

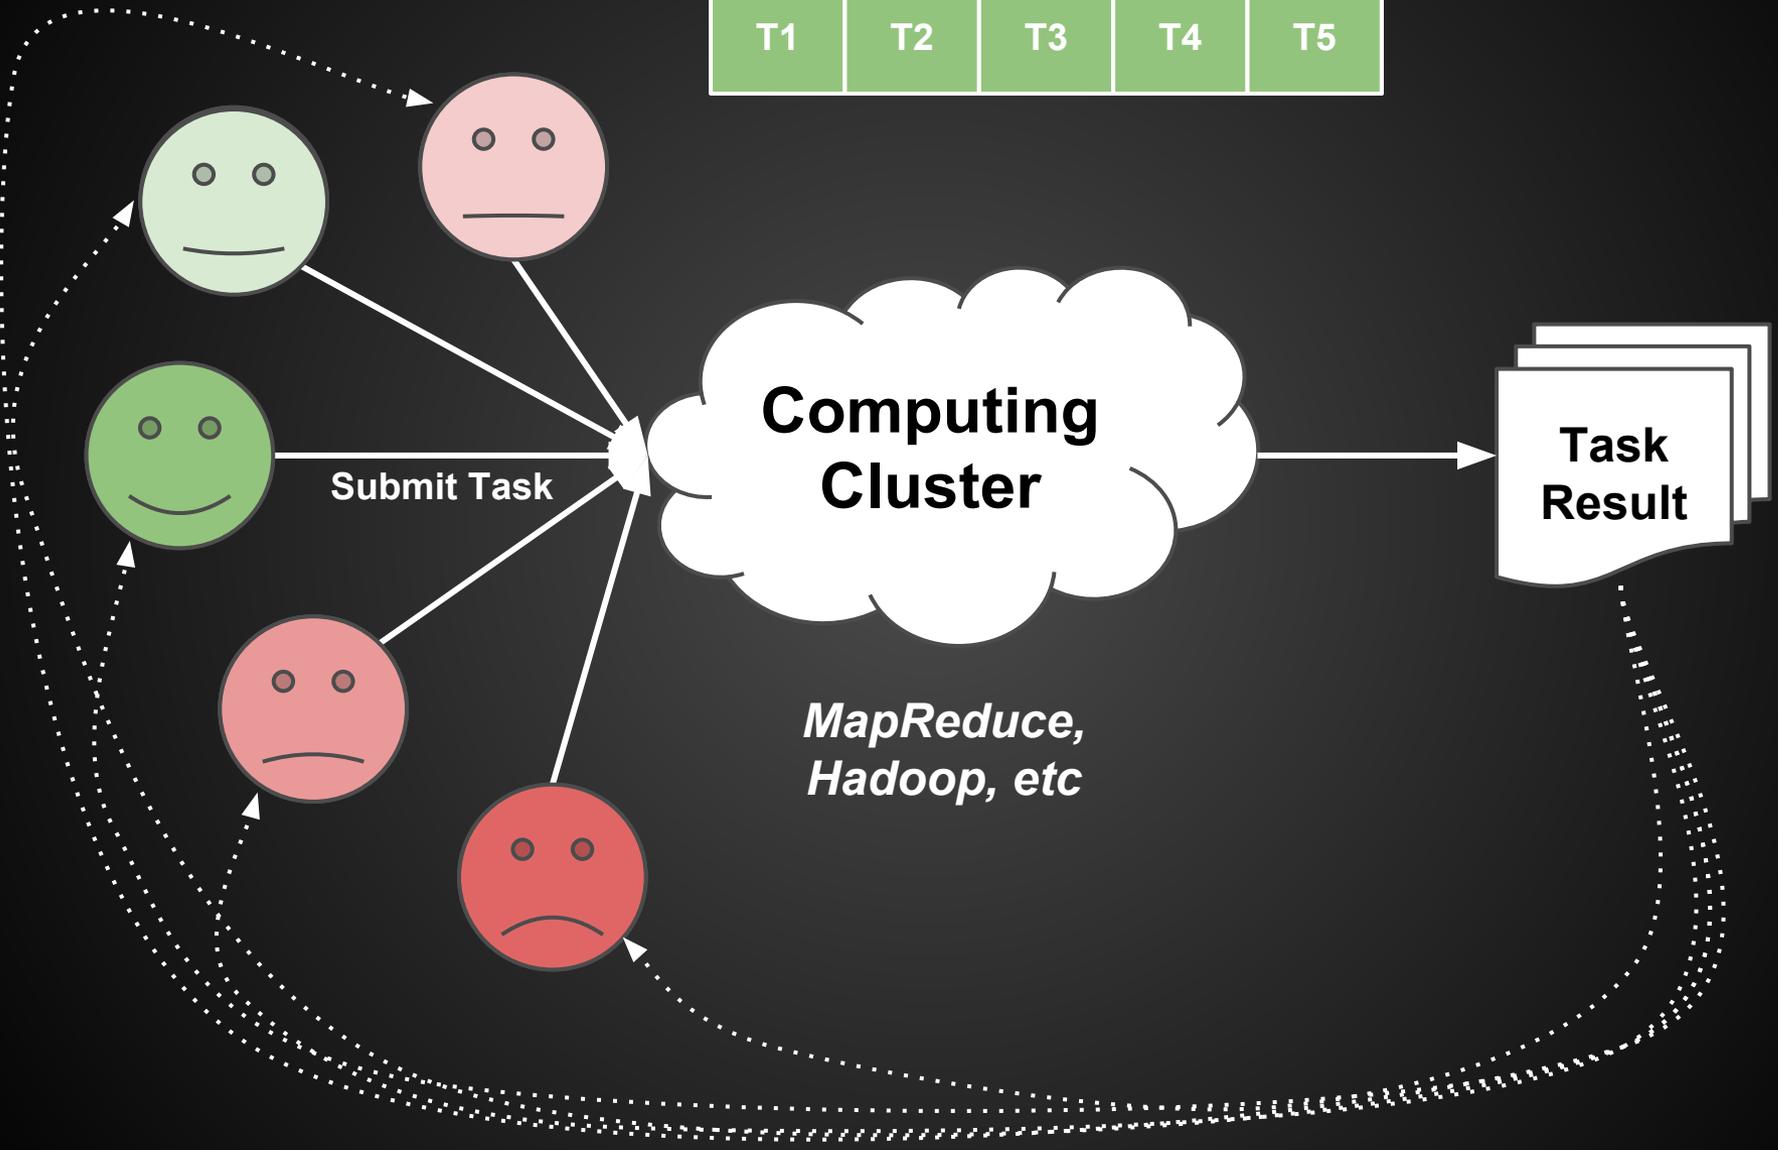
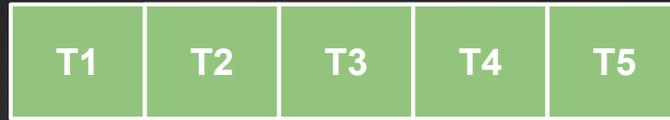
Task Fusion: Improving Utilization of Multi-user Clusters

