

# Online Stochastic Reservation Systems

Pascal Van Hentenryck, Russell Bent, and Yannis Vergados

Department of Computer Science, Brown University,  
Providence, RI 02912, USA

**Abstract.** This paper considers online stochastic reservation problems, where requests come online and must be dynamically allocated to limited resources in order to maximize profit. Multi-knapsack problems with or without overbooking are examples of such online stochastic reservations. The paper studies how to adapt the online stochastic framework and the consensus and regret algorithms proposed earlier to online stochastic reservation systems. On the theoretical side, it presents a constant sub-optimality approximation of multi-knapsack problems, leading to a regret algorithm that evaluates each scenario with a single mathematical programming optimization followed by a small number of dynamic programs for one-dimensional knapsacks. On the experimental side, the paper demonstrates the effectiveness of the regret algorithm on multi-knapsack problems (with and without overloading) based on the benchmarks proposed earlier.

## 1 Introduction

In an increasingly interconnected and integrated world, online optimization problems are quickly becoming pervasive and raise new challenges for optimization software. Moreover, in most applications, historical data or statistical models are available, or can be learned, for sampling. This creates significant opportunities at the intersection of online algorithms, combinatorial and stochastic optimization, and machine learning. In fact, increasing attention has been devoted to these issues in a variety of communities (e.g., [10, 1, 6, 11, 9, 5, 8]).

This paper considers online stochastic reservation systems and, in particular, the stochastic multi-knapsack problems introduced in [1]. Typical applications include, for instance, reservation systems for holiday centers and advertisement placements in web browsers. These problems differ from the stochastic routing and scheduling considered in, say, [10, 6, 9, 5] in that online decisions are not about selecting the best request to serve but rather about how best to serve a request.

The paper shows how to adapt our online stochastic framework, and the consensus and regret algorithms, to online stochastic reservation systems. Moreover, in order to instantiate the regret algorithm, the paper presents a constant-factor suboptimality approximation for multi-knapsack problems using one-dimensional knapsack problems. As a result, on multi-knapsack problems with or without overbooking, each online decision involves solving a mathematical program and a series of dynamic programs. The algorithms were evaluated on the multi-knapsack problems proposed in [1] with and without overbooking. The results indicate that the regret algorithm is particularly effective, providing significant benefits over heuristic, consensus, and expectation approaches. It

also dominates an earlier algorithm proposed in [1] (which applies the best-fit heuristic within the expectation algorithm) as soon as the time constraints allows for 10 optimizations for each online decision or between each two online decisions. The results are particularly interesting in our opinion, because the consensus and regret algorithms have now be applied generically and successfully to online problems in scheduling, routing, and reservation using, at their core, either constraint programming, mathematical programming, or dedicated polynomial algorithms. The rest of the paper introduces online stochastic reservation problems in their simplest form, shows how to adapt our online stochastic algorithms for them, presents the sub-optimality approximation, and describes the experimental results.

## 2 Online Stochastic Reservation Problems

### 2.1 The Offline Problem

The offline problem is defined in terms of  $n$  bins  $B$  and each bin  $b \in B$  has a capacity  $C_b$ . It receives as input a set  $R$  of requests. Each request is typically characterized by its capacity and its reward, which may or may not depend on which bin the request are allocated to. The goal is to find an assignment of a subset  $T \subseteq R$  of requests to the bins satisfying the problem-specific constraints and maximizing the objective function.

**The Multi-Knapsack Problem** The multi-knapsack problem is an example of a reservation problem. Here each request  $r$  is characterized by a reward  $w(r)$  and a capacity  $c(r)$ . The goal is to allocate a subset  $T$  of the requests  $R$  to the bins  $B$  so that the capacities of the bins are not exceeded and the objective function  $w(T) = \sum_{r \in T} w(r)$  is maximized. A mathematical programming formulation of the problem associates with each request  $r$  and bin  $b$  a binary variable  $x[r, b]$  whose value is 1 when the request is allocated to bin  $b$  and 0 otherwise. The integer program can be expressed as:

$$\begin{aligned} \max \quad & \sum_{r \in R, b \in B} w(r) x[r, b] \\ \text{such that} \quad & \\ & \sum_{b \in B} x[r, b] \leq 1 \quad (r \in R) \\ & \sum_{r \in R} c(r) x[r, b] \leq C_b \quad (b \in B) \\ & x[r, b] \in \{0, 1\} \quad (r \in R, b \in B) \end{aligned}$$

**The Multi-Knapsack Problem with Overbooking** In practice, many reservation systems allow for overbooking. The multi-knapsack problem with overbooking allows the bin capacities to be exceeded but overbooking is penalized in the objective function. To adapt the mathematical-programming formulation above, it suffices to introduce a non-negative variable  $y[b]$  representing the excess for each bin  $b$  and to introduce a penalty term  $\alpha \times y[b]$  in the objective function. The integer programming model now becomes

$$\begin{aligned} \max \quad & \sum_{r \in R, b \in B} w(r) x[r, b] - \sum_{b \in B} \alpha y[b] \\ \text{such that} \quad & \\ & \sum_{b \in B} x[r, b] \leq 1 \quad (r \in R) \\ & \sum_{r \in R} c(r) x[r, b] \leq C_b + y[b] \quad (b \in B) \\ & x[r, b] \in \{0, 1\} \quad (r \in R, b \in B) \\ & y[b] \geq 0 \quad (b \in B) \end{aligned}$$

This is the offline problem considered in [1].

**Generic Formalization** To formalize the online algorithms precisely and generically, it is convenient to assume the existence of a dummy bin  $\perp$  with infinite capacity to assign the non-selected requests and to use  $B_\perp$  to denote  $B \cup \{\perp\}$ . A solution  $\sigma$  can then be seen as a function  $R \rightarrow B_\perp$ . The objective function can be specified by a function  $\mathcal{W}$  over assignments and the problem-specific constraints can be specified as a relation over assignments giving us the problem  $\max_{\sigma: c(\sigma)} \mathcal{W}(\sigma)$ . We use  $\sigma[r \leftarrow b]$  to denote the assignment where  $r$  is assigned to bin  $b$ , i.e.,

$$\begin{aligned}\sigma[r \leftarrow b](r) &= r \\ \sigma[r \leftarrow b](r') &= \sigma(r') \text{ if } r' \neq r.\end{aligned}$$

and  $\sigma \downarrow R$  to denote the assignment where the requests in  $R$  are now unassigned, i.e.,

$$\begin{aligned}(\sigma \downarrow R)(r) &= \perp \quad \text{if } r \in R \\ (\sigma \downarrow R)(r) &= \sigma(r) \text{ if } r \notin R.\end{aligned}$$

Finally, we use  $\sigma_\perp$  to denote the assignment satisfying  $\forall r \in R : \sigma(r) = \perp$ .

