Probabilistic Analysis of Multi-Item Capacitated Lot Sizing Problems

Awi Federgruen — Columbia University

Joern Meissner — Lancaster University Management School

November 2004, Revision April 2005

Abstract

This paper conducts a probabilistic analysis of an important class of heuristics for multi-item capacitated lot sizing problems.

We characterize the asymptotic performance of so-called progressive interval heuristics as T, the length of the planning horizon, goes to infinity, assuming the data are realizations of a stochastic process of the following type: the vector of cost parameters follows an arbitrary process with bounded support, while the sequence of aggregate demand and capacity pairs is generated as an independent sequence with a common general bivariate distribution, which may be of *unbounded* support. We show that important subclasses of the class of progressive interval heuristics can be designed to be asymptotically optimal *with probability one*, while running with a complexity bound which grows *linearly* with the number of items N and slightly faster than *quadratically* with T.

We generalize our results for the case where the items' shelf life is uniformly bounded, e.g. because of perishability considerations.

1 Introduction

This paper conducts a probabilistic analysis of an important class of heuristics for multi-item capacitated lot sizing problems. More specifically, we address the following classical problem (P): a family of N items is to be procured from the same production facility or outside supplier. The planning horizon consists of T periods (not necessarily of equal length). Demands are specified for each item and each period of the planning horizon. The aggregate order size, in any given period, is bounded by a capacity limit, which may vary over the course of the planning horizon. The costs consist of inventory carrying, variable and fixed order costs. As to the latter, the fixed order cost in any given period only depends on the period index, but not on the composition of the order. The inventory and variable order costs are proportional with the end-of-period inventories and order sizes, at item- and period-dependent cost rates. The objective is to minimize total costs for the planning horizon while satisfying all demands, without backlogging.

Despite a voluminous literature devoted to the general model (P), it continues to present a major challenge to theoreticians and practitioners alike. The problem is NP complete, even in the special case of a single item (N = 1), as shown by FLORIAN ET AL (1980). Until recently, exact and heuristic solution methods have only been successfully applied to instances with a relatively low number of items and/or periods. Recently, FEDERGRUEN ET AL. (2002) investigated the following class of so-called *progressive interval* heuristics. A progressive interval heuristic consists of J iterations, where, iteration by iteration, the problem is solved, to optimality, over a progressively larger time interval $[1, T_\ell]$, i.e. $T_1 \leq T_2 \leq \cdots \leq T_J = T$. When solving a given interval problem, the necessary and sufficient conditions for a feasible extension to the remainder of the planning horizon are appended as boundary conditions. To ensure that the computational complexity in each iteration remains managable, the heuristic fixes, in iteration ℓ, all integer variables for periods 1 to $T_\ell - \tau$ (for some $\tau > 0$) and all *continuous* variables for periods 1 to some $t_\ell \leq T_{\ell-1}$ at their optimal value after iteration $\ell - 1$. The horizons are chosen such that $0 = t_1 \leq t_2 \leq \cdots \leq t_J$ while $\tau \geq T_\ell - T_{\ell-1}$, the number of periods by which the horizon is expanded in the ℓ-th iteration.

We characterize the asymptotic performance of the progressive interval heuristics as T goes to infinity, assuming the data are realizations of a stochastic process of the following type: the vector of cost parameters follows an arbitrary process with bounded support, while the sequence of aggregate demand and capacity pairs is generated as an independent sequence with a common general bivariate distribution, which may be of *unbounded* support. We show that important subclasses of the class of progressive interval heuristics can be designed to be asymptotically optimal *with probability one*, while running with a complexity bound which grows *linearly* with the number of items N and slightly faster than *quadratically* with T. Our probabilistic analyses complement the worst case analyses in FEDERGRUEN ET AL. (2002) where asymptotic optimalty is shown under conditions which require that all demands and capacities are *uniformly bounded* and that the aggregate capacity over a large enough interval of time exceed the aggregate demand by at least a minimum slack value $\sigma > 0$. *Both* of these assumptions are somewhat restrictive, in the context of an asymptotic analysis where very large planning horizons T are considered.

For many types of complex (NP-complete) logistical planning problems, probabilistic analyses have provided performance guarantees for various classes of heuristics, fostering insights into which algorithmic approaches are effective for large size problems. One such planning area is that of vehicle routing, starting with the seminal papers by KARP (1979) and HAIMOVICH and RINNOOY KAN (1985); see COFFMAN and LUEKER (1996), FEDERGRUEN and SIMCHI-LEVI (1992) and ANILY and BRAMEL (1999) for surveys. (Some of the planning models integrate vehicle routing with inventory planning but, thus far, only in a context of demand processes that occur at *constant* rates.) Other logistical planning areas supported by probabilistic analyses include (hierarchical) facility location and sourcing models (e.g. CHAN and SIMCHI-LEVI (1996), GALLEGO and SIMCHI-LEVI (1997), FISHER and HOCHBAUM (1980), and ROMEIJN and ROMERO MORALES (2001)). See BRAMEL and SIMCHI-LEVI (1997) for a general overview. RHEE and TALAGRAND (1987, 1989) and RHEE (1993) have shown how probabilistic analyses of a variety of logistical planning problems can be based on specific large deviation results. Our analyses, as well, are in part, based on such large deviation techniques. To our knowledge, the probabilistic analyses in this paper represent the first such analyses for inventory planning models with *time-varying* parameters (, otherwise referred to as dynamic lot sizing problems).

We conclude this § with a brief review of the relevant literature beyond the papers mentioned above. The (NP-) complexity of the problem arises from the superpostition of (joint) setup costs and capacity limits. Indeed, the problem is solvable in $O(NT \log T)$ time, if either the capacity constraints are relaxed or in the absense of fixed order costs. In the former case, the problem decomposes into N independent single item lotsizing problems for which one of the $O(T \log T)$ methods by AGGARWAL and PARK (1992), FEDERGRUEN and TZUR (1991) or WAGELMANS ET AL. (1992) can be used. In the latter case, the problem is solvable in $O(NT \log T)$ time with AHUJA and HOCHBAUM (2004)'s recent method.

