

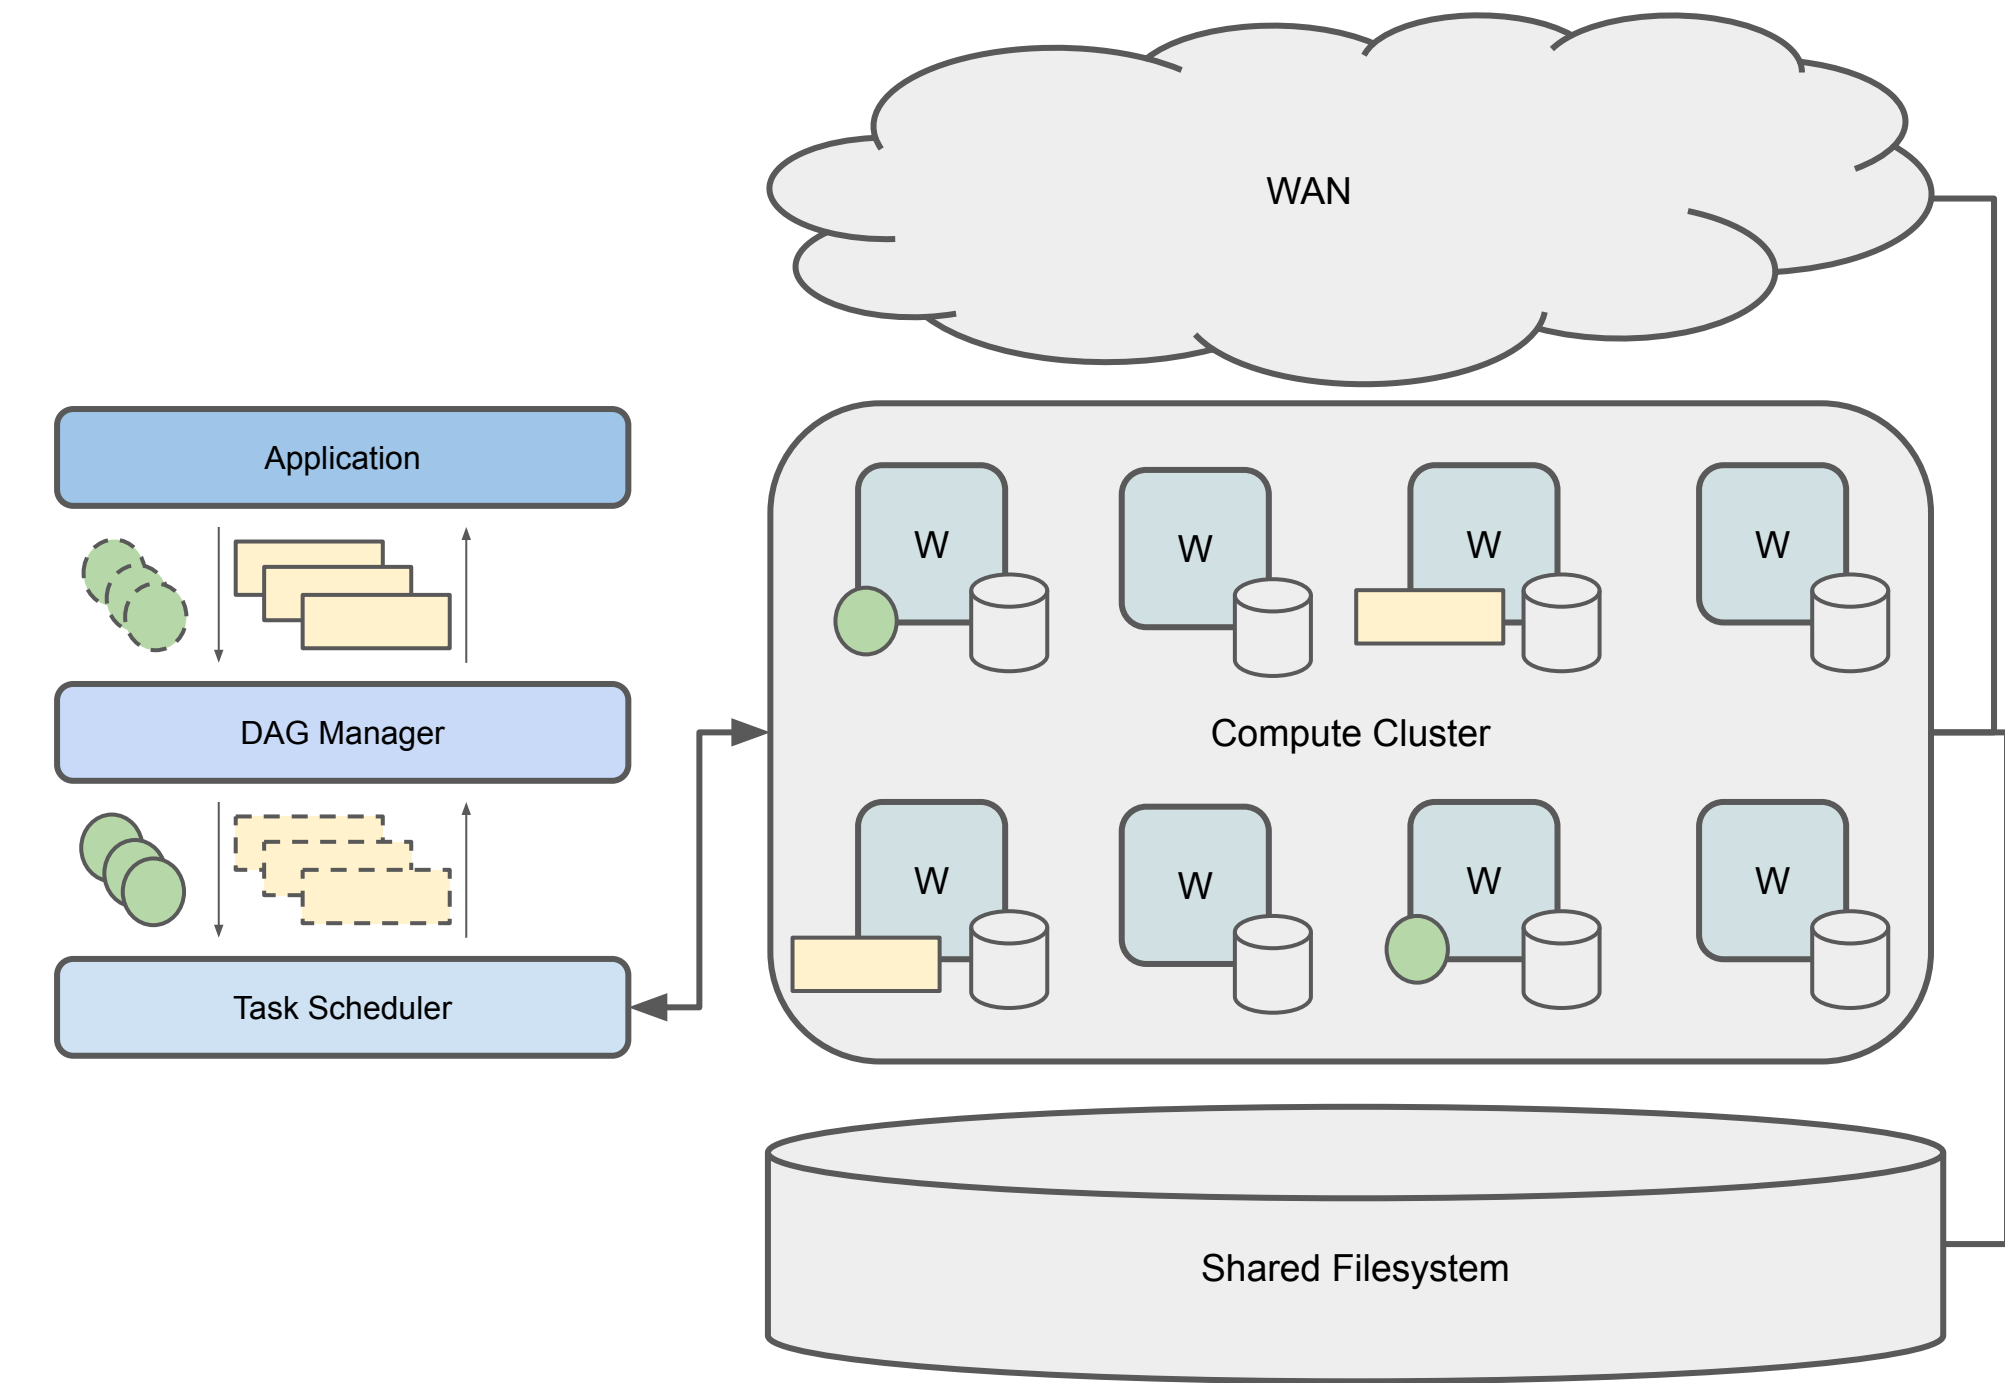
Task Graph Restructuring via Function Based Annotation For Large-Scale Scientific Applications

