

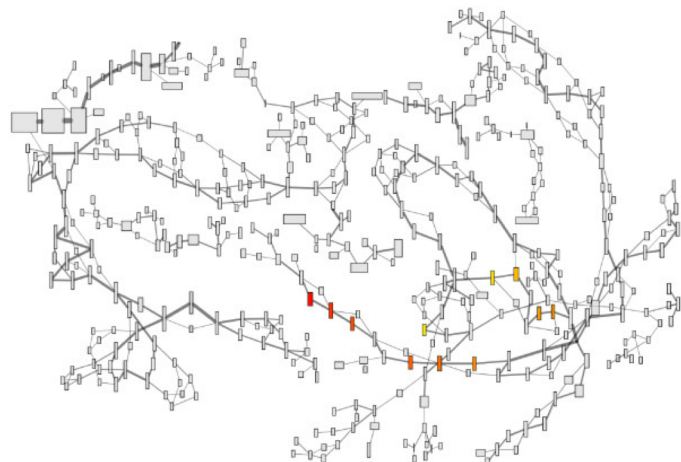

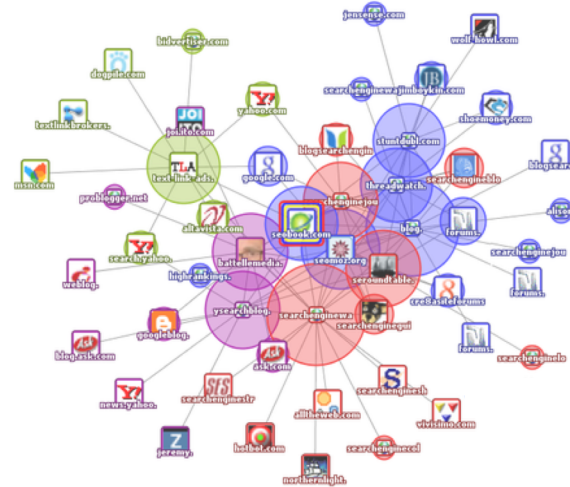
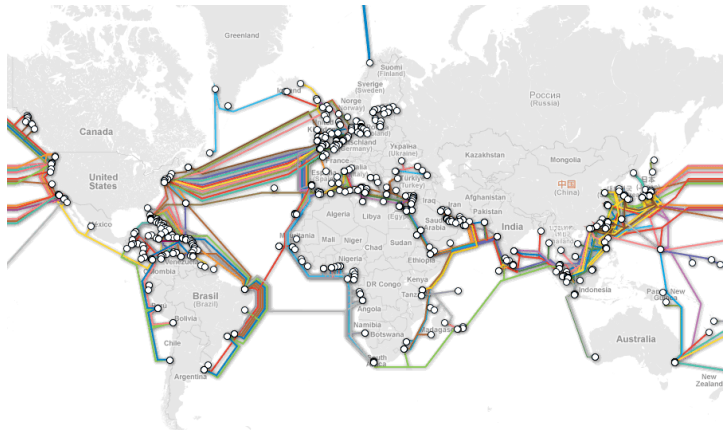


GraphBolt: Dependency-Driven Synchronous Processing of Streaming Graphs

Mugilan Mariappan and Keval Vora
Simon Fraser University

EuroSys'19 – Dresden, Germany
March 27, 2019

Graph Processing



Dynamic Graph Processing

Real-time Processing

- Low Latency


Real-time **Batch** Processing

- High Throughput



THE NEXT PLATFORM

Alipay payments unit of Chinese retailer Alibaba [..] has 120 billion nodes and over 1 trillion relationships [..]; this graph has **2 billion updates each day** and was running at **250,000 transactions per second** on Singles Day [..]



Streaming Graph Processing

Incremental Processing

- Adjust results **incrementally**
- Reuse work that has already been done

Tornado [SIGMOD'16]


GraphIn [EuroPar'16]

KineoGraph [EuroSys'12]

KickStarter [ASPLOS'17]

Tag Propagation
upon mutation

Over 75% values get thrown out



Streaming Graph Processing

Incremental Processing

- Adjust results **incrementally**
- Reuse work that has already been done

Tornado [SIGMOD'16]


GraphIn [EuroPar'16]

KineoGraph [EuroSys'12]

KickStarter [ASPLOS'17]

Maintain **Value Dependences**
Incrementally refine results

Less than 0.0005% values thrown out




Streaming Graph Processing

Incremental Processing

- Adjust results **incrementally**
- Reuse work that has already been done


Monotonic Graph Algorithms



KickStarter [ASPLOS'17]


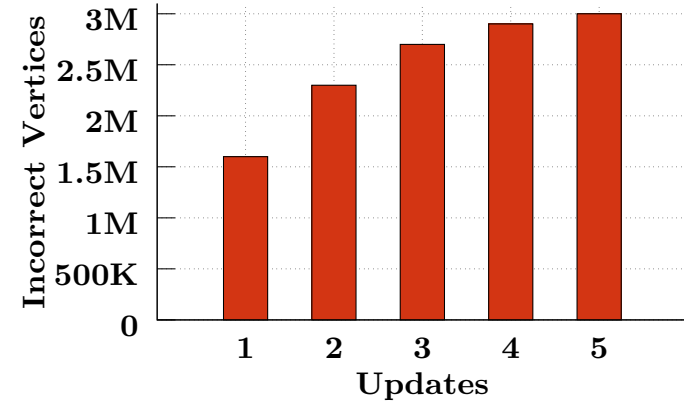
Maintain **Value Dependences**
Incrementally refine results

Less than 0.0005% values thrown out

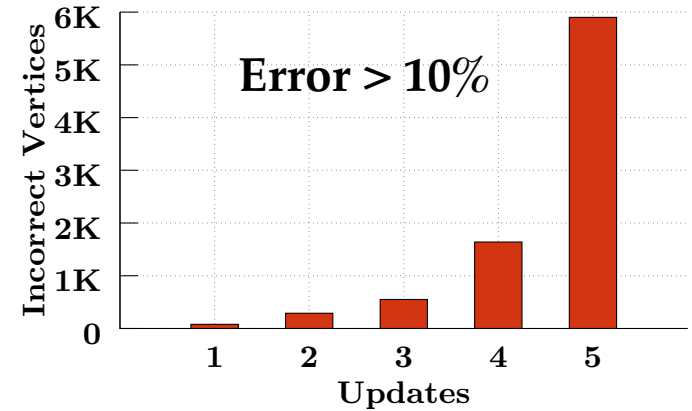



Streaming Graph Processing




- Belief Propagation
- Co-Training Expectation Maximization
- Collaborative Filtering
- Label Propagation
- Triangle Counting
- ...






Streaming Graph Processing




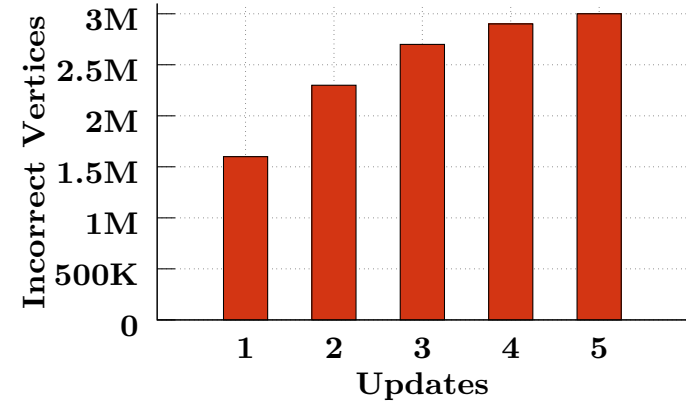
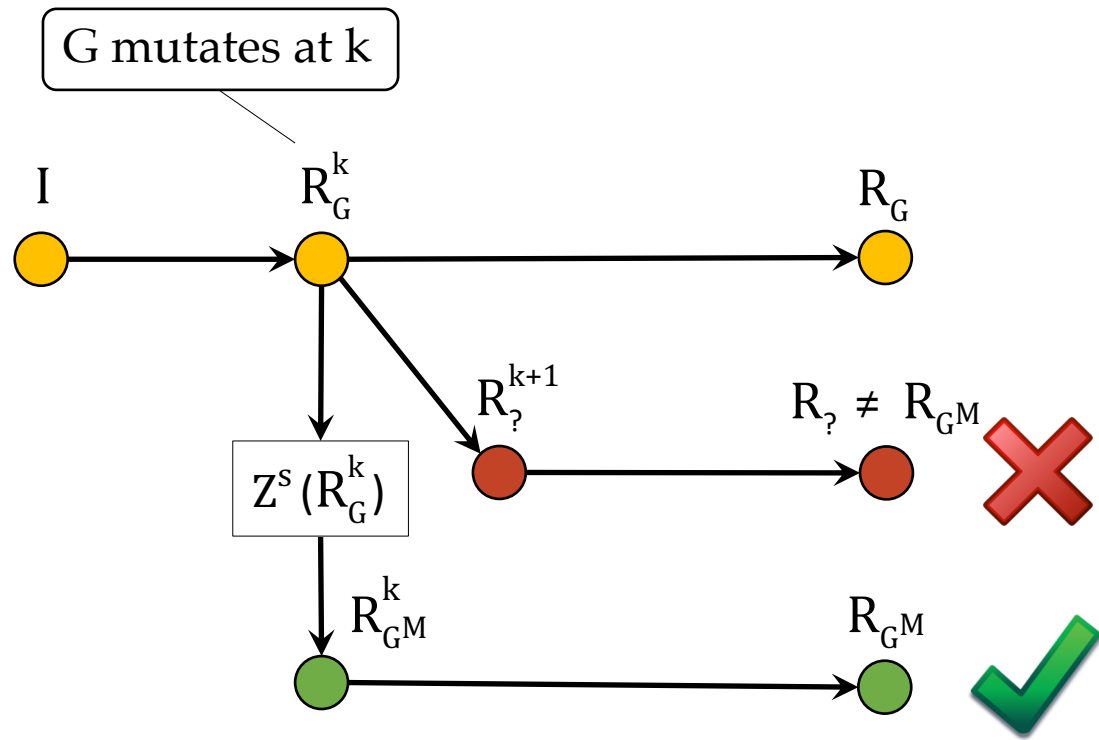
Streaming Graph Processing