## 2.2 The Online Problem

In the online problem, the requests are not known a priori but are revealed online during the execution of the algorithm. For simplicity, we consider a time horizon  $H = [1, h]$  and we assume that a single request arrives at each time  $t \in H$ . (It is easy to relax these assumptions). The algorithm thus receives a sequence of requests  $\xi = \langle \xi_1, \dots, \xi_h \rangle$  over the course of the execution. At time  $i$ , the sequence  $\xi_i = \langle \xi_1, \dots, \xi_i \rangle$  has been revealed, the requests  $\xi_1, \dots, \xi_{i-1}$  have been allocated in the assignment  $\sigma_{i-1}$  and the algorithm must decide how to serve request  $\xi_i$ . More precisely, step  $i$  produces an assignment  $\sigma_i = \sigma_{i-1}[\xi_i \leftarrow b]$  that assigns a bin  $b$  to  $\xi_i$  keeping all other assignments fixed. The requests are assumed to be drawn from a distribution  $\mathcal{I}$  and the goal is to maximize the expected value

$$\mathbf{E}_{\xi}[\mathcal{W}(\sigma_\perp[\xi_1 \leftarrow b_1, \dots, \xi_h \leftarrow b_h])]$$

where the sequence  $\xi = \langle \xi_1, \dots, \xi_h \rangle$  is drawn from  $\mathcal{I}$ .

The online algorithms have at their disposal a procedure to solve, or approximate, the offline problem, and the distribution  $\mathcal{I}$ . The distribution is a black-box available only for sampling. Practical applications often include severe time constraints on the decision time and/or on the time between decisions. To model this requirement, the algorithms may only use the optimization procedure  $\mathcal{O}$  times at each time step.

It is interesting to contrast this online problem with those studied in [7, 5, 3]. In these applications, the key issue was to select which request to serve at each step. Moreover, in the stochastic vehicle routing applications, accepted requests did not have to be assigned a vehicle: the only constraint on the algorithm was the promise to serve every accepted request. The online stochastic reservation problem is different. The key issue is not which request to serve but rather whether and how the incoming request must be served. Indeed, whenever a request is accepted, it must be assigned a specific bin and the algorithm is not allowed to reshuffle the assignments subsequently.

```

ONLINEOPTIMIZATION( $\xi$ )
1  $\sigma_0 \leftarrow \sigma_\perp$ ;
2 for  $t \in H$  do
3    $b \leftarrow \text{CHOOSEALLOCATION}(\sigma_{t-1}, \xi_t)$ ;
4    $\sigma_t \leftarrow \sigma_{t-1}[\xi_t \leftarrow b]$ ;
5 return  $\sigma_h$ ;

```

**Fig. 1.** The Generic Online Algorithm

**The Generic Online Algorithm** The algorithms in this paper share the same online optimization schema depicted in Figure 1. They differ only in the way they implement function CHOOSEALLOCATION. The online optimization schema receives a sequence of online requests  $\xi$  and starts with an empty allocation (line 1). At each decision time  $t$ , the online algorithm considers the current allocation  $\sigma_{t-1}$  and the current request  $\xi_t$  and chooses the bin  $b$  to allocate the request (line 3), which is then included in the new assignment  $\sigma_t$  (line 4). The algorithm returns the last assignment  $\sigma_h$  whose value is  $\mathcal{W}(\sigma_h)$  (line 5). To implement function CHOOSEALLOCATION, the algorithms have at their disposal two black-boxes:

1. a function  $\text{OPTSOL}(\sigma, R)$  that, given an assignment  $\sigma$  and a  $R$  of requests, returns an optimal allocation of the requests in  $R$  given the past decisions in  $\sigma$ . In other words,  $\text{OPTSOL}(\sigma, R)$  solves an online problem where the decision variables for the requests in  $\sigma$  have fixed values.
2. a function  $\text{GETSAMPLE}(t)$  that returns a set of requests over the interval  $[t, h]$  by sampling the arrival distribution.

To illustrate the framework, we specify a best-fit online algorithm as proposed in [1].

**Best Fit (G):** This algorithm assigns the request  $\xi$  to a bin that can accommodate  $\xi$  and has the smallest capacity given the assignment  $\sigma$ :

```

CHOOSEALLOCATION-G( $\sigma, \xi$ )
1 return  $\text{argmin}(b \in B_\perp : \mathcal{C}(\sigma[\xi \leftarrow b])) C_b(\sigma)$ ;

```

where  $C_b(\sigma)$  denotes the remaining capacity of the bin  $b \in B_\perp$  in  $\sigma$ , i.e.,

$$C_b(\sigma) = C_b - \sum_{r \in R: \sigma(r)=b} c(r).$$

### 3 Online Stochastic Algorithms

This section reviews the various online stochastic algorithms. It starts with the expectation algorithm and shows how it can be adapted to incorporate time constraints.

**Expectation (E):** Informally speaking, algorithm E generates future requests by sampling and evaluates each possible allocation against the samples. A simple implementation can be specified as follows:

```

CHOOSEALLOCATION-E( $\sigma_{t-1}, \xi_t$ )
1  for  $b \in B_\perp$  do
2     $f(b) \leftarrow 0$ ;
3  for  $i \leftarrow 1 \dots \mathcal{O}/|B_\perp|$  do
4     $R_{t+1} \leftarrow \text{GETSAMPLE}(t+1)$ ;
5    for  $b \in B_\perp : \mathcal{C}(\sigma_{t-1}[\xi_t \leftarrow b])$  do
6       $\sigma^* \leftarrow \text{OPTSOL}(\sigma_{t-1}[\xi_t \leftarrow b], R_{t+1})$ ;
7       $f(b) \leftarrow f(b) + \mathcal{W}(\sigma^*)$ ;
8  return  $\text{argmax}(b \in B_\perp) f(b)$ ;

```

Lines 1-2 initialize the evaluation  $f(b)$  of each request  $b$ . The algorithm then generates  $\mathcal{O}/|B_\perp|$  samples of future requests (lines 3-4). For each such sample, it successively considers each available bin  $b$  that can accommodate the request  $\xi$  given the assignment  $\sigma_{t-1}$  (line 5). For each such bin  $b$ , it schedules  $\xi_t$  in bin  $b$  and applies the optimization algorithm using the sampled requests  $R_{t+1}$  (line 6). The evaluation of bin  $b$  is incremented in line 7 with the weight of the optimal assignment  $\sigma^*$ . Once all the bin allocations are evaluated over all samples, the algorithm returns the bin  $b$  with the highest evaluation. Algorithm E performs  $\mathcal{O}$  optimizations but uses only  $\mathcal{O}/|B_\perp|$  samples. When  $\mathcal{O}$  is small (due to the time constraints), each request is only evaluated with respect to a small number of samples and algorithm E does not yield much information. To cope with tight time constraints, two approximations of E, consensus and regret, were proposed.

