Virtually
Cool
Ternary
Content Addressable Memory

Suparna Bhattacharya, K Gopinath
IBM, Indian Institute of Science
HotOS XIII, May, 2011

Thanks to Bob Montoye, Vijaylakshmi Srinivasan, Bipin Rajendran, Richard Freitas, John Karidis, C Mohan and Jai Menon

Ternary Content Addressable Memory (TCAM)

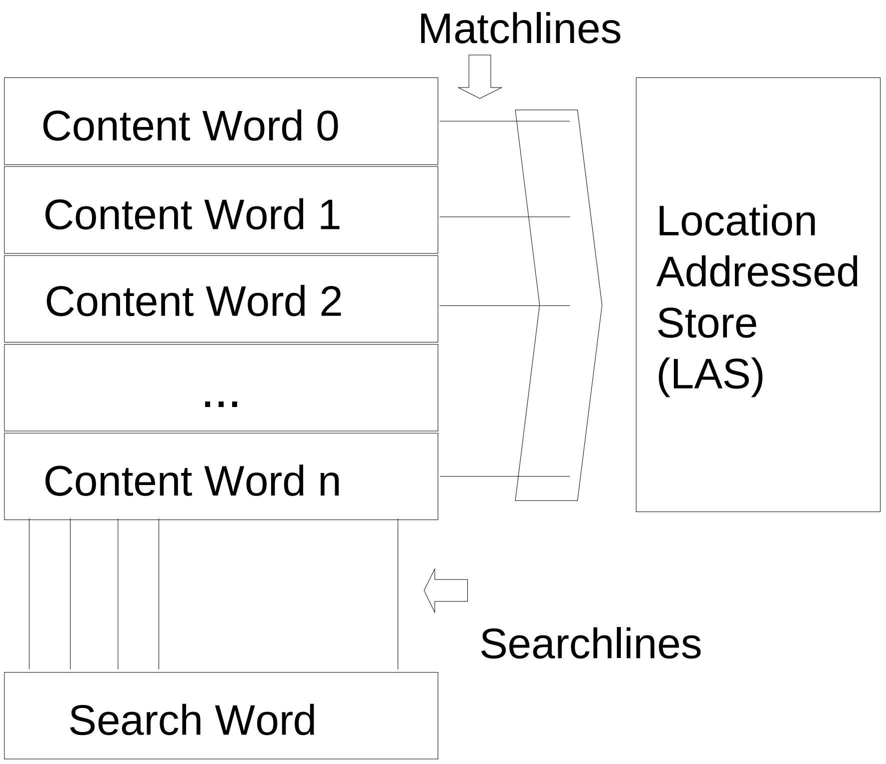

- **Fast (constant time) key lookup**
  - Parallel match on large data array

Ternary Content Addressable Memory (TCAM)

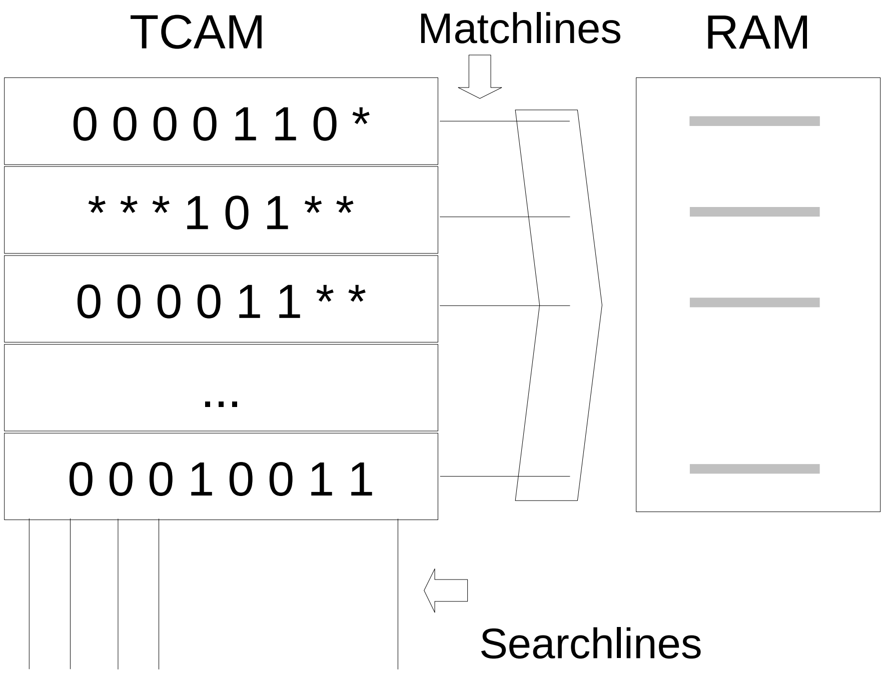

- **Fast (constant time) key lookup**
  - Parallel match on large data array
- **Ternary data: 0, 1, * ("don't care bit")**
  - Binary wild-card storage
- **Used in High Perf. Network routers**

