

# **Evaluating Coprocessor Effectiveness for the Data Assimilation Research Testbed**

**Ye Feng**

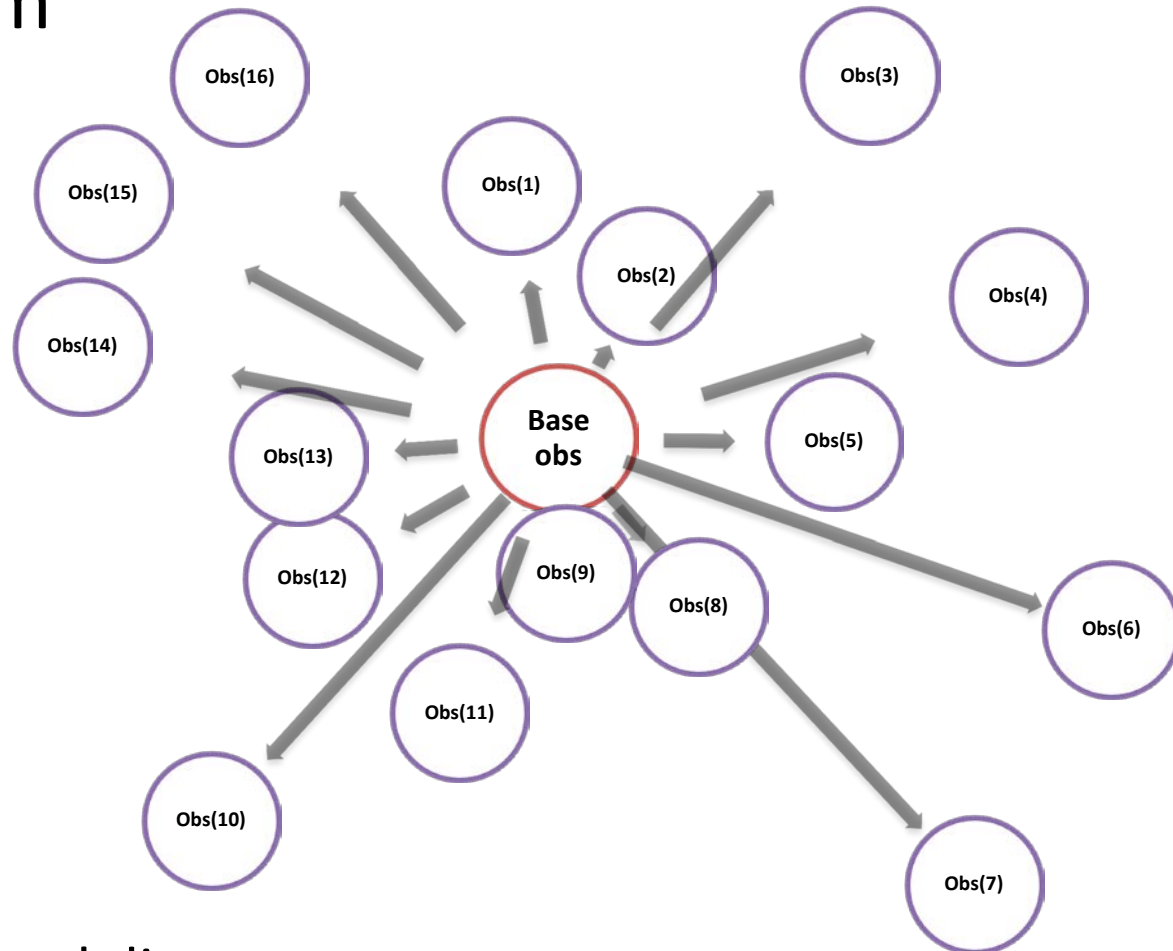
**IMAGe DAREs SIParCS**

**University of Wyoming**

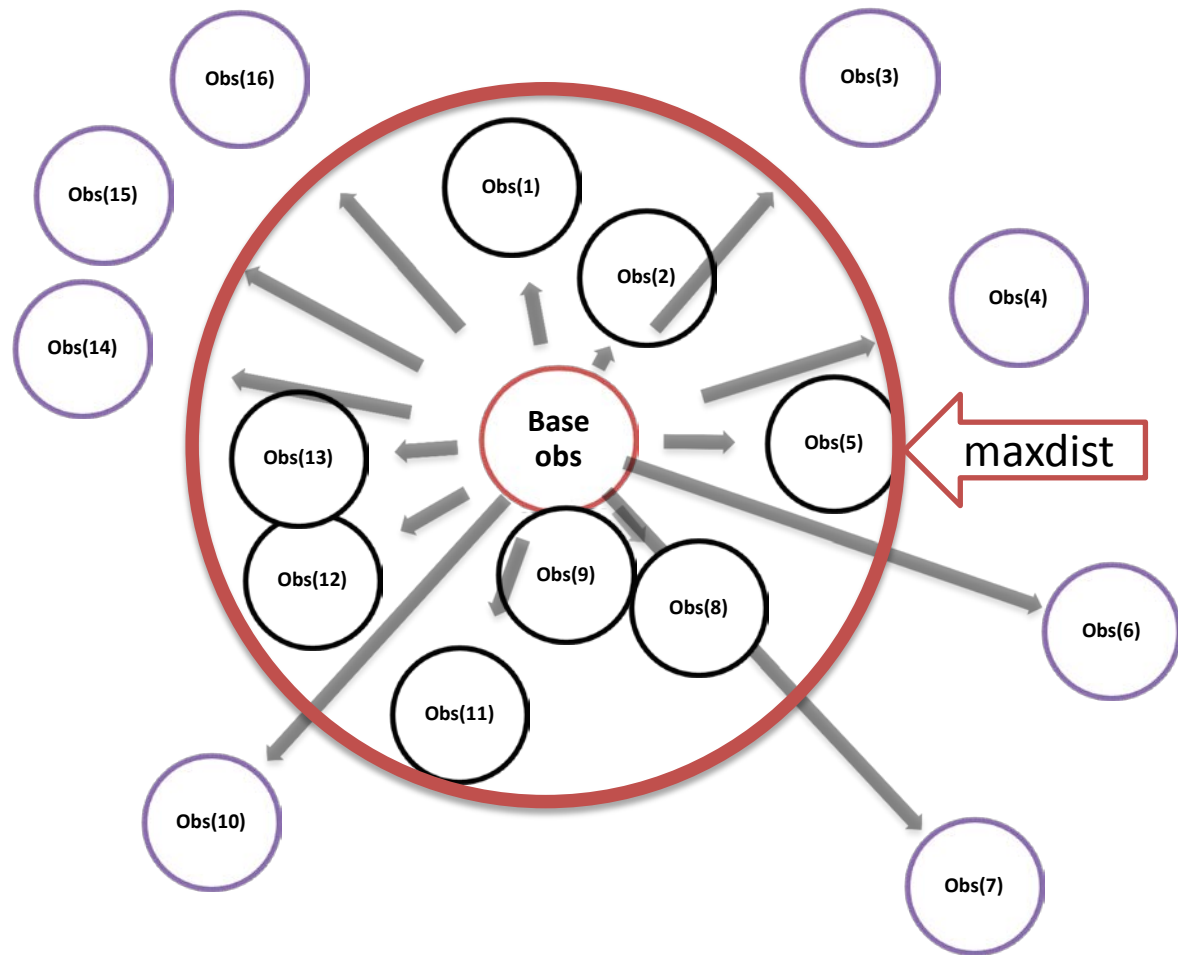
# Introduction

- **Task:** evaluating the feasibility and effectiveness of coprocessor on DART.
- **Target:** `get_close_obs`  
( profiling result: computationally intensive & executed multiple times during a typical DART run.)
- **Coprocessor:** NVIDIA GPUs with CUDA Fortran.
- **Result:** Parallel version of exhaustive search on GPU is faster.

