

Large Scale Matrix Factorization: Systems and Acceleration


Wei Tan

IBM T. J. Watson Research Center

wtan@us.ibm.com

<http://github.com/cumf>

<http://researcher.ibm.com/person/us-wtan>




Agenda


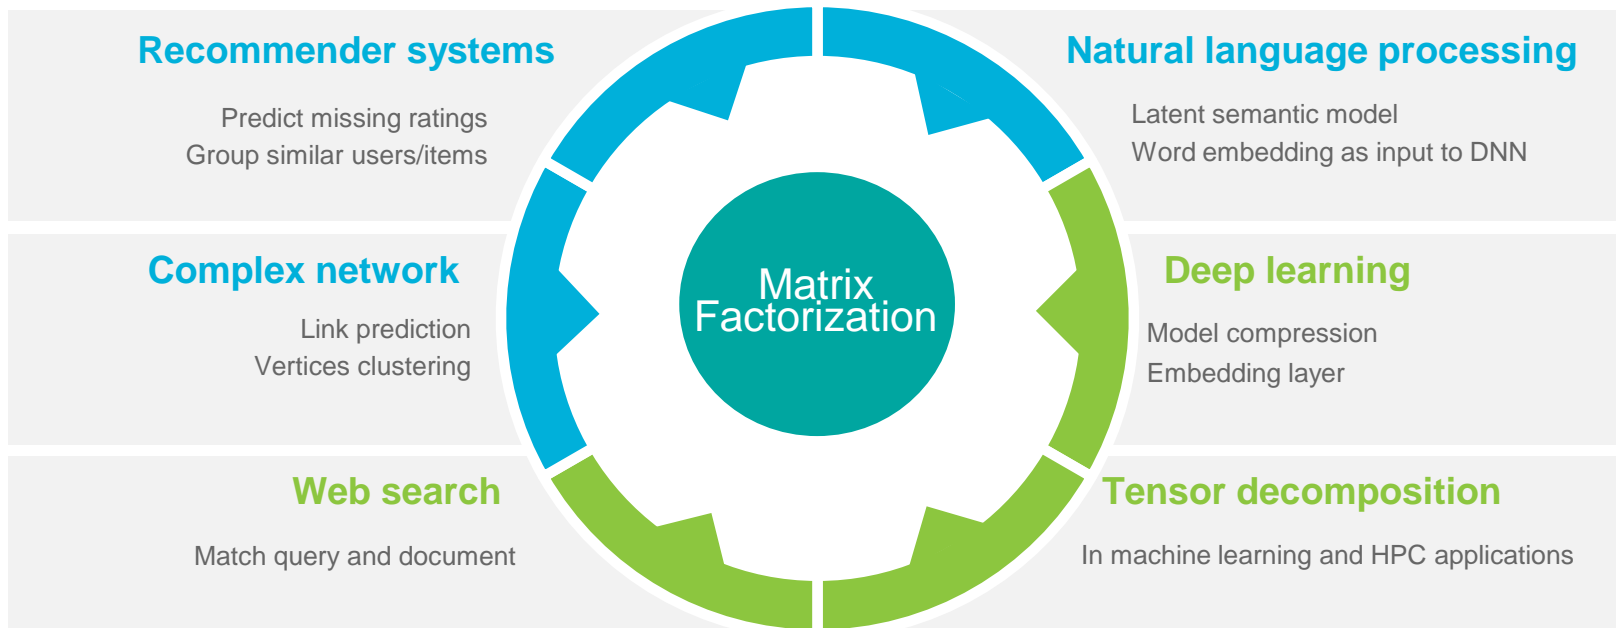
- **Fei's talk covers the formalism/theory/math of MF**
- **My talk focuses on “how to run it fast, scalable and cost-efficient”**
 - Matrix factorization, SGD and ALS (10 min)
 - Parallelize and accelerate SGD and ALS (20 min)
 - GPU accelerated SGD and ALS (20 min)
 - Conclusion and QA (10 min)



