

Bidirectional Best-First Search with Bounded Error: Summary of Results

Andreas L. Koll
Kaisers 13
A-6655 Steeg
Austria, Europe

Hermann Kaindl
Siemens AG Osterreich
Geusaugasse 17
A-1030 Wien
Austria, Europe

Abstract

For more than two decades there has been consensus that bidirectional heuristic search is afflicted by the problem of search wavefronts missing each other. However, our results indicate that a different problem appears to be of primary importance. The fronts typically meet rather early even without using wave-shaping techniques. Especially when aiming for optimal solutions, however, much effort has to be spent for subsequently improving the solution quality, and finally for proving that there is indeed no better solution possible. Therefore, only slightly relaxing the requirements on the solution quality already leads to strong improvements in efficiency.

We describe several new *e-admissible* bidirectional search algorithms which do *not* use wave-shaping techniques. The most efficient of these use a novel termination criterion designed to address the suspected primary problem of bidirectional heuristic search. We prove ϵ -admissibility and a *dominance* result based on this termination criterion. In summary, we show that and how bidirectional best-first search can be more efficient than the corresponding unidirectional counterpart without using computationally very demanding wave-shaping techniques.

Notation

s, t	Start node and goal node, respectively.
$f_s(n)$	Successors of node n in the problem graph.
$f_p(n)$	Parents of node n in the problem graph.
d	Current search direction index; when search is in the forward direction $d = 1$, and when in the backward direction $d = 2$.
d'	$3 - d$; it is the index of the direction opposite to the current search direction.
$c_i(m, n)$	Cost of the direct arc from m to n if $i = 1$, or from n to m if $i = 2$.
$k_i(m, n)$	Cost of an optimal path from m to n if $i = 1$, or from n to m if $i = 2$.
$g_i^*(n)$	Cost of an optimal path from s to n if $i = 1$, or from t to n if $i = 2$.
$h_i^*(n)$	Cost of an optimal path from n to t if $i = 1$,

* Andreas Koll was with Siemens AG Osterreich for a summer internship.

	or from n to s if $i = 2$.
$g_i(n), h_i(n)$	Estimates of $g_i^*(n)$ and $h_i^*(n)$, respectively.
h_F	FOCAL-heuristic.
$f_i(n)$	Static evaluation function.
C^*	Cost of an optimal path from s to t .
L_{min}	Cost of the best (least costly) complete path found so far from s to t .
$TREE_1$	The forward search tree.
$TREE_2$	The backward search tree.
$OPEN_i$	The set of open nodes in $TREE_i$.
$FOCAL_i$	$\{m \mid (m \in OPEN_i) \wedge (g_i(m) + h_i(m)) \leq (1 + \delta) \min_{n \in OPEN_i} (g_i(n) + h_i(n))\}$.
$ OPEN_i $	Number of nodes in $OPEN_i$.
$CLOSED_i$	The set of closed nodes in $TREE_i$.
$p_i(n)$	Parent of node n in $TREE_i$.
$p_d^i(n)$	$p_d(p_d \dots (p_d(n)) \dots)$.
Ω_m	^{i-times} The set of nodes in $OPEN_{d'}$ which are descendants of m in $TREE_{d'}$.
<i>MeetingNode</i>	Node where $TREE_1$ met $TREE_2$ and yielded the best complete path found so far.

1 Introduction and Background

Much research in heuristic search has been devoted to algorithms finding *optimal* solutions. In the unidirectional case A^* [4] and its derivatives received much attention. Especially related to NP-complete problems, an important approach in relaxing the search for optimality is to provide a bound for the worst case quality of a solution compared to an optimal one [5]. A search algorithm is said to be *e-admissible* if it guarantees that the costs of its solutions are bounded by $(1 + \epsilon)C^*$ [14, 13]. An early example is *dynamic weighting* in the form used in [18]. It uses an upper bound N on the depth of the deepest node to be evaluated, which has to be given as a parameter.¹ With $N \rightarrow \infty$, asymptotically the heuristic evaluation function with constant weighting of [15] arises ($f = g + wh$).² Such an evaluation function is used by EPA [16].³ A more recent approach to dynamic weight-

¹In fact, besides using an admissible heuristic component h the algorithm must have the possibility to reopen nodes, i.e., to move nodes back from CLOSED to OPEN whenever a new better path is found to them.

²For constant w , this is equivalent to $f = (1 - w')g + w'h$ with $w' = \frac{w}{1+w}$.

³For our comparisons, we use a slight variant of HPA, which includes the option to move nodes back from CLOSED

ing [7] compensates for the error made by the heuristic component using learnt weights, and it is ϵ -admissible without the need for a given upper bound like N . The ϵ -admissible algorithm A^* [14] is designed to allow for incorporating an additional heuristic to estimate the effort for finishing a search, which is used as FOCAL-heuristic. All these algorithms perform unidirectional search, and for $\epsilon > 0$ they are typically more efficient than A^* .

$BHPA^4$ [15, 17] was the first bidirectional algorithm using a heuristic evaluation function for finding optimal solutions. While BS^* [10] improved $BHPA$ technically, its performance was only nearly as good as the unidirectional A^* . There was consensus that the main reason for the superiority of unidirectional search is that the bidirectional heuristic search is afflicted by the problem of search wavefronts missing each other (cf. the *missile metaphor* in [15]).