# Problem



Calculate:  
the horizontal distances  
between base location and observation locations.




Find:  
the close observations.

cclose\_ind

1
2
5
8
9
11
12
13

cdist

d1
d2
d5
d8
d9
d11
d12
d13



**EASY!**  
or is it?

# It is easy on CPU

```
!-----  
! For validation, it is useful to be able to compare against exact  
! exhaustive search  
if(COMPARE_TO_CORRECT) then  
  cnum_close = 0  
  do i = 1, gc%num  
    this_dist = get_dist(base_obs_loc, obs(i), base_obs_type, obs_kind(i))  
    if(this_dist <= this_maxdist) then  
      ! Add this obs to correct list  
      { cnum_close = cnum_close + 1  
        cclose_ind(cnum_close) = i  
        cdist(cnum_close) = this_dist }  
    endif  
  end do  
endif
```



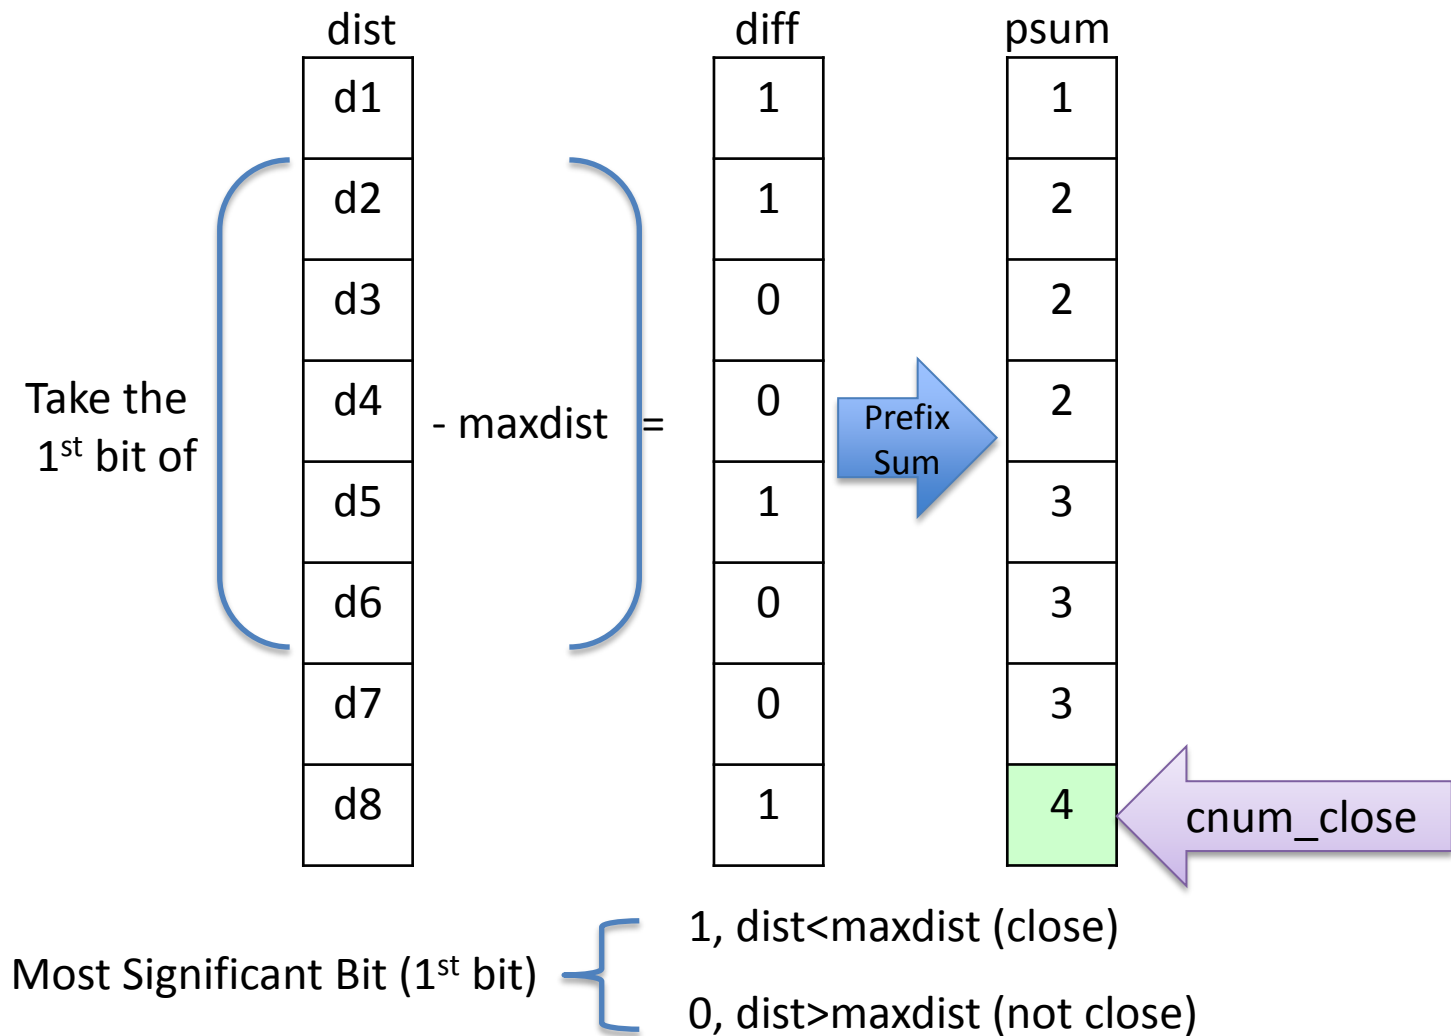
Data  
Dependency

But GPU doesn't work this way!

