

Faster than Weighted A*: An Optimistic Approach to Bounded Suboptimal Search

Jordan Thayer and Wheeler Ruml



UNIVERSITY *of* NEW HAMPSHIRE

{jtd7, ruml} at `cs.unh.edu`

Motivation

- Finding optimal solutions is prohibitively expensive.

Introduction

■ Motivation

Algorithm Overview

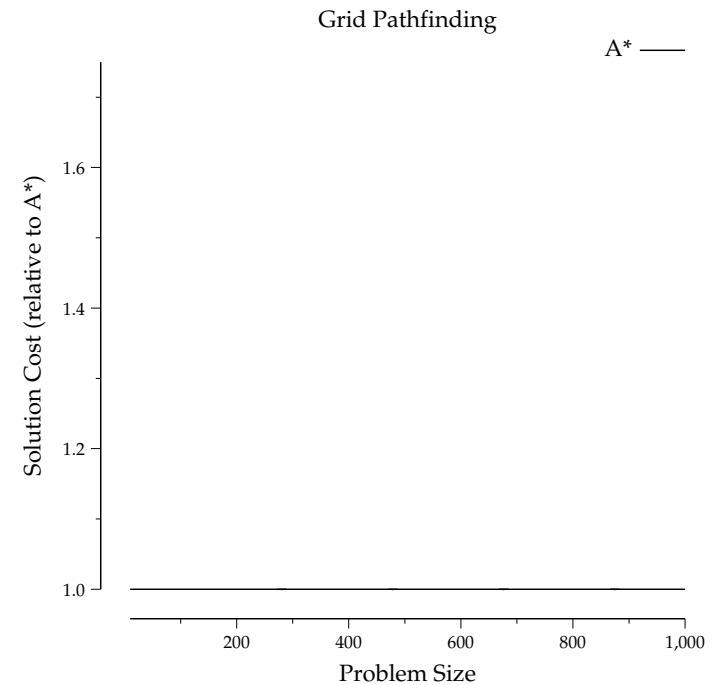
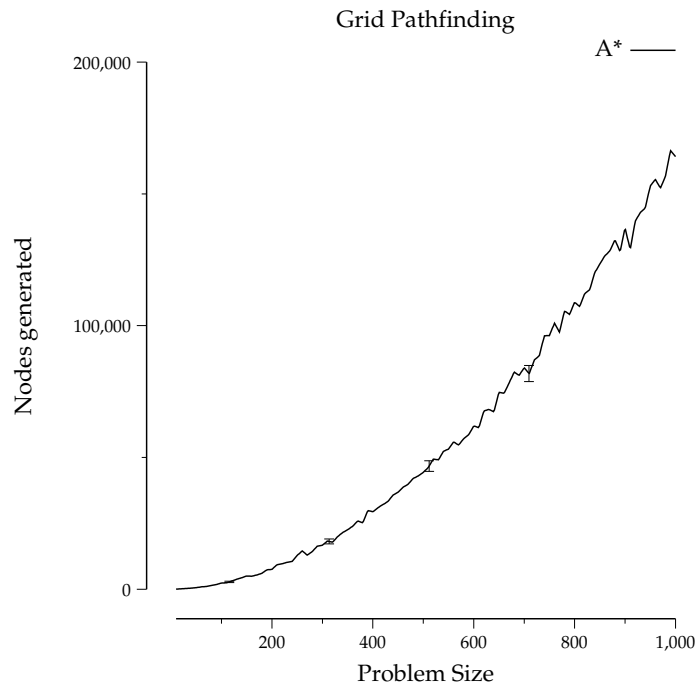
1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion



Motivation

Introduction

■ Motivation

Algorithm Overview

1: Greedy Phase

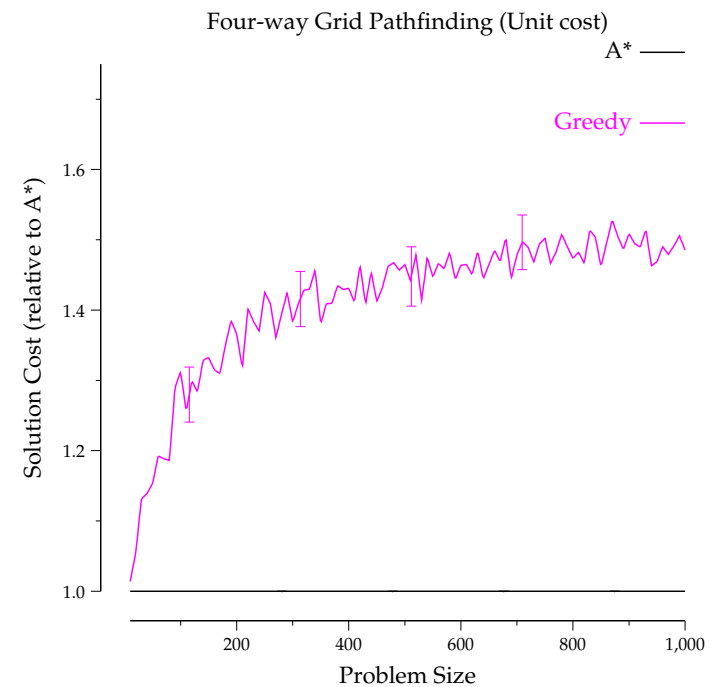
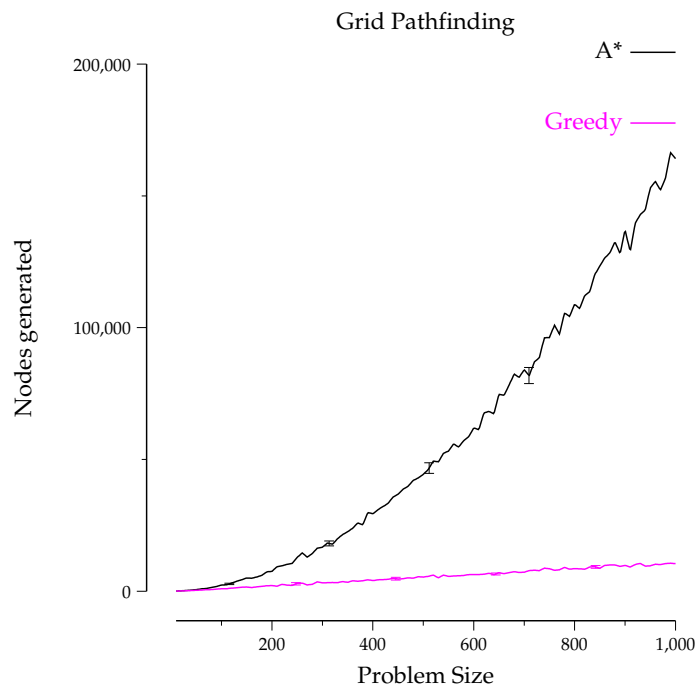
2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

