13. Martinet B (1970) Regularization d'inequations variationelles par approximations successives. Rev Francaise Inf Rech Oper 4:154–159
14. Nemirovskii AS, Yudin DB (1978) Cesaro convergence of the gradient method for approximating saddle points of convex–concave functions. Dokl Akad Nauk SSSR 239(5):1056–1059
15. Nurminskii E (1979) Numerical methods for solving deterministic and stochastic minimax problems. Naukova Dumka, Kiev
16. Rockafellar T, Uryasev S (2000) Optimization of conditional-value-at-risk. J Risk 2:21–41
17. Uryasev S (1990) Adaptive algorithms of stochastic optimization and game theory. Nauka, Moscow

Stochastic Scheduling

JOSÉ NIÑO-MORA
Department of Statistics,
Universidad Carlos III de Madrid, Getafe, Spain

MSC2000: 90B36

Article Outline

Introduction

The field of stochastic scheduling is motivated by problems of priority assignment arising in a variety of systems where jobs with random features (e. g., arrival or processing times) vie over time for access to shared service resources. Two prominent application areas are the dynamic scheduling of flexible manufacturing and computer-communication systems. Think, e. g., of a manufacturing workstation whose capacity is shared by multiple part types. Or consider a packet-switched communication network's channel whose bandwidth is shared by multiple traffic classes. Another rich set of applications is furnished by problems concerning the dynamic scheduling of multiple projects, whose states evolve randomly over time (e. g., research and development projects, or clinical trials).

The performance of such systems, as evaluated by a criterion such as the average time that jobs stay in the system (*flowtime*), can be significantly affected by the *scheduling policy* adopted to prioritize over time the access of jobs to resources. This motivates the interest of finding scheduling policies that optimize performance objectives of concern (e. g., minimizing the average flowtime). Yet, the high degree of discretionality allowed in the design of such policies gives rise to a combinatorial explosion rendering intractable an exhaustive search to determine an optimal policy. Instead, a goal of practical interest is to design relatively simple scheduling policies that achieve an optimal or nearly optimal performance.

The theory of stochastic scheduling addresses such a goal in the idealized setting of stochastic system models. Real-world random features such as job interarrival or processing times are thus modeled by specifying their probability distributions. Model assumptions vary across several dimensions, including the class of scheduling policies considered to be admissible, job interarrival and processing-time distributions, type and arrangement of service resources, and performance objective to be optimized. Typically, admissible policies are required to be nonanticipative, meaning that they cannot make use of future information, such as the unknown total duration of a job whose processing has not yet finished.

Regarding solution methods and techniques, it seems fair to say that no unified and practical approach is yet available to design and analyze optimal or near-optimal policies across the entire range of stochastic scheduling models. Although many such models can be cast in the framework of dynamic programming, straightforward application of this technique typically results in intractable formulations (*curse of dimensionality*). Classical results in the area were obtained through insightful ad hoc ideas, often based on interchange arguments (cf. [41]), whose extension to seemingly close model variations is hard or elusive. Yet the last two decades have witnessed major advances in promising research fronts, such as the use of Brownian or of fluid approximations, the use of mathematical programming formulations, and the development of priority-index policies.

Stochastic scheduling problems can be classified into three broad types, which have evolved with substantial autonomy: problems concerning the scheduling of a batch of stochastic jobs, multi-armed bandit problems, and problems concerning the scheduling of queueing systems.

The historical development of each such area has followed a similar three-stage pattern. In the first, earlier stage, researchers elucidated the optimal policies in relatively simple models. Such policies were often found to be based on priority indices: an index is computed for each job type or project, as a function of its state; then, at each decision epoch jobs or projects with larger index values are awarded higher priority for access to service. In the second stage, research efforts focused on identifying optimal policies in more complex models, often at the expense of introducing rather restrictive conditions on model parameters, such as symmetry assumptions. In the third, current stage, the main focus has shifted to develop computationally tractable methods capable of addressing large-scale models, which yield guidelines for designing good dynamic scheduling policies.

