

Iterative-Refinement for Action Timing Discretization

Todd W. Neller

Department of Computer Science
Gettysburg College
Gettysburg, PA 17325, USA
tneller@gettysburg.edu

Abstract

Artificial Intelligence search algorithms search discrete systems. To apply such algorithms to continuous systems, such systems must first be discretized, i.e. approximated as discrete systems. Action-based discretization requires that both action parameters and action timing be discretized. We focus on the problem of action timing discretization.

After describing an ϵ -admissible variant of Korf's recursive best-first search (ϵ -RBFS), we introduce iterative-refinement ϵ -admissible recursive best-first search (IR ϵ -RBFS) which offers significantly better performance for initial time delays between search states over several orders of magnitude. Lack of knowledge of a good time discretization is compensated for by knowledge of a suitable solution cost upper bound.

Introduction

Artificial Intelligence search algorithms search discrete systems, yet we live and reason in a continuous world. Continuous systems must first be discretized, i.e. approximated as discrete systems, to apply such algorithms. There are two common ways that continuous search problems are discretized: state-based discretization and action-based discretization. State-based discretization (Latombe 1991) becomes infeasible when the state space is highly dimensional. Action-based discretization becomes infeasible when there are too many degrees of freedom. Interestingly, biological high-degree-of-freedom systems are often governed by a much smaller collection of motor primitives (Mataric 2000). We focus here on action-based discretization.

Action-based discretization consists of two parts: (1) action parameter discretization and (2) action timing discretization, i.e. *how* and *when* to act. See Figure 1. For example, consider robot soccer. Search can only sample action parameter continua such as kick force and angle. Similarly, search can only sample infinite possible action timings such as when to kick. The most popular form of discretization is uniform discretization. It is common to sample possible actions and action timings at fixed intervals.

In this paper, we focus on action timing discretization. Experimental evidence of this paper and previous studies (Neller 2000) suggests that a fixed uniform discretization of time is not advisable for search if one has a desired

Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

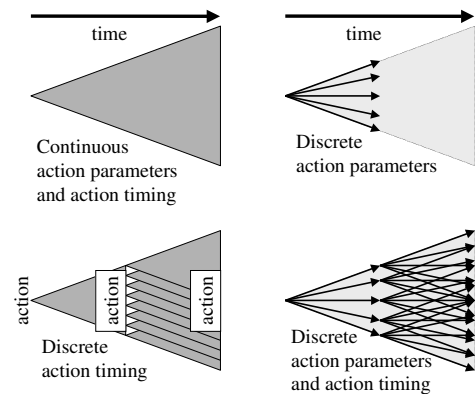


Figure 1: Action-based discretization.

solution cost upper bound. Rather, a new class of algorithms that dynamically adjust action timing discretization can yield significant performance improvements over static action timing discretization.

Iterative-refinement algorithms use a simple means of dynamically adjusting the time interval between search states. This paper presents the results of an empirical study of the performance of different search algorithms as one varies the initial time interval between search states. We formalize our generalization of search, present our chosen class of problems, describe the algorithms compared, and present the experimental results.

The key contributions of this work are experimental insight into the importance of searching with dynamic time discretization, and two new iterative-refinement algorithms, one of which exceeds the performance of ϵ -RBFS across more than four orders of magnitude of the initial time delay between states.

Search Problem Generalization

Henceforth, we will assume that the action discretization, i.e. which action parameters are sampled, is already given. From the perspective of the search algorithm, the action

discretization is static, i.e. cannot be varied by the algorithm. However, action timing discretization is dynamic, i.e. the search algorithm can vary the action timing discretization. For this reason, we will call such searches ‘‘SADAT searches’’ as they have Static Action and Dynamic Action Timing discretization.

We formalize the SADAT search problem as the quadruple:

$$\{S, s_0, A, G\}$$

where

- S is the state space,
- $s_0 \in S$ is the initial state,
- $A = \{a_1, \dots, a_n\}$ is a finite set of action functions $a_i : S \times \mathbb{R}^+ \rightarrow S \times \mathbb{R}$, mapping a state and a positive time duration to a successor state and a transition cost, and
- $G \subset S$ is the set of goal states.

The important difference between this and classical search formulations is the generalization of actions (i.e. operators). Rather than mapping a state to a new state and the associated cost of the action, we additionally take a time duration parameter specifying how much time passes between the state and its successor.

