *soft* cycle-limits. That is, instead of having a hard cycle-limit, one considers a monotone increasing penalty function in the delay in maintenance. The latter is an important modeling extension since in many practical cases the cycle-limits are indeed softer. However, the techniques and results described above do not carry through to the more general model. For several common submodular cost structures, we have been able to develop linear programming (LP) based approximation algorithms that are based on an integer programming (IP) formulation of the problem that is

in turn relaxed to obtain an LP relaxation of the problem. The LP is being solved and then the solution is rounded to a feasible integer solution with cost that is within a constant times the LP solution.

Once a modular system has moved into operations, manpower and transportation scheduling become important considerations when developing maintenance schedules. To address the operations phase, we develop the *modular maintenance and system assembly model* [15, ?] to balance the tradeoffs between inventory, maintenance capacity, and transportation resources. This model explicitly captures the risk-pooling effects of a central repair facility while also modeling the interaction between repair actions at such a facility. The full model is intractable for all but the smallest instances. Accordingly, we decompose the problem into two parts, the *system assembly* portion and *module repair* portion. Even the decomposed models are in general computationally challenging. We first study several practically interesting special cases, and develop optimal algorithms and heuristics to solve them. Finally, we use the output of these models together with the algorithms developed for the modular maintenance scheduling problem to propose an integrated methodology for design and operations of these complex systems.

Maintenance of low observable aircrafts. In joint work with 1st Lt. Phil Cho (former Master student that was co-supervised by the PIs), Vivek Farias, 1st Lt. John Kessler (Master student co-supervised by the PIs) and Major Eric Zarybnisky [9], we have studied the maintenance and flight scheduling of *low observable (LO) aircrafts*. The newest generation of fighter aircrafts in the Air Force, such as the F-22, has low-observable (LO) technologies that make them invisible to radar. This presents unique maintenance issues that did not exist for previous generations of fighter aircraft. In particular, the outer surfaces of the LO aircraft are coated with a metallic paint that is designed to minimize the radar signature of the aircraft. While LO aircraft have many design features that contribute to the LO capability of the aircraft, such as the shape and angles of the aircraft, the special outer metallic coating is the primary contributor to the increased maintenance requirements for LO aircraft. If the coating is damaged in any way, the radar signature of the aircraft can be affected. Since LO aircraft are not considered to be fully mission capable (FMC) unless their radar signature is below a certain level, maintenance personnel must continuously repair the metallic coating on LO aircraft in order to sustain an acceptable FMC rate for a fleet of aircraft. Each time an aircraft flies, maintenance personnel record all new damages into the *signature assessment system* (SAS). Particularly, each aircraft is associated with an evolving SAS number that reflects the estimated cumulative impact of all the damages on its LO capability. (Higher SAS number implies that the radar signature of the aircraft is higher.) Moreover, a SAS number higher than a certain threshold implies that the aircraft is not FMC. Currently, the decision process regarding SAS redux is largely dependent on the personalities of each maintenance unit. Since there is little published guidance regarding LO maintenance, each maintenance unit has the flexibility to make LO maintenance decisions however they see fit. Therefore, the various LO maintenance policies used by flying units throughout the Air Force can vary. In speaking with several experienced maintenance personnel, the LO maintenance decision process was described as being somewhat haphazard.

In this work, we model the LO SAS maintenance scheduling problem based on real data that records the SAS number evolution over 2.5 years (overall 5000 data points). The decision making problem is modeled as a variant of the *restless multi-armed bandit* problem. (This is a well-known model in stochastic control.) In addition, we use index policies that allow maintenance schedulers to quickly rank the aircrafts in the fleet based on each aircrafts current LO capability state, and use the ranking to decide which aircraft to enter into LO maintenance and for how long. We employ two algorithms to compute good index policies. In addition to maintenance scheduling, we explore policies for choosing which aircraft to fly to meet sortie requirements with a focus on the LO implications of those decisions. In extensive computational experiments we have demonstrated the strength of the index policies and the importance of the flight decision. In particular, we show that the index policies paired with good flight decisions perform within 10% of a computational upper bound. Several Air Force personnel are interested in testing the policies that we propose within a fidelity simulation setting they are developing to model various strategic logistics systems of the Air Force. Moreover, our models and policies provide fundamental insight to guide maintenance personnel in designing effective maintenance policies.

