Reactive Control of Dynamic Progressive Processing

Shlomo Zilberstein
Computer Science Department
University of Massachusetts
Amherst, MA 01003 U.S.A.
zilberstein@cs.umass.edu

Abdel-Illah Mouaddib
CRIL-IUT de Lens-Université d'Artois
Rue de l'université, S. P. 16
62307 Lens Cedex France
mouaddib@cril.univ-artois.fr

Abstract

Progressive processing is a model of computation that allows a system to tradeoff computational resources against the quality of results. This paper generalizes the existing model to make it suitable for dynamic composition of information retrieval techniques. The new framework addresses effectively the uncertainty associated with the duration and output quality of each component. We show how to construct an optimal meta-level controller for a single task based on solving a corresponding Markov decision problem, and how to extend the solution to the case of multiple and dynamic tasks using the notion of an opportunity cost.

1 Introduction

This paper is concerned with the design and implementation of an effective computational model for information retrieval search engines that can tradeoff computational resources against the quality of the result. Our approach is based on run-time monitoring of the information retrieval process and on dynamic selection of information retrieval techniques so as to maximize the quality of the answer produced with limited computational resources. Information retrieval from a large collection involves uncertainty regarding the duration of the process and the quality of the result. In addition, there may be large variability in the number of requests that require a response at any given time. By taking a context dependent, dynamic approach to the problem we can significantly improve the average quality of service provided by such systems.

A typical search engine is composed of several information retrieval modules that perform such tasks as query formation, query optimization, query evaluation, precision improvement, recall improvement, clustering, and results visualization. For each one of these phases, there are a wide variety of techniques that have been developed in recent years [Jones and Willett, 1997]. Currently, search engines are built by choosing and integrating a fixed set of modules and techniques. The choices are made off-line by the designer of the system. This static approach excludes techniques that work well in special situations. In addition, current information retrieval systems are optimized for a particular load; they cannot respond dynamically to varying load, availability of computational resources, and to the specific characteristics of a given query.

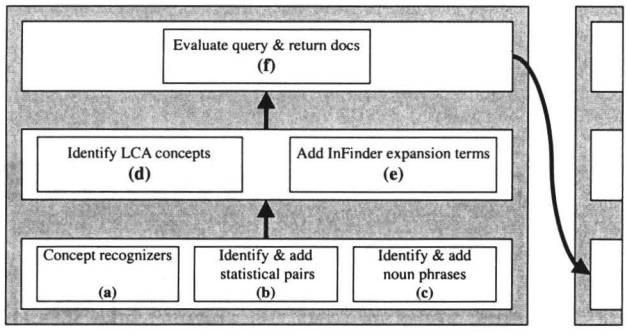

Figure 1: Illustration of a progressive processing task for an information retrieval search engine

The ability to dynamically adjust computational effort based on the availability of computational resources has been studied extensively by the AI community since the mid 1980's. These efforts have led to the development of a variety of techniques such as *anytime algorithms* [Dean and Boddy, 1988; Zilberstein and Russell, 1996], *design-to-time* [Gravey and Lesser, 1993], *flexible computation* [Horvitz, 1988], *imprecise computation* [Liu *et al.*, 1991], and *progressive reasoning* [Mouaddib, 1993; Mouaddib and Zilberstein, 1997].

We adopt the progressive processing framework to formalize and solve the meta-level control problem of a real-time information retrieval application. The technique maps each task to a *progressive processing unit* (PRU). Each PRU is composed of a set of modules that can contribute to the quality of the result. This offers a natural framework to describe the set of information retrieval techniques available to the system. Figure 1 shows a simple task structure whose input is a *query* composed of a list of keywords. The task structure has three processing *levels*. The first level includes three alternative techniques to improve the initial query: (a) scan the query

using concept recognizers to identify company names, dates, locations, personal names, and so on; (b) examine the query for pairs of words that have high statistical likelihood of being related and enhance the query with that information; (c) perform part-of-speech analysis to identify noun phrases within the query. The second level includes two alternative techniques that can improve the query's recall ability by expanding it to include related words and phrases: (d) use of Local Context Analysis (LCA), a statistical method for expanding queries that relies upon in-context analysis of word co-occurrence; (e) use of InFinder, an association thesaurus that is faster than LCA and does not capture context as well. Finally, the third level performs the actual query evaluation and returns the results. Quality in this application is measured by the number of relevant documents within the top n documents retrieved (i.e., precision in the retrieved set).