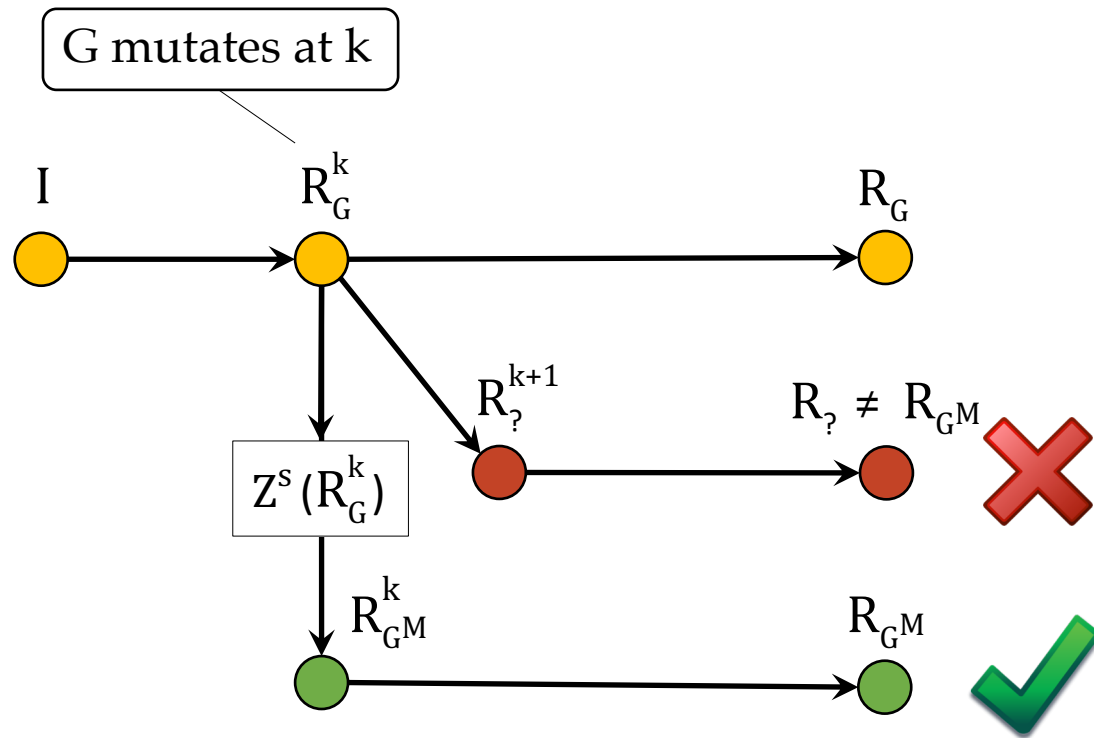
Streaming Graph Processing



Streaming Graph Processing

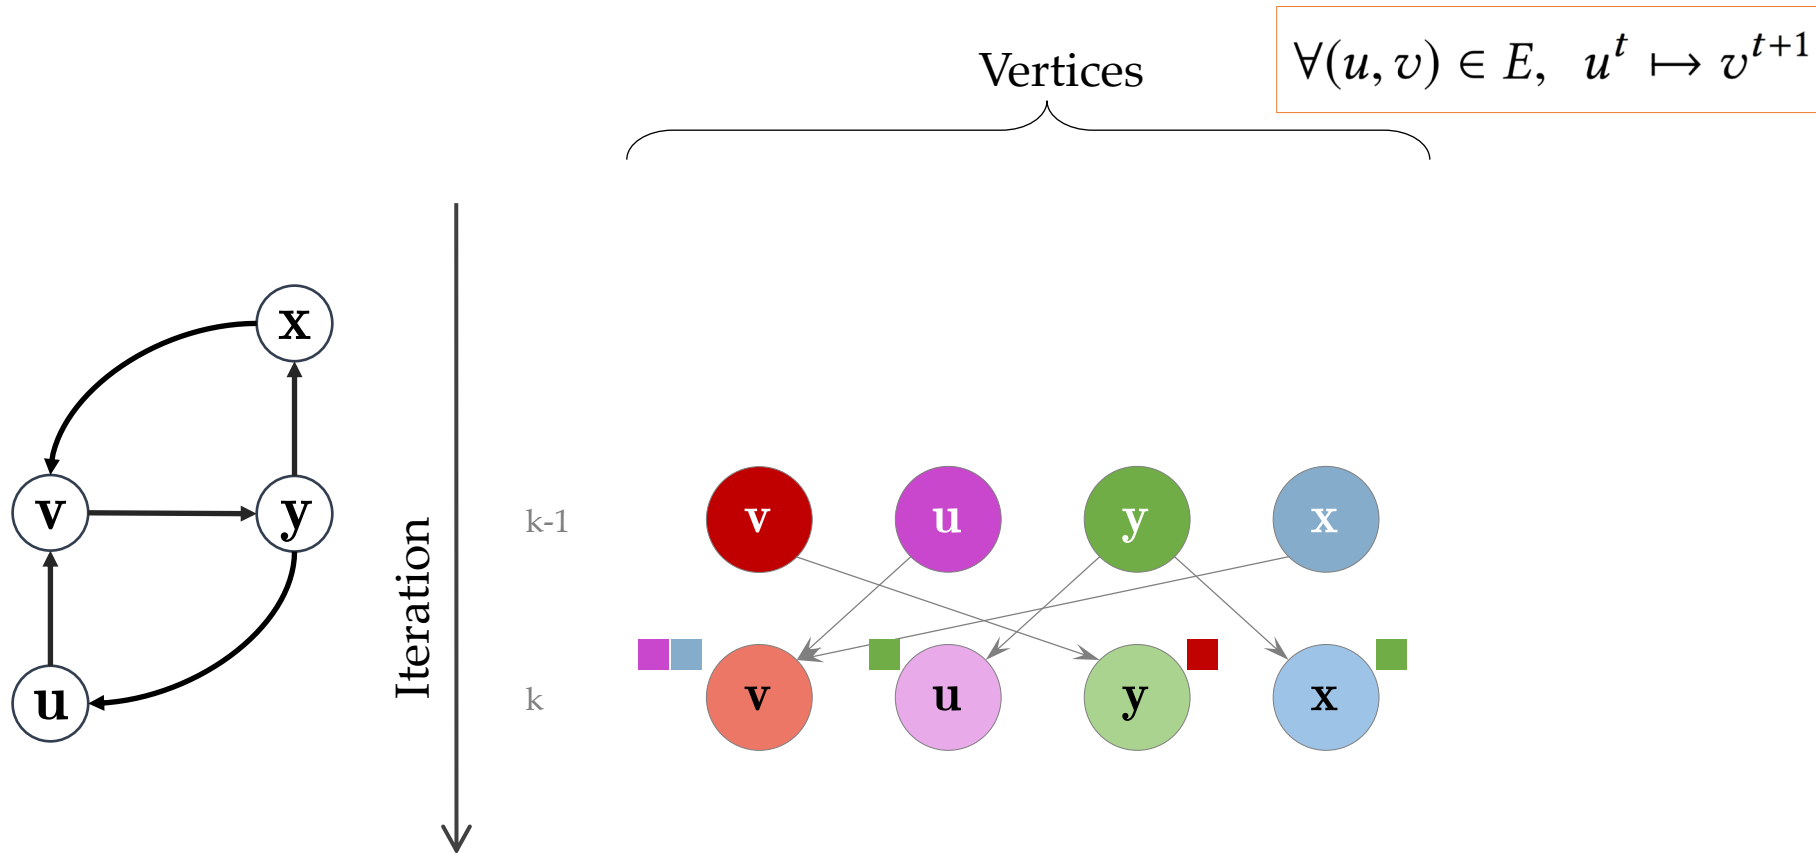


GraphBolt



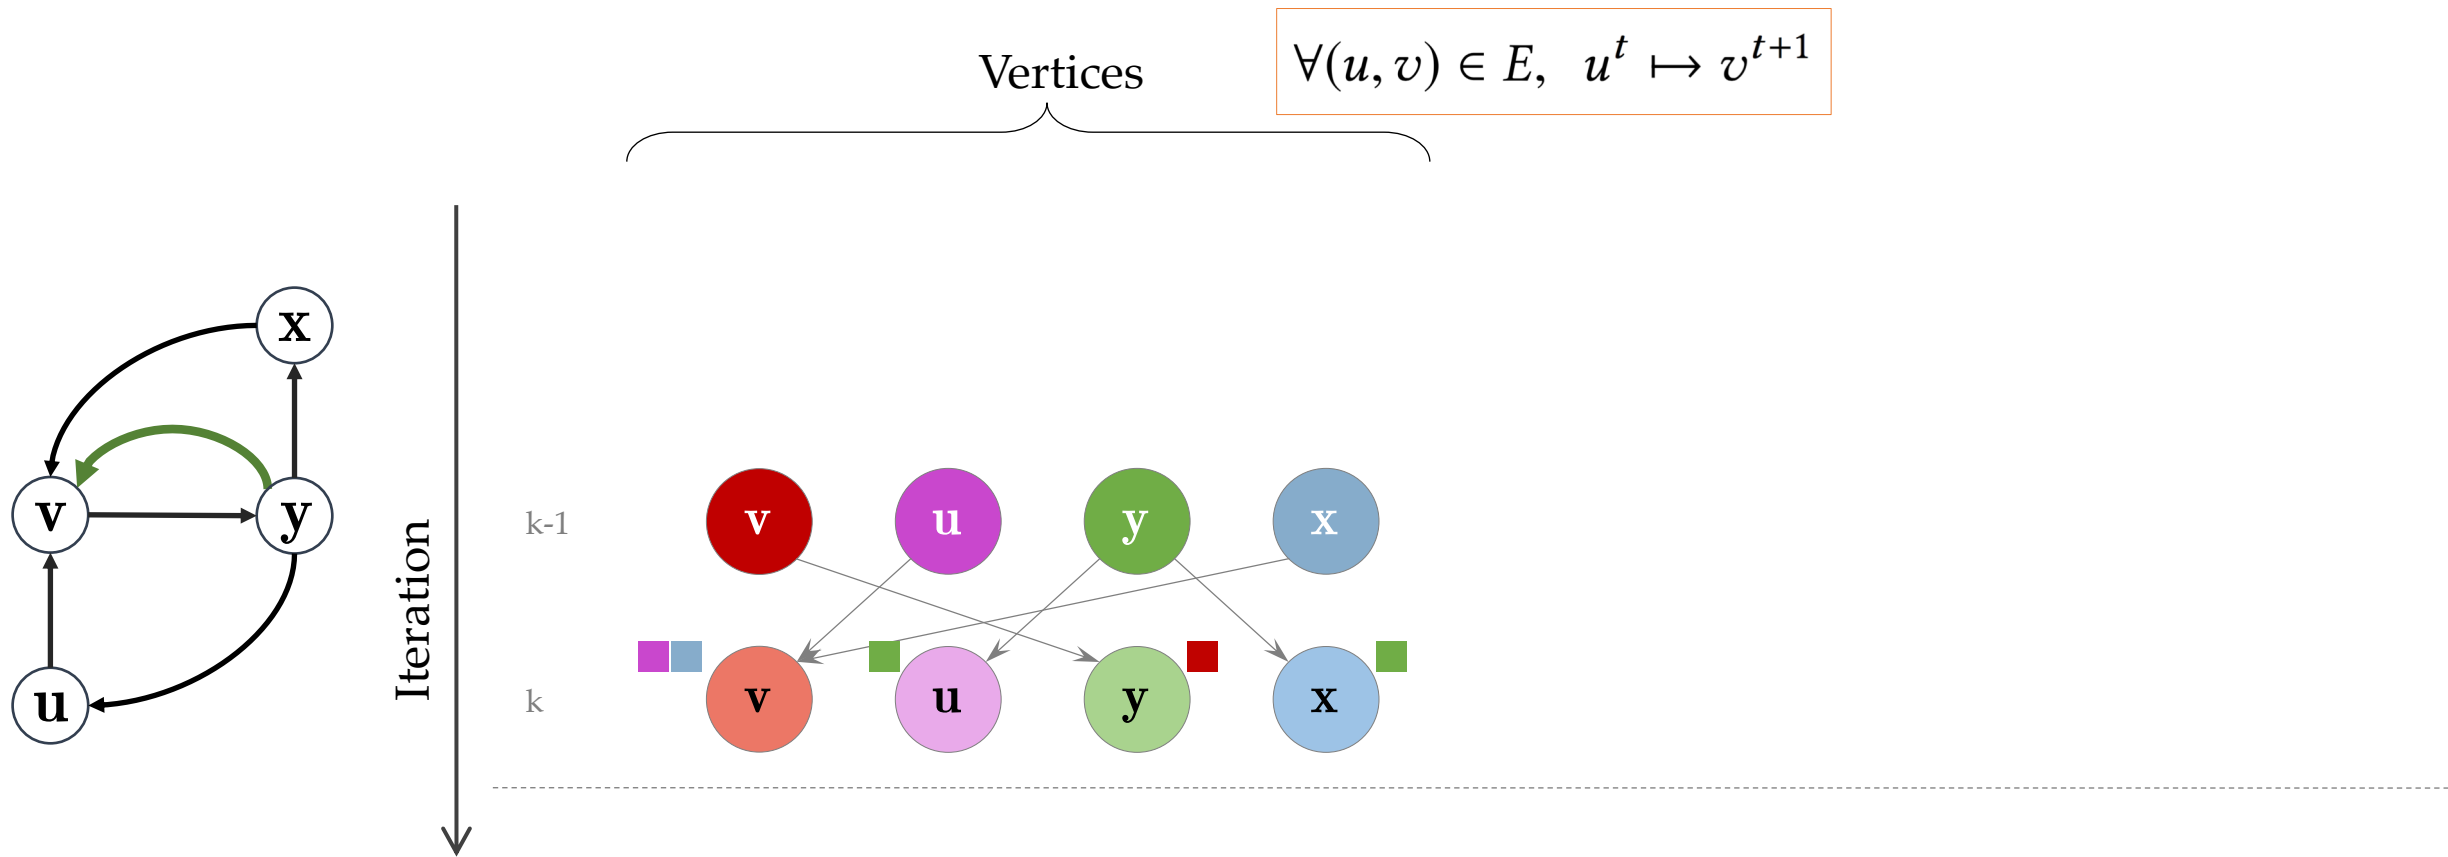
- **Dependency-Driven Incremental** Processing of Streaming Graphs
- Guarantee **Bulk Synchronous Parallel** Semantics
- Lightweight dependence tracking
- Dependency-aware value refinement upon graph mutation

