

# Executing Dynamic Data-Graph Computations Deterministically Using Chromatic Scheduling

Student presenter: Nikola Samardzic

# Motivation

- Data graph computation is useful
- Dynamic data graph computation is even more useful
  - Dynamic data graph := active set depends on data
- Parallelism is good
- Determinism is good
  - Debugging (and even performance!)
- We want it all!

# Key Contribution

- *Deterministic* execution of *dynamic* data graph computations

# Serial Baseline

SERIAL-DDGC( $G, f, Q_0$ )

```
1  for  $v \in Q_0$ 
2    ENQUEUE( $Q, v$ )
3   $r = 0$ 
4  ENQUEUE( $Q, \text{NIL}$ ) // Sentinel NIL denotes the end of a round.
5  while  $Q \neq \{\text{NIL}\}$ 
6     $v = \text{DEQUEUE}(Q)$ 
7    if  $v == \text{NIL}$ 
8       $r += 1$ 
9      ENQUEUE( $Q, \text{NIL}$ )
10   else
11      $S = f(v)$ 
12     for  $u \in S$ 
13       if  $u \notin Q$ 
14         ENQUEUE( $Q, u$ )
```

$f(v)$  returns the subset  
of  $v$ 's neighbors  
activated by the update

# Serial Baseline

SERIAL-DDGC( $G, f, Q_0$ )

```
1  for  $v \in Q_0$ 
2    ENQUEUE( $Q, v$ )
3   $r = 0$ 
4  ENQUEUE( $Q, \text{NIL}$ ) // Sentinel NIL denotes the end of a round.
5  while  $Q \neq \{\text{NIL}\}$ 
6     $v = \text{DEQUEUE}(Q)$ 
7    if  $v == \text{NIL}$ 
8       $r += 1$ 
9      ENQUEUE( $Q, \text{NIL}$ )
10   else
11      $S = f(v)$ 
12     for  $u \in S$ 
13       if  $u \notin Q$ 
14         ENQUEUE( $Q, u$ )
```

OK, so just make  $Q$  thread-safe, enqueue an entire round of active nodes, and run this for loop in parallel for each active node?

# Determinism Aids Parallelism

- Nondeterministic programs are a pain to debug
- Locks suck
  - Implies lock acquisition and contention overheads
- An alternative: chromatic scheduling
  - Gives us determinism
  - And (basically) no locks!

# Prism

PRISM( $G, f, Q_0$ )

```
1   $\chi = \text{COLOR-GRAPH}(G)$ 
2   $r = 0$ 
3   $Q = Q_0$ 
4  while  $Q \neq \emptyset$ 
5       $C = \text{MB-COLLECT}(Q)$ 
6      for  $C \in \mathcal{C}$ 
7          parallel for  $v \in C$ 
8               $\text{active}[v] = \text{FALSE}$ 
9               $S = f(v)$ 
10             parallel for  $u \in S$ 
11                 if  $\text{CAS}(\text{active}[u], \text{FALSE}, \text{TRUE})$ 
12                      $\text{MB-INSERT}(Q, u, \text{color}[u])$ 
13              $r = r + 1$ 
```

# Prism

PRISM( $G, f, Q_0$ )

```
1  $\chi = \text{COLOR-GRAPH}(G)$ 
2  $r = 0$ 
3  $Q = Q_0$ 
4 while  $Q \neq \emptyset$ 
5    $C = \text{MB-COLLECT}(Q)$ 
6   for  $C \in \mathcal{C}$ 
7     parallel for  $v \in C$ 
8        $\text{active}[v] = \text{FALSE}$ 
9        $S = f(v)$ 
10      parallel for  $u \in S$ 
11        if  $\text{CAS}(\text{active}[u], \text{FALSE}, \text{TRUE})$ 
12           $\text{MB-INSERT}(Q, u, \text{color}[u])$ 
13       $r = r + 1$ 
```

Coloring need not be perfect; ~solved problem



# Prism

PRISM( $G, f, Q_0$ )

```
1   $\chi = \text{COLOR-GRAPH}(G)$ 
2   $r = 0$ 
3   $Q = Q_0$ 
4  while  $Q \neq \emptyset$ 
5       $C = \text{MB-COLLECT}(Q)$ 
6      for  $C \in \mathcal{C}$ 
7          parallel for  $v \in C$ 
8               $active[v] = \text{FALSE}$ 
9               $S = f(v)$ 
10             parallel for  $u \in S$ 
11                 if  $\text{CAS}(active[u], \text{FALSE}, \text{TRUE})$ 
12                      $\text{MB-INSERT}(Q, u, color[u])$ 
13              $r = r + 1$ 
```

Multibag  
funny  
business!