The information retrieval application and the resulting task structure raise several fundamental issues that have not been previously addressed.

1. Handling the *duration uncertainty* and *quality uncertainty* associated with each technique.
2. Handling the *dependency* of quality and duration on the quality of intermediate results.
3. Handling a rich task structure in which some levels include several alternatives or optional computational steps; optional steps can be skipped under time pressure, leading to direct evaluation of the input query.
4. Selecting the "best" set of retrieval techniques in a dynamic environment taking into account the entire set of queries waiting for execution.

The rest of this paper offers an efficient solution to the meta-level control problem. Section 2 gives a formal definition of the problem. We then solve the problem in two steps. In Section 3, we develop an optimal solution for a single PRU, ignoring the fact that additional tasks are waiting for processing. Section 4 shows how to handle multiple PRUs using the approach developed in Section 3 and summarizing the effect of the waiting requests using the notion of an opportunity cost. In section 5 we address the issue of reactive control in a highly dynamic environment by estimating the opportunity cost and pre-compiling the control policies. We conclude with a summary of the results and a brief description of related work.

2 The meta-level control problem

This section describes more formally the problem of meta-level control of the (enhanced) progressive processing model. Each information retrieval request is mapped to a task structure described below.

Definition 1 *A* **progressive processing unit** *(PRU) is composed of a sequence of processing levels,* $(l_1, l_2, \ldots, l_L)$. *The first level receives the input query and the last one produces the result.*

Definition 2 *Each processing level,* l_i*, is composed of a set of* p_i **alternative modules,** $\{m_i^1, m_i^2, \ldots, m_i^{p_i}\}$.

Each module can perform the logical function of level l_i, but it has different computational characteristics defined below.

Definition 3 *The* **module descriptor,** $P_i^j((q', \delta)|q)$, *of module* m_i^j *is the probability distribution of output quality and duration for a* given *input quality.*

Note that q is a discrete variable representing quality and δ is a discrete variable representing duration. The module descriptor specifies the probability that module m_i^j takes δ time units and returns a result of quality q' when the quality of the previously executed module is q. Module descriptors are similar to *conditional performance profiles* of anytime algorithms [Zilberstein and Russell, 1996]. They are constructed empirically by collecting performance data for a sample set of inputs.

When the search engine responds to a particular request, it receives an immediate reward defined as follows.

Definition 4 *A* **time-dependent utility function,** $U(q, t)$, *measures the utility of a solution of quality* q *if it is returned* t *time units after the arrival time of the request.*

We assume that there is a given constant T such that $\forall q, t > T \;:\; U(q, t) = 0$. That is, responding to a request more than T time units after its arrival has no value.

Suppose that a system maintains a set of information retrieval requests, W, with arrival times $\{a_1, a_2, \ldots, a_n\}$. The set of requests is updated dynamically as new requests arrive. The system processes the requests in a first-in-first-out order using a progressive processing unit to handle each request.

Given a set of requests, the module descriptors of all the components of the progressive processing unit, and a time-dependent utility function, we define the following control problem.

Definition 5 *The* **reactive control problem** *is the problem of selecting a set of alternative modules so as to maximize the expected utility over the set of information retrieval requests.*

The meta-level control is "reactive" in the sense that we assume that the module selection mechanism is very fast, largely based on off-line analysis of the problem. The rest of the paper provides a solution to this problem.

3 Optimal control of a single PRU

We begin with the problem of meta-level control of a single progressive processing unit corresponding to a single task. This problem can be formulated as a simple Markov decision process (MDP) with states representing the current state of the computation. The state includes the current level of the PRU, the quality produced so far, and the elapsed time since the arrival of the request. The rewards are defined by the utility of the solution which

depends on both quality and time. The possible actions are to *execute* a module of the next processing level or to *skip* that processing level. The transition model is defined by the descriptor of the module selected for execution. The rest of this section gives a formal definition of the MDP and the reactive controller produced by solving it.

