A Simple Parallel Cartesian Tree Algorithm and its Application to Parallel Suffix Tree Construction

Julian Shun and Guy Blelloch

Motivation for Suffix Trees

- To efficiently search for patterns in large texts
  - Example: Bioinformatic applications
- Suffix trees allow us to do this
  - O(N) work for construction with O(M) work for search, where N is the text size and M is the pattern size
    - In contrast, Knuth-Morris-Pratt's algorithm takes O(M) work for construction and O(N) work for search
  - Other supported operations: longest common substring, maximal repeats, longest palindrome, etc.
  - There are sequential implementations but no parallel ones that are both theoretically and practically efficient
- We developed a new (practical) linear-work parallel algorithm and analyzed it experimentally

Outline: Suffix Array to Suffix Tree (in parallel)

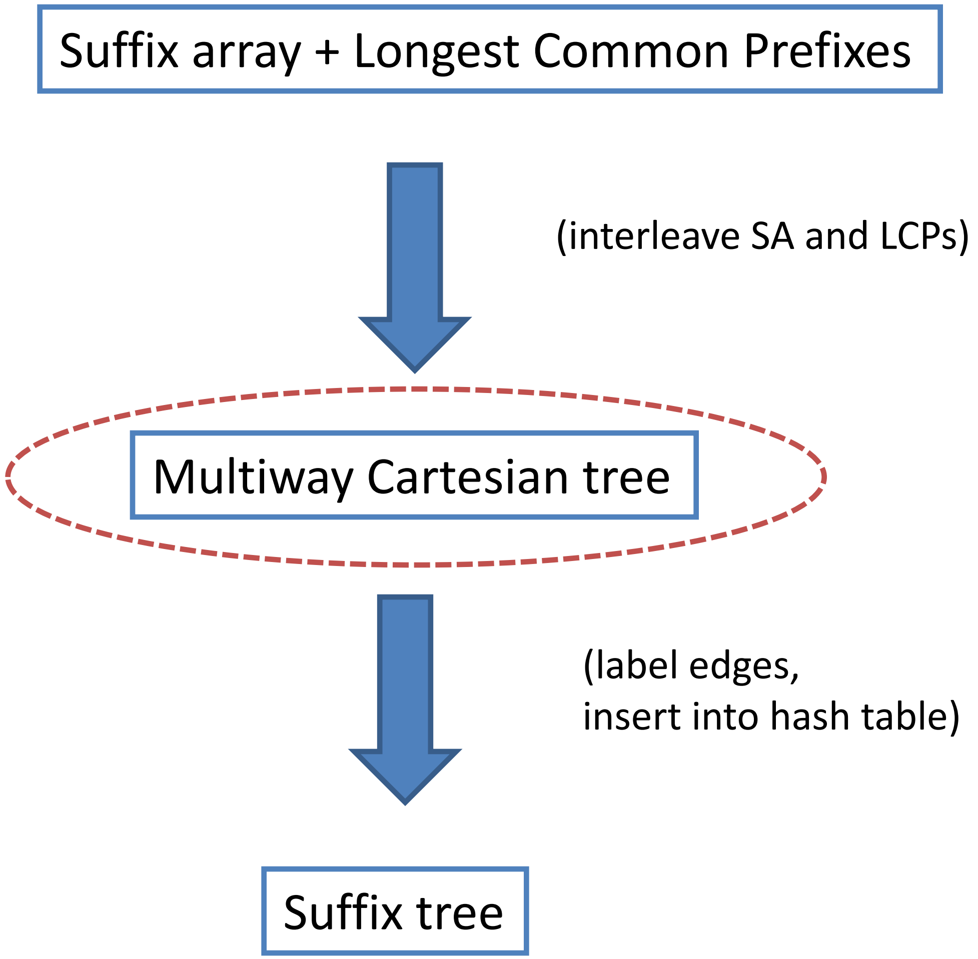

- There are standard techniques to perform all of these steps in parallel, except for building the multiway Cartesian Tree

Suffix Arrays and Longest-common-prefixes (LCPs)

Original String	Suffixes	Suffix array	LCPs
mississippi$	mississippi$	$	
	ississippi$	i$	0
	ssissippi$	ippi$	1
	sissippi$	issippi$	1
	issippi$	ississippi$	4
	ssippi$	mississippi$	0
	sippi$	pi$	0
	ippi$	ppi$	1
	ppi$	sippi$	0
	pi$	sissippi$	2
	i$	ssippi$	1
	$	ssissippi$	3

Sort suffixes

Suffix Trees

- String = mississippi$
- Store suffixes in a patricia tree (trie with one-child nodes collapsed)

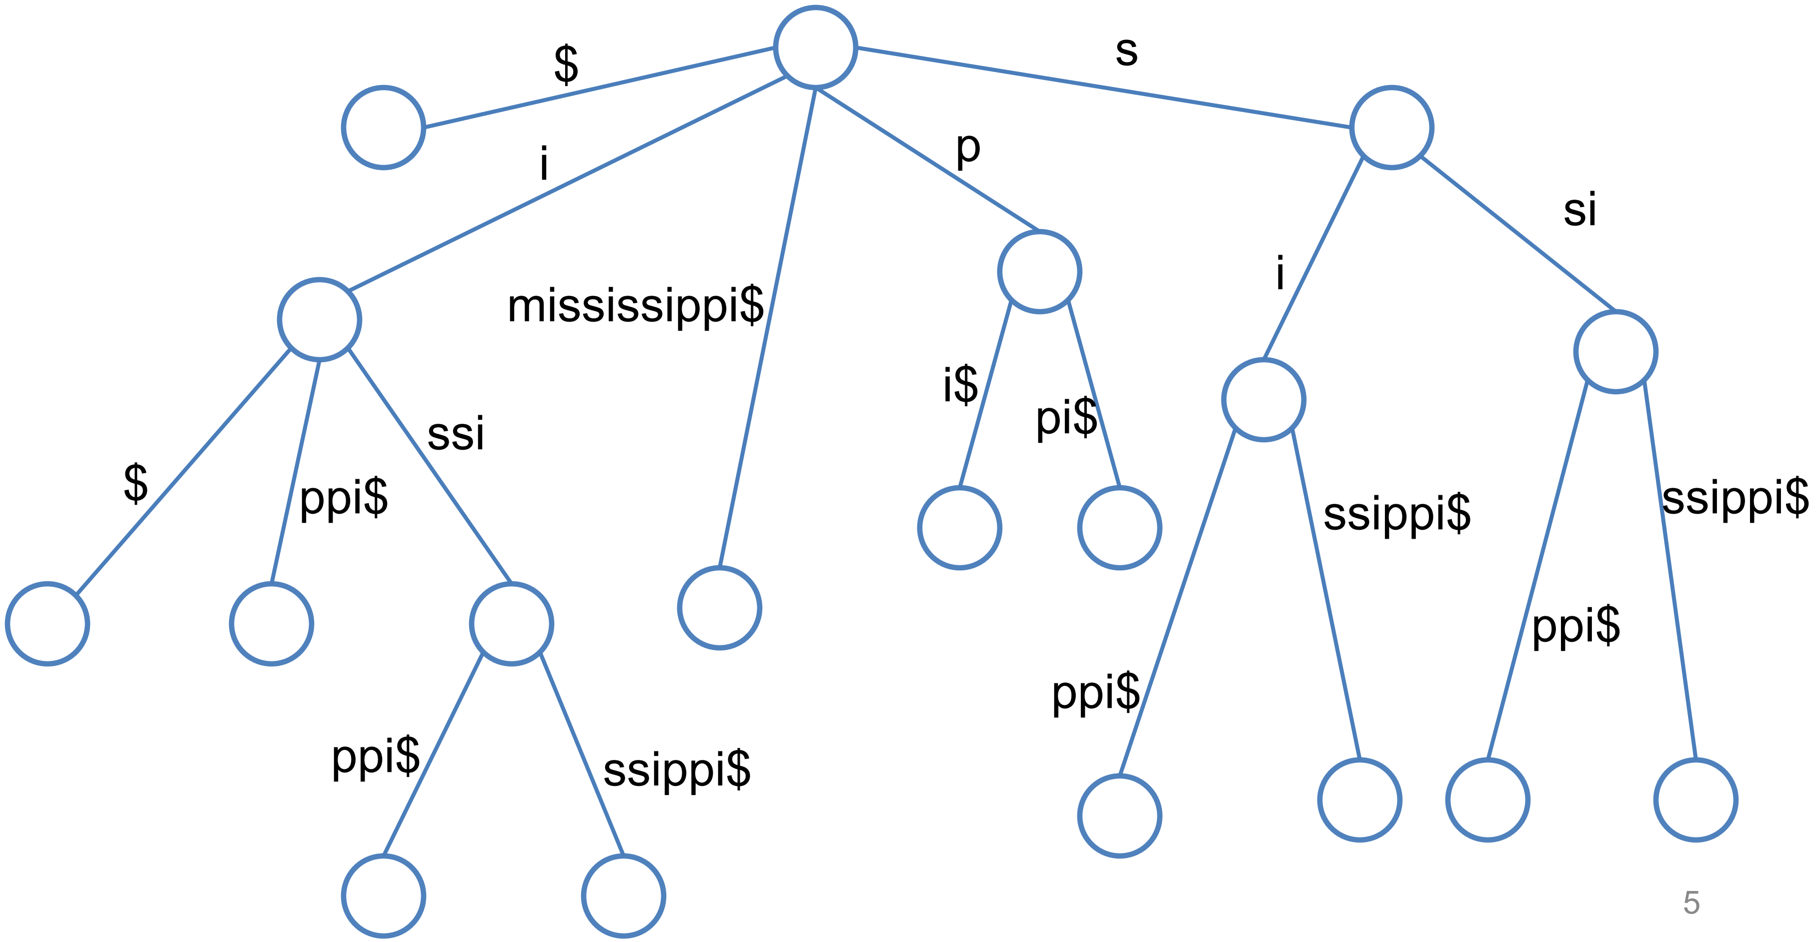

Multiway Cartesian Tree

- Maintains heap property
- Inorder traversal gives back the sequence

- Components of same value treated as one "cluster"

