Sorting and Searching using Ternary-CAMs

Samar Sharma, Rina Panigrahy
(Cisco Systems)

Sorting & Searching in Hardware

- TCAMs: primarily used for forwarding.
- Sorting and searching are classical operations.
  - May be required at high speeds.
- Range search applications.
  - Address range locking
  - Virtual to physical address mapping
  - Port number search
- Priority queues, Sorted lists are common data structures.
  - Special purpose chips are available

Our Results

- We present two stage pipeline using TCAMs
  - for range search, sorting, priority queues
  - O(1) time, for search/update operations.
  - O(1) space for each entry

Disjoint Range Search

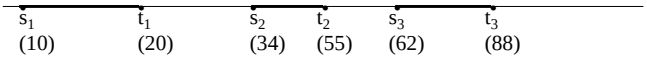

- Numbers are w = 32 bit long
- query point q
- Naïve approach : use upto (2w-2) = 60 prefixes per range
  - A prefix is a special kind of range
- We do it in 2 prefixes per range with two concurrent lookups

Use Longest Common Prefix

- LCP(x, y) : easy to compute by bit operations
  - e.g., LCP(34, 55) is 001*****
  - 34 is 00100010 in binary; 55 is 00110111
- LCP is same as lowest common ancestor
- If q is in a range, q matches LCP of its endpoints