Wave-shaping techniques like those used in $BHFFA$ [3], $BHFFA2$ [2] and d-node retargeting [19] showed that bidirectional heuristic search can be rather efficient, especially in terms of the number of expanded nodes. Davis [1] presented a generalized algorithm (containing $BHPA$ and $BHFFA2$) and analyzed it theoretically. However, these algorithms are either excessively computationally demanding, or they have no restriction on the solution quality.

Considering the missile metaphor of Pohl, the use of ϵ -admissible approaches in bidirectional search should be dangerous without wave-shaping, since the chance of search wavefronts missing each other could be even increased. However, we found that in bidirectional searches the first meeting of the fronts typically occurred early, compared to the effort for finally satisfying the termination condition.

Fig. 1 illustrates this phenomenon.⁵ These data are normalized in the sense that every data point represents the ratio of the number of nodes expanded by the corresponding algorithm to that expanded by A^* . In the average, the fronts meet rather early even without using wave-shaping techniques, requiring only slightly more than half the effort of complete A^* on the 8-Puzzle for $\epsilon = 0$. On the much more difficult 15-Puzzle, the fronts meet already when about 3 orders of magnitude fewer nodes are generated than by complete IDA^* [9] searches (data for A^* are unavailable here).

Especially when aiming for optimal solutions, however, much effort has to be spent for subsequently improving the solution quality. (In the average, this is achieved with less effort than that of a complete A^* search on the 8-Puzzle.) Finally, for proving that there is indeed no better solution possible, many more nodes have to be expanded. Therefore, we were interested in investigating ϵ -admissible bidirectional search and in comparing its efficiency to the corresponding unidirectional case.

to OPEN, if a new better g -value is found. Moreover, we include a check whether OPEN has become empty. Since this is like A^* [4], we call this variant HPA^* .

⁴The earlier version named $VGHA$ in [15] did not move nodes back from CLOSED to OPEN. Hence, it needed a *consistent* heuristic function for being *admissible*.

⁵For $\epsilon = 0$, IBS_ϵ^* — one of our new algorithms as introduced below — is identical to BS^* .




Figure 1: Phases of IBS_ϵ^* (on the 8-Puzzle).

2 Improved Bidirectional Search

In the following, we describe several ϵ -admissible bidirectional search algorithms which do *not* use wave-shaping techniques. We relate them to the earlier known algorithms in pointing out the differences. The origin is clearly $BHPA$ [15, 17], which can be intuitively viewed to consist of two HPA searches in opposing directions. The termination condition compares the length of the best path found so far to an optimistic estimate computed from the nodes in both search frontiers (see also below). While $BHPA$ only considered meeting points between CLOSED nodes from both sides, Kwa [10] showed that meeting of a newly generated node with any node in the opposing front is sufficient. All the algorithms described below use this improvement. This technical change together with the incorporation of checks for empty OPEN lists results in a version we call $BHPA'$.

2.1 Constant weighting ($IBHPA_\epsilon$ and IBS_ϵ^*)

There are two ways of achieving an ϵ -admissible variant of $BHPA'$:

- $BHPA_w$ using an evaluation function of the form $f = g + wh$, where f can be rewritten as $f = g + (1 + \epsilon)h$. The original termination condition of $BHPA$

$$L_{min} \leq \max[\min_{x \in OPEN_1} f_1(x), \min_{x \in OPEN_2} f_2(x)] \quad (1)$$

turns out to include $(1 + \epsilon)h$ implicitly in $BHPA_w$:

$$L_{min} \leq \max[\min_{x \in OPEN_1} (g_1(x) + (1 + \epsilon)h_1(x)), \min_{x \in OPEN_2} (g_2(x) + (1 + \epsilon)h_2(x))] \quad (2)$$

In fact, Pohl [17] proposed such a change in the evaluation function, when not concerned with finding a shortest path. For this reason, he suggested to simplify the termination condition to just meeting. In contrast, $BHPA_w$ using (1) or equivalently (2) as termination condition is ϵ -admissible when h is admissible.

- $BHPA_\epsilon$ using $f = g + h$ (like $BHPA$ or a bidirectional admissible version of A^*), but using the termination criterion

$$L_{min} \leq (1 + \epsilon) \max[\min_{x \in OPEN_1} (g_1(x) + h_1(x)), \min_{x \in OPEN_2} (g_2(x) + h_2(x))] \quad (3)$$

$BHPA_\epsilon$ is also ϵ -admissible when h is admissible.

These ideas can be combined in an improved version as follows: $IBHPA_\epsilon$ uses $f = g + (1 + \epsilon)h$ with $\epsilon \geq 0$ and the termination criterion (3) still guaranteeing e-admissibility.

Another approach is to modify BS^* [10] similarly, in order to utilize its technical improvements. When a node is selected for expansion which is already closed in the opposite search tree, it can just be closed without expansion (nipping). Moreover, it is possible in such a situation to remove the descendants of this node in the opposing open set (pruning). Likewise, it is possible to remove all those open nodes (in both directions) whose l -values are greater or equal than L_{min} (trimming). Of course, such newly generated nodes are not placed in the open sets at all (screening). The interested reader will find a pseudo-code formulation of these improvements in Appendix A.