- Finding optimal solutions is prohibitively expensive.
- Greedy solutions can be arbitrarily bad.



Motivation

Introduction

■ Motivation

Algorithm Overview

1: Greedy Phase

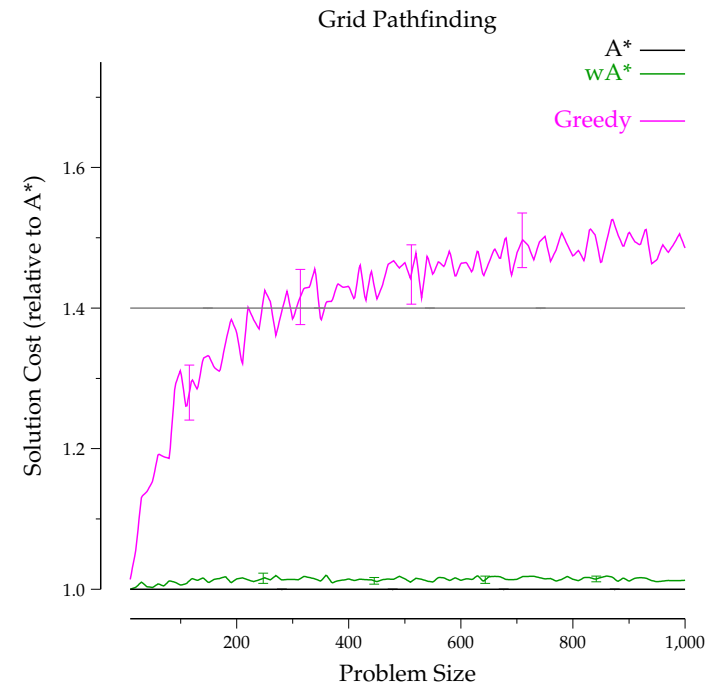

2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

- Finding optimal solutions is prohibitively expensive.
- Greedy solutions can be arbitrarily bad.
- Weighted A* bounds suboptimality.



Motivation

Introduction

■ Motivation

Algorithm Overview

1: Greedy Phase

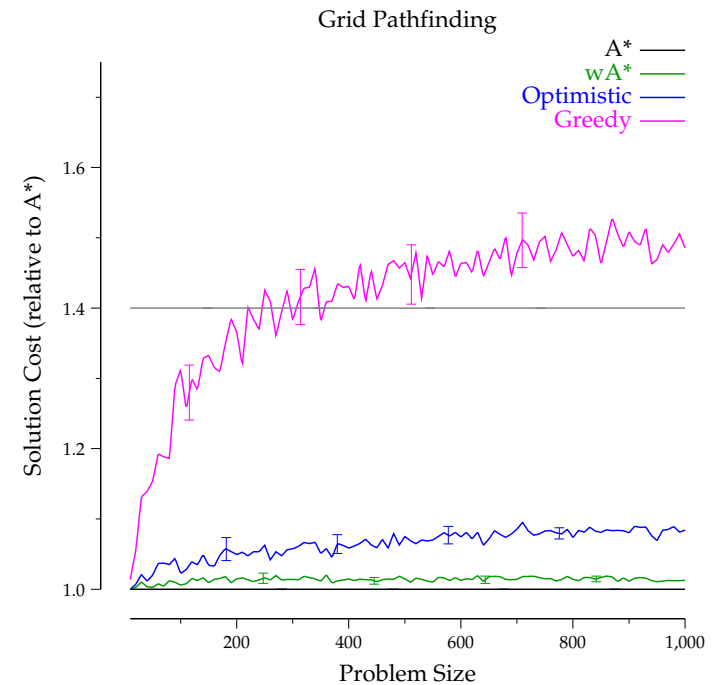
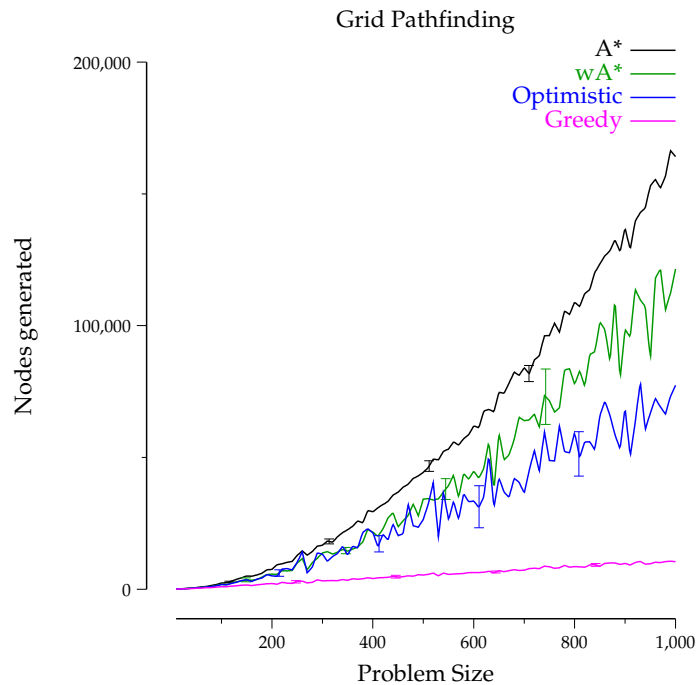
2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

- Finding optimal solutions is prohibitively expensive.
- Greedy solutions can be arbitrarily bad.
- Weighted A* bounds suboptimality.
- Optimistic Search: faster search within the same bound.



