Demonstration Abstract: New Developments in Real-Time Heuristic Search

Carlos Hernández
Depto. de Ingeniería Informática
Univ. Católica de la Ssma. Concepción
Concepción, Chile
chernan@ucsc.cl

Jorge A. Baier
Depto. de Ciencia de la Computación
Pontificia Univ. Católica de Chile
Santiago, Chile
jabaier@ing.puc.cl

Tansel Uras **Sven Koenig**
Department of Computer Science
University of Southern California
Los Angeles, USA
{turas,skoenig}@usc.edu

Abstract

Our demonstration consists of a poster, videos and interactive simulations of real-time heuristic search algorithms for goal-directed navigation on a priori completely or partially unknown grids. It provides a brief introduction to real-time heuristic search by describing LSS-LRTA* and RTAA*. It then illustrates a performance issue of LSS-LRTA* and RTAA* due to depressions in the h-value surface. It describes the real-time heuristic search algorithms aLSS-LRTA*, daLSS-LRTA*, aRTAA*, and daRTAA*—which address this issue—and summarizes their properties. Our demonstration also illustrates a performance issue of LSS-LRTA* and RTAA* due to performing repeated A* searches around the current cells of the agent. It describes RTBA* and TBAA*, two real-time heuristic search algorithms that address this issue, and summarizes their properties.

Motivation

Many applications require agents to act quickly in a priori completely or partially unknown environments. Examples range from autonomous cars to video games, where companies impose time limits on the order of 1 millisecond (ms) for path planning (Bulitko et al. 2011). Such time limits are insufficient for finding complete paths, and an agent thus needs to move before it has found a complete path. We use goal-directed navigation on a priori completely or partially unknown grids with blocked and unblocked cells as running example. The agent always observes the blockage status of its eight neighboring cells and has to move from a given start cell to a given goal cell by repeatedly moving to an unblocked neighboring cell. Performance is measured by the number of moves before the agent reaches the goal cell. We study an agent that uses real-time heuristic search algorithms (Korf 1990) to determine its moves. Real-time heuristic search algorithms interleave A* searches (Hart, Nilsson, and Raphael 1968) with moves. We assume that the reader is familiar with A* and the associated terminology and give two examples of real-time heuristic search algorithms for goal-directed navigation on a priori completely or partially unknown grids in the following, both of which are variants of LRTA* (Korf 1990):

- LSS-LRTA* (Koenig and Sun 2009) assumes that cells with unknown blockage status are unblocked (Zelinsky 1992; Koenig, Tovey, and Smirnov 2003). This free-space assumption allows LSS-LRTA* to perform an A* search from the current cell of the agent to the goal cell until the goal cell is about to be expanded, the open list becomes empty or a given number of cells have been expanded. If the open list becomes empty, the agent stops unsuccessfully. The states in the closed list form the local search space (LSS). LSS-LRTA* sets the h-values of all states in the closed list to the largest possible h-values that maintain the consistency of all h-values. The agent then moves along the shortest path from its current cell to a cell with the smallest f-value found by the A* search and remembers all blocked cells that it observes. If it reaches the goal cell, it stops successfully. If it observes a blocked cell on the current path or reaches the end of the path, it repeats the process.
- RTAA* (Koenig and Likhachev 2006b) is almost identical to LSS-LRTA*; the difference is that it updates the h-values faster than LSS-LRTA*. RTAA* often outperforms LSS-LRTA* even though the h-values of RTAA* are typically not as informed as the ones of LSS-LRTA*. However, this is often compensated for by RTAA* being able to expand more cells within a given time limit (Koenig and Likhachev 2006b; Hernández and Baier 2012).

In the following, we discuss new developments in real-time heuristic search that address two drawbacks of existing real-time heuristic search algorithms such as LSS-LRTA* and RTAA*, namely poor performance due to depressions in the h-value surface and due to performing repeated A* searches around the current cells of the agent. We explain these problems, sketch new real-time heuristic search algorithms that address them and describe their properties.

