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

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

Hermann Kaindl<br>Siemens AG Osterreich<br>Geusaugasse 17<br>A-1030 Wien<br>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 $£$-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 waveshaping techniques.


## Notation

| s,t | St |
| :---: | :---: |
| $\Gamma_{1}(\underline{n})$ | Successors of node $n$ in the problem graph. |
| $r_{p}(n)$ | Parents of node $n$ in the problem graph. |
| $d$ | Current search direction index; when search in the forward direction $d=1$, and when in the backward direction $d=2$. |
| $d^{\prime}$ | $3-d_{i}$ it is the index of the direction opposite to the current search direction. |
| $c_{1}(m, n)$ | 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 $i=2$. |
| $g_{i}^{*}(n)$ | Coat of an optimal path from $s$ to $n$ if $i=1$, or from $t$ to $n$ if $i=2$. |
| $h_{i}^{*}(\underline{n})$ | Coat of an optimal path from $n$ to $t$ if $i=1$, |

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

|  | or |
| :---: | :---: |
| $g_{i}(\underline{n}), h_{1}(n)$ | Estimates of $g_{i}^{*}(n)$ and $h_{1}^{*}(n)$, respectively. |
| $h_{F}$. | Focal-heuristic. |
| $f_{i}(\mathbf{n})$ | Static evaluation function. |
| $C^{*}$ | Cost of an optimal path froms to $t$. |
| $L_{\text {min }}$ | Cost of the best (least costly) complete path found so far from sto $t$. |
| Tree, | The forward search tree. |
| Treer | The backward search tree. |
| Open, | The sei of open nodes in Treest. |
| Focal. | $\left\{m \mid\left(m \in \operatorname{OPEN}_{1}\right) \wedge\left(g_{i}(m)+h_{i}(m)\right) \leq\right.$ |
| \|OPEN, | Number of nodes in Open, |
| Closed, | The set of closed nodes in Trees ${ }_{\text {, }}$. |
| $p_{i}(\underline{n})$ | Parent of node $n$ in Tree,. |
| $p_{d}^{\prime}(n)$ | $\underbrace{p_{d}\left(p_{d} \ldots\left(p_{d}(n)\right) \ldots\right) .}$ |
| $\Omega_{m}$ | -times <br> The set of nodes in OPEN $d^{\prime}$ which are descendants of $m$ in TREE $d^{+}$. |
| MeetingNode | Node where Tree: 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 $(1+\varepsilon) C^{*}[14,13]$. 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 $(f=g+w h)^{2}$ Such an evaluation function is used by EPA [16]. ${ }^{8}$ A more recent approach to dynamic weight-

[^0]ing [7] compensates for the error made by the heuristic component using learnt weights, and it is ^-admissible without the need for a given upper bound like $N$. The e-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 $\boldsymbol{\varepsilon}>0$ they are typically more efficient than $A^{*}$.
$B H P A^{4}[15,17]$ 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 $\varepsilon$-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 $\varepsilon=0$. 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 e-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*.
${ }^{4}$ The 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.
${ }^{3}$ For $\varepsilon=\mathbf{0}$, IBS: - one of our new algorithms as introduced below - is identical to $B S^{*}$.


Figure 1: Phases of $I B S_{\ell}^{*}$ (on the 8-Puzzle).

## 2 Improved Bidirectional Search

In the following, we describe several ${ }^{\wedge}$-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\& and $\left.I B S_{*}^{*}\right)$

There are two ways of achieving an $\wedge$-admissible variant of BHPA:

- $B H P A_{w}$ using an evaluation function of the form $f=g+w h$, where $f$ can be rewritten as $f=g+(1+$ $\varepsilon) h$. The original termination condition of $B H P A$ $L_{m i n} \leq \max \left[\min _{x \in \operatorname{OPEN}_{1}} f_{1}(x), \min _{x \in \operatorname{OPEN}_{2}} f_{2}(x)\right](1)$ turns out to include $(1+\varepsilon) h$ implicitly in $B H P A_{w}$ : $L_{m 1 n} \leq \max \left[\min _{x \in \operatorname{OPEN}_{1}}\left(g_{1}(x)+(1+\varepsilon) h_{1}(x)\right)\right.$,
$\left.\min _{x \in \operatorname{OPEN}_{2}}\left(g_{2}(x)+(1+c) h_{2}(x)\right)\right] \cdot(2)$
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, $\boldsymbol{B H} P A_{w}$ using (1) or equivalently (2) as termination condition is $\varepsilon$-admissible when $h$ is admissible.
- $B H P A_{z}$ using $f=g+h$ (like $B H P A$ or a bidirectional admissible version of $A^{*}$ ), but using the termination criterion