There is a voluminous literature describing various heuristics for the general multi-item model. We refer to SALOMON (1990) and KUIK ET AL. (1994) for surveys of the literature until 1994. State-of-the art solution methods include BELVAUX and WOLSEY (2000, 2001), STADTLER (2003) and SUERIE and STADTLER (2003). These methods are all based on variants of progressive interval heuristics. See FEDERGRUEN ET AL. (2002) for details and a more detailed literature review. Other than the above mentioned worst case analyses in the latter paper, the only performance guarantees for heuristics for capacitated lot-sizing problems with general time-dependent capacity limits are due to GAVISH and JOHNSON (1990) and VAN HOESEL and WAGELMANS (2001). The latter developed a fully polynomial approximation scheme for the general *single-item* model, after the former proposed such a scheme for a more restricted version of the problem.

The remainder of this paper is organized as follows: In §2, we specify the model, the probability model generating its data and the class of progressive interval heuristics. In §3, we establish *almost sure* asymptotic optimality for heuristics in this class as well as their polynomial complexity bound. §4 concludes the paper with a discussion of the case where the items' shelf life is uniformly bounded e.g. because items are perishable.

2 The model and the class of progressive interval heuristics

The model employs the following data, where the index $i \in \{1, \ldots, N\}$ is used to distinguish between items and time periods are indexed by t. (Demands are represented as multiples of the volume that consumes *one* unit of capacity):

c_{it} = variable per unit order cost for item i in period t

h_{it} = cost of carrying a unit of inventory of item i at the end of period t

K_t = setup cost incurred when an order is placed in period t

d_{it} = demand for item i in period t; ($d_{it} \geq 0$)

D_t = aggregate demand in period $t = \sum_{i=1}^{N} d_{it}$

C_t = order capacity, i.e. the maximum number of units which can be ordered in period t.

We use the following set of decisions variables:

x_{it} = order size for item i in period t; $i = 1, \ldots, N$; $t = 1, \ldots, T$

$$Y_t = \begin{cases} 1 & \text{if } \sum_{i=1}^{N} x_{it} > 0 \\ 0 & \text{otherwise} \end{cases} \qquad t = 1, \ldots, T$$

I_{it} = ending inventory of item i in period t; $i = 1, \ldots, N$; $t = 1, \ldots, T$

Let I_t^0 = the *minimum* aggregate inventory at the end of period t, such that a feasible production / inventory plan exists for periods $t+1, \ldots, T$ $(t = 0, 1, \ldots, T)$. These minimum stock levels are easily computed from the following recursion, which can be verified by induction:

$$I_t^0 = \left(D_{t+1} - C_{t+1} + I_{t+1}^0\right)^+, \qquad t = 0, 1, \ldots, T-1, \quad \text{with} \quad I_T^0 = 0 \tag{1}$$

This is the well known Lindley equation, see e.g. ASMUSSEN (1987). The following is a standard formulation:

$$(\mathrm{P}) \qquad z^* = \min \left\{ \sum_{t=1}^{T} \left[K_t Y_t + \sum_{i=1}^{N} (c_{it} x_{it} + h_{it} I_{it}) \right] \right\} \tag{2}$$

s.t.

$$I_{it} = I_{i(t+1)} + x_{it} - d_{it}, \qquad i = 1, \ldots, N, \quad t = 0, \ldots, T \tag{3}$$

$$\sum_{i=1}^{N} x_{it} \leq C_t Y_t \qquad i = 1, \ldots, N, \quad t = 1, \ldots, T \tag{4}$$

$$\sum_{i=1}^{N} I_{it} \geq I_t^0 \qquad t = 1, \ldots, T \tag{5}$$

$$I_{i0} = 0; x_{it} \geq 0; \quad I_{it} \geq 0; \quad Y_t \in \{0, 1\} \tag{6}$$

We conclude, from (1) that:

Lemma 1 *(P) has a feasible solution iff* $I_0^0 = 0$.

We assume that the model data are generated by the following probabilistic model: the $(2N+1)T$ cost parameters $\{K_t, c_{it}, h_{it}\}$ are generated by an *arbitrary* stochastic process with support on a hypercube in the positive orthant of $\mathbb{R}^{(2N+1)T}$. As to the sequence of aggregate demand and capacity pairs $\{(D_t, C_t) : t = 1, \ldots, T\}$, we assume:

(A) $\{(D_t, C_t) : t = 1, \ldots, T\}$ is a sequence of independent pairs of random variables, all distributed like (D, C) with a general bivariate distribution, possibly with *unbounded* support, such that the marginal distribution of D has a moment generating function, i.e. $E(e^{\theta D})$ exists for some $\theta > 0$, $\delta = E(D) > 0$ and the support of the distribution of C is bounded from below by a constant C_*. Moreover, $\mu = E(C) - E(D) > 0$.

The requirement that the demand distribution has a moment generating function which is finite in the neighborhood of the origin covers most of the distributions commonly used in (stochastic) inventory models. (e.g. the Normal, Gamma, Negative Binomial or Weibull distributions). The condition merely precludes heavy-tailed demand distributions which implies heavy-tailed distributions for the steady-state distribution of the reserve-stock variables $\left\{I_t^0\right\}$. The condition $\mu = E(C) - E(D) > 0$ is necessary to ensure that the generated problem instances be *feasible* as T becomes large. Let $\psi(\theta)$ denote the *cumulative* generating function (cgf) of the random variable (D-C) which is the logarithm of its moment generating function:

$$\psi(\theta) = \log E\left[e^{\theta(D-C)}\right].$$

Since D has a finite moment generating function on some interval $[0, \bar{\theta}]$, so does $(D - C)$, so that $\psi(\theta) < \infty$ on $[0, \bar{\theta}]$. Moreover, $\psi(\cdot)$ is differentiable with $\psi(0) = 0$ and $\psi'(0) = -\mu < 0$, by (A), so that $\psi(\theta) < 0$ for all $\theta > 0$, sufficiently small.

When the items have a limited shelf life, we show in §4 that our results continue to apply under generalizations of condition (A), allowing for various forms of intertemporal demand and capacity dependences.

In a progressive interval heuristic, employing J iterations, the ℓ-th iteration consists of solving (P), with T replaced by T_ℓ and all (integer) Y-variables for periods $1, \dots, T_\ell - \tau$ and all (continuous) x- and I-variables for periods $1, \dots, t_\ell \le T_{\ell-1}$ fixed at their optimal value in the $\ell - 1$st iteration. (In the first iteration, no rstrictions apply to any of the variables.) Thus, the number of unrestricted integer variables in each (except for possibly the first) iteration is kept constant at τ. Since the complexity of any mixed integer program is primarily determined by the number of (unrestricted) integer variables, the computational complexity remains managable when choosing τ sufficiently small, and from each iteration to the next it grows only moderately.

