A Worst Case, Constant Time Priority Queue:

Beating a Lower Bound

Ian Munro

University of Waterloo

Joint work with

- Andrej Brodnik (Ljubljana & Luleå)
- Svante Carlsson (Blekinge Inst. Tech.)
- Johan Karlsson (Luleå)

Lower Bounds: What do they mean?

- If upper and lower bounds match, the problem is solved.
  - If lower bound exceeds the time you can take ... give up

Lower Bounds: What do they mean?

- If upper and lower bounds match, the problem is solved.
  - If lower bound exceeds the time you can take ... give up

- Lower bounds, and upper bounds, are proven under a model.
  - So if you have to “get under” a lower bound -focus on the operations the model does not permit
  - It’s time to become imaginative in terms of permitted operations

Beating Lower Bounds: Examples

- Searching: lg n lower bound on comparisons, so hash
  - lg n time becomes constant
- Sorting: n lg n lower bound on comparisons, so use variants of bucketing
  - n lg n time easily linear on average
  - n lg n time becomes n lg lg n even in worst case

Beating the Lower Bound ... Another Case

Self organizing linear search...

- Move to front heuristic (⇒ rheuistic) is within a constant factor of offline optimal for linear search, amortized cost of searching is ~ $1/p_i$ under "exchange adjacent" model
  - But 1 2 3 4 ... n/2 n/2+1 ... n costs $\Theta(n^2)$
- Under "exchange any two" model offline cost is ~$\lg(1/p_i)$ ... comparable to splay trees

The Problem at hand: Extended Priority Queue

- van Emde Boas (SWAT i.e. FOCS 1975)
- Universe integers [1,..m] {n of which are present}
- Operations: insert / delete
  find least value $\geq$ x (or greatest $\leq$)
- Bound: O(lg lg m) time
- Space: Improved to O(m) bits
- Model: Standard RAM, with bit twiddling

Some Subsequent Work

- Kurt Mehlhorn, Stephan Näher and Helmut Alt (SiComp '88): vEB is optimal -

  Lower bound $\Omega(\lg \lg m)$ on pointer machine
- Peter Miltersen (STOC '94):

  Lower bound $\Omega(\sqrt{\lg \lg m})$ on a RAM
- Paul Beame and Faith Fich (STOC '99):

  parameterization by number of values present

  matching upper & lower bounds- $\Theta(\sqrt{\lg n / \lg \lg n})$
- Ram model is rather powerful, how can we extend it for our problem?

Another Model: Rambo

- Random Access Machine with Byte Overlap
  Mike Fredman and Dan Willard
- Several words can share bits:

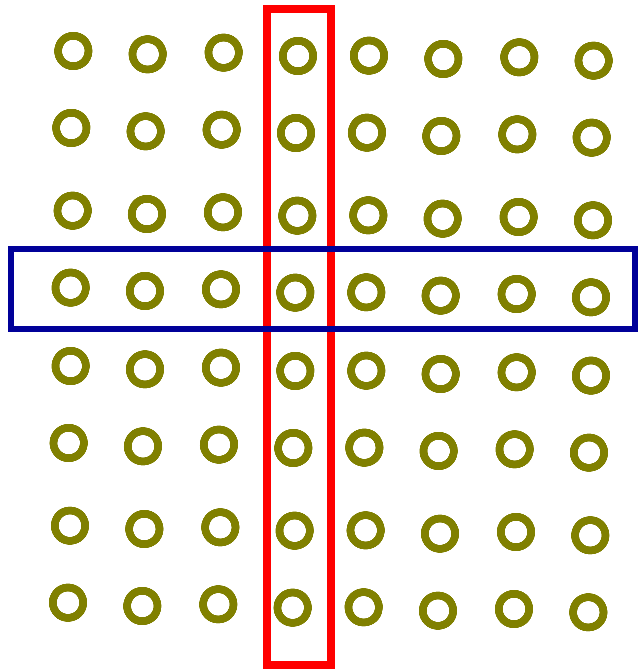

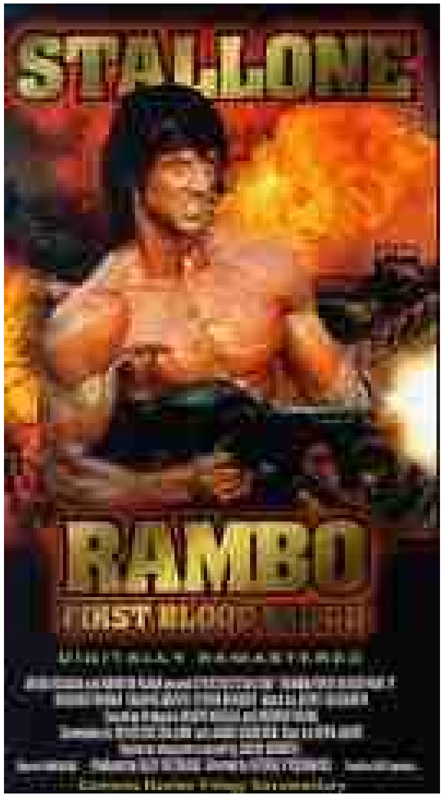

Can we do better under this model?

- Elements are at leaves; an internal node is flagged if it has a descendant
- lg n bit RAMBO word at leaf takes bit pattern of flags on path to root
- Can find lowest ancestor in O(1) time

