A Study Towards Optimal Data Layout for GPU Computing

Eddy Z. Zhang, Han Li and Xipeng Shen

Presenter: Mingzhou Zhou

The College of William and Mary

Problem Description

Irregular Memory References

* A mem. transaction -- read/write a consecutive memory segment at once
* A thread warp -- execute only when all data for all threads in the warp is ready
* Random and complicated patterns
* Example: thread warp size - 4, mem. segment size - 4. **Access A[P[tid]]**

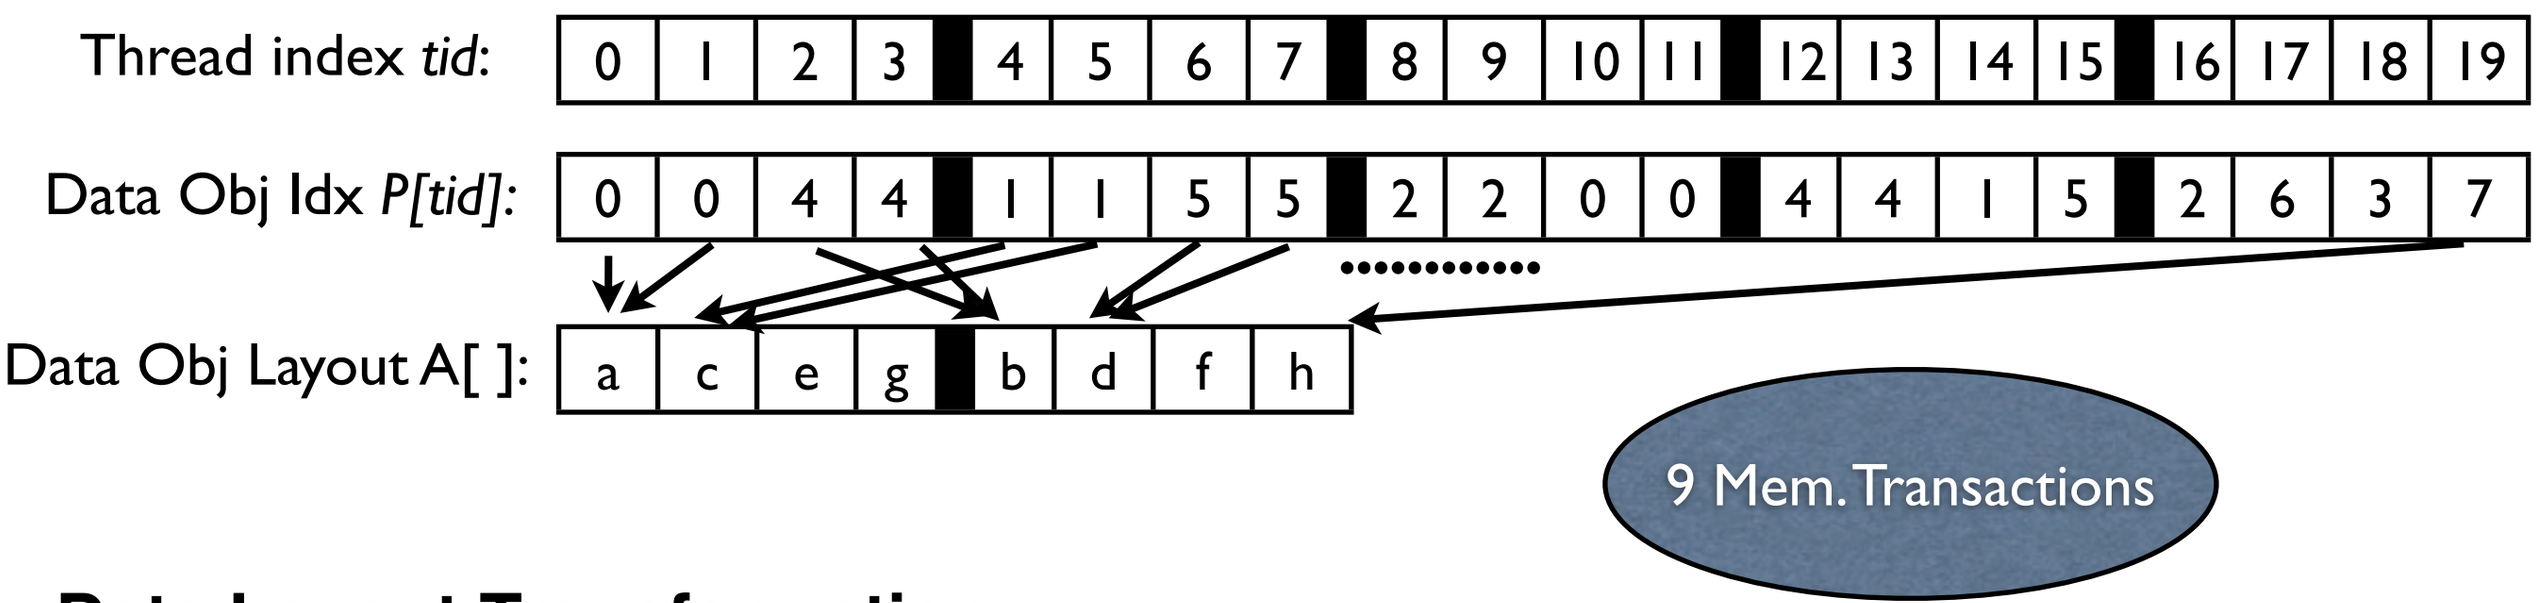

Data Layout Transformation

* Complexity -- NP Completeness if not using any extra memory space or thread relocation

Lack of a study for optimal mapping, previous studies are based on simple heuristics

Duplication Approach

Transform data layout only.

* Duplicate data objects.
* Add space overhead: data size = # threads at a data reference.
* Optimal number of memory transactions
* Adaptive partial duplication [Zhang+:ASOLOS'11]