Ternary Content Addressable Memory (TCAM)

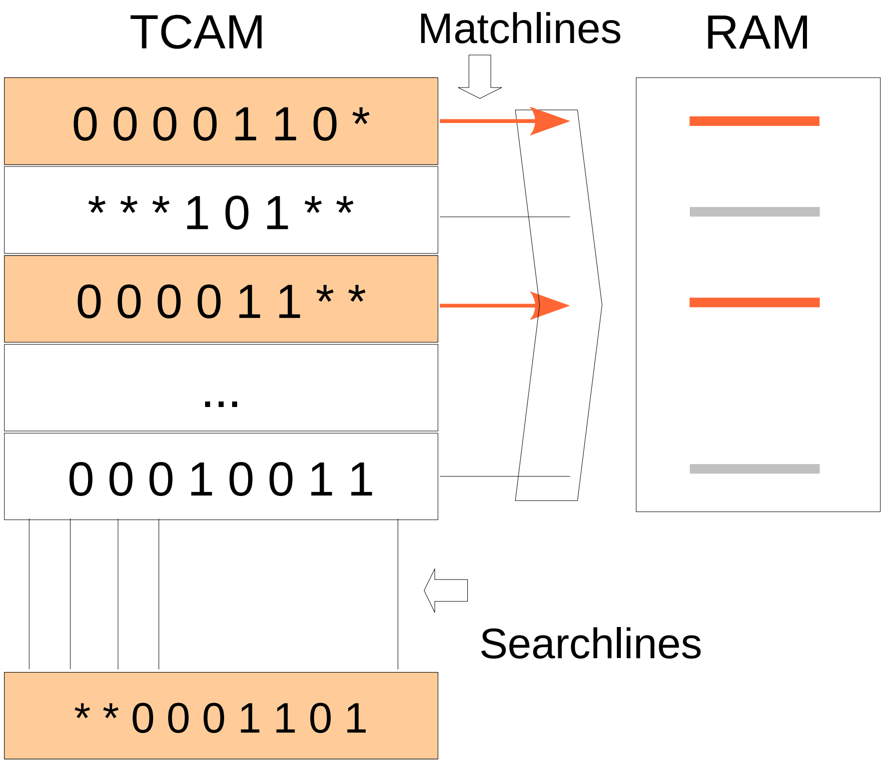

- **Fast (constant time) key lookup**
  - Parallel match on large data array
- **Ternary data: 0, 1, * (“don't care bit”)**
  - Binary wild-card storage
- **Used in High Perf. Network routers**

Ternary Content Addressable Memory (TCAM)