Introduction

Algorithm Overview

■ Predecessors

■ Basic Idea

1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

Algorithm Overview

Talk Outline

Introduction

Algorithm Overview

■ Predecessors

■ Basic Idea

1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

- **Algorithm Overview**

 - Run weighted A^* with a weight higher than the bound.
 - Expand additional nodes to prove solution quality.

- The Greedy Search Phase

- The Cleanup Phase

- Empirical Evaluation

- Further Observations

Previous Algorithms: A^*

Introduction

Algorithm Overview

■ Predecessors

■ Basic Idea

1: Greedy Phase


2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

- A best first search expanding nodes in f order.
- $f(n) = g(n) + h(n)$
If $h(n)$ is admissible, returns optimal solution.



Previous Algorithms: Weighted A^*

Introduction

Algorithm Overview

■ Predecessors

■ Basic Idea

1: Greedy Phase


2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

- A best first search expanding nodes in f' order.
- $f'(n) = g(n) + w \cdot h(n)$
Solution quality bounded by w for admissible $h(n)$.



Optimistic Search: The Basic Idea

Introduction

Algorithm Overview

■ Predecessors

■ Basic Idea

1: Greedy Phase


2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

1. Run weighted A^* with a high weight.
2. Expand node with lowest f value after a solution is found.
Continue until $w \cdot f_{min} > f(sol)$
This 'clean up' guarantees solution quality.



Optimistic Search: The Basic Idea

Introduction

Algorithm Overview

■ Predecessors

■ Basic Idea

1: Greedy Phase


2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

1. Run weighted A^* with a high weight.
2. Expand node with lowest f value after a solution is found.
Continue until $w \cdot f_{min} > f(sol)$
This 'clean up' guarantees solution quality.



Optimistic Search: The Basic Idea

Introduction

Algorithm Overview

■ Predecessors

■ Basic Idea

1: Greedy Phase


2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

1. Run weighted A^* with a high weight.
2. **Expand node with lowest f value after a solution is found.**
Continue until $w \cdot f_{min} > f(sol)$
This 'clean up' guarantees solution quality.



Optimistic Search: The Basic Idea

Introduction

Algorithm Overview

■ Predecessors

■ Basic Idea

1: Greedy Phase


2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

1. Run weighted A^* with a high weight.
2. Expand node with lowest f value after a solution is found.
Continue until $w \cdot f_{min} > f(sol)$
This 'clean up' guarantees solution quality.



Introduction

Algorithm Overview

1: Greedy Phase

■ Weighted A^*

2: Cleanup Phase

Empirical Evaluation

Further Observations

Conclusion

1: Greedy Phase

Talk Outline

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

■ Weighted A^*

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)

- Algorithm Overview
- The Greedy Search Phase
 - Weighted A^* becomes faster as the bound grows.
 - Weighted A^* is often better than the bound.
- The Cleanup Phase
- Empirical Evaluation
- Further Observations

Large Bounds, Faster Solution

- wA^* returns solutions faster as the bound increases.

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[Weighted \$A^*\$](#)


[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)

Pearl and Kim Hard



Weighted A^* is often better than the bound

- wA^* returns solutions better than the bound.

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)


[Weighted \$A^*\$](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)



Introduction

Algorithm Overview

1: Greedy Phase

2: Cleanup Phase

■ w -Admissibility

Empirical Evaluation

Further Observations

Conclusion

2: Cleanup Phase

Talk Outline

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

■ w -Admissibility

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)

- Algorithm Overview
- The Greedy Search Phase
- The Cleanup Phase
 - Expand additional nodes in f order.
 - Quit when the solution is provably within the bound.
- Empirical Evaluation
- Further Observations

Proving w -Admissibility

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)


[2: Cleanup Phase](#)

■ w -Admissibility

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)



- p is the deepest node on an optimal path to opt .
- f_{min} is the node with the smallest f value.

Proving w -Admissibility

Introduction

Algorithm Overview

1: Greedy Phase

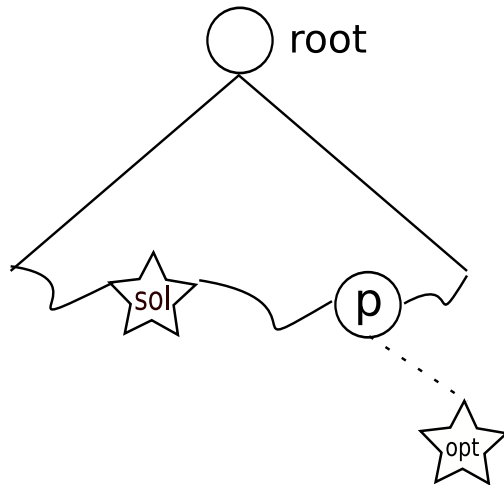
2: Cleanup Phase

■ w -Admissibility

Empirical Evaluation

Further Observations

Conclusion



- p is the deepest node on an optimal path to opt .
- f_{min} is the node with the smallest f value.

$$f(f_{min}) \leq f(p) \leq f(opt)$$

f_{min} provides a lower bound on solution cost.

Determine f_{min} by priority queue sorted on f

Proving w -Admissibility

Introduction

Algorithm Overview

1: Greedy Phase


2: Cleanup Phase

■ w -Admissibility

Empirical Evaluation

Further Observations

Conclusion



- p is the deepest node on an optimal path to opt .
- f_{min} is the node with the smallest f value.

$$f(f_{min}) \leq f(p) \leq f(opt)$$

f_{min} provides a lower bound on solution cost.

