

Stochastic Programming Approaches to Stochastic Scheduling

John R. Birge*

Department of Industrial and Operations Engineering
The University of Michigan
Ann Arbor, Michigan, USA 48109

M. A. H. Dempster†

Department of Mathematics
University of Essex
Colchester, England CO4 3SQ

Abstract : Practical scheduling problems typically require decisions without full information about the outcomes of those decisions. Yields, resource availability, performance, demand, costs, and revenues may all vary. Incorporating these quantities into stochastic scheduling models often produces difficulties in analysis that may be addressed in a variety of ways. In this paper, we present results based on stochastic programming approaches to the hierarchy of decisions in typical stochastic scheduling situations. Our unifying framework allows us to treat all aspects of a decision in a similar framework. We show how views from different levels enable approximations that can overcome nonconvexities and duality gaps that appear in deterministic formulations. In particular, we show that the stochastic program structure leads to a vanishing Lagrangian duality gap in stochastic integer programs as the number of scenarios increases.

Keywords: stochastic programming, hierarchical decision making, stochastic scheduling, queueing networks, large deviations, turnpikes, Lagrangian relaxation.

* This author's work was supported in part by the National Science Foundation under Grants ECS 88-15101, ECS 92-16819, and SES 92-11937.

† This author's work was supported in part by the Natural Sciences and Engineering Research Council of Canada under Grant A-5489 and by the UK Engineering and Physical Sciences Research Council under Grants J90855 and K17897.

1. Introduction

Scheduling decisions involve consideration of many quantities that are not known with certainty. Important characteristics such as process yield, processing time, resource availability, product quality, demand, input costs and eventual revenues may all be random quantities with varying levels of certainty. The inevitability of these uncertainties has promoted interest in stochastic scheduling (see, for example, Dempster *et al* [1982] and Righter [in Shaked and Shanthikumar,1994]).

The great majority of the stochastic scheduling literature concentrates on structural results for optimal solutions that show that, under certain conditions, it may be optimal to follow a fixed permutation schedule or use a form of an index rule (see, for example, Gittins [1981], Glazebrook [1981], Pinedo [1983], Shanthikumar and Yao [1992]). While these results are quite useful in many settings, especially in terms of overall system evaluation, they do not, generally, extend directly to near-term operating conditions. Our interest in this paper is to use stochastic programming formulations to explore the relationship between broad system results for planning purposes and short-term results for actual operations.

Our approach fits into a hierarchical approach to decision making (see also Dempster [1982], Dempster *et al* [1983] and Dempster [1994]) in which we view the decision process as breaking down into: (1) long-term, *strategic* decisions about overall capacity and scope (market); (2) medium-term, *tactical* or planning decisions about aggregate production for known items or services, or fixed orders; and (3) short-term, *operational* decisions for current situations. Scheduling decisions impact on each level of this hierarchy since they may determine overall capacity and fixed term production, as well as decisions on what or how to produce in the present.

Hierarchical stochastic programming provides a unifying framework for this analysis. All levels of the hierarchy can appear in the same model to allow for various methods of decomposition, approximation and solution. In this paper, we explore different types of approximation associated with each level of the hierarchy and show how that approximation may aid the overall decision process.

Our results are consistent with other attempts to unify stochastic and dynamic optimization as in, for example, Bertsimas [1994]. These are most compatible when we take a long-term view of design decisions

for overall performance and capacity. In Section 2, we explore this strategic decision making level. Our approximation in this section is to assume that operational scheduling effects are well approximated by heavy traffic fluid, or Markov-modulated fluid, queueing approximations, Dempster [1994]. We use an example from telecommunications (see Dempster, Key and Medova [1995]) to illustrate how this approximation and large deviation theory can reduce a complicated stochastic program to a simple deterministic problem.

When higher level capacity and scope decisions are set, aggregate production scheduling becomes a dominant concern. In Section 3, we study this level of the decision hierarchy when production decisions can be approximated in a convex model. We give some justification for this in a planning setting. We then provide results from Birge and Dempster [1992] to give optimality conditions for this problem and provide a turnpike result and conditions for optimal cyclic schedules. This also justifies a match-up scheduling approach to disruptions, which was discussed for deterministic systems in Bean *et al* [1991].

At an operational level however, it is often difficult to eliminate the effects of nonconvexities. In this case, other approaches are necessary. In Section 4, we show how to use a Lagrangian relaxation approach to the general stochastic scheduling problem with setups. This method has been used in power system scheduling (see Takriti, Birge and Long [1994a]) with small duality gaps. Extending results from Bertsekas [1982], we show, in fact, that the duality gap in this relaxation approach actually decreases to zero as the number of samples in a discrete approximation procedure to the stochastic problem increases.

2. The Strategic Level and Queueing Approximation.

Following the original ideas of G.B. Dantzig (see. e.g., Dantzig [1963]) in the context of military logistics, from the 1940's management science has been concerned with models at the macro, meso and micro levels of detail to support corporate decision-making in planning, management and control. Integrated hierarchical mathematical and computer models reflect the classical three-level military planning concepts of strategic long-run, tactical medium-run and operational short-run. As decision-making moves down the corporate hierarchy, it becomes increasingly detailed and involves many more, but smaller, uncertainties. The modelling involved at successively lower levels reflects these differences – paralleling the macro, meso and micro scale mathematical models of classical physics – increasing in complexity at each level. In a stationary corporate

environment, operational planning models – involving mainly management and control functions – can become extremely complex. In a highly dynamic uncertain environment, useful mathematical and computer models tend to become simpler, as it is the strategic and tactical decisions involving rarer major uncertainties which are critical for survival (see Table 1).

		Lead Time	Cost	Complexity	Uncertainty
Level 1	Strategic	↑	↑	↓ Macro	↑
Level 2	Tactical	↓	↓	↓ Meso	↓
Level 3	Operational	↓	↓	↓ Micro	↓

Table 1. Strategic and tactical planning levels handle complexities and uncertainties by aggregation in a stable environment in order to focus on rare major environmental change.

Many problems in manufacturing, logistics and telecommunications are modelled at the operational level by open queueing system (see, for example, Buzacott and Shanthikumar [1980], Harrison [1985], Chen and Mandelbaum [1991], Sethi and Zhang [1995] and Kelly [1991]), although the study of the behaviour of these systems with *active* scheduling policies is in its infancy (see, however, Bertsimas and Niño-Mora [1993], Bertsimas [1994] and Garbe and Glazebrook [1995] for situations in which simple index rules can be justified). In this section we first outline the aggregation of minor operational event uncertainties by changes of time and space scales appropriate to macro level strategic decisions and modelling for open queueing systems with finite buffers. (Here we ignore the intervening meso – tactical – level regarding these queueing approximations; the reader is referred to Harrison [1985], Chen and Mandelbaum [1991] and Dempster [1994] for more details.) Using functional central limit theory, the complex operational stochastic model is aggregated to a simple deterministic fluid flow model for strategic planning purposes.

For practical applications of hierarchical planning models, however, this aggregation is often *too* strong in that the resulting approximation ignores the occasional random extraordinary influences of extreme tactical and operational level considerations on strategic planning. For modeling at this meso-macro boundary, Markov-modulated fluid flow models are appropriate. We illustrate strategic planning at this level with a hierarchical telecommunications planning model (treated in detail in Medova [1993], Dempster [1994] and Dempster, Key and Medova [1995]) which uses large deviations theory to produce, from the original chance-

constrained stochastic programming formulation, a simple deterministic mathematical programming model for network dimensioning and traffic balancing which can be used as a framework for fast real-time routing and scheduling algorithms. The section concludes by briefly indicating some extensions of this approach to other scheduling contexts.

Consider now an *open* (i.e. exogenous input) *queueing network* on an arbitrary graph $G(N, A)$ of *nodes* (with finite buffer capacity) and (directed) arcs, with *random routing of particles* from each node on each of its output links to each other node. (Classical sources for queueing networks are Kleinrock [1975] and Kelly [1979].) We assume arbitrary independent identically distributed exogenous input inter-arrival and service time processes at each node subject to the existence of a dynamic equilibrium (long-run) limiting process distribution for the system. This stationary stochastic process is characterized by an *exogenous arrival rate* λ_j^0 , a potential *service rate* μ_j and the *switching fractions* p_{jk} of particles that are routed directly to node k on link (j, k) after service at node j on each of the $J := |N|$ nodes j in the network. The corresponding row vector/matrix triplet (λ_0^0, μ', P) specifies the long-run average performance of the system and – assuming that particles arriving at a node with a full buffer are lost to the system – the *total inflow vector* λ' of *total arrival rates* λ_j at nodes $j = 1, \dots, J$ is the maximum solution of the *traffic equations*

$$\lambda' = \lambda_0^0 + (\lambda' \wedge \mu')P, \tag{1}$$

where \wedge denotes coordinatewise minimum.




Figure 1. Example of a Three Node Network.

An example of this network appears in Figure 1. In this case, external arrivals with rate λ_0^j combine with output particles from the random routing with rates, $p_{ij}\mu_i$, to form the overall arrival rate λ_j at each node. Arrivals which encounter full buffers are lost to the system.

Since the network is open, the *routing transition matrix* P has spectral radius $\sigma(P) < 1$. (A *closed* network has $\sigma(P) = 1$ and $\lambda_0^j := 0$.) The *traffic intensity* at node j is defined to be $\rho_j := \lambda_j/\mu_j$ and node j is defined to be a *bottleneck* if $\rho_j = 1$ and a *nonbottleneck* if $\rho_j < 1$. Let

$$\mathbf{Q}' := \{\mathbf{Q}'(t) := (\mathbf{Q}_1(t), \dots, \mathbf{Q}_J(t)) : t \geq 0\}$$

denote the integral *equilibrium queue length* process representing the network node (i.e. buffer and server) occupancies. Not surprisingly, it may be shown that the macro level fluid approximation of this process is supported on the set of bottleneck nodes.