Matrix Factorization



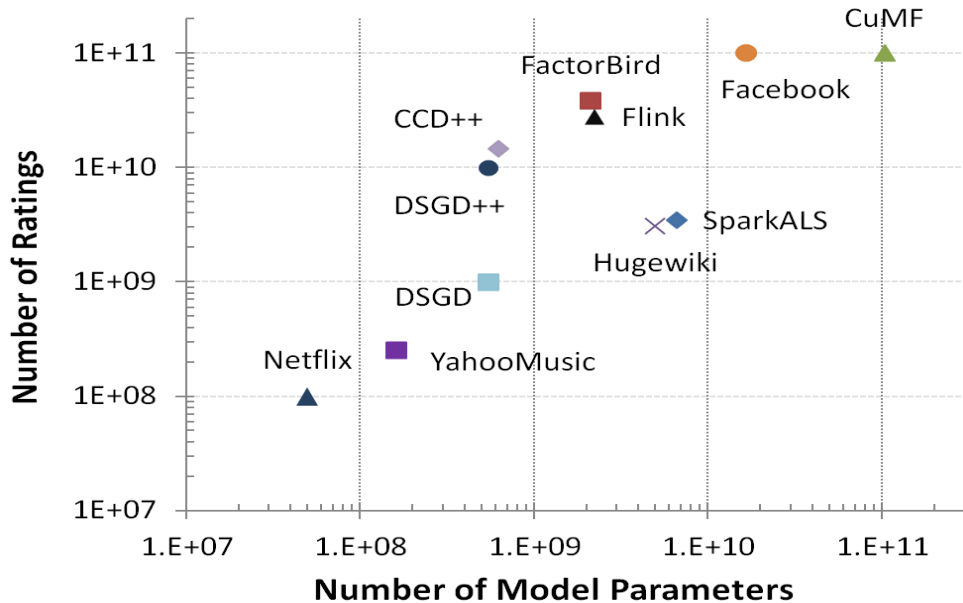
Matrix Factorization is a Key



- Supported in cuMF
- To be supported



Challenge: MF needs to be fast, scalable, economic



- **Fast**
 - Recommend/update timely
- **Scalable**
 - Facebook: 100 B ratings, 1 B users
- **Economic**
 - Avoid large infrastructure



To Solve MF: SGD

$$R \approx X \cdot \Theta^T$$


$$J = \sum_{u,v} (r_{uv} - \mathbf{x}_u^T \boldsymbol{\theta}_v)^2 + \lambda (\sum_u n_{x_u} \|\mathbf{x}_u\|^2 + \sum_v n_{\theta_v} \|\boldsymbol{\theta}_v\|^2)$$

Stochastic gradient descent (**SGD**)

$$\mathbf{x}_u = \mathbf{x}_u - \alpha [(\mathbf{x}_u^T \boldsymbol{\theta}_v - r_{uv}) \boldsymbol{\theta}_v + \lambda \mathbf{x}_u]$$

$$\boldsymbol{\theta}_v = \boldsymbol{\theta}_v - \alpha [(\mathbf{x}_u^T \boldsymbol{\theta}_v - r_{uv}) \mathbf{x}_u + \lambda \boldsymbol{\theta}_v]$$

- Update takes one rating at a time
- Vector inner product: memory bound
- Need many light epochs
- Parallelize: non-trivial
- Handle dense (implicit) ratings: no



To Solve MF: ALS

$$R \approx X \cdot \Theta^T$$


$$J = \sum_{u,v} (r_{uv} - \mathbf{x}_u^T \boldsymbol{\theta}_v)^2 + \lambda (\sum_u n_{x_u} \|\mathbf{x}_u\|^2 + \sum_v n_{\theta_v} \|\boldsymbol{\theta}_v\|^2)$$

Alternating Least Square (ALS)

$$\sum_{r_{uv} \neq 0} (\boldsymbol{\theta}_v \boldsymbol{\theta}_v^T + \lambda I) \cdot \mathbf{x}_u = \Theta^T \cdot R_{u*}^T$$

$$\sum_{r_{uv} \neq 0} (\mathbf{x}_u \mathbf{x}_u^T + \lambda I) \cdot \boldsymbol{\theta}_v = X^T \cdot R_{*v}$$

- Update takes ALL rating at a time
- Vector outer product & solve: compute bound
- Need few heavy epochs
- Parallelize: straightforward
- Handle dense (implicit) ratings: yes




To Solve MF: CD

$$R \approx X \cdot \Theta^T$$

$$J = \sum_{u,v} (r_{uv} - \mathbf{x}_u^T \boldsymbol{\theta}_v)^2 + \lambda (\sum_u n_{x_u} \|\mathbf{x}_u\|^2 + \sum_v n_{\theta_v} \|\boldsymbol{\theta}_v\|^2)$$

Coordinate descent (CD)

- Similar to ALS
- But update one coordinate of \mathbf{x}_u and $\boldsymbol{\theta}_v$ at a time



To Parallelize ALS

$$R \approx X \cdot \Theta^T$$


$$J = \sum_{u,v} (r_{uv} - \mathbf{x}_u^T \boldsymbol{\theta}_v)^2 + \lambda \left(\sum_u n_{x_u} \|\mathbf{x}_u\|^2 + \sum_v n_{\theta_v} \|\boldsymbol{\theta}_v\|^2 \right)$$

Alternating Least Square (ALS)

$$\sum_{r_{uv} \neq 0} (\boldsymbol{\theta}_v \boldsymbol{\theta}_v^T + \lambda I) \cdot \mathbf{x}_u = \Theta^T \cdot R_{u*}^T$$

$$\sum_{r_{uv} \neq 0} (\mathbf{x}_u \mathbf{x}_u^T + \lambda I) \cdot \boldsymbol{\theta}_v = X^T \cdot R_{*v}$$

- Solve x_u s independently (θ_v s thereafter)
- Parallelize the solve of x_u s on multiple nodes
 - Replicate Θ
 - Partially replicate Θ
 - Split Θ on multiple nodes






Parallelize SGD: Hogwild!

