UNIVERSITY OF OREGON

Distributed Tucker for Sparse Tensors

Venkatesan Chakaravarthy[1], **Jee Whan Choi[2]**, et al.
[1]IBM Research
[2]CIS, University of Oregon

SIAM CSE, Spokane, WA, March 1st, 2019

Tucker Decomposition

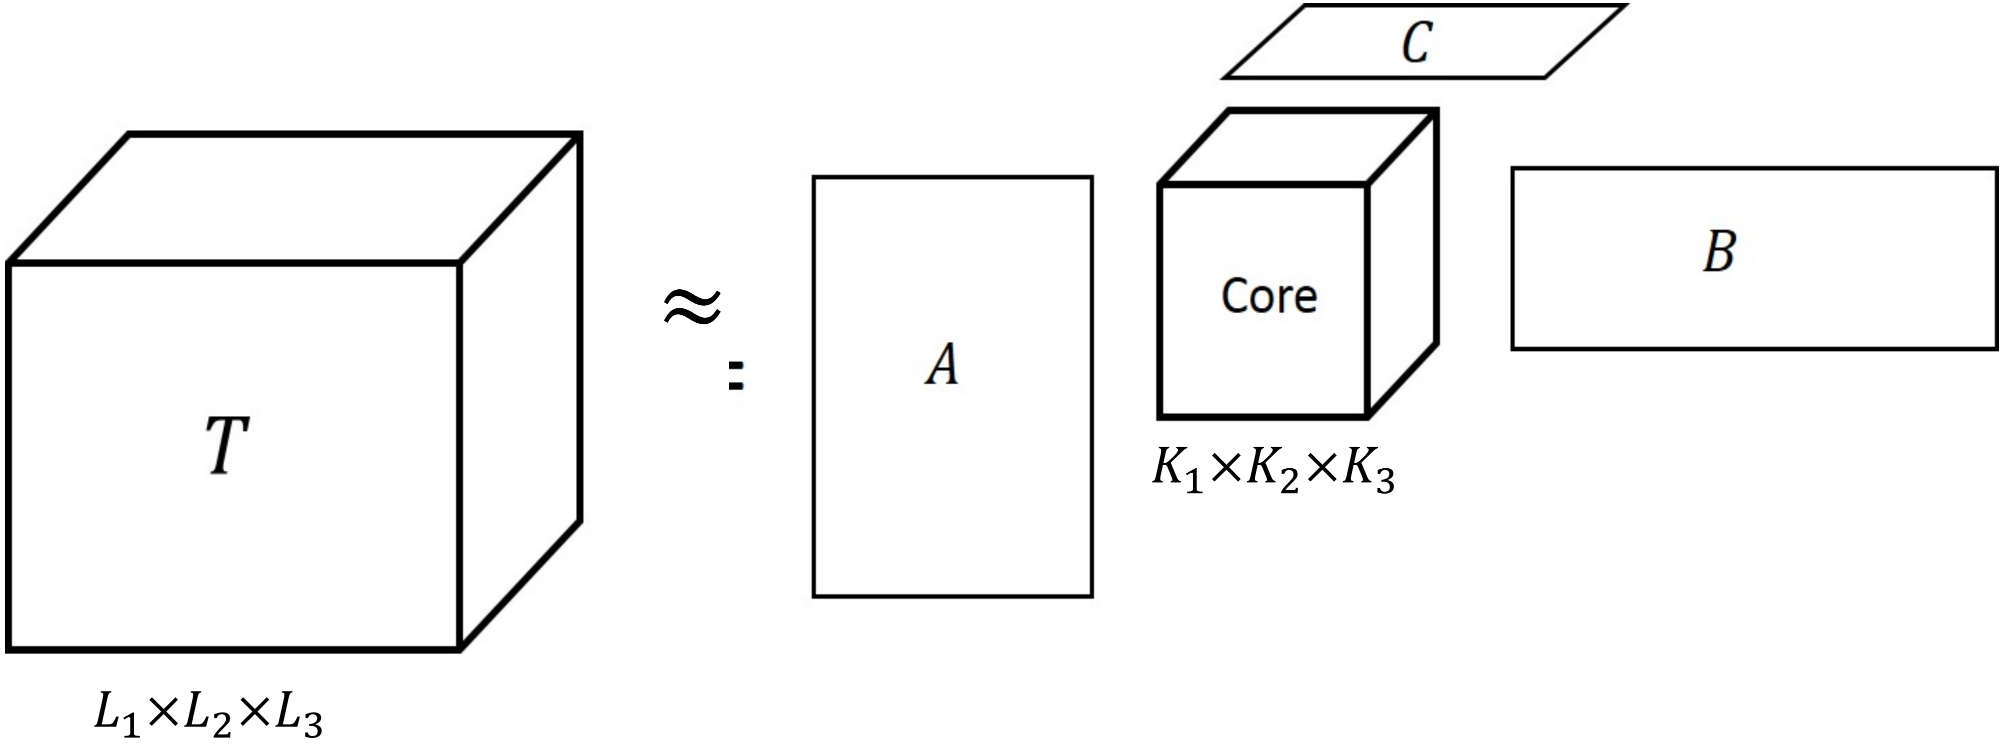

Higher Order Orthogonal Iteration (HOOI) Algorithm

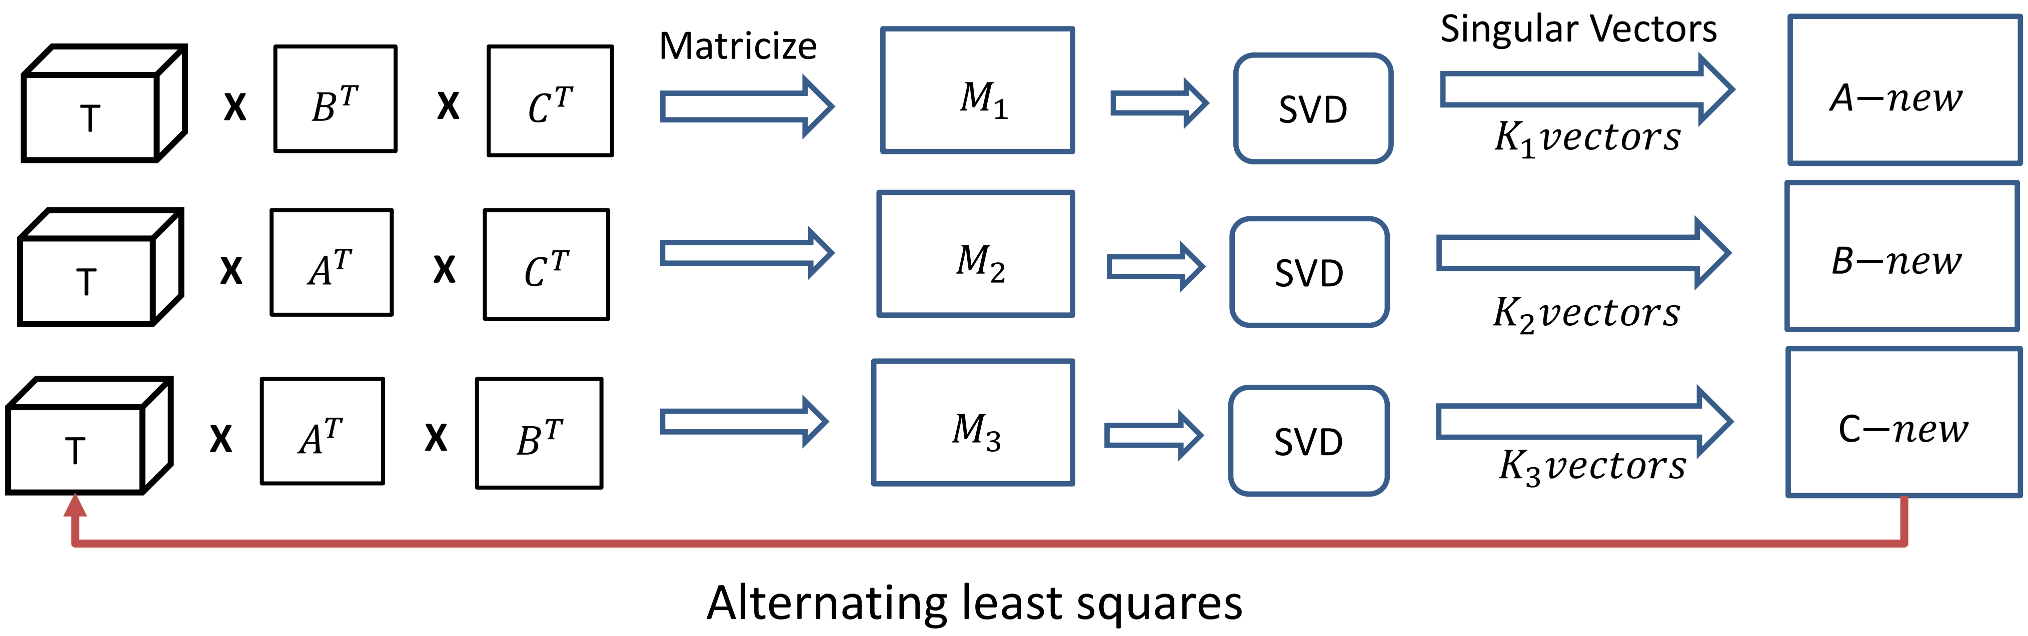

Sparse HOOI – Key Kernels

- TTM
  - Computation only - all schemes have same computational load (i.e., FLOPs)
  - Load balance
- SVD
  - Both computation and communication
  - Both computational load and communication volume are determined by load balance
- Factor Matrix Transfer (FMT)
  - Communication only
  - At the end of each HOOI invocation, factor matrix rows need to be communicated among processors for the next invocation
  - Communication volume

Prior Schemes for Tensor Distribution

- Coarse - Coarse grained schemes [KU'16]
  - Allocate entire "slices" to processors
- Medium - Medium grained scheme [SK '16]
  - Grid based partitioning – similar to block partitioning of matrices
- Fine - Fine grained scheme [KU'16]
  - Allocate individual elements using hypergraph partitioning methods

	TTM	SVD	FMT	Dist. Time
Coarse	Inefficient	Efficient	Inefficient	Fast
Medium	Efficient	Inefficient	Efficient	Fast
Fine	Efficient	Inefficient	Efficient	Slow

Our Scheme – Lite Distribution

- Lite
  - Near optimal on TTM and SVD (both computation and communication)
  - Lightweight (i.e., fast distribution time)
  - Not optimal on FMT (but this is cheap)
  - Performance gain up to 3×