Generally, the provision for reopening nodes must be incorporated, since otherwise e-admissibility cannot be guaranteed. In particular, without this provision there could be no guarantee that the shallowest open node n' on an optimal path always satisfies $g(n') = g^*(n')$. In addition, we have to assume a consistent h because of the BS^* specific actions nipping and pruning.⁶

BS^*_w using $f = g + wh$ with $w \geq 1$ like HPA (in step 14 of [10, p. 101]).⁷

- BS^*_t using $f = g + h$ and the following termination criterion (in step 3 of [10, p. 100]): **OPEN₁** or **OPEN₂** is empty or (3).

Again, these modifications can be combined, resulting in an improved algorithm we call IBS^* .

2.2 Derivatives of A^*_t ($BA^*_{\delta,\epsilon}$ and $BSA^*_{\delta,\epsilon}$)

Especially, when there is another heuristic function available that estimates the effort for completing a search, it can serve as a FoCAL-heuristic in the unidirectional A^*_t [14].⁸ Using A^*_t type of search in both directions and the termination condition (3), BA^*_t arises, which is e-admissible.

In contrast to the bidirectional algorithms $IBHPA_\epsilon$ and $IBSb^*$ described above, BA^*_t is normally less efficient than the corresponding unidirectional algorithm (A^*_t). We found out that its searches are too greedy, leading to comparably bad solutions at the first meeting and consequently to much effort for subsequent satisfaction of the termination condition guaranteeing e-admissible solutions. Generally, we conjecture that bidirectional searches without wave-shaping tend to meet first with a solution quality which is worse than that achieved by the respective evaluation function in a corresponding unidirectional search.

⁶Despite consistency, reopening of nodes is necessary because of the $(1 + \epsilon)$ -factor in the evaluation function.

⁷This variant is only of interest for showing the genesis of the improved algorithm below, since by itself it is very inefficient due to lacking an appropriate termination condition.

⁸Pearl and Kim propose to use the same heuristic for both purposes if no estimator for the effort of completion is available. We wanted to compare all the algorithms with the same domain-specific knowledge, and consequently used the available admissible heuristic h also as a FoCAL-heuristic.

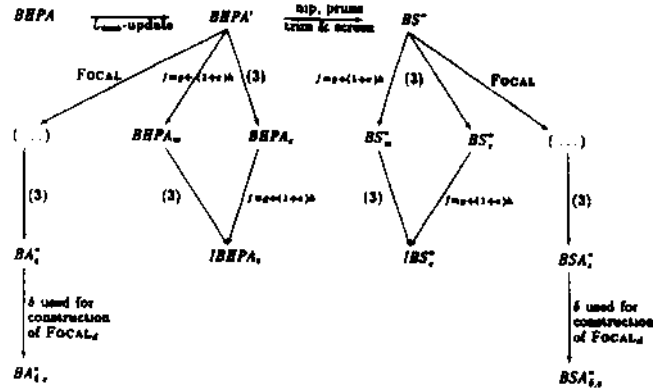


Figure 2: Overview of e-admissible bidirectional search algorithms.

Therefore, we found an improvement by adding another parameter δ . $BA^*_{\delta,\epsilon}$ uses δ for creating the FOCAL-lists, and ϵ for termination. Much better results were achieved, for instance, with $\delta = \frac{\epsilon}{2}$ and $\delta = \frac{\epsilon}{4}$.

Of course, it is also possible to incorporate the approach of A^*_t into the framework of BS^* . BSA^*_t arises from BS^*_t by the following changes:

- reopening of closed nodes if a new better path to them has been found,
- the termination condition (in step 3 of [10, p. 101]) analogously to BS^*_t (see above),
- the selection of a node m for expansion (in step 7 of [10, p. 100]) as usual for A^*_t .

Moreover, another parameter can be used, $BSA^*_{\delta,\epsilon}$ arises from BSA^*_t analogously to $BA^*_{\delta,\epsilon}$ from BA^*_t . Since this algorithm is already somewhat different from BS^* , we present a complete pseudo-code description in Appendix A.⁹

Fig. 2 provides an overview of all the presented algorithms and their genesis.

3 Theoretical Results

Due to lack of space we cannot present here all the proofs for the various algorithms introduced above (the interested reader is referred to [6]). Since the proofs for BSA^*_t and $BSA^*_{\delta,\epsilon}$ are the ones most different from similar proofs for A^*_t , $BHPA$ and BS^* , we selected them for presentation below. Moreover, we will show an interesting result about dominance between bidirectional algorithms which is due to the improved termination condition (3). We assume that the branching degrees of all nodes are finite, the arc costs $c_i(m, n) \geq \epsilon > 0$ for some c , and a consistent heuristic estimator is used.

Definition 3.1 A path P from s to t is ϵ -optimal, if $cost(P) < (1 + \epsilon)C^*$.

Definition 3.2 A search algorithm is said to be ϵ -admissible, if it terminates with an ϵ -optimal solution whenever a solution exists [14,13].

Definition 3.3 An algorithm A_1 is said to dominate A_2 if every node expanded by A_1 is also expanded by A_2 [13].