The models describe above give rise to a range of interesting theoretical questions regarding the performance of index policies, compared to the best optimal policy. We intend to study these questions in the

coming months.

Supply chain management and logistics models with online demand. A major aspect of any stochastic supply chain and logistics optimization model is how to model the uncertainty of the future 'demands'. The most naive models assume that the future demand is fully predictable. While this might lead to more tractable models, it is an unrealistic assumption in many practical settings. A more realistic approach is to assume that the demands follow some known distributions. These distributions are usually obtained based on available historical data. However, specifying these distributions is often very challenging, particularly, in military applications. Moreover, miss-specification could lead to non-robust or even meaningless solutions. Robustness is a very desirable feature in military environments. In online models one assumes that future demands are not only not known, but in fact generated by a worst-case adversary. In turn, one seeks to obtain robust policies that can do well with respect to an oracle that knows the entire future. While this is a stringent benchmark, it leads to algorithms that are very robust against worst-case scenarios.

In joint work with Niv Buchbinder, Tracy Kimbrel, Konstantin Makarychev and Maxim Sviridenko the PI Levi developed several new online algorithms for core logistics models (e.g., the *joint replenishment problem*) that have analytical *competitive ratio* [6]. That is, they are guaranteed to have cost within a constant ratio of an offline optimal solution that is computed based on a-priori knowledge of the future. The algorithms are based on a novel *primal-dual* approach that is based on a linear programming relaxation of the problems and use the dual of the linear program to guide the online algorithm.

In joint work with Adam El-Machtoub the PI Levi consider new online versions of supply chain management and logistics models, where in addition to production decisions, one also has to make decisions regarding which customers (missions) to serve [10, 11]. Specifically, customers (missions) arrive sequentially over time during a *planning phase*, and the decision maker has to permanently decide whether to accept or reject each customer upon arrival. If rejected upon arrival, a lost-sales (or penalty) cost is incurred. Once the selection decisions are all made and the planning phase is over, one has to satisfy (serve) all the accepted customers (missions) with minimum possible production cost. The goal is to minimize the total lost-sales (penalty) and production costs. In contrast to previous work, our assumption is that customers arrive in an online manner. That is, upon arrival of a customer, the decision maker has information on all the customers that arrived prior to the current customer, but very limited or no information about future arrivals. In particular, we assume that the sequence of customer arrivals is generated by a *worst-case adversary*. (This is a very common assumption in the literature on online algorithms that have been applied to many combinatorial optimization models.) We developed several novel algorithms for the respective online models that capture among others many variants of core logistics and supply chain optimization problems, such as the *single lot-sizing problem*, the *joint replenishment problem*, the *facility location problem* and *network design problems*. Our algorithms are based on repeatedly solving offline sub-problems as well as the use of concept from game theory. In particular, upon each customer arrival, we solve a certain offline problem with respect to all of the customers arrived thus far, ignoring future arrivals and previously made decisions by the online algorithm. The solution of this offline problem informs the algorithm whether to accept the current customer. the algorithms can be shown to have competitive ratios that are close or match the best possible ones.