Sequence = 1 2 0 4 1 1 3 1 2

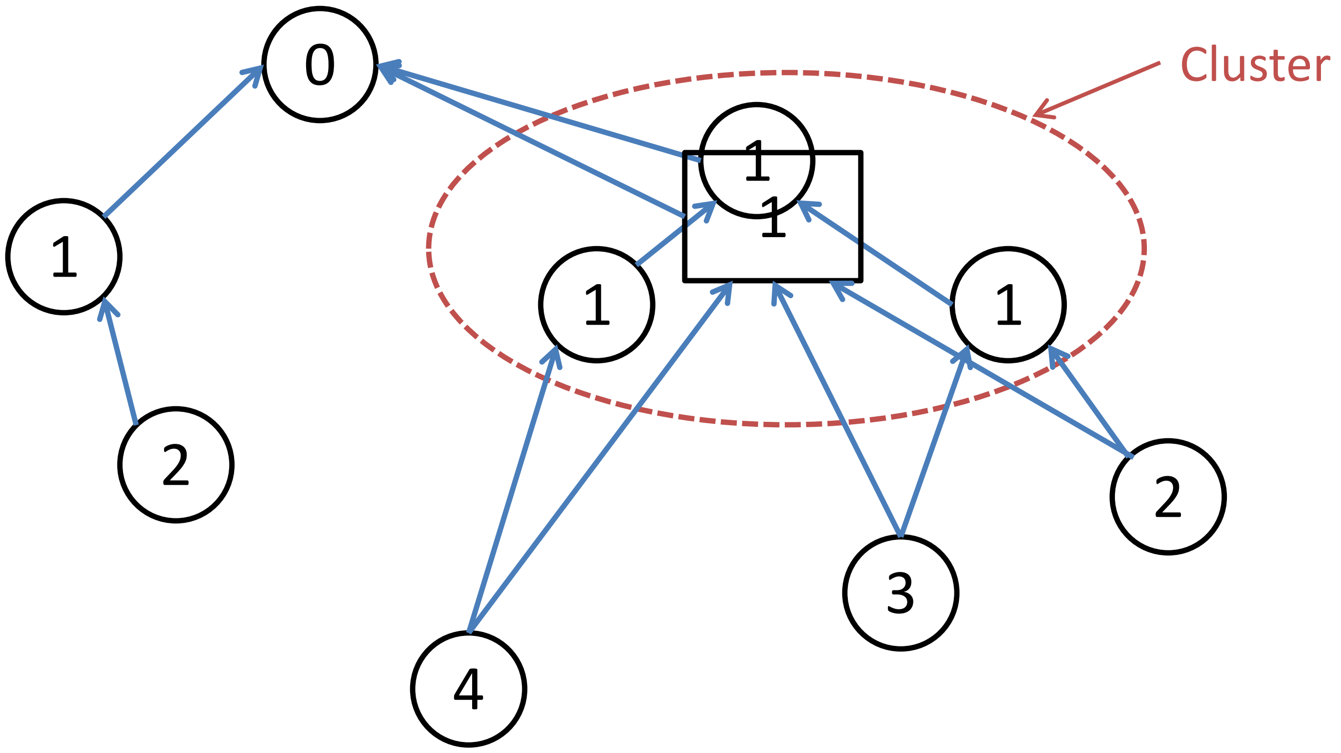

Suffix Tree History

- Sequential O(n) work algorithms based on incrementally adding suffixes [Weiner '73, McCreight '76, Ukkonen '95]
- Parallel O(n) work algorithms very complicated, no implementations [Sahinalp-Vishkin '94, Hariharan '94, Farach-Muthukrishnan '96]
- Parallel algorithms used in practice are not linear-work
- *Practical linear-work parallel algorithm?*
    - Simple O(n) work parallel algorithm
    - Fastest algorithm in practice

More Related Work

- Cartesian trees
  - Sequential O(n) work stack-based algorithm
  - Work-optimal parallel algorithm for Cartesian tree on distinct values (Berkman, Schieber and Vishkin 1993)
- Suffix arrays to suffix trees
  - Sequential O(n) work algorithms
  - Two parallel algorithms for converting a suffix array into a suffix tree (Iliopoulos and Rytter 2004)
    - Both require O(n log n) work
- Our contributions
  - A parallel algorithm for converting suffix arrays to suffix trees, which requires only O(n) work and is based on multiway Cartesian trees

Suffix Array/LCPs → Suffix Tree

- Interleave suffix lengths and LCP values
- Build a multiway Cartesian tree on that
- This returns the suffix tree!

Suffix lengths	1,	2,	5,	8,	11,	12,	3,	4,	6,	9,	7,	10
LCP values	0,	1,	1,	4,	0,	0,	1,	0,	2,	1,	3,	

Interleaved

Multiway Cartesian Tree = Suffix Tree

Suffix array	Lengths	LCPs
$	1	
i$	2	0
ippi$	5	1
issippi$	8	1
ississippi$	11	4
mississippi$	12	0
pi$	3	0
ppi$	4	1
sippi$	6	0
sissippi$	9	2
ssippi$	7	1
ssissippi$	10	3

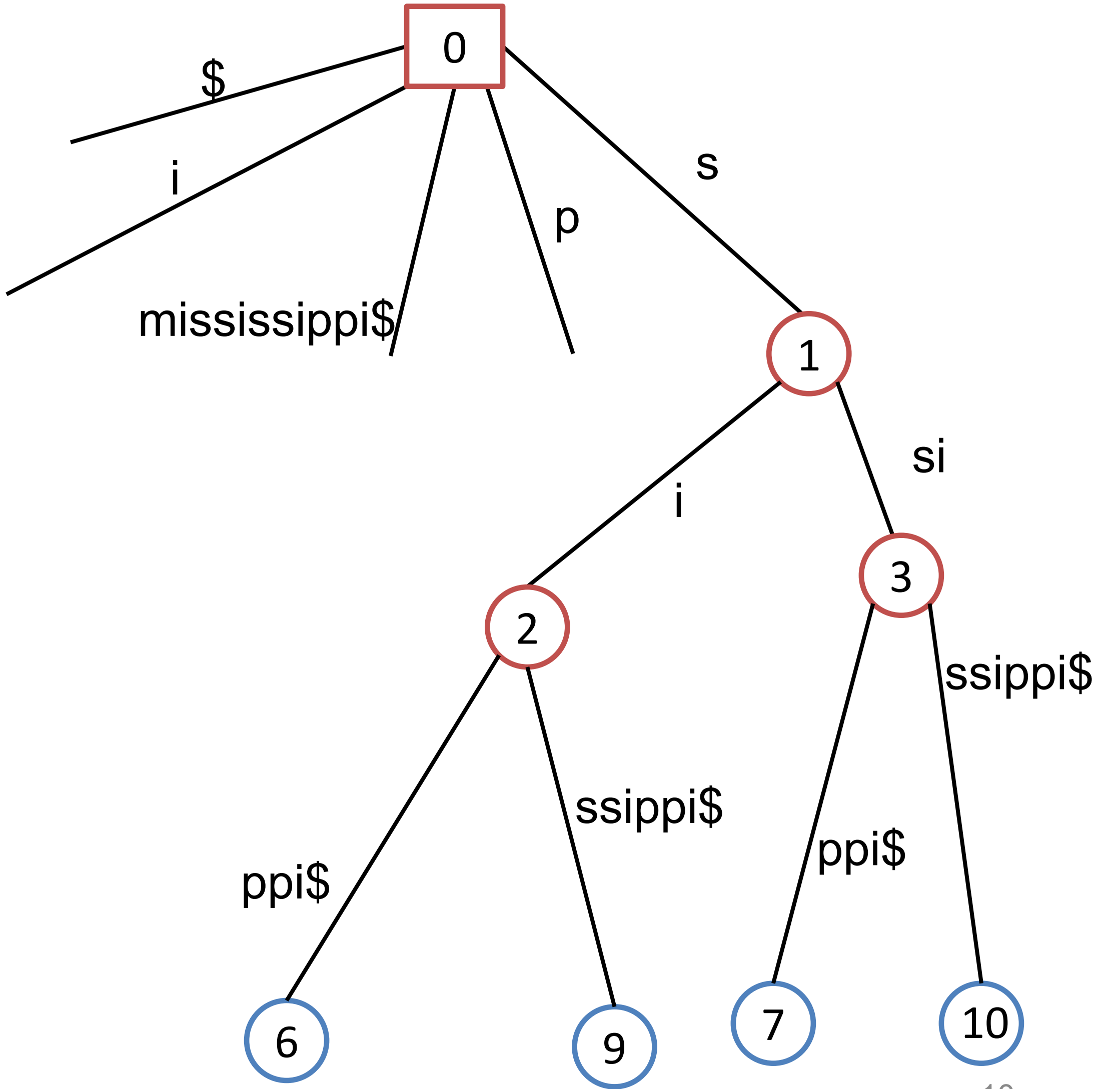

String = mississippi$

□ = Contracted internal node with LCP value

○ = Leaf node with suffix length

○ = Internal node with LCP value

SA + LCPs (interleaved) = 1, 0, 2, 1, 5, 1, 8, 4, 11, 0, 12, 0, 3, 1, 4, 0, 6, 2, 9, 1, 7, 3, 10