Moreover, it may be shown that \mathbf{Q} can be expressed in terms of the *potential net throughput* process

$$\mathbf{X}' := \{\mathbf{X}'(t) := (\mathbf{X}_1(t), \dots, \mathbf{X}_J(t)) : t \geq 0\},$$

which represents the difference between the equilibrium cumulative input and *potential* output processes, and the *equilibrium lost output* process (due to empty nodes)

$$\mathbf{Y}' := \{\mathbf{Y}'(t) := (\mathbf{Y}_1(t), \dots, \mathbf{Y}_J(t)) : t \geq 0\}.$$

Together, \mathbf{X}' and P completely specify the system as *almost surely* (a.s.) the unique solution of

$$\mathbf{Q}' = \mathbf{X}' + [\mathbf{Y}' - \mathbf{Z}'(I - P)] \leq B', \tag{2}$$

where B' is the constant process representing fixed *node capacities* and \mathbf{Z}' represents the *buffer overflow loss* process, subject to the appropriate a.s. nonnegativity conditions and the requirement that \mathbf{Y}' and \mathbf{Z}' are a.s. nondecreasing and increase from 0 (the identically zero process) only when the corresponding coordinates of \mathbf{Q}' and $B' - \mathbf{Q}'$ respectively are 0.

A basic representation of the identity in (2) appears in Figure 2 for a single node with no branching probabilities and a single realization of the processes. In this figure, we assume that the system has no buffer

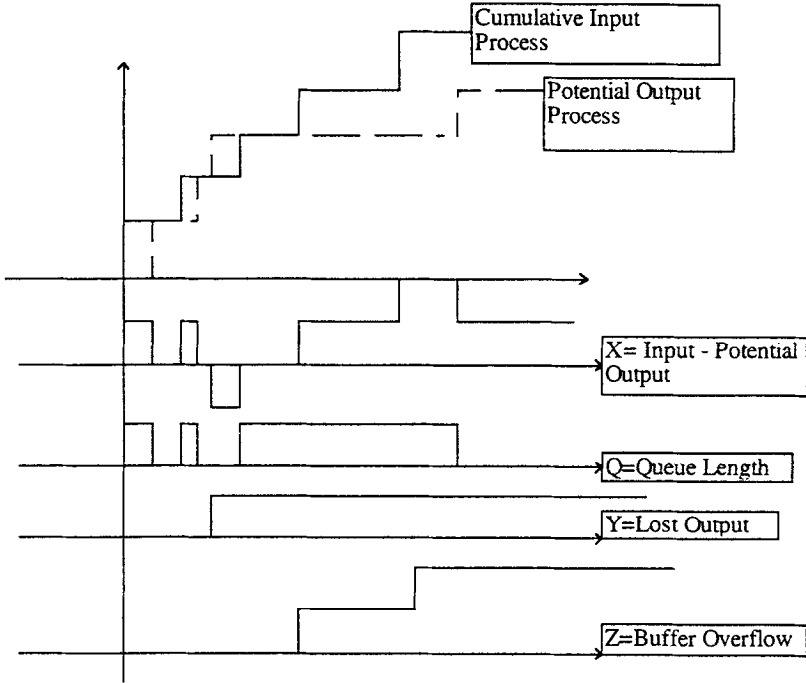


Figure 2. The Queue Length Process in Terms of the Input and Potential Output Processes.

so that the maximum queue length is one. The difference between the input and potential output process is the potential net throughput, X . To obtain the queue length process, Q , the lost output process, Y , must be added whenever the output is limited due to the empty state. The buffer overflow process, Z , must also be subtracted when an arrival is lost due to the buffer limit.

The conditions defining the queue length process may be neatly summarized by observing that Q' must be pathwise the a.s. unique solution of the linear *order complementarity problem* (see e.g. Dempster [1994])

$$(OCP) \quad Q' := X' + [Y' - Z'(I - P)] \leq B' \tag{3}$$

$$Y' \geq 0' \quad Z' \geq 0' \quad Q' \geq 0' \tag{4}$$

$$Q' \wedge \Delta Y' = 0' \quad [B' - Q'] \wedge \Delta Z' = 0 \tag{5}$$

posed in the Skorohod space of left-limited right-continuous functions of time.

Here we note that we may express

$$Y' = \sum_{t_i \leq (\cdot)} \Delta Y'(t_i) \quad \text{and} \quad Z' = \sum_{\tau_i \leq (\cdot)} \Delta Z'(\tau_i)$$

in terms of the jumps $\Delta Y'(t_i)$ and $\Delta Z'(\tau_i)$ of the respective processes at jump epochs t_i, τ_i up to time t . It follows that Y' and Z' are a.s. the least element processes satisfying the inequalities (3) and (4). Thus (OCP) provides a complete description of the equilibrium behaviour of the specified finite buffer open queueing system at the operational level of equilibrium stochastic particle flow.

Aggregating numbers of particles by the scaling X'/n , accelerating time by the scaling nt and applying a suitable functional law of large numbers for stochastic processes (see e.g. Chen and Mandelbaum [1991]), the *expected* queue length process \bar{Q}' emerges at the strategic level in the a.s. limit as $n \rightarrow \infty$ as (equilibrium) *deterministic* fluid flow. This flow is regulated by a deterministic increasing process for *expected* lost output \bar{Y}' and *expected* buffer overflow \bar{Z}' and driven by the deterministic *expected* potential net throughput process.

$$\bar{X}' := \{\bar{X}'(t) := [(\lambda' - \mu') \vee \mu']t : t \geq 0\}, \tag{6}$$

where \vee denotes coordinatewise maximum. Taken together these deterministic fluid flows satisfy a deterministic version of (OCP). As mentioned above, positive coordinates of \bar{Q}' and \bar{Z}' are associated with bottleneck nodes, as is appropriate to strategic decisions such as decreasing exogenous input rates or increasing service rates at bottleneck nodes or changing routing frequencies in the network so as to decrease their input loads. Notice that at this macro *expected* level – which is by no means the case at lower levels – buffer sizes at individual nodes are irrelevant.

To illustrate a situation at the meso-macro level in which such dimensioning becomes relevant to strategic planning, we turn to a hierarchical telecommunications model with extensions to more general scheduling models. Following Dempster[1994], the general 3-level hierarchical telecommunication *network planning*

model is given by

$$\begin{aligned}
 \min_{x_1, x_2} \{ & f(x_1) + && \text{Level 1 : Topology Design} \\
 & g(x_1, x_2) + && \text{Level 2 : Node and Link Capacity Allocation} \\
 & \mathbb{E}h(x_1, x_2, \mathbf{d}) \} && \text{Level 3 : Traffic Management} \quad (7)
 \end{aligned}$$

subject to $x_1 \in X_1, x_2 \in X_2,$

where, in a manufacturing setting, topology design includes plant layout, node and link capacity allocation involves the material handling system design, and traffic management determines routing and sequencing. In this model, x_1 and x_2 represent a choice from usually an exponentially large number of integer design vectors in X_1 and X_2 , representing node and link *placements* and capacity allocations respectively, while $\mathbb{E}(\cdot)$ denotes mathematical expectation in terms of the stationary distribution of the *traffic demand* process with state \mathbf{d} in terms of bandwidth requirements on the allocated directed links. The function given by $h(x_1, x_2, \mathbf{d})$ represents the value of a *multicommodity flow* problem (see e.g. Hu[1969]) which allocates the traffic demand \mathbf{d} by *origin-destination pair* to routing paths over the physical network topology determined by x_1 within the capacities determined by x_2 . Since real telecommunication networks tend to grow incrementally, in practical planning problems the physical topology of the network (determined by x_1 in the model) is usually taken as given, so that the main strategic decisions in telecommunications network management concern node (i.e. switch) and link (i.e. optic fibre) capacity allocation – particularly so as to allow resilience of the network in the face of random node and link failures. In manufacturing systems, the nodes become machines and storage areas while the links are material handling routes, such as conveyors and automatic-guided vehicle paths.

Modern telecommunication networks involve many layers, both physically and with respect to traffic demand time-scales such as private network contracts (months or years), calls (minutes), packet bursts (milliseconds) and bit flow (microseconds). In manufacturing, similar contracts run from product lifetime supply (years) to hourly just-in-time deliveries. To illustrate the influence of meso level rare events on macro level strategic decisions, we shall present a simplified 2-level special case of (7) involving its second and third level. Our model is a special case of a more detailed model of *asynchronous transfer mode* (ATM) network dimensioning and traffic routing and balancing developed in Medova[1993], Dempster[1994] and

Dempster, Key and Medova[1995], to which the reader is referred for more details. In particular, the optimal traffic management cost function h of (7) will result from solving a certainty equivalent form of a *chance-constrained* stochastic program. Applications to telecommunications network planning with the alternative *recourse* stochastic programming formulation have recently been given in Sen *et al* [1992] and Bai *et al* [1994]. In general, the chance-constrained approach is natural to *grade-of-service* (GoS) considerations expressed in terms of acceptable blocking or loss probabilities at various time-scale layers of the network and usually results in a very much more compact computational representation of the problem than the recourse formulation. Similarly, in manufacturing systems, a model can either incorporate penalties for shipping delays for recourse formulations or use service level constraints of the chance-constrained form.

To see this in the telecommunications context, consider once more a fixed network topology represented by a directed graph $G(N, A)$ and consider *stationary stochastic* (packet, cell or bit) *traffic flows* f_w (in packets, cells or megabits per second) between each origin-destination (OD) pair $w \in W$ in the network during a specified traffic regime. The network dimensioning and traffic management problem we shall consider is to allocate *link capacity* C_a , $a \in A$, and route traffic f_p on *paths* $p \in P_w$, a set of specified routes through the network from origin to destination node w , so as to optimize the expected net revenue occurring from carrying the stochastic traffic in terms of total OD pair bit flows

$$f_w = \sum_{p \in P_w} f_p \quad w \in W \tag{8}$$

offered to the network. For simplicity, we treat a linear objective, but the arguments given below can easily be extended to (level) separable convex cost functions and to the provision of node (switch), as well as link capacity (see Dempster, Medova and Moise [1995]).

