#### ALTO: Mode-Agnostic Linearized Storage for High-Performance Sparse Tensor Computations

Ahmed E. Helal,\* <u>Jan Laukemann,\*</u> Fabio Checconi,\* Jesmin Jahan Tithi,\* Teresa Ranadive,† Fabrizio Petrini,\* Jee Choi‡

Intel labs,\* Laboratory for Physical Sciences (LPS),† University of Oregon‡



# Intel PCL

#### Our modus operandi:



# Intel PCL

#### Our modus operandi:



 Application-driven research to improve efficiency and performance of Intel's architectures and tools

# Intel PCL

#### Our modus operandi:



- Application-driven research to improve efficiency and performance of Intel's architectures and tools
- Target domains: machine learning (ML), highperformance computing (HPC), and data analytics

# Outline

- Background & Motivation
  - State-of-the-art sparse tensor formats
- ALTO storage format
- Results & performance evaluation
- Conclusion

 Many important application domains produce and manipulate massive amounts of high-dimensional sparse data

 Many important application domains produce and manipulate massive amounts of high-dimensional sparse data



 Many important application domains produce and manipulate massive amounts of high-dimensional sparse data





Source: Nature Communications and Scientific Reports, Wikimedia Commons

 Many important application domains produce and manipulate massive amounts of high-dimensional sparse data







Source: Nature Communications and Scientific Reports, Wikimedia Commons

9

 Many important application domains produce and manipulate massive amounts of high-dimensional sparse data





Source: Nature Communications and Scientific Reports, Wikimedia Commons

 Many important application domains produce and manipulate massive amounts of high-dimensional sparse data









 High-dimensional datasets exhibit diverse shapes and data distributions as well as unstructured sparsity patterns



- High-dimensional datasets exhibit diverse shapes and data distributions as well as unstructured sparsity patterns
- Challenging to efficiently encode and represent



- Store, group, and organize nonzero elements in sparse tensors is a fundamental problem
- Goals:

- Store, group, and organize nonzero elements in sparse tensors is a fundamental problem
- Goals:
  - reduce tensor storage



- Store, group, and organize nonzero elements in sparse tensors is a fundamental problem
- Goals:
  - reduce tensor storage
  - improve data locality





- Store, group, and organize nonzero elements in sparse tensors is a fundamental problem
- Goals:
  - reduce tensor storage
  - improve data locality
  - ✤ increase parallelism







- Store, group, and organize nonzero elements in sparse tensors is a fundamental problem
- Goals:
  - reduce tensor storage
  - improve data locality
  - ✤ increase parallelism
  - decrease workload imbalance
    & synchronization overhead

		-
ala	nce	

- Store, group, and organize nonzero elements in sparse tensors is a fundamental problem
- Goals:
  - reduce tensor storage
  - improve data locality
  - ✤ increase parallelism
  - decrease workload imbalance
    & synchronization overhead

bala	ance		





MTTKRP (Matricized Tensor Times Khatri-Rao product)

 A key memory-intensive kernel in many sparse tensor computations  $A = X_{(1)}(C \odot B)$ 

MTTKRP (Matricized Tensor Times Khatri-Rao product)

 A key memory-intensive kernel in many sparse tensor computations



Source: Blocking Optimization Strategies for Sparse Tensor Computation, Jee Choi et al.

MTTKRP (Matricized Tensor Times Khatri-Rao product)

- A key memory-intensive kernel in many sparse tensor computations
- Challenge: hard to efficiently utilize the memory system for all dimensions (modes)



Source: Blocking Optimization Strategies for Sparse Tensor Computation, Jee Choi et al.

MTTKRP (Matricized Tensor Times Khatri-Rao product)

- A key memory-intensive kernel in many sparse tensor computations
- Challenge: hard to efficiently utilize the memory system for all dimensions (modes)
  - Different data locality patterns across modes
  - Update conflicts to factor matrices
  - Load imbalance across threads



Source: Blocking Optimization Strategies for Sparse Tensor Computation, Jee Choi et al.



