Determine f_{min} by priority queue sorted on f

Optimistic Search: Run a greedy search

Expand f_{min} until $w \cdot f_{min} \geq f(sol)$

Proving w -Admissibility

Introduction

Algorithm Overview

1: Greedy Phase


2: Cleanup Phase

■ w -Admissibility

Empirical Evaluation

Further Observations

Conclusion



- p is the deepest node on an optimal path to opt .
- f_{min} is the node with the smallest f value.

$$f(f_{min}) \leq f(p) \leq f(opt)$$

f_{min} provides a lower bound on solution cost.

Determine f_{min} by priority queue sorted on f

Optimistic Search: Run a greedy search

Expand f_{min} until $w \cdot f_{min} \geq f(sol)$

Introduction

Algorithm Overview

1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

■ Performance

Further Observations

Conclusion

Empirical Evaluation

Talk Outline

Introduction

Algorithm Overview

1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

■ Performance

Further Observations

Conclusion

- Algorithm Overview
- The Greedy Search
- Guaranteeing solution quality
- **Empirical Evaluation**
Results in several domains.
- Further Observations

Empirical Evaluation

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

■ Performance

[Further Observations](#)

[Conclusion](#)

- Sliding Tile Puzzles
 - Korf's 100 15-puzzle instances (add date)
- Traveling Salesman
 - Unit Square
 - Pearl and Kim Hard (add date)
- Grid world path finding
 - Four-way and Eight-way Movement
 - Unit and Life Cost Models
 - 25%, 30%, 35%, 40%, 45% obstacles
- Temporal Planning
 - Blocksworld, Logistics, Rover, Satellite, Zenotravel

See paper for additional plots.

Performance of Optimistic Search

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

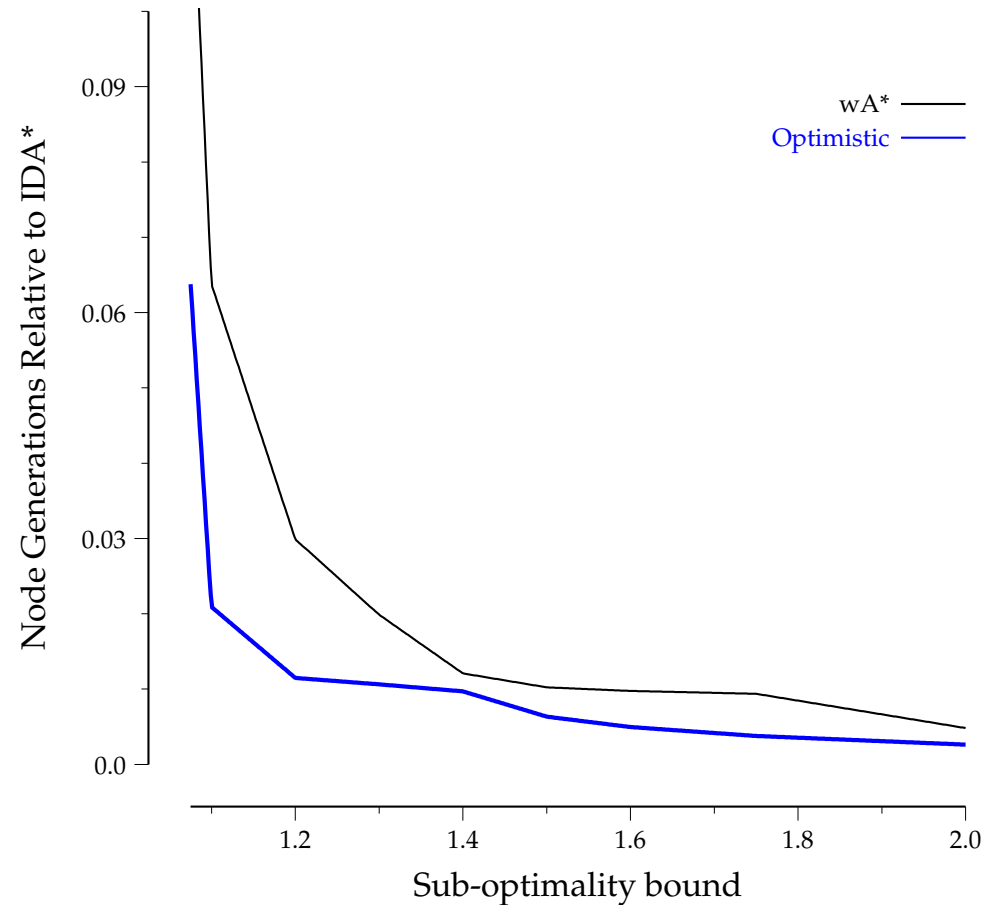
[Empirical Evaluation](#)

■ Performance

[Further Observations](#)

[Conclusion](#)