$$
\begin{aligned}
& L_{\min } \leq(1+\varepsilon) \max \left[\min _{x \in \mathrm{OPEN}_{1}}\left(g_{1}(x)+h_{1}(x)\right),\right. \\
& \left.\min _{x \in \mathrm{OPEN}_{2}}\left(g_{2}(x)+h_{2}(x)\right)\right] \cdot(3) \\
& \text { BHPA, }_{2} \text { is also } \varepsilon \text {-admissible when } h \text { is admissible. }
\end{aligned}
$$

These ideas can be combined in an improved version as follows: $\boldsymbol{I} \boldsymbol{B} \boldsymbol{H} \boldsymbol{P} \boldsymbol{A}_{\boldsymbol{c}}$ uses $\boldsymbol{f}=\boldsymbol{g}+(\mathbf{1}+\boldsymbol{\varepsilon}) \boldsymbol{h}$ with $\boldsymbol{\varepsilon} \geq \mathbf{0}$ and the termination criterion (3) still guaranteeing eadmissibility.

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 /-values are greater or equal than $\mathrm{L}_{\text {min }}$ (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 $\mathrm{n}^{\prime}$ on an optimal path always satisfies $\boldsymbol{g}\left(\boldsymbol{n}^{\prime}\right)=\boldsymbol{g}^{*}\left(\boldsymbol{n}^{\prime}\right)$. In addition, we have to assume a consistent $h$ because of the $B S^{*}$ specific actions nipping and pruning. ${ }^{6}$

## $B S_{w}^{*}$ using $f=g+w h$ with $w \geq 1$ like HPA (in step 14 of $\left[10\right.$, p. 101]). ${ }^{7}$

- $B S_{\varepsilon}^{*}$ using $f=g+h$ and the following termination criterion (in step 3 of [10, p. 100]): $\mathrm{OPEN}_{1}$ or $\mathrm{OPEN}_{2}$ is empty or (3).
Again, these modifications can be combined, resulting in an improved algorithm we call IBS*.


### 2.2 Derivatives of $A_{\boldsymbol{q}}^{*}\left(B A_{\delta, 4}^{*}\right.$ and $\left.B S A_{\delta, k}^{*}\right)$

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 $A_{e}$ * [14]. ${ }^{8}$ Using $A_{\varepsilon}^{*}$ type of search in both directions and the termination condition (3), $\boldsymbol{B} \boldsymbol{A}_{\boldsymbol{e}}^{*}$ arises, which is eadmissible.

In contrast to the bidirectional algorithms IBHPA€ and $I \mathrm{ISSb}^{*}$ described above, $\boldsymbol{B} \boldsymbol{A}_{\boldsymbol{c}}^{*}$ is normally less efficient than the corresponding unidirectional algorithm ( $A_{\epsilon}^{*}$ ). 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-adrnissible 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.

[^1]

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

Therefore, we found an improvement by adding another parameter $\boldsymbol{\delta}, \boldsymbol{B} \boldsymbol{A}_{\delta, \mathrm{E}}^{*}$ uses 6 for creating the FOCALlists, and $£$ for termination. Much better results were achieved, for instance, with $\delta=\frac{\boldsymbol{\varepsilon}}{2}$ and $\delta=\frac{\boldsymbol{\varepsilon}}{4}$.