Can be classified based on their encoding of the indexing metadata into:



- Can be classified based on their encoding of the indexing metadata into:
  - 1. List-based formats



- Can be classified based on their encoding of the indexing metadata into:
  - 1. List-based formats
  - 2. Tree-based formats



- Can be classified based on their encoding of the indexing metadata into:
  - 1. List-based formats
  - 2. Tree-based formats
  - 3. Block-based formats


COO (coordinate):
list-based format

i	j	k
0	3	0
1	0	0
1	6	1
2	2	1
3	1	1
3	4	0



 CSF (compressed sparse fiber): tree-based format





 CSF (compressed sparse fiber): tree-based format





- CSF (compressed sparse fiber): tree-based format
  - 6 possible representations for a 3D tensor



- CSF (compressed sparse fiber): tree-based format
  - 6 possible representations for a 3D tensor
  - In practice, we only keep 3 representations



- CSF (compressed sparse fiber): tree-based format
  - 6 possible representations for a 3D tensor
  - In practice, we only keep 3 representations
    - 1. i -> k -> j
    - 2. k->i->j
    - 3. j -> k -> i



- HiCOO (hierarchical coordinate): block-based format
  - size(e-idx) << size(b-idx)</li>

bptr	b <sub>i</sub>	b <sub>j</sub>	b <sub>k</sub>	e <sub>i</sub>	e <sub>j</sub>	e <sub>k</sub>
0	0	0	0	1	0	0
1	0	1	0	0	1	0
2	0	3	0	1	0	1
3	1	0	0	1	1	1
4	1	1	0	0	0	1
5	1	2	0	1	0	0

#### 2X2X2



#### Spatial data distribution



A box plot of the data (nonzero elements) distribution across multi-dimensional blocks (subspaces). The multi-dimensional subspace size is  $128^{N}$ , where N is the number of dimensions (modes).

#### Spatial data distribution



A box plot of the data (nonzero elements) distribution across multi-dimensional blocks (subspaces). The multi-dimensional subspace size is  $128^{N}$ , where N is the number of dimensions (modes).

increased sparsity  $\rightarrow$  inefficient memory usage



Formats	Granularity	Mode Orientation	Data Locality	Parallelism	Load balance
List-based ( <b>COO</b> )	NNZ element	Mode- agnostic	NA	Suffer from conflicts	Maximized
Tree-based ( <b>CSF</b> )	Compressed sparse fiber	Mode-specific	Improved for specific mode	Improved for specific mode	Work imbalance (especially in short modes)
Block- based ( <b>HiCOO</b> )	Compressed block	Mode- agnostic	Improved	Suffer from conflicts	Work imbalance across blocks

- A new storage format for high performance and scalable sparse tensor computations
- Maps a multi-dimensional space into a linear space, and vice versa, using an adaptive (data-aware) encoding scheme



- A new storage format for high performance and scalable sparse tensor computations
- Maps a multi-dimensional space into a linear space, and vice versa, using an adaptive (data-aware) encoding scheme

mode-agnostic linearized index for each NNZ



- A new storage format for high performance and scalable sparse tensor computations
- Maps a multi-dimensional space into a linear space, and vice versa, using an adaptive (data-aware) encoding scheme
  - mode-agnostic linearized index for each NNZ
  - > compact representation using the minimum number of bits



- A new storage format for high performance and scalable sparse tensor computations
- Maps a multi-dimensional space into a linear space, and vice versa, using an adaptive (data-aware) encoding scheme
  - mode-agnostic linearized index for each NNZ
  - > compact representation using the minimum number of bits
  - > improved data locality across all modes



- A new storage format for high performance and scalable sparse tensor computations
- Maps a multi-dimensional space into a linear space, and vice versa, using an adaptive (data-aware) encoding scheme
  - mode-agnostic linearized index for each NNZ
  - > compact representation using the minimum number of bits
  - > improved data locality across all modes
  - perfectly balanced partitions by NNZs for effective scaling



# <u>A</u>daptive <u>L</u>inearized <u>Tensor</u> <u>O</u>rder (ALTO)

- A new storage format for high performance and scalable sparse tensor computations
- Maps a multi-dimensional space into a linear space, and vice versa, using an adaptive (data-aware) encoding scheme
  - mode-agnostic linearized index for each NNZ
  - > compact representation using the minimum number of bits
  - > improved data locality across all modes
  - perfectly balanced partitions by NNZs for effective scaling



Data-aware synchronization to resolve conflicting updates (writes) across threads





Formats	Granularity	Mode Orientation	Data Locality	Parallelism	Load balance
List-based (COO)	NNZ element	Mode- agnostic	NA	Suffer from conflicts	Maximized
Tree-based ( <b>CSF</b> )	Compressed sparse fiber	Mode-specific	Improved for specific mode	Improved for specific mode	Work imbalance (especially in short modes)
Block- based ( <b>HiCOO</b> )	Compressed block	Mode- agnostic	Improved	Suffer from conflicts	Work imbalance across blocks
ALTO	NNZ element	Mode- agnostic	Improved	Improved	Maximized