Korf's 15 Puzzles: $h = \text{Manhattan Distance}$



Performance of Optimistic Search

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

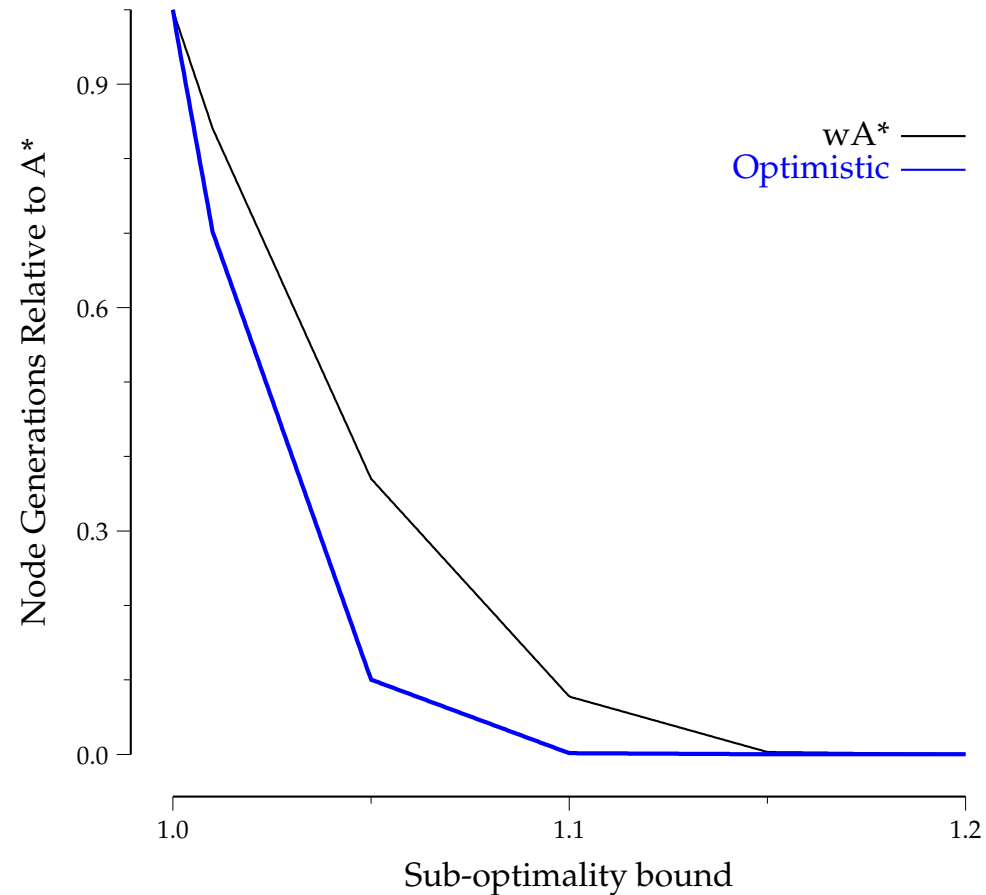
[Empirical Evaluation](#)

Performance

[Further Observations](#)

[Conclusion](#)

TSP: Pearl and Kim Hard



Performance of Optimistic Search

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

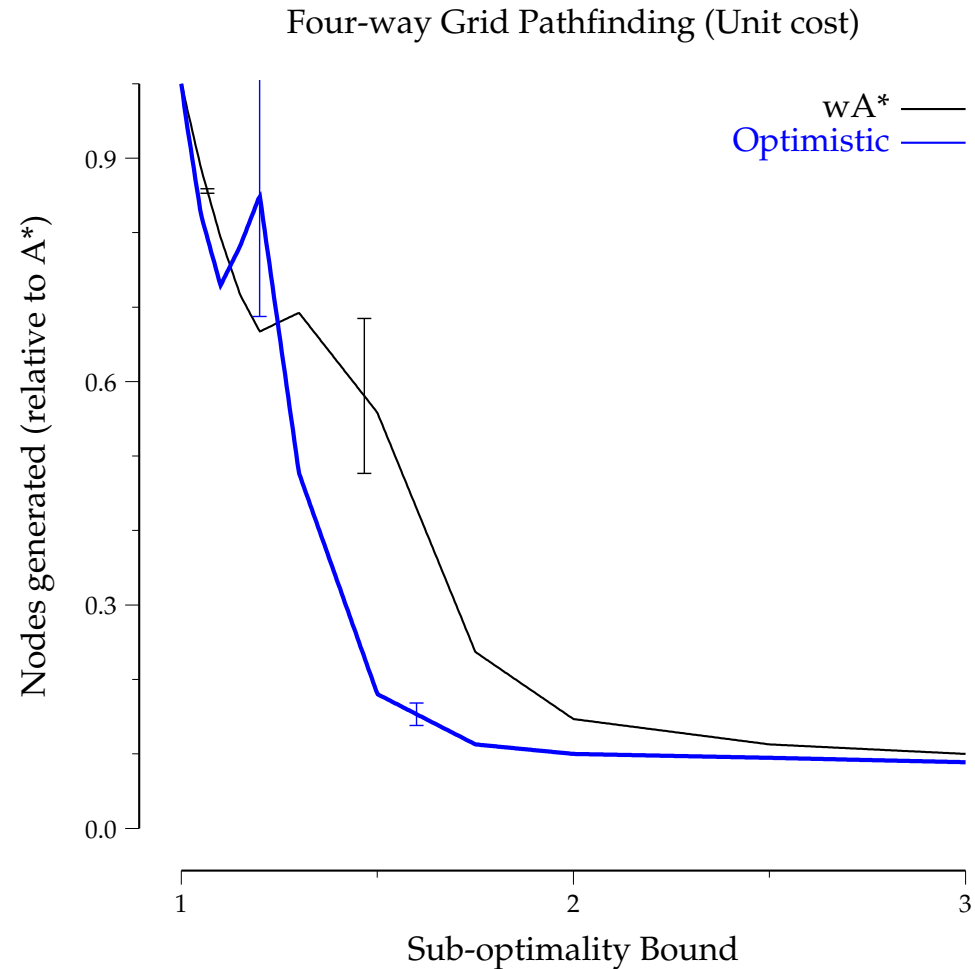
[2: Cleanup Phase](#)

[Empirical Evaluation](#)

Performance

[Further Observations](#)

[Conclusion](#)



