

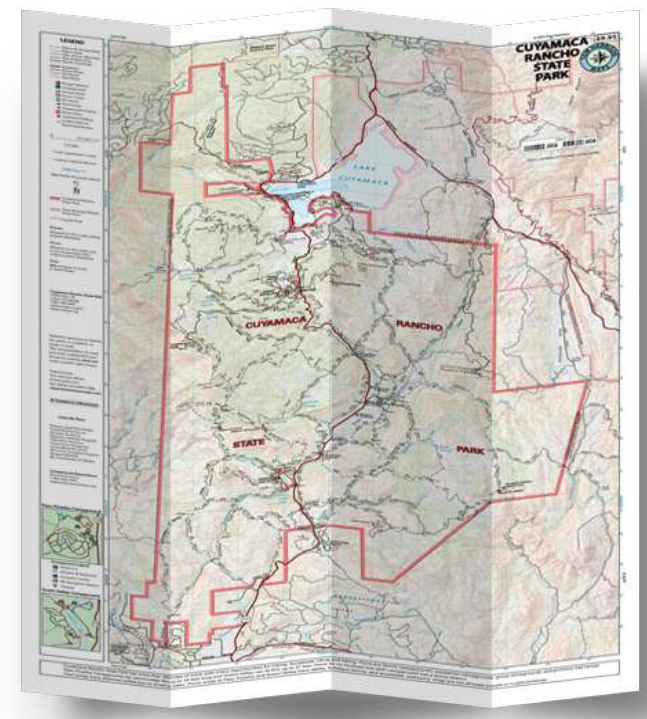
Wave : A Substrate for Distributed Incremental Graph Processing on Commodity Cluster

Swapnil Gandhi, Indian Institute of Science

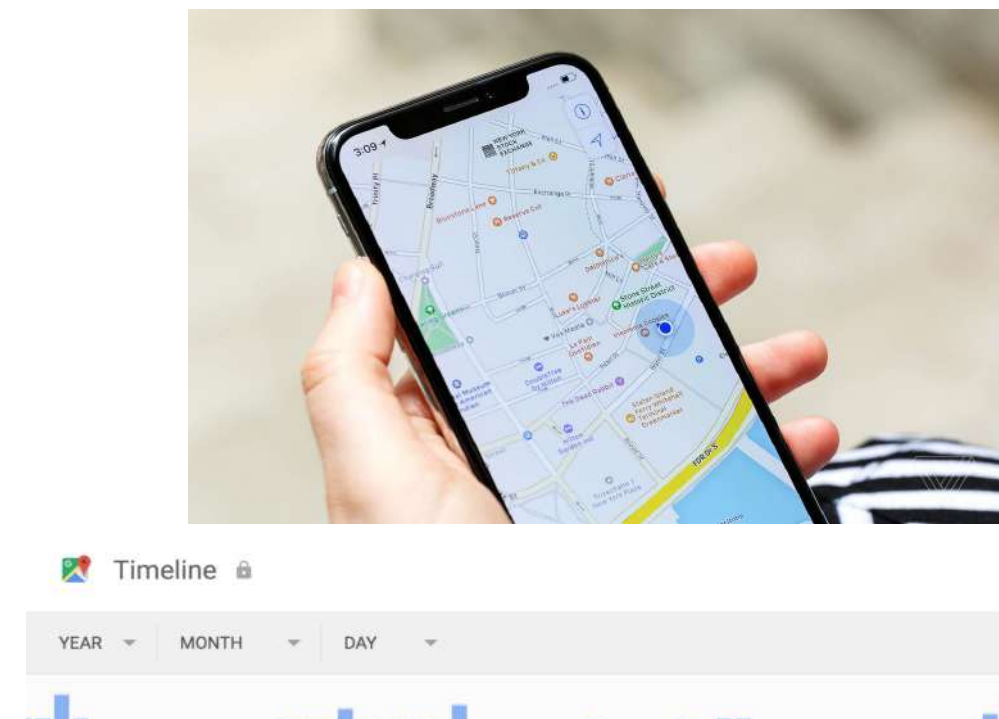
INTRODUCTION

- Graphs are widely considered to be natural means of representation for many networks.
- Real-world networks are often **evolving** with links being added or removed and properties updated over time.

Examples : Social, Citation/Collaboration, Sensor, Financial & Transit Network, Human Connectomes, Internet-of-Things ...