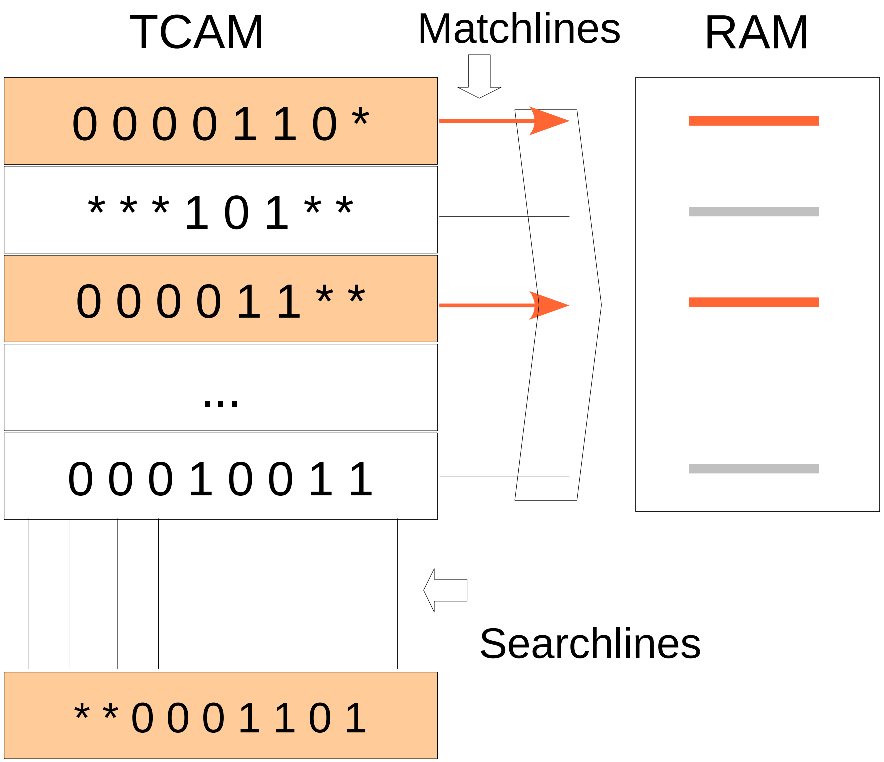

- Fast (constant time) key lookup
  - Parallel match on large data array
- Ternary data: 0, 1, * (“don't care bit”)
  - Binary wild-card storage
- Used in High Perf. Network routers

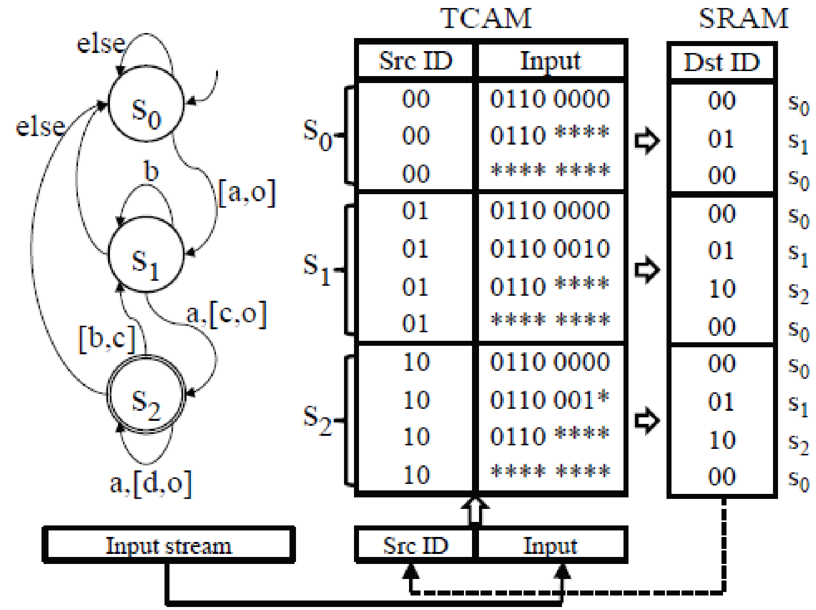

Example: Encoding a DFA in TCAM

Meiners et al, 2010: Fast regular expression matching using small TCAMs for network intrusion detection and prevention

A TCAM is a Natural Candidate for Representation of Space/Time Efficient Associative Search Structures

- Subset query – Ternary Bloom Filter
- Similarity search
  - Ternary Locality Sensitive Hashing (TLSH)
  - Approximate nearest neighbor
- Regular expression pattern matching
  - Compact DFA in TCAM
- Database join
  - Multi-match exploitation
- More flexible than radix tree, grid of tries, hash table
  - different constraints (only power of 2 ranges, not ordered, fixed width)