$$\mathbf{x}_u = \mathbf{x}_u - \alpha[(\mathbf{x}_u^T \boldsymbol{\theta}_v - r_{uv})\boldsymbol{\theta}_v + \lambda \mathbf{x}_u]$$

$$\boldsymbol{\theta}_v = \boldsymbol{\theta}_v - \alpha[(\mathbf{x}_u^T \boldsymbol{\theta}_v - r_{uv})\mathbf{x}_u + \lambda \boldsymbol{\theta}_v]$$

Hogwild! [Niu et al. 2011]


-- parallel SGD converges despite of (occasional) update conflict



worker 1 \mathbf{x}_{u1}


worker 2 \mathbf{x}_{u2}

worker 3 \mathbf{x}_{u3}



Hogwild! Not Good Enough?

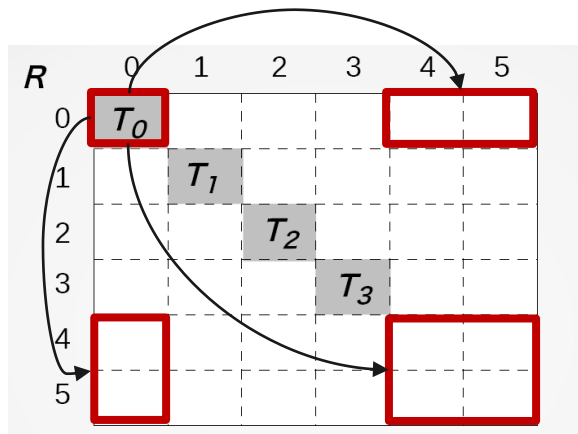


Random sampling hurts cache performance (on GPUs)
-- hardware cannot prefetch



Parallelize SGD: Matrix blocking

1. wave 1

2. wave 2



- Cons: all 4 workers need to complete before the next wave

- Solution: more blocks than workers
– 6*6 blocks, 4 workers
- Worker 0 can immediately pick up another block when T_0 is done
- Cons: scheduling overhead

MF methods with SGD, ALS and CCD


	CPU	GPU
SGD	<p>lock-free: workers independently sample & update single-node: HogWild! [15] multi-nodes: FactorBird [8], Petuum [28]</p> <p>blocking: workers on non-overlapping blocks blockDim=#workers: DSGD [13] blockDim>#workers: LIBMF [9], NOMAD [12], DSGD++ [18] nested blocking: dcMF [29], MLGF-MF [16]</p>	<p>single and multiple GPUs: cuMF_SGD [27] implements lock-free and blocking, and is memory-optimized with reduced precision.</p>
ALS	<p>replicate features: PALS [30], DALs [18] partial replicate: SparkALS [31], GraphLab [32], Sparkler [11] rotate: Facebook [5] approximate ALS: [33]</p>	<p>single GPU: BIDMach [24], ALS-HPC [25] single and multiple GPUs: cuMF [23]</p>
CCD	<p>multi-core and multi node: CCD++ [14]</p>	



Challenge: compute and memory capacity of CPU

Table 2: Compute and memory complexity per epoch: ALS vs. SGD

		Compute (C)	Memory (M)	C/M
ALS	get_hermitian	$\mathcal{O}(N_z f^2)$	$\mathcal{O}(N_z f + (m + n) f^2)$	f
	direct solver	$\mathcal{O}((m + n) f^3)$	$\mathcal{O}((m + n) f^2)$	f
SGD		$\mathcal{O}(N_z f)$	$\mathcal{O}(N_z f)$	1




CPU offers: 1 T flops, 80 GB/s
 $f=100$, per epoch

- ALS floating-point operations
 - Netflix: 1.5 T
 - Hugewiki: 80 T
 - Facebook: 2000 T
- SGD memory transfer
 - Netflix: 80 GB
 - Hugewiki: 2.4 TB
 - Facebook: 80 TB
- **>> CPU flops and BW capacity**




GPU vs. CPU: compute FLOPS and memory bandwidth

Theoretical Peak Performance, Single Precision



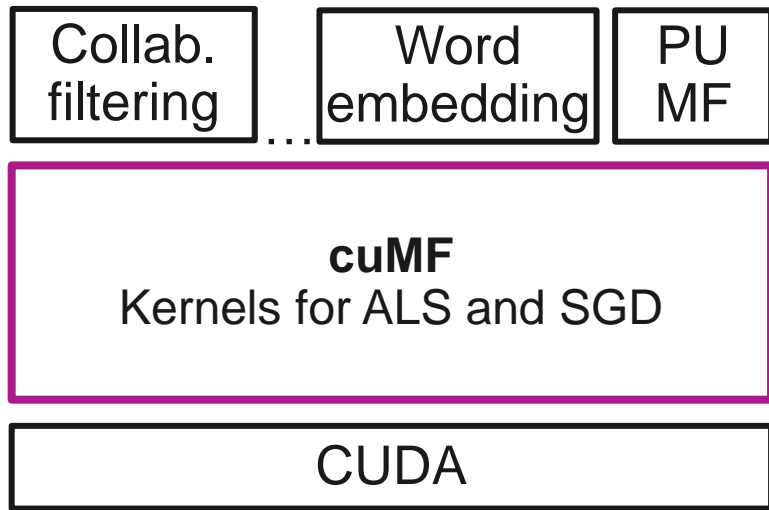
Theoretical Peak Memory Bandwidth Comparison