Robert Dyer
rdyer@iastate.edu

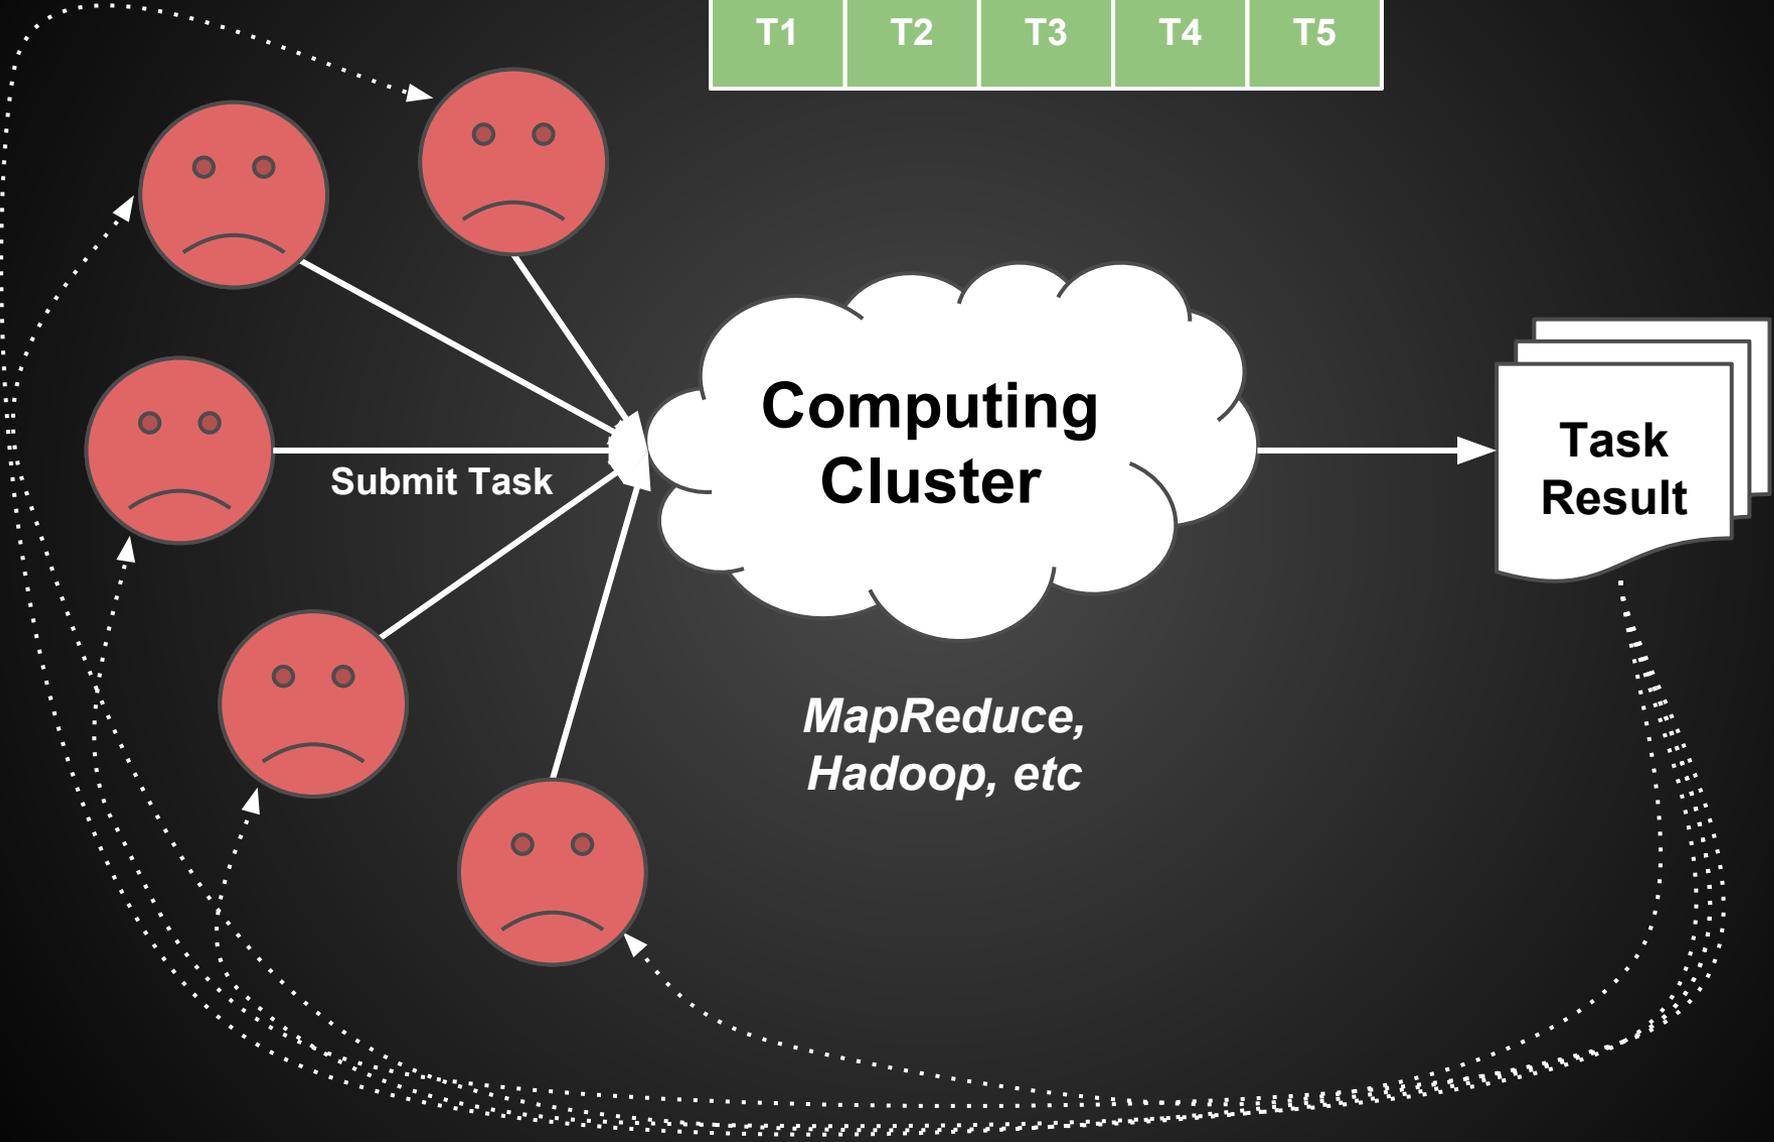
Iowa State University

The research and educational activities described in this talk was supported in part by the US National Science Foundation (NSF) under grants CCF-13-49153, CCF-11-17937, and CCF-08-46059.

FIFO Queue



Time Sharing



Solutions?

- Scale the hardware
 - Expensive
 - Not always feasible (small businesses, MOOCs, researchers, etc)
- Optimize the software
 - Optimize individual tasks
 - standard program optimizations
 - chain folding [MinerShook12], sibling/MSCR fusion [Chambers10]
 - Optimize multiple tasks
 - manual job merging [MinerShook12]

[Chambers10] Craig Chambers et al., “FlumeJava”, PLDI 2010

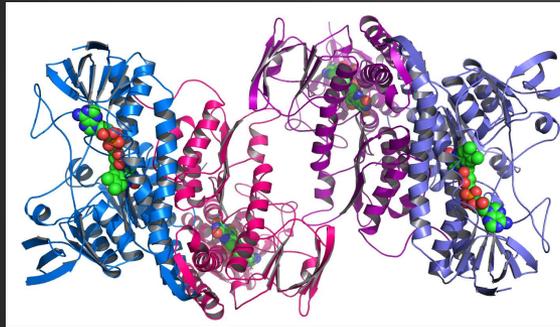
[MinerShook12] Donald Miner and Adam Shook, “MapReduce Design Patterns”, O’Reilly, 2012