Example – Sequential Sparse TTM for Mode 1

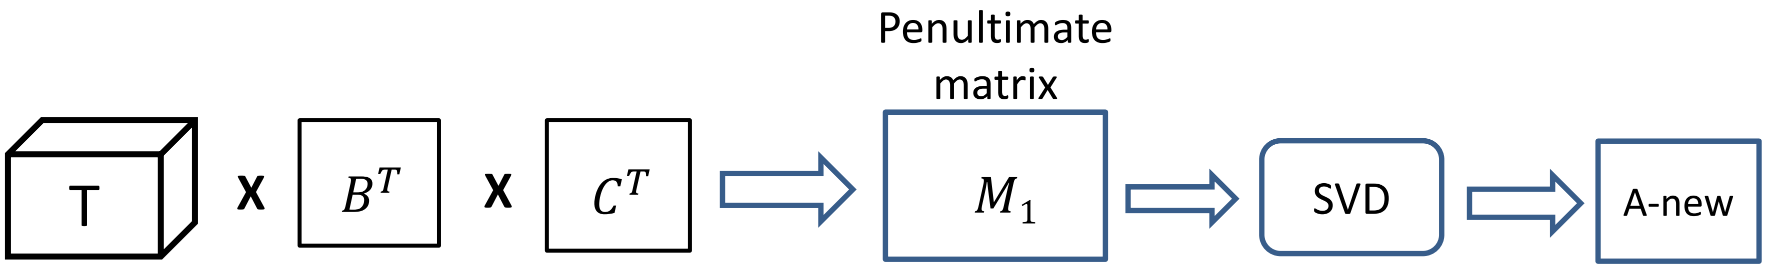

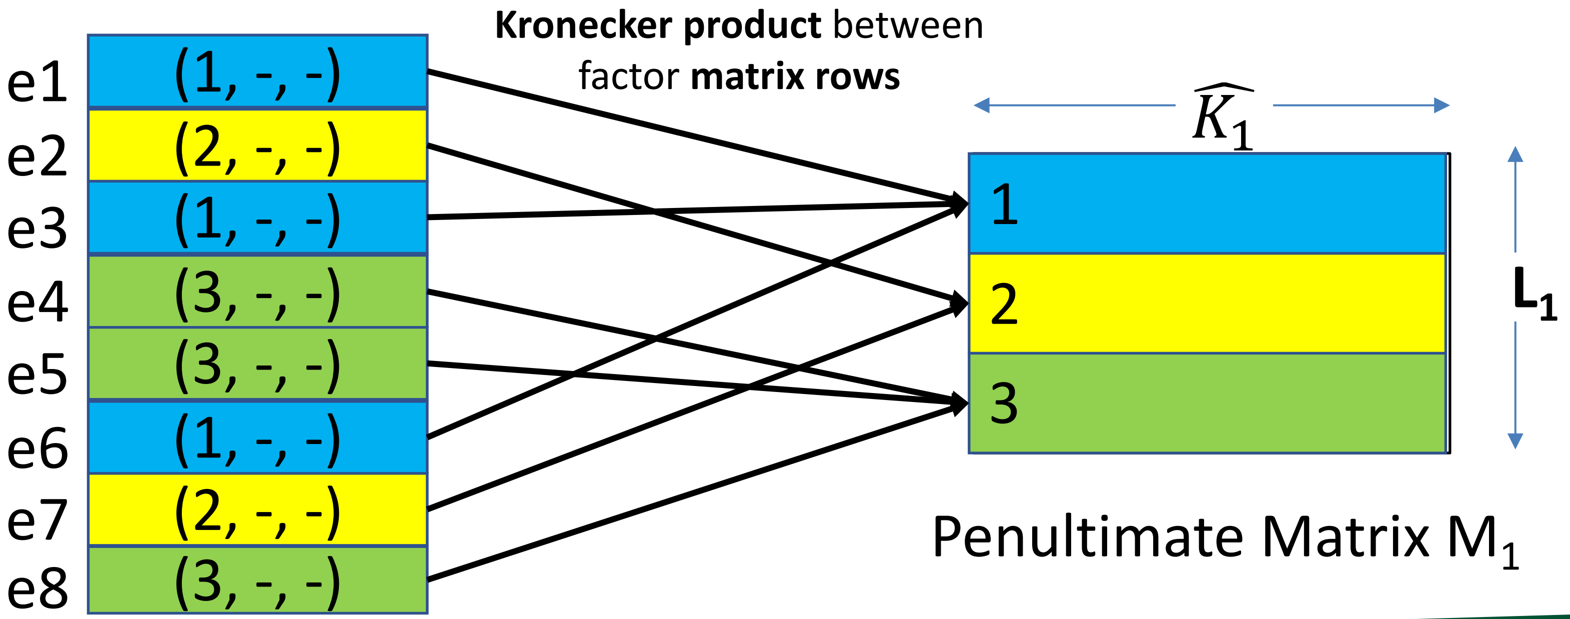

Example – Distributed Sparse TTM for Mode 1

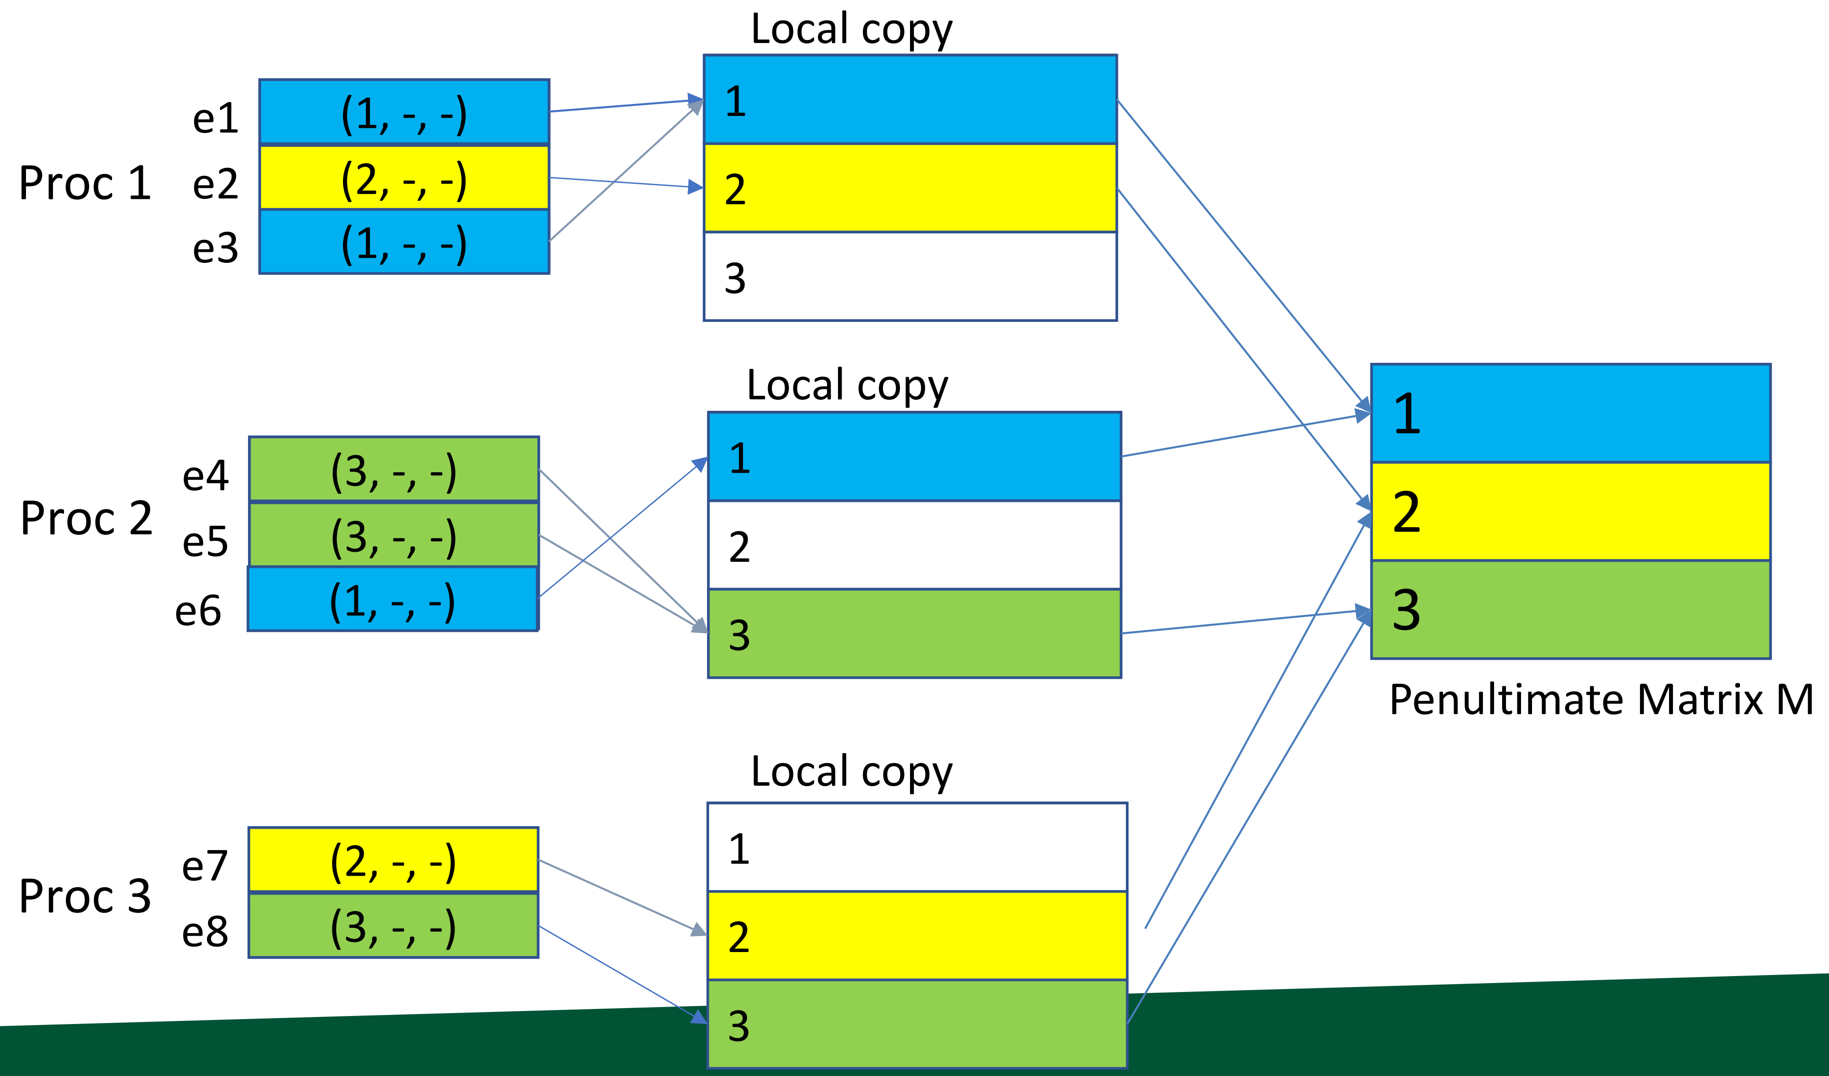

Example – SVD via the Lanczos Method