**Consensus (C):** The consensus algorithm C was introduced in [7] as an abstraction of the sampling method used in online vehicle routing [6]. Its key idea is to solve each sample once and thus to examine  $\mathcal{O}$  samples instead of  $\mathcal{O}/|B_\perp|$ . More precisely, instead of evaluating each possible bin at time  $t$  with respect to each sample, algorithm C executes the optimization algorithm once per sample. The bin to which request  $\xi$  is allocated in optimal solution  $\sigma^*$  is credited  $\mathcal{W}(\sigma^*)$  and all other bins receive no credit. Algorithm C can be specified as follows:

```

CHOOSEALLOCATION-C( $\sigma_{t-1}, \xi_t$ )
1  for  $b \in B_\perp$  do
2     $f(b) \leftarrow 0$ ;
3  for  $i \leftarrow 1 \dots \mathcal{O}$  do
4     $R_t \leftarrow \{\xi_t\} \cup \text{GETSAMPLE}(t+1)$ ;
5     $\sigma^* \leftarrow \text{OPTSOL}(\sigma_{t-1}, R_t)$ ;
6     $f(\sigma^*(\xi_t)) \leftarrow f(\sigma^*(\xi_t)) + \mathcal{W}(\sigma^*)$ ;
7  return  $\text{argmax}(b \in B_\perp) f(b)$ ;

```

The core of the algorithm are once again lines 4-6. Line 4 defines the set  $R_t$  of requests that now includes  $\xi_t$  in addition to the sampled requests. Line 5 calls the optimization algorithm with  $\sigma_{t-1}$  and  $R_t$ . Line 6 increments only the bin  $\sigma^*(\xi_t)$ . The main appeal of Algorithm C is its ability to avoid partitioning the available samples between the requests, which is a significant advantage when  $\mathcal{O}$  is small and/or when the number of bins is large. Its main limitation is its *elitism*. Only the best allocation is given some credit for a given sample, while other bins are simply ignored.

**Regret (R):** The regret algorithm R is the recognition that, in many applications, it is possible to estimate the loss of sub-optimal allocations (called regrets) quickly. In other words, once the optimal solution  $\sigma^*$  of a scenario is computed, algorithm E can be approximated with one optimization [5, 2].

**Definition 1 (Regret).** Let  $\sigma$  be an assignment,  $R$  be a set of requests,  $r$  be a request in  $R$ , and  $b$  be a bin. The regret of a bin allocation  $r \leftarrow b$  wrt  $\sigma$  and  $R$ , denoted by  $\text{DEVIATION}(\sigma, R, r \leftarrow b)$ , is defined as

$$| \mathcal{W}(\text{OPTSOL}(\sigma, R)) - \mathcal{W}(\text{OPTSOL}(\sigma[r \leftarrow b], R \setminus \{r\})) | .$$

**Definition 2 (Sub-Optimality Approximation).** Let  $\sigma$  be an assignment,  $R$  be a set of requests,  $r$  be a request in  $R$ , and  $b$  be a bin. Assume that algorithm  $\text{OPTSOL}(\sigma, R)$  runs in time  $O(f_o(R))$ . A sub-optimality approximation runs in time  $O(f_o(R))$  and, given the solution  $\sigma^* = \text{optSol}(\sigma, R)$ , returns, for each bin  $b \in B_\perp$ , an approximation  $\text{SUBOPT}(\sigma^*, \sigma, R, r \leftarrow b)$  to all regrets  $\text{REGRET}(\sigma, R, r \leftarrow b)$ .

Intuitively, the  $|B_\perp|$  regrets must not take more time than the optimization. We are ready to present the regret algorithm R:

```

CHOOSEALLOCATION-R( $\sigma_{t-1}, \xi_t$ )
1  for  $b \in B_\perp$  do
2     $f(b) \leftarrow 0$ ;
3  for  $i \leftarrow 1 \dots \mathcal{O}$  do
4     $R_t \leftarrow \{\xi_t\} \cup \text{GETSAMPLE}(t+1)$ ;
5     $\sigma^* \leftarrow \text{OPTSOL}(\sigma_{t-1}, R_t)$ ;
6     $f(\sigma^*(\xi_t)) \leftarrow f(\sigma^*(\xi_t)) + \mathcal{W}(\sigma^*)$ ;
7    for  $b \in B_\perp \setminus \{\sigma(\xi_t) : \mathcal{C}(\sigma_{t-1}[\xi_t \leftarrow b])\}$  do
8       $f(b) \leftarrow f(b) + (\mathcal{W}(\sigma^*) - \text{SUBOPT}(\sigma^*, \sigma_{t-1}, R_t, \xi_t \leftarrow b))$ ;
9  return  $\text{argmax}(b \in B_\perp) f(b)$ ;

```

Its basic organization follows algorithm C. However, instead of assigning some credit only to the bin selected by the optimal solution, algorithm R (lines 7-8) uses the sub-optimality approximation to compute, for each available allocation  $\xi_t \leftarrow b$ , an approximation of the best solution that allocates  $\xi_t$  to  $b$ . Hence every available bin is given an evaluation for every sample at time  $t$  for the cost of a single optimization (asymptotically). Observe that algorithm R performs  $\mathcal{O}$  optimizations at time  $t$ .

## 4 Generalizations of the Framework

This section discusses two generalizations: precomputation and cancellations.