⁹Due to the inclusion of reopening of nodes in our algorithms and a technical bug in step 22 of the BS^* description in [10, p. 101], we restructured the code.

Lemma 3.1 Before $BSA_{\delta, \epsilon}^*$ ($\delta \leq \epsilon$) terminates either an ϵ -optimal path P' has already been found or there exist n_i and n_j in $OPEN_1$, $OPEN_2$ respectively, such that n_i and n_j are shallowest open nodes along an optimal path $P = (s = n_1, \dots, n_i, \dots, n_j, \dots, t = n_t)$ with $g_1(n_i) = g_1^*(n_i)$ and $g_2(n_j) = g_2^*(n_j)$.

Proof: Except for the BS^* specific actions such nodes must exist at any time before termination, since we account for re-opening of nodes (cf. [13]), and before the first meeting no such actions can occur. Furthermore, $i < j$. Otherwise the optimal path would have been found. There are four cases where these nodes could be removed from $OPEN_d$ without installing a successor substituting them along P . We consider the case of n_i to be removed from $OPEN_1$ (the case of n_j to be removed from $OPEN_2$ is analogous).

1. **Nipping:** Node n_i is nipped in $TREE_1$, i.e. n_i must be in $CLOSED_2$. n_k is a node in $OPEN_2$ having the smallest f_2 value. Since n_i is closed in $TREE_2$, n_i was selected for expansion before n_j (while searching in the backward direction) and we have

$$\begin{aligned} f_2(n_i) &\leq (1 + \delta)f_2(n_k) & n_i \text{ from } FOCAL_2 \\ &\leq (1 + \epsilon)f_2(n_k) & \delta \leq \epsilon \\ &\leq (1 + \epsilon)f_2(n_j) & f_2(n_k) \leq f_2(n_j) \end{aligned} \quad (4)$$

We distinguish two cases

$$\begin{aligned} (1.1) \quad (1 + \epsilon)[g_2^*(n_j) + k_2(n_i, n_j)] + \epsilon g_1^*(n_i) &\geq g_2(n_i) \\ \text{cost}(P') &= g_1^*(n_i) + g_2(n_i) \\ &\leq g_1^*(n_i) + \epsilon g_1^*(n_i) \\ &\quad + (1 + \epsilon)[g_2^*(n_j) + k_2(n_i, n_j)] \\ &\leq (1 + \epsilon)[g_1^*(n_i) + g_2^*(n_j) + k_2(n_i, n_j)] \\ &\leq (1 + \epsilon)C^* \end{aligned}$$

This implies $L_{min} \leq (1 + \epsilon)C^*$.

$$(1.2) \quad (1 + \epsilon)[g_2^*(n_j) + k_2(n_i, n_j)] + \epsilon g_1^*(n_i) < g_2(n_i)$$

Due to consistency $h_2(n_j) \leq k_2(n_i, n_j) + h_2(n_i)$.

$$\begin{aligned} g_2(n_i) + h_2(n_j) &\leq g_2(n_i) + k_2(n_i, n_j) + h_2(n_i) \\ (1 + \epsilon)(g_2(n_j) + h_2(n_j)) &\leq (1 + \epsilon)[g_2^*(n_j) + k_2(n_i, n_j)] \\ &\quad + (1 + \epsilon)h_2(n_i) \\ &\quad + \epsilon g_1^*(n_i) - \epsilon g_1^*(n_i) \\ (1 + \epsilon)f_2(n_j) &< \underbrace{g_2(n_i) + h_2(n_i)}_{\geq 0} \\ &\quad - \underbrace{\epsilon[g_1^*(n_i) - h_2(n_i)]}_{\geq 0} \\ &< f_2(n_i) \text{ contradicting (4)} \end{aligned}$$

2. **Pruning:** Node n_i has an ancestor node n_x in $TREE_1$ which was nipped in $TREE_2$. n_x must be on an optimal path (say P), otherwise n_i cannot have its parent pointer set to n_x . Therefore, $g_1(n_x) = g_1^*(n_x)$ and the proof is analogous to case 1, substituting n_x for n_i .

3. **Trimming:** If n_i would be trimmed from $OPEN_1$, then

$$g_1(n_i) + h_1(n_i) \geq L_{min} \quad (5)$$

Since n_i is on an optimal path and $g_1(n_i) = g_1^*(n_i)$, and since h_1 is consistent and consequently admissible, $h_1(n_i) \leq h_1^*(n_i)$, this implies $g_1(n_i) + h_1(n_i) \leq g_1^*(n_i) + h_1^*(n_i) = C^*$. If no path P' with $\text{cost}(P') \leq (1 + \epsilon)C^*$ was found yet, $C^* \leq (1 + \epsilon)C^* < L_{min}$, contradicting (5).

4. **Screening:** If n_i is expanded and not nipped, there is one of its successors n_{i+1} substituting it as the shallowest open node on P . If n_{i+1} would be screened and $C^* < L_{min}$, this would lead to an analogous contradiction as in case 3. \square

Lemma 3.2 If a path exists from s to t , $BSA_{\delta, \epsilon}^*$ will not terminate before finding an ϵ -optimal path.