Key Insights

1) People analyze similar data



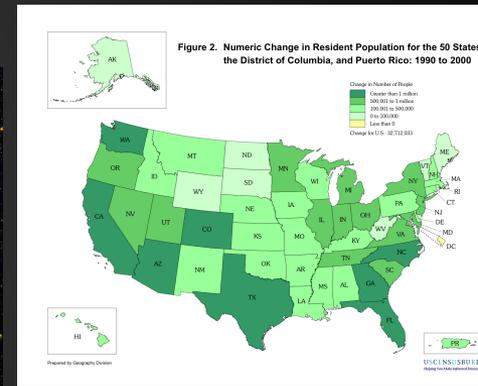
NCDC, NCCS



1k Genomes Project



SDSS



US Census

2) Data-intensive computing

- Loading GB/TB/PB of data takes time

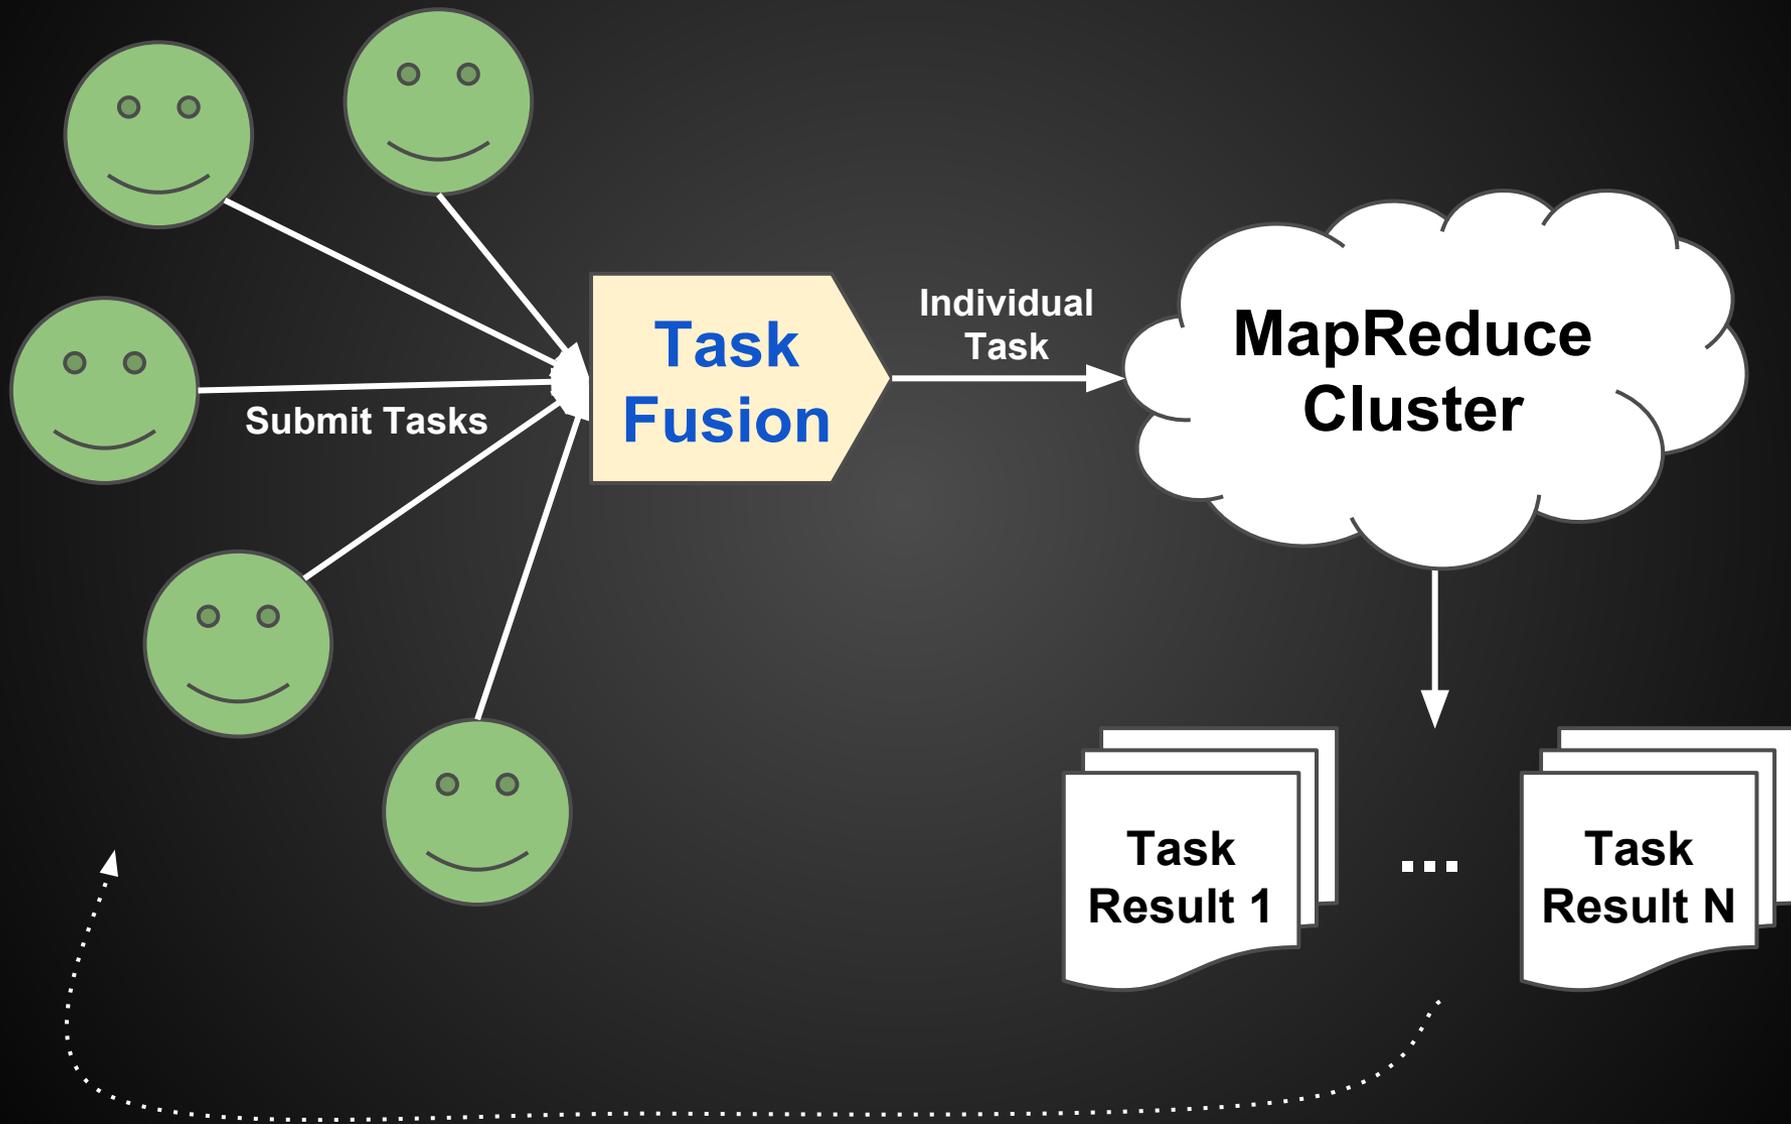
Insight: Load data once, run multiple analyses