Specifically, we consider the chance-constrained 2-level problem

$$\min_{a \in A} \left[\sum_{a \in A} b_a C_a - \mathbb{E} \left(\sum_{w \in W} r_w \sum_{p \in P_w} f_p \right) \right] \tag{9}$$

$$\text{subject to } P \left\{ \sum_{p \in P_w} f_p \geq D_w \right\} \leq g_{\text{call}} \quad w \in W \tag{10}$$

$$P \left\{ \sum_{p \in Q_a} f_p \geq C_a \right\} \leq g_{\text{call}} \quad a \in A \tag{11}$$

$$f_p \geq 0 \quad \text{a.s.} \quad p \in P_w \quad w \in W. \tag{12}$$

The objective (9) of this optimization problem contains coefficients representing the unit *cost of link capacity provision* b_a , $a \in A$, and unit *OD pair net revenue* r_w , $w \in W$, resulting from carrying traffic between OD pairs. The first chance, or *probabilistic*, constraint (10) states that the probability of offered stochastic OD pair traffic flow f_w exceeding a specified *demanded capacity* level D_w must not exceed the GoS for *call refusal probability* g_{call} . The probabilistic constraint (11) states that call GoS g_{call} must be maintained for the sum of offered traffic flows on the set Q_a of *all paths passing through* each (directed) *link* $a \in A$. The constraint (12) requires the a.s. nonnegativity of traffic flows.

Problem (9-12) is completely analogous to the manufacturing situation where W corresponds to a set of products, P_w includes alternate production routings to complete product w , and C_a represents capacity of both paths and machines. At this level, we would assume that operational level sequencing effects would not be significant in reducing overall capacity.

In the result below we shall use a large deviation bound to approximate the OD pair demanded capacities D_w which meet the probabilistic constraints (10) and (11) exactly. It is well known that a linearly chance constrained problem of the above form has a deterministic equivalent problem which is obtained by replacing the probabilistic constraints involving random variables with equivalent constraints written in terms of deterministic variables which meet the probability restrictions as nearly as possible. These equivalent constraints are usually found by inverting the joint distribution of the random variables. However, if D_w is specified (approximately) by a conservative traffic bound derived from large deviation theory, we may replace the stochastic traffic flows f_p by *deterministic fluid flows* f_p , $p \in P_w$, $w \in W$, which represent traffic flows just at the rate above which fluctuations of the stochastic flows f_p would begin to exceed call GoS requirements. Then we can formulate the (approximate) deterministic equivalent of our problem as the *path and capacity allocation* (PCA) problem of Medova [1993], viz.

$$\text{(PCA)} \quad \min \quad \left[\sum_{a \in A} b_a C_a - \sum_{w \in W} r_w \sum_{p \in P_w} f_p \right] \tag{13}$$

$$\text{subject to} \quad \sum_{p \in P_w} f_p = D_w \quad w \in W \tag{14}$$

$$\sum_{p \in Q_a} f_p \leq C_a \quad a \in A \tag{15}$$

$$f_p \geq 0 \quad p \in P_w \quad w \in W. \tag{16}$$

This model dimensions link capacities and routes traffic at the maximal flow rates (bandwidths) consistent with the specified individual call GoS. (In fact, the multilayer multiservice *effective bandwidths* D_w , $w \in W$, developed in Dempster, Key and Medova [1995] for ATM networks are also consistent with more stringent GoS criteria at the burst and cell timescale layers of the network.) The linear programming problem specified by the PCA model (13)-(16) is a 2-level deterministic network planning and traffic management model in which the strategic level (first stage) dimensions link capacities and the lower level (second stage) allocates and routes (approximate) deterministic equivalent traffic between all specified OD pairs in the network so as to maintain GoS for its corresponding stochastic traffic on all paths and links. For *fixed* link capacities, the second stage problem is a multicommodity flow problem in *compact* format instantiated using OD pair-path and arc-path 0-1 incidence matrices.

We summarize the above discussion in the following proposition.

Theorem 1. *The chance-constrained stochastic programming problem (9)-(12) with Poisson flows and random call routing has an approximate deterministic equivalent linear programming problem given by (13)-(16) whose accuracy depends on the tightness of the large deviation bound employed to derive it.*

Remark. We establish the theorem using the least refined of the large deviation bounds, the Chernoff bound, but both reducing constants and, under appropriate stochastic assumptions, tighter bounds, such as the Bahadur-Rao bound, can be obtained (see Dempster, Key and Medova [1995] for references). ■

Proof: First consider the total offered traffic f_w between an arbitrary specified OD pair $w \in W$. Recall that the *Chernoff bound* is given by

$$P\{f_w \geq D\} \leq \exp \{-\max_s [sD - \mu_{f_w}(s)]\}, \tag{17}$$

where the logarithmic generating function μ_{f_w} of f_w is defined as

$$\mu_{f_w}(s) := \ln \mathbb{E}e^{sf_w}. \tag{18}$$

Rearranging (17) yields

$$-\ln P\{\mathbf{f}_w \geq D\} \leq \max_s \{sD - \mu_{\mathbf{f}_w}(s)\}$$

and consequently to approximately invert the constraint (10), taking account of (8), we must solve

$$\max_s \{sD - \mu_{\mathbf{f}_w}(s)\} = -\ln g_{\text{call}} \tag{19}$$

to obtain

$$D_w = \mu'_{\mathbf{f}_w}(s^*), \tag{20}$$

for the optimal s^* in (19), and the corresponding deterministic constraint

$$f_w \leq D_w. \tag{21}$$

If we assume that the stationary traffic flow \mathbf{f}_w has a Poisson distribution with parameter ρ_w , that

$$f_w = \sum_{p \in P_w} f_p$$

and that each arriving particle (call) of the Poisson flow \mathbf{f}_w is independently randomly routed on a path $p \in P_w$ with probability f_p/f_w , then it follows that

$$\sum_{p \in P_w} f_p \leq D_w \tag{22}$$

is the deterministic equivalent of (10), since

$$P\{f_p := f_p \mathbf{f}_w / f_w \geq f_p D_w / f_w := D_p\} \leq g_{\text{call}}. \tag{23}$$

Moreover, since $w \in W$ was arbitrary, we have that

$$P\left\{\sum_{p \in Q_a} f_p \geq C_a\right\} \leq g_{\text{call}} \quad a \in A$$

provided that

$$\sum_{p \in P_a} f_p \leq C_a \quad a \in A, \tag{24}$$

since it may be shown that for sufficiently small g_{call} (typically 10^{-3} or 10^{-4}), the effective bandwidths D_p determined by the Chernoff bound are additive, i.e.

$$\sum_{p \in P_w} D_p = D_w, \tag{25}$$

(see Hui and Karasan [1995]). Hence (24) is the deterministic equivalent of (11) as required. ■

This theorem can be extended in a number of ways. For telecommunications network planning perhaps the most important generalization is to multiservice traffic mixtures of independent Poisson streams of Markov modulated fluid calls with negative exponential holding times (see Dempster, Key and Medova [1995]). In this context the (PCA) model provides a framework for real-time routing and rerouting of actual calls based on stochastic knapsack and bin packing models.

Since Theorem 1 applies to the basic queueing network model treated earlier, it (and its generalizations) can be used to study hierarchical planning problems in, for example, flexible manufacturing. Strategic and tactical decisions might involve shop layout, factory and machine centre product capability, machine centre sizing, setups, etc; while operational decisions would involve work-in-progress scheduling, processing and routing between machine centres. We consider tactical level decisions on overall production in the next section.

3. The Tactical Level and Convex Approximation.

The example of the previous section fits the strategic framework of a manufacturer considering capacity planning. The analysis considered only aggregate production service requirements without consideration of the timeliness of order fulfillment. At a lower decision level, the time element becomes more important. We first consider appropriate methods for aggregate production (for example, daily or monthly) decisions.

Our approximation is based on a convex model in contrast to many scheduling models, in which non-convexities immediately appear through disjunctive constraints. This model gives a good approximation at a tactical level where the production system appears large enough that a continuous approximation to setup conditions can apply. In this section, we follow the model in Birge and Dempster ([1992,1995]). Our model is similar to the continuous time model in Solel [1987] (see also Dempster and Solel [1987]) but we consider a discrete time model that allows us a wider class of problems. Related results for general stochastic optimization models appear in Dempster [1988], Flåm [1983,1985], Kushner [1972], Arkin and Evstigneev [1979], Rockafellar and Wets [1983] and Hiriart-Urruty [1982]. One of our main results in this section con-

cerns turnpike optimality which in the deterministic case may be found in McKenzie [1976] and in stochastic models in, for example, Sethi *et al.* [1992].

To set the stage for our model, we assume a data process, $\omega := \{\omega_t : t = 0, \dots\}$ in a (canonical) probability space (Ω, Σ, μ) . We also assume a decision process $x := \{x_t : t = 0, \dots\}$ such that x is a measurable function $x : \omega \mapsto x(\omega)$. The space of the decision processes is the space of *essentially bounded functions*, $L_\infty^n := L_\infty(\Omega \times N, \Sigma \times \mathcal{P}(N), \mu \times \#; \mathbb{R}^n)$, where \mathcal{P} is the power set and $\#$ is counting measure. Associated with the data process is a filtration $\mathbb{F} := \{\Sigma_t\}_{t=1}^\infty$, where $\Sigma_t := \sigma(\omega^t)$ is the σ -field of the *history process* $\omega^t := \{\omega_0, \dots, \omega_t\}$ and the Σ_t satisfy $\{0, \Omega\} \subset \Sigma_0 \subset \dots \subset \Sigma$.