Performance of Optimistic Search

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

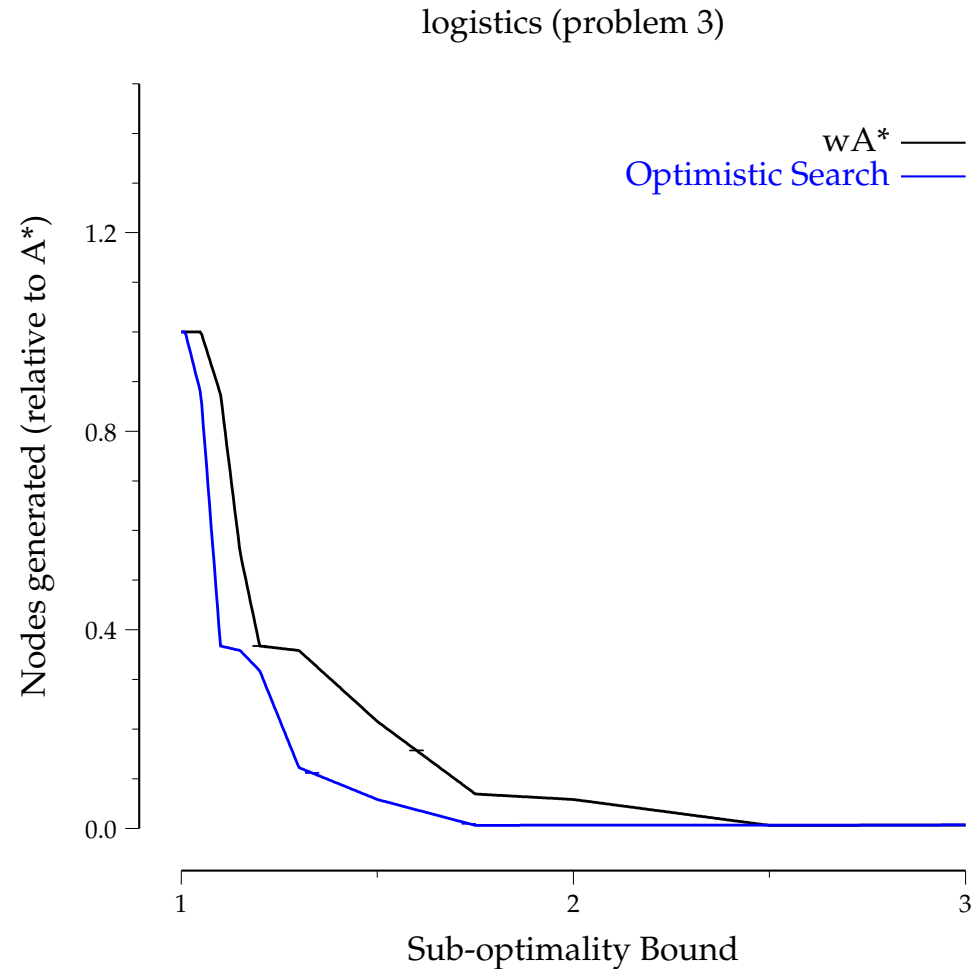
[2: Cleanup Phase](#)

[Empirical Evaluation](#)

Performance

[Further Observations](#)

[Conclusion](#)



Introduction

Algorithm Overview

1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

Further Observations

■ Expansion Policy

■ BAwA*

Conclusion

Further Observations

Talk Outline

Introduction

Algorithm Overview

1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

Further Observations

■ Expansion Policy

■ BAwA*

Conclusion

- Algorithm Overview
- The Greedy Search
- Guaranteeing solution quality
- Empirical Evaluation
- Further Observations
 - Strict vs. Loose Expansion Policy
 - Bounded Anytime Weighted A^*

Expansion Policy

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

■ [Expansion Policy](#)

■ [BAwA*](#)

[Conclusion](#)

Strict Expansion Order:

- Algorithms like wA^* , A_ϵ^* , Dynamically Weighted A^*
- Any expanded node can be shown to be within the bound at the time of their expansion
- Quality bound comes from this

Loose Expansion Order:

- Algorithms like Optimistic Search
- No restriction on the nodes expanded initially.
- Quality bound requires node expansion beyond the initial solution.

Bounded Anytime Weighted A^*

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

■ Expansion Policy

■ **BAwA***

[Conclusion](#)

- Anytime Heuristic Search:
 - Running weighted A^* with a high weight
 - Continue node expansions after a solution is found

Bounded Anytime Weighted A^*

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

■ [Expansion Policy](#)

■ [BAwA*](#)

[Conclusion](#)

- Anytime Heuristic Search:
 - Running weighted A^* with a high weight
 - Continue node expansions after a solution is found
- Bounded Anytime Weighted A^* :
 - Running weighted A^* with a high weight
 - Continue node expansions after a solution is found
 - Add a second priority queue allows us to converge on a bound instead of on optimal.

Optimistic Search expansions

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)


[Further Observations](#)

■ Expansion Policy

■ BAwA*

[Conclusion](#)

1. Run weighted A^* with a high weight.
2. Expand node with lowest f value after a solution is found.
Continue until $w \cdot f_{min} > f(sol)$
This 'clean up' guarantees solution quality.



Bounded Anytime Weighted A^* Expansions

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)


[Further Observations](#)

■ Expansion Policy

■ BAw A^*

[Conclusion](#)

1. Run weighted A^* with a high weight.
2. Expand node with lowest f' value after a solution is found.
Continue until $w \cdot f_{min} > f(sol)$
This 'clean up' guarantees solution quality.



Bounded Anytime Weighted A*

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)


[Further Observations](#)

■ Expansion Policy

■ BAwA*

[Conclusion](#)

Korf's 15 Puzzles



Bounded Anytime Weighted A*

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)


[Further Observations](#)

■ Expansion Policy

■ BAwA*

[Conclusion](#)

Pearl and Kim Hard



Conclusion

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)

■ [Conclusion](#)

■ [Advertising](#)

Optimistic Search:

- Simple to implement.
- Performance is predictable.
- Current results are good, tuning could help.
Optimal greediness is still an open question.
- Consistently better than Weighted A^*
If you currently use wA^* , you should use Optimistic Search.

The University of New Hampshire

Tell your students to apply to grad school in CS at UNH!