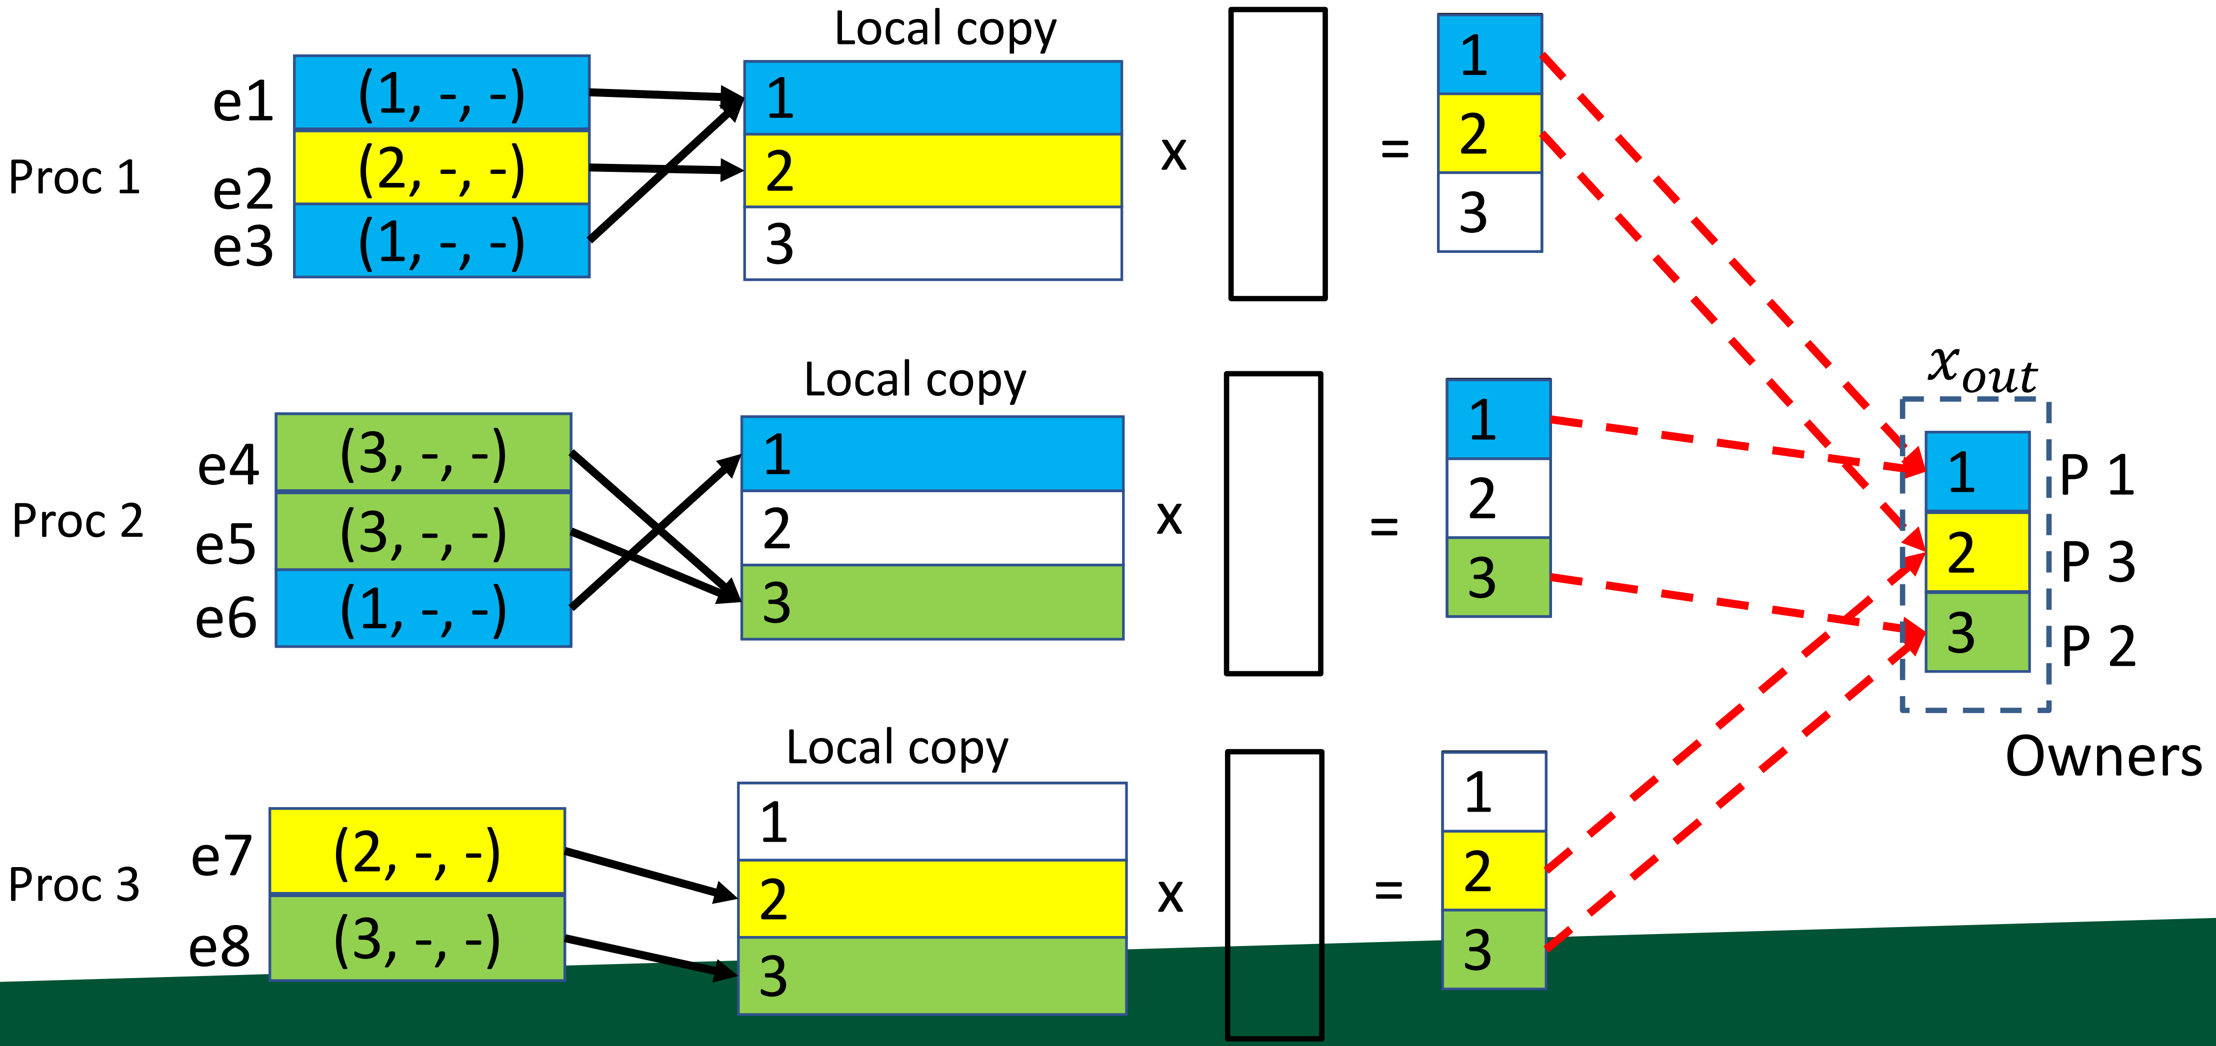

Performance Metrics Along Each Mode

TTM
- TTM-LImb (Load Imbalance)
  - Max number of elements assigned to the processors
  - Optimal value – E / P

SVD
- SVD-Redundancy
  - Total number of times slices are "shared"
  - Measures computational load & comm. volume
  - Optimal value = L (length along the mode, no sharing)
- SVD-LImb:
  - Max number of slices shared by the processors
  - Optimal value = L / P