Models

Scheduling a Batch of Stochastic Jobs

In models of this class, a set of m machines is available to process a batch of n jobs with random processing times having known distributions, in order to optimize a given performance objective. The simplest such problem is to sequence a set of n stochastic jobs on a single machine ($m = 1$) to minimize the expected weighted flowtime. Job processing times are independent random variables, having a general distribution $G_i(\cdot)$ with mean p_i for job i. Admissible scheduling policies are required to be nonanticipative and *nonpreemptive* (processing of a job, once started, must proceed uninterruptedly to completion). Let $w_i \geq 0$ denote the cost rate incurred per unit time in the system (waiting or being processed) for job i, and let $\tilde{C}_i$ denote its random *completion time*. Let Π denote the class of all admissible policies, and let $E^\pi[\cdot]$ denote expectation under policy $\pi \in \Pi$. The problem can be formulated as

$$\min_{\pi \in \Pi} E^\pi \left[\sum_{j=1}^{n} w_j \tilde{C}_j \right] . \tag{1}$$

In the special case where job durations are deterministic, Smith first showed in [60] that it is optimal to sequence jobs in nonincreasing order of the priority index w_i/p_i. Such a rule is also optimal in the general stochastic case (1), as shown in [57]. References [36,37] identify conditions under which such an index rule is optimal when there are multiple identical parallel machines and processing times are exponentially distributed.

The model extension where policies are allowed to be *preemptive* (processing of a job may be interrupted at any time, to be later resumed) was solved by Sevcik in [58]. The optimal policy is again characterized by a priority index for each job, which in this case is a function of the cumulative processing time received so far.

Optimal index policies have also been identified for scheduling a batch of jobs on identical parallel machines, yet only under rather stringent conditions. The main performance objectives investigated in such a setting are: (i) minimize the total expected flowtime,

$$\min_{\pi \in \Pi} E^\pi \left[\sum_{j=1}^{n} \tilde{C}_j \right] ; \tag{2}$$

and (ii) minimize the expected *makespan* (time to finish the last job),

$$\min_{\pi \in \Pi} E^\pi \left[\max_{1 \leq j \leq n} \tilde{C}_j \right] . \tag{3}$$

The index rule that assigns higher priority to jobs with shorter expected processing time (SEPT) has been shown to be optimal for (2) under the following assumptions: when job processing time distributions are exponential [15,29,72]; when jobs have the same general processing time distribution (having possibly received different amounts of processing prior to start) with a nondecreasing hazard rate function [65]; and, more generally, when job processing time distributions are stochastically ordered [67].

As for the expected makespan objective (3), the index rule that assigns higher priority to jobs with longer expected processing times (LEPT) has been shown to be optimal in the following cases: under exponential processing time distributions [15,72]; and when jobs have a common processing time distribution (with possibly different amounts of processing prior to start) with a nonincreasing hazard rate function [65].

Other models incorporate more complex features. Thus, the optimality of the preemptive version of Smith's index rule is extended in [53] to models with stochastic release dates or due dates. Also, in models with *uniform* parallel machines, which differ in speed rates, researchers have characterized optimal policies exhibiting a threshold structure: see [1,55] for the problem of expected flowtime minimization, and [18] for the problem of expected makespan minimization. An optimal policy for the problem of scheduling a batch of stochastic jobs in a *flow shop* (with m machines in series) is identified in [75].

The optimality of the simple policies identified in the work reviewed above typically does not extend to models that violate the required assumptions [19]. Finding an optimal policy in such cases appears to be a computationally intractable goal (see [50] for a study on the complexity of decision-making problems under uncertainty, such as stochastic scheduling). This fact has motivated the analysis of suboptimal heuristic index policies.

For example, it has been shown in [71] that, under mild assumptions, the suboptimality gap for Smith's rule, when used as a heuristic for stochastic scheduling on parallel machines, is bounded above by a quantity that is independent of the number of jobs. Therefore, as the latter grows to infinity the rule's relative suboptimality gap vanishes, yielding a form of asymptotic optimality. An earlier asymptotic optimality result in the same vein for a model of parallel machines stochastic scheduling with in-tree precedence constraints was obtained in [51].