Proof: The proof is analogous to the one in [17, p. 193], but using our termination condition (3). \square

Lemma 3.3 // a path exists from s to t , $BSA_{\delta, \epsilon}^*$ terminates.

Proof: Either the graph is finite and exhausted, or it is infinite and the costs become unbounded (analogous proofs can be found in [4, 12, IS, 10]). \square

Theorem 3.1 If a path exists from s to t , BSA_{6c}^* terminates with a solution whose cost does not exceed the optimal cost by more than a factor $(1 + \epsilon)$, i.e., $BSA_{\delta, \epsilon}^*$ is ϵ -admissible.

Proof: This result follows from the lemmas above. \square

Corollary 3.1 BSA_{ϵ}^* is ϵ -admissible.

Theorem 3.2 A bidirectional algorithm B_1 using the termination condition (S) dominates another bidirectional algorithm B_2 differing from B_1 only in using the condition (2) instead.

Proof: Since the right side of (S) is always greater than or equal to the right side of (2), algorithm B_1 cannot terminate later than B_2 . \square

Corollary 3.2 $IBHPA_{\epsilon}$ dominates $BHPA_w$.

4 Experimental Results

Experiments were conducted to compare the performance of these bidirectional algorithms. The test domains used were sliding-tile puzzles and route planning.¹⁰ These puzzles have been often used in the literature for explanations and comparisons of search algorithms in AI. Therefore, we also used them in order to make it easier to compare our results with published ones. Since finding optimal solutions is NP-complete, it is interesting to study finding approximate solutions, and the effects of scaling up.

However, while the puzzle is illustrative, it is just a game. In contrast, the route planning domain we selected has more real world flavor (though it is easier to scale up). We used randomly generated route planning problems in a map of the Viennese network of public transportation [6] rather than synthetically generated problems as used in [10].

4.1 Constant weighting

We were primarily interested in comparing the bidirectional best-first search algorithms with their corresponding unidirectional counterparts. Let us first summarize the 8-Puzzle results. The case of $\epsilon = 0$ is identical to a comparison of HPA^* with BS^* and $BHPA'$. There, the unidirectional HPA^* is clearly more efficient. However, already for $\epsilon \geq 0.2$ the bidirectional search algorithms are in terms of the number of generated nodes more efficient than the corresponding unidirectional algorithm. The linear-space algorithms $WIDA^*$ (an ϵ -admissible extension of IDA^*) and $RBFS$ [8] showed relatively bad performance in terms of generated nodes. However, since node generation and evaluation is very efficient for the puzzle, they are still competitive in terms of running time, due to avoiding the overhead for managing the lists used for implementing classical best-first search.

These domains are not especially selected to fulfill the conditions under which classical best-first search is best. In particular, the puzzles have uniform cost and the usual evaluation functions there do not have many distinct values.




Figure 3: Comparison on the 15-Puzzle (100 instances).

An important question is how significant the results are. In addition to comparing the numbers of nodes searched, we look at the empirical data from a different viewpoint. Instead of summing up the numbers of generated nodes for each sample, we compare the performance on each instance. We count how often one or the other algorithm "wins" in the sense that it expands fewer nodes than its competitor when solving the same instance. The statistical details of using the *sign test* can be found in [6]. Although it is statistically significant according to this test that for $\epsilon \leq 0.1$ HPA^* expands more often fewer nodes than IBS_ϵ^* , on the same instance, the inverse is true for $\epsilon \geq 0.3$.

Since the sliding-tile puzzle is NP-complete, it is interesting to investigate scaling up. In particular, it is interesting to see the effect of relaxing the solution quality [5]. Fig. 3 compares $IBHPA_\epsilon$ and IBS_ϵ^* with HPA^* , $WIDA^*$ and $RBFS$ on the 15-Puzzle (using the sample of 100 instances in [9]). Due to their memory problem, it is infeasible for $IBHPA_\epsilon$, IBS_ϵ^* and HPA^* to find or even to guarantee optimal solutions here.

For $\epsilon = 1$, HPA^* can solve all the 100 instances of [9], but it requires memory for nearly 1 million nodes in the maximum. In contrast, $IBHPA_\epsilon$ requires only a storage of slightly more than 300k nodes in maximum, and IBS^* slightly more than 250k nodes. Both are clearly more efficient than the others in terms of nodes generated. Interestingly, HPA^* "wins" against both $IBHPA_\epsilon$ and IBS_ϵ^* slightly more often, but the results are not statistically significant according to the sign test. While the worst case bound $\epsilon = 1$ would allow solutions with twice the cost of optimal ones, the solutions actually found by these best-first algorithms are in the average just about 20 percent worse.

For $\epsilon = 0.75$, $IBHPA_\epsilon$ and IBS_ϵ^* require storage of slightly more than 600k nodes in maximum. In contrast, HPA^* requires to store nearly 2 million nodes for this task. The advantage of our bidirectional algorithms over HPA^* in terms of the total number of nodes searched is highly statistically significant (testing the means), but it is not according to the sign test. From this we can conclude that they are better especially on the difficult problems. Generally, the advantage of $IBHPA_\epsilon$ and IBS_ϵ^* is even clearer than for $\epsilon = 1$, generating an order of mag-




Figure 4: Comparison on the 8-Puzzle (500 instances).

nitude fewer nodes than $WIDA^*$.