A fundamental property of the decision process at time t is that it must only depend on the data up to time t , i.e. x_t must be Σ_t -measurable. An alternative characterization of this *nonanticipative* property is that $x_t = \mathbb{E}\{x_t | \Sigma_t\}$ almost surely (a.s.), $t = 0, \dots$, where $\mathbb{E}\{\cdot | \Sigma_t\}$ is conditional expectation with respect to the σ -field Σ_t . Using the projection operator $\Pi_t : z \mapsto \pi_t z := \mathbb{E}\{z | \Sigma_t\}$, $t = 0, \dots$, on L_∞^n , this is equivalent to

$$(I - \Pi_t)x_t = 0, \quad t = 0, \dots \tag{26}$$

We let \mathcal{N} denote the closed linear subspace of nonanticipative processes in L_∞^n and denote by $\Pi := (\Pi_0, \Pi_1, \dots)$ the projection operator from L_∞^n onto \mathcal{N} .

Our general optimization model is to find

$$\inf_{x \in \mathcal{N}} \mathbb{E} \sum_{t=0}^\infty f_t(\omega, x_t(\omega), x_{t+1}(\omega)). \tag{27}$$

where “ \mathbb{E} ” denotes expectation with respect to Σ . We use the notation x_t and f_t to denote respectively x_t and f_t as functions of ω , i.e. as random entities. Expression (27) then becomes

$$\inf_{x \in \mathcal{N}} \mathbb{E} \sum_{t=0}^\infty f_t(x_t, x_{t+1}), \tag{28}$$

with objective $F(x) := \mathbb{E} \sum_{t=0}^\infty f_t(x_t, x_{t+1})$.

We assume in (28) that the objective components f_t are proper convex normal integrands (see Rockafellar [1976]) with the following additional property:

Assumption 1: For any $y = (x_t, x_{t+1})$, there exists $\gamma > 0$ (independent of t) such that for $\pi \in \partial f_t(y) \subset (L_\infty^n)^*$, the Banach dual space of L_∞^n , and for all w , either

(a) $\pi \in \partial f_t(w)$,

or

(b) there exists z such that $\pi \in \partial f_t(z)$ and

$$f_t(z) + \pi(w - z) \leq f_t(w) - \gamma \|z - y\| \text{ a.s.} \tag{29}$$

for all $t \geq 0$. ■

Note that uniform convexity implies Assumption 1, which allows nonstrict convexity involving a von Neumann facet. We use this more general assumption here because it allows us to use the common scheduling objectives which involve linear tardiness and earliness penalties.

We present the main results from Birge and Dempster [1992] on optimality conditions, turnpike results concerning the optimality of cyclic policies and the asymptotic optimality of match-up strategies. The proofs of these results may be found there.

In general, the objective in (28) is infinite. We can avoid this difficulty by defining a policy $x^* := \{x_0, x_1^*, \dots\}$ as (*weakly optimal*), as in McKenzie [1976], if it is not *overtaken* by any other policy, i.e. if there does not exist x' such that

$$\limsup_{T \rightarrow \infty} \mathbb{E} \sum_{t=0}^T [f_t(x'_t, x'_{t+1}) - f_t(x_t^*, x_{t+1}^*)] \leq -\epsilon, \tag{30}$$

where $\epsilon > 0$.

We also assume that the objective functions satisfy a condition ensuring that no infinite terms are present in the sum in (30).

Assumption 2: For any t and $\delta < \infty$, there exists $\epsilon < \infty$ such that $\|x_t\| < \delta$ a.s. implies $\mathbb{E} f_t(x_t, x_{t+1}) > -\epsilon$ and $\|x_{t+1}\| < \epsilon$ a.s. for x_{t+1} feasible. ■

Given Assumption 2, we can subtract a constant from each f_t without changing the weak optimality of x^* . By setting this constant equal to the expected objective value in each period, we obtain an infimum of 0 in (28). Thus without loss of generality we assume a finite infimum in (28) in the sequel.

The first result from Birge and Dempster [1992] is that there exist prices supporting the objective terms in (28). These price supports provide the optimality conditions in the following theorem that allow decomposition of the conditions by time period.

Theorem 2. *Suppose Assumption 2 holds and that x^* is optimal in (28) with finite infimum, and*

(a) *(nonanticipative feasibility) For any $x \in \text{dom } F$ (i.e., such that $\mathbb{E} \sum_{t=0}^{\infty} f_t(x_t, x_{t+1}) < \infty$), the projection of x onto \mathcal{N} , Πx , is such that $\mathbb{E} \sum_{t=0}^{\infty} f_t(\Pi_t x_t, \Pi_{t+1} x_{t+1}) < \infty$,*

(b) *(strict feasibility) For some $x \in \mathcal{N}$, such that $\mathbb{E} \sum_{t=0}^{\infty} f_t(x_t, x_{t+1}) < \infty$, there exists $\delta > 0$ such that for all $\|y - x\| < \delta$, $y \in L_{\infty}^n$, $\mathbb{E} \sum_{t=0}^{\infty} f_t(y_t, y_{t+1}) < \infty$.*

(c) *(finite horizon continuation approximation) There exists x' such that for all T_k in some sequence $\{T_1, T_2, \dots\}$, and, for any $x \in \text{dom } F$, $(x^{T_k}, x'_{T_k+1}, x'_{T_k+2}, \dots)$ is also feasible, and the transition cost to x' is such that $|\mathbb{E}[f_{T_k-1}(x_{T_k-1}, x_{T_k}) + f_{T_k}(x_{T_k}, x'_{T_k+1})]| \rightarrow 0$ as $k \rightarrow \infty$ and $|\mathbb{E}[f_{T_k-1}(x_{T_k-1}, x_{T_k}) + f_{T_k}(x_{T_k}, x'_{T_k+1})]| \geq |\mathbb{E}[f_{T_k-1}(x_{T_k-1}, x_{T_k}) + f_{T_k}(x_{T_k}, x_{T_k+1})]|$ for $k = 1, \dots$*

Then, x^ is optimal with given initial conditions x_0 if and only if there exist $p_t \in L_1^r(\Sigma)$, $t = 0, \dots$, such that*

(i) p_t *is nonanticipative. i.e. $p_t = \mathbb{E}\{p_t | \Sigma_t\}$ a.s. for $t = 0, \dots$,*

(ii) $\mathbb{E}_0(f_0(x_0, x_1) - p_0 x_0 + p_1 x_1)$ *is a.s. minimized by $x_1 := x_1^*$ over $x_1 = \mathbb{E}\{x_1 | \Sigma_1\}$, and, for $t > 0$, $\mathbb{E}_t(f_t(x_t, x_{t+1}) - p_t x_t + p_{t+1} x_{t+1})$ is a.s. minimized by $(x_t, x_{t+1}) := (x_t^*, x_{t+1}^*)$ over $x_t = \mathbb{E}\{x_t | \Sigma_t\}$ and $x_{t+1} = \mathbb{E}\{x_{t+1} | \Sigma_{t+1}\}$, and*

(iii) $\mathbb{E} p_{t_k}(x_{t_k} - x_{t_k}^*) \rightarrow 0$ *as $t_k \rightarrow \infty$. for all $x \in \text{dom } F$. ■*

The optimality conditions (c) characterize optimal solutions which approach a common facet – the *von Neumann facet* – from any given starting condition x_0 . The main implication of this result is that it is asymptotically optimal to match up with a decision process that is optimal for a *specific* initial condition even if that initial condition changes. This result may be elaborated by showing that if the data process is cyclic then it is asymptotically optimal to return to an optimal cyclic policy even if other conditions temporarily obtain. These results justify the match-up scheduling policy in Bean *et al* [1991] and extend the deterministic

results in Bean and Birge [1986]. The following results for the general stochastic optimization model (28) are also from Birge and Dempster [1992] and apply in a variety of contexts.

Proposition 1. *Given Assumptions 1 and 2 and Conditions (a) – (c) in Theorem 2, let X^* be the set of solutions (x_t^*, x_{t+1}^*) that are minimal in (ii) of Theorem 1 for p_t^* a set of optimal supporting prices given the initial condition x_0 and let x' be an optimal decision process given the initial condition x'_0 . Then, for any $\epsilon > 0$ and $\delta > 0$, there exists $T < \infty$, such that, for all $t \geq T$,*

$$P\{\omega : \inf_{(x_t, x_{t+1}) \in X_t} \|(x'_t, x'_{t+1}) - (x_t^*, x_{t+1}^*)\| > \epsilon\} < \delta. \tag{31}$$

■

Theorem 3. *Under the conditions of Proposition 1, we may conclude that as $t \rightarrow \infty$,*

$$\inf_{(x_t, x_{t+1}) \in X_t} \|(x'_t, x'_{t+1}) - (x_t^*, x_{t+1}^*)\| \rightarrow 0 \quad \text{a.s.} \tag{32}$$

■

For Theorem 3 to be fully applicable, we would like to have a method for determining an optimal policy for some initial state so that the policy of matching up to that strategy can be implemented. This determination is simpler if we can show that cyclic policies are optimal. In this case, only a single cycle needs to be analyzed to determine the turnpike optimal policy.