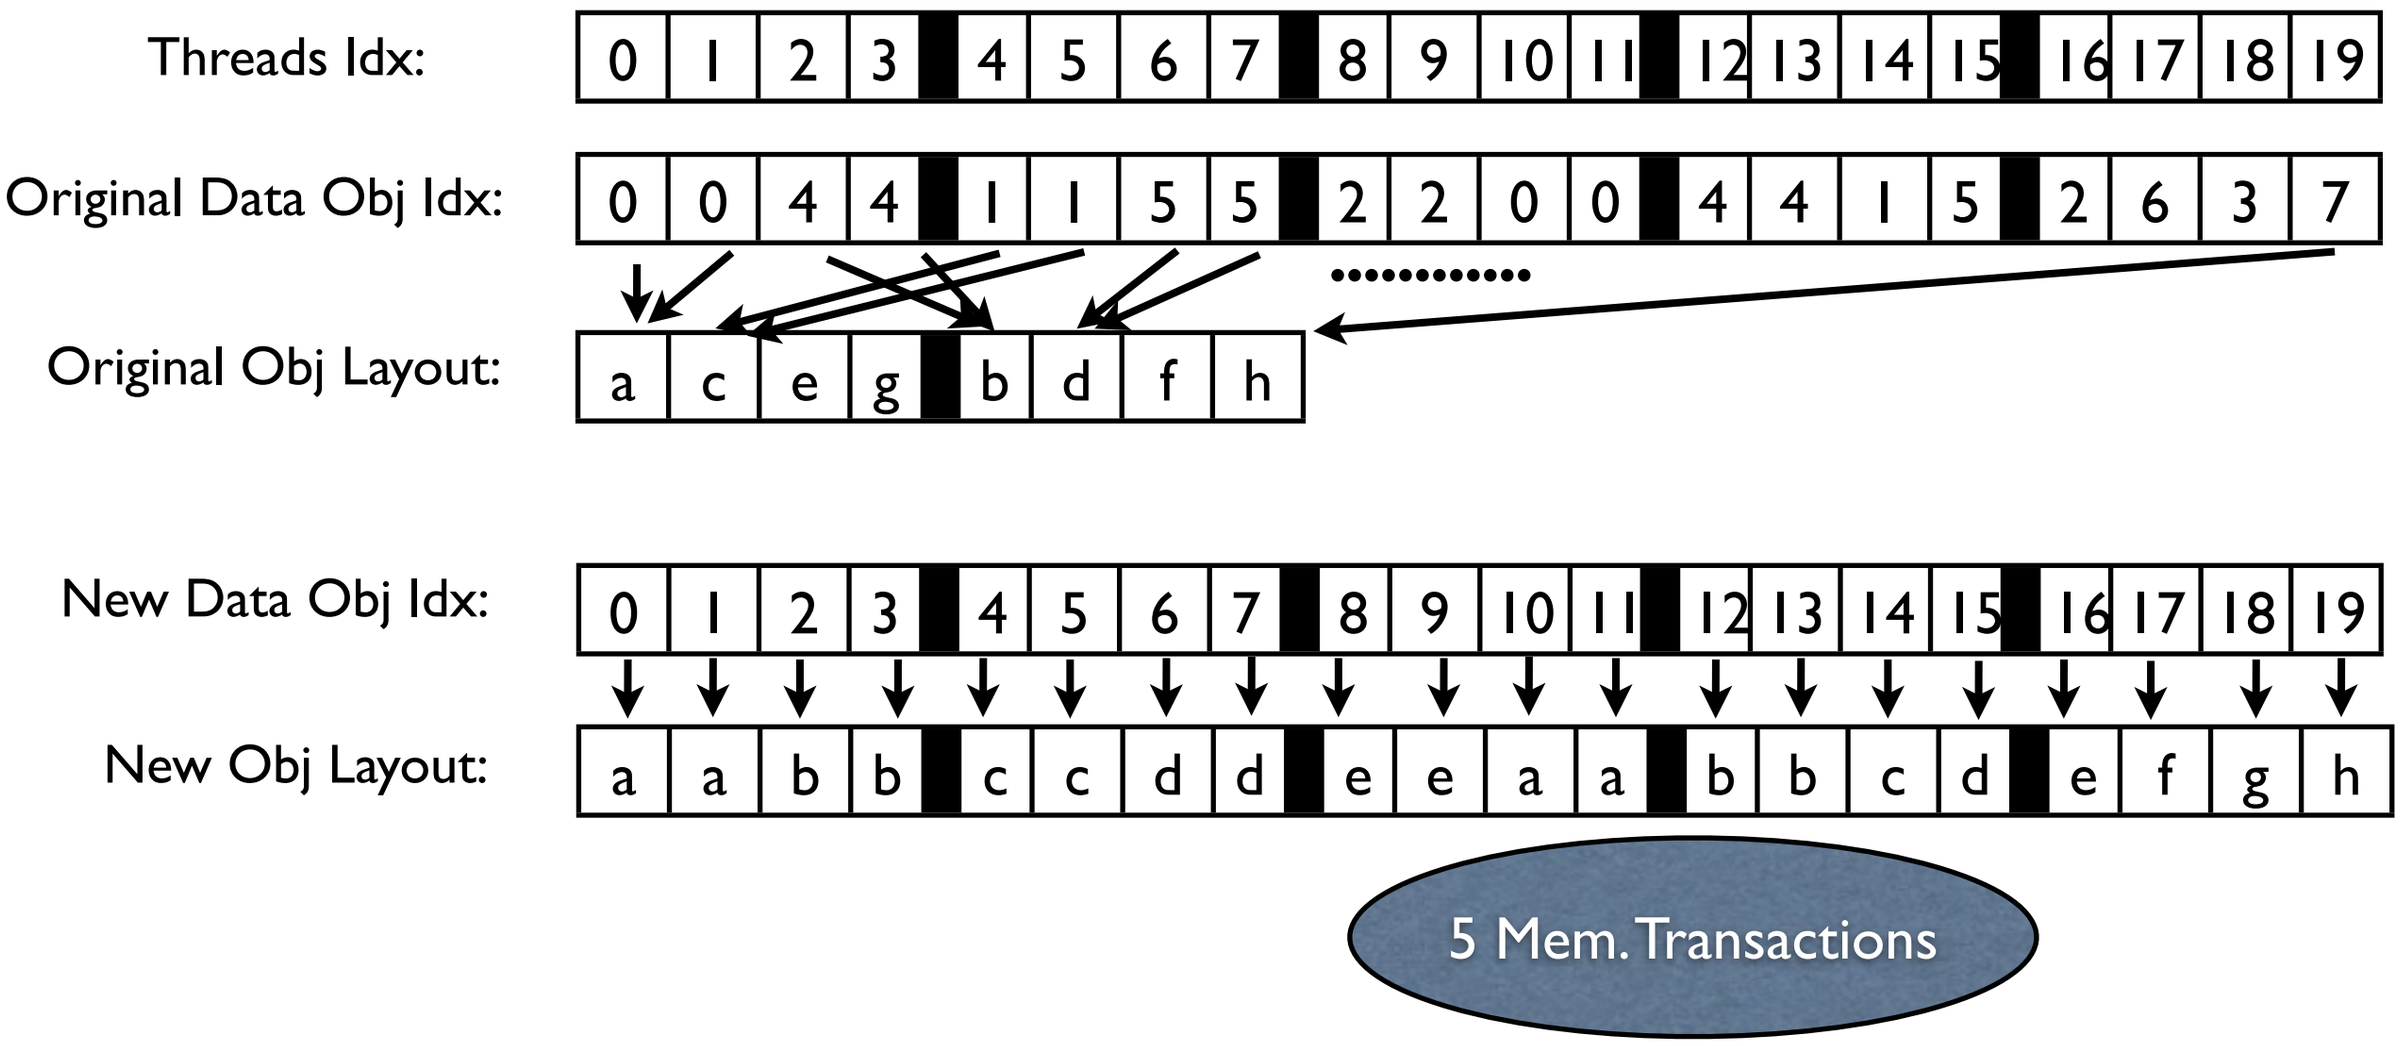

Padding Approach

Reorder Both Threads and Data

* Step 1: Reorder data objects based on access frequencies.
* Step 2: Reorder threads according to their data object order from Step 1.
* Step 3: Put data objects into memory segments. Duplicate or pad dummy objects only when necessary.