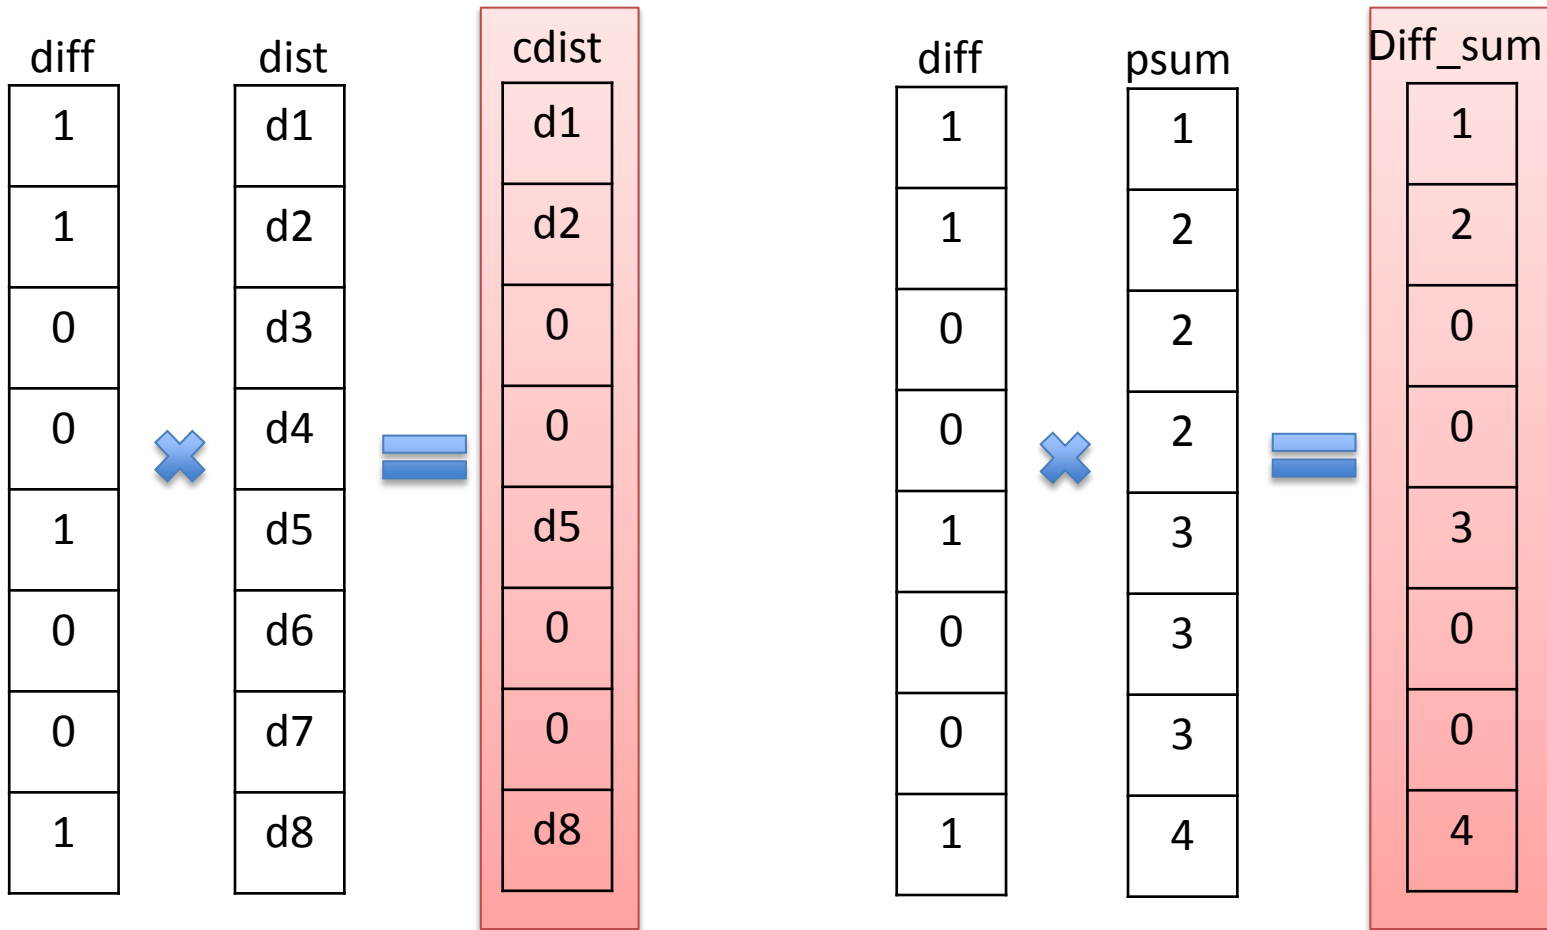
Problems with data dependency usually don't scale so well on GPU.

- *cnum\_close depends on previous cnum\_close value.*
- *cclose\_ind and cdist both depend on cnum\_close.*