In this development, we follow a similar approach to Arkin and Evstigneev [1979]. We first assume that the data process has a left tail. i.e. that ω_0 can be interpreted as $\dots \omega'_{-1}, \omega'_0$. An alternative is to assume some type of Markovian property of the data process (see Arkin and Evstigneev). The data process is assumed to be *cyclic with cycle k* if the measure μ is invariant with respect to the *k -period forward shift operator T_k* where $T_k \omega = \omega'$ such that $\omega'_t := \omega_{t+k}$, i.e., $\omega_t = T_k \omega_t = \omega_{t+k}$ a.s. It follows that we may define $T_k \Sigma_t := \Sigma_t$ for $t = 0, \dots, k - 1$. We also assume that the objective is invariant with respect to T_k so that $f_{t+k}(T_k x_t, T_k x_{t+1}) = f_t(x_t, x_{t+1})$ a.s., where $T_k x_t(\omega) := x_t(T_k \omega)$. In this context, x is a *cyclic policy* if $x_{t+k} = T_k x_t$ a.s. and (28) becomes

$$\inf_{x \in \mathcal{N}} \mathbb{E} \left(\sum_{t=0}^{k-2} f_t(x_t, x_{t+1}) + f_{k-1}(x_{k-1}, T_k x_0) \right). \tag{33}$$

Corollary 1. *Given conditions (a) - (c) of Theorem 2 and a cyclic data process with cycle k . then there exists a weakly optimal policy with cycle k for any initial condition \mathbf{x}_0 .* ■

As an example of extending these general results to a scheduling framework, consider a model with several commodities $i = 1, \dots, n$ processed according to random processing times $\mathbf{p}(i)$, random release dates $\mathbf{r}(i)$ and random due dates $\mathbf{d}(i)$, with a penalty weight of w_i for every period after the due date in which processing is not completed. We wish to model a situation in which orders for each item arrive randomly (according to $\mathbf{r}(i)$), in varying amounts (according to $\mathbf{p}(i)$), and with random due dates ($\mathbf{d}(i)$). The random entities, $\mathbf{r}(i)$, $\mathbf{p}(i)$, $\mathbf{d}(i)$, correspond to sequences of times, $\{r(\omega, i, 1), r(\omega, i, 2), \dots\}$, $\{p(\omega, i, 1), p(\omega, i, 2), \dots\}$, $\{d(\omega, i, 1), d(\omega, i, 2), \dots\}$, for each order number $1, 2, \dots$. The data process is defined so that when the j th order for i arrives at $t = r(\omega, i, j)$, then the processing time $p(\omega, i, j)$ and due date $d(\omega, i, j)$ are also known. Thus Σ_t distinguishes $r(\omega, i, l), p(\omega, i, l), d(\omega, i, l)$ for $1 \leq l \leq j$, but not for $l > j$.

Decisions are the amount of processing performed on each item i in each period t . Since the state of each item is reduced by the processing requirement at each due date, the total processing in period t is $\mathbf{x}_{t+1}(i) - \mathbf{x}_t(i)$ if t is not a due date ($t \neq \mathbf{d}(i, j)$ for any $j = 1, 2, \dots$) or $\mathbf{x}_{t+1}(i) + \mathbf{p}(i, j) - \mathbf{x}_t(i)$ if t is a due date ($t = \mathbf{d}(i, j)$). The decisions are constrained so that no processing can occur if an item is not released ($t < \mathbf{r}(i, j)$) and processing in each period on each item is at most one. Other restrictions on feasible processes appear in an indicator function $\delta(\omega, \mathbf{x}_t, \mathbf{x}_{t+1})$ which considers all resource availabilities.

The only costs in this model are due to tardiness. A penalty w_i is charged in each period for every unit of item i backordered ($\mathbf{x}_t(i) < 0$). The total tardiness cost at time t given ω is then $\sum_{i=1}^n w_i(-\mathbf{x}_t(\omega, i))^+$. The objective is to minimize the expected total tardiness.

The single period objective contribution is thus:

$$\mathbf{f}_t(\mathbf{x}_t, \mathbf{x}_{t+1}) := \sum_{i=1}^n \mathbf{f}_t^i(\mathbf{x}_t(i), \mathbf{x}_{t+1}(i)) + \delta(\omega, \mathbf{x}_t, \mathbf{x}_{t+1}), \tag{34}$$

where

$$f_i^t(\mathbf{x}_t(i), \mathbf{x}_{t+1}(i)) := \begin{cases} -w_i \mathbf{x}_t(i) & \text{if } \sum_{j=1}^{\infty} 1_{\{t=d(i,j)\}} \mathbf{p}(i,j) \leq \mathbf{x}_{t+1}(i) - \mathbf{x}_t(i) \\ & \leq \sum_{j=1}^{\infty} 1_{\{t \geq r(i,j)\}} \mathbf{p}(i,j) - 1_{\{t=d(i,j)\}} \mathbf{p}(i,j). \\ & \mathbf{x}_t(i) < 0 \text{ a.s.,} \\ 0 & \text{if } \sum_{j=1}^{\infty} 1_{\{t=d(i,j)\}} \mathbf{p}(i,j) \leq \mathbf{x}_{t+1}(i) - \mathbf{x}_t(i) \\ & \leq \sum_{j=1}^{\infty} 1_{\{t \geq r(i,j)\}} \mathbf{p}(i,j) - 1_{\{t=d(i,j)\}} \mathbf{p}(i,j). \\ & \mathbf{x}_t(i) \geq 0 \text{ a.s.,} \\ \infty & \text{otherwise.} \end{cases}$$

This noncyclic model is a generalization of the model with *cyclic* data process (with cycle k) in Birge and Dempster [1992]. In that model, it is assumed that $r(i, j + 1) - r(i, j)$, $d(i, j) - r(i, j)$, and $\mathbf{p}(i, j)$ are all identically distributed and that $r(i, j) < d(i, j) < k$ a.s. for all j .

The data process is assumed to determine the availability of the resources (such as machines, labor and tools) for processing all commodities. We allow the δ indicator term to represent feasibility generally by assuming a value of "0" if \mathbf{x}_{t+1} is feasibly reached from \mathbf{x}_t and " ∞ " otherwise. For example, suppose that each process i requires a resource $m(i)$ where $m(i) \in \{1, \dots, M\}$, the set of resources, and each resource can process at most one unit during a time interval if available and cannot process anything if unavailable. In this case, ω_t can be interpreted to have several components such that the first M components form an M -vector of ones and zeroes corresponding to availability and unavailability of resources. We then have

$$\delta(\omega, \mathbf{x}_t, \mathbf{x}_{t+1}) := \begin{cases} 0 & \text{if } \sum_{i=1}^n 1_{\{j=m(i)\}} (\mathbf{x}_{t+1}(i) - \mathbf{x}_t(i) + \mathbf{p}(i) 1_{\{t=d(i)\}}) \\ & \leq \omega_t(j) \text{ for } j = 1, \dots, M, \\ \infty & \text{otherwise.} \end{cases} \tag{35}$$

Other constraints can also be represented in this way. Our only requirement is that $\delta(\omega, \dots)$ is convex.

As an example of an alternative convex function, this constraint on feasibility can be broadened further to a situation with setups by including an additional set of variables, $s_t(i, j)$, which indicate the extent to which resource group j is set up to process commodity i . If resource j is a large number of machines, then certain fractions of the group could be set up for i at different times. By indicating the total single period production capacity of resource j by K_j (which might also vary in time and be an additional decision variable), we would include a constraint that $\sum_{i=1}^n s_t(i, j) \leq K_j$. We can also include a convex constraint on the maximum setup change in a single period as Δ_j , so that $\sum_{i=1}^n |s_{t+1}(i, j) - s_t(i, j)| \leq \Delta_j$ (with a setup change cost also entering the objective if relevant). In this case, the indicator of feasibility becomes:

$$\delta(\omega, \mathbf{x}_t, \mathbf{x}_{t+1}) := \begin{cases} 0 & \text{if } (\mathbf{x}_{t+1}(i) - \mathbf{x}_t(i) + \mathbf{p}(i) 1_{\{t=d(i)\}}) \\ & \leq \omega_t(j) s_t(i, j) \text{ for } i = 1, \dots, n; j = 1, \dots, M, \\ & \sum_{i=1}^n s_{t+1}(i, j) \leq K_j, \\ & \sum_{i=1}^n |s_{t+1}(i, j) - s_t(i, j)| \leq \Delta_j, j = 1, \dots, M, \\ \infty & \text{otherwise.} \end{cases} \tag{36}$$

Note that this function preserves the convexity of the objective function. This model is a basic multiple-processor, minimum expected weighted tardiness problem. With the convexity assumption, it meets the criteria for optimality and asymptotic stability given in Theorems 2 and 3. In some cases, the stationary distribution for this model follows a deterministic path between disruptions and an optimal match point is achieved as quickly as possible (*cf.* Bean *et al* [1991]). To see this, let \mathbf{x}^* be an optimal turnpike schedule in \mathbf{X} , the set of optimal turnpike schedules. Assume that some state $\mathbf{x}'_0 \leq \mathbf{x}^*_0$ a.s. is the initial state instead of \mathbf{x}^*_0 . Let \mathbf{x}' be an optimal trajectory given \mathbf{x}'_0 . In Birge and Dempster [1992], it is shown that a trajectory $\bar{\mathbf{x}}$ that starts at \mathbf{x}'_0 and matches up with \mathbf{x}^* at the earliest feasible t can be constructed with the same objective value as \mathbf{x}' for the cyclic problem. Essentially the same proof *mutatis mutandis* yields this result for the current noncyclic problem.

Theorem 4. *Suppose \mathbf{x}^* is optimal from \mathbf{x}^*_0 above in the tardiness model, find $\inf_{\mathbf{x} \in \mathcal{N}, \mathbf{x}_0 = \mathbf{x}^*_0}$ a.s. $\mathbb{E} \sum_{t=0}^{\infty} f_t(\mathbf{x}_t, \mathbf{x}_{t+1})$ given $\mathbf{x}^*_0 \geq \mathbf{x}'_0$ a.s., with f_t defined in (34) and δ_t in (35), and that there exists a feasible solution $\bar{\mathbf{x}}$ such that $\bar{\mathbf{x}} = \mathbf{x}'_0$ a.s., and $\bar{\mathbf{x}}_t = \mathbf{x}^*_t$ a.s. for some $\tau < \infty$. Then there exists an optimal solution \mathbf{x}' given \mathbf{x}'_0 such that $\mathbf{x}'_t = \mathbf{x}^*_t, t \geq \tau$ a.s. .* ■

