GBFHS: A Generalized *Breadth-First* Heuristic Search Algorithm

Mike Barley[1], Pat Riddle[1], Carlos Linares López[2], Sean Dobson[1], & Ira Pohl[3]

1. University of Auckland, 2. Universidad Carlos III de Madrid, 3. UC Santa Cruz

Introduction: Very Brief Early History

- 1966 – Nicholson: Bidirectional blind search algorithm
  - Exponential reduction in node expansions over unidirectional blind search
- 1968 – Hart, Nilsson, & Raphael: Optimal unidirectional heuristic search algorithm
  - Guaranteed admissible and “optimally efficient”

Introduction: Very Brief Early History cont'd

- 1969 – Pohl: Bidirectional heuristic search algorithm
  - Idea: combine bidirectional blind & unidirectional heuristic search algorithms
  - Approach : Have A* go in both directions
  - Goal: to be better than both
  - Result: didn't work out that way

Introduction: Very Brief Recent History

- 2015: Barker & Korf -
  - "*any front-to-end bidirectional heuristic search algorithm will likely be dominated by unidirectional heuristic search or bidirectional brute-force search*"
- 2016/2017: Holte, Felner, Sharon & Sturtevant/& Chen -
  - **MM/MMe**: first bidirectional heuristic search algorithms that always meet in the middle
- 2017: Chen, Holte, Zilles & Sturtevant -
  - **NBS**: a bidirectional heuristic search algorithm that automatically determines how to split the search between the two directions & "*does no more than 2VC(I) state expansions*"
  - **NBSe** is an unpublished version of **NBS** that uses knowledge of epsilon

State-of-the-art: 2018

- *Surely*, Barker & Korf's 2015 conjecture has been disproved *or has it?*
  - Does A* still win with a strong heuristic?
  - Does bidirectional blind still win over a "weak" heuristic?