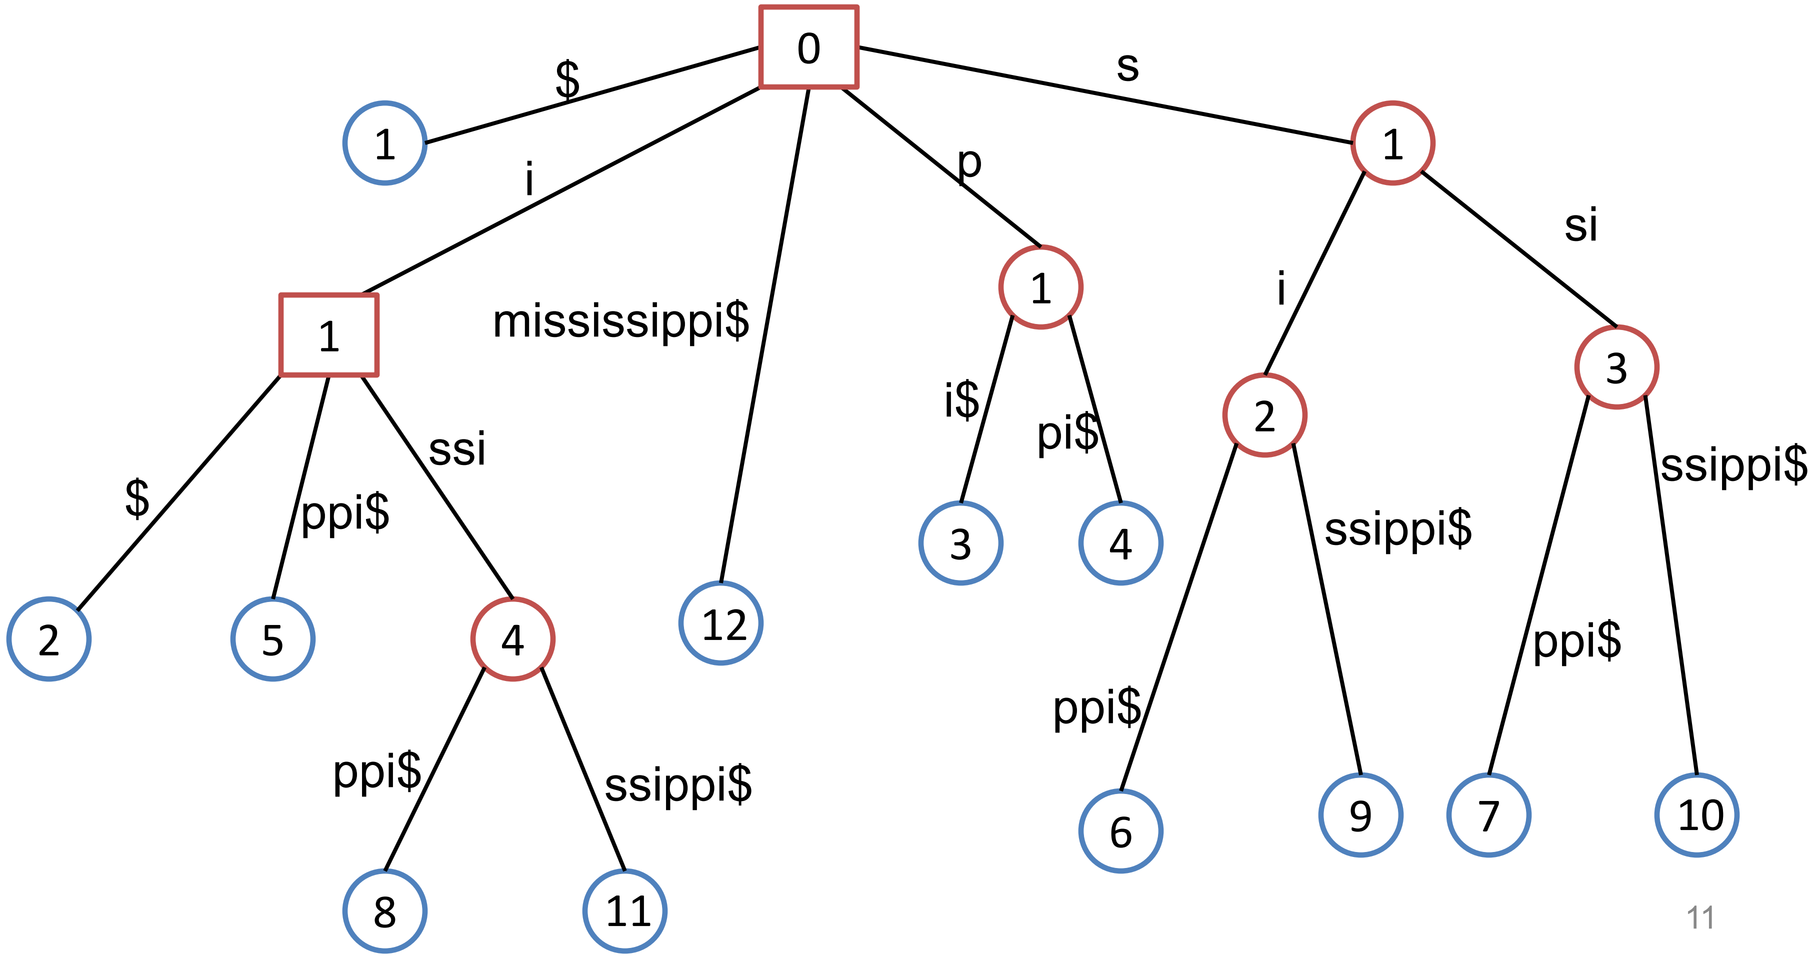

Suffix Array to Suffix Tree (in parallel)

Suffix array + Longest Common Prefixes ✓ Karkkainen and Sander's algorithm O(n) work and O($\log^2$ n) span

↓ (interleave SA and LCPs) ✓

Multiway Cartesian tree ⬅

↓ (label edges, insert into hash table)

Suffix tree

Cartesian Tree (in parallel)

- Divide-and-conquer approach
- Merge spines of subtrees (represented as lists) together using standard techniques

SA + LCPs =

1, 0, 2, 0, 5, 1, 8, 1, 11, 4, 12, 0, 3, | 0, 4, 1, 6, 0, 9, 2, 8, 1, 7, 3, 10

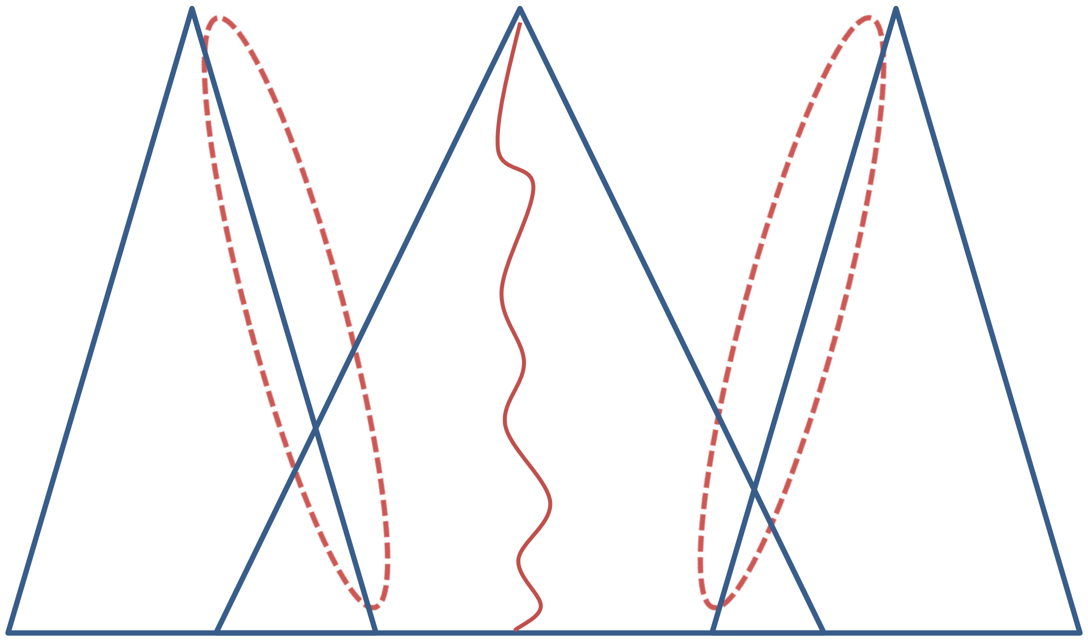

Left subtree Merged tree Right subtree

Cartesian Tree (in parallel)

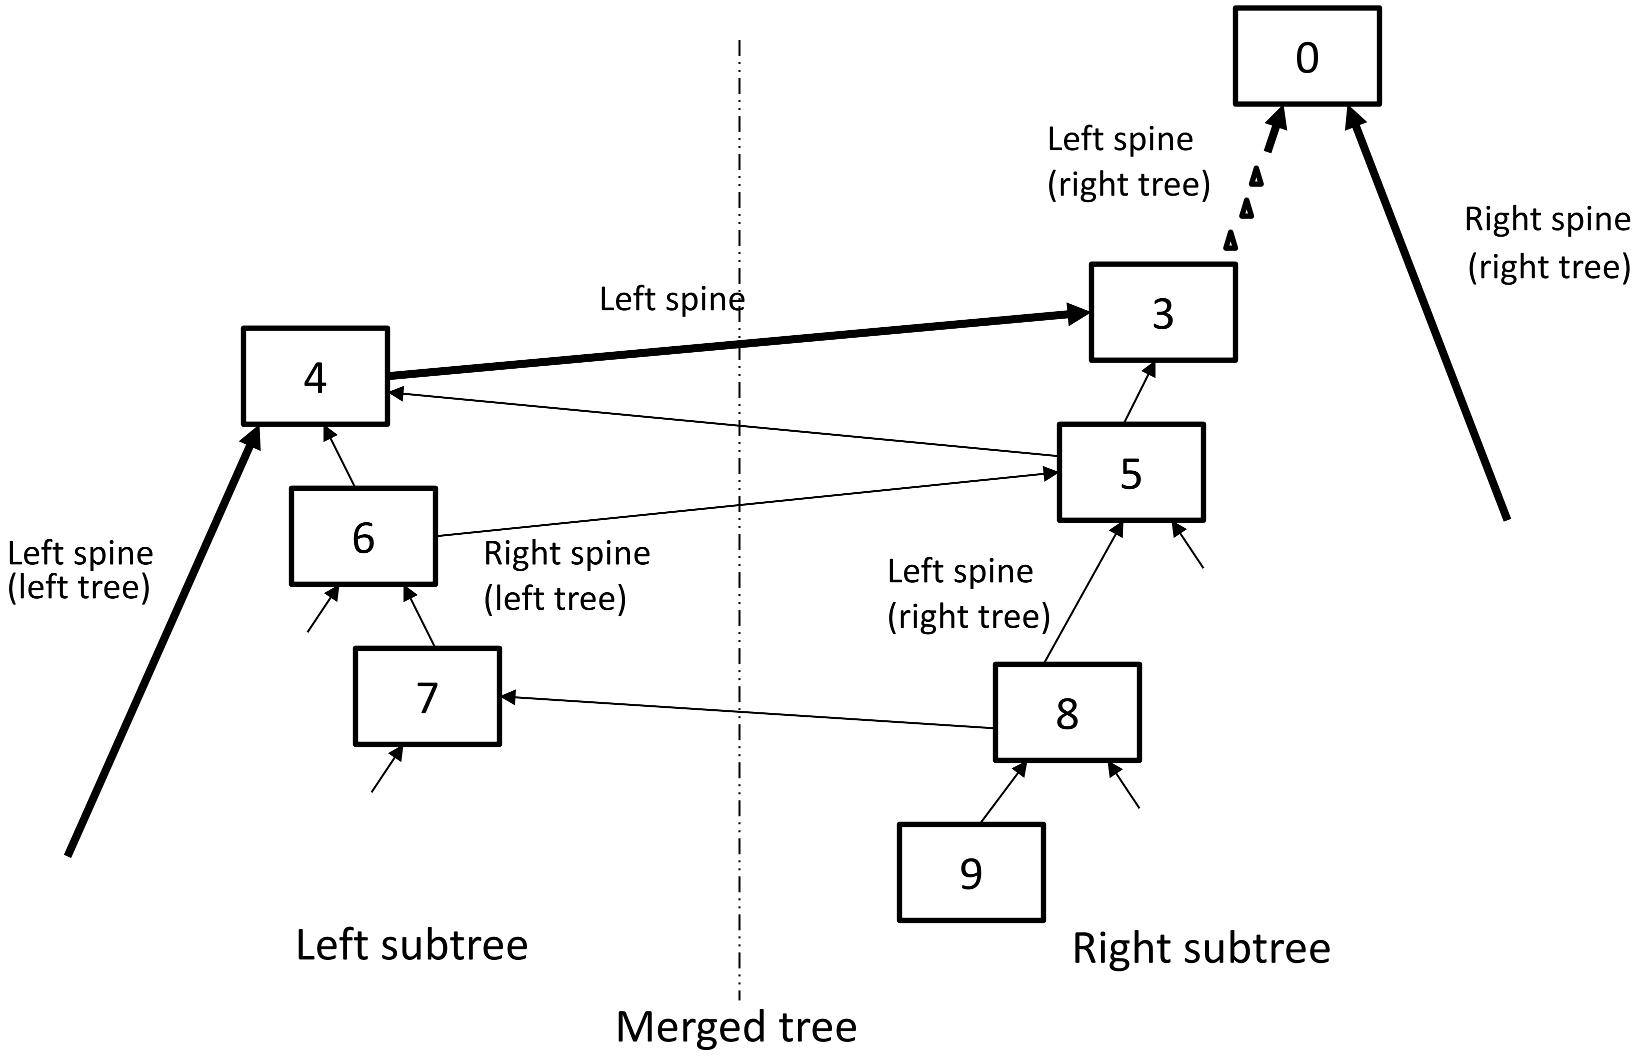

Cartesian Tree (in parallel)

- Input: Array A[1…N]

```
Build(A[1…n]){
    if n < 2 return;
    else in parallel do:
            t1 = Build(A[1…n/2]);
            t2 = Build(A[(n/2)+1…n]);
    Merge(t1, t2);
}
```

```
Merge(t1, t2){
    R-spine = rightmost branch of t1;
    L-spine = leftmost branch of t2;
    use a parallel merge algorithm
    on R-spine and L-spine;
}
```

String = mississippi$

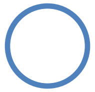 = Leaf node with suffix length 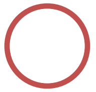 = Internal node with LCP value

SA + LCPs (interleaved) = 1, 0, 2, 1, 5, 1, 8, 4, 11, 0, 12, 0, 3, 1, 4, 0, 6, 2, 9, 1, 7, 3, 10