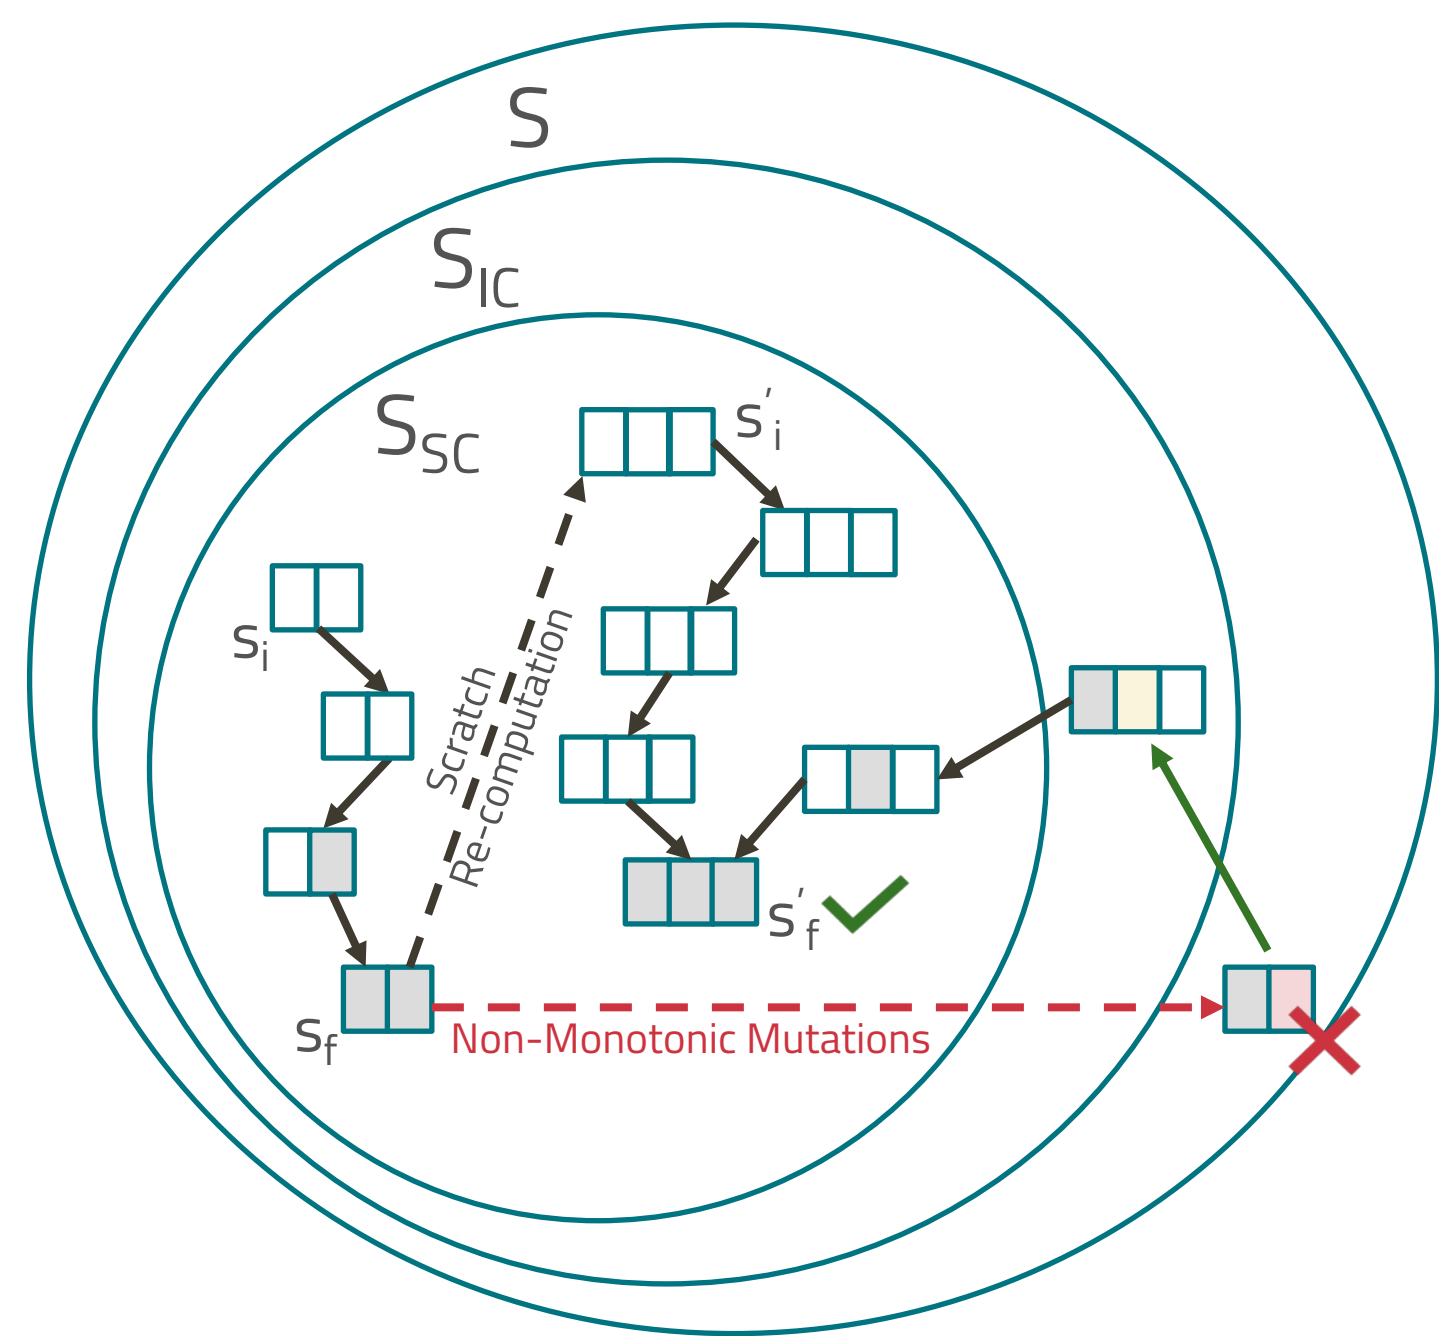
(a) Paper Map




(b) Google Maps

- Incremental computation is used to process such dynamic graphs.


It achieves efficiency by **minimizing redundant** computation and communication compared to complete, from-scratch computation.



(c) Incremental processing of dynamic graphs



(d) #Vertices with incorrect results (WK)



(e) #Vertices recomputed (WK)

- Challenges:

1. **Safety**: Reusing intermediate state naively leads to **incorrect results**.


2. **Profitability**: Conservative tag-propagation guarantees safety, however limits state reuse \Rightarrow ends up **resetting majority** of vertices.

3. Lack of **unifying abstraction** to operate on dynamic graphs.

Existing abstractions either work for a sub-class of algorithms or lack support for non-monotonic updates. Specialized algorithms are designed for single-threaded shared-memory execution.

This Work : General-purpose distributed programming model to support incremental processing over dynamic graphs while leveraging existing vertex-centric semantics.

CONNECTED COMPONENTS



(f) Dynamic graph

	V _A	V _B	V _C	V _D	V _E	V _F	
Before Deletes	A	A	C	A	A	A	1
	V _A \rightarrow V _B and V _B \rightarrow V _F Deleted						
	A	B	C	A	A	F	2
	A	B	C	A	B	F	3
	A	B	C	B	B	C	4
After Deletes	A	B	C	B	B	C	5
	V _E \rightarrow V _F Inserted						
After Insert	A	B	C	B	B	B	6