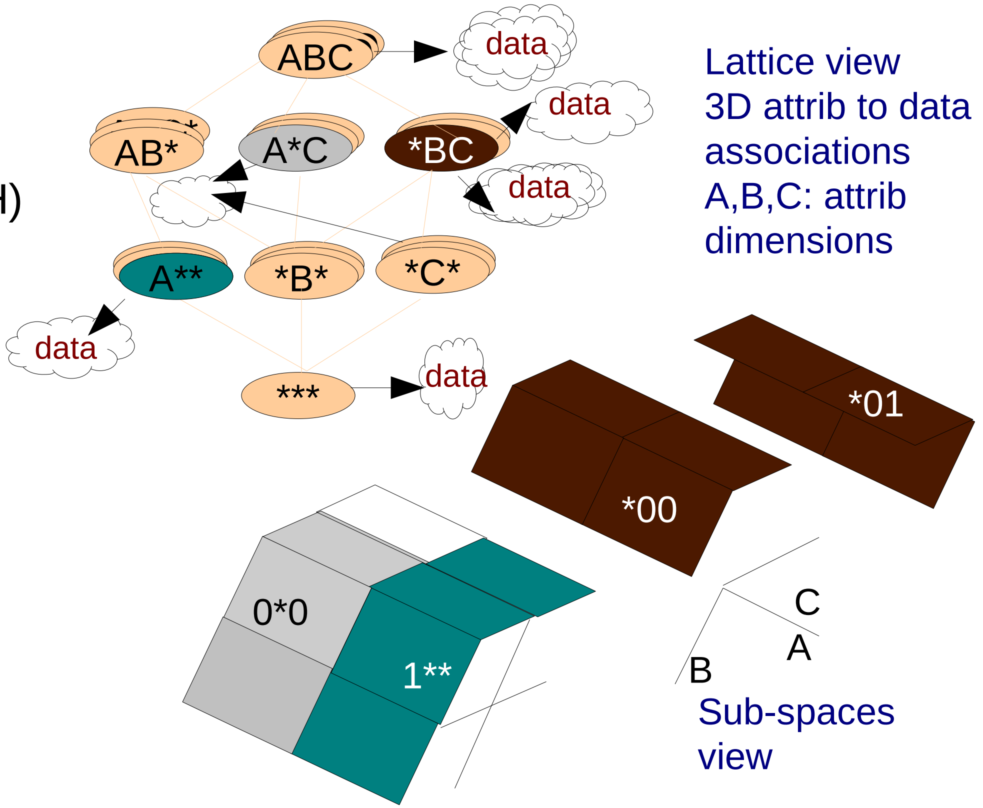

Parallel matching construct on a wild-card storage - powerful abstraction

*Ability to **simultaneously search** through a **large number of sub-spaces** of a (typically sparse) fixed dimensional space.*