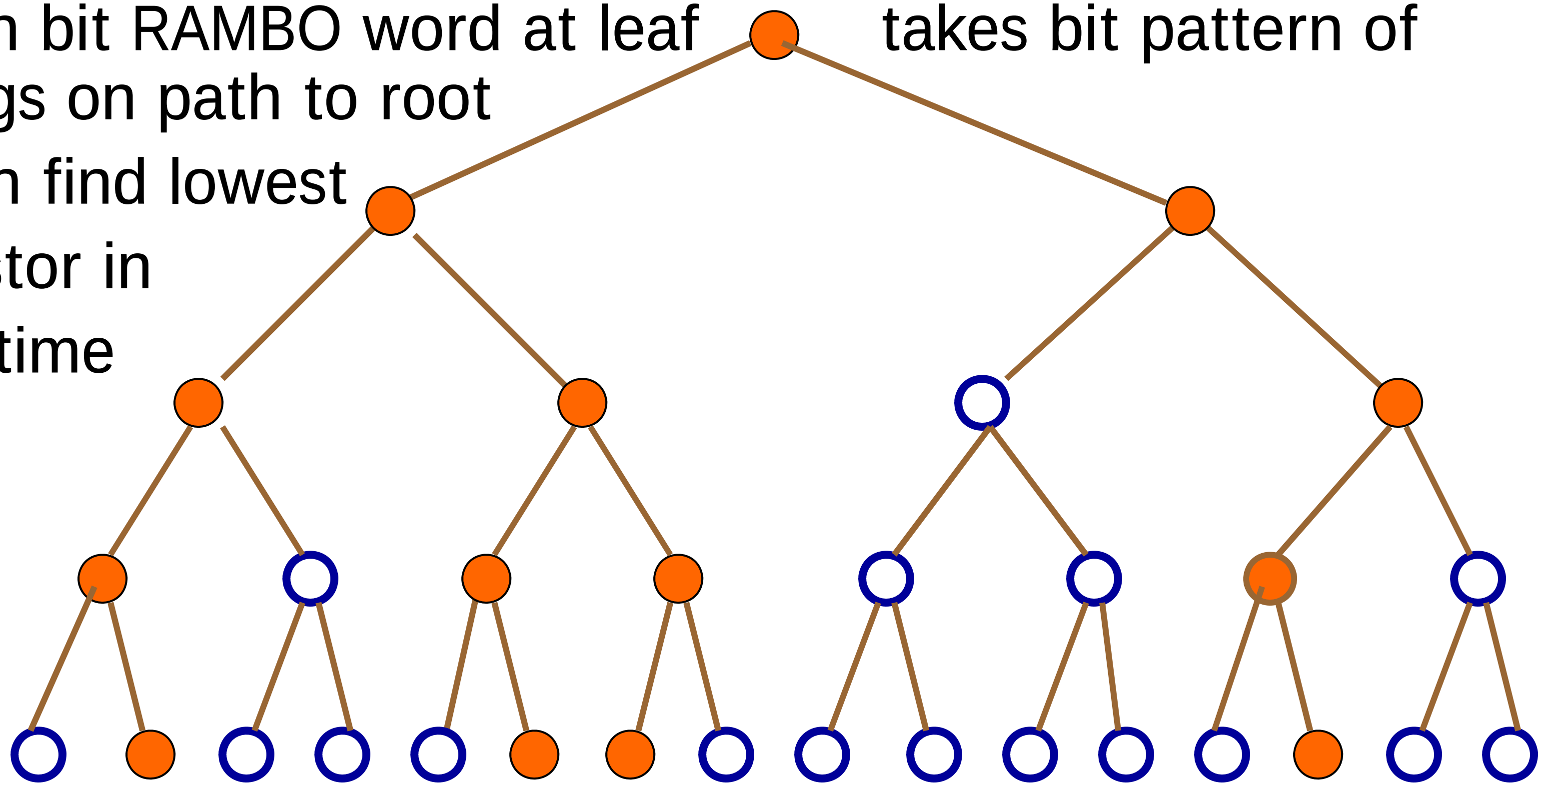

The vEB Stratified Tree

Internal node keeps track of "outside" bottom elements

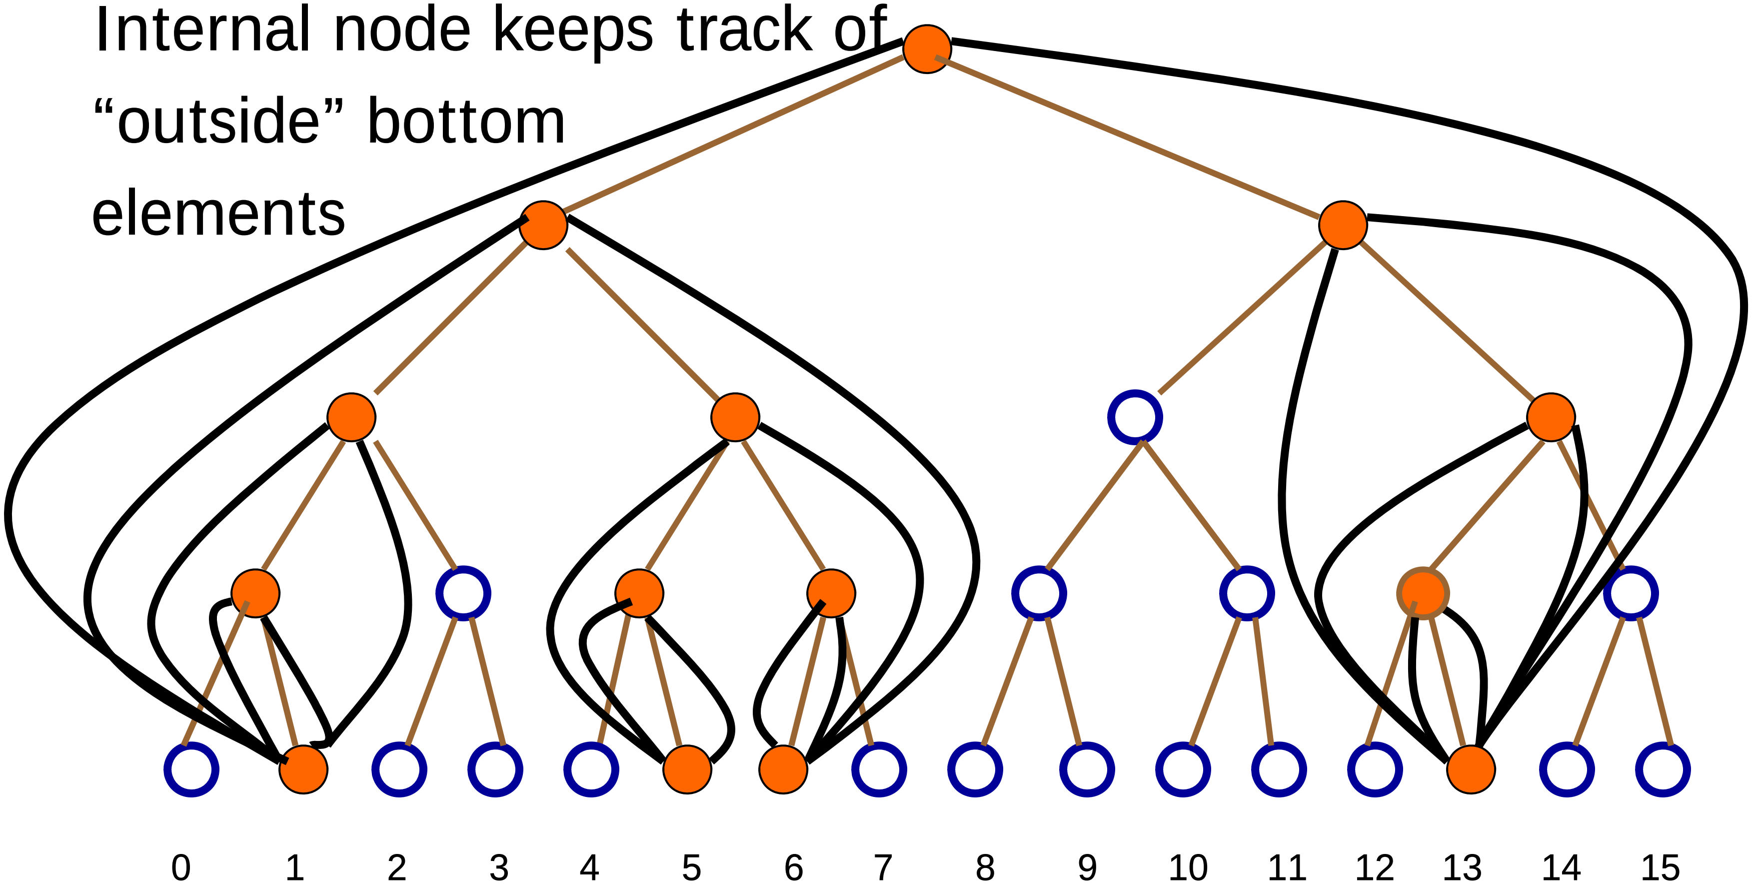

A Problem

- Any individual leaf may be referred to by many ancestors
- So one insertion/deletion can require up to 2lg n reference changes
- Modify the approach