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)

■ Conclusion

■ Advertising



- friendly faculty
- funding
- individual attention
- beautiful campus
- low cost of living
- easy access to Boston, White Mountains
- strong in AI, infoviz, networking, systems, bioinformatics

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)

[Additional Slides](#)

■ Loose Bounds

■ Duplicates

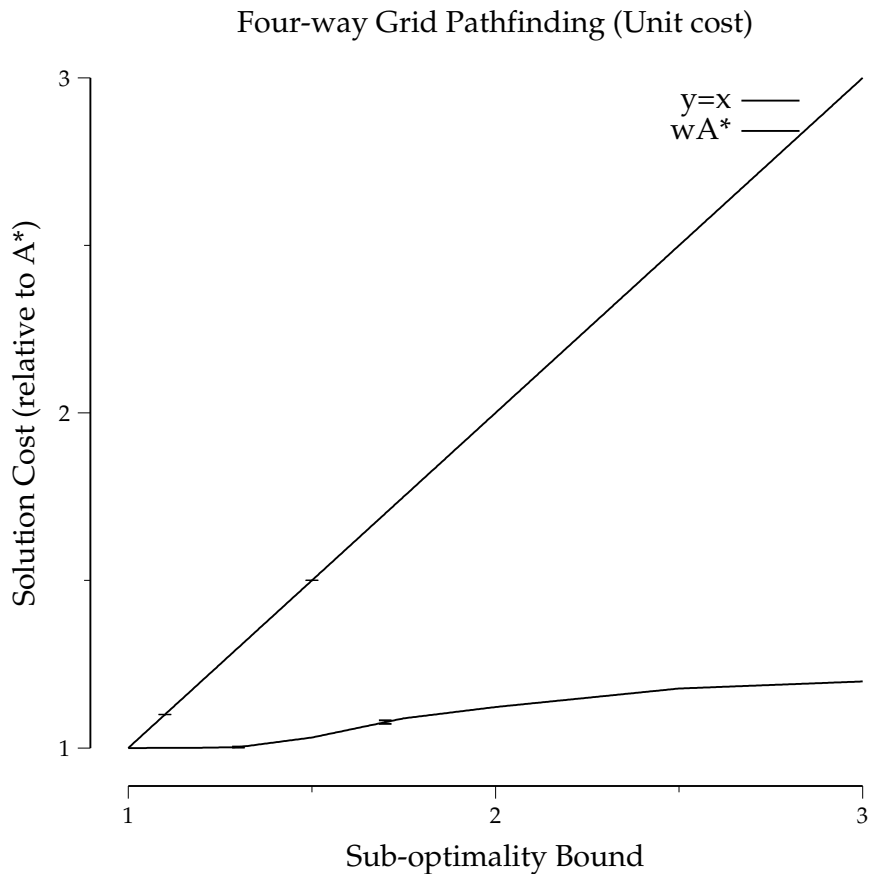
■ Pseudo Code

Additional Slides

Weighted A^* is often better than the bound

- wA^* returns solutions better than the bound.

- Introduction
- Algorithm Overview
- 1: Greedy Phase
- 2: Cleanup Phase
- Empirical Evaluation
- Further Observations
- Conclusion
- Additional Slides
 - Loose Bounds
 - Duplicates
 - Pseudo Code



Weighted A^* Respects a Bound

Introduction

Algorithm Overview

1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

Further Observations


Conclusion

Additional Slides

■ Loose Bounds

■ Duplicates

■ Pseudo Code



$$f(n) = g(n) + h(n)$$

$$f'(n) = g(n) + w \cdot h(n)$$

$$g(sol)$$

$$f'(sol) \leq f'(p)$$

$$g(p) + w \cdot h(p) \leq w \cdot (g(p) + h(p))$$

$$w \cdot f(p) \leq w \cdot f(opt)$$

$$w \cdot g(opt)$$

Therefore, $g(sol) \leq w \cdot g(opt)$

Weighted A^* Respects the Bound and Then Some

Introduction

Algorithm Overview

1: Greedy Phase

2: Cleanup Phase

Empirical Evaluation

Further Observations


Conclusion

Additional Slides

■ Loose Bounds

■ Duplicates

■ Pseudo Code



$$f(n) = g(n) + h(n)$$

$$f'(n) = g(n) + w \cdot h(n)$$

$$g(sol)$$

$$f'(sol) \leq f'(p)$$

$$g(p) + w \cdot h(p) \leq w \cdot (g(p) + h(p))$$

$$w \cdot f(p) \leq w \cdot f(opt)$$

$$w \cdot g(opt)$$

$$g(p) + w \cdot h(p) \leq w \cdot g(p) + w \cdot h(p)$$

Duplicate Dropping can be Important

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)

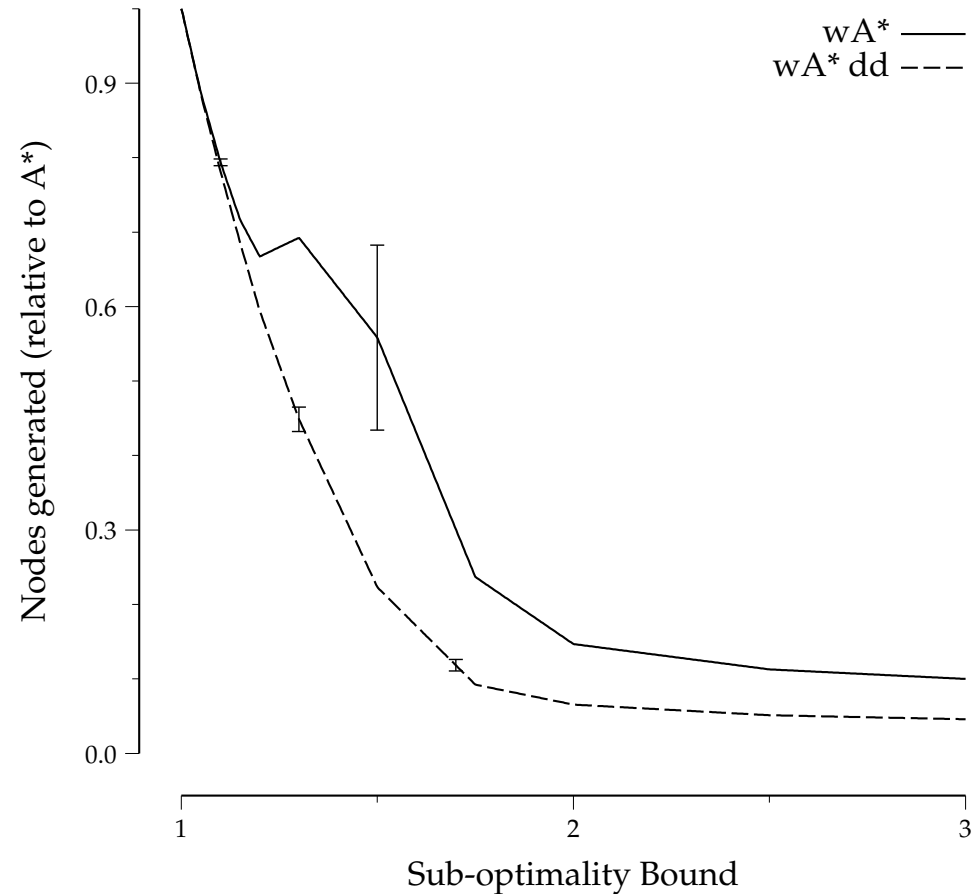
[Additional Slides](#)

