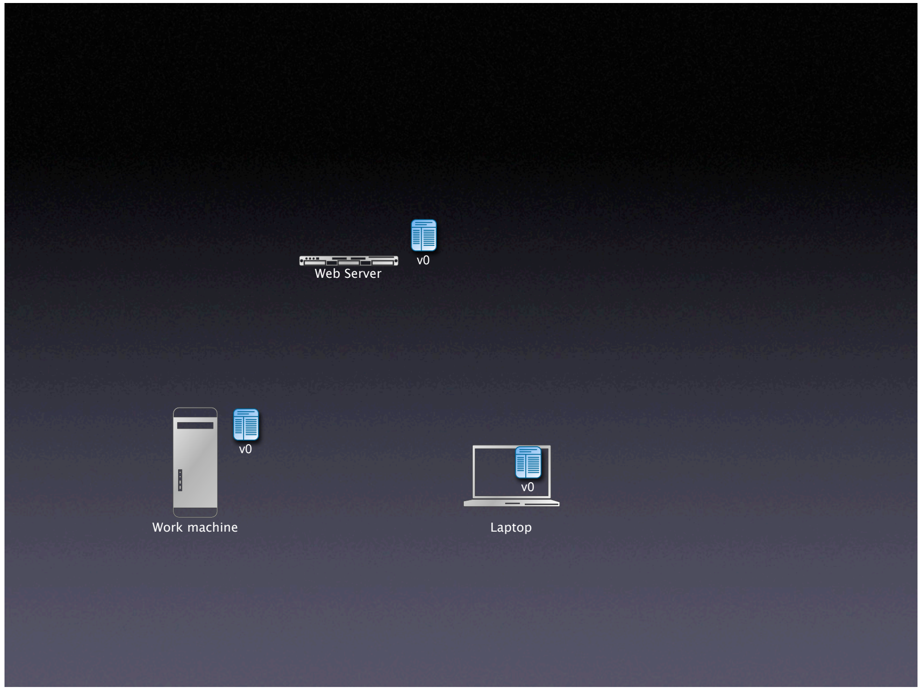

Setup: 3 computers and a website to keep up to date

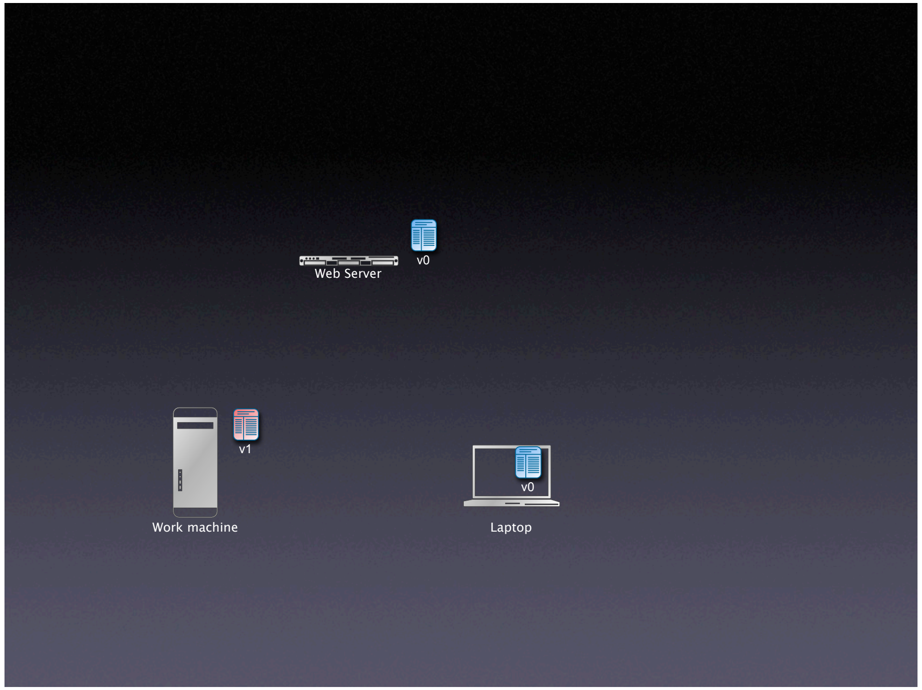

Editing is nicer to do on a personal machine

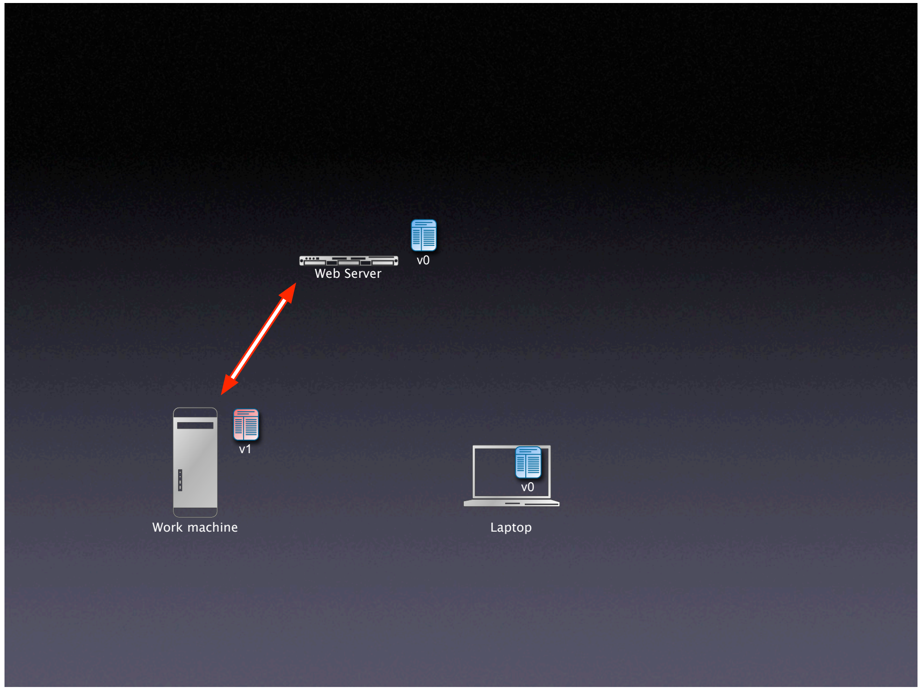

Web publishing tools may not know how to do synchronization
Push and pull changes, may use heterogeneous concrete formats

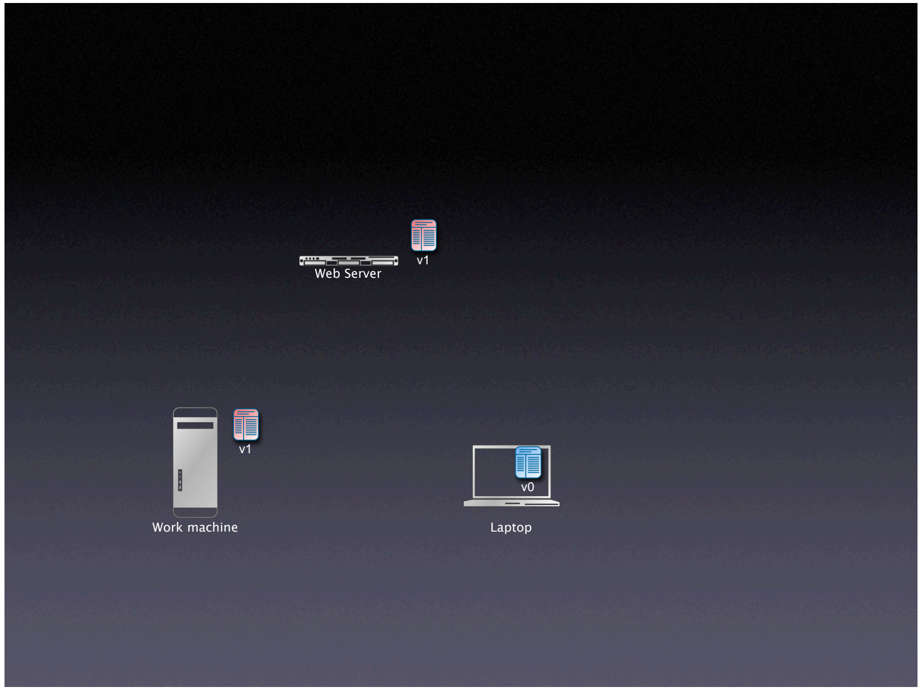
v1
Web Server
v1
v0
Work machine
Laptop

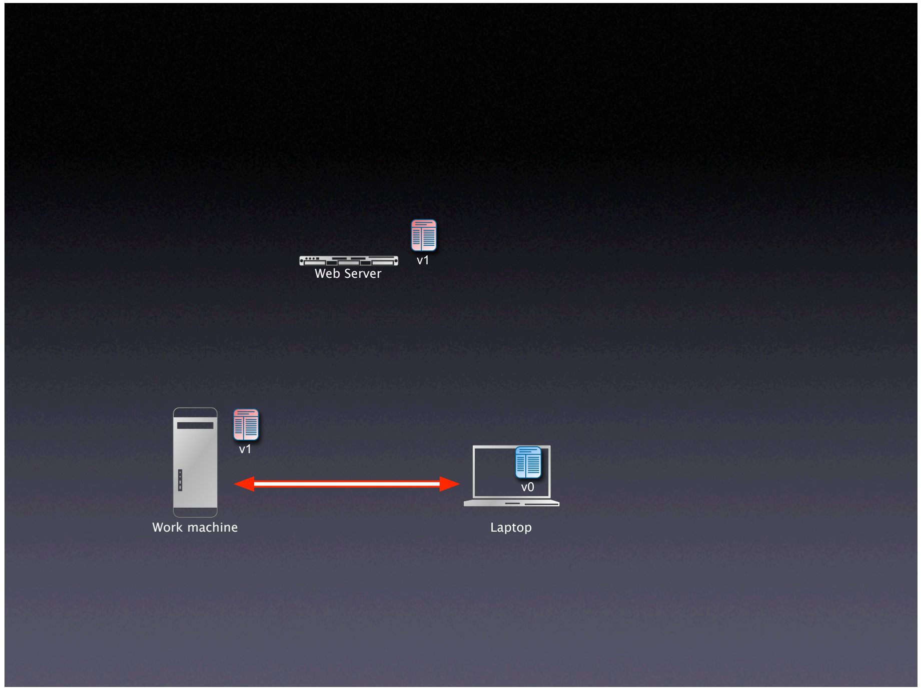

Time to go home, synchronize the laptop

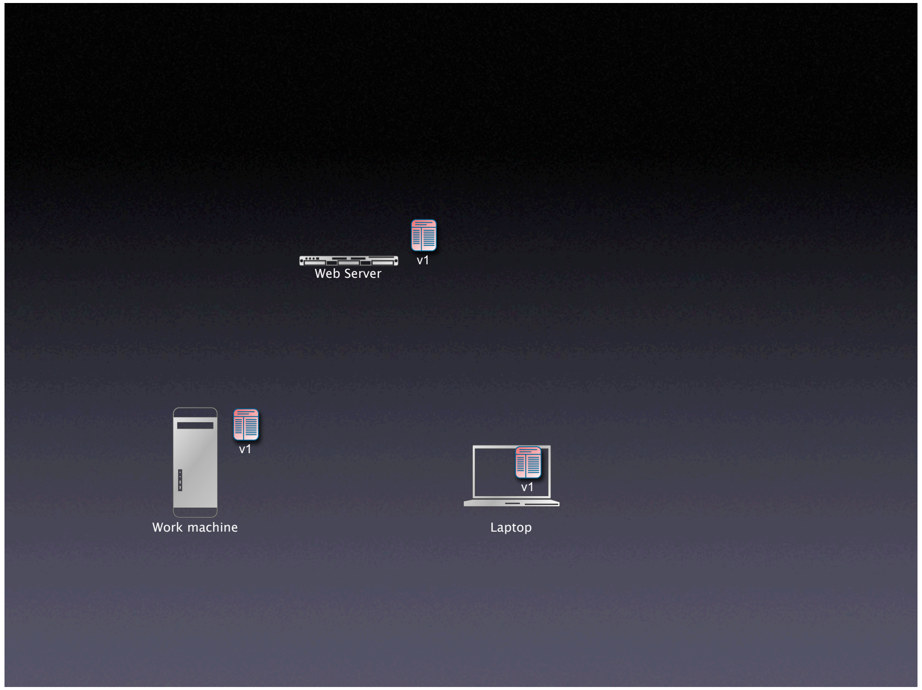

v1
Web Server
v1
Work machine
v1
Laptop

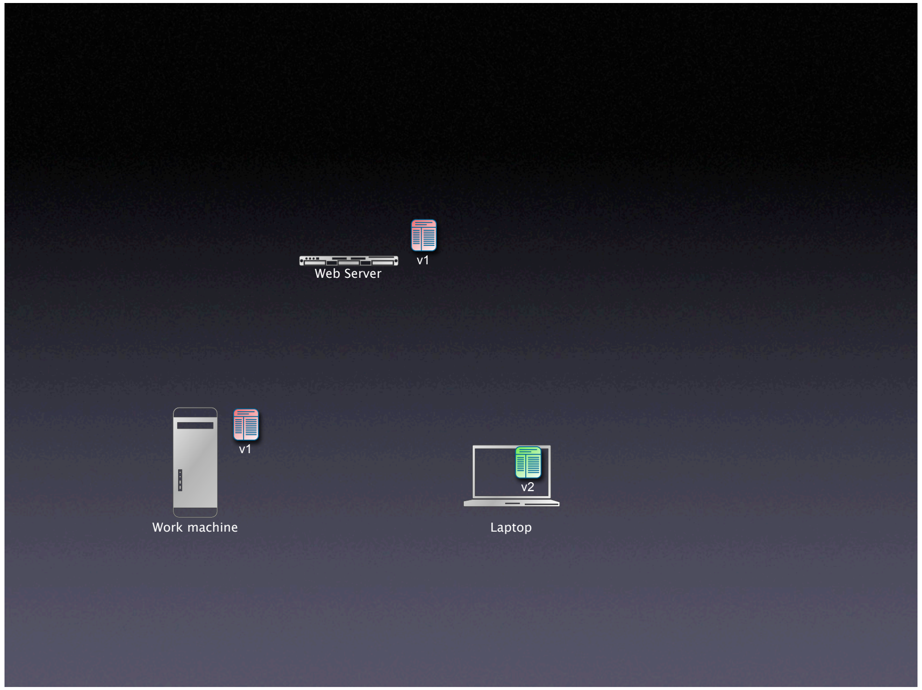

Urgent update needed (contact email needs to be changed), do it directly from home using the notebook and up-to-date local version of the website