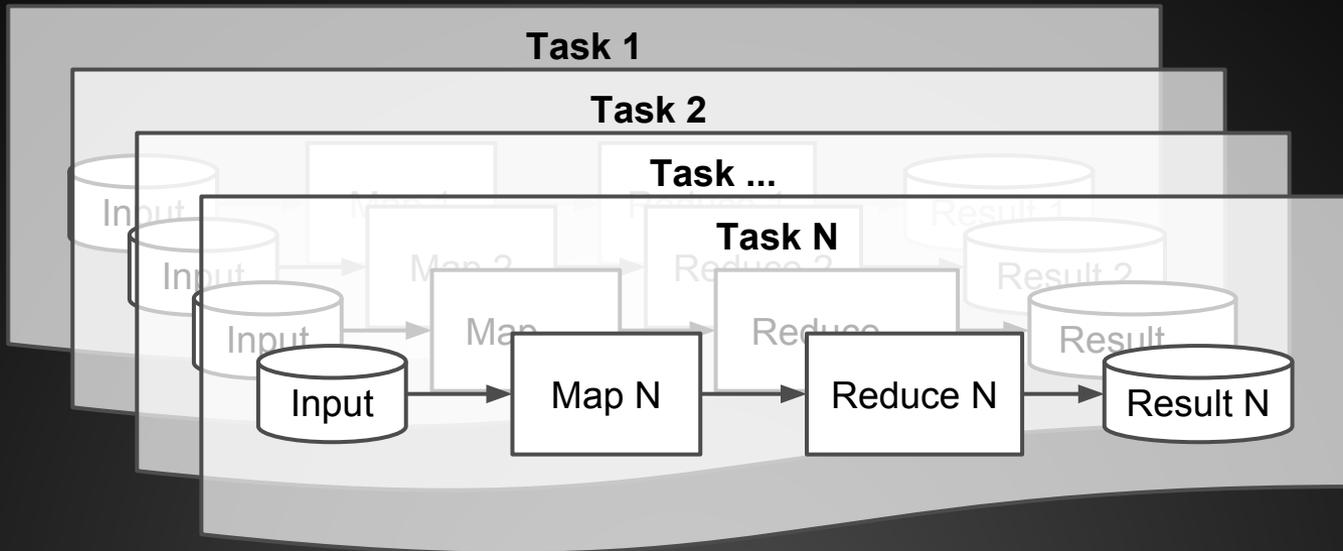
Research Questions

1. Can we **automatically** merge related tasks from **different users**?

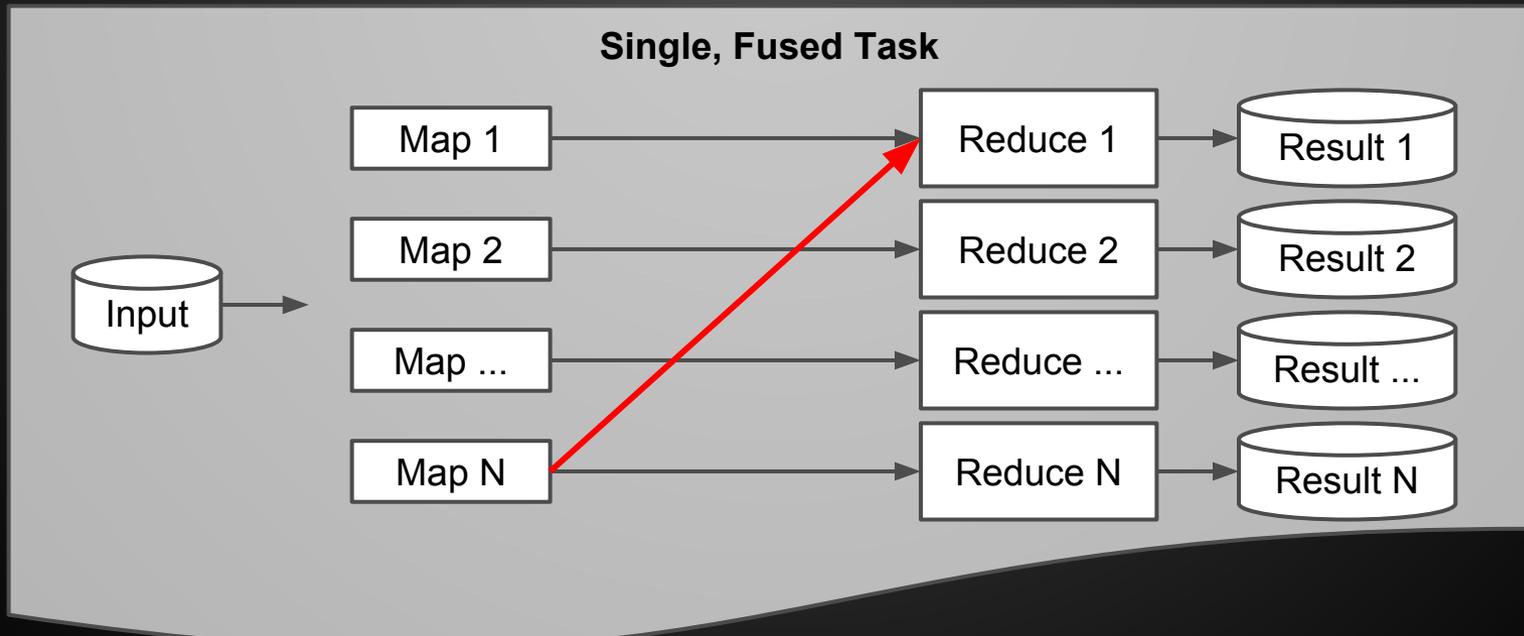