Supply chains models with general demand distributions. The PI Levi continued to develop his stream of work to develop efficient algorithms for some of the core supply chain models with general demand distributions. Many of these core models are computationally challenging (both theoretically and practically) unless one has strong assumptions on the underlying demand distributions. The PI has developed a general framework that is based on new marginal cost accounting schemes as well as cost-balancing techniques. This framework has enabled the development of the first provably near-optimal algorithms for many core supply chain management models. In [23, 7] we study a single item single location model but when the cost structure includes economies of scales, such as fixed ordering costs. Fixed ordering/transportation/production costs arise in many real life scenarios and reflect the fact that ordering, production and transportation in batches lead to economies of scales. The paper studies new algorithmic approaches to several general variants of the classical stochastic lot-sizing problem. This is one of the core models in inventory theory that has challenged researchers and practitioners for decades. The structure of the optimal policies in some of

these variants is not even known. Moreover, even in cases where the structure of optimal policies is well understood, computing them is usually intractable and even finding good policies is often a very challenging task. In particular, to the best of our knowledge, there are no known provably good policies for most of these important models. The contributions of this work are two-fold. First, we propose a new policy that can be applied under the most general assumptions, i.e., with positive lead times and general demand structures. The policy is called randomized cost-balancing policy and has a worst-case performance guarantee of 3. That is, the expected cost of the policy is guaranteed to be at most 3 times the optimal expected cost, regardless of the specific instance. One of the novel aspects of these policies is the use of randomized decision rules. Specifically, the policy randomly chooses among different ordering quantities. While randomized algorithms have been used extensively for many optimization problems, we are not aware of any applications to inventory control models. The worst-case analysis of these algorithms employs several novel ideas that provide new insights on the respective stochastic lot-sizing models. In [21] we extended the framework to the *serial network*, which is one of the core multi-echelon supply chain management models.

Data-driven algorithms. The PI Levi continued to develop his stream *data driven* approaches to various core inventory management and other operational problems [19, 20, 3]. Unlike more traditional approaches the approach in these papers is to study these stochastic optimization models under the assumption that the underlying probability distributions are not given explicitly as part of the input. Instead one has only partial information, such as historical demand or sales data. In [22], Levi, Roundy and Shmoys study non-parametric data-driven variants of the *newsvendor model*, one of the core models in inventory theory, and its multi-period extension. They obtain bounds on the number of samples required to guarantee with high probability that the sampling-based solution is arbitrarily close to the optimal. Surprisingly, the bound is uniform and applies to *all* demand distributions. In [19], Levi, Perakis and a PhD student, Uichanco, obtain a much stronger bound for all log-concave demand distributions that matches empirical computational experiments. In joint work of the PI Levi with Georgia Perakis and Joline Uichanco (a female PhD student), we have studied the problem when historical demand data is available. The analysis characterizes the property of the demand distribution that determines the number of samples required to obtain near-optimal data-driven solution with high confidence. (Near-optimal means that the data-driven solution is close to the optimal solution that could be computed if the demand distribution is actually known.) We call this property the 'weighted spread'. The 'spread' is the difference between the conditional expectations of the demand conditioning on being larger and smaller, respectively, than the newsvendor quantile. The 'weighted spread' is the 'spread' weighted by the probability distribution function value at the newsvendor quantile. Moreover, we were able to characterize a large family of distributions, specifically, logconcave distributions, that captures many of the commonly used distributions in inventory theory and revenue management, for which a tighter theoretical analysis can be obtained. The bound are obtained through a uniform lower bound on the 'weighted spread' for all the distributions within this family that is achieved via a functional optimization program. In particular, the bounds for this class of distributions seem to match empirical experiments. These bounds are more useful since they provide realistic and practical bounds. The underlying analysis is new and sheds new light on several aspects of the problem. Moreover, in [20], we were able to use the weighted spread property to develop a new robust optimization approach to the newsvendor model.

Delayed 2-phase distribution. In joint work of the PI with Assaf Avrahami and Yale Herer from the Technion, Israel [2], we studied the distribution network of one of the largest media group in Israel. In particular, we focused on the distribution network of one of the primary weekly magazines of the group. Currently, this magazine is distributed to many retailing sales points at the beginning of the week. At the end of the week all returns are being collected for full refund to the retailers. We studied the value of information, specifically, a policy that supplies retailers twice a week instead of once. We are using RFID systems and sales people to collect data on the sales and inform the resupply decisions. This gives rise to a fairly complex model. We were able to analyze this model and show that, under very realistic assumptions it is convex. Moreover, a simulation model that we built predicts that the new policy will lead to a 10-15 % reduction in the returns. We are currently piloting our policies in the field. We believe that this new distribution model could have applications in the Air Force context. The paper describing this work was awarded the 2013 Daniel H. Wagner Prize for Excellence in Operations Research Practice.