Factor Matrix Transfer
- Communication volume at each iteration

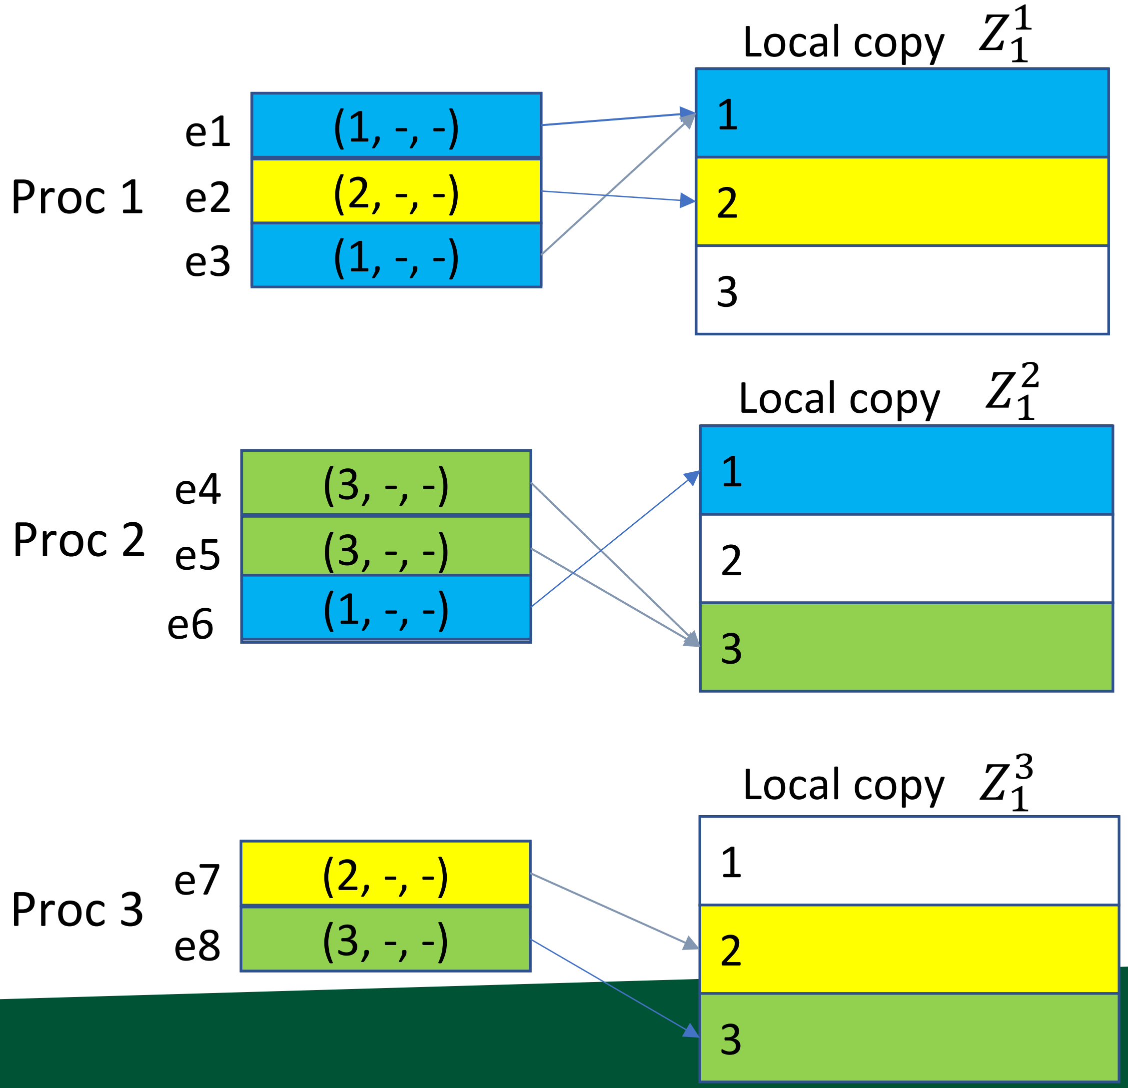

Lite Distribution Scheme

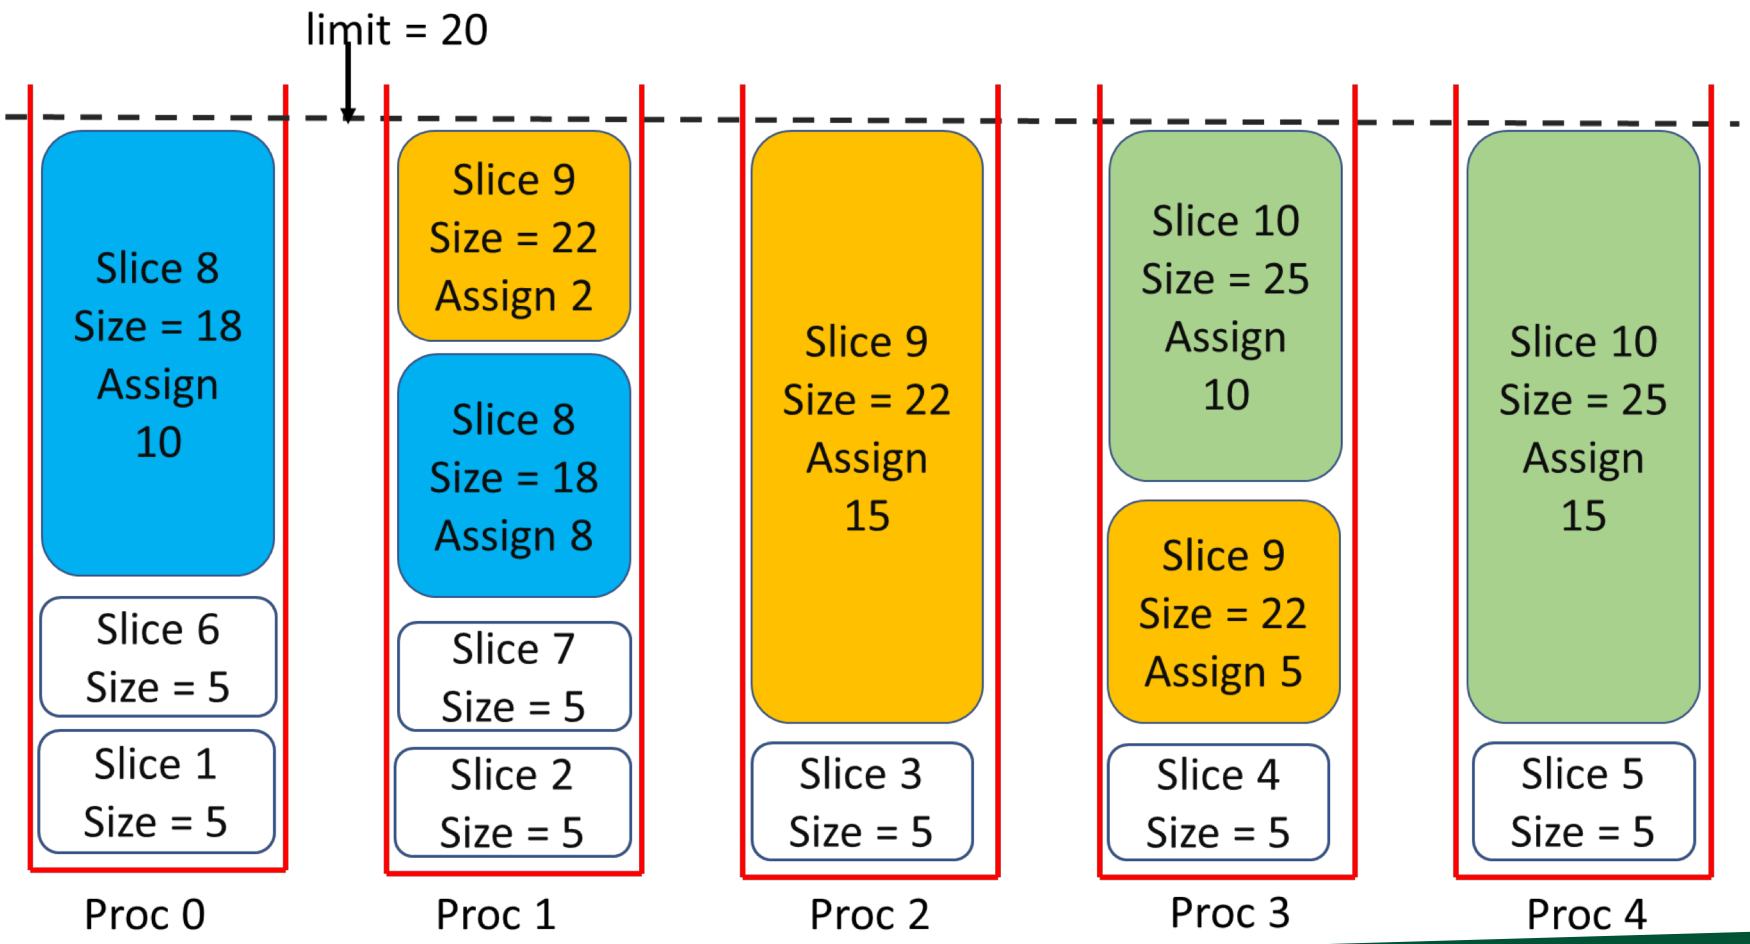

How Does Our Scheme Fare?

- TTM-LImb <= E / P (optimal)
- SVD-Redundancy <= L + P (optimal = L)
- SVD-LImb <= L/P + 2 (optimal = L/P)

- Achieve near optimal
  - TTM computational load
  - SVD computational load, load balance
  - SVD communication volume
- Only issue is high factor matrix transfer volume
  - Computation dominates

Prior Schemes

- **Uni-policy Schemes**
  - A single policy for computation along all modes
  - Only a single copy of the tensor need to be stored
  - Less memory footprint, but less optimization opportunities
- **Multi-policy Schemes**
  - Independent distribution for each mode
  - N copies need to be stored, one along each mode