# Prism

PRISM( $G, f, Q_0$ )

```
1   $\chi = \text{COLOR-GRAPH}(G)$ 
2   $r = 0$ 
3   $Q = Q_0$ 
4  while  $Q \neq \emptyset$ 
5       $C = \text{MB-COLLECT}(Q)$ 
6      for  $C \in \mathcal{C}$ 
7          parallel for  $v \in C$ 
8               $\text{active}[v] = \text{FALSE}$ 
9               $S = f(v)$ 
10             parallel for  $u \in S$ 
11                 if  $\text{CAS}(\text{active}[u], \text{FALSE}, \text{TRUE})$ 
12                      $\text{MB-INSERT}(Q, u, \text{color}[u])$ 
13              $r = r + 1$ 
```

MB-Collect: Empty all bags from multibag  $Q$  into a collection  $C$  that is easy to iterate over

MB-Insert: Insert node  $u$  in bag  $\text{color}[u]$  of multibag  $Q$

# Prism

PRISM( $G, f, Q_0$ )

```
1   $\chi = \text{COLOR-GRAPH}(G)$ 
2   $r = 0$ 
3   $Q = Q_0$ 
4  while  $Q \neq \emptyset$ 
5       $C = \text{MB-COLLECT}(Q)$ 
6      for  $C \in \mathcal{C}$ 
7          parallel for  $v \in C$ 
8               $active[v] = \text{FALSE}$ 
9               $S = f(v)$ 
10             parallel for  $u \in S$ 
11                 if  $\text{CAS}(active[u], \text{FALSE}, \text{TRUE})$ 
12                      $\text{MB-INSERT}(Q, u, color[u])$ 
13              $r = r + 1$ 
```

Parallel compare and swap (CAS) only here to ensure each node receives at most one update per round

# Prism

# Questions?

PRISM( $G, f, Q_0$ )

```
1   $\chi = \text{COLOR-GRAPH}(G)$ 
2   $r = 0$ 
3   $Q = Q_0$ 
4  while  $Q \neq \emptyset$ 
5       $C = \text{MB-COLLECT}(Q)$ 
6      for  $C \in \mathcal{C}$ 
7          parallel for  $v \in C$ 
8               $\text{active}[v] = \text{FALSE}$ 
9               $S = f(v)$ 
10             parallel for  $u \in S$ 
11                 if  $\text{CAS}(\text{active}[u], \text{FALSE}, \text{TRUE})$ 
12                      $\text{MB-INSERT}(Q, u, \text{color}[u])$ 
13              $r = r + 1$ 
```

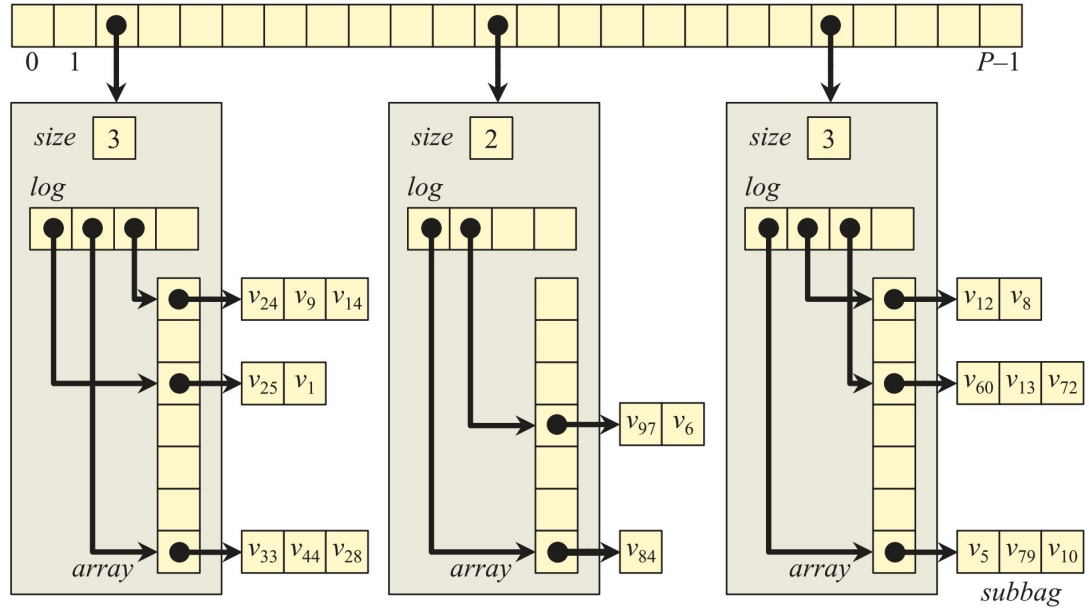
MB-Collect: Empty all bags from multibag  $Q$  into a collection  $C$  that is easy to iterate over