Heuristic Depressions

LSS-LRTA* and RTAA* have a performance issue due to heuristic depressions, that is, valleys in the h-value surface (Ishida 1992). We focus on cost-sensitive heuristic depressions (Hernández and Baier 2012), that is, connected cells in a set D that are completely and immediately surrounded by cells in a set F such that $h(s) < d(s, s') + h(s')$ for all cells

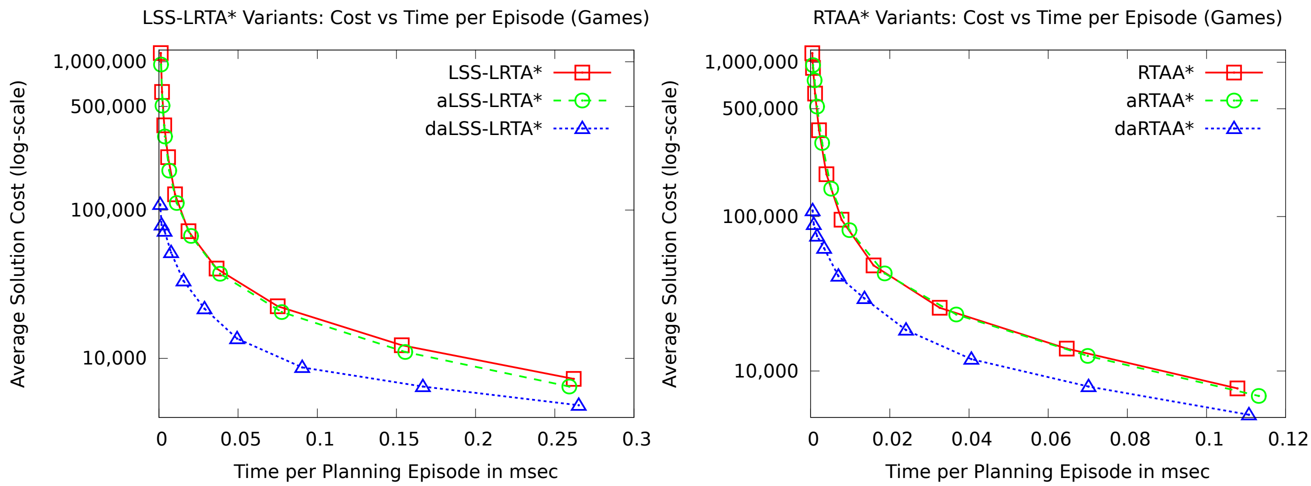

Figure 1: Evaluation of aLSS-LRTA*/daLSS-LRTA* and aRTAA*/daRTAA* on A Priori Completely Unknown Grids

$s \in D$ and $s' \in F$, where $d(s, s')$ is the distance from s to s' and $h(s)$ and $h(s')$ are the h-values of s and s', respectively. When an agent enters such a heuristic depression, it often executes many moves before it leaves the heuristic depression again. We describe real-time heuristic search algorithms for goal-directed navigation in a priori completely or partially unknown grids that address this issue by moving the agent to avoid heuristic depressions, namely aLSS-LRTA*, daLSS-LRTA*, aRTAA* and daRTAA*:

- daLSS-LRTA* (Hernández and Baier 2012) is almost identical to LSS-LRTA* and daRTAA* (Hernández and Baier 2012) is almost identical to RTAA*; the only difference is that daLSS-LRTA* and daRTAA* find a shortest path from the current cell of the agent to a cell with the smallest f-value among all cells for which the h-values have changed the least (rather than to a cell with the smallest f-value). This moves the agent to avoid heuristic depressions for the following reason: Let $\Delta(s)$ be the difference between the length of the shortest path from cell s to the goal cell and the initial h-value of s. If s is a cell close to the border of depression D and s' is a cell more in the interior of D, $\Delta(s') \geq \Delta(s)$.
- aLSS-LRTA* (Hernández and Baier 2012) is almost identical to LSS-LRTA* and aRTAA* (Hernández and Baier 2012) is almost identical to RTAA*; the only difference is that aLSS-LRTA* and aRTAA* find a shortest path from the current cell of the agent to a cell with the smallest f-value among all cells for which the h-values have not changed. If such cells do not exist, aLSS-LRTA* and aRTAA* find a shortest path from the current cell of the agent to a cell with the smallest f-value, like LSS-LRTA* and RTAA*. This is a simpler way of moving the agent to avoid heuristic depressions.

We compared aLSS-LRTA*, daLSS-LRTA*, aRTAA* and daRTAA* with the real-time search algorithms LSS-LRTA* and RTAA*. Our results, shown in Figure 1, indicate that daLSS-LRTA* and daRTAA* outperform LSS-LRTA* and RTAA*, respectively, by one order of magnitude when the time per planning episode is small. Details are given in (Hernández and Baier 2012).