But the Parallel Match Circuit Has a High Power Cost

Pagiamtzis et al 2006: CAM Circuits and Architectures: Tutorial & Survey

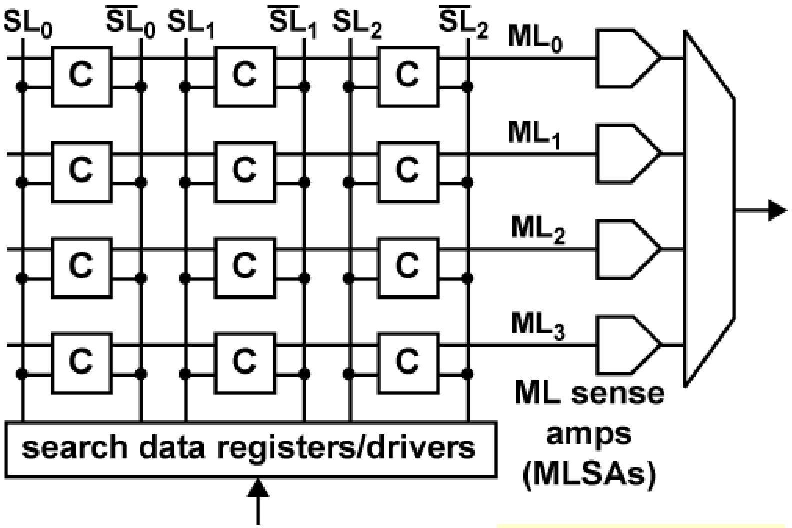

Agarwal & Sherwood 2008: TCAM Power and Delay Model

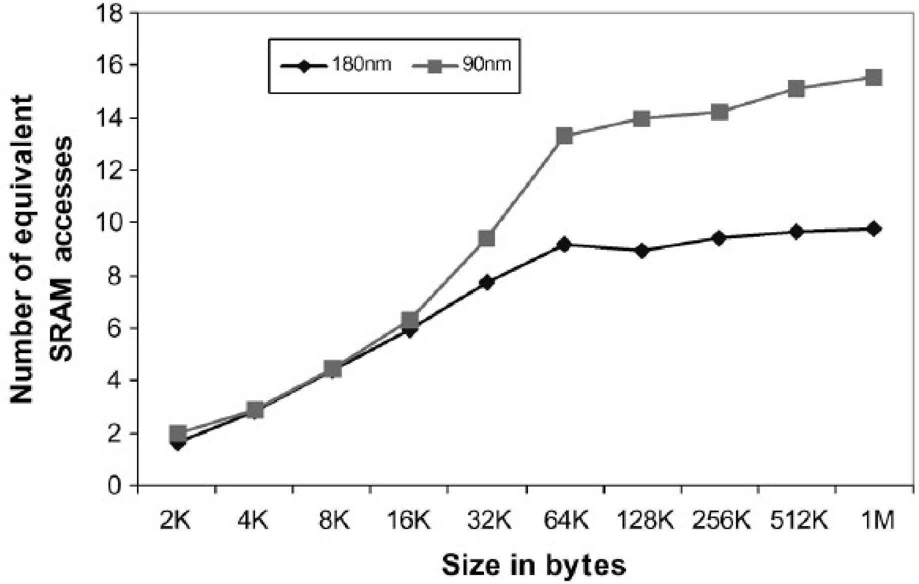

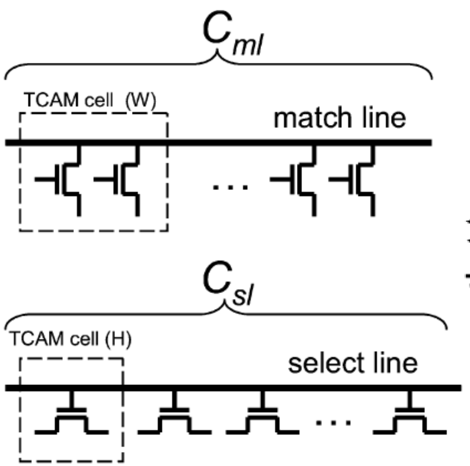

Mismatches are an overhead

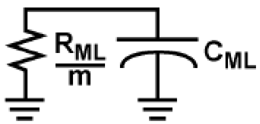

	MB/ chip	\$/chip	\$/MB	Speed (ns)	Watts /chip	Watts/ MB
DRAM	128	10-20	0.08-0.16	40-80	1-2	0.008 - 0.016
SRAM	9	50-70	5.5-7.8	3-5	1.5-3	0.17-0.33
TCAM	4.5	200-300	44.5-66.7	4-5	15-20	3.33 - 4.44

Goel & Gupta, SIGMETRICS'10: Small Subset Queries Using Ternary Bloom Filters

Content Addressable Virtual Memory Hierarchy

- Content Locality
  - Contiguity in content key-space
  - Physically dispersed
- Content-Based Page
  - Sub-space range in content key space
  - Entries may be physically dispersed
  - Different from traditional paging !
- Classifying workload content locality
  - Rare Hits
  - Frequent Item Hits
  - Nearby Item Hits
  - Random Hits