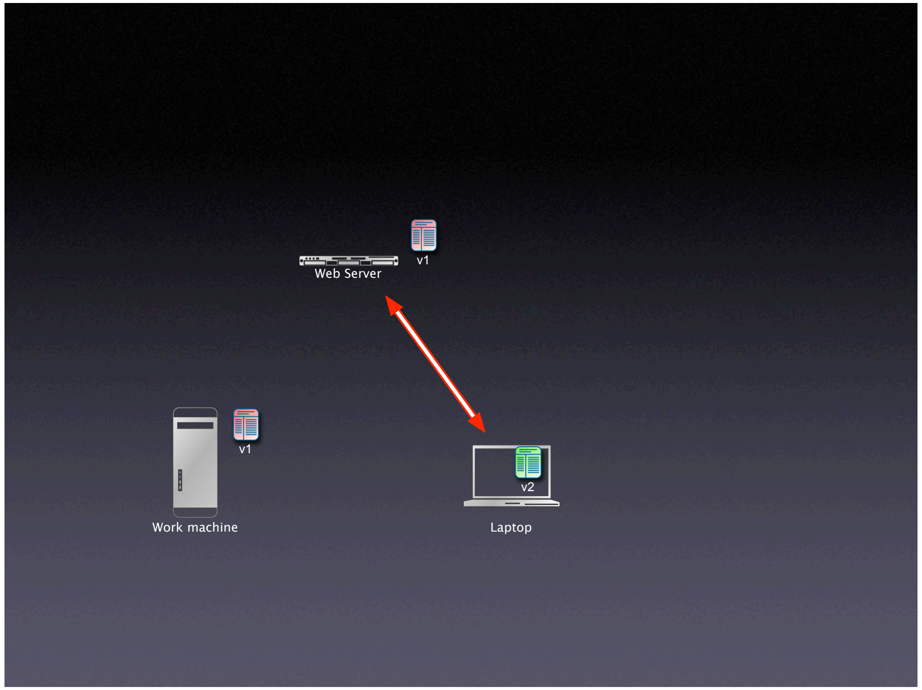

Update done, time to synchronize

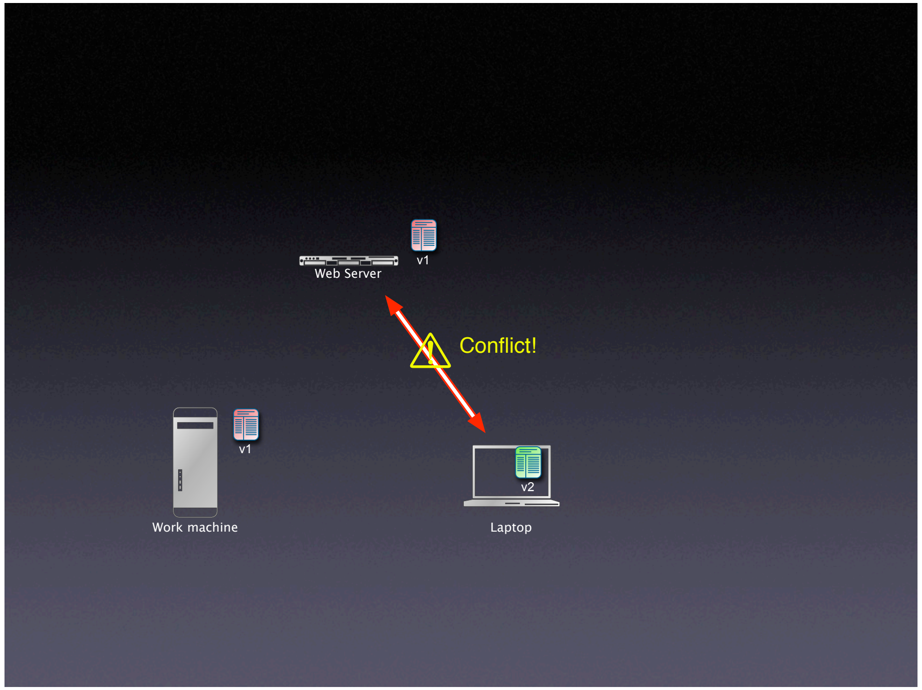

and Synchronization fails... We need a better tool for this.
Will not go into details why this failed, this shows the setting we’re interested in

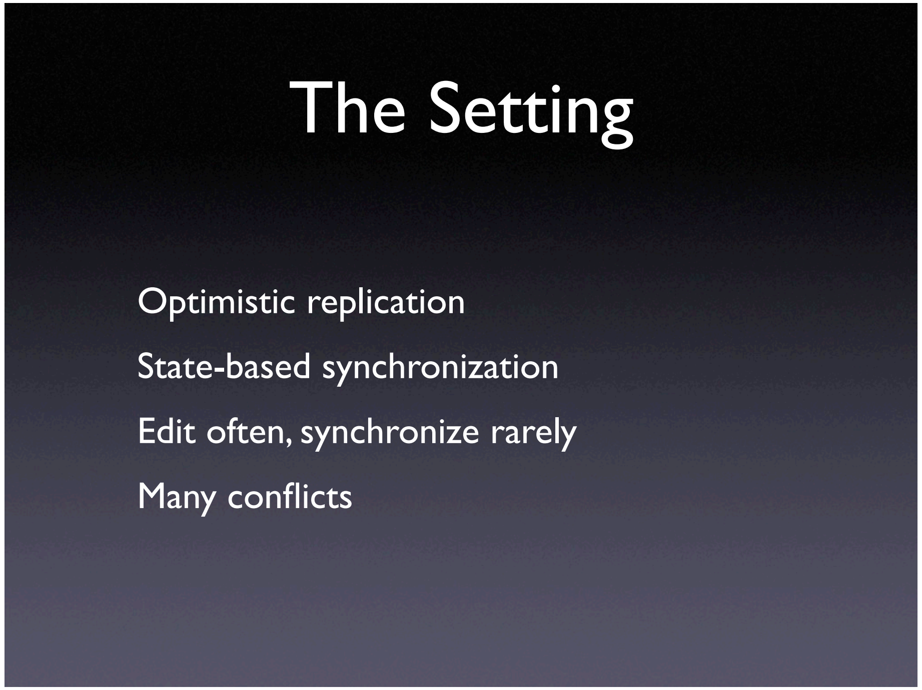

Not well studied in the literature

We propose a better tool for this setting
Joke about many different affiliations and synchronization?

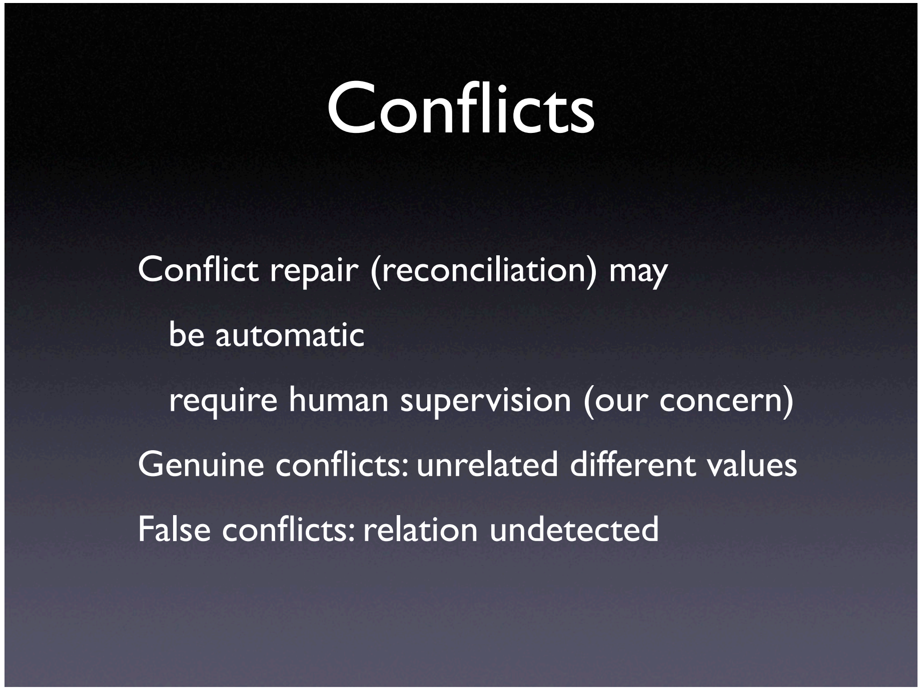

For optimistic replication with many devices that are seldom synchronized, many real conflicts

Dealing with Conflicts

- Ignore them
  - No convergence
- Repair them: choose a better value
  - Create a new value dominating the conflict
  - Convergence?

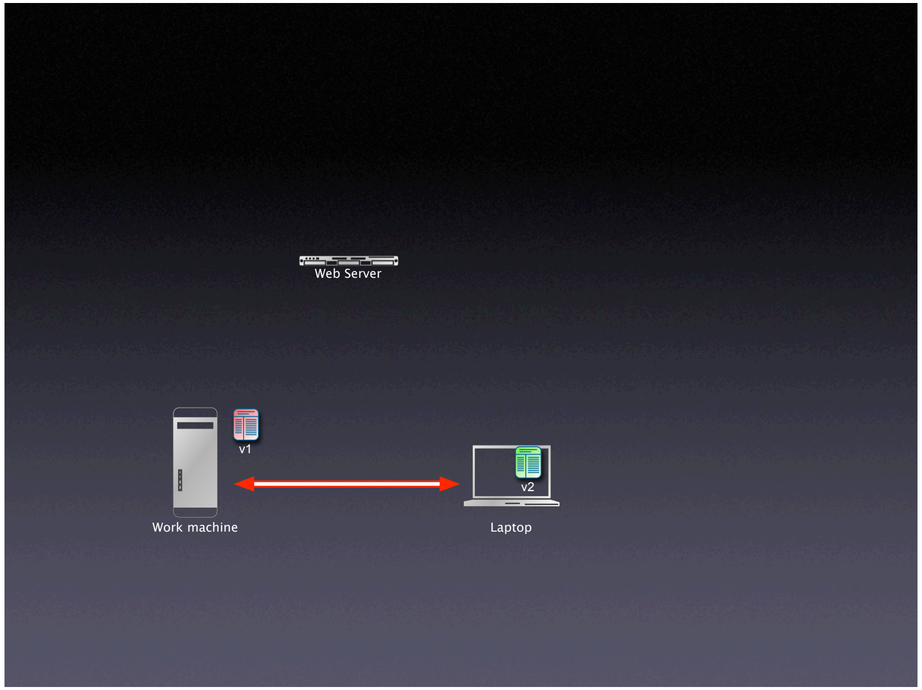

real conflict

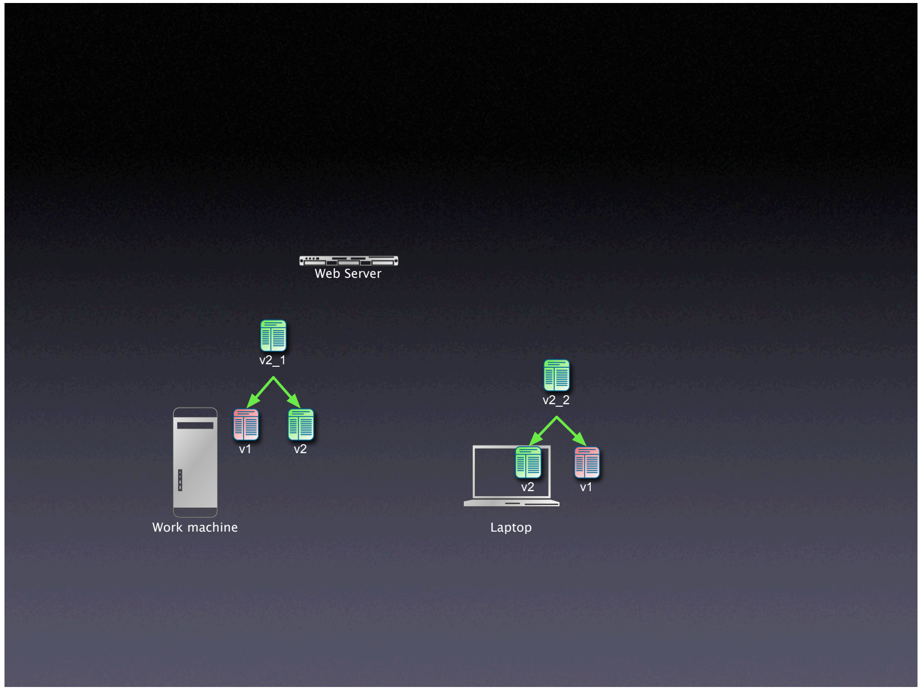

same decision: choose v2’s value

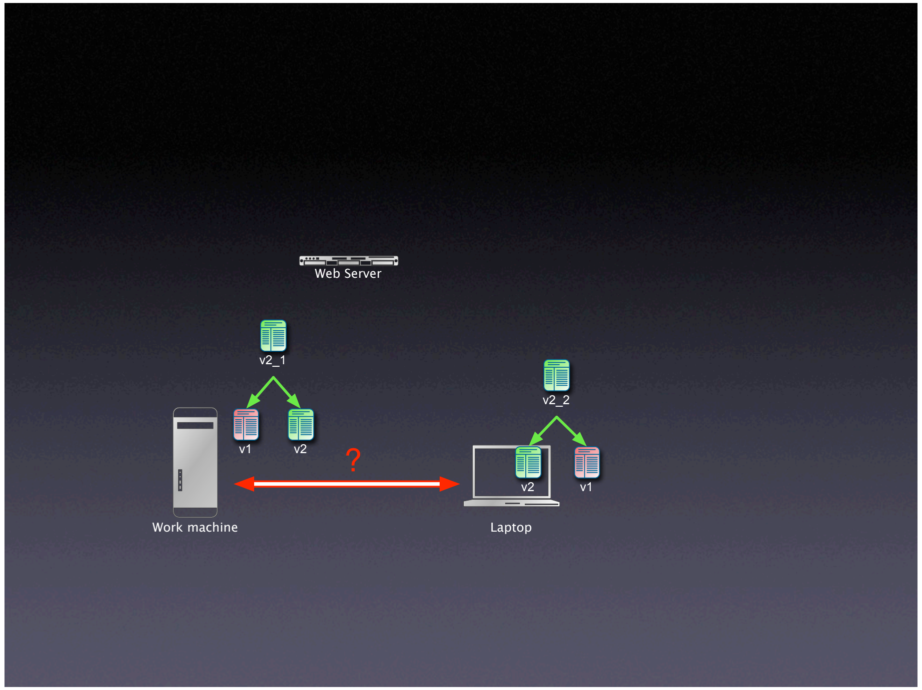

causal history approach does not detect the agreement

Our Approach

- Record agreement explicitly with equivalence
- Equivalence enables optimistic reconciliation
- Bounded representation in $O(n^4)$ of the history graph with equivalence edges