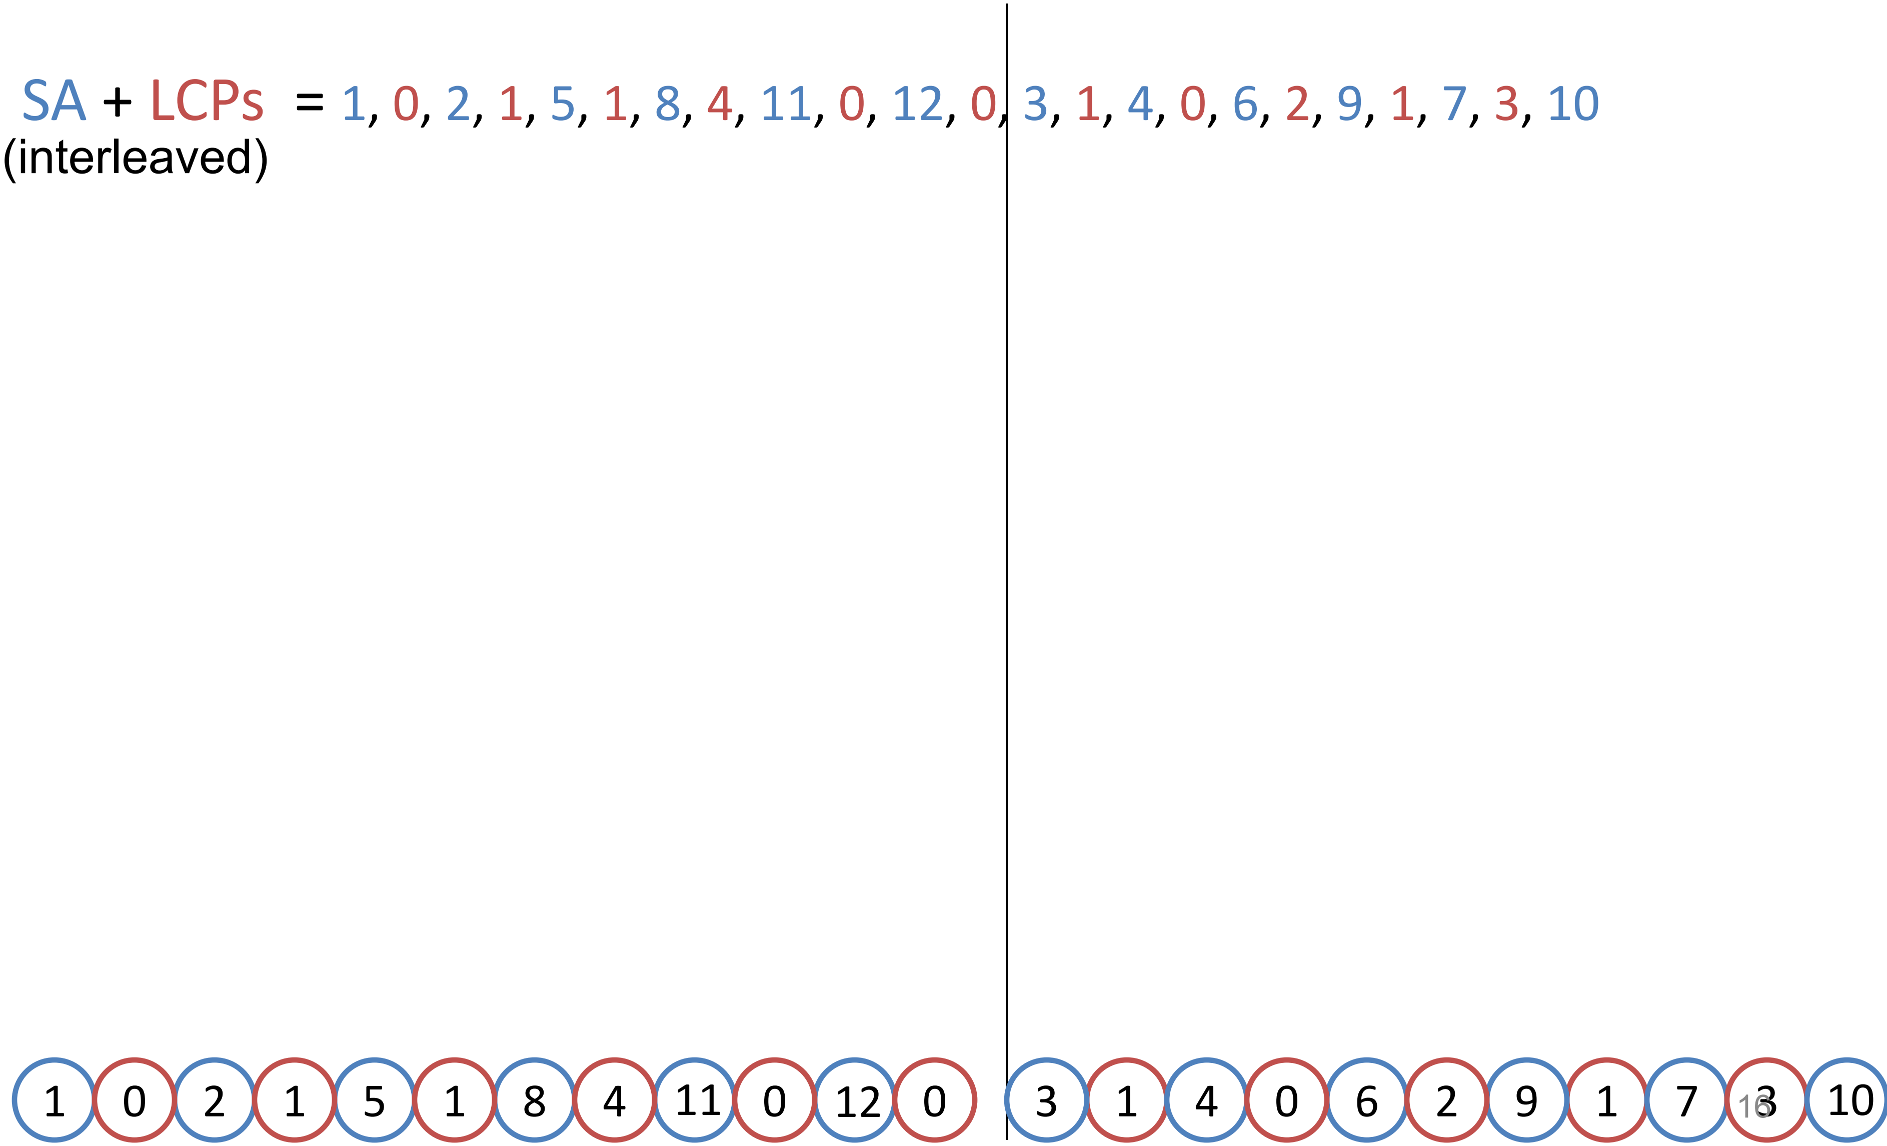

String = mississippi$

○ = Leaf node with suffix length ○ = Internal node with LCP value

SA + LCPs (interleaved) = 1, 0, 2, 1, 5, 1, 8, 4, 11, 0, 12, 0, 3, 1, 4, 0, 6, 2, 9, 1, 7, 3, 10

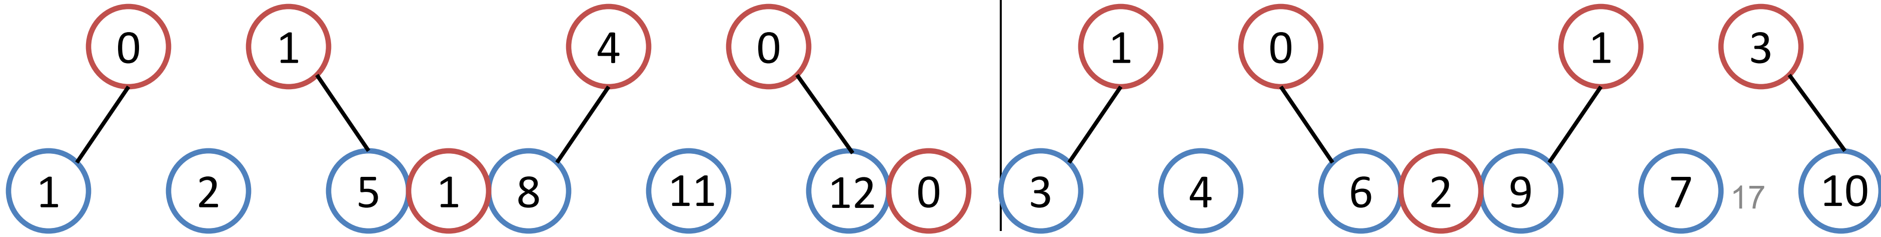

String = mississippi$

◯ = Leaf node with suffix length

◯ = Internal node with LCP value

SA + LCPs (interleaved) = 1, 0, 2, 1, 5, 1, 8, 4, 11, 0, 12, 0, 3, 1, 4, 0, 6, 2, 9, 1, 7, 3, 10

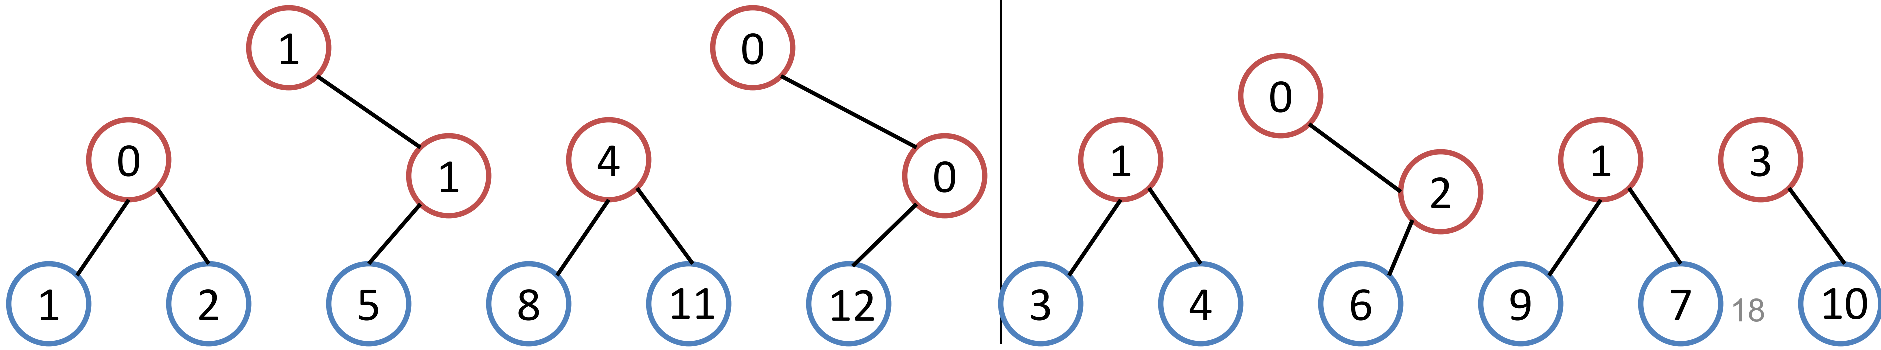

String = mississippi$

◯ = Leaf node with suffix length   ◯ = Internal node with LCP value

SA + LCPs (interleaved) = 1, 0, 2, 1, 5, 1, 8, 4, 11, 0, 12, 0, 3, 1, 4, 0, 6, 2, 9, 1, 7, 3, 10