Bulk Synchronous Processing (BSP)

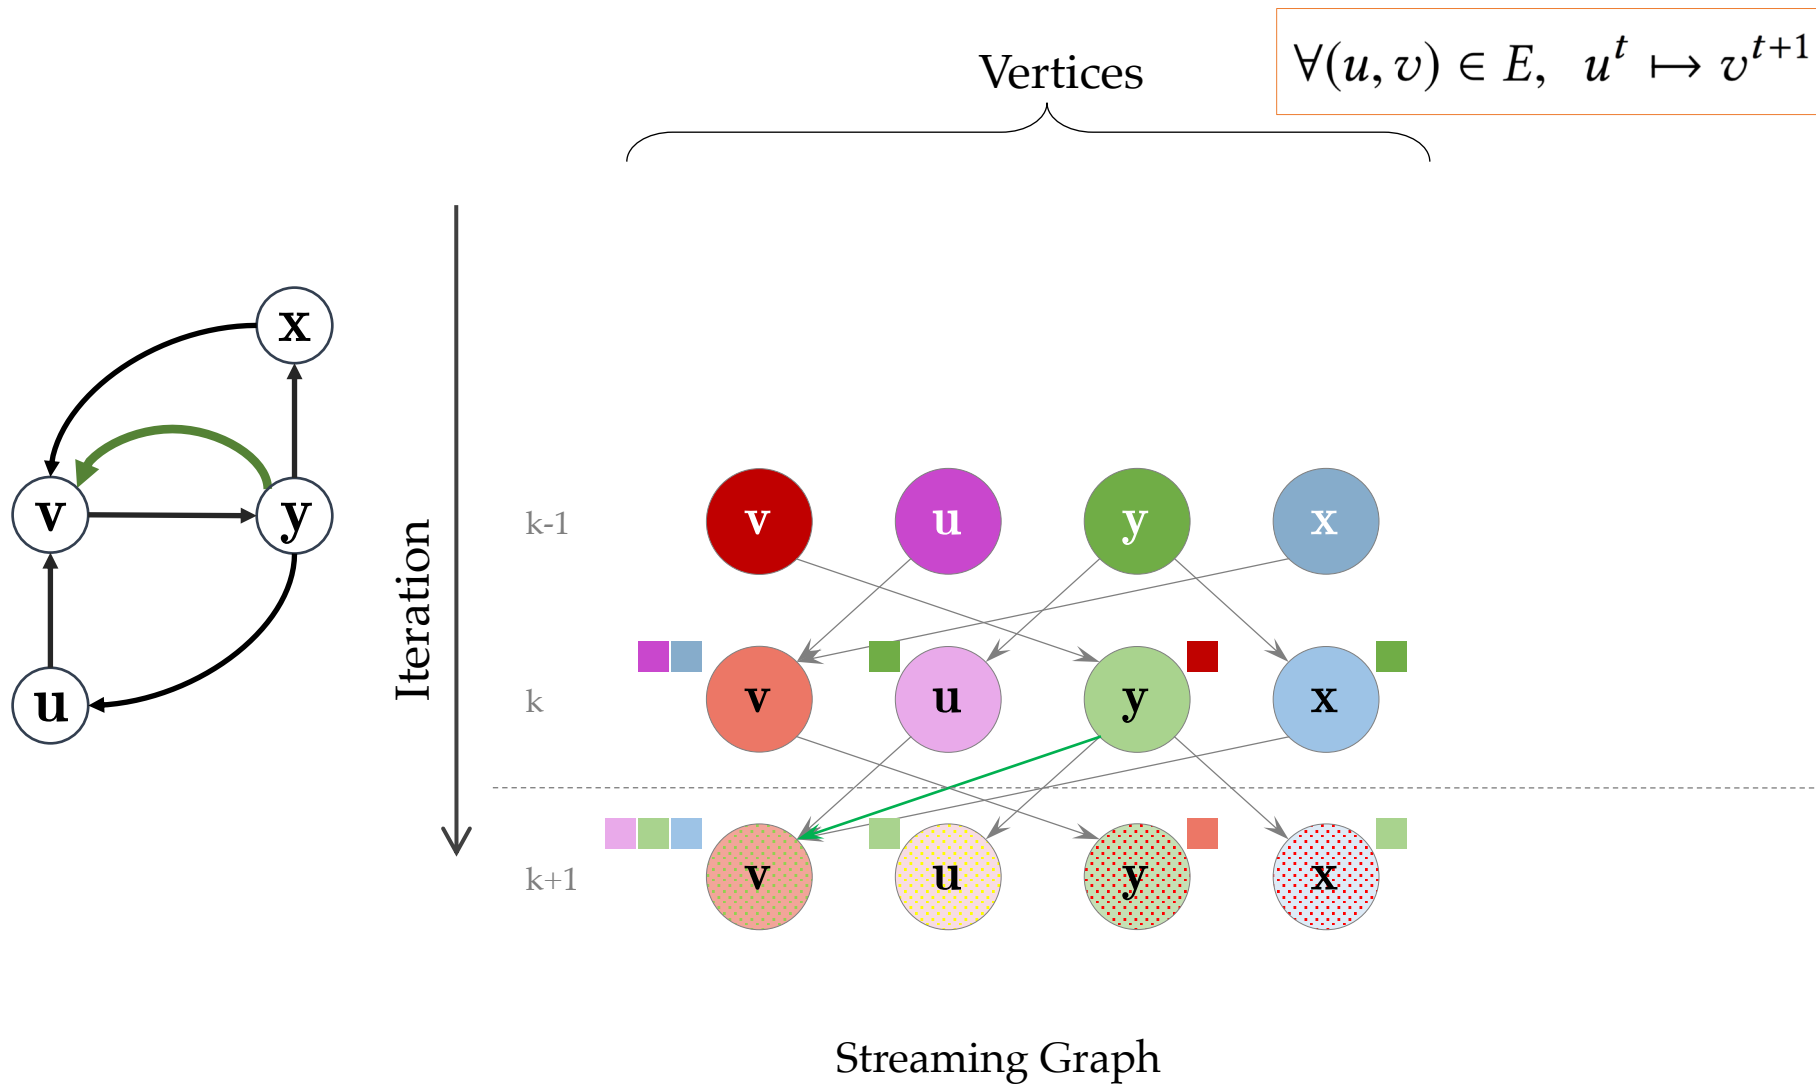


Streaming Graph


Upon Edge Addition




Upon Edge Addition




Upon Edge Addition




Upon Edge Addition




Upon Edge Addition




Upon Edge Addition




Upon Edge Addition



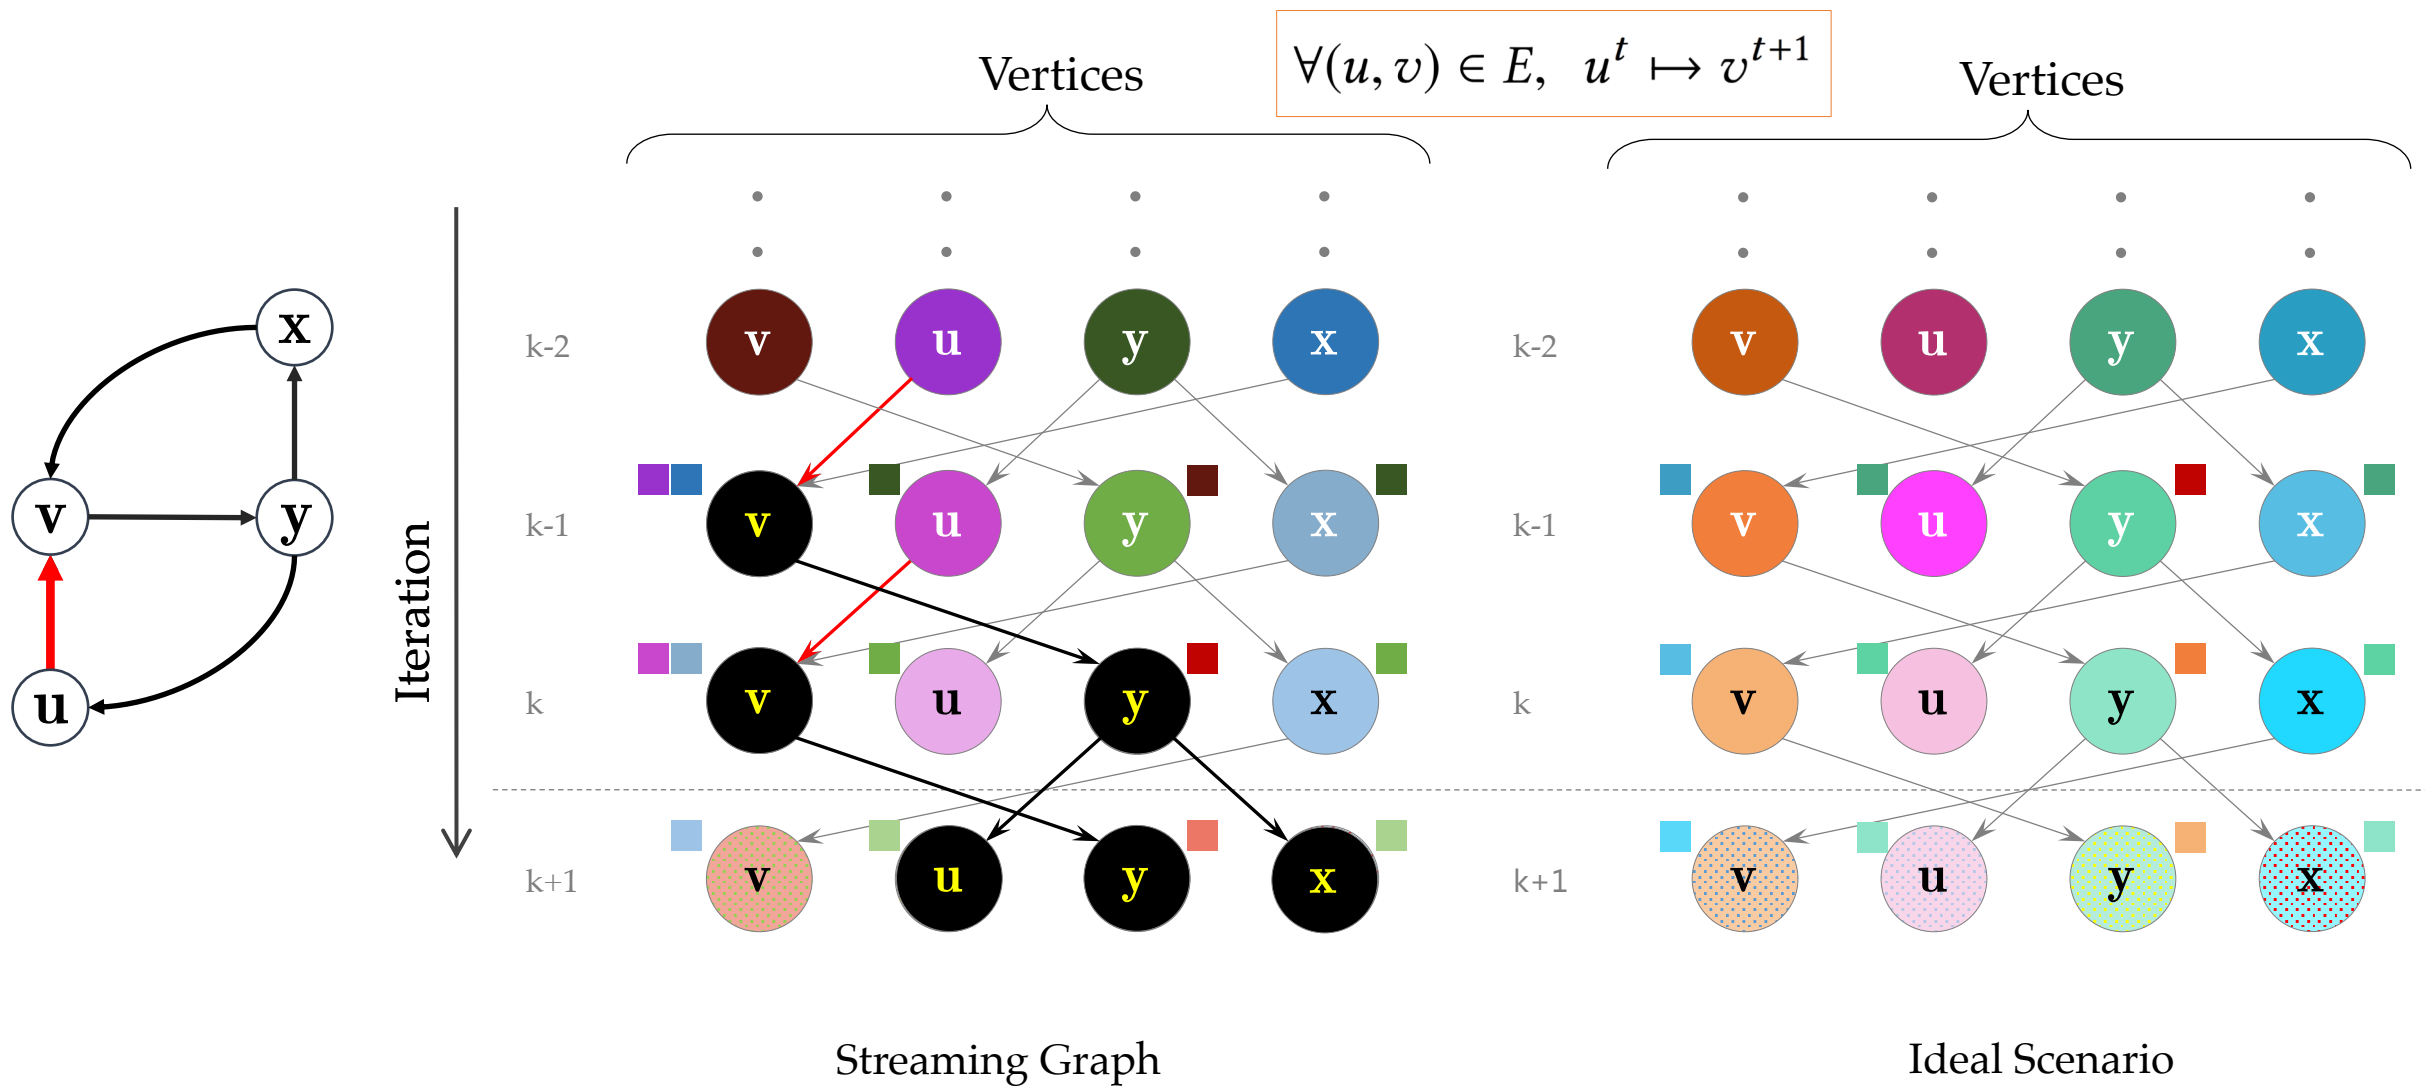
Upon Edge Addition




Upon Edge Addition



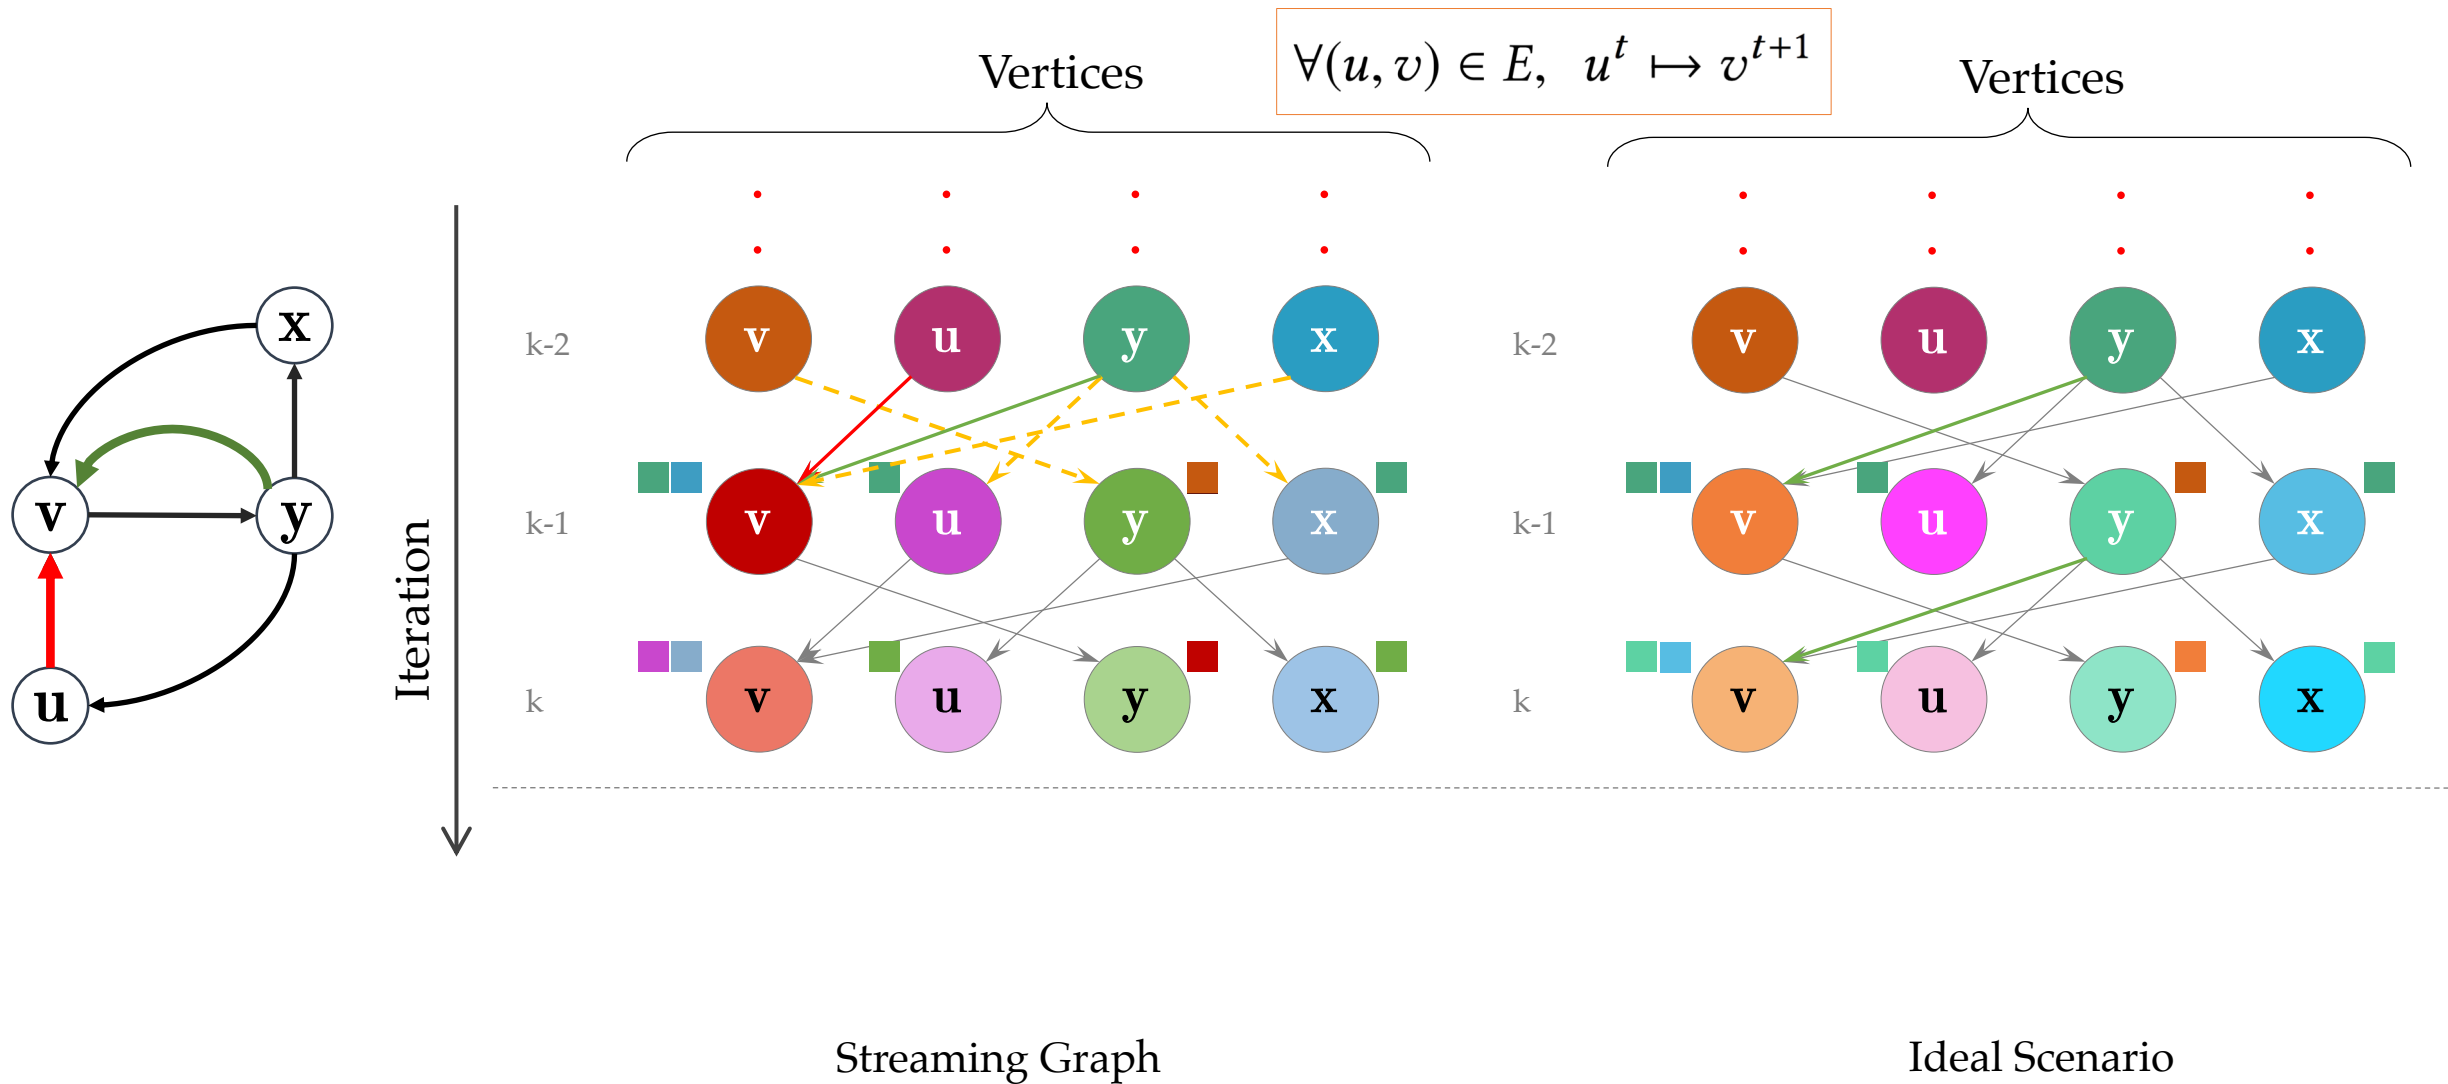
Upon Edge Deletion




GraphBolt: Correcting Dependencies