- Raw performance: **1 GPU ≈ 10 x CPU**
- Practical performance due to slow interconnection
1 GPU > 10 x CPU
4 GPU >> 40 x CPU



Goal: a CUDA library for MF



Fast

- Fast training
- Update model quickly

Scalable

- Deal with big data
- Exploit fast interconnection

Cost efficient

- Fully utilize flops or BW
- Cheaper than CPU solutions



cuMF
cuda Matrix Factorization

CUDA-based matrix factorization libraries.

CUDA-based matrix factorization libraries developed by IBM Research and friends.

<http://researcher.ibm.com/person/us-wtan> ✉ weitan@ieee.org



Challenges of ALS


- ALS needs to solve many:

$$\sum_{r_{uv} \neq 0} (\theta_v \theta_v^T + \lambda I) \cdot x_u = \Theta^T \cdot R_{u*}^T$$

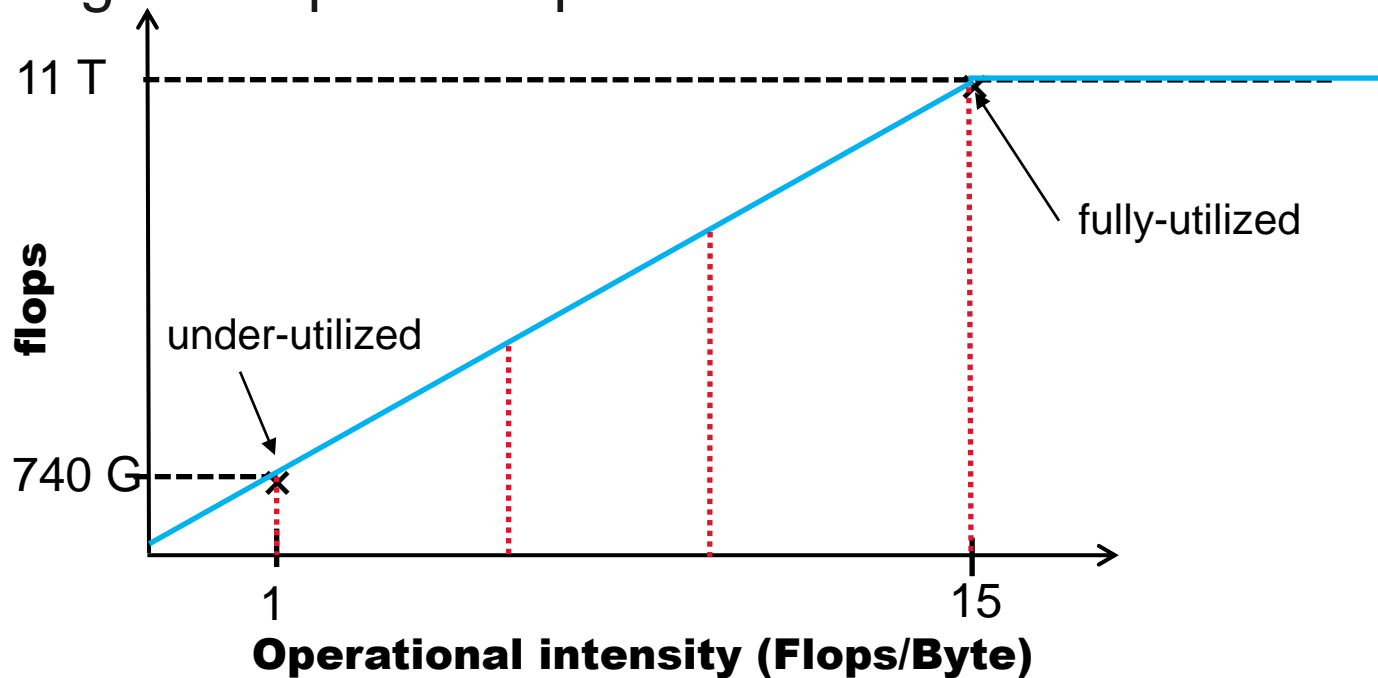
Challenge 2: LU or Cholesky solver: compute intensive

- Challenge 3: Single GPU can NOT handle **big** m , n and Nz

- Challenge 1: access and aggregate many θ_v s: memory **irregular** and compute **intensive**



Challenge 1: improve flops



- Nvidia Pascal: Memory BW: 740 GB/sec, compute: 11 Tflops
- Higher flops \rightarrow higher op intensity (more flops per byte) \rightarrow caching!