The heuristic starts with the recursive computation of the values $\left\{I_t^0\right\}$, via (1). If $I_0^0 > 0$, no feasible solution exists, in which case the optimality gap is defined to be zero, see Lemma 1.

As in Federgruen et al. (2002), we pay special attention to two *extreme* subclasses: (i) the *Strict Partitioning* heuristics (SP), with all (except for possibly the last) interval increment $T_\ell - T_{\ell-1} = \tau$ and $t_\ell = T_{\ell-1} = T_\ell - \tau$; (ii) the *Expanding Horizon* heuristics (EH) with all $t_\ell = 0$. The (SP)-heuristics are related to the Time Partitioning heuristics, see Federgruen and Tzur (1999).

Thus, the (SP)-heuristics minimize the computational complexity of each interval problem at the expense of providing minimal flexibility to the continuous variables. The (EH)-heuristics, while of larger computational complexity, provide *maximal* flexibility for the continuous variables and even for the integer variables, in case interval increments $T_\ell - T_{\ell-1} < \tau$ are chosen. Under such choices, even many of the setup decisions made in one iteration, may be revisited in subsequent iterations, on the basis of additional demand, cost and capacity information pertaining to additional periods. The numerical study in FEDERGRUEN ET AL. indicates that (EH)-heuristics can be used effectively to solve moderate to large size problem instances and that the solutions generated come very close to being optimal. Those gererated by (SP)-heuristics typically exhibit larger optimality gaps.

3 Almost sure asymptotic optimality

In this §, we show that both (SP)- and (EH)-heuristics can be designed to be simultaneously *almost surely asymptotically optimal* as well as of *low polynomial complexity*. As with all (SP)-heuristics, we confine ourselves to (EH)-heuristics in which (with the possible exception of the last iteration) exactly τ periods are appended to the tail of the planning horizon, as we progress from one iteration to the next $(T_\ell - T_{\ell-1} = \tau)$. We show that both heuristics, with cost values z^{SP} and $z^{(EH)}$, respectively, are *almost surely (a.s.) asymptotically optimal* if the interval increment τ is adjusted as a function of T, where

$$\begin{aligned} \tau &= \Omega(\log T), \quad \text{i.e.} \quad \lim_{T\to\infty} \frac{\tau}{\log T} = \infty, \text{e.g.} \\ \tau &= = \eta \lceil \log T \rceil^{\zeta}, \quad \text{for some} \quad \eta > 0 \quad \text{and} \quad \zeta > 1 \end{aligned} \tag{7}$$

To derive a specific complexity bound, we assume that each interval problem in each iteration is solved with a tailored branch-and-bound procedure, i.e. the b&b-procedure in §5 of FEDERGRUEN ET AL. (2002) in which each non-leaf node of the tree is evaluated with lower bound LB_3, ibid.

Theorem 1 *Consider a (SP)-heuristic with $\tau = \eta \lceil \log T \rceil^{\zeta}$ for some $\eta > 0$ and $\zeta > 1$. The heuristic is asymptotically optimal , a.s. and it can be designed to run in O($NT^{\zeta+1} \log\log T$) time as well.*

Proof. Let $c_i^* > 0$ denote the essential infimum of the stochastic process $\{c_{it} : t = 1, \ldots, T\}$. Observe first that by the law of large numbers, with probability one, $\liminf_{T\to\infty} \frac{z^*}{T}$ $\geq \sum_{i=1}^N c_{i*} \liminf_{T\to\infty} \frac{1}{T} \sum_{t=1}^T d_{it} = \sum_{i=1}^N c_{i*}\delta_i > 0$, with $\delta_i = E(d_{it})$. (Note that $\delta_i > 0$ for at least one $i = 1, \ldots, N$ since $\sum_{i=1}^N \delta_i = \delta > 0$.) In other words, with probability one, the numerator in the optimality gap $\frac{z^{SP}-z^*}{z^*}$ grows at least linearly in T. It thus suffices to show,

$$\lim_{T\to\infty} \frac{1}{T}\left[z^{SP} - z^*\right] = 0, \quad \text{a.s.} \tag{8}$$

As in the worst-case analysis of Theorem 2 of FEDERGRUEN ET AL. (2002), we transform an optimal solution for the complete problem in *two* phases into a solution which is achievable by the (SP)-heuristic. (If no feasble solution exists, i.e. $I_0^0 > 0$, the zero optimality gap is bounded by that achieved for the transformed instance in which $I_{i0} = I_0^0/N, \quad i = 1,\dots,N$.) In Phase I, the optimal solution is transformed into one with all intervals' ending *aggregate* inventories equal to their I^0-values. (Note that that the solution generated by the (SP)-heuristic satisfies this property as well.) Let z^I denote its cost value. In Phase II, the composition of the reserve stock at the end of each of the intervals is made identical to that of the solution of the (SP)-heuristic, resulting in a solution with the cost value z^{II}. This solution is one which is among the ones considered by the (SP)-heuristic, i.e. $z^{II} \geq z^{SP}$. Thus,

$$\frac{z^{SP} - z^*}{T} \leq \frac{z^{II} - z^*}{T} = \frac{z^{II} - z^I}{T} + \frac{z^I - z^*}{T} \tag{9}$$

Following the proof of Theorem 2 in FEDERGRUEN ET AL. (2002) and given the (general) assumption about the stochastic process which generates the cost parameters, ones verifies that an integer $\Lambda > 1$ and constants B_1 and B_2 exist such that $z^I - z^* \leq (J-1)B_1 + B_2 \sum_{\ell=1}^{J-1} \sum_{r=T_\ell-\Lambda+1}^{T_\ell} C_r$. If $\Lambda > \tau$, the partial sums $\{\sum_{r=T-\Lambda+1}^{T_\ell} C_r : \ell = 1,\dots,J-1\}$ may overlap. However, $z^I - z^* \leq (J-1)B_1 + B_2\lceil\frac{\Lambda}{\tau}\rceil \sum_{r=T-(J-1)\Lambda+1}^{T} C_r$ is a valid upper bound. Thus,
$\frac{z^I-z^*}{T} \leq \frac{(J-1)}{T}\left[B_1 + B_2\lceil\frac{\Lambda}{\tau}\rceil\Lambda\frac{1}{(J-1)\Lambda}\sum_{r=T-(J-1)\Lambda+1}^{T} C_r\right]$ and $\lim_{T\to\infty}\frac{z^I-z^*}{T} - 0$ a.s., since
$\lim_{T\to\infty}\frac{J-1}{T} = \lim_{T\to\infty}\tau^{-1} = 0$ and since, with probability one,
$\lim_{T\to\infty}\frac{1}{(J-1)\Lambda}\sum_{r=T-(J-1)\Lambda+1}^{T} C_r = E(C_1) < \mu$ by the law of large numbers and the fact that the sequence $\{C_t : t = 1,2,\dots\}$ is an i.i.d. sequence of random variables.