3.1 State representation

The execution of a single progressive processing unit, u, can be seen as an MDP with a finite set of states $\mathcal{S} = \{[l_i, q, t] | l_i \in u\} \bigcup \{[\text{failure}, t]\}$ where $0 \leq i \leq L$ indicates the last executed (or skipped) level, $0 \leq q \leq 1$ is the quality produced by the last executed module, and $0 \leq t \leq T$ is the elapsed time since the arrival time, a_u, of the request. Note that quality is discretized and normalized to be in the range [0..1]. All the intermediate modules use a uniform representation of input and output (a "query" in our application). Note also that T is the maximum delay after which we consider the response to be useless. When the system is in state $[l_i, q, t]$, one module of the i-th level has been executed. (The first level is $i = 1$; $i = 0$ is used to indicate the fact that no level has been executed.) The states $[\text{failure}, t]$ represent termination at time t without any useful result. We distinguish between different failure states because failure can occur before the deadline leaving some remaining time for the execution of other requests in the queue.

3.2 Transition model

The initial state of the MDP is $[l_0, q_{init}, t]$, where t is the elapsed time since the arrival of the request ($t =$ current time $- a_u$) and q_{init} is the initial quality of the request (0 in our application). The initial state indicates that the system is ready to start executing a module of the first level of the PRU. The terminal states are all the states of the form $[l_L, q, t]$ or $[failure, t]$. The former set represents finishing execution of the last level and the latter set represents failure. Other states such as $[l_i, q_{max}, t]$ (reaching maximal intermediate quality) or $[l_i, q, T]$ (reaching the deadline before the execution of the last level) are not considered terminal states. A terminal state can be reached from state $[l_i, q, T]$ by executing a series of *skip* actions until a failure state is reached. Similarly *skip* actions take the automaton from state $[l_i, q_{max}, t]$ to the last level because no *execute* action can improve the intermediate quality.

In every nonterminal state the possible actions are: $\mathbf{E}^j_{i+1}$ (execute the j-th module of the next level) and $\mathbf{S}$ (skip the next level). To complete the transition model, we need to specify the probabilistic outcome of these actions. Equations 1-4 define the transition probabilities for a given nonterminal state $[l_i, q, t]$.

The $\mathbf{S}$ action is deterministic. It skips the next level without affecting the quality or elapsed time. (It can be implemented as an additional "dummy" module whose execution takes no time and has no effect on quality.)

$$Pr([l_{i+1}, q, t] \mid [l_i, q, t], \mathbf{S}) = 1 \quad \text{when } 0 \leq i < L-1 \tag{1}$$

Skipping the last level results in failure.

$$Pr([failure, t] \mid [l_{L-1}, q, t], \mathbf{S}) = 1 \tag{2}$$

The $\mathbf{E}^j_{i+1}$ action is probabilistic. Duration and quality uncertainties define the new state. Equation 3 determines the transitions following successful execution and Equation 4 determines the transition to the failure state when the deadline, T, is reached.

$$Pr([l_{i+1}, q', t+\delta] \mid [l_i, q, t], \mathbf{E}^j_{i+1}) = P^j_{i+1}((q', \delta)|q) \quad \text{when } t + \delta \leq T \tag{3}$$

$$Pr([failure, T] \mid [l_i, q, t], \mathbf{E}^j_{i+1}) = \sum_{q', \delta > T-t} P^j_{i+1}((q', \delta)|q) \tag{4}$$

3.3 Rewards and the value function

Rewards are determined by the given time-dependent utility function applied to the final result (produced by the last level of the PRU). The utility depends on the quality of the result and the elapsed time. Keep in mind that in our application the intermediate results are useless and therefore have no direct rewards associated with them. We now define a value function (expected reward-to-go) over all states. The value of terminal states is defined as follows.