■ Loose Bounds

■ Duplicates

■ Pseudo Code

Four-way Grid Pathfinding (Unit cost)



Sometimes it isn't

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)


[Conclusion](#)

[Additional Slides](#)

■ Loose Bounds

■ Duplicates

■ Pseudo Code



Sometimes it isn't

Duplicates can be delayed during the greedy search phase.

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

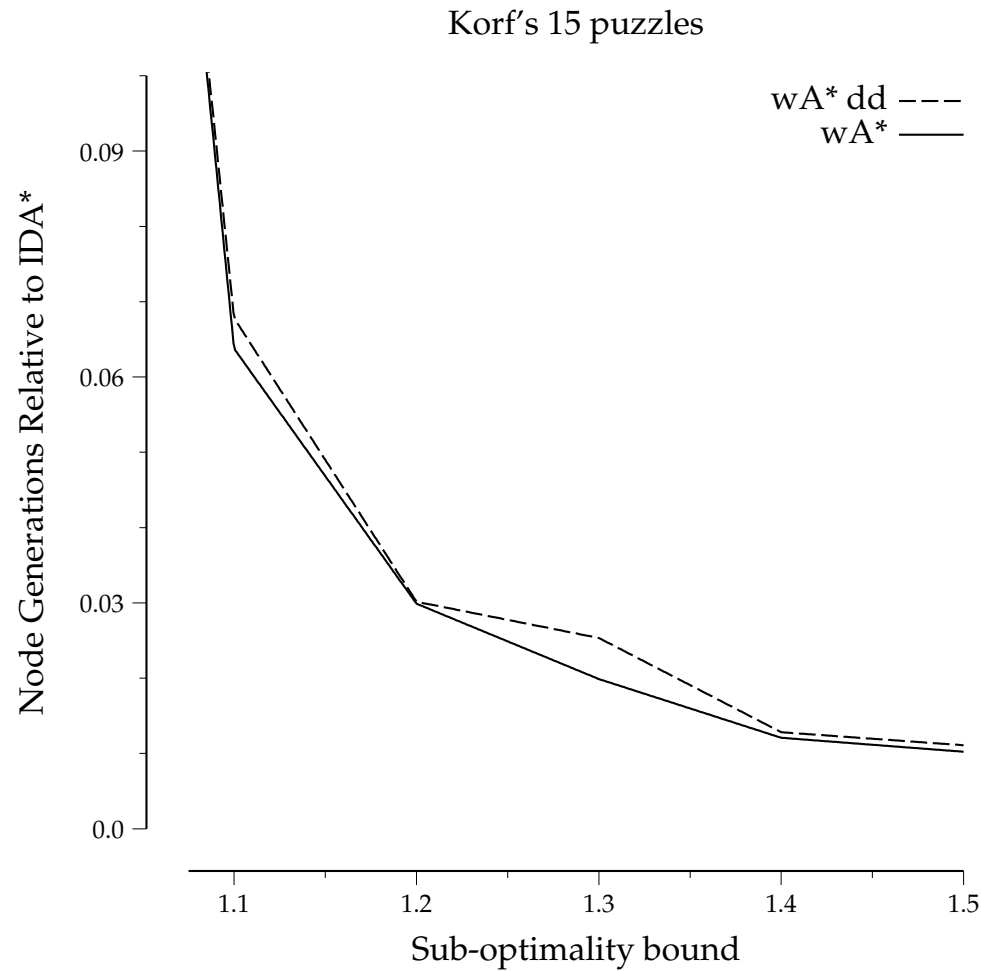
[Conclusion](#)

[Additional Slides](#)

■ Loose Bounds

■ Duplicates

■ Pseudo Code



Pseudo Code

[Introduction](#)

[Algorithm Overview](#)

[1: Greedy Phase](#)

[2: Cleanup Phase](#)

[Empirical Evaluation](#)

[Further Observations](#)

[Conclusion](#)

[Additional Slides](#)

■ Loose Bounds

■ Duplicates

■ Pseudo Code

Optimistic Search(*initial*, *bound*)

1. $open_f \leftarrow \{initial\}$
2. $open_f \leftarrow \{initial\}$
3. $incumbent \leftarrow \infty$
4. repeat until $bound \cdot f(\text{first on } open_f) \geq f(incumbent)$:
5. if $\widehat{f}(\text{first on } open_f) < \widehat{f}(incumbent)$ then
6. $n \leftarrow \text{remove first on } open_f$
7. remove n from $open_f$
8. else $n \leftarrow \text{remove first on } open_f$
9. remove n from $open_f$
10. add n to *closed*
11. if n is a goal then
12. $incumbent \leftarrow n$
13. else for each child c of n
14. if c is duplicated in $open_f$ then
15. if c is better than the duplicate then
16. replace copies in $open_f$ and $open_f$
17. else if c is duplicated in *closed* then
18. if c is better than the duplicate then
19. add c to $open_f$ and $open_f$
20. else add c to $open_f$ and $open_f$