To illustrate the use of this result, we consider a one-machine, multiple commodity version of the model with single period objective (34). In this case, $M := 1$ and $m(i) := 1$ for $i = 1, \dots, n$. For simplicity, we also assume that $w_i := 1$ for $i = 1, \dots, n$. With these assumptions, the familiar *earliest due date scheduling policy* is optimal. To define this policy, suppose that $r(\omega, i, \bar{j}(\omega, t, i)) \leq t < r(\omega, \bar{j}(\omega, t, i) + 1)$ and that $d(\omega, i_1, \bar{j}(\omega, t, i_1)) \leq d(\omega, i_2, \bar{j}(\omega, t, i_2)) \leq \dots \leq d(\omega, i_n, \bar{j}(\omega, t, i_n))$ and define $x_{t+1}(\omega, i_l)$ recursively from $l = 1$ to $l = n$ by:

$$\begin{aligned} x_{t+1}(\omega, i_l) &= x_t(\omega, i_l) - 1_{\{t=d(\omega, i_l, \bar{j}(\omega, t, i_l))\}} p(\omega, i_l, \bar{j}(\omega, t, i_l)) \\ &\quad + 1_{\{\omega_t(1)=1\}} \left(\min \left\{ 1 - \sum_{s=1}^{l-1} (x_{t+1}(\omega, i_s) - x_t(\omega, i_s)) \right. \right. \\ &\quad \left. \left. + 1_{\{t=d(\omega, i_s, \bar{j}(\omega, t, i_s))\}} p(\omega, i_s, \bar{j}(\omega, t, i_s)) \right\}, p(\omega, i_l, \bar{j}(\omega, t, i_l)) \right) \xi. \end{aligned} \tag{37}$$

Equation (37) then forces production to occur up to the machine availability in due date order on any items that have not yet reached the order quantity $p(\omega, i_l, \bar{j}(\omega, t, i_l))$. This definition implies that order j for item i

is always completed before order $j + 1$ is released, i.e. that $x_{r(i,j)}(i) \geq 0$ a.s. This assumption can be relaxed by allowing due date order to apply also to previous orders $j < j(\omega, t, i)$ that are not yet complete.

Theorem 5. *An optimal solution to problem (28) with objective defined by (15), $M := 1$, $x_0 := 0$, and $m(i) := 1, w_i := 1$ for $i = 1, \dots, n$ is to process items according to earliest due date of released items first. i.e. according to (37).*

The result of Theorem 5 for due date order is not valid (even with equal weights) in cases where penalties are charged only when jobs are finished. In such cases, the optimal order follows due dates if all jobs can complete on time, but the optimal order switches to shortest processing time if all jobs are late. The result of Theorem 5 does apply, however, if processing times and due dates follow the same order (see Birge *et al* [1990], Lemma 2.3). It also applies if the weights are ordered in decreasing order from earliest due date to last due date.

Due date order is optimal here regardless of processing time because, according to (37), charges are incurred only on the incomplete portion of each job. This assumption is practical if an order is large and small batches within the large order can be shipped to the customer as they are finished. This ability to break up jobs is the critical factor in our convexity assumptions.

Processing available jobs in due date order according to Theorem 5 provides a long run optimal solution provided the initial system is empty or that we can assume some point in time at which we have nonnegative processing on all released jobs. We would like to show that this policy satisfies the conditions for match-up optimality in Proposition 1 and Theorem 3.

Assumptions 1 and 2 are valid since the costs are just piecewise linear with fixed increment in each period and the set of possible states is at most a unit L_1 -distance from the current state. For condition (a) in Theorem 2 (nonanticipative feasibility), note that the feasibility conditions only depend on information available when an order is accepted. The second condition (b), strict feasibility, is satisfied if we assume that the system has sufficient slack such that the objective can be obtained without completely using all available capacity in some period. We assume this is possible (although an optimal solution may use all capacity).

The third condition (c), finite horizon continuation, is that we can at some point reach a trajectory starting from, for example, the empty inventory state.

With these assumptions and following Theorem 4, the optimal policy for any initial inventory state is to match up with the state from the empty inventory position as quickly as possible. The alternative initial states in Theorem 4 would correspond to entering period 0 with some overdue orders causing initial negative inventories for these items. The optimal match up response then corresponds to processing any items with negative inventories before proceeding to items with zero or positive inventories. The result is that one reaches by time t the same state as in the zero initial inventory state whenever the cumulative excess capacity up to time t is greater than the total negative inventory at time 0.

4. The Operational Level and Lagrangian Relaxation.

At the operational level, distinct setup times and the consideration of discrete variables is often unavoidable. In these situations, however, the stochastic nature of the problems can frequently provide advantageous solution characteristics that are not possible in deterministic problems. For example, consider a single machine scheduling problem with the objective to minimize weighted completion time when the machine is unavailable at some point during processing. If that downtime is known with certainty and occurs sometime strictly between 0 and the time to process all jobs, then finding the optimal order of jobs becomes an NP-complete problem. However, if downtimes occur randomly with negative exponential interarrivals, then the optimal order is simply the same weighted shortest processing time order as in the reliable machine case (Birge et al. [1990]).

Other beneficial effects of random problem parameters include obtaining continuity and other useful objective function characteristics. For example, Van der Vlerk [1995] surveys quite general conditions for continuity in input variables of the expectation over random parameters of the optimal value of optimization problems with integer variables. The applications of these results include problems with fixed order costs.

These types of results are reminiscent of the classical Lyapunov theorem (see, e.g. Young[1969], Theorem 79.1) which implies that the expectation of a nonlinear multifunction of a random (vector) argument with an

absolutely continuous (joint) distribution is convex. Another characteristic of stochastic scheduling problems is that similarities of problem structure across different random outcomes can be exploited in solution procedures. This observation has led to a variety of Lagrangian-based solution methods (see, for example, Dempster [1988] and Rockafellar and Wets [1991]). A particular advantage for scheduling problems is that the integer-variable problems can be solved separately with warm starts provided by solutions to problems with similar data. We show below, in fact, that the duality gap in discrete time stochastic programs of this type actually decreases to zero as the number of scenarios increases.

The basic approach here follows Bertsekas's [1982] original results in bounding duality gaps. Our development here is a generalization of the results in Takriti, Birge and Long [1994b], which apply the Lagrangian relaxation approach to scheduling electric power units. In that paper, it is shown that the duality gap is indeed limited by the number of decision points and the total capacity and that the average gap declines to zero as the number of random scenarios increases. Computational results (Takriti, Birge and Long [1994a]) also show that the solutions are close to optimal with only a few scenarios and that in this case stochastic scheduling problems provide significant savings over deterministic methods.

As a general model, we postulate the model of (27) with a *finite* number of scenarios $\Omega = \{\omega^1, \dots, \omega^N\}$ with probabilities, p^1, \dots, p^N , respectively, and a *finite* time horizon t . For decisions, $x_t^i, i = 1, \dots, N$, we may have integer variables (in addition to continuous decisions) where we use a superscript i to indicate dependence on scenario ω^i . Our problem becomes:

$$\inf_{x \in \mathcal{N}, x \in X^i} \sum_{t=0}^T \sum_{i=1}^N p^i f_t^i(x_t^i, x_{t+1}^i). \tag{38}$$

where X^i includes any feasibility conditions such as integrality on certain components of x^i , f_t^i is a convex function, and the only constraints linking scenarios are represented through \mathcal{N} . As noted earlier, \mathcal{N} can be written as a set of linear constraints, such as

$$x_t^i \left(\sum_{k \in \Sigma_t(i)} p^k \right) - \sum_{k \in \Sigma_t(i)} p^k x_t^k = 0 \quad \forall i, t,$$

where $\Sigma_t(i)$ is the set of scenarios sharing the same history process as i at time t . If there are T periods and $X^i \subset \mathbb{R}^{nT}$, we would typically have $N(T - 1)n$ of these constraints (one for each scenario in every period

from 1 to $T - 1$). If we collectively write these constraints as

$$\sum_{i=1}^N H^i x^i = b, \quad b \in \mathbb{R}^q$$

then (38) becomes

$$\begin{aligned} \min \quad & P \\ \text{subject to} \quad & P = \sum_{i=1}^N p^i F^i(x^i) \\ & x^i \in X^i, \quad i = 1, \dots, N, \\ & \sum_{i=1}^N H^i x^i = b. \end{aligned} \tag{39}$$

For an electric power scheduling problem as in (39), the variables x^i have binary components to represent the set of units which are in service and continuous components to represent the amount of production. We then have constraints in X^i to ensure that a unit can only produce power when it is in service. Objective costs also appear in F^i for setups to place units in service and to remove them from service. The linking constraints ensure that all demand is met.

The Lagrangian relaxation approach is to solve a Lagrangian dual to (39) by relaxing the nonanticipativity constraints. This program has the form:

$$\sup_{\lambda} D \quad \text{subject to} \quad D < D(\lambda)$$

where

$$D(\lambda) := \begin{cases} \min & \sum_{i=1}^N [p^i F^i(x^i) + \lambda^T H^i x^i] - \lambda^T b \\ \text{subject to} & x^i \in X^i, \quad i = 1, \dots, N, \\ & \lambda \geq 0, \end{cases} \tag{40}$$

which then decomposes into completely independent subproblems for each i .

The following result from Bertsekas [1982] provides the basis for the duality gap result in the Lagrangian approach.

Theorem 6. *If problem (39) has a solution, for every i , the set $\{(x^i, F^i(x^i)) | x^i \in X^i\}$ is compact, and, for every $\hat{x} \in \text{conv} X^i$, there exists $x \in X^i$ such that $H^i x \leq H^i \hat{x}$, then*

$$\inf P - \sup D \leq (q + 1)\rho.$$

where $\rho := \max_{i=1, \dots, N} \sup(p^i F^i(x^i) | x^i \in X^i) - \inf(p^i F^i(x^i) | x^i \in X^i)$.

Proof: This proof follows immediately from Bertsekas [Proposition 5.6, 1982]. The only observation is that $p^i F^i$ corresponds to the objective function f_i in Bertsekas' work, and that we use a linear function H^i which preserves compactness in X^i . ■

The significance of this result is that the gap decreases to zero as the probabilities decrease (as p^i , for example set equal to $\frac{1}{N}$, $\rightarrow 0$) without a corresponding increase in q . This tendency occurs in solution procedures that either sample the set of possible outcomes (for statistical results, see, for example, King and Wets [1991]) or include bounding approximations with increasing accuracy (see, for example, Birge and Wets [1986]). In either case, the number of scenarios increases and the probability of any individual scenario decreases with decreasing confidence intervals or bounds on the expected value of a full solution of (39).