$$V([l_L, q, t]) = R(q, t) = U(q, t) \tag{5}$$

$$V([failure, t]) = R(0, t) = U(0, t) \tag{6}$$

The value of nonterminal states of the MDP is defined as follows.

$$V([l_i, q, t]) = \max_a \begin{cases} V([l_{i+1}, q, t]) & \text{If } a = \mathbf{S}, 0 \leq i < L-1 \\ V([failure, t]) & \text{If } a = \mathbf{S}, i = L-1 \\ EV([l_i, q, t] \mid \mathbf{E}^j_{i+1}) & \text{If } a = \mathbf{E}^j_{i+1}, 0 < j \leq p_i \end{cases} \tag{7}$$

Such that $EV([l_i, q, t] \mid \mathbf{E}^j_{i+1}) =$

$$\sum_{q', \delta > T-t} P^j_{i+1}((q', \delta)|q)\ V([failure, T]) + \sum_{q', \delta \leq T-t} P^j_{i+1}((q', \delta)|q)\ V([l_{i+1}, q', t+\delta])$$

The value function is defined as maximum over all actions with the top expression representing the value of a skip action for any level l_i such that $0 \leq i \leq L-1$, the middle expression representing the value of a skip action for level l_{L-1}, and the bottom expression representing the value of an execute action.

This concludes the definition of an MDP. This MDP is a finite-horizon MDP with no cycles. It can be solved easily using standard dynamic programming algorithms or using search algorithms such as AO*.

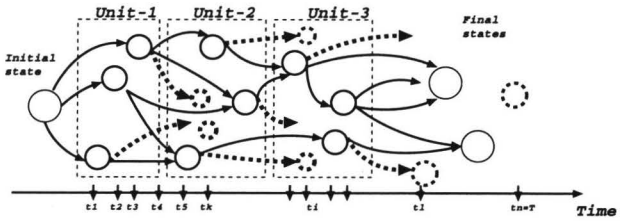

Figure 2: Illustration of the execution paths with multiple PRUs

Theorem 1 *Given* one *progressive processing unit u and a time-dependent utility function $U(q,t)$, the optimal policy for the corresponding MDP is an optimal reactive control for u.*

Proof: Because there is a one-to-one correspondence between the reactive control problem and the MDP (including the fact that the PRU transition model satisfies the Markov assumption), and because of the optimality of the resulting policy, we conclude that it provides optimal reactive control for the progressive processing problem. □

We note that the number of states of the MDP is bounded by the product of the number of levels L, the maximum number of alternative modules per level $\max_i p_i$, the number of discrete quality levels, and the maximum execution time. While the maximum execution time can be quite large, the time unit used for the purpose of meta-level control is an arbitrary system parameter. A small time unit leads to a more effective control at the expense of a larger state-space. The choice of a unit of quality has a similar effect. These units introduce a tradeoff between the size of the policy and its effectiveness. Further empirical evaluation is needed to determine the best choice of time and quality units in practice.

4 Optimal Control of multiple PRUs using opportunity cost

Suppose now that we need to schedule the execution of multiple PRUs. We assume that there are $n+1$ requests whose arrival times are $a_0 \leq a_1 \leq \ldots \leq a_n$. One approach to construct an optimal schedule is to generalize the solution presented in the previous section. We can construct a larger MDP for the combined sequential decision problem including the entire set of $n+1$ PRUs (as illustrated in Figure 2). To do that, each state must also include i, the request number, leading to a general state represented as $[i,l,q,t]$. Note that t is the elapsed time since the arrival of the *first* request.

This rather complex MDP is still a finite-horizon MDP with no loops. Moreover, the only possible transitions between different PRUs are from a terminal state of one PRU to an initial state of a succeeding PRU. Therefore, we can solve this MDP by computing an optimal policy for the *last* PRU for any starting time between 0 and $T + a_{n+1} - a_0$, then use the value of its initial states to compute an optimal policy for the previous PRU and so on.