Other work. The work under the grant has been used to partially fund the work on several additional papers on various resource allocation, assortment, pricing and inventory optimization models [24, 12, 14, 26, 1], as well as issues of cost, sources of variability, risk sharing and resource allocation in healthcare systems and networks [3, 28, 18, 4, 5].

A List of Air Force Contacts

1. HAF:

- Mr. John W McDonald, AF/A4ID, John.McDonald@pentagon.af.mil
- Mr. Eric R Nelson, AF/A4IT, Eric.Nelson@WPAFB.AF.MIL

2. AFLMA:

- Capt. Christopher Arendt, Chief, Operational Analysis Branch, Air Force Logistics Management Agency, Comm: 334-416-4126 DSN: 596-4126, Christopher.Arendt@maxwell.af.mil
- Dr. Doug Blazer, LMI, DSN 596-1406 Comm 334 416 1406, doug.blazer.ctr@maxwell.af.mil
- Maj. Christopher A. Boone, Chief, Wargame Division, Air Force Logistics Management Agency, Maxwell-Gunter Annex, Maxwell AFB, DSN: 596-2325, Comm: 334-416-2325
- 2nd Lt Eric Almeida, AFLMA/LGY, eric.almeida@maxwell.af.mil
- Capt. Brian D Waller, AFLMA/LGLM, Chief, Maintenance Branch, Air Force Logistics Management Agency, DSN 596-4593/Comm 334-416-4593, Brian.Waller@maxwell.af.mil
- Capt. Craig A. Lane, Chief, Wargaming Branch, Air Force Logistics Management Agency (AFLMA), 501 Ward Street, Maxwell AFB (Gunter Annex), AL 36114, DSN: 596.5803 Comm: 334.416.5803, craig.lane@maxwell.af.mil

3. AFMC:

- Mr. Bob McCormick (AFMC/A9), HQ AFMC/A9A, Comm: 937-257-6920 (DSN: 787-6920), FAX: 937-656-1498 (DSN: 986-1498), Bob.McCormick@wpafb.af.mil
- Mr. Mike McClure (AFGLSC), Operations Research Analyst, AFGLSC 402 SCMS/GUSB, Supply Chain Diagnostics and Analysis Flight, Bldg 70, Post 31, 5215 Thurlow Dr, ste. 5, WPAFB, OH 45433, DSN 986-2501 / Comm 937-656-2501, gordon.mcclure@wpafb.af.mil
- Dr. Siva S. Banda, Director, Control Science Center of Excellence, Air Force Research Laboratory, AFRL/RBCA, 2210 Eighth St., Bldg. 146, Wright-Patterson AFB, OH 45433-7531, E-mail: siva.banda@wpafb.af.mil, Phone: (937)255-8677
- Mr. Tim Meixner, ECSS Plan IPT Lead, 937-238-2741 (cell), 937-904-0795 (office)
- Mr. Richard A. Moore, Chief, Studies and Analyses Division, HQ AFMC/A9A, DSN: 787-4044, COMM: 937-257-4044, Richard.Moore@wpafb.af.mil
- Mr. Mike Niklas, Operations Research Analyst, Mike.Niklas@wpafb.af.mil

4. AMC:

- Mr. Donald R. Anderson (AMC/A9), YD-03, DAF, Assistant Director, Analyses, Assessments, and Lessons Learned (AMC/A9), DSN: 770-7629 Comm: (618) 220-7629, Donald.Anderson-02@scott.af.mil
- Dr. Donald R. Erbschloe, Chief Scientist, HQ AMC/ST, 402 Scott Dr., Unit 3EC, Scott AFB IL 62225, Voice: 618-229-4825 (DSN 779), Fax: 618-256-2502 (DSN 576), don.erbschloe@scott.af.mil