A goal path can be specified as a sequence of action-duration pairs that evolve the initial state to a goal state. The cost of a path is the sum of all transition costs. Given this generalization, the state space is generally infinite, and the optimal path is generally only approximable through a sampling of possible paths through the state space.

Sphere Navigation Problem

Since SADAT search algorithms will generally only be able to approximate optimal solutions, it is helpful to test them on problems with known optimal solutions. Richard Korf proposed the problem of navigation between two points on the surface of a sphere as a simple benchmark with a known optimal solution.¹ Our version of the problem is given here.

The shortest path between two points on a sphere is along the great-circle path. Consider the circle formed by the intersection of a sphere and a plane through two points on the surface of the sphere and the center of the sphere. The *great-circle path* between the two points is the shorter part of this circle between the two points. The *great-circle distance* is the length of this path.

The state space S is the set of all positions and headings on the surface of a unit sphere along with all nonnegative time durations for travel. Essentially, we encode path cost (i.e. time) in the state to facilitate definition of G . The initial state s_0 is arbitrarily chosen to have position (1,0,0) and velocity (0,1,0) in spherical coordinates, with no time elapsed initially.

The action $a_i \in A$, $0 \leq i \leq 7$ takes a state and time duration, and returns a new state and the same time duration (i.e. cost = time). The new state is the result of changing the heading $i * \pi/4$ radians and traveling with unit velocity at that heading for the given time duration on the surface of the

unit sphere. If the position reaches a goal state, the system stops evolving (and incurring cost).

The set of goal states G includes all states that are both (1) within ϵ_d great-circle distance from a given position p_g , and (2) within ϵ_t time units of the optimal great-circle duration to reach such positions. Put differently, the first requirement defines the size and location of the destination, and the second requirement defines how directly the destination must be reached. Position p_g is chosen at random from all possible positions on the unit sphere with all positions being equiprobable.

If d is the great-circle distance between (1,0,0) and p_g , then the optimal time to reach a goal position at unit velocity is $d - \epsilon_d$. Then the solution cost upper bound is $d - \epsilon_d + \epsilon_t$. For any position, the great-circle distance between that position and p_g minus ϵ_d is the optimal time to goal at unit velocity. This is used as the admissible heuristic function h for all heuristic search.

Algorithms

In this section we describe the four algorithms used in our experiments. The first pair use fixed time intervals between states. The second pair dynamically refine time intervals between states. The first algorithm, ϵ -admissible iterative-deepening A^* , features an improvement over the standard description. Following that we describe an ϵ -admissible variant of recursive best-first search and two novel iterative-refinement algorithms.

ϵ -Admissible Iterative-Deepening A^*

ϵ -admissible iterative-deepening A^* search, here called ϵ -IDA*, is a version of IDA* (Korf 1985) where the f -cost limit is increased ‘‘by a fixed amount ϵ on each iteration, so that the total number of iterations is proportional to $1/\epsilon$. This can reduce the search cost, at the expense of returning solutions that can be worse than optimal by at most ϵ .’’ (Russell & Norvig 1995).

Actually, our implementation is an improvement on ϵ -IDA* as described above. If Δf is the difference between (1) the minimum f -value of all nodes beyond the current search contour, and (2) the current f -cost limit, then the f -cost limit is increased by $\Delta f + \epsilon$. (Δf is the increase that would occur in IDA*.) This improvement is significant in cases where f -cost limit changes between iterations can significantly exceed ϵ .

To make this point concrete, suppose the current iteration of ϵ -IDA* has an f -cost limit of 1.0 and returns no solution and a new f -cost limit of 2.0. The new f -cost limit is the minimum heuristic f -value of all nodes beyond the current search contour. Let us further assume that ϵ is 0.1. Then increasing the f -cost limit by this fixed ϵ will result in the useless search of the same contour for 9 more iterations before the new node(s) beyond the contour are searched.

It is important to note that when we commit to an action timing discretization, the ϵ -admissibility of search is relative to the optimal solution of this discretization rather than the optimal solution of the original continuous-time SADAT search problem.

¹Personal communication, 23 May 2001.