A recent line of work uses optimal solutions to linear programming relaxations to design and analyze scheduling rules with performance guarantees for hard stochastic scheduling problems [40].

Multi-Armed Bandits

Models in this class are concerned with optimally allocating effort over time to a collection of projects, which change state in a random fashion depending on whether they are engaged or not. A classic example is the multi-armed bandit problem which, in its discrete-time version, can be described as follows: there is a collection of K projects labeled by $k = 1, \ldots, K$, exactly one of which must be engaged at each discrete time period $t = 0, 1, \ldots$. Project k can be in a finite number of states $i_k \in N_k$, where N_k is the project's state space. If at period t project k occupies state i_k and is engaged, an active reward $R_k^1(i_k)$ is earned, geometrically discounted by factor $0 < \beta < 1$; then, the project state changes in a Markovian fashion to j_k with active transition probability $p_k^1(i_k, j_k)$ for $j_k \in N_k$. Projects not engaged do not earn rewards (i. e., passive rewards are $R_k^0(i_k) \equiv 0$) and remain frozen. The problem is to find a nonanticipative scheduling policy for selecting the project to be engaged at each period, so as to maximize the total expected discounted reward earned over an infinite horizon. Denoting by Π the class of such admissible policies, and denoting by $X_k(t)$ and by $a_k(t)$ the state and the action ($a_k(t) = 1$: active; $a_k(t) = 0$: passive) for project k at period t, the problem can be formulated as

$$\max_{\pi \in \Pi} E^{\pi} \left[\sum_{t=0}^{\infty} \beta^t \sum_{k=1}^{K} R_k^{a_k(t)}\big(X_k(t)\big) \right] .$$

Such a classic problem, whose name refers to a slot machine with multiple arms, one of which must be pulled at each time, has its origins in problems of sequential design of experiments [56,62]. After being long considered intractable, the problem was solved in a celebrated result by Gittins and Jones [28]. The optimal policy is given by an index rule: an index $\gamma_k(i_k)$ is defined for each project k as a function of its state i_k; then, at each time a project with currently largest index is engaged, breaking ties arbitrarily. The Gittins index generalizes that introduced in [7] for Bayesian Bernoulli bandits, which in turn was based on the index introduced in [13] for finite-horizon Bayesian bandits.

The optimality of such an index rule, for the original model and extensions, has a rich history of proofs yielding complementary insights. Such proofs are based on different techniques, including interchange arguments [26,28,64,70], dynamic programming [73], intuitive arguments [66], induction arguments [63], and conservation laws/linear programming [8]. See [27] for a comprehensive reference. For efficient methods to compute the Gittins index, see [17,35,46].

The important model extension where projects not engaged continue to evolve, typically with different transition probabilities, was introduced by Whittle in [74]. Its greatly improved modeling power comes, however, at the cost of tractability [52]. In the setting of a time-average version of such a *multi-armed restless bandit problem*, he deployed a Lagrangian approach to obtain a heuristic index rule that reduces to Gittins' in the classic model. His conjecture regarding the asymptotic optimality of such an index rule as both the number of projects and the required number of projects to be engaged grow to infinity in a constant ratio was established, under appropriate conditions, in [68]. Yet his proposed index for restless bandits only exists for a restricted class of bandits, termed indexable, which raises the issue of finding sufficient conditions for indexability.

The results in [74] were based on introduction of a tractable problem relaxation, which also yields useful bounds on the optimal value. Improved bounds based on a hierarchy of linear programming relaxations were introduced in [11].

A framework for the analysis and computation of restless bandit indices, leading to the unifying concept of marginal productivity index, has been recently developed and deployed in several applications in [42,43,44,45]. See [47] for an accessible review of such a line of research.

The incorporation of penalties (costs or delays) for switching projects also yields an important yet intractable model extension of classic bandits [34], as it is no longer solved by index policies [5]. Yet, [3] introduced an intuitive index that partially characterizes optimal policies. An efficient algorithm to compute such an index, based on the natural formulation of classic bandits with switching costs as restless bandits without them, along with extensive computational experience showing that the resulting index policy is nearly optimal, is reported in [49].