  - More memory footprint, but more optimization opportunities

How Does Our Scheme Fare?

	Coarse	Medium	Fine	Lite
Policy Type	Multi-policy	Uni-policy	Uni-policy	Multi-policy
Distribution time	Greedy	Greedy	Complex	Greedy
TTM-Limb	High	Reasonable	Reasonable	Optimal
SVD-Redundancy	Optimal	High	Reasonable	Optimal
SVD-Limb	Reasonable	Reasonable	High	Optimal
SVD-Volume	Optimal	Reasonable	Reasonable	Optimal
FM-Volume	High	Reasonable	Reasonable	High

Experimental Evaluation

- R92 cluster – 2 to 32 nodes.
- 16 MPI ranks per node, each mapped to a core. (32 - 512 MPI ranks)
- Dataset : FROSTT repository (frostt.io)

Tensor	L_1	L_2	L_3	L_4	nnz
delicious	532K	17.2M	2.4M	1.4K	140M
enron	6K	5K	244K	1K	54M
flickr	319K	28M	1.6M	731	112M
nell1	2.9M	2.1M	25.4M	-	143M
nell2	12K	9K	28K	-	77M
amazon	4.8M	1.7M	1.8 M	-	1.7B
patents	46	239 K	239	-	3.5B
reddit	8.2M	176K	8.1M	-	4.6B