State-of-the-art: 2018

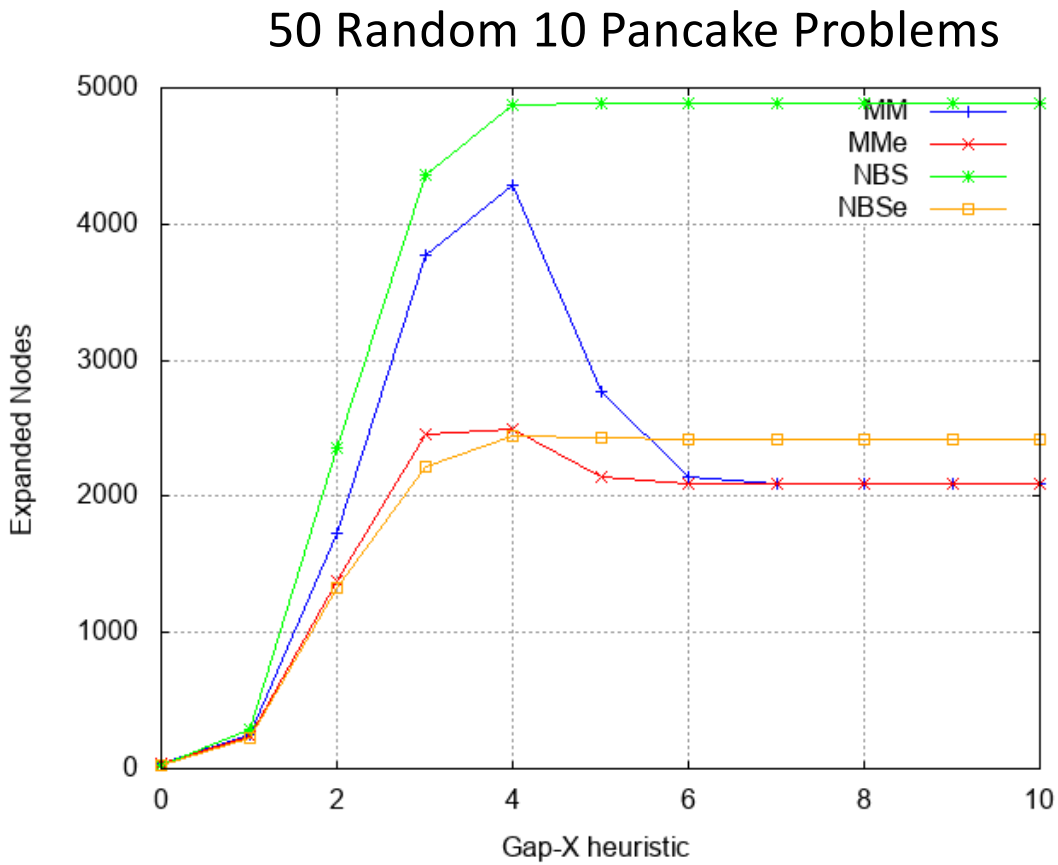

Unit-cost domain

GAP-10 = blind

X>5 -> blind

*MM/MMe GAP-3-5 lose to blind (**hump**)*

NBS/NBSe doesn't seem to lose to blind

NBSe State-of-the-art: 2018

50 Random 10 Pancake Problems

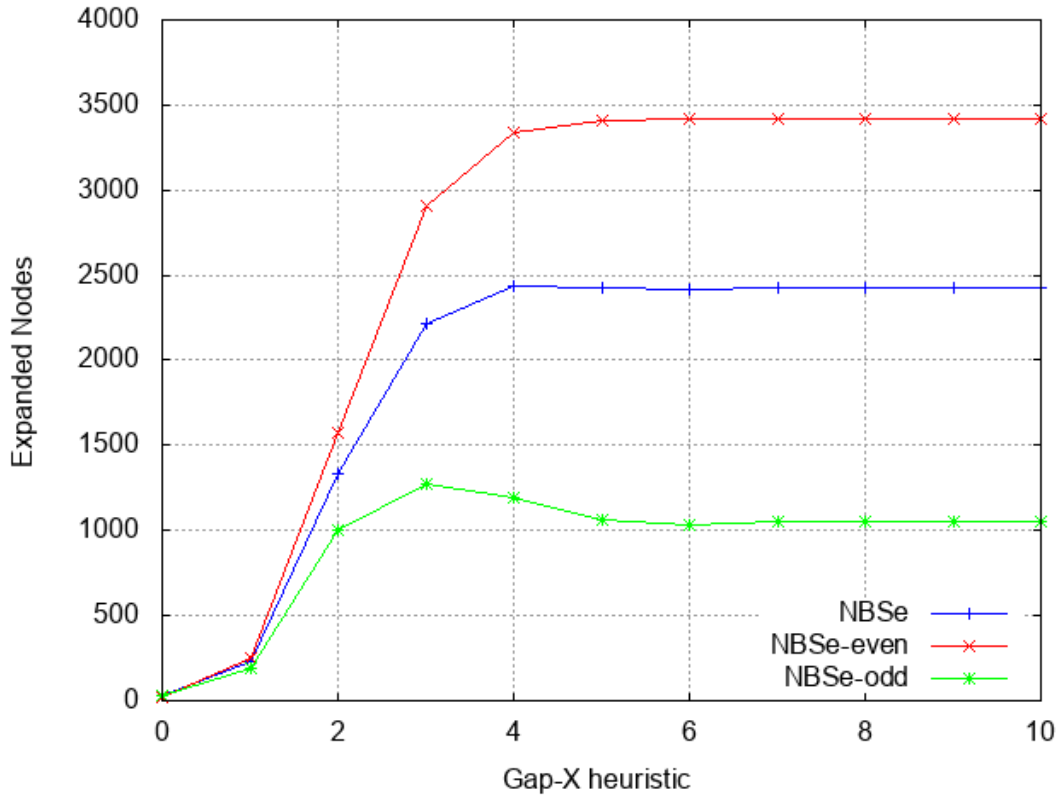

NBSe has no hump

NBSe on odd length soltn loses to blind
NBSe on even length soltn has no hump

NBS similar but has hump on evens & larger hump

Well-behaved heuristic search algorithms

- Want algorithms that given a “stronger” heuristic “never” expand more nodes
- So, informally, an algorithm is *well-behaved* if when given a stronger heuristic, it “never” expands more nodes
- An algorithm that is not well-behaved is *ill-behaved*
- MM/MMe, NBS/NBSe all seem to be **ill-behaved**