To bound the additional cost incurred because of the Phase II transformation, let $\Delta c^* = \max_t \max_{i\neq\ell}[c_{it} - c_{\ell t}]$, $\Delta h^* = \max_t \max_{i\neq\ell}[h_{it} - h_{\ell t}]$, $h_* = \inf_{i,t} h_{it}$ and $K^* = \max_t K_t$ and note that $\Delta c^* = O(1)$, $\Delta h^* = O(1)$ and $K^* = O(1)$ as $T \to \infty$ while $h_* > 0$.

The solution obtained after Phase I and the solution generated by the (SP)-heuristic, both have $I_{T_\ell} = I_{T_\ell}^0$ for all $\ell = 1,\dots,J$. In Phase II, we obtain the desired composition of the ending inventory at the end of the J intervals by changing (only) the item identity of at most (all of the) $I_{T_\ell}^0$ units in the ending inventory of the $\ell - th$ interval, without any additional changes in the order- and inventory plans. The transformed solution remains feasible, incurs no additional fixed order costs and adds at most $\sum_{\ell=1}^{J-1} I^0_{T_\ell}(\Delta c^* + L_\ell \Delta h^*)$ in variable costs, where L_ℓ denotes the shelf life of the oldest unit in the reserve stock at the end of period T_ℓ. Thus, to prove (8) it suffices to show that

$$\lim_{T\to\infty} \frac{1}{T} \sum_{l=1}^{J-1} I^0_{T_\ell} \left(\Delta c^* + L_\ell \Delta h^*\right) = 0, \quad \text{a.s.} \tag{10}$$

Recall that $\{I^0_t\}$ in (1), when traversed *backwards*, is a Lindley process. Since the pairs $\{(D_t, C_t)\}$ are i.i.d, and since $\mu > 0$, $\{I^0_t\}$ has a limiting distribution I^0 (i.e. $\lim_{t\to\infty} I^0(T) \overset{w}{=} I^0$, where the convergence is in distribution.) with $E(I^0) < \infty$, see Asmussen (1987, §8.1). Moreover, in view of the remaining assumption in (A), the distribution of I^0 has an exponential tail, i.e. there exist constants α and $\beta > 0$ such that $Pr[I^0 > x] \sim \alpha e^{-\beta x}$, $x \to \infty$ (i.e. $\lim_{x\to\infty} \frac{Pr[I^0 > x]}{\alpha e^{-\beta x}} = 1$) (see Assmussen (1987, §12.5). Thus, for some $x^0 > 0, Pr[I^0 > x] \le 2\alpha e^{-\beta x}$, for all $x > x^0$. Finally, let $\bar{I}(T) = \max_{t=1,\dots,T} I^0_t$ denote the *largest* minimum reserve stock required over the entire planning horizon. Since I^0 has the same distribution as $[D - C + I^0]$, and since $0 = I^0_T \le_{st} I^0$, one easily verifies by complete induction that $I^0_t \le_{st} I^0$ for *all* $t = 1 \dots, T$.

Let $\tilde{n}(T) = \sqrt{\tau \log T}$ denote the geometric mean of τ and $\log T$ and note from (7) that

$$\lim_{T\to\infty} \frac{\log T}{\tilde{n}(T)} = 0 \quad \text{and} \quad \lim_{T\to\infty} \frac{\tilde{n}(T)}{\tau} = 0. \tag{11}$$

We first show that

$$\lim_{T\to\infty} Pr\left[L_\ell \le \tilde{n}(T) \quad \text{for all} \quad \ell = 1, \dots, J-1\right] = 1 \tag{12}$$

i.e. asymptotically the maximum shelflife of any unit in the reserve stock at the end of any of the intervals (in the Phase I solution) is almost surely bounded by $\tilde{n}(T)$. Under (12) we have almost surely that (8) holds since

$$\begin{aligned}
0 \quad &\le \quad \lim_{T\to\infty} \frac{1}{T} \sum_{\ell=1}^{I-1} I^0_{T_\ell}(\Delta c^* + L_\ell \Delta h^*) \\
&\le \quad \lim_{T\to\infty} \left\{ \frac{(\Delta c^* + \tilde{n}(T)\Delta h^*)}{T} (J-1) \right\} \lim_{T\to\infty} \frac{1}{J-1} \sum_{l=0}^{J-1} I^0_{T_\ell} \\
&= \quad h^* \left(\lim_{T\to\infty} \frac{\tilde{n}(T)}{\tau} \right) E(I^0) = 0, \quad \text{a.s.}
\end{aligned}$$

where the first equality follows from the the fact that the $\{I^0_t\}$-process is ergodic, so that a long-run average, sampled at equidistant epochs, converges with probability one to the expected value of the limiting distribution, while the second equality follows from (11).

It remains to prove (12). Note that $Pr[L_1 > \tilde{n}(T) \quad \text{or} \quad L_2 > \tilde{n}(T) \quad \text{or} \quad \ldots L_{J-1} > \tilde{n}(T)] \leq \sum_{l=1}^{J-1} Pr[L_\ell > \tilde{n}(T)]$. Choose $0 < \theta < \beta$ such that $\psi(\theta) < 0$. To bound each of the terms in the sum, consider first the *conditional* probability $Pr[L_\ell > \tilde{n}(T)]|I_{T_\ell} = i_\ell^0]$, which is bounded by $Pr[\sum_{r=T_\ell-n(T)+1}^{T_\ell}(C_r - D_r) \leq i_\ell^0 | I_{T_\ell} = i_\ell^0]$ with $n(T) = \tilde{n}(T) - \frac{\Delta c^* + K^*}{h_*}$.