For $\epsilon = 0.5$, $IBHPA_\epsilon$ and IBS_ϵ^* require less than 1.5 million nodes in maximum for the whole sample. While this worst case bound would allow solutions that are 50 percent worse than optimal ones, in the average these algorithms find solutions that are about 10 percent worse. HPA^* cannot solve 8 instances with this amount of storage. The number of nodes generated by $IBHPA_\epsilon$ (IBS_ϵ^*) is in the average smaller by a factor of 21 (25) than that generated by $WIDA^*$. Therefore, it appears that the performance advantage of our bidirectional algorithms steadily increases with decreasing values of ϵ (at least in the range where memory suffices). Moreover, their advantage is much more pronounced on the much more difficult 15-Puzzle than on the smaller 8-Puzzle.

4.2 Derivatives of A_ϵ^*

In contrast, the bidirectional algorithms BA_ϵ^* and BSA_ϵ^* are normally less efficient than the corresponding unidirectional A_ϵ^* (see Fig. 4 for the results on a random set of 500 8-Puzzles).¹¹ Therefore, we designed $BA_{\delta,\epsilon}^*$ and $BSA_{\delta,\epsilon}^*$, and Fig. 4 also shows a comparison vs. A_ϵ^* in terms of the total number of expanded nodes, using $\delta = \frac{\epsilon}{4}$. In total, $BA_{\delta,\epsilon}^*$ and $BSA_{\delta,\epsilon}^*$ expanded fewer nodes than A_ϵ^* for $\epsilon \geq 0.2$.

It is statistically significant that for $\epsilon \leq 0.2$, A_ϵ^* expands more often fewer nodes on the same instance than $BSA_{\delta,\epsilon}^*$ with $\delta = \frac{\epsilon}{4}$. However, using $\epsilon = 0.3$ the inverse is true. For $\epsilon \geq 0.5$ the bidirectional search algorithm did significantly better than the corresponding unidirectional one.

Moreover, it is interesting to have a look at the data of the first meeting in Fig. 4. The worse bidirectional algorithms in the average met *earlier* than the better ones (but with worse solutions). After that, the latter could terminate rather quickly (or even immediately), while the former needed much time to satisfy the termination condition.

¹¹The performance of A_ϵ^* with $\epsilon = 0$ is better than that of A^* since it uses its FOCAL-heuristic as a tie-breaking rule.

4.3 Cross-domain comparison

Since we conducted experiments in two different domains, the question arises whether there are major differences in the results. For $\epsilon = 0$, IBS_ϵ^* (equivalent to BS^* there) did relatively much better in the route planning domain than for the 8-Puzzle. Most interestingly, BS^* was even 3 percent more efficient than A^* in terms of overall nodes expanded. However, with high statistical significance it was beaten by the latter in 63 percent of the cases which were not "draws". This observation reinforces the conjecture in [10] that BS^* is relatively more efficient than A^* on difficult problems.

Moreover, we found that BSA_ϵ^* was not that bad in comparison to A_ϵ^* in the route planning domain as it was for the puzzle. Generally, for both domains used for our experiments the better bidirectional search algorithms were more efficient than their corresponding unidirectional counterpart when the goal of finding optimal solutions is slightly relaxed.

4.4 Which is the best algorithm?

Since we investigated several bidirectional algorithms, the question arises which one is the best. Generally, the bidirectional algorithms using the improved termination condition (3) are more efficient than the others.

In particular, let us compare IBS_ϵ^* vs. $IBHPA_\epsilon$. For $\epsilon = 0$, IBS_ϵ^* was clearly superior to $IBHPA_\epsilon$ on the 500 8-Puzzle instances. Although already for $\epsilon \geq 0.1$ the total number of expanded nodes was nearly the same, the BS^* derivative was significantly better for $\epsilon \leq 0.9$ in terms of "winning", but with increasing values of ϵ the vast majority of cases were "draws". On the 100 15-Puzzle instances, the advantage of IBS_ϵ^* over $IBHPA_\epsilon$ was statistically significant over the whole range investigated, and according to both statistic tests. In the route planning domain, IBS_ϵ^* was superior to $IBHPA_\epsilon$ for $\epsilon \leq 0.3$ both regarding the total number of expanded nodes and in terms of "winning". However, for larger values of ϵ the majority of cases were "draws". In most of these cases both algorithms terminate immediately after first meeting, and up to this point they do the same. Generally, the difference appears to be more pronounced for more difficult problems.

Although the BS^* derivatives do not strictly dominate the corresponding derivatives of HPA (especially because of the cardinality criterion for selecting the search direction), they generally "lose" very seldom. Therefore, IBS^* and BSA^* with $\epsilon = e/4$ are candidates for being the best bidirectional heuristic search algorithms. We could not find that one of these is generally better than the other.

In general, it strongly depends on the properties of the given problem which algorithm is "best". Of course, linear-space algorithms must be used when there is insufficient memory available. However, whenever there is, our bidirectional algorithms should be considered. Of course, running time is another important issue. Again, it strongly depends on the domain, whether the overhead of maintaining the lists deteriorates the performance or not. In particular, it depends on the effort for computing heuristic values. Finally, of course, also the efficiency of implementation plays an important role regarding the running time.