# GPU Scan:



# GPU Scan:





# Extract:

Thread ID	Diff_sum	cdist
1	1	d1
2	2	d2
3	0	0
4	0	0
5	3	d5
6	0	0
7	0	0
8	4	d8

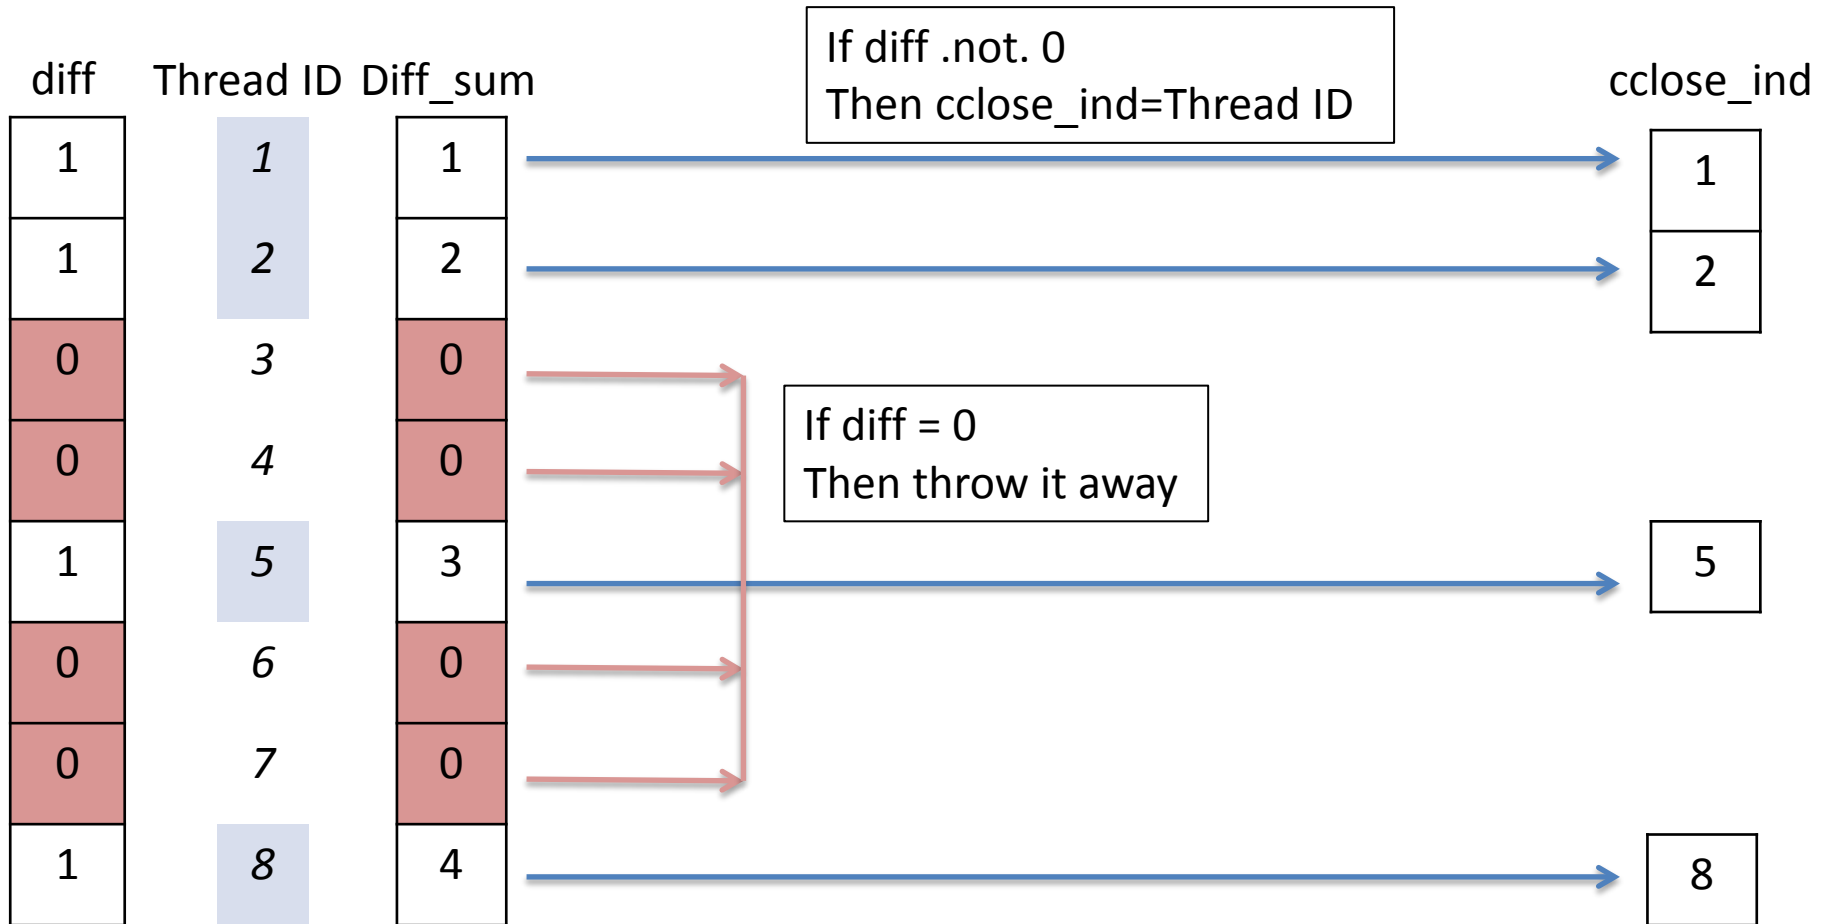



cclose_ind	cdist
1	d1
2	d2
5	d5
8	d8



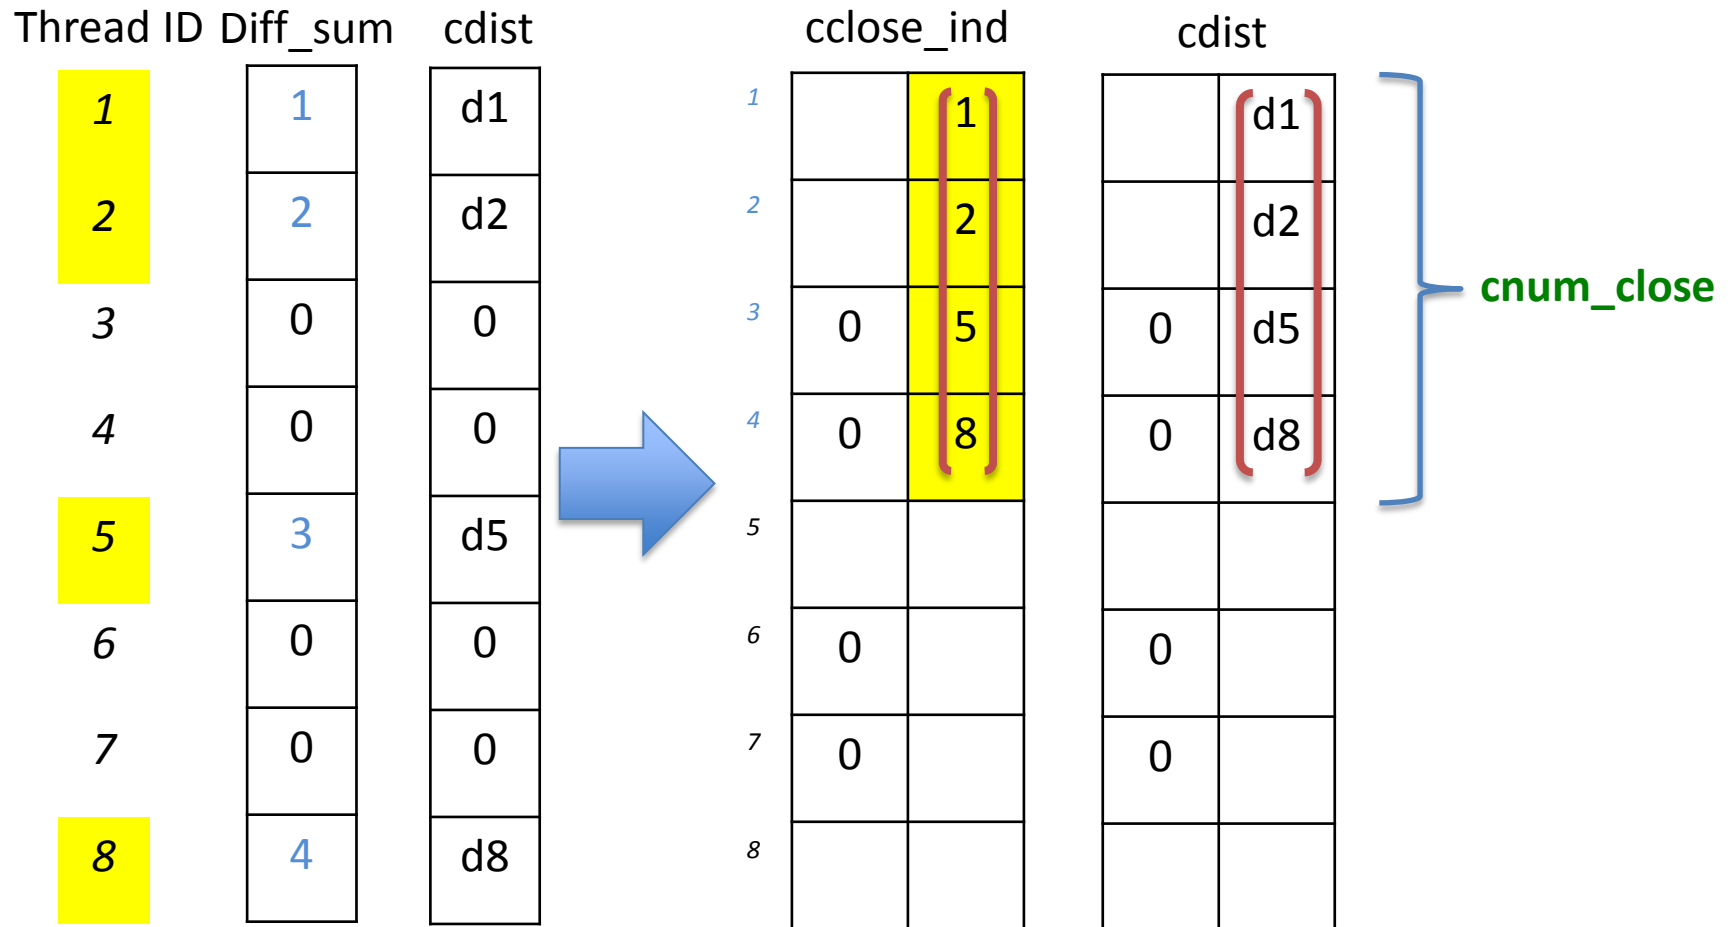
How can we independently eliminate the zeros and extract the indices ?

# Solution?





# Solution!




## Device Functions:

- `gpu_dist`

- `gpu_scan`


si: number of iterations performed in this kernel.




- `extract`

sn: number of `gpu_scan` blocks that each `extract` block in this kernel handles.


Result from  
`gpu_scan`:




# CPU-GPU(si-sn) Execution Time



# GPU(si=1,sn=1) Speedup




# GPU(si-sn) Speedup






# GPU(si-sn) Speedup



# GPU(si-sn) Speedup



# Conclusion

- CUDA Fortran on GPU gave significant speedup vs CPU (10x+).
- Step outside the box (redesign the algorithm).
- In order to get good performance,  $si$  and  $sn$  need to be tuned.
- Be careful with using device memory.
- There's still room to improve the performance of this project.

# Acknowledgements

## **DAReS/IMAGe**

Helen Kershaw (Mentor)

Nancy Collins (Mentor)

Jeff Anderson

Tim Hoar

Kevin Raeder

Richard Loft

Raghu Raj Prasanna Kumar

Kristin Mooney

Silvia Gentile

Carolyn Mueller

**UCAR**

**NCAR**

**University of Wyoming**