Of course, it is also possible to incorporate the approach of $A_{c}^{*}$ into the framework of $\mathrm{BS}^{*}$. BSA* arises from $\boldsymbol{B} S^{*}$ 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 $\boldsymbol{B} \boldsymbol{S}_{\boldsymbol{\varepsilon}}^{*}$ (see above),
- the selection of a node $m$ for expansion (in step 7 of [10, p. 100]) as usual for $\boldsymbol{A}_{\boldsymbol{c}}^{*}$.
Moreover, another parameter can be used, $\boldsymbol{B S} \boldsymbol{A}_{\delta, \varepsilon}^{*}$ arises from $\boldsymbol{B S} \boldsymbol{A}_{\boldsymbol{\varepsilon}}^{*}$ analogously to $\boldsymbol{B} \boldsymbol{A}_{\delta, \boldsymbol{\varepsilon}}^{*}$ from $\boldsymbol{B} \boldsymbol{A}_{\boldsymbol{\varepsilon}}^{*}$. 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 $\boldsymbol{B S} \boldsymbol{A}_{\boldsymbol{\varepsilon}}^{*}$ and $\boldsymbol{B S} \boldsymbol{A}_{\boldsymbol{\delta}, \boldsymbol{e}}^{*}$ are the ones most different from similar proofs for $\boldsymbol{A}_{*}^{*}$, 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 $\boldsymbol{c}_{\boldsymbol{i}}(\boldsymbol{\pi}, \boldsymbol{n}) \geq \boldsymbol{e}>\mathbf{0}$ for some $c$, and a consistent heuristic estimator is used.
Definition 3.1 $A$ path $P$ from $s$ to $t$ is $\varepsilon$-optimal, if $\operatorname{cost}(P)<(1+\varepsilon) C^{*}$.
Definition 3.2 A search algorithm is said to be $\varepsilon$ admissible, if it terminates with an $\varepsilon$-optimal solution whenever a solution exists [14,13].
Definition 3.3 An algorithm $A \backslash$ is said to dominate $A_{2}$ if every node expanded by Al is also expanded by $\mathrm{A}_{2}$ [13].
${ }^{9}$ Due 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 $B S A_{\delta, *}^{*}(\delta \leq \varepsilon)$ terminates either an e-optimal path $P^{\prime}$ hes 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 ahallowest open nodes along an optimal path $P=\left(s=n_{1}, \ldots, n_{i}, \ldots, n_{j}, \ldots, t=n_{t}\right)$ with $g_{1}\left(n_{i}\right)=g_{1}^{*}\left(n_{i}\right)$ and $g_{2}\left(n_{j}\right)=g_{i}^{*}\left(n_{j}\right)$.
Proof: Except for the BS specific actions such nodes must exist at any time before termination, since we account for reopening of nades (cf. (1s)), and before the firat meeting no such actions can occur. Furthermore, $i<j$. Otherwise the optimal path would have been found. There are four cases where theses nodes could be removed from Orend without installing a successor substituting them along $P$. We consider the case of $n_{1}$ to be removed from $\mathrm{OPEN}_{1}$ (the case of $n$, to be removed from $\mathrm{OPEN}_{2}$ is analogous).

1. Nipping: Node $n$, is nipped in Tree $1_{1}$, i.e. $n_{1}$ must be in Closed ${ }_{2} . n_{k}$ is a node in OPEN 2 having the smallest $f_{2}$ value. Since $n_{1}$ is closed in TreE $2, n_{i}$ was selected for expansion before $n$, (while searching in the backward direction) and we have

$$
\begin{array}{rll}
f_{2}\left(n_{i}\right) & \leq\left(\begin{array}{ll}
1+\delta) f_{2}\left(n_{k}\right) & n_{1} \text { from Focal } \\
& \leq(1+\varepsilon) f_{2}\left(n_{k}\right)
\end{array}\right. & \delta \leq \varepsilon \\
& \leq(1+\varepsilon) f_{2}\left(n_{j}\right) & f_{2}\left(n_{k}\right) \leq f_{2}\left(n_{j}\right)
\end{array}
$$