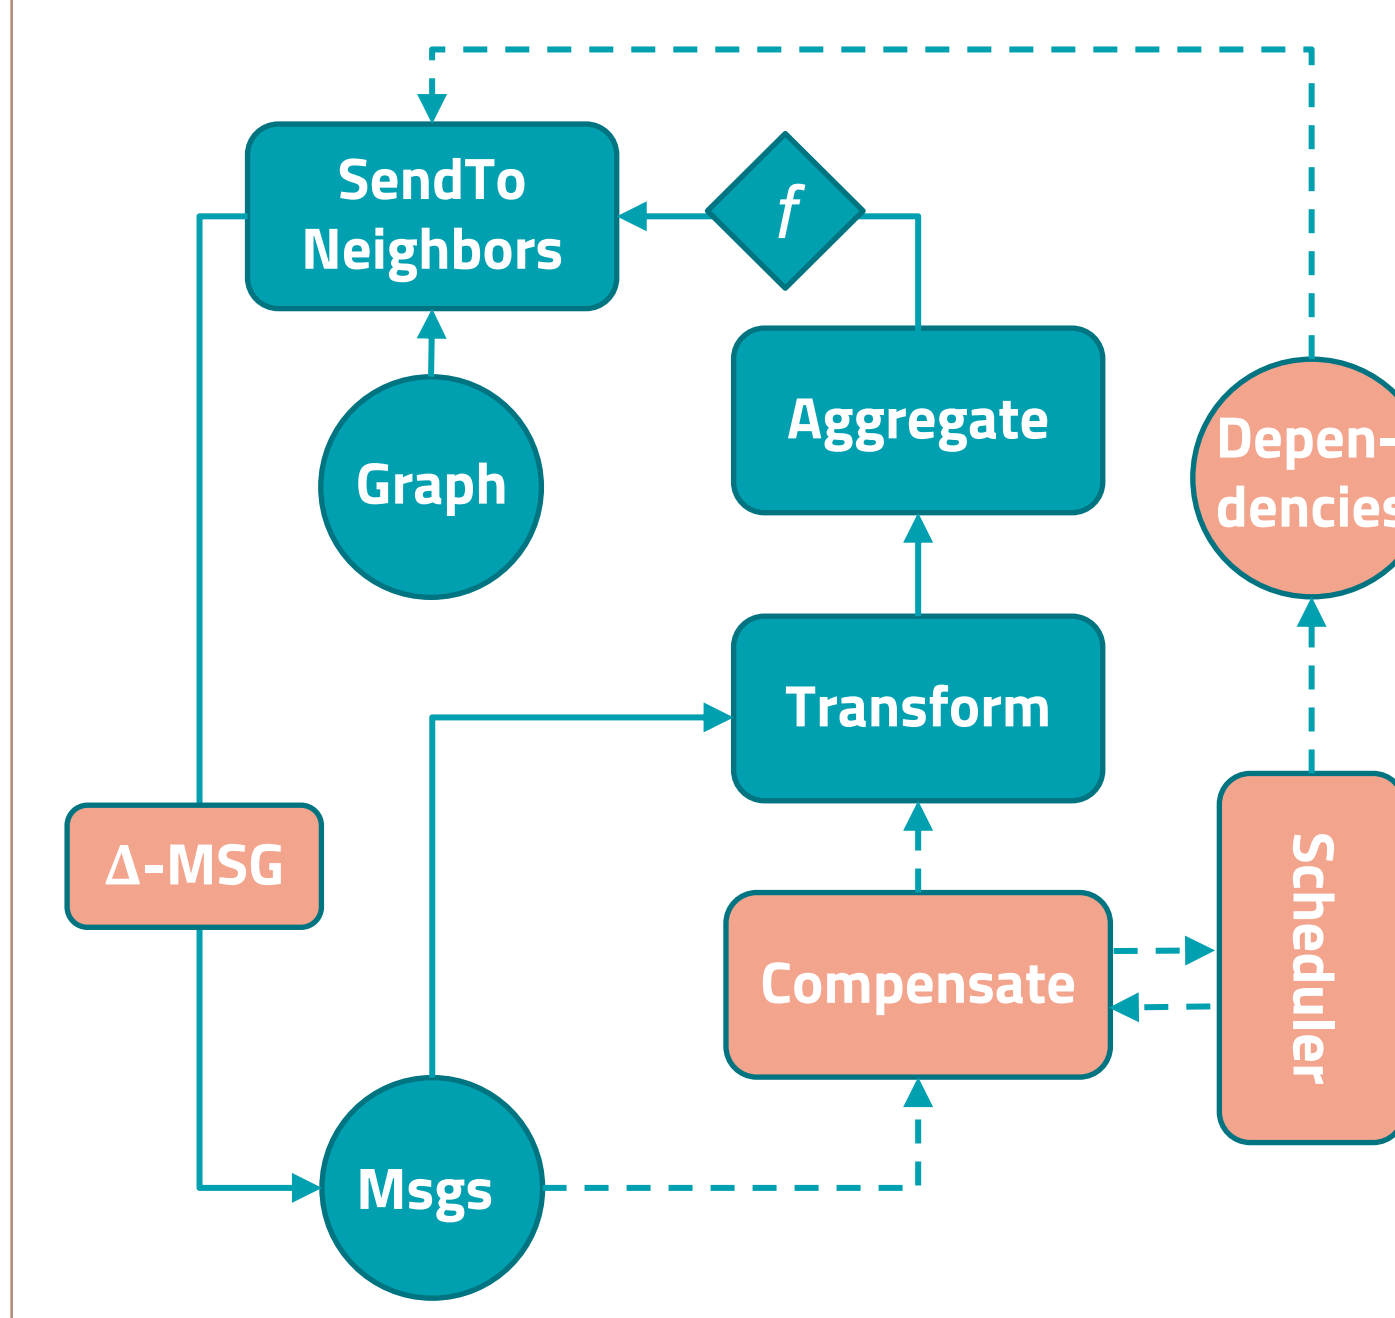
(g) Vertex state updates over supersteps when finding Connected Components. Changes shown in red.