Equivalence

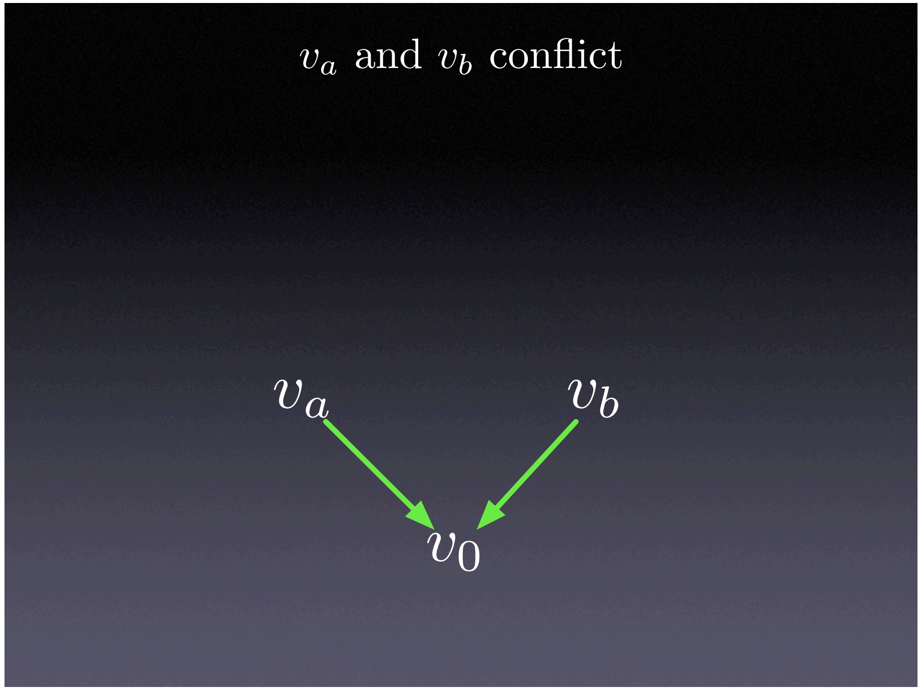

v_a and v_b conflict
v_a
v_b
v_0

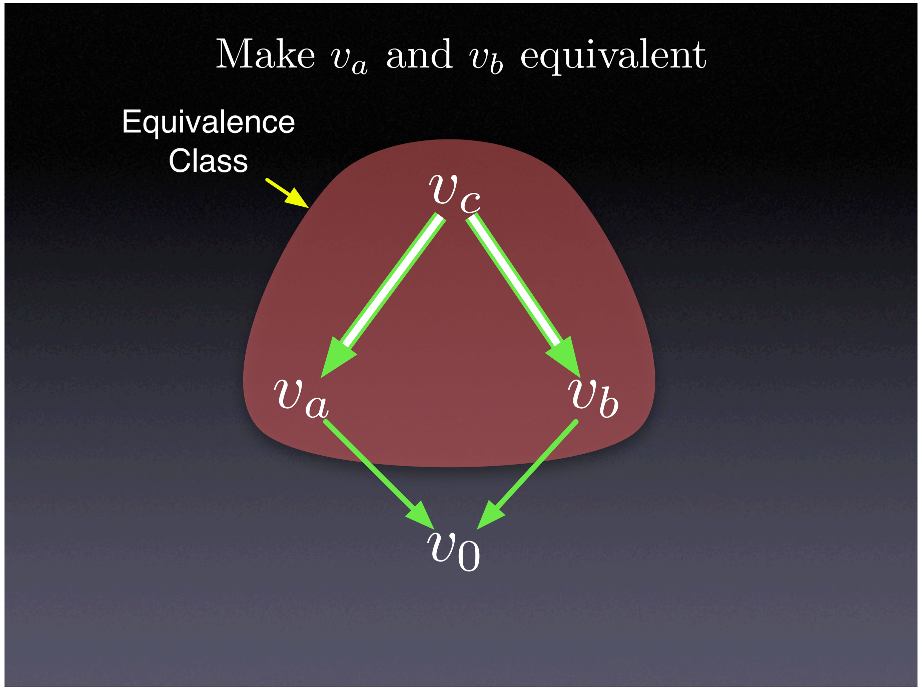

new edges created only from newly introduced values
no edge created between previously existing values

v_a and v_b conflict

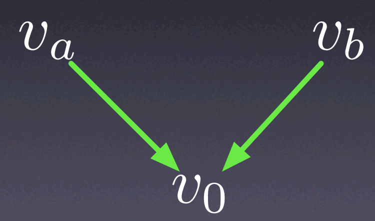

Choose v_b

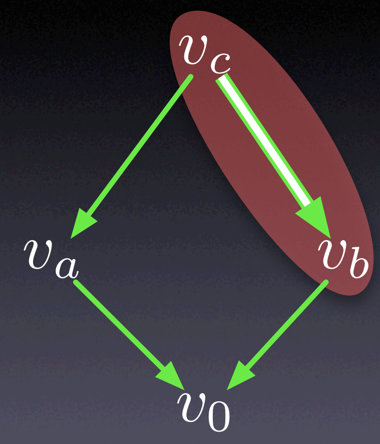

Explicit equivalence edge allow distributed reconciliation

Replicas independently decide events are equivalent

No need for a central decision replica

→ distributed reconciliation

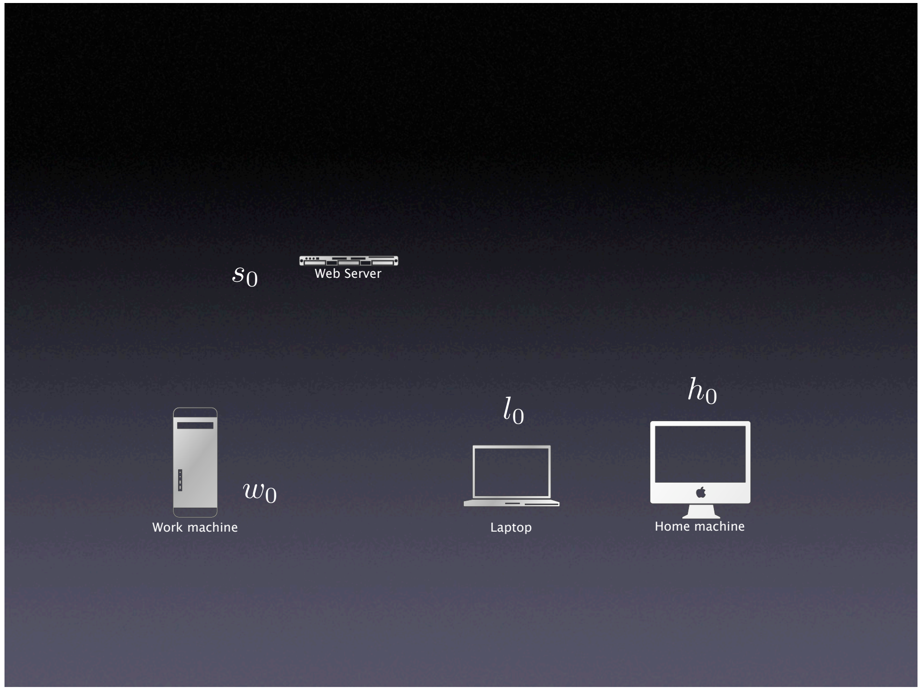

several events, all equivalent from a human point of view (but not yet from the system)