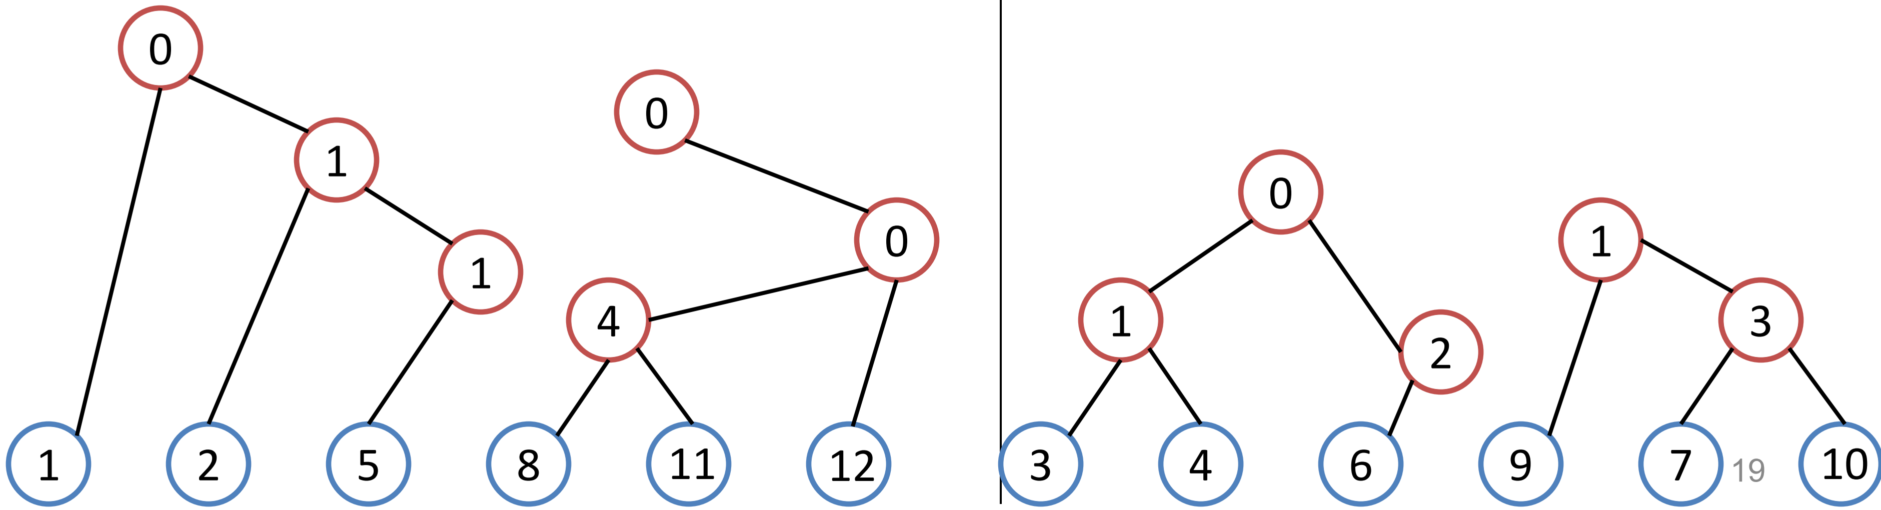

String = mississippi$

◯ = Leaf node with suffix length   ◯ = Internal node with LCP value

SA + LCPs (interleaved) = 1, 0, 2, 1, 5, 1, 8, 4, 11, 0, 12, 0, 3, 1, 4, 0, 6, 2, 9, 1, 7, 3, 10

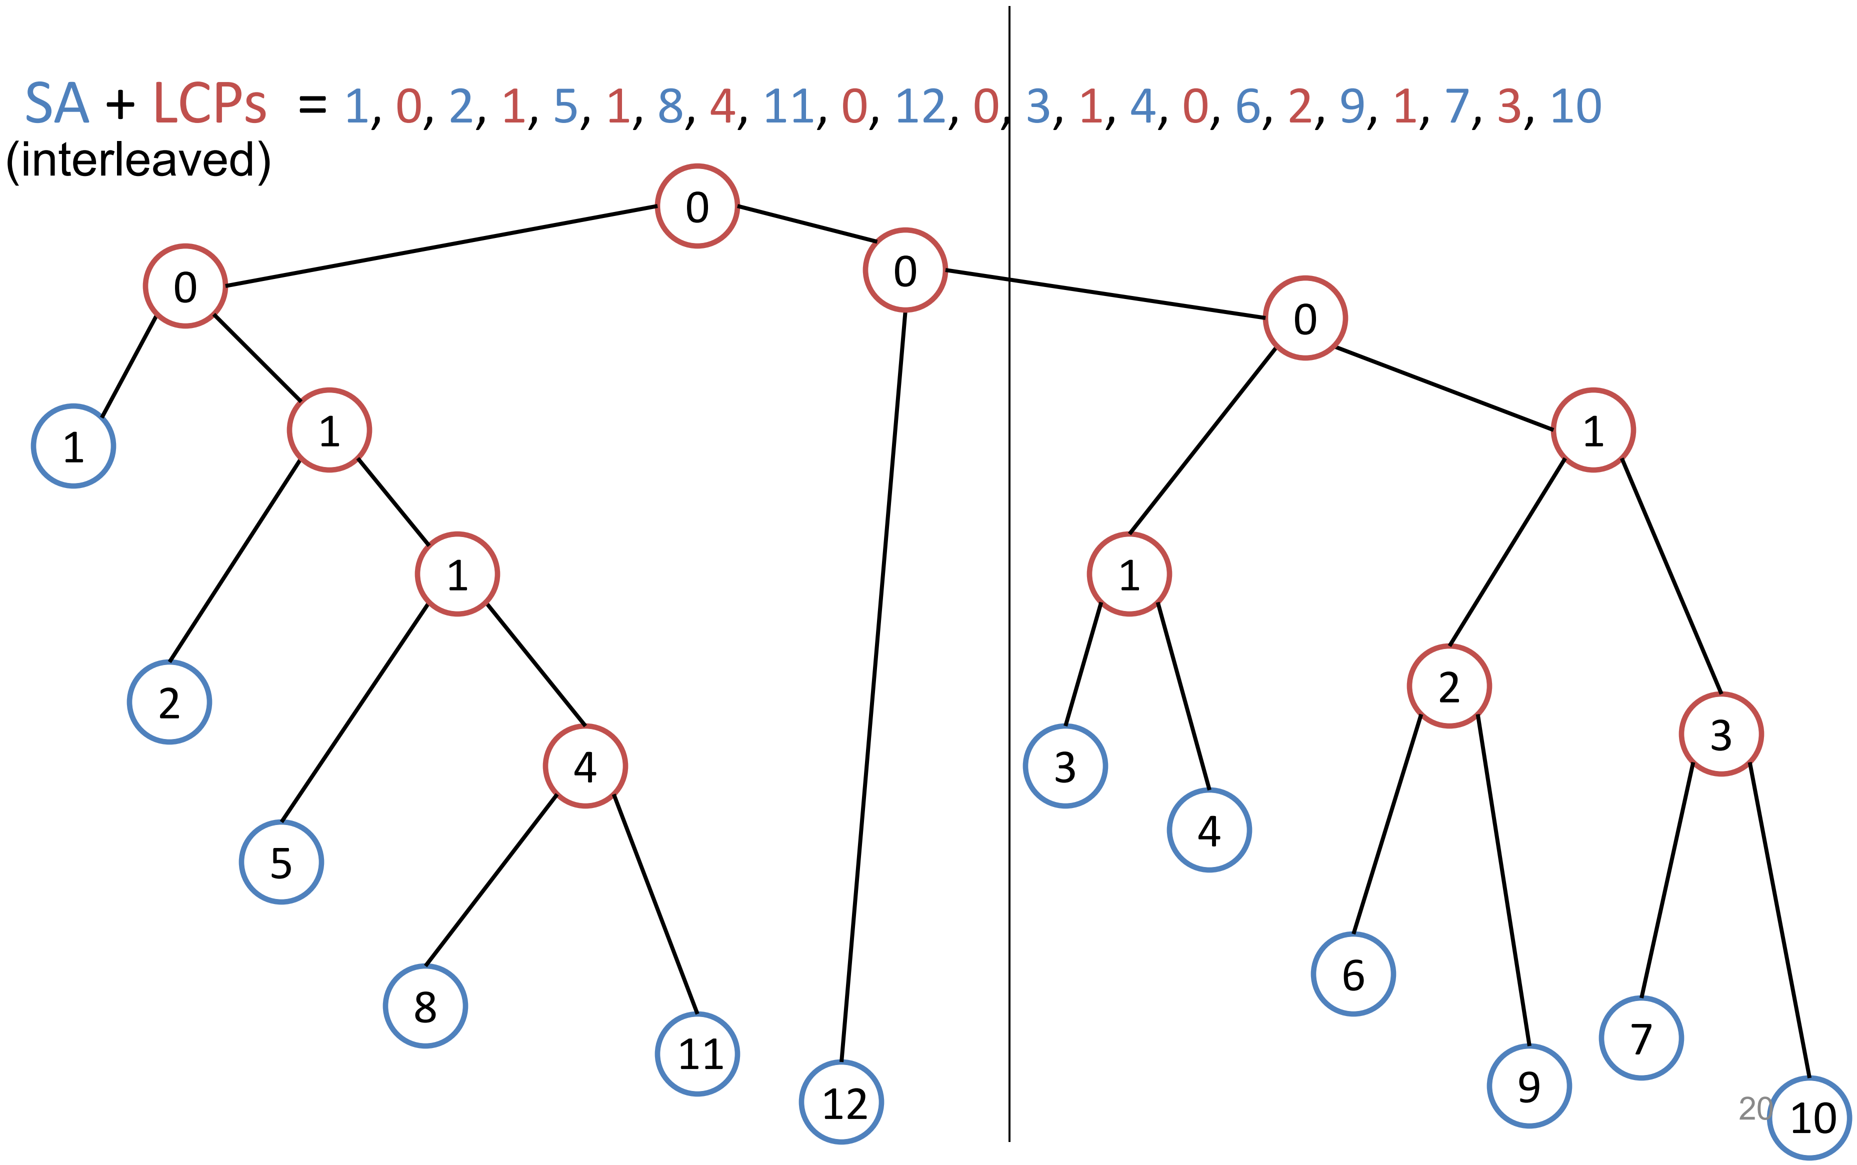

String = mississippi$

◯ = Leaf node with suffix length ◯ = Internal node with LCP value

SA + LCPs = 1, 0, 2, 1, 5, 1, 8, 4, 11, 0, 12, 0, 3, 1, 4, 0, 6, 2, 9, 1, 7, 3, 10
(interleaved)