Much work has been done in discrete search to tradeoff solution optimality for speed. Weighted evaluation functions (e.g. $f(n) = (1 - \omega)g(n) + \omega h(n)$, $0 \leq \omega \leq 1$ or $f(n) = g(n) + Wh(n)$, $W = \omega/(1 - \omega)$) (Pohl 1970; Korf 1993) provide a simple means to find solutions that are suboptimal by no more than a multiplicative factor of ω . For a good comparison of IDA*-styles searches, see (Wah & Shang 1995). For approximation of search trees to exploit phase transitions with a constant relative solution error, see (Pemberton & Zhang 1996).

ϵ -Admissible Recursive Best-First Search

ϵ -admissible recursive best-first search, here called ϵ -RBFS, is an ϵ -admissible variant of recursive best-first search that follows the description of (Korf 1993, §7.3) without further search after a solution is found. As with our implementation of ϵ -IDA*, local search bounds increase by at least ϵ (when not limited by B) to reduce redundant search.

In Korf's style of pseudocode, ϵ -RBFS is as follows:

```
eRBFS (node: N, value: F(N), bound: B)
IF f(N) > B, RETURN f(N)
IF N is a goal, EXIT algorithm
IF N has no children, RETURN infinity
FOR each child Ni of N,
  IF f(N) < F(N), F[i] := MAX(F(N), f(Ni))
  ELSE F[i] := f(Ni)
sort Ni and F[i] in increasing order of F[i]
IF only one child, F[2] := infinity
WHILE (F[1] <= B and F[1] < infinity)
  F[1] := eRBFS(N1, F[1],
                MIN(B, F[2] + epsilon))
  insert Ni and F[1] in sorted order
RETURN F[1]
```

The difference between RBFS and ϵ -RBFS is in the computation of the bound for the recursive call. In RBFS, this is computed as $\text{MIN}(B, F[2])$ whereas in ϵ -RBFS, this is computed as $\text{MIN}(B, F[2] + \text{epsilon})$. $F[1]$ and $F[2]$ are the lowest and second-lowest stored costs of the children, respectively. A correctness proof of ϵ -RBFS is described in the Appendix.

The algorithm's initial call parameters are the root node r , $f(r)$, and ∞ . Actually, both RBFS and ϵ -RBFS can be given a finite bound b if one wishes to restrict search for solutions with a cost of no greater than b and uses an admissible heuristic function. If no solution is found, the algorithm will return the f -value of the minimum open search node beyond the search contour of b .

In the context of SADAT search problems, both ϵ -IDA* and ϵ -RBFS assume a fixed time interval between a node and its child. The following two algorithms do not.

Iterative-Refinement ϵ -RBFS

Iterative-refinement (Neller 2000) is perhaps best described in comparison to iterative-deepening. Iterative-deepening depth-first search (Figure 2(a)) provides both the linear memory complexity benefit of depth-first search and the minimum-length solution-path benefit of breadth-first search at the cost of node re-expansion. Such re-expansion

costs are generally dominated by the cost of the final iteration because of the exponential nature of search time complexity.

Iterative-refinement depth-first search (Figure 2(b)) can be likened to an iterative-deepening search to a fixed time-horizon. In classical search problems, time is not an issue. Actions lead from states to other states. When we generalize such problems to include time, we then have the choice of how much time passes between search states. Assuming that the vertical time interval in Figure 2(b) is Δt , we perform successive searches with delays Δt , $\Delta t/2$, $\Delta t/3$, ... until a goal path is found.

Iterative-deepening addresses our lack of knowledge concerning the proper depth of search. Similarly, iterative-refinement addresses our lack of knowledge concerning the proper time discretization of search. Iterative-deepening performs successive searches that grow exponentially in time complexity. The complexity of previous unsuccessful iterations is generally dominated by that of the final successful iteration. The same is true for iterative-refinement.

However, the concept of iterative-refinement is not limited to the use of depth-first search. Other algorithms such as ϵ -RBFS may be used as well. In general, for each iteration of an iterative-refinement search, a level of (perhaps adaptive) time-discretization granularity is chosen for search and an upper bound on the solution cost is given. If the iteration finds a solution within this cost bound, the algorithm terminates with success. Otherwise, a finer level of time-discretization granularity is chosen, and search is repeated. Search is successively refined with respect to time granularity until a solution is found.