at least 3 synchronizations needed, in many systems the order might matter

in our approach, the order does not matter

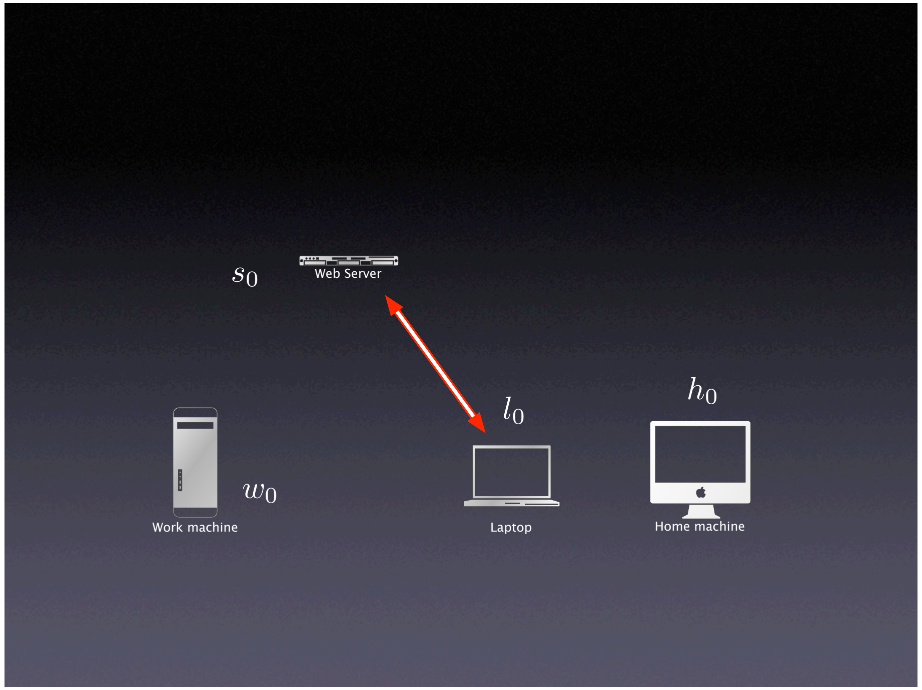

s_0
Web Server
l_0
h_0
w_0
Work machine
Laptop
Home machine

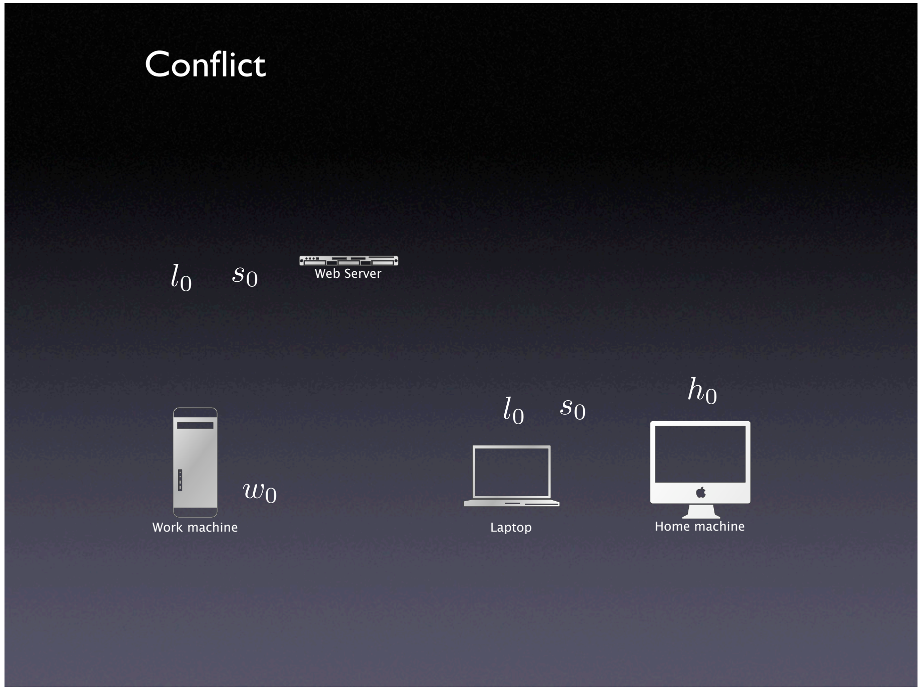

Conflict
l_0 s_0
Web Server
l_0 s_0
h_0
w_0
Work machine
Laptop
Home machine

Conflict repaired by equivalence at laptop

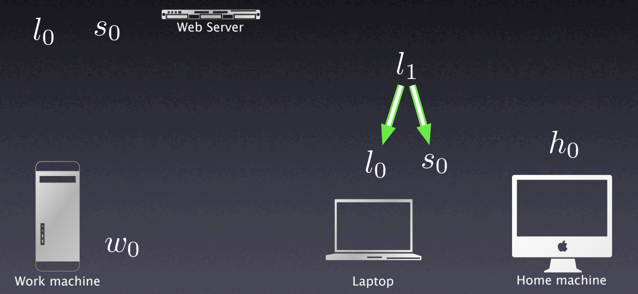

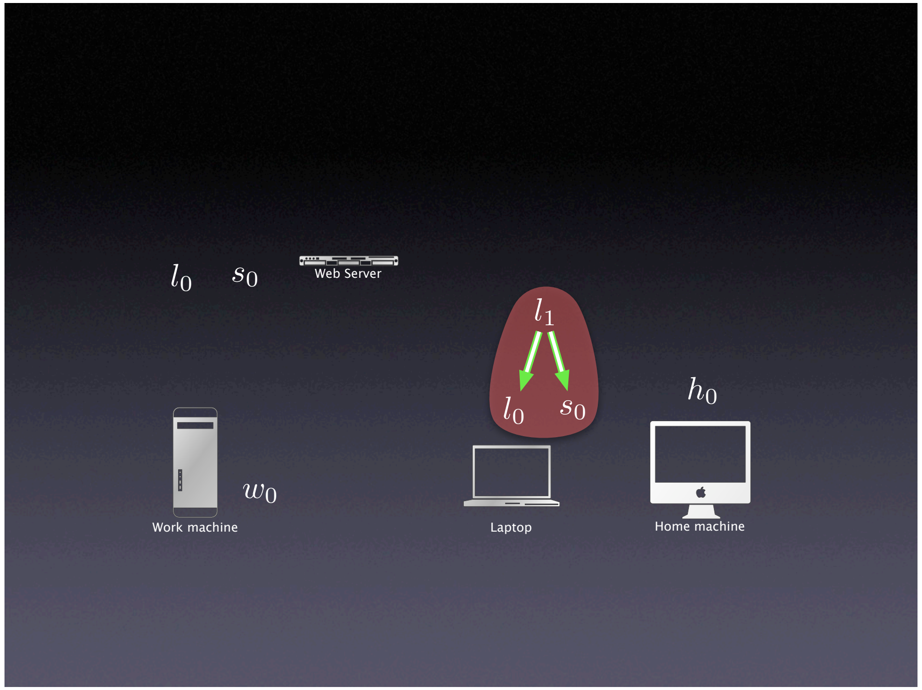

l_0 s_0
Web Server
l_1
l_0 s_0
h_0
w_0
Work machine
Laptop
Home machine

Conflict repaired by equivalence at web server

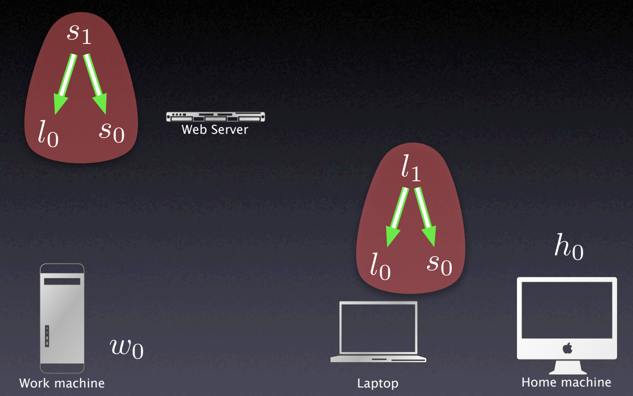

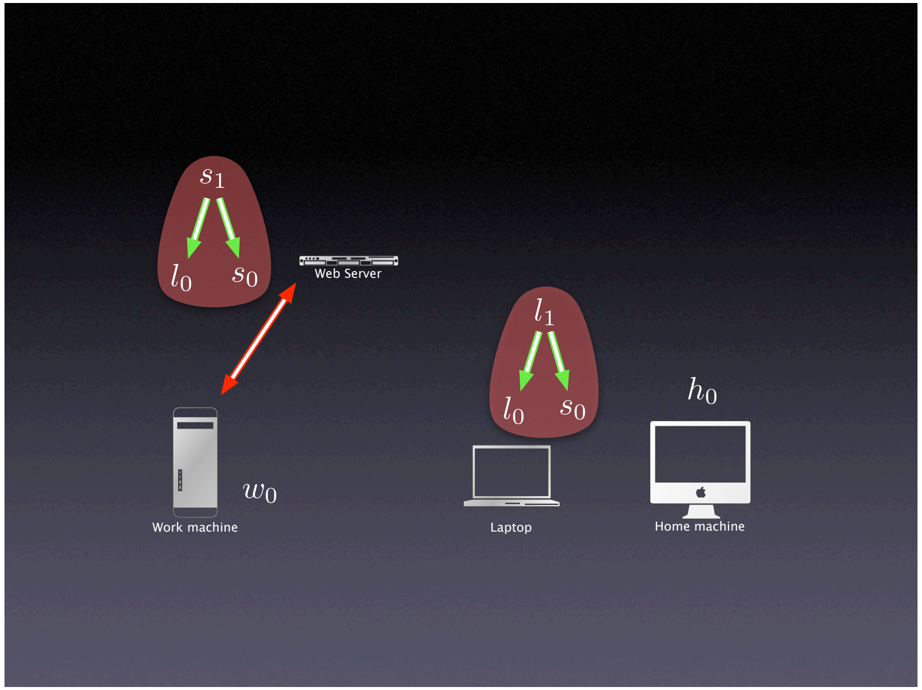

s_1
l_0
s_0
Web Server
l_1
l_0
s_0
h_0
w_0
Work machine
Laptop
Home machine

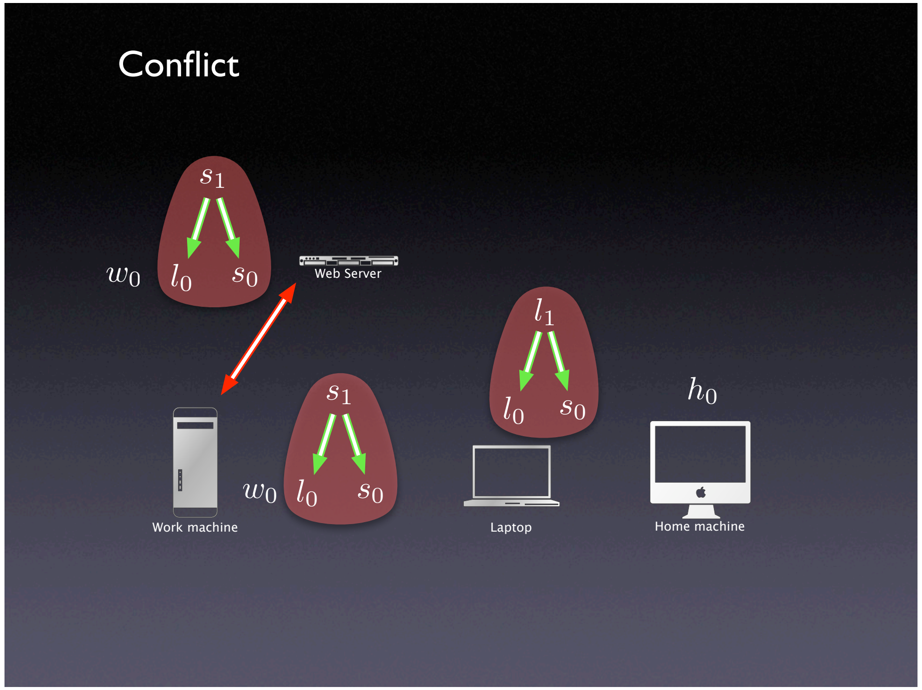
Conflict
s_1
w_0 l_0 s_0
Web Server
l_1
l_0 s_0
h_0
s_1
w_0 l_0 s_0
Work machine
Laptop
Home machine

Conflict repaired by equivalence at both sites

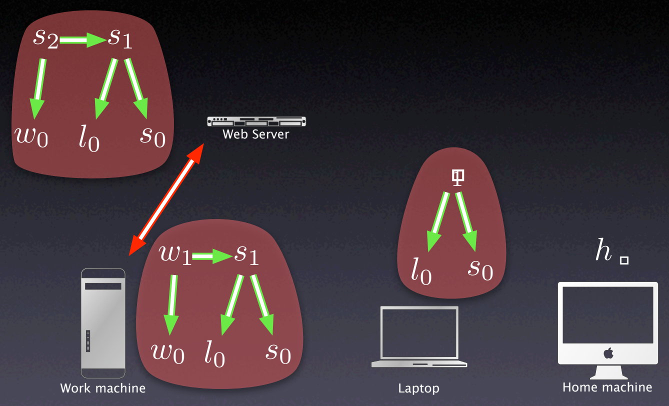

Conflict

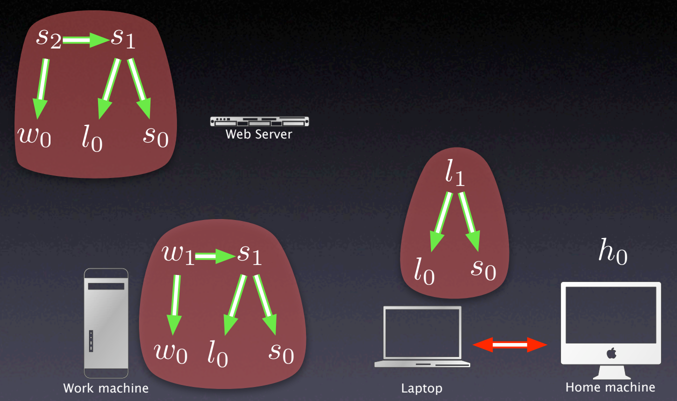

$s_2 \rightarrow s_1$
w_0 l_0 s_0
Web Server
l_1
l_0 s_0
$w_1 \rightarrow s_1$
w_0 l_0 s_0
h_0
Work machine
Laptop
Home machine

Conflict repaired by equivalence at both sites

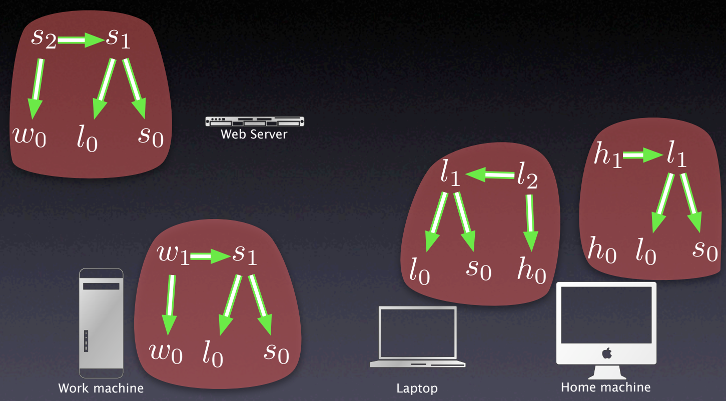

Equivalence reached "globally"; conflict?

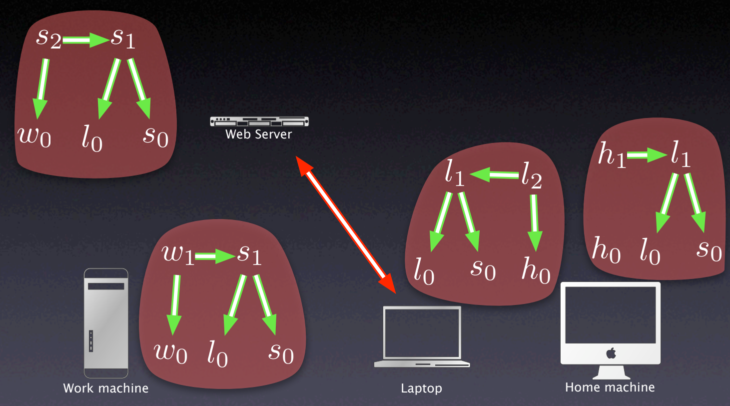

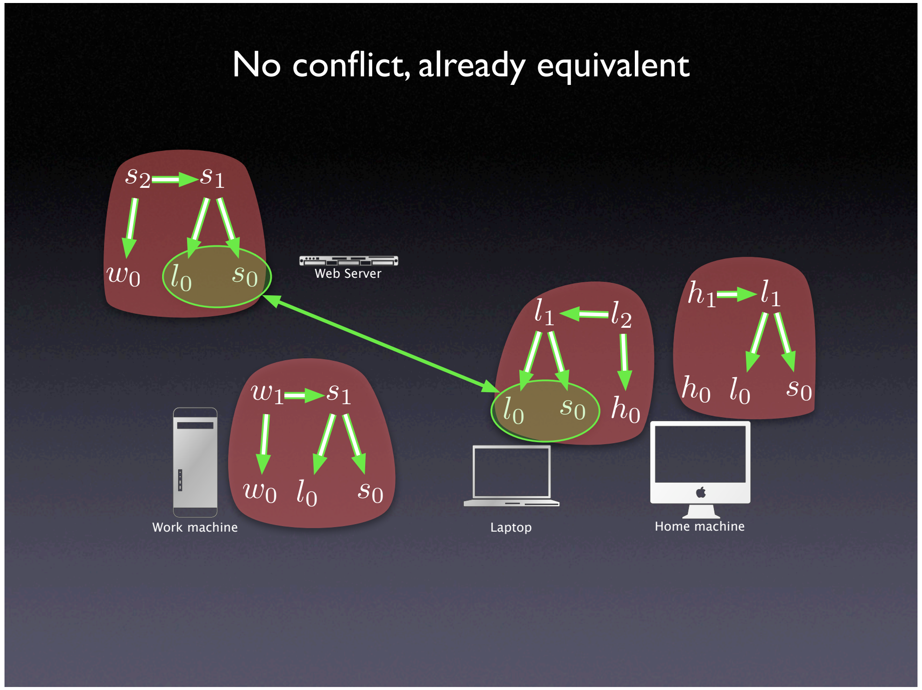
No conflict, already equivalent
s_2 s_1
w_0 l_0 s_0
Web Server
w_1 s_1
w_0 l_0 s_0
Work machine
l_1 l_2
l_0 s_0 h_0
Laptop
h_1 l_1
h_0 l_0 s_0
Home machine

Optimistic Editing

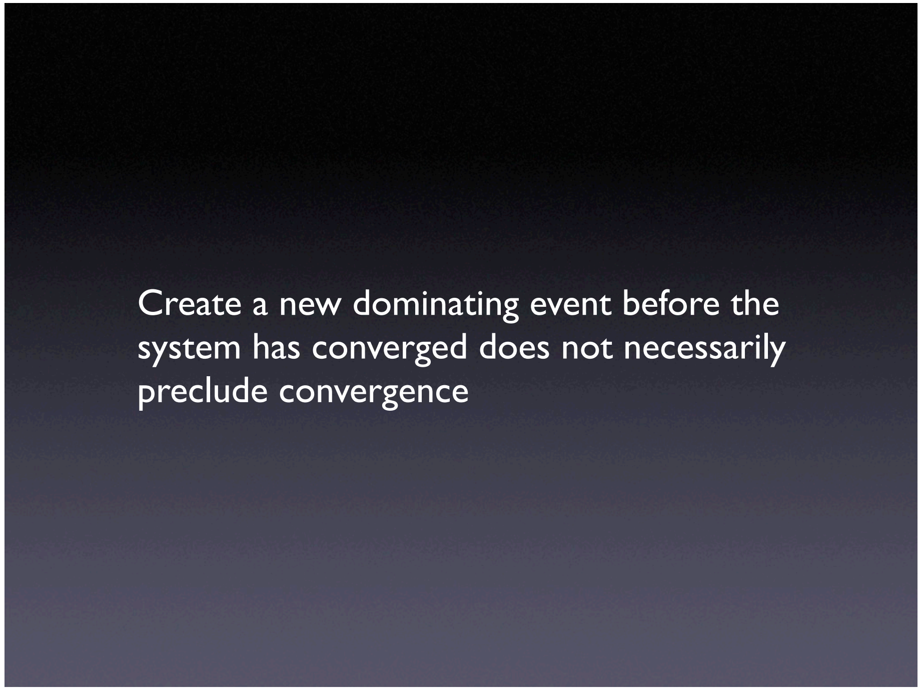

Not only distributed reconciliation but optimistic reconciliation: one may edit before conflict resolution has reached everyone

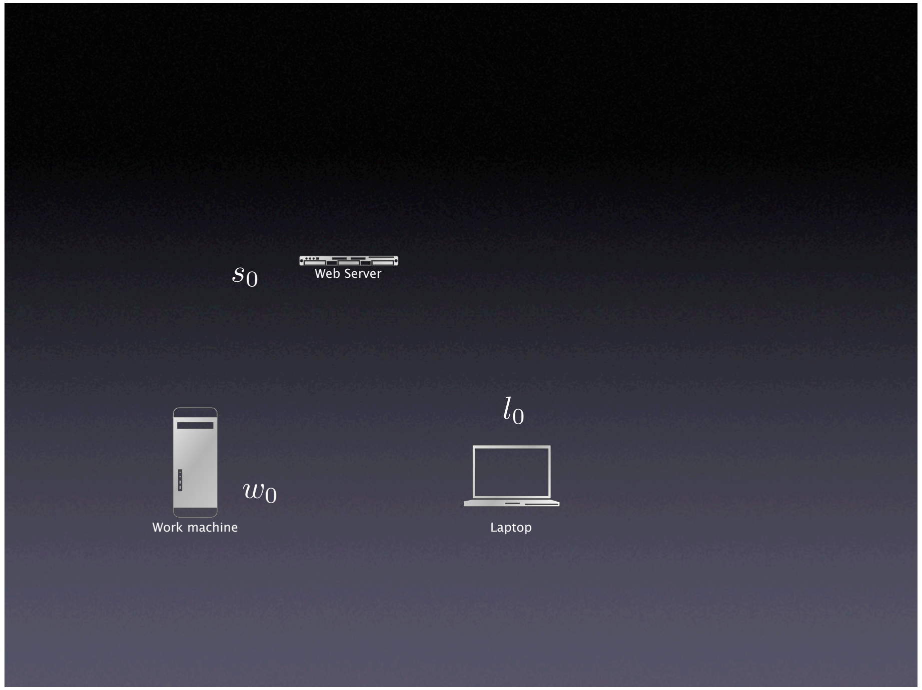
s_0
Web Server
l_0
w_0
Work machine
Laptop