Barry Sly-Delgado(bslydelg@nd.edu) University of Notre Dame - <http://ccl.cse.nd.edu/>

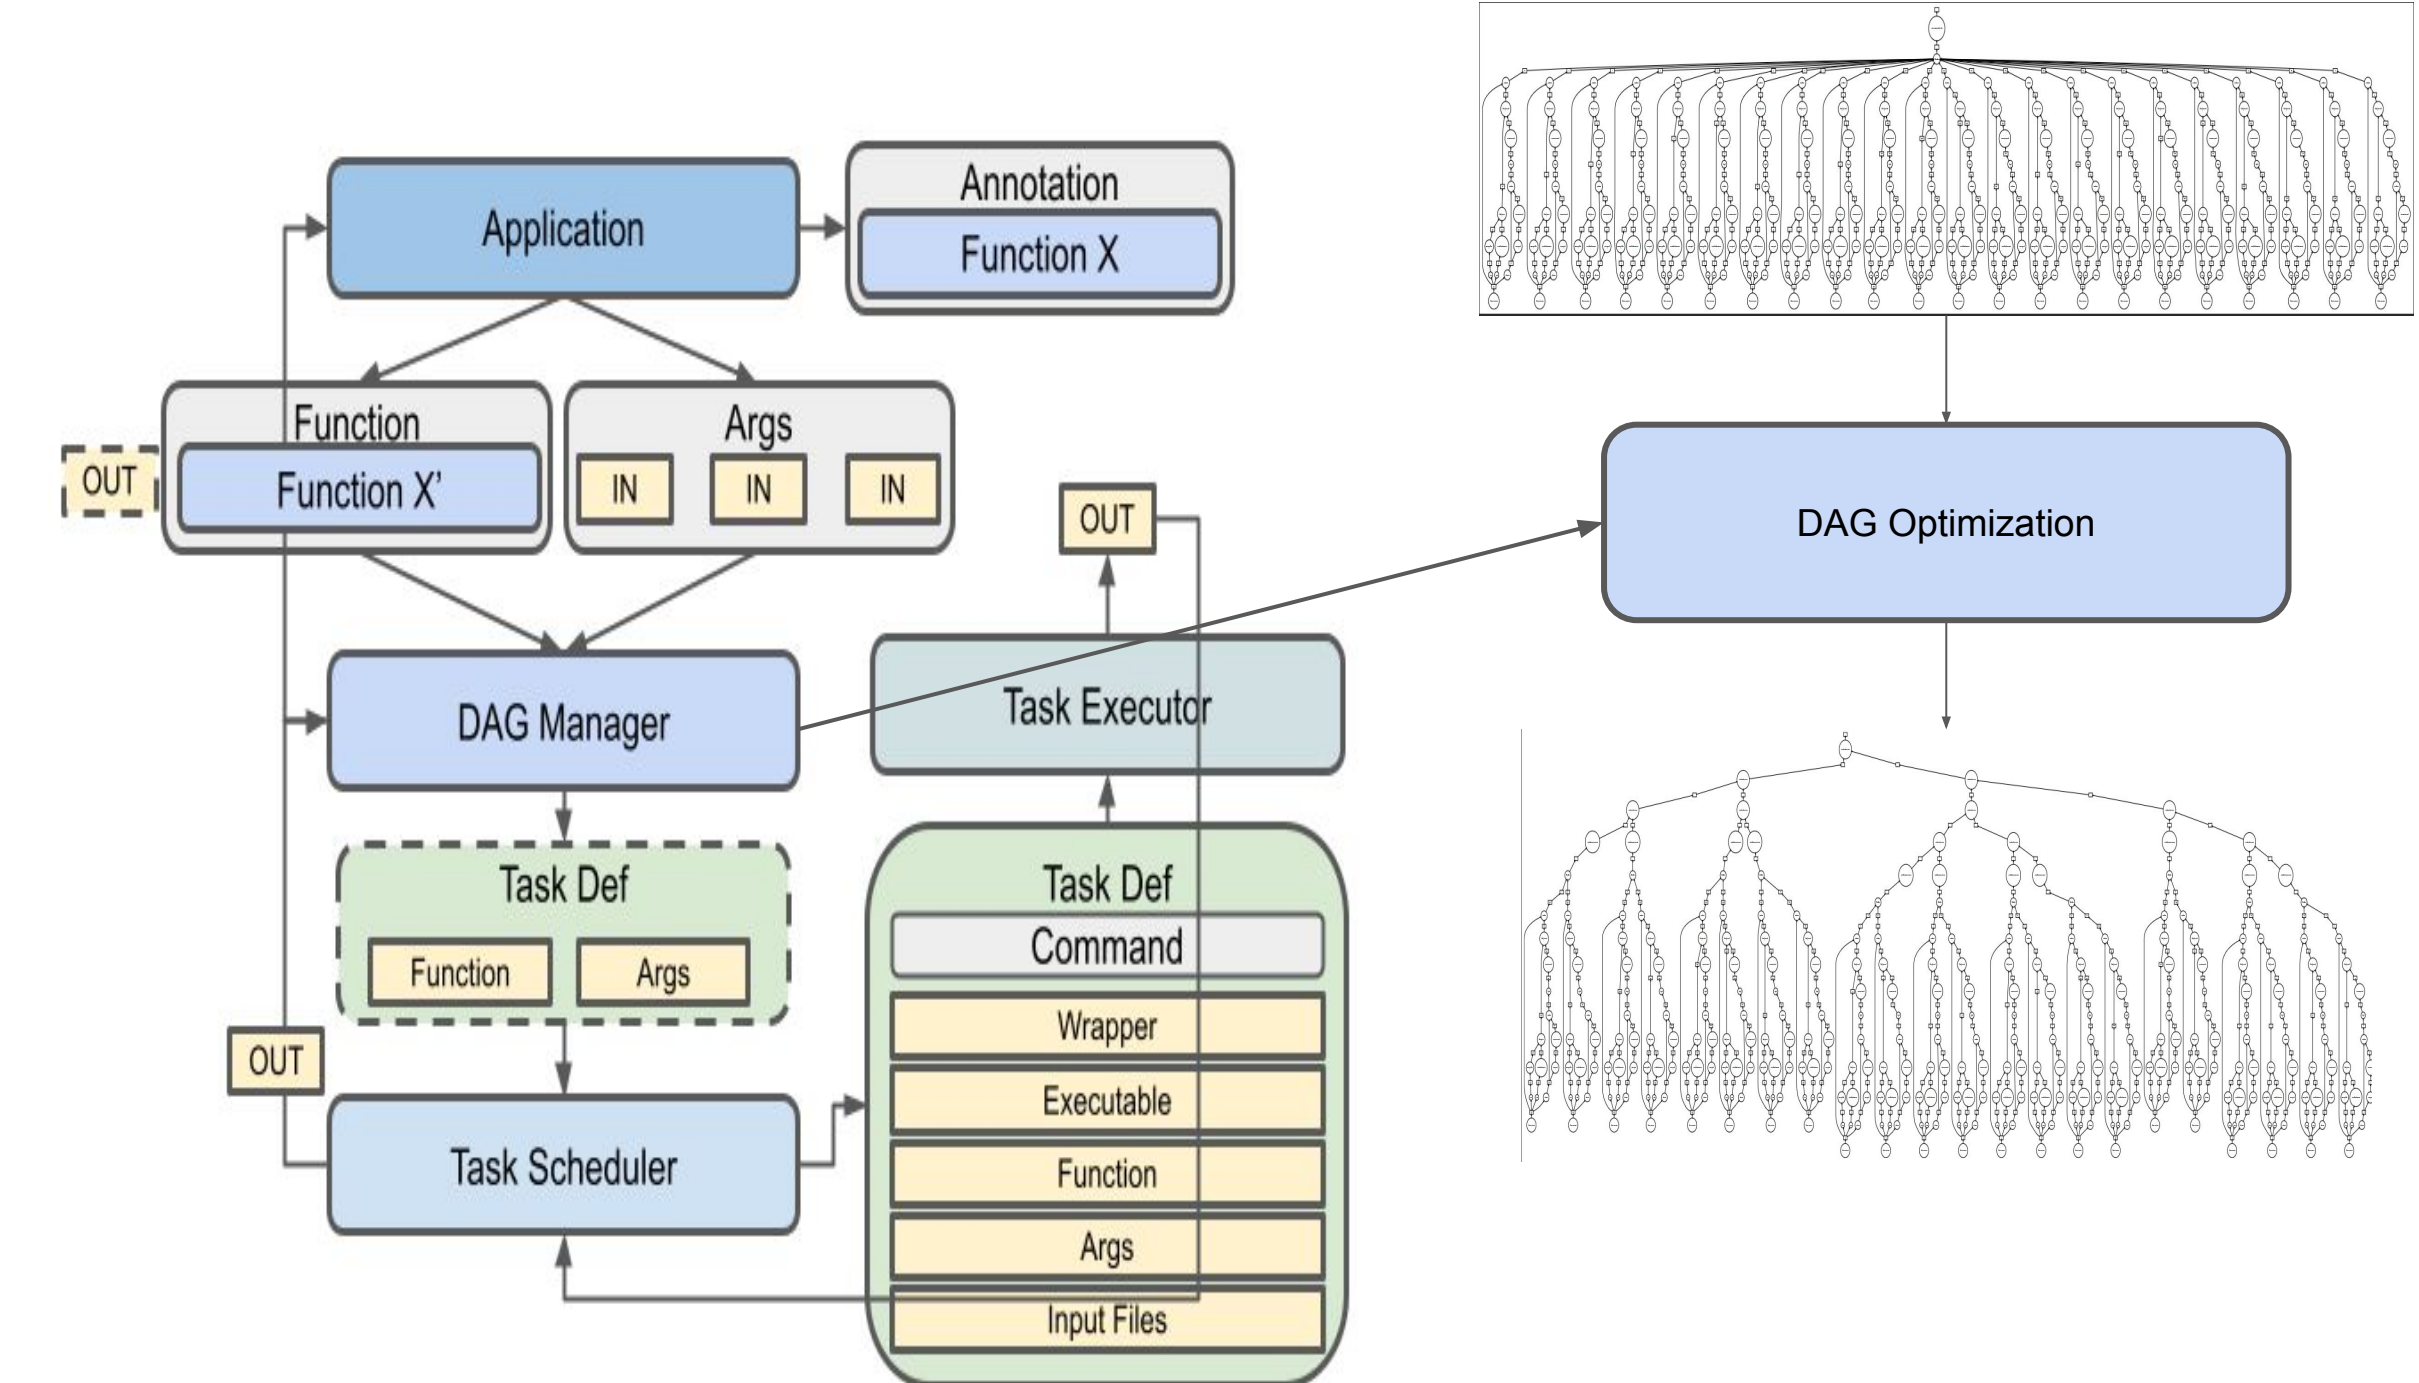
Abstract

Distributed Frameworks allow for users to create large-scale distributed scientific applications within languages such as Python. Within these frameworks users provide functions that represent individual tasks in a broader DAG. Users compose these graphs using the provided API of a given framework. However, the initial composition of the DAG may inhibit the overall performance of the application. That is, the resulting data flow from the initial topology of a graph can become a bottleneck via network and disk utilization. A graph can be reshaped, producing an identical result with a more ideal dataflow that limits network and disk strain. This work provides the methodology used for function annotation, which provides insight to reshape applications and better manage data. The methodology is incorporated into the composition of the frameworks Dask and TaskVine. Additionally, we show improvement with real-world applications in the domain of high energy physics