Iterative-Refinement ϵ -RBFS is one instance of such search. The algorithm can be simply described as follows:

```
IReRBFS (node: N, bound: B, initDelay: DT)
FOR I = 1 to infinity
  Fix the time delay between states at DT/I
  eRBFS(N, f(N), B)
  IF eRBFS exited with success,
    EXIT algorithm
```

Iterative-Refinement ϵ -RBFS does not search to a fixed time-horizon. Rather, each iteration searches within a search contour bounded by B . Successive iterations search to the same bound, but with finer temporal detail. DT/I is assigned to a global variable governing the time interval between successive states in search.

Iterative-Refinement DFS

The algorithm for Iterative-Refinement DFS is given as follows:

```
IRDFS (node: N, bound: B, initDelay: DT)
FOR I = 1 to infinity
  Fix the time delay between states at DT/I
  DFS-NOUB(N, f(N), B)
  IF DFS-NOUB exited with success,
    EXIT algorithm
```

Our depth-first search implementation DFS-NOUB uses a node ordering (NO) heuristic and has a path cost upper-bound (UB). The node-ordering heuristic is as usual: Nodes




Figure 2: Iterative search methods.

are expanded in increasing order of f -value. Nodes are not expanded that exceed a given cost upper bound. Assuming admissibility of the heuristic function h , no solutions within the cost upper-bound will be pruned from search.

Experimental Results

In these experiments, we vary only the initial time delay Δt between search states and observe the performance of the algorithms we have described. For ϵ -IDA* and ϵ -RBFS, the initial Δt is the only Δt for search. The iterative-refinement algorithms search using the harmonic refinement sequence $\Delta t, \Delta t/2, \Delta t/3, \dots$, and are limited to 1000 refinement iterations. ϵ -admissible searches were performed with $\epsilon = .1$.

Experimental results for success rates of search are summarized in Figure 3. Each point represents 500 trials over a fixed, random set of sphere navigation problems with $\epsilon_d = .0001$ and ϵ_t computed as 10% of the optimal time. Thus, the target size for each problem is the same, but the varying requirement for solution quality means that different delays will be appropriate for different search problems. Search was terminated after 10 seconds, so the success rate is the fraction of time a solution was found within the allotted time and refinement iterations.

In this empirical study, means and 90% confidence intervals for the means were computed with 10000 bootstrap resamples.

Let us first compare the performance of iterative-refinement (IR) ϵ -RBFS and ϵ -RBFS. To the left of the graph, where the initial Δt_0 is small, the two algorithms have identical behavior. This region of the graph indicates conditions under which a solution is found within 10 seconds on the first iteration or not at all. There is no iterative-refinement in this region; the time complexity of the first

iteration leaves no time for another.

At about $\Delta t_0 = .1$, we observe that IR ϵ -RBFS begins to have a significantly greater success rate than ϵ -RBFS. At this point, the time complexity of search allows for multiple iterations, and thus we begin to see the benefits of iterative-refinement.

Continuing to the right with greater initial Δt_0 , IR ϵ -RBFS nears a 100% success rate. At this point, the distribution of Δt 's over different iterations allows IR ϵ -RBFS to reliably find a solution within the time constraints. We can see the distribution of Δt 's that most likely yield solutions from the behavior of ϵ -RBFS.

Where the success rate of IR ϵ -RBFS begins to fall, the distribution of first 1000 Δt 's begins to fall outside of the region where solutions can be found. With our refinement limit of 1000, the last iteration uses a minimal $\Delta t = \Delta t_0/1000$. The highest Δt_0 trials fail not because time runs out. Rather, the iteration limit is reached. However, even with a greater refinement limit, we would eventually reach a Δt_0 where the iterative search cost incurred on the way to the good Δt range would exceed 10 seconds.

Comparing IR ϵ -RBFS with IR DFS, we first note that there is little difference between the two for large Δt_0 . For $3.16 \leq \Delta t_0 \leq 100$, the two algorithms are almost always able to perform complete searches of the same search contours through all iterations up to the first iteration with a solution path. The largest statistical difference occurs at $\Delta t_0 = 316$ where IR DFS's success rate is 3.8% higher. We note that our implementation of IR DFS has a faster node-expansion rate, and that ϵ -RBFS's ϵ -admissibility necessitates significant node re-expansion. For these Δt_0 's, the use of IR DFS trades off ϵ -optimality for speed and a slightly higher success rate.