Local Searches

LSS-LRTA* and RTAA* have a performance issue due to performing local searches, that is, repeated A* searches around the current cells of the agent. We describe RTBA* and TBAA*, two real-time heuristic search algorithms for goal-directed navigation in a priori completely or partially unknown grids that address this issue in the context of the game time model. The game time model partitions time into time intervals of a given length of time. During each time interval, the agent is allowed to search for the given length of time and then execute a single move (or pass on the move). Performance is measured by the number of time intervals before the agent reaches the goal cell. This performance measure is more realistic for video games than the number of moves before the agent reaches the goal cell since agents in video games are not allowed to execute moves at arbitrarily high speeds.

TBA* (Björnsson, Bulitko, and Sturtevant 2009) is a real-time heuristic search algorithm for goal-directed navigation on a priori known grids that performs one global search, that is, an A* search around the start cell. It performs an A* search from the start cell to the goal cell until the goal cell is about to be expanded or the open list becomes empty. If the open list becomes empty, the agent stops unsuccessfully. At the end of each time interval, the agent makes one move along a path from its current cell to the cell with the smallest f-value found by the A* search, by either following the shortest path from the start cell to a cell with the smallest f-value (if its current cell is on this path) or by moving to the parent of its current cell in the A* search tree. If it reaches the goal cell, it stops successfully. TBA* often outperforms LSS-LRTA* and RTAA* since a global search increases the chances that the agent follows a short path from the start cell to the goal cell (Hernández et al. 2012). We describe

	RTAA*		**daRTAA***		**RTBA***		**TBAA***		**Repeated A***		**Adaptive A***		**D* Lite**	
Length of Time Intervals (ms)	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves
0.3	3,245	3,244	2,879	2,878	4,613	4,604	2,290	2,286	7,155	2,004	3,230	2,010	**2,203**	2,027
0.6	2,598	2,597	2,472	2,471	3,368	3,360	2,147	2,144	4,487	2,004	2,572	2,010	**2,090**	2,027
0.9	2,451	2,450	2,418	2,417	2,918	2,910	2,101	2,099	3,611	2,004	2,361	2,010	**2,062**	2,027
1.2	2,310	2,309	2,305	2,304	2,695	2,688	2,086	2,083	3,178	2,004	2,260	2,010	**2,051**	2,027
1.5	2,281	2,280	2,272	2,271	2,560	2,553	2,070	2,068	2,920	2,004	2,202	2,010	**2,045**	2,027

Table 1: Evaluation of RTBA* and TBAA* on A Priori Completely Unknown Grids

	RTAA*		**daRTAA***		**RTBA***		**TBAA***		**Repeated A***		**Adaptive A***		**D* Lite**	
Length of Time Intervals (ms)	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves	# Time Intervals	# Moves
0.3	2,694	2,693	2,460	2,459	2,734	2,730	**1,505**	1,504	6,324	1,409	2,430	1,399	1,659	1,418
0.6	2,039	2,038	1,863	1,862	2,037	2,034	**1,442**	1,441	3,812	1,409	1,875	1,399	1,532	1,418
0.9	1,840	1,839	1,779	1,778	1,860	1,857	**1,431**	1,430	2,979	1,409	1,695	1,399	1,490	1,418
1.2	1,707	1,706	1,643	1,642	1,726	1,724	**1,421**	1,420	2,564	1,409	1,608	1,399	1,470	1,418
1.5	1,620	1,619	1,642	1,641	1,668	1,666	**1,415**	1,414	2,316	1,409	1,556	1,399	1,458	1,418

Table 2: Evaluation of RTBA* and TBAA* on A Priori Partially Unknown Grids

two variants of TBA* for goal-directed navigation on a priori completely or partially unknown grids:

- RTBA* (Hernández et al. 2012) is almost identical to TBA*, the only difference is that, if the agent observes a blocked cell on the path from its current cell to the cell with the smallest f-value, RTBA* starts a new A* search from the current cell of the agent to the goal cell.
- TBAA* (Hernández et al. 2012) is almost identical to RTBA*, the only difference is that, like RTAA*, it sets the h-values of all states in the closed list to the largest possible h-values that maintain the consistency of all h-values before it starts a new A* search. To be precise, it actually defers an h-value update until the time when the h-value is needed during a future A* search to avoid computing those h-values that are not needed later. The h-value updates make the h-values more informed and thus focus future A* searches better.