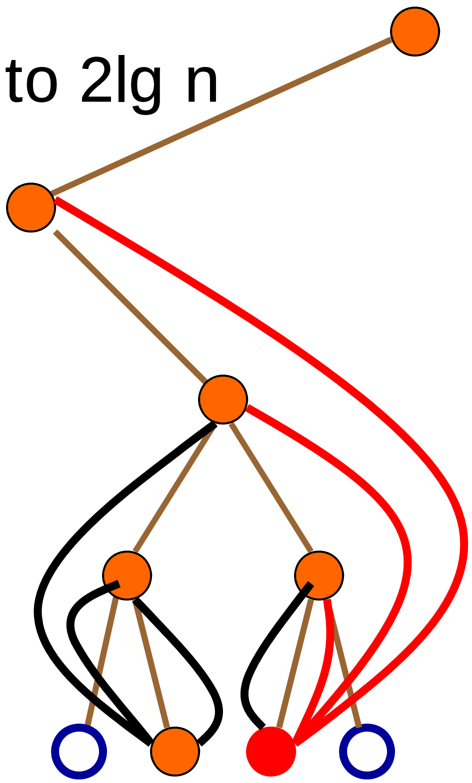

Split Tagged Tree

Keep track of left / right inside leaf ... if different from parent's

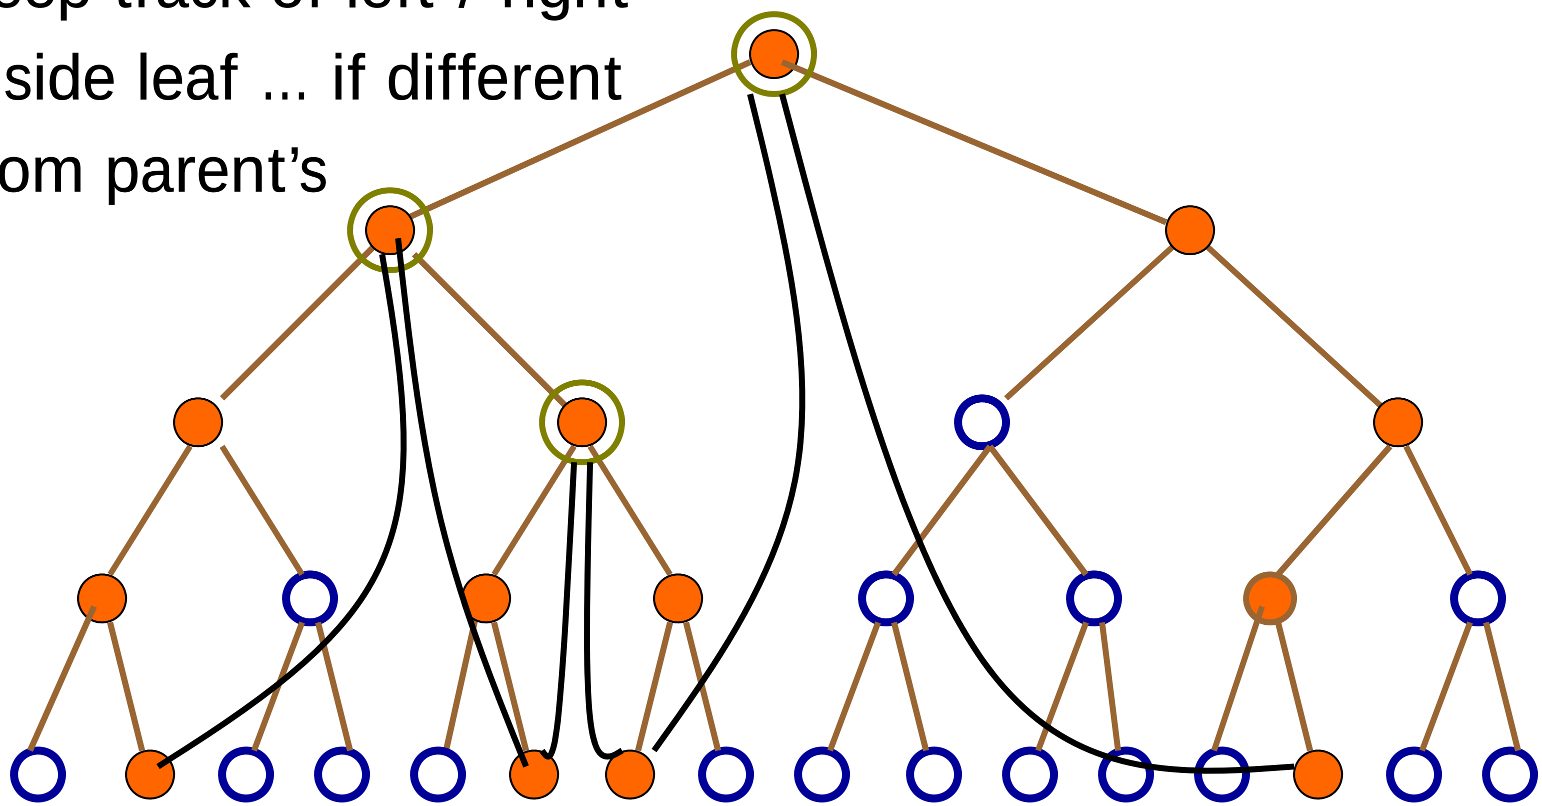

Each leaf can be referred to by at most two ancestors

So what do we have?

- Constant time “extended priority queue” ... two memory accesses for search, three for update on our model
- How much space do we need?
  - 2m + O(lg m) bits of ordinary RAM
  - m bits of RAMBO memory in a particular configuration we call Yggdrasil

- and it has been implemented in hardware