Scheduling Queueing Systems

Models in this class concern the design of optimal policies for dynamic allocation of jobs to servers, where jobs arrive over time according to given stochastic processes The main class of models in this setting is that of *multiclass queueing networks* (MQNs), widely applied as versatile models of computer-communication and manufacturing systems.

The simplest types of MQNs involve scheduling a number of job classes in a single server. Similarly as in the two problem categories discussed above, simple priority-index rules have been shown to be optimal for a variety of such models. Consider the case of a multiclass $M/G/1$ queue, where K job classes vie for the attention of a single server: Jobs of class $k = 1, \ldots, K$ arrive at the system as a Poisson process with rate λ_k, and their service times are drawn independently from a common distribution $G_k(\cdot)$ with mean $1/\mu_k$. Class j jobs incur linear holding costs at rate $c_k \geq 0$ per unit time that a job resides in the system (waiting or in service). The goal is to find a nonanticipative and nonpreemptive scheduling policy prescribing which job class to serve at each decision epoch, in order to minimize the long-run average holding cost rate per unit time. Let Π denote the class of all such admissible policies, and let $E^{\pi}[L_k]$ denote the expected number of class k jobs in the system under policy $\pi \in \Pi$. The problem can be stated as

$$\min_{\pi \in \Pi} \sum_{k=1}^{K} c_k E^{\pi}[L_k] .$$

Its solution is given by the classic $c\mu$-rule [21], which is the same as the Smith index rule discussed above: award highest service priority at each time to a job with largest index $c_k \mu_k$. The $c\mu$-rule is also optimal among preemptive policies when service times are exponential.

The optimality of an index policy for the model extension that incorporates Markovian job feedback (when a class k job completes service it changes class to l with probability p_{kl}, and leaves the system with probability $1 - \sum_{l=1}^{K} p_{kl}$) was established by Klimov in [38]. The optimal priority index is efficiently computed by the K-step Klimov algorithm. The result was extended to the discounted criterion in [61].

An account of these results based on the *achievable region method*, which seeks to characterize the region of achievable performance vectors (e. g., mean queue lengths for each class) by means of linear programming constraints that formulate *conservation laws*, has been given in [8,20,25,59] (in increasing levels of generality). The performance of Klimov's index rule, when used as a heuristic in the model extension that includes identical parallel servers, has been analyzed using such an approach in [31]: a *relaxed* linear programming formulation of the performance region is shown to yield closed-form suboptimality bounds, which imply the rule's asymptotic optimality in heavy-traffic.

More general MQN models involve features such as changeover times for changing service from one job class to another [39], or multiple processing stations, which provide service to corresponding nonoverlapping subsets of job classes. Due to the intractability of such models, researchers have aimed to design relatively simple heuristic policies which achieve a performance close to optimal. The accomplishment of such a goal has been hindered by formidable technical challenges, including the *stability problem* for multiclass queueing networks with multiple stations [14,24]: in general it is not known what conditions on model parameters ensure that a given policy is *stable* (the time-average number of jobs in the system is finite). As a result, computer simulation remains the most widely used tool in applications of these models. Theoretical approaches currently under active development include the development of heuristic scheduling policies based on: diffusion approximations of the original system under heavy-traffic conditions [6,32,33,54,69]; fluid approximations [4,16,23]; mathematical programming formulations [9,10,12,22,30,31]; and restless bandit indexation [2,43,44,45,47,48].