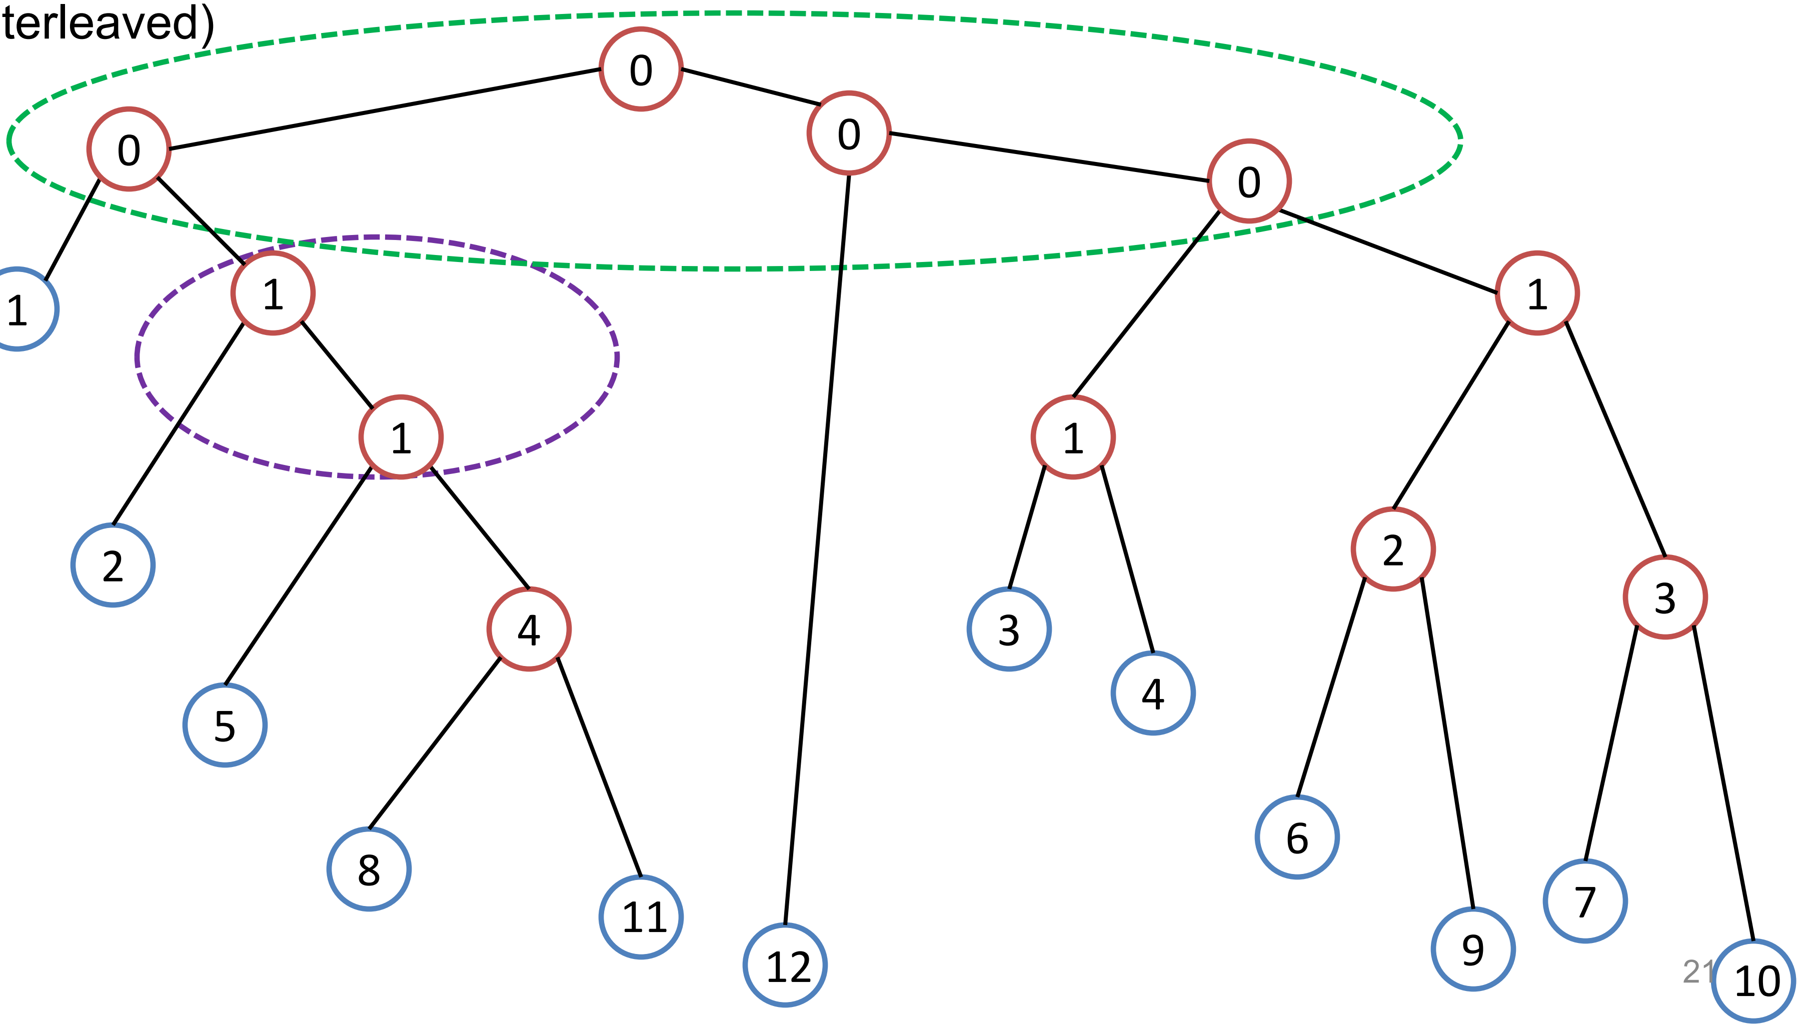

String = mississippi$

□ = Contracted internal node with LCP value

○ = Leaf node with suffix length

○ = Internal node with LCP value

SA + LCPs (interleaved) = 1, 0, 2, 1, 5, 1, 8, 4, 11, 0, 12, 0, 3, 1, 4, 0, 6, 2, 9, 1, 7, 3, 10

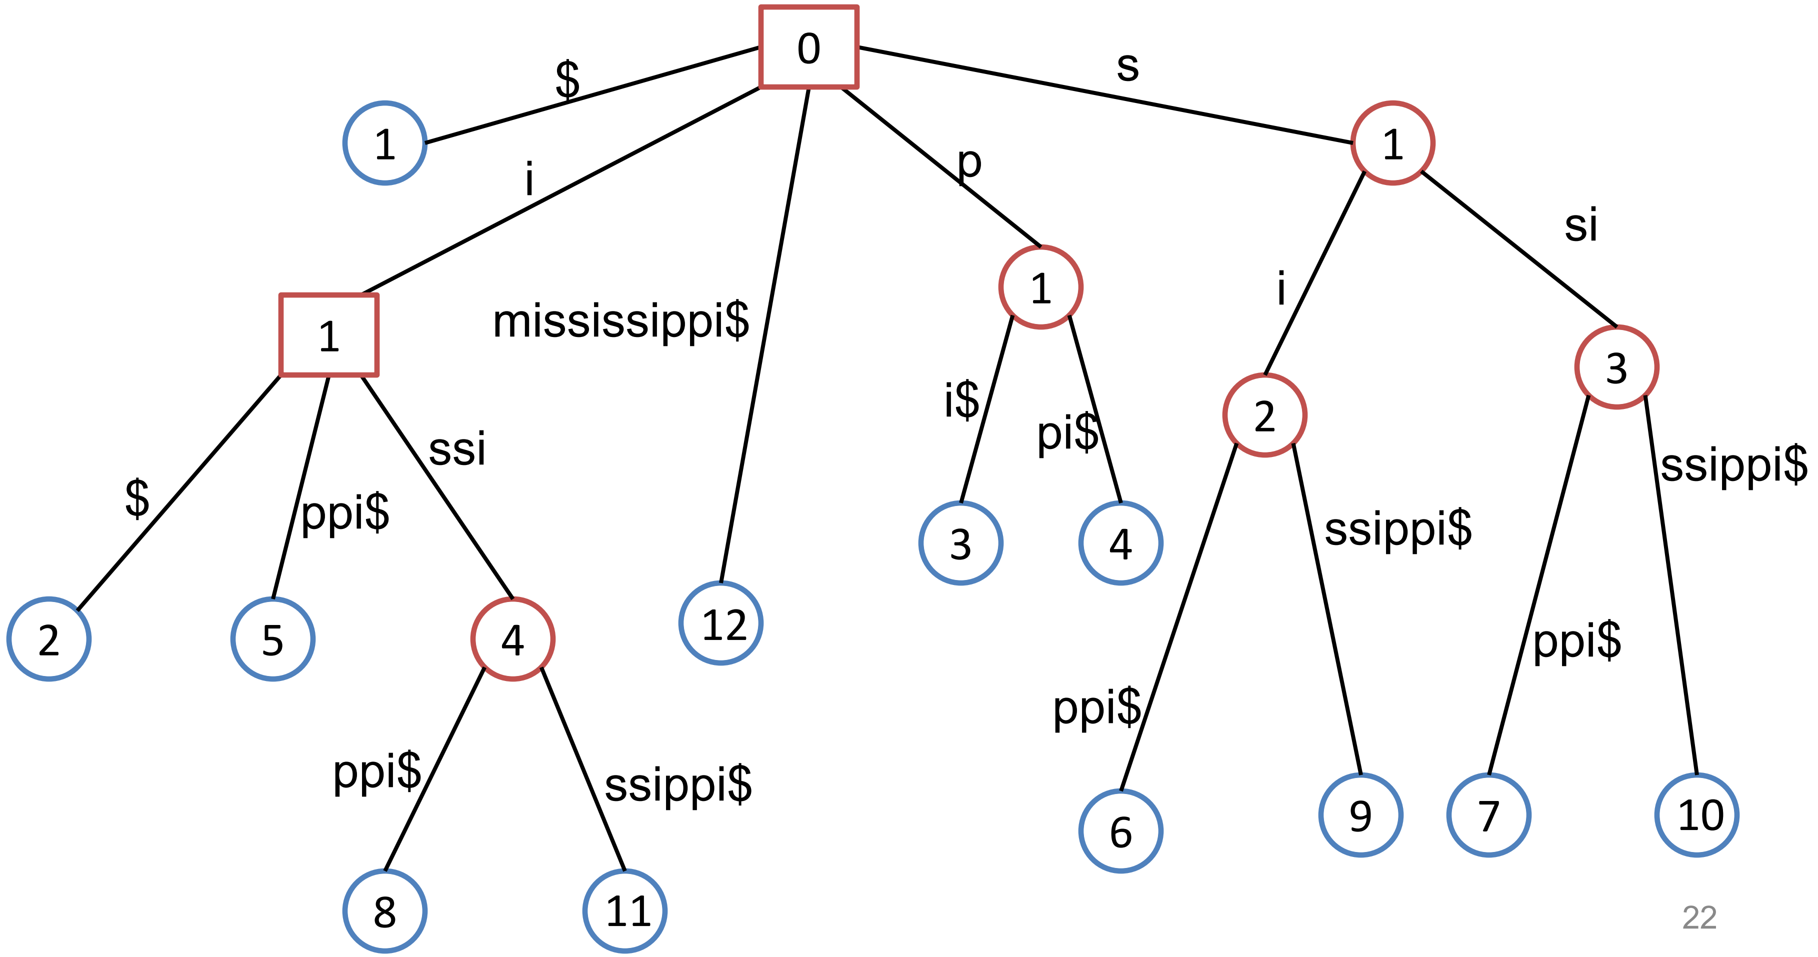

Cartesian Tree (in parallel)

- Almost all merged nodes will never be processed again (they are "protected")