Conflict repaired by equivalence

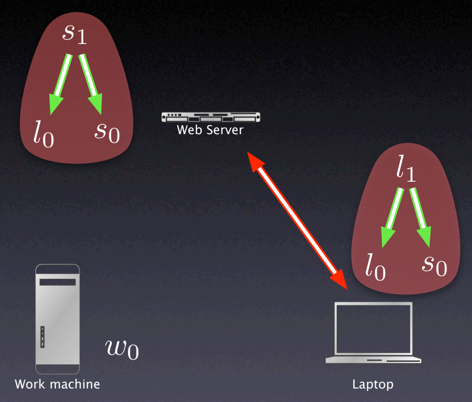

Conflict repaired by equivalence

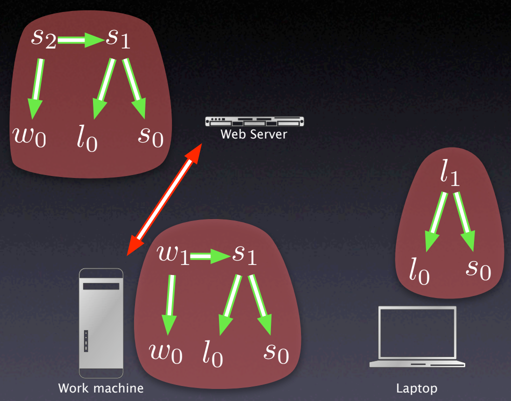

(Optimistic) edit on laptop

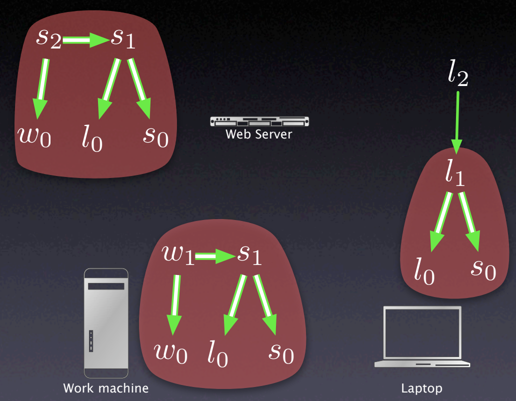

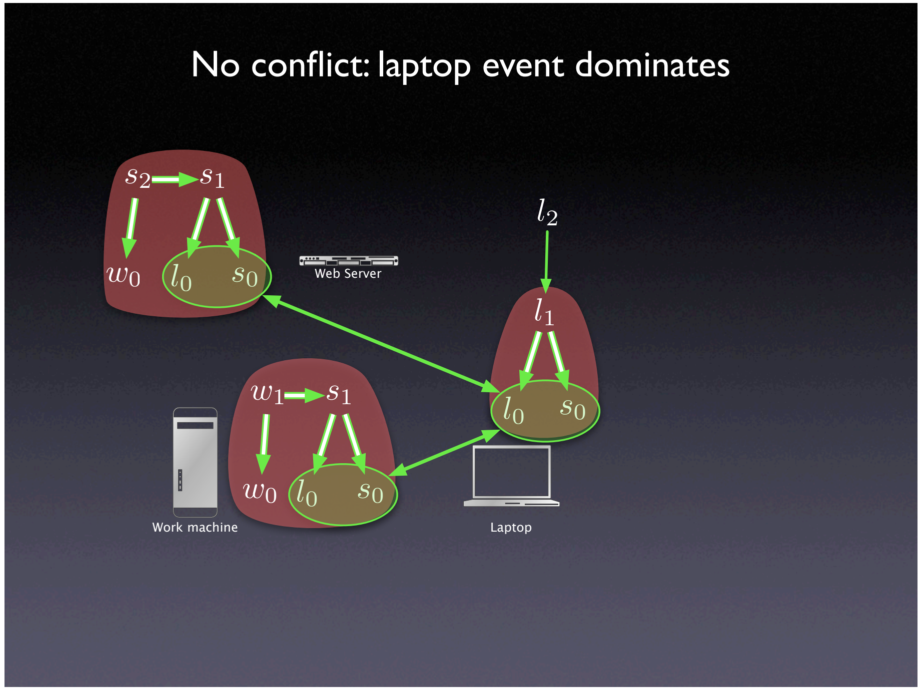

Any further synchronization will propagate v2l safely, even though neither work nor web talked to laptop since synchronizing (optimistic reconciliation)

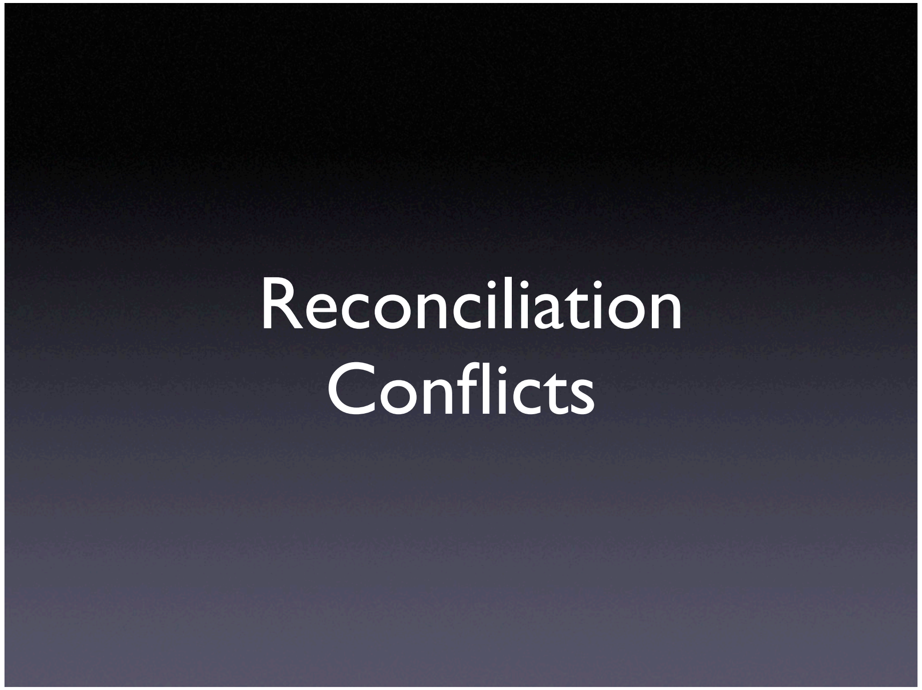

Interesting new kind of conflicts

Conflicts may be solved in incompatible ways

- choose new incompatible events
- choose a different event as the better one

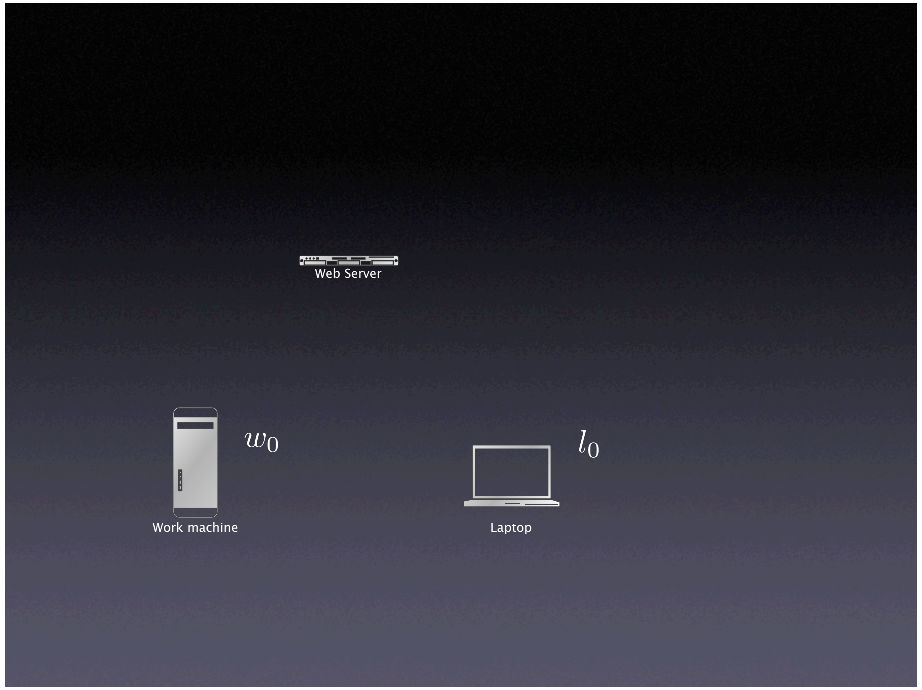
Web Server
w_0
Work machine
l_0
Laptop

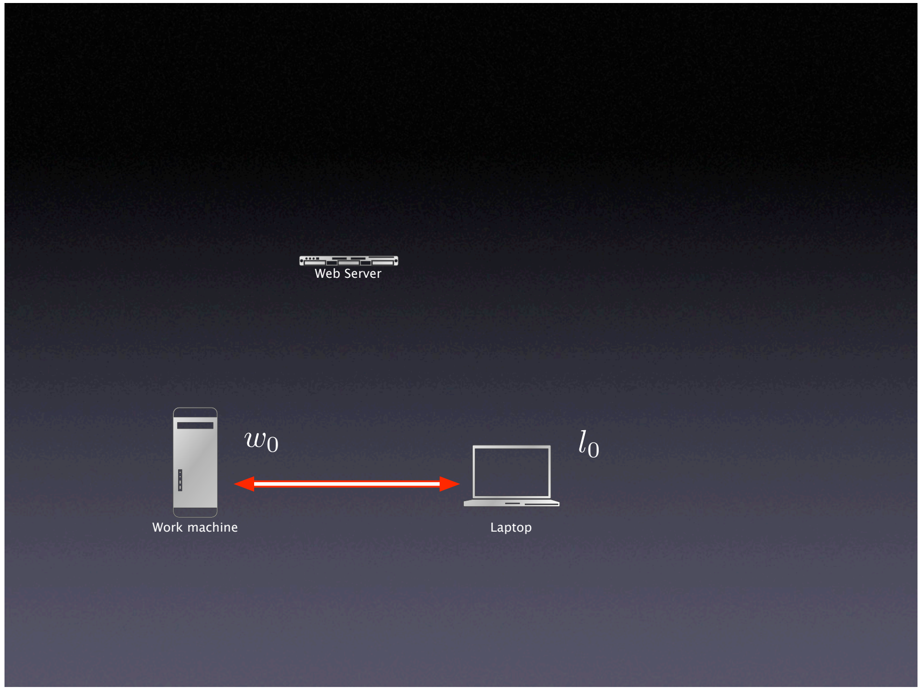
Web Server
w_0
l_0
Work machine
Laptop

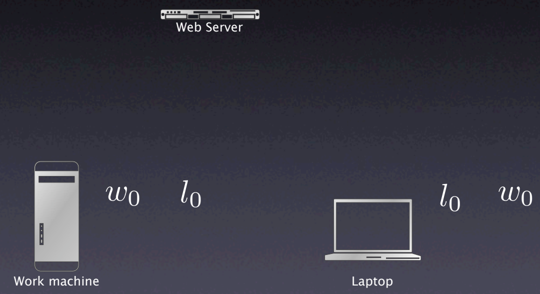
Web Server
Work machine
w_0 l_0
Laptop
l_0 w_0

Conflict: work machine value chosen

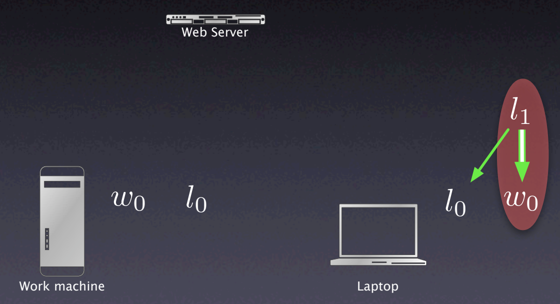

Conflict: laptop value chosen

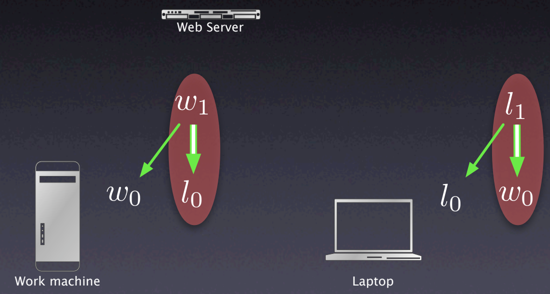

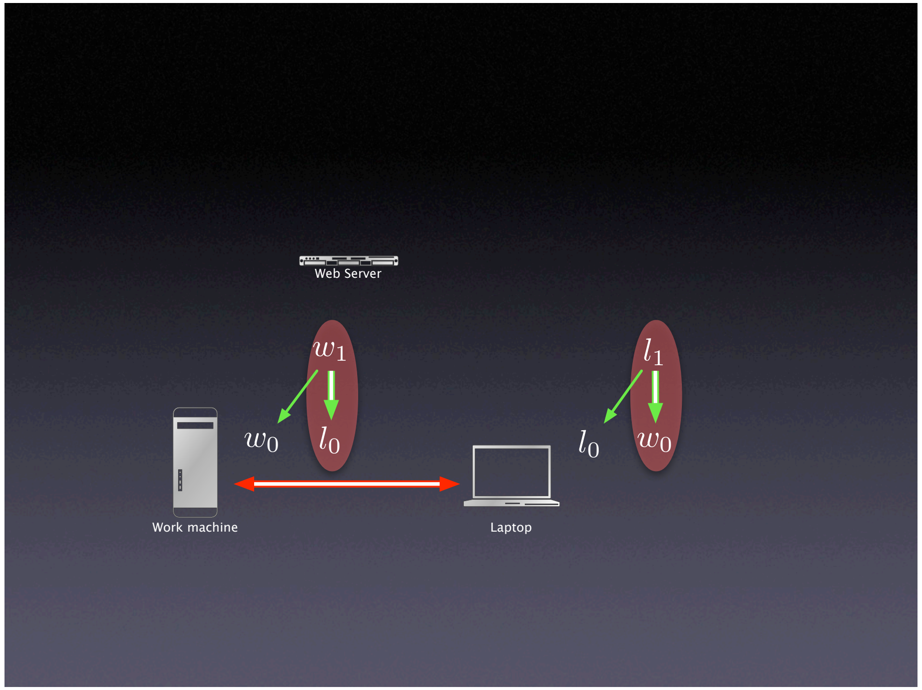
Web Server
w_1
l_0
w_0
l_1
l_0
w_0
Work machine
Laptop