(If a unit, in stock at the end of period T_ℓ, has a shelf life larger than $\tilde{n}(T)$, this implies that a *full capacity* order is placed is in *each* of the periods in the interval $[T_\ell - n(T) + 1, \ldots ..., T_\ell]$, for otherwise the procurement of this unit could be postponed till some period in this interval with slack capacity, saving at least $\tilde{n}(T) - n(T)$ periods' carrying costs, i.e. at least $(\Delta c^* + K^*)$, more than offsetting any additional order costs. However, given the condition $I_{t_\ell} = i_\ell^0$, this situation can only happen if $\sum_{r=T_\ell-n(T)+1}^{T_\ell}(C_r - D_r) \leq i_\ell^0$.) Thus,

$$
\begin{aligned}
Pr[L_\ell > \tilde{n}(T) | i_\ell^0] &\leq Pr[\sum_{r=T_\ell-n(T)+1}^{T_\ell} (C_r - D_r) \leq i_\ell^0 | I_{T_\ell} = i_\ell^0] \\
&= Pr[\sum_{r=T_\ell-n(T)+1}^{T_\ell} (D_r - C_r) \geq -i_\ell^0 | I_{T_\ell} = i_\ell^0] \\
&= Pr[\sum_{l=1}^{n(T)} (D_r - C_r) \geq -i_\ell^0] \qquad (13) \\
&\leq \exp\{-n(T)(\frac{-\theta i_\ell^0}{n(T)} - \psi(\theta))\}
\end{aligned}
$$

where the second equality follows from the fact that $I_{T_\ell}^0$ only depends on the demand and capacity values in periods $T_\ell + 1, \ldots, T$, see (1), so that the conditional distributions of $\{(D_r - C_r | I_{T_\ell}^0) : r = T_\ell - n(T) + 1, \ldots, T_\ell\}$ coincide with the *unconditional* distributions $\{D_r - C_r : r = T_\ell - n(T) + 1, \ldots, T_\ell\}$ and hence those of $\{D_1, \ldots, D_{n(T)}\}$ by the i.i.d. assumption of $\{(D_t, C_t)\}_{t=1}^{\infty}$. The last inequality in (13) follows from Chernoff's inequality. We thus obtain the following bound on the *unconditional* probability

$$
\begin{aligned}
Pr[L_\ell > \tilde{n}(T)] &\leq E_{I_{T_\ell}^0} \left\{ \exp(-n(T)\psi(\theta)) \exp(\theta \bar{I}_{T_\ell}) \right\} \qquad (14) \\
&\leq \exp(-n(T)\psi(\theta)) E_{\bar{I}_r} \exp(\theta \bar{I}_T)
\end{aligned}
$$

Note that

$$
\begin{aligned}
E\exp\{\theta\bar{I}(T)\} &= -\int_0^\infty e^{\theta x} d[1 - Pr(\bar{I}(T) \le x)] = 1 + \int_0^\infty \theta e^{\theta x} Pr[\bar{I}(T) > x] dx \\
&= 1 + \int_0^\infty \theta e^{\theta x} Pr[I_1^0 > x \quad \text{or} \quad I_2^0 > x \quad \text{or} \quad \dots I_T^0 > x] dx \le 1 + \sum_{t=1}^{T} \int_0^\infty \theta e^{\theta x} Pr[I_t^0 > x] dx \\
&\le 1 + \sum_{t=1}^{T} \int_0^\infty \theta e^{\theta x} Pr[I^0 > x] dx = 1 + T \int_0^\infty \theta e^{\theta x} Pr[I^0 > x] dx \qquad (15) \\
&\le 1 + T \left[\int_0^{x^0} \theta e^{\theta x} Pr[I^0 > x] dx + \int_{x^0}^\infty \theta e^{\theta x} 2\alpha e^{-\beta x} dx \right] \\
&\le 1 + T \left[\int_0^{x^0} \theta e^{\theta x} dx + \frac{2\alpha\theta}{(\beta - \theta)} e^{e^{(\theta-\beta)x^0}} \right] = 1 + Tb
\end{aligned}
$$

with $b = e^{\theta x^0} - 1 + \frac{2\alpha\theta}{(\beta-\theta)} e^{(\theta-\beta)x^0}$, where the second inequality follows from $I_t^0 \le_{st} I^0$, for all t. Thus, with $a = -\psi(\theta) > 0$:

$$
\begin{aligned}
0 &\le \lim_{T\to\infty} Pr[L_1 > \tilde{n}(T) \quad \text{or} \quad L_2 > \tilde{n}(T) \quad \text{or} \quad \dots L_{I-1} > \tilde{n}(T)] \\
&\le \lim_{T\to\infty} \sum_{l=1}^{J-1} Pr[L_\ell > \tilde{n}(T)] \\
&\le \lim_{T\to\infty} \{(J-1)\exp\{-an(T)\}(1 + Tb) \\
&= b \lim_{T\to\infty} \{\frac{T^2}{\tau} \exp(-\log T \frac{an(T)}{\log T})\} \\
&= b \lim_{T\to\infty} \frac{T^2}{\tau} \frac{1}{T^{\frac{an(T)}{\log T}}} = 0
\end{aligned}
$$

where the last inequality follows from (14) and (15) and the last equality from (11). This proves (12), hence (10) and (8).