Workflow Stack



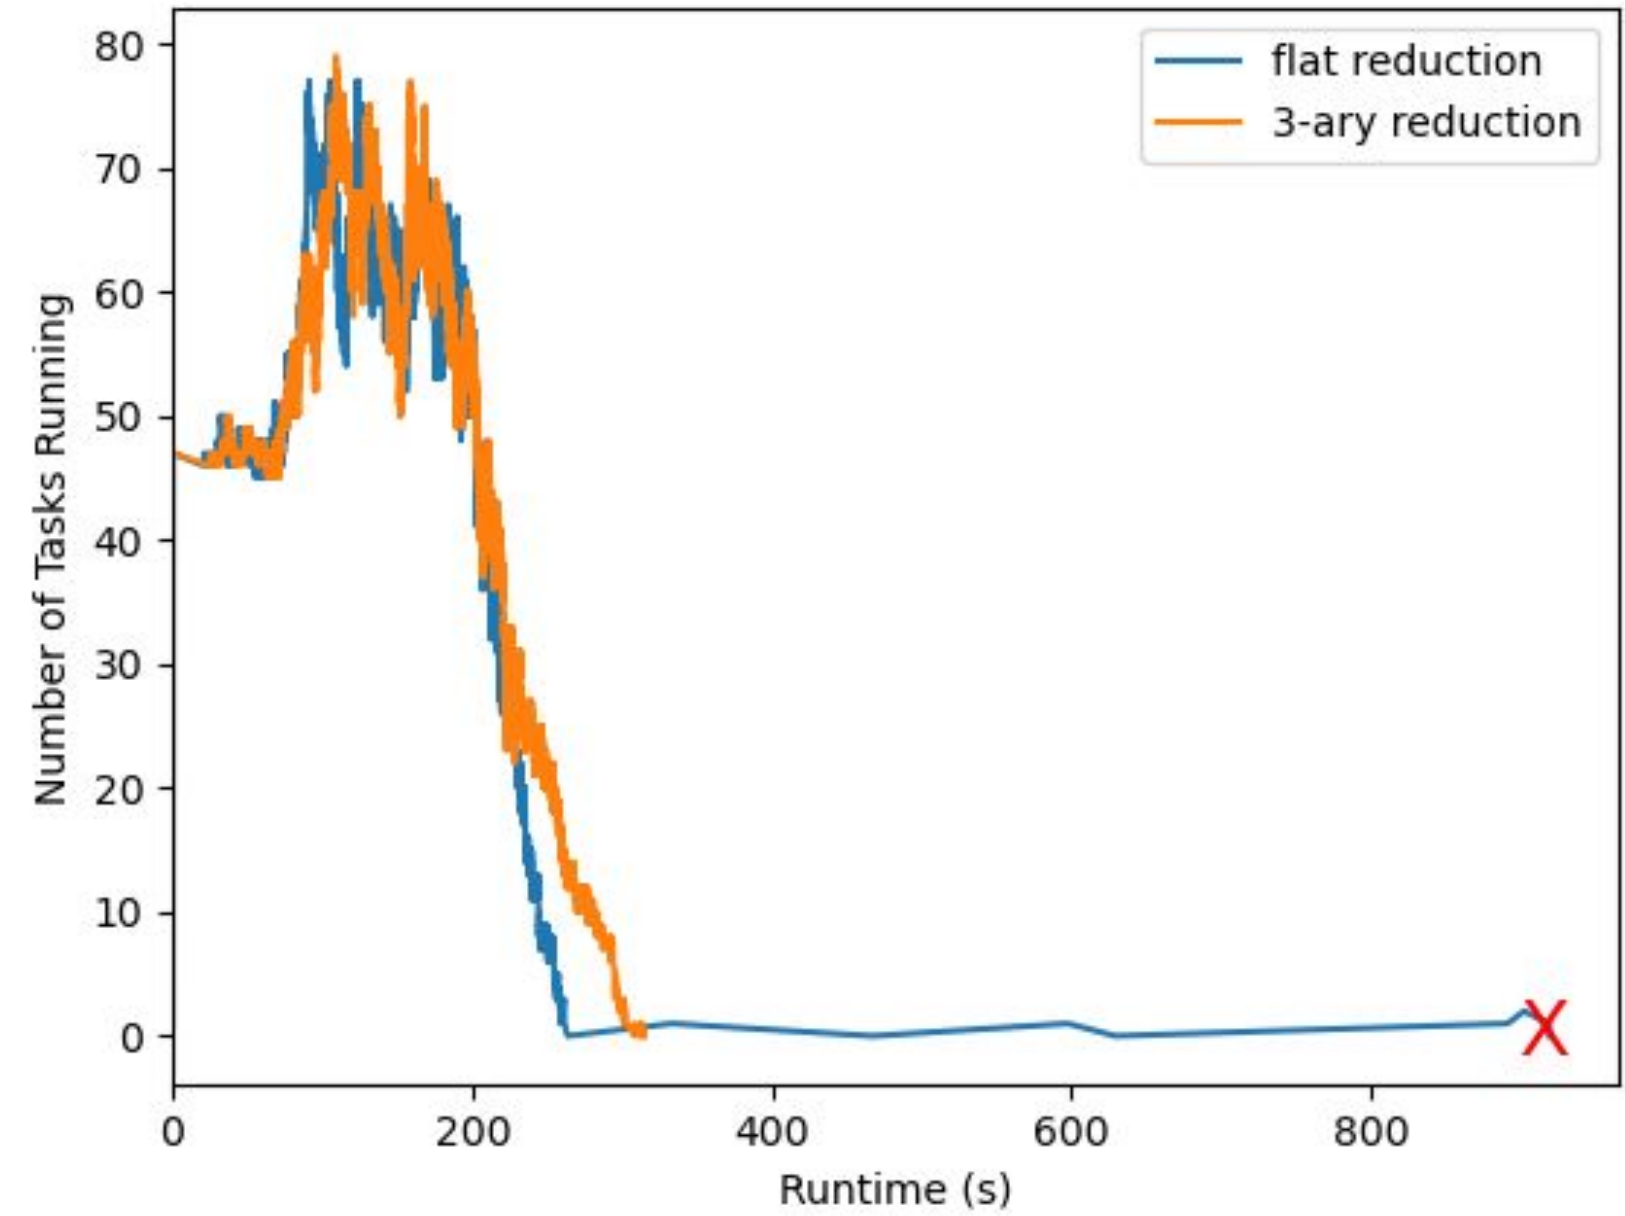

Graph Restructuring



Code

```
1 def write(events):
2     import awkward as ak
3     d = ak.to_parquet(events, "out.parquet")
4     return d
5
6 @daskvine_merge
7 def merge(*events):
8     import awkward as ak
9     return ak.concatenate(events)
10
11 graph = {}
12 events = get_events()
13 proxies = []
14
15 # Tree Reduction
16 for i in range(events.npartitions):
17     d = events.partitions[i:i+1]
18     proxies.append(d)
19 while(len(delayed) > 1):
20     # removes merge_size proxies
21     to_merge = get_proxies(merge_size, proxies)
22     d = dask.delayed(merge)(*to_merge)
23     proxies.append(d)
24 d = dask.delayed(write)(proxies[0])
25 graph[f"{dataset_name}_write"] = d
26
27 # Single Node Reduction
28 d = dask.delayed(write)(d)
29 graph[f"{dataset_name}_write"] = d
30
31 # Auto Merge
32 d = dask.delayed(merge)(events)
33 d = dask.delayed(write)(d)
34 graph[f"{dataset_name}_write"] = d
```

Results




(Left) - Competition times of X% of tasks using a N-ary reductions. Generally, as the size of N increases the time to complete X% of tasks increases as well. Note the increasing disparity between 99% and 100% completion

(Right) - The amount of concurrent tasks executing during a run of RsTriPhoton using a flat reduction and a 3-ary reduction. Larger reduction sizes produce a large tail during execution.



UNIVERSITY OF
NOTRE DAME



TaskVine