**Precomputation** Many reservation systems require immediate responses to requests, giving only limited time to the online algorithm for decision making. However, as is the case in vehicle routing, there is time between decisions to generate scenarios and optimize them. This idea can be accommodated in the framework by separating the optimization phase from the decision-making phase in the online algorithm. This is especially attractive for consensus and regret where each scenario is solved exactly once. Details on this separation can be found in [4] in the context of the original framework.

```

ONLINEOPTIMIZATION( $\xi, \zeta$ )
1  $\sigma_0 \leftarrow \sigma_{\perp}$ ;
2 for  $t \in H$  do
3    $\sigma_{t-1} \leftarrow \sigma_{t-1} \downarrow \zeta_t$ ;
4    $b \leftarrow \text{SELECTALLOCATION}(\sigma_{t-1}, \xi_t)$ ;
5    $\sigma_t \leftarrow \sigma_{t-1}[\xi_t \leftarrow b]$ ;
6 return  $\sigma_h$ ;

```

**Fig. 2.** The Generic Online Algorithm with Cancellations

```

CHOOSEALLOCATION-C( $\sigma_{t-1}, \xi_t$ )
1 for  $b \in B_{\perp}$  do
2    $f(b) \leftarrow 0$ ;
3 for  $i \leftarrow 1 \dots \mathcal{O}$  do
4    $\langle R_{t+1}, Z_{t+1} \rangle \leftarrow \text{GETSAMPLE}(t+1)$ ;
5    $\sigma^* \leftarrow \text{OPTSOL}(\sigma_{t-1} \downarrow Z_{t+1}, \{\xi_t\} \cup R_{t+1})$ ;
6    $f(\sigma^*(\xi_t)) \leftarrow f(\sigma^*(\xi_t)) + \mathcal{W}(\sigma^*)$ ;
7 return  $\text{argmax}(b \in B_{\perp}) f(b)$ ;

```

**Fig. 3.** The Consensus Algorithm with Cancellations

**Cancellations** Most reservation systems allow requests to be cancelled after they are accepted. The online stochastic framework can accommodate cancellations by simple enhancements to the generic online algorithm and the sampling procedure. It suffices to assume that an (often empty) set of cancellations  $\zeta_t$  is revealed at step  $t$  in addition to the request  $\xi_t$  and that the function GETSAMPLE return pairs  $\langle R, Z \rangle$  of future requests  $R$  and cancellations  $Z$ . Figure 2 presents a revised version of the generic online algorithm: its main modification is in line 3 which removes the cancellations  $\zeta_t$  from the current assignment  $\sigma_{t-1}$  before allocating a bin to the new request. Figure 3 shows the consensus algorithm with cancellations, illustrating the enhanced sampling procedure (line 4) and how cancellations are taken into account when calling the optimization.

## 5 The Suboptimality Approximation

This section describes a sub-optimality algorithm approximating multi-knapsack problems within a constant factor. Given a set of requests  $R$ , a request  $r \in R$ , and an optimal solution  $\sigma^*$  to the multi-knapsack problem, the sub-optimality algorithm must return approximations to the regrets of allocating  $r$  to bin  $b \in B_{\perp}$ . The sub-optimality algorithm must run within the time taken by a constant number of optimizations.

The key idea behind the suboptimality algorithm is to solve a small number of one-dimensional knapsack problems (which takes pseudo-polynomial time). There are two main cases to study: either request  $r$  is allocated to a bin in  $B$  in solution  $\sigma^*$  or it is not allocated (i.e., it is allocated to  $\perp$ ). In the first case, the algorithm must approximate the optimal solutions in which  $r$  is allocated to other bins (procedure REGRET-SWAP) or not allocated (procedure REGRET-SWAP-OUT). In the second case, the request must be swapped in all the bins (procedure REGRET-SWAP-IN). The rest of this section presents algorithms for the non-overbooking case; they generalize to the overbooking case.

```

REGRET-SWAP( $i, 1, 2$ )
1  $A \leftarrow bin(1, \sigma^*) \cup bin(2, \sigma^*) \cup U(\sigma^*) \setminus \{i\}$ ;
2 if  $C_1 - c(i) \geq C_2$  then
3    $bin(1, \sigma^a) \leftarrow knapsack(A, C_1 - c(i)) \cup \{i\}$ ;
4    $bin(2, \sigma^a) \leftarrow knapsack(A \setminus bin(1, \sigma^a), C_2)$ ;
5 else
6    $bin(2, \sigma^a) \leftarrow knapsack(A, C_2)$ ;
7    $bin(1, \sigma^a) \leftarrow knapsack(A \setminus bin(2, \sigma^a), C_1 - c(i)) \cup \{i\}$ ;
8  $e \leftarrow argmax(r \in bin(1, \sigma^*) \setminus bin(1..2, \sigma^a) : c(r) > \max(C_1 - c(i), C_2)) w(r)$ ;
9 if  $e$  exists &  $w(e) > \max(w(bin(1, \sigma^a)), w(bin(2, \sigma^a)))$  then
10   $j \leftarrow argmax(j \in 3..n) C_j$ ;
11   $bin(j, \sigma^a) \leftarrow knapsack(bin(j, \sigma^a) \cup \{e\}, C_j)$ ;

```

**Fig. 4.** The Suboptimality Algorithm for the Knapsack Problem: Swapping  $i$  from Bin 2 to Bin 1.

Since the names of the bins have no importance, we assume that they are numbered  $1..n$ . Moreover, without loss of generality, we formalize the algorithms to move request  $i$  from bin 2 to bin 1, to swap request  $i$  out of bin 1, and to swap request  $i$  into bin 1. We use  $\sigma^*$  to represent the optimal solution to the multi-knapsack problem,  $\sigma^s$  to denote the optimal solution in which request  $i$  is assigned to bin 1 (REGRET-SWAP and REGRET-SWAP-OUT) or is not allocated (REGRET-SWAP-IN), and  $\sigma^a$  to denote the sub-optimality approximation. We also use  $bin(b, \sigma)$  to denote the requests allocated to bin  $b$  and generalize the notation to sets of bins. The solution to the one-dimensional knapsack problem on  $R$  for a bin with capacity  $C$  is denoted by  $knapsack(R, C)$ . We also use  $c(R)$  to denote the sum of the capacities of the requests in  $R$  and  $U(\sigma^*)$  the requests that are not allocated in the optimal solution  $\sigma^*$ .

**Swapping a Request Between Two Bins** Figure 4 depicts the algorithm to swap request  $i$  from bin 1 to bin 2. The key idea is to consider all requests allocated to bins 1 and 2 in  $\sigma^*$  and to solve two one-dimensional problems for bin 1 (without the capacity taken by request  $i$ ) and bin 2. The algorithm always starts with the bin whose remaining capacity is largest. After solving these two one-dimensional knapsacks, if there exists a request  $e \in bin(1, \sigma^*)$  not allocated in  $bin(1..2, \sigma^a)$  whose value is higher than the values of these two bins, the algorithm solves a third knapsack problem to place this request in another bin if appropriate. This is important if request  $e$  is of high value but cannot be allocated in bin 1 due to the capacity taken by request  $i$ .