References

1. Agrawala AK, Coffman Jr EG, Garey MR, Tripathi SK (1984) A stochastic optimization algorithm minimizing expected flow times on uniform processors. IEEE Trans Comput c-33:351–356
2. Ansell PS, Glazebrook KD, Niño-Mora J, O'Keeffe M (2003) Whittle's index policy for a multi-class queueing system with convex holding costs. Math Methods Oper Res 57:21–39
3. Asawa M, Teneketzis D (1996) Multi-armed bandits with switching penalties. IEEE Trans Autom Control 41:328–348
4. Atkins D, Chen H (1995) Performance evaluation and scheduling control of queueing networks: fluid model heuristics. Queueing Sys Theory Appl 21:391–413
5. Banks JS, Sundaram RK (1994) Switching costs and the Gittins index. Econometrica 62:687–694
6. Bell SL, and Williams RJ (2001) Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: asymptotic optimality of a threshold policy. Ann Appl Probab 11:608–649
7. Bellman R (1956) A problem in the sequential design of experiments. Sankhyā 16:221–229
8. Bertsimas D, Niño-Mora J (1996) Conservation laws, extended polymatroids and multiarmed bandit problems; a polyhedral approach to indexable systems. Math Oper Res 21:257–306
9. Bertsimas D, Niño-Mora J (1999) Optimization of multiclass queueing networks with changeover times via the achievable region approach: Part I, the single-station case. Math Oper Res 24:306–330
10. Bertsimas D, Niño-Mora J (1999) Optimization of multiclass queueing networks with changeover times via the achievable region approach: Part II, the multi-station case. Math Oper Res 24:331–361
11. Bertsimas D, Niño-Mora J (2000) Restless bandits, linear programming relaxations and a primal-dual index heuristic. Oper Res 48:80–90
12. Bertsimas D, Paschalidis IC, Tsitsiklis JN (1994) Optimization of multiclass queueing networks: Polyhedral and nonlinear characterizations of achievable performance. Ann Appl Probab 4:43–75
13. Bradt RN, Johnson SM, Karlin S (1956) On sequential designs for maximizing the sum of *n* observations. Ann Math Statist 27:1060–1074
14. Bramson M (1994) Instability of FIFO queueing networks. Ann Appl Probab 4:414–431
15. Bruno J, Downey P, Frederickson GN (1981) Sequencing tasks with exponential service times to minimize the expected flow time or makespan. J ACM 28:100–113
16. Chen H, Yao DD (1993) Dynamic scheduling of a multiclass fluid network. Oper Res 41:1104–1115
17. Chen YR, Katehakis MN (1986) Linear programming for finite state multi-armed bandit problems. Math Oper Res 11:180–183
18. Coffman Jr EG, Flatto L, Garey MR, Weber RR (1987) Minimizing expected makespans on uniform processor systems. Adv Appl Probab 19:177–201
19. Coffman Jr EG, Hofri M, Weiss G (1989) Scheduling stochastic jobs with a two-point distribution on two parallel machines. Probab Eng Inform Sci 3:89–116
20. Coffman Jr EG, Mitrani I (1980) A characterization of waiting time performance realizable by single server queues. Oper Res 28:810–821
21. Cox DR, Smith WL (1961) Queues. Methuen, London
22. Dacre M, Glazebrook K, Niño-Mora J (1999) The achievable region approach to the optimal control of stochastic systems (with discussion). J Royal Stat Soc Ser B Methodol 61:747–791
23. Dai JG, Lin W (2005) Maximum pressure policies in stochastic processing networks. Oper Res 53:197–218
24. Dai JG, Weiss G (1996) Stability and instability of fluid models for reentrant lines. Math Oper Res 21:115–134
25. Federgruen A, Groenevelt H (1988) Characterization and optimization of achievable performance in general queueing systems. Oper Res 36:733–741
26. Gittins JC (1979) Bandit processes and dynamic allocation indices (with discussion). J Royal Stat Soc Ser B Methodol 41:148–177
27. Gittins JC (1989) Multi-Armed Bandit Allocation Indices. Wiley, Chichester