Theorem 2 *Given a set, W, of progressive processing units and a time-dependent utility function $U(q,t)$, the optimal policy for the corresponding MDP is an optimal reactive control for W.*

This is an obvious generalization of Theorem 1. The complete proof, by induction on the number of PRUs, is omitted.

We now show how to reformulate the effect of the remaining n requests on the execution of the first task. This reformulation preserves the optimality of the solution, but it suggests a more efficient control structure developed in Section 5.

Definition 6 *Let $V_i^*(t) = V([i,l_0,q_0,t])$ denote the expected value of the optimal policy for the last $n-i+1$ PRUs.*

To compute the optimal policy for the i-th PRU, we can simply use the following reward function.

$$R_i(q,t) = U(q, t + a_0 - a_i) + V([i+1, l_0, q_0, t]) \quad (8)$$

In other words, the reward for responding to the i-th request is composed of the immediate reward (defined by the time-dependent utility function) and the reward-to-go (defined by the remaining PRUs). Alternatively, the reward can be represented as follows.

$$R_i(q,t) = U(q, t + a_0 - a_i) + V_{i+1}^*(t) \quad (9)$$

Therefore, the best policy for the first PRU can be calculated if we use the following reward function for final states:

$$R_0(q,t) = U(q,t) + V_1^*(t) \quad (10)$$

Definition 7 *Let $OC(t) = V_1^*(0) - V_1^*(t)$ be the* **opportunity cost** *at time t.*

The opportunity cost measures the loss of expected value due to delay in the starting point of executing the last n tasks (all the tasks except the first one).

Definition 8 *Let the* **OC-policy** *for the first PRU be the policy computed with the following reward function:*

$$R(q,t) = U(q,t) - OC(t)$$

The OC-policy is the policy computed by deducting from the actual reward for the first task the opportunity cost of its execution time.

Theorem 3 *Controlling the first PRU using the OC-policy is optimal.*

Proof: From the definition of $OC(t)$ we get:

$$V_1^*(t) = V_1^*(0) - OC(t) \quad (11)$$

To compute the optimal schedule we need to use the reward function defined in Equation 9 that can be rewritten as follows.

$$R_0(q,t) = U(q,t) + V_1^*(0) - OC(t) \quad (12)$$

But this reward function is the same as the one used to construct the OC-policy, except for the added constant $V_1^*(0)$. Because adding a constant to a reward function does not affect the policy, the conditions of Theorem 2 are met and the resulting policy is optimal. □

Theorem 3 suggests an optimal approach to scheduling the entire $n+1$ requests by first using an OC-policy for the first request that takes into account the opportunity cost of the remaining n requests. Then the OC-policy for the second request is used taking into account the opportunity cost of the remaining $n-1$ tasks and so on. To be able to implement this approach we need to have the control policies readily available. This issue is addressed in the following section.

5 Reactive control based on estimated opportunity cost

In the previous section, we presented an optimal solution to the control problem of multiple progressive processing units without accounting for its computational complexity. In particular, the opportunity cost must be computed and revised quickly each time a new request arrives. Once the opportunity cost is revised, a new policy for the current PRU must be constructed. Finding the exact opportunity cost requires the construction of an optimal policy for the entire set of tasks. In practice, this may slow down the operation of the information retrieval search engine.

In order to provide an effective reactive controller for dynamic progressive processing, it is necessary to:

1. use a fast approximation scheme to estimate the opportunity cost; and
2. use pre-compiled policies for different opportunity cost functions.

The rest of this section explains this method in more detail.