- ALTO outperforms an oracle that selects the best state-of-the-art format in terms of the parallel performance and tensor storage
- Platform: an Intel Xeon Platinum 8280 (CLX) w/ 2x28 cores
- Mode-agnostic formats: coordinate (COO) and hierarchical COO (HiCOO)
- Mode-specific formats: compressed sparse fiber (CSF) and tiled CSF



"ALTO: Adaptive Linearized Storage of Sparse Tensors", Ahmed Helal, Jan Laukemann, Fabio Checconi, Jesmin Jahan Tithi, Teresa Ranadive, Fabrizio Petrini, Jee Choi, accepted to ACM International Conference on Supercomputing (ICS), 2021

- ALTO outperforms an oracle that selects the best state-of-the-art format in terms of the parallel performance and tensor storage
- Platform: an Intel Xeon Platinum 8280 (CLX) w/ 2x28 cores
- Mode-agnostic formats: coordinate (COO) and hierarchical COO (HiCOO)
- Mode-specific formats: compressed sparse fiber (CSF) and tiled CSF



"ALTO: Adaptive Linearized Storage of Sparse Tensors", Ahmed Helal, Jan Laukemann, Fabio Checconi, Jesmin Jahan Tithi, Teresa Ranadive, Fabrizio Petrini, Jee Choi, accepted to ACM International Conference on Supercomputing (ICS), 2021













#### 63 4X4X2 Subspace





# 4X4X2 Subspace

4X2X2 Subspace

0











#### <u>63</u>



























ALTO Tensor				
Value	Position			
$x_{1,0,0}$	2 (000010)			
<i>x</i> <sub>3,1,1</sub>	15 ( <mark>001111</mark> )			
<i>x</i> <sub>0,3,0</sub>	20 (010100)			
<i>x</i> <sub>2,2,1</sub>	25 ( <mark>011001</mark> )			
<i>x</i> <sub>3,4,0</sub>	42 (101010)			
$x_{1,6,1}$	51 (11 <mark>0011</mark> )			


1 bit 2 bits 3 bits



ALTO Bit Mask

 $b_{j,2}$   $b_{j,1}$   $b_{i,1}$   $b_{j,0}$   $b_{i,0}$   $b_{k,0}$ 

# ALTO: Example

	ALTO Tensor				
	Value	Position			
	$x_{1,0,0}$	2 (000010)			
	<i>x</i> <sub>3,1,1</sub>	15 ( <mark>001111</mark> )			
	<i>x</i> <sub>0,3,0</sub>	20 (010100)			
	<i>x</i> <sub>2,2,1</sub>	25 (011001)	••••		
	<i>x</i> <sub>3,4,0</sub>	42 (101010)			
	$x_{1,6,1}$	51 (11 <mark>0011</mark> )			



ALTO Bit Mask





## ALTO: Example



ALTO Bit Mask





## ALTO: Example







Synchronization based on data reuse of target tensors



Synchronization based on data reuse of target tensors

1. Limited or no data reuse  $\rightarrow$  conflict resolution via **atomic updates** 



Synchronization based on data reuse of target tensors

- 1. Limited or no data reuse → conflict resolution via **atomic updates**
- 2. Otherwise **→** local temporary storage and global parallel reduction

```
Adaptive parallel execution of mode-1 MTTKRP-
Algorithm
ALTO kernel. ALTO automatically uses local storage or atomics,
based on the reuse of output fibers, to resolve the update conflicts.
Input: A third-order ALTO sparse tensor X \in \mathbb{R}^{I \times J \times K} with M
      nonzero elements, dense factor matrices \mathbf{A} \in \mathbb{R}^{I \times R}, \mathbf{B} \in \mathbb{R}^{J \times R},
     and \mathbf{C} \in \mathbb{R}^{K \times R}
Output: Updated dense factor matrix \tilde{\mathbf{A}} \in \mathbb{R}^{I \times R}
  1: for l = 1, \ldots, L in parallel do
                                                          ▶ ALTO line segments.
          for \forall x \in X_I do
  2:
               i = \mathbb{EXTRACT}(pos(x), MASK(1))  > De-linearization.
  3:
              j = \mathbb{EXTRACT}(pos(x), MASK(2))
  4:
              k = \mathbb{EXTRACT}(pos(x), MASK(3))
  5:
               for r = 1, ..., R do
  6:
                     \mathbf{Temp}_{l}(i - T_{l,1}^{s}, r) + = val(x) \times \mathbf{B}(j, r) \times \mathbf{C}(k, r)
  7:
                     \mathbb{ATOMIC}(\tilde{\mathbf{A}}(i,r) + = val(x) \times \mathbf{B}(j,r) \times \mathbf{C}(k,r))
  8:
               end for
  9:
          end for
 10:
 11: end for
```