28. Gittins JC, Jones DM (1974) A dynamic allocation index for the sequential design of experiments. In: Gani J, Sarkadi K, Vincze I (eds) Progress in Statistics (European Meeting of Statisticians, Budapest, 1972). North-Holland, Amsterdam, pp 241–266
29. Glazebrook KD (1979) Scheduling tasks with exponential service times on parallel processors. J Appl Probab 16: 685–689
30. Glazebrook KD, Niño-Mora J (1999) A linear programming approach to stability, optimisation and performance analysis for Markovian multiclass queueing networks. Ann Oper Res 92:1–18
31. Glazebrook KD, Niño-Mora J (2001) Parallel scheduling of multiclass *M*/*M*/*m* queues: approximate and heavy-traffic optimization of achievable performance. Oper Res 49:609–623
32. Harrison JM (1988) Brownian models of queueing networks with heterogeneous customer populations. In: Fleming W, Lions PL (eds) Stochastic Differential Systems, Stochastic Control Theory and Their Applications, vol 10. IMA Volumes in Mathematics and Its Applications. Springer, New York, pp 147–186
33. Harrison JM, Zeevi A (2004) Dynamic scheduling of a multiclass queue in the Halfin-Whitt heavy traffic regime. Oper Res 52:243–257
34. Jun T (2004) A survey on the bandit problem with switching costs. Economist 152:513–541
35. Kallenberg LCM (1986) A note on M. N. Katehakis' and Y.-R. Chen's computation of the Gittins index. Math Oper Res 11:184–186
36. Kämpke T (1987) On the optimality of static priority policies in stochastic scheduling on parallel machines. J Appl Probab 24:430–448
37. Kämpke T (1989) Optimal scheduling of jobs with exponential service times on identical parallel processors. Oper Res 37:126–133
38. Klimov GP (1974) Time-sharing service systems. I. Theory Probab Appl 19:532–551
39. Levy H, Sidi M (1990) Polling systems: applications, modeling, and optimization. IEEE Trans Commun 38:1750–1760
40. Möhring RH, Schulz AS, Uetz M (1999) Approximations in stochastic scheduling: The power of LP-based priority policies. J ACM 46:924–942
41. Nain P, Tsoucas P, Walrand J (1989) Interchange arguments in stochastic scheduling. J Appl Probab 27:815–826
42. Niño-Mora J (2001) Restless bandits, partial conservation laws and indexability. Adv Appl Probab 33:77–98
43. Niño-Mora J (2002) Dynamic allocation indices for restless projects and queueing admission control: a polyhedral approach. Math Program Ser A 93:361–413
44. Niño-Mora J (2006) Marginal productivity index policies for scheduling a multiclass delay-/loss-sensitive queue. Queueing Syst 54:281–312
45. Niño-Mora J (2006) Restless bandit marginal productivity indices, diminishing returns and optimal control of make-to-order/make-to-stock *M*/*G*/1 queues. Math Oper Res 31: 50–84
46. Niño-Mora J (2007) A $(2/3)n^3$ fast-pivoting algorithm for the Gittins index and optimal stopping of a Markov chain. INFORMS J Comput 19:596–606
47. Niño-Mora J (2007) Dynamic priority allocation via restless bandit marginal productivity indices (with discussion). Top 15:161–198
48. Niño-Mora J (2007) Marginal productivity index policies for admission control and routing to parallel multi-server loss queues with reneging. In: Network Control and Optimization: Proceedings of the First EuroFGI International Conference (NET-COOP 2007, Avignon, France), vol 4465. Lecture Notes in Computer Science. Springer, Berlin, pp 138–149
49. Niño-Mora J (2008) A faster index algorithm and a computational study for bandits with switching costs. INFORMS J Comput 20: in press
50. Papadimitriou CH (1985) Games against nature. J Comput Syst Sci 31:288–301
51. Papadimitriou CH, Tsitsiklis JN (1987) On stochastic scheduling with in-tree precedence constraints. SIAM J Comput 16:1–6
52. Papadimitriou CH, Tsitsiklis JN (1999) The complexity of optimal queuing network control. Math Oper Res 24: 293–305
53. Pinedo M (1983) Stochastic scheduling with release dates and due dates. Oper Res 31:559–572
54. Reiman MI, Wein LM (1998) Dynamic scheduling of a two-class queue with setups. Oper Res 4:532–547
55. Righter R (1988) Job scheduling to minimize expected weighted flowtime on uniform processors. Syst Control Lett 10:211–216
56. Robbins H (1952) Some aspects of the sequential design of experiments. Bull Am Math Soc 58:527–535
57. Rothkopf MH (1966) Scheduling with random service times. Manag Sci 12:707–713
58. Sevcik KC (1974) Scheduling for minimum total loss using service time distributions. J ACM 21:66–75
59. Shanthikumar JG, Yao DD (1992) Multiclass queueing systems: Polymatroidal structure and optimal scheduling control. Oper Res 40:S293–S299
60. Smith WE (1956) Various optimizers for single-stage production. Nav Res Logist Quart 3:59–66
61. Tcha DW, Pliska SR (1977) Optimal control of single-server queueing networks and multiclass *M*/*G*/1 queues with feedback. Oper Res 25:248–258
62. Thompson WR (1933) On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika 25:285–294
63. Tsitsiklis JN (1994) A short proof of the Gittins index theorem. Ann Appl Probab 4:194–199
64. Varaiya PP, Walrand JC, Buyukkoc C (1985) Extensions of the multiarmed bandit problem: the discounted case. IEEE Trans Autom Control 30:426–439
65. Weber RR (1992) Scheduling jobs with stochastic pro-