5.1 Estimating the Opportunity Cost

The opportunity cost is defined in terms of the function V_1^* which represents the value of an optimal policy for the remaining tasks in the queue. Thus, it can be estimated by approximating this function. One way to approximate the cumulative value of the remaining tasks is to add the value of each task *without* taking into account the opportunity cost. In this calculation, the start time of each task is the *expected* end time of the previous one. The following set of equations summarizes this approximation scheme.

$$\begin{aligned} V_1^*(t) &\simeq V([l_0, q_0, t + a_0 - a_1]) + V_2^*(t + \tau_1) \\ &\vdots \\ V_i^*(t) &\simeq V([l_0, q_0, t + a_0 - a_i]) + V_{i+1}^*(t + \textstyle\sum_{j=1}^{j=i} \tau_j) \\ &\vdots \\ V_n^*(t) &= V([l_0, q_0, t + a_0 - a_n]) \end{aligned} \tag{13}$$

where $V[l, q, t]$ is the value function defined in Section 3 for a single PRU. Therefore, V_1^* can be approximated as follows.

$$V_1^*(t) \simeq \sum_{i=1}^{i=n} V[l_0, q_0, t + a_0 + (\sum_{j<i} \tau_j) - a_i] \tag{14}$$

The expected duration of task i, τ_i, depends on the duration of the previous tasks. Let $\tau(t)$ be the expected duration of the optimal single-PRU policy when starting at time t (relative to the arrival time of the request). Then τ_i is computed using τ with the expected starting time of task i relative to *its* arrival time.

$$\begin{aligned} \tau_0 &= 0 \\ \tau_i &= \tau(t + a_0 + (\textstyle\sum_{j<i} \tau_j) - a_i) \end{aligned} \tag{15}$$

The function τ (expected duration) can be computed for any finite-horizon MDP once the optimal policy is available by simply using durations as rewards. The function can be computed once off-line, making it easy to revise the opportunity cost when a new request is added.

We are also examining an alternative learning technique to approximate the opportunity cost function based on such features as total queue size and total waiting time. Standard function approximation techniques can be used resulting in a compact representation of the opportunity cost. It is not clear at this point which approximation technique works better in practice. This will be determined based on empirical evaluation.

5.2 Pre-compiled control policies

To make the meta-level control truly reactive for large task structures, one may want to avoid computing a new policy (for a single PRU) each time the opportunity cost is revised. To avoid this, the space of opportunity cost can be divided into a small set of regions representing typical situations. For example, there could be just three regions that capture *low*, *medium*, and *high* loads. For each region, an optimal policy would be computed off-line and stored in a library. At run-time, the system will first estimate the opportunity cost and then use the appropriate pre-compiled policy from the library. These policies remain valid as long as the overall task structure and the utility function are fixed. Because the dependency of the control decisions on the opportunity cost is monotonic (higher costs imply less time for execution), we anticipate that a small set of classes that correspond to *qualitatively* different action selection will be sufficient.

Another advantage of the use of pre-compiled policies is the ability to react quickly to dynamic changes. Control policies can be switched *during* the execution of a single request if the opportunity cost changes. This is possible because the policies share the same state space.

6 Conclusion

We present an innovative approach to meta-level control of progressive processing based on reformulating it as a

Markov decision problem. It is shown that an optimal policy for a set of tasks can be constructed by controlling a single PRU, taking into account the opportunity cost of the remaining tasks. To apply this model to control the operation of an information retrieval search engine, a fast approximation of the opportunity cost is developed. Finally, a highly reactive controller is described that uses a library of pre-compiled control policies to operate in a dynamic environment.

A less complex model of progressive processing that relies on heuristic scheduling has been developed [Mouaddib and Zilberstein, 1997]. The task structure, however, is limited to a linear set of levels with one module per level and no quality uncertainty or quality dependency. The heuristic scheduler is fast, but it cannot solve the more complex task structure presented in this paper and it does not provide optimal control. Heuristic scheduling of computational tasks has also been studied by Garvey and Lesser [1993] for the *design-to-time* problem-solving framework. The latter framework represents explicitly non-local interactions between sub-tasks.

The progressive processing framework relates to a large body of work within the systems community on *imprecise computation* [Liu *et al.*, 1991]. Each task in that model is decomposed into a *mandatory* subtask and an *optional* subtask. A variety of scheduling algorithms have been developed for imprecise computation under different assumptions about the optional part. Our model allows for a richer representation of quality and duration uncertainty and quality dependency. Unlike imprecise computation, the schedule constructed by the MDP scheduler is a *conditional schedule*; the selection of modules is conditioned on the *actual* execution time and outcome of previous modules.