For mid-to-low-range Δt_0 values, however, we begin to see the efficiency of ϵ -RBFS over DFS with node ordering




Figure 3: Effect of varying initial Δt .

as the first iteration with a solution path presents a more computationally costly search. Since the target destination is so small, the route that actually leads through the target destination is not necessarily the most direct route. Without a perfect heuristic where complex search is necessary, ϵ -RBFS shows its strength relative to DFS. Rarely will problems be so unconstrained and offer such an easy heuristic as this benchmark problem, so IR ϵ -RBFS will be generally be better suited for all but the simplest search problems.

Comparing IR ϵ -RBFS with ϵ -IDA*, we note that ϵ -IDA* performs relatively poorly over all Δt_0 . What is particularly interesting is the performance of ϵ -IDA* over the range where IR ϵ -RBFS behaves as ϵ -RBFS, i.e. where no iterative-refinement takes place. Here we have empirical confirmation of the significant efficiency of ϵ -RBFS over ϵ -IDA*.

In summary, iterative-refinement algorithms are statistically the same as or superior to the other searches over the range of Δt_0 values tested. IR ϵ -RBFS offers the greatest average success rate across all Δt_0 . With respect to ϵ -RBFS, IR ϵ -RBFS offers significantly better performance for Δt_0 spanning more than four orders of magnitude. These findings are in agreement with previous empirical studies concerning a submarine detection avoidance problem (Neller 2000).

This is significant for search problems where reasonable values for Δt are unknown. This is also significant for

search problems where reasonable values for Δt are known and one wishes to find a solution more quickly and reliably. This performance comes at a reasonable price for many applications. Lack of knowledge of a good time discretization is compensated for by knowledge of a suitable solution cost upper bound.

Conclusions

This empirical study concerning sphere navigation provides insight into the importance of searching with dynamic time discretization. Iterative-refinement algorithms are given an initial time delay Δt_0 between search states and a solution cost upper bound. Such algorithms iteratively search to this bound with successively smaller Δt until a solution is found.

Iterative-refinement ϵ -admissible recursive best-first search (IR ϵ -RBFS) was shown to be similar to or superior to all other searches studied for Δt_0 spanning over five orders of magnitude. With respect to ϵ -RBFS (without iterative-refinement), a new ϵ -admissible variant of Korf's recursive best-first search, IR ϵ -RBFS offers significantly better performance for Δt_0 spanning over four orders of magnitude.

Iterative-refinement algorithms are important for search problems where reasonable values for Δt are (1) unknown or (2) known and one wishes to find a solution more quickly and reliably. The key tradeoff is that of knowledge. Lack of knowledge of a good time discretization is compensated for by knowledge of a suitable solution cost upper bound. If one

knows a suitable solution cost upper bound for a problem where continuous time is relevant, an iterative-refinement algorithm such as IR ϵ -RBFS is recommended.

Future Work

The reason that our iterative-refinement algorithms made use of a harmonic Δt refinement sequence (i.e. Δt , $\Delta t/2$, $\Delta t/3$, ...) was to facilitate comparison to iterative-deepening. It would be interesting to see the performance of different Δt refinement sequences. For example, a geometric refinement sequence Δt , $c\Delta t$, $c^2\Delta t$, ... with $0 < c < 1$ would yield a uniform distribution of Δt 's on the logarithmic scale.

Even more interesting would be a machine learning approach to the problem in which a mapping was learned between problem initial conditions and Δt refinement sequences expected to maximize the utility of search. The process could be viewed as an optimization of searches over Δt . Assuming that both time and the success of search have known utilities, one would want to choose the next Δt so as to maximize expected success in minimal time across future iterations.

Acknowledgements

The author is grateful to Richard Korf for suggesting the sphere navigation problem, and to the anonymous reviewers for good insight and suggestions. This research was done both at the Stanford Knowledge Systems Laboratory with support by NASA Grant NAG2-1337, and at Gettysburg College.

Appendix: ϵ -RBFS Proof of Correctness

Proof of the correctness of ϵ -RBFS is very similar to the proof of the correctness of RBFS in (Korf 1993, pp. 52–57). For brevity, we here include the changes necessary to make the correctness proof of (Korf 1993) applicable to ϵ -RBFS. It will be necessary for the reader to have the proof available to follow these changes.