We distinguish two cases

$$
\begin{aligned}
&(1.1)(1+e)\left[g_{2}^{*}\left(n_{j}\right)+k_{2}\left(n_{i}, n_{j}\right)\right]+\varepsilon g_{1}^{*}\left(n_{i}\right) \geq g_{2}\left(n_{i}\right) \\
& \operatorname{cost}\left(P^{\prime}\right)=g_{i}^{*}\left(n_{i}\right)+g_{2}\left(n_{i}\right) \\
& \leq g_{1}\left(n_{i}\right)+\varepsilon g_{1}\left(n_{i}\right) \\
& \leq(1+\varepsilon)\left[g_{2}^{*}\left(n_{j}\right)+k_{2}\left(n_{i}, n_{j}\right)\right] \\
& \leq(1+\varepsilon)\left[g_{1}^{\prime}\left(n_{i}\right)+g_{2}^{*}\left(n_{j}\right)+k_{2}\left(n_{i}, n_{j}\right)\right]
\end{aligned}
$$

This implies $L_{\min } \leq(1+\varepsilon) C^{*}$.
(1.2) $(1+e)\left[g_{2}^{*}\left(n_{j}\right)+k_{2}\left(n_{1}, n_{j}\right)\right]+\varepsilon g_{1}^{*}\left(n_{1}\right)<g_{2}\left(n_{1}\right)$

Due to consistency $h_{2}\left(n_{j}\right) \leq k_{2}\left(n_{i}, n_{j}\right)+h_{2}\left(n_{1}\right)$.

$$
\begin{aligned}
g_{2}\left(n_{j}\right)+h_{2}\left(n_{2}\right) \leq & g_{2}\left(n_{j}\right)+k_{2}\left(n_{i}, n_{j}\right)+h_{2}\left(n_{1}\right) \\
(1+c)\left(g_{2}\left(n_{j}\right)+h_{2}\left(n_{j}\right)\right) \leq & (1+\varepsilon)\left[g_{2}^{*}\left(n_{j}\right)+k_{2}\left(n_{i}, n_{j}\right)\right] \\
& +(1+\varepsilon) h_{2}\left(n_{1}\right) \\
& +\varepsilon g_{1}^{*}\left(n_{i}\right)-\varepsilon g_{1}^{\prime}\left(n_{i}\right) \\
(1+\varepsilon) f_{2}\left(n_{j}\right)< & g_{2}\left(n_{j}\right)+h_{2}\left(n_{i}\right) \\
& -\varepsilon[\underbrace{\geq 0}_{\left.\sum_{1}\left(n_{i}\right)-h_{2}\left(n_{i}\right)\right]} \\
< & f_{2}\left(n_{i}\right) \text { contradicting }(4)
\end{aligned}
$$

2. Pruning: Node $n_{i}$ has an ancestor node $n_{s}$ in Tree ${ }_{1}$ which was nipped in TreE $2 . n_{2}$ must be on an optimal path (say $P$ ), otherwise $n$, cannot have its parent pointer set to $n_{x}$. Therefore, $g_{1}\left(n_{x}\right)=g_{1}^{*}\left(n_{x}\right)$ and the proof is analogous to case 1 , substituting $n_{s}$ for $n_{4}$.
s. Trimming: If $n_{1}$ would be trimmed from OPEN ${ }_{1}$, then

$$
\begin{equation*}
g_{1}\left(n_{i}\right)+h_{1}\left(n_{i}\right) \geq L_{\min } \tag{5}
\end{equation*}
$$

Since $n_{1}$ is on an optimal path and $g_{1}\left(n_{i}\right)=g_{i}^{*}\left(n_{i}\right)$, and since $h_{1}$ is consistent and consequently admissible, $h_{1}\left(n_{i}\right) \leq h_{1}^{1}\left(n_{i}\right)$, this implies $g_{1}\left(n_{i}\right)+h_{1}\left(n_{i}\right) \leq g_{i}^{i}\left(n_{i}\right)+$ $h_{1}^{*}\left(n_{i}\right)=C^{*}$. If no path $P^{\prime}$ with $\operatorname{cost}\left(P^{\prime}\right) \leq(1+\varepsilon) C^{*}$ was found yet, $C^{*} \leq(1+\varepsilon) C^{*}<L_{\text {min }}$, 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 $C^{\bullet}<$ $L_{\text {min }}$, this would lead to an analogous contradiction as in case 3 .
Lemma 3.2 If a path exists from s to $t, B S A_{6, c}^{*}$ will not terminate before finding an e-optimal path.
Proof: The proof is analogous to the one in [17, p. 199], but using our termination condition (3).