We compared RTBA* and TBAA* with the real-time heuristic search algorithms RTAA* and daRTAA* as well as the heuristic search algorithms Repeated A*, Adaptive A* and D* Lite (using a different experimental setup from the previous section). Repeated A* performs a complete A* search from the current cell of the agent to the goal cell until the goal cell is about to be expanded or the open list becomes empty. If the open list becomes empty, the agent stops unsuccessfully. The agent then moves along the shortest path from its current cell to the goal cell and remembers all blocked cells that it observes. If it reaches the goal cell, it stops successfully. If it observes a blocked cell on the path, it repeats the process. Incremental heuristic search algorithms, such as Adaptive A* (Koenig and Likhachev 2006a) and D* Lite (Koenig and Likhachev 2002), are almost identical to Repeated A*, the difference is that they speed up the A* searches by using their experience with prior A* searches to speed up future ones. Adaptive A* performs A* searches from the current cell of the agent to the goal cell, while D* Lite performs searches in the opposite direction. Our results, shown in Tables 1 and 2, indicate that TBAA* has a slight performance advantage over D* Lite in a priori partially unknown grids and vice versa in a priori completely unknown grids, although the differences might not be statistically significant. In both cases, TBAA* has the advantage over D* Lite that the agent starts to move right away. Details are given in (Hernández et al. 2012).

Objectives of the Demonstration

Our demonstration consists of a poster, videos and interactive simulations of real-time heuristic search algorithms. It has the following objectives:

1. Our demonstration provides a brief introduction to real-time heuristic search by describing LSS-LRTA* and RTAA*.
2. Our demonstration illustrates a performance issue of LSS-LRTA* and RTAA* due to depressions in the h-value surface. It describes aLSS-LRTA*, daLSS-LRTA*, two real-time heuristic search algorithms that address this issue, and summarizes their properties described in (Hernández and Baier 2012).
3. Our demonstration illustrates a performance issue of LSS-LRTA* and RTAA* due to performing repeated A* searches around the current cells of the agent. It describes real-time heuristic search algorithms that address this issue, namely RTBA* and TBAA*, and summarizes their properties described in (Hernández et al. 2012).

Acknowledgments

This material is based upon research supported by NSF (while Sven Koenig was serving at NSF). It is also based upon research supported by ARL/ARO under contract/grant number W911NF-08-1-0468 and ONR in form of a MURI under contract/grant number N00014-09-1-1031. Jorge Baier was partly funded by Fondecyt grant number

11110321. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies or the U.S. government.

References

Björnsson, Y.; Bulitko, V.; and Sturtevant, N. 2009. TBA*: Time-bounded A*. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 431–436.

Bulitko, V.; Björnsson, Y.; Sturtevant, N.; and Lawrence, R. 2011. *Real-time Heuristic Search for Game Pathfinding*. Applied Research in Artificial Intelligence for Computer Games. Springer.

Hart, P.; Nilsson, N.; and Raphael, B. 1968. A formal basis for the heuristic determination of minimal cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4(2).

Hernández, C., and Baier, J. A. 2012. Avoiding and escaping depressions in real-time heuristic search. *Journal of Artificial Intelligence Research* 43:523–570.

Hernández, C.; Baier, J. A.; Uras, T.; and Koenig, S. 2012. Time-Bounded Adaptive A*. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

Ishida, T. 1992. Moving target search with intelligence. In *Proceedings of the 10th National Conference on Artificial Intelligence (AAAI)*, 525–532.

Koenig, S., and Likhachev, M. 2002. D* Lite. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI)*, 476–483.

Koenig, S., and Likhachev, M. 2006a. A new principle for incremental heuristic search: Theoretical results. In *Proceedings of the 16th International Conference on Automated Planning and Scheduling (ICAPS)*, 402–405.

Koenig, S., and Likhachev, M. 2006b. Real-Time Adaptive A*. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 281–288.

Koenig, S., and Sun, X. 2009. Comparing real-time and incremental heuristic search for real-time situated agents. *Autonomous Agents and Multi-Agent Systems* 18(3):313–341.

Koenig, S.; Tovey, C.; and Smirnov, Y. 2003. Performance bounds for planning in unknown terrain. *Artificial Intelligence* 147(1-2):253–279.

Korf, R. 1990. Real-time heuristic search. *Artificial Intelligence* 42(2-3):189–211.

Zelinsky, A. 1992. A mobile robot exploration algorithm. *IEEE Transactions on Robotics and Automation* 8(6):707–717.