Algorithm Adaptive parallel execution of mode-1 MTTKRP-								
ALTO kernel. ALTO automatically uses local storage or atomics,								
based on the reuse of output fibers, to resolve the update conflicts.								
<b>Input:</b> A third-order ALTO sparse tensor $X \in \mathbb{R}^{I \times J \times K}$ with $M$ nonzero elements, dense factor matrices $\mathbf{A} \in \mathbb{R}^{I \times R}$ , $\mathbf{B} \in \mathbb{R}^{J \times R}$ , and $\mathbf{C} \in \mathbb{R}^{K \times R}$								
<b>Output:</b> Updated dense factor matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$								
1: for $l = 1,, L$ in parallel do $\triangleright$ ALTO line segments								
2: for $\forall x \in X_l$ do								
3: $i = \mathbb{EXTRACT}(pos(x), MASK(1)) \rightarrow De-linearization$								
4: $j = \mathbb{EXTRACT}(pos(x), MASK(2))$								
5: $k = \mathbb{EXTRACT}(pos(x), MASK(3))$								
6: <b>for</b> $r = 1,, R$ <b>do</b>								
7: $\mathbf{Temp}_{l}(i - T_{l,1}^{s}, r) + = val(x) \times \mathbf{B}(j, r) \times \mathbf{C}(k, r)$								
8: $\mathbb{ATOMIC}(\tilde{\mathbf{A}}(i,r) + = val(x) \times \mathbf{B}(j,r) \times \mathbf{C}(k,r))$								
9: end for								
10: <b>end for</b>								
11: end for								

Algorithm Adaptive parallel execution of mode-1 MTTKRP-									
ALTO kernel. ALTO automatically uses local storage or atomics,									
based on the reuse of output fibers, to resolve the update conflicts.									
<b>Input:</b> A third-order ALTO sparse tensor $X \in \mathbb{R}^{I \times J \times K}$ with $M$ nonzero elements, dense factor matrices $\mathbf{A} \in \mathbb{R}^{I \times R}$ , $\mathbf{B} \in \mathbb{R}^{J \times R}$ , and $\mathbf{C} \in \mathbb{R}^{K \times R}$									
<b>Output:</b> Updated dense factor matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{I \times R}$									
1: for $l = 1,, L$ in parallel do $\triangleright$ ALTO line segments.									
2: for $\forall x \in X_l$ do									
3: $i = \mathbb{EXTRACT}(pos(x), MASK(1)) $ ▷ De-linearization.									
4: $j = \mathbb{EXTRACT}(pos(x), MASK(2))$									
5: $k = \mathbb{EXTRACT}(pos(x), MASK(3))$									
6: <b>for</b> $r = 1,, R$ <b>do</b>									
7: $\mathbf{Temp}_l(i - T^s_{l,1}, r) + = val(x) \times \mathbf{B}(j, r) \times \mathbf{C}(k, r)$									
8: $\mathbb{ATOMIC}(\tilde{\mathbf{A}}(i,r) + = val(x) \times \mathbf{B}(j,r) \times \mathbf{C}(k,r))$									
9: end for									
10: <b>end for</b>									
11: end for									

Algorithm Adaptive parallel execution of mode-1 MTTKRP-								
ALTO kernel. ALTO automatically uses local storage or atomics,								
based on the reuse of output fibers, to resolve the update conflicts.								
<b>Input:</b> A third-order ALTO sparse tensor $X \in \mathbb{R}^{I \times J \times K}$ with $M$ nonzero elements, dense factor matrices $\mathbf{A} \in \mathbb{R}^{I \times R}$ , $\mathbf{B} \in \mathbb{R}^{J \times R}$ , and $\mathbf{C} \in \mathbb{R}^{K \times R}$								
<b>Output:</b> Updated dense factor matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{I \times R}$								
1: for $l = 1,, L$ in parallel do $\triangleright$ ALTO line segments.								
for $\forall x \in X_l$ do								
3: $i = \mathbb{EXTRACT}(pos(x), MASK(1)) $ ▷ De-linearization.								
4: $j = \mathbb{EXTRACT}(pos(x), MASK(2))$								
5: $k = \mathbb{EXTRACT}(pos(x), MASK(3))$								
6: <b>for</b> $r = 1,, R$ <b>do</b>								
7: $\mathbf{Temp}_{l}(i - T_{l,1}^{s}, r) + = val(x) \times \mathbf{B}(j, r) \times \mathbf{C}(k, r)$								
8: $\mathbb{ATOMIC}(\tilde{\mathbf{A}}(i,r) + = val(x) \times \mathbf{B}(j,r) \times \mathbf{C}(k,r))$								
9: end for								
10: <b>end for</b>								
11: end for								