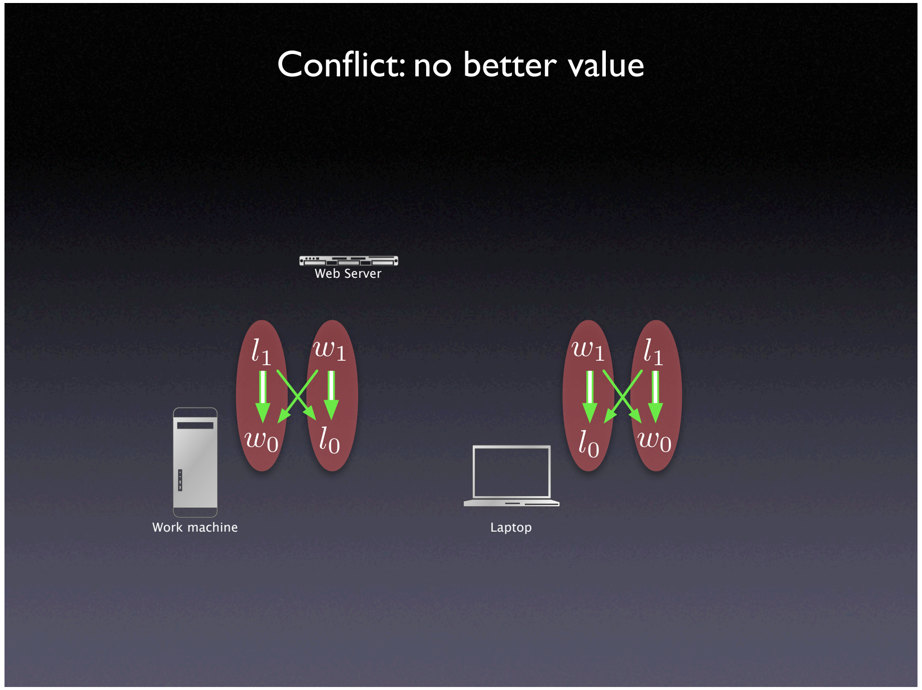

No dominant (maximal) event

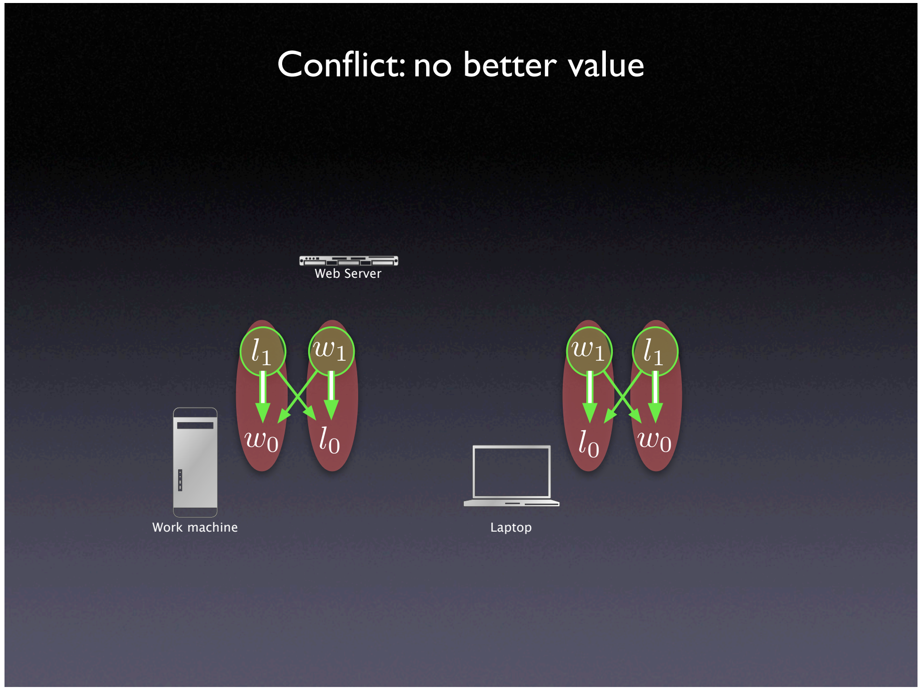

Latest events
Notion of maximal event

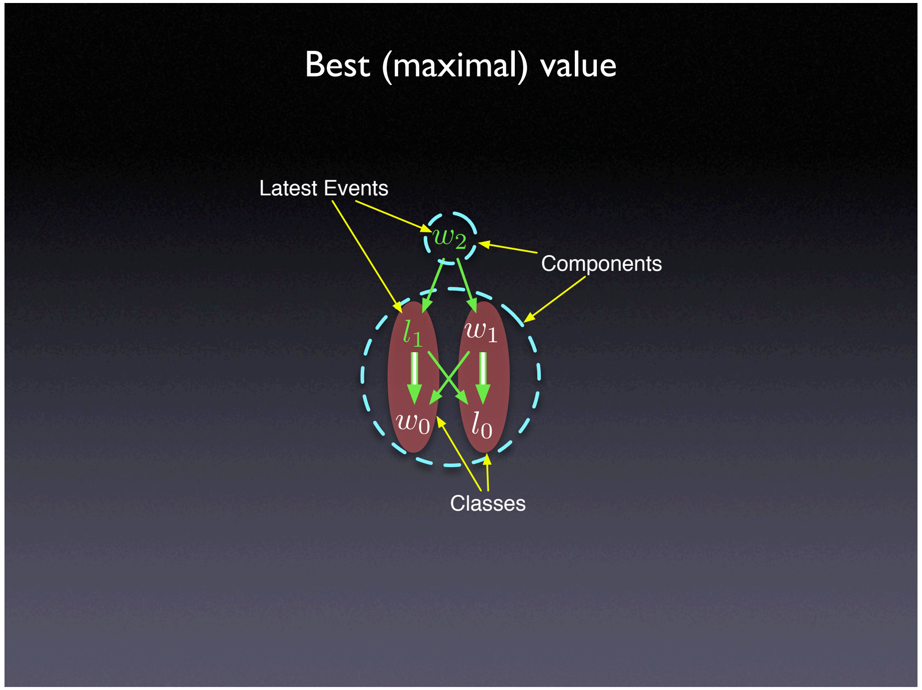

classes, components, latest events, maximal events
new edges only created to maximal events

A Bounded Representation

Slight generalization to simplify the examples and consistent with the paper
One simple way to restrict communication, otherwise unbounded storage is needed
n: number of replicas

Reciprocal Communication

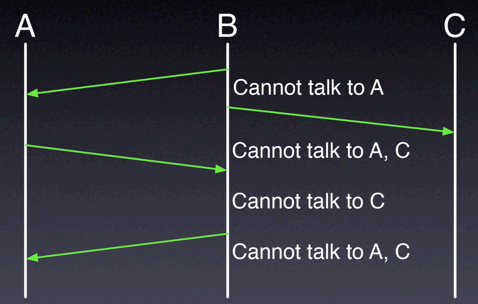

Open Events

New edges only created to maximal events

Event ‘v’ considered open when

- sent a as latest event to replica A
- no communication from A since
- ⇒ A may have created edges to ‘v’

Events that matter: equivalent to open events

Sparse Representation

1. Events considered open
2. Most recent events taken into account by latest events
3. Oldest events equivalent to or dominating open events
4. Equivalence class of open events

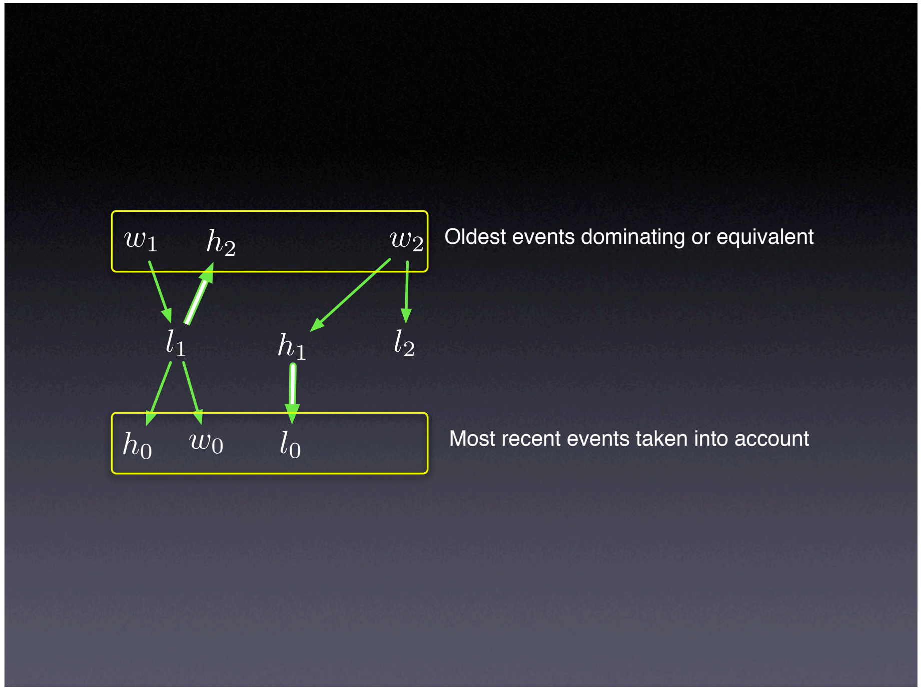
w_1
h_2
w_2
Oldest events dominating or equivalent
l_1
h_1
l_2
h_0
w_0
l_0
Most recent events taken into account

Bound on Open Events

- At most n^2 events from each replica considered open
- This is a tight limit

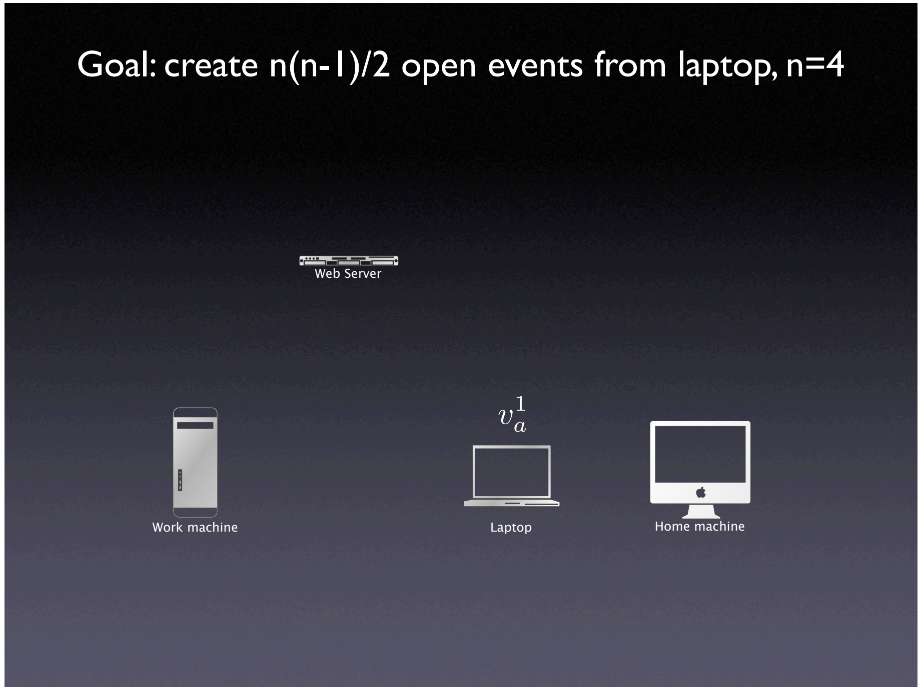

Events considered open shown

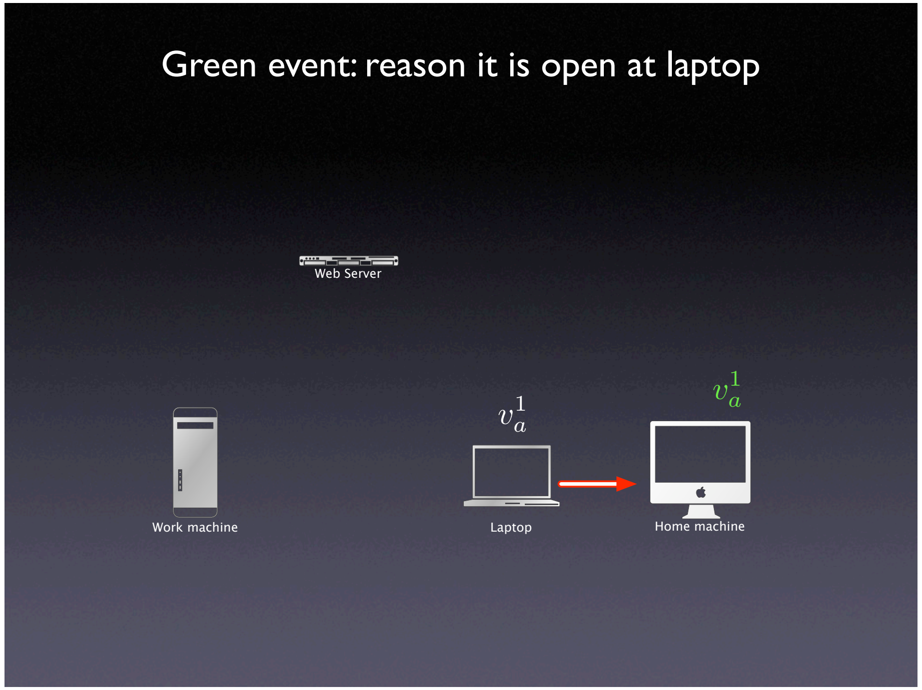

Green: events that have been latest and not communicated back to Laptop

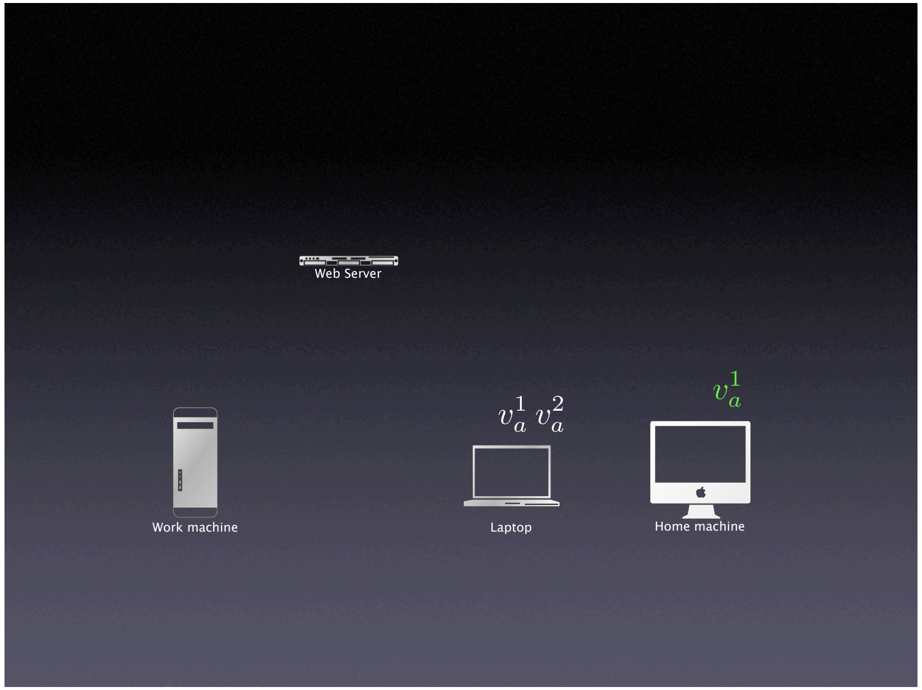
Web Server
v_a^1 v_a^2
v_a^1
Work machine
Laptop
Home machine

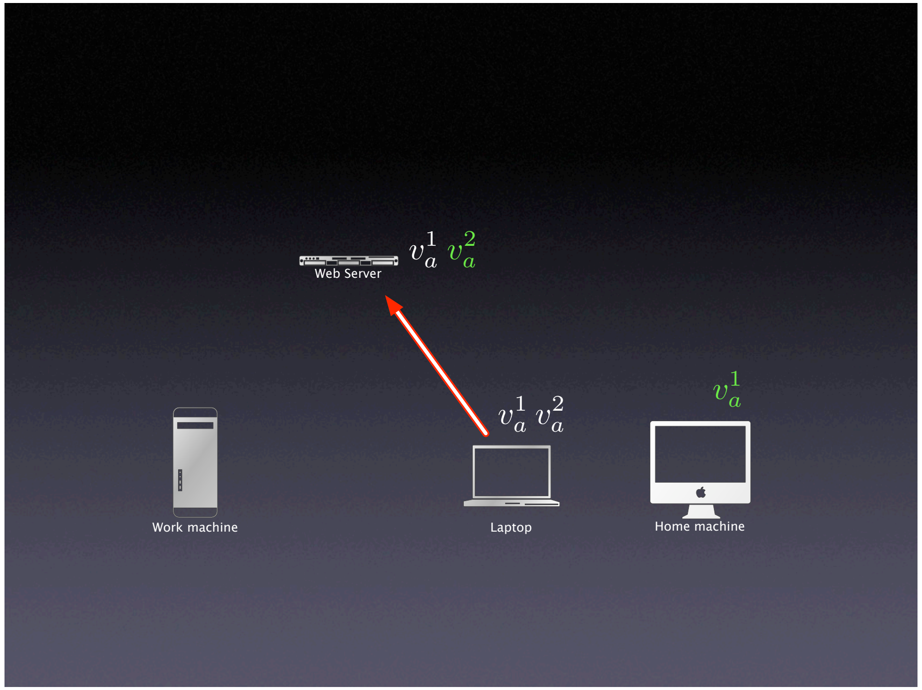
v_a^1 v_a^2
Web Server
v_a^1 v_a^2
v_a^1
Work machine
Laptop
Home machine