Address challenge 1: memory-optimized ALS

- To obtain


$$\sum_{r_{uv} \neq 0} (\theta_v \theta_v^T + \lambda I)$$

Register


Shared memory



1. non-coalesced read



2. stage into smem



Block 30
 $(x, y) = (3, 6) = \theta^{(6)}(\theta^{(3)})^T$

3. tile and aggregate in register



Address Challenge 2: exact solver is compute intensive

$$\sum_{r_{uv} \neq 0} (\theta_v \theta_v^T + \lambda I) \cdot \mathbf{x}_u = \Theta^T \cdot R_{u*}^T$$

ALS iter 1

Exactly solve X
Exactly solve Θ

ALS iter 2

Exactly solve X
Exactly solve Θ

...

ALS iter n

Exactly solve X
Exactly solve Θ

$O(f^3)$

ALS iter 1

Approx. solve X
Approx. solve Θ

ALS iter 2

Approx. solve X
Approx. solve Θ

...

ALS iter n

Approx. solve X
Approx. solve Θ

$O(f^2)$



Algorithm 1 The CG Solver.

```


1: procedure CGSOLVE( $A, \mathbf{b}, \mathbf{x}, f_s, \epsilon$ )
2:    $\mathbf{r} = \mathbf{b} - A \cdot \mathbf{x}$ 
3:    $\mathbf{p} = \mathbf{r}$ 
4:    $r_{old} = \mathbf{r}^T \cdot \mathbf{r}$ 
5:   for  $j = 1 : f_s$  do
6:      $\mathbf{a}_p = A \cdot \mathbf{p}$ 
7:      $\alpha = r_{old} / (\mathbf{p}^T \cdot \mathbf{a}_p)$ 
8:      $\mathbf{x} = \mathbf{x} + \alpha \mathbf{p}$ 
9:      $\mathbf{r} = \mathbf{r} - \alpha \mathbf{p}$ 
10:     $r_{new} = \mathbf{r}^T \cdot \mathbf{r}$ 
11:    if  $\sqrt{r_{new}} < \epsilon$  then
12:      break
13:    end if
14:     $\mathbf{p} = \mathbf{r} + (r_{new} / r_{old}) \mathbf{p}$ 
15:     $r_{old} = r_{new}$ 
16:  end for
17:  return  $\mathbf{x}$ 
18: end procedure

```

use $f_s \ll f$



Address Challenge 2: use CG solver



- Solver time: CG = $\frac{1}{4}$ LU
- CG solver is memory- (instead of compute-) bound
 - CG w/ FP16 = $\frac{1}{2}$ CG w/ FP32

Algorithm 1 The CG Solver.


```
1: procedure CGSOLVE( $A, \mathbf{b}, \mathbf{x}, f_s, \epsilon$ )
2:    $\mathbf{r} = \mathbf{b} - A \cdot \mathbf{x}$ 
3:    $\mathbf{p} = \mathbf{r}$ 
4:    $r_{old} = \mathbf{r}^T \cdot \mathbf{r}$ 
5:   for  $j = 1 : f_s$  do
6:      $\mathbf{a}_p = A \cdot \mathbf{p}$ 
7:      $\alpha = r_{old} / (\mathbf{p}^T \cdot \mathbf{a}_p)$ 
8:      $\mathbf{x} = \mathbf{x} + \alpha \mathbf{p}$ 
9:      $\mathbf{r} = \mathbf{r} - \alpha \mathbf{a}_p$ 
10:     $r_{new} = \mathbf{r}^T \cdot \mathbf{r}$ 
11:    if  $\sqrt{r_{new}} < \epsilon$  then
12:      break
13:    end if
14:     $\mathbf{p} = \mathbf{r} + (r_{new} / r_{old}) \mathbf{p}$ 
15:     $r_{old} = r_{new}$ 
16:  end for
17:  return  $\mathbf{x}$ 
18: end procedure
```