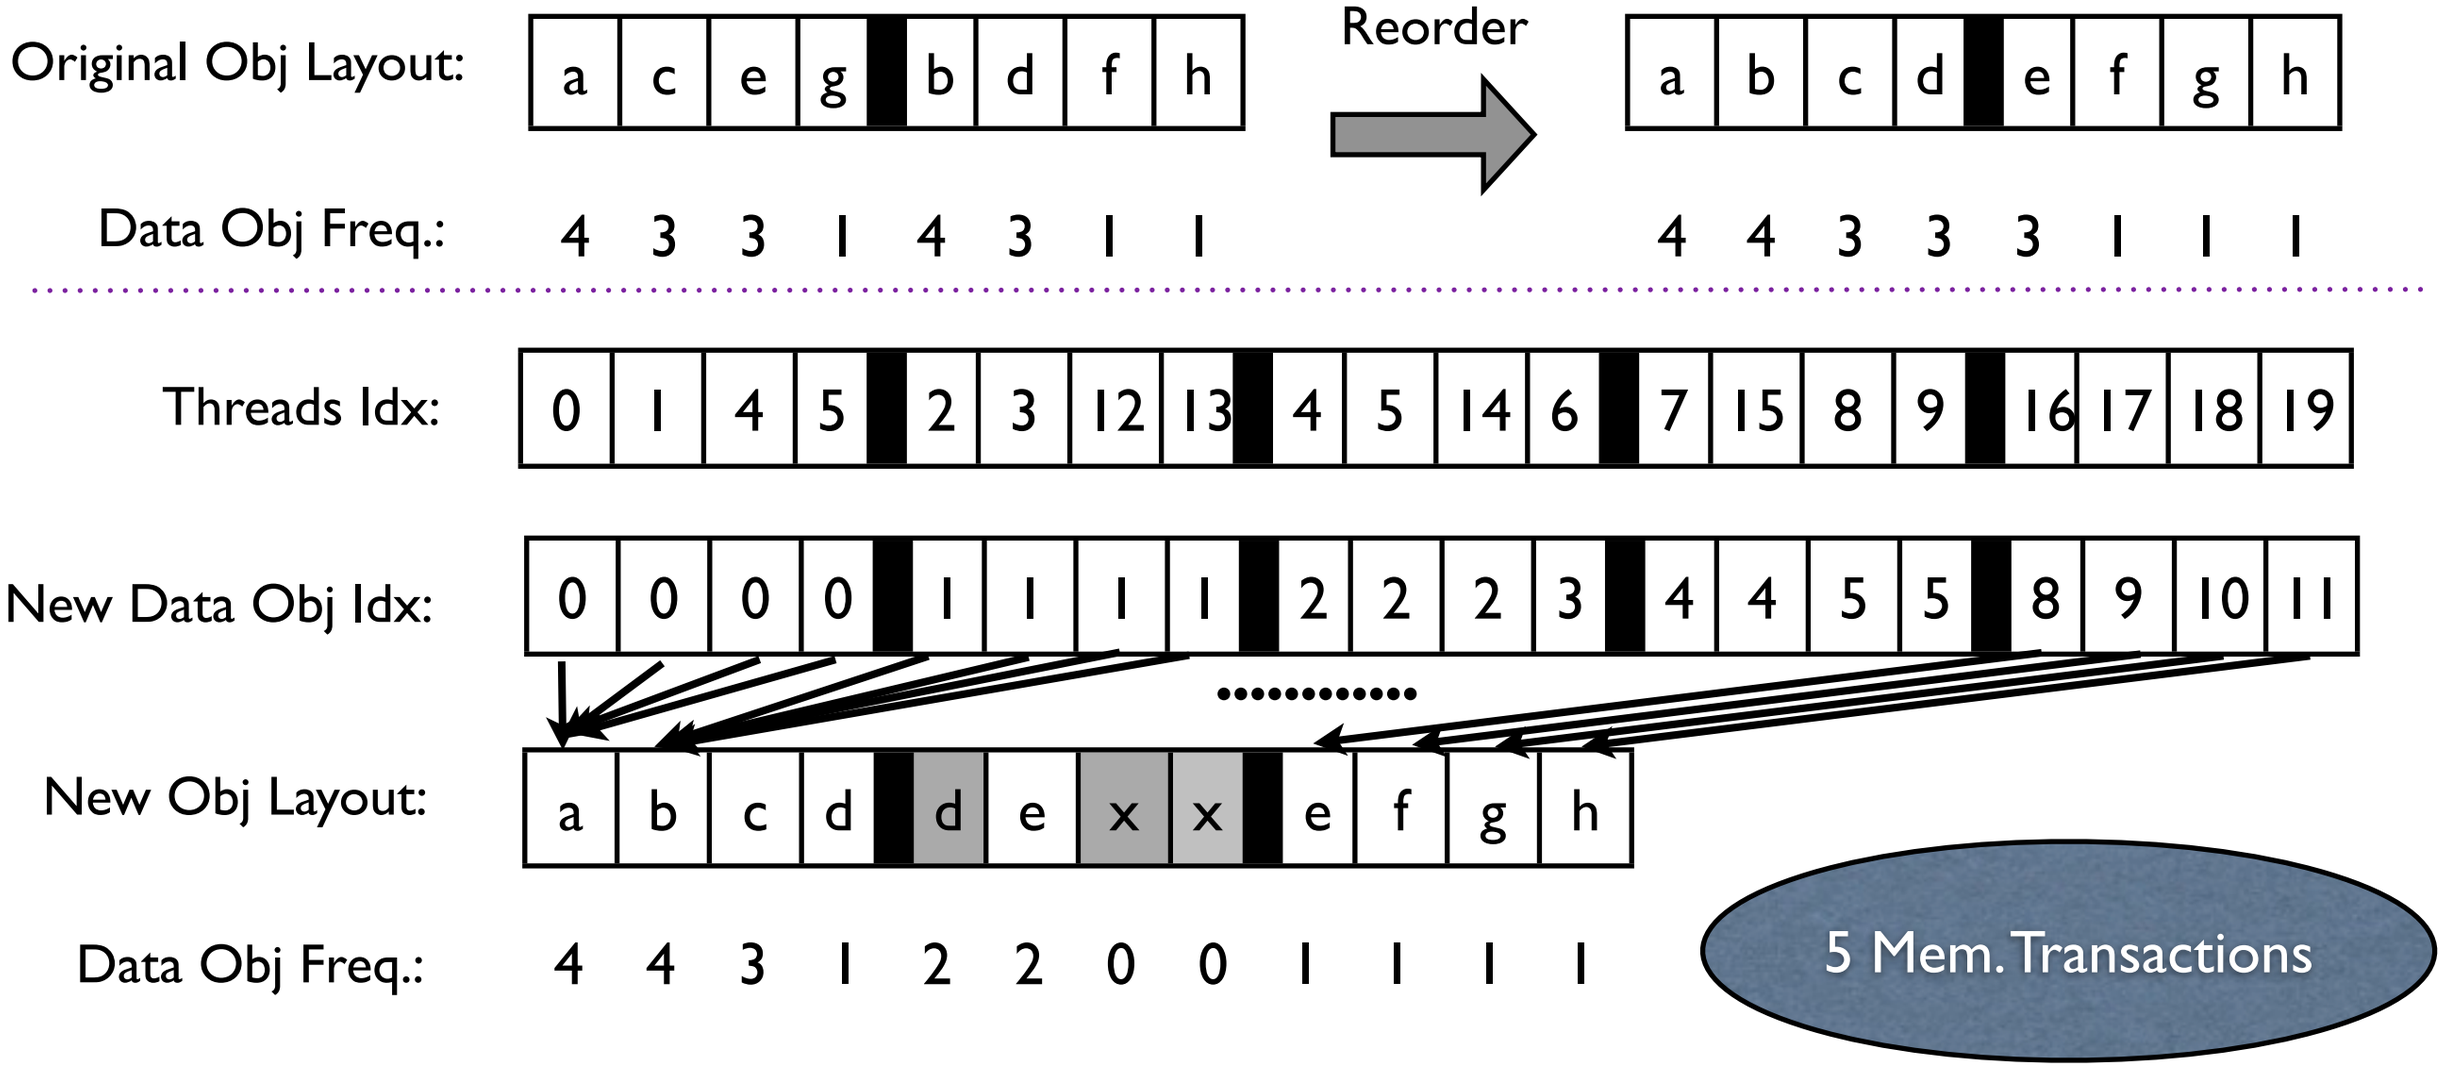

Approximation with Confidence

Reorder Threads More Aggressively

* Step 1: Group threads:
  (a). those accessing a popular object (accesses >= warp size)
  (b). those accessing alone object (1 access only)
  (c). others
* Step 2: Form warps for (a) and (b). Order their objects accordingly.
* Step 3: Form warps for (c) and the remainders of (a) and (b). Order objects.