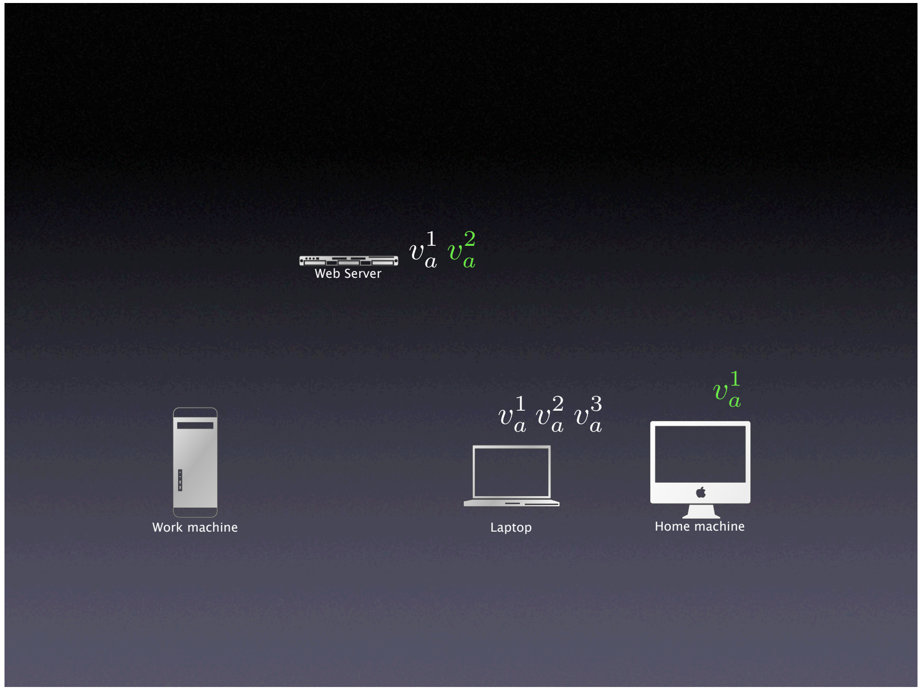

Web Server
v_a^1 v_a^2
v_a^1 v_a^2 v_a^3
v_a^1
Work machine
Laptop
Home machine

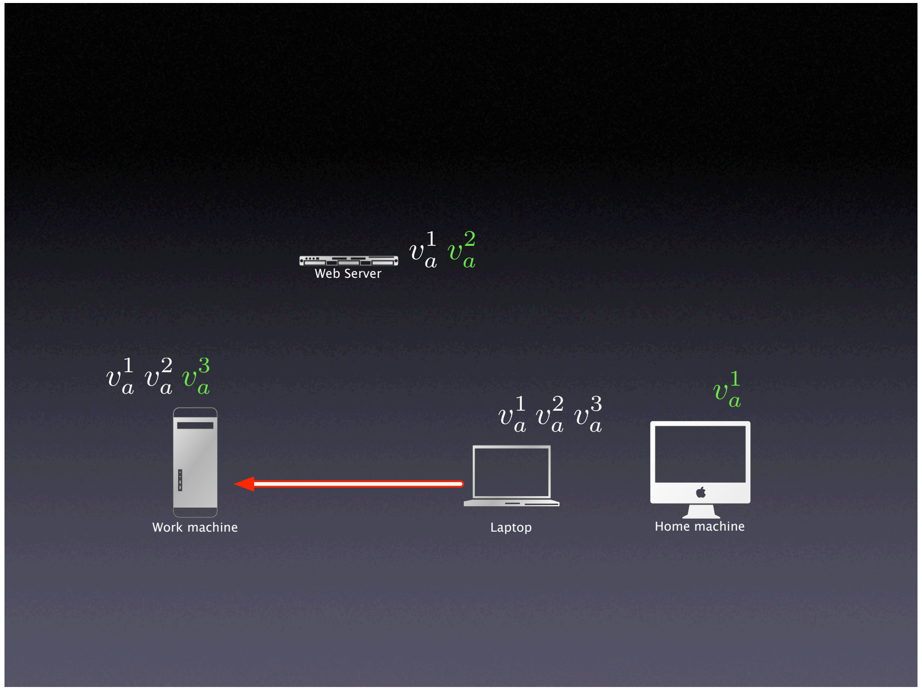

Web Server
v_a^1 v_a^2
v_a^1 v_a^2 v_a^3
Work machine
v_a^1 v_a^2 v_a^3
Laptop
v_a^1
Home machine

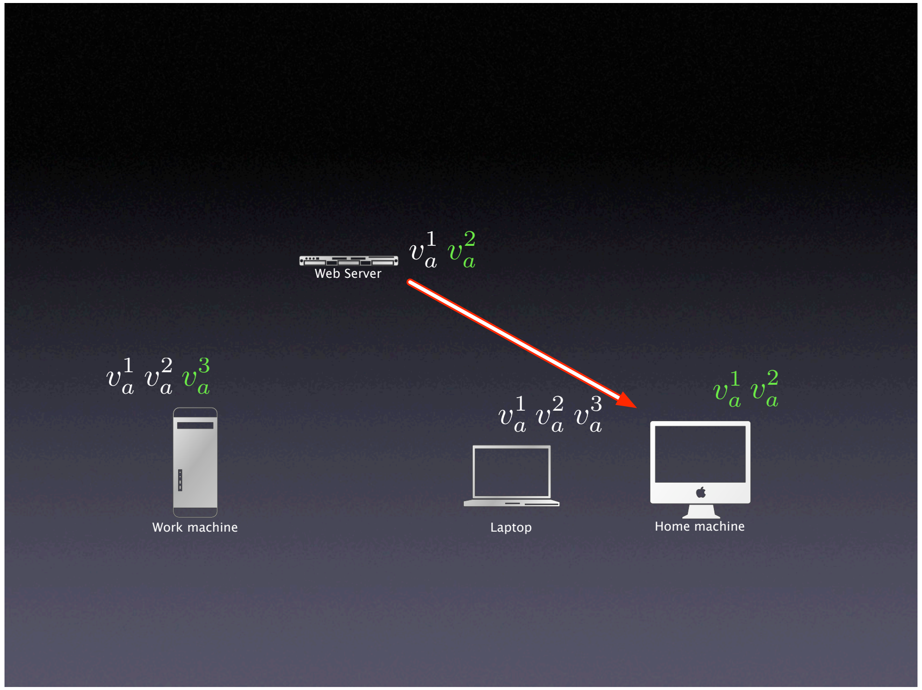

Web Server
v_a^1 v_a^2
v_a^1 v_a^2 v_a^3
Work machine
v_a^1 v_a^2 v_a^3
Laptop
v_a^1 v_a^2
Home machine

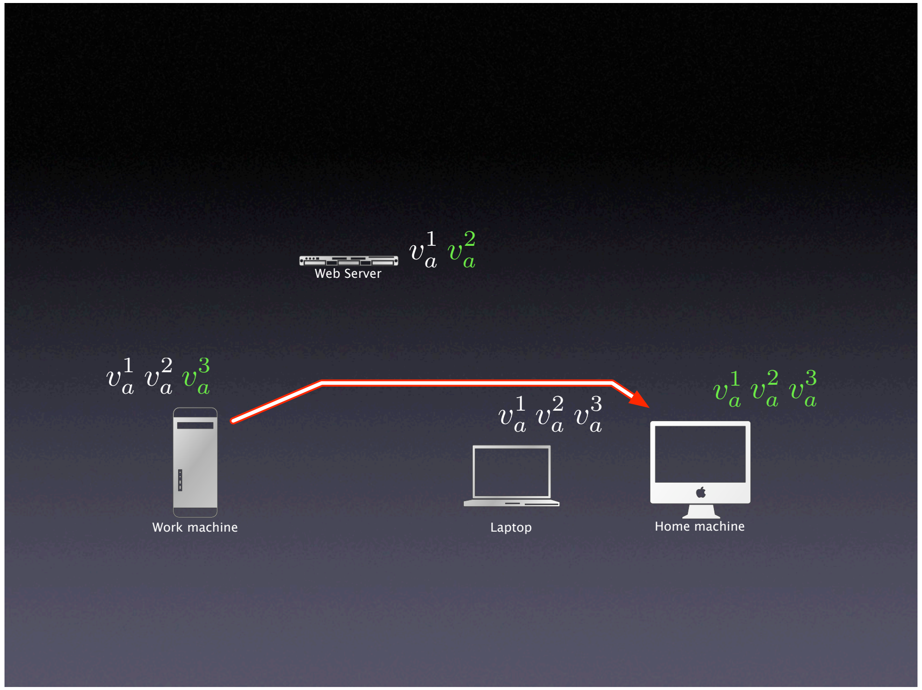

Web Server
v_a^1 v_a^2
v_a^1 v_a^2 v_a^3
Work machine
v_a^1 v_a^2 v_a^3
Laptop
v_a^1 v_a^2 v_a^3
Home machine

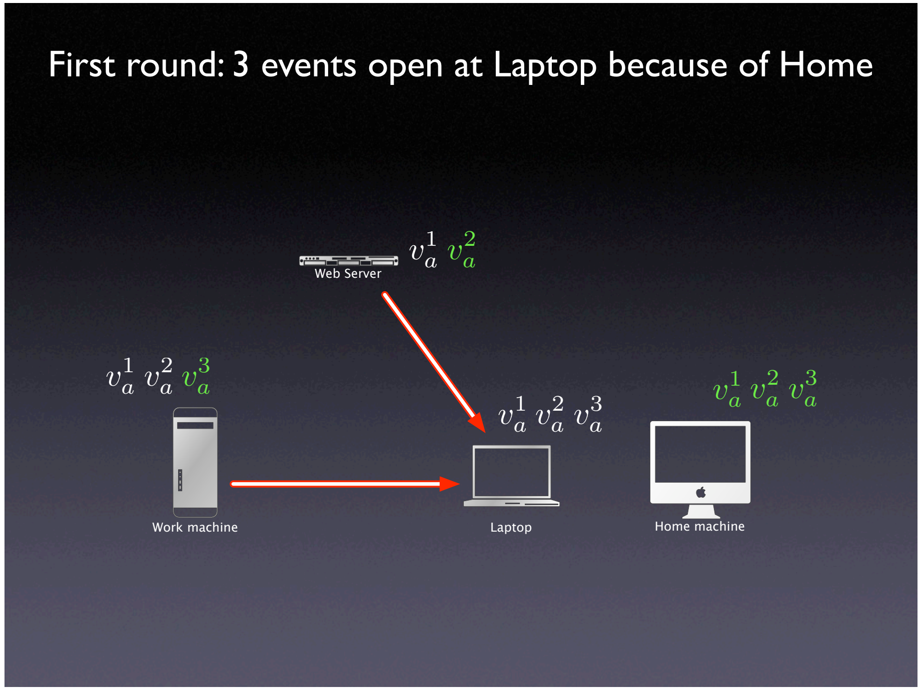

1, 2, and 3 considered open because they have been latest events at the Home machine and the laptop has not (transitively) heard back from it.

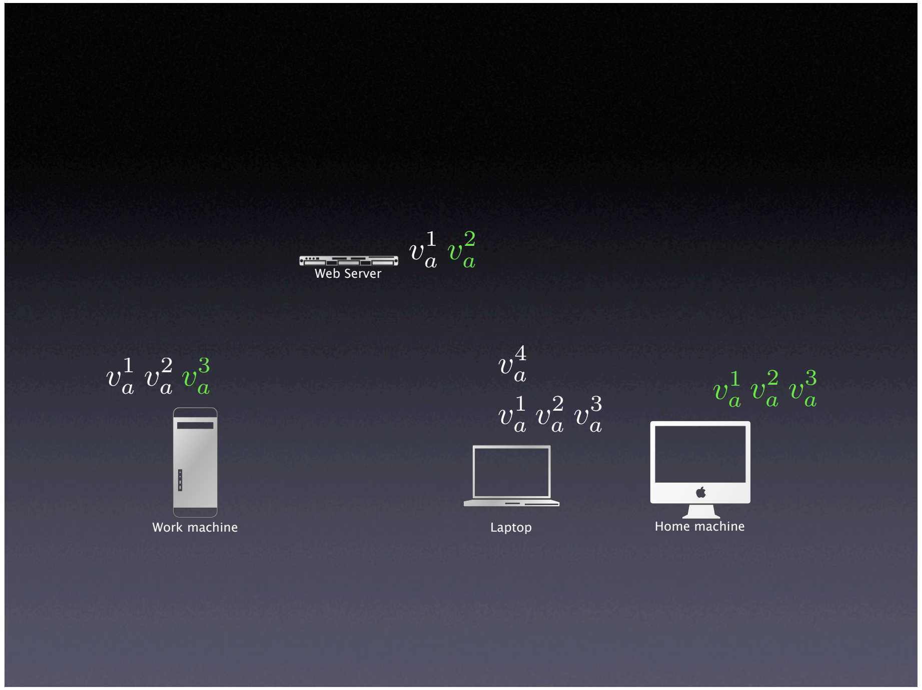

Second round
Cannot send v4 to Home because we have not heard from it

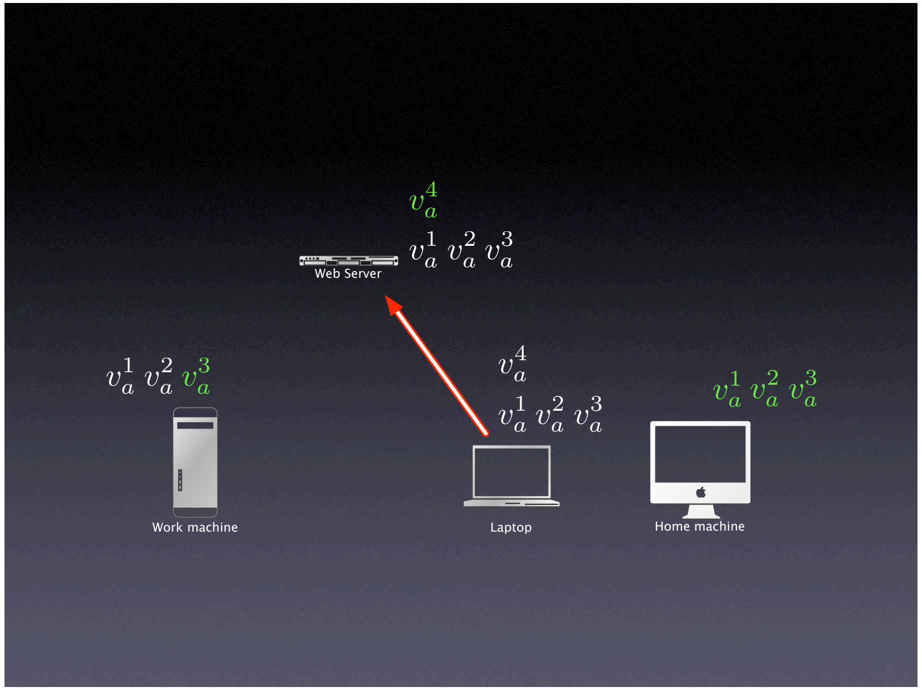

v_a^4
$v_a^1\ v_a^2\ v_a^3$
Web Server
$v_a^1\ v_a^2\ v_a^3$
Work machine
v_a^4
$v_a^1\ v_a^2\ v_a^3$
Laptop
$v_a^1\ v_a^2\ v_a^3$
Home machine