**Algorithm** Adaptive parallel execution of mode-1 MTTKRP-ALTO kernel. ALTO automatically uses local storage or atomics , based on the reuse of output fibers, to resolve the update conflicts.

**Input:** A third-order ALTO sparse tensor  $X \in \mathbb{R}^{I \times J \times K}$  with M nonzero elements, dense factor matrices  $\mathbf{A} \in \mathbb{R}^{I \times R}$ ,  $\mathbf{B} \in \mathbb{R}^{J \times R}$ , and  $\mathbf{C} \in \mathbb{R}^{K \times R}$ 

**Output:** Updated dense factor matrix  $\tilde{\mathbf{A}} \in \mathbb{R}^{I \times R}$ 

1: for l = 1, ..., L in parallel do  $\triangleright$  ALTO line segments.

```
2: for \forall x \in X_l do
```

3:  $i = \mathbb{EXTRACT}(pos(x), MASK(1)) \rightarrow De-linearization.$ 

- 4:  $j = \mathbb{EXTRACT}(pos(x), MASK(2))$
- 5:  $k = \mathbb{EXTRACT}(pos(x), MASK(3))$
- 6: **for** r = 1, ..., R **do** 
  - $\mathbf{Temp}_{l}(i T_{l,1}^{s}, r) + = val(x) \times \mathbf{B}(j, r) \times \mathbf{C}(k, r)$  $\mathbb{ATOMIC}(\tilde{\mathbf{A}}(i, r) + = val(x) \times \mathbf{B}(j, r) \times \mathbf{C}(k, r))$
- 8:
- 9: end for
- 10: **end for**

11: **end for** 

7:

12:for b = 1, ..., I in parallel do> Pull-based accumulation.13:for  $\forall l$  where  $b \in [T_{l,1}^s, T_{l,1}^e]$  do14:for r = 1, ..., R do15: $\tilde{A}(b, r) + = \operatorname{Temp}_l(b - T_{l,1}^s, r)$ 16:end for17:end for18:end for19:return  $\tilde{A}$ 

#### Platform

- Cascade Lake Platinum 8280
  - 2 sockets, 28 cores each, fixed @1.8 GHz
- Intel compiler and runtime
  - "-O3 -xCORE-AVX512 -qopt-zmm-usage=high" flags to fully utilize vector units
  - Parallel runs utilize all hardware threads (112) with thread pinning via LIKWID

Configurations

- MTTKRP with double-precision and native word (64-bit) data types
- Number of iterations = 100
- 15 tensors from FROSTT and HaTen2 datasets
  - NNZ range: [3 M, 4.7 B], Density range: [1.5E<sup>-02</sup>, 4.3E<sup>-15</sup>]
- We consider the state-of-the-art sparse tensor libraries for CPUs
  - ParTI: coordinate (COO) and hierarchical COO (HiCOO) formats
  - SPLATT: compressed sparse fiber (CSF) and tiled CSF (CSF-tile)

## Parallel performance (all-modes MTTKRP)

- 47x geomean speedup for high reuse cases (84% of realizable theoretical speedup)
- 22x geomean speedup for memory-bound cases
- Prior formats' performance varies widely due to sensitivity to irregularities in shape and data



## Parallel performance (all-modes MTTKRP)

- 47x geomean speedup for high reuse cases (84% of realizable theoretical speedup)
- 22x geomean speedup for memory-bound cases
- Prior formats' performance varies widely due to sensitivity to irregularities in shape and data





#### Tensor storage

ALTO format storage < COO</p>



- ALTO format storage < COO</p>
- Mode-specific formats' storage > COO
- imposing a tilling over CSF increases storage requirements



- ALTO format storage < COO</p>
- Mode-specific formats' storage > COO
- imposing a tilling over CSF increases storage requirements
- HiCOO's memory consumption depending on
  - a) block/superblock sizes
  - b) spatial distribution of nonzero elements