Address Challenge 3: scale-up ALS on multiple GPUs

- **Model** parallel: solve a portion of the model

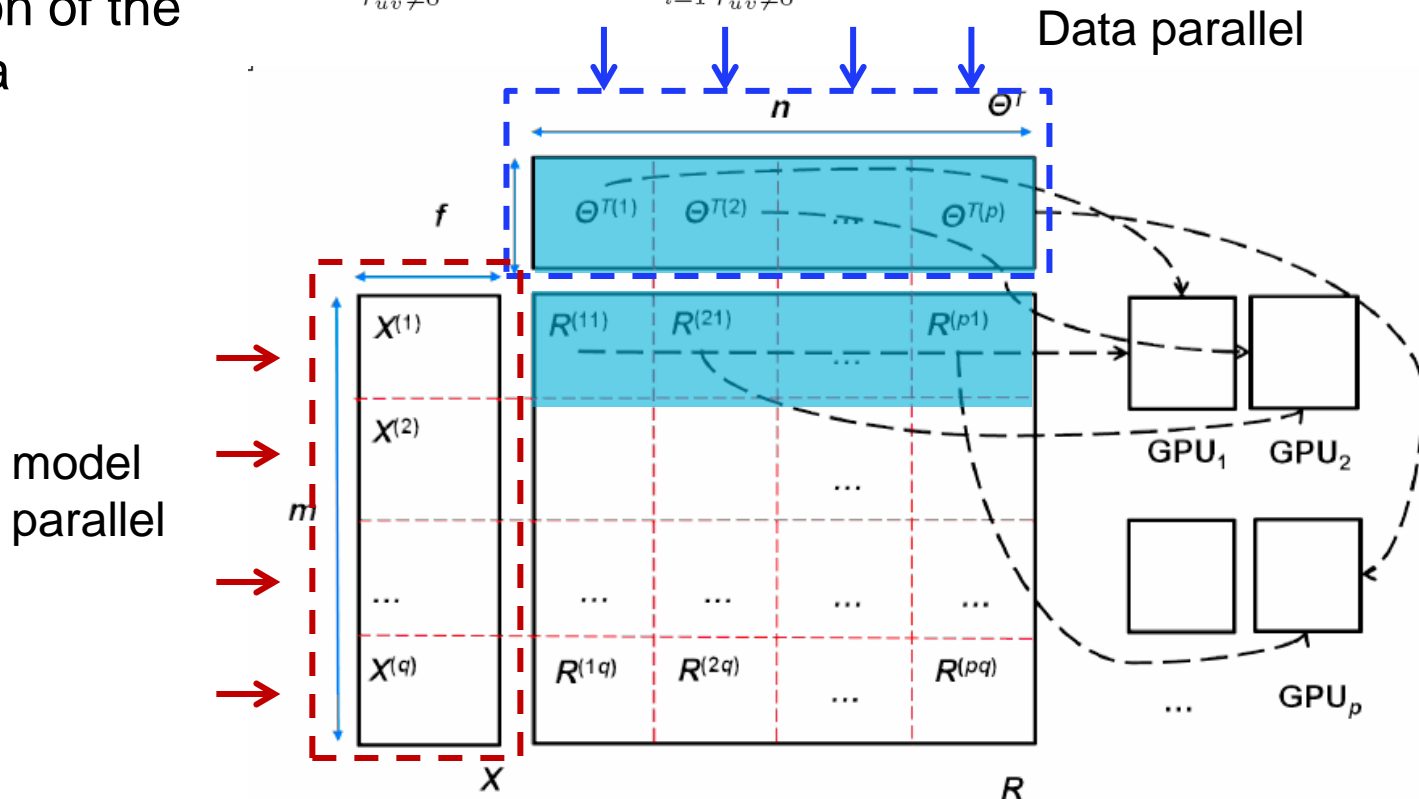
$$\sum_{r_{uv} \neq 0} (\theta_v \theta_v^T + \lambda I) \cdot x_u = \Theta^T \cdot R_{u*}^T$$



Address Challenge 3: scale-up ALS on multiple GPUs

- **Data parallel:** solve with a portion of the training data

$$A_u = \sum_{r_{uv} \neq 0} (\theta_v \theta_v^T + \lambda I) = \sum_{i=1}^p \sum_{r_{uv} \neq 0}^{GPU_i} (\theta_v \theta_v^T + \lambda I)$$



Recap: challenges of ALS

- ALS needs to solve many:

$$\sum_{r_{uv} \neq 0} (\theta_v \theta_v^T + \lambda I) \cdot x_u = \Theta^T \cdot R_{u*}^T$$

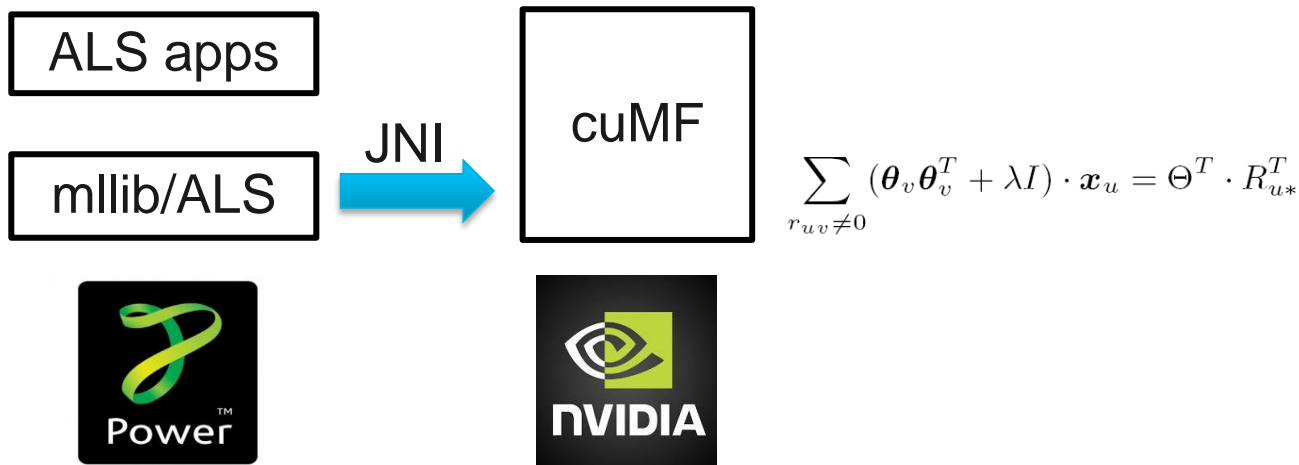
- Challenge 1: access and aggregate many θ_v s: memory **irregular** and compute **intensive**
-- use register, smem and non-coalesced read