GraphBolt: Correcting Dependencies




GraphBolt: Correcting Dependencies




GraphBolt: Correcting Dependencies




GraphBolt: Correcting Dependencies




GraphBolt: Correcting Dependencies



GraphBolt: Dependency Tracking



Iteration
 \downarrow
 k-1
 k




$$v = \phi(\text{[purple][blue]})$$


$$c_i(v) = \phi \left(\bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u)) \right)$$

$$g_i(v) = \bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u))$$

GraphBolt: Dependency Tracking



Iteration
↓
k-1
k



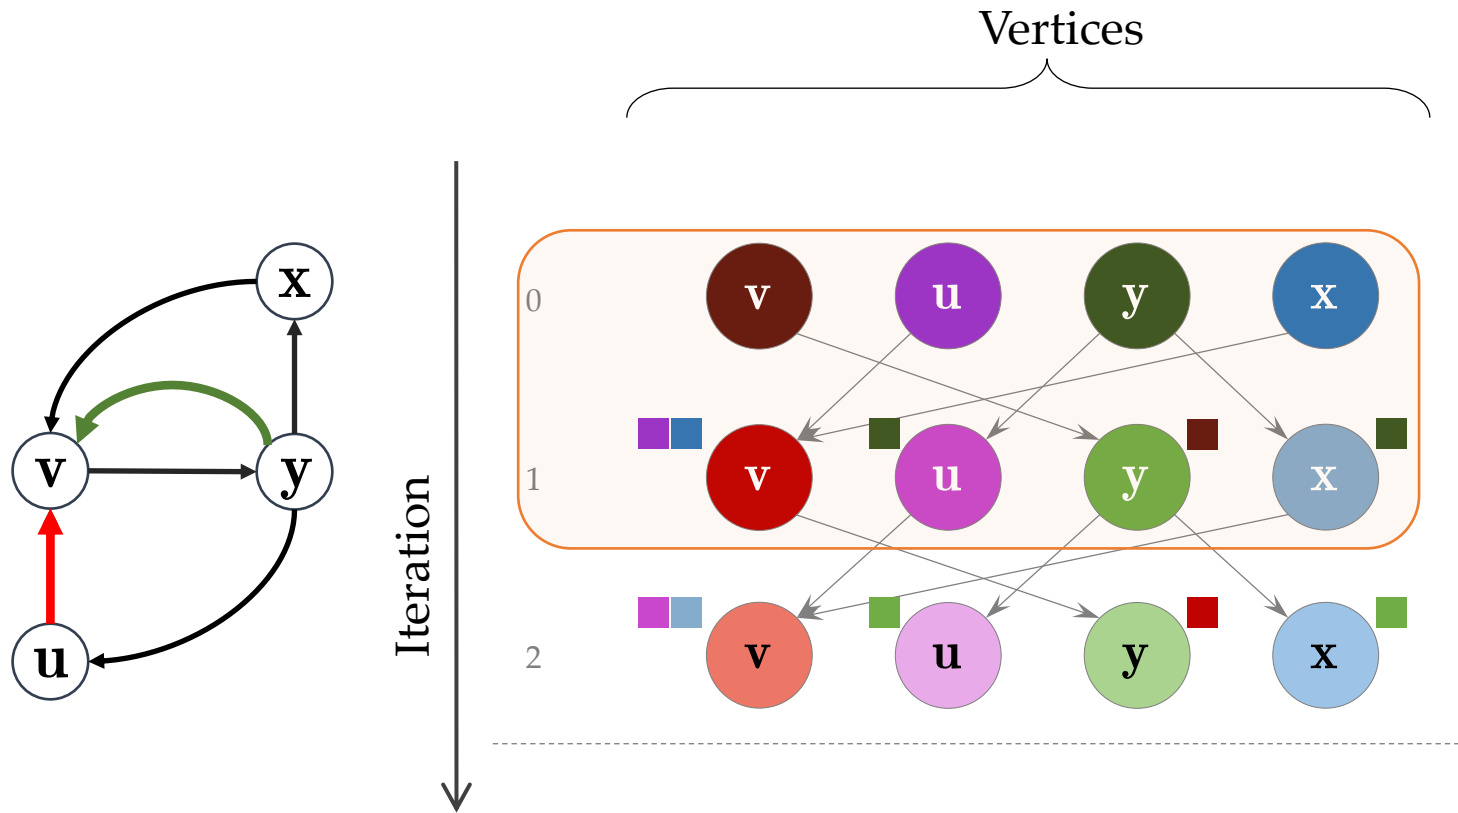
$$v = \phi(\text{purple bar, blue bar})$$

$$c_i(v) = \phi \left(\bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u)) \right)$$

$$g_i(v) = \bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u))$$

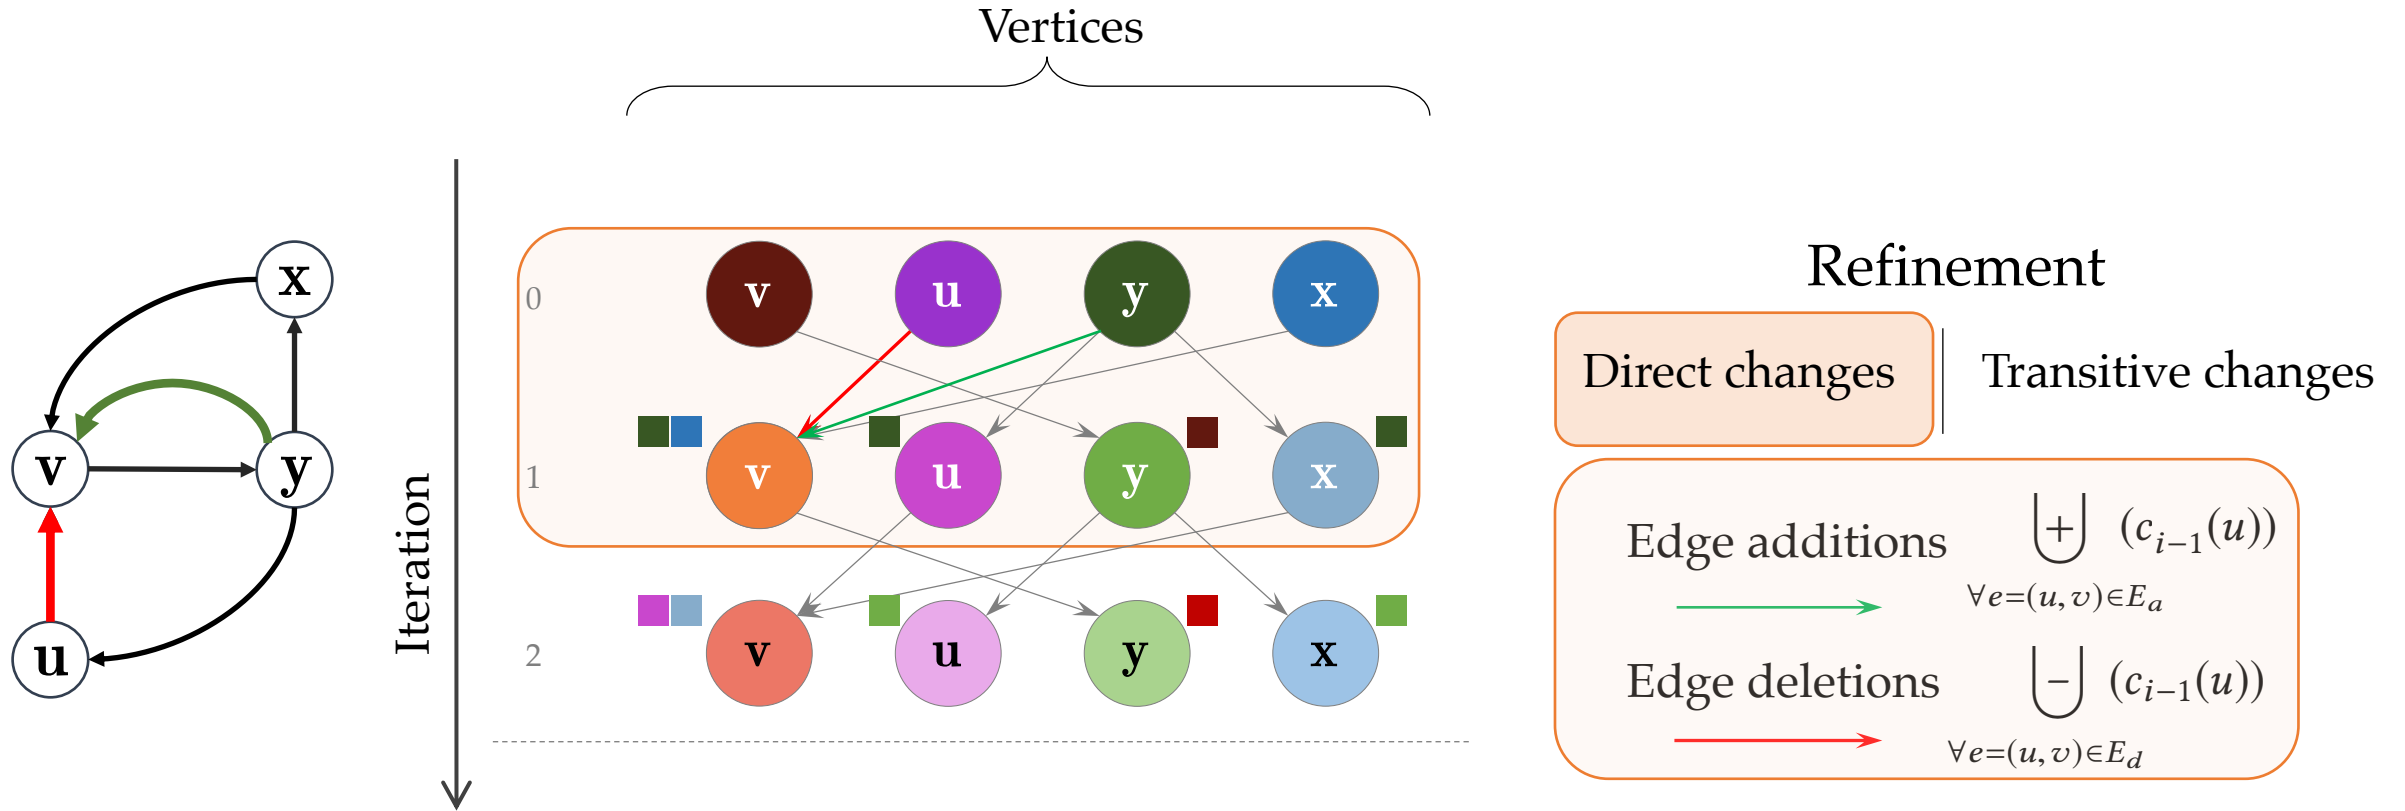
- Dependency relations translate across aggregation points
- Structure of dependencies inferred from input graphs