INCREMENTAL GRAPH COMPUTATION

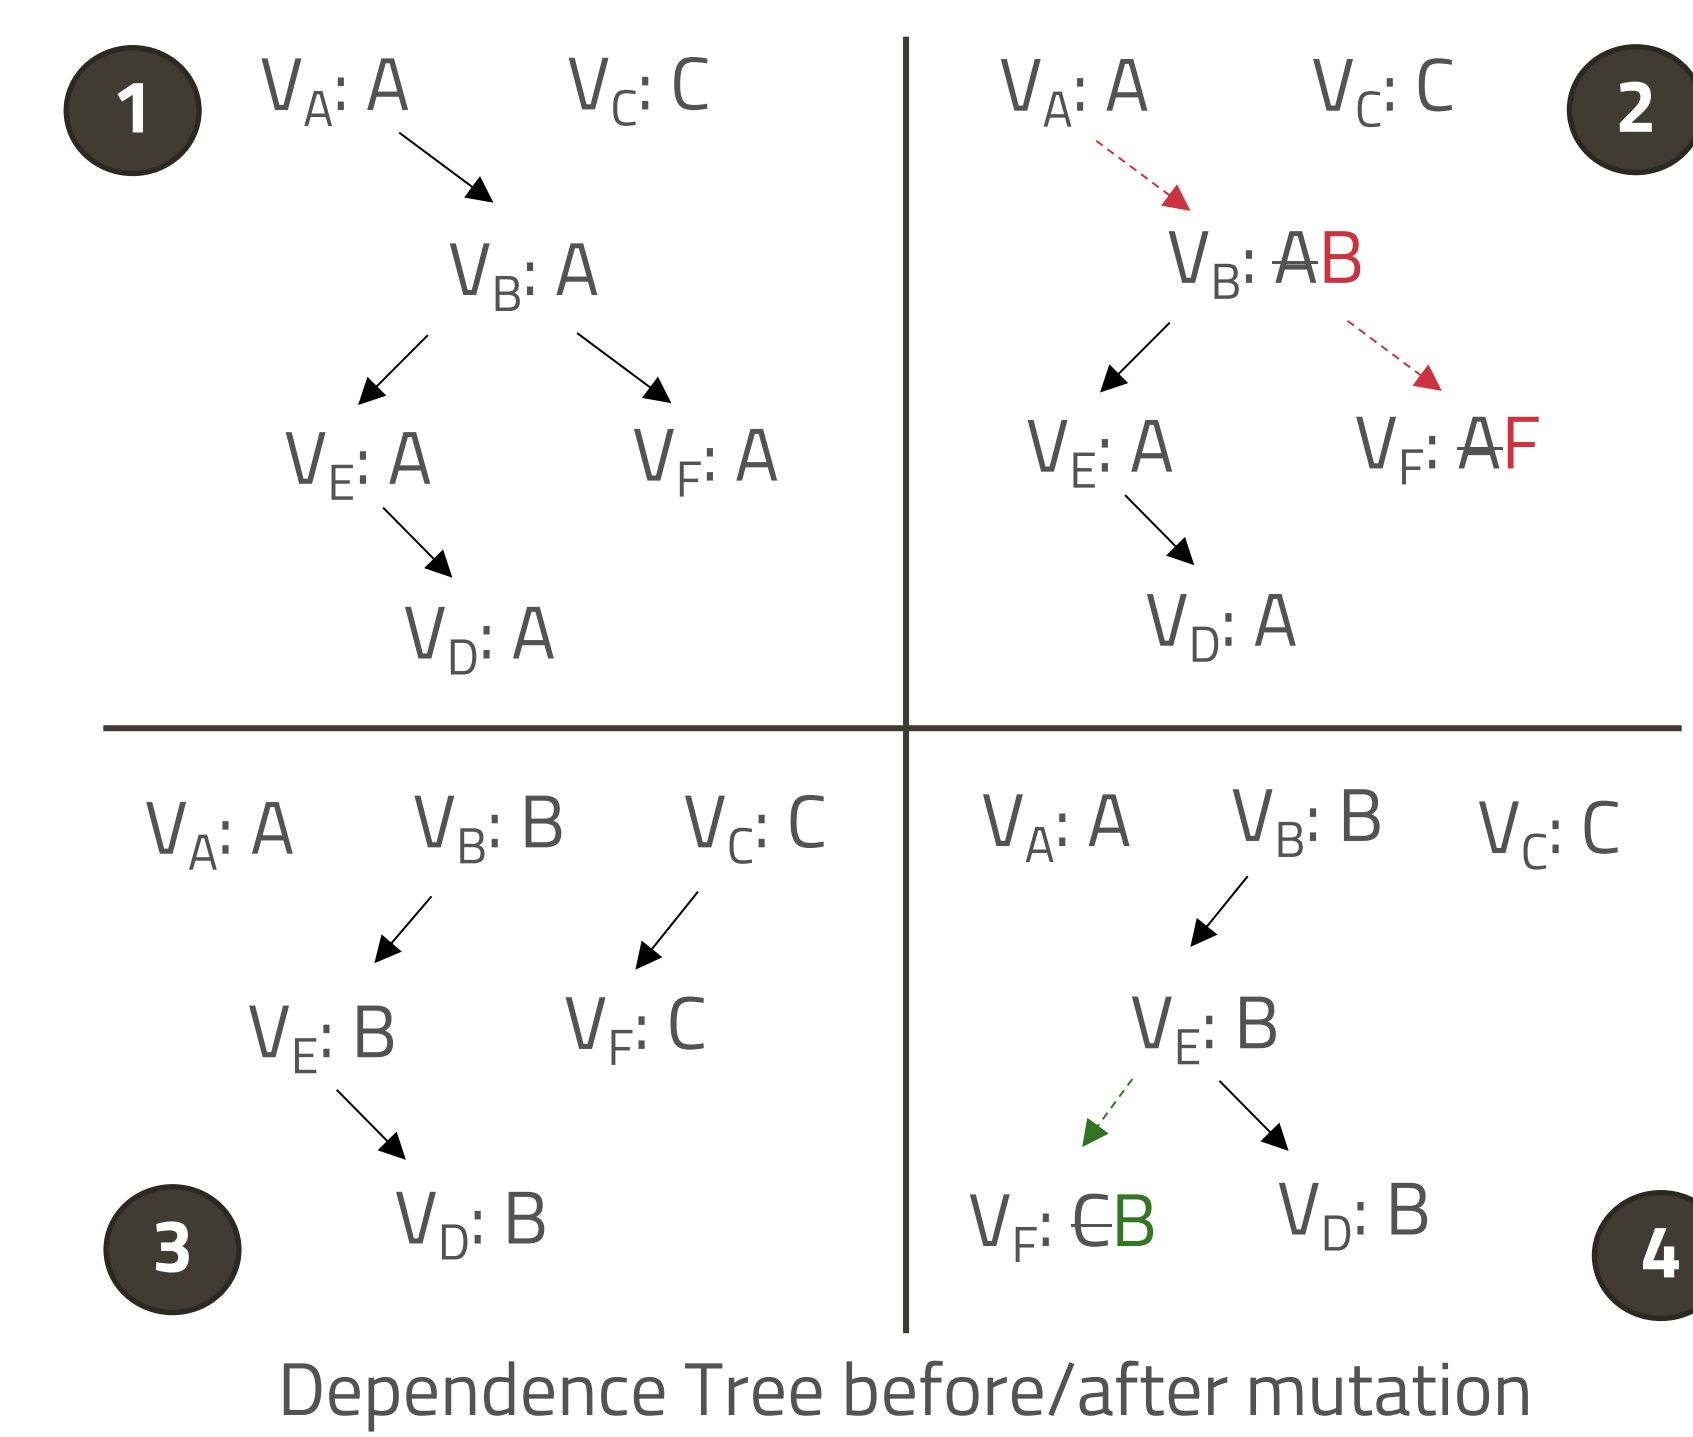
- Transparency**: We do not require the vertex program to be made incremental, but instead **incrementalize its boundary**.
 - Makes the approach applicable to all existing vertex programs
 - Requires minimal additional effort from the programmer to devise an incremental algorithm