Challenge 2: LU or Cholesky solver: compute intensive
-- use approximate CG solver and FP16

- Challenge 3: Single GPU can NOT handle **big** m , n and Nz
-- use model and data parallelism, and topology aware reduction



Connect cuMF to Spark MLlib



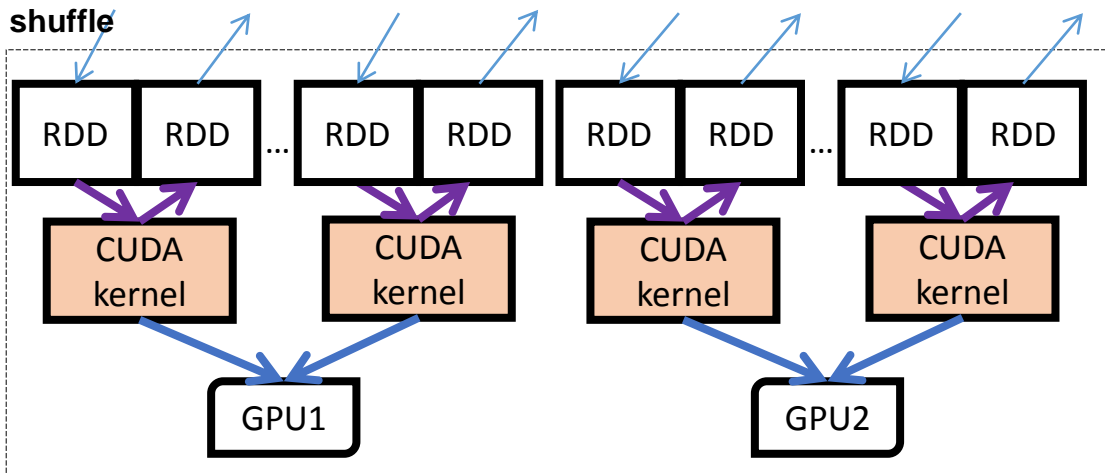
- Spark applications relying on mllib/ALS need no change
- Modified mllib/ALS detects GPU and offload maxtix computation
- Leverage the best of Spark (scale-out) and GPU (scale-up)

<https://github.com/IBMSparkGPU/CUDA-MLlib>
<http://www-01.ibm.com/support/docview.wss?uid=swg21983421>

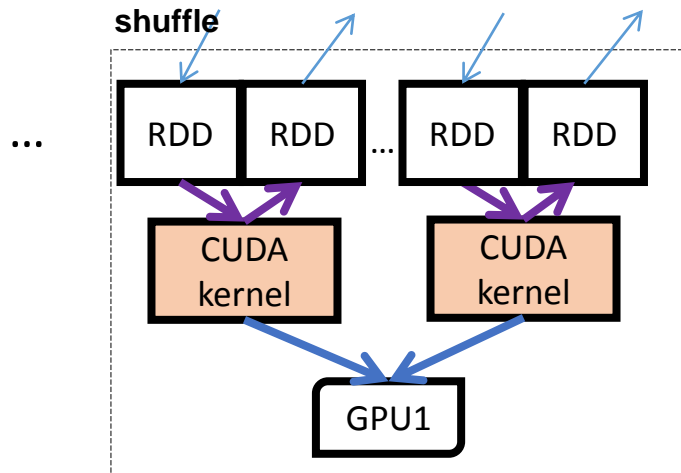


Connect cuMF to Spark MLlib

- RDD on CPU: to distribute rating data and shuffle parameters
- Solver on GPU: to form and solve $\sum_{r_{uv} \neq 0} (\theta_v \theta_v^T + \lambda I) \cdot x_u = \Theta^T \cdot R_{u*}^T$
- Able to run on multiple nodes, and multiple GPUs per node



1 Power 8 node + 2 K40



1 Power 8 node + 2 K40

Challenges of SGD

- Iterate over all ratings and do this in sequence:

$$\mathbf{x}_u = \mathbf{x}_u - \alpha[(\mathbf{x}_u^T \boldsymbol{\theta}_v - r_{uv})\boldsymbol{\theta}_v + \lambda \mathbf{x}_u]$$

$$\boldsymbol{\theta}_v = \boldsymbol{\theta}_v - \alpha[(\mathbf{x}_u^T \boldsymbol{\theta}_v - r_{uv})\mathbf{x}_u + \lambda \boldsymbol{\theta}_v]$$