5. ACC:

- Capt. John V. Miller, A9 Executive Officer, ACC DSN 574-5751, (757) 764-5751, 61X FAM, Air Combat Command, UFPD, ACC A9, John.Miller@langley.af.mil
- Capt. Eric Gons, Operations Research Analyst, HQ ACC/A9A, 574-7684, (757) 764-7684, Eric.Gons@langley.af.mil
- Capt. Elim H. Sniady, VaANG, Maintenance Operations Officer, 192d Maintenance Squadron, DSN: 574-0310, COMM: (757) 764-0310, CELL: 757-813-1903, elim.sniady@langley.af.mil

- 2nd Lt Alex J. Shamp, VaANG, OIC, LO CRF, 192d MXS, Langley AFB, VA 23665, DSN 575-5700, COMM (757) 225-5700, CELL (757) 869-0805, alex.shamp@langley.af.mil
- Maj. Lucius A. Cattles, 1 MOS, MX Operations Officer, Langley AFB VA, DSN: 574-6238 COMM: 757-764-6238, Lucius.Cattles@Langley.AF.Mil

6. PM-JAIT:

- Jerry Rodgers, PM J-AIT Support Contractor, 2461 Eisenhower Ave., Hoffman Building 1, Room 1140/1148, Alexandria, VA 22331-3011, (703) 325-2988, jerry.d.rodgers@us.army.mil, www.ait.army.mil
- Sumra Manning, CISSP, PMP, Engineering Manager, RFID, Federal Systems, Unisys, 11109 Sunset Hills Road, Reston VA 20190, Office 703 439 6580, Cell 215-527-8914, sumra.manning@unisys.com, sumra.manning@us.army.mil
- Lee James III, PM J-AIT, 703.373.1355, lee.jamesiii@us.army.mil

7. Langley personnel

- 1LT Alex Shamp, Maintenance Officer for F-22 LO.
- Capt Eric Gons, Air Force Operations Research Analyst.

8. Hickam personnel

- Mr. Greg Stonelake.

References

- [1] J. Arts, R. Levi, G.-J. van Houtum, and B. Zwart. Base-stock policies for lost-sales models: State space aggregation and limiting transition probabilities. Working paper, 2013.
- [2] A. Avrahami, Y. Herer, and R. Levi. Matching supply and demand via delayed 2-phase distribution at Yedioth Group - Models, algorithms and IT. *Interfaces*, 44 (5):445–460, 2014. Paper was awarded the first prize in the 2013 Daniel H. Wagner Prize for Excellence in Operations Research Practice.
- [3] M. Begen, R. Levi, and M. Queyranne. A sampling-based approach to appointment scheduling. *Operations Research*, 60(3):675–681, 2012.
- [4] F. Bravo, M. Braun, V. Farias, and R. Levi. Data-driven optimization models to cost capturing and resource allocation in healthcare networks. Working paper, 2013.
- [5] F. Bravo, L. Ferrari, R. Levi, and M. McManus. Nature and sources of variability in surgical case duration. Submitted to *Pediatric Anesthesiology*, 2013.
- [6] N. Buchbinder, T. Kimberl, R. Levi, K. Makarychev, and M. Sviridenko. Online make-to-order joint replenishment model: Primal-Dual competitive algorithms. *Operations Research*, 61 (4):1014–1029, 2013.
- [7] X. Chao, R. Levi, C. Shi, and H. Zhang. Approximation algorithms for capacitated stochastic inventory systems with setup cost. To appear in *Naval Research Logistics*, 2009.
- [8] M. Cheung, A. ElMachtoub, R. Levi, and D. Shmoys. The submodular joint replenishment problem. Conditionally accepted to *Mathematical Programming*, 2012.
- [9] P. Cho, V. Farias, J. Kessler, R. Levi, T. Magnanti, and E. Zarybnisky. Maintenance and flight scheduling of low observable aircraft. Conditionally accepted to *Naval Research Logistics*, 2011.
- [10] A. ElMachtoub and R. Levi. Supply chain management with online customer selection. Submitted to *Operations Research* - second minor revision requested, 2010.