- ALTO format storage < COO</p>
- Mode-specific formats' storage > COO
- imposing a tilling over CSF increases storage requirements
- HiCOO's memory consumption depending on
  - a) block/superblock sizes
  - b) spatial distribution of nonzero elements



#### ALTO generation cost



## ALTO generation cost

ALTO substantially decreases sorting costs



## ALTO generation cost

- ALTO substantially decreases sorting costs
- Block-based formats require expensive clustering & scheduling of NNZs



## ALTO generation cost

- ALTO substantially decreases sorting costs
- Block-based formats require expensive clustering & scheduling of NNZs
- CSF formats generated from *presorted* tensors; still slower than ALTO



## Conclusion

 The ALTO format efficiently encodes sparse tensors of arbitrary dimensionality and spatial distributions using a single mode-agnostic representation

## Conclusion

- The ALTO format efficiently encodes sparse tensors of arbitrary dimensionality and spatial distributions using a single mode-agnostic representation
- ALTO-based tensor decomposition operations proved to *outperform* an oracle that selects the best state-of-the-art format due to
  - 1. superior workload balance,
  - 2. adaptive synchronization,
  - 3. compact encoding

# Conclusion

- The ALTO format efficiently encodes sparse tensors of arbitrary dimensionality and spatial distributions using a single mode-agnostic representation
- ALTO-based tensor decomposition operations proved to *outperform* an oracle that selects the best state-of-the-art format due to
  - 1. superior workload balance,
  - 2. adaptive synchronization,
  - 3. compact encoding
- ALTO achieves **84% parallel efficiency** in high reuse cases, 63% for all cases
  - 7.9x (13.2x w/ rank specialization) geometric-mean speedup over the best general formats
  - 4.3x geometric-mean compression ratio over the best mode-specific formats

# Future Work

 Working on distributed-memory execution of key tensor decomposition algorithms

Exploring other massively parallel platforms (GPGPUs)

Investigating additional sparse tensor kernels/algorithms
 provide an ALTO library

# 

ALTO pre-print: arxiv.org/abs/2102.10245



# Backup

Tensor	Dimensions	#NNZs	Density	Fib. reuse
LBNL	$1.6K \times 4.2K \times 1.6K \times$	1.7 <i>M</i>	$4.2 \times 10^{-14}$	Limited
	$4.2K \times 868.1K$			
NIPS	$2.5K \times 2.9K \times 14K \times 17$	3.1 <i>M</i>	$1.8 \times 10^{-06}$	High
UBER	$183 \times 24 \times 1.1 K \times 1.7 K$	3.3 <i>M</i>	$3.8 \times 10^{-04}$	High
CHICAGO	$6.2K \times 24 \times 77 \times 32$	5.3 <i>M</i>	$1.5 \times 10^{-02}$	High
VAST	$165.4K \times 11.4K \times 2 \times$	26 <i>M</i>	$7.8 \times 10^{-07}$	High
	$100 \times 89$			
DARPA	22.5 <i>K</i> ×22.5 <i>K</i> ×23.8 <i>M</i>	28.4 <i>M</i>	$2.4 \times 10^{-09}$	Limited
ENRON	$6K \times 5.7K \times 244.3K \times$	54.2 <i>M</i>	$5.5 \times 10^{-09}$	High
	1.2K			
NELL-2	$12.1K \times 9.2K \times 28.8K$	76.9 <i>M</i>	$2.4 \times 10^{-05}$	High
FB-M	$23.3M \times 23.3M \times 166$	99.6 <i>M</i>	$1.1 \times 10^{-09}$	Limited
FLICKR	319.7K $ imes$ $28.2M$ $ imes$	112.9 <i>M</i>	$1.1 \times 10^{-14}$	Limited
	$1.6M \times 731$			
DELI	532.9 $K$ $\times$ 17.3 $M$ $\times$	140.1 <i>M</i>	$4.3 \times 10^{-15}$	Medium
	$2.5M \times 1.4K$			
NELL-1	$2.9M \times 2.1M \times 25.5M$	143.6 <i>M</i>	$9.1 \times 10^{-13}$	Medium
AMAZON	$4.8M \times 1.8M \times 1.8M$	1.7 <i>B</i>	$1.1 \times 10^{-10}$	High
PATENTS	$46 \times 239.2K \times 239.2K$	3.6 <i>B</i>	$1.4 \times 10^{-03}$	High
REDDIT	$8.2M \times 177K \times 8.1M$	4.7 <i>B</i>	$4.0 \times 10^{-10}$	High