It remains to he shown that when $\tau = \eta\lceil(\log T)^\zeta\rceil$, with $\zeta > 1$, the progressive interval heuristic runs in $O(NT^{1+\zeta}\log T)$ time, when each interval problem is solved with the above described b&b method. The discussion in §5 of Federgruen et al. (2002) shows that evaluation of any node of a b&b tree requires $O(N\tau\log\tau)$ time. Since this needs to be done at most 2^τ times to evaluate the complete tree, and since $\tau = O(\frac{T}{\zeta})$ interval problems need to be solved, the complexity bound follows immediately. ■

The same simulaneous (almost sure) asymptotic optimality and polynomial complexity can be obtained for the above (EH)-heuristic, under the same choice for the interval increment τ as in (11). The complexity of this (EH)-heuristic is $O(T\log\log T/(\log T)^{\zeta-1})$ larger than that of the (SP)-heuristic. Nevertheless, complexity grows only *linearly* with N and (only) slighly faster than cubically with T:

Theorem 2 *Consider an (EH)-heuristic with $\tau = \eta\lceil(\log T)^{\zeta}\rceil$ for some $\eta > 0$ and $\zeta > 1$. The heuristic is asymptotically optimal a.s. and can be designed to run in $O(NT^{2+\zeta}/(\log T)^{\zeta-1})$ time.*

Proof: The proof is analogous to that of Theorem 1, with only the following modifications: The Phase II transformation should modify the composition of the reserve stock at the end of periods $T_1, T_2, \dots, T_{J=1}$ to that prevailing at the end of the ℓ-th iteration of the (EH)-heuristic. (In the case of the (EH)-heuristic, this composition may change in subsequent iterations). As shown in Theorem 2 (b) of FEDERGRUEN ET AL. (2002), a *third* Phase transformation is necessary to obtain a solution which is is among the ones considered by the (EH)-heuristic, but this third transformation only *reduces* the cost value. The derivation of the complexity bound is again analogous, except that the evaluation of a single node in *one* of the b&b trees now requires $O(NT \log T)$ time. ■

4 Products with limited shelf life

Thus far, we have assumed that items can be kept in stock for un umlimited amount of time. In this §, we address the situation where the shelf life of each item is bounded by an (integer) constant λ, perhaps because the items are perishable. We refer to the survey paper by NAHMIAS (1982) for a review of inventory models with limited shelf lives. Within the context of dynamic lot sizing models, the complication of a fixed shelf life has not been addressed until HSU (2000) who showed that the single item uncapacitated model can be solved in $O(T^2)$ time. (For this case, HSU addresses, in addition, more general life time models and more general order and inventory cost functions then those used in (P).)

We show that, in the presence of a limited shelf life, almost sure asymptotic optimality of (SP)- and (EH)-heuristics can be established under conditions even more general than (A), for example:

(A^f) The sequence $\{(D_t, C_t)\}$ is strongly ergodic, i.e. for any Lipschitz continuous function $g : \mathbb{R} \to \mathbb{R}$, there exists a constant G such that

$$\lim_{t\to\infty} \sum_{t=1}^{T} g((D_t, C_t); (D_{t+1}, C_{t+1}); (D_{t+\Lambda}, C_{t+\Lambda})) = G \qquad \text{a.s.} \tag{16}$$

Moreover, $0 < \mu \stackrel{def}{=} \lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} C_t - \lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} D_t$ (a.s.).

The condition is related to that of asymptotic mean stationarity, see e.g. GRAY (1990). Beyond the case of i.i.d. aggregate capacity and demand pairs considered under (A), (A^f) encompasses a large variety of processes, for example:

(I) $\{(D_t, C_t)\}$ is stationary and ergodic

(II) $\{(D_t, C_t)\}$ is a so-called "world driven" process. Here, the distribution of (D_t, C_t) is time invariant but it depends on the state of the world W_t, with $\{W_t\}$ a Markov process with a finite or countable state space which is ergodic (i.e., the Markov chain has a single positive recurrent set of states). Thus the conditional distributions $\{(D_t, C_t)|W_t = w\}$ are time-invariant. Moreover, $\lim_{t\to\infty} W_t \overset{w}{=} W$ and $\lim_{t\to\infty}(D_t, C_t) \overset{w}{=} ((D, C)|W)$. See ZIPKIN (2000) for a detailed discussion of the use of world driven demand processes in inventory models.

(III) A third type of process satisfying (A^f) and modeling different types of intertemporal correlations, is where the process $\{(D_t, C_t)\}$ is autoregressive, e.g. a stable ARMA(p,q) process, i.e.

$$
\begin{aligned}
D_t &= \sum_{i=1}^{p} \varphi_i D_{t-i} + \sum_{j=1}^{q} \psi_j \epsilon_{t-j} + \epsilon_t \quad \forall t && (17)\\
C_t &= \sum_{i=1}^{p} \hat{\varphi}_i C_{t-i} + \sum_{j=1}^{q} \hat{\psi}_j \hat{\epsilon}_{t-j} + \hat{\epsilon}_t \quad \forall t && (18)
\end{aligned}
$$

where $\{\epsilon_t\}_{t=-\infty}^{+\infty}$ and $\{\hat{\epsilon}_t\}_{t=-\infty}^{+\infty}$ are independent sequences of i.i.d. random variables with finite second moments. A sufficient condition for the processes to be stable is that the characteristic polynomials $\Phi(z) = \sum_{i=1}^{p} \varphi_i z^i$ $[\hat{\Phi}(z) = \sum_{i=1}^{p} \hat{\varphi}_i z^i]$ and $\Psi(z) = \sum_{i=1}^{q} \psi_i z^i$ $[\hat{\Psi}(z) = \sum_{i=1}^{q} \hat{\psi}_i z^i]$ do not have common (complex) roots and that the roots of the former are outside the unit circle.

Lemma 2 *Assume the process $\{(D_t, C_t)\}$ is of type (I)–(III). Then $\{(D_t, C_t)\}$ is strongly ergodic.*

Proof: (I) Immediate, see e.g. Proposition 6.31 in BREIMAN (1992).

(II) The process $\{(W_t, W_{t+1}, \ldots, W_{t+\lambda})\}$ is a Markov process, whose Markov chain has a single positive recurrent set of states, i.e. there exists a state of the process with a finite expected recurrence time. Almost sure convergence of the limit to the left of (16) then follows from the renewal reward theorem.