Execution time

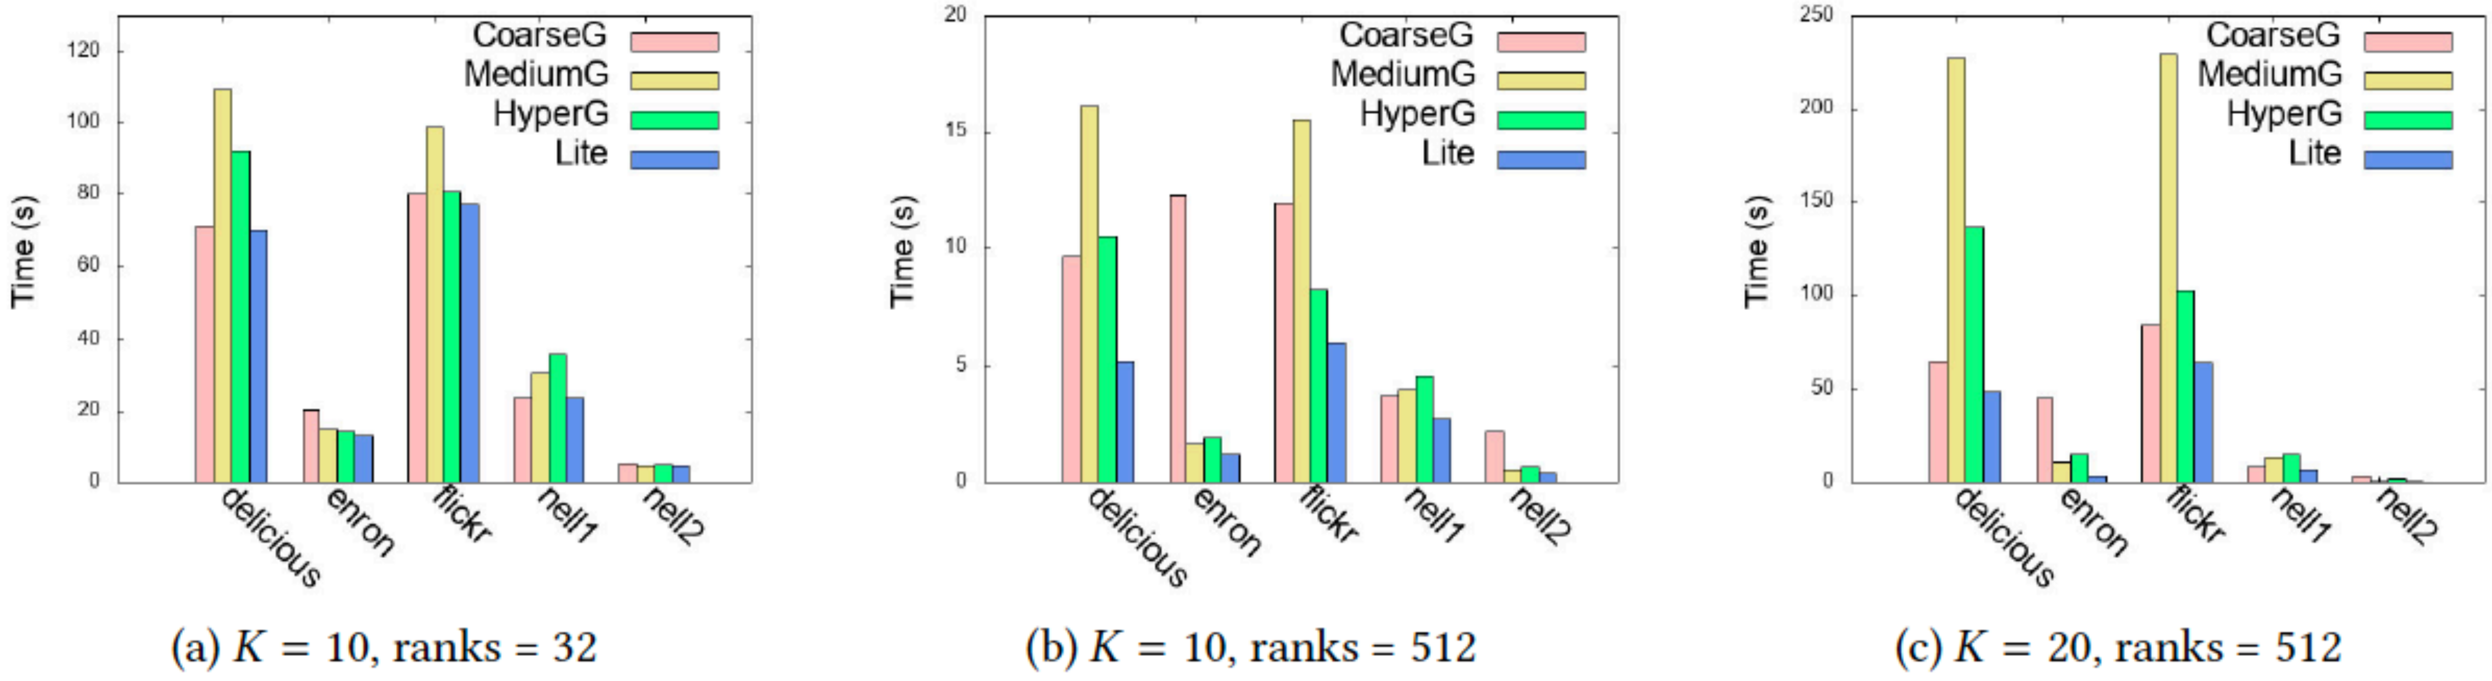

Speedup

Coarse – **12×** Medium – **4.5×** Hyper – **4. 1×** Best Prior – **3×**

Breakdown – Flickr @ rank = 512 & K = 10

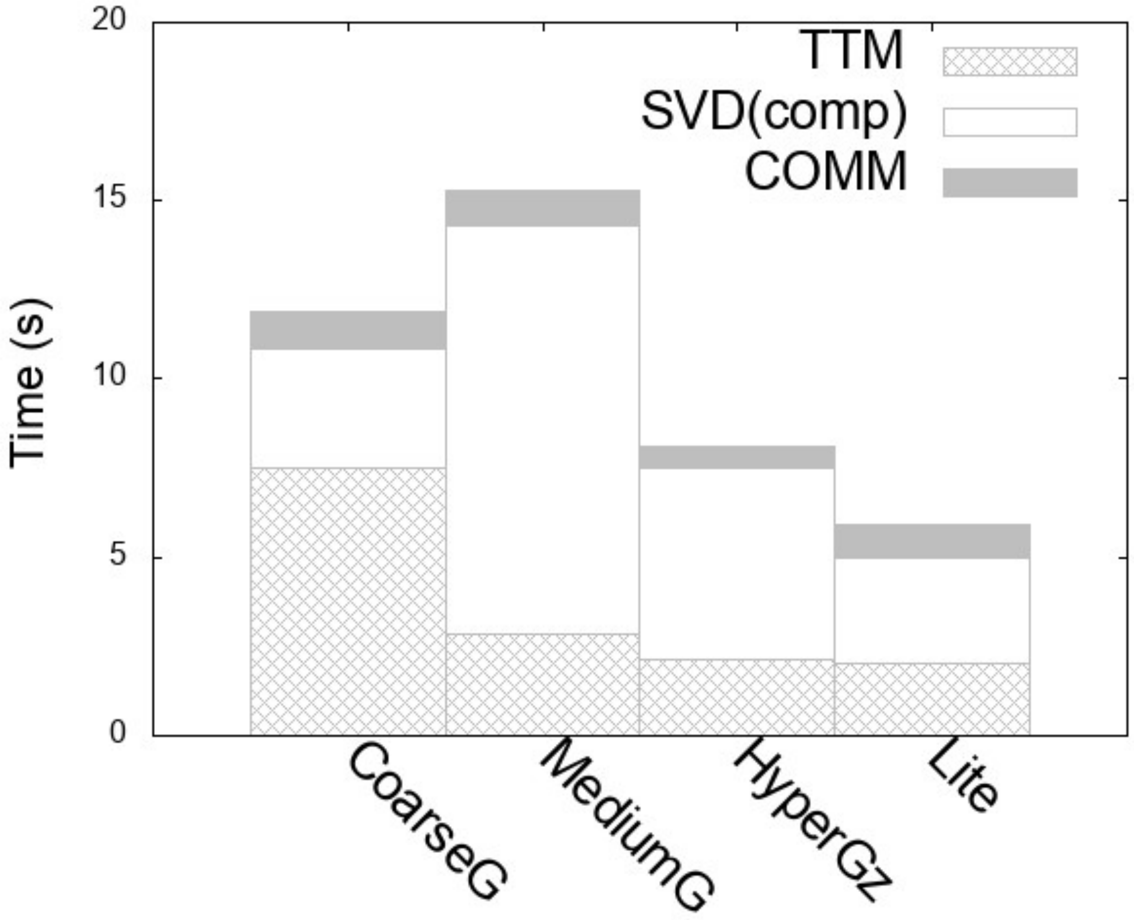

Comparison of the Performance Metrics

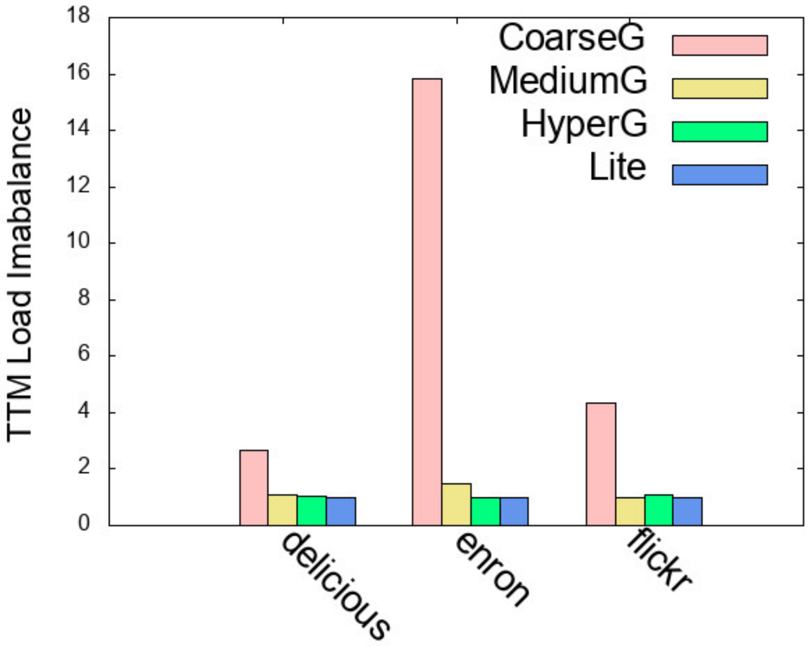

TTM Load Imbalance
(TTM-Limb)

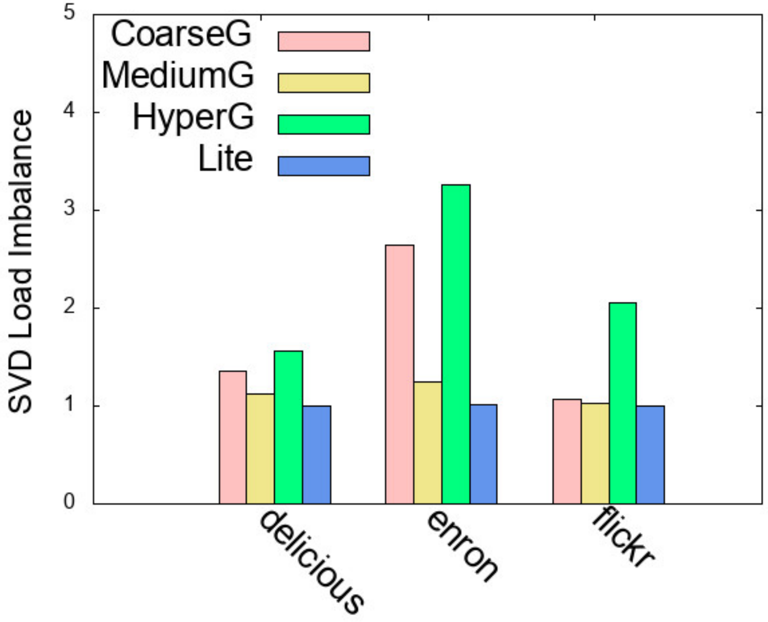

SVD Load Imbalance
(SVD-Limb)

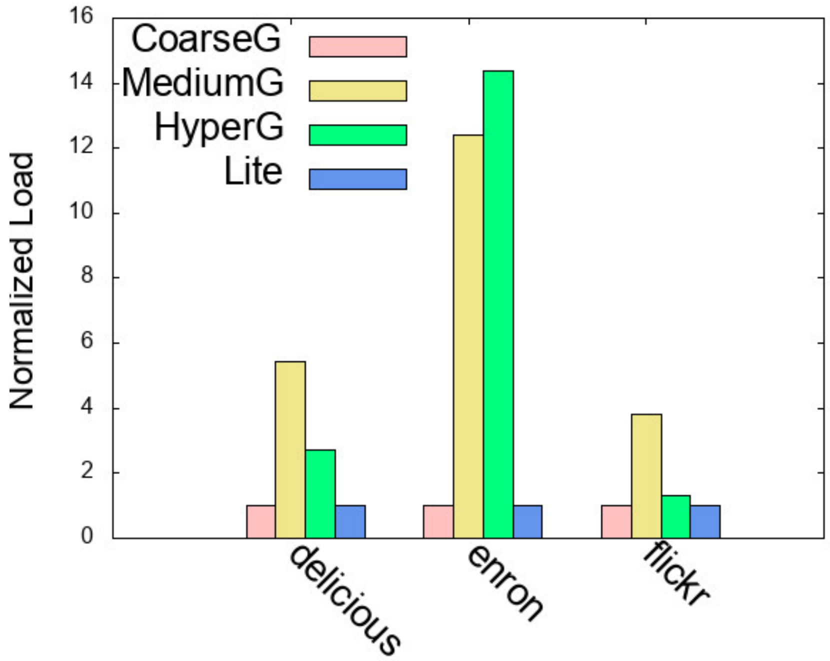

SVD Load
(SVD-Redundancy)

Strong Scaling Results (32 – 512 ranks)

Speedup	CG	MG	HG	Lite
delicious	7.4	6.8	8.8	13.4
enron	1.7	9.0	7.4	11.1
flickr	6.7	6.4	9.8	12.9
nell1	6.4	7.6	7.9	8.6
nell2	2.4	8.4	7.5	12.2
amazon	1.8	11.0	x	13.5
patents	2.7	14.5	x	15.5
reddit	1.8	14.2	x	14.6

Tensor Distribution Time

Time (s)	CG	MG	HG	Lite	HOOI
delicious	6.8	9.3	345	3.9	5.2
enron	0.1	0.08	125	0.1	1.1
flickr	10.9	14.0	203	5.5	6.0
nell1	10.5	13.9	356	6.2	2.7
nell2	0.07	0.05	91	0.07	0.3
amazon	2.9	5.5	x	2.5	8.7
patents	3.2	0.9	x	2.0	14.2
reddit	7.8	11.6	x	5.7	21.6