- Framework identifies vertices directly and transitively affected by graph mutations.

- Only parts of the graph **affected** by input changes are re-computed
- Graph structure used to **actively deduce** value dependencies



Incremental Vertex Program in Wave



Dependence Tree before/after mutation

```

1 void compensate(vertex v, Message[] msgs,
2 boolean reScheduledCompute() {
3 Message[] cMsgs;
4 for(Message msg : msgs) {
5 if(msg.getType() ==
6 MessageType.RETRACT) {
7 retract(vertex v, msg);
8 } else {
9 cMsgs.append(msg);
10 v.dequeueDeferredMsg(msg.getSrc());
11 }
12 }
13 for(Message deferredMsg :
14 v.getDeferredMsgs()) {
15 if(deferredMsg.getSuperstepTillDeferred()
16 == getCurrentSuperstep()) {
17 cMsgs.append(deferredMsg);
18 v.dequeueDeferredMsg(deferredMsg);
19 }
20 }
21 if(isVertexStateAffected()) {
22 deferMsgs(v, cMsgs);
23 if(isInvertible() || reScheduledCompute())
24 compute(v, cMsgs);
25 } else {
26 gatherStateFromInNeighbors(v);
27 rescheduleCompute();
28 }
29 }

```

Master Program

```

1 void init(vertex v) {
2 v.setState(v.getId());
3 v.sendToAllNeighbors(v.getState());
4 }
5
6 void compute(vertex v, Message[] msgs) {
7 minCC = ∞;
8 for(Message msg : msgs) {
9 minCC = min(msg, minCC);
10 if(minCC < v.getState()) {
11 v.setState(minCC);
12 v.sendToAllNeighbors(v.getState());
13 }
14 }
15
16 void retract() {
17 if(v.getState().getSRC() == msg.getSRC()) {
18 init();
19 }
20
21 Message repropagate(vertex v, int dst) {
22 return new Message(v.getState());
23 }