GraphBolt: Incremental Refinement




$$g_i^T(v) = g_i(v) + ?$$

GraphBolt: Incremental Refinement



$$g_i^T(v) = g_i(v) \bigcup_{\forall e=(u,v) \in E_a} (c_{i-1}(u)) \bigcup_{\forall e=(u,v) \in E_d} (c_{i-1}(u)) \quad + ?$$

GraphBolt: Incremental Refinement

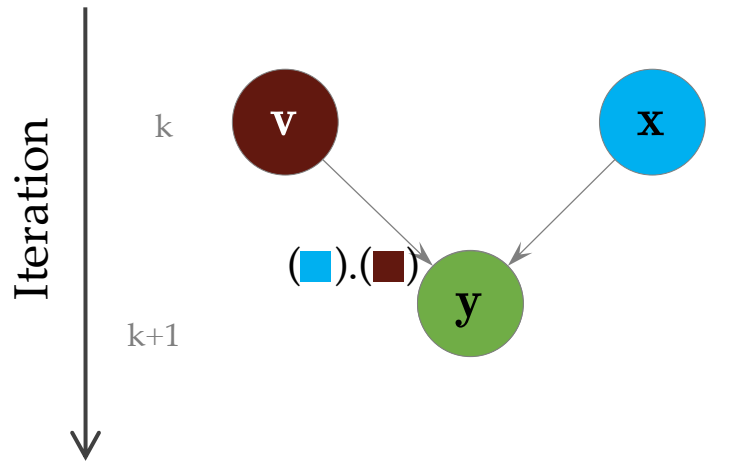


$$g_i^T(v) = g_i(v) \quad \bigcup_{\forall e=(u,v) \in E_a} (c_{i-1}(u)) \quad \bigcup_{\forall e=(u,v) \in E_d} (c_{i-1}(u)) \quad \bigcup_{\forall e=(u,v) \in E^T} (c_{i-1}^T(u))$$

$$s.t. c_{i-1}(u) \neq c_{i-1}^T(u)$$

Refinement: Transitive Changes

Vertex aggregation

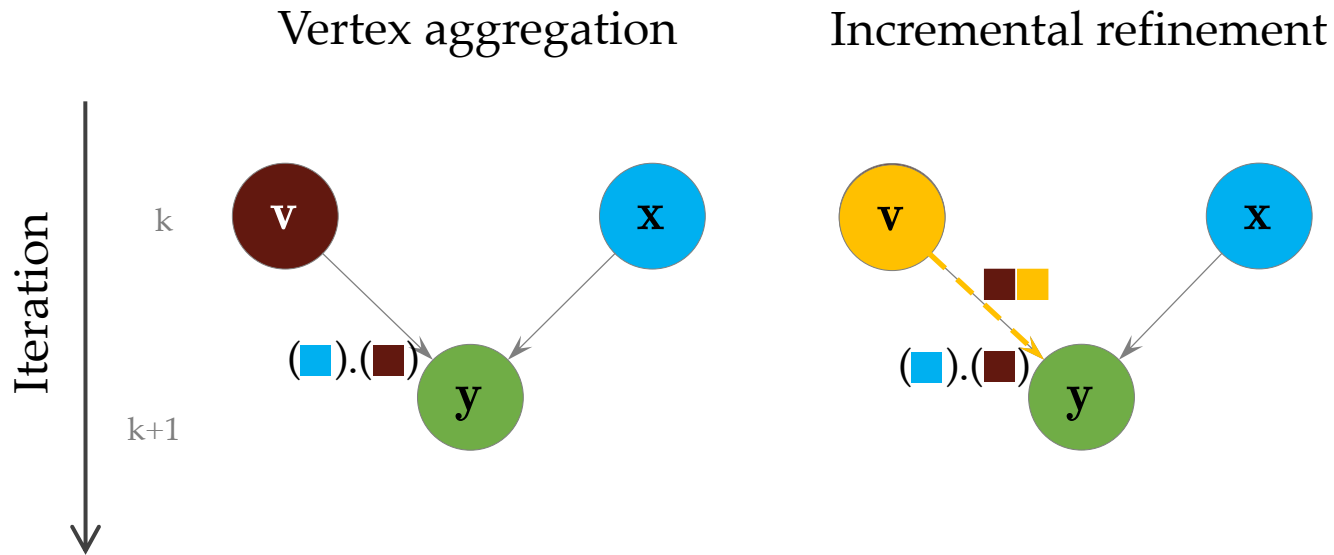


(■) - Transform

(■).(■) - Combine

$$g_i(v) = \bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u))$$

Refinement: Transitive Changes



(■) - Transform
 (■).(■) - Combine

$$g_i(v) = \bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u))$$

Aggregation

- Complex
- Retract : Old value
 - Propagate : New value

Belief Propagation

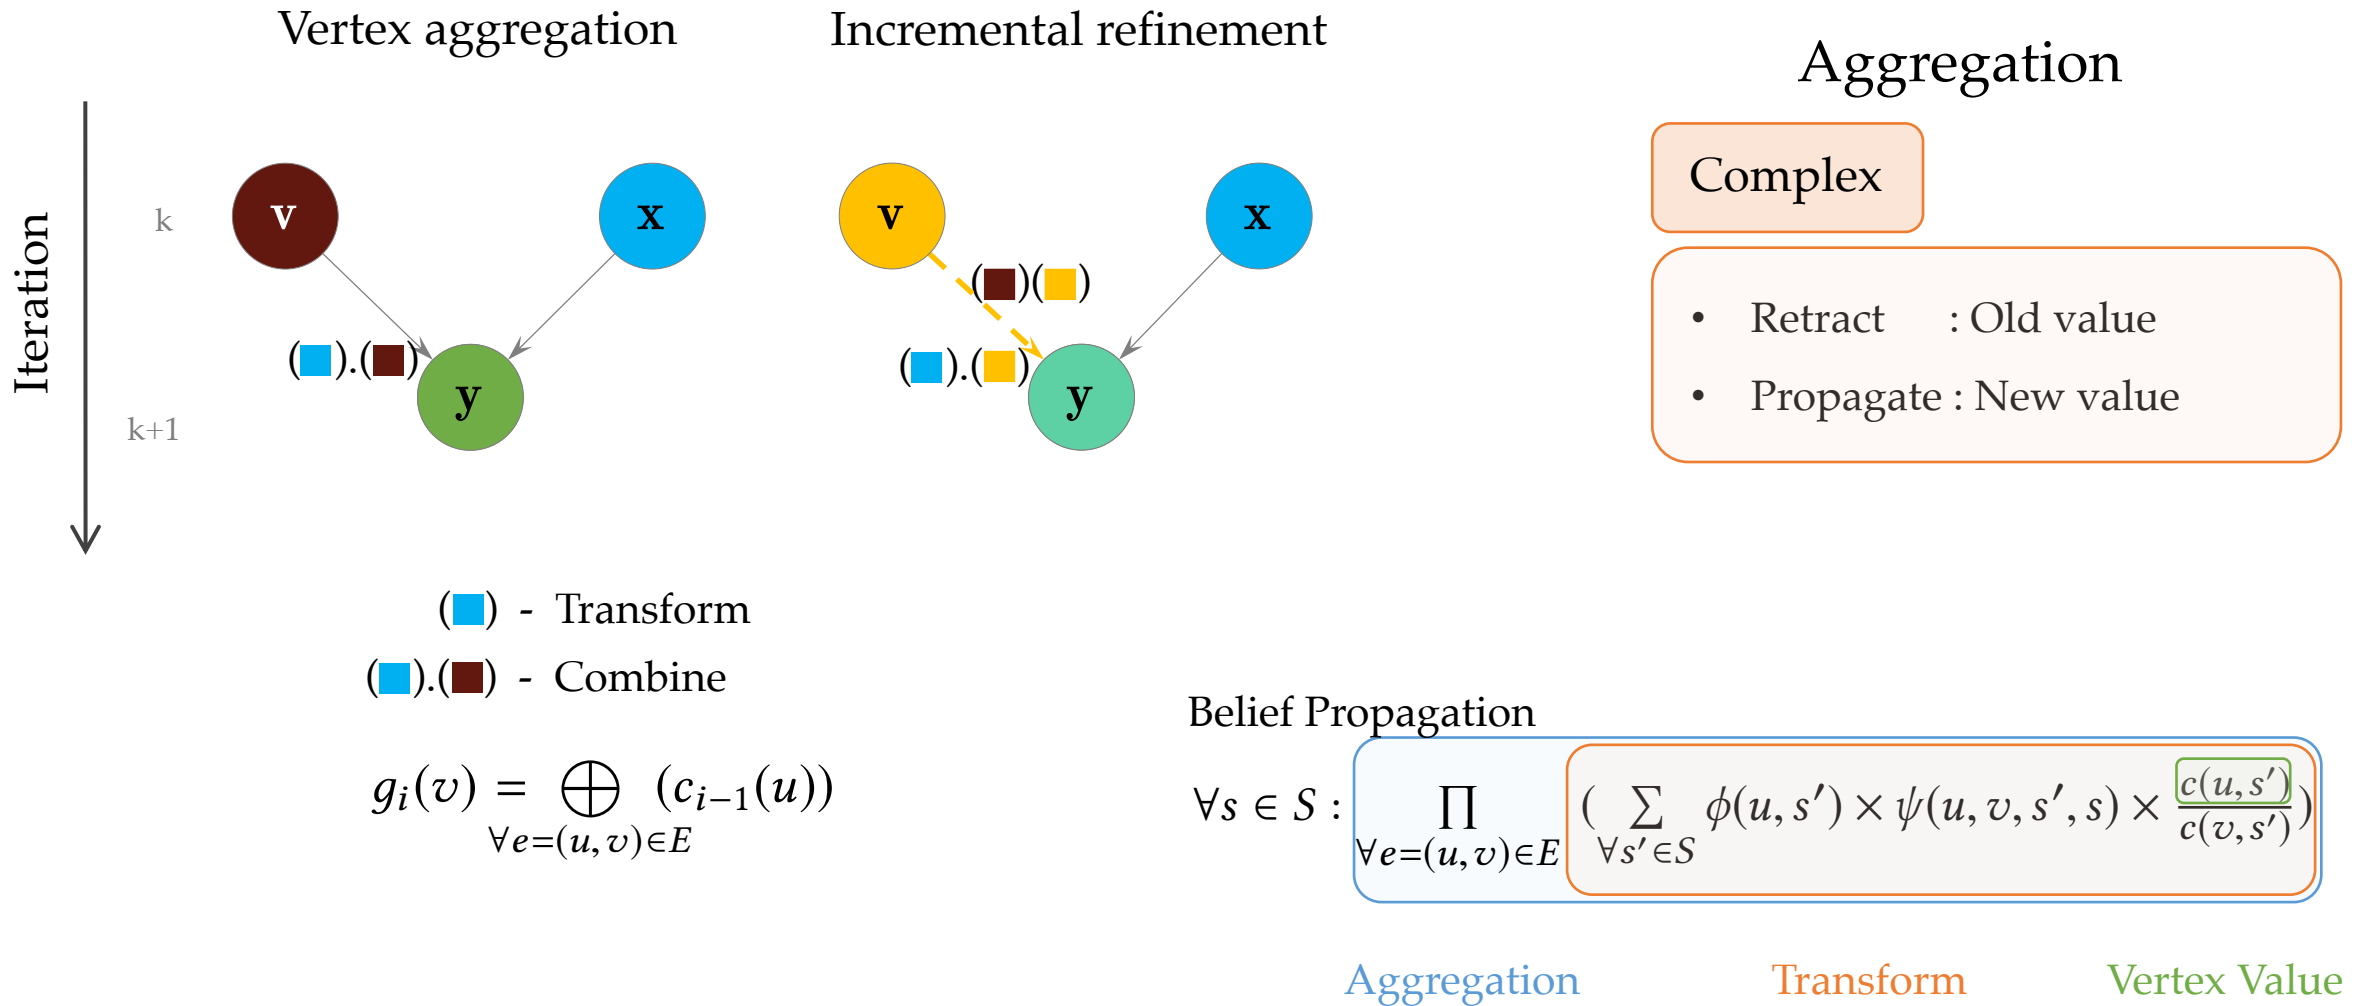
$$\forall s \in S : \prod_{\forall e=(u,v) \in E} \left(\sum_{\forall s' \in S} \phi(u, s') \times \psi(u, v, s', s) \times \frac{c(u, s')}{c(v, s')} \right)$$