MB-Insert: Insert node  $u$  in bag  $\text{color}[u]$  of multibag  $Q$

# Key Idea: Parallel Lock-free Multibags

- Problem: How to make MB-Collect work-efficient?
- Just using over a bitmap of active nodes won't do
  - $\Theta(V * \text{Color})$  work to do a round; not work-efficient
- Just having each P maintain Color many local arrays of active elements also won't do
  - $\Theta(P * \text{Color})$  work to do a round; not work-efficient

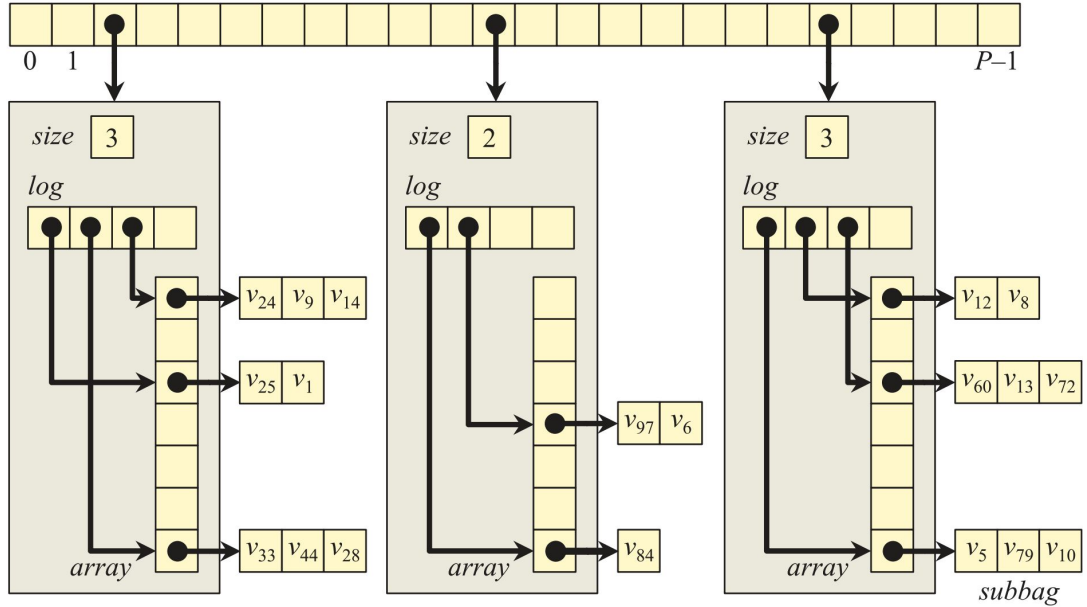
# Key Idea: Parallel Lock-free Multibags



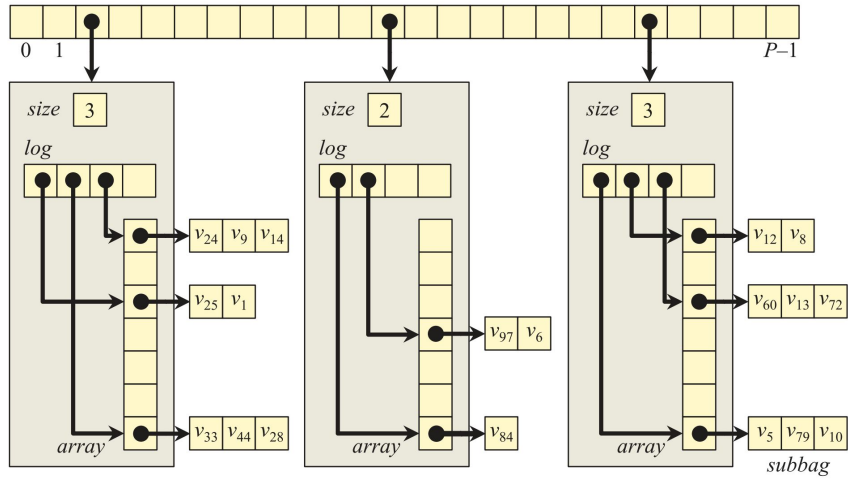
# MB-Insert

MB-INSERT( $Q, v, k$ )