Our goal is thus to show that the number q of constraints necessary to maintain nonanticipativity need not increase at the same rate when we consider integer variable values separately. The burden of proof is in fact less than this since we really need consider only constraints involving integer variables. We show this in the following theorem. For this result, assume that x^i has a continuous part and a binary part, so that $x^i := (y^i, s^i)$ and $X^i := Y^i \times S^i$, where Y^i is a convex set and S^i is the intersection of a convex set with $\times_{j=1}^J \{0, 1\}$, where $J = mT$ if we have T periods with m integer variables in each period. In the power scheduling problem, Y^i gives production levels while S^i provides service availability. The scenarios correspond to alternative demand paths. We suppose that the N total scenarios are combined by the nonanticipativity constraints into a tree with K non-leaf nodes or branching points. We then have the following result.

Theorem 7. *Suppose the conditions of Theorem 6, that $X^i = Y^i \times S^i$ as defined above, and that the equivalent decision tree for problem (39) includes K non-leaf nodes, then*

$$\inf P - \sup D \leq (Km + 1)\rho, \tag{41}$$

where $\rho := \max_{i=1, \dots, N} \sup(p^i F^i(x^i) | x^i \in X^i) - \inf(p^i F^i(x^i) | x^i \in X^i)$.

Proof: The proof of this result follows the lines of the proof in Bertsekas [1982] except that we recognize the difference between the continuous and integer parts of the variables in order to reduce the first factor of

the bound from the general case of $q = N(T - 1)$, where T is the number of periods, to Km .

First observe that we can replace all nonanticipativity constraints in P for variables s by the following:

$$s_t^{i^*} \left(\sum_{k \in \Sigma_t(i)} p^k \right) - \sum_{k \in \Sigma_t(i)} p^k s_t^k = 0, \forall t, \tag{42}$$

where we have only chosen a single representative, $i^* \in \Sigma_t(i)$, for each set of scenarios sharing the same history process as i at time t . We therefore need only a single constraint (42) for every node of the tree. To see this note that the s_t^i components are all binary. If $s_t^{i^*}(j) = 0$, then (42) is only satisfied when $s_t^k(j) = 0$ for all $k \in \Sigma_t(i)$. If $s_t^{i^*}(j) = 1$, then, of course, the same condition holds. Thus, constraints (42) can substitute for all s_t^i nonanticipativity constraints. If there are K non-leaf nodes and $s_t^{i^*} \in \mathbb{R}^m$, then we have Km constraints of the form (42).

If we let the constraints of form (42) be denoted collectively by

$$\sum_i G^i s^i = g,$$

and let the nonanticipativity constraints on the y^i variables be of the form

$$\sum_i L^i y^i = l,$$

then the primal problem becomes

$$\begin{aligned} \inf \quad & \sum_{i=1}^N p^i (C^i(y^i) + D^i(s^i)) \\ \text{subject to} \quad & y^i \in Y^i \quad i = 1, \dots, N. \\ & s^i \in S^i \quad i = 1, \dots, N, \\ & \sum_{i=1}^N G^i s^i = g, \\ & \sum_{i=1}^N L^i y^i = l, \end{aligned} \tag{43}$$

where we have written $F^i(x^i)$ as $C^i(y^i) + D^i(s^i)$ and note that $g \in \mathbb{R}^{2Km}$ and $l \in \mathbb{R}^{\bar{q}}$, for example.

Now, following Bertsekas, consider $W^i := \{w^i : w^i = [L^i y^i, C^i(y^i)], y^i \in Y^i\}$ and $Z^i := \{z^i : z^i = [G^i s^i, D^i(s^i)], s^i \in S^i\}$, with $W := \oplus_{i=1}^N W^i$ and $Z := \oplus_{i=1}^N Z^i$. Then we have:

$$\inf P = \min \{u + v | \exists ((w, u), (z, v)) \in W \times Z \text{ such that } w = g, z = l\}. \tag{44}$$

From duality theory (see, for example, Magnanti *et al.* [1976]), we have that

$$\sup D = \min\{u + v \mid \exists((w, u), (z, v)) \in \text{conv}(W \times Z) \text{ such that } w = g, z = l\}, \tag{45}$$

where *conv* denotes the convex hull.

Next observe that $\text{conv}(W \times Z) = W \times \text{conv}(Z)$, since Y^i is convex. Now, following Bertsekas, we can use the Shapley-Folkman theorem to write every $z \in \text{conv}(Z)$ using a subset $I(z) \subset \{1, \dots, N\}$ with at most $Km + 1$ indices such that

$$z \in \oplus_{i \in I(z)} Z^i. \tag{46}$$

Now, suppose $((\bar{w}, \bar{u}), (\bar{z}, \bar{v})) \in W \times \text{conv}(Z)$ with $\bar{u} + \bar{v} = \sup D$ and $\bar{w} = g, \bar{z} = l$. Then we have $\bar{y}^i \in Y^i$ such that $\sum_{i=1}^N L^i \bar{y}^i = g$ and $\sum_{i=1}^N C^i(\bar{y}^i) = \bar{u}$.

From Shapley-Folkman, we also have some $\bar{I} \subset \{1, \dots, N\}$, $|\bar{I}| \leq Km + 1$, with $(\bar{l}^i, \bar{v}^i) \in \text{conv}(Z^i)$ and $\bar{s}^i \in S^i, i \notin \bar{I}$, such that

$$\sum_{i \notin \bar{I}} G^i(\bar{s}^i) + \sum_{i \in \bar{I}} l^i = \bar{z} = l, \tag{47}$$

and

$$\sum_{i=1}^N C^i(\bar{y}^i) + \sum_{i \notin \bar{I}} D^i \bar{s}^i + \sum_{i \in \bar{I}} \bar{v}^i = \sup D. \tag{48}$$

Now, using the approach in Bertsekas and our assumptions, we can obtain for every $i \in \bar{I}$, some \bar{s}^i such that $G^i \bar{s}^i = l^i$ and $f_i(\bar{s}^i) \leq \bar{v}^i + \rho^i + \epsilon$ for any $\epsilon > 0$. Thus, we have found a feasible solution (\bar{y}, \bar{s}) for (P) such that

$$\inf P \leq \sum_{i=1}^N C^i(\bar{y}^i) + \sum_{i=1}^N D^i \bar{s}^i \leq \sup D + \sum_{i \in \bar{I}} \rho^i. \tag{49}$$

which, since $|\bar{I}| \leq 2Km + 1$, yields the result. ■

This result is only valuable in the present context if we can restrict the growth of K to $o(N)$ so that the gap in Theorem 6 goes to zero. In fact, this is the general case. Suppose, for example, that we employ a sampling scheme where every scenario path has weight $p^i := \frac{1}{N}$. If each path has ν branches at each node in each time period (see Figure 3 for an example), we then have $N = \nu^T$ scenarios and $K = (\nu^T - 1)/(\nu - 1)$ branching nodes. The duality gap vanishes in the limit since $K/N \rightarrow 0$ as $\nu \rightarrow \infty$. We state the result as follows.

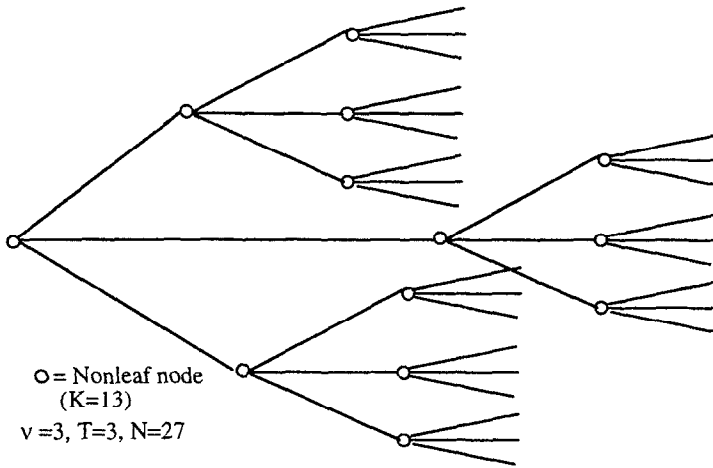


Figure 3. Scenario tree where with $T = 3$ periods (excluding $t = 0$), $\nu = 3$ branches at each non-leaf nodes. $N = 27$ scenarios, and $K = 13$ non-leaf nodes.

Corollary 2. If the conditions of Theorem 7 hold, $N = \nu^T$, $p^i := \frac{1}{N}$, $K = (\nu^T - 1)/(\nu - 1)$, and $\sup_i (F^i(x^i)|x^i \in X^i) - \inf(F^i(x^i)|x^i \in X^i) < M < \infty$ uniformly for all ν , then

$$\lim_{\nu \rightarrow \infty} (\inf P - \sup D) \rightarrow 0. \tag{50}$$

■

The vanishing duality gap result holds for other general sampling schemes. The result also indicates that the Lagrangian duality gap for a problem with any distribution obtained in the limit of the sampling procedure must also have zero duality gap. It implies in general that randomness – absolutely continuous random variables in this case – can indeed simplify some of the combinatorial difficulties in problems. Of course, one must still solve the separate deterministic subproblems optimally, but no additional gap is present in this case.

This result also indicates that parallel processing may offer distinct advantages for stochastic scheduling. Repeated deterministic solutions of deterministic Lagrangian subproblems on separate processors will yield increasingly accurate solutions as the number of scenarios increases.

Conclusion.

This paper has explored three methods for incorporating optimization procedures into stochastic scheduling problems as part of a decision hierarchy. At the capacity planning level, we showed how queueing approximations yield a convenient model when lower level timing effects are not – or are only occasionally – significant. At the aggregate production level, we demonstrated that a convex approximation can yield useful characterizations of optimal policies. At the lowest level of detailed scheduling, we showed that Lagrangian relaxation methods obtain smaller duality gaps as stochastic program formulation sizes grow.