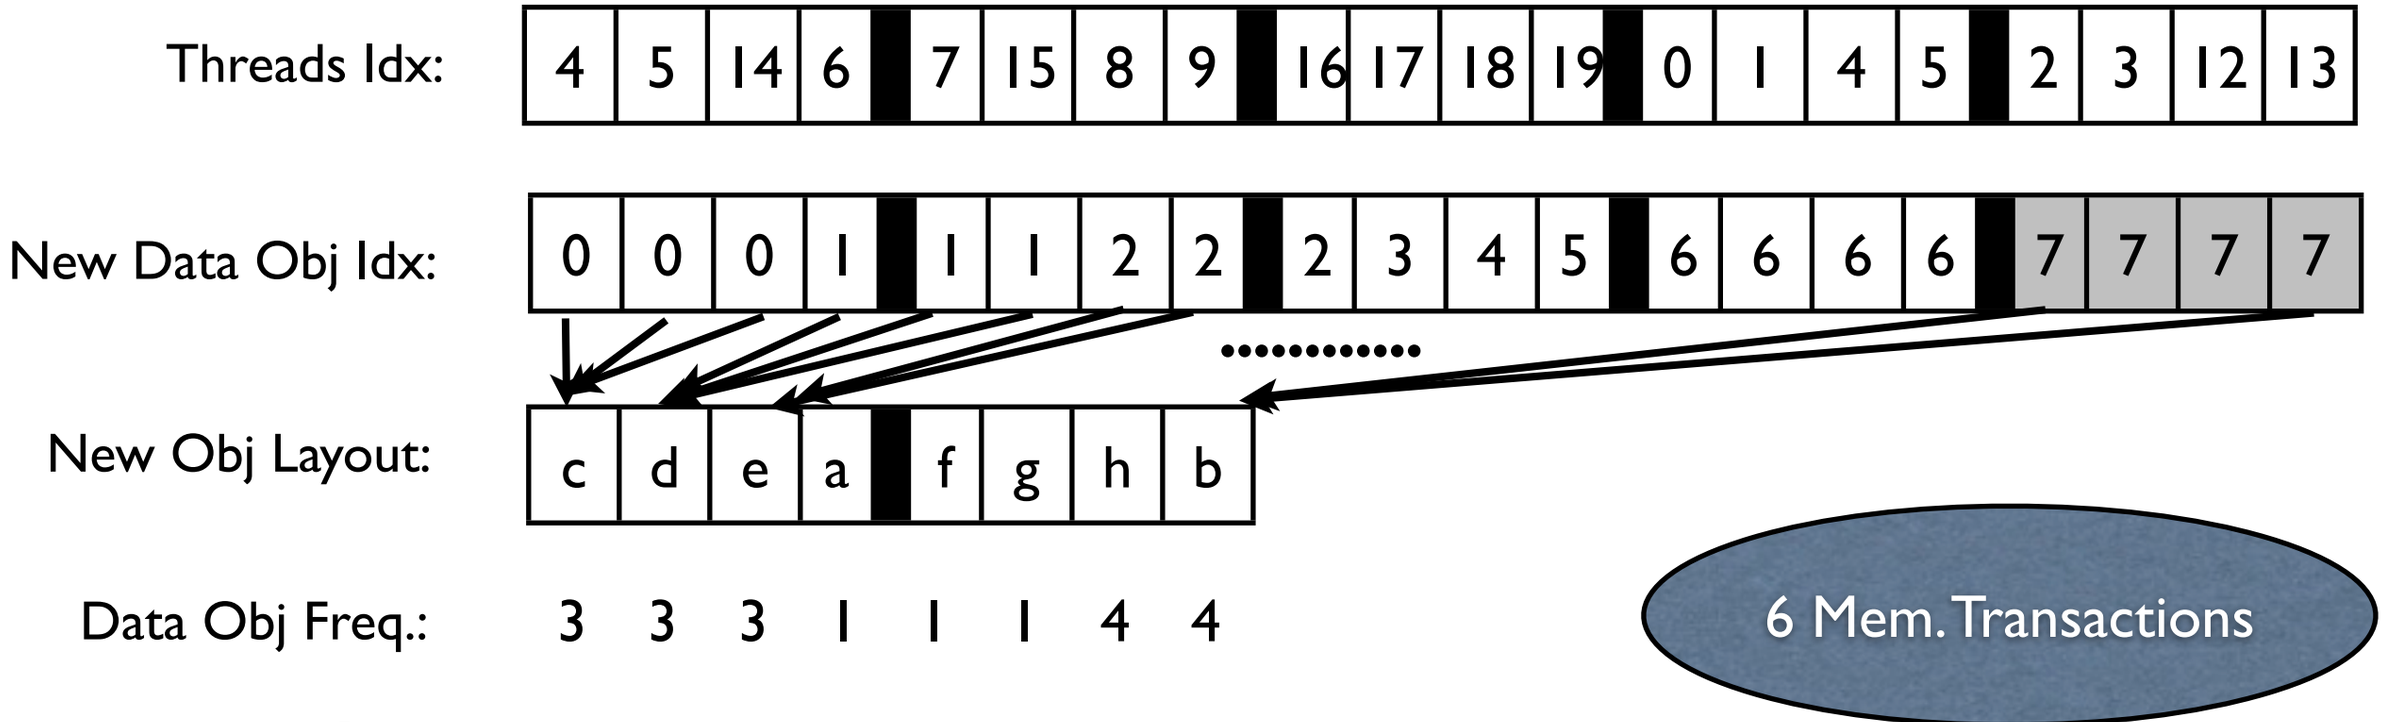

Analytical Bound

* Optimal case: Total number of memory transactions = Total number of warps.
* Upper bound: Optimal + R (Icl + #remainder threads) / W (warp size)

Conclusion

- Data layout transformation is critical for GPU
  - Two algorithms to achieve the optimal
    - duplication & padding (less space)
    - One algorithm to approx. with guarantees
  - A first step to reveal the limit
  - Future
    - Testing and refining them for practical usage

Questions?