**Theorem 1.** *Algorithm REGRET-SWAP is a constant-factor approximation.*

*Proof.* Let  $\sigma^s$  be the sub-optimal solution,  $\sigma^a$  be the regret solution, and  $\sigma^*$  be the optimal solution. Consider the following sets

$$\begin{aligned}
I_1 &= \sigma^s \cap \sigma^a & I_7 &= (bin(2, \sigma^s) \setminus \sigma^a) \cap bin(1, \sigma^*) \\
I_2 &= (bin(1, \sigma^s) \setminus \sigma^a) \cap U(\sigma^*) & I_8 &= (bin(2, \sigma^s) \setminus \sigma^a) \cap bin(2, \sigma^*) \\
I_3 &= (bin(2, \sigma^s) \setminus \sigma^a) \cap U(\sigma^*) & I_9 &= (bin(3..n, \sigma^s) \setminus \sigma^a) \cap bin(1, \sigma^*) \\
I_4 &= (bin(3..n, \sigma^s) \setminus \sigma^a) \cap U(\sigma^*) & I_{10} &= (bin(3..n, \sigma^s) \setminus \sigma^a) \cap bin(2, \sigma^*) \\
I_5 &= (bin(1, \sigma^s) \setminus \sigma^a) \cap bin(1, \sigma^*) & I_{11} &= (bin(1..n, \sigma^s) \setminus \sigma^a) \cap bin(3..n, \sigma^*) \\
I_6 &= (bin(1, \sigma^s) \setminus \sigma^a) \cap bin(2, \sigma^*).
\end{aligned}$$

The suboptimal solution  $\sigma^s$  can be partitioned into  $\sigma^s = \bigcup_{k=1}^{11} I_k$  and the proof shows that  $w(I_k) \leq c_k w(\sigma^a)$  ( $1 \leq k \leq 13$ ) which implies that  $w(\sigma^s) \leq c w(\sigma^a)$  for some constant  $c = c_1 + \dots + c_{11}$ . The proof of each inequality typically separates two cases:

- A:**  $C_1 - c(i) \geq C_2$ ;  
**B:**  $C_1 - c(i) < C_2$ .

Observe also that the proof that  $K(I_1) \leq K(\sigma^a)$  is immediate. We now give the proofs for the remaining sets. In the proofs,  $C'_1$  denotes  $C_1 - c(i)$  and  $K(E, C)$  is defined as follows:

$$K(E, C) = w(\text{knapsack}(E, C)).$$

$I_{2.A}$  : By definition of  $I_2$  and by definition of  $\text{bin}(1, \sigma^a)$  in line 3,

$$K(I_2, C'_1) \leq K(U(\sigma^*), C'_1) \leq K(\text{bin}(1, \sigma^a), C'_1) \leq w(\sigma^a).$$

$I_{2.B}$  : By definition of  $I_2$ ,  $C'_1 < C_2$ , and by definition of  $\text{bin}(2, \sigma^a)$  in line 6

$$K(I_2, C'_1) \leq K(U(\sigma^*), C'_1) \leq K(U(\sigma^*), C_2) \leq K(\text{bin}(2, \sigma^a), C_2) \leq w(\sigma^a).$$

$I_{3.A}$  : By definition of  $I_3$ ,  $C'_1 \geq C_2$ , and by definition of  $\text{bin}(1, \sigma^a)$  in line 3

$$K(I_3, C_2) \leq K(U(\sigma^*), C_2) \leq K(U(\sigma^*), C'_1) \leq K(\text{bin}(1, \sigma^a), C'_1) \leq w(\sigma^a).$$

$I_{3.B}$  : By definition of  $I_3$  and by definition of  $\text{bin}(2, \sigma^a)$  in line 6

$$K(I_3, C_2) \leq K(U(\sigma^*), C_2) \leq K(\text{bin}(2, \sigma^a), C_2) \leq w(\sigma^a).$$

$I_4$  : Assume that  $w(I_4) > w(\sigma^a)$ . This implies

$$\begin{aligned} w(I_4) &> w(\text{bin}(1, \sigma^a)) + w(\text{bin}(2, \sigma^a)) + w(\text{bin}(3..n, \sigma^a)) \\ &> w(\text{bin}(3..n, \sigma^a)) > w(\text{bin}(3..n, \sigma^*)) \end{aligned}$$

which contradicts the optimality of  $\sigma^*$  since  $I_4 \subseteq U(\sigma^*)$ .

$I_{5.A}$  : By definition of  $I_5$  and line 3 of the algorithm

$$K(I_5, C'_1) \leq K(\text{bin}(1, \sigma^*), C'_1) \leq K(A, C'_1) \leq w(\text{bin}(1, \sigma^a)) \leq w(\sigma^a).$$

$I_{5.B}$  : By definition of  $I_5$ ,  $C'_1 \geq C_2$ , and line 6 of the algorithm

$$\begin{aligned} K(I_5, C'_1) &\leq K(\text{bin}(1, \sigma^*), C'_1) \leq K(\text{bin}(1, \sigma^*), C_2) \leq K(A, C_2) \\ &\leq K(\text{bin}(2, \sigma^a), C_2) \leq w(\sigma^a) \end{aligned}$$

$I_{6.A}$  : By definition of  $I_6$  and line 3 of the algorithm

$$K(I_6, C'_1) \leq K(\text{bin}(2, \sigma^*) \setminus \{i\}, C'_1) \leq K(\text{bin}(1, \sigma^a), C'_1) \leq w(\sigma^a)$$

$I_{6.B}$  : By definition of  $I_6$  and line 6 of the algorithm.

$$K(I_6, C'_1) \leq K(\text{bin}(2, \sigma^*) \setminus \{i\}, C_2) \leq K(\text{bin}(2, \sigma^a), C_2) \leq w(\sigma^a)$$

$I_7.A$  : by definition of  $I_7$ ,  $C_2 \leq C'_1$ , and line 3 of the algorithm,

$$K(I_7, C_2) \leq K(I_7, C'_1) \leq K(\text{bin}(1, \sigma^*), C'_1) \leq K(\text{bin}(1, \sigma^a), C'_1) \leq w(\sigma^a).$$

$I_7.B$  : By definition of  $I_7$ ,  $C_2 > C'_1$ , and line 6 of the algorithm

$$K(I_7, C_2) \leq K(\text{bin}(1, \sigma^*), C_2) \leq K(\text{bin}(2, \sigma^a), C_2) \leq w(\sigma^a).$$