5 Conclusion

For more than two decades there has been consensus that the main reason for the superiority of unidirectional search is that the bidirectional heuristic search is afflicted by the problem of search wavefronts missing each other. However, our results show that the missile metaphor of Pohl [15] is not really applicable. The primary problem appears to be that after the first meeting especially an algorithm aiming for optimal solutions has to spend much effort for subsequently improving the solution quality, and finally for proving that no other open node can give a better solution. Therefore, only a slight relaxation of solution quality already leads to strong improvements in efficiency, and in particular to even more than in the unidirectional case. Our novel termination criterion contributes strongly to this result because of the importance of this problem.

However, a serious limitation of classical best-first search is its memory problem. In contrast, $WIDA^*$ and $RBFS$ have only linear space requirements. Rao *et al.* [20] compared depth-first with best-first search generally. While the latter has the serious disadvantage of its memory problem, there are conditions, where classical best-first search is still most efficient. In particular, it is useful when both density of solutions and heuristic branching factor are very low (and there is sufficient memory available). Moreover, both IDA^* and $RBFS$ have large search overhead when there are many distinct values, and when the problem graph is not a tree [11].

Our bidirectional algorithms are designed to improve the performance under these conditions. While there are some differences, for both domains used for our experiments the main result about these algorithms is the same. We showed that and how bidirectional best-first search can be more efficient than the corresponding unidirectional counterpart without expensive wave-shaping. Therefore, even for only slightly relaxing the requirements on the solution quality, the best of our algorithms are the most efficient e-admissible algorithms known for those domains where sufficient memory is available. Such domains are not the most difficult ones addressable by search techniques like $WIDA^*$ and $RBFS$. But in those areas dominated by classical best-first search we found even better algorithms, based on a paradigm that was thought for long to be bad.

Currently, we are investigating a bidirectional combination of best-first search with linear-space search, in order to mitigate the issue of storage requirements. We hope that this will renew the interest in bidirectional search, which has additional potential for efficient implementation on parallel hardware.

Acknowledgements

The original implementation of our algorithms was based on a program written by Max Hofferer. A more efficient implementation was done by Angelika Gangl-Ehrenwerth and Harald Smetana, and the most efficient implementation based on code provided by Richard Korf is due to Ah Khorsand. The help of Harald Smetana and Ah Khorsand in running the 15-Puzzle experiments is highly appreciated. The experiments requiring very large memory were run on a machine of the computing center of the Technical University of Vienna.

A Pseudo Code of $BSA_{\delta, \epsilon}^*$

```

procedure  $BSA_{\delta, \epsilon}^*(s, t)$ 
1.  $g_1(s) \leftarrow g_2(t) \leftarrow 0; f_1(s) \leftarrow f_2(t) \leftarrow h_1(s); L_{min} \leftarrow \infty;$ 
2.  $OPEN_1 \leftarrow \{s\}; OPEN_2 \leftarrow \{t\};$ 
3.  $CLOSED_1 \leftarrow CLOSED_2 \leftarrow \emptyset;$ 
4. until  $OPEN_1 = \emptyset$  or  $OPEN_2 = \emptyset$  or
    $L_{min} \leq (1 + \epsilon) \max[\min_{x \in OPEN_1} (g_1(x) + h_1(x)),$ 
    $\min_{x \in OPEN_2} (g_2(x) + h_2(x))]$  do
5.   if  $|OPEN_1| \leq |OPEN_2|$  /* cardinality criterion */
6.     then  $d \leftarrow 1$  else  $d \leftarrow 2$ 
7.     endif
8.      $d' \leftarrow 3 - d;$  /* set the opposite search direction */
9.     Select  $m \in FOCAL_d$  with lowest  $h_{F_d}$  -value;
10.     $OPEN_d \leftarrow OPEN_d \setminus \{m\};$ 
11.     $CLOSED_d \leftarrow CLOSED_d \cup \{m\};$ 
12.    if  $m \in CLOSED_{d'}$ 
13.      then /* nip  $m$  in  $TREE_{d'}$  and prune  $TREE_{d'}$  */
14.         $\Omega_m \leftarrow \{n \mid n \in F_{d'}(m) \wedge p_{d'}(n) = m\};$ 
15.         $OPEN_{d'} \leftarrow OPEN_{d'} \setminus \Omega_m;$ 
16.      else
17.         $EXPAND(m);$ 
18.      endif
19.    enduntil
20.    if  $L_{min} = \infty$ 
21.      then no path exists
22.    else
23.      the solution path with costs  $L_{min}$  is
24.       $(s, \dots, p_1^*(MeetingNode), p_1(MeetingNode),$ 
25.       $MeetingNode, p_2(MeetingNode), \dots, t).$ 
26.    endif
27.  endprocedure.

procedure  $EXPAND(m)$ 
1.  $TrimFlag \leftarrow false;$ 
2. foreach  $n \in F_d(m)$  do
3.    $g \leftarrow g_d(m) + c_d(m, n);$ 
4.   if  $g + h_d(n) < L_{min}$ 
5.     then /* no screening */
6.        $f \leftarrow g + h_d(n);$ 
7.       if  $n \notin TREE_d$  /*  $n$  is a new node */
8.         then
9.            $g_d(n) \leftarrow g; f_d(n) \leftarrow f; p_d(n) \leftarrow m;$ 
10.           $OPEN_d \leftarrow OPEN_d \cup \{n\};$ 
11.          elseif  $g < g_d(n)$  /* better path to  $n$  found */
12.            then
13.               $g_d(n) \leftarrow g; f_d(n) \leftarrow f; p_d(n) \leftarrow m;$ 
14.              if  $n \in CLOSED_d$ 
15.                then /* reopen  $n$  */
16.                   $CLOSED_d \leftarrow CLOSED_d \setminus \{n\};$ 
17.                   $OPEN_d \leftarrow OPEN_d \cup \{n\};$ 
18.                endif
19.              endif
20.            if  $n \in TREE_{d'}$  and  $g_1(n) + g_2(n) < L_{min}$ 
21.              then /* update  $L_{min}$  */
22.                 $L_{min} \leftarrow g_1(n) + g_2(n);$ 
23.                 $MeetingNode \leftarrow n;$ 
24.                 $TrimFlag \leftarrow true;$ 
25.              endif
26.            endif
27.          endforeach
28.          if  $TrimFlag$ 
29.            then /* trim the open lists */
30.              Remove from  $OPEN_1$  and  $OPEN_2$  those nodes  $n$ 
31.              with  $(g + h)$ -values  $\geq L_{min}$   $(d = 1, 2)$  and which
32.              are not source nodes (for  $OPEN_1$  the source node
33.              is  $s$ ; for  $OPEN_2$  it is  $t)$ 
34.            endif
35.          endprocedure.