Lemma 3.3 // a path exists from s to $\mathrm{t}, \boldsymbol{B S} \boldsymbol{A}_{\delta, \%}^{*}$ 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]).
Theorem 3.1 If a path exists from $s$ to $t$, $\mathrm{BSA}^{*}{ }_{6, \mathrm{c}}$ terminates with a solution whose cost does not exceed the optimal cost by more than a factor $(1+\varepsilon)$, i.e., $B S A_{\delta, \varepsilon}^{*}$ is $\boldsymbol{\varepsilon}$-admissible.
Proof: This result follows from the lemmas above.
Corollary 3.1 BSA* is $\varepsilon$-admissible.
Theorem 3.2 A bidirectional algorithm $\mathrm{B} \backslash$ using the termination condition (S) dominates another bidirectional algorithm $\mathrm{B}_{2}$ differing from $\mathrm{B} \backslash$ only in using the condition (2) instead.
Proof: Since the right side of $(\mathrm{S})$ is always greater than or equal to the right side of (2), algorithm $B \backslash$ cannot terminate later than $\mathbf{B}_{\mathbf{2}}$.
Corollary 3.2 $I B H P A_{\text {, }}$ dominates $B H P A_{\boldsymbol{w}}$.

## 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 $\varepsilon=0$ is identical to a comparison of $H P A^{*}$ with $B S^{*}$ and $B H P A^{\prime}$. There, the unidirectional HPA* is clearly more efficient. However, already for $\varepsilon \geq 0.2$ 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 e-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 $\varepsilon \leq 0.1 H P A^{*}$ expands more often fewer nodes than $I B S_{s}^{*}$ on the same instance, the inverse is true for $\varepsilon \geq 0.3$.

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 $1 B H P A_{\varepsilon}$ and $/ B S_{\varepsilon}^{*}$ with $H P A^{*}$, WIDA* and RBFS on the 15 -Puzzle (using the sample of 100 instances in [9]). Due to their memory problem, it is infeasible for $I B H P A_{\varepsilon}, I B S_{\varepsilon}^{*}$ and $H P A^{*}$ to find or even to guarantee optimal solutions here.

For $\varepsilon=1, H P A^{*}$ can solve all the 100 instances of [9], but it requires memory for nearly 1 million nodes in the maximum. In contrast, $\boldsymbol{I B H} P \boldsymbol{A}_{\boldsymbol{\varepsilon}}$ requires only a storage of slightly more than 300 k nodes in maximum, and IBS* slightly more than 250 k nodes. Both are clearly more efficient than the others in terms of nodes generated. Interestingly, $H P A^{*}$ "wins" against both $I B H P A_{t}$ and $I B S_{c}^{*}$ slightly more often, but the results arenot statisticaily significant according to the sign test. While the worst case bound $\varepsilon=1$ 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 $\varepsilon=0.75, I B H P A_{\varepsilon}$ and $I B S_{\epsilon}^{*}$ require storage of slightly more than $6000 \bar{k}$ 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 $I B H P A_{\boldsymbol{e}}$ and $I B S_{\boldsymbol{e}}^{*}$ is even clearer than for $\boldsymbol{\varepsilon}=\mathrm{I}$, generating an order of mag-


Figure 4: Comparison on the 8-Puzzle (500 instances).
nitude fewer nodes than $W I D A^{*}$.
For $\varepsilon=0.5, I B H P A_{\varepsilon}$ and $J B S_{*}^{*}$ 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. $H P A^{*}$ cannot solve 8 instances with this amount of storage. The number of nodes generated by $I B H P A_{c}\left(I B S_{z}^{*}\right)$ is in the average smaller by a factor of 21 (25) than that generated by $W I D A^{*}$. Therefore, it appears that the performance advantage of our bidirectional algorithms steadily increases with decreasing values of $\varepsilon$ (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 $A_{e}^{*}$

In contrast, the bidirectional algorithms $B A_{c}^{*}$ and $B S A_{z}^{*}$ are normally less efficient than the corresponding unidirectional $A_{\text {; }}^{*}$ (see Fig. 4 for the results on a random set of 5008 -Puzzles). ${ }^{11}$ Therefore, we designed $B A_{\delta,}^{*}$ and $B S A_{\delta, \varepsilon}^{*}$, and Fig. 4 also shows a comparison vs. $A_{t}^{*}$ in terms of the total number of expanded nodes, using $\delta=\frac{\ell}{4}$. In total, $B A_{\delta, c}^{*}$ and $B S A_{\delta, c}^{*}$ expanded fewer nodes than $A_{e}^{*}$ for $\varepsilon \geq 0.2$.