Answer: **Task Fusion**

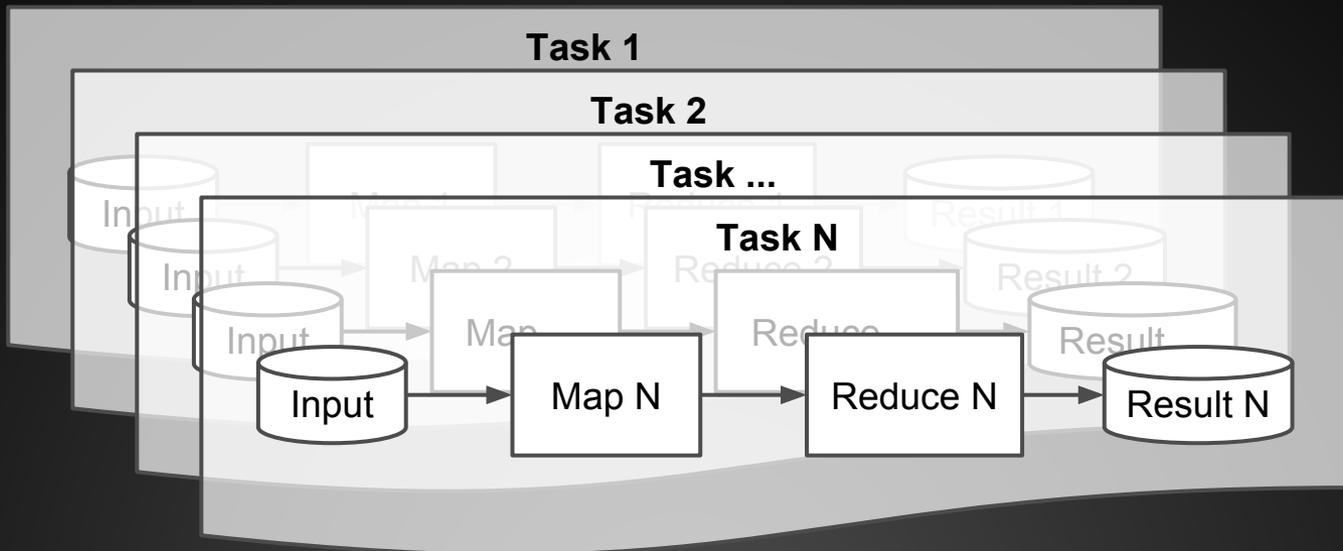
2. Does *Task Fusion* decrease user wait times in shared computing clusters?