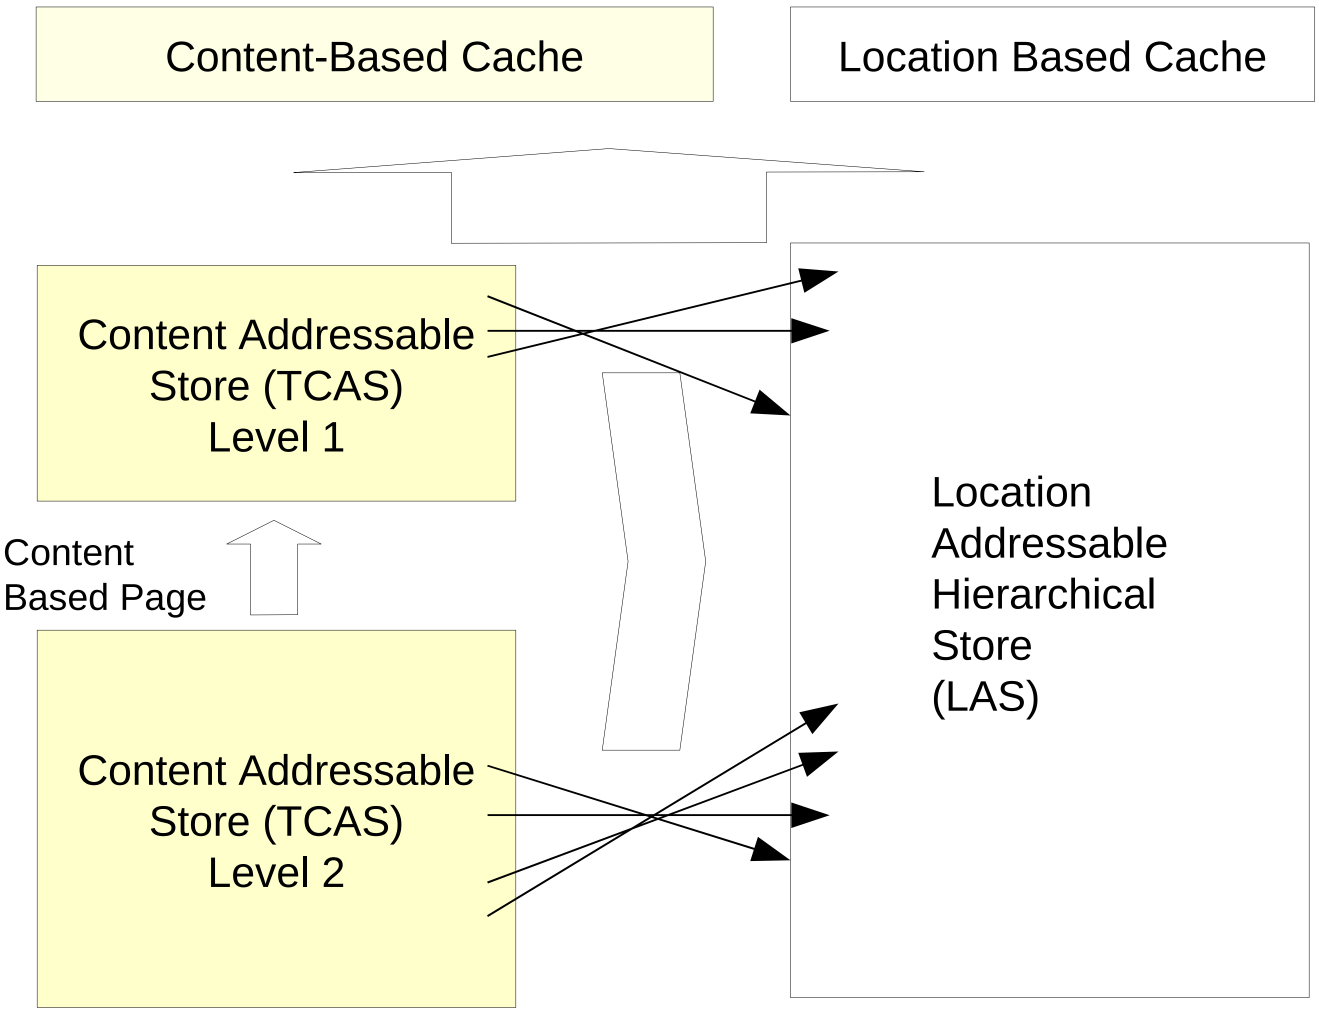

Example: Virtual Content Space to TCAS Mapping

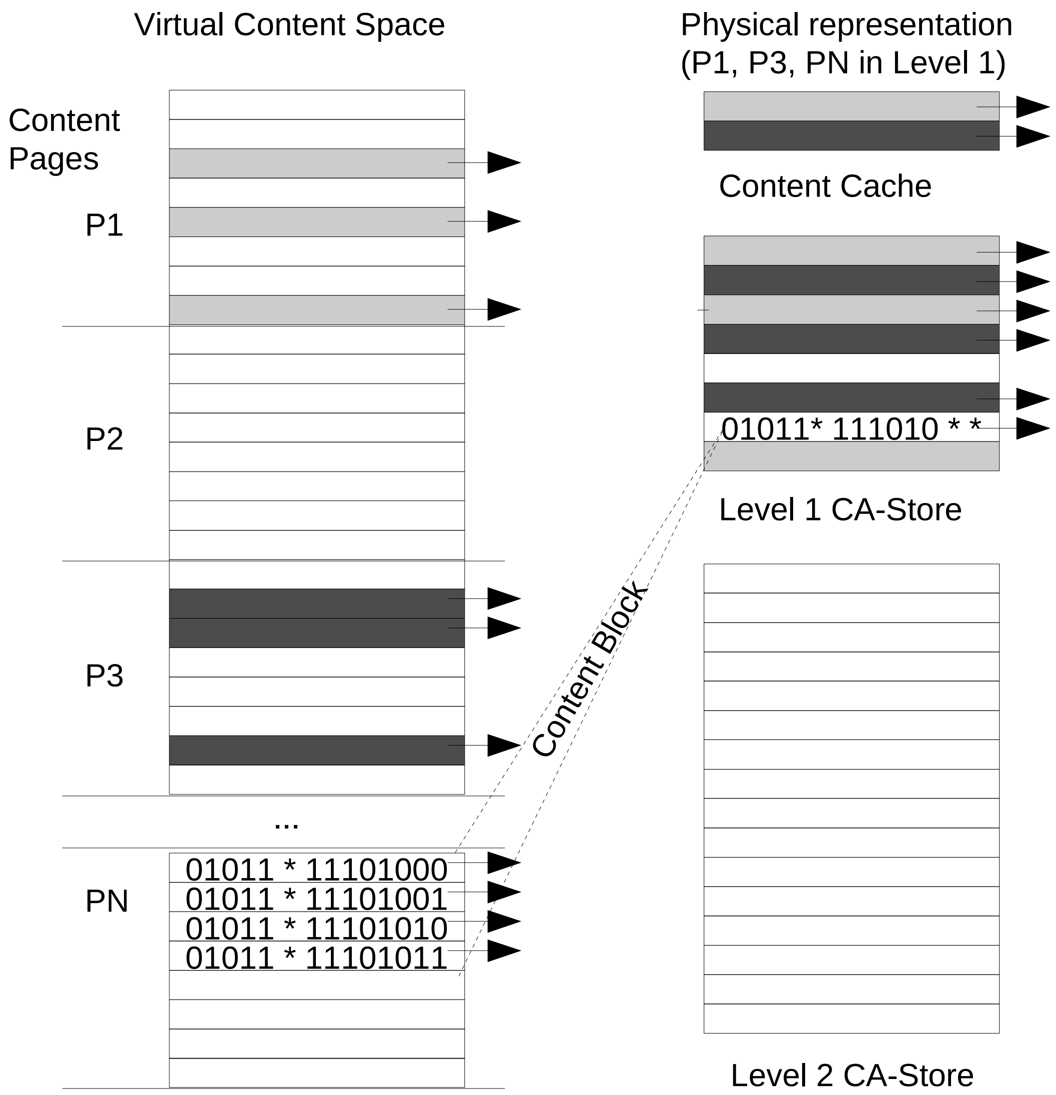

Many interesting questions arise, let us explore one of them in a little more detail...

- How do we save and find TCAM entries that have been paged out to DRAM ?
  - Representing ternary content words in a binary store
    - Easy: with extra bits
  - Indexing ternary content words in a location addressable store
    - What in-memory data structures should we use ?
      - Hash tables ?
      - Integer radix tree ?
      - ??