(III) It suffices to prove strong ergodity of $\{D_t\}$ and $\{C_t\}$ separately. We prove the former; the proof of the latter is identical. Since the ARMA process is stable, there exists a con-

stant $0 < a < 1$ such that $D_t = \sum_{j=0}^{t} \alpha_j \epsilon_{t-j}$, with $|\alpha_j| < a^j$, see e.g. SAMORODNITSKY and TAQQU (1994). The sequence $\{D_t\}$ is non-stationary. Let $D_t^0 \stackrel{def}{=} \sum_{j=0}^{\infty} \alpha_j \epsilon_{t-j}$. $\left\{D_t^0\right\}$ is clearly stationary and it is well known to be ergodic. Fix a function $g : \mathbb{R}^\lambda \to \mathbb{R}$ that is Lipschitz continuous. By the argument for (I), there exists a constant G such that $\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} g(D_t^0, D_{t+1}^0, \ldots, D_{t+\lambda}^0) = G$ a.s.. To show that
$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} g(D_t, D_{t+1}, \ldots, D_{t+\Lambda}) = G$ a.s., as well, it suffices to show that for any $\delta > 0$,

$$| \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} [g(D_t, D_{t+1}, \ldots, D_{t+\lambda}) - g(D_t^0, \ldots, D_{t+\lambda}^0)]| < \delta \quad \text{a.s.} \tag{19}$$

Since for any integer $n \geq 1$,

$$\begin{aligned} & | \lim_{T\to\infty} \frac{1}{T} \sum_{t=n}^{T} [g(D_t, D_{t+1}, \ldots, D_{t+\Lambda}) - g(D_t^0, D_{t+1}^0, \ldots, D_{t+\Lambda}^0)]| \\ \leq \quad & \lim_{T\to\infty} \frac{1}{T} \sum_{t=n}^{T} |g(D_t, D_{t+1}, \ldots, D_{t+\Lambda}) - g(D_t^0, D_{t+1}^0, \ldots, D_{t+\Lambda}^0)| \end{aligned}$$

it follows from the Lipschitz continuity of $g(\cdot)$ that it suffices to show, for any $\delta > 0$, that an integer $n \geq 1$ exists such that

$\lim_{T\to\infty} \frac{1}{T} \sum_{t=n}^{T} |D_t - D_t^0| \leq \lim_{T\to\infty} \frac{1}{T} \sum_{t=n}^{T} \sum_{j=t+1}^{\infty} |\alpha_j||\epsilon_{t-j}| \leq \lim_{T\to\infty} \frac{1}{T} \sum_{t=n}^{T} \sum_{j=t+1}^{\infty} a^j |\epsilon_{t-j}|$
$= \lim_{T\to\infty} \frac{1}{T} \sum_{t=n}^{T} a^{t+1} \sum_{j=0}^{\infty} a^j |\epsilon_{-j-1}| \leq \lim_{T\to\infty} \sum_{t=n}^{T} a^{t+1} \{\lim_{r\uparrow 1} \sum_{j=0}^{\infty} r^j |\epsilon_{-j-1}|\}$
$= \lim_{T\to\infty} \frac{1}{T} \sum_{t=n}^{T} a^{t+1} \{\lim_{M\to\infty} \frac{1}{M+1} \sum_{j=0}^{M} |\epsilon_{-j-1}|\} \leq a^{n+1} E|\epsilon| < \delta \quad$ a.s.

where the equality follows from the Abel-Tauberian theorem and the next to last inequality from the law of large numbers. (Since the random variable ϵ has a finite second moment, E $|\epsilon| < \infty$.) The last inequality is satisfied for all $n \geq \lfloor \log(\rho/E|\epsilon|)/\log a \rfloor$. ■

To pursue the algorithm's performance analysis, note that I_t^0, the minimum reserve stock at the end of period t is now given by:

$$I_t^{0,f} = \max_{t+1 \leq s \leq t+\lambda} \sum_{r=t+1}^{s} (D_r - C_r) \tag{20}$$

instead of (1). (Using repeated substitutions in (1), note that $I_t^{0,f} = I_t^0$ when $\lambda = \infty$) In other words, assuming that a feasable solution exists, to ensure that a solution for the first t periods $[1, \ldots, t]$ can be extended into a feasible solution over the complete horizon $[1, \ldots, T]$, it is *necessary* and *sufficient* that $I_t \geq I_t^{0,f}$. (If $I_t < I_t^{0,f}$, aggregate demand in the periods $t+1, \ldots, s$ (for some $t+1 \leq s \leq t+\lambda$) exceeds $(\sum_{r=t+1}^{s} C_r + I_t)$, so demand in

$[t+1, s]$ can not be satisfied even when placing a full capacity order in each of the periods of this interval. At the same time, if $I_t \geq I_t^{0,f}$, the first period whose demand can not be met, has a period index greater than $t+\lambda$ and any additional inventory at the end of period t is of no use to meet this demand.)

Theorem 3 *Assume items have a fixed shelf life time $\lambda > 0$ and (A^f) holds.*

(a) Consider an (SP)-heuristic with $\tau = \eta\lceil \log T \rceil$ for some $\eta > 0$. The heuristic is asymptotically optimal, a.s., and it can be designed to run in $O(N^2T^2 \log T(\log N + \log\log T)^2)$ time as well.

(b) Consider an (EH)-heuristic with $\tau = \eta\lceil \log T \rceil$ for some $\eta > 0$. The heuristics is asymptotically optimal, a.s., and it can be designed to run in $O(N^2T^4(\log N + \log T + \log^2 N/\log T))$ time as well.

Proof: The proof is analogous to that of Theorem 1 and 2 and is, in fact, simpler. $\lim_{T\to\infty} \frac{z^I - z^*}{T} = 0$ a.s., is verified as in the proof of Theorem 1. Moreover, it was shown there that for $\lim_{T\to\infty} \frac{z^{II} - z^I}{T} = 0$ a.s., it is sufficient to verify that (10) holds. Since $L_\ell \leq \lambda$, (10) reduces to showing that $\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} I_t^{0,f}$ converges to a constant a.s.. This, however, follows from (A^f), since, by (16), $I_t^{0,f}$ is a Lipschitz continuous function of $\{(D_{t+1}, C_{t+1}), (D_{t+2}, C_{t+2}), \ldots, (D_{t+\lambda}, C_{t+\lambda})\}$.

We now verify the complexity bounds. Under a fixed life time, the minimum cost network flow problem to be solved in each node of the b&b-trees, associated with the different interval instances, now needs to be solved by a standard method, rather than the AHUJA and HOCHBAUM (2004) method. The best strongly polynomial time algorithm to solve minimum cost network flow problems is due to ORLIN (1989). The network flow model has a source, a sink and two sets of nodes; the first set has a node for every period and the second one has a node for every period / item combination. Thus, the model has O($N\tau$) nodes and O($N\tau$) arcs in the (SP)-heuristic and O(NT) nodes and arcs in the (EH)-heuristic. ORLIN's method solves the problem therefore in O($N^2\tau^2 \log^2 N\tau$) and O($N^2T^2 \log^2 NT$), respectively. Since $J = T/\tau$ interval instances are solved and since in each, in the worst case, all 2^τ nodes of the b&b tree need to be evaluated, the complexity bounds follow readily. ■

Acknowledgement: We are greatly indebted to ASSAF ZEEVI for many helpful suggestions and comments regarding earlier versions of this paper.