It is statistically significant that for $\varepsilon \leq 0.2, A_{\varepsilon}^{*}$ expands more often fewer nodes on the same instance than $B S A_{b, \varepsilon}^{*}$ with $\delta=\frac{e}{4}$. However, using $\varepsilon=0.3$ the inverse is true. For $\varepsilon \geq 0.5$ 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 carlier 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.

[^2]
### 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 $\varepsilon=0, I B S_{\varepsilon}^{*}$ (equivalent to $B S^{*}$ there) did relatively much better in the route planning domain than for the 8-Puzzle. Most interestingly, $B S^{*}$ 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 $B S^{*}$ is relatively more efficient than $A^{*}$ on difficult problems.

Moreover, we found that $B S A_{\varepsilon}^{*}$ was not that bad in comparison to $A_{\varepsilon}^{*}$ 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 $J B S_{c}^{*}$ vs. $I B H P A_{f}$. For $\varepsilon=0, I B S_{\varepsilon}^{*}$ was clearly superior to $I B H P A_{\varepsilon}$ on the 500 8-Puzzle instances. Although already for $\varepsilon \geq 0.1$ the total number of expanded nodes was nearly the same, the $B S^{\prime}$ derivative was significantly better for $\varepsilon \leq 0.9$ in terms of "winning", but with increasing values of $\varepsilon$ the vast majority of cases were "draws". On the 100 15Puzzle instances, the advantage of $I B S_{\varepsilon}^{*}$ over $I B H P A_{\varepsilon}$ was statistically significant over the whole range investigated, and according to both statistic tests. In the route planning domain, $I B S_{\varepsilon}^{*}$ was superior to $I B H P A_{\varepsilon}$ for $\varepsilon \leq 0.3$ both regarding the total number of expanded nodes and in terms of "winning". However, for larger values of $\varepsilon$ 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 $B S$ derivatives do not strictly dominate the corresponding derivatives of $H P A$ (especially because of the cardinality criterion for selecting the search direction), they generally "lose" very seldom. Therefore, $I B S^{*}$ and $B S A^{*}$ e with $6=\mathrm{e} / 4$ 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 15Puzzle 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 $B S A_{\delta, e}^{*}$

```
    procedure \(B S A_{b, s}^{*}(s, t)\)
        \(g_{1}(s) \leftarrow g_{2}(t) \leftarrow 0 ; f_{1}(s) \leftarrow f_{2}(t)-h_{1}(s) ; L_{m i n} \leftarrow \infty ;\)
        \(\mathrm{OPEN}_{1} \leftarrow\{s\} ; \mathrm{OPEN}_{2} \leftarrow\{t\} ;\)
        Closed \(_{1} \leftarrow\) Closed \(_{2}-\varnothing_{\text {; }}\)
        until OPEN \(_{1}=\varnothing\) or OPEN \(_{2}=\varnothing\) or
            \(L_{\text {min }} \leq(1+\varepsilon) \max \left[\min _{x \in \operatorname{OPEN}_{1}}\left(g_{1}(x)+h_{1}(x)\right)\right.\),
                        \(\left.\min _{x \in \mathrm{OPEN}_{3}}\left(g_{2}(x)+h_{2}(x)\right)\right]\) do
            if \(\left|\mathrm{OPEN}_{1}\right| \leq\left|\mathrm{OPEN}_{2}\right| \boldsymbol{P R}^{*}\) cardinality criterion */
                then \(d \approx 1\) else \(d \leftarrow 2\)
            endif
            \(d^{\prime}-3-d ; / *\) set the opposite search direction */
            Select \(m \in\) Focal \(d_{d}\) with lowest \(h_{F_{d}}-\) value;
            OPEN \(_{d} \leftarrow\) OPEN \(_{d} \backslash\{m\} ;\)
```



```
            if \(m \in \operatorname{CLOSED}_{\mathrm{d}^{\prime}}\)
                then /* nip \(m\) in Treeed and prune Tree \(d^{*} * /\)
                    \(\Omega_{m} \leftarrow\left\{n \mid n \in \Gamma_{d^{\prime}}(m) \wedge p_{d^{\prime}}(n)=m\right\} ;\)
                    OPEN \(_{d^{\prime}}-\) OPEN \(_{d^{\prime}} \backslash \Omega_{m} ;\)
                else
                    EXPAND(m);
            endif
        enduntil
15. if \(L_{m i n}=\infty\)
            then no path exists
            else
```