Example: Virtual Content Space to TCAS Mapping

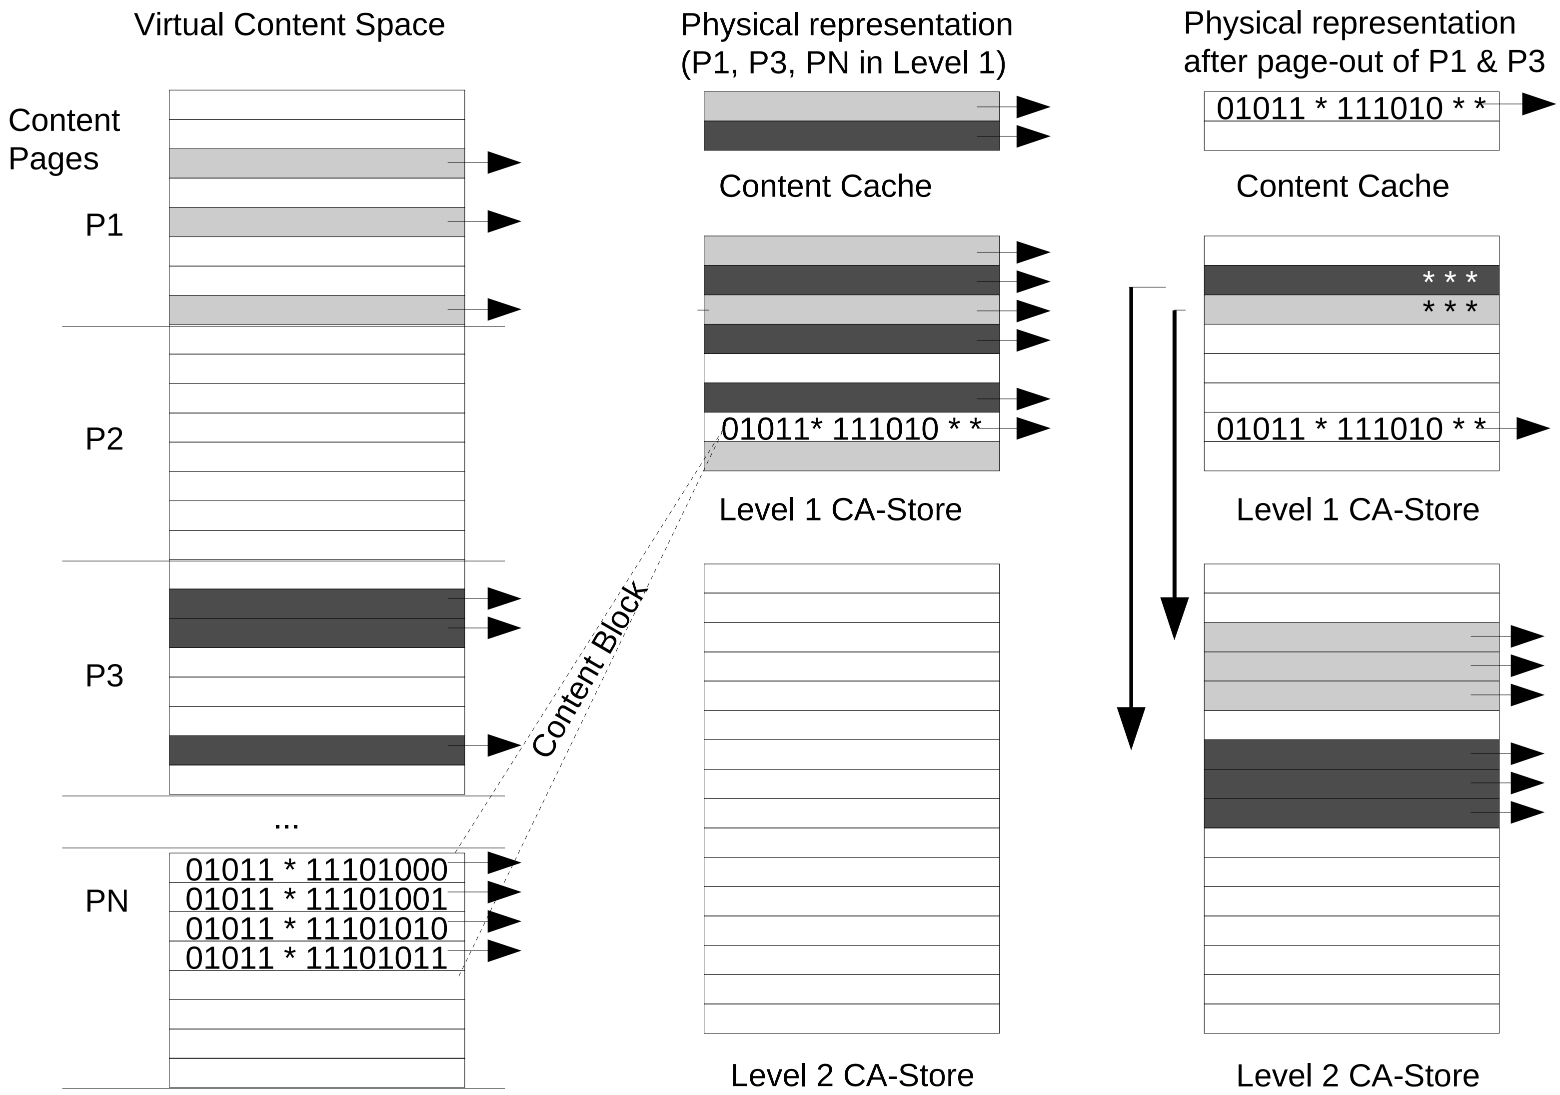

Implementation Challenges, Design Issues, Debates ...

- Feasibility and Potential: e.g. Power-perf-cost trade-off
  - Understand content locality/working sets of existing workloads
  - TCAM extensions for efficient multi-match ?
  - PCM(Phase Change Memory) based TCAM ?
- TCAS(Ternary Content Addr Store) & LAS(Location Addr Store) management
  - Esp. concurrency, sharing ...
- Choice of interface: How should the abstraction be exposed to applications ?
  - Fully transparent vs Exposed interface ?

- What new possibilities could be opened up if we make content addressability a first class abstraction in virtual memory design ?
  - Too radical or outrageous to be worth it ?
  - Or so crazy that it just might work ?
  - The good news is that it doesn't need to be **that** radical unless it makes sense
    - e.g. compatibility with location based addressing straightforward

Today is the 150th birth-anniversary of Rabindranath Tagore

Bengali Poet, writer, philosopher, Nobel laureate

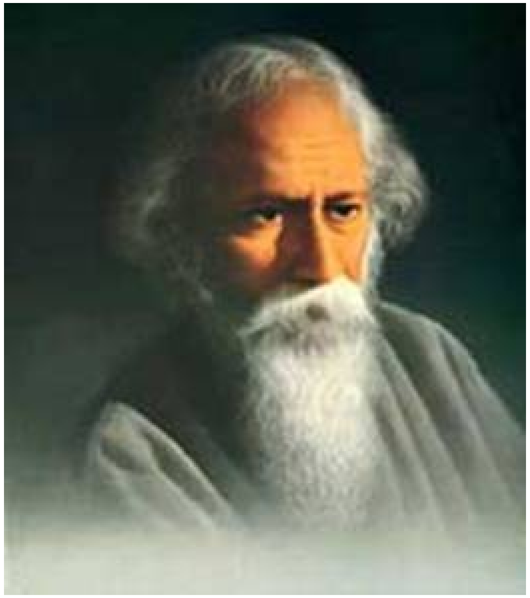

“... where words come out from the depth of truth,

where tireless striving stretches its arms towards perfection,

where the clear stream of reason has not lost its way into the dreary desert sand of dead habit ...”

“ ... Pearl fishers dive for pearls, merchants sail in their ships, while children gather pebbles and scatter them again.

They seek not for hidden treasures, they know not how to cast nets.”