

Ultra-Fast Similarity Search Using Ternary Content Addressable Memory

Yotam Harchol

The Hebrew University of Jerusalem, Israel

Joint work with:

Anat Bremner-Barr

IDC Herzliya,
Israel

David Hay

Hebrew University,
Israel

Yacov Hel-Or


IDC Herzliya,
Israel

DaMoN 2015, Melbourne, VIC, Australia

This research was supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement no 259085.




Hebrew University

The image features a light gray background with a series of vertical stripes in varying shades of gray on the left side. A solid orange horizontal bar spans across the middle of the image, containing the word "Background" in white text.

Background

The Nearest Neighbor Problem

- Classic problem in computer science
 - Many applications in reality
 - Machine learning
 - Computer vision
 - Image processing
 - Spatial databases
 - Network traffic classification and QoS
 - Performance bottleneck




The Nearest Neighbor Problem


- Input:
 - A set of data points
 - A query point (or a series of those)

(points are in a discrete space)
- Output:


Data point closest to the query point
- ***The Curse of Dimensionality***: Hard problem for high dimensions (even over 10 or so)




Known Solutions




Brute force search
Too long to compute



Space partitioning techniques
Time exponential in dimension



Locality sensitive hashing (LSH)
More efficient, but probabilistic
and still takes long time



Ternary locality sensitive hashing
First to show similarity search on TCAM
Probabilistic, requires more time and
space than our technique [SIGMOD '10]

Our Contribution

- ***Single TCAM lookup → Approximate solution***

- ***Four orders of magnitude faster than LSH***

- ***Also faster than Ternary LSH***

- ***Deterministic***

- ***Smaller (and practical) TCAM space required***



- **Ultra-fast algorithms for:**

- Exact nearest neighbor search

- k-nearest neighbors search

- More related problems...

(Not covered in this talk but shown in the paper)




- 
- Memory chip
or
Coprocessor?

Example: 2-dimensional IP lookup:




General Idea

- Encode cubes on data points
 - Given a query point, single TCAM lookup returns the smallest cube that contains it
 - Goal: no TCAM entry expansion in high dimension



General Idea

- Or instead, save TCAM space
 - Encode cubes on the query point
 - Query TCAM with growing cubes
 - Find first data point to match





Encoding Scheme

Ternary Match

$$\begin{array}{c} 01^*101^{**}1 \\ \approx \\ 0^*11^*101^* \end{array}$$

Match

$$\begin{array}{c} 01^*101^{**}1 \\ \not\approx \\ 0011^*101^* \end{array}$$

No match

Formally, $a=\{a_1,\dots,a_m\}$ matches the word $b=\{b_1,\dots,b_m\}$, denoted $a\approx b$, if and only if:
For any i , $a_i=b_i$ or a_i is $*$ or b_i is $*$


Encoding Function

- Goal:
Admissible encoding function *tcode*:


for points and cubes:

$$tcode(p) \approx tcode(C)$$

if and only if p is in C



- First step:
Solution in a single
dimensional space:



Encoding Function for Intervals

Ternary encoding function *tcode* for points and intervals:

Binary
encoding

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0000	0001	0010	0011	0100	0101	0110	0111	1000	1001	1010	1011	1100	1101	1110	1111

Requires multiple TCAM entries:

001*

010*

Encoding Function for Intervals

Ternary encoding function *tcode* for points and intervals:

Binary reflected
Gray code

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0000	0001	0011	0010	0110	0111	0101	0100	1100	1101	1111	1110	1010	1011	1001	1000
00**				01**				11**				10**			
????	????			????			????			????			????		
*00**		0*1*			*10*			1*1*			*00*				
????		????			????			????			????		????		

Encoding Function for Intervals

Ternary encoding function *tcode* for points and intervals:

Binary reflected Gray code

+

Extra bits

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0000 00	0011 10	0010 11	0110 11	0111 01	0101 01	0100 00	1100 00	1101 10	1111 10	1110 11	1010 11	1011 01	1001 01	1000 00	
00** **		01** **		11** **		10** **									
*0** 0*	0*** 1*		*1** 0*		1*** 1*		*0** 0*								
00 **	0*1* **		*10* **		1*1* **		*00* **								
*0** *0	0*** *1		*1** *0		1*** *1		*0** *0								

Encoding Function for Intervals

Ternary encoding function $tcode$ for points and intervals:

Final code
After eliminating
redundant bits

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
00000	00010	00110	00111	01111	01101	01001	01000	11000	11010	11110	11111	10111	10101	10001	10000
00***				01***				11***				10***			
*0*0*	0**1*				*1*0*				1**1*				*0*0*		
*00**		0*1**				*10**				1*1**				*00**	
*0**0				0***1				*1**0				1***1		*0**0	