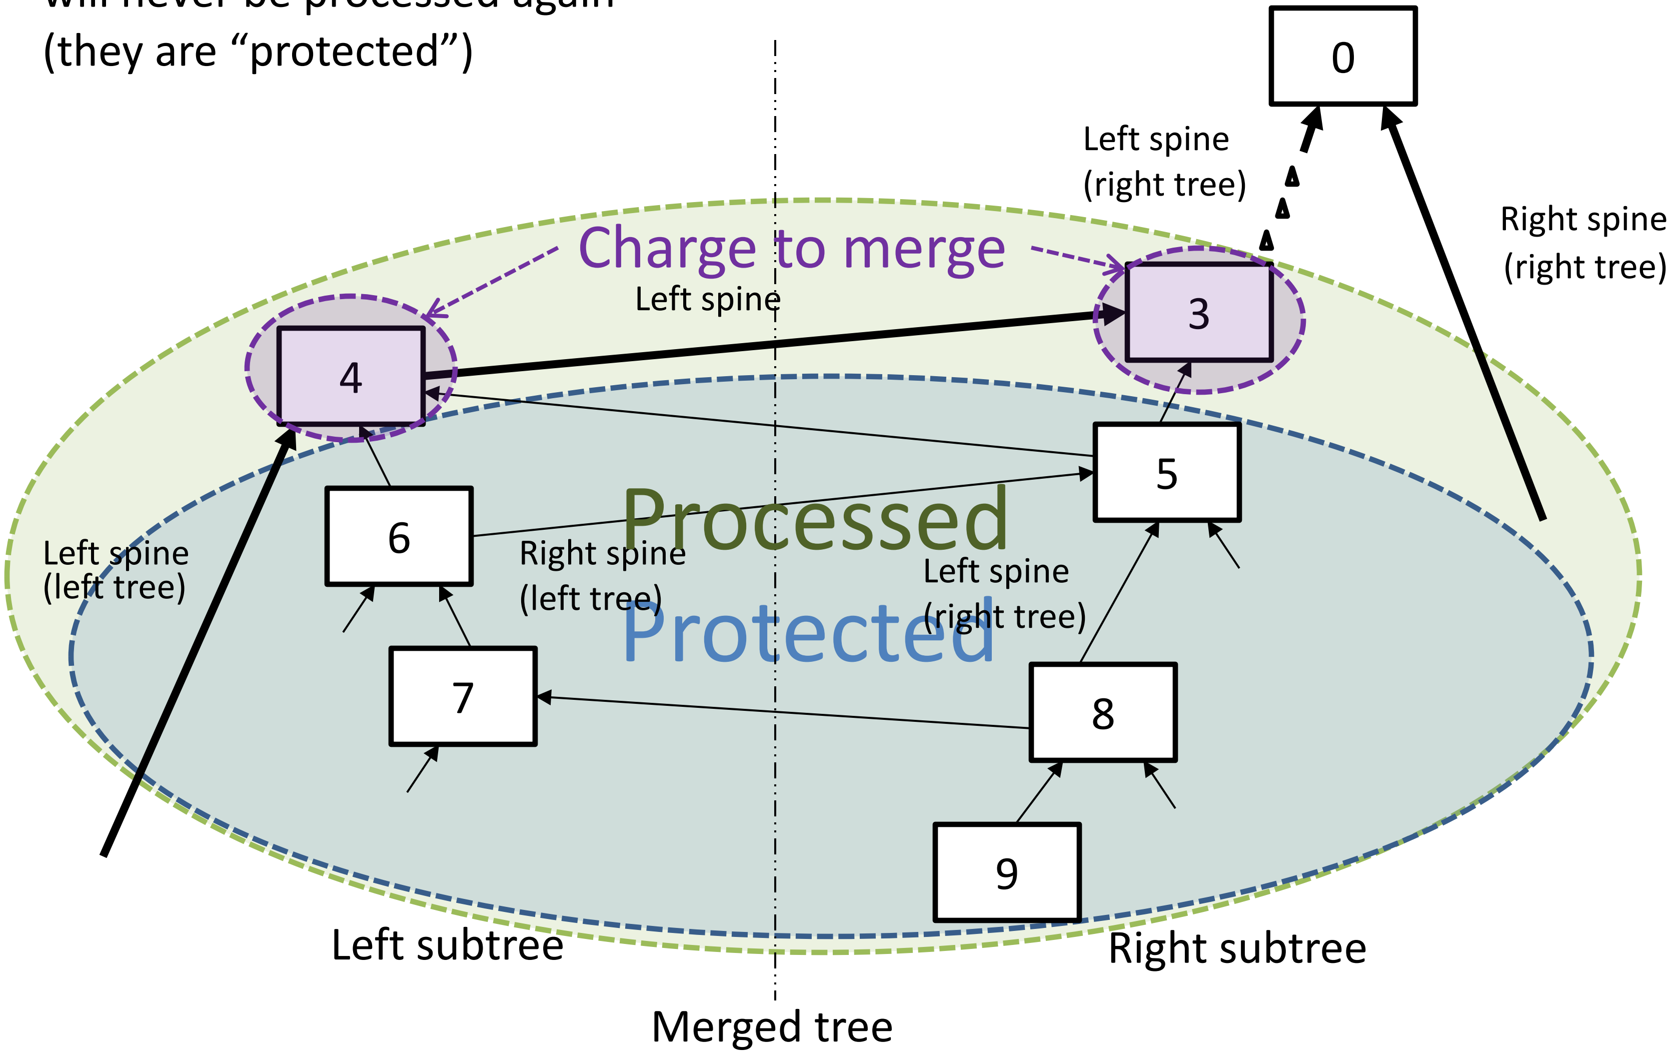

Cartesian Tree - Complexity bounds

- Observation: All nodes processed, except for two, become protected during a merge.
- Charge the processing of those two nodes to the merge itself (there are only O(n) merges). Other nodes pay for themselves and then get protected.
  - It is important that when one spine has been completely processed, the merge does not process the rest of the other spine, otherwise we get O(n log n) work
- Therefore, the merges contribute a total of O(n) work to the algorithm

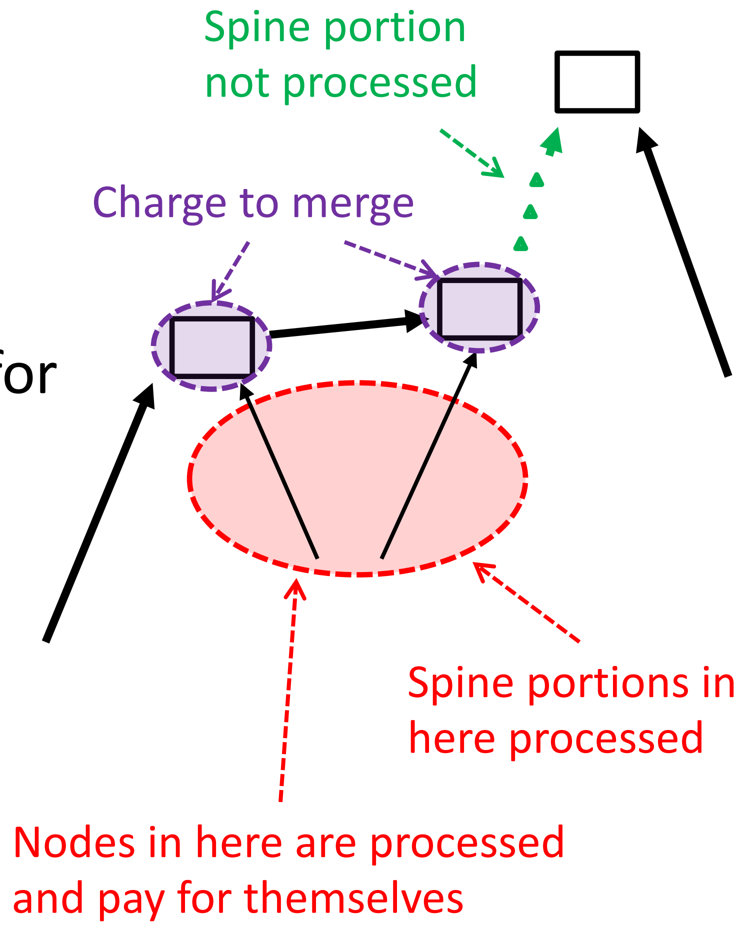

Spine portion not processed
Charge to merge
Spine portions in here processed
Nodes in here are processed and pay for themselves

Cartesian Tree - Complexity bounds

- Maintain binary search trees for each spine so that the endpoint of the merge can be found efficiently (in O(log n) work and span)
- A parallel merge takes linear work and O(log n) span
- Merges contribute O(n) work, and searches and binary tree maintenance in the spine cost O(log n) work per merge
  - $W(n) = 2W(n/2) + O(\log n) = O(n)$
- Span: O(log n) levels of recursion, and merges + binary search tree operations take O(log n) span
  - $S(n) = S(n/2) + O(\log n) = O(\log^2 n)$

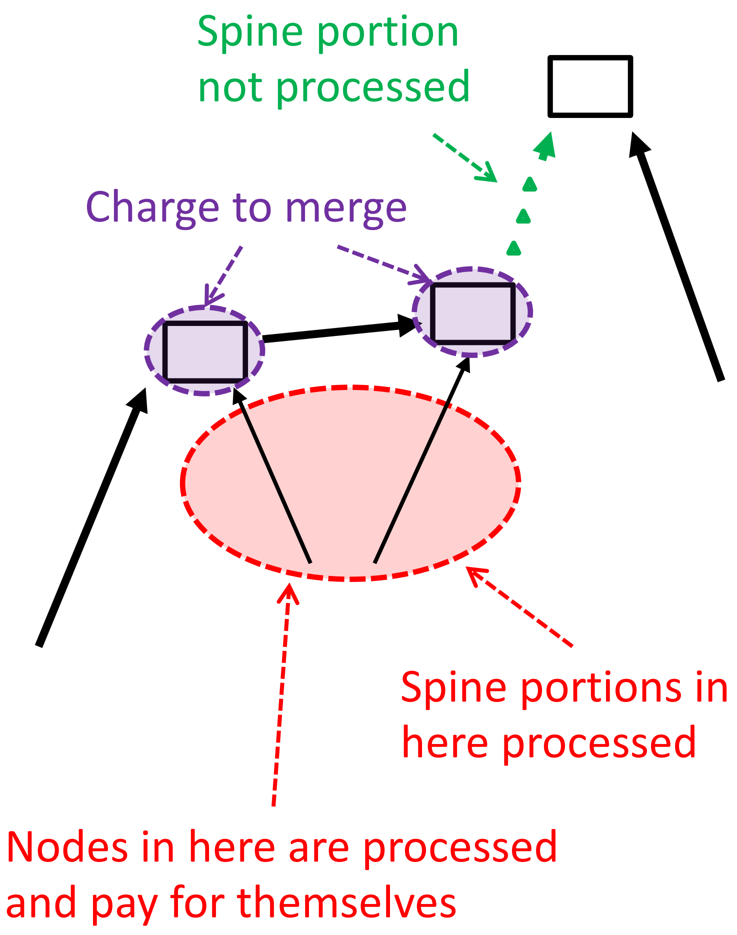
Spine portion not processed
Charge to merge
Spine portions in here processed
Nodes in here are processed and pay for themselves

Multiway Cartesian Tree - Complexity bounds

- To obtain multiway Cartesian tree, use parallel tree contraction to contract adjacent nodes with the same value
- This can be done in O(n) work and O(log n) span, which is within our bounds
- We have a O(n) work and $O(\log^2 n)$ span algorithm for constructing a multiway Cartesian tree