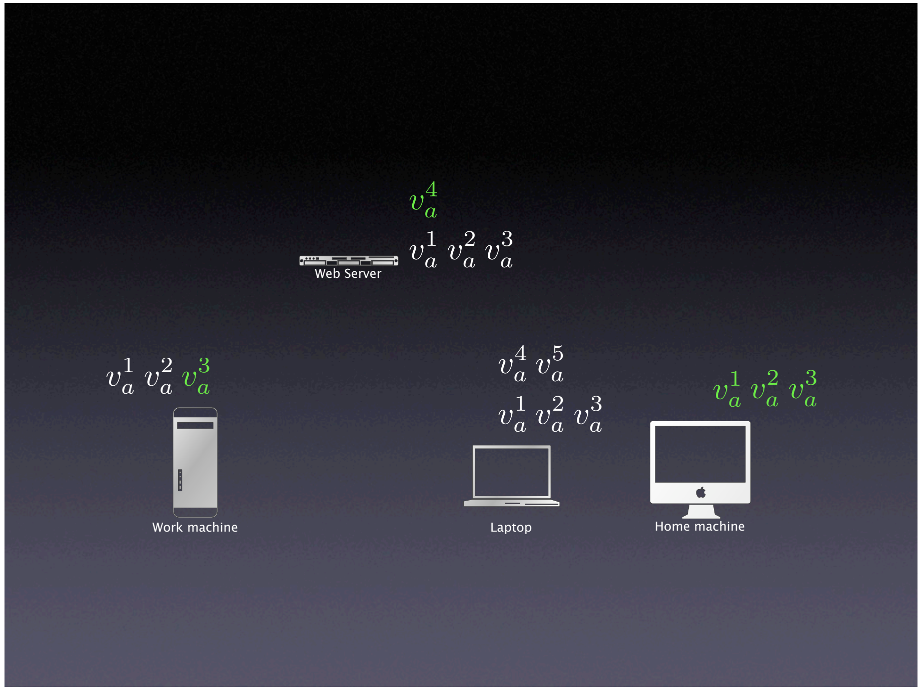

v_a^4
$v_a^1 \; v_a^2 \; v_a^3$
Web Server
$v_a^1 \; v_a^2 \; v_a^3$
Work machine
$v_a^4 \; v_a^5$
$v_a^1 \; v_a^2 \; v_a^3$
Laptop
$v_a^1 \; v_a^2 \; v_a^3$
Home machine

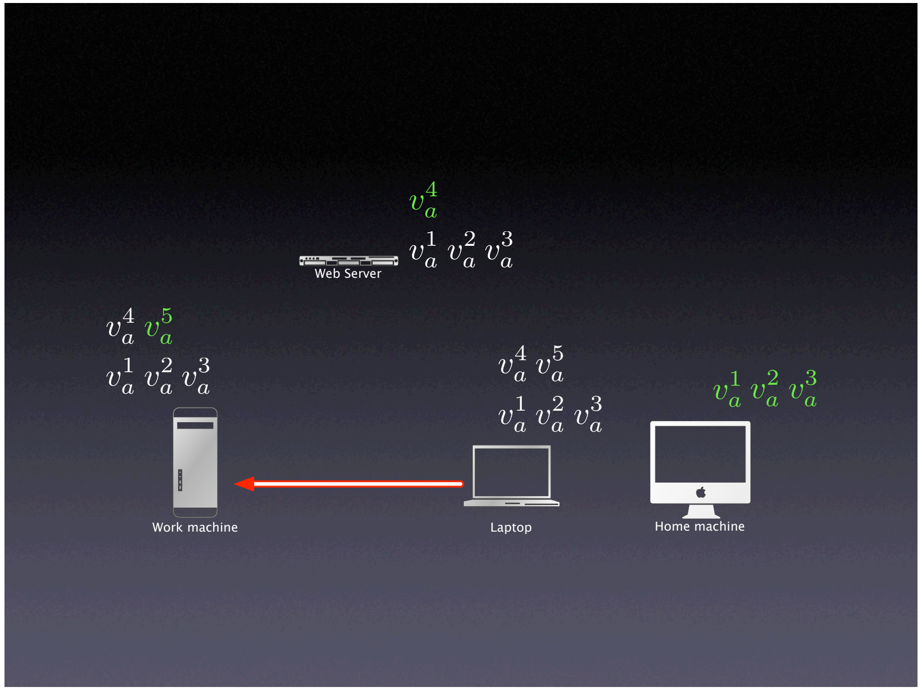
v_a^4
v_a^1 v_a^2 v_a^3
Web Server
v_a^4 v_a^5
v_a^1 v_a^2 v_a^3
Work machine
v_a^4 v_a^5
v_a^1 v_a^2 v_a^3
Laptop
v_a^1 v_a^2 v_a^3
Home machine

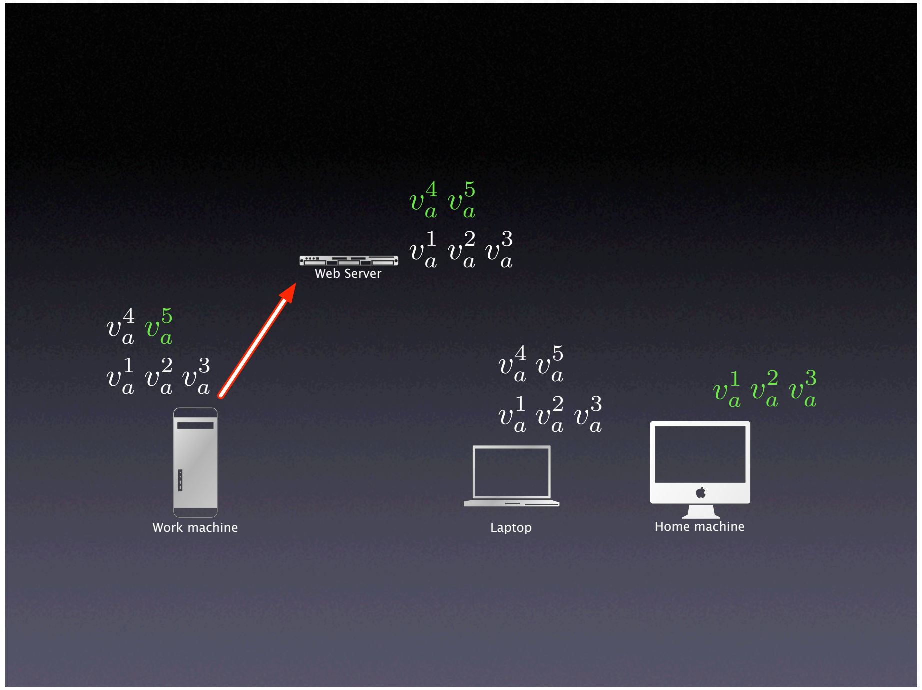

v_a^4 v_a^5
v_a^1 v_a^2 v_a^3
Web Server
v_a^4 v_a^5
v_a^1 v_a^2 v_a^3
Work machine
v_a^4 v_a^5
v_a^1 v_a^2 v_a^3
Laptop
v_a^1 v_a^2 v_a^3
Home machine

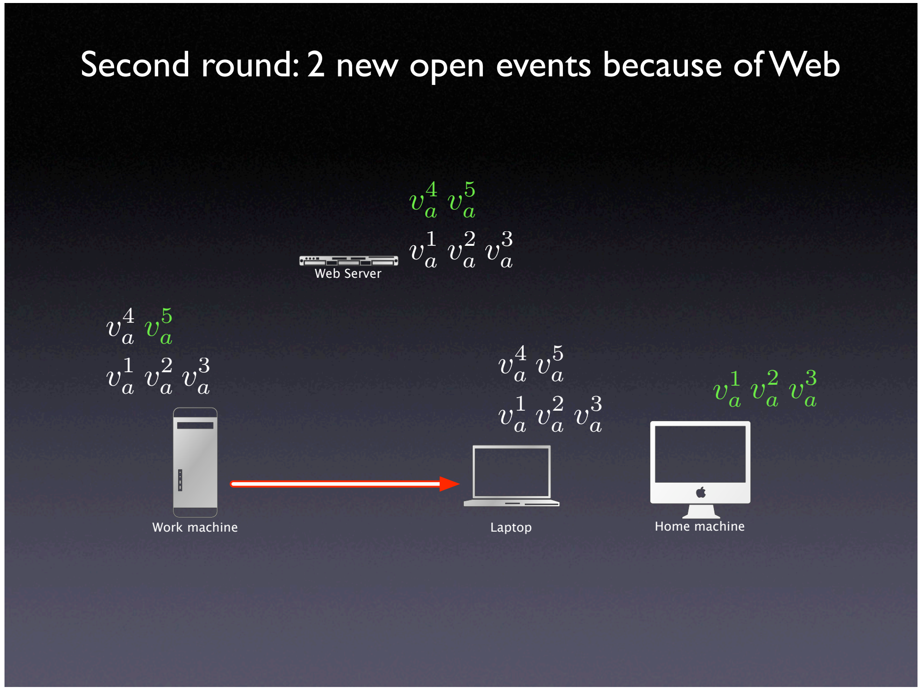
Second round: 2 new open events because of Web
v_a^4 v_a^5
v_a^1 v_a^2 v_a^3
Web Server
v_a^4 v_a^5
v_a^1 v_a^2 v_a^3
Work machine
v_a^4 v_a^5
v_a^1 v_a^2 v_a^3
Laptop
v_a^1 v_a^2 v_a^3
Home machine

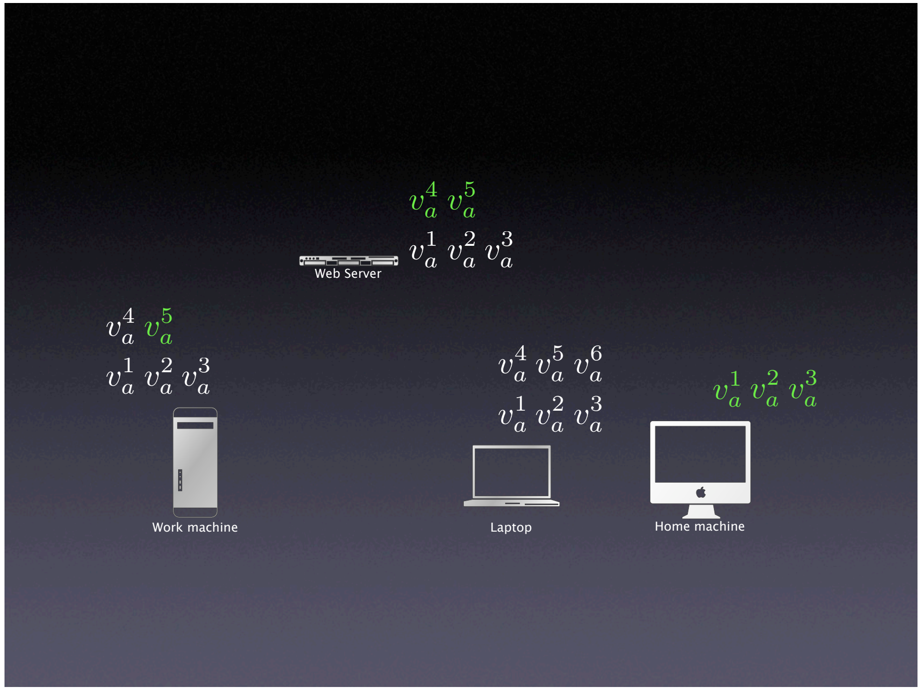

Last round: can only send event to Work

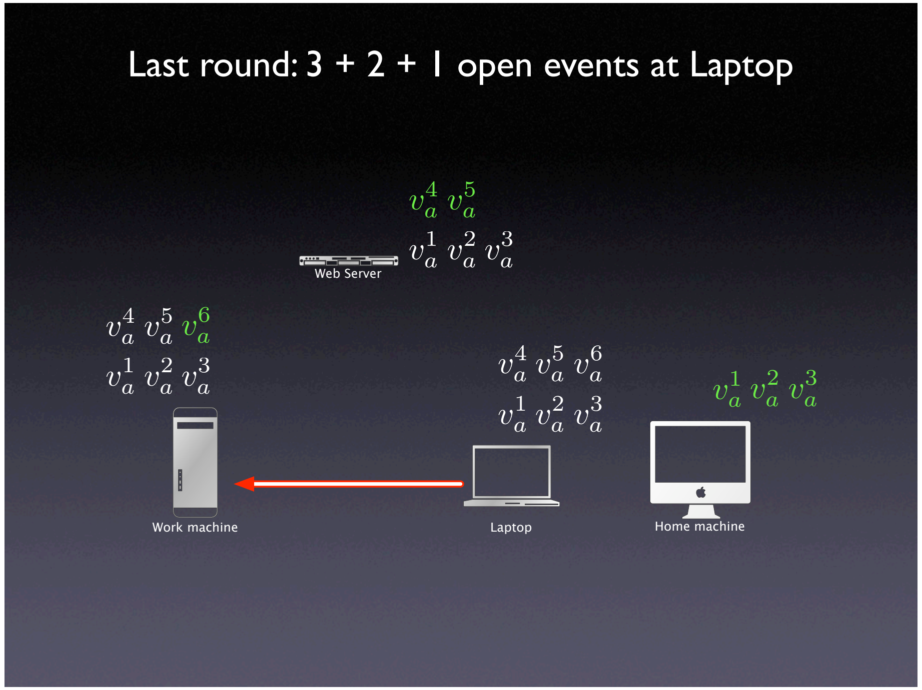

n*(n-1)/2 events considered open at Laptop

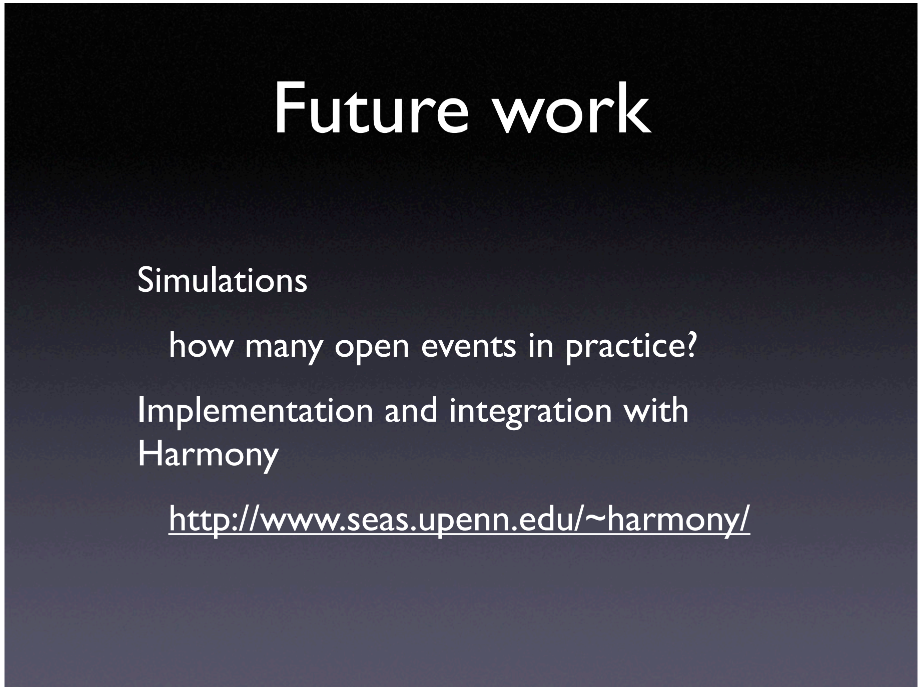

Simulations: how many open events?

Take home points

Very optimistic replication: many conflicts

Equivalence enables optimistic reconciliation

Correct bounded representation in $O(n^4)$ (as worst case)