- [11] A. Elmachtoub and R. Levi. From cost sharing mechanisms to online selection problems. To appear in *Mathematics of Operations Research*, 2012.
- [12] V. Goyal, R. Levi, and D. Segev. Near-optimal algorithms for the assortment planning under dynamic substitution and stochastic demand. Submitted to *Operations Research*- under third round of revision. Extended abstract appeared in *MSOM* 2009, 2010.
- [13] D. W. Johnson, U. Schmidt, E. A. Bittner, B. Christensen, R. Levi, and R. M. Pino. Delay of transfer from the ICU: A prospective observational study of incidence, causes and financial impact. *Critical Care*, 17 (4), 2013.
- [14] P. Keller, R. Levi, and G. Perakis. Efficient formulations for multi-product construed pricing under attraction demand models. To appear in *Mathematical Programming*, 2009.
- [15] R. Levi, T. Magnanti, J. Muckstadt, D. Segev, and E. Zarybnisky. Multi-echelon modular maintenance and system assembly with transportation and maintenance capacities. Working paper, 2011.
- [16] R. Levi, T. Magnanti, J. Muckstadt, D. Segev, and E. Zarybnisky. Maintenance scheduling for modular systems modeling and algorithms. *Naval Research Logistics*, 61 (6):472488, 2014.
- [17] R. Levi, T. Magnanti, D. Segev, and E. Zarybnisky. The graph visiting problem. Submitted to *Mathematical Programming* - revision requested, 2011.
- [18] R. Levi, G. Perakis, C. Shi, and W. Sun. The efficiency of revenue sharing contracts in joint ventures in operations management. Submitted to *Management Science*, 2013.
- [19] R. Levi, G. Perakis, and J. Uichanco. The data driven newsvendor problem new bounds and insights. Fifth revision requested from *Operations Research*, 2009.
- [20] R. Levi, G. Perakis, and J. Uichanco. Regret optimization for stochastic inventory models with spread information. Submitted to *Operations Research*, revision requested, 2011.
- [21] R. Levi, R. Roundy, and V. A. Truong. Provably near-optimal balancing policies for stochastic multi-echelon inventory control models. Under second revision in *Mathematics of Operations Research*, 2012.
- [22] R. Levi, R. O. Roundy, and D. B. Shmoys. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research*, 32(4):821–838, 2007.
- [23] R. Levi and C. Shi. Approximation algorithms for the stochastic lot-sizing problem. To appear in *Operations Research*. The paper won the 2009 Nicholson Student Competition, 2009.
- [24] R. Levi and C. Shi. Revenue management of reusable resources with advanced reservations. Submitted to *Mathematic of Operations Research*, 2011.
- [25] R. Levi, D. B. Shmoys, and C. Swamy. LP-based approximation algorithms for capacitated facility location. *Mathematical Programming*, 131(1):365–379, 2012.
- [26] R. Levi and L. Yedidsion. The assembly network problem is NP-Hard. *Operations Research Letters*, 41(2):134137, 2013.
- [27] D. Segev, R. Levi, P. Dunn, and W. Sandberg. Modeling the impact of changing patient transportation system on perioperative process performance in a large hospital: Insights from a computer simulation study. *Healthcare Management Science*, 2010. To appear in *Healthcare Management Science*.
- [28] A. C. Zenteno, T. Carnes, R. Levi, B. Daily, D. Price, S. Moss, and P. Dunn. Pooled open blocks shorten wait times for non-elective surgical cases. To appear in *Anal of Surgery*, 2013.