- 1  $p = \text{GET-WORKER-ID}()$
- 2 **if**  $Q[p].\text{array}[k] == \text{NIL}$
- 3     APPEND( $Q[p].\text{log}, k$ )
- 4      $Q[p].\text{array}[k] = \text{new subbag}$
- 5     APPEND( $Q[p].\text{array}[k], v$ )

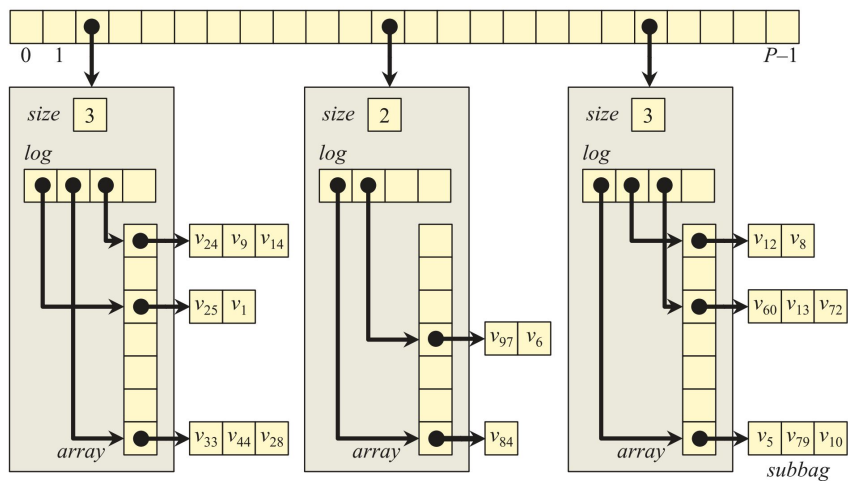


# MB-Collect (Board)

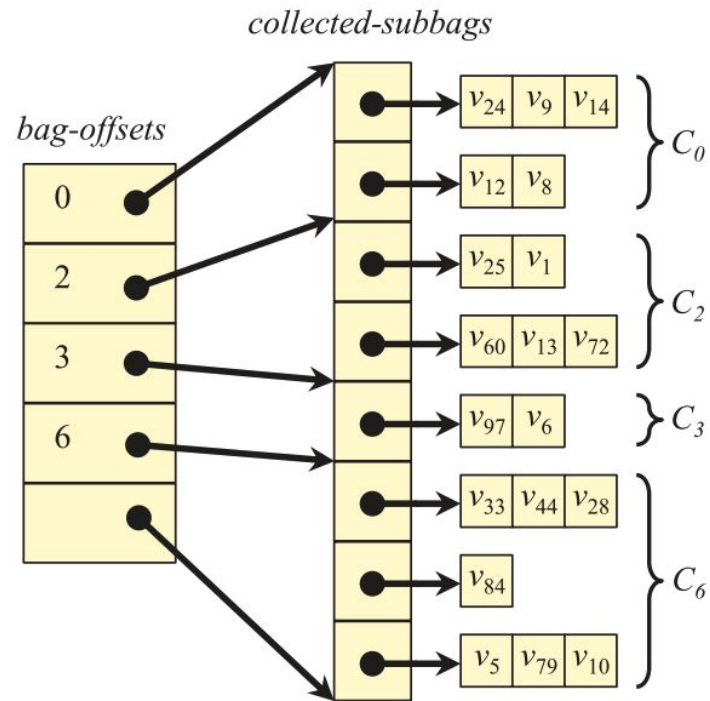




# MB-Collect (Board)



MB-Collect



# Multibags Performance

- MB-insert is  $\Theta(1)$
- For  $m$  number of total subbags
- MB-Collect is  $O(m + \text{Color} + P)$  work and  $O(\log(m) + \text{Color} + \log(P))$  span
- For all the obvious reasons... (use prefix sum and friends)

# Final Theoretical Result

- Suppose that Prism colors a degree- $D$  data graph  $G=(V, E)$  using  $\xi$  colors and then executes the data-graph computation  $(G, f, Q_0)$ . Then, on  $P$  processors, Prism executes updates on all vertices in the activation set  $Q_r$  for a round  $r$  using

Work:  $O(\text{size}(Q_r) + P)$

Span:  $O(\xi * (\log(Q_r/\xi) + \log(D)) + \log(P))$ ,

where  $\text{size}(Q) := |Q| + \sum_{v \in Q} \text{deg}(v)$

PRISM( $G, f, Q_0$ )