$I_8.A$  : By definition of  $I_8$ ,  $C_2 \leq C'_1$ , and line 3 of the algorithm

$$K(I_8, C_2) \leq K(I_8, C'_1) \leq K(\text{bin}(2, \sigma^*), C'_1) \leq K(\text{bin}(1, \sigma^a), C'_1) \leq w(\sigma^a)$$

$I_8.B$  : by definition of  $I_8$ ,  $C_2 > C'_1$ , and line 6 of the algorithm,

$$K(I_8, C_2) \leq K(\text{bin}(2, \sigma^*), C_2) \leq K(\text{bin}(2, \sigma^a), C_2) \leq w(\sigma^a).$$

$I_9.A$  : Consider

$$T = \text{knapsack}(\text{bin}(1, \sigma^*), C'_1);$$

$$L = \text{bin}(1, \sigma^*) \setminus T$$

and let  $e = \arg\text{-max}_{e \in L} w(e)$ . By optimality of  $T$ , we know that  $c(T) + c(e) > C'_1$  and, since  $\text{bin}(1, \sigma^*) = T \cup L$ , we have that  $c(L \setminus \{e\}) < c(i)$ .

If  $w(e) \leq \max(w(\text{bin}(1, \sigma^a)), w(\text{bin}(2, \sigma^a)))$ , then

$$\begin{aligned} w(I_9) &\leq w(T) + w(L \setminus \{e\}) + w(e) \\ &\leq w(\text{bin}(1, \sigma^a)) + w(\text{bin}(2, \sigma^a)) + w(e) \\ &\leq 2(w(\text{bin}(1, \sigma^a)) + w(\text{bin}(2, \sigma^a))) \leq 2w(\sigma^a). \end{aligned}$$

Otherwise, by optimality of  $\text{bin}(1, \sigma^a)$  and  $\text{bin}(2, \sigma^a)$ , we have that

$$c(e) > C'_1 \ \& \ c(e) > C_2$$

and the algorithm executes lines 10–11. If  $c(e) \leq C_j$ , then

$$\begin{aligned} w(I_9) &\leq w(T) + w(L \setminus \{e\}) + w(e) \\ &\leq w(\text{bin}(1, \sigma^a)) + w(\text{bin}(2, \sigma^a)) + w(\text{bin}(j, \sigma^a)) \leq w(\sigma^a). \end{aligned}$$

Otherwise, if  $c(e) > C_j$ ,  $e \notin \sigma^s$  and

$$w(I_9) \leq w(T) + w(L \setminus \{e\}) \leq w(\text{bin}(1, \sigma^a)) + w(\text{bin}(2, \sigma^a)) \leq w(\sigma^a).$$

$I_9.B$  : Consider

$$T = \text{knapsack}(\text{bin}(1, \sigma^*), C_2);$$

$$L = \text{bin}(1, \sigma^*) \setminus T$$

and let  $e = \arg\text{-max}_{e \in L} w(e)$ . If  $w(T) \geq w(L)$ , we have that

$$w(\text{bin}(1, \sigma^*)) \leq 2w(T) \leq 2w(\text{bin}(2, \sigma^a)) \leq 2w(\sigma^a).$$

REGRET-SWAP-OUT( $i, 1$ )  
1  $A \leftarrow \text{bin}(1, \sigma^*) \cup U(\sigma^*) \setminus \{i\}$ ;  
2  $\text{bin}(1, \sigma^a) \leftarrow \text{knapsack}(A, C_1)$ ;

**Fig. 5.** The Suboptimality Algorithm for the Knapsack Problem: Swapping  $i$  out of Bin 1.

Otherwise,  $c(L) > C_2$  by optimality of  $T$  and thus  $c(L) > c(i)$  since  $C_2 \geq c(i)$ . By optimality of  $T$ ,  $c(T \cup \{e\}) > C_2 > C'_1$  and, since  $\text{bin}(1, \sigma^*) = T \cup L$ , it follows that  $c(L \setminus \{e\}) \leq c(i)$ . Hence  $w(L \setminus \{e\}) \leq w(T)$  by optimality of  $T$  and  $w(I_9) \leq w(T) + w(L \setminus \{e\}) + w(e) \leq 2w(T) + w(e) \leq 2w(\text{bin}(2, \sigma^a)) + w(e)$ .

If  $w(e) \leq w(\text{bin}(2, \sigma^a))$ ,  $w(I_9) \leq 3w(\text{bin}(2, \sigma^a)) \leq 3w(\sigma^a)$  and the result follows. Otherwise, by optimality of  $\text{bin}(2, \sigma^a)$ ,  $c(e) > C_2 \geq C'_1$  and the algorithm executes lines 10–11. If  $c(e) \leq C_j$ , then

$$w(I_9) \leq 2w(\text{bin}(1, \sigma^a)) + w(\text{bin}(j, \sigma^a)) \leq w(\sigma^a).$$

Otherwise, if  $c(e) > C_j$ ,  $e \notin \sigma^s$  and

$$w(I_9) \leq w(T) + w(L \setminus \{e\}) \leq 2w(\text{bin}(2, \sigma^a)) \leq 2w(\sigma^a).$$

$I_{10}.A$  : By definition of  $I_{10}$ ,  $C'_1 \geq C_2$ , and line 3 of the algorithm

$$w(I_{10}) \leq w(\text{bin}(2, \sigma^*)) - w(i) \leq w(\text{bin}(1, \sigma^a)) \leq w(\sigma^a).$$

$I_{10}.B$  : By definition of  $I_{10}$  and by line 6 of the algorithm

$$w(I_{10}) \leq w(\text{bin}(2, \sigma^*)) - w(i) \leq w(\text{bin}(2, \sigma^a)) \leq w(\sigma^a).$$

$I_{11}$  : By definition of the algorithm,  $K(\text{bin}(3..n, \sigma^*)) \leq K(3..n, \sigma^a)$ .

□

**Swapping a Request Out of a Bin** The algorithm to swap a request  $i$  out of bin 1 is depicted in Figure 5. It consists of solving a one-dimensional knapsack with the requests already in that bin and the unallocated requests. The proof is similar, but simpler, to the proof of Theorem 1.