Lemma 4.1 *All calls to ϵ -RBFS are of the form ϵ RBFS($n, F(n), b$), where $F(n) \leq b$.*

Substitute “ ϵ RBFS” for “RBFS” through all proofs. For the second to last sentence of this lemma proof, substitute: “Thus, $F[1] < F[2] + \epsilon$. Thus, $F[1] \leq \min(b, F[2] + \epsilon)$.”

Lemma 4.2 *If b is finite, and $T(n, b)$ does not contain an interior goal node, then ϵ RBFS($n, F(n), b$) explores $T(n, b)$ and returns $MF(n, b)$.*

In the induction step's first and fourth paragraphs, substitute “ $\min(b, F[2] + \epsilon)$ ” for “ $\min(b, F[2])$ ”. In the last sentence of induction step paragraph two, the assumption of “no infinitely increasing cost sequences” is not necessary because the $F[2] + \epsilon$ term forces a minimum increment of ϵ while less than b .

Lemma 4.3 *For all calls ϵ RBFS($n, F(n), b$), $F(n) \leq OD(n)$ and $b \leq ON(n) + \epsilon$.*

Note the addition of “ $+$ ϵ ” to the lemma and make a similar addition everywhere a bound is compared to an ON term. For the third sentence of the second to last paragraph, substitute “Since $b' = \min(b, F[2] + \epsilon)$, then $b' \leq F[2] + \epsilon$. Because $F[2] \leq OD(n/2)$ and nodes are sorted by F value, $b' \leq OD(m') + \epsilon$ for all siblings m' of n' .”

Lemma 4.4 *When a node is expanded by ϵ -RBFS, its f value does not exceed the f values of all open nodes at the time by more than ϵ .*

Note the lemma change. “ v ” in the second paragraph is the “value” parameter. Wherever “ $ON(n)$ ” occurs, substitute “ $ON(n) + \epsilon$ ”. In the last sentence, substitute “... $f(n)$ does not exceed the f values of all open nodes in the tree when n is expanded by more than ϵ .”

Theorem 4.5 *ϵ RBFS($r, f(r), \infty$) will perform a complete ϵ -admissible search of the tree rooted at node r , exiting after finding the first goal node chosen for expansion.*

For the first sentence, substitute “Lemma 4.4 shows that ϵ -RBFS performs an ϵ -admissible search.” In the second to last sentence, substitute “Since the upper bound on each of these calls is the next lowest F value plus ϵ , the upper bounds must also increase continually, ...”.

References

- Korf, R. E. 1985. Depth-first iterative-deepening: an optimal admissible tree search. *Artificial Intelligence* 27(1):97–109.
- Korf, R. E. 1993. Linear-space best-first search. *Artificial Intelligence* 62:41–78.
- Latombe, J.-C. 1991. *Robot Motion Planning*. Boston, MA, USA: Kluwer Academic Publishers.
- Mataric, M. J. 2000. Sensory-motor primitives as a basis for imitation: Linking perception to action and biology to robotics. In Nehaniv, C., and Dautenhahn, K., eds., *Imitation in Animals and Artifacts*. Cambridge, MA, USA: MIT Press. See also USC technical report IRIS-99-377.
- Neller, T. W. 2000. *Simulation-Based Search for Hybrid System Control and Analysis*. Ph.D. Dissertation, Stanford University, Palo Alto, California, USA. Available as Stanford Knowledge Systems Laboratory technical report KSL-00-15 at www.ksl.stanford.edu.
- Pemberton, J. C., and Zhang, W. 1996. Epsilon-transformation: Exploiting phase transitions to solve combinatorial optimization problems. *Artificial Intelligence* 81(1–2):297–325.
- Pohl, I. 1970. Heuristic search viewed as path finding in a graph. *Artificial Intelligence* 1:193–204.
- Russell, S., and Norvig, P. 1995. *Artificial Intelligence: a modern approach*. Upper Saddle River, NJ, USA: Prentice Hall.
- Wah, B. W., and Shang, Y. 1995. Comparison and evaluation of a class of IDA* algorithms. *Int'l Journal of Tools with Artificial Intelligence* 3(4):493–523.