Create code for intervals of length $h_{\max}$ $\rightarrow$ Encode all intervals of length $\leq h_{\max}$
(By intersecting two intervals of length $h_{\max}$)

* Full details and correction proofs are in the paper

Multiple Dimensions

Trivial expansion to d dimensions:

$$Y=[2,4)$$

$$tcode(Y)=001^{**}$$


$$Z=[4,6)$$

$$tcode(Z)=011^{**}$$

$$X=[1,3)$$

$$tcode(X)=00^{*}10$$

$$tcode(C_1)=00^{*}10001^{**}011^{**}$$



Multiple Dimensions

Trivial expansion to d dimensions:


$Y=[0,4)$
 $tcode(Y)=00***$

$Z=[0,4)$
 $tcode(Z)=00***$

$X=[4,8)$
 $tcode(X)=01***$

$tcode(C_1)=00*10001**011**$

$tcode(C_2)=01***00***00***$



Multiple Dimensions

Trivial expansion to d dimensions: Query $q=(5,2,3)$ $tcode(5)=01101$
 $tcode(2)=00110$
 $tcode(3)=00111$

**No TCAM entry expansion at all,
even in high dimensions!**

$tcode(p)=011010011000111$

$tcode(C_1)=00*10001**011**$ ✗

$tcode(C_2)=01***00***00***$ ✓


Cube C_2



Solving Nearest Neighbor Search on TCAM

Nearest Neighbor Search on TCAM

- Time efficient method:
TCAM entries are cubes around data points
 - Sort cubes by size – smaller first
 - Solve in a single TCAM lookup!



$tcode(q)$: 011101001110001011

$tcode(p_1, 2)$:
 $tcode(p_2, 2)$:
 $tcode(p_1, 4)$:
 $tcode(p_2, 4)$:

00*1**001***011***	0
01*1**00*1**00*1**	1
00*****0***1*0***1	2
01*****00***00***	3
	4
*****	5


TCAM

$p_1, 2$	0
$p_2, 2$	1
$p_1, 4$	2
$p_2, 4$	3
	4
None	5


SRAM
21

Nearest Neighbor Search on TCAM

- Space efficient method:
TCAM entries are the data points
Query is split to a set of cubes
 - Solve in a several TCAM lookups
 - Requires less TCAM space
 - Binary search on cube size



$tcode(q,2):$ 011***00**10001***
 $tcode(q,4):$ 0*****100*****0***1*



Optimal Results in ℓ_∞

- Space is discrete
- Encoding all cubes of odd edge lengths 1,3,5,... provides an exact solution in ℓ_∞


1=The point itself

3=Distance 1 in ℓ_∞

5=Distance 2 in ℓ_∞

...

- Less cubes:
 c -approximation in ℓ_∞
- In ℓ_p : $c \cdot \sqrt[p]{d}$ -approximation
($c \geq 1$)






Experiments and Evaluation

Dataset




- Use case: Image similarity search (LabelMe 22K)
 - Using GIST descriptors (vectors in $\mathbb{R}^{512}$)
 - Downsampled to $\mathbb{R}^{40}$ using PCA
 - Quantified uniformly to $\{0, \dots, 255\}^{40}$
 - TCAM Entry width: 40×12 bits = 480 bits (contemporary TCAMs support up to 640 bits per entry)
 - Maximal cube edge length: 8
- 22,019 images
 - Data set: up to 21,019 images
 - Total size on TCAM: 9.6Mbit to 80Mbit (contemporary TCAMs have 40Mbit-80Mbit space)
 - Queries: 1000 images



Experiment on a Real TCAM

- Currently - no evaluation board for TCAM
- Instead, we used a commodity network switch that contains a TCAM
 - Switch has 48 ports of 1Gbps each (1.5 Million packet per second)
 - TCAM is 192 bits wide

d-Dimensional Cube Representation:



OpenFlow
Counters

**Single port:
1.5 Million Queries
Per Second!**

Experiment on a Real TCAM


- Can we do more?
 - Yes!



Proof of concept experiment:

Forwarding Rules:


- From port 1 to port 2
- From port 3 to port 4
- From port 5 to port 6
- ...
- From port 47 to port 48



24 ports:
35.69 Million
Queries
Per Second!
(almost $24 \times 1.5M$)


Performance Analysis & Simulation

- Contemporary TCAMs provide rates of **360 million** to more than **1 billion** queries per second
- We were *limited by the number of ports*.
TCAM could do more



Quality of Results

- In ℓ_∞ : optimal results
- In ℓ_2 : approximated results:



Conclusions

- TCAM provides high parallelism
- Our encoding scheme and algorithms solve the Nearest Neighbor problem in high dimensional spaces using a single (or a few) TCAM lookups
 - Very high search rate (many millions per second)
- We believe TCAM can be useful for more purposes, also outside the networking field
 - Use TCAM as a coprocessor / search accelerator



Questions?

Thank you.