Technical Challenge:
map output == side effect

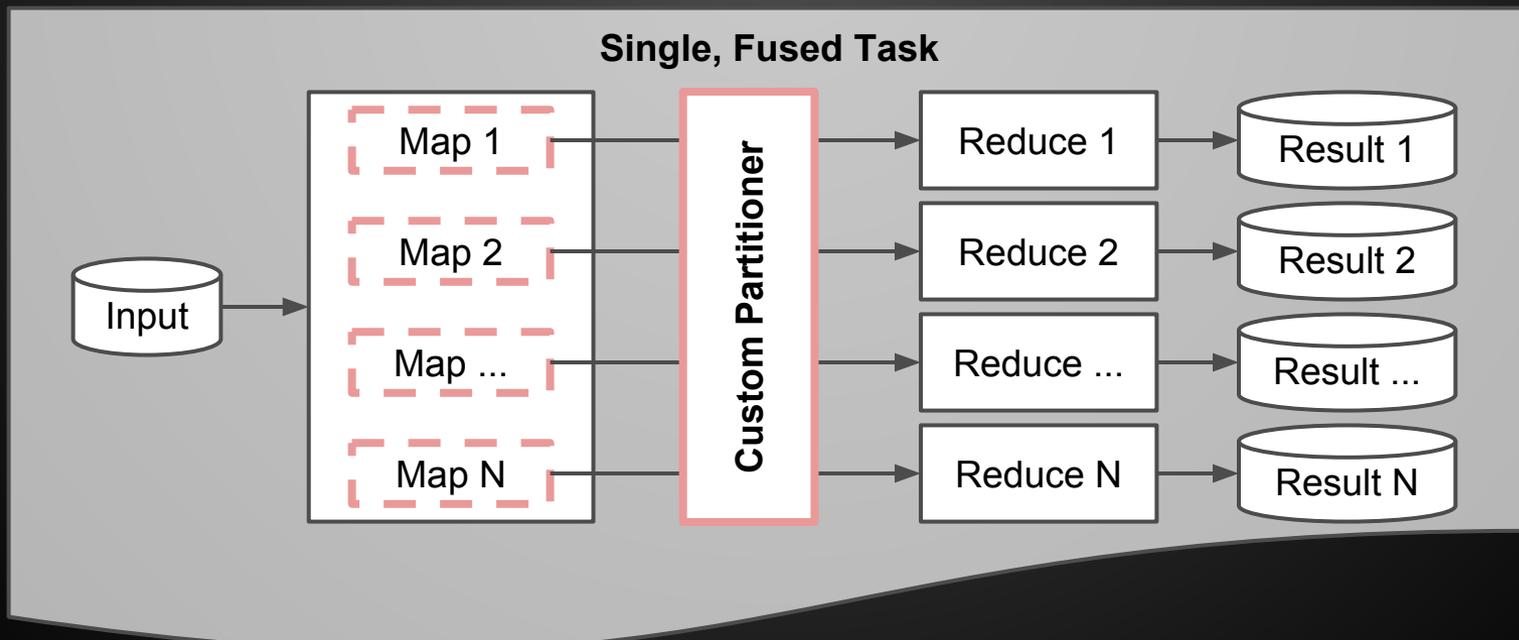




Solution: **modify maps to output composite keys**



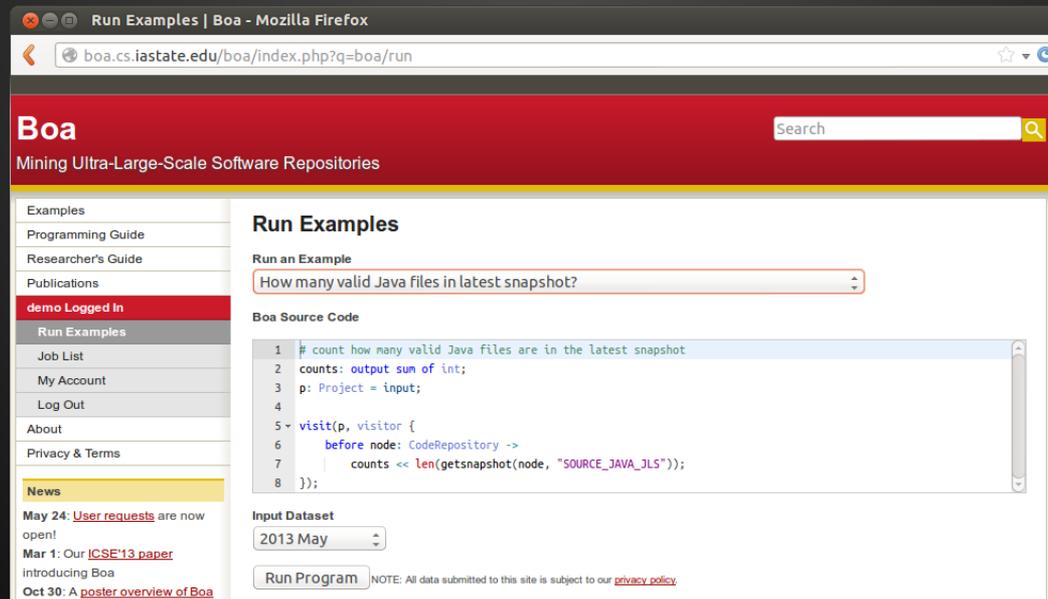
Custom partitioner ensures proper routing



Research Prototype

Task Fusion implemented for Boa

- Large-scale software repository mining
- SourceForge data (700k projects)
- Automatically parallelizes queries



The screenshot shows a web browser window titled "Run Examples | Boa - Mozilla Firefox" with the URL `boa.cs.iastate.edu/boa/index.php?q=boa/run`. The page features a red header with the "Boa" logo and the tagline "Mining Ultra-Large-Scale Software Repositories". A search bar is located in the top right corner. On the left side, there is a navigation menu with items like "Examples", "Programming Guide", "Researcher's Guide", "Publications", "demo Logged In", "Run Examples", "Job List", "My Account", "Log Out", "About", "Privacy & Terms", and "News". The main content area is titled "Run Examples" and contains a "Run an Example" dropdown menu with the selected option "How many valid Java files in latest snapshot?". Below this is a "Boa Source Code" section with a code editor showing the following code:

```
1 # count how many valid Java files are in the latest snapshot
2 counts: output sum of int;
3 p: Project = input;
4
5 visit(p, visitor {
6   before node: CodeRepository ->
7     counts << len(getsnapshot(node, "SOURCE_JAVA_JLS"));
8 });
```

Below the code editor is an "Input Dataset" dropdown menu set to "2013 May" and a "Run Program" button. A note at the bottom right states: "NOTE: All data submitted to this site is subject to our [privacy policy](#)."