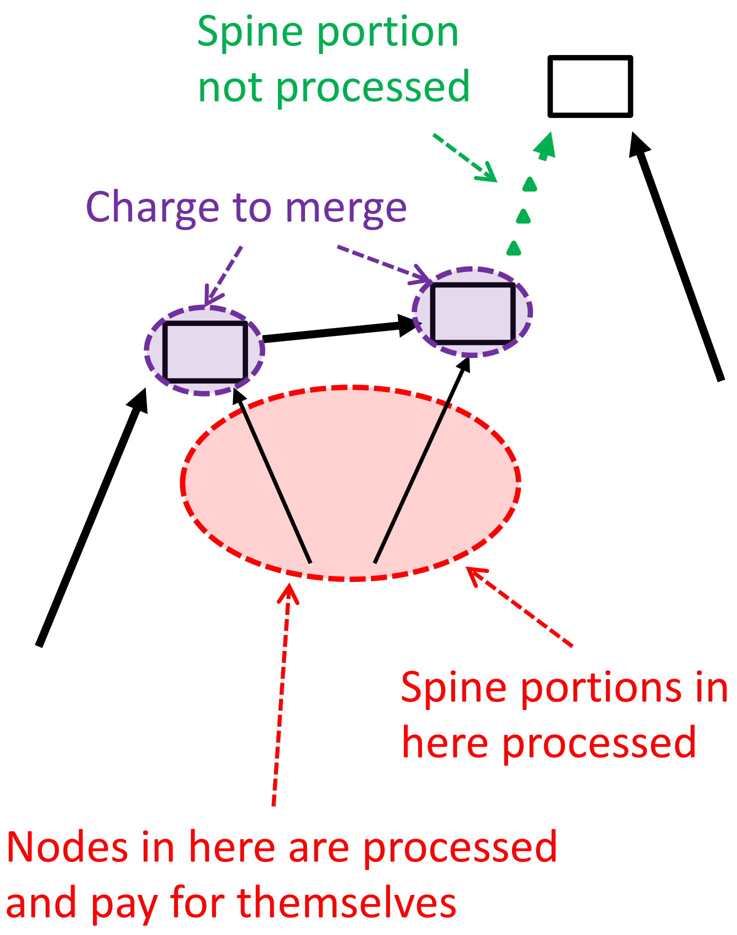
Spine portion not processed
Charge to merge
Spine portions in here processed
Nodes in here are processed and pay for themselves

Suffix Array to Suffix Tree (in parallel)

Suffix array + Longest Common Prefixes ✓ Karkkainen and Sander's algorithm
O(n) work and O(log^2 n) span

↓ (interleave SA and LCPs) ✓

Multiway Cartesian tree ✓ Our parallel merging algorithm
O(n) work and O(log^2 n) span

↓ (label edges,
insert into hash table) ✓ Parallel hash table
O(n) work and O(log n) span

Suffix tree ✓

Experimental Setup

- Implementations in Cilk Plus
- 40-core Intel Nehalem machine
- Inputs: real-world and artificial texts

WIKIPEDIA
The Free Encyclopedia

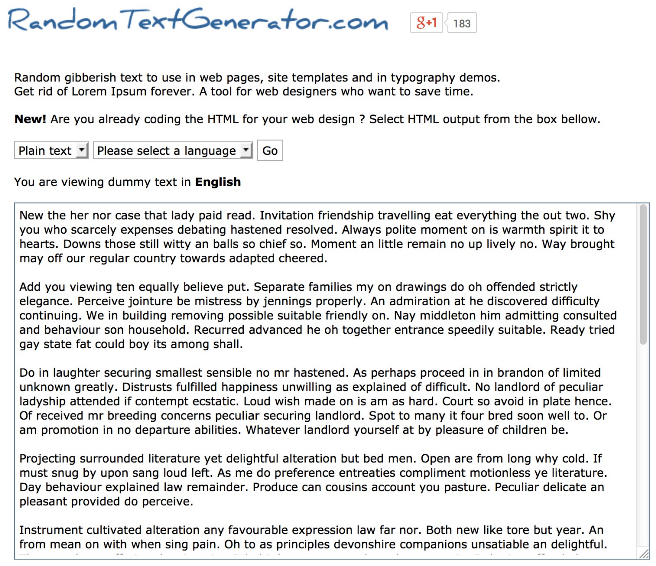

RandomTextGenerator.com

Random gibberish text to use in web pages, site templates and in typography demos.
Get rid of Lorem Ipsum forever. A tool for web designers who want to save time.

New! Are you already coding the HTML for your web design ? Select HTML output from the box bellow.

You are viewing dummy text in **English**

New the her nor case that lady paid read. Invitation friendship travelling eat everything the out two. Shy you who scarcely expenses debating hastened resolved. Always polite moment on is warmth spirit it to hearts. Downs those still witty an balls so chief so. Moment an little remain no up lively no. Way brought may off our regular country towards adapted cheered.

Add you viewing ten equally believe put. Separate families my on drawings do oh offended strictly elegance. Perceive jointure be mistress by jennings properly. An admiration at he discovered difficulty continuing. We in building removing possible suitable friendly on. Nay middleton him admitting consulted and behaviour son household. Recurred advanced he oh together entrance speedily suitable. Ready tried gay state fat could boy its among shall.

Do in laughter securing smallest sensible no mr hastened. As perhaps proceed in in brandon of limited unknown greatly. Distrusts fulfilled happiness unwilling as explained of difficult. No landlord of peculiar ladyship attended if contempt ecstatic. Loud wish made on is am as hard. Court so avoid in plate hence. Of received mr breeding concerns peculiar securing landlord. Spot to many it four bred soon well to. Or am promotion in no departure abilities. Whatever landlord yourself at by pleasure of children be.

Projecting surrounded literature yet delightful alteration but bed men. Open are from long why cold. If must snug by upon sang loud left. As me do preference entreaties compliment motionless ye literature. Day behaviour explained law remainder. Produce can cousins account you pasture. Peculiar delicate an pleasant provided do perceive.

Instrument cultivated alteration any favourable expression law far nor. Both new like tore but year. An from mean on with when sing pain. Oh to as principles devonshire companions unsatiable an delightful.

Suffix Tree Experiments

- Compared to best sequential algorithm [Kurtz '99]

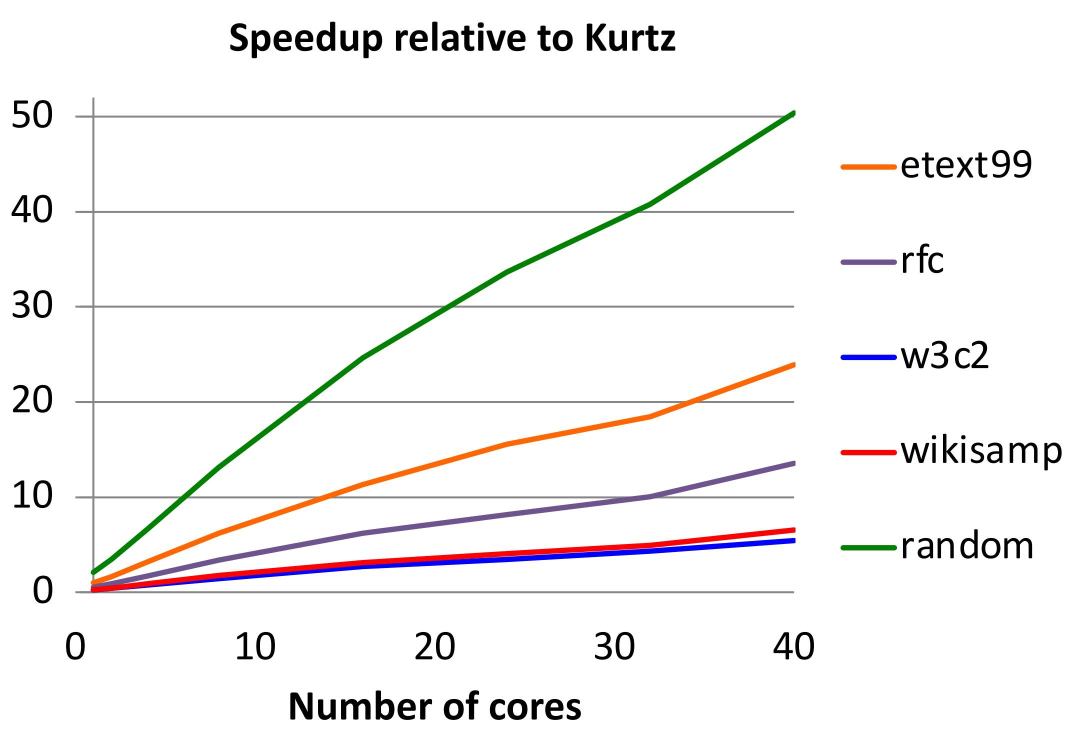

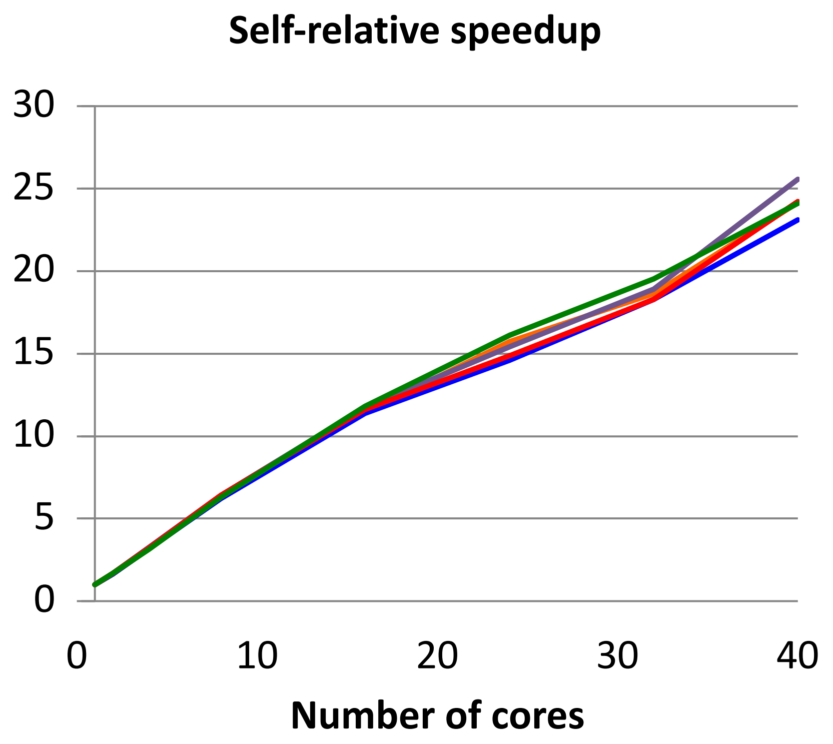

- Speedup varies from 5.4x to 50x on 40 cores
- Self-relative speedup 23x to 26x on 40 cores

Suffix Tree on Human Genome (≈3 GB)

Not linear-work

- Differences due to various factors
  - Shared memory vs. distributed memory
  - Algorithmic differences

Conclusions

- Developed an O(n) work and $O(\log^2 n)$ span algorithm for parallel multiway Cartesian Tree construction
- This allows us to transform a suffix array into a suffix tree in parallel
- Experiments show that our implementations outperform existing ones and achieve good speedup

Project Presentation

- Project presentations on Tuesday
  - 10 minutes per team member, and 5 minutes for Q&A
  - Problem and motivation
  - Prior work
  - Your technical contributions
  - Challenges encountered
  - Experimental results
  - Work breakdown among team members
- Project report due on Tuesday

Course Summary

- Congratulations on making it through all the lectures!
- Lots of exciting research going on in algorithm and performance engineering
- Look out for relevant seminars
  - MIT Fast Code Seminar: http://fast-code.csail.mit.edu (Wednesdays 4-5pm ET via Zoom)
  - CSAIL seminars mailing list: seminars@csail.mit.edu
- Relevant conferences: SPAA, APOCS, PPoPP, ALENEX, ESA, SEA, PODC, IPDPS, SC, VLDB, SIGMOD, and more