                the solution path with costs \(L_{\text {min }}\) is
                ( \(s_{1} \ldots, p_{1}^{2}\) (MeetingNode), \(p_{1}\) (MeetingNode),
                MeetingNode, \(p_{2}\) (MeetingNode), ..., \(\mathbf{t}\) ).
        endif
    endprocedure.
    procedure EXPAND(m)
        TrimFlag-false,
        foreach \(n \in \Gamma_{d}(m)\) do
            \(g+g_{d}(m)+c_{d}(m, n) ;\)
            if \(g+h_{d}(n)<L_{\text {min }}\)
                then /* no screening */
                \(f \leftarrow g+h_{d}(n)\);
                if \(n \notin\) TRFE, \(^{\prime} /{ }^{*} n\) is a new node */
                    then
                    \(g_{d}(n)-g ; f_{d}(n) \leftarrow f_{i} p_{d}(n) \leftarrow m ;\)
                    OPENd \(_{d}-\) Opend \(_{d} \cup\{n\} ;\)
                    elsif \(g<g_{d}(n) / *\) better path to \(n\) found *,
                    then
                    \(g_{d}(n)-g_{i} f_{d}(n) \leftarrow f ; p_{d}(n) \leftarrow m ;\)
                    if \(n \in\) Closed \(_{d}\)
                                    then /* reopen \(n\) */
                                    Closed \(_{d}-\) Closed \(_{d} \backslash\{n\} ;\)
                                    \(\mathrm{OPEN}_{d}-\mathrm{OPEN}_{d} \cup\{n\} ;\)
                    endif
                endif
    14. if $n \in$ TrEE $_{d^{\prime}}$ and $g_{1}(n)+g_{2}(n)<L_{m i n}$
then /* update $L_{\text {min }}{ }^{*} /$
$L_{m, n} \leftharpoondown g_{1}(n)+g_{2}(n) ;$
MeetingNode $-n$;
TrimFlag - true;
endif
endif
endforeach
15. If TrimFlag
then /* trim the open lists */
16. Remove from OPEN ${ }_{1}$ and OPEN 2 those nodes $n$
with $(g+h)$-values $\geq L_{\text {min }}(d=1,2)$ and which
are not source nodes (for OPEN, the source node
is $s$; for $\mathrm{OPEN}_{2}$ it is $t$ )
endif
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 Universitat Wien, 1992.
[7] A.L. K611 and H. KaindI. 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):97109, 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 tress 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 heunsttc 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.


[^0]:    ${ }^{2}$ In fact, besides using an admissible heuristic component $h$ the algorithm muat 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.
    ${ }^{2}$ For constant $w$, this is equivalent to $f=\left(1-w^{\prime}\right) g+w^{\prime} h$ with $w^{\prime}=\frac{w}{1+w}$.
    ${ }^{3}$ For our comparisons, we use a slight variant of $H P A$, which includes the option to move nodes back from Closed

[^1]:    ${ }^{6}$ Despite consistency, reopening of nodes is necessary because of the $(1+\boldsymbol{c})$-factor in the evaluation function.
    ${ }^{7}$ This 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.
    ${ }^{8}$ Pearl 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.

[^2]:    ${ }^{11}$ The performance of $A_{\varepsilon}^{*}$ with $\epsilon=0$ is better than that of $A^{*}$ since it uses its Focal-heuristic as a tie-breaking rule.