These results can all be viewed as separate solutions of parts of the 3-level model in (7) with each higher level incorporating some approximation of the lower level objectives. Our results give some indication of structural properties and uses of optimization procedures. We anticipate that this general framework will yield many other characterizations and procedures for alternate problem structures.

References.

- V.I. Arkin and I.V. Evstigneev (1979). *Stochastic Models of Control and Economic Dynamics*. Nauka, Moscow. (In Russian). English translation by E.A. Medova and M.A.H. Dempster. Academic, Orlando (1986).
- D. Bai, T. Carpenter and J.M. Mulvey (1994). Stochastic programming to promote network survivability. Technical Report SOR-94-14, Statistics and Operations Research, Princeton University, August 1994.
- J.C. Bean and J.R. Birge (1986). Match-up real-time scheduling. In *Real-Time Optimization in Automated Manufacturing Facilities*, R.H.F. Jackson and A.W.T. Jones, eds., NBS Special Publication 724, National Bureau of Standards, 197–212.
- J.C. Bean, J.R. Birge, J. Mittenthal and C.E. Noon (1991). Match-up scheduling with multiple resources, release dates and disruptions *Operations Research* 39, 470–483.
- D.P. Bertsekas (1982). *Constrained Optimization and Lagrange Multiplier Methods*. Academic Press, New York.
- D. Bertsimas (1994). A mathematical programming approach to stochastic and dynamic optimization prob-

- lems. Technical Report, Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, March 1994.
- D. Bertsimas and J. Niño-Mora (1993). Conservation laws, extended polymatroids and multi-armed bandit problems: A unified approach to indexable systems. Technical Report, Sloan School of Management, Massachusetts Institute of Technology, March 1993.
- J.R. Birge (1982). The value of the stochastic solution in stochastic linear programs with fixed recourse. *Mathematical Programming* 24, 314-325.
- J.R. Birge and M.A.H. Dempster (1992). Optimality for match-up strategies in stochastic scheduling. Technical Report, Department of Industrial and Operations Engineering, The University of Michigan. Submitted to: *Mathematics of Operations Research*.
- J.R. Birge and M.A.H. Dempster (1995). Optimal match-up strategies in stochastic scheduling. *Discrete Applied Mathematics* 57, 105-120.
- J.R. Birge, J.B.G. Frenk, J. Mittenthal and A.H.G. Rinnooy Kan (1990). Single machine scheduling subject to stochastic breakdowns. *Naval Research Logistics* 37, 661-677.
- J.R. Birge and K.D. Glazebrook (1988). Assessing the effects of machine breakdowns in stochastic scheduling. *Operations Research Letters* 7, 267-271.
- J.R. Birge and R. J-B. Wets (1986). Designing approximation schemes for stochastic optimization problems, in particular, for stochastic programs with recourse. *Mathematical Programming Study* 27, 54-102.
- J.A. Buzacott and J.G. Shanthikumar (1980). Models for understanding flexible manufacturing systems. *AIIE Transactions* 12, 339-350.
- H. Chen and A. Mandelbaum (1991). Stochastic discrete flow networks: Diffusion approximations and bottlenecks. *Annals of Probability* 19, 1463-1510.
- G.B. Dantzig (1963). *Linear Programming and Extensions*. Princeton University Press, Princeton, NJ.
- M.A.H. Dempster (1988). On stochastic programming: II. Dynamic problems under risk. *Stochastics* 25, 15-42.
- M.A.H. Dempster (1994). Hierarchical approximation of telecommunications networks. *BT Technology J.*

12, 40-49.

- M.A.H. Dempster, M.L. Fisher, L. Jansen, B.J. Lageweg, J.K. Lenstra and A.H.G. Rinnooy Kan. (1983). Analysis of heuristics for stochastic programming: Results for hierarchical scheduling problems. *Mathematics of Operations Research* 8, 525-537.
- M.A.H. Dempster, P.B. Key and E.A. Medova (1995). Integrated system modelling for ATM. Research Report 95-7, Department of Mathematics, University of Essex. Submitted to: *Telecommunications Research*.
- M.A.H. Dempster, J.K. Lenstra and A.H.G. Rinnooy Kan, eds. (1982). *Deterministic and Stochastic Scheduling*. Reidel, Dordrecht.
- M.A.H. Dempster and E. Solel (1987). Stochastic scheduling via stochastic control. *Proceedings 1st World Conference of the Bernoulli Society*. K. Krickeberg and A. Shiryaev, eds. VNU Science Press, Utrecht, 783-788.
- S.D. Flåm (1983). Asymptotically stable solutions to stochastic problems of Bolza. Chr. Michelsen Institute Report, Fantoft, Norway.
- S.D. Flåm (1985). Nonanticipativity in stochastic programming. *Journal of Optimization Theory and Applications* 46, 23- 30.
- R. Garbe and K.D. Glazebrook (1995). Conservation laws, submodular returns and greedy heuristics for job selection and scheduling problems. Research Report, Department of Mathematics and Statistics, University of Newcastle. April 1995.
- J.C. Gittins (1981). Multiserver scheduling of jobs with increasing completion times. *Journals of Applied Probability* 16, 321-324.
- K.D. Glazebrook (1981). On nonpreemptive strategies in stochastic scheduling. *Naval Research Logistics Quarterly* 28, 289-300.
- K.D. Glazebrook (1984). Scheduling stochastic jobs on a single machine subject to breakdowns. *Naval Research Logistics Quarterly* 31, 251-264.
- K.D. Glazebrook (1987). Evaluating the effects of machine breakdowns in stochastic scheduling problems.

Naval Research Logistics **34**, 319-335.

K.D. Glazebrook (1991). On non-preemptive policies for stochastic single machine scheduling with breakdowns. *Probability in Engineering and Informational Sciences* **5**, 77-87.

S.C. Graves (1981). A review of production scheduling. *Operations Research* **29**, 646-675.

J.M. Harrison (1985). *Brownian Motion and Stochastic Flow Systems*. John Wiley, New York.

J-B. Hiriart-Urruty (1982). Extension of Lipschitz integrands and minimization of nonconvex integral functionals: Applications to the optimal recourse problem in discrete time. *Probability and Mathematical Statistics* **3**, 19-36.

J.Y. Hui and E. Karasan (1995). A thermodynamic theory for broadband networks. *Proceedings of the Third INFORMS Telecommunications Conference*, Boca Raton, Florida, March 1995. F.P. Kelly (1991). Loss networks. *Annals Applied Probability* **1**, 319-378.

A.J. King and R. J-B Wets (1991). Epiconsistency of convex stochastic programs. *Stochastics and Stochastics Reports* **34**, 83-92.

H.J. Kushner (1972). Necessary conditions for discrete parameter stochastic optimization problems. In: *Proceedings Sixth Berkeley Symposium Math. Stats. and Probability*, J. Neyman, ed. U. California, Berkeley, 667-685.

L. McKenzie (1976). Turnpike theory. *Econometrica* **44**, 841-864.

T.L. Magnanti, J.F. Shapiro, and M. H. Wagner (1976). Generalized linear programming solves the dual. *Management Science* **22**, 1195-1203.

E.A. Medova (1993). ATM admission control and routing. BT Laboratories Performance Engineering Report DS DZ/783-03/DP/00029, December 1993.

J. Mittenthal (1986). Single Machine Scheduling Subject to Random Breakdowns. Ph.D. dissertation, The University of Michigan, Ann Arbor, Michigan.

M.L. Pinedo (1983). Stochastic scheduling with release dates and due dates. *Operations Research* **31**, 559-572.

- M.L. Pinedo and S.M. Ross (1980). Scheduling jobs subject to nonhomogeneous Poisson shocks. *Management Science* **26**, 1250–1257.
- R. Righter (1994). Scheduling. Chapter 13 in M. Shaked and J.G. Shanthikumar, *Stochastic Orders and their Applications*. Academic Press, San Diego, CA.
- A.H.G. Rinnooy Kan (1976). *Machine Scheduling Problems: Classification, Complexity and Computations*. Martinus Nijhoff, The Hague.
- R.T. Rockafellar (1976). *Integral Functionals, Normal Integrands and Measurable Selections*. Lecture Notes in Mathematics **543**. Springer-Verlag, Berlin.
- R.T. Rockafellar and R. J-B. Wets (1983). Deterministic and stochastic optimization problems of Bolza type in discrete time. *Stochastics* **10**, 273–312.
- S. Sen, R.D. Doverspike and S. Cosares (1992). Network planning with random demand. Technical Report, Systems and Industrial Engineering Department, University of Arizona, December 1992.
- S.P. Sethi, H.M. Soner, Q. Zhang and J. Jiang (1992). Turnpike sets and their analysis in stochastic production planning problems. *Mathematics of Operations Research* **17**, 932–950.
- S.P. Sethi and Q. Zhang (1995). Multilevel hierarchical decision making in stochastic marketing-production systems. *SIAM J. Control and Optimization* **33**, 529–553.
- M. Shaked and J.G. Shanthikumar (1994). *Stochastic Orders and their Applications*. Academic Press. San Diego, CA.
- J.G. Shanthikumar and D. Yao (1992). Multiclass queueing systems: Polymatroid structure and optimal scheduling and control, *Operations Research* **40**, S293–S299.
- E. Solel (1986). A Dynamic Approach to Stochastic Scheduling via Stochastic Control. Ph.D. dissertation, Dalhousie University, Halifax, Nova Scotia.
- S. Takriti, J.R. Birge and E. Long (1994 a). A stochastic model for the unit commitment problem. *IEEE Transactions on Power Systems*, to appear.
- S. Takriti, J.R. Birge and E. Long (1994 b). Lagrangian solution techniques and bounds for loosely-coupled mixed-interger stochastic programs. Technical Report, Department of Industrial and Operations Engi-

neering, University of Michigan, December 1994.

M.H. van der Vlerk (1995). *Stochastic Programming with Integer Recourse*. Ph.D. Thesis, University of Groningen.

L.C. Young (1969). *Lectures on the Calculus of Variations and Optimal Control Theory*. W.B. Saunders, Philadelphia.