```

References

- [1] H.W. Davis, R.B. Pollack, and T. Sudkamp. Towards a better understanding of bidirectional search. In Proc. of AAAI-84, pages 68-72, Austin, TX, 1984. Los Altos, CA.: Kaufmann.
- [2] D. de Champeaux. Bidirectional heuristic search again. J. ACM, 30:22-32, 1983.
- [3] D. de Champeaux and L. Sint. An improved bidirectional heuristic search algorithm. J. ACM, 24:177-191, 1977.
- [4] P.E. Hart, N.J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. IEEE Transactions on Systems Science and Cybernetics, SSC-4(2):100-107, 1968.
- [5] E. Horowitz and S. Sahni. Fundamentals of Computer Algorithms. Springer-Verlag, New York, 1978.
- [6] A.L. K611. Heuristische Suche mit begrenzbarem Fehler. Diplomarbeit, Technische Universität Wien, 1992.
- [7] A.L. K611 and H. Kaindl. A new approach to dynamic weighting. In Proc. EC A1-92, pages 16-17, Vienna, 1992. Chichester, Wiley.
- [8] R.E. Korf. Linear-space best-first search. Artificial Intelligence. To appear.
- [9] R.E. Korf. Depth-first iterative deepening: An optimal admissible tree search. Artificial Intelligence, 27(1):97-109, 1985.
- [10] J.B.H. Kwa. BS*: An Admissible Bidirectional Staged Heuristic Search Algorithm. Artificial Intelligence, 38(2):95-109, 1989.
- [11] A. Mahanti, S. Gosh, D.S. Nau, A.K. Pal, and L. Kanal. Performance of IDA* on trees and graphs. In Proc. Tenth National Conference on Artificial Intelligence, pages 549-544, San Jose, 1992. Los Altos, CA.: Kaufmann.
- [12] N.J. Nilsson. Principles of Artificial Intelligence. Palo Alto, CA, Tioga, 1980.
- [13] J. Pearl. Heuristics: Intelligent Search Strategies for Computer Problem Solving. Addison-Wesley, Reading, MA, 1984.
- [14] J. Pearl and J.H. Kim. Studies in semi-admissible heuristics. IEEE Transactions on Pattern Analysis and Machine Intelligence, 4(4):392-399, 1982.
- [15] I. Pohl. Bi-directional and heuristics search in path problems. PhD thesis, Stanford University, 1969. SLAC Rept. No. 104.
- [16] I. Pohl. First results on the effect of error in heuristic search. In B. Meltzer and D. Michie, editors, Machine Intelligence 5, pages 219-236. Edinburgh University Press, Edinburgh, 1970.
- [17] I. Pohl. Bi-directional search. In Machine Intelligence 6, pages 127-140, Edinburgh, 1971. Edinburgh University Press.
- [18] I. Pohl. The avoidance of (relative) catastrophe, heuristic competence, genuine dynamic weighting and computational issues in heuristic problem solving. In Proc. International Joint Conference on Artificial Intelligence 1973, pages 12-17, Los Altos, CA, 1973.
- [19] G. Politowski and I. Pohl. D-node retargeting in bidirectional heuristic search. In Proc. of AAAI-84, pages 274-277, Austin, TX, 1984. Los Altos, CA.: Kaufmann.
- [20] V.N. Rao, V. Kumar, and R.E. Korf. Depth-first vs best-first search. In Proc. Ninth National Conference on Artificial Intelligence, pages 434-440, Anaheim, 1991. Los Altos, CA.: Kaufmann.