Splits & Lower Bounds

- *Split* is the division between forward & backward searches & it's precise definition is algorithm dependent
  - E.g., the split for MM/MMe is *(ceiling(C*/2), ceiling(C*/2))*

- Search algorithm's *lowerBound* when using a heuristic & split is the smallest possible number of nodes it expands to stop with an optimal solution regardless of tie-breaking strategy used

Formal Definition of Well-Behaved

- Algorithm **A** is *well-behaved* if given 2 admissible & consistent heuristics h1 & h2, such that h1(n)⩾h2(n), and $\mathbf{A}_{h1}$ & $\mathbf{A}_{h2}$ both stop with same split then lowerBound($\mathbf{A}_{h2}$) ⩾ lowerBound($\mathbf{A}_{h1}$)

Generalized Breadth-First Heuristic Search (GBFHS)

- Inputs: problem/domain, least edge cost(e), split fn, heuristic fn
- Works with any split, even unidirectional, determined by split fn
- Works with uniform-cost or arbitrary-cost domains
- Requires edge costs to be non-negative integers
- Admissible

GBFHS Claims

1. First bidirectional heuristic search algorithm *proven to be well-behaved*
2. It can expand substantially fewer nodes than existing algorithms
3. It can be “directed” to meet anywhere
4. In uniform-cost domains, first collision results in optimal solution
5. When going unidirectionally it can expand fewer nodes than textbook A*

GBFHS logic

- *fLim* : is there a solution with this length? 0 -> *?*
- $fLim = gLim_F + gLim_B + e - 1$, $gLim_F$ & $gLim_B$ is called a *split*
- Node *n* is *expandable* in direction *d* if $f_d(n) \leqslant fLim$ & $g_d(n) < gLim_d$
- Stops as soon as best current solution has cost $\leqslant fLim$
- *fLim* incremented by 1 when no solution with cost $\leqslant fLim$ is found

Comparisons

Heuristic/Algorithm	MMe	NBS	NBSe	GBFHS
16 Pancake: GAP	283	383	335	**279**
16 Pancake: GAP-2	587,283	836,251	625,900	**391,277**
16 Pancake: GAP-3	7,100,998	11,843,943	6,682,497	**3,778,017**
15 Puzzle: Man Dist	13,162,312	13,543,385	12,709,777	**12,507,393**