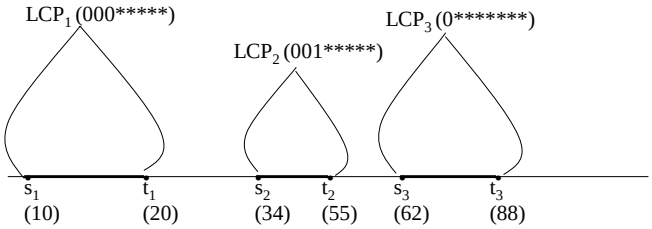

Store LCPs in TCAM

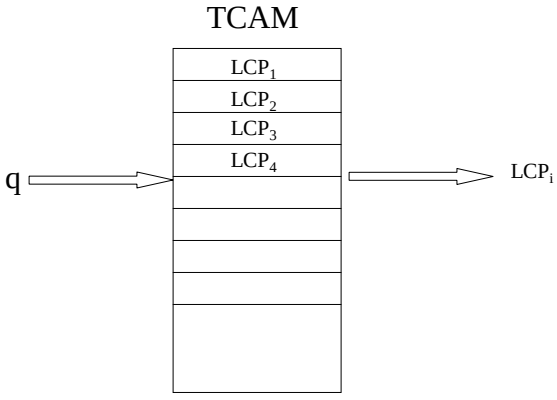

•Doesn't work. Matching range may not be longest prefix

Why LCPs don't work

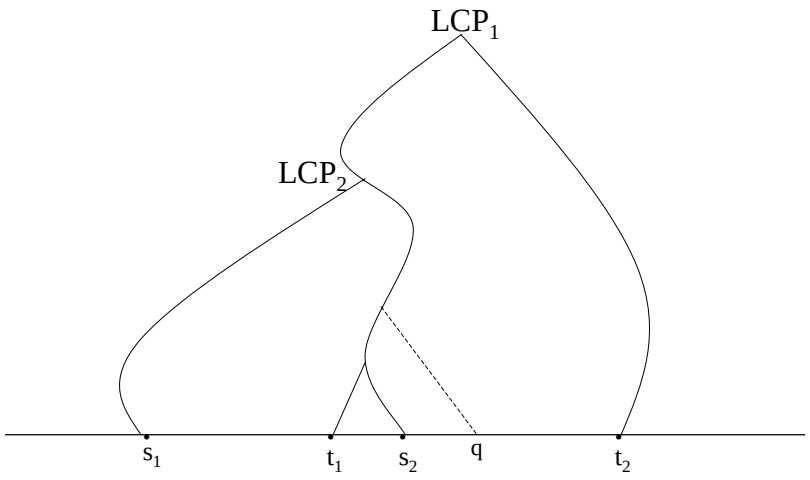

Store Extended LCPs in TCAM

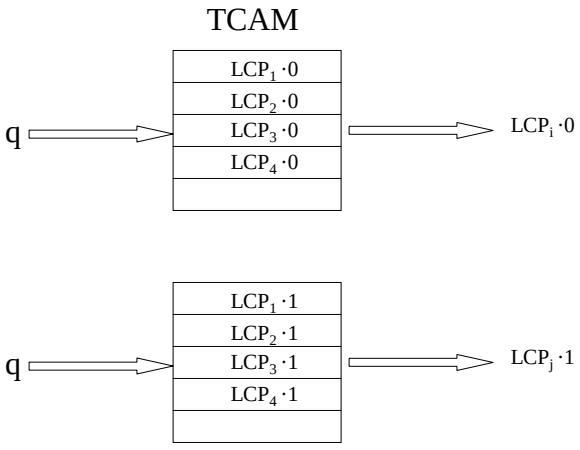

•Only the range i and range j can potentially match q

Store Extended LCPs in TCAM

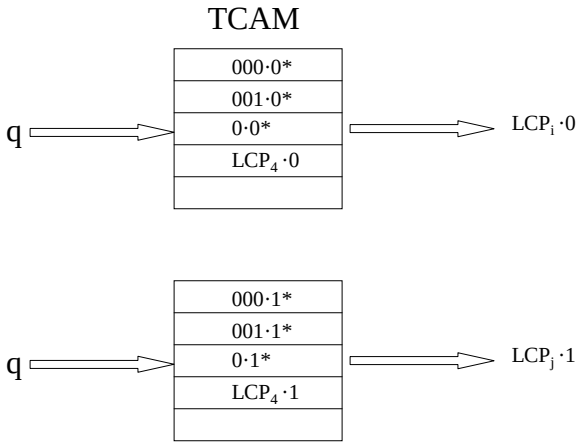

•Only the range i and range j can potentially match q

Proof

- Theorem : If q is in the range [s, t], then either LCP.0 or LCP.1 is longest matching prefix in its TCAM
  - where LCP = longest common prefix of (s, t)
- Proof :
  - Assume otherwise
  - Find that the ranges overlap
  - contradiction.

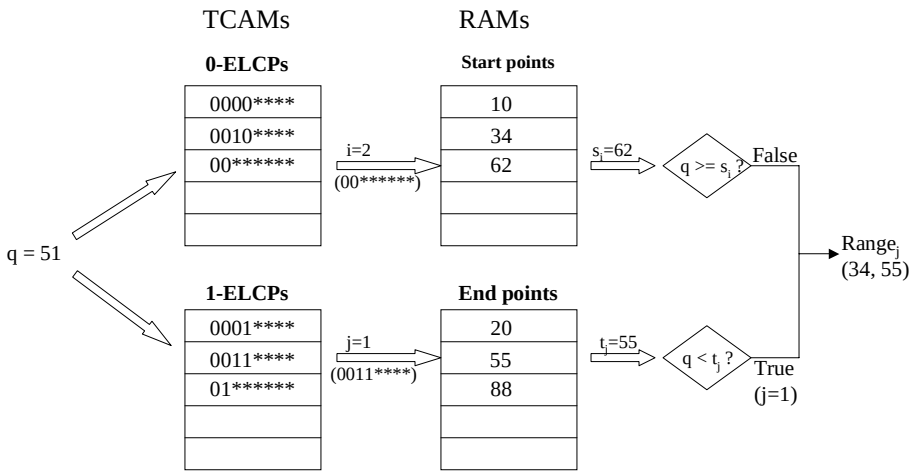

Results

- Search : two parallel lookups, followed by two parallel compares : O(1)
- Updates : insert 2 Extended LCP of endpoints : O(1)
- Can be extended to overlapping ranges Sorting : special case of disjoint range search. Updates to a sorted list is fast : O(1)
- Priority queue (findmin) : special case of sorting
- Range intersection (useful for locking)

Applications

- High speed sorting in O(n)
- Virtual to physical address translation
  - TLB can be compressed by viewing it as range search and using a TCAM
  - storage networks
- Locking, memory protection
- Graphics co-processors