**Theorem 2.** *Algorithm REGRET-SWAP-OUT is a constant-factor approximation.*

**Swapping a Request Into a Bin** Figure 6 depicts the algorithm for swapping a request  $i$  in bin 1, which is essentially similar REGRET-SWAP but only uses one bin. It assumes that request  $i$  can be placed in at least two bins since otherwise a single additional optimization suffices to compute all the regrets. Once again, it solves a one-dimensional knapsack for bin 1 (after having allocated request  $i$ ) with all the requests in  $\text{bin}(1, \sigma^*)$  and the unallocated requests. If the resulting knapsack is of low quality (i.e., the remaining requests from  $\text{bin}(1, \sigma^*)$  have a higher value than  $\text{bin}(1, \sigma^a)$ ), REGRET-SWAP-IN solves an additional knapsack problem for the largest available bin. The proof is once again similar to the proof of Theorem 1.

**Theorem 3.** *Assuming that item  $i$  can be placed in at least two bins, Algorithm REGRET-SWAP-IN is a constant-factor approximation.*

```

REGRET-SWAP-IN( $i, 1$ )
1  $A \leftarrow \text{bin}(1, \sigma^*) \cup U(\sigma^*)$ ;
2  $\text{bin}(1, R) \leftarrow \text{knapsack}(A, C_1 - c(i)) \cup \{i\}$ ;
3  $L \leftarrow \text{bin}(1, \sigma^*) \setminus \text{bin}(1, \sigma^a)$ ;
4 if  $w(L) > w(\text{bin}(1, \sigma^a))$  then
5    $j \leftarrow \text{argmax}(j \in 2..n) C_j$ ;
6    $\text{bin}(j, \sigma^a) \leftarrow \text{knapsack}(\text{bin}(j, \sigma^a) \cup L, C_j)$ ;

```

**Fig. 6.** The Suboptimality Algorithm for the Knapsack Problem: Swapping  $i$  into Bin 1.

## 6 Experimental Results

**The Instances** The experimental results use the benchmarks proposed in [1]. Requests are classified in  $k$  types. Each type is characterized by a weight, a value, two exponential distributions indicating how frequently requests of that type arrive and are cancelled, and an overbooking penalty. We generated ten instances based on the master problem proposed in [1]. The goal was to try to produce a diverse set of problems revealing strengths and weaknesses of the various algorithms. The ten problems are named (A-J) here. Problem A scales the master problem by doubling the weight and value of the request types in the master problem, as well as halving the number of items that arrive. Problem B further scales problem A by increasing the weight and value of the types. Problem C considers 7 types of items whose cost ratio takes the form of a bell shape. Problem D looks at the master problem and doubles the number of bins while dividing their capacity by 2. Problem E considers a version of the master problem with bins of variable capacity. Problem F depicts a version of the master problem whose items arrive three times as often and cancel three times as often. Problem G considers a much larger problem with 35 requests types who cost ratio is also shaped in a bell. Problem H is like problem G, the main difference is that the cost ratio shape is reversed. Problem I is a version of G with an extra bin. Problem J is a version of H with fewer bins.

The mathematical programs are solved with CPLEX 9.0 with a time limit of 10 seconds. The optimal solutions can be found within the time limit for all instances but I and J. Every instance is executed under various time constraints, i.e.,  $\mathcal{O} = 1, 5, 10, 25, 50, \text{ or } 100$ , and the results are the average of 10 executions.

It is important to highlight that, on the master problem and its variations, the best-fit heuristic performs quite well. On the offline problems, it is 5% off the optimum in the average and is never worse than 10% off. This will be discussed again when the regret algorithm is compared to earlier results.

**Comparison of the Algorithms** Figure 7 describes the average profit (left) and loss (right) of the various online algorithms as a percentage of the optimal offline solution. The loss sums the weights of the rejected requests and the overbooking penalty (if any); it is often used in comparing online algorithms as it gives a sense of the “price” of uncertainty. The results clearly show the value of stochastic information as algorithms R, C, E recovers most of the gap between the online best-fit heuristic (G) and the offline optimum (which cannot typically be achieved in an online setting). Moreover, they show that algorithms R and C achieve excellent results even with small number of available optimizations (tight time constraints). In particular, algorithm R achieves

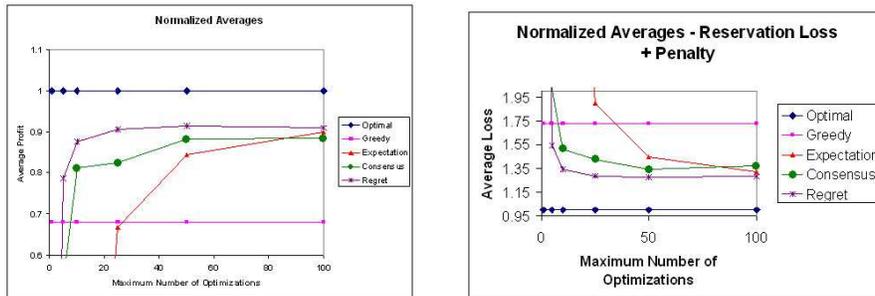


Fig. 7. Experimental Results over All Instances with Overbooking Allowed.

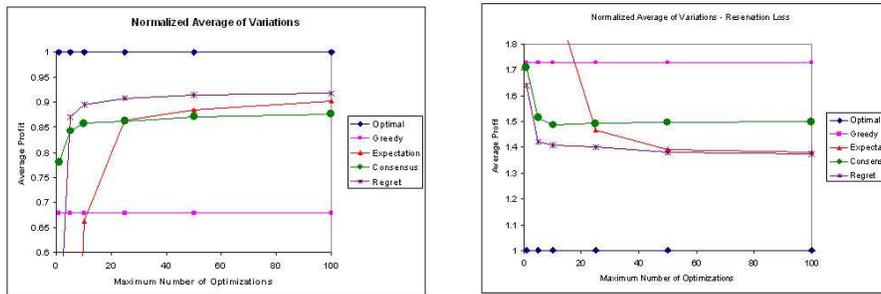


Fig. 8. Experimental Results over All Instances with Overbooking Disallowed.

about 89% of the offline optimum with only 10 samples and 91% with 50 optimizations. It also achieves a loss of 28% over the offline optimum for 25 optimizations and 34% for 10 optimizations. The regret algorithm clearly dominates the expectation algorithm E which performs poorly for tight time constraints. It becomes reasonable for 50 optimizations and reaches the quality of the regret algorithm for 100 optimizations.

Figure 8 shows the same results when no overbooking is allowed. These instances are easier in the sense that fewer optimizations are necessary for the algorithms to converge. But they exhibit the same pattern as when overbooking is allowed. These results are quite interesting and shows that the benefits of the regret algorithm increase with the problem complexity but are significant even on easier instances.