1  $\chi = \text{COLOR-GRAPH}(G)$

2  $r = 0$

3  $Q = Q_0$

4 **while**  $Q \neq \emptyset$

5      $C = \text{MB-COLLECT}(Q)$

6     **for**  $C \in \mathcal{C}$

7         **parallel for**  $v \in C$

8              $active[v] = \text{FALSE}$

Work: Theta(deg(v)); Span: Theta(log(deg(v)))

9              $S = f(v)$

10             **parallel for**  $u \in S$

11                 **if** CAS( $active[u], \text{FALSE}, \text{TRUE}$ ) Work: Theta(S); Span: Theta(log(S))

12                     MB-INSERT( $Q, u, color[u]$ )

13      $r = r + 1$

PRISM( $G, f, Q_0$ )

1  $\chi = \text{COLOR-GRAPH}(G)$

2  $r = 0$

3  $Q = Q_0$

4 **while**  $Q \neq \emptyset$

5      $C = \text{MB-COLLECT}(Q)$

6     **for**  $C \in \mathcal{C}$

7         **parallel for**  $v \in C$

8              $active[v] = \text{FALSE}$

9              $S = f(v)$

10             **parallel for**  $u \in S$

11                 **if**  $\text{CAS}(active[u], \text{FALSE}, \text{TRUE})$

12                      $\text{MB-INSERT}(Q, u, color[u])$

13      $r = r + 1$

Work: Theta(size(C));

Span: Theta(log(C) + log(D))

PRISM( $G, f, Q_0$ )

1  $\chi = \text{COLOR-GRAPH}(G)$

2  $r = 0$

3  $Q = Q_0$

4 **while**  $Q \neq \emptyset$

5      $C = \text{MB-COLLECT}(Q)$

6     **for**  $C \in \mathcal{C}$

7         **parallel for**  $v \in C$

8              $active[v] = \text{FALSE}$

9              $S = f(v)$

10             **parallel for**  $u \in S$

11                 **if**  $\text{CAS}(active[u], \text{FALSE}, \text{TRUE})$

12                      $\text{MB-INSERT}(Q, u, color[u])$

13      $r = r + 1$

Work:  $\Theta(\text{size}(Q_r) + \xi)$

Span:  $\Theta(\xi \cdot (\log(Q_r/\xi) + \log(D)))$

PRISM( $G, f, Q_0$ )


```
1  $\chi = \text{COLOR-GRAPH}(G)$ 
2  $r = 0$ 
3  $Q = Q_0$ 
4 while  $Q \neq \emptyset$ 
5    $C = \text{MB-COLLECT}(Q)$   $O(m + \text{Color} + P)$  work and  $O(\log(m) + \text{Color} + \log(P))$  span
6   for  $C \in \mathcal{C}$ 
7     parallel for  $v \in C$ 
8        $\text{active}[v] = \text{FALSE}$ 
9        $S = f(v)$ 
10      parallel for  $u \in S$ 
11        if  $\text{CAS}(\text{active}[u], \text{FALSE}, \text{TRUE})$ 
12           $\text{MB-INSERT}(Q, u, \text{color}[u])$ 
13       $r = r + 1$ 
```

# Experimental Results

<i>Graph</i>	$ V $	$ E $	$\chi$	CILK+LOCKS	PRISM	PRISM-R	<i>Coloring</i>
cage15	5,154,860	94,044,700	17	36.9	35.5	35.6	12%
liveJournal	4,847,570	68,475,400	333	36.8	21.7	22.3	12%
randLocalDim25	1,000,000	49,992,400	36	26.7	14.4	14.6	18%
randLocalDim4	1,000,000	41,817,000	47	19.5	12.5	13.7	14%
rmat2Million	2,097,120	39,912,600	72	22.5	16.6	16.8	12%
powerGraph2M	2,000,000	29,108,100	15	12.1	9.8	10.1	13%
3dgrid5m	5,000,210	15,000,600	6	10.3	10.3	10.4	7%
2dgrid5m	4,999,700	9,999,390	4	17.7	8.9	9.0	4%
web-Google	916,428	5,105,040	43	3.9	2.4	2.4	8%
web-BerkStan	685,231	7,600,600	200	3.9	2.4	2.7	8%
web-Stanford	281,904	2,312,500	62	1.9	0.9	1.0	11%
web-NotreDame	325,729	1,469,680	154	1.1	0.8	0.8	12%



# Coloring Need Not Be Perfect



Thanks! Questions?