- Memory bound

1. update kernel

```

1 //read the sample
2 int index = get_index_of_sample();
3 float r = __ldg(&samples[index].r);
4 int u = __ldg(&samples[index].u);
5 int v = __ldg(&samples[index].v);
6
7 //read the p & q
8 int p_offset = u*k, q_offset = v*k;
9 float tmp_p1 = __half2float(p[p_offset + threadIdx.x]), tmp_p2 = __half2float(p[p_offset + threadIdx.x + 32]);
10 float tmp_q1 = __half2float(q[q_offset + threadIdx.x]), tmp_q2 = __half2float(q[q_offset + threadIdx.x + 32]);
11
12 //get dot product.
13 float tmp_product = tmp_p1*tmp_q1 + tmp_p2*tmp_q2;
14 tmp_product += __shfl_down(tmp_product, 16);
15 tmp_product += __shfl_down(tmp_product, 8);
16 tmp_product += __shfl_down(tmp_product, 4);
17 tmp_product += __shfl_down(tmp_product, 2);
18 tmp_product += __shfl_down(tmp_product, 1);
19 tmp_product = __shfl(tmp_product, 0);
20 float der = e - tmp_product;
21
22 //compute the derivative and update
23 p[p_offset + threadIdx.x + 0] = __float2half(tmp_p1 + tmp_lrate*(der*tmp_q1 - lambda_p*tmp_p1));
24 q[q_offset + threadIdx.x + 0] = __float2half(tmp_q1 + tmp_lrate*(der*tmp_p1 - lambda_q*tmp_q1));
25 p[p_offset + threadIdx.x + 32] = __float2half(tmp_p2 + tmp_lrate*(der*tmp_q2 - lambda_p*tmp_p2));
26 q[q_offset + threadIdx.x + 32] = __float2half(tmp_q2 + tmp_lrate*(der*tmp_p2 - lambda_q*tmp_q2));
    
```

Cache

ILP



Half precision

Memory Coalescing

Warp shuffle

Register Reuse

2. how to parallelize






Parallel worker 0
Parallel worker 1

(a) Hogwild

(b) Matrix Blocking


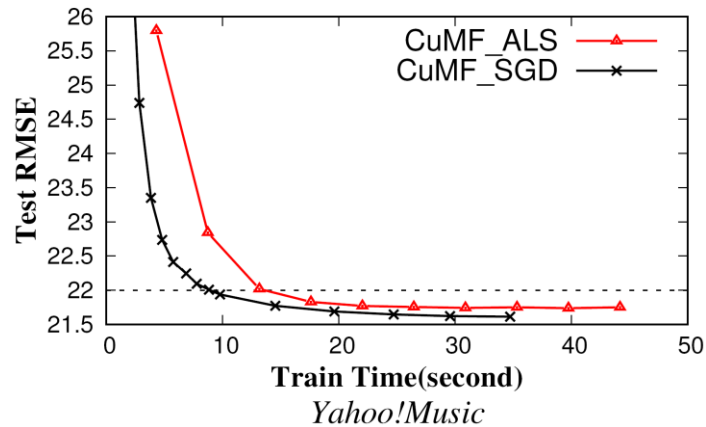



Experiment 1: is cuMF fast and scalable?



- cuMF_ALS w/ FP16 on **Maxwell** and **Pascal**
 - 1 GPU for Netflix and Yahoo
- LIBMF: 1 CPU w/ 40 threads
- NOMAD
 - 32 nodes for Netflix and Yahoo
 - 64 HPC nodes for Hugewiki
- **2-10x as fast**


Experiment 2: cuMF_ALS and cuMF_SGD (on Maxwell)



- ALS slightly slower than SGD on single GPU
- On big data set Hugewiki, ALS@4 GPU performs best -- SGD harder to parallel to multiple GPUs!



Experiment 3: is cuMF cost efficient?



- cuMF_ALS @4 Maxwell
≈ **1/10** SparkALS @50 nodes
≈ \$2.5/hr → **1/10** of 50 nodes
≈ **1%** of SparkALS's cost

Baseline	baseline config	#nodes	price /node/hr	cuMF speed	cuMF cost
NOMAD	m3.xlarge	32	\$0.27	10x	3%
SparkALS	m3.2xlarge	50	\$0.53	10x	1%
Factorbird	c3.2xlarge	50	\$0.42	6x	2%



Conclusion

- **Why** accelerate matrix factorization using GPU?
 - MF need to be fast, scalable, and economic
 - GPU offers ~10x flops, memory BW, and fast interconnect
- **How** cuMF tackles the challenges?
 - Optimize memory access, parallelism and communication
 - Approximate computing
 - Reduced precision
- **What** is the result?
 - Implemented ALS and SGD
 - Up to 10x as fast, 100x as cost-efficient
 - Use cuMF standalone, with Spark or Tensorflow



Thank you, questions?

- Faster and Cheaper: Parallelizing Large-Scale Matrix Factorization on GPUs. [HPDC](#) 2016
- CuMF_SGD: Fast and Scalable Matrix Factorization. [CoRR abs/1610.05838](#), 2016
- Code: <http://github.com/cuMF/>
- Blog: <http://ibm.biz/cumf-blog>
- Contact: Wei Tan, wtan@us.ibm.com