**Comparison with Earlier Results** As mentioned earlier, the best-fit algorithm is only 5% below the optimal offline solution in these problems. It is thus tempting to replace the IP solver in algorithm E by the best-fit heuristic to evaluate more samples. The algorithm, denoted by BF EXP, was proposed in [1] and was shown to be superior to several approaches including yield management and an hybridization with Markov Models [12]. Because the best-fit algorithm is so fast, BF EXP can easily be run with 10,000 samples and remedies the limitations of algorithm E under tight time constraints.

Figure 9 compares algorithms BF EXP, R, and C when overbooking is allowed. The results show that BF EXP indeed produces excellent results but is quickly dominated by R as time increases. In particular, the loss of BF EXP is above 40%, although it goes down to 34% for 10 optimizations and 28% for 25 optimizations in algorithm R.

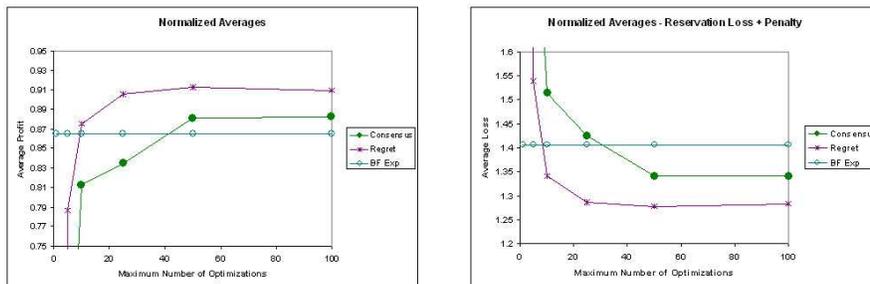


Fig. 9. Comparison with Earlier Results: Average Results for Instances with Overbooking

Similarly, the profit increases by 4% in the average starting at 25 optimizations. BF EXP is also dominated by algorithm C but only for 50 optimizations or more.

What is quite remarkable here is that the 5% difference in quality between the best-fit heuristic and the offline algorithm translates into a similar difference in quality in the online setting. Moreover, when looking at specific instances, one can see that BF EXP is often comparable to R but its loss (resp. profit) may be significantly higher (resp. lower) on instances that seem particularly difficult. This is the case for instances E and G, where the gap between the offline solutions and the solutions by algorithm R is larger. This seems to indicate that the harder the problems the more beneficial algorithm R becomes. This in fact confirms our earlier results on stochastic vehicle routing where the algorithms use a large neighborhood heuristic [3, 13]. Indeed, using a simpler, lower-quality, heuristic on more samples did not produce high-quality results in an online setting. The results presented here also show that the additional information produced by a more sophisticated solver quickly amortizes its computational cost, making algorithm R particularly effective and robust for many problems.

## 7 Conclusion

This paper adapted our online stochastic framework and algorithms to online stochastic reservations initially proposed in [1]. These problems, whose core can be modelled as multi-knapsacks, are significant in practice and are also different from the scheduling and routing applications we studied earlier. Indeed the main decision is not which request to select next but rather how best to serve a request given limited resources. The paper shows that the framework and its associated algorithms naturally apply to online reservation systems. It also presented a constant-factor sub-optimality approximation of multi-knapsack problems that only solves one-dimensional knapsack problems, leading to a regret algorithm that uses both mathematical programming and dynamic programming algorithms. The algorithms were evaluated on the multi-knapsack problems proposed in [1] with and without overbooking. The results indicate that the regret algorithm is particularly effective, providing significant benefits over heuristic, consensus, and expectation approaches. It also dominates an earlier algorithm proposed in [1] (which applies the best-fit heuristic with algorithm E) as soon as the time constraints allows for 10 optimizations at decision time or between decisions. Even more interesting

perhaps, the regret algorithm has now been applied to online stochastic problems where the offline problem is solved by either constraint programming, integer programming, or (special-purpose) polynomial algorithms, indicating its versatility and benefits for a wide variety of applications.

## References

1. T. Benoist, E. Bourreau, Y. Caseau, and B. Rottembourg. Towards stochastic constraint programming: A study of online multi-choice knapsack with deadlines. In *Proceedings of the Seventh International Conference on Principles and Practice of Constraint Programming (CP'01)*, pages 61–76, London, UK, 2001. Springer-Verlag.
2. R. Bent, I. Katriel, and P. Van Hentenryck. Sub-Optimality Approximation. In *Eleventh International Conference on Principles and Practice of Constraint Programming*, Stiges, Spain, 2005.
3. R. Bent and P. Van Hentenryck. A Two-Stage Hybrid Local Search for the Vehicle Routing Problem with Time Windows. *Transportation Science*, 8(4):515–530, 2004.
4. R. Bent and P. Van Hentenryck. Online Stochastic and Robust Optimization. In *Proceeding of the 9th Asian Computing Science Conference (ASIAN'04)*, Chiang Mai University, Thailand, December 2004.
5. R. Bent and P. Van Hentenryck. Regrets Only. Online Stochastic Optimization under Time Constraints. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI'04)*, San Jose, CA, July 2004.
6. R. Bent and P. Van Hentenryck. Scenario Based Planning for Partially Dynamic Vehicle Routing Problems with Stochastic Customers. *Operations Research*, 52(6), 2004.
7. R. Bent and P. Van Hentenryck. The Value of Consensus in Online Stochastic Scheduling. In *Proceedings of the 14th International Conference on Automated Planning & Scheduling (ICAPS 2004)*, Whistler, British Columbia, Canada, 2004.
8. R. Bent and P. Van Hentenryck. Online Stochastic Optimization without Distributions . In *Proceedings of the 15th International Conference on Automated Planning & Scheduling (ICAPS 2005)*, Monterey, CA, 2005.
9. A. Campbell and M. Savelsbergh. Decision Support for Consumer Direct Grocery Initiatives. *Report TLI-02-09, Georgia Institute of Technology*, 2002.
10. H. Chang, R. Givan, and E. Chong. On-line Scheduling Via Sampling. *Artificial Intelligence Planning and Scheduling (AIPS'00)*, pages 62–71, 2000.
11. B. Dean, M.X. Goemans, and J. Vondrak. Approximating the Stochastic Knapsack Problem: The Benefit of Adaptivity. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 208–217, Rome, Italy, 2004.
12. M. Puterman. *Markov Decision Processes*. John Wiley & Sons, New York, 1994.
13. P. Shaw. Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems. In *Proceedings of Fourth International Conference on the Principles and Practice of Constraint Programming (CP'98)*, pages 417–431, Pisa, October 1998.