Aggregation

Transform

Vertex Value

Refinement: Transitive Changes



(■) - Transform

(■).(■) - Combine

$$g_i(v) = \bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u))$$

Belief Propagation

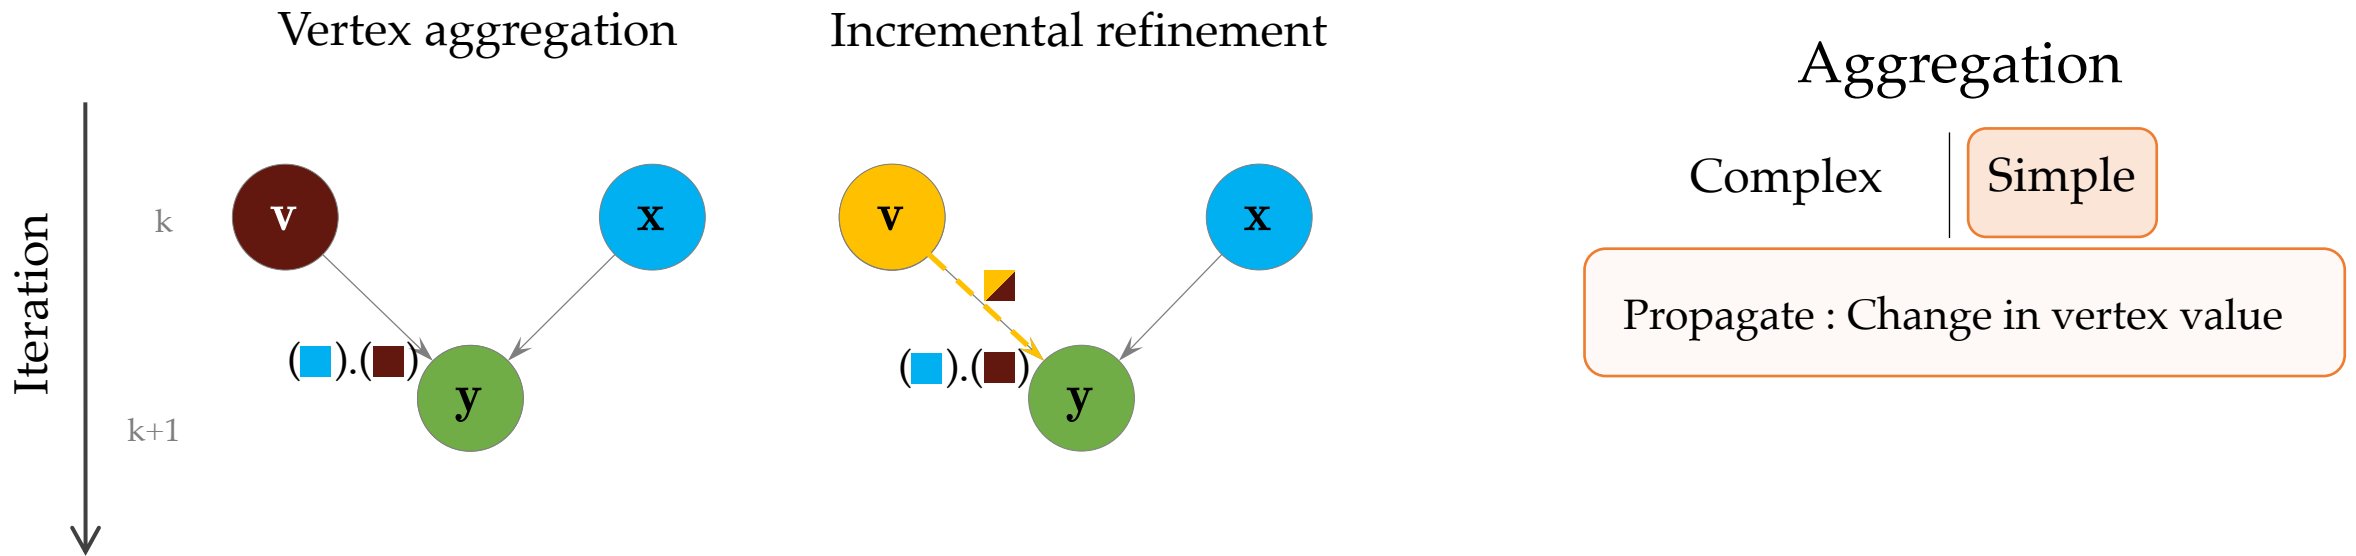
$$\forall s \in S : \prod_{\forall e=(u,v) \in E} \left(\sum_{\forall s' \in S} \phi(u, s') \times \psi(u, v, s', s) \times \frac{c(u, s')}{c(v, s')} \right)$$

Aggregation

Transform

Vertex Value

Refinement: Aggregation Types



(■) - Transform

(■).(■) - Combine

$$g_i(v) = \bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u))$$

Aggregation

Complex

Simple

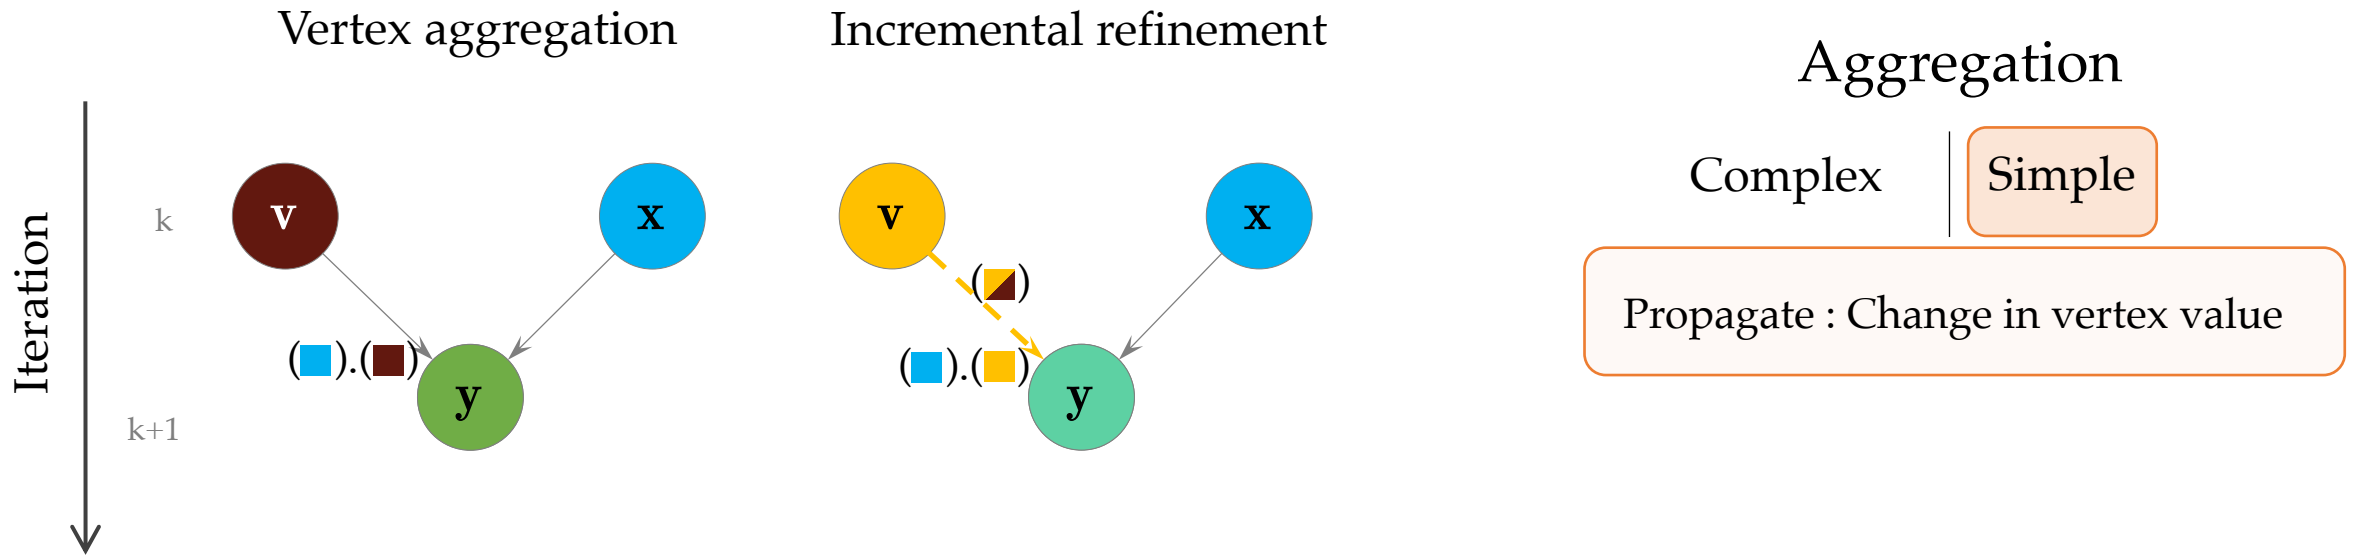
Propagate : Change in vertex value

PageRank

$$\sum_{\forall e=(u,v) \in E} \frac{c(u)}{out_degree(u)}$$

Aggregation Transform Vertex Value

Refinement: Aggregation Types



(■) - Transform

(■).(■) - Combine

$$g_i(v) = \bigoplus_{\forall e=(u,v) \in E} (c_{i-1}(u))$$

PageRank

$$\sum_{\forall e=(u,v) \in E} \frac{c(u)}{out_degree(u)}$$

Aggregation Transform Vertex Value

GraphBolt: Programming Model

```
function REPROPAGATE( $e = (u, v), i$ )  
    ATOMICADD(&sum[v][i + 1],  $\frac{oldpr[u][i]}{old\_degree[u]}$ )  
end function
```

```
function RETRACT( $e = (u, v), i$ )  
    ATOMICSUB(&sum[v][i + 1],  $\frac{oldpr[u][i]}{old\_degree[u]}$ )  
end function
```

```
function PROPAGATE( $e = (u, v), i$ )  
    ATOMICADD(&sum[v][i + 1],  $\frac{newpr[u][i]}{new\_degree[u]}$ )  
end function
```

```
function PAGERANK()
```