Early Results

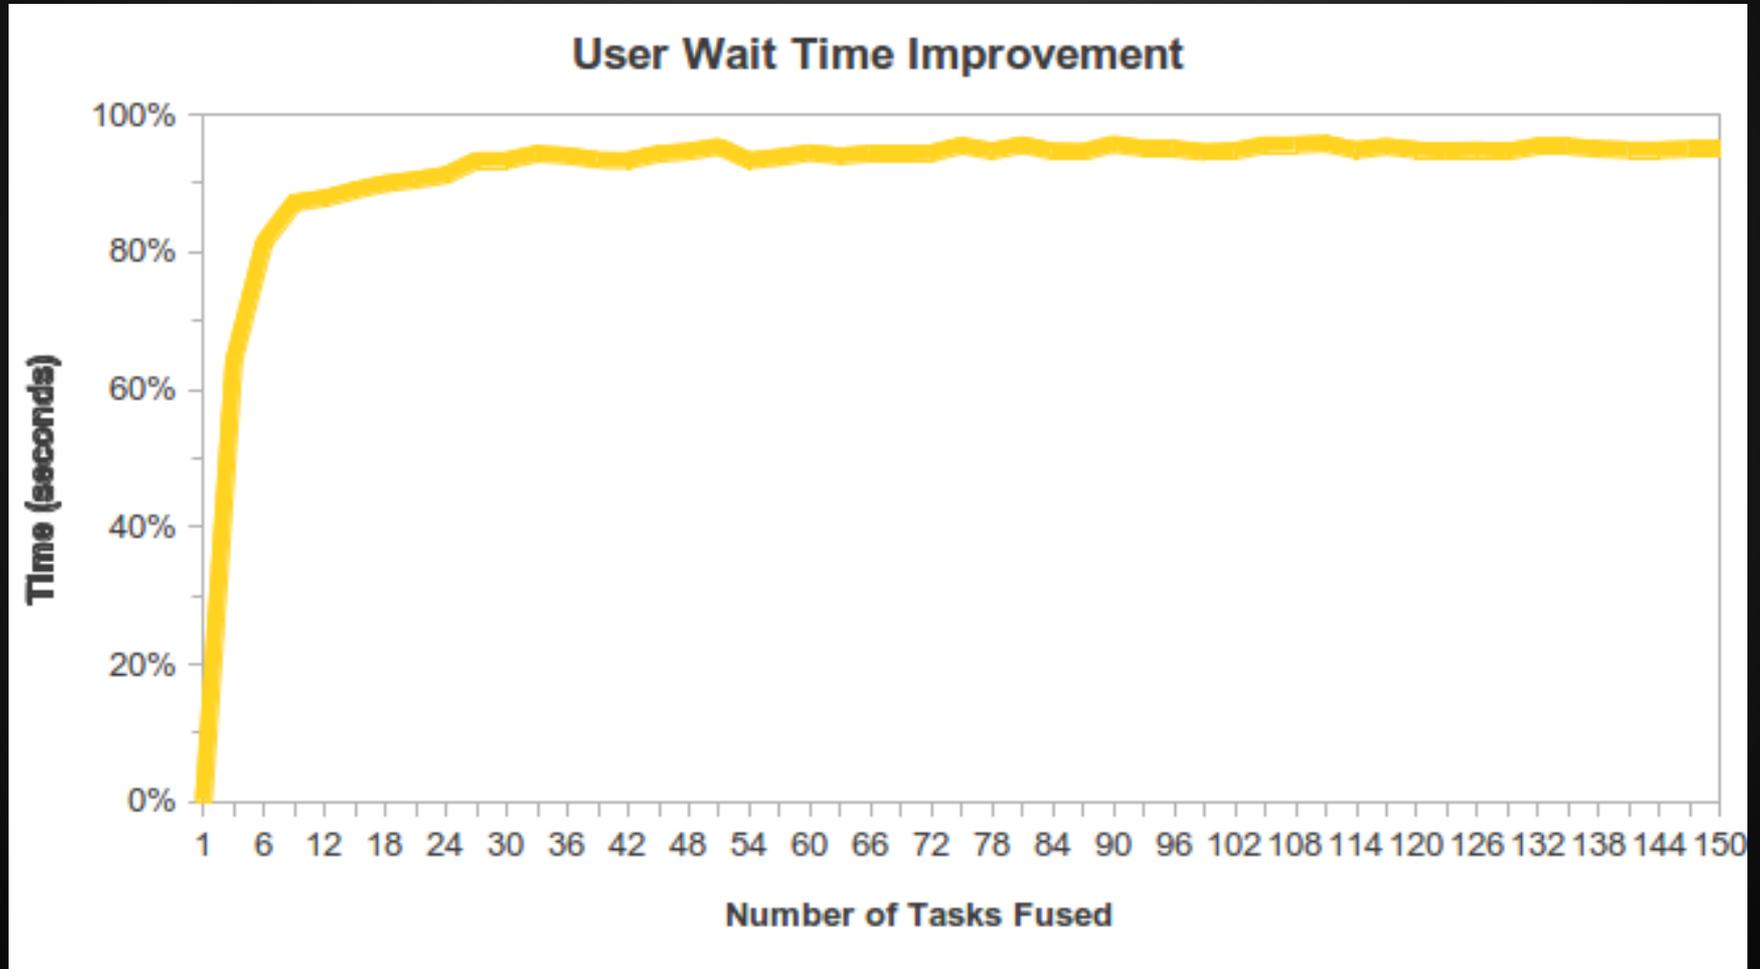
Task Size	# of Tasks	Times		Speedup
		No Task Fusion	Task Fusion	
Small ¹	21	8.1m	0.8m	10.8X
Medium ²	22	2.3h	1.8h	1.3X
Large ²	18	4.6h	3.9h	1.2X
Mixed ³	9	1.3h	0.9h	1.4X

[1] queries on project and revision metadata only

[2] queries on metadata and millions of source files

[3] 3 small, 3 medium, 3 large

Early Results



Assumptions - Relax Assumptions

1. No shared state
2. No dependency conflicts

Idea: Separate class spaces (a la OSGi)

3. Controllable side effects

Idea: Automated program transformations