```


User Program

EXPERIMENTAL EVALUATION

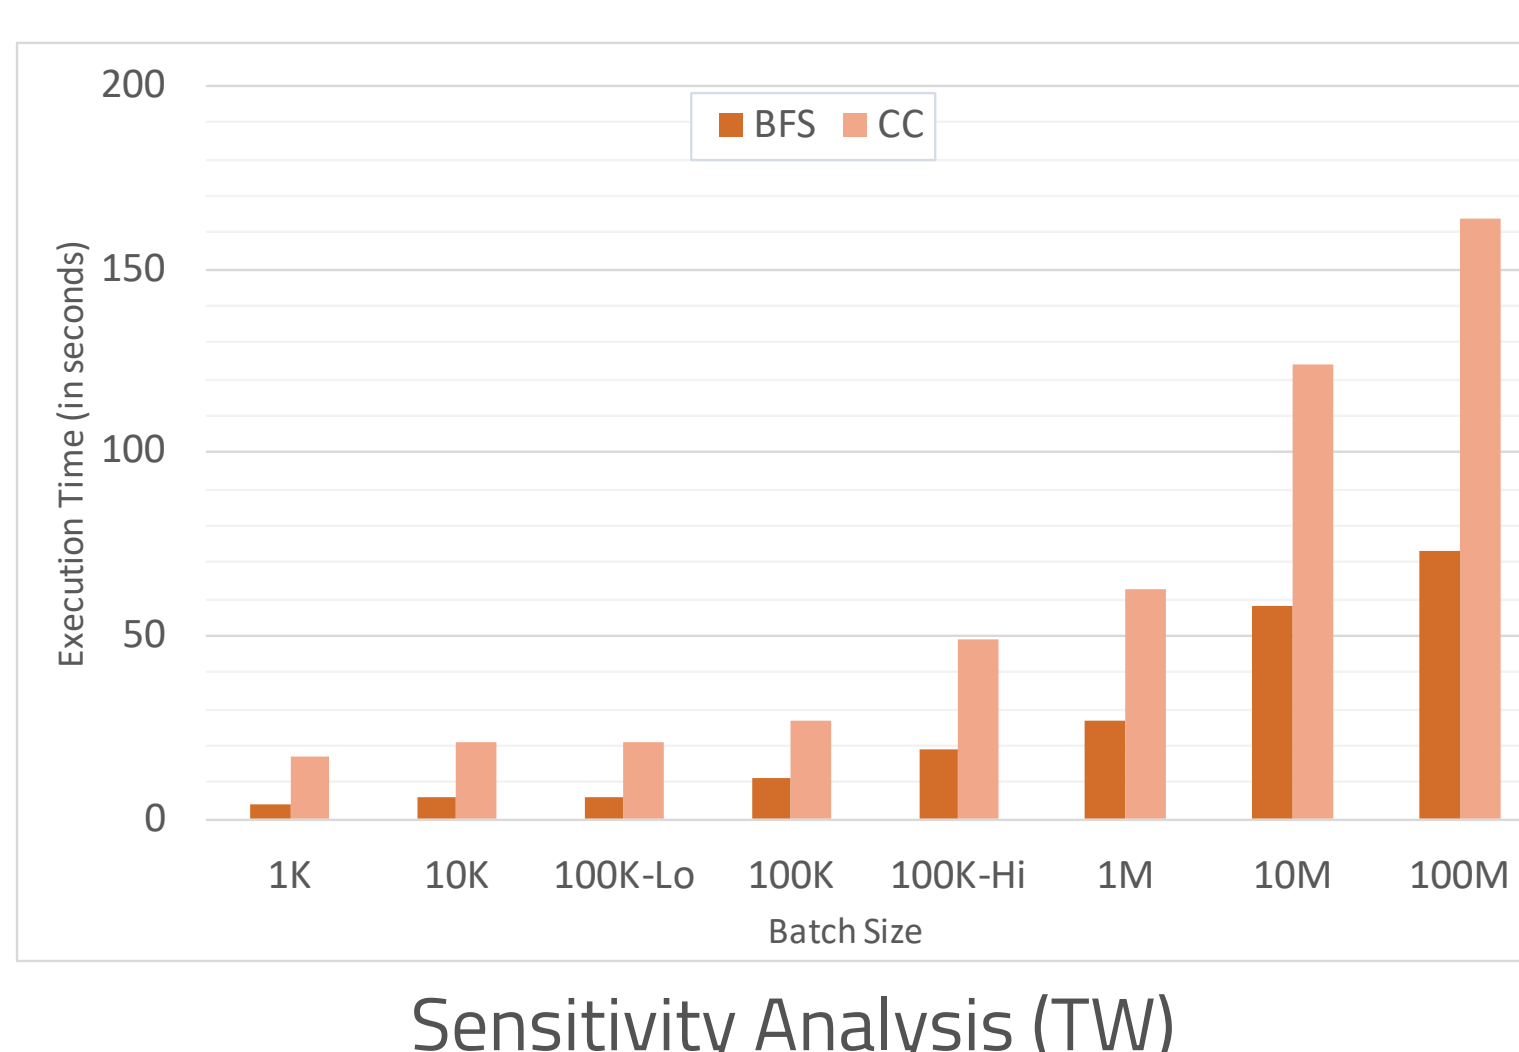
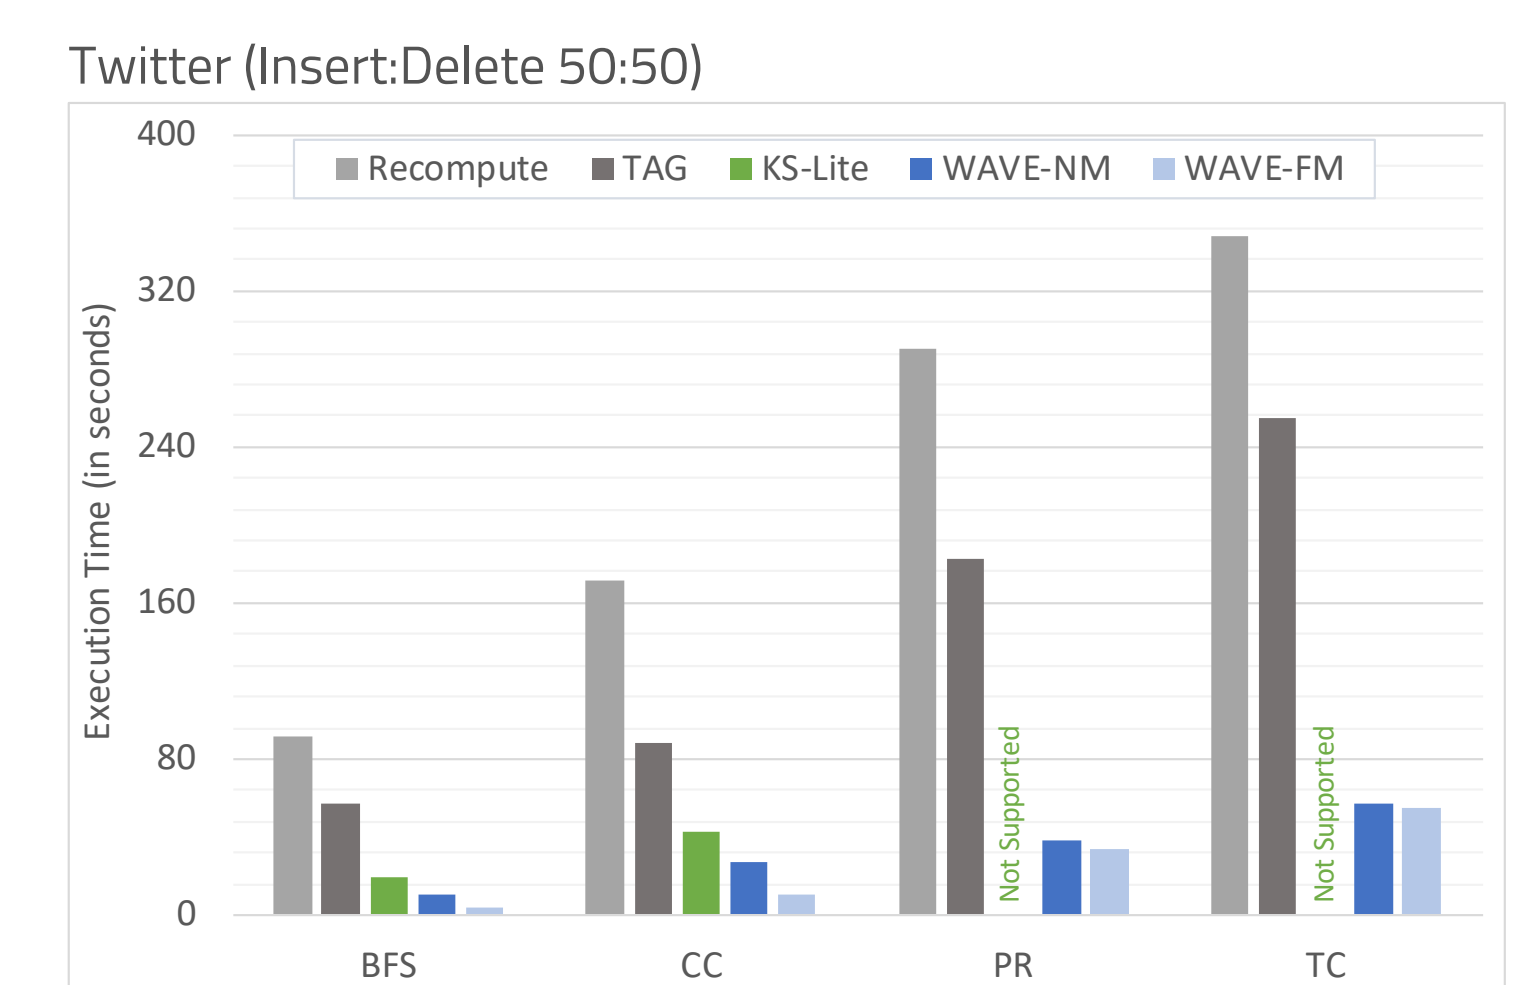
Algorithm	Aggregation
BFS	$\min_{v \in (u,v) \in E} (s(u) + 1)$
CC	$\min(v, \min_{v \in (u,v) \in E} u)$
PR	$\sum_{v \in (u,v) \in E} \frac{s(u)}{\text{out_degree}(u)}$
TC	$\sum_{v \in (u,v) \in E} \text{in_neighbors}(u) \cap \text{out_neighbors}(v) $

Graph	Vertices	Edges
Wiki (WK)	12M	378M
Twitter (TW)	41M	1.4B


8 Servers (Gigabit Ethernet)
Server : 16 threads / 2.1 GHz / 64 GBs
Apache Giraph 1.3 / Java 8.0



Execution time (in seconds) for Recompute, TAG, Kickstarter-Lite, Wave for 100K mutations
Wave is 6-23x faster than recomputation, 4-14x faster than TAG, and 2-6x faster than KS-Lite



Sensitivity Analysis (TW)



#Vertices recomputed (TW)

SUMMARY

Incremental Graph Processing \Rightarrow Better state re-use \Rightarrow Minimized redundant computation and communication \Rightarrow Faster convergence time