Direct changes

```
    for  $i \in [0...k]$  do
```

```
        EDMAP( $E\_add$ , REPROPAGATE,  $i$ )
```

```
        EDMAP( $E\_delete$ , RETRACT,  $i$ )
```

```
    end for
```

```
     $V\_updated = GETSOURCES(E\_add \cup E\_delete)$ 
```

```
     $V\_change = GETTARGETS(E\_add \cup E\_delete)$ 
```

```
    for  $i \in [0...k]$  do
```

```
         $E\_update = \{(u, v) : u \in V\_updated\}$ 
```

```
        EDMAP( $E\_update$ , RETRACT,  $i$ )
```

```
        EDMAP( $E\_update$ , PROPAGATE,  $i$ )
```

```
         $V\_dest = GETTARGETS(E\_update)$ 
```

```
         $V\_change = V\_change \cup V\_dest$ 
```

```
         $V\_updated = VERTEXMAP(V\_change, COMPUTE, i)$ 
```

```
    end for
```

```
end function
```

GraphBolt: Programming Model - Complex aggregations

```
function REPROPAGATE( $e = (u, v)$ ,  $i$ )  
    ATOMICADD(&sum[v][i + 1],  $\frac{oldpr[u][i]}{old\_degree[u]}$ )  
end function
```

```
function RETRACT( $e = (u, v)$ ,  $i$ )  
    ATOMICSUB(&sum[v][i + 1],  $\frac{oldpr[u][i]}{old\_degree[u]}$ )  
end function
```

```
function PROPAGATE( $e = (u, v)$ ,  $i$ )  
    ATOMICADD(&sum[v][i + 1],  $\frac{newpr[u][i]}{new\_degree[u]}$ )  
end function
```

```
function PAGERANK()  
    for  $i \in [0...k]$  do  
        EDGEMAP( $E\_add$ , REPROPAGATE,  $i$ )  
        EDGEMAP( $E\_delete$ , RETRACT,  $i$ )  
    end for  
     $V\_updated = GETSOURCES(E\_add \cup E\_delete)$   
     $V\_change = GETTARGETS(E\_add \cup E\_delete)$   
    for  $i \in [0...k]$  do  
         $E\_update = \{(u, v) : u \in V\_updated\}$  Transitive changes  
        EDGEMAP( $E\_update$ , RETRACT,  $i$ )  
        EDGEMAP( $E\_update$ , PROPAGATE,  $i$ )  
         $V\_dest = GETTARGETS(E\_update)$   
         $V\_change = V\_change \cup V\_dest$   
         $V\_updated = VERTEXMAP(V\_change, COMPUTE, i)$   
    end for  
end function
```


GraphBolt: Programming Model - Simple aggregations

```
function REPROPAGATE( $e = (u, v)$ ,  $i$ )  
    ATOMICADD(&sum[v][i + 1],  $\frac{oldpr[u][i]}{old\_degree[u]}$ )  
end function
```

```
function RETRACT( $e = (u, v)$ ,  $i$ )  
    ATOMICSUB(&sum[v][i + 1],  $\frac{oldpr[u][i]}{old\_degree[u]}$ )  
end function
```

```
function PROPAGATEDELTA( $e = (u, v)$ ,  $i$ )  
    ATOMICADD(&sum[v][i + 1],  $\frac{newpr[u][i]}{new\_degree[u]} - \frac{oldpr[u][i]}{old\_degree[u]}$ )  
end function
```

```
function PAGERANK()  
    for  $i \in [0 \dots k]$  do  
        EDGEMAP( $E\_add$ , REPROPAGATE,  $i$ )  
        EDGEMAP( $E\_delete$ , RETRACT,  $i$ )  
    end for  
     $V\_updated = GETSOURCES(E\_add \cup E\_delete)$   
     $V\_change = GETTARGETS(E\_add \cup E\_delete)$   
    for  $i \in [0 \dots k]$  do  
         $E\_update = \{(u, v) : u \in V\_updated\}$  Transitive changes  
        EDGEMAP( $E\_update$ , PROPAGATEDELTA,  $i$ )  
         $V\_dest = GETTARGETS(E\_update)$   
         $V\_change = V\_change \cup V\_dest$   
         $V\_updated = VERTEXMAP(V\_change, COMPUTE, i)$   
    end for  
end function
```

GraphBolt Details ...

- **Aggregation properties** & **non-decomposable** aggregations
- **Pruning dependencies** for light-weight tracking
 - Vertical pruning
 - Horizontal pruning
- **Hybrid incremental** execution
 - With & without dependency information
- Graph data structure & parallelization model








Experimental Setup

Algorithm	Aggregation (\oplus)
PageRank (PR)	$\sum_{\forall e=(u,v) \in E} \frac{c(u)}{out_degree(u)}$
Belief Propagation (BP)	$\forall s \in S : \prod_{\forall e=(u,v) \in E} (\sum_{\forall s' \in S} \phi(u, s') \times \psi(u, v, s', s) \times c(u, s'))$
Label Propagation (LP)	$\forall f \in F : \sum_{\forall e=(u,v) \in E} c(u, f) \times weight(u, v)$
Co-Training Expectation Maximization (CoEM)	$\sum_{\forall e=(u,v) \in E} \frac{c(u) \times weight(u,v)}{\sum_{\forall e=(w,v) \in E} weight(w,v)}$
Collaborative Filtering (CF)	$\langle \sum_{\forall e=(u,v) \in E} c_i(u) \cdot c_i(u)^{tr}, \sum_{\forall e=(u,v) \in E} c_i(u) \cdot weight(u, v) \rangle$
Triangle Counting (TC)	$\sum_{\forall e=(u,v) \in E} in_neighbors(u) \cap out_neighbors(v) $

Graphs	Edges	Vertices
Wiki (WK) [47]	378M	12M
UKDomain (UK) [7]	1.0B	39.5M
Twitter (TW) [21]	1.5B	41.7M
TwitterMPI (TT) [8]	2.0B	52.6M
Friendster (FT) [14]	2.5B	68.3M
Yahoo (YH) [49]	6.6B	1.4B






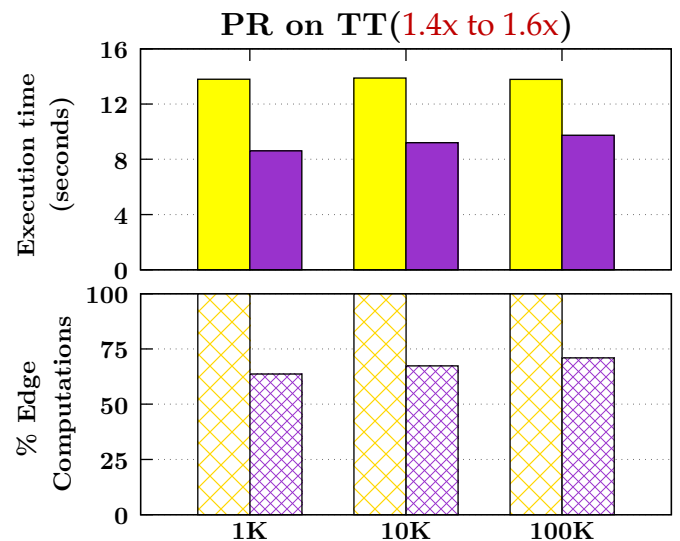
Server: 32-core / 2 GHz / 231 GB

GraphBolt Performance






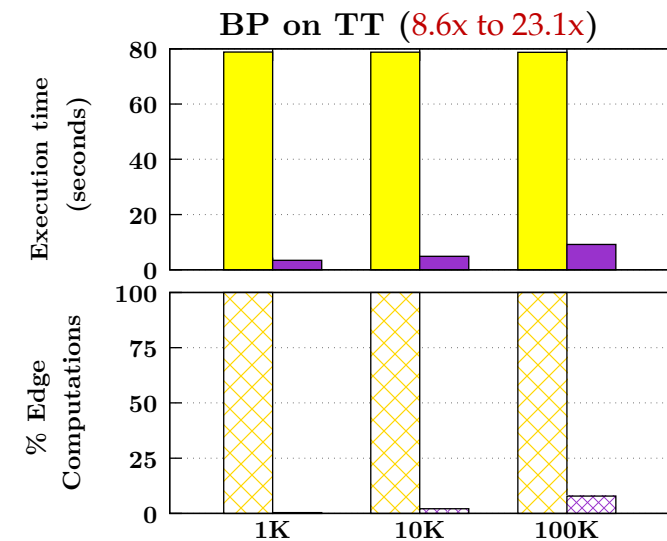


GraphBolt Performance

GraphBolt-Reset  
 GraphBolt  







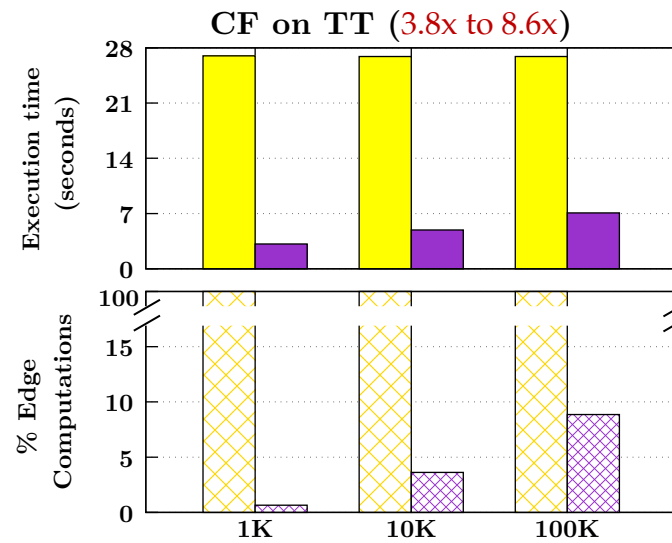

GraphBolt Performance

GraphBolt-Reset  
 GraphBolt  



GraphBolt Performance


GraphBolt-Reset  
 GraphBolt  




Detailed Experiments ...

- Varying workload type
 - Mutations affecting majority of the values v/s affecting small values
- Varying graph mutation rate
 - Single edge update (reactiveness) v/s 1 million edge updates (throughput)
- Scaling to large graph (Yahoo) over large system (96 core/748 GB)

GraphBolt v/s Differential Dataflow



PR on TT



PR on TT
(100 single edge mutations)

Summary

- Efficient **incremental processing** of **streaming graphs**
- Guarantee **Bulk Synchronous Parallel** semantics
- Lightweight **dependence tracking** at aggregation level
- **Dependency-aware value refinement** upon graph mutation
 - Programming model to support **incremental complex aggregation** types

Acknowledgements