State-of-the-art: 2018 revisited

50 Random 10 Pancake Problems: *Total* Expanded Nodes

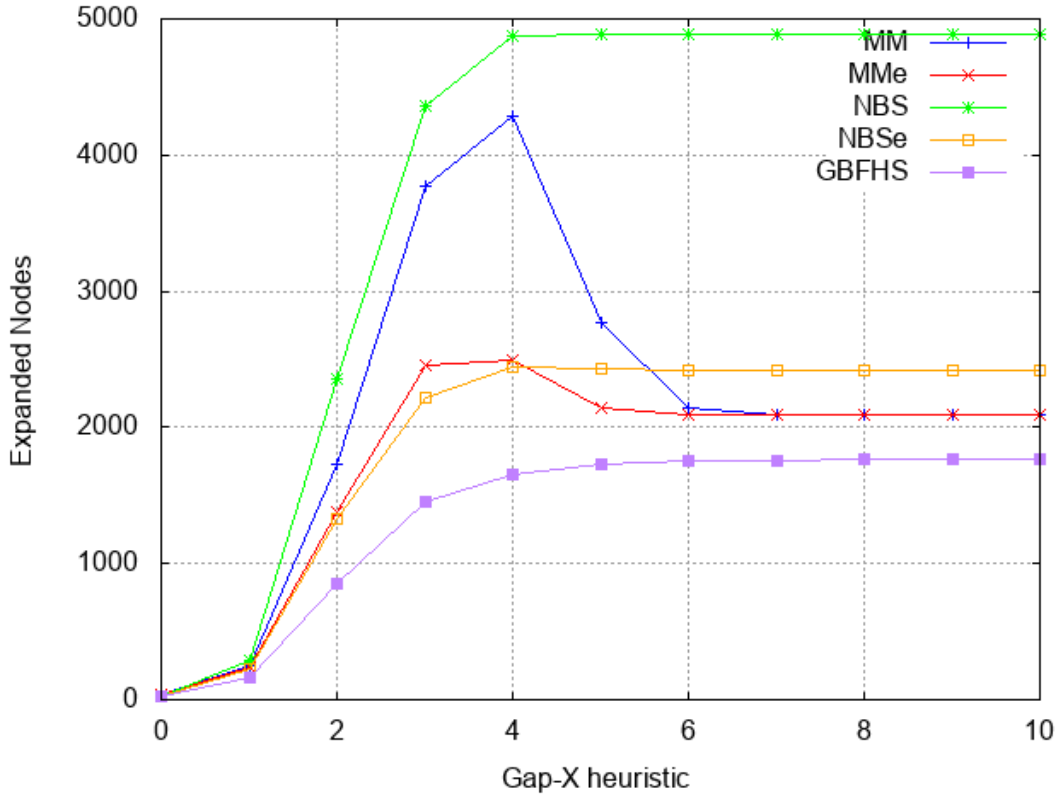

At GAP-3,4, everyone expands on average at least 50% more nodes than GBFHS

Even on blind, GBFHS expands fewer nodes

*Using Lower Bounds***

50 Random 10 Pancake Problems

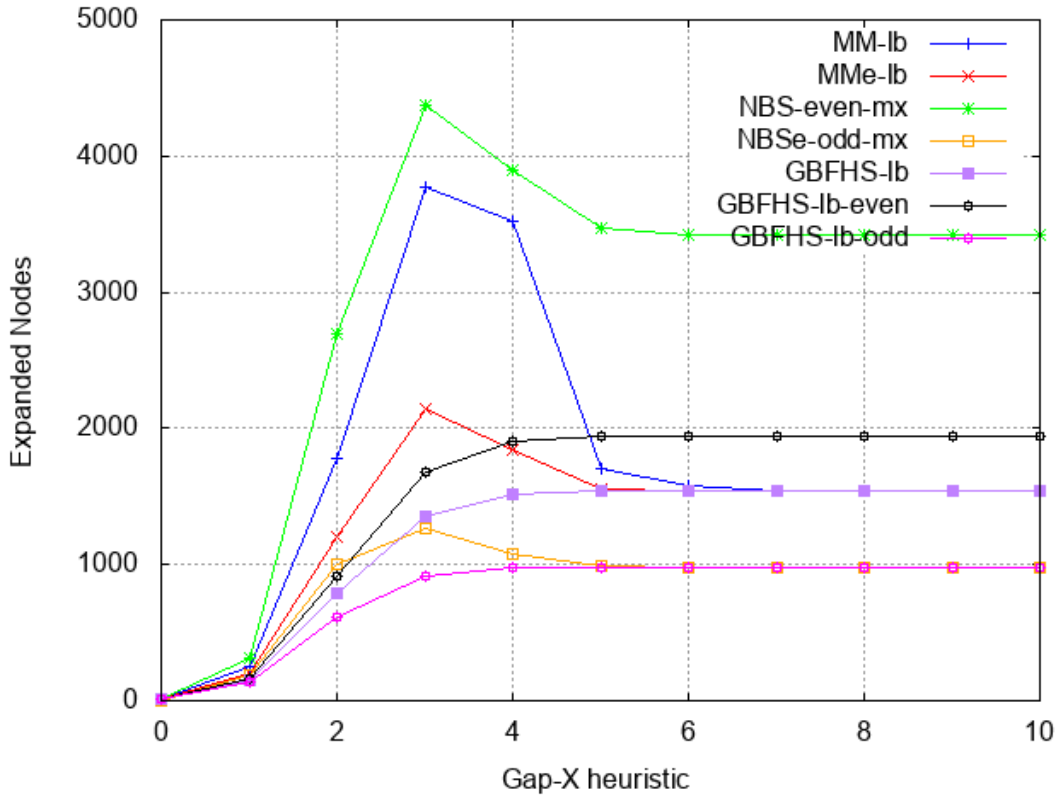

*** Don't have lower bounds for NBS/NBSe, used their "must expand" counts ⩽ lower bound*

No GBFHS has an hump

Humps for everyone else

for blind,
NBSe-odd-mx = GBFHS-lb-odd,
NBSe-even-mx = ~3500
NBSe-mx = ~2500

Unidirectional GBFHS vs A*

50 Random 10 Pancake Problems

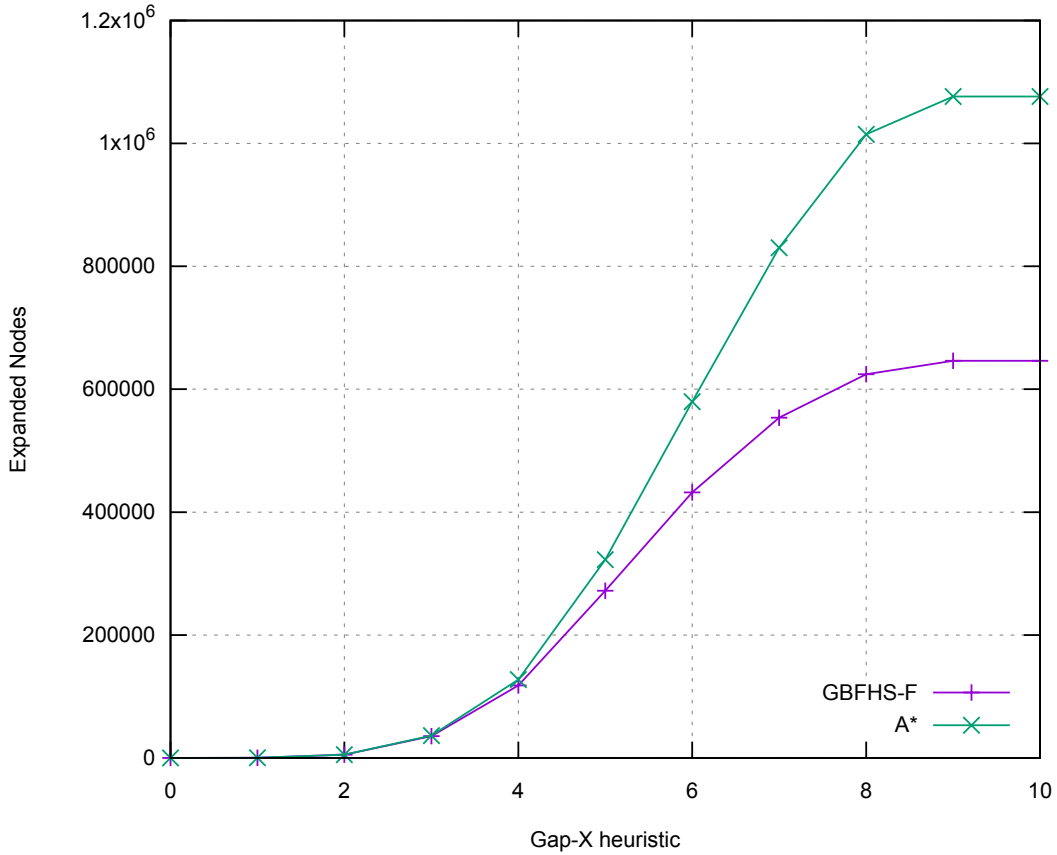

Summary

- **GBFHS**: new bidirectional heuristic search algorithm
- Split determined by input parameter
- Unifies unidirectional and bidirectional search
- First well-behaved bidirectional heuristic search algorithm
- Can perform substantially better than current algorithms
- First bidirectional heuristic search algorithm that, in uniform-cost domains, guarantees first collision results in optimal solution
- Unidirectional, can expand substantially fewer nodes than textbook A*

Future Research

- Now have bidirectional heuristic algorithm never worse than blind
- Next Challenge: Extend GBFHS so that it is “never worse” than A*
  - Automatically Finding Good Splits without losing first property
- More Domains

Questions?

The End