References

AGGARWAL, A., J. PARK. 1993. Improved algorithms for economic lot size problems. *Operations Research* **41** 549–571

AHUJA, R., D. HOCHBAUM. 2004. Solving Linear Cost Dynamic Lot Sizing Problems in O($n \log n$) time. Working paper, University of California, Berkeley.

ANILY, S., J. BRAMEL 1999. Vehicle Routing and the Supply Chain. In: S. TAYUR, R. GANESHAN, M. MAGAZINE (Edts): *Quantitative Models for Supply Chain Management*. Kluwer, Boston.

ASMUSSEN, S.. 1987. *Applied Probability and Queues*, Wiley, New York.

BELVAUX, G., L. WOLSEY. 2000. *bc-prod* A specialized branch-and-cut system for lot-sizing problems. *Management Science* **46** 724–738

BELVAUX, G., L. WOLSEY. 2001. Modelling practical lot-sizing problems as mixed-integer programs. *Management Science* **47** 993–1007

BRAMEL, J., D. SIMCHI-LEVI. 1997. *The Logic of Logistics*. Springer, New York.

BREIMAN, L. 1992. *Probability*. SIAM's Classics in Applied Mathematics series. SIAM, Philadelphia.

CHAN, L., D. SIMCHI-LEVI. 1998. Probabilistic Analyses and Algorithms for Three-Level Distribution Systems. *Management Science* **44** 1562–1576

COFFMAN, E., G. LUEKER. 1991. Probabilistic Analysis of Packingand Partitioning Heuristics. Wiley, New York.

FEDERGRUEN, A., D. SIMCHI-LEVI. (1991). Analytical Analysis of Vehicle Routing and Inventory Routing problems. In: M. BALL, T. MAGNANTI, C. MONMA, G. NEMHAUSER (eds): *Handbooks in Operations Research and Management Science, Vol. Network Routing*. North-Holland, Amsterdam.

FEDERGRUEN, A., M. TZUR. 1991. A Simple Forward Algorithm to solve General Dynamic Lot-Sizing Models with n Periods in O(n log n) or O(n) time. *Management Science* **37** 909–925

FEDERGRUEN, A., M. TZUR. 1999. Time-Partitioning Heuristics: Application to One-Warehouse, Multiitem, Multiretailer Lot-Sizing Problems. *Naval Research Logistics* **46** 463–486

FEDERGRUEN, A., J. MEISSNER, M. TZUR. 2002. Progressive Interval Heuristics for Multi-Item Capacitated Lot Sizing Problems. (Conditionally accepted in *Operations Research*)

FISHER, M., D. HOCHBAUM. 1980. Probabilistic Analysis of the K-median problem. *Mathematics* *of Operations Research* **5** 27–34

FLORIAN, M., J. LENSTRA, A. RINNOY KAN. 1980. Deterministic production planning: algorithms and complexity. *Management Science* **26** 669–679

GALLEGO, G., D. SIMCHI-LEVI 1996. On the Effectiveness of Direct Shipping Strategy for the One-Warehouse Multi-Retailer R-Systems. *Management Science* **36** 240–248

GAVISH, B., R. JOHNSON. 1990. A fully polynomial approximation scheme for single-product scheduling in a finite capacity facility. *Operations Research* **38** 70–83

GRAY, R.. 1990. *Entropy and Information Theory*. Springer Verlag, Berlin.

HAIMOVICH, M., A. RINNOOY KAN.1985. Bounds and Heuristics for Capacitated Routing Problems. *Mathematics of Operations Research* **10** 527–542

HSU, V.. 2000. Dynamic Economic Lot Size Model with Perishable Inventory. *Management Science* **46** 1159–1169

KARP, R. 1977. Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman Problem. *Mathematics of Operations Research* **2** 209–224

KUIK, R., M. SOLOMON, L. VAN WASSENHOVE. 1994. Batching decisions: structures and models. *European Journal Of Operational Research* **18** 213–263

NAHMIAS, S.. 1982. Perishable Inventory Theory: A Review. *Operations Research* **30** 680–708

ORLIN, J.. 1989. A faster strongly polynomial minimum cost flow algorithm. *Journal of the ACM* **36** 873–886

RHEE, W.. 1993. Inequalities for Bin Packing – II. *Mathematics of Operations Research* **18** 685–693

RHEE, W. and M. TALAGRAND. 1989. Concentration Inequalities for the K-Median Problem. *Mathematics of Operations Research* **14** 189–202

RHEE, W. and M. TALAGRAND. 1987. Martingale Inequalities and NP-Complete Problems. *Mathematics of Operations Research* **12** 177–181

ROMEIJN, H., D. ROMERO MORALES. 2001. A probabilistic analysis of the multi-period sourcing problem. *Discrete Applied Mathematics* **42** 301–328

SALOMON, M.. 1990. Deterministic lotsizing models for production planning. Ph.D. dissertation, Erasmus University.

SAMORODNITSKY, G., M. TAQQU. 1994. *Stable Non-Gaussian Random Processes*. Chapman Hall, Boca Raton, Florida.

Stadtler, H.. 2003. Multilevel Lot Sizing with Setup Times and Multiple Constrained Resources: Internally Rolling Schedules with Lot-Sizing Windows. *Operations Research* **51** 487–502

Suerie, C., H. Stadtler. 2003. The Capacitated Lot-Sizing Problem with Linked Lot Sizes. *Management Science* **49** 1039–1054

Van Hoesel, C., A. Wagelmans. 2001. Fully Polynomial Approximation Schemes for Single-Item capacitated lot-Sizing Problems. *Mathematics of Operations Research* **26** 339–357

Wagelmans, A., C. van Hoesel, A. Kolen. 1992. Economic lot-sizing: an O(n log n)-algorithm that runs in linear time in the Wagner-Whitin case. *Operations Research* **40** 145–156

Zipkin, P.. 2000. *Foundations of Inventory Management*. McGraw Hill, New York.