The application of dynamic programming to solve meta-level control problems have been previously used by Hansen and Zilberstein [1996] to control interruptible anytime algorithms. Optimal monitoring of progressive processing tasks using a corresponding MDP has been studied by Mouaddib and Zilberstein [1998] with respect to a simpler task structure and without the notion of quality uncertainty and quality dependency.

The notion of *opportunity cost* is borrowed from economics. It has been used previously in meta-level reasoning by Russell and Wefald [1991]. Horvitz [1997] uses a similar notion to develop a model of *continual computation* in which idle time is used to solve anticipated future problems.

The use of pre-compiled control policies to construct a highly reactive real-time system has been studied by several researchers. For example, Greenwald and Dean [1998] show how a real-time avionics control system can use a library of schedules that cover all possible situations. Each schedule is conditioned on the state of the flight operation.

In collaboration with the Information Retrieval Center at UMass we are currently developing the stochastic module descriptors for the components of the search engine. By definition, IR tasks involve large collections and a substantial amount of test data allowing us to test the applicability and scalability of this resource-bounded reasoning technique.

Acknowledgments

We thank James Allan and Victor Lavrenko for their contribution to the problem formulation and to the construction of the information retrieval testbed.

This work was supported in part by the National Science Foundation under grants No. IRI-9624992, IRI-9634938, and INT-9612092, by the GanymedeII Project of Plan Etat/Nord-Pas-De-Calais, and by IUT de Lens.

References

[Dean and Boddy, 1988] Thomas Dean and Mark Boddy. An analysis of time-dependent planning. *Seventh National Conference on Artificial Intelligence*, 49–54, 1988.

[Gravey and Lesser, 1993] Alan Garvey and Victor Lesser. Design-to-time real-time scheduling. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(6):1491–1502, 1993.

[Greenwald and Dean, 1998] Lloyd Greenwald and Thomas Dean. A conditional scheduling approach to designing real-time systems. *AI Planning systems*, 1229–1234, 1998.

[Hansen and Zilberstein, 1996] Eric A. Hansen and Shlomo Zilberstein. Monitoring the progress of anytime problem-solving. *Thirteenth National Conference on Artificial Intelligence*, 1229–1234, 1996.

[Horvitz, 1988] Eric Horvitz. Reasoning under varying and uncertain resource constraints. *Seventh National Conference on Artificial Intelligence*, 111–116, 1988.

[Horvitz, 1997] Eric Horvitz. Models of continual computation. *Fourteenth National Conference on Artificial Intelligence*, 286–293, 1997.

[Liu *et al.*, 1991] J. Liu, K. Lin, W. Shih, A. Yu, J. Chung, and W. Zao. Algorithms for scheduling imprecise computations. *IEEE Transactions on Computers*, 24(5):58–68, 1991.

[Mouaddib, 1993] Abdel-Illah Mouaddib. *Contribution au raisonnement progressif et temps réel dans un univers multi-agents.* PhD thesis, University of Nancy I, (in French), 1993.

[Mouaddib and Zilberstein, 1997] Abdel-Illah Mouaddib and Shlomo Zilberstein. Handling duration uncertainty in meta-level control of progressive reasoning. *Fifteenth International Joint Conference on Artificial Intelligence*, 1201–1206, 1997.

[Mouaddib and Zilberstein, 1998] Abdel-Illah Mouaddib and Shlomo Zilberstein. Optimal scheduling of dynamic progressive processing. *Thirteenth Biennial European Conference on Artificial Intelligence*, 449–503, 1998.

[Russell and Wefald, 1991] Stuart Russell and Eric Wefald. *Do the Right Thing: Studies in Limited Rationality* MIT Press, 1991.

[Jones and Willett, 1997] Karen S. Jones and Peter Willett. (eds.) *Readings in Information Retrieval.* Morgan Kaufmann Publishers, 1997.

[Zilberstein and Russell, 1996] Shlomo Zilberstein and Stuart Russell. Optimal composition of real-time systems. *Artificial Intelligence* 82(1-2):181–213, 1996.