cessing requirements on parallel machines to minimize makespan or flowtime. J Appl Probab 19:167–182
66. Weber RR (1992) On the Gittins index for multiarmed bandits. Ann Appl Probab 2:1024–1033
67. Weber RR, Varaiya P, Walrand J (1986) Scheduling jobs with stochastically ordered processing times on parallel machines to minimize expected flowtime. J Appl Probab 23:841–847
68. Weber RR, Weiss G (1990) On an index policy for restless bandits. J Appl Probab 27:637–648
69. Wein LM (1992) Dynamic scheduling of a multiclass make-to-stock queue. Oper Res 40:724–735
70. Weiss G (1988) Branching bandit processes. Probab Eng Inf Sci 2:269–278
71. Weiss G (1992) Turnpike optimality of Smith's rule in parallel machines stochastic scheduling. Math Oper Res 17:255–270
72. Weiss G, Pinedo M (1980) Scheduling tasks with exponential service times on non-identical processors to minimize various cost functions. J Appl Probab 17:197–202
73. Whittle P (1980) Multiarmed bandits and the Gittins index. J Royal Stat Soc Ser B Method 42:143–149
74. Whittle P (1988) Restless bandits: Activity allocation in a changing world. In: Gani J (ed) A celebration of applied probability, spec vol 25A. Appl Probab Trust, Sheffield, pp 287–298
75. Wie SH, Pinedo M (1986) On minimizing the expected makespan and flow time in stochastic flow shops with blocking. Math Oper Res 11:336–342

Stochastic Transportation and Location Problems

KAJ HOLMBERG
Department Math.,
Linköping Institute Technol.,
Linköping, Sweden

MSC2000: 90B80, 90C11

Article Outline

Keywords

Location; Transportation; Fixed charge; Branch and bound

The basic *transportation problem* (a minimal cost network flow problem in a bipartite graph) is a very well-known problem, which can be efficiently solved with existing methods. It is also an important problem in practice; transporting goods from a set of supply points (factories) to a set of demand points (customers) so as to minimize transportation costs is a situation that often faces planners.

However, in practice the demand of the customers is often not known exactly. In many cases it is best seen as a stochastic amount, with a certain probability function and a certain expected value. Models of this situation also exist in the literature, and quite efficient solution methods have been developed, see for example [1,10] and [9]. The problem is called the *stochastic transportation problem*, (STP), and is a transportation problem with the demand constraints replaced by nonlinear convex costs.

Considering the other end of the transportation problem, there are *facility location* models, which deal with the question of whether or not a certain supply point should be utilized. In such problems, a supply point (facility) is available only if a certain fixed cost is paid. Such models, with linear costs for transportation, are also well known andseveral efficient solution methods exist, see for example [3,6,17], and [21]. (Other variants of location problems are treated in ► Facility location with staircase costs and ► Facility location problems with spatial interaction.)

Obviously both these aspects can be interesting to consider simultaneously, i. e. planning the location of supply facilities and transportation of goods to the customers, while considering the demand as stochastic.This is what we call the *stochastic transportation and location problem*. This problem has received little attention until now. Only a few suggestions for solution methods can be found in the literature, [2,11,12,14]. The latter two papers actually consider a further generalization with generalconcave costs at the supply points, together with the convex costs at thedemand points.